

Geodesic Distance Descriptors

Gil Shamai and Ron Kimmel

Technion - Israel Institute of Technologies

sgils@tx.technion.ac.il ron@cs.technion.ac.il

Abstract

The Gromov-Hausdorff (GH) distance is traditionally used for measuring distances between metric spaces. It was adapted for non-rigid shape comparison and matching of isometric surfaces, and is defined as the minimal distortion of embedding one surface into the other; while the optimal correspondence can be described as the map that minimizes this distortion. Solving such a minimization is a hard combinatorial problem that requires pre-computation and storing of all pairwise geodesic distances for the matched surfaces. A popular way for compact representation of functions on surfaces is by projecting them into the leading eigenfunctions of the Laplace-Beltrami Operator (LBO). When truncated, the basis of the LBO is known to be the optimal for representing functions with bounded gradient in a min-max sense. Methods such as Spectral-GMDS exploit this idea to simplify and efficiently approximate a minimization related to the GH distance by operating in the truncated spectral domain, and obtain state of the art results for matching of nearly isometric shapes. However, when considering only a specific set of functions on the surface, such as geodesic distances, an optimized basis could be considered as an even better alternative. Moreover, current simplifications of approximating the GH distance introduce errors due to low rank approximations and relaxations of the permutation matrices.

Here, we define the geodesic distance basis, which is optimal for compact approximation of geodesic distances, in terms of Frobenius norm. We use the suggested basis to extract the Geodesic Distance Descriptor (GDD), which encodes the geodesic distances information as a linear combination of the basis functions. We then show how these ideas can be used to efficiently and accurately approximate the metric spaces matching problem with almost no loss of information. We incorporate recent methods for efficient approximation of the proposed basis and descriptor without actually computing and storing all geodesic distances. These observations are used to construct a very simple and efficient procedure for shape correspondence. Experimental results show that the GDD improves both accuracy and

efficiency of state of the art shape matching procedures.

1. Introduction

One line of thought in shape analysis considers an object as a metric space, and object matching, classification, and comparison as the operation of measuring the discrepancies and similarities between such metric spaces, see, for example, [13, 33, 27, 23, 8, 3, 24].

Although theoretically appealing, the computation of distances between metric spaces poses complexity challenges as far as direct computation and memory requirements are involved. As a remedy, alternative representation spaces were proposed [26, 22, 15, 10, 31, 30, 19, 20]. The question of which representation to use in order to best represent the metric space that define each form we deal with, and yet allow for an accurate representation of the mapping from one metric space to another occupied the attention of some recent efforts, see for example [2] and [1]. Indeed, some compact spaces, in the case of matching metric spaces, allow to reduce the complexity and consequently improve the accuracy of the resulting correspondence maps between surfaces. For example, the *Spectral generalized multidimensional scaling* (SGMDS) [3], *functional maps* (FMaps) [26], and the *Iterative Closest Spectral Kernel Maps* (ICSKM) [30] all try to find a linear mapping in a dual space that encodes the minimal distance distortion mapping between two shapes. ICSKM extends the idea of FMaps by finding the best alignment between the spectral kernels of the shapes using iterative closest point (ICP). There, each pair of points produces a constraint, such that much more information is used in the ICP alignment step. ICSKM is initialized using Spectral Gradient fields (SGF) [31], where the normal field of the shape is incorporated, so that some intrinsic symmetries can be identified and resolved.

By trying to match geodesic distances, the SGMDS efficient procedure provides accurate correspondence maps between nearly isometric shapes. Nevertheless, casting the Gromov-Hausdorff related minimization as is into the spectral domain forces a low rank representation of a relaxed

version of the permutation matrix that encodes the correspondence and introduces errors.

When considering the set of all gradient bounded functions on a given manifold, it can be shown that the eigenfunctions of the Laplace-Beltrami operator (LBO) provide an optimal and unique, in a min-max sense, representation for truncated bases [2]. Although not explicitly acknowledged at the time, it motivated its usage in shape matching methods like the SGMDS and functional maps. When dealing with a smaller subset of these functions, such as geodesic distances, there could be a basis that would provide a better representation. We introduce the geodesic distance basis for optimal representation of geodesic distances. We then construct the Geodesic Distance Descriptor (GDD) which encodes the geodesic distances information as a linear combination of the basis functions. The GDD can be seen as a canonical form without the metrication error. It is shown that an approximated metric space matching minimization can be reduced to comparing the GDD of the shapes using ICP procedures [6], [11], without truncation or relaxation. The result is an accurate correspondence permutation matrix. The new linear formulation significantly improves the computational complexity required to solve the shape matching problem. Moreover, when casting the problem on the truncated geodesic distance basis, almost no information is lost. As a stand-alone method, GDD outperforms both SGMDS and functional maps in accuracy, efficiency, and simplicity. When combined with any of the state of the art methods, superior results are obtained.

In Section 2 we define the Gromov-Hausdorff distance and its use for shape correspondence. In Section 3 we define the optimal basis for geodesic distance representation, and relate it to the LBO basis. Next, Section 4 deals with efficiently approximating the geodesic distance basis without actually computing all pairwise geodesic distances. The induced geodesic distance descriptor is defined in Section 5, where we show how it can be used to approximate the solution of the metric space matching minimization problem. In Section 6 we discuss a few initialization and post-processing alternatives to our final shape correspondence procedure that provides state-of-the-art results and presented in Section 7.

2. The Gromov-Hausdorff Distance

Given two shapes \mathcal{S}_1 and \mathcal{S}_2 , consider the map that best preserves the inter-geodesic distances while embedding one shape into the other. The *Gromov-Hausdorff* (GH) distance is defined as the distortion of that embedding. Let, $d_1(s, s')$ and $d_2(q, q')$ represent the inter-geodesic distances between $s, s' \in \mathcal{S}_1$ and $q, q' \in \mathcal{S}_2$, respectively. The GH distance is

defined as

$$d_{GH}(\mathcal{S}_1, \mathcal{S}_2) = \frac{1}{2} \min_{\mathcal{C}} \max_{(s,q) \in \mathcal{C}, (s',q') \in \mathcal{C}} |d_1(s, s') - d_2(q, q')| \quad (1)$$

where

$$\forall s \in \mathcal{S}_1 \exists q \in \mathcal{S}_2 \text{ s.t. } (s, q) \in \mathcal{C}, \quad (2)$$

and

$$\forall q \in \mathcal{S}_2 \exists s \in \mathcal{S}_1 \text{ s.t. } (s, q) \in \mathcal{C}. \quad (3)$$

The set of corresponding points is represented by \mathcal{C} . The set \mathcal{C} could be defined through an indicator function $p(s, q)$ such that $p(s, q) = 1$ if $(s, q) \in \mathcal{C}$ and $p(s, q) = 0$ for $(s, q) \notin \mathcal{C}$. In practice, we detect correspondences between well sampled manifolds, for which we can re-write the optimization problem in matrix notation that reads

$$\arg \min_{P \in \pi(n)} \|PD_1P^T - D_2\|_{\infty}, \quad (4)$$

where $\pi(n)$ is the set of $n \times n$ permutation matrices, and P, D_1, D_2 are the discretizations of $p(s, q), d_1(s, s'), d_2(q, q')$.

Several variations were proposed to reduce the complexity of the problem [8], [9], [21], [25]. In the Generalized Multi-Dimensional Scaling (GMDS) [8], the L_{∞} Hausdorff distance was replaced by an L_2 norm.

SGMDS [3] further simplifies this minimization by relaxing the permutation matrix P and reformulating it as

$$\begin{aligned} & \arg \min_P \|PA_1D_1 - D_2A_2P\|_{\mathcal{S}_1, \mathcal{S}_2} \\ & \text{s.t.} \\ & PA_1\mathbb{1} = \mathbb{1} \\ & P^TA_2\mathbb{1} = \mathbb{1} \end{aligned} \quad (5)$$

where $\|F\|_{\mathcal{S}_1, \mathcal{S}_2} = \text{trace}(F^TA_2FA_1)$, and A_1, A_2 hold the infinitesimal areas about each sample point of $\mathcal{S}_1, \mathcal{S}_2$ along their diagonals. Then, in order to be able to solve this minimization in a practical computational complexity, P, D_1 and D_2 are represented in the truncated spectral domain of the LBO, such that only the first k eigenfunctions are considered. The number of variables in the simplified SGMDS minimization is quadratic in the number of eigenfunctions, so the optimization becomes relatively slow when considering more than $k = 20 - 30$ eigenfunctions. In addition, using only k eigenfunctions to represent the permutation matrix P forces it to be represented as a k -rank matrix, which introduces significant errors to the minimization.

In the following sections of this paper, we will propose an alternative to 4 that does not require any relaxation or truncation of P , while allowing much more eigenfunctions to be incorporated. In addition, we will work with a basis that is optimized for geodesic distance representation, for which truncating the eigenspace almost does not affect the solution.

3. Geodesic Distances Basis

The set of eigenfunctions of the Laplace-Beltrami operator of some manifold \mathcal{M} form a basis, which generalizes the Fourier basis to surfaces and is widely used for representation of functions on manifolds. When considering the set of all gradient bounded functions on \mathcal{M} , the LBO eigenfunctions is the best basis for a truncated representation of this set in a min-max sense [2]. These eigenfunctions were used for a truncated representation of descriptors, permutation functions [26], and geodesic distance functions [3]. However, when considering a subset of these functions, such as geodesic distances, an optimized basis could provide an even better representation. In this section, we define the *Geodesic Distance Basis* and show its usefulness for compactly representing geodesic distance functions. For simplicity, we limit our discussion to the discrete domain with the standard inner product, assuming the shapes were sampled uniformly. All definitions and derivations could be easily extended to the more general case of non-uniform sampling.

Assume we are given a shape \mathcal{S} with n vertices, sampled from a smooth manifold. Let the $n \times n$ symmetric matrix D hold all geodesic distances of \mathcal{S} , such that D_{ij} holds the geodesic distance between $i, j \in \mathcal{S}$. Denote by $D = Q\Lambda Q^T$ the eigenvalue decomposition of D , where the columns of Q are orthonormal and Λ is a real diagonal matrix. Assume that the set of columns of Q are ordered by the magnitude of their corresponding eigenvectors, in a descending order. The k -truncated eigenvalue decomposition of D is defined by $\hat{D}_Q = Q_k \Lambda_k Q_k^T$, where Λ_k holds the k first eigenvalues along its diagonal and Q_k holds the first k corresponding eigenvectors. In general, it is known that the best k -rank approximation of a matrix, in terms of Frobenius norm, is given by computing its k -truncated singular value decomposition [32]. For symmetric matrices, it is equivalent to the k -truncated eigenvalue decomposition. This can be formulated as

$$\hat{D}_Q = \arg \min_{\hat{D} \in \mathcal{K}(n)} \|D - \hat{D}\|_F, \quad (6)$$

where $\mathcal{K}(n)$ is the space of $n \times n$ matrices with rank k .

We term the set of columns of Q as the basis of geodesic distances of the shape \mathcal{S} . Notice that $\hat{D}_Q = Q_k Q_k^T D$. Let the matrix B_k hold k vectors of some other basis. The truncated representation of D in the new basis is obtained by $\hat{D}_B = B_k B_k^T D$, where $B_k^T D$ are the coefficients of representation. The rank of \hat{D}_B is at most k as a product of $n \times k$ matrices. Hence, it cannot approximate D better than \hat{D}_Q . In other words, the truncated reconstruction of D using the geodesic distance basis has the lowest approximation error, in terms of Frobenius norm, among all other bases, independent of the number of vertices n .

The computation of Q is actually not practical when dealing with more than a few thousand points. However, it can be efficiently approximated. In the next section we discuss on how to compute an approximation to Q . In Figure 1 we compare the truncated reconstruction error of the geodesic distances obtained using the basis of the LBO to that of the proposed Geodesic Distance Basis. Here, we computed Q on the Wolf shape from TOSCA, which is the only shape with less than ten thousand vertices. It can be seen that the proposed basis supplies a compact representation with a much better reconstruction.

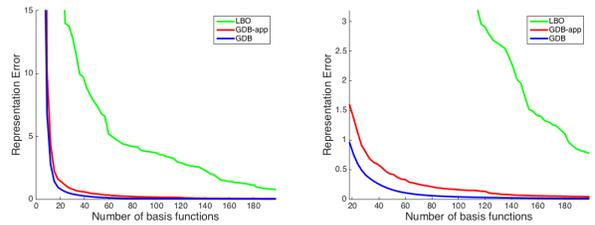


Figure 1. Left - Comparing the reconstruction errors $\|D - \hat{D}\|_F$, where \hat{D} is obtained using the basis of the LBO (LBO), the suggested optimal basis Q (GDB), and its approximation \hat{Q} (GDB-app). The right image is a zoom of the left one.

4. Computing the basis

In order to compute Q and Λ , one would have to compute all pairwise geodesic distances in the large $n \times n$ matrix D , and then perform eigenvalue decomposition to obtain the largest magnitude k eigenvalues and corresponding eigenvectors. The task of computing all pairwise geodesic distances is time consuming and impractical when dealing with more than a few thousand points, even when using efficient methods such as *Fast Marching* [18]. In SGMDs, Spectral MDS (SMDS) [4] was adopted for an efficient computation of D . There, the geodesic distances were computed between roughly 2000 samples of the shape, and the rest of the distances were interpolated by minimizing a Dirichlet energy term while working in the truncated spectral domain of the LBO. The geodesic distances were computed using fast marching and the samples were chosen using the *Fastest Point Sampling* procedure [16].

Recently, an alternative efficient implementation to Multidimensional Scaling (MDS) was suggested in [28, 29]. There, geodesic distances were interpolated from a few of them, similar to SMDS, and with the same complexity. However, no truncated representation in any basis was used, avoiding the errors caused by the truncation, and significantly increasing the accuracy of the approximation. It was shown that only $p = 100$ samples were enough to reconstruct the geodesic distances up to negligible errors. Eventually, the approximation to the pairwise geodesic dis-

tances matrix was written as a product of smaller matrices $S_{n \times k} T_{k \times k} S_{n \times k}^T$, where k is half the size of the number of samples p .

Here, we adopt this idea to compute an approximation to the k first basis functions of the geodesic distance basis. Assume we have a decomposition $\hat{D}_{\tilde{Q}} = STS^T$, obtained from [28], that well approximates D , where S is an $n \times k$ matrix and T is a $k \times k$ matrix. Denote by QR the QR factorization of S , where Q is orthonormal and R is upper triangular. Denote by $V\tilde{\Lambda}V^T$ the eigenvalue decomposition of $RT R^T$. Define $\tilde{Q} = QV$. Then, we have obtained $\hat{D}_{\tilde{Q}} = \tilde{Q}\tilde{\Lambda}\tilde{Q}^T$. Moreover, $\tilde{Q}\tilde{\Lambda}\tilde{Q}^T$ is the k truncated eigenvalue decomposition of $\hat{D}_{\tilde{Q}}$, since $\tilde{\Lambda}$ is diagonal and \tilde{Q} is orthonormal as a product of orthonormal matrices. Finally, notice that $\hat{D}_{\tilde{Q}}$ is the reconstruction of D using the approximated basis is \tilde{Q} .

Next, we measure how well \tilde{Q} approximates the geodesic distances basis Q by comparing the reconstructions $\hat{D}_{\tilde{Q}}$ and \hat{D}_Q . Figure 1 demonstrates that \tilde{Q} can be used instead of the optimal basis Q with almost no effect on \hat{D}_Q . Note that the above procedure supplies us the approximated basis \tilde{Q} and the coefficients $\tilde{\Lambda}$ without the need to compute or store the entire matrix D , but only up to $p = 2k$ columns of it for computing S and T with the method from [28]. The geodesic basis vectors are no more than linear combinations of geodesic distances computed on the surface. In Figure 2 we visualize the first 10 basis vectors (columns of \tilde{Q}) on the Cat shape from TOSCA dataset.

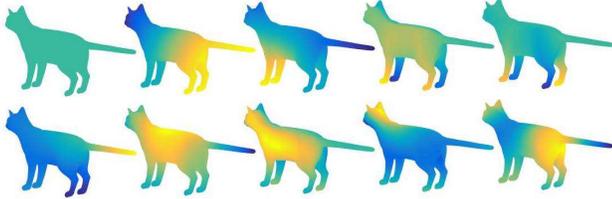


Figure 2. The first 10 eigenvectors of the basis.

5. Geodesic descriptors

Denote by the diagonal matrix W the square root of the diagonal matrix Λ , such that

$$W_{ii} = \sqrt{\Lambda_{ii}}. \quad (7)$$

Define

$$X = QW, \quad (8)$$

such that,

$$D = XX^T. \quad (9)$$

We term X as the *Geodesic Distance Descriptor* (GDD). Note, that in general X is complex. X can be used as a

point descriptor that encodes the geodesic distances. Since it stems from geodesic distances, it is not affected by isometric deformations of the shape. X holds all information of the geodesic distances D , and therefore could be used instead of D to find the solution for 4. Moreover, X is already represented in the basis Q since $QQ^T X = X$, so only a few columns of X contain almost all information encapsulated in D . This was demonstrated in Figure 1 by the reconstruction error of D as a function of number of eigenvectors. The GDD can be therefore used for dimensionality reduction and simplification tasks that involve geodesic distances. In a sense, the GDD can be thought of a canonical form obtained using Multidimensional scaling [7], but without the embedding errors that occur because of flattening a curved surface.

Each point i in the shape corresponds to a row x^i in X . x^i can be referred to as the descriptor of point i . The descriptor of a point i is invariant to vertex ordering of rest of the points. Hence, the GDD can be used as a point descriptor to find the correspondence between two shapes. Denote by $E_{ij} = \|x^i - x^j\|_2$ the Euclidean distance between the descriptors of points i and j . Figure 3 shows the relation between E_{ij} and the geodesic distance D_{ij} . For comparison, we show the same analysis for Φ instead of X , where Φ holds the eigenvectors of the LBO. The upper cat shapes in Figure 4 visualize $E_i(j)$ for a selected vertex i marked in red, for both X and Φ . It seems that the geodesic descrip-

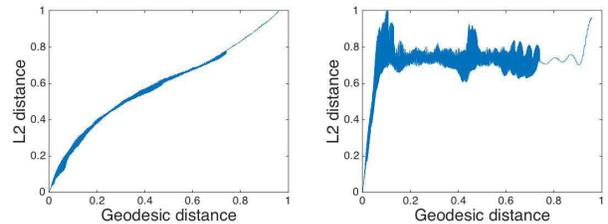


Figure 3. E_{ij} w.r.t D_{ij} . The left figure corresponds to X , and the right one corresponds to Φ .

tors have monotonic relation with the geodesic distances, and thus more robust to large correspondence errors.

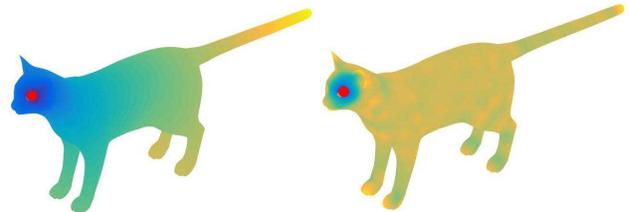


Figure 4. The cats display the function $E_i(j)$ derived from X (left) and Φ (right), for a selected point i marked in red.

An alternative formulation

Note that the eigenvalue decomposition has ambiguities. When assuming no repeating eigenvalues, the eigenvectors in an eigenvalue decomposition are unique up to sign flips. This can be formulated as $Q = Q_0 C$, where C is a diagonal sign matrix, and Q_0 and Q are two possible eigenvectors matrices. More generally, C is an orthogonal matrix, representing rotation ambiguities that correspond to repeating eigenvalues. In any case, the non-diagonal elements of C correspond to the locations of the repeating eigenvalues on the diagonal of W . It is therefore possible to swap between C and W , such that

$$X = QW = Q_0 C W = Q_0 W C = X_0 C, \quad (10)$$

where X_0 and X are two possible derived GGD-s. Therefore, the proposed geodesic distance descriptor is invariant to isometric deformations up to a rotation and reflection ambiguity XC .

Consider two isometric shapes with corresponding descriptors X_1 and X_2 , and some correspondence encoded by the permutation matrix P . Plugging $D_1 = X_1 X_1^T$ and $D_2 = X_2 X_2^T$ in 4, we obtain

$$\arg \min_{P \in \pi(n)} \|P X_1 X_1^T P^T - X_2 X_2^T\|_\infty. \quad (11)$$

It appears that this minimization can be reduced to solving

$$\arg \min_{P \in \pi(n), C \in \mathcal{O}(n)} \|P X_1 C - X_2\|_{2,\infty}, \quad (12)$$

where $\|F\|_{2,\infty}$ stands for the maximal L_2 norm of any row in F , and $\mathcal{O}(n)$ is the set of $n \times n$ orthogonal matrices. For isometric shapes, the two minimizations are equivalent. For nearly isometric shapes, the solution of 12 approximates the one of 4, up to some bound. If we manage to find a good solution to 12, it guarantees some bound on the minimizer of 4. More details regarding the bounds can be found in the supplementary material. In other words, solving the complex GH related minimization in 4 is nothing but matching the geodesic descriptors of two shapes, under the best rotation.

The minimization in 12 can be efficiently solved using methods like Iterative Closest Point (ICP), with a quasi-linear complexity in the number of points, using efficient approximations such as kd-tree. Moreover, since X_1 and X_2 are already represented in the geodesic distance basis, it is enough to consider only their few k first columns, with almost no effect on the solution. In our experiments, $k = 50$ were enough for this task. Note that ICP finds for each vertex in one shape a matching vertex in the other shape, rather than a bijective map. However, this makes more sense when dealing with two discrete shapes that might have been sampled differently from their corresponding manifolds.

6. Initializations and post processing

Suppose we treat the descriptors X_1 and X_2 as two point clouds, where each row is a point, and the orthogonal matrix C is a rotation and reflection of the points. ICP ([6], [11]) is an efficient optimization that attempts to compute the best match between two point clouds, under any rotation and reflection. In practice, ICP iterates between point match and rotation alignment:

1. For each point in the first cloud, find its nearest point in the second cloud, in terms of Euclidean distance.
2. Find the rotation and reflection that best aligns the matched points. This step has a closed form by using the Procrustes Theorem [14].

We suggest to find the solution of 12 by applying ICP to X_1 and X_2 . ICP usually requires a good initialization. We propose the following alternatives.

Initialization using correspondence

Assume we have some initial correspondence given as an output from another shape matching procedure. Plugging this initialization, it is possible to start with step 2 of ICP, and continue to iterate. In our experiments, we used correspondences found by other methods as initializations, and managed to outperform any state of the art method.

Initialization using descriptors

An alternative way is to start with some initial estimate of C , denoted here as C_0 . In functional maps, ICP between the eigenfunctions of the LBO was performed as a post processing step, while computing the initialization C_0 was the core of the method. First, different descriptors were computed for each of the shapes. Then, C_0 was treated as a linear map between the coefficients of the descriptors in the LBO basis, and obtained using a simple least-squares minimization. At first sight, translating these steps to our problem is direct - compute C_0 using the coefficient of the descriptors in the basis of geodesic distances Q instead of the LBO basis, and proceed similarly. However, unlike functional maps, in our case the matrix C encodes the deformation between the descriptors X_1 and X_2 which are not orthonormal bases (the columns are orthogonal but not normalized). To that end, suppose that f_1 and f_2 are column vectors corresponding to some point descriptors on shapes S_1 and S_2 . The orthogonal permutation matrix P encodes a mapping between the shapes and can be therefore used for denoting $f_2 = P f_1$. Since P is orthogonal we can write instead $P^T f_2 = f_1$ or $f_2^T P = f_1^T$. Denote by F_1 and F_2 the coefficients of f_1 and f_2 in the geodesic distance basis representation, i.e. $F_1 = Q_1^T f_1$ and $F_2 = Q_2^T f_2$. Suppose that we seek for some matrix C that encodes the relation

between PX_1 and X_2 as

$$PX_1C = X_2. \quad (13)$$

By multiplying both sides of the equation by f_2^T we obtain

$$f_2^T PX_1C = f_2^T X_2. \quad (14)$$

Notice that

$$f_2^T PX_1 = f_1^T X_1 = f_1^T Q_1 W_1 = F_1^T W_1. \quad (15)$$

Hence, Equation 14 can be further reduced to

$$F_1^T W_1 C = F_2^T W_2. \quad (16)$$

Then, we could instead search for C that encodes the relation between $F_1^T W_1$ and $F_2^T W_2$, which is independent of the mapping P . This can be defined as the minimization

$$\arg \min_C \|F_1^T W_1 C - F_2^T W_2\|. \quad (17)$$

In other words, we propose to find an approximation to C by repeating the procedure of functional maps, while using the basis Q instead of the LBO basis, and multiplying the coefficients of the descriptors by the square-root of the eigenvalues, W .

Initialization using feature points

Assume we have an initial set of m points in shape \mathcal{S}_1 that correspond to m points in shape \mathcal{S}_2 . To find the initial m point correspondence, for example, in SGMDS it was suggested to first find in each shape a set of points that are locally farthest from the rest of the points, and then match the candidates using descriptors such as WKS [5]. It was noted that $m = 5$ points were enough for a good initialization of SGMDS.

m point correspondences can be considered as m rows in X_1 that correspond to m rows in X_2 . Denote the submatrices that correspond to these rows by \hat{X}_1 and \hat{X}_2 . Assume that the columns of X_1 and X_2 are ordered by the size of their corresponding eigenvalues, in a descending order.

For isometric shapes, assuming non repeating eigenvalues, C would be a diagonal matrix. If the shapes are approximately isometric, C would have a sparse and diagonally dominant structure. This effect was already demonstrated in functional maps. Hence, it is possible to estimate C by solving

$$\arg \min_{C \in \mathcal{U}(n)} \|\hat{X}_1 C - \hat{X}_2\|, \quad (18)$$

while adding some off-diagonal penalty. In fact, it is enough to estimate only the first rows and columns of C , and then obtain an initial correspondence using only the first corresponding columns of X_1 and X_2 . In our experimental setting, we used the same $m = 5$ point correspondences that were used in SGMDS, with which we approximated the 20×20 first rows and columns of C .

Post Processing

The correspondence obtained using minimization 12 is robust to large geodesic distances errors, since they would penalize the objective function. This was demonstrated earlier by showing that the GDD has a point signature that is unique to the point (Figures 3 and 4). Other bases or descriptors, however, could produce a signature with a more local nature. These descriptors can further improve the solution by combining them with the GDD. One simple way used in our experimental results is to refine the correspondence by performing ICP on the LBO basis, initialized with the correspondence obtained by our method. As this basis appears to be better localizer than the GDD, superior results are obtained for correspondence when combining the two. However, note that while the correspondence improves, this post-processing harms the approximation of the minimizer of 4 (see experimental results). This is not surprising as the GDD is related to the minimization of 4.

7. Results

Throughout this section, we refer to our proposed method as the Geodesic Distance Descriptor (GDD), and compare it to methods discussed in the introduction and throughout the paper. In our experiments we used shapes from the publicly available datasets TOSCA [9] and SCAPE [12] that contain real and synthetic human and animal poses. For accuracy comparison of shape correspondence we use the evaluation procedure proposed by Kim et al. [17]. To the best of our knowledge, the state of the art methods for efficiently computing correspondences of nearly isometric shapes are Spectral-GMDS [3], functional maps [26], and Spectral-GF with ICSKM refinement [31, 30]. For GDD, we used $k = 50$ basis functions computed from 100 samples in each surface.

In our first experiment, given two shapes, we computed their point correspondence by applying ICP to their geodesic descriptors, using the same 5 points initializations used for SGMDS. This was done with and without the post-processing step suggested in section 6 (GDD+post and GDD-5pt). We compared the results to functional maps (FMaps), SGMDS and Spectral-GF (SGF). We repeated the same experiment using the LBO basis instead of the X (Phi-5pt). The results are shown in Figure 5. It can be seen that GDD performs better than FMaps with a much simpler initialization, without the need of the descriptors. Compared to SGMDS, It can be seen that GDD performs better for the same initialization and with a much simpler procedure, without the spectral formulation for the optimization of the GH related minimization. In addition, our procedure required computing much less geodesic distances than SGMDS (see Section 4, and Figure 6). The performance of SGF is better than GDD when comparing large geodesic er-

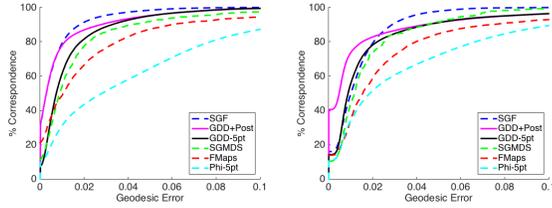


Figure 5. Quantitative evaluation of shape correspondence methods applied to the shapes from the TOSCA and SCAPE datasets, using the protocol from Kim et al.

rors, probably because the initialization step of SGF is more sophisticated and incorporates more information about the shape, such as the normal vector at each point. A better initialization for GDD can significantly improve the results, see Figure 9.

Next, Figure 6 presents the computation time of the GDD, SGMDS, FMaps and the SGF methods applied to four shapes from the TOSCA database with different numbers of vertices. The computations were evaluated on a 2.8 GHz i7 Intel computer with 16GB RAM. In this experiment, the evaluated GDD is the same as that used for the results in figures 5, 7 and 8, as a stand-alone method, including all initialization techniques. In practice, the post processing step for the GDD should not add time as it is done by simply replacing the GDD with the eigen-functions of the LBO within the ICP iterations.

In the benchmark of Kim et al. the correspondence between shapes is assumed to be provided. The geodesic distance of each point mapped by the method we evaluate from what is referred to as true location is computed. The distortion curves describe the relative number of points falling within a relative geodesic distance from what is assumed to be their exact location. Notice that the given “exact” location is, in fact, a subjective measure. The distortion curves contain an intrinsic ambiguity of up to about 25% as there is no exact isometry between objects at different poses.

In the next experiment, we compute the correspondence P as

$$\min_P \|PD_1P^T - D_2\|_F^2. \quad (19)$$

Since we cannot really compute all pairwise geodesic distances, we evaluated the result by considering only 1000 randomly sampled rows and columns of the matrix inside the above norm. We state the results in Figure 7. GDD and Post stand for GDD-5pt and GDD+post were discussed in the previous experiment. It can be seen that the proposed GDD performs best as a method for approximating the minimizer of 19. Note that the post processing step damages the approximation, since it involves the LBO basis for the sake of localizing the correspondence. Surprisingly, GDD approximates 19 even better than the “ground truth” corre-

# of vertices	4344	19248	27894	52565
Spectral-GMDS				
5% G-Distances	28.6	476.1	998.3	4981.2
100 LBO eigs	1.8	5.5	9.7	22.9
SGMDS	14.5	15.9	17.7	24.8
Total	45	498	1022	4961
Functional Maps				
100 LBO eigs	2.5	9.9	14.8	28.9
FMaps	6.6	76.9	46.3	34.4
ICP	1.5	15.2	17.4	25.3
Total	11	102	79	88
Spectral-GF				
200 LBO eigs	5.1	12	17.2	33.9
SGF	109.7	122.7	123.2	131.2
Total	115	135	140	165
GDD				
100 G-Distances	1.9	11.8	16.4	44.7
GDD	0.3	0.7	0.8	1.3
ICP	0.8	3.3	4.9	7.9
Total	3	16	22	54

Figure 6. Computation times (in seconds) of different methods. G-Distances stands for the geodesic distance computation. In SGMDS 5% of the shape points were used, whereas in GDD 100 points were used for this step. LBO-eigs stands for the LBO eigenvectors computation. Note that for each method, different versions of the LBO and different procedures for computing the eigenvectors were chosen. The ICP step of the GDD is faster than the one of the FMaps since only 1000 random samples of the shapes were enough to refine the matrix C . In FMaps, some of the descriptors, such as finding segments matches, required iterative optimization, and the computation times significantly varied for different shapes.

	GT	SGF	SGMDS	FMaps	GDD	Post
Horses	470.1	518.7	707.4	1418.2	397.3	483.5
Victorias	146.2	149.4	202.1	147.6	128.5	150.7
Cats	160.4	178.9	178.4	189.7	123.5	154.5
Wolfs	9.32	9.35	9.04	9.42	7.84	9.36
Centauris	153.5	174	151.5	771.9	122.7	154.4
Dauids	58.6	58.6	66.2	62.4	50	58.5

Figure 7. Comparing correspondences as minimizers of 19.

spondence provided by Kim et al. It implies that these two measures might not always align for non-isometric shapes.

Next, we computed point-to-point correspondences between five Michael shapes from the TOSCA dataset. We then colored each shape according to the Voronoi regions of a set of 20 points. Note that the Voronoi diagram was generated separately for each shape after mapping the set of 20 points. The results are shown in Figure 8. Notice that some of the shapes were mapped into their intrinsic symmetries, as the objective 4 cannot differentiate between them and both solutions are optimal.

In our final experiment, we used the GDD to find the correspondence while initializing it with the correspondences computed by SGMDS (GDD+SGMDS), functional maps (GDD+FMaps) and Spectral-GF (GDD+SGF). Figure

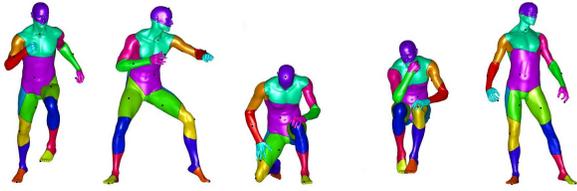


Figure 8. Visualization of a Voronoi diagram on different poses of the Michael shape, as computed by the proposed method.

9 shows the results. It can be seen that using the GDD it is

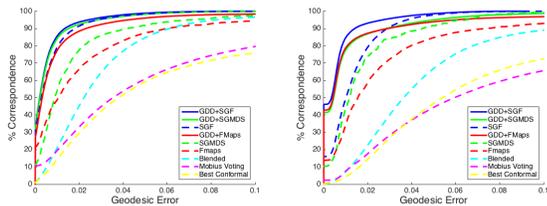


Figure 9. Quantitative evaluation of shape correspondence methods applied to the shapes from the TOSCA and SCAPE datasets, using the protocol of Kim et al.

possible to successfully improve each of the methods, and thus reach state of the art results for the nearly isometric shape correspondence challenge.

8. Conclusions

The main contributions of this paper can be summarized by:

1. Definition of a new basis that is optimized for geodesic distances representation. We also showed how to efficiently approximate it.
2. Definition of a generalized canonical form that does not suffer from embedding errors, and contains all information about the geodesic distances. We termed it as the geodesic distance descriptor.
3. An alternative formulation for the approximated GH distance related minimization of nearly isometric shapes using geodesic distance descriptors, that is both efficient and does not require relaxation or truncation of the permutation matrices.
4. Introduction of a shape correspondence procedure that obtains state of the art results for matching nearly isometric shapes.

We introduced an efficient and accurate model for finding the best correspondence between two metric spaces. The proposed method does not involve any relaxation or

truncation of the eigenspace in which the permutation matrix is encoded. The new formulation bridges the gaps between Spectral-GMDS, functional maps, and canonical forms, by introducing the Geodesic Distance Descriptors. The geodesic distance descriptor can be used for dimensionality reduction of tasks that involve geodesic distances. These distances are translated into a compact representation which is invariant to the order of vertices. An optimal basis is proposed whose computation is based on recent methods for geodesic distance approximations. Experimental results show that while the accuracy of the metric space matching minimizer improves, the accuracy of the correspondence, as evaluated by a given manually-labeled pairs of corresponding points, does not necessarily improve. This finding suggests that the two measures are not necessarily the same. Finally, it was shown that geodesic distance descriptor can be used to obtain state of the art matching results for nearly isometric shapes.

9. Acknowledgments

The research leading to these results was supported by European Communitys FP7- ERC program, grant agreement no. 267414.

References

- [1] Y. Aflalo, H. Brezis, A. Bruckstein, R. Kimmel, and N. Sochen. Best bases for signal spaces. *Comptes Rendus Mathematique*, 2016. 1
- [2] Y. Aflalo, H. Brezis, and R. Kimmel. On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences*, 8(2):1141–1160, 2015. 1, 2, 3
- [3] Y. Aflalo, A. Dubrovina, and R. Kimmel. Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision*, 118(3):380–392, 2016. 1, 2, 3, 6
- [4] Y. Aflalo and R. Kimmel. Spectral multidimensional scaling. *Proceedings of the National Academy of Sciences*, 110(45):18052–18057, 2013. 3
- [5] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 1626–1633. IEEE, 2011. 6
- [6] P. J. Besl and N. D. McKay. Method for registration of 3-d shapes. In *Robotics-DL tentative*, pages 586–606. International Society for Optics and Photonics, 1992. 2, 5
- [7] I. Borg and P. J. Groenen. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2005. 4
- [8] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5):1168–1172, 2006. 1, 2
- [9] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. 2, 6

- [10] Q. Chen and V. Koltun. Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2039–2047, 2015. **1**
- [11] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992. **2, 5**
- [12] P. S. Dragomir Anguelov, H.-C. Pang, D. Koller, and J. D. Sebastian Thrun. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Advances in Neural Information Processing Systems 17: Proceedings of the 2004 Conference*, volume 17, page 33. MIT Press, 2005. **6**
- [13] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE Transactions on pattern analysis and machine intelligence*, 25(10):1285–1295, 2003. **1**
- [14] J. C. Gower and G. B. Dijkstra. *Procrustes problems*. Number 30. Oxford University Press on Demand, 2004. **5**
- [15] X. Gu, Y. Wang, T. F. Chan, P. M. Thompson, and S.-T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, 23(8):949–958, 2004. **1**
- [16] 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. **3**
- [17] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. In *ACM Transactions on Graphics (TOG)*, volume 30, page 79. ACM, 2011. **6**
- [18] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences*, 95(15):8431–8435, 1998. **3**
- [19] R. Lai and H. Zhao. Multi-scale non-rigid point cloud registration using robust sliced-wasserstein distance via laplace-beltrami eigenmap. *arXiv preprint arXiv:1406.3758*, 2014. **1**
- [20] H. Ling and D. W. Jacobs. Shape classification using the inner-distance. *IEEE transactions on pattern analysis and machine intelligence*, 29(2):286–299, 2007. **1**
- [21] Y. Lipman and I. Daubechies. Surface comparison with mass transportation. *arXiv preprint arXiv:0912.3488*, 2009. **2**
- [22] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. In *ACM Transactions on Graphics (TOG)*, volume 28, page 72. ACM, 2009. **1**
- [23] F. Memoli. On the use of gromov-hausdorff distances for shape comparison. 2007. **1**
- [24] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5(3):313–347, 2005. **1**
- [25] I. Navazo, P. Poulin, J. Pokrass, A. Bronstein, M. Bronstein, P. Sprechmann, and G. Sapiro. Sparse modeling of intrinsic correspondences. **2**
- [26] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30, 2012. **1, 3, 6**
- [27] E. L. Schwartz, A. Shaw, and E. Wolfson. A numerical solution to the generalized mapmaker’s problem: flattening non-convex polyhedral surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(9):1005–1008, 1989. **1**
- [28] G. Shamai, Y. Aflalo, M. Zibulevsky, and R. Kimmel. Classical scaling revisited. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2255–2263, 2015. **3, 4**
- [29] G. Shamai, M. Zibulevsky, and R. Kimmel. Fast classical scaling. *arXiv preprint arXiv:1611.07356*, 2016. **3**
- [30] A. Shtern and R. Kimmel. Iterative closest spectral kernel maps. In *3D Vision (3DV), 2014 2nd International Conference on*, volume 1, pages 499–505. IEEE, 2014. **1, 6**
- [31] A. Shtern and R. Kimmel. Spectral gradient fields embedding for nonrigid shape matching. *Computer Vision and Image Understanding*, 140:21–29, 2015. **1, 6**
- [32] G. W. Stewart. On the early history of the singular value decomposition. *SIAM review*, 35(4):551–566, 1993. **3**
- [33] G. Zigelman, R. Kimmel, and N. Kiryati. Texture mapping using surface flattening via multidimensional scaling. *IEEE Transactions on Visualization and Computer Graphics*, 8(2):198–207, 2002. **1**