

Hierarchical Framework for Shape Correspondence

Dan Raviv*, Anastasia Dubrovina and Ron Kimmel

Technion, Israel Institute of Technology, Technion City, Haifa 32000, Israel.

Received 18 December 2011; Accepted (in revised version) 1 July 2012

Available online 11 January 2013

Abstract. Detecting similarity between non-rigid shapes is one of the fundamental problems in computer vision. In order to measure the similarity the shapes must first be aligned. As opposite to rigid alignment that can be parameterized using a small number of unknowns representing rotations, reflections and translations, non-rigid alignment is not easily parameterized. Majority of the methods addressing this problem boil down to a minimization of a certain distortion measure. The complexity of a matching process is exponential by nature, but it can be heuristically reduced to a quadratic or even linear for shapes which are smooth two-manifolds. Here we model the shapes using both local and global structures, employ these to construct a quadratic dissimilarity measure, and provide a hierarchical framework for minimizing it to obtain sparse set of corresponding points. These correspondences may serve as an initialization for dense linear correspondence search.

AMS subject classifications: 65M50, 65M60, 68T45

Key words: Shape correspondence, Laplace-Beltrami, diffusion geometry, local signatures.

1. Introduction

Knowing correspondence between shapes is required for various applications, such as shape retrieval, registration, deformation, shape morphing, symmetry, self-similarity detection, etc. Detecting accurate correspondence between non-rigid shapes is a hard problem, since in general it cannot be parameterized by a finite number of unknowns. It can be cast as an assignment problem, and as such is NP-hard. A common approach for detecting correspondence between shapes differing by a certain class of transformations consists of employing shape properties that remain invariant under these transformations. These invariant surface properties are used to formulate a measure of dissimilarity between the shapes. By minimizing it one finds the correct matching. Here we use a matching scheme based on local and global surface properties, namely, local surface descriptors and global metric structures. The proposed method is demonstrated with two different types of

*Corresponding author. *Email addresses:* darav@cs.technion.ac.il (D. Raviv),
nastyad@tx.technion.ac.il (A. Dubrovina), ron@cs.technion.ac.il (R. Kimmel)

metrics - geodesic and diffusion, and different surface descriptors that include histograms of geodesic and diffusion distances, heat kernel signatures [39], and related descriptors based on the Laplace-Beltrami operator [12].

The main issue addressed in this paper is the complexity of the matching. Direct comparison of their pointwise surface descriptors and metric structures of two shape given by sampled surfaces is combinatorial in nature (the metric comparison problem was addressed in [25]). The main contribution of this work is a multi-resolution matching algorithm that can handle large number of points, and is able to produce correspondence consistent in terms of both pointwise and pairwise surface properties. According to the proposed scheme, at the lowest resolution the algorithm solves the correspondence problem exactly. The correspondence information is then propagated to a higher resolution, and refined using small neighborhoods of the matched points - thus we effectively reduce the size of the matching problem. The algorithm iterates between correspondence propagation and refinement, until a desired number of matches is found.

The rest of the paper is organized as follows: a brief review of the previous work is presented in the next section. Section 3 presents the correspondence problem formulation, followed by Section 4 where we present the hierarchical framework. In Section 5 we elaborate on distances and descriptors, and in Section 6 we discuss the numerical aspects. Section 7 contains the matching results, and comparison to the state-of-art algorithms, followed by Section 8 that concludes the paper.

2. Previous work

A wide range of methods were suggested in recent years for matching non-rigid shapes. Zigelman *et al.* [44], and Elad and Kimmel [15] suggested a method for matching isometric shapes by embedding them into a Euclidian space using multidimensional scaling (MDS), thus obtaining isometry invariant representations, followed by rigid shape matching in that space. Bronstein *et al.* [5] showed that for some surfaces, such as faces, a spherical domain better captures intrinsic properties. Mapping onto a sphere was also used for cortex alignment in medical imaging, for instance by Glaunès *et al.* [16], Tosun *et al.* [40] and Durrleman *et al.* [13]. Since it is generally impossible to embed a non-flat 2D manifold into a flat Euclidean domain without introducing some errors, the inherited embedding error affects the matching accuracy of all methods of this type. In order to eliminate this embedding error, Memoli and Sapiro [25] introduced the Gromov-Hausdorff distance [8] into the shape matching arena. Soon after, Bronstein *et al.* [6] formulized the Gromov-Hausdorff distance as a solution of a continuous optimization problem, which they called the generalized multi-dimensional scaling (GMDS). It performs a direct embedding between two non-rigid shapes, which does not suffer from the unbounded distortion of an intermediate ambient space. Lipman and Funkhouser [23] used the fact that isometric transformation between two shapes is equivalent to a Möbius transformation between their conformal mappings, and obtained this transformation by comparing the respective conformal factors. However, there is no guarantee that this result minimizes the cumulative difference between geodesic distances measure between pairs of matched points.

A different approach was suggested by Berard *et al.* [2], who considered embedding of Riemannian manifolds into an infinite dimensional Euclidean space defined by their heat kernels. Coifman and Lafon [9] introduced *diffusion maps*, also related to diffusion processes on manifolds. In the past several years we witness a dramatic prosperity in this field [19,24,35]. The Global Point Signature (GPS) suggested by Rustamov [35] for shape comparison employs the discrete Laplace-Beltrami operator, which, at least theoretically, captures the shape's geometry more faithfully than graph Laplacian. It was followed by the works of Sun *et al.* [39], and Ovsjanikov *et al.* [28], that suggested surface descriptors named Heat Kernel Signature (HKS) and Heat Kernel Maps (HKM), respectively, and used them for correspondence search. These and other methods, such as [43], employ linear matching problem formulations, which produces good matching results for near-isometric shapes. Once distortions start to appear, the descriptors become less discriminative and the matching accuracy deteriorates, thus alternative matching problem formulations should be thought of. On the other hand, the GMDS algorithm [6], uses a quadratic matching problem formulation, and results in a non-convex optimization problem. Therefore it requires good initializations in order to obtain meaningful solutions. It can be used as a refinement step for other shape matching algorithms.

Matching shapes by combining the surface descriptor information and geodesic distance in a quadratic optimization formulation of the matching problem was proposed by Hu and Hua [18], who used the Laplace-Beltrami operator for matching using prominent features, and by Dubrovina and Kimmel [12], who suggested employing surface descriptors based on the Laplace-Beltrami operator eigendecomposition together with geodesic distances. The above methods, incorporating pairwise constraints, tend to be slow due to high computational complexity, and efficient algorithms are required to match. Bronstein *et al.* [6] solved a relaxed problem in continuous domain. Wang *et al.* [42] used a graph labeling scheme at different scales, and thus were able to obtain large number of correspondences. The framework we discuss in this paper solves the quadratic matching problem posed by Dubrovina and Kimmel [12], by applying to it a hierarchical matching scheme, motivated by Wang *et al.* [42], that speeds up the optimization and allows detection of significantly higher number of correspondences.

Other important instances of non-rigid matching are self-similarity and symmetry detection. Instead of detecting the non-rigid mapping between two shapes, [29,30,32] search for a mapping from the shape to itself, and thus are able to detect intrinsic symmetries. The problem of non-rigid shape alignment also occurs in medical imaging. Work in that area was performed by Miller *et al.* [26], Leow *et al.* [22], Lai *et al.* [21], and Shi *et al.* [36], to mention a few. Specifically, Lai *et al.* [21] showed how the previously mentioned diffusion geometry-based signatures can be used for analysis of cortical surfaces.

3. Problem formulation

The problem formulation used in this work is based on comparison of pointwise and pairwise surface properties that remain approximately invariant under non-rigid ϵ -isometric transformations. Given a shape X , represented by a sampled surface, we assume that it

is endowed with a metric $d_X : X \times X \rightarrow \mathbb{R}_+ \cup \{0\}$, and a set of pointwise d -dimensional descriptors denoted $f_X : X \rightarrow \mathbb{R}^d$.

Given two shapes X and Y , endowed with metrics d_X, d_Y and descriptors f_X, f_Y , the correspondence between the shapes should preserve these properties. Formally, let us denote the correspondence between X and Y by a mapping $\mathcal{C} : X \times Y \rightarrow \{0, 1\}$ such that

$$\mathcal{C}(x, y) = \begin{cases} 1, & x \in X \text{ corresponds to } y \in Y, \\ 0, & \text{otherwise.} \end{cases} \quad (3.1)$$

In order to measure how well the correspondence \mathcal{C} preserves the above geometric properties of the shapes we use the following dissimilarity function

$$\text{dis}(\mathcal{C}) = \text{dis}_{lin}(\mathcal{C}) + \lambda \cdot \text{dis}_{quad}(\mathcal{C}). \quad (3.2)$$

The first term, $\text{dis}_{lin}(\mathcal{C})$, measures the dissimilarity between the pointwise descriptors of the two shapes

$$\text{dis}_{lin}(\mathcal{C}) = \sum_{x \in X, y \in Y} d_F(f_X(x), f_Y(y)) \mathcal{C}(x, y), \quad (3.3)$$

where d_F is some metric in the descriptor space. Here, $\text{dis}_{lin}(\mathcal{C})$ is a linear function of the correspondence \mathcal{C} . The second term, $\text{dis}_{quad}(\mathcal{C})$, measures the dissimilarity between the metric structures of the two shapes

$$\text{dis}_{quad}(\mathcal{C}) = \sum_{\substack{x, \tilde{x} \in X \\ y, \tilde{y} \in Y}} (d_X(x, \tilde{x}) - d_Y(y, \tilde{y}))^2 \mathcal{C}(x, y) \mathcal{C}(\tilde{x}, \tilde{y}), \quad (3.4)$$

and it is a quadratic function of \mathcal{C} . The parameter $\lambda \geq 0$ (Eq. (3.2)) determines the relative weight of the linear and the quadratic terms in the total dissimilarity measure. The optimal matching, denoted here by \mathcal{C}^* , is obtained by minimizing the dissimilarity measure $\text{dis}(\mathcal{C})$. Note that by changing λ to be 0 or $\gg 1$, we can transform the problem into a pure linear matching form, or into a pure quadratic programming problem, similar to these in [6, 25].

In order to avoid a trivial solution $\mathcal{C}^*(x, y) = 0, \forall x, y$, we constraint the optimal solution. The constraints are determined by the type of the correspondence we are looking for. For example, when a bijective mapping from X to Y is required, the appropriate constraints on \mathcal{C} are

$$\sum_{x \in X} \mathcal{C}(x, y) = 1, \quad \forall y \in Y, \quad \sum_{y \in Y} \mathcal{C}(x, y) = 1, \quad \forall x \in X. \quad (3.5)$$

The resulting optimization problem can be written as

$$\min_{\mathcal{C}} \left\{ \text{dis}_{lin}(\mathcal{C}) + \lambda \cdot \text{dis}_{quad}(\mathcal{C}) \right\} \quad \text{s.t.} \quad (3.5). \quad (3.6)$$

In [12], it was shown how to formulate (3.6) as a *quadratic programming* (QP) problem with binary variables $\mathcal{C}(x, y)$. A general QP problem is an optimization problem with quadratic objective function and affine constraint functions

$$\min_z \quad \frac{1}{2} z^T E z + q^T z + r \quad \text{s.t.} \quad Gz \preceq h, \quad Az = b. \quad (3.7)$$

The above problem is called *convex* when the matrix E is positive semi-definite. The *Integer Quadratic Programming* (IQP) has similar form, with the additional constraint on the variables x : $x_i \in \{0, 1\}$ (binary variables). While convex QP has one global minimum and can be solved efficiently, IQP is an NP-Hard problem. Two common methods are used to solve QAP problems [3]. The first is a heuristic approach based on a search procedure. For example, [12] used a branch-and-bound procedure to solve the optimization problem in Eq. (3.6). This approach usually provides good results assuming the local structures are both robust and unique, and there is no intrinsic symmetry. The second approach is based on relaxation. It is a three step solution, consisting of relaxing the integer constraints, solving a continuous optimization problem and projecting the solution back into integers. As expected, this procedure is highly influenced by the initial conditions. As for complexity, the relaxed IQP problem remains NP-Hard. While matching points using local structures alone (by setting $\lambda = 0$, for instance) is a linear problem, and thus can be solved efficiently, it can not guarantee global invariance in the presence of noise and symmetries. A better solution can be found by considering global structures. Unfortunately, solving the quadratic assignment problem for a large number of variables is almost infeasible, even after relaxation. In the next section we introduce a hierarchical approach for calculating an approximate solution of the above optimization problem. We use branch-and-bound for initial alignment, and then refine it using a continuous optimization technique.

4. Hierarchical formulation

Solving (3.6) reveals the main drawback of the quadratic problem formulation. As noted in [12], the dimensionality of the problem allows us to handle up to several dozens of points. Let us assume that X and Y have N and M vertices, respectively. The number of possible correspondences between X and Y is therefore NM , and thus, the dimension of the matrix E in the quadratic problem (3.7) is $NM \times NM$. Even for a small number of points, e.g. 30, the problem becomes almost infeasible.

Since the problem is not strictly combinatorial by nature, but it is derived from a smooth geometric measure, the latter can be exploited to reduce the complexity of the problem. We suggest using in the following iterative scheme. At the first step we follow [12] and solve (3.6) using a branch-and-bound procedure [1], to obtain a small number of correspondences. Each point $x \in X$ is now matched to a point $c(x) \in Y$ by the mapping c . We denote $y = c(x)$ if $\mathcal{C}(x, y) = 1$. In each iteration we search for the best correspondence between x and $c(x)$ neighborhood, instead of all points $y \in Y$, in a manner similar to [42]. Between iterations we add points $x \in X$ and $y \in Y$ using the 2-optimal *Farthest Point Sampling* (FPS) strategy [17], evaluate the neighborhood in Y of the new points, reevaluate the neighborhood of the old points, and continue until convergence. As already mentioned, the parameter λ controls the relation between the linear and quadratic parts of the dissimilarity function. In our implementation, to obtain the initial small number of corresponding points we used a small λ (0.1), giving more weight to the descriptor dissimilarity. In the consecutive iterations, we increased the value of λ , thus giving more weight to the quadratic part of the dissimilarity measure, based on distances.

Let us further analyze the complexity of the proposed method. We consider the first step to be $\mathcal{O}(N + M)$ as we use a constant number of points (usually around 20) from each mesh, sampled using the farthest point sampling in a linear time. At every consecutive iteration, we match N_j points sampled from X and KN_j points sampled from Y (K is constant), such that N_j is linear in the number of iterations. Thus, the size of the problem we solve is KN_j^2 (as the size of the distance dissimilarity matrix in the quadratic term). To solve it and obtain the N_j correspondences one can use the branch-and-bound algorithm from the first stage. Note that, in contrast to [12] who solved one problem of size $\mathcal{O}(N^2M^2)$, in the proposed formulation we solve a set of problems of size $\mathcal{O}(K^2)$ to $\mathcal{O}(N^2K^2)$ where $K \ll M$, thus effectively reducing the complexity of the problem. A different approach to solving the quadratic optimization problem (3.6) in the j 'th iteration is to solve its relaxed version. That is, instead of searching for a binary solution in $\{0, 1\}$, use a continuous optimization technique (for instance, the quasi-Newton method) to find a solution in range $[0, 1]$, and project it back to $\{0, 1\}$, using thresholding, for instance. The overall complexity of the problem: in the j 'th iteration the number of unknowns are KN_j which leads to $\mathcal{O}(K^2N_j^2)$ complexity, and for the entire iterative procedure the complexity is $\mathcal{O}(\sum_j N_j^2 K^2) = \mathcal{O}(N^3 K^2)$. In Fig. 1 we show a diagram of the process, and summarize the framework in Algorithm 4.1.

Algorithm 4.1: Hierarchical matching algorithm

-
- 1: Find a coarse matching for a small number of points on X and Y , using branch-and-bound algorithm.
 - 2: Add pairs of potential correspondences: add points from X using the farthest point sampling, and calculate their corresponding points on Y using the previously obtained coarse matching.
 - 3: For every potential match $(x_i, y_i), x_i \in X, y_i \in Y$, also use K nearest neighbors of y_i in Y as possible matches for x_i .
 - 4: Formulate the quadratic optimization problem (3.6) in terms of the potential correspondences obtained above (N_j points in X and KN_j points in Y).
 - 5: Solve the quadratic problem using branch-and-bound, or solve its relaxed version using quasi-Newton method and project the solution back into integers.
 - 6: Adjust λ if needed.
 - 7: Stop when obtain the desired number of correspondences; otherwise, return to step 2.
-

5. Distances and descriptors

5.1. Choice of metric

Differential geometry: Smooth surfaces, also known as *Riemannian manifolds*, are *differential manifolds* equipped with an inner product in the tangent space, which provides geometric notions such as angles, lengths, areas and curvatures without resorting to the ambient space, and are referred to as *intrinsic* measures. We further assume that X is embedded into $\mathbb{E} = \mathbb{R}^3$ by means of a regular map $\mathbf{x} : U \subseteq \mathbb{R}^2 \rightarrow \mathbb{R}^3$, so that the metric tensor




Figure 1: In the first step (left) we construct a quadratic correspondence matrix from all points in X into all points in Y . In each iteration (right) we search for possible matches between points in X from the previous iteration (blue circle) and new sampled points in X (green Xs) and their corresponding neighborhoods (black circles) in Y .

can be expressed in coordinates as

$$g_{ij} = \left\langle \frac{\partial \mathbf{x}}{\partial u_i}, \frac{\partial \mathbf{x}}{\partial u_j} \right\rangle, \quad (5.1)$$

where u_i are the coordinates of U , which yields the infinitesimal displacement dp

$$dp^2 = g_{11}du_1^2 + 2g_{12}du_1du_2 + g_{22}du_2^2. \quad (5.2)$$

The simplest example of an intrinsic metric is the *geodesic metric*, defined by the length of the shortest path on the surface of a shape,

$$d_X(x, x') = \inf_{\gamma \in \Gamma(x, x')} \ell(\gamma), \quad (5.3)$$

where $\Gamma(x, x')$ is the set of all admissible paths between the points x and x' on the surface X , and $\ell(\gamma)$ is the length of the path γ . There exist several numerical methods to evaluate (5.3) [4, 20, 37]. We use *fast marching method*, that simulates a wavefront propagation on a triangular mesh, associating the time of arrival of the front with the distance it traveled. While a naive implementation of one-to-many has the complexity of $\mathcal{O}(N \log N)$, it was shown that linear complexity can be achieved for the same accuracy [41]. On parametric surfaces, the fast marching can be carried out by means of a raster scan and efficiently parallelized, which makes it especially attractive for GPU-based computation [4, 38].

Diffusion geometry: Heat diffusion on the surface X is described by the heat equation,

$$\left(\Delta_g + \frac{\partial}{\partial t} \right) u(t, x) = 0, \quad (5.4)$$

where a scalar field $u : X \times [0, \infty) \rightarrow \mathbb{R}$ is the heat profile at location x and time t , and Δ_g is the *Laplace-Beltrami Operator* given a metric g .

Laplace Beltrami Operator (LBO), named after Eugenio Beltrami, is the generalization of the Laplace operator. It is a linear operator, defined as the divergence of the gradient of a function on a manifold

$$\Delta_g f = \text{div}_g \text{grad}_g f, \quad (5.5)$$

where $f : X \rightarrow \mathbb{R}$ is a scalar function.

A well known formulation of (5.5) in local coordinates u is

$$\Delta_g f = \frac{1}{\sqrt{|g|}} \frac{\partial}{\partial u^\alpha} \left(\sqrt{|g|} g^{\alpha\beta} \frac{\partial}{\partial u^\beta} f \right), \quad (5.6)$$

where $X(u^1, u^2, \dots, u^n) = (X^1, X^2, \dots, X^n)$ is an n dimensional manifold.

Since our focus will be on two dimensional affine invariants, we constrain ourself to two dimensions

$$X(u^1, u^2) = (x(u^1, u^2), y(u^1, u^2), z(u^1, u^2)). \quad (5.7)$$

For compact manifolds, the Laplace-Beltrami Operator has a discrete eigendecomposition of the form

$$\Delta_X \phi_i = \lambda_i \phi_i, \quad (5.8)$$

where $\lambda_0, \lambda_1, \dots$ are eigenvalues and ϕ_0, ϕ_1, \dots are the corresponding eigenfunctions, which construct the heat kernel

$$h_t(x, z) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x) \phi_i(z). \quad (5.9)$$

The diffusion distance is defined as a cross-talk between two heat kernels [7]

$$\begin{aligned} d_{X,t}^2(x, y) &= \|h_t(x, \cdot) - h_t(y, \cdot)\|_{L_2(X)}^2 = \int_X |h_t(x, z) - h_t(y, z)|^2 dz \\ &= \sum_{i=0}^{\infty} e^{-2\lambda_i t} (\phi_i(x) - \phi_i(y))^2. \end{aligned} \quad (5.10)$$

Since diffusion distances are derived from the Laplace Beltrami operator, they are also intrinsic properties, and, according to [2, 10, 11], also fulfill the metric axioms.

5.2. Choice of descriptors

Distance histograms: Given two surfaces X and Y and their metrics d_X and d_Y respectively, we can evaluate the distances between any two points on each one of the shapes using either choices of metrics. For isometries, a good candidate that matches point $x \in X$ to $y \in Y$ will have similar distances to all other corresponding points. Assuming the surface is well sampled, the distance histograms of corresponding points $x \in X$ and $y \in Y$ have to be similar. Comparison of histograms is a well studied operation. While straight forward bin-to-bin comparison may work, we refer the reader to more robust algorithms such as the *earth moving distance* (EMD) [34].

Heat kernel signatures: Another local descriptor based on the *heat equation*, was presented by Sun *et al.* [39]. They proposed using the diagonal of the heat kernel $k_t(x, x)$

(5.9) at multiple scales as a local descriptor. It was employed by Bronstein *et al.* [27] for shape retrieval, and was recently adapted to volumes by Raviv *et al.* [31]. The HKS remains invariant under isometric deformations of X , and it is insensitive to topological noise at small scales. It is also informative in the sense that under certain assumptions one could reconstruct the surface (up to an isometry) from it. Furthermore, the HKS descriptor can be efficiently computed from the eigenfunctions and eigenvalues of the Laplace-Beltrami operator.

Intrinsic symmetry-aware descriptors: Another possible choice for a surface descriptor is one based on the eigendecomposition of the Laplace-Beltrami operator, as suggested in [12] where the focus was on matching intrinsically symmetric non-rigid shapes. The solution proposed in [12] consists of defining distinct sets of descriptors for several possible correspondences, and minimizing the distortion $\text{dis}(\mathcal{C})$ separately for each of them, to obtain distinct matchings. Thus, when using these descriptors within an hierarchical framework, we can also find more than a single matching of the two shapes, while obtaining denser correspondence.

6. Numerical considerations

As we are interested in the eigendecomposition of the Laplace Beltrami Operator, we followed [33], and used the FEM numerical scheme [14]. For that purpose, we consider the weak form of $\Delta\phi = \lambda\phi$

$$\int \psi_k \Delta\phi \, da = \lambda \int \psi_k \phi \, da \tag{6.1}$$

with respect to some basis $\{\psi_k\}$ functions. Specifically, we choose the ψ_k 's to be the first-order finite element function obtaining the value of one at a vertex k and decaying linearly to zero in its 1-ring. Substituting these functions into (6.1), we obtain

$$\begin{aligned} \int \psi_k \Delta\phi \, da &= \int \langle \nabla\psi_k, \nabla\phi \rangle_x \, da \\ &= \int g^{ij}(\partial_i\phi)(\partial_j\psi_k) \, da = \lambda \int \psi_k \phi \, da, \end{aligned} \tag{6.2}$$

for vanishing boundary conditions. Next, we approximate the eigenfunction ϕ in the finite element basis by $\phi = \sum_{l=1} \alpha_l \psi_l$, which yields

$$\int g^{ij}(\partial_i \sum_l \alpha_l \psi_l)(\partial_j \psi_k) \, da = \lambda \int \psi_k \sum_l \alpha_l \psi_l \, da,$$

or, equivalently,

$$\sum_l \alpha_l \int g^{ij}(\partial_i \psi_l)(\partial_j \psi_k) \, da = \lambda \sum_l \alpha_l \int \psi_k \psi_l \, da.$$

The last equation can be rewritten in matrix form as a generalized eigendecomposition problem $\mathbf{A}\alpha = \lambda\mathbf{B}\alpha$ solved for the coefficients α_l , where

$$a_{kl} = \int g^{ij}(\partial_i\psi_l)(\partial_j\psi_k) da, \quad b_{kl} = \int \psi_k\psi_l da,$$

and the local surface area is expressed in parametrization coordinates as $da = \sqrt{g}du_1du_2$.


The eigenfunctions and the eigenvalues of the discretized Laplace-Beltrami operator obtained by solving the generalized eigenvalue problem are used to calculate the diffusion distances (see Eq. (5.10)), and the surface descriptors – HKS, histograms of diffusion distances, and the intrinsic symmetry-aware descriptors.

7. Results


In this section we provide several matching results obtained using our hierarchical procedure. Figs. 2 and 3 show the matching obtained with the proposed framework, combined with different descriptors and metrics, at several hierarchies. The correspondences are shown by Voronoi cells of the matched points, with corresponding patches having the same color. The matching was performed using 10 points at the coarse scale, and 30-64 points at the finest scale. Figs. 2(a), 2(b) show the result of matching shapes using geodesic distance histogram descriptors and geodesic distance metric. Fig. 2(c) shows the matching result obtained using diffusion distances instead of geodesic ones, and Fig. 2(d) – the result obtained using Heat Kernel Signatures [39] and diffusion distances. Note that the last two matchings are in fact reflected ones (follows from the intrinsically symmetric shape matching ambiguity described in [12]). Fig. 3 shows the result of matching shapes using the Laplace-Beltrami operator-based descriptors [12] and geodesic distances. There we were able to obtain the both possible correspondences between two cat shapes – the true correspondence and the reflected one. As can be seen, all setups provide good results, and we can conclude that the proposed hierarchical framework is independent of the choice of descriptors.

We compared the hierarchical method to [12]’s quadratic matching and [6]’s GMDS framework. Both are based on global structures. Since we followed [12] formulation as our first step, our initial matchings are the same. But, since the complexity of [12] rises rapidly, it can not be used to match more than a few dozen points. In addition, even for a low number of points we have a major quality advantage over [12], since the matched points on the second mesh can move, and are not restricted to the initial sampling. In Fig. 4 we see that the ear and the nose of the cat were matched using 10 points, and relocated after several iterations.


In the following test we compared the integer quadratic matching and the proposed hierarchical matching calculation times. We used the two pairs of shapes shown in Figs. 3(a) and (b), each consisting of 3400 vertices. The hierarchical scheme was tested with geodesic distances and the symmetry-aware descriptors of [12]. We used pre-computed distances and descriptors, as we were primarily interested in the matching calculation times. Evaluation of different metric and descriptor calculation times is beyond the scope of this work.




(a) Geodesic distance histogram descriptors and geodesic distance metric



(b) Geodesic distance histogram descriptors and geodesic distance metric



(c) Diffusion distance histogram and diffusion distance metric



(d) Heat Kernel Signatures and diffusion distance metric

Figure 2: Matching results obtained with the proposed framework combined with different descriptors and metrics, at several hierarchies. The hierarchical framework works well with all setups, and it performs equally well with all types of descriptors.

For the hierarchical matching, we used both branch-and-bound and quasi-Newton methods, in order to obtain the correspondence at each hierarchy level. The results are shown in Figs. 5 and 7, for different numbers of matched points. All the calculations were performed on desktop computer with Intel Core i7 2.93 GHz CPU and 8 GB RAM. The descriptors and distance calculation times were incorporated in the total calculation times.




Figure 3: Matching results obtained with the proposed framework combined with the intrinsic symmetry-aware surface descriptors and geodesic distance metric. (a), (b) The upper row – same orientation correspondence; the lower row – the reflected one.




Figure 4: Using geodesic distances as a global structure, and geodesic based histograms as a local one, the wrong ear-to-nose match gets closer to the correct one during subsequent iterations.

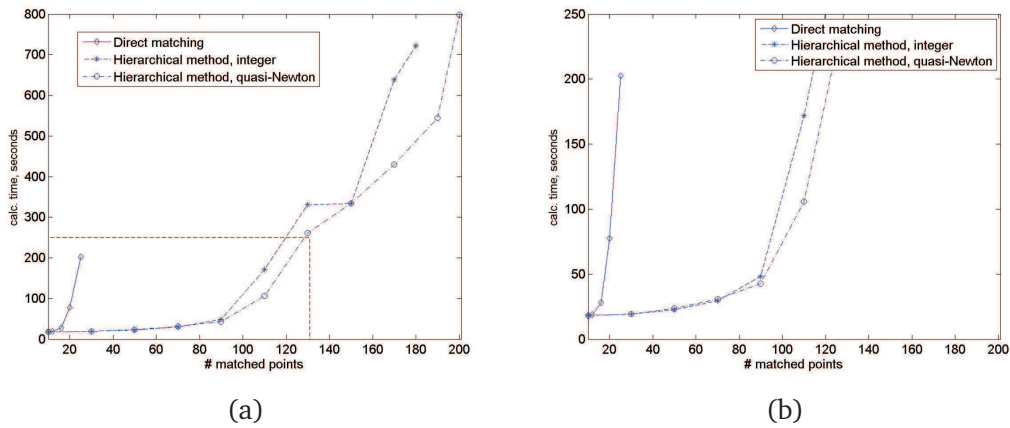


Figure 5: The total matching calculation time as a function of the number of matched points, for the integer quadratic matching of [12], and the proposed hierarchical scheme, used with geodesic distances and symmetry-aware descriptors [12], and tested on the cat shapes given in Fig. 3(a). The distance and descriptor calculation time was 17.33 seconds. (a) Up to 200 matched points. (b) Zoom into the previous graph, up to 120 matched points.

In Fig. 5(a) we show the calculation times for up to 210 matches. In Fig. 5(b) we show a zoom on the previous graph, for better comparison of the proposed method and the direct matching of [12]. We see that the proposed hierarchical method, with both branch-and-bound and quasi-Newton solvers, was able to obtain approximately 120 correspondences, while the integer quadratic matching calculated only 25 correspondences during the same amount of time. Also, for higher number of correspondences we can no longer use the branch-and-bound optimization within the proposed hierarchical framework, because of its high complexity, while using the proposed relaxation with quasi-Newton minimization they can still be processed in a reasonable time.

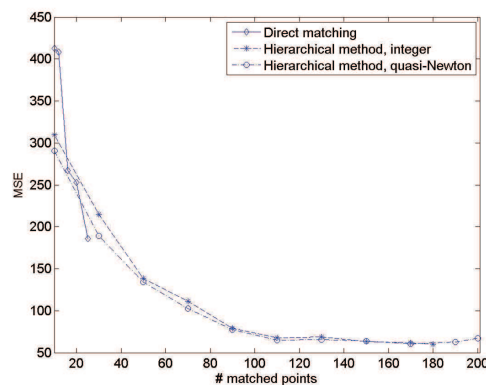


Figure 6: Mean square geodesic distance error, calculated between the correspondences obtained with the integer quadratic matching and the proposed method, and the ground truth correspondences; corresponds to the calculation times given in Fig. 5.

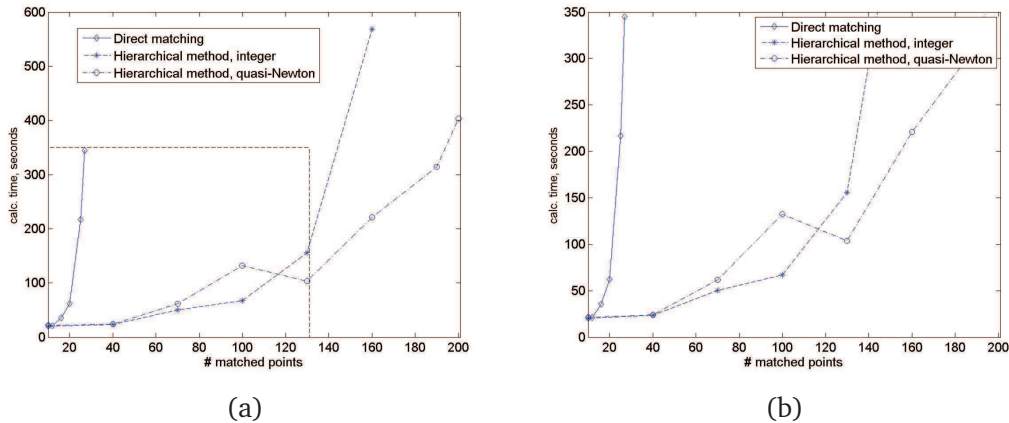


Figure 7: The total matching calculation time as a function of the number of matched points, for the integer quadratic matching of [12], and the proposed hierarchical scheme, used with geodesic distances and symmetry-aware descriptors [12], and tested on the horse shapes given in Fig. 3(b). The distance and descriptor calculation time was 19.10 seconds. (a) Up to 200 matched points. (b) Zoom into the previous graph, up to 120 matched points.

In order to evaluate the accuracy of the matching obtained using each one of the methods above, we calculated the mean square geodesic distance error (MSE) between the positions of the detected matchings and the ground-truth correspondences. The results are shown in Fig. 6. We see that adding correspondences reduces the matching error, and that the matching problem relaxation has little or no effect on the matching accuracy, compared to the solution obtained using branch-and-bound. Fig. 7 shows similar results obtained for a pair of the horse shapes from Fig. 3(b).

Bronstein *et al.* [6] proposed to minimize the Gromov-Hausdorff distance between shapes, which in theory provides the best correspondence between approximate isometries. Since their framework is based on non-convex optimization, the first alignment is critical. We evaluated GMDS results using its own initializer and our quadratic first step, which provided better results. We repeated the experiments shown in Fig. 2(a) and measured the geodesic distances between the corresponding points versus the ground truth correspondence. We improved the L_∞ error by 26% and the mean error by 6.25%. It is not surprising, since usually the best correspondence can not be originated from a global structure alone. One can think, for example, on a trivial experiment where only the head rotates. The best correspondence will suffer a distortion in the neck alone, but GMDS will suffer from a distortion in all points.

8. Conclusions

We presented a hierarchical framework, based on quadratic programming, to find matching between non-rigid shapes. While the problem is NP-Hard in general, taking into account the smooth structure of our shapes and using an iterative scheme allows us to find large number of corresponding points in reasonable time. In order to evaluate the pro-

posed scheme we combined it with different metrics and descriptors, tested it on various non-rigid shapes, and compared the results to the state of art methods.

Acknowledgments This research was supported by European Community's FP7-ERC program, grant agreement no. 267414.

References

- [1] A. Bemporad. Hybrid Toolbox - User's Guide, 2004.
- [2] P. Bérard, G. Besson, and S. Gallot. Embedding riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, 4(4):373–398, 1994.
- [3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2006.
- [4] A. M. Bronstein, M. M. Bronstein, Y. S. Devir, R. Kimmel, and O. Weber. Parallel algorithms for approximation of distance maps on parametric surfaces. In *Proc. ACM Transactions on Graphics (SIGGRAPH)*, volume 27, 2008.
- [5] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Expression-invariant face recognition via spherical embedding. In *Proc. Int'l Conf. Image Processing (ICIP)*, volume 3, pages 756–759, 2005.
- [6] 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.
- [7] M. M. Bronstein and A. M. Bronstein. Shape recognition with spectral distances. *Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2010.
- [8] D. Burago, Y. Burago, S. Ivanov, and American Mathematical Society. *A course in metric geometry*. American Mathematical Society Providence, 2001.
- [9] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21:5–30, July 2006.
- [10] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. Definition of diffusion distance.
- [11] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS*, 102(21):7426–7431, 2005.
- [12] A. Dubrovina and R. Kimmel. Matching shapes by eigendecomposition of the laplace_beltrami operator. In *Proc. Symposium on 3D Data Processing Visualization and Transmission (3DPVT)*, 2010.
- [13] S. Durrleman, X. Pennec, A. Trounev, and N. Ayache. Measuring brain variability via sulcal lines registration: a diffeomorphic approach. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 675–682, 2007.
- [14] G. Dziuk. Finite elements for the Beltrami operator on arbitrary surfaces. In S. Hildebrandt and R. Leis, editors, *Partial differential equations and calculus of variations*, pages 142–155. 1988.
- [15] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 25(10):1285–1295, 2003.
- [16] J. Glaunès, M. Vaillant, and M. I. Miller. Landmark matching via large deformation diffeomorphisms on the sphere. *Journal of Mathematical Imaging and Vision*, 20:179–200, 2004.
- [17] D.S. Hochbaum and D.B. Shmoys. A best possible heuristic for the k -center problem. *Mathematics of Operations Research*, pages 180–184, 1985.
- [18] J. Hu and J. Hua. Salient spectral geometric features for shape matching and retrieval. *Vis. Comput.*, 25(5-7):667–675, 2009.

- [19] V. Jain and H. Zhang. A spectral approach to shape-based retrieval of articulated 3D models. *Computer-Aided Design*, 39:398–407, 2007.
- [20] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proc. National Academy of Sciences (PNAS)*, 95(15):8431–8435, 1998.
- [21] R. Lai, Y. Shi, K. Scheibel, S. Fears, R. Woods, A. W. Toga, and T. F. Chan. Metric-induced optimal embedding for intrinsic 3d shape analysis. In *Proc. International Conference of Computer Vision (CVPR)*, 2010.
- [22] A. Leow, C. L. Yu, S. J. Lee, S. C. Huang, H. Protas, R. Nicolson, K. M. Hayashi, A. W. Toga, and P. M. Thompson. Brain structural mapping using a novel hybrid implicit/explicit framework based on the level-set method. *NeuroImage*, 24(3):910–927, 2005.
- [23] Y. Lipman and T. Funkhouser. Mobius voting for surface correspondence. In *Proc. ACM Transactions on Graphics (SIGGRAPH)*, volume 28, 2009.
- [24] D. Mateus, R. P. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [25] 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.
- [26] M. I. Miller, M. Faisal Beg, C. Ceritoglu, and C. Stark. Increasing the power of functional maps of the medial temporal lobe by using large deformation diffeomorphic metric mapping. *Proc. National Academy of Science (PNAS)*, 102(27):9685–9690, 2005.
- [27] M. Ovsjanikov, A. M. Bronstein, M. M. Bronstein, and L. J. Guibas. Shape Google: a computer vision approach to invariant shape retrieval. In *Proc. Non-Rigid Shape Analysis and Deformable Image Alignment (NORDIA)*, 2009.
- [28] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. J. Guibas. One point isometric matching with the heat kernel. *Proc. Symposium on Geometry Processing (SGP)*, 29(5):1555–1564, Jul 2010.
- [29] M. Ovsjanikov, J. Sun, and L. J. Guibas. Global intrinsic symmetries of shapes. In *Computer Graphics Forum*, volume 27, pages 1341–1348, 2008.
- [30] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Full and partial symmetries of non-rigid shapes. *International Journal of Computer Vision (IJCV)*, 2009.
- [31] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Volumetric heat kernel signatures. In *Proc. 3D Object recognition (3DOR), part of ACM Multimedia.*, 2010.
- [32] D. Raviv, A. M. Bronstein, M. M. Bronstein, R. Kimmel, and G. Sapiro. Diffusion symmetries of non-rigid shapes. In *Proc. International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)*, 2010.
- [33] M. Reuter, S. Biasotti, D. Giorgi, G. Patanè, and M. Spagnuolo. Discrete Laplace–Beltrami operators for shape analysis and segmentation. *Computers & Graphics*, 33(3):381–390, 2009.
- [34] Y. Rubner, C. Tomasi, and L.J. Guibas. The earth mover’s distance as a metric for image retrieval. *IJCV*, 40(2):99–121, 2000.
- [35] R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. Symposium on Geometry Processing (SGP)*, pages 225–233, 2007.
- [36] Y. Shi, R. Lai, R. Gill, D. Pelletier, D. Mohr, N. Sicotte, and A. W. Toga. Conformal metric optimization on surface (cmos) for deformation and mapping in laplace-beltrami embedding space. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2011.
- [37] A. Spira and R. Kimmel. An efficient solution to the eikonal equation on parametric manifolds. *Interfaces and Free Boundaries*, 6(4):315–327, 2004.
- [38] A. Spira and R. Kimmel. An efficient solution to the eikonal equation on parametric manifolds. *Interfaces and Free Boundaries*, 6(3):315–327, 2004.

- [39] J. Sun, M. Ovsjanikov, and L. J. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. Symposium on Geometry Processing (SGP)*, 2009.
- [40] D. Tosun, M. E. Rettmann, and J. L. Prince. Mapping techniques for aligning sulci across multiple brains. *Medical Image Analysis*, 8:295–309, 2004.
- [41] J. N. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Trans. Automatic Control*, 40(9):1528–1538, 1995.
- [42] C. Wang, M. M. Bronstein, and N. Paragios. Discrete minimum distortion correspondence problems for non-rigid shape matching. Research Report 7333, INRIA, 2010.
- [43] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [44] G. Zigelman, R. Kimmel, and N. Kiryati. Texture mapping using surface flattening via multi-dimensional scaling. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 9(2):198–207, 2002.