Advances in Adaptive Data Analysis
Vol. 3, Nos. 1 & 2 (2011) 203–228

DOI: 10.1142/S1793536911000829

World Scientific
www.worldscientific.com

APPROXIMATELY ISOMETRIC SHAPE CORRESPONDENCE BY MATCHING POINTWISE SPECTRAL FEATURES AND GLOBAL GEODESIC STRUCTURES

ANASTASIA DUBROVINA* and RON KIMMEL†

Technion - Israel Institute of Technology,
Haifa, 32000, Israel
**nastyad@tx.technion.ac.il*
†*ron@cs.technion.ac.il*

A practical method for finding correspondences between nonrigid isometric shapes is presented. It utilizes both pointwise surface descriptors, and metric structures defined on the shapes to perform the matching task, which is formulated as a quadratic minimization problem. Further, the paper explores the correspondence ambiguity problem arising when matching intrinsically symmetric shapes using only intrinsic surface properties. It is shown that when using isometry invariant surface descriptors based on eigendecomposition of the Laplace–Beltrami operator, it is possible to construct distinctive sets of surface descriptors for different possible correspondences. When used in a proper minimization problem, those descriptors allow us to explore number of possible correspondences between two given shapes.

Keywords: Shape correspondence; Laplace–Beltrami; surface descriptors.

1. Introduction

Correspondence detection is an important part of many three dimensional (3D) shape-processing applications, such as shape retrieval, registration–either as a stand-alone task, or as an initialization for deformation and morphing algorithms, symmetry, self-similarity detection, etc. Unlike rigid shape correspondence, where the transformation connecting the two shapes can be modeled using six parameters, detecting nonrigid shape correspondence is a far more challenging task. In this paper we assume that the two shapes we try to match are approximately isometric in terms of corresponding geodesic distances measured between matched surface points.

We suggest a framework for an unsupervised nonrigid isometric shape correspondence detection, based on matching isometry invariant surface descriptors, and the metric structures of the shapes. We formulate the matching task as a problem of minimizing a quadratic dissimilarity objective function, using the surface properties mentioned above, and solve it to find the correspondence.

The suggested framework is general, and can be used with any type of descriptors and metric. As in this paper we consider the problem of matching isometric shapes, we employ isometry invariant surface descriptors based on eigendecomposition of the Laplace–Beltrami operator [Reuter *et al.* (2006); Levy (2006); Levy and Zhang (2009)], and isometry invariant geodesic distance metric.

We also explore the correspondence ambiguity problem that arises when matching a pair of intrinsically symmetric shapes using only intrinsic surface properties. In this case, there exist more than one correspondence between the shapes. We show that when using the proposed isometry invariant surface descriptors, it is possible to find several alignments of two symmetric shapes. Specifically, we employ the proposed approach to detect correspondences between shapes belonging to the large class of shapes with one intrinsic symmetry, such as human and animal shapes. To the best of our knowledge, this is the first attempt to deal with this correspondence ambiguity problem.

Thus, the two main contributions of the paper can be summarized as follows.

- We suggest a general framework for matching nonrigid shapes that utilizes pointwise surface descriptors and metric information. The matching is cast to a quadratic optimization problem, that can be minimized to find the optimal correspondence. We show empirical evidence that the proposed approach, which combines the two surface properties mentioned above, is superior to using each of them separately.
- We describe the correspondence ambiguity problem that arises when matching intrinsically symmetric shapes, and suggest a method for dealing with that ambiguity. We propose to use surface descriptors based on eigendecomposition of the Laplace–Beltrami operator, and show how they can be used to find several possible correspondences between the shapes.

The work presented in this paper expands the concepts earlier presented in [Dubrovina and Kimmel (2010)].

The rest of the paper is organized as follows: we start with a review of related efforts in the field of nonrigid shape matching, presented in Sec. 2. In Sec. 3 we present the correspondence detection problem, and in Sec. 4 we show how it can be formulated as an optimization problem. In Sec. 5 we describe isometry invariant surface descriptors, and an algorithm for constructing distinct sets of descriptors corresponding to each possible alignment. In Sec. 6 we discuss the complexity of the proposed algorithm, present the matching results obtained with it, and show empirical evidence that minimizing the combined measure provides better results than using each of its parts separately. Section 7 concludes the paper.

2. Previous Work

Feature-based surface matching algorithms are an active research area. Zhang *et al.* [2008] found features by examining extremities of the geodesic distance field defined

on the mesh. Zaharescu *et al.* [2009] proposed an extension of an existing method for salient feature detection in 2D images. Hu and Hua [2009] and Sun *et al.* [2009] used the eigenfunctions of the Laplace–Beltrami operator and the heat kernel operator, respectively, to find the feature points. d'Amico *et al.* [2006] suggested a theoretical analysis of descriptor-based matching (denoted there as *measuring functions*), and applied it for planar shape comparison.

A different approach to nonrigid shape matching is based on embedding the shapes into some canonical domain, where the matching can be easily performed. Elad and Kimmel [2003] embedded the shapes into a (flat) Euclidean domain using the multidimensional scaling (MDS) method [Borg and Groenen (2005)], and performed rigid matching. Different embedding domains spanned by eigenfunctions of either affinity matrix or a graph Laplacian operator defined on triangulated shapes were suggested by Jain *et al.* [2007], Mateus *et al.* [2008], and Knossow *et al.* [2009]. Rustamov [2007] suggested using eigendecomposition of the Laplace–Beltrami operator, which is more consistent with the shape's intrinsic geometry. Recently, Lipman and Funkhouser [2009] suggested conformal surface flattening to a complex plane, and matched the shapes based on their corresponding conformal factors, thereby simplifying the set of nonrigid isometric deformations to the Möbius group.

Memoli and Sapiro [2005], and Memoli [2007] suggested using the Gromov–Hausdorf distance [Gromov (1981)] to compare the metric structures of the shapes. Bronstein *et al.* [2006] formulated the metric comparison as a problem of embedding one shape into another with minimal geodesic distance dissimilarity by introducing the generalized MDS (GMDS). Thorstensen and Keriven [2009] extended the GMDS to textured shapes. Tevs *et al.* [2009] suggested randomized geodesic distance preserving matching algorithm. Anguelov *et al.* [2004] proposed matching shapes by minimizing a probabilistic model based on geodesic distances between all pairs of corresponding points. Leordeanu and Hebert [2005] employed both local descriptors and global pairwise similarity for the matching. Chang and Zwicker [2008], and Huang *et al.* [2008] approximated nonrigid transformations that align the two shapes by a finite set of rigid transformations, which were claimed to be simple to calculate. Self-similarity detection is a particular case of the shape-matching problem. Algorithms for detecting intrinsic symmetries of nonrigid isometric shapes were proposed in [Raviv *et al.* (2007); Ovsjanikov *et al.* (2008); Raviv *et al.* (2009); Lipman *et al* (2010); Kim *et al* (2010)].

Many of the above algorithms share a common denominator, namely the correspondence detection is performed by minimization of a certain shape dissimilarity measure. The definition of the dissimilarity is usually based on surfaces properties that remain approximately invariant under possible transformations of the shapes. Roughly speaking, descriptor-based methods [Zaharescu *et al.* (2009); Sun *et al.* (2009)] measure the dissimilarity between some local signatures (or, descriptors) associated with the shapes, while metric based approaches

[Memoli and Sapiro (2005); Bronstein *et al.* (2006); Memoli (2007); Bronstein *et al.* (2009); Tevs *et al.* (2009)] find correspondences by minimizing the difference between the metric structures of the two shapes. There is also a family of methods that measure dissimilarity using a mixture of several common quantities, e.g. [Leordeanu and Hebert (2005); Hu and Hua (2009); Thorstensen and Keriven (2009)]. The approach we present in this paper belongs to the latter group. That is, we try to match both pointwise surface descriptors, and the metric structures of the shapes, and define a dissimilarity measure accordingly.

3. Problem Formulation

Let us denote by X and Y the two shapes we would like to match. We represent the correspondence between X and Y by a bijective mapping $\varphi : X \to Y$, such that for each point $x \in X$ its corresponding point is $\varphi(x) \in Y$. We seek for correspondence that preserves both pointwise surface properties, and global pairwise relationships between corresponding points — those that remain approximately invariant under a given class of transformations. In order to measure the pointwise dissimilarity between X and Y we associate with each point $x \in X$ a surface descriptor $f^X(x)$, and, correspondingly, with each point $y \in Y$ — a descriptor $f^Y(y)$. The pointwise dissimilarity measure is defined as a sum of distances between the descriptors of all pairs of corresponding points

$$\mathrm{Dis}_l(\varphi) = \sum_{x \in X} d_F\left(f^X(x), f^Y(\varphi(x))\right). \tag{1}$$

The distance measure d_F is defined in the descriptor space. It is chosen according to the type of the descriptors $f^X(x)$, $f^Y(y)$.

In order to compare the pairwise relationships between the corresponding points we adopt the metric space shape representation approach, along the line of [Elad and Kimmel (2003); Memoli and Sapiro (2005); Bronstein *et al.* (2006); Bronstein *et al.* (2008)]. According to it, the shape is represented by a pair (X, d_X), where X is a smooth compact connected Riemannian manifold, with associated distance measure $d_X : X \times X \to \mathbb{R}_+ \bigcup \{0\}$.

Next, following [Memoli and Sapiro (2005); Bronstein *et al.* (2006), and Bronstein *et al.* (2008)], we measure the metric dissimilarity induced by the correspondence given by φ. Specifically, given two pairs of matched points $(x, \varphi(x))$ and $(\tilde{x}, \varphi(\tilde{x}))$, we can compare the corresponding distances measured on X and Y by $|d_X(x, \tilde{x}) - d_Y(\varphi(x), \varphi(\tilde{x}))|$. Thus, the overall pairwise dissimilarity induced by the correspondence set φ, is given by

$$\mathrm{Dis}_q(\varphi) = \sum_{x, \tilde{x} \in X} |d_X(x, \tilde{x}) - d_Y(\varphi(x), \varphi(\tilde{x}))|. \tag{2}$$

The overall cost of the correspondence is a combination of the dissimilarity measures (1) and (2):

$$\mathrm{Dis}(\varphi) = \sum_{x \in X} d_F(f^X(x), f^Y(\varphi(x))) + \lambda \cdot \sum_{x,\tilde{x} \in X} |d_X(x,\tilde{x}) - d_Y(\varphi(x),\varphi(\tilde{x}))|. \quad (3)$$

The scalar parameter λ determines the relative weight of the second term in the overall dissimilarity measure. The optimal correspondence φ^* would minimize $\mathrm{Dis}(\varphi)$.

The above formulation is general and can be used with any choice of descriptors and distance measures, and with different minimization techniques. Here, we would like to match surfaces that differ by approximately isometric transformations, that is nonrigid transformations that do not stretch or tear the surfaces. Figure 1 shows several instances of a human body differing by approximately isometric transformations.

The descriptors and the metric chosen for this application have to be an isometry invariant as possible. In Sec. 5, we describe such a descriptor based on the eigendecomposition of the Laplace–Beltrami operator. It is related to the global point signature (GPS) proposed by Rustamov [2007]. Other descriptors that can be employed in the proposed framework include the Gaussian curvature of the surface, histograms of geodesic distances like those used by [Raviv *et al.* (2009)] and [Ruggeri and Saupe (2008)], Gaussian curvature of the surface, or the heat kernel based descriptors proposed in [Bronstein and Kokkinos (2010)].

For nonrigid isometric transformation that does not tear or stretch the surface, a good choice of a metric is based on geodesic distance. Given the points x_1, x_2 on X, the geodesic distance between them, denoted by $d_X(x_1, x_2)$, is equal to the length of the shortest path on the surface X connecting the points x_1 and x_2. In the presence of topological changes, a better choice for the metric would be the diffusion distance [Bérard *et al.* (1994); Coifman and Lafon (2006)], or the commute time distance [Qiu and Hancock (2007)], that are less sensitive to these kind of transformations.

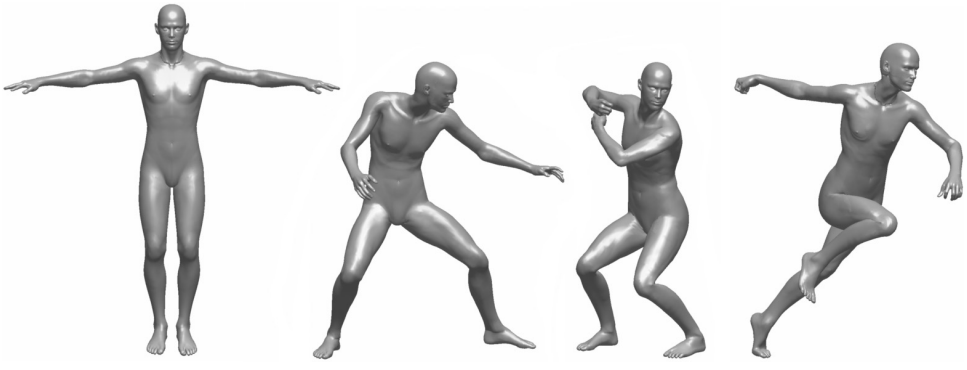

Fig. 1. Isometric transformations example.

4. Correspondence Detection as a Quadratic Optimization Problem

In order to express the correspondence detection as an optimization problem we need to redefine the correspondence between X and Y. Consider the set of all possible mappings P over the space of all pairs (x, y), where $x \in X$, and $y \in Y$, such that

$$P(x,y) = \begin{cases} 1, & x \text{ corresponds to } y, \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

The correspondence cost presented in Eq. (3) can be written equivalently in terms of P as

$$\begin{aligned} \mathrm{Dis}(P) = &\sum_{\substack{x\in X \\ y\in Y}} d_F(f^X(x), f^Y(y))P(x,y) \\ &+ \lambda \cdot \sum_{\substack{x,\tilde{x}\in X \\ y,\tilde{y}\in Y}} \|d_X(x,\tilde{x}) - d_Y(y,\tilde{y}))\|P(x,y)P(\tilde{x},\tilde{y}), \end{aligned} \tag{5}$$

where P^* that minimizes $\mathrm{Dis}(P)$ represents the optimal correspondence between the shapes X and Y.

Next, we would like to define the optimization problem for the discrete setting, i.e. when X and Y are both discretized and are given either as triangulated meshes, or clouds of points. The mapping P could be approximated as a binary matrix. The correspondence between some $x_i \in X$ and $y_j \in Y$ is given by the (i, j) entry of P, namely $P_{ij} = P(x_i, y_j) \in \{0, 1\}$.

The cost (5) can then be discretized as follows

$$\begin{aligned} \mathrm{Dis}(P) = &\sum_{i,j} d_F(f^X(x_i), f^Y(y_j))P_{ij} \\ &+ \lambda \cdot \sum_{i,j,m,n} \|d_X(x_i,x_m) - d_Y(y_j,y_n))\|P_{ij}P_{mn}. \end{aligned} \tag{6}$$

To avoid a trivial solution, we add constraints on P to Eq. (6). Those constraints are subject to the type of correspondence we look for. We briefly review several of them here. For an in-depth discussion on related constraint optimization the reader is referred to [Maciel and Costeira (2003)].

When a bijective (one-to-one and onto) correspondence is required, the constraints are given by

$$\sum_i P_{ij} = 1, \quad \sum_j P_{ij} = 1, \quad \forall i, j \tag{7}$$

In this case, the solution P is a permutation matrix.

This constraint may be too restrictive. If the shapes have different number of points, or significantly different triangulations, the optimal correspondence is not

necessarily a bijection. In this case, we may fix the points on X, and for each $x \in X$ look for a correspondence $y \in Y$ that minimizes the dissimilarity measure

$$\sum_i P_{ij} = 1, \quad \forall\, i. \tag{8}$$

Each $x \in X$ may have one or more corresponding points according to this constraint. Finally, we rewrite the dissimilarity measure (6) in matrix form, as a quadratic function of the correspondence P. We denote each double index (i, j), as in (x_i, y_j) (or ij in P_{ij}) by a single index k. The pairwise dissimilarity term is converted into a vector with entries

$$b_k = d_F(f^X(x_{i_k}), f^Y(y_{j_k})), \tag{9}$$

and the metric dissimilarity term — into a matrix with entries

$$Q_{kl} = \|d_X(x_{i_k}, x_{m_k}) - d_Y(y_{j_l}, y_{n_l}))\|. \tag{10}$$

We readily obtain the following quadratic optimization problem

$$P^* = \arg\min_P \{b^T P + \lambda \cdot P^T Q P\} \quad \text{s.t.} \quad SP = \mathbb{1}, \tag{11}$$

where the sparse matrix S represents the matrix form of the chosen constraints; $\mathbb{1}$ is a vector with all entries equal to 1, of an appropriate size. Implementation details: in our experiments we normalized the pairwise and the metric dissimilarity terms by their maximal values, in order to obtain meaningful contributions of the two terms to the cost function, and used λ in the range of $[0.1, 0.5]$.

5. Isometry Invariant Surface Descriptors

In this section, we describe a method for construction of isometry invariant surface descriptors based on eigendecomposition of the Laplace–Beltrami operator. We start with a brief review of the Laplace–Beltrami operator, and the associated eigendecomposition theory, and then describe the proposed descriptor.

5.1. *Laplace–Beltrami operator*

The Laplace–Beltrami operator is a generalization of the Laplacian operator from flat domain to compact Riemannian manifolds. Given a manifold M, its Laplace–Beltrami operator Δ_M is given by

$$\Delta_M f = -\mathrm{div}_M(\nabla_M f), \tag{12}$$

for any function $f : M \to \mathbb{R}$. The divergence and the gradient operators, div_M and ∇_M, respectively, are defined by the intrinsic geometry of the manifold M. Hence, the operator Δ_M is invariant to isometric transformations of the manifold M.

Consider the Laplace–Beltrami operator eigenvalue problem given by

$$\Delta_M \phi = \lambda \phi. \tag{13}$$

ϕ is the eigenfunction of Δ_M, corresponding to the eigenvalue λ. The spectrum of the Laplace–Beltrami operator consists of positive eigenvalues (see, e.g. [Rosenberg (1997)]). When M is a connected manifold without boundary, then Δ_M has additional eigenvalue equal to zero, with corresponding constant eigenfunction. We can order the eigenvalues as follows

$$0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \cdots \tag{14}$$

The set of corresponding eigenfunctions given by

$$\{\phi_1, \phi_2, \phi_3, \ldots\} \tag{15}$$

forms an orthonormal basis of functions defined on M (see [Rosenberg (1997)]).

5.2. *Surface descriptors*

Like the Laplace–Beltrami operator, the eigenvalues and the eigenfunctions are defined by the intrinsic geometry of the manifold, and thus remain invariant under its isometric transformations. This fact has been exploited for nonrigid isometric shape recognition [Reuter *et al.* (2006); Rustamov (2007)] and registration [Zhang *et al.* (2008); Mateus *et al.* (2008); Knossow *et al.* (2009)].

Let us consider a candidate surface descriptor constructed from the values of the eigenfunctions of Δ_M

$$f^M(q) = \{\phi_1^M(q), \phi_2^M(q), \ldots, \phi_K^M(q)\}, \quad q \in M. \tag{16}$$

$\phi_k^M(q)$ is the value of the kth eigenfunction at point $q \in M$.

Here, we choose the dimension K of the descriptor to be small. As we have already mentioned, the eigenfunction corresponding to the zero eigenvalue is a constant function. As we increase the value of the eigenvalue, the corresponding eigenfunction (or eigenfunctions) varies more rapidly over the manifold. Eigenfunctions corresponding to large eigenvalues are therefore more sensitive to the discretization. On the other hand, too small K would reduce the discriminative power of the descriptor. In our experiments we used K in the range of 5–10.

The descriptor $f^M(q)$ is defined only by the intrinsic properties of M and is thus suitable for isometry invariant matching. The descriptor $f^M(q)$ can be viewed as an embedding of the point q into a K-dimensional Euclidean space spanned by the eigenfunctions $\{\phi_1, \phi_2, \ldots, \phi_K\}$. Hence, we can measure the dissimilarity between the descriptors in this space using L_p-norm (in our experiments we used $p = 2$).

Now, given two isometric shapes X and Y, we denote the sets of eigenvalues and eigenfunctions of their Laplace–Beltrami operators by

$$\{\lambda_k^X\}_{k\geq 1}^K, \quad \{\phi_k^X\}_{k\geq 1}^K \tag{17}$$

and

$$\{\lambda_k^Y\}_{k\geq 1}^K, \quad \{\phi_k^Y\}_{k\geq 1}^K, \tag{18}$$

respectively. Despite the isometry invariancy of the Laplace–Beltrami operator, the sets $\{\phi_k^X\}_{k\geq 1}^K$ and $\{\phi_k^Y\}_{k\geq 1}^K$ are not necessarily identical (up to a perturbation defined by correspondence). There are several factors that explain this loss of identity.

(1) The eigenvalues of the Laplace–Beltrami operator may have multiplicity greater than one, with several eigenfunctions corresponding to each such eigenvalue. Each set of eigenfunctions corresponding to the eigenvalue with multiplicity greater than one spans a subspace of the space of functions defined on the manifold M, or an N-dimensional Euclidean space in case of triangulated surface with N vertices. Thus, in order to measure the distance between two such sets calculated for X and Y we need some distance measure between two subspaces, other than an L_p-norm. Additionally, since we work with shapes represented by discrete triangular meshes, the calculation of the Laplace–Beltrami operator and its eigendecomposition suffers from approximation and numerical errors. This sometimes leads to switching of eigenfunctions corresponding to eigenvalues with similar values.

(2) In addition, the eigenfunctions of the Laplace–Beltrami operator corresponding to eigenvalues without multiplicity are defined up to a sign. Therefore, the sets $\{\phi_k^X\}_{k\geq 1}$ and $\{\phi_k^Y\}_{k\geq 1}$, restricted to eigenfunctions corresponding to eigenvalues without multiplicity, are connected by some arbitrary sign sequence, that has to be estimated prior to the matching. Figure 2 presents an example of this sign ambiguity. It shows two articulations of a human body, colored according to the values of the first four eigenfunctions of their corresponding Laplace–Beltrami operators. In order to obtain coherent eigenfunctions for the two shapes, we need to multiply the first eigenfunction of the lower shape by -1.

(3) An additional problem of eigenfunction sign ambiguity arises when matching *intrinsically symmetric* shapes (for the definition of intrinsic symmetry the reader is referred to [Raviv *et al.* (2007), (2009)]), such as the two shapes of the human body in Fig. 2. It was observed by Ovsjanikov *et al.* [2008] that eigenfunctions of the Laplace–Beltrami operator of an intrinsically symmetric shape M are symmetric functions as well, with respect to the intrinsic symmetry of M (Theorem 3.1 in [Ovsjanikov *et al.* (2008)]). Eigenfunctions corresponding to eigenvalues without multiplicity exhibit reflection symmetries, whereas eigenfunctions corresponding to eigenvalues with multiplicity greater than one may also exhibit rotation symmetries. That is, for two intrinsically symmetric points $p, q \in M$, the eigenfunctions corresponding to eigenvalues without multiplicity are related by

$$\phi_k^M(p) = \phi_k^M(q) \quad \text{or} \quad \phi_k^M(p) = -\phi_k^M(q). \tag{19}$$

In case of eigenfunctions corresponding to eigenvalues with multiplicity greater than one, the connection between the eigenfunctions is more complex (see [Ovsjanikov

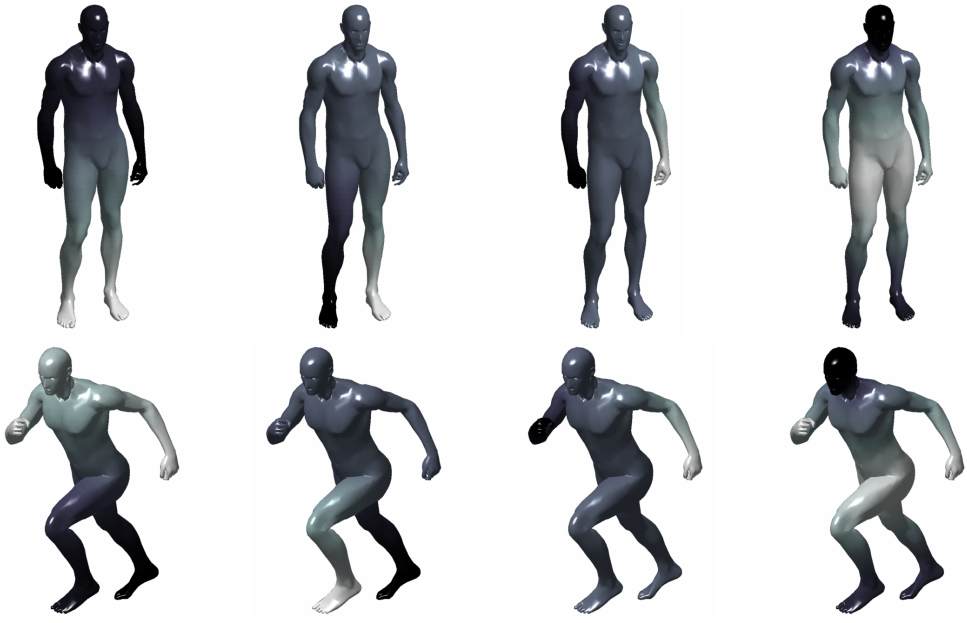

Fig. 2. Two articulations of a human shape, colored according to the values of the first five eigenfunctions of their Laplace–Beltrami operators, from left to right. The two possible sign sequence relating the two groups of the eigenfunctions are $[-,+,+,+]$ and $[-,-,-,+]$.

et al. (2008)]). In this paper, we limit the discussion only to eigenfunctions corresponding to eigenvalues without multiplicity.

We would like to note that there may exist nontrivial intrinsic symmetries that necessarily imply appearance of repeated eigenvalues, as described in Sec. 3.4 of [Ovsjanikov *et al.* (2008)]), and the above restriction to non-repeating eigenvalues can rule out these symmetries (e.g., continuous rotational symmetries discussed in [Ben-Chen *et al.* (2010)]). In particular, with the proposed method we can find only those intrinsic symmetries that result in reflection symmetries of the eigenfunctions of the Laplace–Beltrami operator, and not their rotation symmetries. This is a major limitation to the proposed method, and further research is required to determine whether it is possible to employ eigenvectors corresponding to these eigenvalues, to detect all possible symmetries and symmetrical correspondences.

We would like to display how the intrinsic symmetry of the matched shapes affects the matching problem using the example in Fig. 2. As we have already mentioned, we need to multiply the eigenfunctions of the lower shape by the sign sequence $[-,+,+,+]$ in order to align the two sets of eigenfunctions. We call the resulting alignment a *primary alignment*, or a *primary correspondence*. But, since the shapes are intrinsically symmetric, there exists another sign sequence that produces matching with the same dissimilarity value (calculated according to Eq. (1)).

It is $[-,-,-,+]$, and results in what we call a *symmetric correspondence*. Roughly speaking, it matches each point on the first shape not to its corresponding point on the second shape, but to a point intrinsically symmetric to it. When the shape has more than one intrinsic symmetry, the number of the possible *symmetric correspondences* equals the number of the symmetries. Moreover, it is impossible to distinguish between the different correspondences, since they result in equal dissimilarity value, according to Eq. (1).

From the discussion above it follows that we must preprocess the candidate descriptors $f^X(x), f^Y(y)$ before we can actually use them for matching. Two of the above problems — eigenfunction ordering and sign ambiguity, were previously addressed with respect to the spectral decomposition-based shape matching. Several authors, among them Shapiro and Brady [1992], and Jain *et al.* [2007], proposed using either exhaustive search or greedy approach for the eigenvalue ordering and sign detection. Umeyama [1988] proposed using a combination of the absolute values of the eigenfunctions and an exhaustive search. Mateus *et al.* [2007] expressed the connection between the eigenfunctions of the two shapes by an orthogonal matrix. They formulated the matching as a global optimization problem, optimizing over the space of orthogonal matrices, and solved it using the expectation minimization approach. Later, Mateus *et al.* [2008] and Knossow *et al.* [2009] suggested using histograms of eigenfunction values to detect their ordering and signs.

In this paper, we suggest an algorithm that, in the presence of intrinsic symmetries, aims to estimate several possible sign sequences connecting the two sets of eigenfunctions, and not only the sign sequence corresponding to the *primary correspondence*. To make this estimation possible we suggest constructing the descriptor $f^M(p)$ using the values of K eigenfunctions of the Laplace–Beltrami operator corresponding to the first K eigenvalues with no multiplicity, similar to [Ovsjanikov *et al.* (2008)]. In practice, we used the eigenfunctions corresponding to all the K first eigenvalues of the Laplace–Beltrami operator, assuming them to be of unit multiplicity. This assumption held true for most of the shapes we tested the algorithm on, and did not cause the algorithm to miss major primary or symmetrical correspondences. We did not find the scheme employing relative difference between eigenvalues to eliminate repeating eigenvalues, as suggested by [Ovsjanikov *et al.* (2008)], or the eigenfunction histogram-based method for eigenfunction order detection suggested in [Knossow *et al.* (2009)], to be efficient enough for our test cases. Currently this inability to cope with eigenvalue multiplicity and eigenfunction switching problems poses the main limitation on the proposed method, and must be eliminated in the future research.

5.3. *Multiple sign sequence estimation*

In order to estimate the sign sequences relating the two sets of eigenfunctions we suggest using a coarse matching based on absolute values of eigenfunctions together with geodesic distances measured on the two shapes.

It follows from Eq. (19) that the absolute values of the eigenfunction ϕ_k^M at two symmetric points $p, q \in M$ are equal, for all k

$$|\phi_k^M(p)| = |\phi_k^M(q)|. \tag{20}$$

Equivalently, the absolute values of the descriptors $f^M(p)$, $f^M(q)$ are equal

$$|f^M(p)| = |f^M(q)|. \tag{21}$$

In order to find the coarse matching between X and Y we minimize the pointwise dissimilarity cost, similar to Eq. (1), but using the absolute values of the descriptors f^X, f^Y

$$\tilde{\varphi}^* = \min_{\varphi} \sum_{x \in X} \| |f^X(x)| - |f^Y(\varphi(x))| \|, \tag{22}$$

$\tilde{\varphi}^*$ is the optimal coarse correspondence, and is less accurate than the correct correspondence. In order to simplify the problem, we fix the points on the first shape, X, and search for their corresponding points on Y

$$\tilde{\varphi}^*(x) = \arg\min_{\tilde{y} \in Y} \| |f^X(x)| - |f^Y(\tilde{y})| \|, \quad \forall\, x \in X \tag{23}$$

The sign sequence relating the two sets of descriptors can be calculated using the above correspondence set as follows

$$S_k = \arg\min_{\{+,-\}} \sum_{x \in X} \| f_k^X(x) - S_k f_k^Y(\tilde{\varphi}^*(x)) \|, \quad 1 \le k \le K, \tag{24}$$

where by S_k we denote the kth entry of the sign sequence. Somewhat similar approach to the sign sequence estimation was used by Ovsjanikov *et al.* [2008].

Now, let us analyze the coarse correspondence set. If the shapes are intrinsically symmetric, for each $x \in X$ there may exist several (symmetrical) points on Y, with descriptors equal up to signs, that minimize the expression

$$\| |f^X(x)| - |f^Y(y)| \|. \tag{25}$$

Thus, the set $\tilde{\varphi}^*$ would include both primary and symmetrical correspondences. In practice, as we work with sampled surfaces, the minimizer of Eq. (25) is usually unique, and may correspond to either primary or symmetrical alignment. We would like to estimate the sign sequences induced by all these alignments using Eq. (24). In order to do that we have to cluster the correspondences into groups according to the alignment they correspond to, and perform the sign sequence estimation for each group separately. The clustering procedure we suggest is based on comparing geodesic distances between pairs of corresponding points. Suppose we are given two such pairs, (x, y) and $(\tilde{x}, \tilde{y})$. If (x, y) and $(\tilde{x}, \tilde{y})$ represent the same alignment of X and Y, the corresponding geodesic distances $d_X(x, \tilde{x})$ and $d_Y(y, \tilde{y})$ must be similar. Otherwise, the geodesic distance error $|d_X(x, \tilde{x}) - d_Y(y, \tilde{y})|$ could be large.

We construct a matrix of the geodesic distance errors of all pairs of corresponding points: for all $(x_m, \tilde{\varphi}^*(x_m)), (x_n, \tilde{\varphi}^*(x_n))$

$$A_{mn} = |d_X(x_m, x_n) - d_Y(\tilde{\varphi}^*(x_m), \tilde{\varphi}^*(x_n))|. \tag{26}$$

The matrix A is used to cluster the correspondences, so that the sum of geodesic distance errors of all pair of corresponding points belonging to the same cluster is maintained low. In order to find the clusters we apply a dimensionality reduction algorithm, namely multidimensional scaling [Borg and Groenen (2005)], to A to obtain a set of points in Euclidean space, and cluster them using the K-means algorithm [Macqueen (1967)]. Given the clusters, we detect the corresponding sign sequences using Eq. (24). Moreover, we suggest clustering the correspondence into a large number of clusters, greater than the expected number of intrinsic symmetries. Thus, each sign sequence is supported by several clusters, improving the robustness to imprecise coarse correspondences and clustering errors. If the number of intrinsic symmetries of the shapes is known *a priori*, we choose a suitable number of sign sequences among sequences that were induced by the highest number of clusters. Alternatively, if we do not know the number of the intrinsic symmetries, we can choose the sign sequences supported by sufficiently high number of clusters. In our experiments, we clustered a set of 1,000 initial correspondences into 50 clusters. Figure 3 presents an example of two correspondence clusters obtained for a cat model. Note that, knowing the exact number of intrinsic symmetries, it is also possible to cluster the correspondences into the number of the symmetries plus one cluster. The sign sequences induced by those clusters will reflect the possible alignments of the two shapes.

The sign sequence estimation algorithm is summarized below.

(1) Coarse correspondence detection: for each $x \in X$, find its corresponding point using

$$y = \arg\min_{\tilde{y} \in Y} \| |f^X(x)| - |f^Y(\tilde{y})| \|. \tag{27}$$

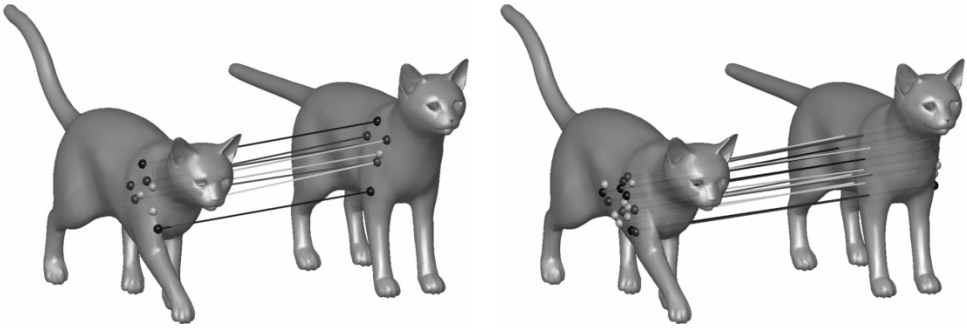

Primary correspondence Symmetrical correspondence

Fig. 3. Correspondence clustering example.

(2) Clustering: construct a matrix A

$$A_{mn} = |d_X(x_m, x_n) - d_Y(\tilde{\varphi}^*(x_m), \tilde{\varphi}^*(x_n))|, \quad \forall\, x_m, x_n \in X. \tag{28}$$

Apply to A MDS and K-means clustering to obtain J clusters of correspondences. We denote the set of clusters by $\{\mathcal{C}_j\}_{j=1}^{J}$.

(3) Sign sequence estimation: for each cluster, estimate the sign sequence S^j it induces by

$$S_k^j = \underset{\{+,-\}}{\arg\min} \sum_{(x,y)\in\mathcal{C}_j} \|f_k^X(x) - S_k f_k^Y(y)\|, \quad 1 \le k \le K. \tag{29}$$

Select the sign sequences induced by the highest number of clusters, according to the number of the intrinsic symmetries of the shape.

5.4. *Discretization of the Laplace–Beltrami operator*

In this work we used the cotangent weight scheme for the Laplace–Beltrami operator discretization, proposed by [Pinkall and Polthier (1993); Meyer *et al.* (2002)], with Neumann boundary conditions, to better account for small holes in some of the shapes we worked with. As was noted by Reuter *et al.* [2006], the Laplace–Beltrami operator calculated with Neumann boundary conditions is less altered by insertion of small holes than when Dirichlet boundary conditions are used. In order to calculate the eigendecomposition of the Laplace–Beltrami operator we solved the generalized eigendecomposition problem, as suggested by Rustamov [2007]. Different methods for discretization of the Laplace–Beltrami operator can be found in [Xu (2004); Reuter *et al.* (2006)].

6. Complexity Analysis and Results

The proposed algorithm was implemented in C++ and MATLAB©, and tested on various nonrigid isometric shapes from the TOSCA high resolution database [Bronstein *et al.* (2008)].

6.1. *Complexity analysis*

The algorithm consists of the following stages:

(1) Descriptors calculation. This part requires the calculation of the Laplace–Beltrami operator with complexity $O(N)$, where N is the number of vertices of the shape. The high-resolution shapes we tested the algorithm on had $27K$ to $52K$ vertices. The eigendecomposition of the Laplace–Beltrami operator was efficiently performed using the ARPACK package (available within MATLAB). The worst performance complexity is $O(N^3)$.

(2) Sign sequence estimation. First, we subsample the shapes at $N_1 = 1000$ vertices, using the farthest point sampling algorithm [Gonzalez (1985); Hochbaum and

Shmoys (1985)], and estimated the geodesic distances between them using the fast marching method [Kimmel and Sethian (1998)]. The overall complexity of the subsampling is therefore $O(N_1 N \log N)$. In order to reduce the calculation time, we could simplify the meshes prior to the subsampling [Garland (1999)], or approximate the distances using a fast approximation method that exploits parallel architectures [Weber *et al.* (2008)]. Since the subsampling is not a part of the proposed matching algorithm, we did not include it in the calculation times presented in Table 1. The coarse correspondence detection, the clustering using MDS and K-means, and the sigh sequence calculation can be performed in $O(JN_1^2)$, where J is the number of clusters.

(3) The highest complexity stage of the algorithm is solving the quadratic programming problem. It has been shown in [Pardalos *et al.* (1994)] that the quadratic assignment problem, or integer quadratic problem , is NP-hard. Hence, there exists no algorithm that can calculate the exact solution for the above optimization problem in polynomial time, unless we exploit some additional properties of the shapes' structure, e.g., as suggested in [Wang *et al.* (2010), and Raviv *et al.* (2011)]. We leave the discussion of possible approximation techniques to future research. In this work we used the mixed integer quadratic programming (MIQP) solver to approximate the solution of the quadratic problem. The MIQP solver is distributed as a part of the hybrid toolbox by Bemporad [2004]. Its calculation time is exponential in the number of variables, due to branch-and-bound method it employs. Therefore, the number of correspondences we were able to find was relatively small (it was chosen so that it would be possible to solve the IQP problem in reasonable time). We subsampled 20 points from one shape, and 40 candidate corresponding points from the second shape, using farthest point sampling with random seed point, and solved the optimization problem (11) with constraints given by Eq. (8), thus obtaining a total of 20 correspondences. As we consider all possible correspondences between the 20 and 40 points sampled from the two shapes, respectively, the order of the points in the two sets is not important. This subsampling, however, affects the accuracy of the matching, in terms of proximity of the matched points and the true correspondences — the denser the subsampling is the more accurate the obtained correspondences potentially are. According to [Memoli and Sapiro (2005)], the L_∞-norm of the geodesic distance differences converges to

Table 1. Calculation time, in seconds, for different models from the TOSCA database, on a laptop with Intel Core 2 Duo T7500 processor and 2 G Bytes memory.

Model	#vertices	Total	IQP	IQP (%)
Human	53K	256.7	222.3	86.6
Horse	19K	241.9	222.5	92.0
Cat	28K	241.3	222.0	92.0
Hand	7K	142.3	110.9	78.0

the Gromov–Hausdorff distance as the sampling is refined, with an approximation error bounded by the sampling rate. In the future research we would like to extend their result for the problem we formulated.

Other optimization techniques can be used to approximate the solution of the above IQP, for instance GMDS-like continuous optimization techniques, or discrete optimization algorithms developed for labeling problems, e.g. [Torresani *et al.* (2008)]. In particular, the GMDS algorithm requires good initialization, and thus it can be used for refinement of correspondences found by the proposed method, or for adding more correspondences. In this work we used the general integer quadratic solver, as described above, and we intend to experiment with different solvers in the future.

Table 1 presents typical algorithm calculation times. The column named "Total" presents the total computation time, and the column "IQP" — computation time of the IQP solver. The rightmost column, named "IQP (%)", presents the computation time of the IQP solver, as a fraction of the total calculation time. Note that the first three results in Table 1 correspond to shapes with two intrinsic symmetries, therefore two quadratic problems were solved for each. Future attempts to improve the algorithm must include a thoughtful analysis of the quadratic optimization problem, and suggestions for reducing its complexity.

6.2. *Correspondence detection*

We tested the proposed algorithm on shapes that underwent different types of transformations.

Rigid transformations: both the Laplace–Beltrami operator and the geodesic distances remain constant under rigid transformations, therefore in this case the matching accuracy is determined by the subsampling and by the accuracy of the approximation of the quadratic problem solution. We empirically examined the effect of the sampling on the correspondence accuracy by matching two shapes differing by rigid transformations, and sampled using the farthest point sampling algorithm with random initial sample choice. Figure 4 presents the correspondence obtained with the proposed method. The sampled points are shown with small spheres. The shape on the right was subsampled denser than the left shape, and the matched points were connected by lines. It can be seen that the algorithm converged to a local minimum, which is due to both the subsampling and the approximation of the IQP solution. Note that the subsampling affects the accuracy of our matching algorithm in all the following settings as well.

Different sampling and triangulation: the algorithm succeeded to match models with different number of vertices and different triangulations. It can be explained by the fact that both geodesic distances and eigenfunctions of the Laplace–Beltrami operator approximations that we used remain approximately constant under this kind of change in shape representation. Figures 5(a) and 5(b) show the

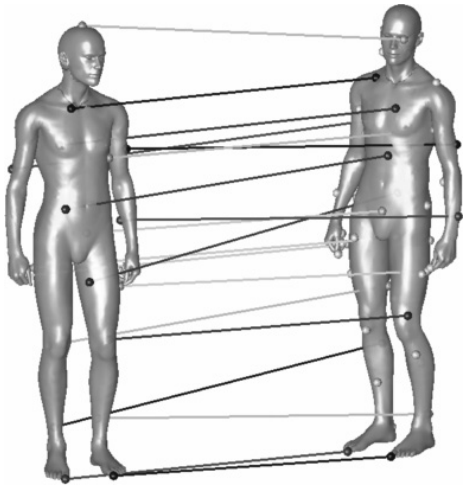

Fig. 4. Correspondence results obtained for two shapes differing by rigid transformation, and sampled using the farthest point sampling algorithm with random initialization.

correspondence detected between the shapes having vertex number ratio of approximately 1:2 and 1:10, respectively.

The correspondence obtained for pairs of *isometric shapes* from the TOSCA database are shown in Figs. 6–10. Figures 6–8 and 10 show both primary and symmetrical correspondence results, as all these shapes have exactly one intrinsic symmetry. The results presented in Fig. 10 were calculated using the commute time

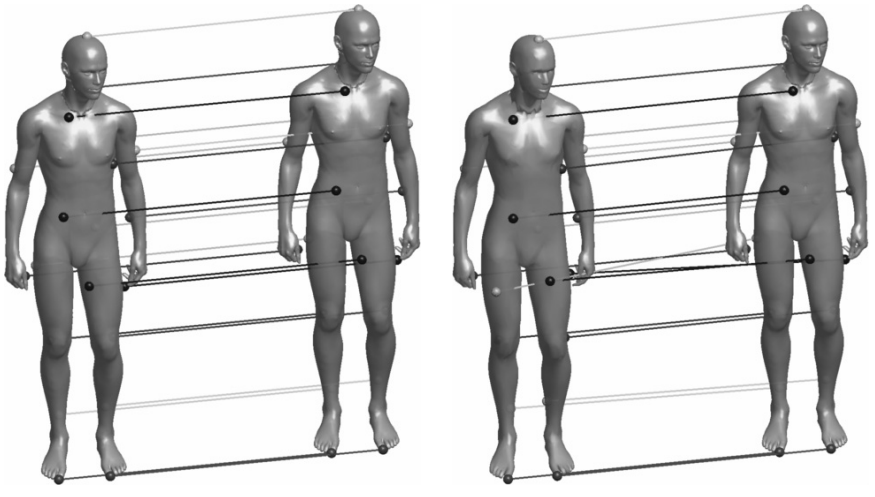

(a) (b)

Fig. 5. Correspondence results obtained with the proposed method for a human body with different number of vertices and triangulations. (a) A model with 25K vertices vs. original model (53K vertices) and (b) A model with 5K vertices vs. original model.

distances [Qiu and Hancock (2007)], and are similar to those obtained with geodesic distances (Fig. 7). Figure 9 shows two examples of correspondences obtained for hand shapes.

In order to be able to match shapes that differ by *scaling*, we normalized the vertices of the two shapes to obtain equal maximal geodesic distances for both shapes. Other descriptors and metric can be employed here as well: e.g., the Global Point Signature [Rustamov (2007)], which is invariant to uniform scaling of the surface, and diffusion scale-space distance described in [Bronstein and Bronstein (2010)], respectively. Figure 11 presents the correspondence detected between shapes that differ by both isometric transformation and scaling.

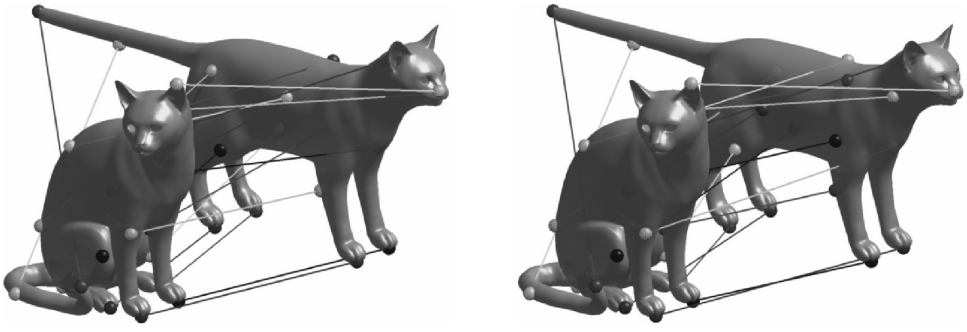

Primary correspondence Symmetrical correspondence

Fig. 6. Correspondence results obtained with the proposed method for a cat shape.

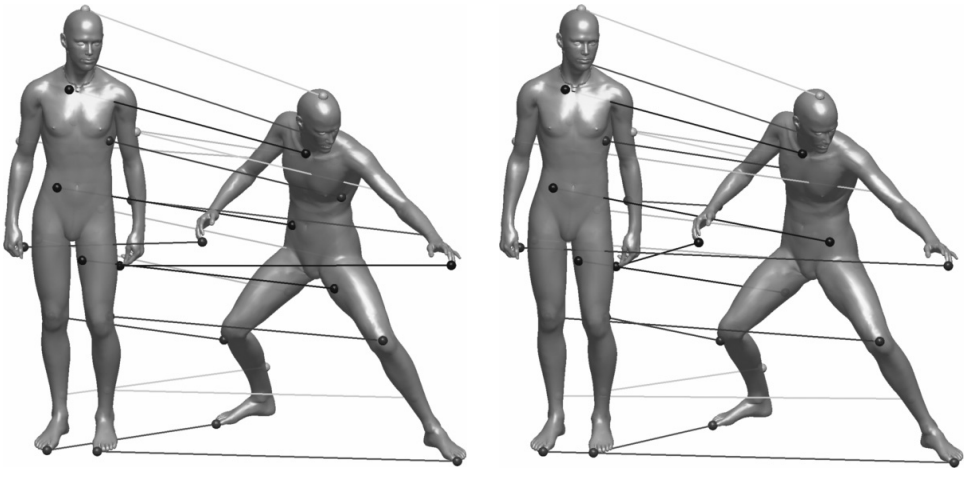

Primary correspondence Symmetrical correspondence

Fig. 7. Correspondence results obtained with the proposed method for a human shape.

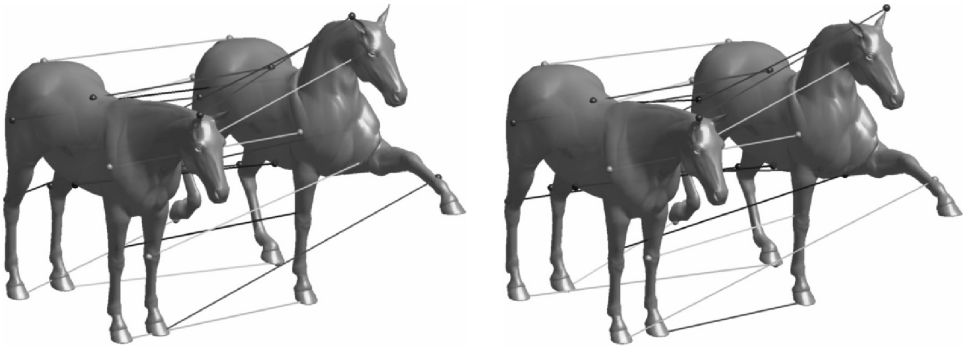

Primary correspondence Symmetrical correspondence

Fig. 8. Correspondence results obtained with the proposed method for a horse shape.

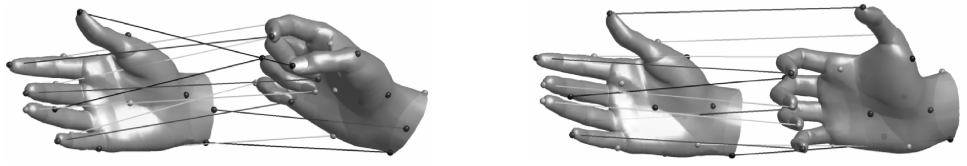

Fig. 9. Two correspondence results obtained with the proposed method for hand shapes.

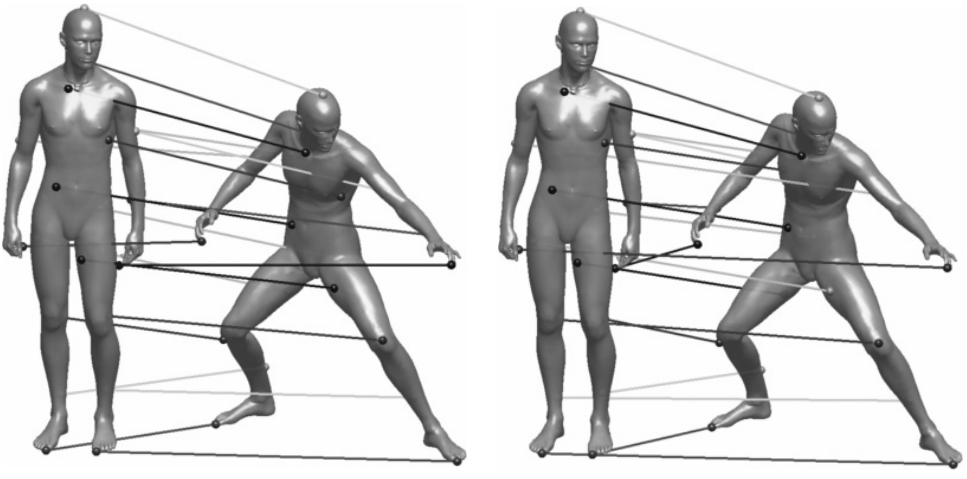

Primary correspondence Symmetrical correspondence

Fig. 10. Correspondence results obtained with the proposed method, using commute time distances to calculate the pairwise dissimilarity.

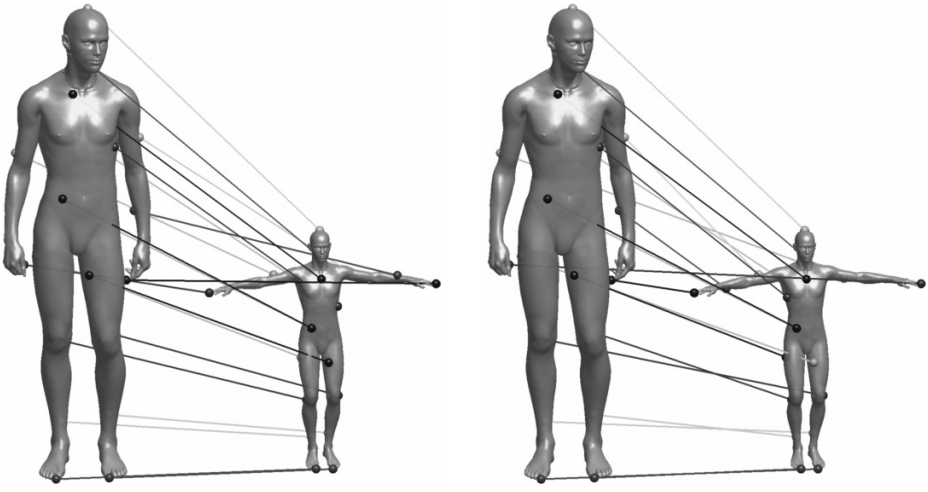

Primary correspondence Symmetrical correspondence

Fig. 11. Two correspondence result obtained for shapes at different scales.

*Algorithm evaluation on the SHREC'*10 *benchmark*: additional results of the algorithm performance can be found in [Bronstein *et al.* (2010)], presenting the algorithm evaluation on the SHREC'10 robust correspondence benchmark. The benchmark included shapes that underwent various types of transformations: isometric and topological transformations, insertion of microholes and big holes, scaling, local scaling, additive Gaussian noise, shot noise, and down sampling. Each one of the transformations was presented at five different strengths. The algorithm performance was evaluated based on the quality of the correspondence it provided. That is, for each pair of shapes, denoted here by X and Y, the algorithm provided a set of $M < |Y|$ correspondences $\mathcal{C}(X,Y) = \{(x_k, y_k)\}_{k=1}^{M}$. Those correspondence were compared to the groundtruth correspondence set as $\mathcal{C}_0(X,Y) = \{(x'_k, y_k)\}_{k=1}^{|Y|}$, in the following manner

$$D(\mathcal{C}, \mathcal{C}_0) = \frac{1}{M} \sum_{k=1}^{M} d_X(x_k, x'_k), \tag{30}$$

where d_X is the geodesic distances measured on the shape X. The results (see Tables 3 and 4 in [Bronstein *et al.* (2010)]) indicate that the proposed method is robust to isometric transformations, holes, global scaling, sampling (the cotangent weight scheme), and moderate noise. As expected, since neither the Laplace–Beltrami operator nor the geodesic distance measure is invariant to changes of topology or local scale, the algorithm performs poorly for those transformations.

To compare the performance of the proposed method with the GMDS algorithm (used with geodesic distances), we summarized the averaged values of the correspondence errors $D(X,Y)$ in Table 2. The results show that in both settings

Table 2. SHREC'10, algorithm performance comparison: average geodesic distances from the groundtruth correspondence.

	Transformation strength				
Algorithm	1	≤2	≤3	≤4	≤5
GMDS	39.92	36.77	35.24	37.40	39.10
The proposed method, cotangent weight scheme	15.51	18.21	22.99	25.26	28.69
The proposed method, graph Laplacian	10.61	15.48	19.01	23.22	23.88

the average performance of the proposed method is better than that of the GMDS algorithm, for all transformation strengths.

For a more detailed description of the SHREC'10 correspondence benchmark, the compared algorithms and the comparison results the reader is referred to [Bronstein *et al.* (2010)].

6.3. *Combined dissimilarity vs. other dissimilarity measures*

In this section, we would like to present an empirical evidence for the fact that minimizing the proposed combined dissimilarity measure $\mathrm{Dis}(P)$ produces better correspondence than minimizing each of its parts (linear and quadratic) separately. For this purpose, we conducted the following experiment: we compared the correspondence obtained by minimizing the dissimilarity measure $\mathrm{Dis}(P)$, which we denote by P_0^*, with correspondences obtained by

(1) Minimizing the linear part of the dissimilarity $\mathrm{Dis}(P)$ (Eq. (6)), thus only minimizing the difference between the pointwise surface descriptors.
(2) Minimizing the quadratic part of the dissimilarity $\mathrm{Dis}(P)$ (Eq. (6)), thus only minimizing the difference between the metric structures of the shapes. This matching problem formulation is similar to that used in [Memoli and Sapiro (2005); and Bronstein *et al.* (2006)], and thus is likely to suffer from the same problem, namely convergence to a local minimum without proper initialization.
(3) Using the combined dissimilarity measure $\mathrm{Dis}(P)$, with different surface descriptors. That is, we constructed the surface descriptors using the absolute values of the eigenvectors of the Laplace–Beltrami operator, thus we had only one set of descriptors for each shape.

Figure 12 shows the primary correspondences obtained as described above, for two shapes of a cat. We suggest two types of evaluation of the obtained correspondence quality: numerical and visual evaluation (Fig. 12). For the numerical evaluation we calculated three dissimilarity measures, for each one of the obtained correspondences: the combined dissimilarity $\mathrm{Dis}(P)$ given by Eq. (6), linear dissimilarity $\mathrm{Dis}_{\mathrm{l}}(P)$, and quadratic dissimilarity $\mathrm{Dis}_{\mathrm{q}}(P)$. The results can be found in Table 3. Naturally, the correspondence obtained using the proposed method is the most accurate in terms of $\mathrm{Dis}(P)$. The correspondence obtained by minimizing the linear part of the dissimilarity measure, $\mathrm{Dis}_{\mathrm{l}}(P)$ is very similar to P_0^*, but is

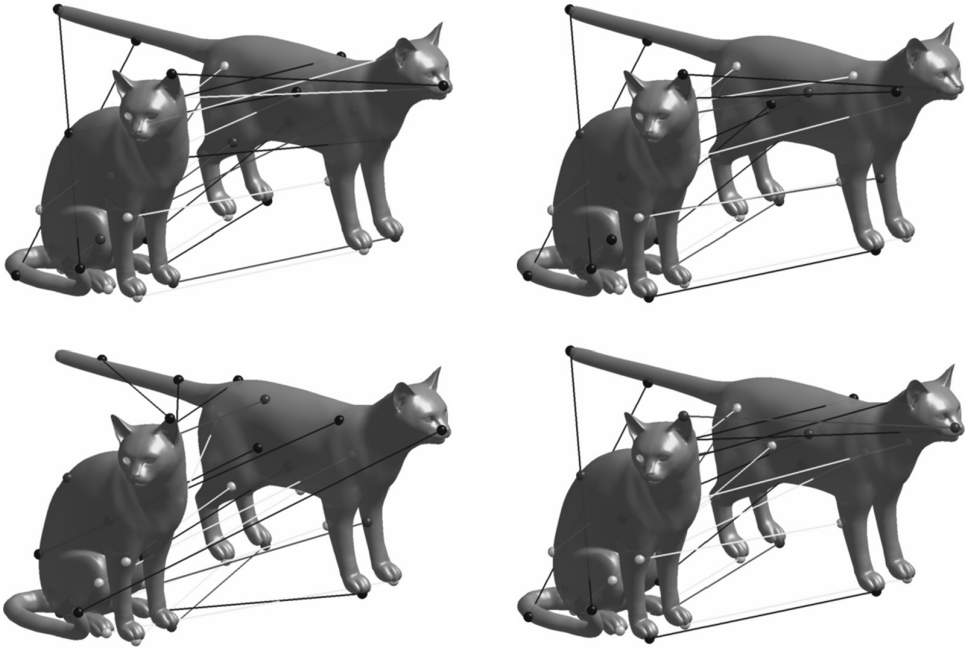

Fig. 12. Correspondences obtained by minimizing different distortion measures: (a) combined distortion measure, denoted by P_0^*; (b) descriptor based distortion; (c) metric structure based distortion; and (d) combined distortion with descriptors constructed of absolute values of eigenvectors of Δ_M.

less accurate in terms of pairwise point relationships $\mathrm{Dis}_{\mathrm{q}}(P)$. Since in this case we do not have to solve the quadratic problem, its computation time is the shortest. The correspondence obtained by minimizing the quadratic part alone is clearly incorrect. We think it is the result of the minimization algorithm getting stuck at a local minima, with high dissimilarity values. The calculation time is also the longest for this type of optimization. The fourth correspondence is the closest to P_0^* in terms of the dissimilarity $\mathrm{Dis}(P)$. The main disadvantage of using the absolute values of the eigenvectors as the descriptors is the fact that one can only find a

Table 3. Dissimilarity values $\mathrm{Dis}(P)$, $\mathrm{Dis}_l(P)$, and $\mathrm{Dis}_q(P)$ calculated using the correspondences obtained using the four methods described in Sec. 6.3.

	Evaluation measure		
Minimized dissimilarity	$\mathrm{Dis_l}(P)$	$\mathrm{Dis_q}(P)$	$\mathrm{Dis}(P)$
$\mathrm{Dis}(P)$	1.48	7.18	8.66
$\mathrm{Dis_l}(P)$	1.28	8.37	9.66
$\mathrm{Dis_q}(P)$	9.44	10.49	19.93
$\mathrm{Dis}(P)$, with $\|\phi\|$	1.62	7.64	9.26

single correspondence this way. Thus, the descriptors calculated using the absolute values of the eigenvectors should be considered when one is interested in a single correspondence, either primary or symmetrical one.

7. Conclusions

The paper presents a general framework for finding correspondences between non-rigid isometric shapes. We formulated the correspondence detection problem as a minimization of dissimilarity between pointwise surface properties and pairwise structures. We described a method for construction of isometry invariant surface descriptors based on eigendecomposition of the Laplace–Beltrami operator. Those descriptors, coupled with geodesic distances were used to calculate the dissimilarity measure. We also discussed the problem of correspondence ambiguity that occurred when matching intrinsically symmetric shapes. We showed that by employing the above surface descriptors we could find several possible correspondences between the shapes. We then demonstrated the algorithm's performance for different types of transformations, including isometric transformations, scaling, and remeshing.

One of the current limitations of the proposed method is the inability to cope with complex symmetry groups, that are represented by rotational symmetries in the Laplace–Beltrami eigenfunction space, as well as the problem of coping with eigenfunction switching. The algorithm can also benefit from a more efficient solution of the minimization problem, which currently constitutes the bottle-neck of the algorithm's computational complexity. This also requires severe subsampling of the shapes, thus affecting the matching accuracy. The possible improvements in this direction include utilizing the problem structure to diminish the space of possible solutions, and applying continuous relaxation techniques to speed up the minimization. Future work may also include extending the algorithm to handle other important cases of shape matching, such as partial shape matching, and matching in the presence of topological and local scale changes.

Acknowledgment

We thank Drs. Alexander and Michael Bronstein for helpful discussions. This research was supported by the USA Office of Naval Research (ONR) grant.

References

Anguelov, D., Srinivasan, P., Pang, H.-C., Koller, D. and Thrun, S. (2004). The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces, *Proceedings of the Neural Information Processing Systems* (*NIPS*) *Conference*, pp. 33–40.

Bemporad, A. (2004). Hybrid Toolbox — User's Guide. http://www.dii.unisi.it/hybrid/toolbox.

Ben-Chen, M., Butscher, A., Solomon, J. and Guibas, L. J. (2010). On discrete killing vector fields and patterns on surfaces. *Comput. Graph. Forum (Proc. SGP)*, **29**(5): 1701–1711.

Bérard, P., Besson, G. and Gallot, S. (1994). Embedding Riemannian manifolds by their heat kernel, in *Geometric Find Functional Analysis*, **4**(4): 373–398.

Borg, I. and Groenen, P. J. F. (2005). *Modern Multidimensional Scaling: Theory and Applications* (*Springer Series in Statistics*), 2nd Edition, Springer, Berlin.

Bronstein, A. M. and Bronstein, M. M. (2010). Analysis of diffusion geometry methods for shape recognition. *IEEE Trans. Patt. Anal. Mach. Intell.* (*PAMI*) (submitted).

Bronstein, A. M., Bronstein, M. M., Castellani, U., Dubrovina, A., Guibas, L. J., Horaud, R. P., Kimmel, R., Knossow, D., von Lavante, E., Mateus, D., Ovsjanikov, M. and Sharma, A. (2010 (to appear)). Shrec 2010: Robust correspondence benchmark. *Proceedings of the EUROGRAPHICS Workshop on 3D Object Retrieval* (*3DOR*).

Bronstein, A. M., Bronstein, M. M. and Kimmel, R. (2006). Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching, *Proceedings National Academy of Sciences* (*PNAS*), pp. 1168–1172.

Bronstein, A. M., Bronstein, M. M. and Kimmel, R. (2008). *Numerical Geometry of Non-Rigid Shapes*, Springer, New York.

Bronstein, M. M. and Kokkinos, I. (2010). Scale-invariant heat kernel signatures for non-rigid shape recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*).

Bronstein, A. M., Bronstein, M. M., Kimmel, R., Mahmoudi, M. and Sapiro, G. (2009). A Gromov–Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *Int. J. Comput. Vis.* (*IJCV*).

Chang, W. and Zwicker, M. (2008). Automatic registration for articulated shapes. *Comput. Graph. Forum* (*Proc. SGP*), **27**: 1459–1468.

Coifman, R. R. and Lafon, S. (2006). Diffusion maps. *Appl. Comput. Harmonic Anal.*, **21**: 5–30.

Dubrovina, A. and Kimmel, R. (2010). Matching shapes by eigendecomposition of the Laplace–Beltrami operator. *International Symposium on 3D Data Processing Visualization and Transmission*.

d'Amico, M., Frosini, P. and Landi, C. (2006). Using matching distance in size theory: A survey. *Int. J. Imag. Syst. Technol.*, **16**(5): 154–161.

Elad, A. and Kimmel, R. (2003). On bending invariant signatures for surfaces. *IEEE Trans. Patt. Anal. Mach. Intell.* (*PAMI*), **25**: 1285–1295.

Garland, M. (1999). *Quadric-Based Polygonal Surface Simplification*. PhD thesis, School of Computer Science, Carnegie Mellon University.

Gonzalez, T. F. (1985). Clustering to minimize the maximum intercluster distance. *Theoret. Comput. Sci.*, **38**: 293–306.

Gromov, M. (1981). *Structures Metriques Pour les Varietes Riemanniennes*, Textes Mathematiques, no. 1.

Hochbaum, D. S. and Shmoys, D. B. (1985). A best possible heuristic for the k-center problem. *Math. Oper. Res.*, **10**(2): 180–184.

Hu, J. and Hua, J. (2009). Salient spectral geometric features for shape matching and retrieval. *Vis. Comput.*, **25**: 667–675.

Huang, Q.-X., Adams, B., Wicke, M. and Guibas, L. J. (2008). Non-rigid registration under isometric deformations. *Comput. Graph. Forum*, **27**: 1449–1457.

Jain, V., Zhang, H. and Kaick, O. V. (2007). Non-rigid spectral correspondence of triangle meshes. *Int. J. Shape Model.*, 101–124.

Kimmel, R. and Sethian, J. A. (1998). Computing geodesic paths on manifolds. In *Proceedings of National Academy of Sciences* (*PNAS*), pp. 8431–8435.

Kim, V., Lipman, Y., Chen, X. and Funkhouser, T. (2010). Mobius transformations for global intrinsic symmetry analysis. *Proceedings of Eurographics Symposium on Geometry Processing* (*SGP*), **29**(5).

Knossow, D., Sharma, A., Mateus, D. and Horaud, R. P. (2009). Inexact matching of large and sparse graphs using Laplacian eigenvectors. *Proceedings of 7th Workshop on Graph-based Representations in Pattern Recognition*.

Leordeanu, M. and Hebert, M. (2005). A spectral technique for correspondence problems using pairwise constraints, *Proceedings of International Conference of Computer Vision* (*ICCV*), pp. 1482–1489.

Levy, B. (2006). Laplace–Beltrami eigenfunctions towards an algorithm that "understands" geometry. *IEEE International Conference on Shape Modeling and Applications*, 13.

Levy, B. and Zhang, R. H. (2009). Spectral Geometry Processing. *ACM SIGGRAPH ASIA Course Notes*.

Lipman, Y. and Funkhouser, T. (2009). Mobius voting for surface correspondence. *ACM Trans. Graph.* (*Proc. SIGGRAPH*), **28**.

Lipman, Y., Chen, X., Daubechies, I. and Funkhouser, T. (2010). Symmetry factored embedding and distance. *ACM Trans. Graph.* (*Proc. SIGGRAPH*).

Maciel, J. and Costeira, J. P. (2003). A global solution to sparse correspondence problems. *IEEE Trans. Patt. Anal. Mach. Intell.* (*PAMI*), **25**: 187–199.

Macqueen, J. B. (1967). Some methods of classification and analysis of multivariate observations. *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, pp. 281–297.

Mateus, D., Cuzzolin, F., Horaud, R. P. and Boyer, E. (2007). Articulated shape matching using locally linear embedding and orthogonal alignment. *Proceedings of the IEEE Workshop on Non-Rigid Registration and Tracking Through Learning* (*NRTL*).

Mateus, D., Horaud, R. P., Knossow, D., Cuzzolin, F. and Boyer, E. (2008). Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*).

Memoli, F. and Sapiro, G. (2005). A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, **5**: 313–347.

Memoli, F. (2007). On the use of Gromov–Hausdorff distances for shape comparison. *Point Based Graph.*, 81–90.

Meyer, M., Desbrun, M., Schröder, P. and Barr, A. (2002). Discrete differential geometry operators for triangulated 2-manifolds. *International Workshop on Visualization and Mathematics*.

Ovsjanikov, M., Sun, J. and Guibas, L. (2008). Global intrinsic symmetries of shapes. *Comput. Graph. Forum*, **27**: 1341–1348.

Pardalos, P. M., Rendl, F. and Wolkowicz, H. (1994). The quadratic assignment problem: A survey and recent developments. *Proceedings of the DIMACS Workshop on Quadratic Assignment Problems, Vol. 16 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, American Mathematical Society, pp. 1–42.

Pinkall, U. and Polthier, K. (1993). Computing discrete minimal surfaces and their conjugates. *Exp. Math.*, **2**: 15–36.

Qiu, H. and Hancock, E. (2007). Clustering and embedding using commute times. *IEEE Trans. Patt. Anal. Mach. Intell.* (*PAMI*).

Raviv, D., Bronstein, A. M., Bronstein, M. M. and Kimmel, R. (2007). Symmetries of non-rigid shapes. *Proceedings of Workshop on Non-Rigid Registration and Tracking Through Learning* (*NRTL*).

Raviv, D., Bronstein, A. M., Bronstein, M. M. and Kimmel, R. (2009). Full and partial symmetries of non-rigid shapes. *Int. J. Comput. Vis.* (*IJCV*).

Raviv, D., Dubrovina, A. and Kimmel, R. (2011). Hierarchical matching of non-rigid shapes. *International Conference on Scale Space and Variational Methods in Computer Vision* (*SSVM'11*), submitted.

Reuter, M., Wolter, F.-E. and Peinecke, N. (2006). Laplace–Beltrami spectra as "Shape-DNA" of surfaces and solids, *Computer-Aided Design*, **38**: 342–366.

Rosenberg, S. (1997). *The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds*, Cambridge University Press, Cambridge.

Ruggeri, M. R. and Saupe, D. (2008). Isometry-invariant matching of point set surfaces, *Proceedings of the Eurographics 2008 Workshop on 3D Object Retrieval*.

Rustamov, R. M. (2007). Laplace–Beltrami eigenfunctions for deformation invariant shape representation, *Proceedings of SGP*, pp. 225–233.

Shapiro, L. S. and Brady, J. M. (1992). Feature-based correspondence: An eigenvector approach. *Image Vision Comput.*, **10**: 283–288.

Sun, J., Ovsjanikov, M. and Guibas, L. (2009). A concise and provably informative multi-scale signature based on heat diffusion. *Proceedings of the Eurographics Symposium on Geometry Processing* (*SGP*).

Tevs, A., Bokeloh, M., Wand, M., Schilling, A. and Seidel, H.-P. (2009). Isometric registration of ambiguous and partial data. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*).

Thorstensen, N. and Keriven, R. (2009). Non-rigid shape matching using geometry and photometry. *Asian Conference on Computer Vision*, pp. 1–12.

Torresani, L. and Kolmogorov, V. and Rother, C. (2008). Feature correspondence via graph matching: Models and global optimization. *Proceedings of the 10th European Conference on Computer Vision* (*ECCV '08*), pp. 596–609.

Umeyama, S. (1988). An eigendecomposition approach to weighted graph matching problems. *IEEE Trans. Patt. Anal. Mach. Intell.* (*PAMI*), **10**: 695–703.

Wang, C., Bronstein, M. M. and Paragios, N. (2010). Discrete minimum distortion correspondence problems for non-rigid shape matching. *Technical Report, INRIA Research Report 7333, Mathématiques Appliquées aux Systèmes, École Centrale Paris*.

Weber, O., Devir, Y., Bronstein, A. M., Bronstein, M. M. and Kimmel, R. (2008). Parallel algorithms for approximation of distance maps on parametric surfaces. *ACM Trans. Graph.*, **27**: 1–16.

Xu, G. (2004). Discrete Laplace–Beltrami operators and their convergence. *Comput. Aided Geom. Design*, **21**: 767–784.

Zaharescu, A., Boyer, E., Varanasi, K. and Horaud, R. P. (2009). Surface feature detection and description with applications to mesh matching. in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*).

Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., van Kaick, O. and Tagliasacchi, A. (2008). Deformation-driven shape correspondence. *Comput. Graph. Forum* (*Proc. SGP*), **27**: 1431–1439.

Zigelman, G., Kimmel, R. and Kiryati, N. (2008). Texture mapping using surface flattening via multi-dimensional scaling. *IEEE Trans. Visual. Comput. Graph.*, **2**(2): 198–207.