Matching shapes by eigendecomposition of the Laplace-Beltrami operator

Anastasia Dubrovina
Technion - Israel Institute of Technology
nastyad@tx.technion.ac.il

Ron Kimmel
Technion - Israel Institute of Technology
ron@cs.technion.ac.il

Abstract

We present a method for detecting correspondences between non-rigid shapes, that utilizes surface descriptors based on the eigenfunctions of the Laplace-Beltrami operator. We use clusters of probable matched descriptors to resolve the sign ambiguity in matching the eigenfunctions. We then define a matching cost that measures both the descriptor similarity, and the similarity between corresponding geodesic distances measured on the two shapes. We seek for correspondence by minimizing the above cost. The resulting combinatorial problem is then reduced to the problem of matching a small number of feature points using quadratic integer programming.

1. Introduction

Correspondence detection is an important component in various applications in the field of three dimensional shape processing and analysis, amongst them registration, similarity measurement for recognition and retrieval, shape deformation and morphing, symmetry detection, surface completion, etc. The problem of detecting correspondence between rigid shapes has been widely addressed in the literature (see *e.g.* [14, 31, 34]). As for non-rigid shapes, correspondence detection is a challenging problem.

In this paper we propose a framework for an unsupervised correspondence detection between non-rigid shapes that differ by nearly isometrical transformations. It is based on intrinsic invariant surface descriptors. We show that such descriptors can be derived from the eigendecomposition of the Laplace-Beltrami operator [18, 19, 29].

We then formulate the correspondence detection as a quadratic optimization problem, over the space of all possible mappings between the two shapes. Its objective function consists of a linear term, that measures similarity between the descriptors of the corresponding points on the two shapes, and a quadratic regularization term that measures the quality of pairwise correspondence assignment.

In addition we elaborate on the correspondence ambiguity problem that arises when matching two isometries of an intrinsically symmetrical shape. In this case, there is more than one correspondence that preserves the metric structure of the shapes (geodesic distances). We show that it is possible to find a distinct set of descriptors for each of those correspondences, thus resolving to a set of distinct correspondence detection problems. We would like to mention that, to our best knowledge, there is no alternative algorithm that can handle this correspondence ambiguity problem.

The paper is organized as follows: some related results in the field of non-rigid surface matching are reviewed in the next section. In Section 3 we describe the isometric invariant surface descriptors based on eigendecomposition of the Laplace-Beltrami operator. Additionally, we describe a new method for detection of multiple potential alignments of the two surfaces. In Section 4 we formulate the correspondence detection as a quadratic optimization problem, and elaborate on its solution. In Section 5 we present the possible matching results, including intrinsic symmetry detection. Section 6 concludes the paper and presents several directions for future study.

2. Previous work

A common approach to non-rigid shape matching consists in finding transformation invariant representations of the shapes, and performing the matching in the representation domain. One of the first methods, proposed by Elad and Kimmel [7], embeds the surface into a (flat) Euclidian space, such that the Euclidian distances between points in the flat space approximate the geodesic distances on the surface. Jain and Zhang [12] proposed matching the embeddings of the shapes into a spectral domain, using a non-rigid variant of the ICP (Iterative Closest Point) algorithm. Mateus *et al*. [23], and later Knossow *et al*. [15], used spectral embedding as well, performing registration in the embedding domain using a probabilistic framework. Rustamov [32] suggested using the eigendecomposition of the Laplace-Beltrami operator to construct an isometric invariant surface representation, though aiming rather to measure similarity between non-rigid shapes, than for correspondence detection. In a recent paper, Lipman and Funkhouser [20], treated isometrical transformations as a sub-set of

Möbius transformations, and detected correspondences in the canonical domain of Möbius-invariant representations. The above methods produce good correspondence results, but none of them deals with possible ambiguity due to intrinsic symmetries.

Another group of methods for non-rigid shape comparison [4, 24, 25] measures the similarity between the metric spaces of the shapes using the Gromov-Hausdorf distance [8]. The Generalized Multidimensional Scaling (GMDS) algorithm [4] also produces correspondence between two given shapes. Thorstensen and Keriven [36] extended the GMDS framework to surfaces with textures. Though mathematically sound, the above methods require significant computation efforts.

The two methods proposed by by Hu and Hua [10] and Zaharescu *et al*. [38] detect correspondences between non-trivial feature points on the surfaces. In order to find features they employ the Laplace-Beltrami decomposition, and an extension of 2D feature detection and description approach to meshes, accordingly.

Other correspondence detection methods include the algorithm proposed by Zhang *et al*. [39], based on combinatorial tree traversal for correspondence search; a method by Anguelov *et al*. [1], who defined a probabilistic model over the set of all possible correspondences, and used belief propagation technique to match the shapes. Chang and Zwicker [6] and Huang *et al*. [11] proposed modeling the non-rigid transformation between the shapes by a set of rigid transformations applied to parts of the shapes. Tevs *et al*. [35] proposed a new algorithm based on a geodesic distance-preserving randomized feature matching.

Another, more general, approach to the correspondence detection was proposed by Leordeanu and Hebert in [17]. It is based on a spectral decomposition technique, and employs both local descriptor similarity and global pairwise correspondence assignment quality to measure the correspondence cost.

Intrinsic symmetry detection can be viewed as an extension of correspondence detection, where one needs to find the best possible mapping of a shape to itself. Raviv *et al*. [28] formulated the symmetry detection as a problem of embedding a shape into itself, and used GMDS [4] to solve it. Ovsjanikov *et al*. [27] showed that intrinsic symmetry detection can be reduced to extrinsic symmetry detection in the domain of the Global Point Signature (GPS) embedding [32] of the surface.

Our main contributions can be summarized as follows:

- We propose simple yet effective isometric invariant surface descriptors.

- Introduce the problem of correspondence ambiguity and show how it can be solved, resulting in distinct sets of descriptors for each possible correspondence.

- Formulate the correspondence problem as an Integer Quadratic Programming (IQP) problem, so that both the similarity between the descriptors and the pairwise geodesic distances on the two shapes would be preserved by the matching.

3. Isometric invariant descriptors

In this section we introduce surface descriptors based on the eigenfunctions of the Laplace-Beltrami operator.

3.1. Laplace-Beltrami operator

The Laplace-Beltrami operator Δ_M of a compact Riemannian manifold M is a second order differential operator defined by the metric tensor of M (see [30]). Hence, the Δ_M operator is invariant to isometries of the manifold.

The Laplacian eigenvalue problem is given by

$$\Delta_M \phi = \lambda \phi, \tag{1}$$

where λ is an eigenvalue of Δ_M, and ϕ is the eigenfunction associated with λ. By definition, Δ_M is a positive semi-definite operator. Therefore, all its eigenvalues $\lambda_i, i \geq 1$, are positive, with corresponding orthogonal set of eigenfunctions $\{\phi_k\}$. Let us order the eigenvalues according to their magnitude: $0 < \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq ....$ Note that when the manifold M has no boundary, Δ_M has an additional eigenvalue $\lambda_0 = 0$. Its corresponding eigenfunction is a constant function on M. Generally speaking, as the value of the eigenvalue increases, its corresponding eigenfunction varies more rapidly over the mesh (see Theorem 2.1 [13], for example).

3.2. Surface descriptors

Like the Laplace-Beltrami operator, the eigenfunctions ϕ_k are invariant to the isometries of the manifold M. Therefore, they could serve as descriptors

$$f^M(\mathbf{p}) = \{\phi_1^M(\mathbf{p}), \phi_2^M(\mathbf{p}), \phi_3^M(\mathbf{p}), ...\}, \quad \mathbf{p} \in M, \tag{2}$$

where $\phi_k^M(\mathbf{p})$ is the value of the k-th eigenfunction of Δ_M at the point $\mathbf{p}$.

The descriptor $f^M(\mathbf{p})$ is in fact the embedding of $\mathbf{p}$ into the vector space spanned by the eigenfunctions of the Δ_M operator. There are, though, several difficulties preventing such a straight forward usage.

- The computational complexity of the matching depends on the dimension of the descriptors. Instead of using all the eigenfunctions of the Laplace-Beltrami operator, we restrict ourselves to the first K eigenfunctions

$$f^M(\mathbf{p}) = \{\phi_1^M(\mathbf{p}), \phi_2^M(\mathbf{p}), ..., \phi_K^M(\mathbf{p})\}. \tag{3}$$

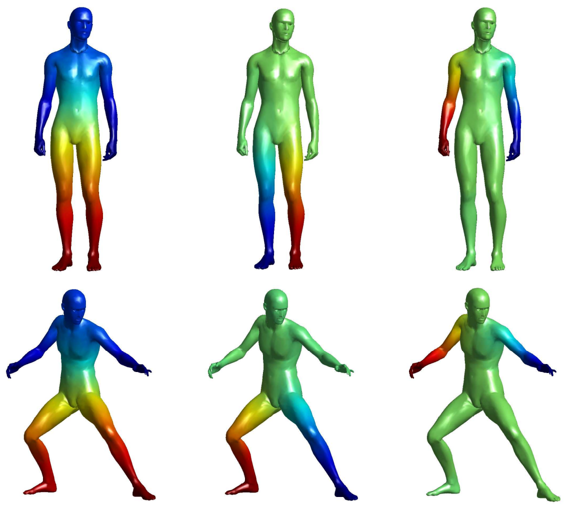

Figure 1. Two articulations of a human shape, colored according to the values of the first three eigenfunctions of their Laplace-Beltrami operators, from left to right.

In our experiments we were able to perform a quality matching with a relatively small number of eigenfunctions ($K = 5$ to 10).

- The multiplicity of the eigenvalues λ_k can be greater than one. The set of eigenfunctions corresponding to each such eigenvalue spans some sub-space of the eigenfunction domain. Matching such eigenfunctions calculated for two different shapes must include measuring distances between their spanned subspaces. Additionally, eigenfunctions belonging to close eigenvalues may switch their places, because of numerical errors.

- The eigenfunctions of the Laplace-Beltrami operators Δ_X and Δ_Y are defined up to a sign. Therefore, we have to find a sign sequence that relates between the corresponding eigenfunctions, such that the following holds for all corresponding pairs $(\mathbf{x}, \mathbf{y})$, $\mathbf{x} \in X, \mathbf{y} \in Y$:

$$\left\{\phi_1^X(\mathbf{x}), ..., \phi_K^X(\mathbf{x})\right\} = \left\{S_1\phi_1^Y(\mathbf{y}), ..., S_K\phi_K^Y(\mathbf{y})\right\}, \tag{4}$$

  where $S_k \in \{+, -\}$ is the k-th entry of the sign sequence. Following [22], we will refer to finding S as the sign ambiguity problem. The example given in Figure 1 shows two articulations of a human body, colored according to the values of the first three eigenfunctions of their Laplace-Beltrami operator. The correct sign sequence in this case is $[+, -, +]$.

All of the above problems were addressed before, and numerous attempts were made to overcome them. Shapiro and Brady [33], and later Jain and Zhang [12], suggested either exhaustive search or some greedy approach for the eigenfunction ordering and the sign sequence calculation, without accounting for the eigenvalue multiplicity problem. Umeyama [37] used absolute values of the eigenfunctions to solve the sign ambiguity problem. Mateus *et. al.* [22] suggested formulating the matching task as a global optimization problem, in which the eigenfunctions ordering and the sign sequence are part of the unknowns. In their later papers Mateus *et. al.* [23] and Knossow *et. al.* [15] suggested comparing histograms of the values of eigenfunctions to estimate the correct signs of the eigenvectors, prior to the alignment.

Here, we propose to solve the matching problem using only the eigenfunctions that correspond to eigenvalues without multiplicity. The distances between those eigenfunctions can be measured using the Euclidian norm. This considerably simplifies the matching task, and, according to our experiments, yields good results. We therefore restrict the eigenfunctions used in the descriptors to those that correspond to the first K non-repeating eigenvalues. In practice, we discard pairs of eigenfunctions if the difference between their corresponding eigenvalues is below some predefined threshold. This also accounts for possible eigenfunction flipping.

Finally, we have to estimate the sign sequence connecting the eigenvalues of the two operators Δ_X and Δ_Y. Before we introduce the estimation algorithm, we would like to note that there may be more than one sign sequence that align the two surfaces.

The human body in the example given in Figure 1 is intrinsically symmetric. Hence, there are two possible alignments of the two surfaces. The two sign sequences induced by those alignments are $[+, -, +]$ and $[+, +, -]$. Loosely speaking, the first sign sequence aligns the right side of the upper human body to the right side of the lower one. The second sequence aligns the right side of the upper human body to the left side of the lower one. We call the alignment produced by the first sequence the *primary alignment*, or *primary correspondence*, and the alignment produced by the second sequence *secondary*, or *symmetrical correspondence*.

Ovsjanikov *et al.* [27] showed that intrinsic symmetry of a surface induces reflection symmetry of the eigenfunctions of its Laplace-Belrami operator (to be precise, eigenfunctions corresponding to non-repeating eigenvalues). In other words, given two symmetric points $\mathbf{p}$ and $\mathbf{q}$ on M, one of the following relations holds for each ϕ_k^M corresponding to eigenvalue λ_k of unit multiplicity

$$\phi_k^M(\mathbf{p}) = \phi_k^M(\mathbf{q}) \quad \text{or} \quad \phi_k^M(\mathbf{p}) = -\phi_k^M(\mathbf{q}). \tag{5}$$

Thus, the sign sequences induced by the primary and the symmetrical correspondences might produce equally good alignments, in terms of distances between the descriptors.

In order to estimate the sign sequences induced by all possible alignments, we suggest a new algorithm based on the distances between the absolute values of the descriptors, as described in the next section.

3.3. Sign sequence estimation

The motivation for the proposed algorithm came from the following observation: if the shape has no intrinsic symmetries, we can use the absolute values of the descriptors (equivalently, eigenfunctions) for matching. Let us denote by X and Y the two surfaces that we wish to align. We say that a point $\mathbf{y} \in Y$ corresponds to the point $\mathbf{x} \in X$ if

$$\mathbf{y} = \underset{\forall \tilde{\mathbf{y}} \in Y}{\arg\min} \left\| \left| f^X(\mathbf{x}) \right| - \left| f^Y(\tilde{\mathbf{y}}) \right| \right\| \tag{6}$$

Or, $\mathbf{y} \in Y$ corresponds to $\mathbf{x} \in X$ if it is its nearest neighbor, in terms of distance between absolute values of their descriptors.

The correspondences obtained in such way are less accurate, compared to those obtained using the descriptors plus the correct sign sequence. We call them the approximate correspondences. They allow us to approximate the sign sequence that relates between the descriptors of the two shape, as follows

$$S_k = \underset{\{+,-\}}{\arg\min} \sum_{\mathcal{C}} \left\| f_k^X(\mathbf{x}) - S_k f_k^Y(\mathbf{y}) \right\|, \forall k, \tag{7}$$

where $\mathcal{C}$ is the space of all $(\mathbf{x}, \mathbf{y})$ correspondences, and S_k is the k-th entry of the sign sequence.

In case where the shape is intrinsically symmetric, each point $\mathbf{x} \in X$ may have several corresponding points on Y, with possibly equal matching costs. In practice, because of numerical errors, the correspondences obtained from Equation (6) include both the primary correspondence, and the symmetrical ones.

By dividing the correspondences into groups according to the symmetry they reflect, we can estimate the sign sequence induced by each group using Equation (7). For that goal, we suggest using the geodesic distances measured on the two surfaces. Let us define the group of all the primary correspondences by $\mathcal{C}_1$. The geodesic distances measured between the points on the first shape and between their correspondences on the second shape should be similar. This should hold for the symmetrical correspondences as well. Therefore, we can cluster the correspondences using the differences between geodesic distances measured between corresponding points on X and Y, and find the sign sequence induced by each group.

These observations can be formalized into an algorithm for the sign sequence estimation.

1. Initial correspondence detection: For each $\mathbf{x} \in X$ find its corresponding point $\mathbf{y}$ on Y using

$$\mathbf{y} = \underset{\forall \tilde{\mathbf{y}} \in Y}{\arg\min} \left\| \left| f^X(\mathbf{x}) \right| - \left| f^Y(\tilde{\mathbf{y}}) \right| \right\|. \tag{8}$$

   To make the algorithm more robust, it is possible to choose more than one nearest neighbor on Y. In our experiments we used only one corresponding candidate for each point.

2. Clustering: calculate the affinity matrix A of pairs of the corresponding points $(\mathbf{x}_m, \mathbf{y}_m), (\mathbf{x}_n, \mathbf{y}_n) \in \mathcal{C}$, in the following way

$$A_{mn} = \left| d_X(\mathbf{x}_m, \mathbf{x}_n) - d_Y(\mathbf{y}_m, \mathbf{y}_n) \right|. \tag{9}$$

   Then, use a clustering algorithm to divide the correspondences into several clusters, based on A. We denote the correspondence clusters by $\{\mathcal{C}_j\}_{j=1}^J$, where J is the number of clusters. Each cluster $\mathcal{C}_j$ induces a sign sequence. The number of the clusters J has to be sufficiently large so that the correct sign sequences would be represented by multiple clusters. Typically, we would divide a set of 1000 initial correspondences into 50 to 100 clusters.

   In our experiments we performed the clustering using the K-means algorithm [21]. It requires a set of vectors, which we obtained from the affinity matrix A using the multidimesional scaling algorithm [3]. An example of two correspondence clusters, reflecting two possible alignments, is presented in Figure 2.

3. Sign sequences estimation: For each cluster $\mathcal{C}_j$, find the sign sequence $S^{(j)}$ it induces, using

$$S_k^{(j)} = \underset{\{+,-\}}{\arg\min} \sum_{\mathcal{C}_j} \left\| f_k^X(\mathbf{x}) - S_k^{(j)} f_k^Y(\mathbf{y}) \right\|, \tag{10}$$

   where $S_k^{(j)}$ is the k-th entry of sign sequence $S^{(j)}$.

   The set $\left\{S^{(j)}\right\}_{j=1}^J$ includes multiple instances of the correct sign sequences. Therefore, we construct a reduced set of the candidate sign sequences using the sequences $S^{(j)}$ that were induced by the largest number of clusters. The size of this reduced set equals to the number of the intrinsic symmetries of the shape, plus 1. In the example showed in Figure 2 there are two such sequences.

   We would like to note that a related approach to the sign sequence estimation was proposed by Ovsjanikov *et al.* in [27], with regard to intrinsic symmetry detection, and by Mateus *et al.* [23] and Knossow *et. al. et al.*, with regard to surface matching. Our approach differs from [15, 23] by the fact that we aim to find the sign sequences induced by all possible alignments of the two surfaces. In this sense, it reminds the intrinsic symmetry detection of [27]. Indeed, the proposed sign sequence estimation algorithm, combined with the framework in [27], can be used for intrinsic symmetry detection.

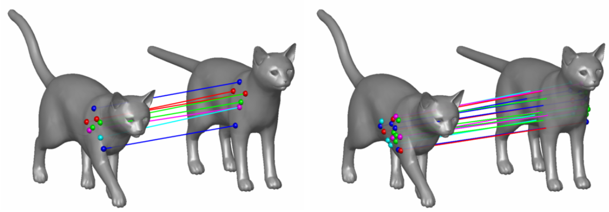

Figure 2. Correspondence clustering example.

Once we obtained the candidate sign sequences, we can use them to calculate the correspondences between the surfaces, using the matching algorithm described in the next section. To simplify notations, we re-define the descriptor f^Y to include the sign sequence information

$$f^Y(\mathbf{y}_j) = \{S_1\phi_1^Y(j), S_2\phi_2^Y(j), ..., S_3\phi_K^Y(j)\}, \quad (11)$$

where S is a sign sequence from the candidate set.

We would like to mention that another possible approach to correspondence type detection (primary or secondary) would be to explore the relative surface orientation changes implied by different alignments. We do not elaborate further on this approach in the current paper, yet we think it is a possible direction for future research.

4. Matching algorithm

We formulate the matching task as a minimization problem. Its objective function consists of a linear cost term, based on the isometry invariant descriptors described earlier, and a following geodesic distance preserving regularization term.

We denote by X and Y the two surfaces we wish to align. We express the correspondence between points on X and Y by a mapping $P : X \times Y \to \{0, 1\}$:

$$P(\mathbf{x}, \mathbf{y}) = \begin{cases} 1, & \mathbf{x} \text{ and } \mathbf{y} \text{ are in correspondence,} \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

When X and Y are finite, P is a matrix with binary entries. In this case, we can write $P_{ij} = P(\mathbf{x}_i, \mathbf{y}_j)$, for some $\mathbf{x}_i \in X$ and $\mathbf{y}_j \in Y$. The optimal correspondence, denoted by P^*, minimizes some cost function $C(P)$ over the space of all mappings P.

The following simple cost function is based on the isometric invariant descriptors introduced in the previous section

$$C_l(P) = \sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{y}_j \in Y}} P_{ij} b_{ij}, \quad (13)$$

where $b_{ij} = \left\| f^X(\mathbf{x}_i) - f^Y(\mathbf{y}_j) \right\|$. We construct a different cost function for each possible alignment, using the descriptors defined by Equation (11).

Note that the cost $C_l(P)$ is linear in P, and contains no information about how well pairs of correspondences in P coincide. Therefore, we regularize it with the following quadratic term

$$C_q(P) = \sum_{\substack{\mathbf{x}_i, \mathbf{x}_m \in X \\ \mathbf{y}_j, \mathbf{y}_n \in Y}} P_{ij} P_{mn} Q_{ijmn}, \quad (14)$$

where $Q_{ijmn} = |d_X(\mathbf{x}_i, \mathbf{x}_m) - d_Y(\mathbf{y}_j, \mathbf{y}_n)|$, and d_X and d_Y denote geodesic distances measured on X and Y, respectively. The motivation behind the definition of C_q is similar to that mentioned in the previous section; that is, if $(\mathbf{x}_i, \mathbf{y}_j)$ and $(\mathbf{x}_m, \mathbf{y}_n)$ are two correspondences, the geodesic distances $d_X(\mathbf{x}_i, \mathbf{x}_m)$ and $d_Y(\mathbf{y}_j, \mathbf{y}_n)$ must be similar.

The total correspondence cost is given by

$$C(P) = \sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{y}_j \in Y}} P_{ij} b_{ij} + \lambda \cdot \sum_{\substack{\mathbf{x}_i, \mathbf{x}_m \in X \\ \mathbf{y}_j, \mathbf{y}_n \in Y}} P_{ij} P_{mn} Q_{ijmn} \quad (15)$$

The parameter λ determines the weight of the quadratic term in the total cost function. The quadratic term alone does not allow to distinguish between primary and symmetric correspondences. Therefore, for $\lambda \gg 1$, for each set of descriptors defined in the previous section we will obtain the same optimal solution P^* that minimizes the cost $C_q(P)$. In our experiments, in order to find all correspondences, we used λ in the range of $[0.1, 0.5]$.

In order to avoid a trivial solution when minimizing $C(P)$ we must impose certain constraints on the matrix P. Those constraints depend on the type of the required correspondence. If X and Y consist of the same number of corresponding points, and the correspondence is required to be one-to-one, the constraints are given by

$$\sum_i P_{ij} = 1, \quad \sum_j P_{ij} = 1, \quad \forall i, j \quad (16)$$

Such P is called a permutation matrix.

The one-to-one requirement may be too strict when matching two sampled surfaces, or when the numbers of the points sampled from the two surfaces are different. In this case, we suggest fixing the points on one surface, say X, and demand that each $\mathbf{x} \in X$ has a correspondence on Y. Or, in terms of P,

$$\sum_i P_{ij} = 1, \quad \forall i \quad (17)$$

Note that we can re-write the problem in a matrix form, by converting the double correspondence index (i, j), as in $(\mathbf{x}_i, \mathbf{y}_j)$, to a single index k. Consequently, the correspondence between X and Y is represented by a vector $\{P_k\}$. The matrix $\{b_{ij}\}$ and the tensor $\{Q_{ijmn}\}$ are converted into a vector and a matrix $\{b_k\}$ and $\{Q_{kl}\}$, accordingly. The objective function and the constraints can be re-written as follows

$$P^* = \min_{\forall P} \left\{ b^T P + \lambda P^T Q P \right\} \quad \text{s.t.} \quad SP = \mathbb{1}, \qquad (18)$$

where $P_k \in \{0, 1\}$, $\forall k$. The matrix S is sparse and it is constructed according to the constraints being used (Equations (16) or (17), or different constraints); $\mathbb{1}$ is a vector with all entries equal to 1, of an appropriate size. In our experiments we normalized the vector b and the matrix Q by their maximal values, to obtain meaningful contributions of the two terms to the cost function.

Equation (18) describes a combinatorial problem called an Integer Quadratic Programming problem (IQP). Similar to [10], we solve it using an MIQP solver, which is distributed as a part of the Hybrid Toolbox by Bemporad *et al.* [2].

5. Results

We tested the proposed method on pairs of shapes represented by triangulated meshes, from the TOSCA high-resolution database [5]. All the models were downsampled to about 4500 vertices ($+/-10\%$).

All calculations were performed using the MATLAB© software. To calculate the descretized Laplace-Beltrami operator we used the cotangent-weights scheme proposed in [26]. For the sign sequence estimation we used sets of 1000 points sub-sampled from each model using Farthest Point Sampling [9]. For the correspondence detection we used a relatively small number of points (20 points) sampled from one of the models using [9], and a larger set of candidate corresponding points sampled from the second model (40 points), thus obtaining 20 pairs of corresponding points. This sparse set may later serve as an initialization for other shape processing algorithms, *e.g.* the GMDS [4], or the dense cross-parameterization algorithm [16].

Correspondence detection: We applied the proposed method for correspondence detection to several pairs of shapes, with and without intrinsic symmetries. The results are presented in Figures 3, 4, 5, and are quite accurate, in most cases. The correspondences detected for the cat and the wolf shapes, presented in Figure 3, (b) and (c), respectively, exhibit limitations of the proposed method. The correspondences between the points on the heads of the cats and the wolfs are not accurate enough. This is due to the combinatorial nature of the problem, that limits the number of the matches examined and detected by our method, and may also result in some locally minimal solution.

Symmetry detection: We applied the proposed eigenfunction sign estimation algorithm for intrinsic symmetry detection. Figure 6, captures the of intrinsic symmetry detected for a human shape, using the nearest neighbor search

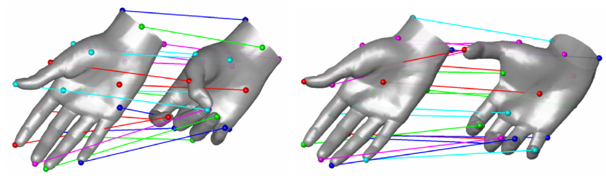

Figure 5. Correspondences obtained for two pairs of hand models.

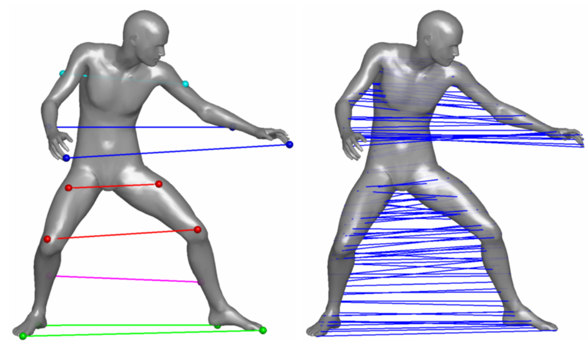

(a) (b)

Figure 6. Intrinsic symmetry detected in the human shape, with (a) and without (b) geodesic distance preservation term.

(Figure 6, (b)) as proposed in [27], and using the proposed method (Figure 6, (a)).

6. Conclusions

We presented a framework for detecting correspondence between non-rigid, approximately isometric shapes, which relays on surface descriptors calculated using the eigenfunctions of the Laplace-Beltrami operator. We showed that when matching two sets of such descriptors, one encounters a problem of ambiguity in the signs of the eigenfunctions. Our key observation is that the above ambiguity can be exploited to detect different possible alignments of the two shapes. This happens when the matched shapes are intrinsically symmetric. We demonstrated how to obtain those alignments using Integer Quadratic Programming.

Possible limitations: at its current form, partial matching poses a challenge to the proposed framework. The same applies to matching shapes with topological changes or significant changes of the geodesic distances. We assume that by combining the proposed framework with local surface descriptors and, possibly, other distance measures, we should be able to overcome these limitations. It is also possible to combine the proposed methods with salient feature detection methods, for robustness and better matching precision. Finally, the complexity of the algorithm is currently determined by the integer programming solver used to find the minimum of the quadratic problem. The overall complexity is therefore exponential in the number of the possible cor-

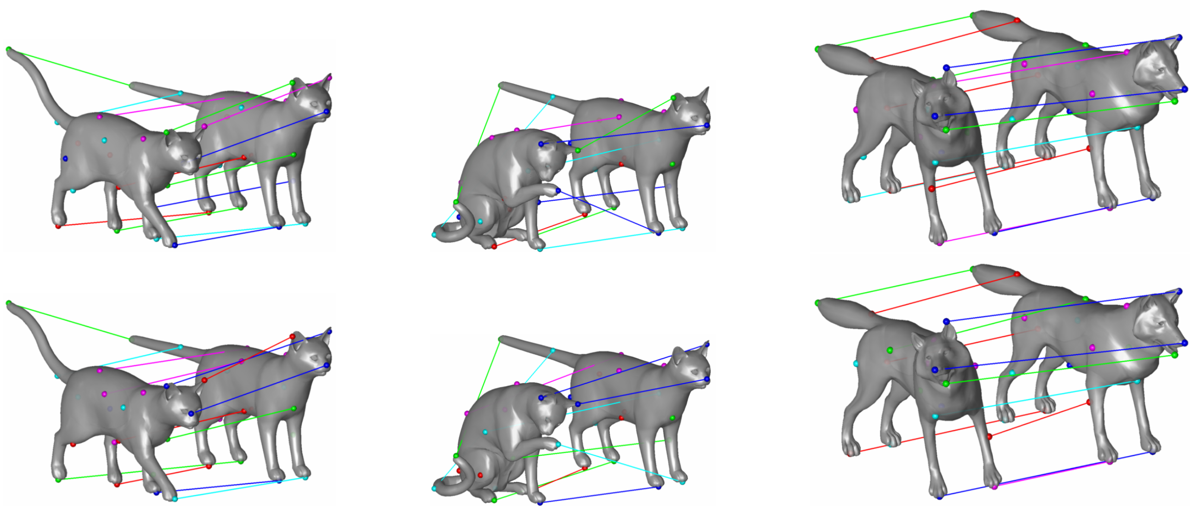

(a) (b) (c)

Figure 3. Primary (first row) and the symmetrical (second row) correspondences obtained for pairs of different models using the proposed method.

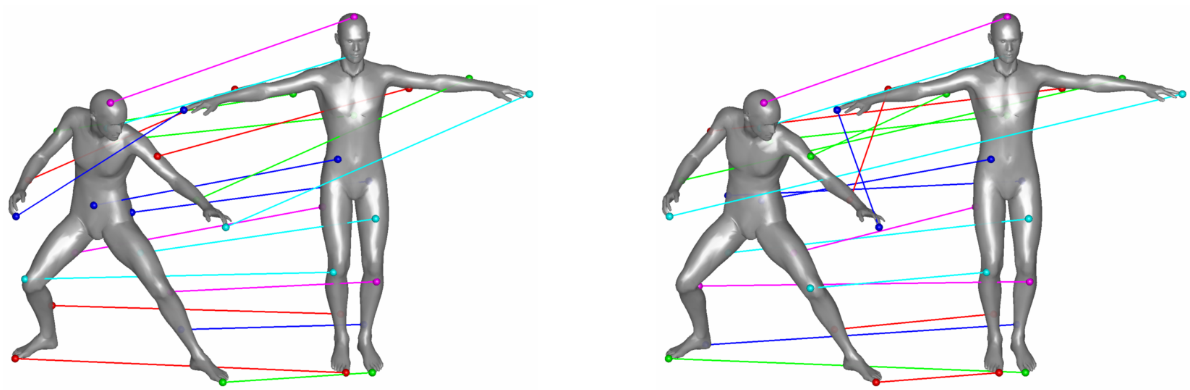

Figure 4. Primary (left) and symmetrical (right) correspondence between two human shapes, obtained with the proposed method.

respondences. By employing other optimization techniques we could possibly reduce the computation time and find the correspondence between larger sets of points.

Acknowledgments

This research was partially supported by ONR and ISF grants. We also wish to thank Drs. Alexander and Michael Bronstein for helpful discussions during the progress of this work.

References

[1] D. Anguelov, P. Srinivasan, H.-C. Pang, and D. K. S. Thrun. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Proceedings of the Neural Information Processing Systems (NIPS) Conference*. 2

[2] A. Bemporad. Hybrid Toolbox - User's Guide, 2004. http://www.dii.unisi.it/hybrid/toolbox. 6

[3] I. Borg and P. J. F. Groenen. *Modern Multidimensional Scaling: Theory and Applications (Springer Series in Statistics)*. Springer, Berlin, 2nd ed. edition, 2005. 4

[4] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. In *Proc. of National Academy of Sciences (PNAS)*, pages 1168–1172, 2006. 2, 6

[5] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer, 2008. 6

[6] W. Chang and M. Zwicker. Automatic registration for articulated shapes. *Computer Graphics Forum (Proceedings of SGP 2008)*, 27(5):1459–1468, 2008. 2

[7] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE Trans. PAMI*, 25(10):1285–1295, 2003. 1

[8] M. Gromov. *Structures metriques pour les varietes riemanniennes*. Textes Mathematiques, no. 1, 1981. 2

[9] D. S. Hochbaum and D. B. Shmoys. A best possible heuristic for the k-center problem. *Mathematics of Operations Research*, 10(2):180–184, 1985. 6

[10] J. Hu and J. Hua. Salient spectral geometric features for shape matching and retrieval. *Vis. Comput.*, 25(5-7):667–675, 2009. 2, 6

[11] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Comput. Graph. Forum*, 27(5):1449–1457, 2008. 2

[12] V. Jain, H. Zhang, and O. V. Kaick. Non-rigid spectral correspondence of triangle meshes. In *International Journal on Shape Modeling*, volume 13, pages 101–124, 2007. 1, 3

[13] D. Jakobson, N. Nadirashvili, and J. Toth. Geometric properties of eigenfunctions. *Russian Math. Surveys*, 56:67–88, 2001. 2

[14] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In *SGP '03: Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 156–164, 2003. 1

[15] D. Knossow, A. Sharma, D. Mateus, and R. P. Horaud. Inexact matching of large and sparse graphs using laplacian eigenvectors. In *Proceedings 7th Workshop on Graph-based Representations in Pattern Recognition*, 2009. 1, 3, 4

[16] V. Kraevoy and A. Sheffer. Cross-parameterization and compatible remeshing of 3d models. *ACM Transactions on Graphics*, 23(3):861 – 869, Aug. 2004. 6

[17] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *International Conference of Computer Vision (ICCV)*, volume 2, pages 1482 – 1489, 2005. 2

[18] B. Levy. Laplace-beltrami eigenfunctions: Towards an algorithm that understands geometry. In *IEEE International Conference on Shape Modeling and Applications, invited talk*, 2006. 1

[19] B. Levy and R. H. Zhang. Spectral geometry processing. In *ACM SIGGRAPH ASIA Course Notes*, 2009. 1

[20] Y. Lipman and T. Funkhouser. Mobius voting for surface correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 28(3), Aug. 2009. 1

[21] J. B. Macqueen. Some methods of classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297, 1967. 4

[22] D. Mateus, F. Cuzzolin, R. P. Horaud, and E. Boyer. Articulated shape matching using locally linear embedding and orthogonal alignment. In *IEEE Workshop on Non-rigid Registration and Tracking through Learning - NRTL 2007*, 2007. 3

[23] D. Mateus, R. P. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2008. 1, 3, 4

[24] F. Memoli. On the use of gromov-hausdorff distances for shape comparison. In *Point Based Graphics 2007*, pages 81–90, Sept. 2007. 2

[25] F. Memoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5(3):313–347, 2005. 2

[26] M. Meyer, M. Desbrun, S. P., and A. Barr. Discrete differential geometry operators for triangulated 2-manifolds. In *International Workshop on Visualization and Mathematics*, 2002. 6

[27] M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. *Computer Graphics Forum*, 27(5):1341–1348, July 2008. 2, 3, 4, 6

[28] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Symmetries of non-rigid shapes. In *Proc. Workshop on Non-rigid Registration and Tracking through Learning (NRTL)*, 2005. 2

[29] M. Reuter. *Laplace Spectra for Shape Recognition*. Books on Demand GmbH, 2006. 1

[30] S. Rosenberg. *The Laplacian on a Riemannian manifold: an introduction to analysis on manifolds*. Cambridge University Press, 1997. 2

[31] S. Rusinkiewicz and M. Levoy. Efficient variants of the icp algorithm. In *Proc. Int. Conf. Recent Advances 3-D Digital Imaging Modeling*, pages 145–152, 2001. 1

[32] R. M. Rustamov. Laplace-beltrami eigenfunctions for deformation inavriant shape representation. In *Proceedings of SGP*, pages 225–233, 2007. 1, 2

[33] L. S. Shapiro and J. M. Brady. Feature-based correspondence: an eigenvector approach. *Image Vision Comput.*, 10(5):283–288, 1992. 3

[34] J. Tangelder and R. Veltkamp. A survey of content-based 3d shape retrieval methods. In *SMI '04: Proceedings of the Shape Modeling International 2004*, pages 145–156, 2004. 1

[35] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. Isometric registration of ambiguous and partial data. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009)*. 2

[36] N. Thorstensen and R. Keriven. Non-rigid shape matching using geometry and photometry. In *Asian Conference on Computer Vision*, 2009. 2

[37] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(5):695–703, 1988. 3

[38] A. Zaharescu, E. Boyer, K. Varanasi, and R. P. Horaud. Surface feature detection and description with applications to mesh matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 2

[39] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. van Kaick, and A. Tagliasacchi. Deformation-driven shape correspondence. *Computer Graphics Forum (Proc. of SGP 2008)*, 27(5):1431–1439, 2008. 2