

Non-rigid shape correspondence using pointwise surface descriptors and metric structures

Anastasia Dubrovina, Dan Raviv and Ron Kimmel

Abstract Finding a correspondence between two non-rigid shapes is one of the cornerstone problems in the field of three-dimensional shape processing. We describe a framework for marker-less non-rigid shape correspondence, based on matching intrinsic invariant surface descriptors, and the metric structures of the shapes. The matching task is formulated as a quadratic optimization problem that can be used with any type of descriptors and metric. We minimize it using a hierarchical matching algorithm, to obtain a set of accurate correspondences. Further, we present the correspondence ambiguity problem arising when matching intrinsically symmetric shapes using only intrinsic surface properties. We show 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 a number of possible correspondences between two given shapes.

1 Introduction

Three-dimensional shape processing became increasingly popular in the last decade. One of its corner-stone tasks is detecting a correspondence between two given shapes. It is essential for shape comparison, retrieval, shape morphing and deformation, or shape calculus [5], etc. The most interesting yet complex task is automatic

Anastasia Dubrovina
Technion, Israel Institute of Technology, e-mail: nastyad@cs.technion.ac.il

Dan Raviv
Technion, Israel Institute of Technology, e-mail: darav@cs.technion.ac.il

Ron Kimmel
Technion, Israel Institute of Technology e-mail: ron@cs.technion.ac.il

non-rigid shape matching. In this work we address the problem of matching non-rigid approximately isometric shapes. We perform the matching using certain surface properties that remain invariant under isometric transformations. In particular, we use two types of such properties - pointwise surface descriptors, and distances measured between pairs of points on the surface. We show how these two properties can be incorporated into a measure of dissimilarity between the shapes, which can be written as a quadratic function of the correspondence. We then minimize this dissimilarity measure in order to find the minimal dissimilarity correspondence.

Another important issue we address here is the correspondence ambiguity present when matching intrinsically symmetric shapes. In this case, there may exist several correspondences minimizing the proposed dissimilarity measure. We show that this ambiguity can be resolved by constructing distinct sets of symmetry-aware surface descriptors. By employing them within the proposed framework it is possible to find several matchings between the shapes.

The rest of the paper is organized as follows: in the next section we review the related work on matching non-rigid shapes. In Section 3 we describe the proposed problem formulation. In Section 4 we describe the possible choices of metric and descriptors. In Section 5 we describe the correspondence ambiguity problem and the construction of the symmetry-aware surface descriptors. In Section 6 we present the matching results obtained with the proposed framework combined with different descriptors and distances measures. We summarize the paper and discuss future research directions in Section 7.

2 Related work

Zigelman *et al.* [43], and Elad and Kimmel [9] suggested a method for matching isometric shapes by embedding them into a Euclidian space using multidimensional scaling (MDS), thus obtaining isometry invariant representations, followed by rigid shape matching in that space. Since it is generally impossible to embed a non-flat 2D manifold into a flat Euclidean domain without introducing some errors, the inherited embedding error affects the matching accuracy of all methods of this type. For that end, Jain *et al.* [13], Mateus *et al.* [19] and Sharma and Horaud [34] suggested alternative isometry-invariant shape representations, obtained by using eigendecomposition of discrete Laplace operators. The Global Point Signature (GPS) suggested by Rustamov [33] for shape comparison employs the discrete Laplace-Beltrami operator, which, at least theoretically, captures the shape's geometry more faithfully. The Laplace-Beltrami operator was later employed by Sun *et al.* [35], and Ovsjanikov *et al.* [24], to construct their Heat Kernel Signature (HKS) and Heat Kernel Maps, respectively. Zaharescu *et al.* [41] suggested an extension of 2D descriptors for surfaces, and used them to perform the matching. While linear methods, such as [41, 24] produce good results, once distortions start to appear, ambiguity increases, and alternative formulations should be thought of. Adding the proposed approach as a first step in one of the above linear dense matching algorithms can improve the fi-

nal results. Hu and Hua [12] used the Laplace-Beltrami operator for matching using prominent features, and Dubrovina and Kimmel [8] suggested employing surface descriptors based on its eigendecomposition, combined with geodesic distances, in a quadratic optimization formulation of the matching problem. The above methods, incorporating pairwise constraints, tend to be slow due to high computational complexity. Wang *et al.* [40] used a similar problem formulation, casted as a graph labeling problem, and experimented with different surface descriptors and metrics.

Memoli and Sapiro [22], Bronstein *et al.* [4], and Memoli [20, 21] compared shapes using different approximations of the Gromov-Hausdorff distance [10]. Bronstein *et al.* [6] used the approach suggested in [4] with diffusion geometry, in order to match shapes with topological noise, and Thorstensen and Keriven [37] extended it to handle surfaces with textures. The methods in [22, 20, 21] were intended for surface comparison rather than matching, and as such they do not produce correspondence between shapes. At the other end, the GMDS algorithm [6] results in a non-convex optimization problem, therefore it requires good initializations in order to obtain meaningful solutions, and can be used as a refinement step for most other shape matching algorithms. Other algorithms employing geodesic distances to perform the matching were suggested by Anguelov *et al.* [1], who optimized a joint probabilistic model over the set of all possible correspondences to obtain a sparse set of corresponding points, and by Tevs *et al.* [36] who proposed a randomized algorithm for matching feature points based on geodesic distances between them. Zhang *et al.* [42] performed the matching using extremal curvature feature points and a combinatorial tree traversal algorithm, but its high complexity allowed them to match only a small number of points. Lipman and Funkhouser [18] used the fact that isometric transformation between two shapes is equivalent to a Möbius transformation between their conformal mappings, and obtained this transformation by comparing the respective conformal factors. However, there is no guarantee that this result minimizes the difference between pairwise geodesic distances of matched points.

Self-similarity and symmetry detection are particular cases of the correspondence detection problem. Instead of detecting the non-rigid mapping between two shapes, [28, 25, 17, 14] search for a mapping from the shape to itself, and thus are able to detect intrinsic symmetries.

3 Matching problem formulation

The suggested problem formulation is based on comparison of shape properties that remain approximately invariant under non-rigid ϵ -isometric transformations, specifically - distances between the points on the shape, and pointwise surface descriptors defined at every point of the shape. We assume to be given shapes represented by sampled surfaces, which is one of the common 3D shape representations. In this work shapes were represented by triangular meshes, but the following discussion is not limited to some specific sampled surface representation.

Given a shape X , we assume that it is endowed with a distance measure $d_X : X \times X \rightarrow \mathbb{R}^+ \cup \{0\}$, and a set of pointwise d -dimensional surface descriptors $f_X : X \rightarrow \mathbb{R}^d$. Given two shapes X and Y , we define a correspondence between them by a mapping $P : X \times Y \rightarrow \{0, 1\}$, such that

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

We can measure the dissimilarity introduced by the mapping P into the surface descriptors and the metric structures by

$$\text{dis}(P) = \sum_{\substack{x \in X \\ y \in Y}} \|f_X(x) - f_Y(y)\|_F P(x, y) + \alpha \cdot \sum_{\substack{x, \tilde{x} \in X \\ y, \tilde{y} \in Y}} |d_X(x, \tilde{x}) - d_Y(y, \tilde{y})| P(x, y) P(\tilde{x}, \tilde{y}), \quad (2)$$

where $\|\cdot\|_F$ is a norm in the descriptor space. The first term of the dissimilarity measure is a linear function of the mapping P , and it expresses the pointwise surface descriptor dissimilarity. The second term of $\text{dis}(P)$ is a quadratic function of the mapping P , and it expresses the metric structure dissimilarity. The parameter $\alpha \geq 0$ determines the relative weight of the second term in the total dissimilarity measure.

Note that by setting $\alpha = 0$ we obtain the linear matching method used by [41, 24]. When the descriptors of different points on the shape are not unique or sufficiently distinct, say due to numerical inaccuracies, the correspondences obtained by minimizing only the linear part of the dissimilarity measure may not be consistent in terms of pairwise relationships between the matched points. By adding the quadratic term in Eq. (2) we ensure that the optimal correspondence preserves also these pairwise relationships. On the other hand, by choosing $\alpha \gg 1$ we obtain a problem similar to the one addressed in [22, 4, 20], and, since the problem is non-convex, a good initialization is required in order to obtain a close-to-global minimizer. This is achieved by adding the linear term as in Eq. (2).

The optimal matching, which we denote by P^* , is obtained by minimizing the dissimilarity

$$P^* = \operatorname{argmin}_{P: X \times Y \rightarrow \{0, 1\}} \{\text{dis}(P)\}. \quad (3)$$

In order to avoid a trivial solution $P^* \equiv 0$, we constrain P to the space of valid correspondences. Note that the above problem formulation allows us to consider different types of possible correspondences between the two shapes. For example, when a bijective mapping from X to Y is required, the constraints on P are

$$\sum_{x \in X} P(x, y) = 1, \forall y \in Y, \quad \sum_{y \in Y} P(x, y) = 1, \forall x \in X. \quad (4)$$

For a surjective mapping we relax the constraints to be

$$\sum_{x \in X} P(x, y) = 1, \forall y \in Y. \quad (5)$$

Thus, the resulting optimization problem is

$$P^* = \operatorname{argmin}_{P: X \times Y \rightarrow \{0,1\}} \{\operatorname{dis}(P)\} \quad \text{s.t. suitable constrains on } P. \quad (6)$$

3.1 Quadratic programming formulation

When the two shapes X and Y are represented by a finite number of points N , the mapping P is a binary matrix of size N^2 . In order to convert the problem into a more convenient quadratic programming form we reshape the matrix P by taking its columns and concatenating them, thus obtaining a vector p of dimension N^2 , referred to as a *correspondence vector*. Thus, k -th entry of the vector p corresponds to some entry (i, j) in the matrix P - we will denote these corresponding indices by (i_k, j_k) . The vector entry p_k encodes the correspondence between the points x_{i_k} and y_{j_k} .

Similarly, we introduce the following notations for the metric and the descriptor dissimilarity

$$b_k = \|f_X(x_{i_k}) - f_Y(y_{j_k})\|_F, \quad Q_{kl} = \|d_X(x_{i_k}, x_{m_k}) - d_Y(y_{j_l}, y_{n_l})\|. \quad (7)$$

The vector $b \in \mathbb{R}^{N^2}$ represents the pointwise descriptor dissimilarity, and the matrix $Q \in \mathbb{R}^{N^2 \times N^2}$ represents the metric structure dissimilarity.

Lastly, we re-write optimization problem (6) in the quadratic programming form

$$p^* = \operatorname{argmin}_{p \in \{0,1\}^{N^2}} \{b^T p + \lambda \cdot p^T Q p\} \quad \text{s.t. } Sp = \mathbf{1}, \quad (8)$$

where $Sp = \mathbf{1}$ is the matrix form of the constraints in Eq. (4) or Eq. (5).

3.2 Hierarchical matching

The optimization problem in Eq. (8) belongs to the class of NP -hard Integer Quadratic Programming (IQP) problems. There exist different techniques for approximating its solution, [2, 38] among them, which are able to solve only moderate size IQPs. The implication on the matching problem is that the algorithm will be able to find only small number of correspondences - up to several tens of points. In addition, prior to the matching the algorithm has to choose the initial set of N candidate points on each one of the shapes X and Y . The simplest way to choose these points is by using the Farthest Point Sampling technique [11], and the sampling density will determine the accuracy of the matching.

In order to overcome these limitations we use the hierarchical matching technique introduced in [29]. It exploits the shapes' geometric structures to reduce the number of potential correspondences, and thus is able to find a denser matching,




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

with improved accuracy. Since the problem is not strictly combinatorial by nature, but rather derived from a smooth geometric measure. At the first step we follow [8] and solve (6) using a branch-and-bound procedure [2]. Each point $x \in X$ is now matched to a point $p(x) \in Y$ by the mapping P . We denote $y = p(x)$ if $P(x, y) = 1$. In each iteration we search for the best correspondence between x and $p(x)$ neighborhood, instead of all points $y \in Y$, in a manner similar to [39]. Between iterations we add points $x \in X$ and $y \in Y$ using the 2-optimal *Farthest Point Sampling* (FPS) strategy [11], evaluate the neighborhood in Y of the new points, reevaluate the neighborhood of the old points, and continue until convergence. In Figure 1 we show a diagram of the process.

We solve the relaxed version of (6), using quazi-Newton optimization, and project the solution to integers between iterations. Convergence is guaranteed, but only to a local minimum, as for all QAP problems.

A different approach for approximating the solution of the IQP in Eq. (8) can be, for instance, using the relaxation technique of Bronstein *et al.* [4] and solving the problem on a continuous domain. The optimization problem can also be solved using the approach for graph matching by Torresani *et al.* [38]. Both can reduce the complexity of the solution. We will explore these directions in the future research.

4 On the choice of metric and descriptors

The above formulation of the matching problem can be used with any type of surface descriptors or distance measure. Below we describe different descriptors and metrics that can be employed in the proposed framework. We start with a brief review of the Laplace-Beltrami operator, and later use concepts related to it for both metric and descriptor definition. Note that both metric definitions and some of the descriptor definitions are given in terms of continuous surface representation (or 2D Riemannian manifolds). For each one of them we state the discrete approximation we used for numerical evaluation.

4.1 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

$$\Delta_X f = -\operatorname{div}_X (\nabla_X f), \quad f : X \rightarrow \mathbb{R}. \quad (9)$$

The divergence and the gradient operators, div_X and ∇_X respectively, are defined by the intrinsic geometry of the manifold X . Explicitly, the Laplace-Beltrami operator of a function $f : X \rightarrow \mathbb{R}$ defined on the manifold X equipped with a Riemannian metric g is given by

$$\Delta_X f = -\frac{1}{\sqrt{\det g}} \sum_{j,k} \frac{\partial}{\partial x_j} \left(g^{jk} \sqrt{\det g} \frac{\partial f}{\partial x_k} \right). \quad (10)$$

In the above equation, $\det g = \det(g_{ij})$ and the g^{jk} are the elements of g^{-1} . For more details see [30].

Consider the Laplace-Beltrami operator eigenvalue problem given by

$$\Delta_X \phi_i = \lambda_i \phi_i. \quad (11)$$

$\{\phi_i\}$ are the eigenfunctions of Δ_X , corresponding to the eigenvalues $\{\lambda_i\}$. The spectrum of the Laplace-Beltrami operator consists of positive eigenvalues (see, for example, [30]). When X is a connected manifold without boundary, then Δ_X has additional eigenvalue equal to zero, with corresponding constant eigenfunction. We can order the eigenvalues as follows

$$0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots \quad (12)$$

The set of corresponding eigenfunctions given by

$$\{\phi_1, \phi_2, \phi_3, \dots\} \quad (13)$$

forms an orthonormal basis defined on X with inner product induced by the metric g .

There exist various approximations for the Laplace-Beltrami operator. In this work we used the cotangent weight scheme [26, 23].

4.2 Choice of metric

Geodesic distance: The simplest *intrinsic* metric defined on a surface X is the *geodesic metric*. It measures the lengths of the shortest paths on the surface X

$$d_X(x, x') = \inf_{\gamma \in \Gamma(x, x')} \ell(\gamma). \quad (14)$$

$\Gamma(x, x')$ is the set of all admissible paths between the points x and x' on the surface X , with a length of a path γ given by $\ell(\gamma)$. In order to calculate the geodesic distances we used the *fast marching method* [15], which simulates a wavefront propagation on a triangular mesh, and associates the front arrival time with the distance traveled by it.

Diffusion geometry: The diffusion of heat on surface X is governed by the heat equation,

$$\left(\Delta_X + \frac{\partial}{\partial t} \right) u(x, t) = 0, \quad (15)$$

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

The *heat kernel* $h_t(x, z)$ describes the amount of heat transferred from a point heat source located at x to another point z at time t , and can be written as

$$h_t(x, z) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x) \phi_i(z). \quad (16)$$

The *diffusion distance* can then be defined as a cross-talk between two heat kernels [3, 7]

$$\begin{aligned} d_{X,t}^2(x, y) &= \|h_t(x, \cdot) - h_t(y, \cdot)\|_{L_2(X)}^2 \\ &= \int_X |h_t(x, z) - h_t(y, z)|^2 dz \\ &= \sum_{i=0}^{\infty} e^{-2\lambda_i t} (\phi_i(x) - \phi_i(y))^2. \end{aligned} \quad (17)$$

Since the heat flow on the surface is governed entirely by its intrinsic geometry, the diffusion distance defined above is an intrinsic property of the surface, and, according to [3, 7], also fulfills the metric axioms.

We approximate the diffusion distances using a finite number of the eigenvalues and the eigenvectors of the discretized Laplace-Beltrami operator. Specifically, we used several hundred eigenvalues with the smallest magnitude and their corresponding eigenvectors.

4.3 Choice of descriptors

Distance histograms: Given a shape X and the corresponding distance d_X , the distance histogram descriptor [27, 32, 28] is constructed as follows

$$f_X(x) = \text{hist} \{d_X(x, \tilde{x}) \mid d_X(x, \tilde{x}) \leq d_{max}, \tilde{x} \in X\}, \quad (18)$$

where d_{max} controls the local support of the descriptor. If two shapes are represented by differently sampled surfaces, the descriptors can be normalized to have L_1 -norm equal to one. The descriptor comparison can be performed using either an L_p -norm, or some measure of distances between histograms, such as the *earth moving distances* (EMD) [31].

Heat kernel signatures and heat kernel maps: Local descriptors based on the *heat equation* were presented by Sun *et al.* in [35] and Ovsjanikov *et al.* in [24]. The *heat kernel signature* (HKS) is constructed using the diagonal of the heat kernel $h_t(x, x)$ (16) at multiple times t

$$f_X(x) = [h_{t_1}(x, x), h_{t_2}(x, x), \dots, h_{t_d}(x, x)]. \quad (19)$$

The *heat kernel map* (HKM) is constructed using the heat kernel values with a pre-specified heat source x_0

$$f_X(x) = [h_{t_1}(x_0, x), h_{t_2}(x_0, x), \dots, h_{t_d}(x_0, x)]. \quad (20)$$

For the latter descriptors, the heat sources chosen for the two shapes we want to match must be in correspondence in order to produce consistent descriptors. One can choose the heat source x_0 either as proposed by the authors of [24], or by some different method. Both HKS and HKM remains invariant under isometric deformations of X , and are insensitive to topological noise at small scales.

To compute HKS and HKM we used eigenvalues and eigenfunction of the discretized Laplace-Beltrami operator, similar to the diffusion distance calculation.

5 Matching ambiguity problem

The matching ambiguity problem arises when matching intrinsically symmetric shapes [27, 25, 28]. Given a shape X , we say that it is intrinsically symmetric if there exists a mapping $S : X \rightarrow X$ that preserves all the geodesic distances between the corresponding points

$$d_X(x, \tilde{x}) = d_X(S(x), S(\tilde{x})), \quad \forall x, \tilde{x} \in X. \quad (21)$$

If the shape X is intrinsically symmetric, and $S : X \rightarrow X$ is its intrinsic symmetry, then the surface descriptors mentioned in the previous section are also symmetric functions with respect to S . That is, for each of their components $f_X^{(i)}$ the following holds

$$f_X^{(i)}(x) = f_X^{(i)}(S(x)). \quad (22)$$

From the Equation (21) and the above property of the descriptors it follows that if $P^*(x, y) = \text{argmin} \{ \text{dis}(P) \}$, then $P^*(S(x), y)$ also minimizes the dissimilarity $\text{dis}(P)$,

with the same minimal value. Thus, when matching intrinsically symmetric shapes, the optimization problem (8) has multiple solutions, and by minimizing $\text{dis}(P)$ we can obtain only one of them.

In order to overcome the above problem, a technique for construction of symmetry-aware surface descriptors was suggested in [8]. These descriptors are based on the eigendecomposition of the Laplace-Beltrami operator, and exploit the important property of the eigenfunctions of Δ_X , described by Ovsjanikov *et al.* in [25]. As stated in Theorem 3.1. of [25], eigenfunctions corresponding to non-repeating eigenvalues of the Laplace-Beltrami operator of an intrinsically symmetric shape exhibit reflection symmetry, with respect to the shape's intrinsic symmetry. That is, such an eigenfunction ϕ can be either symmetric or anti-symmetric with respect to S

$$\phi(x) = \phi(S(x)) \text{ or } \phi(x) = -\phi(S(x)). \quad (23)$$

As described in [8], the symmetry-aware surface descriptors are constructed as follows

$$f_X(x) = [\phi_1^X(x), \phi_2^X(x), \dots, \phi_d^X(x)], \quad (24)$$

and

$$f_Y(y) = [s_1\phi_1^Y(y), s_2\phi_2^Y(y), \dots, s_3\phi_d^Y(y)]. \quad (25)$$

In the above, $\{\phi_i^X\}_{i=1}^d$ and $\{\phi_j^Y\}_{j=1}^d$ are the eigenfunctions corresponding to the first d non-repeating eigenvalues of the Laplace-Beltrami operators of the two shapes, respectively. The values of $\{\phi_j^Y\}_{j=1}^d$ are then multiplied by the sign sequence $\{s_j\}_{j=1}^d$, to obtain consistent descriptors for X and Y .

Figure 2 shows an example of two human shapes colored according to the values of the first 3 eigenfunctions of their corresponding Laplace-Beltrami operators. It is easy to see that the eigenfunctions of the lower shape have to be multiplied by a sequence $[+, -, +]$, in order to be equal to the eigenfunctions of the upper shape in the corresponding points. But it is also possible to multiply them by a sequence $[+, +, -]$, and thus obtain eigenfunctions reflectionally symmetric to the eigenfunctions of the upper shape. In general, the number of different sign sequences, and thus different sets of descriptors for the shape Y , is determined by the number of intrinsic symmetries of the shape. Using these sets of descriptors in the optimization problem (6) allows us to find several different correspondences between the two shapes. The exact algorithm for the sign sequence detection and its limitations are presented in details in [8].

6 Results

In this section we provide several matching results obtained with the proposed framework. All the shapes we used in our tests were represented by triangulated




Fig. 2 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. The two possible sign sequence relating the two groups of the eigenfunctions are $[+,-,+]$ and $[+,+,-]$.

meshes with several thousand vertices. We further sub-sampled the shapes using the Farthest Points Sampling algorithm [11], to obtain sets of matching candidate points. In each one of our tests, we performed the matching using 10 points at the coarse scale, and 30 - 64 points at the finest scale. Note that the later sub-sampling affects the accuracy of the matching, and the denser the sub-sampling is the more accurate the obtained correspondences are.

Figures 3 and 4 present the results of matching ε -isometric shapes using the proposed framework combined with different distance measured and descriptors, at several hierarchies, where the correspondences are shown by Voronoi cells of the matched points, corresponding patches having the same color. The matches in Figure 3(a, b, c) are the symmetrical ones, which is one of the possible matchings in this case, as explained in Section 5. Some matching inaccuracies, *e.g.* inaccurate correspondences between the cats' ears in Figure 4(a), appear when the algorithm converges to local minima.

In order to find dense correspondence between all the points on the shapes, the above matching results can be used as a input for algorithms such as described in [24] or [16].

7 Conclusions

In this paper we have presented a method for automatic correspondence of non-rigid ε -isometric shapes, based on comparison of surface descriptors and distance

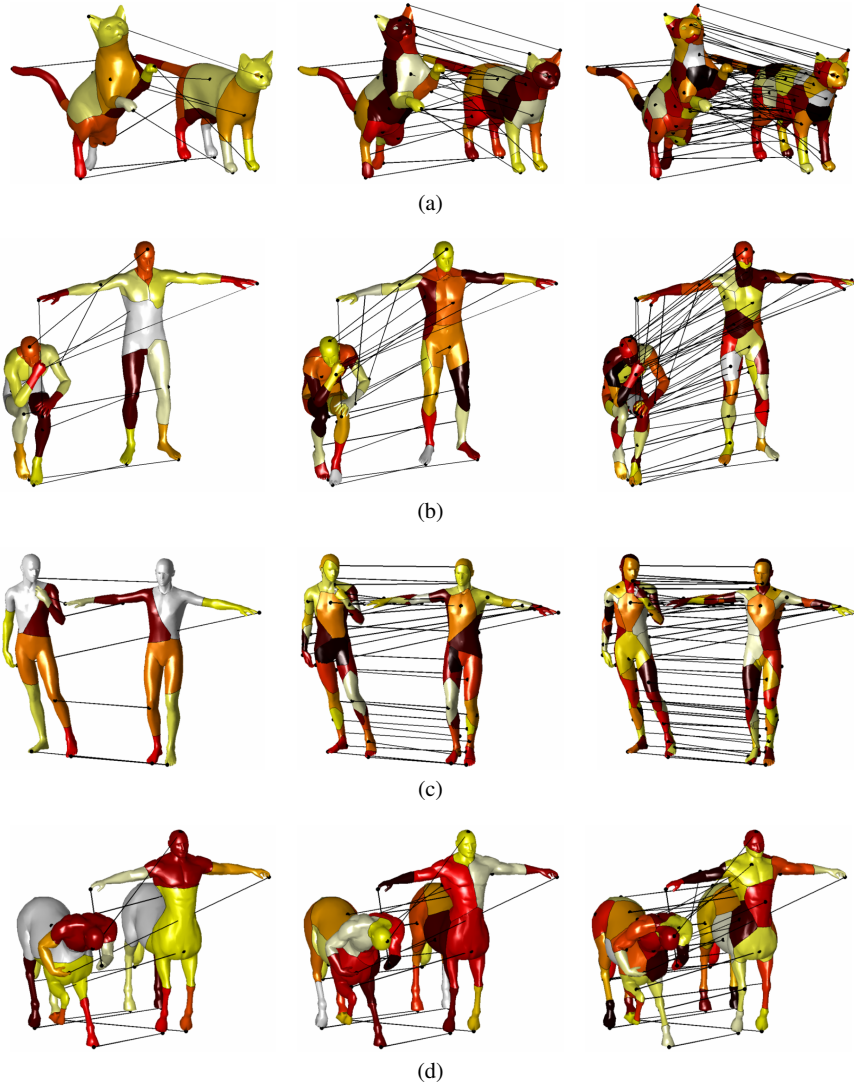


Fig. 3 Matching results obtained with the proposed framework combined with different descriptors and metrics, at several hierarchies. (a), (b) Geodesic distance and geodesic distance-based histogram descriptor; (c) diffusion distance and diffusion distance-based histogram descriptor; (d) Diffusion distance and Heat Kernel Signatures.

structures, and tested it with different choices of the latter. In addition, we have formulated the matching ambiguity problem arising when matching intrinsically symmetric shapes, and showed that the proposed framework combined with certain descriptors allows us to detect multiple possible correspondences. In future work

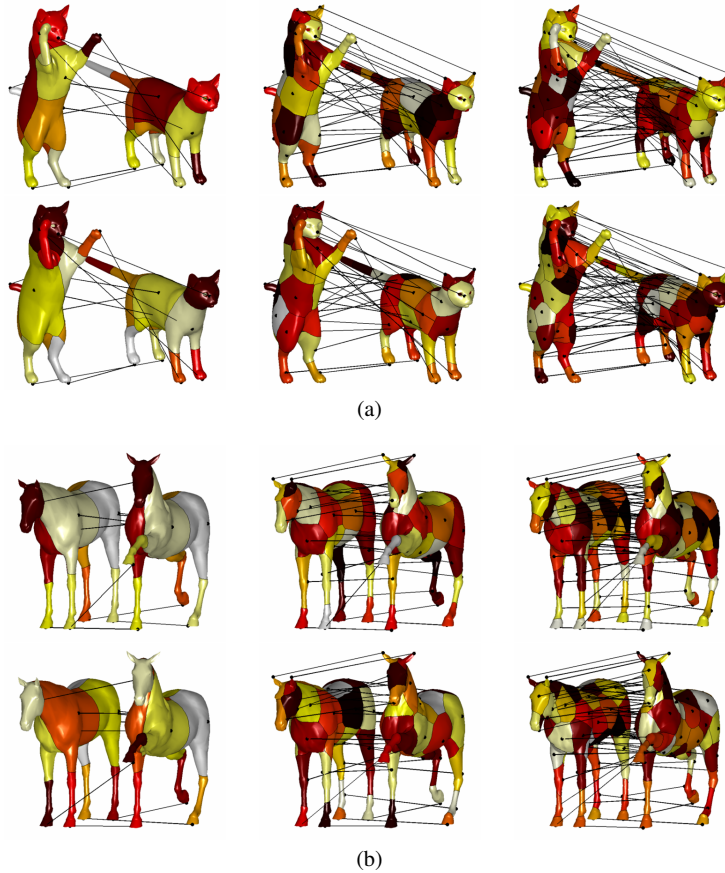


Fig. 4 Matching results obtained with the proposed framework combined with geodesic distance metric and the Laplace-Beltrami operator-based descriptors; upper row - same orientation correspondence, lower row - the reflected one.

we intend to adapt the proposed framework for partial shape matching and extend it to shapes that are not necessarily ε -isometric.

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

References

1. D. Anguelov, P. Srinivasan, H.-C. Pang, D. Koller, and S. Thrun. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Proc. of the Neural Information Processing Systems (NIPS) Conference*, volume 17, pages 33–40, 2004.
2. A. Bemporad. Hybrid Toolbox - User's Guide, 2004. <http://www.dii.unisi.it/hybrid/toolbox>.
3. P. Bérard, G. Besson, and S. Gallot. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, 4(4):373–398, 1994.
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.
5. A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer, 2008.
6. A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision (IJCV)*, 2009.
7. R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21:5–30, 2006.
8. A. Dubrovina and R. Kimmel. Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In *International Symposium on 3D Data Processing Visualization and Transmission (3DPVT)*, 2010.
9. A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*, 25(10):1285–1295, 2003.
10. M. Gromov. *Structures métriques pour les variétés riemanniennes*. Textes Mathématiques, no. 1, 1981.
11. 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.
12. J. Hu and J. Hua. Salient spectral geometric features for shape matching and retrieval. *Vis. Comput.*, 25(5-7):667–675, 2009.
13. V. Jain, H. Zhang, and O. Van Kaick. Non-rigid spectral correspondence of triangle meshes. *International Journal on Shape Modeling*, 13(1):101–124, 2007.
14. V. Kim, Y. Lipman, X. Chen, and T. Funkhouser. Mobius transformations for global intrinsic symmetry analysis. In *Proc. Eurographics Symposium on Geometry Processing (SGP)*, 2010.
15. R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. In *Proc. of National Academy of Sciences (PNAS)*, volume 95, pages 8431–8435, 1998.
16. V. Kraevoy and A. Sheffer. Cross-parameterization and compatible remeshing of 3D models. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 23(3):861 – 869, August 2004.
17. Y. Lipman, X. Chen, I. Daubechies, and T. Funkhouser. Symmetry factored embedding and distance. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 2010.
18. Y. Lipman and T. Funkhouser. Mobius voting for surface correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 28(3), August 2009.
19. D. Mateus, R. P. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
20. F. Mémoli. On the use of Gromov-Hausdorff distances for shape comparison. In *Point Based Graphics 2007*, pages 81–90, September 2007.
21. F. Mémoli. Spectral Gromov-Wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment (ICCV workshop, NORDIA'09)*, 2009.
22. F. Mémoli and G. G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5(3):313–347, 2005.
23. M. Meyer, M. Desbrun, P. Schröder, and A. Barr. Discrete differential geometry operators for triangulated 2-manifolds. In *International Workshop on Visualization and Mathematics*, 2002.

24. M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas. One point isometric matching with the heat kernel. In *Eurographics Symposium on Geometry Processing (SGP)*, 2010.
25. M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. *Computer Graphics Forum*, 27(5):1341–1348, July 2008.
26. Ulrich Pinkall and Konrad Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2:15–36, 1993.
27. 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.
28. D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Full and partial symmetries of non-rigid shapes. *International Journal of Computer Vision (IJCV)*, 2010.
29. D. Raviv, A. Dubrovina, and R. Kimmel. Hierarchical matching of non-rigid shapes. In *Int. Conf. Scale Space and Variational Methods (SSVM)*, 2011.
30. Steven Rosenberg. *The Laplacian on a Riemannian manifold: an introduction to analysis on manifolds*. Cambridge University Press, 1997.
31. Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *Int. J. Comput. Vision*, 40:99–121, November 2000.
32. M. R. Ruggeri and D. Saupe. Isometry-invariant matching of point set surfaces. In *Proc. Eurographics 2008 Workshop on 3D Object Retrieval*, 2008.
33. R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. of SGP*, pages 225–233, 2007.
34. A. Sharma and R. P. Horaud. Shape matching based on diffusion embedding and on mutual isometric consistency. In *Proc. Workshop on Nonrigid Shape Analysis and Deformable Image Alignment (NORDIA)*, June 2010.
35. J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. Eurographics Symposium on Geometry Processing (SGP)*, 2009.
36. A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. Isometric registration of ambiguous and partial data. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, page to appear, 2009.
37. N. Thorstensen and R. Keriven. Non-rigid shape matching using geometry and photometry. In *Asian Conference on Computer Vision*, pages 1–12, 2009.
38. L. Torresani, V. Kolmogorov, and C. Rother. Feature correspondence via graph matching: Models and global optimization. In *Proc. of the 10th European Conference on Computer Vision (ECCV ’08)*, pages 596–609, 2008.
39. C. Wang, M. M. Bronstein, and N. Paragios. Discrete minimum distortion correspondence problems for non-rigid shape matching. Technical report, INRIA Research Report 7333, Mathématiques Appliquées aux Systèmes, École Centrale Paris, 2010.
40. C. Wang, M. M. Bronstein, and N. Paragios. Discrete minimum distortion correspondence problems for non-rigid shape matching. Technical report, Int. Conf. Scale Space and Variational Methods (SSVM), 2011.
41. A. Zaharescu, E. Boyer, K. Varanasi, and R. P. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
42. 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)*, 27(5):1431–1439, 2008.
43. G. Zigelman, R. Kimmel, and N. Kiryati. Texture mapping using surface flattening via multi-dimensional scaling. *IEEE Trans. on Visualization and Computer Graphics*, 8(2):198–207, 2002.