

Discrete Minimum Distortion Correspondence Problems for Non-rigid Shape Matching

Chaohui Wang^{1,2}, Michael M. Bronstein³, Alexander M. Bronstein⁴, and Nikos Paragios^{1,2}

¹ Laboratoire MAS, Ecole Centrale de Paris, Châtenay-Malabry, France

² Equipe GALEN, INRIA Saclay - Île de France, Orsay, France

³ Institute of Computational Science, Faculty of Informatics
Università della Svizzera Italiana, Lugano, Switzerland

⁴ Department of Electrical Engineering, Tel Aviv University, Israel

Abstract. Similarity and correspondence are two fundamental archetype problems in shape analysis, encountered in numerous application in computer vision and pattern recognition. Many methods for shape similarity and correspondence boil down to the minimum-distortion correspondence problem, in which two shapes are endowed with certain structure, and one attempts to find the matching with smallest structure distortion between them. Defining structures invariant to some class of shape transformations results in an invariant minimum-distortion correspondence or similarity. In this paper, we model shapes using local and global structures, formulate the invariant correspondence problem as binary graph labeling, and show how different choice of structure results in invariance under various classes of deformations.

1 Introduction

Recent works in computer vision and shape analysis [1–4] have shown that different approaches to shape similarity and correspondence can be considered as instances of the *minimum distortion correspondence problem*, in which two shapes are endowed with certain structure, and one attempts to find the best (least distorting) matching between these structures. Examples of such structures include multiscale heat kernel signatures [5–7], local photometric properties [8, 9], eigenfunctions of the Laplace-Beltrami operator [10–13], triplets of points [14, 15], and geodesic [2, 3, 16], diffusion [17], and commute time [10, 18] distances. By defining a structure invariant under certain class of transformations (e.g. non-rigid deformations), one obtains correspondence invariant under that class (in the above example, deformation invariant matching). The Gromov-Hausdorff distance [19] is an important particular case of the minimum distortion correspondence problem, in which the matched structures are metric spaces, invariant to isometries of the metric structures.

Some settings of the minimum distortion correspondence problem can be reformulated as labeling problems [20], such that the objective function can be

optimized efficiently using the recently developed discrete optimization algorithms. For example, the dual-decomposition strategy [21], introduced by [22] to perform pairwise Markov random field (MRF) inferences, provides a powerful technique to solve such labeling problems. Based on such a strategy, Torresani *et al.* proposed a pairwise graph matching algorithm [20] to compute correspondence between images using a criterion combining local features and Euclidean distances between nearby features. Such an approach showed better performance than feature-only based methods in deformable 2D object tracking. The increased performance attributed to the use of inter-feature distances as a geometric consistency constraint. However, Euclidean distances are not deformation invariant and can be applied only locally, thus limiting the usefulness of such a constraint.

Main contribution. In this paper, we study the minimum distortion correspondence problem in the context of non-rigid shape analysis. We formulate *invariant correspondence* as a minimizer of a distortion criterion based on structures invariant to some classes of transformations. In particular, we use local and global structures invariant to important classes of transformations such as non-rigid deformations, changes in topology, and scaling. By such an axiomatic construction of invariant structures, we obtain invariant correspondence. In particular, we show scale invariant shape matching using only singleton and pairwise interactions without higher-order terms. Compared to Torresani *et al.* [20], our use of global structures in non-rigid shapes provides a better regularization to the problem and is better motivated geometrically. Yet, it also increases the computational complexity of the optimization. To address this problem, we use hierarchical matching, in which candidate correspondences are restricted to neighborhoods of matching points from coarser levels.

While the described axiomatic approach is suitable for modeling geometric shape transformations such as bendings, it is not applicable to intra-class shape variations (e.g. different appearances of a human shape). To cope with this case, we show a probabilistic extension of our framework, in which local and global structures are replaced with respective multidimensional distributions, accounting for shape variability.

Related work. Feature-based shape matching methods for non-rigid shapes were used in numerous recent works [8, 5, 12]. Tree-based [23] and branch-and-bound techniques [24] were used to find the matches between the feature points. Elad and Kimmel [16] used multidimensional scaling (MDS) to represent shapes in a low-dimensional Euclidean space and compare them as rigid objects. The use of an intermediate embedding space was eliminated in [2] using the Gromov-Hausdorff formalism [19]. Bronstein *et al.* [3] proposed an MDS-like algorithm referred to as generalized MDS (GMDS) for the computation of the Gromov-Hausdorff distance and deformation invariant correspondence between shapes. This framework was extended in [17] using diffusion geometry instead of the geodesic one. In [25], Mémoli extended [2, 3] by modeling shapes as metric-measure spaces. He introduced the Gromov-Wasserstein distance based on measure coupling between two metric-measure spaces, and formulated it as a quadratic

assignment problem (QAP). Thorstensen and Keriven [9] extended the GMDS framework to textured shapes introducing *photometric stress* as a local matching term in addition to geodesic distance distortion. Dubrovina and Kimmel [13] generalized this approach for the matching of textureless shapes using Laplace-Beltrami eigenfunctions as local geometric descriptors. Mateus *et al.* [11] showed a non-rigid shape correspondence approach with inexact graph matching based on spectral embedding. In the image domain, Torresani *et al.* [20] used graph labeling problem to match 2D images.

2 Problem formulation

Our shape model is an extension of the *metric model* used in [2, 3, 16]. We assume that the shapes are endowed with *local* and *global* structure, and try to find such a correspondence between the shapes that best preserves these structures. The structures are defined having in mind certain invariance properties required in the particular problem, as discussed in Section 3. Given a shape X , modeled as a connected surface (possibly with boundary) embedded into \mathbb{R}^3 (or \mathbb{R}^2 in case of planar shapes), its local structure is modeled by a vector field $\mathbf{f}_X : X \rightarrow \mathbb{R}^m$ referred to as a *local descriptor*. The global structure of the shape is modeled as a *metric* $d_X : X \times X \rightarrow \mathbb{R}$, defined as a positive-definite subadditive function between pairs of points on X .

Given two shapes X and Y with the local descriptors \mathbf{f}_X and \mathbf{f}_Y and metrics d_X and d_Y , respectively, we define a *bijective correspondence* between X and Y as $\mathcal{C} \subset X \times Y$ satisfying $\forall x \in X \exists! y \in Y$ such that $(x, y) \in \mathcal{C}$ and $\forall y \in Y \exists! x \in X$ such that $(x, y) \in \mathcal{C}$. A good correspondence should match similar descriptors between corresponding points and similar metrics between corresponding pairs of points. This can be quantified using *first-* and *second-order distortion* terms, $\text{dis}(\mathcal{C}) = \|\mathbf{f}(\mathcal{C})\|$ and $\text{dis}(\mathcal{C} \times \mathcal{C}) = \|\mathbf{d}(\mathcal{C} \times \mathcal{C})\|$, measuring the quality of correspondence of local and global structures, respectively. (here, $\mathbf{f}(\mathcal{C})$ is a $|\mathcal{C}| \times 1$ vector with elements $\|\mathbf{f}_X(x) - \mathbf{f}_Y(y)\|$ for all $(x, y) \in \mathcal{C}$; $\mathbf{d}(\mathcal{C} \times \mathcal{C})$ is a $|\mathcal{C}|^2 \times 1$ vector with elements $|d_X(x, x') - d_Y(y, y')|$ for all $(x, y), (x', y') \in \mathcal{C}$; and $\|\cdot\|$ is some norm). In particular,

$$\begin{aligned} \text{dis}_2(\mathcal{C}) &= \sum_{(x,y) \in \mathcal{C}} \|\mathbf{f}_X(x) - \mathbf{f}_Y(y)\|^2; \\ \text{dis}_2(\mathcal{C} \times \mathcal{C}) &= \sum_{(x,y),(x',y') \in \mathcal{C}} (d_X(x, x') - d_Y(y, y'))^2. \end{aligned}$$

The optimal correspondence is found by minimizing a combination of first- and second-order distortion terms,

$$\min_{\mathcal{C}} \text{dis}(\mathcal{C}) + \beta \text{dis}(\mathcal{C} \times \mathcal{C}), \quad \beta \geq 0. \quad (1)$$

The minimizer of problem (1) is the *minimum distortion correspondence* between X and Y . The minimum of problem (1) can be interpreted as the similarity of

X and Y .⁵ A particular theoretically important case is a minimum-distortion correspondence with an L_∞ second-order distortion term, referred to as the *Gromov-Hausdorff distance* [19]:

$$d_{\text{GH}}(X, Y) = \frac{1}{2} \min_{\mathcal{C}} \max_{(x,y),(x',y') \in \mathcal{C}} |d_X(x, x') - d_Y(y, y')|.$$

3 Invariance

The choice of the local and global structures ($\mathbf{f}_X, \mathbf{f}_Y$ and d_X, d_Y) defines the invariance properties of the correspondence. Assume that the shape $Y = \tau(X)$ is obtained from X by means of some transformation τ from a class \mathcal{T} . If $\mathbf{f}_X \circ \tau = \mathbf{f}_Y$ and $d_X \circ (\tau \times \tau) = d_Y$ for all $\tau \in \mathcal{T}$, our structures are *invariant* under the class of transformations \mathcal{T} . As a result, correspondence obtained by the solution of problem (1) is also invariant under \mathcal{T} . Important invariance classes can be addressed by appropriate definition of the descriptors and the metric. In particular, we are interested in *inelastic deformations* (bendings), changing the embedding of the shape without changing its intrinsic structure; *topological transformations*, resulting in local changes in the connectivity of the shape, appearing as holes or “gluing” two points on the surface; and *scaling*.

3.1 Choice of the metric

Geodesic metric. One of the most straightforward definitions of a metric on a surface is the *geodesic metric*, measuring the length of a shortest path between points x and x' ,

$$d_X(x, x') = \min_{\gamma \in \Gamma(x, x')} \ell(\gamma),$$

where $\Gamma(x, x')$ denotes the set of all admissible paths between x and x' , γ is some admissible path, and $\ell(\gamma)$ is its length. The geodesic metric is intrinsic, dependent only on local distance structure of the shape, and is thus invariant to inelastic deformations [2, 3, 16]. A notable drawback of the geodesic distance is its sensitivity to topological transformations. Connectivity changes alter the admissible paths Γ (e.g., gluing the fingers of the hand creates new paths that have not existed before), and, since the geodesic distance takes the minimum over all path lengths, sometimes the change in the geodesic metric can be very significant.

Diffusion metric. A more robust definition of an intrinsic metric based on heat diffusion properties has been recently popularized by Lafon *et al.* [27]. Heat

⁵ The minimizer of problem (1) is not necessarily unique, i.e., there may be two different correspondences $\mathcal{C} \neq \mathcal{C}'$ with $\text{dis}(\mathcal{C}) = \text{dis}(\mathcal{C}')$. Such situations are typical when the shapes have *intrinsic symmetries*. Intrinsic symmetry is manifested by the existence of a self-isometry of X with respect to the metric d_X , i.e., an automorphism $g : X \rightarrow X$ satisfying $d_X = d_X \circ (g \times g)$ [24, 26].

diffusion on manifolds is governed by the *heat equation* $(\Delta_X + \frac{\partial}{\partial t})u = 0$, where u is the heat distribution and Δ_X is the positive semi-definite *Laplace-Beltrami operator* (LBO), which can be roughly thought of as a generalization of the Laplacian to non-Euclidean domains. The *heat kernel* $h_{X,t}(x, z)$ is the solution of the heat equation with a point heat source at point x at time $t = 0$. For compact manifolds, the Laplace-Beltrami operator has discrete eigendecomposition of the form $\Delta_X \phi_i = \lambda_i \phi_i$, where $\lambda_0 = 0, \lambda_1, \dots \geq 0$ are eigenvalues and ϕ_0, ϕ_1, \dots are eigenfunctions. Using the eigenbasis of Δ_X , the heat kernel can be presented as

$$h_{X,t}(x, z) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x) \phi_i(z). \quad (2)$$

A family of metrics

$$d_{X,t}(x, y) = \|h_{X,t}(x, \cdot) - h_{X,t}(y, \cdot)\|_{L_2(X)} = \sum_{i=1}^{\infty} e^{-2\lambda_i t} (\phi_i(x) - \phi_i(y))^2, \quad (3)$$

parameterized by the time scale t , is referred to as *diffusion metrics*. Diffusion metric is inversely related to the connectivity of points x and y by paths of length t . Unlike the geodesic distance which measures the length of the shortest path, the diffusion metric has an averaging effect over all paths connecting two points. As a result, diffusion metric is less sensitive to topology and connectivity changes [17]. With an appropriate selection of the time scale t , the effect of topological noise can be reduced.

Commute-time metric. At the same time, the need to select the scale parameter is a disadvantage, as it depends on the shape scale. Moreover, the diffusion metric is not scale invariant, since scale change affects the eigenvalues λ_i and eigenfunctions ϕ_i . A different metric,

$$\delta_X(x, y) = \sum_{i=1}^{\infty} \frac{1}{\lambda_i} (\phi_i(x) - \phi_i(y))^2, \quad (4)$$

called the *commute time* (or *resistance* [28]) *distance*, is similar in its spirit to the diffusion metric, while being scale-invariant. The commute time metric measures the connectivity of points by paths of any length and is related to the expected time it takes a random walk initiating at point x go through point y and return to x .

3.2 Choice of the descriptor

Similarly to our motivation in the selection of the metric, the choice of the local descriptor is also dictated by the desired invariance properties. Due to their locality, many types of descriptors are usually less susceptible to changes as a result of non-rigid deformations. However, some descriptors have explicit invariance properties by construction.

Structure	Bending Topology Scale		
<i>Local histogram</i> [24]	Yes	No	No
<i>LBO eigenfunctions</i> [12, 13]	Yes	No	Yes
<i>HKS</i> [5]	Yes	Approx	No
<i>SI-HKS</i> [7]	Yes	Approx	Yes
<i>Geodesic metric</i> [2, 3]	Yes	No	No
<i>Diffusion metric</i> [17]	Yes	Approx	No
<i>Commute-time metric</i> [18]	Yes	Approx	Yes

Table 1. Invariance properties of local (top rows) and global (bottom rows) structures.

Heat kernel signature. Sun *et al.* [5] introduced intrinsic descriptors based on multi-scale heat kernels, referred to as *heat kernel signatures* (HKS). The HKS is constructed at every point of the shape by considering the values of the heat kernel diagonal at multiple time scales, $\mathbf{f}_X(x) = (h_{X,t_1}(x,x), \dots, h_{X,t_n}(x,x))$, where t_1, \dots, t_n are some time scale. The HKS is invariant to inelastic deformations and was also shown to be insensitive to topological transformations [6].

Scale-invariant heat kernel signature. The disadvantage of HKS is the lack of scale invariance. In a follow-up work, Bronstein and Kokkinos [7] introduced a scale-invariant modification of HKS, referred to as SI-HKS. The main idea is to sample the time scales logarithmically ($t = \alpha^\tau$) such that shape scaling corresponds to a scale-space shift. Such a shift is then undone by taking the magnitude of the Fourier transform w.r.t. τ . The SI-HKS enjoys the invariance properties of HKS, while in addition also being scale-invariant.

4 Correspondence as a graph labeling problem

Our minimum-distortion correspondence problem can be formulated as a *binary labeling* problem with uniqueness constraints [20] in a graph with vertices defined as pairs of points and edges defined as quadruplets. More formally, let $\mathcal{V} = \{(x, y) : x \in X, y \in Y\} = X \times Y$ be the set of pairs of points from X and Y , and let $\mathcal{E} = \{((x, y), (x', y')) \in \mathcal{V} \times \mathcal{V} \text{ and } (x, y) \neq (x', y')\}$. Let $\mathcal{L} = \{0, 1\}$ further denote the set of binary labels. We can represent a correspondence $\mathcal{C} \subseteq \mathcal{V}$ as binary labeling $u \in \mathcal{L}^{\mathcal{V}}$ of the graph $(\mathcal{V}, \mathcal{E})$, as follows: $u(x, y) = 1$ iff $(x, y) \in \mathcal{C}$ and 0 otherwise. When using L_2 distortions, the correspondence problem (1) can be reformulated as:

$$\begin{aligned}
\min_{u \in \mathcal{L}^{\mathcal{V}}} & \sum_{(x,y) \in \mathcal{V}} u_{x,y} (\|\mathbf{f}_X(x) - \mathbf{f}_Y(y)\| - \gamma) + \\
& \beta \sum_{((x,y),(x',y')) \in \mathcal{E}} u_{x,y} u_{x',y'} |d_X(x, x') - d_Y(y, y')|^2 \\
\text{s.t.} & \sum_y u_{x,y} \leq 1 \quad \forall x \in X; \quad \sum_x u_{x,y} \leq 1 \quad \forall y \in Y.
\end{aligned} \tag{5}$$

where $\gamma > 0$ is an occlusion term [20] to penalize unmatched points. We can choose a sufficiently large γ to ensure the bijective correspondence and the equivalence of the two problems.

In general, optimization of this energy is NP-hard [29]. Here, we adopt the graph matching algorithm [20] based on dual-decomposition to perform the optimization of (5). The key idea of this approach is, instead of minimizing directly the energy (5) of the original problem, to maximize a lower bound on it by solving the dual to the linear programming (LP) relaxation of (5). Such approaches demonstrate good global convergence behavior [22]. We first decompose the original problem, which is too complex to solve directly, into a series of sub-problems, each of which is smaller and solvable. After getting the solution of the sub-problems, we combine them using a projected-subgradient scheme to get the solution of the original problem. In the numerical experiments, following [20], we decomposed problem (5) into a linear subproblem, a maxflow subproblem and a set of local subproblems.

4.1 Hierarchical matching

Assuming for simplicity $|X| = |Y| = N$, the number of vertices in the graph is $|\mathcal{V}| = N^2$ and the number of edges, assuming full connectivity, is $\mathcal{O}(N^4)$. The complexity of problem (5) is $\mathcal{O}(|\mathcal{V}|^2|\mathcal{E}|)$ multiplied by the number of iterations, i.e., $\mathcal{O}(N^8)$. This complexity can be reduced by adopting a hierarchical matching strategy: after finding a coarse correspondence between a small number of points, correspondence between nearby points only is looked for. This allows to significantly reduce the graph size.

Let x_1, x_2, \dots denote a progressive sampling of the shape X , such that $X_n = \{x_1, \dots, x_n\}$ constitutes an r_n -covering of X (i.e., $d_X(X, X_n) \leq r_n$, where d_X is some metric on X). Such a sequence of points can be found using e.g. farthest point sampling (FPS) strategy [30], in which x_1 is selected arbitrarily and the next point is selected as $x_{k+1} = \arg \max_{x \in X} \min_{i=1, \dots, k} d_X(x, x_i)$. Same way, $Y_n = \{y_1, \dots, y_n\}$ will denote an r'_n -covering of Y .

At the first stage of hierarchical matching, correspondence is found between X_{N_1} and Y_{N_1} , where N_1 is some small number (in our experiments, it varied between 4 and 10), solving the labeling problem (5) on the full graph ($\mathcal{V}_1 = X_{N_1} \times Y_{N_1}$, $\mathcal{E}_1 = \{((x, y), (x', y')) \in \mathcal{V}_1 \times \mathcal{V}_1 \text{ and } (x, y) \neq (x', y')\}$). The solution provides a coarse correspondence $\mathcal{C}_1 \subset X_{N_1} \times Y_{N_1}$.

At the $(k+1)$ st level, correspondence is found between $X_{N_{k+1}}$ and $Y_{N_{k+1}}$ (the number of points is increased by a factor typically $2 \leq q = N_{k+1}/N_k \leq 4$), restricting the correspondence candidates for points within a certain radius around x to points within a certain radius around y , where $(x, y) \in \mathcal{C}_k$. This way, the $(k+1)$ st level labeling problem is solved on the graph with vertices

$$\mathcal{V}_{k+1} = \{(x_i, y_i) \in X_{N_{k+1}} \times Y_{N_{k+1}} : \exists (x, y) \in \mathcal{C}_k \text{ s.t. } d_X(x, x_i) < \rho r_k, d_Y(y, y_i) < \rho r'_k\},$$

where $\rho > 1$, and $\mathcal{E}_{k+1} = \{((x, y), (x', y')) \in \mathcal{V}_{k+1} \times \mathcal{V}_{k+1} \text{ and } (x, y) \neq (x', y')\}$. For $\rho \approx 1$, the size of the ρr_k -neighborhood in $X_{N_{k+1}}$ of a point from X_{N_k} contains $\mathcal{O}(q)$ points. Thus, $|\mathcal{V}_{k+1}| = \mathcal{O}(q^2 N_k)$, and $|\mathcal{E}_{k+1}| = \mathcal{O}(q^4 N_k^2)$ points, a significant reduction compared to $\mathcal{O}(N_{k+1}^2)$ vertices and $\mathcal{O}(N_{k+1}^4)$ edges in a full graph. As a result, the complexity of the optimization becomes $\mathcal{O}(N^4)$.

5 Probabilistic matching and shape prototypes

While invariance to geometric transformations such as bending can be accounted by the selection of local (descriptor) and global (metric) structures, many types of shape variability cannot be accounted for in this way. For example, variability within the shape class (e.g. fat or thin man) results in different local and global structures that cannot be modeled geometrically. At the same time, such a variability can be modeled statistically. Instead of \mathbf{f}_X and d_X , we now have *distributions* $\mathbf{f}_X \sim \mathcal{F}_X$ and $d_X \sim \mathcal{D}_X$, e.g., *Gaussian mixture model* for distances,

$$p_{xx'}(d) = \sum_{k=1}^K \pi_{xx'k} \frac{1}{\sqrt{2\pi}\sigma_{xx'k}} \exp\left\{-\frac{(d - \mu_{xx'k})^2}{2\sigma_{xx'k}^2}\right\}, \quad \sum_{k=1}^K \pi_{xx'k} = 1; \quad \forall x \neq x' \in X,$$

and descriptors,

$$p_x(\mathbf{f}) = \sum_{k=1}^K \pi_{xk} \frac{\exp\left\{-\frac{1}{2}(\mathbf{f} - \boldsymbol{\mu}_{xk})^T \boldsymbol{\Sigma}_{xk}^{-1}(\mathbf{f} - \boldsymbol{\mu}_{xk})\right\}}{(2\pi)^{m/2}(\det \boldsymbol{\Sigma}_{xk})^{1/2}}, \quad \sum_{i=1}^K \pi_{xk} = 1; \quad \forall x \in X,$$

where p denotes probability density. The distance distribution between points x and x' is parameterized by $\mathcal{D}_{xx'} = \{\mu_{xx'k}, \sigma_{xx'k}^2, \pi_{xx'k}\}_{k=1}^K$; the distribution of descriptors at each points x is parameterized by $\mathcal{F}_x = \{\boldsymbol{\mu}_{xk}, \boldsymbol{\Sigma}_{xk}, \pi_{xk}\}_{i=1}^K$, where $\boldsymbol{\mu}_{xk}$ are $m \times 1$ vectors and $\boldsymbol{\Sigma}_{xk}$ are $m \times m$ matrices. We call $\mathcal{X} = ((\mathcal{F}_x)_{x \in X}, (\mathcal{D}_{xx'})_{x \neq x' \in X})$ a *shape prototype*.

In this probabilistic setting, given a shape Y , we determine the correspondence between Y and the prototype \mathcal{X} by solving a problem similar to (5), with the distortion terms replaced by negative log-likelihood functions,

$$\begin{aligned} \min_{u \in \mathcal{L}^{\mathcal{V}}} & - \sum_{(x,y) \in \mathcal{V}} u_{x,y} (\log p_x(\mathbf{f}_Y(y)) + \gamma) - \\ & \beta \sum_{((x,y),(x',y')) \in \mathcal{E}} u_{x,y} u_{x',y'} \log p_{xx'}(d_Y(y, y')) \\ \text{s.t.} & \sum_y u_{x,y} \leq 1 \quad \forall x \in X; \quad \sum_x u_{x,y} \leq 1 \quad \forall y \in Y. \end{aligned} \quad (6)$$

6 Results

To assess the performance of the presented approach, we performed multiple experiments of shape correspondence and similarity computation under a variety of transformations. Shapes from the TOSCA [31] and Princeton [32] datasets were used in our experiments. Textured shapes acquired with a multicamera system were taken from the INRIA Grenoble dataset [8]. The shapes were represented as triangular meshes with 2000-10000 vertices. Textures were given as RGB values for each vertex. Geodesic distances were computed using fast marching [33]. Diffusion and commute time metrics were computed using the spectral formulae (3) and (4) taking the first 100 eigenvalues. The Laplace-Beltrami operator

was approximated using cotangent weights [34]. The heat kernel was approximated using formula (2). Hierarchical matching was implemented in MATLAB with discrete optimization module in C++. Typical running times for pairwise shape matching in the following experiments were about 10 – 20 sec.

Invariance and the choice of the metric/descriptor. In the first experiment, matching was performed between eight points with equal weight given to the local and global distortion terms in the optimization problem. Three combinations of first- and second-order structures were used: geodesic metric/HKS descriptor, diffusion metric/HKS descriptor, and commute time metric/SI-HKS descriptor. Figure 1 shows the result of correspondence computation between shapes with different transformations for different choice of metric/descriptor. All three methods are invariant to bendings (first row; note that correspondence is defined up to an intrinsic symmetry). The combination geodesic metric/HKS descriptor is sensitive to topology (a human with hands glued to legs, second row) and scale. The combination diffusion metric/HKS descriptor is insensitive to topology but sensitive to scale. Finally, commute time metric with SI-HKS descriptor are invariant to all of the above.

Shape prototypes. In the fourth experiment, a shape prototype was created based on 64 examples of a human shape, in which the length of the hands and legs and the size of the head varied. Distance and descriptor distributions were represented using Gaussian mixtures with 5 components. Figure 2 shows a comparison of deterministic and probabilistic matching. Using deterministic matching, the shape of a humanoid alien from the Princeton database is matched to the human shape from TOSCA dataset incorrectly (second column from left), because of different proportions of the head, legs, and hands. On the other hand, matching to the human shape prototype using probabilistic matching produced correct symmetric correspondence (third column). Figure 2 (columns four and five) shows additional examples of shape prototype matching. These results show that the probabilistic matching framework allows to address shape variability that cannot be simply accommodated into the metric model by choosing the metric.

7 Conclusions

We presented a generic framework for invariant matching between shapes, in which matching is performed by minimizing the distortion of local and global geometric structures under the correspondence. Using structures invariant to pre-defined classes of transformations (or, using their statistical distributions if such transformations cannot be modeled explicitly) allows obtaining invariant matching between shapes. Our approach generalizes many previous works in the field, in particular, methods based on metric distortion minimization [2, 3, 17] and global and local features [9, 13, 20], allowing incorporating many existing geometries and local descriptors [5, 7, 8]. In particular, it extends the Gromov-Hausdorff framework [2, 3, 19]. Formulating the problem as graph labeling, we use powerful optimization method recently developed for this class of problems which

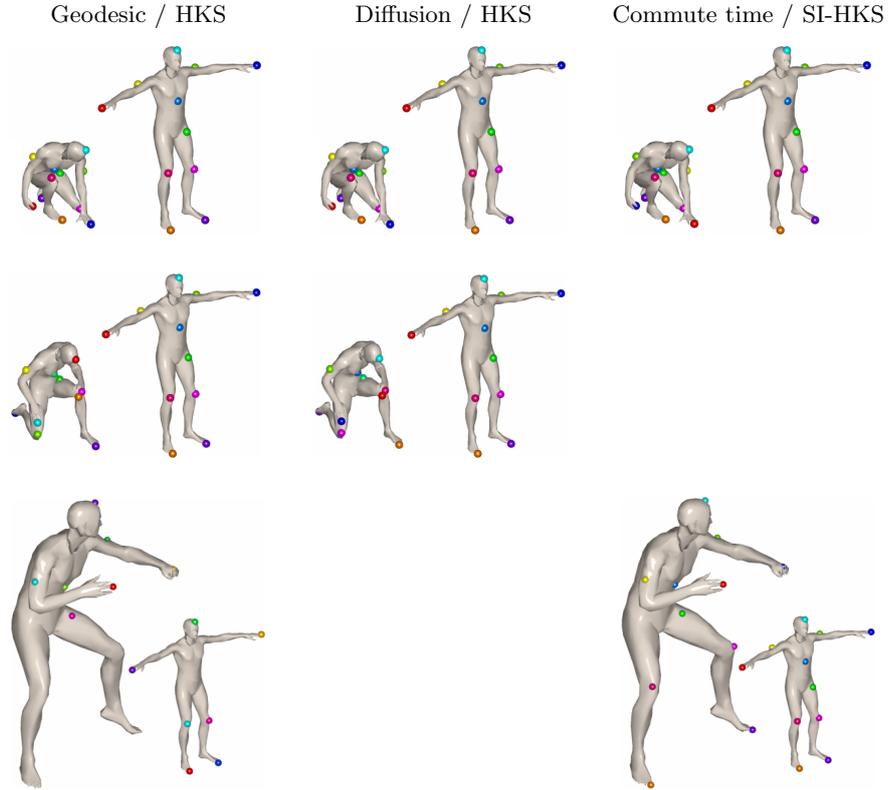


Fig. 1. Invariance to different types of transformations and the choice of the metric/descriptor. Shown is matching between isometric deformations (first row), shapes with different topology (second row), and shapes with different scale (third row), using geodesic metric and HKS descriptors (left), diffusion metric and HKS descriptors (middle), and commute time metric and scale-invariant HKS descriptors (right).

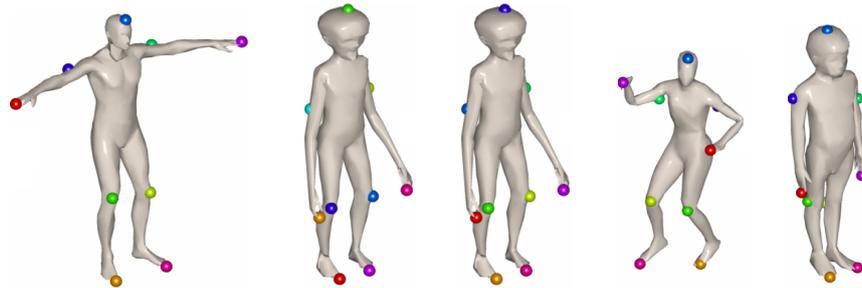


Fig. 2. Matching of an alien shape to the human shape (first column from left) using deterministic (second column) and probabilistic (third column) approaches. Columns four and five: additional probabilistic matching examples.

are known to have favorable convergence properties. Our approach is especially appropriate for the challenging problems of finding similarity and correspondence between non-rigid shapes.

Limitations and extensions. The problem of symmetric correspondences, inherent to all approaches based on intrinsic structures, cannot be resolved without resorting to some side information. There are a few potential cures to this problem. First, providing some initial correspondence between the shapes could be used to restrict the vertex set, ruling out symmetric correspondences. Second, exploiting shape orientation could be used to find orientation-consistent matches. Finally, higher-order distortions (in particular, third-order between triplets of points) can be combined to resolve the symmetry problem [15].

References

1. Berg, A., Berg, T., Malik, J.: Shape matching and object recognition using low distortion correspondences. In: Proc. CVPR. Volume 1. (2005)
2. Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics* **5** (2005) 313–346
3. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* **103** (2006) 1168–1172
4. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Calculus of non-rigid surfaces for geometry and texture manipulation. *Trans. Visualization and Computer Graphics* **13** (2007) 902–913
5. Sun, J., Ovsjanikov, M., Guibas, L.J.: A concise and provably informative multi-scale signature based on heat diffusion. In: *Computer Graphics Forum*. Volume 28. (2009) 1383–1392
6. Ovsjanikov, M., Bronstein, A.M., Bronstein, M.M., Guibas, L.J.: Shape Google: a computer vision approach to invariant shape retrieval. *Proc. NORDIA* (2009)
7. Bronstein, M.M., Kokkinos, I.: Scale-invariant heat kernel signatures for shape recognition. *INRIA Technical Report 7161* (2009)
8. Zaharescu, A., Boyer, E., Varanasi, K., Horaud, R.: Surface Feature Detection and Description with Applications to Mesh Matching. In: *Proc. CVPR*. (2009)
9. Thorstensen, N., Keriven, R.: Non-rigid Shape matching using Geometry and Photometry. *Proc. CVPR* (2009)
10. Rustomov, R.M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In: *Proc. SGP*. (2007) 225–233
11. Mateus, D., Horaud, R.P., Knossow, D., Cuzzolin, F., Boyer, E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. *Proc. CVPR* (2008)
12. Hu, J., Hua, J.: Salient spectral geometric features for shape matching and retrieval. *Visual Computer* **25** (2009) 667–675
13. Dubrovina, A., Kimmel, R.: Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In: *Proc. 3DPVT*. (2010)
14. Lipman, Y., Funkhouser, T.: Möbius voting for surface correspondence. *TOG* **28** (2009)
15. Zeng, Y., Wang, C., Wang, Y., Gu, X., Samaras, D., Paragios, N.: Dense non-rigid surface registration using high-order graph matching. In: *Proc. CVPR*. (2010)

16. Elad, A., Kimmel, R.: Bending invariant representations for surfaces. In: Proc. Computer Vision and Pattern Recognition (CVPR). (2001) 168–174
17. Bronstein, A.M., Bronstein, M.M., Kimmel, R., Mahmoudi, M., Sapiro, G.: A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *IJCV* (2010) 1–21
18. Qiu, H., Hancock, E.R.: Clustering and embedding using commute times. *Trans. PAMI* **29** (2007) 1873–1890
19. Gromov, M.: Structures Métriques Pour les Variétés Riemanniennes. Number 1 in Textes Mathématiques. (1981)
20. Torresani, L., Kolmogorov, V., Rother, C.: Feature correspondence via graph matching: Models and global optimization. In: Proc. ECCV. (2008) 596–609
21. Bertsekas, D.P.: Nonlinear Programming. Athena Scientific (1999)
22. Komodakis, N., Paragios, N., Tziritas, G.: MRF optimization via dual decomposition: Message-passing revisited. In: Proc. ICCV. (2007)
23. Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., van Kaick, O., Tagliasacchi, A.: Deformation-driven shape correspondence. In: Computer Graphics Forum. Volume 27. (2008) 1431–1439
24. Raviv, D., Bronstein, A., Bronstein, M., Kimmel, R.: Symmetries of non-rigid shapes. In: Proc. NRTL. (2007)
25. Memoli, F.: (On the use of gromov-hausdorff distances for shape comparison)
26. Ovsjanikov, M., Sun, J., Guibas, L.J.: Global intrinsic symmetries of shapes. In: Computer Graphics Forum. Volume 27. (2008) 1341–1348
27. Coifman, R.R., Lafon, S., Lee, A.B., Maggioni, M., Nadler, B., Warner, F., Zucker, S.W.: Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS* **102** (2005) 7426–7431
28. Grinstead, C.M., Snell, L.J.: Introduction to Probability. AMS (1998)
29. Gold, S., Rangarajan, A.: A graduated assignment algorithm for graph matching. *Trans. PAMI* **18** (1996) 377–388
30. Hochbaum, D., Shmoys, D.: A best possible heuristic for the k-center problem. *Mathematics of Operations Research* **10:2** (1985) 180–184
31. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Numerical geometry of non-rigid shapes. Springer (2008)
32. Shilane, P., Min, P., Kazhdan, M., Funkhouser, T.: The Princeton shape benchmark. In: Proc. SMI. (2004) 167–178
33. Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. *PNAS* **95** (1998) 8431–8435
34. Meyer, M., Desbrun, M., Schroder, P., Barr, A.H.: Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and Mathematics III* (2003) 35–57