Hierarchical Matching of Non-rigid Shapes

Dan Raviv*, Anastasia Dubrovina*, and Ron Kimmel

Technion - Israel Institute of Technology

Abstract. Detecting similarity between non-rigid shapes is one of the fundamental problems in computer vision. While rigid alignment can be parameterized using a small number of unknowns representing rotations, reflections and translations, non-rigid alignment does not have this advantage. The majority of the methods addressing this problem boil down to a minimization of a distortion measure. The complexity of a matching process is exponential by nature, but it can be heuristically reduced to a quadratic or even linear for shapes which are smooth two-manifolds. Here we model shapes using both local and global structures, and provide a hierarchical framework for the quadratic matching problem.

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

1 Introduction

The paper addresses the problem of finding point-correspondences between non-rigid almost isometric shapes. The correspondence is required for various applications such as shape retrieval, registration, deformation, shape morphing, symmetry, self-similarity detection, to name a few.

A common approach to detect correspondence between shapes differing by a certain class of transformations consists of employing invariant properties under those transformations to formulate a measure of dissimilarity between the shapes, and minimize it in order to find the correct matching. Here we use a matching scheme based on local and global surface properties, namely, local surface descriptors and global metric structures. The proposed method is demonstrated with two different types of metrics - geodesic and diffusion, and different surface descriptors, that include histograms of geodesic and diffusion distances, heat kernel signatures [33], and related descriptors based on the Laplace-Beltrami operator [12].

The main issue we address is the matching complexity. Given two shapes represented by triangular meshes, direct comparison of their pointwise surface descriptors and metric structures is combinatorial in nature (see [24] for the metric comparison problem). Our main contribution is a multi-resolution matching algorithm that can handle a large number of points, and still produces a correspondence consistent in terms of both pointwise and pairwise surface properties.

* Equal contributors.

A.M. Bruckstein et al. (Eds.): SSVM 2011, LNCS 6667, pp. 604–615, 2012.

According to the proposed scheme, at the lowest resolution we solve the exact correspondence problem, up to the approximation introduced by the optimization algorithm. We then propagate this information to higher resolutions thus refining the solution.

The rest of the paper is organized as follows: a brief review of some previous efforts is presented in the next section. Section 2 presents the correspondence problem formulation, followed by Section 3, reviewing relevant mathematical background. Section 4 presents the suggested multi-resolution algorithm. Section 5 contains numerical results, and comparison to the state-of-art algorithms, followed by Section 6 that concludes the paper.

1.1 Non-Rigid Correspondence in a Brief

Zigelman *et al*. [39], and Elad and Kimmel [13] 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*. [17] and Mateus *et al*. [21] suggested alternative isometry-invariant shape representations, obtained by using eigendecomposition of discrete Laplace operators. The Global Point Signature (GPS) suggested by Rustamov [31] 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*. [33], and Ovsjanikov *et al*. [25], to construct their Heat Kernel Signature (HKS) and Heat Kernel Maps, respectively. Zaharescu *et al*. [37] suggested an extension of 2D descriptors for surfaces, and used them to perform the matching. While linear methods, such as [37,25] 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 final results. Hu and Hua [16] used the Laplace-Beltrami operator for matching using prominent features, and Dubrovina and Kimmel [12] 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*. [36] used a similar problem formulation, casted as a graph labeling problem, and experimented with different surface descriptors and metrics.

Memoli and Sapiro [24], Bronstein *et al*. [6], and Memoli [22,23] compared shapes using different approximations of the Gromov-Hausdorff distance [14]. Bronstein *et al*. [7] used the approach suggested in [6] with diffusion geometry, in order to match shapes with topological noise, and Thorstensen and Keriven [35] extended it to handle surfaces with textures. The methods in [24,22,23] 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 [7] 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.* [34] who proposed a randomized algorithm for matching feature points based on geodesic distances between them. Zhang *et al.* [38] 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 [20] 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, [28,26,18] search for a mapping from the shape to itself, and thus are able to detect intrinsic symmetries.

2 Problem Formulation

The problem formulation we use is based on comparison of local and global surface properties that remain approximately invariant under non-rigid ϵ-isometric transformations. Given a shape X, we assume that it is endowed with a *metric* $d_X : X \times X \to \mathbb{R}_+ \cup \{0\}$, measuring distances on X, and pointwise structure $f_X : X \to \mathbb{R}^d$, which is represented by a set of d-dimensional descriptors.

Given two shapes X and Y, endowed with metrics d_X, d_Y and descriptors f_X, f_Y, we would like to find correspondence that best preserves these properties. We denote the correspondence between X and Y by a mapping $\mathcal{C} : X \times Y \to \{0, 1\}$ such that

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

In order to measure how well the correspondence $\mathcal{C}$ preserves the geometric structures of the shapes we use the following dissimilarity function based on global and local shape properties,

$$\mathrm{dis}\left(\mathcal{C}\right) = \mathrm{dis}_{lin}\left(\mathcal{C}\right) + \lambda \cdot \mathrm{dis}_{quad}\left(\mathcal{C}\right). \tag{2}$$

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

$$\mathrm{dis}_{lin}\left(\mathcal{C}\right) = \sum_{x \in X, y \in Y} d_F\left(f_X(x), f_Y(y)\right)\mathcal{C}(x,y), \tag{3}$$

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

$$\text{dis}_{quad}(\mathcal{C}) = \sum_{\substack{x,\tilde{x}\in X \\ y,\tilde{y}\in Y}} (d_X(x,\tilde{x}) - d_Y(y,\tilde{y}))^2 \mathcal{C}(x,y)\mathcal{C}(\tilde{x},\tilde{y}), \tag{4}$$

and it is a quadratic function of $\mathcal{C}$. The parameter $\lambda \geq 0$ (Eq. (2)) determines the relative weight of the linear and the quadratic terms in the total dissimilarity measure. The optimal matching, denoted here by $\mathcal{C}^*$, is obtained by minimizing the dissimilarity measure $\text{dis}(\mathcal{C})$. In order to avoid a trivial solution $\mathcal{C}^*(x,y) = 0, \forall x, y$, we introduce constraints defined by the type of the correspondence we are looking for. For example, when a bijective mapping from X to Y is required, the appropriate constraints on $\mathcal{C}$ are

$$\sum_{x\in X} \mathcal{C}(x,y) = 1, \forall y \in Y, \ \sum_{y\in Y} \mathcal{C}(x,y) = 1, \forall x \in X \tag{5}$$

The resulting optimization problem can be written as

$$\min_{\mathcal{C}} \ \{\text{dis}_{lin}(\mathcal{C}) + \lambda \cdot \text{dis}_{quad}(\mathcal{C})\} \quad s.t. \ (5) \tag{6}$$

Note that the dissimilarity measure $\text{dis}(\mathcal{C})$ is a quadratic function of the correspondence $\mathcal{C}$. In [12], it was shown how to formulate (6) as a quadratic programming problem with binary variables $\mathcal{C}(x,y)$. The optimization problem described above belongs to the class of Integer Quadratic Programming (IQP) problems, also referred to as Quadratic Assignment Problems (QAP), when used with (5). In general, IQP and QAP problems are NP-Hard. Therefore, in order to minimize $\text{dis}(\mathcal{C})$, one has to resort to either some relaxation technique or a heuristic approach (see *e.g.* [27]). While matching points using local structures alone (by setting $\lambda = 0$, for instance) is a linear problem, and thus can be solved efficiently, it can not guarantee global invariance in the presence of noise and symmetries. A better solution can be found by considering global structures. Unfortunately, solving the quadratic assignment problem for a large number of variables is almost infeasible, even after relaxation. In Section 4 we apply a hierarchical approach for calculating an approximate solution of the above optimization problem.

3 Mathematical Background

3.1 Choice of Metric

Differential geometry: Smooth surfaces, also known as *Riemannian manifolds*, are *differential manifolds* equipped with an inner product in the tangent space, which provides geometric notions such as angels, lengths, areas and curvatures without resorting to the ambient space, and are referred to as *intrinsic* measures.

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

$$d_X(x, x') = \inf_{\gamma \in \Gamma(x,x')} \ell(\gamma), \tag{7}$$

where $\Gamma(x, x')$ is the set of all admissible paths between the points x and x' on the surface X,and $\ell(\gamma)$ is the length of the path γ. There exist several numerical methods to evaluate (7) [19,32,5]. We use *fast marching method*, that simulates a wavefront propagation on a triangular mesh, associating the time of arrival of the front with the distance it traveled.

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

$$\left(\Delta_X + \frac{\partial}{\partial t}\right) u(t, x) = 0, \tag{8}$$

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

For compact manifolds, the Laplace-Beltrami operator has a discrete eigen-decomposition of the form

$$\Delta_X \phi_i = \lambda_i \phi_i, \tag{9}$$

where $\lambda_0, \lambda_1, ...$ are eigenvalues and $\phi_0, \phi_1, ...$ are the corresponding eigenfunctions, which construct the heat kernel

$$h_t(x, z) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x) \phi_i(z). \tag{10}$$

The diffusion distance is defined as a cross-talk between two heat kernels [9]

$$\begin{aligned} d^2_{X,t}(x, y) &= \|h_t(x, \cdot) - h_t(y, \cdot)\|^2_{L_2(X)} = \int_X |h_t(x, z) - h_t(y, z)|^2 dz \\ &= \sum_{i=0}^{\infty} e^{-2\lambda_i t} \left(\phi_i(x) - \phi_i(y)\right)^2 . \end{aligned} \tag{11}$$

Since diffusion distances are derived from the Laplace Beltrami operator, they are also intrinsic properties, and, according to [3,11,10], also fulfill the metric axioms.

3.2 Choice of Descriptors

Distance histograms: Given two surfaces X and Y and their metrics d_X and d_Y respectfully, we can evaluate the distances between any two points on each one of the shapes using either choices of metrics. For isometries, a good candidate that matches point $x \in X$ to $y \in Y$ will have similar distances to all other corresponding points. Assuming the surface is well sampled, the distance histograms of corresponding points $x \in X$ and $y \in Y$ have to be similar. Comparison of histograms is a well studied operation. While straight forward bin-to-bin comparison may work, we refer the reader to more robust algorithms such as the *earth moving distances* (EMD) [30].

Heat kernel signatures: Another local descriptor based on the *heat equation*, was presented by Sun *et al*. [33]. It was employed by Bronstein *et al*. [8] for shape retrieval, and was recently adapted to volumes by Raviv *et al.* [29]. Sun *et al.* [33] proposed using the diagonal of the heat kernel $k_t(x, x)$ (10) at multiple scales as a local descriptor, and referred to it as *heat kernel signatures* (HKS). The HKS remains invariant under isometric deformations of X, and it is insensitive to topological noise at small scales. It is also informative in the sense that under certain assumptions one could reconstruct the surface (up to an isometry) from it. Furthermore, the HKS descriptor can be efficiently computed from the eigenfunctions and eigenvalues of the Laplace-Beltrami operator.

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

3.3 Integer Quadratic Programming

A *quadratic program* (QP) is an optimization problem with quadratic objective function and affine constraint functions

$$\min \ \tfrac{1}{2}x^T Ex + q^T x + r \ \ s.t. \ \ Gx \preceq h, \ Ax = b. \tag{12}$$

The above problem is called *convex* when the matrix E is positive semi-definite. The *Integer Quadratic Programming* (IQP) has similar form, with the additional constraint on the variables x: $x_i \in \{0, 1\}$ (binary variables). While convex QP has one global minimum and can be solved efficiently, IQP is an NP-Hard problem. Two common methods are used to solve an QAP problem [4]. The first is a heuristic approach based on a search procedure. For example, [12] used a branch-and-bound procedure to solve the optimization problem in Eq. (6). This approach usually provides good results assuming the local structures are both robust and unique, and there is no intrinsic symmetry. The second approach is based on relaxation. It is a three step solution, consisting of relaxing the integer constraints, solving a continuous optimization problem and projecting the solution back into integers. As expected, this procedure is highly influenced by the initial conditions. As for complexity, the relaxed IQP problem remains NP-Hard. We use branch-and-bound for initial alignment, and then refine it using a continuous optimization technique.

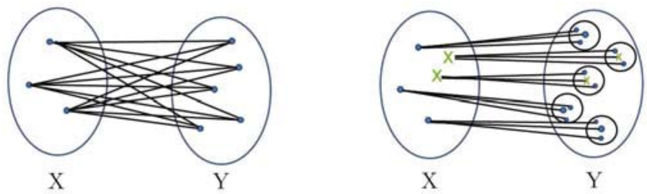

Fig. 1. In the first step (left) we construct a quadratic correspomdence matrix from all points in X into all points in Y. In each iteration (right) we search for possible matches between points in X from the previous iteration (blue circle) and new sampled points in X (green Xs) and their corresponding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

4 Hierarchical Formulation

Solving (6) reveals the main drawback of the quadratic problem formulation. As noted in [12], the dimensionality of the problem allows us to handle up to several dozens of points. Let us assume that X and Y have N and M vertices, respectively. The number of possible correspondences between X and Y is therefore NM, and thus, the dimension of the matrix E in the quadratic problem (12) is $NM \times NM$. Even for a small number of points, e.g. 30, the problem becomes almost infeasible.

Since the problem is not strictly combinatorial by nature, but rather derived from a smooth geometric measure, there should be a way to reduce the complexity. We suggest reducing the high dimensionality of the problem using an iterative scheme. At the first step we follow [12] and solve (6) using a branch-and-bound procedure [2]. Each point $x \in X$ is now matched to a point $c(x) \in Y$ by the mapping c. We denote $y = c(x)$ if $\mathcal{C}(x, y) = 1$. In each iteration we search for the best correspondence between x and $c(x)$ neighborhood, instead of all points $y \in Y$, in a manner similar to [36]. Between iterations we add points $x \in X$ and $y \in Y$ using the 2-optimal *Farthest Point Sampling* (FPS) strategy [15], 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. The solver can now handle up to several hundred of points. Let us further analyze the complexity. We consider the first step to be $\mathbb{O}(N+M)$ as we use a constant (usually around 20) points from each mesh, and only FPS is required, which can be evaluated in linear time. Assuming that each neighborhood in Y consists of K vertices, and a linear growth in each iteration of the matched points from X, then, for the j'th iteration, the quadratic correlation matrix has $jK \times jK$ members which has a complexity of $\mathbb{O}(j^2K^2)$, and the entire iterative framework takes $\mathbb{O}(\Sigma_{j=1}^{N} j^2K^2) = \mathbb{O}(N^3K^2)$. Since each iteration requires a correlation matrix of size j^2K^2, the number of matched points can be significantly higher than the results shown in [12].

5 Results

In this section we provide several matching results obtained using our hierarchical procedure. Figure 2 shows the matching obtained with the proposed framework, combined with different descriptors and metrics, at several hierarchies. The matching was performed using 10 points at the coarse scale, and 30 - 64 points at the finest scale. Figure 2(a) shows the result of matching two cat shapes using geodesic distance histogram descriptors and geodesic distance metric. Figure 2(b) shows the matching result obtained using diffusion distances instead of geodesic ones, and Figure 2(c) - the result obtained using Heat Kernel Signatures [33] and diffusion distances. Note that the last two matchings are in fact reflected ones (follows from the intrinsically symmetric shape matching ambiguity described in [12]). When using the proposed algorithm with Laplace-Beltrami operator-based descriptors [12] and geodesic distances, we were able to obtain both possible correspondences between two cat shapes - the true correspondence and the reflected one. The results are shown in Figure 2(d). As can be seen, all setups provide good results, and we can conclude that the proposed hierarchical framework is independent of the choice of descriptors.

We compared the hierarchical method to [12]'s quadratic matching and [6]'s GMDS framework. Both are based on global structures. Since we followed [12] formulation as our first step, our initial matchings are the same. But, since the complexity of [12] rises rapidly, it can not be used to match more then a few dozen points. In addition, even for a low number of points we have a major quality advantage over [12], since the matched points on the second mesh can move, and are not restricted to the initial sampling. In Figure 3 we see that the ear and the nose of the cat were matched using 10 points, and relocated after several iterations. We also compared the hierarchical matching and the quadratic matching calculation times. The result are shown in Figure 4, for different number of matched points. The quadratic matching succeeded to match only up to 22 points in a reasonable time - less than 4 minutes, while the proposed hierarchical method was able to find 60 matches in shorter time.

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

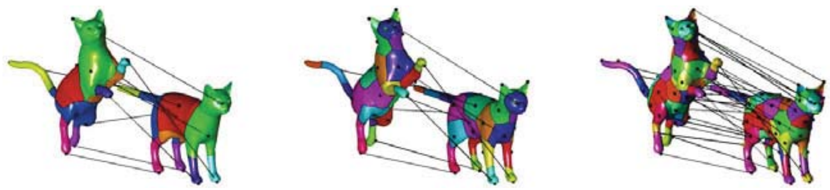

(a) Geodesic distance histograms descriptors and geodesic distance metric

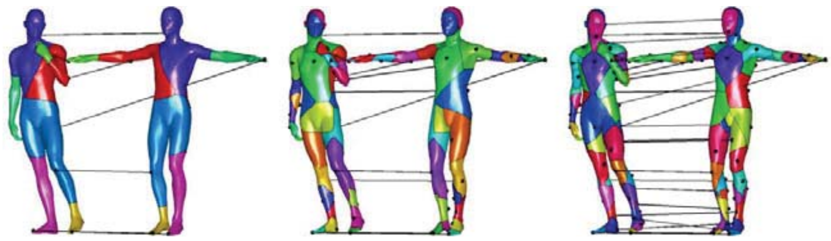

(b) Diffusion distance histograms and diffusion distance metric

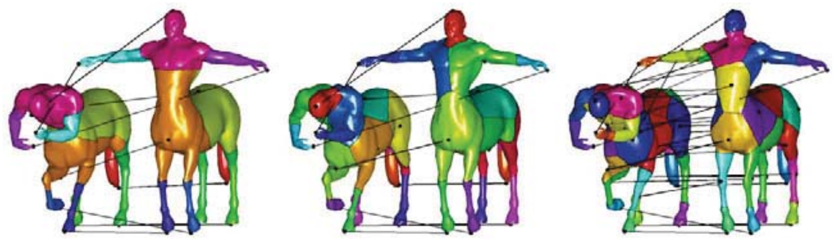

(c) Heat Kernel Signatures and diffusion distance metric

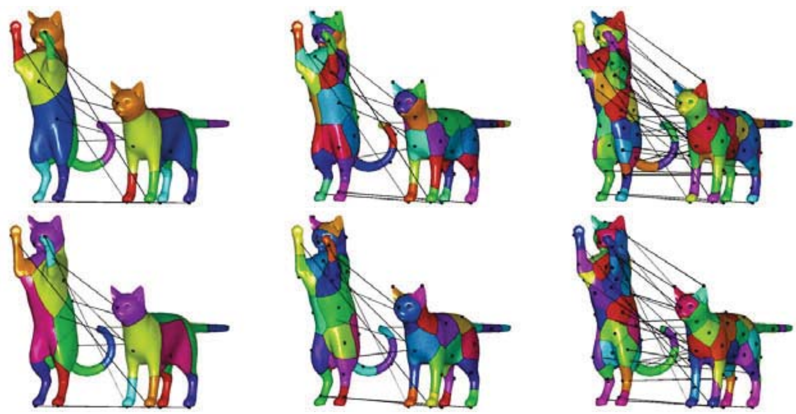

(d) The Laplace-Beltrami operator-based descriptors and geodesic distance metric; the upper row - same orientation correspondence, the lower row - the reflected one.

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

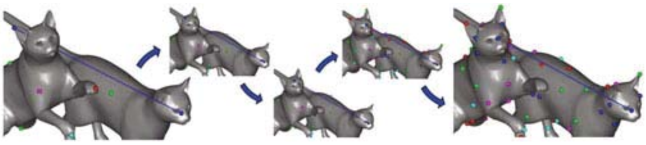

Fig. 3. Using geodesic distances as a global structure, and geodesic based histograms as a local one, the wrong ear-to-nose match gets closer to the correct one during subsequent iterations.

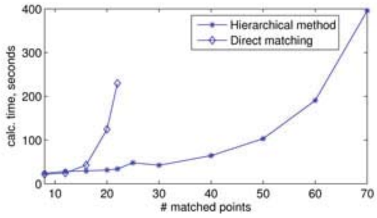

Fig. 4. Graph of calculation time as a function of number of matched points, showing results of the proposed hierarchical method, alongside the quadratic matching algorithm.

6 Conclusions

We presented a hierarchical framework, based on quadratic programming, that solves non-rigid matchings between shapes. While being NP-Hard in general, we solve the assignment problem by taking into account the smooth structure of our shapes using an iterative scheme. We provided numerical results, and compared it to state of art methods.

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

References

1. Anguelov, D., Srinivasan, P., Pang, H.-C., Koller, D., Thrun, S.: The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In: Proc. of the Neural Information Processing Systems (NIPS) Conference, pp. 33–40 (2004)
2. Bemporad, A.: Hybrid Toolbox - User's Guide (2004), `http://www.dii.unisi.it/hybrid/toolbox`
3. Bérard, P., Besson, G., Gallot, S.: Embedding riemannian manifolds by their heat kernel. Geometric and Functional Analysis 4(4), 373–398 (1994)
4. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, Cambridge (2006)

5. Bronstein, A.M., Bronstein, M.M., Devir, Y.S., Kimmel, R., Weber, O.: Parallel algorithms for approximation of distance maps on parametric surfaces. In: Proc. International Conference and Exhibition on Computer Graphics and Interactive Techniques, SIGGRAPH (2007) (submitted)
6. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. Proc. of National Academy of Sciences (PNAS), 1168–1172 (2006)
7. 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. International Journal of Computer Vision, IJCV (2009)
8. Bronstein, A.M., Bronstein, M.M., Ovsjanikov, M., Guibas, L.J.: Shape google: a computer vision approach to invariant shape retrieval. In: Proc. NORDIA (2009)
9. Bronstein, M.M., Bronstein, A.M.: Shape recognition with spectral distances. IEEE Transactions on Pattern Analysis and Machine Intelligence 99 (2010) (preprints)
10. Coifman, R.R., Lafon, S.: Diffusion maps. Applied and Computational Harmonic Analysis 21, 5–30 (2006)
11. 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. Proc. National Academy of Sciences 102(21), 7426–7431 (2005)
12. Dubrovina, A., Kimmel, R.: Matching shapes by eigendecomposition of the laplace-beltrami operator (2010)
13. Elad, A., Kimmel, R.: On bending invariant signatures for surfaces. IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI) 25(10), 1285–1295 (2003)
14. Gromov, M.: Structures metriques pour les varietes riemanniennes. Textes Mathematiques (1) (1981)
15. Hochbaum, D.S., Shmoys, D.B.: A best possible heuristic for the k-center problem. Mathematics of Operations Research 10(2), 180–184 (1985)
16. Hu, J., Hua, J.: Salient spectral geometric features for shape matching and retrieval. Vis. Comput. 25(5-7), 667–675 (2009)
17. Jain, V., Zhang, H., Van Kaick, O.: Non-rigid spectral correspondence of triangle meshes. International Journal on Shape Modeling 13, 101–124 (2007)
18. Kim, V., Lipman, Y., Chen, X., Funkhouser, T.: Mobius transformations for global intrinsic symmetry analysis. In: Proc. Eurographics Symposium on Geometry Processing, SGP (2010)
19. Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. Proceedings of the National Academy of Sciences (PNAS) 95(15), 8431–8435 (1998)
20. Lipman, Y., Funkhouser, T.: Mobius voting for surface correspondence. ACM Transactions on Graphics (Proc. SIGGRAPH) 28(3) (August 2009)
21. Mateus, D., Horaud, R.P., Knossow, D., Cuzzolin, F., Boyer, E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2008)
22. Mémoli, F.: On the use of gromov-hausdorff distances for shape comparison. In: Point Based Graphics 2007, pp. 81–90 (September 2007)
23. Mémoli, F.: Spectral gromov-wasserstein distances for shape matching. In: Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment (ICCV Workshop, NORDIA 2009) (2009)
24. Mémoli, F., Sapiro, G.G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. Found. Comput. Math. 5(3), 313–347 (2005)

25. Ovsjanikov, M., Mérigot, Q., Mémoli, F., Guibas, L.: One point isometric matching with the heat kernel. In: Eurographics Symposium on Geometry Processing, SGP (2010)
26. Ovsjanikov, M., Sun, J., Guibas, L.: Global intrinsic symmetries of shapes. Computer Graphics Forum 27(5), 1341–1348 (2008)
27. Pardalos, P.M., Rendl, F., Wolkowicz, H.: The quadratic assignment problem: A survey and recent developments. In: Proc. of the DIMACS Workshop on Quadratic Assignment Problems. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 16, pp. 1–42 (1994)
28. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Full and partial symmetries of non-rigid shapes. International Journal of Computer Vision, IJCV (2010)
29. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Volumetric heat kernel signatures. In: Proc. 3D Object recognition (3DOR), Part of ACM Multimedia (2010)
30. Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover's distance as a metric for image retrieval. International Journal of Computer Vision 40(2), 99–121 (2000)
31. Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation inavriant shape representation. In: Proc. of SGP, pp. 225–233 (2007)
32. Spira, A., Kimmel, R.: An efficient solution to the eikonal equation on parametric manifolds. Interfaces and Free Boundaries 6(4), 315–327 (2004)
33. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: Proc. Eurographics Symposium on Geometry Processing, SGP (2009)
34. Tevs, A., Bokeloh, M., Wand, M., Schilling, A., Seidel, H.-P.: Isometric registration of ambiguous and partial data. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2009)
35. Thorstensen, N., Keriven, R.: Non-rigid shape matching using geometry and photometry. In: Zha, H., Taniguchi, R.-i., Maybank, S. (eds.) ACCV 2009. LNCS, vol. 5996, pp. 644–654. Springer, Heidelberg (2010)
36. Wang, C., Bronstein, M.M., Paragios, N.: Discrete minimum distortion correspondence problems for non-rigid shape matching. Technical report, INRIA Research Report 7333, Mathématiques Appliquées aux Systèmes, École Centrale Paris (2010)
37. Zaharescu, A., Boyer, E., Varanasi, K., Horaud, R.P.: Surface feature detection and description with applications to mesh matching. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2009)
38. Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., van Kaick, O., Tagliasacchi, A.: Deformation-driven shape correspondence. Computer Graphics Forum (Proc. of SGP) 27(5), 1431–1439 (2008)
39. Zigelman, G., Kimmel, R., Kiryati, N.: Texture mapping using surface flattening via multi-dimensional scaling. IEEE Trans. on Visualization and Computer Graphics 8(2), 198–207 (2002)