

SHREC 2010: robust correspondence benchmark

A. M. Bronstein^{†1}, M. M. Bronstein^{†1}, U. Castellani^{†3}, A. Dubrovina⁶, L. J. Guibas^{†2}, R. P. Horaud⁴, R. Kimmel¹,
D. Knossow⁴, E. von Lavante⁴, D. Mateus⁴, M. Ovsjanikov^{†5}, A. Sharma⁴

¹Department of Computer Science, Technion – Israel Institute of Technology

²Department of Computer Science, Stanford University

³Department of Computer Science, University of Verona

⁴INRIA Grenoble Rhône-Alpes

⁵Institute for Computational and Mathematical Engineering, Stanford University ⁶Department of Electrical Engineering, Technion – Israel Institute of Technology

Abstract

SHREC'10 robust correspondence benchmark simulates a one-to-one shape matching scenario, in which one of the shapes undergoes multiple modifications and transformations. The benchmark allows evaluating how correspondence algorithms cope with certain classes of transformations and what is the strength of the transformations that can be dealt with. The present paper is a report of the SHREC'10 robust correspondence benchmark results.

Categories and Subject Descriptors (according to ACM CCS): H.3.2 [Information storage and retrieval]: Information Search and Retrieval—Retrieval models I.2.10 [Artificial intelligence]: Vision and Scene Understanding—Shape

1. Introduction

Correspondence and similarity are two intimately related problems in shape analysis. Defining optimal correspondence based on some structure preservation criterion, one can obtain a criterion of shape similarity as the amount of structure distortion. Finding correspondence between two shapes that would be invariant to a wide variety of transformations is thus a cornerstone problem in many approaches for shape similarity and retrieval.

SHREC'10 robust correspondence benchmark simulates one-to-one shape matching, in which the shapes to be matched are modifications of the same shape. The benchmark allows evaluating how algorithms cope with certain classes of transformations and what is the strength of the transformations that can be dealt with.

2. Data

The dataset used in this benchmark was from the TOSCA shapes [BBK08], available in the public domain. The shapes were represented as triangular meshes with approximately 10,000–50,000 vertices.

The dataset consisted of 3 shapes, with simulated transformations applied to them. For each null shape, transformations were split into 9 classes shown in Figure 1: *isometry* (non-rigid almost inelastic deformations), *topology* (welding of shape vertices resulting in different triangulation), *micro holes* and *big holes*, *global* and *local scaling*, additive Gaussian *noise*, *shot noise*, *down sampling* (less than 20% of original points). Shapes in classes *isometry*, *scaling*, *local scaling*, *noise*, and *shot noise* has identical triangulation.

In each class, the transformation appeared in five different versions numbered 1–5. In all shape categories except *scale* and *isometry*, the version number corresponded to the transformation strength levels: the higher the number, the stronger the transformation (e.g., in noise transformation, the noise variance was proportional to the strength number). For *scale* transformations, the levels 1–5 corresponded to scaling by the factor of 0.5, 0.875, 1.25, 1.625, and 2. For the *isometry* class, the

[†] Organizer of the SHREC track. All organizers and participants are listed in alphabetical order. For any information about the benchmark, contact mbron@cs.technion.ac.il

numbers did not reflect transformation strength. The total number of transformations per shape was 45, and the total dataset size was 138. The dataset is available at http://tosca.cs.technion.ac.il/book/shrec_correspondence.html.

3. Evaluation methodology

The participants were asked to provide a set of corresponding pairs $\mathcal{C}(X, Y) = \{(y_k, x_k)\}_{k=1}^M$; $M \leq |Y|$ between all the transformed shapes (denoted here by Y) and the null shapes (X); a total of 135 correspondence sets). Correspondence was referred to as *dense* if $M = |Y|$ and *sparse* if $M \ll |Y|$.

For each transformed shape Y in the dataset, groundtruth dense correspondence to the null shape X was given in the form of pairs of points $\mathcal{C}_0(X, Y) = \{(y'_k, x_k)\}_{k=1}^{|Y|}$. Since all the shapes has reflection intrinsic symmetries (e.g. flipping left and right sides), a set of symmetric corresponding points $\bar{\mathcal{C}}_0(X, Y) = \{(y''_k, x_k)\}_{k=1}^{|Y|}$ was also computed.

The quality of the correspondence \mathcal{C} was measured as the average geodesic distance from the groundtruth correspondence, taking into consideration possible intrinsic symmetry,

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

where $(x_k, y_k) \in \mathcal{C}(X, Y)$, $(x'_k, y_k) \in \mathcal{C}_0(X, Y)$, $(x''_k, y_k) \in \bar{\mathcal{C}}_0(X, Y)$, and d_X denotes the *geodesic distance* measured on the shape X .

4. Methods

Three families of methods were evaluated in this benchmark: spectral matching [MHK*08] (denoted hereinafter as SM for notation brevity); shape matching using Laplace-Beltrami eigenfunctions [DK10] (denoted LB), and generalized multidimensional scaling [BBK06b] (GMDS).

SM1–2: Spectral matching

The algorithm based on the Laplacian eigenvector alignment and unsupervised point registration algorithms by Mateus et al. [MHK*08], with adaptations for the registration of mesh data and point datasets from the work by Knossow et al. [KSMH09] and Horaud et al. [HFY*10]. The method consists of three steps: at first the Laplacian embedding of the two shapes to be matched is computed; then a permutation matrix \mathbf{P} and a sign matrix \mathbf{S} which matches the first k eigenvectors of these shapes to each other are computed. This is done by computing the distance of the eigenvector *signatures* between the shapes and finding the minimal cost assignment between these vectors. The final step is the point registration of these two aligned embeddings by using the EM-algorithm.

Laplacian embedding: The correspondence problem is

cast as graph matching. A shape X is represented as a connected undirected weighted graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathbf{A}\}$ where $\mathcal{V}(\mathcal{G}) = \{x_1, \dots, x_N\}$ is the vertex set, $\mathcal{E}(\mathcal{G}) = \{e_{ij}\}$ is the edge set, and the entries of the weighted adjacency matrix \mathbf{A} are $a_{ij} > 0$ whenever two vertices are linked by an edge, and $a_{ij} = 0$ otherwise. A Gaussian kernel $a_{ij} = \exp(-d^2(x_i, x_j)/\sigma^2)$ is used with d being the local shortest path distances between vertices and σ is selected in such a way that weights lie in a small interval around $a = 0.5$. The degree matrix $\mathbf{D} = \text{Diag}[d_i]$ is defined as $d_i = \sum_j a_{ij}$. The Laplacian of \mathcal{G} is computed following two methods. In SM1, the *unnormalized Laplacian*: $\mathbf{L} = \mathbf{D} - \mathbf{A}$ is used, while in SM2, the *normalized Laplacian* $\mathbf{L} = \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$ is employed. Denoting by $\mathbf{L}u = \lambda u$ the eigenvalue problem of \mathbf{L} with $\Lambda = \text{Diag}[\lambda_2 \dots \lambda_{k+1}]$, and $\mathbf{U}^k = [u_2 \dots u_{k+1}]$, the $N \times k$ matrix formed with the k smallest non-zero eigenvectors of \mathbf{L} . Given these definitions, Laplacian embeddings \mathbf{U}_X^k and \mathbf{U}_Y^k are computed for shapes X and Y , respectively.

Eigenvector matching: For the eigenvector matching the combined permutation and sign flip matrix \mathbf{R}^k is sought for which aligns \mathbf{U}_X^k and \mathbf{U}_Y^k . For each eigenvector of \mathbf{U}_X^k and \mathbf{U}_Y^k a *signature* $h(\mathbf{u}_X^i)$ is computed, and given such signatures a measure of dissimilarity between them $C(h(\mathbf{u}), h(\mathbf{v}))$ is evaluated. The minimum of this cost function for all pairs of eigenfunctions $(\mathbf{u}_X^i, \pm \mathbf{u}_X^j)$ will give \mathbf{R}^k . The optimal solution to this assignment problem is found with the Hungarian algorithm. For the eigenvector signature function $h(\mathbf{u})$ for each eigenvector its density function is computed with a Gaussian kernel. The dissimilarity function $C(a, b)$ is simply the L_2 norm between the two given density functions.

Point registration: In the present point registration formulation vertices of X are treated as observations and vertices of Y as centers of normally distributed clusters. In addition to these Gaussian clusters, there is an outlier cluster with uniform distribution. By introducing latent variables that assign each observation to a cluster, the point registration problem is formulated in the framework of the EM-algorithm solving for a global orthogonal transformation \mathbf{R}^* . The details for the formulation of this framework can be found in [MHK*08].

LB1–2: Laplace-Beltrami matching

The proposed method for finding correspondence between non-rigid isometric shapes utilizes surface descriptors based on the eigendecomposition of the Laplace-Beltrami operator [DK10]. Like the Laplace-Beltrami operator, those descriptors are invariant to isometric transformations, and as such are suitable for the correspondence detection. In order to make the correspondence robust to small perturbations in the values of the eigenfunctions, due to numerical and approximation errors, the algorithm combines the surface descriptors and the geodesic distances measured on the shapes when calculating the correspondence quality. The above results in a quadratic optimization problem formulation for




Figure 1: Transformations of the human shape used as queries (shown in strength 5, left to right): null, isometry, topology, sampling, local scale, scale, holes, micro holes, noise, shot noise.

correspondence detection, and its minimizer is the best possible correspondence.

When using the eigenfunctions of the Laplace-Beltrami operator, one encounters a sign ambiguity problem that follows from the fact that the eigenfunctions are defined up to a sign. Hence, additional sign estimation phase is required prior to the construction of the descriptors. Moreover, Ovsjanikov *et al.* [OSG08] showed that the eigenfunctions of the Laplace-Beltrami operator defined on intrinsically symmetric shape are also symmetric functions. Hence, during the sign estimation phase it is possible to find more than one sign sequence that aligns the two sets of the eigenfunctions corresponding to two such shapes. Since the geodesic distances do not provide us with information regarding shape orientation, the minimizer of the optimization problem described above is not unique. This implies existence of multiple equally good correspondences between two instances of intrinsically symmetric shapes. Specifically, all the shapes in the proposed data set have one intrinsic symmetry, therefore there exist two possible correspondences. As showed in [DK10], the proposed algorithm is able to find both.

Two settings of the method were used in the benchmark: LB1 used the eigenfunctions of the graph Laplacian, and LB2 used the cotangent weight scheme [MDPB02] instead. The latter discretization is known to be less sensitive to sampling and triangulation changes.

In order to make the descriptors robust to scaling, they were normalized by the square roots of the corresponding eigenvalues, as proposed by Rustamov in [Rus07]. The geodesic distances measured on the two shapes were also normalized by their maximal values. The results with high correspondence error are due to eigenfunctions switching that follows from discretization and numerical errors - a feature that the current algorithm was not designed to deal with.

GMDS: Generalized multidimensional scaling

The generalized multidimensional scaling (GMDS) algorithm was introduced in [BBK06b] being one of the few state-of-the-art methods for deformable shape matching at

that time. GMDS is an MDS-like problem, computing correspondence between two shapes by trying to embed one shape into another with minimum distortion (referred to as *stress* in MDS literature [BG97]). Given a fixed set of points $x_1, \dots, x_N \in X$, GMDS attempts to find a set of corresponding points on Y in barycentric coordinates $y_i = (t_i, \mathbf{u}_i)$ (where $t_i \in T(Y)$ is a triangle index, and $u_{ij} \in [0, 1]$, $\sum_j u_{ij} = 1$ is a vector of barycentric weights) by minimizing the stress

$$\min_{\{t_i; \mathbf{u}_i\}_{i=1}^N} \sum_{i>j} (d_X(x_i, x_j) - d_Y((t_i; \mathbf{u}_i), (t_j; \mathbf{u}_j)))^2, \quad (1)$$

where $d_X(x_i, x_j)$ is a pre-computed geodesic distance between x_i and x_j on X , and d_Y is interpolated from pre-computed geodesic distances between points on Y .

While the stress function is highly non-convex, GMDS optimization is performed in a multi-resolution manner and in practice shows good convergence is initialized sufficiently close to the global minimizer [BBK06a]. In this benchmark, branch-and-bound initialization proposed in [RBBK07] was used.

5. Results

Tables 1–5 show the performance of the tested algorithms, given in terms of average geodesic distance to ground truth correspondence. SM1–2 (Tables 1–2) provided dense correspondence between all the shape points; other methods provided sparse correspondence between 10–50 points.

SM1–2 show ideal performance (0 average geodesic distance) in isometry, scale, local scale, noise, and shot noise classes. This is explained by the fact that shapes in these classes have the same triangulation, resulting in error-free graph matching. The performance of SM drops significantly (to average geodesic distance of 22.63 in SM1 and 19.17 in SM2) in the sampling class, which indicates potential sensitivity of the graph Laplacians to shape triangulation. This phenomenon is consistently observed in LB1–2, which shows dramatically better performance when using geometrically consistent discretization of the Laplacian using cotangent weights (LB2, average geodesic distance of 16.65) compared to graph Laplacian (LB1, 67.2). LB2 performs the

Transform.	Strength				
	1	≤2	≤3	≤4	≤5
Isometry	0.00	0.00	0.00	0.00	0.00
Topology	6.89	7.92	7.92	8.04	8.41
Holes	7.26	8.39	9.34	9.47	12.47
Micro holes	0.37	0.39	0.44	0.45	0.49
Scale	0.00	0.00	0.00	0.00	0.00
Local scale	0.00	0.00	0.00	0.00	0.00
Sampling	11.43	13.32	15.70	18.76	22.63
Noise	0.00	0.00	0.00	0.00	0.00
Shot noise	0.00	0.00	0.00	0.00	0.00
Average	2.88	3.34	3.71	4.08	4.89

Table 1: Performance of SM1: spectral graph matching dense correspondence algorithm using unnormalized graph Laplacian (average geodesic distance to groundtruth correspondence).

Transform.	Strength				
	1	≤2	≤3	≤4	≤5
Isometry	0.00	0.00	0.00	0.00	0.00
Topology	5.96	6.76	7.14	7.55	8.13
Holes	5.17	5.55	6.05	6.44	10.32
Micro holes	0.68	0.70	0.79	0.79	0.83
Scale	0.00	0.00	0.00	0.00	0.00
Local scale	0.00	0.00	0.00	0.00	0.00
Sampling	10.51	12.08	13.65	15.58	19.17
Noise	0.00	0.00	0.00	0.00	0.00
Shot noise	0.00	0.00	0.00	0.00	0.00
Average	2.48	2.79	3.07	3.37	4.27

Table 2: Performance of SM2: spectral graph matching dense correspondence algorithm using normalized graph Laplacian (average geodesic distance to groundtruth correspondence).

best in the sampling transformation class. Smallest sensitivity to topology, holes is achieved by SM2. Smallest sensitivity to micro holes is achieved by SM1. Significantly higher error of LB1–2 and GMDS in these transformation classes point to the known fact of lower sensitivity of commute-time and diffusion distances compared to geodesic ones.

6. Conclusions

On the average, spectral graph matching approaches (SM1–2) show the best performance. Some cases are idealistic due to the use of identical triangulations in the dataset. Best resilience to sampling density change is obtained by LB2 using cotangent weight discretization of the Laplace-Beltrami operator. As a general conclusion, we note the methods based on diffusion geometry (SM1–2) are less sensitive to topological noise and scaling compared to those based on geodesic distances (LB, GMDS).

A more detailed version of this report presenting additional details and experiments will be published separately.

Transform.	Strength				
	1	≤2	≤3	≤4	≤5
Isometry	8.36	9.01	8.73	9.09	8.68
Topology	32.22	36.20	38.26	64.64	72.24
Holes	7.89	9.10	24.78	20.90	23.41
Micro holes	8.77	8.45	7.88	7.75	7.85
Scale	9.13	9.50	10.63	10.01	9.53
Local scale	8.03	7.66	8.70	8.93	9.15
Sampling	6.77	44.47	57.05	72.11	67.20
Noise	6.67	6.69	6.97	7.59	7.88
Shot noise	7.67	8.24	8.09	7.95	8.99
Average	10.61	15.48	19.01	23.22	23.88

Table 3: Performance of LB1: correspondence algorithm based on graph Laplacian eigenfunctions (average geodesic distance to groundtruth correspondence). Average number of corresponding points: 10.

Transform.	Strength				
	1	≤2	≤3	≤4	≤5
Isometry	15.43	12.55	11.10	11.41	12.78
Topology	18.15	26.20	46.36	52.80	54.34
Holes	7.83	18.83	17.88	17.47	20.64
Micro holes	16.68	13.16	13.42	12.19	11.32
Scale	18.47	17.54	18.42	17.86	16.24
Local scale	10.40	27.02	29.21	40.00	65.38
Sampling	26.75	20.14	21.56	18.87	16.65
Noise	15.91	11.70	28.98	36.72	41.40
Shot noise	9.93	16.77	20.02	20.03	19.49
Average	15.51	18.21	22.99	25.26	28.69

Table 4: Performance of LB2: correspondence algorithm based on Laplace-Beltrami eigenfunctions computed using cotangent weight discretization (average geodesic distance to groundtruth correspondence). Average number of corresponding points: 10.

Transform.	Strength				
	1	≤2	≤3	≤4	≤5
Isometry	15.73	13.67	28.12	31.83	31.40
Topology	77.68	70.60	60.94	62.03	68.34
Holes	13.28	32.74	26.41	32.54	36.23
Micro holes	38.20	24.63	28.95	30.91	32.53
Scale	75.18	60.50	52.06	48.83	50.33
Local scale	41.88	42.93	32.10	39.94	40.81
Sampling	38.41	27.62	21.69	25.44	27.94
Noise	16.25	28.95	32.84	35.70	32.47
Shot noise	42.65	29.32	34.03	29.34	31.86
Average	39.92	36.77	35.24	37.40	39.10

Table 5: Performance of GMDS: generalized multidimensional scaling (average geodesic distance to groundtruth correspondence). Average number of corresponding points: 50.

Transform.	Strength		
	1	≤3	≤5
<i>Isometry</i>	SM2,SM1	SM2,SM1	SM2,SM1
<i>Topology</i>	SM2	SM2	SM2
<i>Holes</i>	SM2	SM2	SM2
<i>Micro holes</i>	SM1	SM1	SM1
<i>Scale</i>	SM2,SM1	SM2,SM1	SM2,SM1
<i>Local scale</i>	SM2,SM1	SM2,SM1	SM2,SM1
<i>Sampling</i>	LB1	SM2	LB2
<i>Noise</i>	SM2,SM1	SM2,SM1	SM2,SM1
<i>Shot noise</i>	SM2,SM1	SM2,SM1	SM2,SM1
Average	SM2,SM1	SM2,SM1	SM2,SM1

Table 6: Winning algorithms across transformation classes and strengths. *LB1*=correspondence algorithm based on graph Laplacian eigenfunction, *LB2*=correspondence algorithm based on eigenfunction of Laplace-Beltrami operator computed using cotangent weights, *SM2*=spectral graph matching using normalized graph Laplacian, *SM1*=spectral graph matching using unnormalized graph Laplacian.

- [RBBK07] RAVIV D., BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Symmetries of non-rigid shapes. In *Proc. NRTL* (2007). 3
- [Rus07] RUSTAMOV R. M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP* (2007), pp. 225–233. 3

References

- [BBK06a] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Sci. Comp.* 28, 5 (2006), 1812–1836. 3
- [BBK06b] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103, 5 (2006), 1168–1172. 2, 3
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical geometry of non-rigid shapes*. Springer, 2008. 1
- [BG97] BORG I., GROENEN P.: *Modern multidimensional scaling - theory and applications*. Springer, 1997. 3
- [DK10] DUBROVINA A., KIMMEL R.: Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In *Proc. 3DPVT* (2010). 2, 3
- [HFY*10] HORAUD R. P., FORBES F., YGUEL M., DEWAELE G., ZHANG J.: Rigid and articulated point registration with expectation conditional maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2010). in press. 2
- [KSMH09] KNOSSOW D., SHARMA A., MATEUS D., HORAUD R. P.: Inexact matching of large and sparse graphs using laplacian eigenvectors. In *Proc. Workshop on Graph-based Representations in Pattern Recognition* (Venice, Italy, May 2009), LNCS 5534, Springer. 2
- [MDPB02] MEYER M., DESBRUN M., P. S., BARR A.: Discrete differential geometry operators for triangulated 2-manifolds. In *International Workshop on Visualization and Mathematics* (2002). 3
- [MHK*08] MATEUS D., HORAUD R. P., KNOSSOW D., CUZZOLIN F., BOYER E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR* (2008). 2
- [OSG08] OVSIANIKOV M., SUN J., GUIBAS L. J.: Global intrinsic symmetries of shapes. *Computer Graphics Forum* 27, 5 (July 2008), 1341–1348. 3