Eurographics Workshop on 3D Object Retrieval (2012)
M. Spagnuolo, M. Bronstein, A. Bronstein, and A. Ferreira (Editors)

Spectral Analysis Driven Sparse Matching of 3D Shapes

Tal Darom[†] and Yosi Keller[‡]

Abstract

In this work we present an approach for matching three-dimensional mesh objects related by isometric transformations and scaling. We propose to utilize the Scale invariant Scale-DoG detector and Local Depth SIFT mesh descriptor, to derive a statistical voting-based scheme to robustly estimate the scale ratio between the registered meshes. This paves the way to formulating a novel non-rigid mesh registration scheme, by matching sets of sparse salient feature points using spectral graph matching. The resulting approach is shown to compare favorably with previous state-of-the-art approaches in registering meshes related by partial alignment, while being a few orders of magnitude faster.

1. Introduction

The registration of three-dimensional (3D) objects is a contemporary research challenge that has been given considerable attention in recent years [DK10, ZWW*10, SHM12]. A class of 3D registration schemes relates to point matching in images by assuming that the registered meshes are related by parametric transformations models (similarity, affine, projective), that can be estimated by robust statistical schemes such as RANSAC. Others assume local isometry that can be encoded by graph algorithms [LH05].

Energy minimization schemes formulate the points correspondence via an objective functions, resulting in discrete optimization problems [DK10], while other apply high-order potentials and the corresponding graph matching techniques [ZWW*10]. In dense matching we aim to align all of the common vertices in the registered meshes. Such schemes may operate in the mesh domain by formulating continuous variational problems [BBK06b], or embed the meshes into a space where they can be aligned [SHM12, BBK*10]. Others extend the sparse matching results to compute a dense matching [ZWW*10].

In this work we propose two core contributions:

First, we propose to utilize spectral graph matching to encode and recover local isometries agglomerated over the mesh objects. By applying our approach to isometries with respect to geodesic distances, we derive a non-parametric mesh matching scheme that is able to recover non-rigid alignments. The downside of this approach is its inability to handle transformations consisting of scale changes. For that we propose our **Second** contribution, that is a straightforward approach for estimating the global scaling between a pair of mesh objects, based on an analysis of the local scales of weakly-corresponding feature points. In particular, the proposed scheme is applicable to meshes related by partial matching.

This paper is organized as follows. In section 2, we survey previous results on mesh registration, and propose our computationally efficient registration approach in Section 3. In Section 4 we show how to estimate global scale differences between 3D shapes using a histogram-based scale estimation approach, and apply it to mesh registration to derive a unified scheme. We verify and exemplify the use of our approach in Section 5, while concluding remarks and future extensions are discussed in Section 6.

2. Related work

The registration of images, 3D models or more generally, sets of points is a fundamental problem in Computer Vision and has been extensively researched [LH05, CFSV04].

The Iterative Closest Point algorithm proposed by Besl and Mckay [BM92], matches each point on one surface to its closest point on the other mesh, and computes a parametric motion model. Its draw back is the lack of global convergence, thus requiring an initial estimate of the relative transformation.

Spectral embedding was also used by Sharma *et.*

[†] Faculty of Engineering, Bar Ilan University, Ramat Gan, Israel. tal.darom@gmail.com.

[‡] Faculty of Engineering, Bar Ilan University, Ramat Gan, Israel. yosi.keller@gmail.com.

DOI: 10.2312/3DOR/3DOR12/059-062

al. [SHM12] that first align the embeddings using histograms related to the eigenvectors of the spectral embedding. Dense matching is computed using an EM based point matching algorithm. Bronstein *et. al.* [BBK06b] proposed to compute the dense alignment between meshes by extending the MDS algorithm using the Generalized Multi-Dimensional scaling (GMDS). A different approach to surface registration is the minimization of the Gromov-Hausdorff distance between the surfaces, that quantifies the discrepancy in pairwise distances between corresponding vertices in meshes. It was used by Mémpli and Sapiro [MS05] to derive a scale invariant geometry-based mesh similarity measure. Bronstein *et. al.* [BBK06a] applied the Gromov-Hausdorff within a multi-resolution formulation to compute it efficiently and robustly, and extended it by utilizing the diffusion distances instead of geodesic ones. In this work we follow these methods in making use of the fact that geodesic distances between matching points on a surface are maintained by isometric transforms.

The minimization of the Gromov-Hausdorff distance is the equivalent of the pairwise matching problem that was analyzed by Leordeanu and Hebert [LH05], who formulated it as a discrete optimization problem that is NP-hard. The solution is approximated by spectral relaxation, followed by a greedy discetization scheme. The drawback of this approach is its sensitivity to scale changes between the sets of points. Dubrovina and Kimmel [DK10] applied a pairwise matching approach to mesh alignment, by using the eigenvectors of the Laplace-Beltrami operator as local descriptors, and defining a matching cost that quantifies both descriptor similarity, and the difference of corresponding geodesic distances in both shapes. Our work relates to pairwise matching schemes, as only we utilize pairwise affinities that encode purely geometric constraints.

3. Spectral Point Matching

In this section we propose a novel approach to non-parametric mesh alignment that is based on spectral graph matching, for which we utilize geodesic distances that are invariant to isometric transformations.

Let $M=(V,F)$ be a 3D mesh object, with a set of N_V vertices $V=\{v_i\}$, and a set of N_F triangle faces $F=\{f_i\}$, where a shape is represented by a sparse set of local LD-SIFT features $\mathbf{x}=\left\{\mathbf{x}_j\right\}_1^J$, each characterizing the vicinity of a particular feature point $v_i \in M$ [DK12]. d_{ij}^k denotes the geodesic distance between the points v_i and v_j in a model M_k.

We consider the joint matching of *two pairs* of points $v_i \in M_0$ to $v_{i'} \in M_1$ and $v_j \in M_0$ to $v_{j'} \in M_1$. In order to make the matching problem computationally tractable, the possible assignments per each vertex $\{v_i\}_{i=1}^N \in M_0$, are reduced to the L most similar feature points in M_1, in terms of LD-SIFT descriptor distance. Figure 1 depicts v_i and v_j in M_0 and three matching candidates for each of them.

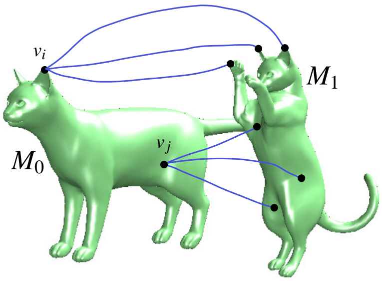

Figure 1: *Mesh alignment as a matching problem. The feature points $v_i, v_j \in M_0$ can each be matched to $L=3$ candidates, resulting in* 9 *affinity values.*

For each such pair matching candidate we define a pairwise affinity measure

$$\tilde{a}_{i,i',j,j'} = \exp\left(-\frac{|d_{ij}^0 - d_{i'j'}^1|}{\sigma}\right), \tag{1}$$

where the parameter $\sigma > 0$ sets the tolerance of deviation from perfect isometric matching, and we used $\sigma = 6$ throughout this work. The affinity function adheres to $\tilde{a}_{i,j,i',j'} \approx 1$ for mutually valid point assignments, and $\tilde{a}_{i,j,i',j'} \approx 0$ for invalid ones. Thus, the mesh alignment is formulated as a pairwise matching problem and is solved by spectral matching [LH05]. An affinity matrix $\mathbf{A} \in \mathbb{R}^{N\cdot L\times N\cdot L}$ is computed using Eq. 1, such that

$$\begin{aligned} & a_{(i-1)L+r_{i,i'},(j'-1)L+r_{j,j'}} = \tilde{a}_{i,i',j,j'}, \\ & \forall i,j = 1..N, \quad r_{i,i'}, r_{j,j'} = 1..L \end{aligned} \tag{2}$$

where N is the number of feature points detected in M_0 and $r_{i,i'}$ is the ranking of the descriptor similarity between v_i and $v_{i'}$. This also implies that this formulation of the assignment problem is asymmetric, such that we explicitly match M_0 to M_1.

Thus, the mesh alignment is formulated as a quadratic assignment problem, that can be solved via spectral relaxation [LH05], where we compute φ_0, the leading eigenvector of $\mathbf{A}$, and apply the Hungarian algorithm to discretize φ_0 [EKG12], and recover the most probable hard assignments. The proposed scheme is summarized in Algorithm 1:

Algorithm 1

```
1. Compute the sets of local descriptors $\{v_i\} \in M_0$ and $\{v'_i\} \in M_1$.
2. $\forall v_i \in M_0$ find the subset of $M_1$ containing $L$ closest points in terms of descriptor similarity.
3. Apply Eq. 2 to compute the pairwise affinity matrix $\mathbf{A}$.
4. Compute $\varphi_0$, the leading eigenvector of $\mathbf{A}$.
5. Apply the Hungarian Algorithm to $\varphi_0$ to compute the discrete assignment result.
```

4. Robust scale estimation

One of the downsides of the mesh alignment scheme presented in the previous section, is its inability to align meshes related by significant scalings. Namely, Eq. 1 assumes an isometric transformation that is scaling-free. To alleviate this limitation, we propose a straightforward approach for estimating the relative global scalings between pairs of meshes. For that we utilize the Scale-DoG interest point detector, and the corresponding LD-SIFT descriptor proposed by Darom and Keller [DK12].

Given a pair of 3D models M_0 and M_1 we apply the Scale-DoG detector to detect the interest points and corresponding local scales $\{v_i, s_i\} \in M_0$ and $\{v_{i'}, s_{i'}\} \in M_1$, respectively. Let $v_i \in M_0$ and $v_{i'} \in M_1$ be a pair of corresponding points, matched by their descriptors similarity

$$i' = \arg\min_k |\mathbf{x}_i - \mathbf{x}_{k'}|_2, \ \forall v_i \in M_0, \forall v_{k'} \in M_1 \qquad (3)$$

The ratio of their local scales $\Delta s_{ij} = s_i / s_j$ is an estimate of the global scaling between the two models. Thus, we propose to compute the histogram of $\log(\Delta s_{ij})$ for all pairs of corresponding interest points detected in M_0 and M_1.

Using Eq. 3 might result in a significant number of outlier matches. But the distribution of their scaling ratios is uniform over the range of scale ratios, while the inlier scaling measurements concentrate over a short interval centered at the true scaling value. This approach does not utilize global attributes of the meshes and can thus be used to recover the relative scaling of partial matches.

Given the scale ratio Δs between the two models Eq. 1 is reformulated as

$$\tilde{a}_{i,i',j,j'} = \exp\left(-\frac{|d^0_{ik} - \Delta s \cdot d^1_{jl}|}{\sigma}\right), \qquad (4)$$

and as in Section 3.

5. Experimental results

In this section we experimentally verify the proposed mesh matching and scale estimation schemes. We applied both of the proposed methods to models containing arbitrary scale ratios, and considered full and partial matchings†.

5.1. Sparse matching

We tested the sparse matching scheme on models taken from the TOSCA [BBK08] and SHREC'11 [BBB*11] datasets. For computational reasons we used graph distance computed by the Dijkstra algorithm as an approximation for geodesic distances, limited the number of matched vertices in M_0 to 1,000, and considered $L = 5$ matching candidates for each such vertex. Figure 2 shows successful sparse matching of hundreds of points between three model pairs related by scaling and an isometric transformation.

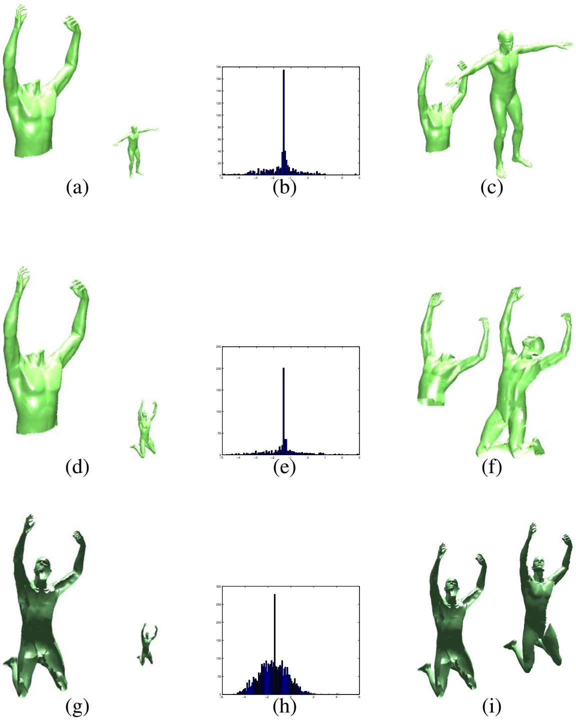
(a) (b) (c)
(d) (e) (f)
(g) (h) (i)

Figure 3: *Global scale detection for partial models: (a, d, g) Scaled partial man models. (b, e, h) the 100 bin Log-Scale histogram. (c, f, i) Left man model scaled by the detected* Δs.

5.2. Scale detection

We tested the proposed scale detection scheme using a 100 bins Log-Scale histogram on models taken from the SHREC'11 challenge dataset including one model consisting of 52,565 vertices under different transformations to test our method on a more realistic and difficult scenario, by scaling one of the meshes in each pair. Figure 3 shows partial man models for which scaling has been applied, and the corresponding Log-Scale histogram and the rescaled models.

5.3. Timing results

The proposed scheme was implemented in Matlab with some critical parts coded in C. In Table 1 we report the timing results of the proposed scheme. Our experiments were carried out on a computer running an 2.6GHz Intel i5 processor with 4GBytes of memory. It follows that the proposed methods requires tens of seconds to align hundreds of points

---

† additional results for the proposed methods can be found at: `http://sites.google.com/site/taldarom`

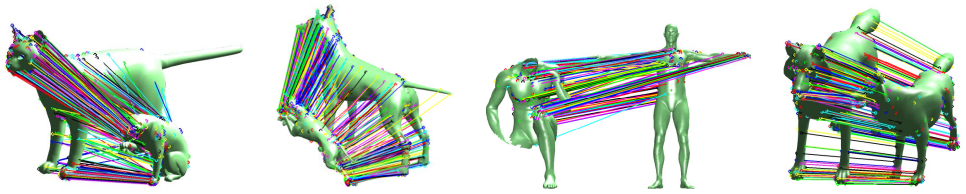

(a) (b) (c) (d)

Figure 2: *Sparse matching of various scaled models. (a)* $|V| = 27,894$*, 505 matches,* $\Delta s = 0.53$*. (b)* $|V| = 25,290$*, 541 matches,* $\Delta s = 0.53$*. (c)* $|V| = 52,565$*, 176 matches,* $\Delta s = 1.2$*. (d)* $|V| = 4,344$*, 329 matches,* $\Delta s = 0.68$*.*

Models (Vertices)	Points matched	Features extraction & Geodesic distances	Scale detection	Affinity	Spectral matching
Cat (27,894)	505	68.18 sec	0.33 sec	4.16 sec	1.53 sec
Dog (25,290)	541	57.07 s.	0.11 sec	4.04 sec	1.11 sec
Wolf (4,344)	329	4.95 s.	0.16 sec	0.50 sec	0.16 sec
Man (52,565)	176	123.98 s.	0.32 sec	4.91 sec	1.34 sec

Table 1: *Running times of the proposed methods.*

in meshes with O(10^4) vertices, where most of the computational time is used by the extraction of feature points and the computation of the graph distances, compared with computational time used for spectral analysis.

6. Conclusions

In this work we presented a computational approach to aligning meshes related by isometric transformations, by formulating the matching as a combinatorial optimization problem solved via spectral graph matching. Using local detectors and descriptors we match sparse sets of salient points in both shapes, making our scheme robust to partial matchings. The assumption of isometric transformation does not hold under significant relative scalings, and we propose a robust approach for estimating the relative scale between meshes based on the local scale estimate of local features. Thus, we derive a unified mesh alignment scheme that allows to estimate both isometric and scaling transformations, that is shown to be accurate and robust. In particular, compared to state-of-the art methods our method is better able to process large meshes and is faster by orders of magnitude.

References

[BBB*11] BOYER E., BRONSTEIN A. M., BRONSTEIN M. M., BUSTOS B., DAROM T., HORAUD R., HOTZ I., KELLER Y., KEUSTERMANS J., KOVNATSKY A., LITMAN R., REININGHAUS J., SIPIRAN I., SMEETS D., SUETENS P., VANDERMEULEN D., ZAHARESCU A., ZOBEL V.: Shrec 2011: robust feature detection and description benchmark. In *Eurographics Workshop on 3D Object Retrieval* (2011). 3

[BBK06a] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Scientific Computing 28*, 5 (2006), 1812–1836. 2

[BBK06b] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. National Academy of Sciences (PNAS) 103*, 5 (January 2006), 1168–1172. 1, 2

[BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical geometry of non-rigid shapes*. Springer, 2008. 3

[BBK*10] 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. *Int. J. Comput. Vision 89* (September 2010), 266–286. 1

[BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell. 14* (February 1992), 239–256. 1

[CFSV04] CONTE D., FOGGIA P., SANSONE C., VENTO M.: Thirty years of graph matching in pattern recognition. *IJPRAI 18*, 3 (2004), 265–298. 1

[DK10] DUBROVINA A., KIMMEL R.: Matching shapes by eigendecomposition of the laplace-beltrami operator. In *3DPVT* (2010). 1, 2

[DK12] DAROM T., KELLER Y.: Scale invariant features for 3D mesh models. *Accepted for publication, IEEE Transactions on Image Processing* (2012). 2, 3

[EKG12] EGOZI A., KELLER Y., GUTERMAN H.: A probabilistic approach to spectral graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2012). 2

[LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *International Conference of Computer Vision (ICCV)* (October 2005), vol. 2, pp. 1482 – 1489. 1, 2

[MS05] MÉMOLI F., SAPIRO G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics 5*, 3 (2005), 313–347. 2

[SHM12] SHARMA A., HORAUD R. P., MATEUS D.: 3D shape registration using spectral graph embedding and probabilistic matching. In *Image Processing and Analysing With Graphs: Theory and Practice*, Lezoray O., Grady L., (Eds.). CRC Press, 2012, ch. 8. 1, 2

[ZWW*10] ZENG Y., WANG C., WANG Y., GU X., SAMARAS D., PARAGIOS N.: Dense non-rigid surface registration using high-order graph matching. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), 382–389. 1