A Unified Spectral Embedding for Shape Correspondence

Zizhao Wu, Ruyang Shou and Xinguo Liu
State Key Lab of CAD&CG, Zhejiang University, HangZhou, China

Keywords: Geometry Processing, Shape Correspondence, Spectral Embedding.

Abstract: Spectral embedding, as one of shape representative techniques, takes hold of many researchers' attention in field of shape correspondence. One of the biggest challenges of spectral correspondence method is that embeddings of different shapes need to be aligned in the embedding space in order to eliminate sign flip and ordering ambiguity of their eigenfunctions, before seeking for correspondence. In this paper, we introduce a spectral correspondence method by embedding shapes in a unified space simultaneously. In the unified embedding space, the sample points of the same shape with small intrinsic distances, and from different shapes with high similarity, are close to each other. Our unified embedding can be used for correspondence directly, without need of alignment. Furthermore, the unified embedding captures both the spatial arrangement and the feature similarity. Shape correspondence is achieved with such embedding by minimizing an objective function. Results show the efficiency of our method.

1 INTRODUCTION

Shape correspondence (also called shape matching) is an important prerequisite to many geometry processing applications such as statistical shape modeling, shape morphing, deformation transfer, shape registration, and sequential meshes analysis, etc. Existed shape correspondence methods could be roughly categorized into rigid and non-rigid cases, regarding to the types of transformations they involve. In this paper, we suppose the shapes to be isometric, we make such restriction based on the observation that many real world deformations, such as articulated shapes in different poses, preserve pairwise geodesic distances up to minor errors, which approximately satisfies the definition of isometry.

Spectral correspondence method as it represents the shape in a spectral representation, is an important tool in finding non-rigid shape correspondence. It mainly comprises two key ingredients: representing shapes using spectral embeddings and matching the shapes based on their embeddings. To represent an individual shape using spectral embedding, some methods usually first construct an affinity matrix based on the geometric information of the shape, and then perform eigen decomposition on the affinity matrix that will produces low dimensional embedding whose entirety is also called the spectrum. While as some researchers (Jain and Zhang, 2006; Mateus et al., 2008) pointed out, directly utilizing the embeddings to find correspondences between shapes causes fallacious inaccurate results, because there may exist sign flip and ordering ambiguity in their spectral representations. To address this issue, several authors (Jain and Zhang, 2006; Mateus et al., 2008) employed an additional alignment step, using either brute force search or greedy approach, which is a reasonable way but a complicated and exhausting task.

In this paper, we introduce a novel framework for shape correspondence. We simultaneously embed the target shapes in one unified space. The unified embedding encodes both the spatial arrangement and feature similarity as they are captured in the unified affinity matrix, where the spatial arrangement is preserved for sample points of each shape and simultaneously the sample points from different shapes with high similarity are close to each other. Given such embedding, the matching results can be achieved by minimizing an objective function, which is expected to be solved in polynomial time.

Comparing to state-of-the-art shape correspondence methods based on spectral embedding, our main contributions and advantages are summarized as follows:

- Our method takes advantages of the unified representation of spectral embedding for different shapes, which doesn't need a further alignment step as done by some traditional spectral methods.
- To our knowledge, we are the first to encode both

Wu Z., Shou R. and Liu X..
A Unified Spectral Embedding for Shape Correspondence.
DOI: 10.5220/0004278700940099
In *Proceedings of the International Conference on Computer Graphics Theory and Applications and International Conference on Information Visualization Theory and Applications* (GRAPP-2013), pages 94-99
ISBN: 978-989-8565-46-4

the spatial arrangement and feature similarity in one embedding space. The unified spectral embedding is experimented suitable for shape correspondence.

2 RELATED WORK

Descriptor based Approaches. Dealing with the non-rigid shapes, the appealing descriptor should preserve intrinsic (e.g., geodesic) distances which are invariant to inelastic deformations of the shape. Gal et al. (Gal et al., 2007) extended the shape context to the non-rigid cases. Elad and Kimmel (Elad and Kimmel, 2001) developed multidimensional scaling (MDS) to represent shape in a low-dimensional embedding space and compared them using Euclidean distance.

However, descriptors based on geodesic distances are very sensitive to topological noise, since a small modification on the connectivity of the mesh will change geodesic distances dramatically over a large part of the shape. Spectral method however is stable under small changes in the intrinsic metric, as some authors proved (Jain and Zhang, 2006; Sun et al., 2009; Ovsjanikov et al., 2010). Among them, the Heat Kernel Signature (HKS) is a very efficient tool in extracting features and comparing shapes at different scales. Sun et al. (Sun et al., 2009) developed the HKS which inherits many properties of the heat kernel including isometric invariance, multi-scale geometric information encoding, and robustness to small perturbations.

Due to the strong properties of the HKS, we employed it as one of the components of our descriptor measuring the feature similarity between a pair of shapes.

Spectral Embedding. Besides the HKS, other spectral embedding methods have been well studied. Since the affinity matrix encodes pairwise geometric invariants which can also be preserved in the spectral embedding space, in order to describe articulated motion, Jain and Zhang (Jain and Zhang, 2006) proposed to use geodesic distance to construct the matrix. Mateus et al. (Mateus et al., 2007) took use of LLE (Roweis and Saul, 2000) and Laplace embedding (Mateus et al., 2008).

The embeddings of different shapes, obtained through techniques mentioned above, are incompatible and cannot be used to generate correspondences directly, as their proposer pointed out the reasons that:

- There exists a sign flip problem when computing eigenvectors ($Av = \lambda v \Leftrightarrow A(-v) = \lambda(-v)$).
- There exists an unreliable ordering of its eigenvalues due to large algebraic multiplicity of an eigenvalue and numerical instabilities in calculation process.

To deal with the sign flip and the eigenfunction reordering problems. Jain and Zhang (Jain and Zhang, 2006) suggested to align the embeddings by minimizing a cost function. Sahillioğlu and Yemez (Sahillioğlu and Yemez, 2010) aligned the spectral embeddings using the the same way as (Jain and Zhang, 2006), and found initialized correspondences by solving a bipartite graph matching problem. Mateus et al. (Mateus et al., 2008) proposed to find an orthogonal transformation that best aligns pairwise embeddings, followed by introducing an unsupervised clustering and the EM algorithm for seeking correspondences of embeddings.

3 ALGORITHM OVERVIEW

In this paper, we focus on the shape matching between two approximately isometric shapes. We denote by $\mathcal{S}$ and $\mathcal{T}$ the discrete version of both shapes which are differentiable 2-manifolds in $\mathbb{R}^3$, indicating the source and target shape respectively. Furthermore, we sample a few feature points on them based on the importance sampling technique (Hilaga et al., 2001), the sets of sample points of $\mathcal{S}$ and $\mathcal{T}$ are designated as $K^{\mathcal{S}}$ and $K^{\mathcal{T}}$. By embedding the shapes into a d-dimensional metric space, we use $U^{\mathcal{S}} \in \mathbb{R}^d$ and $U^{\mathcal{T}} \in \mathbb{R}^d$ to represent the sets of embeddings of sample points $K^{\mathcal{S}}$ and $K^{\mathcal{T}}$ in the embedding space, respectively. In this respect, the innovative point of our framework is that we achieve $U^{\mathcal{S}}$ and $U^{\mathcal{T}}$ at once by simultaneously embed $K^{\mathcal{S}}$ and $K^{\mathcal{T}}$ into a unified embedding space. Our final goal is to find a matching function ϕ that maps $U^{\mathcal{S}}$ to $U^{\mathcal{T}}$. ϕ should be meaningful and accurate, at the same time cover the embedding as much as possible, and can be efficiently determined.

Figure 1 illustrates the pipeline of our 3-stage framework:

In the "**Affinity Matrices**" stage, we introduce the spatial affinities which represent spatial constraints of sample points of each shape and the similarity affinities which represent similarity constraints of sample points between different shapes. These affinities form submatrices that constitute the unified affinity matrix. We will show it in details in Section 6.

In the "**Unified Embedding**" stage, we perform eigen decomposition on the unified matrix and obtain the embeddings $U^{\mathcal{S}}$ and $U^{\mathcal{T}}$. Since the unified affinity matrix are composed of the spatial affinities

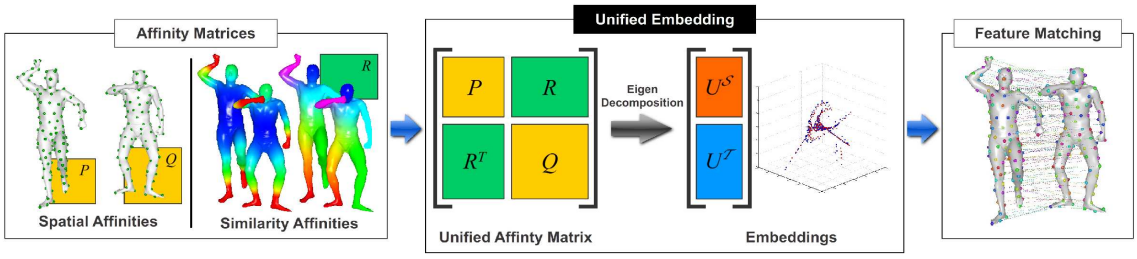

Figure 1: Three main stages of the proposed method. We first construct the spatial affinities and similarity affinities for both input shapes, then achieve unified embedding through eigen decomposition of unified affinity matrix built in the way of using spatial and similarity matrices as its block submatrices. Finally, the matching result is obtained by minimizing an objective function based on the unified embedding coordinates.

and the similarity affinities, our unified embedding intrinsically possesses the spatial and similarity properties of significance. Both properties are applicable for shape matching. The spatial arrangement maintains global consistency of correspondences and the similarity constraints improve matching accuracy. We will offer a detailed discussion in the next section.

In the "**Feature Matching**" stage, we determine the optimal matching ϕ by minimizing an objective function with regard to the unified embedding, which will be particularized in Section 5.

4 THE UNIFIED SPECTRAL EMBEDDING

To formulate properties of the embedding mentioned in Section 3, inspired from the work (Torki and Elgammal, 2010), for the source shape $\mathcal{S}$, we denote the spatial affinity matrix P, with the entry p_{ij} measures the spatial distances of i-th and j-th sample point of $\mathcal{S}$. Similarly, we define Q with entry q_{ij} for the target shape $\mathcal{T}$. We next define the similarity affinity matrix for $\mathcal{S}$ as R with respect to $\mathcal{T}$, the entry r_{ij} measures similarity between i-th sample point on $\mathcal{S}$ and j-th on $\mathcal{T}$. The construction of these affinities will be discussed in Section 6.

Suppose we are given the matrices above, then the embeddings $U^{\mathcal{S}}$ and $U^{\mathcal{T}}$ are determined aiming at minimizing the following objective function

$$\begin{aligned}\arg\min \quad & \sum_{i,j}\|U_i^{\mathcal{S}}-U_j^{\mathcal{S}}\|^2 p_{ij}+\sum_{i,j}\|U_i^{\mathcal{T}}-U_j^{\mathcal{T}}\|^2 q_{ij} \\ & +\sum_{i,j}\|U_i^{\mathcal{S}}-U_j^{\mathcal{T}}\|^2 r_{ij}+\sum_{i,j}\|U_i^{\mathcal{T}}-U_j^{\mathcal{S}}\|^2 r_{ji},\end{aligned} \tag{1}$$

where $U_i^{\mathcal{S}}\,(i\le|K^{\mathcal{S}}|)$ denotes the embedding representation of i-th sample point on shape $\mathcal{S}$. It can be seen from the objective function that the first two terms encode the spatial arrangement of two point sets, which guarantees that two points will be close to each other in the embedding space if their corresponding affinity entry is relatively high. The last two terms aiming at keeping points pairs of different point sets close to each other if value of their corresponding similarity affinity entry is large.

We simplify the Equation 1 by introducing a matrix U stacks the desired embedding coordinates,

$$\begin{aligned}U &= \left[(U^{\mathcal{S}})^T,(U^{\mathcal{T}})^T\right]^T \\ &= \left[u_1^{\mathcal{S}},\ldots,u_{|K^{\mathcal{S}}|}^{\mathcal{S}},u_1^{\mathcal{T}},\ldots,u_{|K^{\mathcal{T}}|}^{\mathcal{T}}\right]^T,\end{aligned}$$

and a unified affinity matrix A defined as

$$A=\begin{bmatrix}P & R\\ R^T & Q\end{bmatrix},$$

let a_{ij} denotes the element on i-th row and j-th column of A, the Equation 1 can then be simplified to

$$\arg\min\sum_{i,j}\|U_i-U_j\|^2 a_{ij}, \tag{2}$$

where U_i is i-th element of U.

Note that the Equation 2 can then be rewritten in the following form

$$\arg\min_U(U^T L U),$$

where L is the Laplacian of the matrix A, i.e., $L=D-A$, and D is the diagonal matrix defined as $D_{ii}=\sum_j a_{ij}$. Since each element of U is a multi-dimensional representation, according to (Belkin and Niyogi, 2003), the above objective function reduces to

$$U^*=\arg\min_{U^TDU=I} tr(U^T L U), \tag{3}$$

the constraint $U^TDU=I$ removes the arbitrary scaling. The unified embedding can be reached through a generalized eigen decomposition $Lu=\lambda Du$. We make use of the bottom d nonzero eigenvectors to form the optimal solution of U.

5 FEATURE MATCHING

In this Section, we model the matching problem as a Linear Assignment Problem, which can be solved in polynomial time.

Let $\phi : K^S \rightarrow K^T$ denote the bijective mapping function representing the correspondence between K^S and K^T, we achieve the matching results by minimizing the following matching cost

$$\arg\min_{\phi} \sum_i \psi(U_i^S, U_{\phi(i)}^T), \tag{4}$$

where $\psi(a,b)$ is the cost function measuring similarity between the embeddings a and b, which in practice we use L_2 distance.

Similar to (Jiang and Yu, 2009), we rewrite the Equation 4 in a succinct matrix form. Let X be the assignment matrix and C be the cost matrix related to cost function ψ respectively. Thus, we have

$$X_{ij} = \begin{cases} 1, & \phi(i) = j \\ 0, & \text{otherwise} \end{cases}, C_{ij} = \psi(U_i^S, U_j^T).$$

Without loss of generality we assume that $|K^S| \geq |K^T|$, and let tr denotes the trace of a matrix, then Equation 4 can be solved by

$$\min f(X) = tr(C^T X), \tag{5}$$

subject to

$$X 1^{|K^S|} = 1^{|K^T|},$$
$$X^T 1^{|K^T|} \leq 1^{|K^S|},$$
$$X \in \{0,1\}^{|K^T| \times |K^S|}.$$

where 1^n denotes a column vector of size n with entries equal to 1, $X \in \{0,1\}^{|K^S| \times |K^T|}$ denotes that the matching between the candidate pair (i,j) is either true or false. For the $|K^S| = |K^T|$ case we force each row and each column of X containing exactly a single one, that means all sample points from source shape should be matched into target sample points and vise versa. For more general case where there exists $\min(|K^S|, |K^T|)$ injections in the matching result, we allow some undetected matchings, which is represented as zero elements in $X 1^{|K^S|}$ or $X^T 1^{|K^T|}$.

In our implementation we employed the lp_solve (lp_solve,) to solve this linear assignment problem.

6 AFFINITIES

6.1 Spatial Arrangement Affinity

The spatial affinity should reflect the spatial arrangement of the sample points of each shape, which means in our case two sample points are close to each other in the embedding space if the intrinsic distance between them is small. In order to attain bending invariant representation, we use geodesic distances to define the relationships between the sample points of each shape. We further apply a Gaussian kernel filtering on the geodesic distances so as to remove scale relevance. In summary, for shape S, we construct the spatial affinity matrix P using the following definition

$$p_{ij} = \exp\left(\frac{-d_g^2(K_i^S, K_j^S)}{2\sigma_S^2}\right),$$

where $d_g(a,b)$ is the geodesic distance between sample points a and b, K_i^S is i-th sample point of shape S, $\sigma_S = w_S d_{max}$ is the Gaussian kernel width with respect to the maximal geodesic distance d_{max} among sample points, with w_S usually set to 0.005. We also construct the spatial affinity matrix Q for $\mathcal{T}$ similarly.

6.2 Feature Similarity Affinity

The similarity affinity matrix should reflect the feature similarity of embeddings from different shapes. In this work, we adopt the HKS (Sun et al., 2009) as our choice with consideration to its strong properties as mentioned previously. However, due to the geometry discretization errors and sampling distortion, it's hard to preserve the consistency of signatures in the diffusion embedding space. Moreover, there also exists ambiguity problem when dealing with intrinsically symmetric shapes of the HKS. In order to measure the feature similarity following a continuity constrained way, we employed a geodesic field and combine it with HKS. Hence, our similarity affinity matrix is constructed through the following steps

- Calculate the HKS for each shape, based on which we then detect anchor correspondences.
- Compute the geodesic field using the anchors. Construct new signature as combination of the HKS and the geodesic field.
- Compute optimized matrix through singular value decomposition (SVD) on primal affinity matrix measured on Gaussian kernel on squared Euclidean distance in the signature space.

We follow the work proposed by (Sun et al., 2009) to calculate the HKS, and extract the anchor correspondences using the method of (Ovsjanikov et al., 2010).

Given a set of anchors on shape S, denoted by $\kappa_1^S, \cdots, \kappa_k^S, k < |K^S|$, for each sampling point we establish the geodesic field defined as

$$y_i = \left[\frac{d_g(K_i^S, \kappa_1^S)}{d_{max}}, \cdots, \frac{d_g(K_i^S, \kappa_k^S)}{d_{max}}\right]^T,$$

where d_{max} is the maximal geodesic distance among sampling points, K_i^S is i-th feature point of shape $\mathcal{S}$.

The combination of the HKS and geodesic field is a straightforward dimensional concatenation, that's to say our new signature can be written as

$$f_i = \left[h_i^T, \lambda y_i^T\right]^T, \quad (6)$$

where f_i is the combined signature for i-th sample point of the shape and h_i is the HKS. We believe that a balanced weight $\lambda = 1$ is ideal for most cases.

Now we can compute feature similarity affinity R measured on combined signatures. The affinity should exclude most of features, while at the same time, avoid to make any hard decision, i.e., a zero-one permutation matrix is not the one we pursuit.

To address this issue, according to (Scott and Longuet-Higgins, 1991), we first define an affinity matrix $\tilde{R}$, whose entry is given as

$$\tilde{r}_{ij} = \exp\left(\frac{-\| f_i^S - f_j^T \|^2}{2\sigma^2}\right),$$

where we set the width of Gaussian kernel $\sigma = 0.005$. Then we run the singular value decomposition of the affinity matrix $\tilde{R}$, which we express it in the form $\tilde{R} = \tilde{U}\tilde{S}\tilde{V}^T$, where $\tilde{U}$ and $\tilde{V}$ are orthogonal matrices, and $\tilde{S}$ is a non-negative diagonal matrix. Next we convert the matrix $\tilde{S}$ into an identity matrix E by replacing the singular values of $\tilde{S}$ by ones, and obtain an optimized matrix $R^* = \tilde{U}E\tilde{V}^T$, where each element indicates the extent of pairing between two sample points of different shapes. The optimized matrix R^* correlates best with $\tilde{R}$ as is shown in (Scott and Longuet-Higgins, 1991).

Finally, we use the optimized matrix R^* to construct the similarity affinity matrix R by setting the negative values to zero.

7 RESULTS

Time Evaluation. We implemented our framework in C++ and evaluated the performance on a computer with Intel Core™ 2 Quad processor and 4GB of RAM. The running time of our pipeline mainly depended on the construction of the spatial and similarity affinity matrices. For a mesh with about 6000 vertices, these preprocessing stages totally spent about one minute, with the computation of HKS taking about more than half of the duration. Fortunately, finding that the complexity of HKS can be reduced through a multi-resolution optimization approach (Vaxman et al., 2010), we consider employing it in future to extend capability of our framework to extremely large meshes.

Table 1: Time consumption of the alignment using brute force search. #f denotes the number of sample points, k represents the dimension of embeddings. Times are given in seconds.

#f	$k=3$	$k=4$	$k=5$	$k=6$
100	0.11	0.72	6.8	73.2
200	0.21	1.8	16.25	180
300	0.38	2.5	20.0	208
400	0.42	2.99	31.6	367.8

Table 2: The matching distortion of our method in comparison with (Sahillioğlu and Yemez, 2010) on the TOSCA nonrigid world dataset (Bronstein et al., 2008).

	Sahillioğlu et al.	Our method
pair	D_{avg},D_{max}	D_{avg},D_{max}
Cat	0.032, 0.117	0.024, 0.167
Michael	0.031, 0.119	0.035, 0.165
Victoria	0.014, 0.094	0.021, 0.117

In contrast with the state-of-the-art spectral based techniques, our method doesn't entail an alignment step to eliminate sign flips and orderings between embeddings of different shapes. The alignment is quite time consuming (Jain and Zhang, 2006; Mateus et al., 2007), for two k-dimensional embeddings, there are $2^k k!$ possible cases to be verified. Table 1 shows the time consumption by using exhaustive search. Although the author of (Mateus et al., 2007) addressed this problem by searching for an orthogonal transformation, still the algorithm is low performance.

Matching Results Evaluation. We evaluated matching qualities using triangular meshes from the TOSCA nonrigid world dataset (Bronstein et al., 2008) and human bodies generated using the SCAPE model (Anguelov et al., 2005). We sampled about 200 feature points on each shape that originally have about 3K to 20K vertices.

We used the average matching error D_{avg} and the maximal matching error D_{max}, as defined by Sahillioğlu and Yemez (Sahillioğlu and Yemez, 2011) to quantify the matching distortion. The quantified distortions are provided in Table 2, and visualized matching results are shown in Figure **??**, where we can observe that our method achieves good results with less time-consuming.

8 CONCLUSIONS

We proposed a unified spectral embedding for shape correspondence. The embedding is achieved through eigen decomposition on a unified affinity matrix composed of the spatial affinities and the similarity affini-

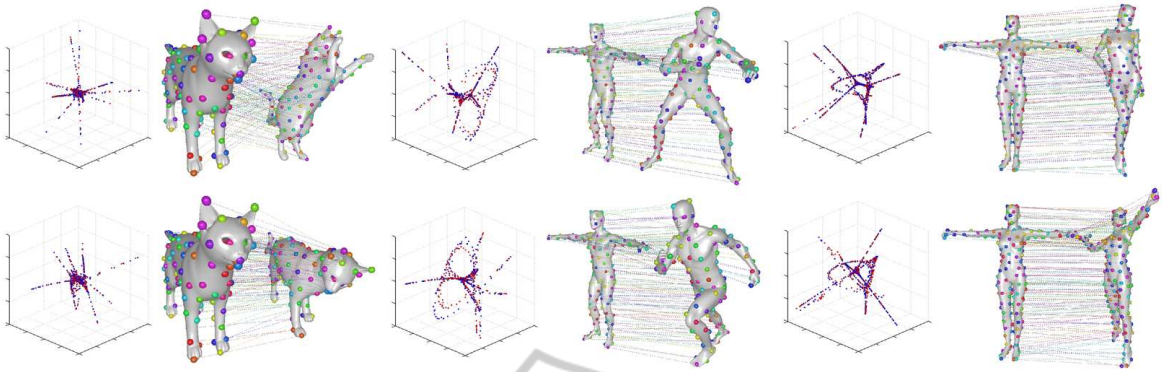

Figure 2: Some experimental results by our method tested on shapes from the TOSCA nonrigid world dataset (Bronstein et al., 2008).

ties, which means that the obtained embedding encodes both the spatial arrangement and the feature similarity. Comparing with other existing spectral techniques, our embeddings of different shapes can be directly used as shape descriptors for correspondence, without necessity alignment between them.

In this paper, we only consider isometric shapes matching in this paper because the construction of affinity matrices are based on isometric invariant geometric properties. In the future we plan to explore the possibilities for rigid matching and partial matching of our framework.

REFERENCES

Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., and Davis, J. (2005). Scape: shape completion and animation of people. *ACM Trans. Graph.*, 24(3):408–416.

Belkin, M. and Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.*, 15(6):1373–1396.

Bronstein, A., Bronstein, M., and Kimmel, R. (2008). Numerical geometry of non-rigid shapes. *Springer-Verlag New York Inc*.

Elad, A. and Kimmel, R. (2001). Bending invariant representations for surfaces. In *CVPR (1)*, pages 168–174.

Gal, R., Shamir, A., and Cohen-Or, D. (2007). Pose-oblivious shape signature. *IEEE Trans. Vis. Comput. Graph.*, 13(2):261–271.

Hilaga, M., Shinagawa, Y., Komura, T., and Kunii, T. L. (2001). Topology matching for fully automatic similarity estimation of 3d shapes. In *SIGGRAPH*, pages 203–212.

Jain, V. and Zhang, H. (2006). Robust 3d shape correspondence in the spectral domain. In *Proc. of Shape Modeling International*, pages 118–129.

Jiang, H. and Yu, S. X. (2009). Linear solution to scale and rotation invariant object matching. In *CVPR*, pages 2474–2481.

lp_solve. http://sourceforge.net/projects/lpsolve.

Mateus, D., Cuzzolin, F., Horaud, R. P., and Boyer, E. (2007). Articulated shape matching by robust alignment of embedded representations. In *IEEE Workshop on 3D Representation for Recognition (3DRR 2007)*. IEEE Computer Society Press.

Mateus, D., Horaud, R. P., Knossow, D., Cuzzolin, F., and Boyer, E. (2008). Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Ovsjanikov, M., Mérigot, Q., Mémoli, F., and Guibas, L. J. (2010). One point isometric matching with the heat kernel. *Comput. Graph. Forum*, 29(5):1555–1564.

Roweis, S. and Saul, L. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326.

Sahillioğlu, Y. and Yemez, Y. (2010). 3d shape correspondence by isometry-driven greedy optimization. In *CVPR*, pages 453–458.

Sahillioğlu, Y. and Yemez, Y. (2011). Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Comput. Graph. Forum*, 30(5):1461–1470.

Scott, G. L. and Longuet-Higgins, H. C. (1991). an algorithm for associating the features of two images. In *Proc. Biol Sci*.

Sun, J., Ovsjanikov, M., and Guibas, L. (2009). A concise and provably informative multi-scale signature based on heat diffusion. In *Eurographics Symposium on Geometry Processing (SGP)*.

Torki, M. and Elgammal, A. M. (2010). One-shot multi-set non-rigid feature-spatial matching. In *CVPR*, pages 3058–3065.

Vaxman, A., Ben-Chen, M., and Gotsman, C. (2010). A multi-resolution approach to heat kernels on discrete surfaces. *ACM Trans. Graph.*, 29(4).