Kernel Functions for Robust 3D Surface Registration with Spectral Embeddings

Xiuwen Liu, Arturo Donate
Department of Computer Science
Matthew Jemison, Washington Mio
Department of Mathematics
Florida State University, Tallahassee, FL 32306
{liux,donate}@cs.fsu.edu mio@math.fsu.edu

Abstract

Registration of 3D surfaces is a critical step for shape analysis. Recent studies show that spectral representations based on intrinsic pairwise geodesic distances between points on surfaces are effective for registration and alignment due to their invariance under rigid transformations and articulations. Kernel functions are often applied to the pairwise geodesic distances to make the registration process based on spectral embedding robust to elastic deformations. The Gaussian kernel is most commonly used, but the effect of the choice of the kernel function has not been studied in the previous works. In this paper, we compare the results obtained with several different choices and show empirically that significant improvements can be achieved in shape registration with appropriate choices.

1. Introduction

Registration of surfaces, that is, establishing meaningful correspondences between their points, is a fundamental problem in shape analysis [6, 8, 14], texture mapping [10], and cross parametrization [9]. Since surfaces are often represented by triangular meshes, a registration is frequently described by correspondences between vertices of meshes. As meshes representing similar shapes can differ significantly due to deformations (see Fig. 4), sampling and positioning in space, one seeks representations that capture the essential geometry of the shapes, are invariant to rigid transformations and articulations, and are robust to elastic deformations.

Several studies show that spectral representations based on pairwise geodesic distances are effective for shape registration [13, 6, 8]. These intrinsic distances are clearly invariant to rigid transformations and articulations [7]. In order to make the spectral representation more robust, a kernel function is typically applied to the pairwise distances [8]. While the choice of kernel function clearly affects the representation, this issue has not been investigated in previous studies. In this paper, we design different kernel functions and study their effect on registration. We show that kernels that stretch small distances and compress the relatively large ones are particularly attractive. We demonstrate empirically that appropriate choices lead to registration results that are significantly more accurate than those obtained with the Gaussian kernel, which is the most common choice.

The paper is organized as follows. In Section 2, we give an overview of shape registration based on spectral representations. Kernel functions are discussed in Section 3 and experimental results with three data sets are reported in Section 4. Section 5 concludes the paper with a summary and some discussion.

2. Spectral Representation and Registration

The shape registration problem is often presented as follows: given two 3D shapes $\mathcal{M}_1$ and $\mathcal{M}_2$ as triangular meshes, construct a continuous mapping between the meshes that match the corresponding geometric features in the given shapes. Here, we assume that $\mathcal{M}_1$ and $\mathcal{M}_2$ are related by some unknown deformation and that a meaningful correspondence exists. For example, given two horses, we would like to produce a mapping such that the physical features (such as legs, head and so on) in $\mathcal{M}_1$ map to the corresponding features in $\mathcal{M}_2$. Here, we consider a slightly simpler variant, known as the point-correspondence problem: the goal is to compute a mapping from the vertex set $\mathcal{V}_1$ of $\mathcal{M}_1$ to the vertex set $\mathcal{V}_2$ of $\mathcal{M}_2$. For $i = 1, 2$, we let $n_i = |\mathcal{V}_i|$. Without loss of generality, we assume that $n_1 \leq n_2$.

The point-correspondence problem has been studied for point clouds in Euclidean space, e.g., [2, 6, 1] with the iterative closest point algorithm (ICP) [2] and several of its variants. In a typical implementation of ICP, an initial correspondence is computed by mapping each point on the first shape to the closest point on the second. The second shape is then warped to best align it with the first with respect

978-1-4244-2175-6/08/$25.00 ©2008 IEEE

to the estimated correspondences. This process is iterated until some matching criterion is satisfied. If points that are expected to correspond are initially close enough, ICP often performs well. Otherwise, the process may fail to converge to a meaningful correspondence (see [8] for an example). One of the ways to improve the performance of ICP is to replace the original point clouds with representations that retain geometric information that is meaningful for registration and are invariant to rigid transformations and articulations. Given a surface, the geodesic distance (i.e., the length of the shortest path on the surface) between two points has the desired invariance properties [7]. This observation leads to shape registration methods that use representations derived from geodesic distances. One of the most commonly used methods is based on spectral embeddings [13, 8] obtained from all pairwise geodesic distances between vertices of the mesh, as described next.

Order the vertices of $\mathcal{M}_1$ arbitrarily and let $\mathcal{P}_1$ be the symmetric $n_1 \times n_1$ matrix of pairwise geodesic distances for the vertices of $\mathcal{M}_1$. The (i, j)-entry of $\mathcal{P}_1$ is the geodesic distance between the ith and jth vertices. Similarly, construct $\mathcal{P}_2$ for $\mathcal{M}_2$. The goal is to derive new k-dimensional Euclidean representations (k small) for $\mathcal{M}_1$ and $\mathcal{M}_2$ with corresponding vertices lying relatively close to each other so that ICP-type algorithms with respect to the Euclidean distance can be applied effectively for registration. We can view the jth column of $\mathcal{P}_1$ as a new high-dimensional representation of the jth vertex of $\mathcal{M}_1$ and similarly for $\mathcal{M}_2$. A common approach is to use principal component analysis to reduce the dimension of this representation using the coordinates of the orthogonal projections onto the eigenvectors associated with the k largest eigenvalues. Note that even under the assumption that all eigenvalues are distinct (this holds generically), the directions of the principal components are inherently ambiguous, as we can change the sign of the eigenvectors arbitrarily. Thus, there are several possible choices of spectral representations and the optimal choice of signs needs to be determined [8]. Also, due to other deformations – such as local stretching and compression – that are not accounted for by the model, the principal axes may not naturally correspond so that further rigid alignment may be necessary [11]. In this paper, we first use ICP with rigid transformations to align the surfaces in the low-dimensional representation obtained for each choice of signs and select the representation that yields the best alignment. Subsequently, the full shape alignment is done in the chosen spectral representation using ICP with non-linear warps estimated with thin-plate splines (TPS) [3, 6]. For computational efficiency, we work at two different resolutions. We first register the surfaces and compute the TPS transformation at a lower resolution. Then, we apply the estimated warp to the original mesh and use ICP again at full resolution to establish the correspondences.

3. Kernel Spectral Representation

For typical meshes representing well sampled surfaces, local pairwise geodesic distances are rather small so that distances to nearby vertices tend to have little effect on spectral representations. This makes the matching algorithm somewhat insensitive to local geometry. Thus, we seek to improve the matching accuracy with a representation more sensitive to small distances. The goal is to stretch the small distances and relatively compress the large ones, an objective that can be achieved with appropriately chosen kernel functions. To produce such kernels, we first normalize the geodesic distances and the distances in the embedding space to make the largest distance unitary. To attain the desired relative stretching and compression in the normalized representation, we first design first derivatives that generate the targeted effects and then integrate them to obtain suitable kernel functions. Fig. 1 shows several examples. To illustrate the effect of such kernels, we plot the distribution of geodesic distances in the original mesh versus the corresponding Euclidean distances in a 3-dimensional spectral representation for a mesh representing the first horse of Fig. 4. Fig. 2(a) shows a Shepard diagram [4] for a spectral embedding generated with a Gaussian kernel applied to the pairwise distances (e.g. [12, 13, 5]). Fig. 2(b) shows the Shepard diagram for the square-root kernel. Clearly, the

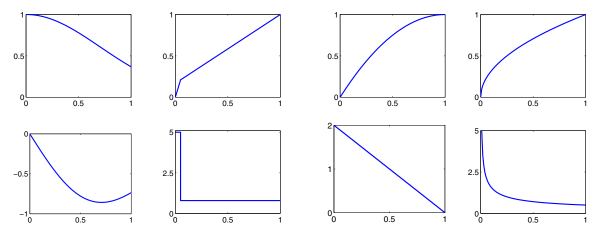

(a) e^{-x^2} (b) piecewise linear (c) $2x - x^2$ (d) $\sqrt{x}$

Figure 1. Kernel functions with their first derivatives shown on the 2nd row.

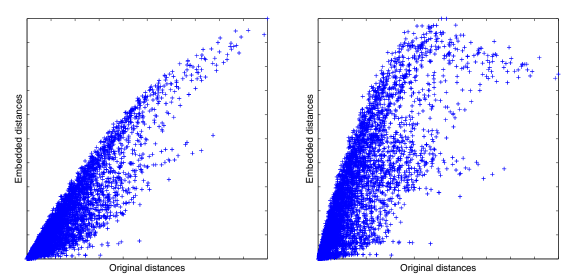

(a) (b)

Figure 2. Shepard diagrams for (a) a Gaussian kernel and (b) the square-root function.

small distances are stretched by this kernel and thus play a more significant role in the spectral representation.

4. Experimental Results

We used the spectral embedding method with the proposed kernel functions to register shapes in 3 families of surfaces consisting of 6 horses, 4 octopuses and 3 representations of the word “Happy”. In each case, we started with a single mesh from the AIM@SHAPE Shape Repository and “manually” deformed it in different ways with the graphics software Blender to produce sizable bending and stretching deformations, as shown in Fig. 4. Thus, ground truth was available for these data sets (i.e., the underlying correspondences were known) to evaluate the performance of the registration algorithm.

In our experiments, we applied several kernel functions to obtain different spectral representations. Table 1 shows a comparison of the results for the Gaussian kernel and the square-root function for a pair of horses at two different resolutions: 200 and 1,000 vertices. A match for a vertex is considered correct with tolerance k if it falls within the k-ring of the correct vertex; in other words, there is a path formed by at most k edges joining the point to the correct match. Thus, tolerance 0 means that the matching is perfect. As explained above, in all experiments, non-linear warps were estimated with thin-plate splines at a low resolution. Consistent with the fact that the square-root is more sensitive to small distances and also compresses larger distances, the results show a significant improvement in matching accuracy over the Gaussian kernel at low resolution. High accuracy at a low resolution is very important to improve the overall computational efficiency of the algorithm in a multi-resolution approach to registration, as proposed.

Table 1. Comparison of matching accuracy (%) for the Gaussian and square-root kernels.

Tolerance	Gaussian		Square root	
	Low res.	Full res.	Low res.	Full res.
0	82.0	80.8	100.0	99.9
1	82.5	87.3	100.0	99.9
2	94.5	96.7	100.0	100.0
3	99.0	99.4	100.0	100.0
4	99.5	99.8	100.0	100.0
5	100.0	100.0	100.0	100.0

Figure 3 shows the average matching accuracy with tolerance zero for 11 pairs of horses with ten different kernel functions in the following order: the Gaussian kernel, the identity map, three piecewise linear functions with different slopes, $f(x) = 2x - x^2$, $\sqrt{x}$, $\sqrt[3]{x}$, $\sqrt[4]{x}$, x^2. The solid (blue) curve shows the matching accuracy on the low resolution meshes (with 200 vertices) and the dashed (red) curve on the high resolution meshes (with 1000 vertices). The kernel functions that stretch small distances the most while compressing the larger ones give higher matching accuracy. The square-root function gives the highest average on the low resolution meshes and the cubic-root function yields the highest average on the high resolution ones. An improvement of as much as 9% over the Gaussian kernel has been observed.

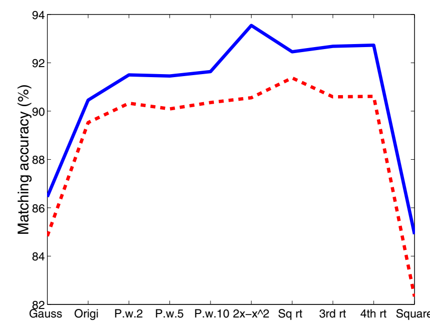

Figure 3. Average matching accuracy on 11 pairs of horses with ten kernel functions.

Figure 4 shows the results of several experiments with the square-root kernel. Each shape was registered with the one on the leftmost position. The alignment of the horses was calculated as described above, with no-linear warps estimated at a lower resolution. The mesh representing the octopuses and Happy have 386 and 706 vertices, respectively. Since the original meshes already are rather sparse, all computations were performed with the full meshes.

5. Discussion and Conclusion

Due to the invariance of spectral representations to rigid transformations and articulations, shape registration based on spectral methods using intrinsic geodesic distances on surfaces is robust and can achieve high accuracy. In this paper, we showed that the choice of kernel functions for spectral embedding has a significant effect on the matching accuracy. We showed experimentally that an appropriate choice of kernel function can lead to significant improvements over the most commonly used Gaussian kernel. While the proposed kernels noticeably improve performance, they may not be optimal. Thus, techniques to optimize the choices as well as variants of the registration method are being further studied.

The shape matching method used in this paper does not impose continuity or smoothness on the mapping induced on the underlying triangular meshes. One way of enforcing some degree of smoothness is to add a regularization term to prevent large jumps from occurring in the vertex correspon-

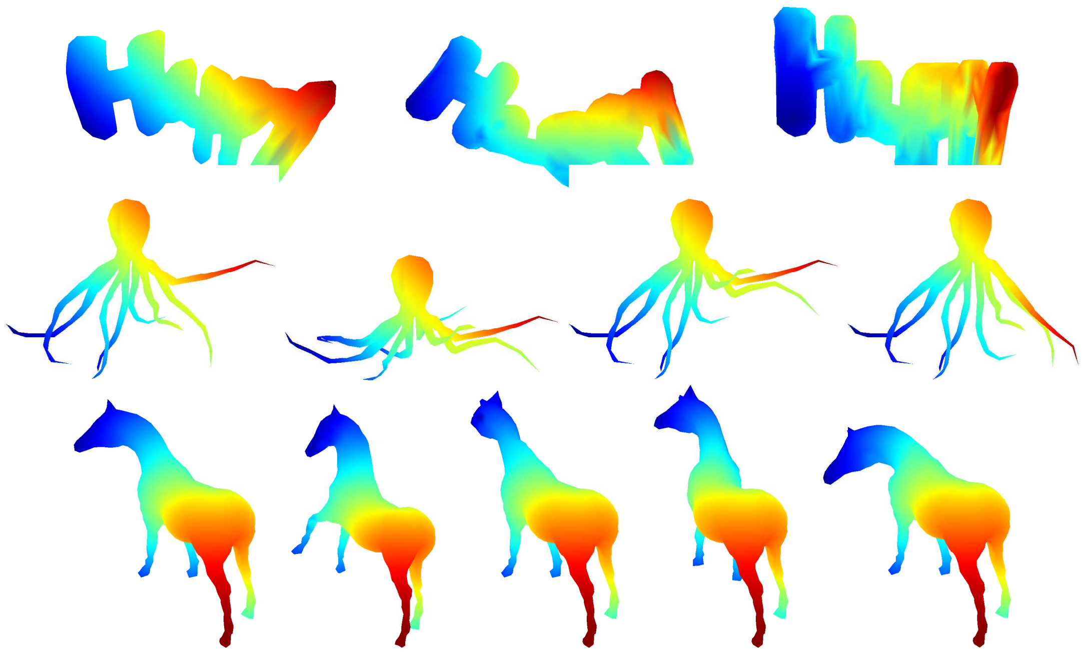

Figure 4. Examples of automatic shape registration: each shape is matched with the leftmost one.

dences. With this additional element, the proposed method may lead to an automated procedure to achieve accurate shape matchings that are better suited to the computational analysis of shapes of surfaces, as needed in applications to problems in medical image analysis and computer vision. These problems are being considered in our ongoing investigations.

Acknowledgments

This research was supported in part by the National Science Foundation grants CCF-0514743 and DMS-0713012. The original shape models were provided courtesy of INRIA and MPII by the AIM@SHAPE Shape Repository.

References

[1] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):509–522, 2002.

[2] P. J. Besl and N. D. Mckay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[3] F. L. Bookstein. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.

[4] I. Borg and P. Groenen. *Modern Multidimensional Scaling*. Springer, New York, 1997.

[5] M. Carcassoni and E. R. Hancock. Correspondence matching with modal clusters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(12):1609–1615, 2003.

[6] H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89:114–141, 2003.

[7] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10):1285–1295, 2003.

[8] V. Jain, R. Zhang, and O. V. Kaick. Non-rigid spectral correspondence of triangle meshes. *International Journal for Shape Modeling*, 2008.

[9] V. Kraevoy and A. Sheffer. Cross-parameterization and compatible remeshing of 3d models. *ACM Transations on Graphics*, 23(3):861–869, 2004.

[10] V. Kraevoy, A. Sheffer, and C. Gotsman. Matchmaker: Constructing constrained texture maps. In *ACM SIGGRAPH*, pages 326–333, 2003.

[11] D. Mateus, F. Cuzzolin, R. P. Horaud, and E. Boyer. Articulated shape matching by robust alignment of embedded representations. In *IEEE International Conference on Computer Vision*, pages 1–8, 2007.

[12] G. L. Scott and H. C. Longuet-Higgins. An algorithm for associating the features of two images. *Proceedings: Biological Sciences*, 244(1309):21–26, 1991.

[13] L. S. Shapiro and J. M. Brady. Feature based correspondence: an eigenvector approach. *Computer Vision and Image Understanding*, 10(5):283–288, 1992.

[14] Y. Shi, P. M. Thompson, G. I. de Zubicaray, S. E. Rose, Z. Tu, I. Dinov, and A. W. Toga. Direct mapping of hippocampal surfaces with intrinsic shape context. *NeuroImage*, 37:792–807, 2007.