

Geometric-Feature-Based Spectral Graph Matching in Pharyngeal Surface Registration

Qingyu Zhao¹, Stephen Pizer¹, Marc Niethammer¹, and Julian Rosenman²

¹ Computer Science, UNC Chapel Hill, NC, United States

² Radiation Oncology, UNC Chapel Hill, NC, United States

Abstract. Fusion between an endoscopic movie and a CT can aid specifying the tumor target volume for radiotherapy. That requires a deformable pharyngeal surface registration between a 3D endoscope reconstruction and a CT segmentation. In this paper, we propose to use local geometric features for deriving a set of initial correspondences between two surfaces, with which an association graph can be constructed for registration by spectral graph matching. We also define a new similarity measurement to provide a meaningful way for computing inter-surface affinities in the association graph. Our registration method can deal with large non-rigid anatomical deformation, as well as missing data and topology change. We tested the robustness of our method with synthetic deformations and showed registration results on real data.

1 Introduction

The goal of surface registration is to find a dense set of corresponding points between two surfaces. Usually this is challenging because the surface may undergo large deformations, and sometimes there might be missing data, such as unexpected holes and different boundary locations, in the surface.

For example, in our application of the fusion between an endoscope movie and a CT image for head and neck cancer, one can acquire an endoscopic video clip of the pharyngeal region at radiation treatment planning time, from which a 3D reconstruction of the pharyngeal surface is derived. On the other hand, we can also segment a 3D pharyngeal surface from a head and neck CT image of the same patient. A registration between these two surfaces will permit fusion of the endoscopically available information about the tumor extent on the pharyngeal surface with the tumor information seen in the CT, thereby improving the radiation plan. As shown in Figs. 5a, 5b, a large deformation between the two surfaces is caused by the swallowing process and posture change of a patient. Due to the limitation of endoscope procedure, a part of the pharyngeal anatomy is visually inaccessible by the camera. Therefore, the reconstruction surface is only a partial surface with respect to the CT surface. The reconstruction artifacts will also create many holes in the surface. The fusion between endoscope and CT has been tried in other anatomies, but they consider only rigid or small deformations between the modalities and thus cannot be applied to the pharyngeal region.

For non-rigid surface registration some approaches directly solve for the deformation parameters [1–3] by minimizing the closeness of two surfaces in the original 3D domain, but they usually involve a non-convex optimization. Conformal mappings [4] and Möbius transformations [5] have also been proposed to map the surfaces onto a canonical domain conformally, and seek the matching in that space. Other methods [6, 7] for matching isometric shapes embed the surfaces into a Euclidean space to obtain isometry-invariant representations. Among those, spectral graph theory offers a nice mathematical framework for matching shapes in the spectral domain. Many registration methods adopt the idea of feature matching. To that end, spin images [8] and Heat Kernel Signature (HKS) [9] are the two most widely used surface features. However, spin images are variant to large deformations, and HKS can not handle missing data in the surface because of different boundary conditions.

Lombaert [10] found that by combining the two surfaces into one graph with some initial links, spectral decomposition could yield consistent eigenvectors, which he used for surface matching. In his application of cortical surface matching, he chose to use a conventional spectral matching to provide initial links. Expectedly, that method has been shown in our results not to be suitable for finding initial correspondences when facing large deformations and different topology.

We made the following contributions in our paper. First, we design a novel geometry-based feature descriptor that can be applied on any surface with notable geometric structures. Second, we define a correspondence confidence score based on feature comparison, with which an effective initial correspondence set can be derived to work with Lombaert’s new form of spectral method. Third, we discuss the advantage of our method in the context of partial surface matching, which has not been studied before in any spectral matching framework.

2 Geometric Feature Extraction and Use

2.1 Feature Descriptor

In our application, a surface is represented by a triangulated mesh with a set of vertices $\{V\}$ and a set of edges $\{E\}$. We compute geometric feature descriptors at each vertex, on which vertex correspondences are based.

We design a special feature descriptor $f(v)$ to create a signature for each vertex. In order to describe the local shape around a vertex, we collect geometric information on both the vertex itself and a number of surrounding vertices. As shown in Fig. 1a, for each vertex v , we find 8 surrounding vertices $\{v_i|i = 1..8\}$ by going along 8 equally angularly spaced geodesic directions $\{g_i|i = 1..8\}$ from v by a certain distance d . We choose the nearest vertices at the end points of the 8 paths as the surrounding vertices. g_1 and g_3 overlap with the two principal directions p_1, p_2 . Since the local shape can be captured by curvatures measured at different scales, the descriptor is defined as $f(v) = \{\mathbf{C}, \mathbf{S}, \Delta\mathbf{N}, \Delta\mathbf{F}, \Delta\mathbf{N}_{1,5}, \Delta\mathbf{N}_{3,7}\}$.

Koenderink’s [11] informative curvature measures c, s derivable from the two principal curvatures k_1, k_2 , are computed at the center vertex and 8 surrounding vertices to describe local curvatures by the tuple \mathbf{C}, \mathbf{S} . Larger scale measures

of curvature between each of the surrounding vertices and the center vertex are computed as the normal direction difference magnitudes $\Delta\mathbf{N}$, as well as by the local coordinate frame rotation quaternions $\Delta\mathbf{F}$. The local coordinate frame is constructed as the two principal directions plus the normal direction. Also, normal direction differences between two extreme endpoint pairs (v_1, v_5) and (v_3, v_7) are computed to describe the general shape structure.

We used LMNN (Large Margin Nearest Neighbor) to learn the weights for different features using a set of ground truth corresponding vertices with their features and deleted features with near-zero weights, but we found the algorithm performed noticeably better when all features are used.

2.2 Computing Correspondence Confidence

We propose a similarity measurement between vertices from two different surfaces. This measurement is defined by a confidence score $\Delta_{i,j}$, indicating how likely $v_i \in S_1$ and $v_j \in S_2$ are corresponding. Define the two surfaces to be S_1, S_2 with N, M vertices respectively and v_i to be the i th vertex in a surface. Under the assumption that S_1 and S_2 are rigidly aligned first, the feature distance between $v_i \in S_1$ and $v_j \in S_2$ is defined as

$$\delta(i, j) = \|f(v_i) - f(v_j)\|^2 + \alpha(1 + e^{-\|x_i - x_j\| - \tau})^{-1} \tag{1}$$

where the second part is a sigmoid function penalizing a too large Euclidean distance between two corresponding vertices. Based on this feature distance function, we propose an efficient method to compute the confidence score $\Delta_{i,j}$.

A confidence score considers both-way corresponding likelihoods, namely v_i being the closest vertex to v_j and v_j being the closest vertex to v_i . $\kappa_{i,j}^1$ is defined as the likelihood of $v_j \in S_2$ being the closest vertex of $v_i \in S_1$, compared to all other vertices in S_2 . It is computed by normalizing $\delta(i, j)$ to $[0, 1]$ using $\{\delta(i, k) | k = 1 \dots M\}$ (Eq. 2). $\kappa_{i,j}^2$ is defined and computed vice versa (Eq. 3):

$$\kappa_{i,j}^1 = 1 - (\delta(i, j) - \min_k \delta(i, k)) / (\max_k \delta(i, k) - \min_k \delta(i, k)) \tag{2}$$

$$\kappa_{i,j}^2 = 1 - (\delta(i, j) - \min_k \delta(k, j)) / (\max_k \delta(k, j) - \min_k \delta(k, j)) \tag{3}$$




Fig. 1. (a) Local geometry from which $f(v)$ is computed. (b) A vertex is selected in S_1 , indicated as the cross point. (c) The value of Δ 's i th row (red indicates large value).

Because the two likelihoods are now at the same scale, the confidence score $\Delta_{i,j}$ is computed by taking the sum of $\kappa_{i,j}^1$ and $\kappa_{i,j}^2$. All the confidence scores will form a $N \times M$ confidence score matrix Δ . As shown in Figs. 1b, 1c, for a vertex $v_i \in S_1$, Δ 's i th row is color-coded in S_2 . The vertex with the largest value is selected as the corresponding point. The overall dense correspondences based on this strategy are color-coded as shown in Fig. 2b.

3 Spectral Graph Matching

3.1 Spectral Graph Matching on an Association Graph

We build two graphs $G_1 = \{V_1, E_1\}$ and $G_2 = \{V_2, E_2\}$ from the two surfaces S_1 and S_2 with the vertices and edges of the triangulated surface meshes. An association graph $G = \{V, E\}$ is built by connecting G_1 and G_2 with a set of initial links. Lombaert in his work defined the $|N + M| \times |N + M|$ affinity matrix W by the Euclidean distance between two vertices in the original 3D space for both intra-surface links and inter-surface links, i.e., $w_{i,j} = \|x_i - x_j\|^{-2}$ if $\exists e_{i,j} \in E$. The graph Laplacian operator L is defined as $L = D - W$, where D is a diagonal matrix with $d_i = \sum_j w_{i,j}$.

The spectral decomposition of L provides an orthogonal set of eigenvectors $[u^1, u^2, \dots, u^{N+M}]$ with the corresponding non-decreasing eigenvalues, the first of which is zero for appropriate boundary conditions. Each of the eigenvectors u^i can be separated into two functions: u_1^i , the first N values of u^i , representing the i th vibration mode of G_1 , and u_2^i , the last M values of u^i , representing the i th vibration mode of G_2 . The inter-surface links ensure that they represent a consistent vibration mode. Moreover, the spectral embedding of the graph into a k -dimensional Euclidean space, also known as the spectral domain, is given by $[u^2, u^3, \dots, u^{k+1}]$. In other words, we define $F = [f_1, f_2, \dots, f_k]$ as an $n \times k$ matrix, and the first k eigenmodes with non-zero eigenvalues provide the solution to the problem:

$$\arg \min_{f_1, f_2, \dots, f_k} \sum_{i,j=1}^n w_{i,j} \|f^{(i)} - f^{(j)}\|^2, \text{ with } F^T F = I \quad (4)$$

where $f^{(i)}$ is the i^{th} row of F , representing the embedded Euclidean coordinates of the i^{th} vertex. Intuitively, the k eigenmodes define an embedding into a k -dimensional Euclidean space that tries to respect the edge lengths of the graph. The final matching is accomplished by a nearest-neighbor search in the k -dimensional spectral domain.

3.2 Finding Initial Links

The inter-surface affinity in the Lombaert paper was defined according to the Euclidean distance between two corresponding vertices, which is conceptually unnatural, because in most large deformation situations, two corresponding vertices might have a large Euclidean distance, ending up with a small affinity,




Fig. 2. (a) S_2 is uniformly colored. The overall correspondences are indicated by the corresponding color in S_1 . (b) Correspondences derived from the confidence scores. (c) Correspondences derived from spectral graph matching.

even though there is a clear evidence showing the correspondence is correct and should have a high affinity. Therefore, we propose to compute the inter-surface affinity based on the confidence score of the initial correspondences.

We use an iterative max-row-column approach described in [5] to construct a set of t initial correspondences based on the confidence score matrix Δ . In each iteration, we select the largest non-zero element $\Delta_{i,j}$ and add (v_i, v_j) to the initial correspondence set. To avert non-one-to-one correspondences, we zero out the i th row and j th column of Δ . We repeat this procedure t times to select t most credible correspondences. The affinity matrix W is now defined as

$$w_{i,j} = \begin{cases} \|x_i - x_j\|^{-2} & \text{if } v_i, v_j \text{ are in same the surface,} \\ \Delta_{i,j} & \text{if } (v_i, v_j) \text{ is in the initial correspondence set,} \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

The final matching result is shown in Fig. 2c. As we can see, the correspondences are smoother than from the confidence scores directly.

4 Different Intrinsic Geometry

In our application, we have to register two surfaces with different intrinsic geometry, such as different boundary locations and holes (Fig. 5). Conventional separated spectral decompositions [12] in this situation will yield two totally different sets of eigenmodes. Just think of the simplest partial surface problem in Fig. 3a, in which one surface is a half of the other one. The first eigenmodes have distinct patterns, because surfaces with different sizes have different vibration modes. However, if we randomly assign only 5% initial correspondences between the two surfaces, as shown in Fig. 3b, the first eigenmodes become consistent with each other. Intuitively, we can achieve a joint vibration by patching the partial surface onto the other one using the initial links, so that the partial surface is forced to vibrate together with the other. Moreover, we can see this happening from the objective function (Eq. 4), because the energy is minimized when both intra-surface and inter-surface affinities are preserved in the spectral domain, which means corresponding vertices have similar embedded coordinates, as well as vibration properties.




Fig. 3. The first eigenmodes derived from (a) separated spectral decompositions. (b) a spectral decomposition on an association graph. (c) a spectral decomposition on an association graph with initial links only on one side.




Fig. 4. The color-coded correspondences (a,b) between a complete surface and a partial surface with a hole and truncation. (c,d) between surfaces with a bridge.

We find that the initial links have to be scattered all over the surface, but not necessarily dense. For example, the first vibration modes are shown in Fig. 3c, if the initial links are only on one side of the surface. Intuitively, two pieces of paper won't be stuck together if there is only one piece glued together.

Therefore, it is essential to find a credible set of initial links. As shown before, conventional spectral matching is not able to provide correct correspondences. However, our geometric feature descriptor has the advantage of providing robust initial links regardless of whether the overall surface being partial or not, because the correspondences are derived only using local geometric features. For the same reason, in most situations where the partial surface has holes in it, the joint vibration can still be achieved. Figs. 4a, 4b show the final matching result for a partial surface with a hole and a truncated boundary.

Our method can also handle some other simple topology changes. However, in many cases, regions with complicated topology changes usually yield inconsistent geometric features, which makes the initial correspondences there unstable. For example, as shown in Figs. 4c, 4d, there is a bridge connecting the epiglottis and the pharyngeal wall, and the correspondence there is not reasonable.

5 Results

We tested our method on 12 surface pairs created from 6 patients. The pharyngeal surface from the pharynx down to the vocal cord was automatically

segmented for the patient’s CT. Each surface has 2K-6K vertices, with an approximately $2\text{cm}\times 3\text{cm}$ elliptical cross section. We manually applied synthetic deformations to the surfaces with the help of a medical physicist, ending up with 12 surface pairs, two for each patient. The synthetic deformation includes the distortion and contraction of the pharyngeal wall and the closing and opening of the laryngeal region and of the epiglottis. We measured the registration error of each vertex as the Euclidean distance between the resulting corresponding vertex and the ground truth. The registration error for each surface pair is defined as the average registration error over all vertices.

We studied the optimal choice of different parameters. 15 eigenmodes were used to perform the final matching. The size of the initial correspondence set was chosen as half the number of vertices. We set the geodesic path distance d to 4mm and the Euclidean distance threshold τ to 1cm. All the parameters were chosen according to a different set of surface pairs.

The average registration error for the 12 surface pairs using initial links derived by different options is provided in Table 1. In the first option, we used conventional spectral matching to compute a dense set of initial correspondences. In the second option, we used our method except that the inter-surface affinity was computed by Euclidean distance between two corresponding vertices in the original 3D space. The third option was exactly our method. We tested all options in both scenarios: complete surface matching and partial surface matching. In the partial surface matching context, we picked one surface from each pair and manually created holes in large deformation regions, such as the epiglottis, and truncated the surface in a different location. The registration error was only measured for boundary vertices for partial surface matching. We also ran the algorithm on several real data cases with large topology change. One of the results is shown in Figs. 5c, 5d.

Table 1. Registration error for complete surface and partial surface matching

	Complete Surface (mm)	Partial Surface (mm)
Initial Error	3.09 ± 1.73	3.48 ± 1.79
1. Conventional Spectral Matching	1.83 ± 2.37	3.26 ± 6.71
2. Feature + Euclidean Distance	1.38 ± 2.55	1.90 ± 2.15
3. Feature + Confidence Score	0.67 ± 0.96	1.15 ± 1.36




Fig. 5. The pharynx. (a) A CT segmentation. (b) An endoscopic video reconstruction. (c,d) Color-coded correspondences between a CT surface and a real reconstruction.

6 Conclusion

We have presented a non-rigid surface registration method based on spectral graph matching with the application of registering pharyngeal surfaces in CT/Endoscope fusion. We proposed an efficient approach for extracting initial correspondences using our novel geometric feature descriptor. The association graph based on this kind of initial correspondences produces better registrations. We showed the method's potential to handle partial surface matching and discussed its disadvantages when dealing with complicated topology change. Our results suggest that this approach might be applicable to other surface registrations with large deformations, holes and different boundary locations.

Acknowledgements. This work was supported by NIH grant R01 CA158925. The authors would like to thank Dr. Ron Alterovitz, Dr. Jan-Michael Frahm, Dr. Bisham Chera, Hina Shah and Federico Menozzi for the meaningful discussion and the preprocessing and 3D reconstruction work.

References

1. Allen, B., Curless, B., Popovic, Z.: The space of human body shapes: reconstruction and parameterization from range scans. In: ACM SIGGRAPH, vol. 22, pp. 587–594 (2003)
2. Li, H., Luo, L., Vlastic, D., Peers, P., Popovic, J., Pauly, M., Rusinkiewicz, S.: Temporally coherent completion of dynamic shapes. *ACM Transactions on Graphics* 31(2) (2012)
3. Huang, Q., Adams, B., Wicke, M., Guibas, L.: Non-rigid registration under isometric deformations. In: Symposium on Geometry Processing, pp. 1449–1457 (2008)
4. Zeng, W., Gu, X.D.: Registration for 3d surfaces with large deformations using quasi-conformal curvature flow. In: Computer Vision and Pattern Recognition, pp. 2457–2464 (2011)
5. Lipman, Y., Funkhouser, T.: Möbius voting for surface correspondence. *ACM Transactions on Graphics* 28(72) (2009)
6. Zigelman, G., Kimmel, R., Kiryati, N.: On bending invariant signatures for surfaces. *IEEE Trans. Vis. Comput. Graph.* 8(2), 198–207 (2002)
7. Sharma, A., Horaud, R.: Shape matching based on diffusion embedding and on mutual isometric consistency. In: IEEE CVPRW, pp. 29–36 (2010)
8. Johnson, A.E., Hebert, M.: Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Transactions on PAMI* 21(5), 433–449 (1999)
9. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: Eurographics Symposium on Geometry Processing 2009, vol. 28, pp. 1383–1392 (2009)
10. Lombaert, H., Sporring, J., Siddiqi, K.: Diffeomorphic spectral matching of cortical surfaces. In: Gee, J.C., Joshi, S., Pohl, K.M., Wells, W.M., Zöllei, L. (eds.) *IPMI 2013*. LNCS, vol. 7917, pp. 376–389. Springer, Heidelberg (2013)
11. Koenderink, J.: *Solid Shape*. The MIT Press (1990)
12. Lombaert, H., Grady, L., Polimeni, J., Cheriet, F.: Fast brain matching with spectral correspondence. In: Székely, G., Hahn, H.K. (eds.) *IPMI 2011*. LNCS, vol. 6801, pp. 660–673. Springer, Heidelberg (2011)