Surface Registration in the Presence of Missing Patches and Topology Change

Qingyu Zhao[1]
zenyo@cs.unc.edu
True Price[1]
jtprice@cs.unc.edu
Stephen Pizer[1]
smp@cs.unc.edu
Marc Niethammer[1]
mn@cs.unc.edu
Ron Alterovitz[1]
ron@cs.unc.edu
Julian Rosenman[2]
julian_rosenman@med.unc.edu

[1] Department of Computer Science, University of North Carolina at Chapel Hill, North Carolina, USA
[2] Department of Radiation Oncology, University of North Carolina at Chapel Hill, North Carolina, USA

Abstract

The fusion between an endoscopic movie and a CT poses a special surface registration problem. The surface extracted from CT is complete and accurate, whereas the surface extracted from endoscopy suffers from serious missing patches and topology change. We propose a surface registration method, Thin Shell Demons, that is robust under these two situations. Motivated by Thirion's Demons idea, the partial surface can provide virtual forces to attract the complete surface, which is equipped with a novel physics-based deformation energy. This energy can help preserve the correct surface topology while producing realistic deformation for the regions that don't have any attracting counterpart regions. The attraction direction assures the deformation is not affected by the surface completeness. Moreover, we propose to use geometric feature matching for computing virtual forces to handle inaccurate 3D point positions and large deformations. We test our method for CT/endoscope fusion and show its potential to achieve successful registration.

1 Introduction

Partial matching is a special problem in surface registration. In the context of CT/endoscope fusion for head-and-neck radiation treatment, a deformable registration has to be carried out between a CT segmentation surface (Fig. 1c) and a surface reconstructed from endoscopy. The challenge is that the endoscopic reconstruction has much missing data, including holes and truncations (Fig. 1b), due to some anatomy being not able to show up in the camera view (Fig. 1a). Moreover, the apparent topology of two surfaces may differ, due to two different anatomical regions touching each other. For example, a bridge is created between the epiglottis tip and the pharyngeal wall in Fig. 1b.

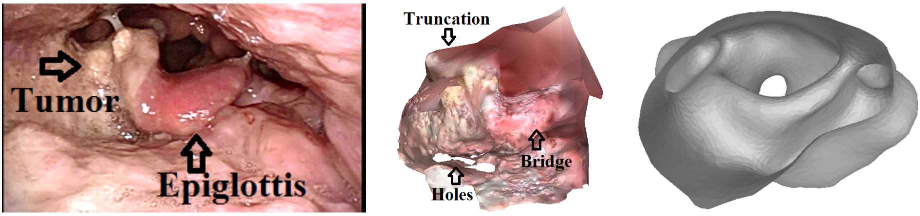

(a) (b) (c)

Figure 1: (a) An endoscopic movie frame. (b) A surface reconstruction derived from the Structure-from-Motion pipeline. (c) A CT segmentation surface.

Many current methods can deal with large non-rigid deformations under different contexts, but they do not handle missing data and topology change well because these properties violate the methods' underlying assumptions, such as near-isometric deformation [6, 7], angle preservation [5], or uniform compact support (surface completeness) [10]. Our work is motivated by Thirion's Demons algorithm [9], which was originally proposed for image registration and then used for surface registration [1].

The Demons idea has minimum assumptions about deformation/surface properties; it does not require surface completeness and identical topology. Based on this observation, we propose a physics-based surface registration method, Thin Shell Demons, that is robust to these two problems. In the fixed partial surface S, we define a set of demons $\{v\}$ that can provide virtual attraction forces $\{F(v)\}$, so the deformable complete surface M can be gradually attracted to S in a smooth fashion while preserving a correct topology. The direction of this attraction is essential. Without considering the topology change, the mapping $\Phi : S \to M$ is injective and non-surjective, which means for any $v \in S$, there must exist a uniquely defined corresponding point $\Phi(v)$ that should be attracted, whereas the opposite is not true. This justifies that the attraction direction we choose will not be affected by the missing patches. Moreover, the physics-based energy guarantees a realistic topology-preserving deformation, especially for the regions that are not being attracted. The next section will discuss the method pipeline and the proposed physics-based deformation model and geometry-based attraction forces.

2 Thin Shell Demons Algorithm

In Thirion's Demons algorithm, the notion is that the object boundary S in a static image contains a set of so-called demons that can produce virtual forces to attract the object boundary M in a moving image until the two images are close. This idea is naturally extendable to surface registration in 3D. However, with no intensity information available in surface registration, there are no easy ways to compute Demons forces and to regularize surface deformation.

In the context of CT/endoscope surface registration, surface deformation is mostly induced by an underlying physical process caused by the muscles. To produce physically realistic deformations, we regard the moving surface M as an elastic thin shell. Its deformation energy consists of local elastic energy and novel non-local structural energies that can preserve important 3D structures' shape. Another special property in this registration problem is that endoscopic reconstruction surfaces frequently produce inaccurate 3D point depth but relatively densely produce local geometry, including curvatures and normal directions, so we propose to use geometric feature matching to produce attraction forces, which is more robust than the closest-point strategy under large deformations. An additional novel feature is that our method works directly on discrete triangle meshes without using any implicit level-set representation [1, 3]. This allows more flexible triangulation and higher surface resolution.

2.1 Thin Shell Deformation Model

We first study the deformation energy of a thin shell model. 3D thin shells are bounded by two curved surfaces that are bisected by a middle surface $\mathbf{M}$, where the distance between the surfaces is much smaller than their overall dimension. The deformation energy of the thin shell is usually approximated by integrating local membrane and bending energy of $\mathbf{M}$'s deformation, but we add the aforementioned non-local structural energy.

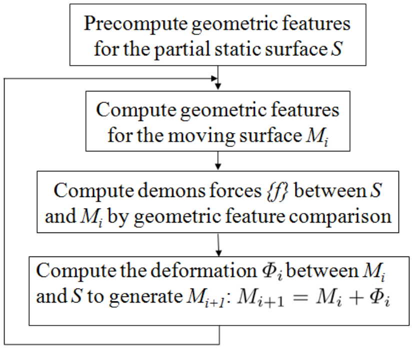

Figure 2: Thin Shell Demons workflow.

Membrane Energy. Membrane strains represent stretching and shearing effects of the local deformation and can be fully characterized by the tangential Cauchy-Green strain tensor: $\sigma_{mem}(p) = J_\phi(p)^T J_\phi(p)$, where $J_\phi(p)$ is the Jacobian of the tangential transformation $\phi : T_p \to T_{\Phi(p)}$ at local point p. The total membrane energy is given by $E_{mem} = \int_{\mathbf{M}} W(\sigma_{mem}(p))$, where

$$W(\sigma) = \frac{\mu}{2}\mathrm{tr}(\sigma) + \frac{\lambda - 2\mu}{8}\det(\sigma) + \frac{\lambda + 2\mu}{8}\det(\sigma)^{-1}. \tag{1}$$

λ and μ are the Lamé parameters of the tissue, which can be determined experimentally. We adopt the same discretization of the membrane strain used in [11], in which the local tangential transformation ϕ is just the linear shape transformation of the triangle.

Bending Energy. The bending strain measures local curvature changes. We know that the local curvature information is fully characterized by its second fundamental form expressed in tangent plane coordinates, which can be written as a 2×2 tensor operator (shape operator) Λ_p such that $v^T \Lambda_p v$ gives the normal derivative (curvature) in the direction of tangent vector v. The bending strain σ_{bend} is computed by the difference between two shape operators under deformation: $\sigma_{bend}(p) = \phi^T \widetilde{\Lambda}_{\Phi(p)} \phi - \Lambda_p$. The total bending energy is given by the $E_{bend} = \int_{\mathbf{M}} \|\sigma_{bend}(p)\|_F^2$. To be consistent with the discrete membrane strain, we use a triangle-ring stencil to compute a discrete shape operator for each triangle. [2]

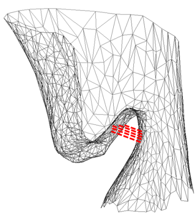

Figure 3: Cross-object links (red) are placed between the frontal and posterior wall of the epiglottis.

Structural Energy. In our problem the surface is a boundary representation of solid tissues that show shape integrity at certain regions. For example, the epiglottis does not change its thickness. To handle this, we propose to add manually placed structural links into the thin shell model (Fig. 3). In our problem, five cross-object structural links $\{L_c^i | i = 1,..,5\}$ are connected between two subregions of the surface, namely the frontal and posterior wall of the epiglottis, to preserve the epiglottal thickness. The related structural energy is defined as $E_{L_c} = \sum_i (\Delta|L_c^i|)^2$, where $\Delta|\cdot|$ is the length change of the link.

Finally, the augmented thin shell energy E_{shell} is the weighted sum of all the energy terms: $E_{mem}, E_{bend}, E_{L_c}$.

2.2 Geometric-Feature-Based Demons Force

In order to effectively match the shapes of the two surfaces, we want the demons forces to attract similar geometric structures towards each other. Therefore, we use a feature descriptor to produce local geometry information, which will be used later to compute demons forces.

It has been shown that the feature descriptor introduced in [12] is effective in finding similar geometrical structures due to its ability to capture the local shape by measuring curvatures at different scales. For each vertex pair $\{v,u\}$, the feature distance $\delta(v,u)$ is computed. Then for each vertex $v \in M_i$, where M_i is the moving surface in the i^{th} iteration, we find the most geometrically similar vertex $m(v) \in S$ as its matching point, such that $\delta(v,m(v)) = \min\{\delta(v,u)|u \in S\}$. Then the force vector $F(v) = (m(v) - v)$ defines a virtual force applied on v induced by a demon $m(v) \in S$. Forces computed in this way may contain outliers, so we associate each force vector $F(v)$ with a confidence score $c(v) = e^{-\delta(v,m(v))}$ based on the feature distance, indicating how likely the force vector is accurate.

2.3 Computing Deformation

Given $\{F\}$, we need to compute the deformation ϕ_i that brings M_i closer to S. Then the problem becomes to solve the thin shell deformation induced by the external force $\{F\}$.

In order to incorporate the confidence score information of $\{F\}$, we formulated an optimization framework for Thin Shell Demons. The deformation Φ_i can be approximated by minimizing the objective function

$$E(\Phi) = \int_M c(v)(\Phi(v) - F(v))^2 dv + E_{shell}(\Phi). \tag{2}$$

The first part penalizes the inconsistency between the deformation vector and the force vector applied on a point and uses the confidence score to weight the penalization. The second part minimizes the augmented thin shell deformation energy. This optimization problem is solved using the L-BFGS method.

3 Experiments

Synthetic Deformation. We first tested Thin Shell Demons with synthetic deformations. We extracted 6 patients' CT surfaces, each of which has approximately 3500 vertices. For each surface, we manually applied 4 synthetic deformations, two of which contained truncations and holes. Thus we had 24 surface pairs: 12 for complete surface registration, and 12 for partial matching. The synthetic deformations contain the most realistic deformations expected to be seen in real data, such as the stretching of the pharyngeal wall and the bending of the epiglottis. The registration error for a surface pair is defined as the average error over all vertices. The error for a single vertex is measured by the Euclidean distance between its resulting corresponding vertex and the ground truth. The Lamé parameters in Eq. 1 and the energy weighting parameters were chosen based on the behavior of two separately created synthetic deformations.

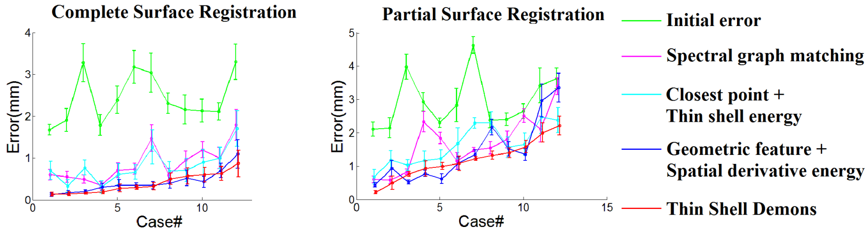

Figure 4: The case indices are reordered according to the final error produced by our method. Each vertical bar shows the magnitude of the standard deviation for each case.

We compared Thin Shell Demons with three other strategies suggested in other works to indicate the importance of the choices made in our method. The spectral graph matching method was proposed in [12]. Closest point + thin shell energy uses a closest point search [4] to drive the deformation instead of using geometric features. Geometric feature + spatial derivative energy [8] uses the spatial derivative of the deformation field to compute the deformation energy instead of using the thin shell energy. Since the three comparison methods are fully automatic, we left out the manually constructed structural links in the synthetic test to make the comparison fair. Also, the energy weighting parameters were re-tuned for each method to produce the best results. From Fig. 4, we can see that Thin Shell Demons outperforms the other methods in most cases.

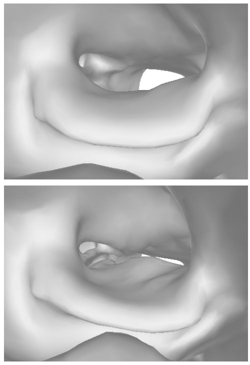

Figure 5: Registration results w/ (top) and w/o (bottom) structural links.

Real Reconstruction. Figs. 5&6 show some qualitative results of registration of real data. Two algorithms were used to produce 3D reconstructions from endoscopic movies. The Structure-from-Motion (SFM) pipeline can produce a 3D surface from successive movie frames (Fig. 5 top), and the Shape-from-Shading (SFS) algorithm can produce a depth map for each single frame (Fig. 5 bottom). The cross-object structural links were added in real data registration. We can see that the CT and reconstruction surfaces were aligned reasonably well after registration. Fig. 6 shows that the structural links can help preserve the epiglottis shape in the registration.

4 Conclusion

To handle missing patches and topology change, we have proposed a physics-motivated surface registration method: Thin Shell Demons. We experimented with this method for pharyngeal surface registration and showed its potential to achieve successful registration. In our method, we computed the virtual attraction forces by geometric feature matching. Experiments have shown that this strategy is more robust than the closest-point rule. We proposed an augmented thin shell model as the physical model of the pharyngeal surface. The model keeps the advantage of being simple to analyze while incorporating 3D integrity information without modeling the entire 3D structure.

Acknowledgements This work was supported by NIH grant R01 CA158925.

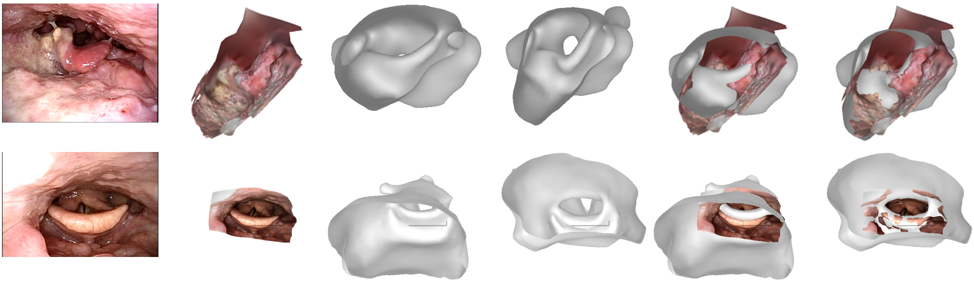

Figure 6: From left to right: An endoscopic movie frame; Endoscopic reconstruction; CT surface; Registered CT; Initial overlay; Final overlay.

References

[1] A. Dedner, M. Lüthi, T. Albrecht, and T. Vetter. Curvature guided level set registration using adaptive finite elements. *Pattern Recognition*, 4713:527–536, 2007.

[2] E. Grinspun, Y. Gingold, J. Reisman, and D. Zorin. Computing discrete shape operators on general meshes. *Comput. Graph. Forum*, 25(3):547–556, 2006.

[3] J. A. Iglesias, B. Berkels, M. Rumpf, and O. Scherzer. A thin shell approach to the registration of implicit surfaces. In *Vision, Modeling, and Visualization*, pages 89–96, 2013.

[4] H. Li, R. W. Sumner, and M. Pauly. Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum, SGP 08*, 27(5), 2008.

[5] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. *ACM Transactions on Graphics*, 28(72), 2009.

[6] H. Lombaert, J. Sporring, and K. Siddiqi. Diffeomorphic spectral matching of cortical surfaces. In *IPMI 2013*, volume 7917, pages 376–389, 2013.

[7] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG) - SIGGRAPH 2012 Conference Proceedings*, 31(4), 2012.

[8] M. Pauly, N.J. Mitra, J. Giesen, L. Guibas, and M. Gross. Example-based 3D scan completion. In *Symposium on Geometry Processing*, pages 23–32, 2005.

[9] J.P. Thirion. Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis*, 2(3):243–260, 1998.

[10] M. Vaillant and J. Glaunès. Surface matching via currents. In *IPMI*, volume 3565, pages 381–392, 2005.

[11] Y. Zeng, C. Wang, X. Gu, D. Samaras, and N. Paragios. A generic deformation model for dense non-rigid surface registration: A higher-order MRF-based approach. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3360–3367, 2013.

[12] Q. Zhao, S. Pizer, M. Niethammer, and J. Rosenman. Geometric-feature-based spectral graph matching in pharyngeal surface registration. In *MICCAI*, volume 1, pages 259–266, 2014.