

Improving Deformable Surface Meshes through Omni-directional Displacements and MRFs

D. Kainmueller, H. Lamecker, H. Seim, S. Zachow, and H.-C. Hege
kainmueller@zib.de

Zuse Institute Berlin, Germany

Abstract. Deformable surface models are often represented as triangular meshes in image segmentation applications. For a fast and easily regularized deformation onto the target object boundary, the vertices of the mesh are commonly moved along line segments (typically surface normals). However, in case of high mesh curvature, these lines may intersect with the target boundary at “non-corresponding” positions, or even not at all. Consequently, certain deformations cannot be achieved. We propose an approach that allows each vertex to move not only along a line segment, but within a surrounding sphere. We achieve globally regularized deformations via Markov Random Field optimization. We demonstrate the potential of our approach with experiments on synthetic data, as well as an evaluation on 2x106 coronoid processes of the mandible in Cone-Beam CTs, and 56 coccyxes (tailbones) in low-resolution CTs.

1 Introduction

Deformable surface models are widely used for image segmentation [1]. Among the many different representations of surfaces, polygonal meshes are advantageous in many respects, such as flexibility or topology preservation [2]. The deformation of the model is often driven by minimizing an energy that consists of an image term that measures how well the model is aligned with features in the image, and a regularization term that controls the smoothness of the model. During energy minimization, the vertices of the mesh are iteratively displaced.

A fundamental challenge of this approach is how to keep the mesh valid, i.e. as regular as possible [3]. One way to confront this problem is to allow vertex displacements only in surface normal direction. Additional measures are adaptive step-size control, adaptive remeshing, special regularization or mesh surgery. The drawbacks of these approaches are that they are often difficult to implement, slow down the convergence of the method drastically, or even lead to situations where the deformation gets stuck. This is especially the case in regions of high curvature. Hierarchical approaches may be able to circumvent this problem: First the model is adapted with strong regularization, e.g. by using shape priors such as statistical shape models [4], which is then gradually relaxed. Although this increases robustness, the last bit of accuracy may only be achieved with a deformation model free of a-priori shape constraints (see e.g. [5]). This is the type of deformation model we focus on.




Fig. 1. Normals on a cube (a) and on a tip-like structure (c). Unregularized deformations along normals: No target boundary points (b) or non-corresponding target boundary points (d) found for most vertices.

Normal (or other one-directional) displacements of mesh vertices may either lead to very few features being detected in the image (low *visibility*), or many “wrong” features (in terms of *correspondence*), as depicted in Fig. 1. This induces large distortions of the mesh. Remeshing may restore mesh regularity, while implausible shape deformations are to be remedied by regularization. In summary, certain deformations can hardly be achieved with one-directional displacements, given that directions are not known a-priori. E.g., parallel movements of tip-like structures cannot be achieved along surface normals.

This paper contributes a solution to this problem: We propose a method that allows arbitrary displacement directions at each mesh vertex. This enlarges visibility, while correspondence is likely to be improved, too. However, a larger set of image features may be found per vertex, so global regularization is required to cope with highly inconsistent neighboring displacements. Our method allows displacements to a discrete set of points within a sphere around each vertex. This discrete nature allows us to formulate the segmentation problem as a Markov Random Field (MRF), as shown in Sec. 2. The MRF can be solved efficiently [6], yielding a globally regularized deformation of the mesh. MRFs have been applied to many problems, e.g. image-to-image registration [7], due to their capability of finding good optima. Global regularization has been proposed for one-directional displacements using graph cuts [8]. However, this approach is not applicable in our case as it requires a total order on displacements in terms of “above/below”.

We illustrate on synthetic and real data that omni-directional displacements combined with MRF optimization can handle parallel movements of meshes with high curvature, where previous approaches based on normal displacements fail.

2 Method

We denote the set of vertices $v \in \mathbf{R}^3$ of the deformable surface mesh as V , and the set of pairs of adjacent (i.e. edge-connected) vertices as $N \subset V \times V$. Each vertex can be moved by adding a vector $l \in L$, where $L \subset \mathbf{R}^3$ is a discrete set




Fig. 2. Three vertices (black dots) on a 2D contour. (a) Omni-directional displacements to yellow/gray dots. (b) Exemplary “same” displacements shown by black arrows with corresponding numbers. (c) Same displacement for all vertices leads to parallel translation.

of possible displacements. We call a position $v + l$ *sample point*, and a mapping $V \rightarrow L, v \mapsto l_v$ that assigns a displacement to each vertex *displacement field*.

2.1 Omni-directional Displacements

We propose to extend the range of motion for a vertex from a line segment to a sphere centered at its current position. Therefore, we define L as a set of displacements that are uniformly distributed within a sphere (see Fig. 2a). The sphere radius is a parameter of our method. The set of displacements is interpreted in world coordinates for all vertices (see Fig. 2b and 2c). As discussed in Sec. 2.2, this is important for regularization.

2.2 Objective Function

For each displacement $l \in L$ and vertex $v \in V$, a scalar cost $c(v, l) \geq 0$ encodes whether sample point $v + l$ is believed to lie on the object boundary. The stronger the belief, the lower should be the cost. In other words, $c(v, l)$ serves as a penalty for the case that v is displaced by l . We calculate $c(v, l)$ from the image $I : \mathbf{R}^3 \rightarrow \mathbf{R}$. It depends on $I(v + l)$, $\nabla_{n_v} I(v + l)$ (where n_v denotes the surface normal at v), and application-specific parameters (see Sec. 3). In general, our objective function accepts any $c : V \times L \rightarrow \mathbf{R}_0^+$, so c can be defined as appropriate. E.g., if the surface mesh might locally lie perpendicular to the object boundary, derivatives in directions other than n_v may be considered as well.

For each two displacements l_1, l_2 , a scalar distance value $d(l_1, l_2) \geq 0$ serves as a penalty for the case that l_1 and l_2 occur on adjacent vertices. The distance function $d : L \times L \rightarrow \mathbf{R}_0^+$ is supposed to take care of regularization. It has to satisfy $d(l_1, l_2) = 0 \Leftrightarrow l_1 = l_2$, but does not have to be a metric (see Sec. 2.3). It can, e.g., be the Euclidean norm $\|l_2 - l_1\|$ to some power.

We define the objective function of the mesh adaptation problem as follows:

$$\sum_{v \in V} c(v, l_v) + \sum_{(v, w) \in N} d(l_v, l_w) \quad (1)$$

We are looking for the displacement field that minimizes Eq. 1. Note that interpreting displacements in world coordinates yields distance-penalties for scaling the mesh, while parallel translations are not penalized (see Fig. 2c). We consider this beneficial as we expect our initial meshes (as well as its local features) to have approximately correct scale. Alternatively, if scaling should not be penalized, one could interpret displacements in local coordinate systems per vertex.

2.3 Optimal Displacement Field

We encode the objective function in Eq. 1 as an MRF, with vertices being represented by MRF-nodes, mesh edges by MRF-edges, and displacements by the possible states (also called *labels*) of the nodes. Cost $c(v, l)$ defines the unary potential of node v in state l , and distance $d(l_1, l_2)$ defines the binary potential of two adjacent nodes in states l_1, l_2 . The MRF-state with minimal sum of potentials yields the desired displacement field. We optimize the MRF by a method named FastPD [6]. FastPD can deal with non-metric distance functions d as specified in Sec. 2.2. It is guaranteed to find an approximately optimal solution.

3 Results

To evaluate our MRF-based method, we applied it to three types of 3D data: (1) Synthetic binary images, (2) synthetic binary images with various amount of noise, and (3) clinical image data. On synthetic binary images and clinical image data, we also computed results with a globally regularized method [8] (GraphCuts) and a locally regularized method [5] (FreeForm), both employing vertex normals as displacement directions. GraphCuts computes the displacement field with minimum sum of costs while respecting a hard constraint on the difference between the lengths of adjacent displacements. FreeForm takes the minimum cost displacement for each vertex and regularizes locally via a small displacement toward the centroid of the respective adjacent vertices.

For the computation of the costs $c(v, l)$, we used the strategy proposed in [9]: If the intensity $I(v + l)$ lies within a certain window $[i_0, i_1]$, costs are inversely proportional to $\nabla_{n_v} I(v + l)$. Otherwise costs are set to a constant, high value. The thresholds i_0 and i_1 are parameters of the strategy. As distance function d , we used $d(l_1, l_2) = ||l_2 - l_1||^3$ in all experiments. Whenever we employ GraphCuts or FreeForm, we use the same cost function as for the respective MRF experiment, and normal segments with the length of the respective sphere diameter. In contrast to the GraphCuts- and MRF based adaptations, all FreeForm adaptations were performed iteratively, with 30 steps. All GraphCuts experiments were performed with a difference constraint of 2 sample points (i.e. lengths of adjacent displacements can be at most 2 sample points apart).

MRF optimization with FastPD took less than 10 seconds in all our experiments. The computation of the MRF’s unary potentials $c(v, l)$ was more time-consuming, taking up to 10 minutes depending on the number of vertices and labels. FastPD required memory up to 8GB for the experiments on clinical data. All experiments were performed on a 3GHz core with 8GB main memory.




Fig. 3. Deformable mesh (red/dark grey mesh) and target object (transparent gray surface). (a) Initial situation. (b) Displacements along normals without regularization. (c) FreeForm approach. (d) GraphCuts approach. (e) MRF approach.




Fig. 4. Performance of MRF approach in the presence of noise. Random noise with range (a) $[-0.5..0.5]$, (b) $[-2.5..2.5]$ and (c) $[-5..5]$. Slices of the image data and respective adaptation result (red/dark grey mesh). Grey surface: ideal target object.

Synthetic Images. We performed experiments on binary images (i.e. intensities $\in \{0, 1\}$) of a cube and a thin ellipsoid. As initial meshes, we used triangulated cube and tip surfaces with ideal shape, but shifted pose (see Fig. 3(a)). The cube mesh had 770, the ellipsoid 1797 vertices. One sphere diameter (or surface normal segment, respectively) was covered by 53 sampling points for the cube, and 63 for the ellipsoid. For all experiments on synthetic data, we chose sphere radii such that the target object boundary was located completely within a band around the initial mesh that has this radius. We set the cost function parameters to $i_0 = 0.1$ and $i_1 = 1.1$. The results of MRF-, FreeForm- and GraphCuts adaptation are shown in Fig. 3(c), 3(d) and 3(e), respectively. Fig. 3(b) shows the results of adding normal displacements without any regularization.

We added various amounts of random noise to the binary cube image and performed MRF based adaptation as before. The cube was detected correctly for noise with ranges $[-0.5..0.5]$ and $[-2.5..2.5]$, and failed for $[-5..5]$. Fig. 4 shows slices of the noisy image data and the respective adaptation results.

Clinical Data. In a quantitative evaluation on 106 mandible Cone-Beam CTs and 50 pelvis CTs we compared MRF, FreeForm and GraphCuts results to gold standard surfaces obtained from manual segmentations. Initial meshes were

automatically generated by Statistical Shape Model adaptation [9, 10]. For the mandible surfaces (gold standard as well as adaptation result), we extracted the right coronoid processes as the region of the mesh that lies above 1/2 of the extension of the mandible in transversal direction, between 1/3 and 2/3 of extension in dorsoventral direction, and above 2/3 in longitudinal direction. Extraction of the left coronoid process worked analogously. We identified the tip point as the upmost vertex in longitudinal direction. For the sacrum, we extracted the coccyx as the region of the mesh that lies below 1/3 of the extension of the sacrum in longitudinal direction. We identified the coccyx tip as the vertex with minimum $3 \cdot \text{longitudinal} - 1 \cdot \text{dorsoventral}$ coordinate. As error measures we assessed the tip-to-tip distances (`tip2tip`), tip-to-surface distances from result tip to gold standard surface (`tip2surfRtG`) and the other way round (`tip2surfGtR`), and roots mean square surface distances (`rmsRtG` and `rmsGtR`).

All mandible meshes had 8561 vertices, all sacrum meshes 6161. The number of samples along a sphere diameter (or normal segment) was 39 for the mandible and 45 for the sacrum. The sphere diameters were 15 mm (mandible) and 25 mm (sacrum). We set the cost function parameters to $i_0 = 350$ and $i_1 = 800$ (mandible) and $i_0 = 120$ and $i_1 = 520$ (sacrum). For the mandible, we performed a second MRF based adaptation (`mrfZ`) with a slightly different cost function: We added a small extra cost to sample points with lower longitudinal coordinate, thus slightly preferring the a motion in upward direction.

Evaluation results are plotted as histograms in Fig. 5. As measurements are not normally distributed, we performed the Wilcoxon signed-rank test to asses significant differences (level 0.05). For the mandible, both MRF results are significantly better than FreeForm and GraphCuts in terms of `tip2tip`, `rmsGtR` and `rmsRtG`, while `mrfZ` performs significantly better in terms of `tip2surfGtR`, too. For the sacrum, MRF results are significantly better than GraphCuts in terms of all but the `rmsGtR` measure. However, compared to FreeForm, the error difference is significant only for the two rms measures.

4 Discussion

Experiments on synthetic binary images show that our MRF-based approach is able to handle parallel translations, in contrast to a globally and a locally regularized approach (GraphCuts and FreeForm) that employ normal displacements. Experiments on noisy synthetic images show that the MRF approach is able to produce well-regularized displacement fields in the presence of noise. However, for a very low signal-to-noise ratio, it may fail to detect the target object. Real-world experiments show that the MRF approach is able to produce very accurate segmentations of tip-like structures. On the mandible tips, the MRF approach clearly outperforms the GraphCuts and FreeForm approach. Here, normal directions often exhibit the “wrong-visibility” problem, see Fig. 6(a-c), which our new method resolves. However, on the sacrum tips, more experiments need to be performed to draw decisive conclusions. At least, the “non-visibility” problem for normals can be resolved in a few exemplary cases, see Fig. 6(d,e).




Fig. 5. Row 1: Coronoid process results on 106 individuals (i.e. 212 cases). Row 2: Coccyx results on 50 individuals. X-axis: Error measure in mm (see Sec. 3). Y-axis: Frequency. MRF (mrf, mrfZ), FreeForm (free) and GraphCuts (gc).

5 Conclusion

We proposed a method that allows omni-directional displacements for all vertices of a surface mesh during deformable model adaptation. We achieve global regularization by encoding the adaptation problem as a Markov Random Field, which we then optimize with a fast approximate solver. In an evaluation on synthetic as well as clinical data, we showed that this approach can outperform traditional mesh adaptation along line segments (normals) in regions with high curvature (tips) in terms of segmentation accuracy. In this paper, we focused on parallel translations and tip-like structures. A closer investigation of our MRF based method in situations where scaling and rotation is desired, as well as an extended quantitative evaluation that considers all regions on mandible and pelvis shall be performed in future work. Furthermore, computational performance (both in terms of memory and speed) shall be enhanced, possibly involving a hybrid deformation model that employs omni-directional displacements in regions with high curvature and normal displacements elsewhere. Last but not least, a question of interest is if the MRF based method can be extended to simultaneous adaptation of multiple, adjacent meshes, i.e. multi-object segmentation.




Fig. 6. Exemplary results on clinical data. (a-c) Coronoid processes of the mandible. (d,e) Coccyx tips of the sacrum. Contours: Black: Gold standard. White: Initial mesh. Green/gray: MRF result. Blue/light gray: FreeForm result.

Acknowledgements D. Kainmueller is funded by the DFG Collaborative Research Center SFB760. H. Lamecker is funded by the German Research Center MATHEON. Thanks to M. Zinser (Universitätsklinikum Köln, Germany) and M. Heller (Julius Wolff Institut, Charité Berlin, Germany) for providing image data and manual segmentations.

References

1. McInerney, T., Terzopoulos, D.: Deformable models in medical image analysis: A survey. *Medical Image Analysis* **1** (1996) 91–108
2. Montagnat, J., Delingette, H., Ayache, N.: A Review of Deformable Surfaces: Topology, Geometry and Deformation. *Image Vis. Comput.* **19** (2001) 1023–1040
3. Park, J.Y., McInerney, T., Terzopoulos, D., Kim, M.H.: A non-self-intersecting adaptive deformable surface for complex boundary extraction from volumetric images. *Computers & Graphics* **25**(3) (2001) 421–440
4. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active Shape Models - Their Training and Application. *Comput. Vis. Image Underst.* **61**(1) (1995) 38–59
5. Kainmueller, D., Lange, T., Lamecker, H.: Shape Constrained Automatic Segmentation of the Liver based on a Heuristic Intensity Model. In: *3D Segmentation in the Clinic: A Grand Challenge*. (2007) 109–116
6. Komodakis, N., Tziritas, G., Paragios, N.: Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal-dual strategies. *Comput. Vis. Image Underst.* **112**(1) (2008) 14–29
7. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. *MIA* **12**(6) (2008) 731 – 741 Special issue on information processing in medical imaging 2007.
8. Li, K., Wu, X., Chen, D.Z., Sonka, M.: Optimal Surface Segmentation in Volumetric Images - A Graph-Theoretic Approach. *IEEE TPAMI* **28**(1) (2006) 119–134
9. Seim, H., Kainmueller, D., Heller, M., Lamecker, H., Zachow, S., Hege, H.C.: Automatic Segmentation of the Pelvic Bones from CT Data Based on a Statistical Shape Model. In: *Proc. VCBM*. (2008) 93–100
10. Kainmueller, D., Lamecker, H., Seim, H., Zinser, M., Zachow, S.: Automatic Extraction of Mandibular Nerve and Bone from Cone-Beam CT Data. In: *MICCAI* (1). Volume 5762 of LNCS., Springer (2009) 76–83