Medical Image Analysis 17 (2013) 429–441

Contents lists available at SciVerse ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Omnidirectional displacements for deformable surfaces

Dagmar Kainmueller [a,*], Hans Lamecker [a], Markus O. Heller [b], Britta Weber [a], Hans-Christian Hege [a], Stefan Zachow [a]

[a] Zuse Institute Berlin, Berlin, Germany
[b] Julius Wolff Institute, Charité – Universitätsmedizin, Berlin, Germany

ARTICLE INFO

Article history:
Received 13 January 2012
Received in revised form 25 July 2012
Accepted 30 November 2012
Available online 14 December 2012

Keywords:
Segmentation
Deformable surfaces
Markov random field
Mandible
Acetabulum

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 not intersect with the target boundary at all. Consequently, certain deformations cannot be achieved. We propose *omnidirectional displacements for deformable surfaces* (*ODDS*) to overcome this limitation. ODDS allow each vertex to move not only along a line segment but within the volumetric inside of a surrounding sphere, and achieve globally optimal deformations subject to local regularization constraints. However, allowing a ball-shaped instead of a linear range of motion per vertex significantly increases runtime and memory. To alleviate this drawback, we propose a hybrid approach, *fastODDS*, with improved runtime and reduced memory requirements. Furthermore, fastODDS can also cope with simultaneous segmentation of multiple objects. We show the theoretical benefits of ODDS with experiments on synthetic data, and evaluate ODDS and fastODDS quantitatively on clinical image data of the mandible and the hip bones. There, we assess both the global segmentation accuracy as well as local accuracy in high curvature regions, such as the tip-shaped mandibular coronoid processes and the ridge-shaped acetabular rims of the hip bones.

1. Introduction

In this paper, we address the issue of segmenting highly curved anatomical structures in three-dimensional medical image data. The aim is to improve segmentation accuracy. Segmentation methods based on *deformable models* (Terzopoulos, 1988; Xu et al., 2000; He et al., 2008) have been shown to cope in a highly robust manner with typical imaging deficiencies, such as noise, artifacts, partial volume effects, low or no contrast due to adjacent anatomical structures with similar appearance, etc. The basic idea is to deform a given (template) shape in such a way that the deformed shape provides an optimal geometric representation of the corresponding structure in the image.

Among many different types of deformable models, meshes are advantageous in many respects, such as flexibility and topology preservation (Montagnat et al., 2001). Typically, the degrees of freedom of the deformable mesh are increased in a multi-level fashion (Okada et al., 2007; Ma et al., 2010; Yin et al., 2010; Zhang et al., 2010; Seim et al., 2008; Kainmueller et al., 2007). At first, only global deformations like rigid transformations or statistical variations (Cootes et al., 1995; Heimann and Meinzer, 2009) are allowed. This robustly produces initial deformed shapes that roughly capture the structure sought-after in the image. On the finer levels, more local assumptions are made on deformations (Okada et al., 2007; Ma et al., 2010), in order to allow for more flexibility and thus capture the specific details of the structure in the given image data. On the finest level, each vertex position of the mesh can move "freely", subject only to regularity constraints that consider its direct neighborhood (Yin et al., 2010; Zhang et al., 2010; Seim et al., 2008; Kainmueller et al., 2007). We refer to such kind of deformations as *free deformations*.

Usually, the deformable mesh *probes* the image information at each vertex position: The image data is evaluated within a certain *search space* to assess suitable image features. Given these probes, a new shape is computed by *displacing* the vertices of the mesh, following a trade-off between image fidelity and anatomically plausible deformation. Note that for free deformations, search space and resulting displacement of an individual vertex are closely related, while this is in general not the case for global deformations, where individual resulting vertex displacements may deviate arbitrarily from the respective search space.

The details of the image probing play a crucial role in the segmentation process. To this end, *unidirectional* (i.e. linear, one-dimensional) search spaces per vertex of the deformable mesh are commonly used (Heimann and Delingette, 2011) due to a

* Corresponding author.
E-mail address: kainmueller@zib.de (D. Kainmueller).

1361-8415/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.media.2012.11.006

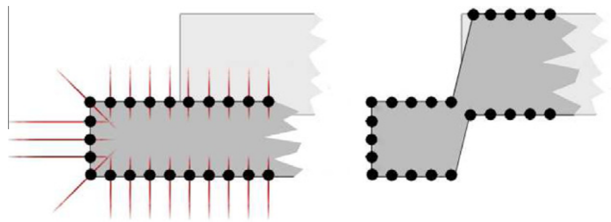

Fig. 1. 2D sketch of an exemplary deformable mesh (dark grey, with vertices as black dots) and target object (light grey). (a) Normal search spaces (directions indicated by lines through vertices) on a tip-like structure detect no target boundary points for a large set of vertices. (b) Resulting unregularized deformation onto target object boundary. Avoiding self-intersection of the mesh suppresses displacement of bottom left-most vertex.

number of benefits: (1) Feature assessment is fast within one-dimensional subsets of the image; (2) It is easy to select the "best" feature, as required by many methods (Cootes et al., 1995; Kainmueller et al., 2007), because a one-dimensional search space is likely to hit the target surface at only one single point (or at most a finite number of points), and hence the set of suitable features is likely to be small; (3) Free deformations can be computed in a globally optimal way for unidirectional search spaces (Li et al., 2006); (4) Normal vertex displacements implicitly restrict the deformation of the surface in a way that reduces (but does not prevent) the risk of generating mesh inconsistencies like self-intersections or fold-overs.

However, unidirectional search spaces suffer from *restricted visibility*: They are prone to miss features in the image data (Fig. 1). In case global deformations are employed, this problem may be alleviated by the fact that individual vertex displacements are not tightly coupled to their respective search spaces. On free deformations, however, the problem has a severe impact: E.g., local translations of highly curved surface regions such as tips or ridges can hardly be achieved (cf. Figs. 1 and 4). This holds true independently of the chosen mesh resolution.

One approach to confront the visibility problem is repeated – i.e. *iterative* – search for image features and respective deformation, where the hope is that visibility will improve in the next iteration. There is, however, no guarantee to this end. Furthermore, iterative deformation of meshes may easily lead to mesh inconsistencies such as self-intersections (Park et al., 2001). This requires additional remedial actions such as adaptive step-size control, adaptive remeshing or mesh surgery (Bucki et al., 2010).

In this paper, we propose a method to overcome the visibility problem for free deformations.[1] The basic idea is to enlarge the search space for image features to allow not only unidirectional but *omnidirectional* displacements at each point of the deformable model. On a deformable mesh, we asses features at – and allow displacements to – a discrete set of points within a ball[2] around each vertex, thus guaranteeing visibility within some radius. Free deformations are modeled by penalizing differences of displacements on edge-connected mesh vertices. This discrete formulation enables us to frame the segmentation problem as a Markov Random Field (MRF), as will be explained in Section 2. MRFs can be optimized efficiently (Komodakis et al., 2008), which has made them attractive for many applications in image processing and computer graphics (see e.g. Glocker et al., 2008; Paulsen et al., 2010). We denote the method of ball-shaped search spaces combined with MRF optimization for surface mesh deformation as *omnidirectional displacements for deformable surfaces*, or *ODDS*.

Allowing a three-dimensional search space per mesh vertex has the drawback of significantly increased run-time and memory requirements as compared to unidirectional search spaces. Therefore, we also propose an extension to ODDS that is faster and less memory-intensive – denoted as *fastODDS*. The key idea for fastODDS, presented in detail in Section 3, is to allow omnidirectional displacements only in regions of high curvature, while restricting displacements to surface normals in "flat" surface regions.

Section 4 provides an extensive evaluation of ODDS and fastODDS on synthetic and clinical data. In Section 5 we will analyze and discuss these results in depth. Here, we will also address the influence of mesh resolution and mesh consistency.

In summary, our results indicate that

1. ODDS can handle free deformations of meshes with high curvature where previous approaches based on normal displacements fail.
2. fastODDS keep all the benefits of ODDS for highly curved surface regions, while being twice as fast and requiring 50% less memory.
3. In contrast to ODDS, fastODDS can also be applied successfully for simultaneous segmentation of multiple objects.

2. ODDS

For a more thorough search for image features in terms of the visibility problem (see Section 1), we propose to extend the search space at each vertex of a deformable surface mesh from a line segment to a ball centered at the respective vertex position. We define the segmentation problem as a trade-off between finding suitable image features within these ball-shaped search spaces and simultaneously considering local regularization.

Volumetric (three-dimensional) ball-shaped search spaces of neighboring vertices overlap heavily in case the ball radius is bigger than the distance between the respective vertices; furthermore, individual search spaces most probably contain a whole region (two-dimensional manifold) of the target surface. Hence highly inconsistent (dissimilar) displacements on neighboring vertices may point to the target surface. The type of local regularization we employ must be able to avoid highly inconsistent displacements of adjacent vertices. We achieve this in a discrete setting (Sections 2.1 and 2.2) via Markov Random Field (MRF) energy minimization (Section 2.3).

We denote the set of vertices v of the deformable surface mesh as $V = \{v_i \in \mathbf{R}^3 | i = 1 \cdots n_V\}$, and the set of pairs of adjacent (i.e. edge-connected) vertices (v, w) as $E \subset V \times V$. Each vertex v can be moved by adding a vector, or *displacement*, $s \in S$, where $S = \{s_i \in \mathbf{R}^3 | i = 1 \cdots n_S\}$ is a discrete set of possible displacements. We refer to a mapping $d : V \to S,\ v \mapsto d(v) =: d_v$ that assigns a displacement to each vertex as *displacement field*. We call a position $v + s$ *sample point*. The set of sample points $v + S$ defines the *search space* for vertex v. Note that this definition has the effect that the search space of a vertex equals its *range of motion*.

2.1. Omnidirectional displacements

We define S as a set of displacements that are uniformly distributed within a ball of radius r_S, i.e. $\forall s \in S : \|s\| < r_S$, where r_S is a parameter of our method. Displacements in S are arranged as a cubic close-packed lattice (Conway et al., 1999); see Fig. 2a for a 2D sketch. We denote the minimum Euclidean distance between unequal displacements $s_i, s_j \in S$ as *sampling distance* $\delta_S := min_{s_i \neq s_j} \|s_i - s_j\|$.

[1] This work extends the authors' paper presented at MICCAI 2010 (Kainmueller et al., 2010), from which some text passages and figures are reused with kind permission from Springer Science + Business Media.

[2] Note that we use the term *ball* to refer to the volumetric (three-dimensional) interior of a sphere, while with the term *sphere* we refer to the surface of a ball, i.e. a two-dimensional manifold embedded in 3d.