Medical Image Analysis 18 (2014) 531–541

Contents lists available at ScienceDirect

Medical Image Analysis

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

Optimal surface segmentation using flow lines to quantify airway abnormalities in chronic obstructive pulmonary disease

Jens Petersen [a,*], Mads Nielsen [a], Pechin Lo [a,b], Lars Haug Nordenmark [c], Jesper Holst Pedersen [d], Mathilde Marie Winkler Wille [f], Asger Dirksen [f], Marleen de Bruijne [a,e]

[a] *Image Group, Department of Computer Science, University of Copenhagen, Denmark*
[b] *Thoracic Imaging Research Group, University of California, Los Angeles, CA, United States*
[c] *Personalized Healthcare & Biomarkers, AstraZeneca R&D, Mölndal, Sweden*
[d] *Department of Thoracic Surgery RT, Rigshospitalet, University of Copenhagen, Denmark*
[e] *Biomedical Imaging Group Rotterdam, Departments of Radiology & Medical Informatics, Erasmus MC, Rotterdam, The Netherlands*
[f] *Department of Respiratory Medicine, Gentofte Hospital, Denmark*

ARTICLE INFO

Article history:
Received 24 May 2013
Received in revised form 4 February 2014
Accepted 7 February 2014
Available online 17 February 2014

Keywords:
Segmentation
Graph
Flow lines
Airways
Computed tomography

ABSTRACT

This paper introduces a graph construction method for multi-dimensional and multi-surface segmentation problems. Such problems can be solved by searching for the optimal separating surfaces given the space of graph columns defined by an initial coarse surface. Conventional straight graph columns are not well suited for surfaces with high curvature, we therefore propose to derive columns from properly generated, non-intersecting flow lines. This guarantees solutions that do not self-intersect.

The method is applied to segment human airway walls in computed tomography images in three-dimensions. Phantom measurements show that the inner and outer radii are estimated with sub-voxel accuracy. Two-dimensional manually annotated cross-sectional images were used to compare the results with those of another recently published graph based method. The proposed approach had an average overlap of 89.3 ± 5.8%, and was on average within 0.096 ± 0.097 mm of the manually annotated surfaces, which is significantly better than what the previously published approach achieved. A medical expert visually evaluated 499 randomly extracted cross-sectional images from 499 scans and preferred the proposed approach in 68.5%, the alternative approach in 11.2%, and in 20.3% no method was favoured. Airway abnormality measurements obtained with the method on 490 scan pairs from a lung cancer screening trial correlate significantly with lung function and are reproducible; repeat scan R^2 of measures of the airway lumen diameter and wall area percentage in the airways from generation 0 (trachea) to 5 range from 0.96 to 0.73.

1. Introduction

Optimal net surface methods (Wu and Chen, 2002) have seen a growing use within medical image segmentation in the last couple of years, likely due to their ability to find the globally optimal solution of multiple interacting surfaces in multiple dimensions given surface cost functions and a useful range of geometric constraints and penalties in polynomial time using minimum cut algorithms (Li et al., 2006; Liu et al., 2012; Petersen et al., 2010; Yin et al., 2009; Abràmoff et al., 2008; Petersen et al., 2011b; Arias et al., 2012; Kainmueller et al., 2013). In order to use these methods, the segmentation problem needs to be transformed from the space defined by the image voxel grid to some graph representation defined by a set of columns. Each column is associated with a point on the sought surface and represents the set of possible solutions, or positions, the surface can take. A suitable graph should be able to represent all plausible solutions in the image space. For instance, if a graph column does not cross the sought surface or if it crosses it multiple times, then this surface cannot be represented by the graph. Similarly, admissible solutions in the space defined by the graph representation should represent valid surfaces in image space, that is, the graph space should for instance not allow self-intersecting surfaces. It is also important that the graph structure allows for a meaningful representation of the surface cost functions and geometric constraints and penalties. Surface non-smoothness can, for instance be reduced, by increasing the cost of solutions in proportion to how much they vary in neighbouring

* Corresponding author. Tel.: +45 60687733.
E-mail addresses: phup@diku.dk (J. Petersen), marleen@diku.dk (M. de Bruijne).

http://dx.doi.org/10.1016/j.media.2014.02.004
1361-8415/

columns. However this is only meaningful if the relative variation within the columns is somehow related to the associated relative variation within image space.

In some cases the sought surfaces are expected to be oriented along an image axis and the voxel columns of the image itself may be used. This has for instance been used in the case of the intraretinal layers in macular optical coherence tomography images (Abràmoff et al., 2008). Other approaches have used simple mathematical transformations, such as those of Li et al. (2006) and Petersen et al. (2010), in which images of tubular airway segments were unfolded using polar transforms in two or three-dimensions. The graph columns were oriented perpendicular to the resulting contours or terrain like surfaces allowing for an easy representation of surface smoothness constraints and penalties. In many cases, however, the surfaces are much too complicated for such an approach and/or the prior knowledge of the surfaces' shape and position required is not available. In these cases such prior knowledge may be gained by employing an initial method to roughly estimate the position of the surfaces and then use an optimal surface graph to refine this estimate. This was done in Liu et al. (2012) by placing columns at points of the initial surface and oriented along the surface normals inward and outward. Problems with intersecting columns and thus self-intersecting surfaces were avoided, by limiting the length of each column to the minimum distance to the initial surface inner and outer medial axes. This approach can result in columns that are too short to reach the desired solution, as shown in Fig. 1(a). Yin et al. (2009) suggested columns inspired by the non-intersecting property of electric lines of force. The columns were constructed by simulating electrical charges at surface points of the initial segmentation and tracing the electric lines of force within the field inward and outward. This method is computationally infeasible for large scale problems, as every surface point charge influences the computation of every electric line of force. Furthermore the electric lines of force can behave erratically if the initial segmentation contains small scale errors or noise. Recently Kainmueller et al. (2013) proposed to use omnidirectional displacements, which allow each initial surface mesh vertex to move to uniformly distributed positions within a ball shaped region around it. Self-intersections are minimised by using regularisation and the solution is found using Markov Random Field energy minimisation. The approach is too computationally expensive for larger problems and so Kainmueller et al. (2013) also shows how it can be combined with an optimal surface – unidirectional column type approach. This makes the method practical for larger problems by using omnidirectional displacements in high curvature regions and unidirectional columns in low curvature regions. The two problems are solved sequentially and so the method does not guarantee global optimality.

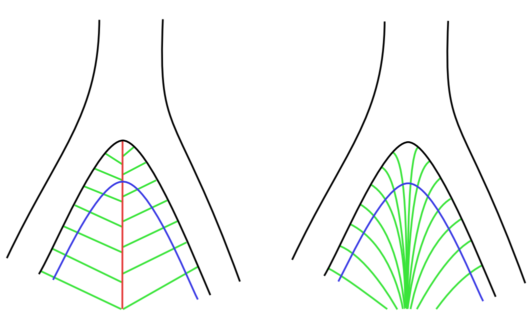

(a) Medial Axes (b) Flow Lines

Fig. 1. (a) illustrates the fish-bone like structure of surface normal direction columns (*green*) based on the distance to the medial axis (*red*) in areas where the initial segmentation (*black*) has high curvature. Notice that the four inner-most columns do not cross the sought surface border (*blue*), which means that the desired solution cannot be represented by the graph and the segmented surface will be wrong in these positions. (b) shows the advantage of columns based on flow lines (*green*), notice that all columns cross the sought surface. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In Petersen et al. (2011b) we proposed to use graph columns defined from flow lines within a regularized version of the initial segmentation. Flow lines are non-intersecting and are uniquely defined if the regularisation is smooth, and noise and small errors in the segmentation are naturally dealt with by the same regularisation. Moreover, fast approximations can be computed using image convolution. Fig. 1(b) illustrates the concept. The method was originally applied to the problem of segmenting human airway walls in CT images and has since then been used for segmenting the carotid artery bifurcation in magnetic resonance imaging (Arias et al., 2012).

Assessing the dimensions of the airway walls is important in the study of airway remodelling diseases such as Chronic Obstructive Pulmonary Disease (COPD) (Hackx et al., 2012). It is a dual surface problem, consisting of an inner and an outer wall surface, where bifurcations form regions of high curvature that would cause problems for conventional graph construction approaches. The vast majority of previous airway wall segmentation methods have been one- or two-dimensional in nature. The one-dimensional techniques work by casting rays from the centre of the airways outwards looking for the wall surfaces using the full width at half maximum edge detection principle (Nakano et al., 2000), by phase congruency (Estépar et al., 2006), or more complex models of the scanning point spread function (Weinheimer et al., 2008). The airway wall surfaces resemble concentric circles when seen in a cross-sectional view centred on and perpendicular to the airway centreline. This is what two-dimensional methods typically exploit to impose some degree of regularity on the solution (Petersen et al., 2010; Saragaglia et al., 2006). Three-dimensional methods, however, may use more of the information present in the image, allowing surfaces to be found more accurately when they are close to other structures such as blood vessels. Moreover, bifurcation and carina regions, which typically cannot be segmented with previous two-dimensional approaches, can be analysed (Liu et al., 2012). Besides the already mentioned methods of Liu et al. (2012), Petersen et al. (2011b), a three-dimensional method is also described in Saragaglia et al. (2006), which evolves a deformable mesh, constructed from an initial segmentation of the lumen. The evolution is done with force constraints computed from intensity and gradient magnitude values; elastic forces penalising local wall thickness variations; and regularisation forces, locally smoothing the result. The method does not guarantee a global optimal solution and unlike the approaches of Liu et al. (2012), Petersen et al. (2011b) the two surfaces are not estimated simultaneously, and thus the added knowledge of the position of the exterior surface is not used to improve the inner surface. Ortner et al. (2010) also proposed to use a deformable mesh. Their mesh is built from an initial segmentation of the lumen and its evolution is governed by gradient vector flow and simplified Lagrangian dynamics and so avoids self-intersections. The approach was evaluated on simulated CT data and 15 clinical cases of mild and severe asthmatics, showing good agreement with segmentation result and clinical expertise.

This paper is an extension of the work presented in Petersen et al. (2011b). The main differences are the addition of a constraint, that forces the outer surface to be outside the inner; improvements in the parameter tuning, such that all involved parameters are automatically estimated using a manually annotated training set; adjustment of parameters and evaluation of results according to the COPDGene phantom (Sieren et al., 2012) to account for a possible bias present in the manual annotations; and finally the addition of an extensive medical expert visual evaluation comparing