Keuhkot: A Method for Lung Segmentation

Rômulo Pinho[2,1], Vivien Delmon[1], Jef Vandemeulebroucke[1], Simon Rit[1], and David Sarrut[1]

CREATIS, CNRS UMR 5220, Inserm U1044, INSA-Lyon, University of Lyon, France
Léon Bérard Cancer Centre, University of Lyon, France
vivien.delmon@creatis.insa-lyon.fr; pinho@lyon.fnclcc.fr
jefvdmb@gmail.com; {simon.rit; david.sarrut}@creatis.insa-lyon.fr

Abstract. Lung segmentation is an important operation in the analysis of medical images. It can be very challenging in circumstances where image artefacts and abnormalities deteriorate the lung boundaries, therefore hindering the lung delineation process. In this work, we propose a combination of algorithms to achieve good segmentation results. These algorithms essentially comprise region growing, morphological operations, and graph-cut segmentation, which together become a powerful tool to automate the segmentation. The proposed method was evaluated on a database of 55 patients, with varying image quality and drastic anatomical changes.

1 Introduction

Lung segmentation in 3D CT images is a very important process in pneumology. Delineation of the organs directly determines, for instance, their volumes, which has a strong correlation with cronic obstructive pulmonary diseases (COPD) [1,2]. The correct segmentation is also invaluable in registration techniques for respiratory motion modelling, in which the segmented lungs define masks that constrain the region whose motion is to be determined [3].

In normal patients, the lungs tend to appear as a region of high contrast with respect to surrounding structures. This usually facilitates the segmentation and delineation processes. Often, however, different types of artefacts may bring extra challenges to the segmentation. Automatic or even semi-automatic algorithms must then be robust enough to cope with low signal to noise ratios, partial volume effects, anatomical variations, morphological abnormalities, and the like.

In this paper, we propose to solve the lung segmentation problem with a set of operations involving different algorithms. The proposed method executes these operations repeatedly with different parameters and analyze the results in other to chose the best possible segmentation. Experiments were carried out in the context of the segmentation challenge *LOLA11: LObe and Lung Analysis 2011*, in which several state of the art algorithms were executed on a database of 55 patients, so as to verify which of them gives the best results. In the following, we describe all the steps and algorithms employed in the present work.

2 Method

The proposed procedure for lung segmentation consists in executing a sequence of steps comprising 6 types of operations: patient extraction, localization of the trachea, adaptive region growing for airway tree segmentation, graph-cut based segmentation, region growing for lung segmentation, and morphological operations. Of these, the latter two are carried out in their classical implementations. Namely, in region growing, a connected region is formed out of a set of seed points. Pixel (or voxel) neighbours that pass a similarity test are recursively aggregated to the region, with neighbourhood and similarity being application-dependent. In this work, 26-connected neighbourhoods are employed, while similarity means that two neighbour voxels are within the same range of pre-defined grey values. Morphological operations, in turn, encompass combinations of dilations and erosions with a spherical structuring element of varying radius. For more details on region growing and morphological operations, we refer the reader to [4]. The other three types of operations will be discussed in the following subsections.

The reasoning behind the proposed procedure is to obtain two segmentations that later will be subtracted from each other. Let $\mathcal{S}_1$ be the set of voxels representing lungs and airways altogether (Fig. 1(a)). Let $\mathcal{S}_2$ be the set of voxels representing the airways only (Fig. 1(b)). The result is then $\mathcal{S} = \mathcal{S}_1 - \mathcal{S}_2$ (Fig. 1(c)). In general, $\mathcal{S}$ will contain remnants of $\mathcal{S}_2$, which can easily be removed with a morphological *Open* operation. In theory, all that would be missing is the labelling of the two lungs in the segmented image and, *voila!*, the segmentation would be done. However, image artefacts and anatomical abnormalities may hinder the procedure described, calling for a more robust approach, such as graph-cut segmentation. To cope with the difficulties, the operations above are put together as shown in **Algorithm 1** below.

(a) lungs and airways (b) airways (c) lungs mask (d) graph-cut

Fig. 1. Illustration of intermediate results

Algorithm 1 segment(image, L_p, U_p, L_a, U_a, nlevels, L_l, U_l)

```
flip(image) {if not in superior-inferior direction}
patient ← extract_patient(image, L_p, U_p) {segment the patient}
seed ← localize_trachea(patient)
airways ← segment_airways(image, seed, L_a, U_a, nlevels) {compute S_2}
lungs ← segment_lungs(image, seed, L_l, U_l) {compute S_1}
diff ← difference(lungs, airways) {compute S = S_1 − S_2}
opened ← open(diff, radius = 2) {clean-up}
radius ← 5
closed ← close(opened, radius) {collapse holes}
while largest_component_size(closed)/second_component_size(closed) < 0.99 do
    radius ← radius + 3
    closed ← close(opened, radius)
end while
closed ← keep_largest_component(closed)
radius ← 3
eroded ← erode(closed, radius)
while largest_component_size(eroded)/second_component_size(eroded) > 0.99 do
    radius ← radius + 3
    eroded ← erode(closed, radius)
end while
seed_right, seed_left ← lunglabel(opened)
right_lung, left_lung ← graphcut(image, closed, seed_right, seed_left)
finalmask ← open(diff, radius = 2) {clean-up}
finalmask ← close(finalmask, radius = 6) {collapse holes}
right_lung ← mask(right_lung, finalmask)
left_lung ← mask(left_lung, finalmask)
```

2.1 Patient Extraction

The patient extraction step of **Algorithm 1** is a pre-processing operation intended to clean the image and facilitate the trachea finding stage that comes right afterwards. It starts with threshold operations that segment all areas outside the patient, which generally comprise air. The second stage finds the two largest connected components on the thresholded image and remove all areas outside the patient, assuming that they correspond to the largest region. The complement of what was removed is equivalent to the patient's soft tissues, which are aggregated to the lungs and airways, giving the patient mask. Finally all areas of the original image outside the computed patient mask are set to a fixed background value.

2.2 Trachea Localization

As mentioned earlier, region growing algorithms depend on a set of seed points from which the aggregation of voxels starts. In the case of lung and airway tree segmentation, a typical choice for seed point is a voxel somewhere inside the trachea. Many approaches rely on the manual setting of such point, since

a human observer can easily identify the trachea on a CT image. Automating this task, however, may be challenging, since the trachea might be confused with other structures in the image, such as the oesophagus, parts of the upper airways, areas of the lungs, image artefacts, etc.

To solve the problem of automatically choosing a seed point inside the trachea, we start by finding in each axial slice of the input image a candidate region for the trachea. This region is obtained after a pre-propcessing step that thresholds air, removes noise by *Opening* morphological operations, and removes regions that are too small, too big, or too thin. We assume that the trachea is a relatively elliptical region centrally located with respect to the sagittal plane. The candidate is then found by solving Eq. (1)

$$\begin{aligned} f(e_j, d_j) &= e_j + (1 - d_j) \\ \mathrm{R}_{e_i} &= \arg\,\max_{\mathrm{R}_j} f(.), \end{aligned} \tag{1}$$

where $j = 1 \ldots n_{r_i}$ is the index of a region in slice i, with n_{r_i} the number of labelled regions in the slice after the pre-processing step, $e_j < e_{max}$ is the eccentricity of region R_j, and d_j is the normalized distance between the barycentre of R_j and the sagittal plane. To guarantee that the patient is relatively centred in the RL and AP directions (assuming that scans are never truncated), the image is first cropped at the bounding box of the patient mask (obtained in the patient extraction step above).

With the candidate for the trachea determined in each slice of the volume, the challenge is now to choose the slice that contains the best candidate. This is done by grouping candidate regions that do not lie far appart from each other in two subsequent slices. The result is thus a collection of sequences of regions, and the first region of the longest sequence is taken as the candidate for the trachea. The seed point for the segmentation is finally computed as the centre of gravity of the chosen candidate region.

2.3 Airway Tree Segmentation

The region growing algorithm employed in the airway tree segmentation is that proposed in [5]. The main concept behind this algorithm is to bound the segmentation with cylindrical regions of interest (ROIs), a technique introduced in [6]. Starting from a seed point inside the trachea, the algorithm iteratively carries out the segmentation inside one ROI of pre-computed height, orientation, and radius. The intersection between the segmentation and the walls of the ROI determines the location, radius, and orientation of the ROI of the next iteration, with bifurcations being handled through the computation of an approximate skeleton of the segmentation inside the ROI. The cylindrical ROIs facilitate the detection of leaks and limit their occurence. Leaks are a common problem in region growing based segmentations, especially in the case of the airways, where partial volume effects may cause the voxels in the lungs to be aggregated to the segmented region.

A mechanism to detect leaks based on anatomical information is also used in this algorithm. The general idea is to check at each iteration whether the diameter of the segmented region is decreasing, which is expected as we go down the airway tree, and if bifurcations are limited to a small number, which also depends on the tree level. Whenever leaks are detected, the segmentation inside the ROI is repeated with a neighbour affinity technique. In this technique, voxels are only aggregated to the segmentation if all its neighbours inside a connectivity mask also pass the similarity test. Each time the segmentation is repeated, a mask of larger radius is used, until no leaks are detected. This technique is based on the assumption that leaks occur through small holes in the structure being segmented. The algorithm thus tries to find a mask size that stops the voxel aggregation from passing through these small holes.

2.4 Graph-cut Based Segmentation

At this step, the lung mask constructed by removing the airway mask from the segmented lungs can be a single connected component Fig. 1(c), for instance when lungs are closed in some regions. To find where to cut our mask, we used a graph-cut algorithm that separates two seed regions on the optimal border over a certain neighbourhood constraint Fig. 1(d).

The graph-cut implementation used in this work is that proposed in [7], which finds on a weighted graph which edges to cut in order to separate the source seeds and the sink seeds. The input graph is composed of all voxels of the connected lung label, which is constructed by repeatedly closing the lung mask until having one single large component (which, in practice, represents more than 99% of the size of the union of all components in the image). Each of these voxels becomes a node of the graph and is connected to its 26 neighbours by edges labelled with the inverse of the gradient of the initial image. Source and sink nodes are obtained by eroding the closed mask until having two distinct regions, which are labelled as left and right seeds.

The graph-cut result cuts a closed mask which can be larger than the lungs in some places. The exceeding voxels are removed by masking the graph-cut result with the lung mask opened with a ball of radius 2 voxels and then closing the result again with a ball of radius 6 voxels. Both lung masks are summed to obtain the final result.

3 Experiments

The proposed method was evaluated with a dataset of 55 patients, provided as part of the workshop and lung segmentation challenge *LOLA11: LObe and Lung Analysis 2011*. The segmentation was evaluated by a team of trained observers. The aim of the workshop was to compare the performance of different state of art algorithms. For this, a ground truth was constructed from all submitted segmentations and all algorithms were evaluated with respect to this ground truth.

The algorithms presented here were implemented in C++, using the ITK framework, and the programs executions were carried out on the cluster of the Institut national de physique nuclaire et de physique des particules (IN2P3) of the Centre National de la Recherche Scientifique (CNRS), in Lyon, France. The correct parameters for the proposed method were chosen empirically after several runs over a training image database. In the end, the threshold values for the patient extraction step were L_p = -1024HU and U_p = -300HU. In the trachea localization, $N = 50$ slices were used in the search, $e_{max} = 0.5$. For the airway tree segmentation, L_a = -1024HU, U_a = -600HU, and nlevels = 4. In the lung segmentation, L_l = -1024HU, U_l = -300HU. We used the same parameters for all patients to provide a fully automatic algorithm.

3.1 Results

The results obtained with the proposed method [1] are shown in Table 1, with the execution times given in Table 2. The main difficulty in the segmentation of image database was the existence of images with various artefacts, either due to anatomical variations or acquisition issues. Situations in which the lungs appeared together after the segmentation were particularly challenging, especially when combined with patient abonormalities such as atelectasis. Higher upper thresholds for the lung segmentation can overcome the granularity observed in the presence of atelectasis, but it can also increase the frequency in which lungs will appear together in the segmentation result. Eroding the result to the point where the lungs would be separated could remove shape information to a point that the dilations could not recover the original lung volume. Finding a balance between threshold choice and number of erosions/dilations was thus arduous.

The use of the graph-cut algorithm was very effective in dealing with the "connected lungs versus atelectasis" problem, since it could find the cutting region between the two lungs with good precision. The masking and the sequence of morphological operations carried out after the graph-cut could then easily take care of cleaning the segmentation result, by removing those parts that were considered outside the lungs.

In any case, our application assumes that the lungs are filled with air, on which the threshold-based operations depend. In severe atelectasis, pleural effusion may take large parts of the lungs, thus changing the density of regions that would otherwise contain air. Although in these situations the lung contours may still be visible on CT, our threshold operations will not be able to completely recover them, since only those regions still containing air will be segmented.

As future improvements, we would like first and foremost to increase execution speed. Our application is still in prototypical stages and certain parts can immediately be optimized. For example, forcing the lungs to be glued together before the graph-cut is done with repeated closing operations with an increasing ball. We observed this was often the most time-critical operation among all steps of the algorithm, since closing with large structural elements can take long to

[1] Provided by the organizers of the workshop.

Table 1. Results of lung segmentation for the 55 scans in LOLA11.

	mean	SD	min	Q1	median	Q3	max
left lung	0.935	0.209	0	0.978	0.99	0.995	0.998
right lung	0.961	0.147	0	0.986	0.993	0.996	0.998
score	0.948						

Table 2. Execution times of lung segmentation for the 55 scans in LOLA11.

	mean	SD	min	Q1	median	Q3	max
seconds	2425.51	2574.69	464	839	1333	3270.5	12804
minutes	40.42	42.91	7.73	13.98	22.21	54.51	213.4

run. In addition, finding the correct ball radius is a trial-and-error process. We believe the optimal ball size can be found at one go with a chamfer distance map, computed in $O(2n)$ time in the number of pixels of the image, requiring then only one closing operation.

Concerning the graph-cut implementation, optimizations in terms of memory use are necessary, since the implementation used in this work was more on the naive side with respect to the use of data structures. Implicit graph representations, such as adjacency lists or the image structure itself, will probably work better.

We would also like to improve the trachea localization algorithm to find a seed point as up in the trachea as possible. Currently, due to the existence of misleading regions, the sequences of regions discussed in Section 2.2 might be broken, eventually causing the seed point to be chosen at a lower part of trachea.

Fig. 2 shows some of the results obtained as visualized with the VV platform (http://vv.creatis.insa-lyon.fr), and the source code of all applications will be available at the same address.

4 Conclusions

This paper presented a method for lung segmentation in 3D CT images. The method is based on the combination of different algorithms consisting of region growing, morphological operations, ang graph-cut segmentation. This combination has proved to be powerful enough to cope with even the most difficult cases, as was observed after the evaluation of the method on an image database with 55 patients.

References

1. Hedlund, L.W., Anderson, R.F., Goulding, P.L., Beck, J.W., Effmann, E.L., Putman, C.E.: Two methods for isolating the lung area of a ct scan for density information. Radiology **144** (1982) 353357

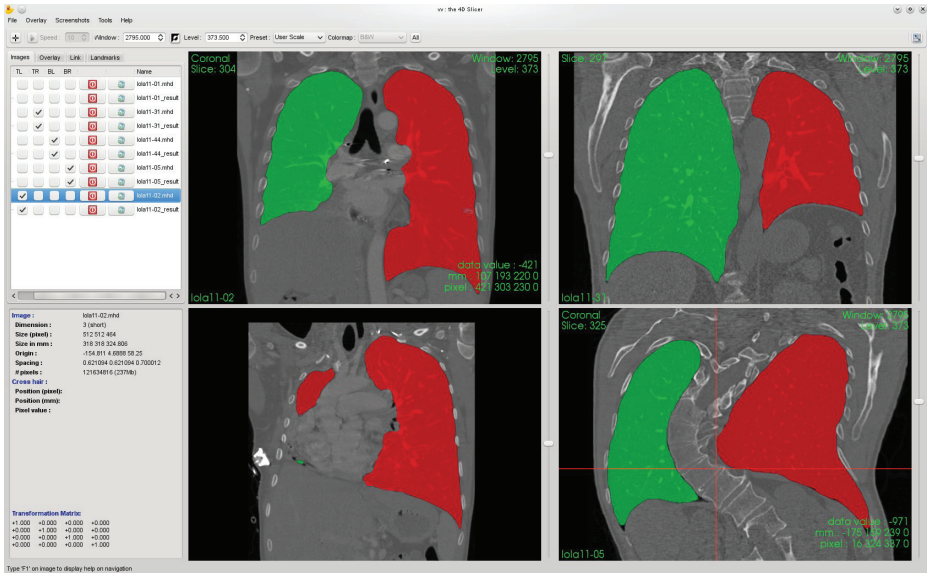

Fig. 2. Segementation results.

2. Uppaluri, R., Mitsa, T., Sonka, M., Hoffman, E.A., McLennan, G.: Quantification of pulmonary emphysema from lung ct images using texture analysis. Amer. J. Resp. Crit. Care Med. **156** (1997) 248254
3. Vandemeulebroucke, J., Rit, S., Kybic, J., Clarysse, P., Sarrut, D.: Spatio-temporal motion estimation for respiratory-correlated imaging of the lungs. Med Phys **38** (2011)
4. Gonzalez, R.C., Woods, R.E.: Digital Image Processing. Addison-Wesley Longman Publishing Co., Inc. (2001)
5. Pinho, R., Luyckx, S., Sijbers, J.: Robust region growing based intrathoracic airway tree segmentation. In: 2nd International Workshop on Pulmonary Image Analysis, London, England (2009) 261–271
6. Tschirren, J., Hoffman, E., McLennan, G., Sonka, M.: Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose CT scans. Medical Imaging, IEEE Transactions on **24** (2005) 1529– 1539
7. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. Pattern Analysis and Machine Intelligence, IEEE Transactions on **26** (2004) 1124–1137