Segmentation of Juxtapleural Lung Nodules in CT Scans Based on Ellipsoid Approximation

Jan Hendrik Moltz, Jan-Martin Kuhnigk, Lars Bornemann and Heinz-Otto Peitgen

MeVis Research GmbH – Center for Medical Image Computing, Bremen, Germany
jan.moltz@mevis.de

Abstract. This article presents a new algorithm for segmenting juxtapleural lung nodules in CT scans. Segmentation is an essential part of volumetric therapy monitoring for cancer patients. Pulmonary nodules that have extensive contact to the chest wall or other structures of similar density are a special challenge for automatic segmentation. We propose a ray casting approach to identify points at the visible boundary of the nodule and then approximate its shape by an ellipsoid that is a least squares fit of these points. The adjacent structures are cut off by morphological processing within a dilated version of this ellipsoid. Evaluation on 333 juxtapleural nodules showed that this method yields good results and can be integrated easily into a general segmentation algorithm for lung nodules with no substantial increase in computation time.

1 Introduction

In oncological therapy monitoring, the estimation of tumor growth from consecutive CT scans is an important aspect in deciding whether the given treatment is adequate for the patient. Traditionally, this is done by measuring and comparing the largest axial diameter of each lesion manually, but this approach implies several problems. First, manual examinations are always subjective, error-prone and time-consuming. Second and even more importantly, a 3d quantity (volume) is estimated based on a 1d measurement (diameter). This simplification would be valid if tumors were perfectly spherical and grew symmetrically but in practice it leads to inaccurate results.

Although volumetry has the potential to enhance the accuracy and reproducibility of growth estimation, measuring the lesion volume manually would take too much time in the workflow of a radiologist. This is the motivation for employing software assistants in oncological therapy monitoring since they are able to perform automatic volume measurements in the 3d data. In order to be accepted in clinical routine, they have to work both fast and accurately. Lesion segmentation is an essential prerequisite for volumetry and efficient algorithms are needed for different kinds of tumors.

In this article, we focus on the segmentation of juxtapleural lung nodules. Pulmonary nodules are mostly located centrally within the lung parenchyma, but they can also be attached to the pleura, a thin membrane that covers the

lungs. In CT scans of the thorax, the voxels can basically be divided into two density classes: while the dark ones represent the lung parenchyma, bright voxels may be nodules, but also blood vessels or structures adjacent to the lungs such as the chest wall, the heart, or the diaphragm (Fig. 1(a)). Since the pleura itself is invisible in CT images the boundary between a juxtapleural nodule and any of these structures shows little or no contrast and it is sometimes impossible even for a radiologist to determine the exact boundary of a nodule. Therefore the segmentation of this kind of lesions is particularly challenging; a mere threshold-based method or simple morphological processing is not sufficient.

In the following, all computations are restricted to a region of interest (ROI) whose center and size are determined by a user-defined stroke that is drawn across the nodule. We assume that the ROI contains the lesion completely and its center is close to the center of the lesion. Furthermore, we supersample the ROI to isotropic voxels if necessary.

2 State of the Art

Several authors that worked on solid lung nodule segmentation have also proposed solutions for juxtapleural nodules. An obvious idea is to compute a lung segmentation in order to separate the nodule from structures outside the lungs as it is done by Fetita *et al.* [1]. However, we decided not to incorporate any global information in order to keep computation times as low as possible and to be able to integrate the algorithm into an existing workstation. Van Ginneken [2] uses a local 2d lung field segmentation but no evaluation for juxtapleural nodules is given.

One of the first dedicated segmentation algorithms for juxtapleural nodules was presented by Shen *et al.* [3]. Assuming that the chest wall is physiologically smooth and that a nodule creates a "bump" with a high local curvature, the nodule can be separated by smoothing the wall surface. This is implemented by projecting the surface to a plane whose normal is the mean of all surface voxel normals. On the projection image, the nodule appears as a region of high values which are replaced by a cubic polynomial interpolation of the other values. The smoothed 3d surface is then computed via backprojection. This idea is promising but there will be problems when the wall itself has points of high curvature as in Fig. 2(c).

An algorithm proposed by Wiemker *et al.* [4] starts with region growing and determines the optimal cut-off value retrospectively by means of an objective function that separates the wall at its strongest inflection. Unfortunately, no results for juxtapleural nodules are shown.

Reeves *et al.* [5] approximate the pleural surface by a clipping plane that is iteratively refined until a leap in volume change is observed when it actually reaches the pleura. This works for small nodules where the actual convex or concave shape of the lungs can be ignored in a local view, but in other cases a plane is not suitable and the algorithm will fail.

Okada *et al.* [6] presented an approach that uses morphological opening similar to ours, but applied it in an inverse way: the size of the structure element is chosen such that the nodule is removed and the adjacent structure is retained. This presumes, however, that in the ROI the wall region is significantly larger than the nodule. Therefore the approach will fail when the nodule is large or when the wall has a concave shape. Furthermore, it is implemented such that segmenting a juxtapleural nodule takes more than twice as long as for a central one.

In a recent publication, Dehmeshki *et al.* [7] describe an algorithm that uses sphericity-constrained region growing on a fuzzy connectedness map. Successful results are reported for nodules "very close to lung wall or diaphragm" but the method always needs a visible contrast between the nodule and the adjacent structure.

The methods described in this article are based on the algorithm by Kuhnigk *et al.* [8] which has been designed to segment small and large nodules even with extensive vascularization. This is done by an initial region growing, followed by morphological opening with an adaptive erosion strength to remove blood vessels. As a preliminary solution for juxtapleural nodules, the algorithm makes use of the fact that the lungs are convex in most parts and that juxtapleural nodules create a concavity in this shape. The idea is to reconstruct the lung shape of the tumor-free state by computing the convex hull of the lung parenchyma within the ROI and cut off the nodule along the boundary of the convex hull. However, in regions where the lungs are not convex, such as the boundaries to the heart or the diaphragm, the convex hull is not suitable for this purpose since it does not remove convex parts of the attached structure completely. Furthermore, the results depend on the size of the ROI. This leads to poor segmentation results as shown in Fig. 2(d,e,f). In our contribution, we present an improvement to the algorithm that can handle this case as well. Parts of this work have previously been published in German [9].

3 Segmentation Method

The goal of our extension of Kuhnigk's algorithm [8] was to improve the segmentation of nodules located at concave parts of the pleura while changing the original method as little as possible in order to get consistent results. The convex hull operation is obviously not suitable for reconstructing the shape of a concave object but we observed that the error decreases when the ROI is made smaller since the convex hull is basically determined by the most distant points of the lung boundary that are contained in the ROI. Therefore our approach is to make the ROI as small as possible so that the disturbing effect of the concavity is minimized. As a minimal ROI, we choose a dilated ellipsoid that is computed as an approximation of the nodule shape. Ellipsoid approximation of lung nodules has been used with different goals and methods in the literature [6, 10].

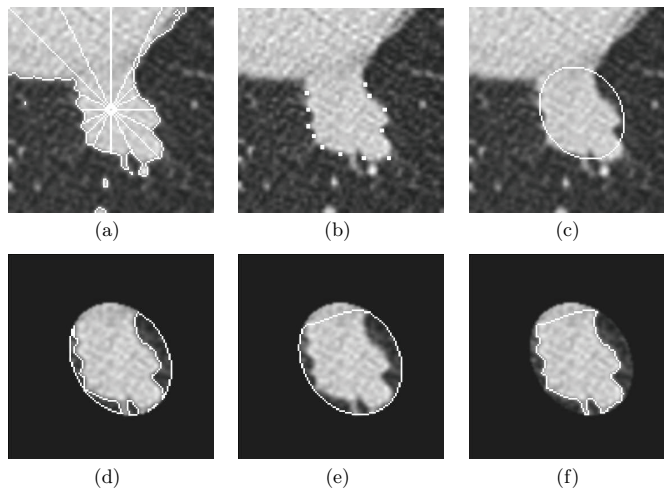

Fig. 1. Step-by-step illustration of the segmentation algorithm for juxtapleural lung nodules, exemplified by the central axial slice. (a) Result mask of the initial region growing and rays cast towards the boundary of the mask. (b) Valid ray endpoints on the nodule boundary. (c) Ellipsoid fitted to the boundary points. (d) Part of the mask within the dilated ellipsoid. (e) Convex hull of the inverse mask within the dilated ellipsoid. (f) Difference of convex hull and inverse mask.

Our method is a preprocessing step for Kuhnigk's algorithm [8] and consists of three parts which are described in the following sections and illustrated in Fig. 1:

1. identification of points on the nodule boundary by region growing and subsequent ray casting from the seed point;
2. calculation of an ellipsoid that approximates the shape of the nodule;
3. convex hull operation as in [8], but restricted to the dilated ellipsoid.

3.1 Region Growing and Ray Casting

Initially, region growing is performed, using the ROI center as a seed point. Since we only need to separate the nodule and attached high-density structures from the lung parenchyma in this first step, we can use -400 HU as a fixed threshold. In order to find points on the boundary between the nodule and the parenchyma, we apply a ray casting approach. Starting from the seed point, rays are sent out through all surface voxels of a $5 \times 5 \times 5$ cube around the seed point. This ensures a symmetric distribution of rays and an alignment to the voxel grid. These $5^3 - 3^3 = 98$ rays are traced until they reach either the boundary

of the region growing mask or leave the ROI (Fig. 1(a)). In the former case, the endpoints are stored, otherwise discarded (Fig. 1(b)). Since some false boundary points may be found due to noise or other structures in the outer parts of the ROI, ray endpoints above a certain distance from the seed point should also be discarded. We found the 95% quantile of the distances of all points to provide a good threshold.

3.2 Ellipsoid Approximation

Typically, the points found by the ray casting procedure cover a major part of the actual nodule surface. Assuming that the nodule has approximately an ellipsoid shape, we aim at reconstructing this shape by fitting an ellipsoid to the points (Fig. 1(c)).

A 3d ellipsoid is defined as a conic section

$$\{\mathbf{x} \in \mathbb{R}^3 \mid \mathbf{x}^T\mathbf{A}\mathbf{x} + \mathbf{b}^T\mathbf{x} + c = 0\}, \tag{1}$$

where the symmetric matrix $\mathbf{A} \in \mathbb{R}^{3\times3}$ is positive or negative definite. Due to its symmetry, $\mathbf{A}$ has only six degrees of freedom, plus a total of four for $\mathbf{b} \in \mathbb{R}^3$ and $c \in \mathbb{R}$. From the valid endpoints of the 98 rays we want to determine those ellipsoid parameters which are optimal in a least squares sense. This establishes a non-linear equation system which can be reduced to a generalized eigenvalue problem and solved efficiently with a method proposed by Grammalidis and Strintzis [11]. It does not guarantee $\mathbf{A}$ to be definite, but our experiments showed that this is almost always the case. If the points are distributed in a way such that it is not possible to fit an ellipsoid to them – if the nodule, for example, has an irregular shape or very extensive contact to other structures – a sphere can be computed instead with the radius as the only free parameter. Although this is a coarser approximation, it can still yield acceptable results in most of these rare cases.

It should be noted that the center of the ellipsoid is included in the optimization. The user-defined seed point influences only the distribution of the boundary points. Since the equation system is highly overdetermined it is robust against variations caused by different user interactions.

3.3 Convex Hull

For the following computations we use a slightly dilated version of the ellipsoid as a new minimal ROI (Fig. 1(d)). The dilation is necessary to ensure that the nodule is covered completely. At its margin, the ellipsoid contains some parenchyma voxels as well, so the convex hull operation can now be applied to reconstruct the original lung shape within this ROI (Fig. 1(e)). This is sufficient for determining the boundary of the nodule and it works in concave parts of the lungs as well because these concavities are no longer visible inside the ellipsoid. For performance reasons, the convex hull has been implemented as the union of slice-wise convex hulls in axial, sagittal, and coronal views.

Subsequently, the original algorithm [8] is executed and the adaptive opening procedure removes attached vessels. The final result is shown in (Fig. 1(f)).

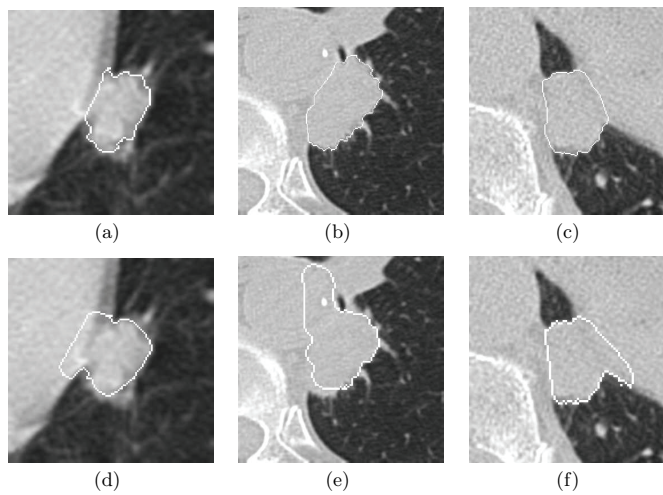

Fig. 2. Upper row: Examples of successful segmentations with the extended algorithm. Lower row: Corresponding results of the original version.

4 Results and Discussion

For our evaluation, we used a database of 333 ROIs of juxtapleural lung nodules from various patients, clinics and CT scanners with seed points set manually by radiologists. Since extensive studies have been conducted for the original version [8], we focussed on the effects of the extensions presented above. It is often impossible to determine the exact boundary between a nodule and a structure it is attached to. In cases like these where no reliable ground truth is available it is most important to produce consistent results so that volumes can be compared over time. Therefore we evaluated the segmentation results visually and examined if they were consistent with our approach to reconstruct the tumor-free shape of the lung parenchyma.

While in 71% of the cases the result of the original algorithm was classified as good, our extension could increase this proportion to 89%. For an additional 5%, a good result was obtained after applying the interactive correction procedure of [8]. Fig. 2 shows some examples of successful segmentations and reveals a significant improvement over the previous results. Most of the nodules that could not be segmented had complex shapes or a very extensive connection to high-density structures that made it impossible to fit ellipsoids to them. Examples can be seen in Fig. 3.

The experiments show that we have developed an algorithm which is able to segment most juxtapleural lung nodules, even if the surrounding parenchyma

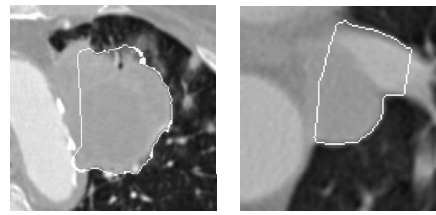

(a) (b)

Fig. 3. Examples of difficult cases where the new algorithm did not succeed. (a) The ellipsoid does not cover the entire lesion since its 3d shape is hard to estimate. (b) The extensive vasculature is not cut off completely.

is not convex. Thus we closed a gap in Kuhnigk's algorithm [8] that had systematic problems in these cases. Our proposed method can easily be integrated there and has no effect on central nodules. No significant increase in runtime could be found. Although some additional operations are performed, subsequent computations can be restricted to the dilated ellipsoid. The computation time for the complete segmentation is around 2 s for large nodules and often below 1 s for small ones on a PC with a 2 GHz DualCore processor.

All of the nodules in this study had direct contact to the pleura and constituted a sample of very difficult cases. Considering that [8] already reported a success rate of 91% on a representative collection of central and juxtapleural lung nodules, we can conclude that with our extension the algorithm allows a successful segmentation in almost all cases.

Acknowledgement. This work was supported by a research grant from Siemens Healthcare, Computed Tomography, Forchheim, Germany.

References

1. Fetita, C.I., Prêteux, F., Beigelman-Aubry, C., Grenier, P.: 3d automated lung nodule segmentation in HRCT. In: Proc. MICCAI. (2003) 626–634
2. van Ginneken, B.: Supervised probabilistic segmentation of pulmonary nodules in CT scans. In: Proc. MICCAI. (2006) 912–919
3. Shen, H., Goebel, B., Odry, B.: A new algorithm for local surface smoothing with application to chest wall nodule segmentation in lung CT data. In: Proc. SPIE. Volume 5370. (2004) 1519–1526
4. Wiemker, R., Rogalla, P., Blaffert, T., Sifri, D., Hay, O., Shah, E., Truyen, R., Fleiter, T.: Aspects of computer-aided detection (CAD) and volumetry of pulmonary nodules using multislice CT. Br. J. Radiol. **78** (2005) S46–S56
5. Reeves, A.P., Chan, A.B., Yankelevitz, D.F., Henschke, C.I., Kressler, B., Kostis, W.J.: On measuring the change in size of pulmonary nodules. IEEE Trans. Med. Imag. **25**(4) (2006) 435–450

6. Okada, K., Ramesh, V., Krishnan, A., Singh, M., Akdemir, U.: Robust pulmonary nodule segmentation in CT: Improving performance for juxtapleural cases. In: Proc. MICCAI. (2005) 781–789
7. Dehmeshki, J., Amin, H., Valdivieso, M., Ye, X.: Segmentation of pulmonary nodules in thoracic CT scans: A region growing approach. IEEE Trans. Med. Imag. **27**(4) (2008) 467–480
8. Kuhnigk, J.M., Dicken, V., Bornemann, L., Bakai, A., Wormanns, D., Krass, S., Peitgen, H.O.: Morphological segmentation and partial volume analysis for volumetry of solid pulmonary lesions in thoracic CT scans. IEEE Trans. Med. Imag. **25**(4) (2006) 417–434
9. Moltz, J.H., Kuhnigk, J.M., Bornemann, L., Peitgen, H.O.: Segmentierung pleuraständiger Lungenrundherde in CT-Bildern mittels Ellipsoidapproximation. In: Proc. BVM. (2008) 173–177
10. Fan, L., Qian, J., Odry, B.L., Shen, H., Naidich, D., Kohl, G., Klotz, E.: Automatic segmentation of pulmonary nodules by using dynamic 3d cross-correlation for interactive CAD systems. In: Proc. SPIE. Volume 4684. (2002) 1362–1369
11. Grammalidis, N., Strintzis, M.G.: Head detection and tracking by 2-d and 3-d ellipsoid fitting. In: Proc. CGI. (2000) 221–226