

Robust segmentation of pulmonary nodules of various densities: from ground-glass opacities to solid nodules

Toshiro Kubota¹, Anna Jerebko², Marcos Salganicoff², Maneesh Dewan², and Arun Krishnan²

¹ Department of Mathematics and Computer Science, Susquehanna University

² CAD and Knowledge Solutions, Siemens Medical Solutions, Inc. USA

Abstract. Accurate segmentation of a pulmonary nodule is an important and active area of research in medical image processing. Although many algorithms have been reported in literature for this problem, those that are applicable to various density types have not been available until recently. In this paper, we propose a new algorithm that is applicable to solid, non-solid and part-solid types and solitary, vascularized, and juxtapleural types. The algorithm works on the distance map computed from a foreground image. It first locates the core of a nodule in a manner that is robust against the presence of attached structures such as pleura and vessels, and then detaches the nodule from the attached structures by a variant of region growing and region partition.

The performance of the technique is evaluated using 23 data sets with manual segmentation and 1521 data sets with manual diameter measurements. The experiments show that the algorithm is highly reliable in segmenting convex nodules of various types.

1 Introduction

Pulmonary nodules are potential manifestations of lung cancer, and their detection and inspection are essential for screening and diagnosis of the disease. The growth of a nodule is considered the most important cue for assessing its malignancy.

Various segmentation methods targeted for pulmonary nodules have been developed, and some have been deployed in commercial applications. Many technical issues still remain, including accuracy and handling of non-solid and part-solid nodules. Most of existing segmentation algorithms are targeted toward nodules composed only with solid components, although studies have shown that nodules of non-solid and part-solid types are frequent and have higher risks of being malignant than solid ones [1]. More recent segmentation methods reported in literature claim to handle these types of nodules[2-4]. However, we believe that the field is relatively new and requires further investigation.

The goal of our work is to develop a semi-automated nodule segmentation algorithm that is applicable to not only the solid type but also the non-solid

and part-solid types. Figure 1 shows examples of pulmonary nodules of different density types. As exemplified in the figure, the appearance varies among different nodule types. In particular, non-solid nodules are extremely subtle with fuzzy boundaries, and part-solid nodules exhibit highly irregular intensity variations and boundary shapes. Thus, handling them under a single framework presents a great challenge to the segmentation problem. The inputs to the algorithm are a thin-slice thoracic CT volume and a click point. The output is a segmentation map of a nodule found in the vicinity of the click point. One assumption we make is that the nodule is convex. Thus, we are interested in segmenting a convex part of the nodule.




Fig. 1. Examples of pulmonary nodules of various densities. (a)-(b): Solid nodules, (c)-(d): Non-solid nodules, (e)-(f): Part-solid nodules.

The important feature of our algorithm is twofold. First, it locates the core of a nodule using a new approach that is highly robust against the presence of attached structures such as pleura and vessels and is invariant to the size of the nodule. The problem of segmenting a juxtapleural nodule lies in detecting a boundary between the nodule and the lung wall. We take the problem in two steps: locating the core of the nodule and growing a region from the core. Insensitivity to attached structures is a key in achieving the first step. Many previous techniques search both location and size of a nodule either in iterative manners or in scale-space representations [5, 6]. The size-invariance property of our approach makes it unnecessary to estimate the size of the nodule, thus renders the localization process computationally more efficient. Second, it employs a two-step region growing process as described in Sections 2.3 and 2.4, where the first step over-segments the nodule and the second step refines the segmentation via a competition process. The approach extracts a convex nodule from attached structures such as lung walls and vessels.

Throughout the paper, the following notations are used. A lower bold letter (\mathbf{x}) is used for a vector, an upper bold letter (\mathbf{X}) is used for a 3D volume, and a non-bold letter (x) is used for a scalar. Without ambiguity, we also use an upper bold letter associated with a binary volume (eg. segmentation) to indicate those voxels whose binary values are non-zero. For example, we use \mathbf{L} to denote a set of voxels that are foreground. Furthermore, we use $\neg\mathbf{L}$ to denote those voxels whose binary values are zero in \mathbf{L} .

2 Algorithm

2.1 Overview of the algorithm

The segmentation algorithm consists of six stages: pre-processing, figure-ground separation, localization of a nodule core, region growing, region partition, and post-processing. The first two stages are briefly discussed in this section. The other four are discussed, respectively, in the following sub-sections.

At the pre-processing module, a sub-volume that is large enough to contain a nodule of interest is extracted. The sub-volume is then resampled at a fixed interval to produce an isotropic sub-volume, called a region of interest (*ROI*). We denote a set of voxels in ROI as Ω . The figure-ground separation applies a bi-class segmentation algorithm previously developed by the authors to the ROI³. The result of the operator is a binary map denoted as $\mathbf{L} : \Omega \rightarrow \{0, 1\}$, where 0 and 1 voxels indicate backgrounds and foregrounds, respectively.

Figure 2 shows results of the figure-ground separation process applied to 23 data sets provided by the NIH Lung Imaging Database Consortium (LIDC)[7]. For each data, five images are shown. The left image is a representative axial slice of the nodule in the original CT data. The second image from the left is the result of the figure-ground separation. The other three images are results of seed point localization and segmentation processes, which are discussed in detail below. A bold number accompanying each set of images is a unique number for identifying the data, and two non-bold numbers indicate the accuracy of the segmentation. We will describe the accuracy measures in Section 3. Note that the images are shown at the resolution of the original CT data. Thus, all results shown are interpolated back to the resolution of the original CT data.

2.2 Localization of Nodule Center

After the figure-ground separation, our next task is to locate a core of a nodule inside the ROI. We first apply Euclidean distance transform to \mathbf{L} to compute for each location in \mathbf{L} the shortest Euclidean distance to the background. Call the resulting distance map \mathbf{D} . Our motivation for working on \mathbf{D} is twofold. First, \mathbf{D} is free from signal noise present in the original intensity volume. Thus, no noise removal and signal restoration processes are needed in the subsequent processes. Furthermore, it renders the rest of the algorithm insensitive to intensity variations of images. This is important as solid, non-solid and part-solid nodules present significantly different intensity distributions. Second, \mathbf{D} compactly encodes information regarding the shape of the foreground. For example, \mathbf{D} tells how close each voxel is to the background. It provides, by means of the gradient vector computed on \mathbf{D} , a rough direction to the background.

A limitation of \mathbf{D} is that it is extremely sensitive to a small hole present in the foreground and partially sensitive to an attached structure. The foreground extraction procedure described above effectively removes small holes in \mathbf{L} . We will discuss how to handle the attachment next.

³ For double blind review, details of our previous work are intentionally omitted.




Fig. 2. Segmentation results on LIDC data set. For each data set, five images are shown. They are from left to right: original slice, figure-ground separation map, sphericity map, segmentation result, and comparison with manual segmentation. In the rightmost image, a dark gray color shows the area selected by at least one segmentation, a light gray color shows the area selected by 90% of the methods, and a white line shows the boundary of our segmentation result. If no light gray area is shown, the slice did not contain an area agreed by 90% of the methods. Three numbers on the side of each image set are a unique identification number, under-segmentation fraction percentage, and over-segmentation fraction percentage.

When the nodule is modeled as a sphere and more than a half of the nodule is buried in the lung wall, the local maximum of \mathbf{D} no longer resides inside the nodule. To mitigate the problem, we propose another transformation built on top of \mathbf{D} as defined below.

$$\mathbf{S}(\mathbf{x}) = \int_{\mathbf{y} \in N(\mathbf{x})} \chi(\mathbf{D}(\mathbf{x}) - \mathbf{D}(\mathbf{y})) d\mathbf{y} \quad (1)$$

where $N(\mathbf{x})$ is a ball of radius δ centered at \mathbf{x} and $\chi(x) = 1$ if $x > 0$ and 0 otherwise. We call \mathbf{S} *sphericity*.

It is convenient to introduce the following definition of $\mathbf{C}(d)$:

$$\mathbf{C}(d) = \{\mathbf{x} \in \Omega | \mathbf{D}(\mathbf{x}) \leq d\}. \quad (2)$$

We call $\mathbf{C}(d)$ the *level core* of \mathbf{D} at d . The next lemmas indicate that \mathbf{S} can replace \mathbf{D} while providing more robust protection against a pleural attachment.

Lemma 1. *For a convex foreground, \mathbf{D} has a single local maximum component, and for every d , $\mathbf{C}(d)$ is convex.*

Lemma 2. *For a convex foreground, \mathbf{S} has a single local maximum component, whose position is identical to the local maximum of \mathbf{D} . Furthermore, with a sufficiently small δ , \mathbf{x} is a local maximum of \mathbf{S} if \mathbf{x} is a local maximum of \mathbf{D} .*

Lemma 3. *For a spherical nodule, its center remains a local maximum of \mathbf{S} under any degree of partial occlusion by a half plane.*

The first lemma justifies our approach in segmenting a nodule by delineating every level set around the local maximum. The second lemma indicates that \mathbf{S} can be used instead of \mathbf{D} to locate the core of a nodule. The third lemma suggests that the local maximum of \mathbf{S} is more stable against partial occlusion by the half plane than that of \mathbf{D} .

Figure 3 compares \mathbf{D} and \mathbf{S} of a sphere under various degrees of occlusion by a half plane. The radius of the sphere is 100 voxel unit. The top row shows instances of \mathbf{D} and the bottom row shows instances of \mathbf{S} . For computational simplicity, $N(\mathbf{x})$ is set as a 3x3x3 cube instead of a sphere. From left to right, columns (a-e) show a solid plane penetrating the sphere at (a) 30 voxels above, (b) 0 voxel above, (c) 30 voxels below, (d) 60 voxels below, and (e) 90 voxels below the sphere center, respectively. When the plane penetrates at below the center, the local maximum of \mathbf{D} shifts to inside the wall while the local maximum of \mathbf{S} retains its position regardless of the degree of partial occlusion. Although Lemma 3 is limited to a sphere, robustness of \mathbf{S} in detecting the nodule center has been observed for other more general shapes.

Note that the sphericity value at the nodule center is 1 regardless of the size of the nodule when $N(\mathbf{x})$ is completely contained inside the nodule. With our setting of 3x3x3 voxels for $N(\mathbf{x})$, this implies that the sphericity value is invariant to the nodule size no smaller than 3x3x3 voxels.

To locate a nodule center, we compute \mathbf{S} and locate a local maximum of \mathbf{S} in a neighborhood of the click point. In our experiment, a $7 \times 7 \times 7$ sub-volume centered at the click point constitutes the neighborhood. Figure 2 shows sphericity maps and seed points resulted from applying the above procedure to 23 LIDC data sets. The middle image in each set of five images is the axial slice of \mathbf{S} at the detected seed point. The seed point is shown by a black cross.




Fig. 3. Distance (top) and sphericity (bottom) maps at various degrees of occlusion.

2.3 Region Growing

Once we locate the core of a nodule, the next task is to extract the nodule from the foreground. We apply region growing on \mathbf{D} starting from the seed point detected in Section 2.2 and incrementally include each level set to the growing region. The approach can be implemented in the following algorithm. Note that \mathbf{R} denotes a resulting segmentation map, which is initialized to constant zero.

Input: \mathbf{D} : Distance map, \mathbf{s} : seed point

Output: \mathbf{R} : Segmentation map

$d \leftarrow \mathbf{D}(\mathbf{s})$;

while $d > 0$ **do**

foreach z that is adjacent to the current region in \mathbf{R} **do**

if $D(z) = d$ **then**

$\mathbf{R}(z) \leftarrow 1$;

 decrement d to the next possible distance value

Algorithm 1: Region-growing algorithm

The approach is mainly justified by Lemma 1; For a convex isolated nodule, there is a single local maximum component, from which each level set can be included into the region incrementally. For a non-isolated nodule, the foreground may not be convex. As exemplified by Lemma 3, the nodule center can be located by the local maximum of \mathbf{S} , and we can extract the core of the nodule by Algorithm 1. However, at some point of d , $\mathbf{C}(d)$ becomes non-convex. Including the entire $\mathbf{C}(d)$ into the region results in significant over-segmentation. By

approaching the inclusion of each level set in a region growing manner (that is to include only a single layer around the current region into the region), we can limit the over-segmentation.

Figures 4(a) and (d) show results of Algorithm 1 applied to two artificial data which illustrate juxtapleural and vascularized cases. In both cases, the nodule is a sphere, and the foreground is non-convex due to the attached structure. The black, gray, and white regions are $\neg\mathbf{L}$, \mathbf{L} and $\neg\mathbf{R}$, and \mathbf{R} , respectively. Although the algorithm successfully segmented the entire nodule, it also included a portion of the attached structure near the nodule. We leave it to the next step to remove the oversegmented portion.




Fig. 4. Segmentation of simple artificial data.

2.4 Region Partition

The objective of this stage is to remove over-segmented parts of the region-growing segmentation. We treat the problem as follows. The segmentation of Algorithm 1 consists of multiple convex regions, among which the one including the seed point is the nodule of interest. Thus, the objective is to isolate the nodule from other convex regions. We take the same strategy as in Algorithm 1, as each convex part has a local maximum of \mathbf{D} and a convex core around it. We need to compute a new \mathbf{D} using \mathbf{R} , which we denote as $\tilde{\mathbf{D}}$. Since \mathbf{R} is already isolated from a pleural wall, it is not necessary to carry out computation of \mathbf{S} and find its local maxima; $\tilde{\mathbf{D}}$ and its local maxima are sufficient.

Local maxima in $\tilde{\mathbf{D}}$ are treated as a new set of seed points, which are either *positive* or *negative*. The positive seed points are in the vicinity to the original seed point and contribute toward delineating the nodule part. The negative seed points are those that are not positive, and contribute toward delineating non-nodule components. We use the following steps to select positive seed points. First, the location of the original seed point is moved to the local maximum of $\tilde{\mathbf{D}}$ by tracing the gradient of $\tilde{\mathbf{D}}$. Call this relocated seed point $\tilde{\mathbf{s}}$. Then, those local maxima of $\tilde{\mathbf{D}}$ that are within $\tilde{\mathbf{D}}(\tilde{\mathbf{s}})$ away from $\tilde{\mathbf{s}}$ are classified positive.

Once we collected both positive and negative seed points, we perform a region growing on $\tilde{\mathbf{D}}$ from each seed point in down-hill directions. A region grown from a positive seed point is marked as nodule, while a region grown from a negative seed point is marked non-nodule. The process continues until no further growth is possible. Algorithm 2 summarizes the process. \mathbf{P} denotes the result, which is initialized to constant zero.

Input: \tilde{D} : distance map, \mathbf{s}^+ : positive seed points, \mathbf{s}^- : negative seed points
Output: \mathbf{P} : Partition map
 $\mathbf{P}(\mathbf{s}^+)=1$; $\mathbf{P}(\mathbf{s}^-)=-1$;
repeat
 foreach x in $\{z | \mathbf{P}(z) = 0\}$ **do**
 if x has a positive neighbor y & $\tilde{D}(x) \leq \tilde{D}(y)$ **then**
 $\mathbf{P}(x)=1$;
 if x has a negative neighbor y & $\tilde{D}(x) \leq \tilde{D}(y)$ **then**
 $\mathbf{P}(x)=-1$;
 until No change to \mathbf{P} ;
Set all negative voxels in \mathbf{P} to 0;
Algorithm 2: Region-partition algorithm

Results of Algorithm 2 applied to the artificial data of Figure 4(a) and (d) are shown in Figure 4(b) and (e), in which black, gray, and white regions are $-\mathbf{R}$, \mathbf{R} and $-\mathbf{P}$, and \mathbf{P} , respectively. In both examples, most of the nodule is included in \mathbf{P} while non-nodule components are successfully eliminated.

2.5 Post-processing

Segmentation after Algorithm 2 is not necessarily convex. At this final stage, \mathbf{P} is first made convex by convex hull of \mathbf{P} , and then taken intersection with \mathbf{L} . Figure 4(c) and (f) show results of this stage applied to the artificial data of Figure 4(a) and (d). We call the segmentation after this stage \mathbf{F} .

3 Experiments

3.1 LIDC Segmentation

First, we apply the segmentation algorithm to the LIDC data set. The study came from two cancer cases with collected segmentation done by six radiologists using three methods: one fully manual and two with automated programs. Thus, it collected a total of 18 segmentation maps. A sub-volume of $81 \times 81 \times L$ voxels is extracted around the nodule where L is the number of slices in the data set. A click point is set at the center of the sub-volume. We also derive two different segmentation maps from the manual segmentation. One is a set of voxels included in at least one out of 18 methods. The other is a set of voxels included in 90% of the methods. The former is denoted as \mathbf{G}_1 and the latter \mathbf{G}_2 .

Figure 2 shows the result of our segmentation for each LIDC nodule. As stated in Section 2.1, each nodule corresponds to a set of horizontally aligned five images where the first is a representative axial slice of the original sub-volume, the second shows \mathbf{L} at the slice, the third shows \mathbf{S} and \mathbf{s} . The fourth image shows \mathbf{F} super-imposed on top of the original sub-volume. The fifth is made by first painting \mathbf{G}_1 in dark gray, then painting \mathbf{G}_2 in light gray, and finally painting the boundary of \mathbf{F} in white. Shown to the left of the image set are

three numbers, which are, from top to the bottom, the unique data identification number (in bold), under-segmentation fraction (p_U) and over-segmentation fraction (p_O). The under/over segmentation fraction percentages are defined as $p_U = 100|\neg\mathbf{F} \cap \mathbf{G}_2| |\mathbf{G}_2|^{-1}$ and $p_O = 100|\mathbf{F} \cap \neg\mathbf{G}_1| |\mathbf{G}_1|^{-1}$, respectively, where $|\cdot|$ is the cardinality of the set.

3.2 Automated Diameter Measurement

We apply our segmentation algorithm to a larger set of CT data. Since it is difficult to obtain a reliable segmentation on a large set of data, we use the diameter of a nodule to test the accuracy of our segmentation. The test data consist of 1521 nodules from 253 cases, where 1237 are solid, 206 are non-solid, and 77 are part-solid. Each sub-volume after isotropic resampling is 41x41x41 voxels. The ELCAP protocol ([8]) is used to measure the nodule diameter both manually and automatically from the segmentation obtained by our algorithm.

Figure 5(a) shows a scatter plot of the manual and automated diameter measurements. The horizontal axis is the manual diameter measurement and the vertical axis is the automated diameter measurement. Figure 5(b) shows a histogram of normalized errors ($\bar{\epsilon}$) defined as $\bar{\epsilon} = |D_m - D_a| D_m^{-1}$ where D_m and D_a are manual and automated diameter measurements, respectively. Table 3.2 shows descriptive statistics of estimates.




Fig. 5. Diameter measurement results on 1521 nodule data.

4 Discussion and Conclusion

The paper presented a new general purpose segmentation algorithm for a pulmonary nodule. The algorithm is applicable to solid, non-solid and part-solid types and handles juxtapleural and vascularized ones without having the lung separated from the pleural walls. The results on 23 LIDC data and 1521 data with manual diameter measurements are highly encouraging.

A limitation of the approach is that it cannot handle spiculated nodules due to its underlying assumption of nodule convexity. For a spiculated nodule, we

Table 1. Descriptive statistics of the diameter measurement experiment

	ρ^1	μ^2	m^3	σ^4	τ^5
All	0.755	0.359	0.254	0.397	0.607
Solid	0.759	0.372	0.256	0.422	0.603
Non-solid	0.624	0.334	0.285	0.266	0.611
Part-solid	0.864	0.214	0.177	0.163	0.664

¹ ρ : Pearson correlation coefficient.

² μ : mean normalized absolute error.

³ m : median normalized absolute error.

⁴ σ : std normalized absolute error.

⁵ τ : mean computation time (sec) on 2.8GHz PC with 1G memory.

may be able to treat it as one comprised of multiple convex parts, and extract each part in succession. Another possible solution is to perform segmentation at a coarse resolution where the degree of spicularity is reduced. Another limitation is that it tends to over-segment non-solid nodules. The limitation can be circumvented by referring to the ROI intensity volume and carefully trimming voxels that are likely to be partial volumes.

References

1. Henschke, C., Yankelevitz, D.F., Mirtcheva, R., McGuinness, G., McCauley, D., Miettinen, O.S.: CT screening for lung cancer: Frequency and significance of part-solid and nonsolid nodules. *Am. J. of Roent.* **178** (2002) 1053–1057
2. Kauczor, H.U., Heitmann, K., Heussel, C.P., Marwede, D., Uthmann, T., Thelen, M.: Automatic detection and quantification of ground-glass opacities on high-resolution ct using multiple neural networks: Comparison with a density mask. *Am. J. of Roent.* **175** (2000) 1329–1334
3. v. Ginneken, B.: Supervised probabilistic segmentation of pulmonary nodules in CT scans. In: MICCAI. (2006) 912–919
4. Wang, J., Engelmann, R., Li, Q.: Segmentation of pulmonary nodules in three-dimensional CT images by use of a spiral-scanning technique. *Med. Phys.* **34**(1) (2007) 4678–89
5. Diciotti, S., Picozzi, G., Falchini, M., Mascalchi, M., Villari, N., Valli, G.: 3-d segmentation algorithm of small lung nodules in spiral CT images. *IEEE Trans. Inf. Technol. Biomed.* **12**(1) (2008) 7–19
6. Okada, K., Comaniciu, D., Krishnan, A.: Robust anisotropic gaussian fitting for volumetric characterization of pulmonary nodules in multislice CT. *IEEE Trans. Med. Imaging* **24**(3) (Mar. 2005) 409–423
7. Institute, N.C.: Lids: Datasets as a public resource. <http://imaging.cancer.gov/reportsandpublications/reportsandpresentations/firstdataset>
8. Henschke, C.: International early lung cancer action program: Protocol. [Online], <http://www.ielcap.org/ielcap.pdf> (Dec. 2007)