

Fast and Memory Efficient Segmentation of Lung Tumors Using Graph Cuts

Nicolas Lermé^{1,2}, François Malgouyres¹, and Jean-Marie Rocchisani^{3,4}

(1) LAGA CNRS UMR 7539, (2) LIPN CNRS UMR 7030, (3) SMBH
Université Paris 13 –Avenue J.B. Clément
93430 Villetaneuse - France

(4) Hôpital Avicenne, 93009 Bobigny - France
`nicolas.lerme@lipn.univ-paris13.fr`,
`malgouy@math.univ-paris13.fr`,
`jean-marie.rocchisani@univ-paris13.fr`

Abstract. In medical imaging, segmenting accurately lung tumors remains a quite challenging task when they are directly in contact with healthy tissues. In this paper, we address the problem of extracting interactively these tumors with graph cuts. The originality of this work consists in (1) reducing input graphs to decrease drastically memory consumption when segmenting a large volume of data and (2) introducing a novel energy formulation to inhibit the propagation of the object seeds. We detail our strategy to achieve relevant segmentations of lung tumors and compare our results to hand made segmentations provided by an expert. Comprehensive experiments show how our method can give solutions near from ground truth in a fast and memory efficient way.

Keywords: segmentation, lung tumor, graph cut, reduction.

1 Introduction

Since last years, accurate measurements of lung tumors sizes has become a challenging task for staging and assessing tumor response to treatments or its progression. Revised RECIST criterions, largely used by radiologists, are based on the measurement of one diameter on a few number of lesions [24], and suffer from a lack of reproducibility [22]. Alternatively, tumor volumetry has been proposed to overcome those difficulties in order to improve the staging of nodules [5], the evaluation of tumor aggressiveness [18], tumor response to chemotherapy [3,26] or to radiotherapy [16] and the progression rate of tumors [18] or metastases [15]. Moreover, it becomes a necessary tool for the automatic screening of lung nodules on CT scans, and is currently on evaluation on ongoing trials [23]. Several methods have been proposed to deal with the different kind of objects to segment. Nodules are homogeneous spheroid of small size. Masses and tumors have larger sizes and irregular shapes, and may be necrotic. All may be connected to some extent to vessels, to the pleura wall, or to the mediastinum. To tackle this issue,

methods make often use of morphological operators [9,10,17]. A classification of those methods can be found for instance in [21] and [5].

Among semi-automatic approaches of segmentation based on level-sets and (geodesic) active contours, graph cuts have become in few years a leading method since the introduction of a fast maximum-flow/minimum-cut algorithm [2]. In contrast to other methods, graph cuts have the ability to solve quickly a wide range of problems in computer graphics such as image segmentation [1], while achieving a global minimum of the energy function.

Recently, Ye *et al.* have used this technique for automatically segmenting lung nodules using a volumetric shape index [25]. Since nodules have presumably an elliptical shape, they can select the appropriate range of index values for segmenting nodules. However, segmenting lung tumors of various shapes is a much more difficult task. Tumors might indeed be connected to healthy tissues and it is not possible to distinguish the tumor and the healthy tissues by only using simple features like the gray levels. The correct segmentation can therefore only be achieved thanks to the interaction of an expert. To our knowledge, this is the first paper to tackle this problem using graph cuts. In this paper, we propose a semi-interactive graph cut-based method for segmenting lung tumors. An overview of the approach is given on Figure 1. First, we compute a distance




Fig. 1: Flow diagram of our approach.

map from the object seeds for lowering the “seeds propagation”. Then during the graph construction, we reduce the input graph by deciding locally which nodes are really useful for the minimum-cut computation according to [12]. Typically, the nodes are located around the contours of the object to segment. Finally, we compute the minimum-cut and get the final solution.

The rest of this paper is organized as follows. In section 2, we review the graph-cuts framework. We detail our strategy for reducing graphs in section 3 while we introduce our novel energy formulation in section 4. Finally, we validate experimentally our algorithm on several CT images in section 5.

2 Graph cuts framework

Let us first review the graph cuts framework. In this setting, an image I is a function defined over a finite discrete set $\mathcal{P} \subset \mathbb{Z}^d$ ($d > 0$) that maps each point $p \in \mathcal{P}$ to a value $I(p)$. Usually, \mathcal{P} correspond to a square when $d = 2$, a cube when $d = 3$ and a cube during a time interval when $d = 4$. A binary segmentation

of the image is defined by a mapping u that assigns to each element of \mathcal{P} the value 0 for the background and 1 for the object. We write $u \in \{0, 1\}^{\mathcal{P}}$.

In [1], Boykov and Jolly showed that the image segmentation problem can be efficiently solved by minimizing a Markov Random Field of the form:

$$E(u) = \beta \cdot \sum_{p \in \mathcal{P}} E_p(u_p) + \sum_{\substack{p, q \in \mathcal{P} \\ q \in \mathcal{N}(p)}} E_{p,q}(u_p, u_q), \quad (1)$$

among $u \in \{0, 1\}^{\mathcal{P}}$ and for $\beta > 0$. The neighborhood system $\mathcal{N}(p)$ is in practice either

$$\begin{aligned} \mathcal{N}_0(p) &= \{q : \sum_{i=1}^d |q_i - p_i| = 1\} & \forall p \in \mathcal{P}, \text{ or} \\ \mathcal{N}_1(p) &= \{q : |q_i - p_i| \leq 1, \forall 1 \leq i \leq d\} & \forall p \in \mathcal{P}, \end{aligned}$$

where p_i denote the i^{th} coordinate of the point p and $|\cdot|$ denotes the modulus. (in this paper $|\cdot|$ also denotes the cardinality of a set, the notations will not be ambiguous once in context). The above neighborhood systems correspond to the classical 4-connectivity and 8-connectivity when $d = 2$. Beside on the border of the image/volume, we have for any d and any $p \in \mathcal{P}$: $|\mathcal{N}_0(p)| = (2d)$ and $|\mathcal{N}_1(p)| = 3^d - 1$. In practice, larger neighborhood systems (i.e. \mathcal{N}_1) yield better results but increase running time and memory consumption. Typically, we have $|\mathcal{E}_n| \sim |\mathcal{P}| \cdot |\mathcal{N}|$, where $|\cdot|$ denotes cardinality. In the sequel, the terms ‘‘connectivity 0’’ and ‘‘connectivity 1’’ will denote the use of a \mathcal{N}_0 and \mathcal{N}_1 neighborhood, respectively.

As usually, the region term $E_p(\cdot)$ in (1) favors the belonging of each pixel/voxel to either the background or to the object. It is deduced from the input data, an object seed \mathcal{O} and a background seed \mathcal{B} . The regularity term $E_{p,q}(\cdot)$ penalizes neighboring pixels p and q having different labels. The weight of the penalization depends on the difference $|I(p) - I(q)|$ and favors boundaries located at pixels/voxels with a strong gradient. Generally speaking, the definition of E_p and $E_{p,q}$ depends on the considered application.

According to [8], the minimizer of the energy (1) corresponds to a minimum-cut in a graph that can be efficiently computed by the algorithm proposed in [2]. In this context, the directed weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, c)$ consists of a set of nodes $\mathcal{V} = \mathcal{P} \cup \{s, t\}$, a set of edges $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ and a positive weighting function $c : \mathcal{E} \rightarrow \mathbb{R}^+$ defining the edge capacity. Notice that two special nodes are distinguished from \mathcal{V} : the source node s (‘‘object’’ terminal) and the sink node t (‘‘background’’ terminal). After the computation of the minimum cut we set $u_p = 1$, if p is connected to s and $u_p = 0$ otherwise. Moreover, the set of edges \mathcal{E} is split into two disjoint sets \mathcal{E}_n and \mathcal{E}_t denoting respectively n-links and t-links. The t-links are the edges connecting the terminal nodes s or t to the pixels/voxels and the n-links are the edges connecting pixels/voxels.

3 Reducing graphs

To obtain high-resolution output, graph cuts must build huge graphs containing several billions of nodes and even more edges. Such graphs may sometimes do not

fit in central memory. To solve this issue, some authors have recently proposed heuristics [13,14,20,7]. However, these algorithms can easily get trapped in local minima of the energy. Also, these algorithms often fail to recover details. This is a real drawback since thin structures like blood vessels or nodules are ubiquitous in medical imaging. The only exact alternative is [11], but it has not been developed for the purpose of image segmentation.

Thus, segmenting high-resolution data using graph cuts require a prohibitive amount of memory. For instance, the maximum-flow algorithm described in [2] allocates $24|\mathcal{P}| + 14|\mathcal{E}_n|$ bytes¹. Table 1 shows that for a fixed amount of RAM, the maximum volume size decreases quickly as dimension d increases. Nevertheless,

	Connectivity 0	Connectivity 1
2D	6426	4459
3D	319	219
4D	68	45

Table 1: Maximum size of a square image for which the graph fits in 2GB of RAM.

as showed in a previous paper [12], most of the nodes in the graph are useless during the maximum-flow computation. They are indeed not traversed by any flow. Then, one would like to extract the smallest possible graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}', c)$ from \mathcal{G} while keeping a minimum cut u' identical (or very close) to u . In other words, we want to minimize $|\mathcal{V}'|$ under the constraint that $u \simeq u'$. In fact, this is an ideal optimization problem which we will not try to solve, because the method for determining \mathcal{G}' also needs to be (very) fast. We will rather consider heuristics aiming at that goal.

First, let us introduce some definitions before describing our method for building \mathcal{G}' . In accordance with the graph construction given in [8], we consider (without loss of generality) that a node is linked to at most one terminal:

$$(s, p) \in \mathcal{E}_t \Rightarrow (p, t) \notin \mathcal{E}_t, \quad \forall p \in \mathcal{P}.$$

We also summarize the capacities on the t-links connected to any node $p \in \mathcal{P}$:

$$c(p) = c(s, p) - c(p, t).$$

Let us consider a square window B of size $(2r + 1)$ ($r > 0$) centered at the origin. We denote by \tilde{B}_p the translation of B at a point $p \in \mathcal{P}$: $\tilde{B}_p = \{b + p \mid b \in B\}$. For $Z \subset \mathcal{P}$, we also denote by $\tilde{Z}_B = \bigcup_{p \in Z} \tilde{B}_p$ the dilation of Z by B .

The intuitive idea for building \mathcal{G}' is the following: removing the nodes in any $Z \subset \mathcal{P}$ such that pixels/voxels in Z are not directly connected to the sink t and the flow that might come into the region $\tilde{Z}_B \setminus Z$ suffices to saturate the edges located around \tilde{Z}_B (see Figure 2). Building such sets Z is done by testing each pixel p of Z . Thus, the nodes in \mathcal{G}' are typically located around the contours of the object to segment. Assuming that all capacities on n-links are smaller than

¹ This corresponds to the max-flow algorithm v2.2 freely available at <http://www.cs.cornell.edu/People/vnk/software.html>




Fig. 2: Principle of the reduction. The nodes from Z are removed because every node $p \in Z$ satisfy (2). Remaining nodes are typically located in the narrow band $\tilde{Z}_B \setminus Z$.

one (which remains true for all the energy models in segmentation), we use a more conservative condition for testing each individual pixel $p \in Z$ [12]:

$$\left\{ \begin{array}{l} (\forall q \in \tilde{B}_p, c(q) \geq \delta) \\ (\forall q \in \tilde{B}_p, c(q) \leq -\delta) \end{array} \right\} \quad \text{or} \quad (2)$$

where $\delta = \frac{P(B)}{(2r+1)^2-1}$, with

$$P(B) = \max(|\{(p, q) : p \in B, q \notin B \text{ and } p \in \mathcal{N}(q)\}|, |\{(p, q) : p \in B, q \notin B \text{ and } q \in \mathcal{N}(p)\}|).$$

For any p satisfying (2), p is only connected to s (respectively t) and the flow that might come in (respectively come out) through t-links in $\tilde{B}_p \setminus \{p\}$ suffices to saturate the n-links going out (respectively going in) of \tilde{B}_p . The pixel/voxel p is not needed and can be removed from \mathcal{G} . The subgraph \mathcal{G}' is now fully determined by the set of nodes

$$\mathcal{V}' = \{p \in \mathcal{P} \text{ not satisfying (2)}\} \cup \{s, t\}.$$

Experiments presented in [12] confirm the intuitive dependence between the reduction rates and the parameters of the model. For instance, the capacities $c(q)$ are obtained by multiplying a quantity by the parameter β of (1). Looking at (2), it is straightforward to see that the test is satisfied on a smaller set of pixels/voxels if β decreases. In fact, β small corresponds to a strong regularization. In such a situation, we need a larger window radius to obtain a smaller δ . The latter results in wide bands around the object contours. Conversely, this results in narrow bands around the object contours when β is large. The result of such a reduction is illustrated in Figure 3. In our experiments, we always take $\beta = 3$ and $r = 1$. Additionally, the condition (2) can be tested through an easy to implement “non-optimized” algorithm with a worst-case complexity of $O(|B|)$. However for large window radii, such an algorithm cannot handle images of large size and large dimension d . Decomposing (2) along the dimensions d speed up significantly the previous algorithm. This yields a test whose computation is of complexity $O(1)$ (except for image borders). In particular, its complexity is independent of the window radius. Finally, we have both theoretical and empirical evidence suggesting that this reduction scheme provides an exact solution (see [12] for details).




Fig. 3: Illustration of the reduction for segmenting a CT image ($r = 1$). Light gray pixels correspond to the nodes belonging to \mathcal{G}' (middle). Object and background seeds are superimposed on the original image (left). On the right image, the segmentation is superimposed in blue.

4 Energy function

The most famous graph cut-based energy model for image segmentation was proposed by Boykov and Jolly in [1] (see below). Total Variation-based models have also been proposed (see [19]). To obtain good results, those models require the colors in the object to be different from the colors of the background. This requirement is not efficient when segmenting lung tumors in CT images, because tumors and healthy tissues appear in the same range of intensities. Moreover, in many cases (and in our experiments), the tumor is attached to the healthy tissues and the corresponding area of the image has a uniform color.

To solve this issue, we propose to add in our energy a prior on the location of the tumor. The prior is obtained from the location of the object seeds. This leads to a modification of the original Boykov/Jolly's energy model [1]. We take the same regularity criterion:

$$E_{p,q}(u_p, u_q) = g(p, q) \cdot |u_p - u_q| \quad \text{and} \quad g(p, q) = \frac{1}{d_e(p, q)} \cdot \exp\left(-\frac{|I(p) - I(q)|^2}{2\sigma^2}\right),$$

where d_e is the Euclidean distance between p and q , I is the original image and $\sigma > 0$. The region term is defined in Table 2. The sets \mathcal{O} and \mathcal{B} correspond

	$p \in A_{\sigma_a}$	$p \notin A_{\sigma_a}$
$E_p(u_p = \text{"bkg"})$	$-\log [Pr(I(p) p \in \mathcal{O}) \times \exp(-\frac{d(p, \mathcal{O})^2}{\sigma_a^2})]$	$+\infty$
$E_p(u_p = \text{"obj"})$	$-\log [Pr(I(p) p \in \mathcal{B})]$	0

Table 2: Definition of the region term.

respectively to object and background seeds provided by the user, the probability distributions are estimated according to [1], $d(p, \mathcal{O})$ is a distance function between the point $p \in \mathcal{P}$ and the set $\mathcal{O} \subset \mathcal{P}$ and $\sigma_a > 0$ is a parameter. The parameter σ_a controls how far the object seeds propagate from their location and then defines an area of influence A_{σ_a} . Beyond this area, the nodes are only

linked to the background terminal with a large weight. This ensures both that the algorithm categorizes them as background pixels or voxels and that the capacity of the corresponding t-link is sufficiently high for removing the node from the graph. Although the parameter σ_a is an important parameter that impact the way of positioning the seeds in the image, we always take $\sigma_a = 10$ in our experiments.

The main difference between the proposed energy and [1] lies in the distance term. The function d is defined as $d(p, \mathcal{O}) = \min\{dist(p, q) \mid q \in \mathcal{O}\}$, where $dist$ denotes the distance between two points. We have made two attempts for $dist$:

- The Euclidean distance. In this case, the distance between a set and a point is efficiently computed with the algorithm described in [4]. We mostly use it for the purpose of illustration.
- The geodesic distance is according to the graph metric where the distance between a node $p \in \mathcal{P}$ and a node $q \in \mathcal{P}$ is defined as:

$$dist(p, q) = \begin{cases} \sqrt{(I(p) - I(q))^2 + |p - q|^2} & \text{if } q \in \mathcal{N}(p), \\ 0 & \text{otherwise.} \end{cases}$$

In this latter case, the distance transform is computed with [6].

The area of influence for the above two metrics is displayed through an example on Figure 4. The green color corresponds to the region where the exponential in the region term is greater than some $\epsilon \simeq 0$. Observe how the geodesic distance better sticks to the tumor boundaries than the Euclidean distance. In particular, it only has a limited overflow on the healthy tissues.




Fig. 4: Area of influence for an Euclidean and a geodesic distance. Here, we set $\sigma_a = 40$.

5 Experimental results

In this section, we present experiments for segmenting a set of ten 3D CT images consisting both of nodules, masses and tumors (see Table 3 for more information). Each volume has a size of $512 \times 512 \times 50$ except T8 which has a size of $512 \times 512 \times 316$. All experiments are performed in connectivity 1. Objects to segment may present a very different contrast with their surrounding structures

among the images. Since the parameter σ is contrast-sensitive, we are constrained to use different values for this parameter. For example, when the average gradient around the object become lower, we need a smaller σ . Then, the edges around the object are more likely to belong to the minimum-cut because they become cheaper to cut. The automatic tuning of σ is left for future work. In this setting, we use $\sigma = 0.2$ for all images except for T8 where $\sigma = 2$ and T7 where $\sigma = 0.05$. Note that a sub-volume is automatically extracted for all images (except for T8 where the border is sufficiently high to encompass the whole volume) by considering an extra border of 60 pixels around the object seeds for speeding up the segmentation.

First, we evaluate our algorithm with hand made segmentations provided by an expert, for all CT images. Table 4 contains statistics on the differences between the segmentation and the ground truth. We use several evaluation measures ². Table 4 shows promising results. For all images, we always get a Dice Coefficient greater than 70% while having a mean maximum distance less than 20mm between the ground truth and the segmentation.

We also evaluate our method in a qualitative manner. Figure 7 shows the segmentations obtained at equally spaced z-values for images T1, T8 and T9 (see Figure 5). For illustrating the propagation of seeds, the seeds in the Figure were chosen on equally spaced on z but for different values. Thus, one can observe how the seeds propagate around object seeds, avoiding us to mark every slice. Compared to the ground truth, the segmentation of T1 is very close, while the segmentations obtained for T8 and T9 differ slightly. This also illustrate the difficulty to extract tumors/masses with a large connection to healthy tissues and the ability of our method to segment such objects.

Secondly, we compare the performance of standard graph cuts against our method in terms of speed and memory consumption (see Table 5) for segmenting the CT images using the same set of seeds and parameters as previously. Experiments were performed on an Athlon Dual Core 6000+ 3GHz with 2GB RAM. Times are averaged over 10 runs. Table 5 also indicates the proportion of object seeds with respect to the tumor volume in the ground truth. This provides an objective measure of the interaction for assessing the effort required by the user for positioning the seeds. The results obtained show that our method performs a little bit faster using 7 to 500x less memory while getting exactly the same solution. Note that a relatively small amount of seeds is necessary for segmenting all images.

Generally, the segmentation time depends both on the image size and the skill of the user for positioning the seeds not too far from the contours of the object to segment. The segmentation accuracy also depend directly on the seeds location but additionnal corrections can be done quickly if some parts of the object are uncorrectly labeled. The computation of the distance map, the building of the graph and the computation of the minimum-cut take only few seconds. Thus, our method demonstrates its ability to segment lung tumors quickly without requiring much effort if it is supported by a good graphical user interface.

² A detailed view of these measures is available at <http://lts08.bigr.nl/about.php>

Tumor	Type	Resolutions (x,y,z)	Description
T1	Mass	$0.68 \times 0.68 \times 3$	Mass of the upper right lobe (CT)
T2	Nodule	$0.70 \times 0.70 \times 1$	Nodule of the right apex (CT)
T3	Nodule	$0.68 \times 0.68 \times 3$	Nodule of the lower right lobe (CT)
T4	Tumor	$1.17 \times 1.17 \times 1.5$	Marge left hilar tumor inducing a peripheral atelectasia (CT)
T5	Tumor	$1.17 \times 1.17 \times 1.5$	Same as T4 (dosimetric CT scanner)
T6	Mass	$0.77 \times 0.77 \times 1.25$	Mass of the lower left lobe appended to the pleura (CT)
T7	Mass	$0.69 \times 0.69 \times 1.25$	Same as T6, after four months of treatment (CT)
T8	Tumor	$0.63 \times 0.63 \times 1$	Large left hilar tumor and peripheral atelectasia, before treatment (contrast enhanced CT)
T9	Tumor	$0.70 \times 0.70 \times 1$	Same as T8, after chemo-radiotherapy (CE-CT)
T10	Mass	$1.17 \times 1.17 \times 1.5$	Right hilar lymph node mass

Table 3: Characteristics of images containing lung tumors. Resolutions are given in millimeters.

Tumor	Dice Coefficient (%)	Volume Overlap (%)	Volume Difference (%)	Average Surface Distance (mm)	RMS Surface Distance (mm)	Maximum Surface Distance (mm)
T1	90.97	83.45	7.39	0.86	0.92	4.42
T2	80.95	67.99	4.98	1.25	1.54	6.63
T3	72.95	57.42	15.76	1.26	1.50	6.87
T4	71.33	55.44	42.31	3.30	4.01	14.34
T5	80.53	67.41	29.22	3.63	4.55	16.56
T6	86.63	76.42	18.02	1.30	1.49	5.90
T7	82.49	70.21	22.28	1.34	1.56	5.16
T8	89.25	80.59	9.59	1.20	1.47	9.32
T9	72.66	57.07	34.17	1.75	2.09	7.36
T10	74.04	58.79	41.09	4.97	5.55	15.99
Average	80.18	67.47	22.48	2.08	2.46	9.25

Table 4: Comparison between our method and the segmentations provided by the expert.



Fig. 5: Overall context of lung tumors T1 (left), T8 (middle) and T9 (right).

Tumor	Standard graph cuts		Our method		Amount of object seeds (%)
	Time	Memory	Time	Memory	
T1	4.08	472.34	2.00	24.71	2.52
T2	4.89	573.05	2.71	83.42	2.71
T3	4.90	580.78	2.87	83.42	2.47
T4	5.34	729.72	2.14	37.07	12.99
T5	5.36	737.41	2.10	36.37	10.22
T6	10.18	1476.30	3.80	37.07	3.15
T7	5.16	544.74	3.21	83.42	8.86
T8	MP	43091.25	68.93	83.42	2.45
T9	4.24	496.79	2.39	37.07	8.01
T10	10.36	1151.74	5.42	125.13	9.19

Table 5: Speed (secs) an memory usage (Mb) for our method and the graph cuts without reduction. The label MP means there is not enough memory for allocating the graph.




Fig. 6: Seeds location for segmenting lung tumors T1 (top row), T8 (middle row) and T9 (bottom row). Object seeds (cyan) and background seeds (red) are superimposed on successive slices of the original image.

References

1. Y. Boykov and M-P. Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images. In *ICCV*, volume 1, pages 105–112, 2001.
2. Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on PAMI*, 26(9):1124–1137, 2004.




Fig. 7: Segmentation of lung tumors T1 (top row), T8 (middle row) and T9 (bottom row). Ground truth (red) and segmentation (cyan) are superimposed on the original image. Yellow color corresponds to the intersection.

3. A. J. Buckler, J. L. Mulshine, R. Gottlieb, B. Zhao, P. D. Mozley, and L. Schwartz. The use of volumetric CT as an imaging biomarker in lung cancer. *Academic Radiology*, 17(1):100–106, January 2010.
4. P. F. Felzenszwalb and D. P. Huttenlocher. Distance transforms of sampled functions. Technical report, Cornell Computing and Information Science, 2004.
5. M. A. Gavrielides, L. M. Kinnard, K. J. Myers, and N. Petrick. Noncalcified lung nodules: Volumetric assessment with thoracic CT. *Radiology*, 251(1):26–37, April 2009.
6. L. Ikonen. Pixel queue algorithm for geodesic distance transforms. In *Discrete Geometry for Computer Imagery*, volume 3429, pages 228–239, April 2005.
7. P. Kohli, V. Lempitsky, and C. Rother. Uncertainty driven multi-scale energy minimization. Technical report, april 2010.
8. V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Transactions on PAMI*, 26(2):147–159, 2004.
9. W. J. Kostis, A. P. Reeves, D. F. Yankelevitz, and C. I. Henschke. Three-dimensional segmentation and growth-rate estimation of small pulmonary nodules in helical CT images. *IEEE Transactions on Medical Imaging*, 22(10):1259–1274, October 2003.
10. J-M. Kuhnigk, V. Dicken, L. Bornemann, A. Bakai, D. Wormanns, S. Krass, and H-O. Peitgen. Morphological segmentation and partial volume analysis for volumetry of solid pulmonary lesions in thoracic CT scans. *IEEE Transactions on Medical Imaging*, 25(4):417–434, April 2006.
11. V. Lempitsky and Y. Boykov. Global optimization for shape fitting. In *IEEE, Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2007.

12. N. Lermé, F. Malgouyres, and L. Létocart. Reducing graphs in graph cut segmentation. In *To appear in ICIP*, 2010.
13. Y. Li, J. Sun, CK. Tang, and HY. Shum. Lazy Snapping. *ACM Transactions on Graphics*, 23(3):303–308, 2004.
14. H. Lombaert, Y.Y. Sun, L. Grady, and C.Y. Xu. A multilevel banded graph cuts method for fast image segmentation. In *ICCV*, volume 1, pages 259–265, 2005.
15. K. Marten, F. Auer, S. Schmidt, E. J. Rummeny, and C. Engelke. Automated CT volumetry of pulmonary metastases: The effect of a reduced growth threshold and target lesion number on the reliability of therapy response assessment using RECIST criteria. *European Radiology*, 17(10):2561–2571, October 2007.
16. M. F. McNitt-Gray, L. M. Bidaut, S. G. Armato, C. R. Meyer, M. A. Gavrielides, C. Fenimore, G. McLennan, N. Petrick, B. Zhao, A. P. Reeves, R. Beichel, H-J. Kim, and L. Kinnard. Computed tomography assessment of response to therapy: Tumor volume change measurement, truth data, and error. *Translational Oncology*, 2(4):216–222, December 2009.
17. J. H. Moltz, Bornemann L., Kuhnigk J-M., Dicken V., Peitgen E., S. Meier, H. Bolte, M. Fabel, H. C. Bauknecht, M. Hittinger, A. Kiessling, M. Pusken, and H. O. Peitgen. Advanced segmentation techniques for lung nodules, liver metastases, and enlarged lymph nodes in CT scans. *Selected Topics in Signal Processing*, Volume: 3 , Issue: 1:122–134, 2009.
18. L. E. Quint, J. Cheng, M. Schipper, A. C. Chang, and G. Kalemkerian. Lung lesion doubling times: Values and variability based on method of volume determination. *Clinical Radiology*, 63(1):41–48, January 2008.
19. F. Ranchin, A. Chambolle, and F. Dibos. Total variation and graph cuts approaches for variational segmentation. In *Proceedings of SSVM*, pages 743–753, June 2007.
20. A.K. Sinop and L. Grady. Accurate banded graph cut segmentation of thin structures using laplacian pyramids. In *MICCAI*, volume 9, pages 896–903, 2006.
21. I. Sluimer, A. Schilham, M. Prokop, and B. van Ginneken. Computer analysis of computed tomography scans of the lung: a survey. *IEEE Transactions on Medical Imaging*, 25(4):385–405, April 2006.
22. C. Suzuki, M. R. Torkzad, H. Jacobsson, G. Aström, A. Sundin, T. Hatschek, H. Fujii, and L. Blomqvist. Interobserver and intraobserver variability in the response evaluation of cancer therapy according to RECIST and WHO-criteria. *Acta Oncologica*, 49(4):509–514, May 2010.
23. R. J. van Klaveren, M. Oudkerk, M. Prokop, E. T. Scholten, K. Nackaerts, R. Vernhout, C. A. van Iersel, K. A. van den Bergh, S. van 't Westeinde, C. van der Aalst, E. Thunnissen, D. M. Xu, Y. Wang, Zhao Y., H. A. Gietema, B. J. de Hoop, H. J. Groen, G. H. de Bock, P. van Ooijen, C. Weenink, J. Verschakelen, J. W. Lammers, W. Timens, D. Willebrand, A. Vink, W. Mali, and H. J. de Konin. Management of lung nodules detected by volume CT scanning. *The New England Journal of Medicine*, 361(23):2221–2229, December 2009.
24. E. L. van Persijn van Meerten, H. Gelderblom, and J. L. Bloem. RECIST revised: Implications for the radiologist. a review article on the modified RECIST guideline. *European Radiology*, 20(6):1456–1467, June 2010.
25. X. Ye, G. Beddoe, and G. Slabaugh. Graph cut-based automatic segmentation of lung nodules using shape, intensity, and spatial features. In *Workshop on Pulmonary Image Analysis, MICCAI*, 2009.
26. B. Zhao, L. H. Schwartz, C. S. Moskowitz, M. S. Ginsberg, N. A. Rizvi, and M. G. Kris. Lung cancer: Computerized quantification of tumor response—initial results. *Radiology*, 241(3):892–898, December 2006.