



Fast and Memory Efficient Segmentation of Lung Tumors Using Graph Cuts

Nicolas Lermé, François Malgouyres, Jean-Marie Rocchisani

► To cite this version:

Nicolas Lermé, François Malgouyres, Jean-Marie Rocchisani. Fast and Memory Efficient Segmentation of Lung Tumors Using Graph Cuts. 2010. hal-00495260

HAL Id: hal-00495260

<https://hal.science/hal-00495260>

Preprint submitted on 25 Jun 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fast and Memory Efficient Segmentation of Lung Tumors Using Graph Cuts

Nicolas Lermé^{1,2}, François Malgouyres¹, and Jean-Marie Rocchisani³

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

(4) Hopital 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 stay a quite challenging task when touching directly 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 reduce resource consumption when segmenting large volume 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 get 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, for the 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 [23], and suffer of a reproducibility [22]. Alternatively, tumor volumetry has been proposed to overcome those difficulties that would improve the staging of nodules [6], the evaluation of tumor aggressiveness [19,7], tumor response to chemotherapy [4,18,25] or to radiotherapy [3,17] and progression rate of tumors [19,14] or metastases [16].

Among semi-automatic approaches 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 as N-D 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 [24]. 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, based on gray level in a CT image to distinguish the tumor and the healthy tissue. 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. We propose a semi-interactive graph cut-based method for lung tumors segmentation. An overview of the approach appears on Figure 1. First, we compute a distance map from the object seeds




Fig. 1: Flow diagram of our approach.

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, our algorithm is validated experimentally 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 is common, in (1) the region term $E_p(\cdot)$ 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 $E_{p,q}(\cdot)$ penalize neighboring pixels p and q having different labels. The weight of the penalization depend on the difference $|I(p) - I(q)|$ and favor boundaries located at pixels/voxels with a strong gradient. Generally speaking, the definition of E_p and $E_{p,q}$ depend on the considered application.

According to [10], 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. 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 get round this problem, some authors have recently proposed heuristics [13,5,15,21,9]. Nevertheless, these algorithms can easily get trapped in local minima of the energy and sometimes fail to recover details. In medical imaging, this is a real drawback since such thin structures like blood vessels or nodules are ubiquitous. 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] al-

locates $24|\mathcal{P}| + 14|\mathcal{E}_n|$ 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 [10], 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 $p \in \mathcal{P}$ by p : $\tilde{B}_p = \{b + p \mid b \in B\}$. For $Z \subset \mathcal{P}$, we also denote $\tilde{Z}_B = \bigcup_{p \in Z} \tilde{B}_p$.

The intuitive idea to build \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 $\tilde{Z}_B \setminus Z$ suffice 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 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]:

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

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 $\tilde{Z}_B \setminus Z$.

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 out) through $\tilde{B}_p \setminus \{p\}$ suffices to saturate the n -links going out (respectively in) of \tilde{B}_p . The pixel/voxel p is not needed in the graph. 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 model's parameters. 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 dimensions d . Decomposing the condition 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 (see [12] for details). Finally, we have both theoretical and empirical evidence suggesting that this reduction scheme provides an exact solution (see [12]).

4 Energy function

The most famous energy model for image segmentation with graph cut is the model proposed by Boykov and Jolly in [1] (see below). Total Variation-based




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.

models have also been proposed (see [20]). 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 tissue and the corresponding zone of the image has a uniform color.

To get round this problem, 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 seed. This leads to a modification of the original Boykov/Jolly's energy model [1]. We take the same regularity criterion:

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

where I is the original image and $\sigma > 0$. The region term is defined as

$$\begin{cases} E_p(u_p = \text{"bkg"}) = -\log [Pr(I(p) \mid p \in \mathcal{O}) \times \exp\left(-\left(\frac{d(p, \mathcal{O})}{\sigma_a}\right)^2\right)] \\ E_p(u_p = \text{"obj"}) = -\log [Pr(I(p) \mid p \in \mathcal{B})], \end{cases} \quad (3)$$

where the sets \mathcal{O} and \mathcal{B} correspond respectively to object and background seeds provided by the user, the probability distributions are estimated according to [1], the function $d(p, \mathcal{O})$ is a distance 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 define an area of influence A_{σ_a} . Although it is an important parameter that impact the way the seed construction, we always take $\sigma_a = 10$ in our experiments.

The main difference when compared to the energy in [1] is the distance term. The distance d is according to $d(p, \mathcal{O}) = \min\{d(p, q) \mid q \in \mathcal{O}\}$, where d is a distance between two points ². We have made two attempts for the distance d :

² By an abuse of language we write d for both the distance between a point and a set and between two points.

- The euclidean distance. In this case, the distance between a set and a point is computed with the algorithm described in [8]. 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

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

In this latter case we compute the distance with the algorithm described in [8].

We display on Figure 4, the zone of influence of the object seed for the above two metrics. We draw the region where the exponential in (3) is greater than some $\epsilon \simeq 0$. Beyond this area, the nodes are only linked to the sink t , ensuring that the algorithm always categorizes them as background pixels/voxels. Notice that the geodesic distance better fit the the tumor boundaries. In particular, it only has a limited overflow on the healthy tissue.




Fig. 4: Area of influence of the object seed for the Euclidean and the geodesic distances. In this experience, we set $\sigma_a = 40$.

5 Experimental results

In this section, we present some experiments for segmenting a set of 10 3D CT images of size $512 \times 512 \times 50$ consisting both of nodules/masses and tumors (see Table 2 for image characteristics). All experiments presented are performed in connectivity 1. Since the images haven't been acquired in the same way, the contrast is different. Thus, we are constrained to use different values of σ for some images. Here, we use $\sigma = 0.2$ for all images except for T8 where $\sigma = 2$ and for T7 where $\sigma = 0.05$. Note that we extract automatically a sub-volume by considering an extra border of 100 pixels around the object seeds for speeding up the segmentation.

First, we evaluate our algorithm for segmentation algorithm against hand made segmentations provided by an expert, for all CT images. Table 3 contains statistics on the differences between our segmentation and the ground truth. We

use several evaluation measures ³. Table 3 shows encouraging results. For all images, we get a Dice Coefficient always greater than 70% while having a mean maximum distance of about 10 millimeters between the ground truth and the segmentation.

Moreover, 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 z. While the segmentation of T1 is very close to the ground truth, the segmentations obtained for T8 and T9 also illustrate the difficulty of extracting tumors/nodules with a large connection to healthy tissues.

Secondly, we compare the performance of standard graph cuts against our method in terms of speed and memory consumption (see Table 4) for segmenting the CT images using the same set of seeds and parameters as in the previous experiments. Experiments were performed on an Athlon Dual Core 6000+ 3GHz with 2GB RAM. Times are averaged over 10 runs. Table 4 shows that our method performs a little bit faster using 6 to 20x less memory while getting exactly the same solution. We also indicate the proportion of object seeds with respect to the tumor volume in Table 4. This provides an objective measure of the interaction for assessing the effort required by the user for positioning the seeds. Observe that a relatively small amount of seeds is necessary for all images.

Generally, the segmentation time depends on the image size and the skill of the user placing the seeds. The computation of the distance map, the building of the graph and the computation of the minimum-cut takes 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.




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

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.

³ A detailed view of these measures is available at <http://lts08.bigr.nl/about.php>

Tumor	Type	Resolutions (x,y,z)	Description
T1	Nodule	$0.683 \times 0.683 \times 3$	Mass of the upper right lobe (CT)
T2	Nodule	$0.707 \times 0.707 \times 1$	Nodule of the right apex (CT)
T3	Nodule	$0.683 \times 0.683 \times 3$	Nodule of the lower right lobe (CT)
T4	Tumor	$1.171 \times 1.171 \times 1.5$	Marge left hilar tumor inducing a peripheral atelectasia (CT)
T5	Tumor	$1.171 \times 1.171 \times 1.5$	Same as T4 (dosimetric CT scanner)
T6	Nodule	$0.771 \times 0.771 \times 1.25$	Mass of the lower left lobe appended to the pleura (CT)
T7	Nodule	$0.699 \times 0.699 \times 1.25$	Same as T6, after four months of treatment (CT)
T8	Tumor	$0.636 \times 0.636 \times 1$	Large left hilar tumor and peripheral atelectasia, before treatment (contrast enhanced CT)
T9	Tumor	$0.707 \times 0.707 \times 1$	Same as T8, after chemo-radiotherapy (CE-CT)
T10	Nodule	$1.171 \times 1.171 \times 1.5$	Right hilar lymph node mass

Table 2: 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	90.24	82.22	2.57	1.11	1.48	7.39
T2	81.24	68.41	7.59	1.19	1.46	6.63
T3	72.25	56.56	16.99	1.26	1.49	5.86
T4	71.33	55.44	42.31	3.28	3.98	14.34
T5	80.53	67.40	29.23	3.61	4.53	16.56
T6	87.40	77.62	17.85	1.31	1.52	5.90
T7	80.01	66.69	28.72	1.48	1.75	6.29
T8	89.25	80.60	9.60	1.20	1.46	9.32
T9	72.66	57.07	34.17	1.74	2.08	7.36
T10	74.04	58.79	41.09	4.95	5.54	15.99
Average	79.89	67.08	23.01	2.11	2.52	9.56

Table 3: Comparison between our method and the segmentations provided by the expert.

Tumor	Standard graph cuts		Our method		Object seeds (%)
	Time	Memory	Time	Memory	
T1	5.01	472.34	3.23	24.71	2.46
T2	6.07	573.07	4.31	83.42	2.73
T3	6.01	580.78	4.25	83.42	2.51
T4	8.07	729.72	5.39	37.07	12.99
T5	7.98	737.41	5.38	36.37	10.22
T6	5.12	489.97	3.47	37.07	3.19
T7	6.35	544.74	4.48	83.42	9.66
T8	11.65	793.95	9.10	125.13	2.45
T9	4.94	496.79	3.46	37.07	8.01
T10	14.62	1150.97	10.46	125.13	9.01

Table 4: Speed (secs) an memory usage (Mb) for our method and the graph cuts without reduction.




Fig. 6: Seeds placement for segmenting lung tumors T1 (top), T8 (middle) and T9 (bottom). Object seeds (blue) and background seeds (red) are superimposed on the original image.




Fig. 7: Segmentation of lung tumors T1 (top), T8 (middle) and T9 (bottom). Ground truth is superimposed in red while segmentation is superimposed in blue. Purple color corresponds to the intersection.

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.
3. Samuel Bral, Michaël Duchateau, Mark De Ridder, Hendrik Everaert, Koen Tournel, Denis Schallier, Dirk Verellen, and Guy Storme. Volumetric response analysis during chemoradiation as predictive tool for optimizing treatment strategy in locally advanced unresectable nscl. *Radiother Oncol*, 91(3):438–442, June 2009.
4. Andrew J Buckler, James L Mulshine, Ronald Gottlieb, Binsheng Zhao, P. David Mozley, and Lawrence Schwartz. The use of volumetric ct as an imaging biomarker in lung cancer. *Acad Radiol*, 17(1):100–106, Jan 2010.
5. C. Cigla and A.A. Alatan. Region-based image segmentation via graph cuts. In *ICIP*, pages 2272–2275, 2008.
6. Marios A Gavrielides, Lisa M Kinnard, Kyle J Myers, and Nicholas Petrick. Noncalcified lung nodules: volumetric assessment with thoracic ct. *Radiology*, 251(1):26–37, Apr 2009.
7. Osamu Honda, Takeshi Johkoh, Junko Sekiguchi, Noriyuki Tomiyama, Naoki Mihara, Hiromitsu Sumikawa, Atsuo Inoue, Masahiro Yanagawa, Tadahisa Daimon, Meinoshin Okumura, and Hironobu Nakamura. Doubling time of lung cancer determined using three-dimensional volumetric software: comparison of squamous cell carcinoma and adenocarcinoma. *Lung Cancer*, 66(2):211–217, Nov 2009.
8. L. Ikonen. Pixel queue algorithm for geodesic distance transforms. pages 228–239, April 2005.
9. P. Kohli, V. Lempitsky, and C. Rother. Uncertainty driven multi-scale energy minimization. Technical report, april 2010.

10. V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Transactions on PAMI*, 26(2):147–159, 2004.
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. Yin. Li, Jian. Sun, Chi-Keung. Tang, and Heung-Yeung. Shum. Lazy snapping. *ACM Transactions on Graphics*, 23(3):303–308, 2004.
14. Rebecca M Lindell, Thomas E Hartman, Stephen J Swensen, James R Jett, David E Midthun, and Jayawant N Mandrekar. 5-year lung cancer screening experience: growth curves of 18 lung cancers compared to histologic type, ct attenuation, stage, survival, and size. *Chest*, 136(6):1586–1595, Dec 2009.
15. 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.
16. Katharina Marten, Florian Auer, Stefan Schmidt, Ernst J Rummeny, and Christoph 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. *Eur Radiol*, 17(10):2561–2571, Oct 2007.
17. Michael F McNitt-Gray, Luc M Bidaut, Samuel G Armato, Charles R Meyer, Marios A Gavrielides, Charles Fenimore, Geoffrey McLennan, Nicholas Petrick, Binsheng Zhao, Anthony P Reeves, Reinhard Beichel, Hyun-Jung Grace Kim, and Lisa Kinnard. Computed tomography assessment of response to therapy: tumor volume change measurement, truth data, and error. *Transl Oncol*, 2(4):216–222, Dec 2009.
18. P. D. Mozley, L. H. Schwartz, C. Bendtsen, B. Zhao, N. Petrick, and A. J. Buckler. Change in lung tumor volume as a biomarker of treatment response: a critical review of the evidence. *Ann Oncol*, Mar 2010.
19. Leslie Eisenbud Quint, Joan Cheng, Matthew Schipper, Andrew C Chang, and Gregory Kalemkerian. Lung lesion doubling times: values and variability based on method of volume determination. *Clin Radiol*, 63(1):41–48, Jan 2008.
20. 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.
21. 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.
22. Chikako Suzuki, Michael R Torkzad, Hans Jacobsson, Gunnar Aström, Anders Sundin, Thomas Hatschek, Hirofumi Fujii, and Lennart Blomqvist. Interobserver and intraobserver variability in the response evaluation of cancer therapy according to recist and who-criteria. *Acta Oncol*, 49(4):509–514, May 2010.
23. Els L van Persijn van Meerten, Hans Gelderblom, and Johan L Bloem. Recist revised: implications for the radiologist. a review article on the modified recist guideline. *Eur Radiol*, 20(6):1456–1467, Jun 2010.
24. 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.
25. Binsheng Zhao, Lawrence H Schwartz, Chaya S Moskowitz, Michelle S Ginsberg, Naiyer A Rizvi, and Mark G Kris. Lung cancer: computerized quantification of tumor response—initial results. *Radiology*, 241(3):892–898, Dec 2006.