Segmentation of Liver Metastases in CT Scans by Adaptive Thresholding and Morphological Processing

Release 0.00

Jan Hendrik Moltz, Lars Bornemann, Volker Dicken, and Heinz-Otto Peitgen

July 7, 2008

MeVis Research GmbH – Center for Medical Image Computing, Bremen, Germany

Abstract

This article presents an algorithm for the segmentation of liver metastases in CT scans. It is a hybrid method that combines adaptive thresholding based on a gray value analysis of the ROI with model-based morphological processing. We show the results of the MICCAI liver tumor segmentation competition 2008 where successful results were obtained for all ten lesions.

Contents

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 diameters 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 liver metastases in CT scans. In contrast to other lesion types such as lung nodules, they are very diverse in their appearance. The gray values of parenchyma and metastases in the liver are not fixed but depend on primary cancer, contrast agent, contrast timing, scan parameters, and patient conditions. Lesions can be brighter (hyperdense) or darker (hypodense) than the surrounding tissue or they can appear inhomogeneous if they have a contrast-enhanced rim around them or if they are partly necrotic or calcified. Primary liver tumors, so called hepatocellular carcinomas, are not targeted by our algorithm.

In the following, all computations are restricted to a region of interest (ROI) whose center and size is determined by a user-defined stroke that is drawn across the tumor. 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

The segmentation of liver metastases, unlike that of lung nodules, has not been an area of intensive research so far. Few algorithms have been designed for 3d CT data. Lu *et al.* [7] use 2d active contours that are initialized manually. This is relatively slow and not very accurate. Other authors focus on the detection of metastases in a segmented liver and apply rather simple methods for lesion segmentation. Park *et al.* [8] and Ciecholewski and Ogiela [2] both try to find optimal thresholds based on statistical analysis of the histogram under the liver mask, but this will fail on poorly contrasted or inhomogeneous lesions. Li *et al.* [6] apply a machine learning technique to classify possible boundary positions on 1d intensity profiles. Again, poor contrasts and especially lesions with contact to a structure of the same density will pose problems.

One of the most promising methods was proposed by Li and Jolly [5]. They use a graph-theoretic approach, representing voxels by vertices and surfaces by paths along the edges. The optimal path minimizes an objective function that incorporates boundary, regional and elasticity constraints. Notably, the algorithm is able to detect multiple surfaces simultaneously, making it possible to segment tumors with necroses and calcifications in one pass. Its rather abstract nature, however, makes it difficult to integrate an intuitive interactive correction procedure.

Another interesting approach was presented by Jolly and Grady [3]. Their algorithm is suitable for different kinds of lesion including liver metastases. It learns the gray value distribution of a lesion from user interaction and several 2d segmentations on orthogonal planes and finally applies the random walker algorithm. Due to its general character, the method lacks a handling of typical difficult cases such as liver metastases adjoining an isodense structure.

A rough description and a preliminary evaluation of our method was given in [1].

3 Segmentation Method

3.1 Adaptive Thresholding

Our method is based on the lung nodule segmentation algorithm by Kuhnigk *et al.* [4]. The method proved useful for other tumor types as well because it tackles a problem that is common to all of them: the contact to vessels or other thin, elongated structures that feature a similar density in CT scans. It is obvious that a mere threshold-based method is not sufficient for segmentation. Since the lesions are mostly homogeneous it can, however, be used as a first step to obtain a superset of voxels that may be part of the lesion. This can be implemented efficiently as a 3d region growing starting from the center of the ROI.

Due to the high diversity in their appearance, the thresholds for liver metastases have to be determined adaptively. This is based on an analysis of the density distribution in the ROI. The size and center of the ROI are determined from the stroke that the user draws across the lesion. This stroke also allows the detection of inhomogeneous metastases because the density profile under the stroke provides information about the "relative density" of the lesion compared to the parenchyma.

As a first step, a typical parenchyma value is estimated which allows us to decide whether the lesion is hypodense, hyperdense or inhomogeneous. Since the size of the ROI is chosen as a multiple of the stroke length, we can assume that the largest part of the ROI contains healthy liver parenchyma and that the maximum peak of the ROI histogram represents a typical parenchyma density P. There are two cases where this can lead to a wrong conclusion. First, for metastases close to the liver surface, the ROI contains not only liver tissue and the maximum peak may correspond to a structure that is actually outside the liver. To prevent this, the histogram is restricted to voxels with at least 40 HU, so as to exclude less dense materials such as fluids or air. Second, if several metastases lie close to each other, they can contribute to the maximum peak rather than the parenchyma. This can easily be detected and is handled in the following.

Next, a histogram of the values under the stroke is computed. The stroke is dilated with a $3 \times 3 \times 3$ kernel in order to get a larger amount of representative values. From the histogram, we compute the ratio of voxels covered by the dilated stroke that have a density less than P. This ratio is a measure of the "hypodensity" H of the lesion compared to the healthy parenchyma. A value close to 100% indicates that almost all voxels under the stroke are below P and that the lesion is hypodense. If H is near 0%, we have a hyperdense lesion. We compute a Gaussian-weighted mean value in a small area around the center of the ROI and store it as the typical lesion value L. Some lesions have hyperdense and hypodense parts which is reflected by H having a value between 20% and 80%. They require a special handling that is not discussed here.

Given P and L, the thresholds T_0 and T_1 for region growing are determined as follows (all values are given in HU):

- Hypodense lesions: $T_0 = 10$, $T_1 = (P+L)/2$.
- Hyperdense lesions: $T_0 = (P+L)/2$, $T_1 = 180$.

The "inner" thresholds are set to the average of P and L, the "outer" thresholds to empirical values in order to remove materials that can be excluded a priori such as air or bones.

If L and P have a distance of less than 25 HU, we suspect that P is not representative of the parenchyma because it is not the largest region in the ROI. In this case, we try to find a better estimate than P by searching for a quantile of the ROI histogram that has a reasonable distance to L.

- Ambiguous cases with $L < Q_{50}$: $T_0 = L - 25$,

$$T_1 = \begin{cases} (L+Q_{25})/2 & \text{if } L < Q_{25} - 25 \\ (L+Q_{50})/2 & \text{if } L < Q_{50} - 25 \\ (Q_{50}+Q_{75})/2 & \text{otherwise} \end{cases}$$

- Ambiguous cases with $L > Q_{50}$: $T_1 = L + 25$,

$$T_0 = \begin{cases} (L+Q_{75})/2 & \text{if } L < Q_{75} + 25 \\ (L+Q_{50})/2 & \text{if } L < Q_{50} + 25 \\ (Q_{25}+Q_{50})/2 & \text{otherwise} \end{cases}$$

where Q_n denotes the $n\%$ quantile of the restricted ROI histogram and all values are given in HU.

Finally, we enlarge the threshold range so that most of the values under the stroke are covered but a safety distance to P is preserved:

$$T_0 = \min(T_0, \max(S_5, P - 10)),$$
$$T_1 = \max(T_1, \min(S_{95}, P + 10))$$

where S_n is the $n\%$ quantile of the stroke histogram. For ambiguous cases, P is replaced by Q_{50}.

3.2 Morphological Postprocessing

In a typical case, the region growing result contains the complete lesion and additionally parts of the attached vasculature (Fig. 1). A morphological opening operation is an obvious choice to remove the vessels, but the challenge lies in determining the optimal erosion strength. Since tumors and their supplying vessels can differ in size significantly, an erosion with a fixed-size kernel cannot be used since this would either maintain thick vessels that should be removed or erode the lesion too strongly so that details of its boundary would be lost.

A solution to this problem was already presented in Kuhnigk's algorithm for lung nodules. It was described elaborately in [4] and will be reviewed briefly in the following. The basic idea is to choose the erosion strength adaptively. To facilitate the computations, erosion is implemented by thresholding on a distance map that contains the distance of each mask voxel to the closest background voxel. The underlying assumption is that all vessels are connected to the boundary of the ROI and that the diameter of a vessel decreases monotonically in its course. In order to disconnect the mask, all paths from the ROI center to the boundary are considered. The maximum of all minimum path diameters is the cut-off value that removes all vessels (Fig. 1).

The second step of the opening operation is a dilation that reconstructs the lesion in its original size without regrowing the vessels. The dilation is again implemented by thresholding on a distance map, but this time it shows the distances of all background voxels to the eroded mask. The threshold is chosen slightly higher than the erosion threshold so that in a final refinement step all boundary details can be reproduced by intersection with the region growing result.

The algorithm contains an efficient and intuitive mechanism for interactive correction of segmentation results. The user can change the erosion strength and thereby slightly alter the overall shape of the mask in cases where the underlying anatomical assumptions are not fulfilled or where the result is not optimal due to imaging artifacts. The possible values for the actual erosion threshold are mapped to the range from 0% to 100%. After a change, the segmentation is partly recomputed and the result is available almost immediately. This concept improves the flexibility of the algorithm while preserving reproducibility.

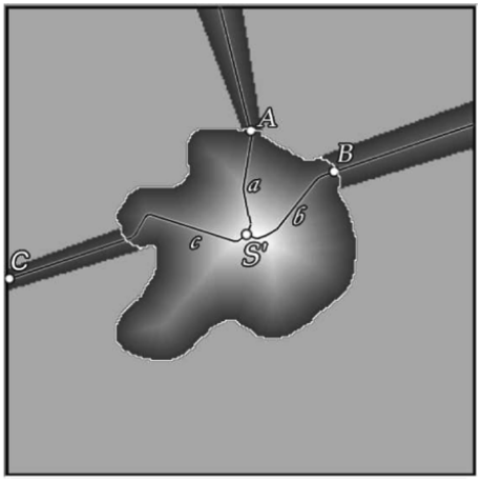

Figure 1: Finding the optimal erosion strength. The image shows the distance map of the initial region growing result, containing a tumor and adjacent vessels. S' is the maximum of the distance map; a, b, and c are paths from S' to the ROI boundary along "distance ridges"; A, B, and C are the minima on these paths, corresponding to the points with minimum vessel diameter. The maximum distance value of these points (in this case B) is the cut-off value for optimal erosion.

4 Results and Discussion

Our algorithms were evaluated on ten tumors from five different datasets in the course of the MICCAI Liver Tumor Segmentation Challenge 2008. Fig. 2 shows the computed segmentation masks for all tumors. The evaluation results are summarized in Table 1. It can be seen that our algorithm was successful in all cases: both hyperdense and hypodense lesions could be segmented, independent of their location inside the liver or in its periphery. It should be noted that in our opinion tumor 5 is rim-enhancing and the bright area is a part of the tumor but we adapted the stroke according to the reference segmentation.

From the point of view of our algorithm, tumors 6 and 8 were most difficult. Their inhomogeneous density distribution made it hard to determine a threshold range that covered the tumor as exactly as possible. In these two cases, we tried several different initialization strokes to produce the final results. A goal for future work is to reduce the dependency from the user interaction.

Typical problem classes for segmentation are peripheral and rim-enhancing liver metastases. The former can have the same density as an adjacent structure outside the liver such as the intercostal muscles or the parenchyma of other abdominal organs such as the kidneys or the stomach. The latter are hyperdense and have a hypodense rim around them so that a single thresholding interval does not suffice. We are developing special solutions for these cases but since they were not necessary for the competition tumors we omit the details here.

We performed the computations on a standard PC with a 2 GHz dual core processor and needed around 2s for each tumor. In three cases (tumors 4, 5, 7) the result was refined with the interactive correction was applied which takes another second since a part of the segmentation is recomputed. Our discussions with radiologists confirmed that this is an acceptable runtime for the clinical routine.

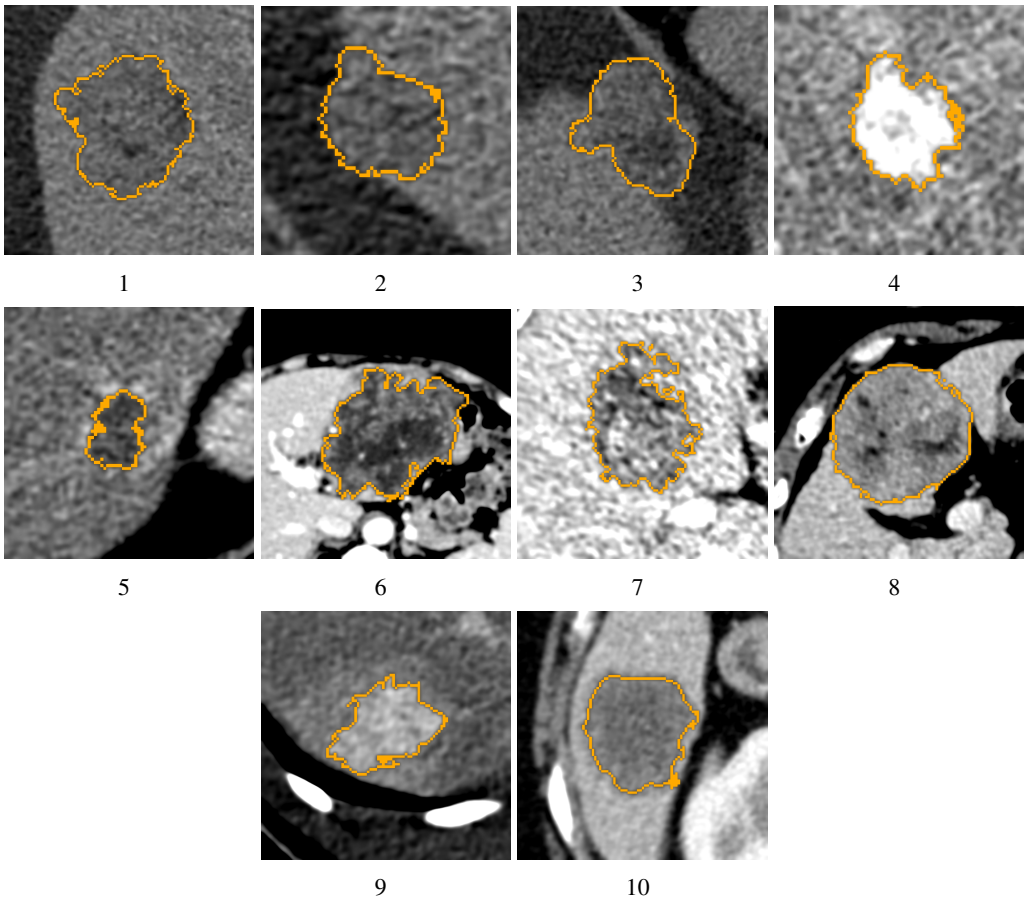

Figure 2: Segmentation results for the ten tumors of the MICCAI competition. A representative axial slice is shown for each lesion.

Tumor	Overlap Error		Volume Diff.		Avg. Dist.		RMS Dist.		Max. Dist.		Total
	%	Score	%	Score	%	Score	%	Score	%	Score	Score
1	24.79	81	7.68	92	1.72	56	2.27	68	8.00	80	76
2	37.44	71	41.65	57	1.42	64	2.06	71	8.76	78	68
3	38.36	70	53.28	45	1.77	55	2.72	62	9.45	76	62
4	46.64	64	67.67	30	1.51	62	2.02	72	6.60	83	62
5	30.32	77	3.34	97	0.64	84	0.97	86	4.61	88	86
6	26.86	79	11.89	88	3.28	17	4.60	36	19.91	50	54
7	32.38	75	22.62	77	1.46	63	2.02	72	9.44	76	73
8	11.13	91	4.94	95	1.14	71	1.72	76	8.88	78	82
9	27.47	79	10.95	89	0.91	77	1.46	80	7.24	82	81
10	30.14	77	29.13	70	1.67	58	2.14	70	8.39	79	71
Average	30.55	76	25.32	74	1.55	61	2.20	69	9.13	77	72

Table 1: Evaluation results for the ten tumors of the MICCAI competition. The measures are volumetric overlap error, relative absolute volume difference, average symmetric absolute surface distance, symmetric RMS surface distance, and maximum symmetric absolute surface distance. A score of 100 was given for the reference segmentation, a score of 90 corresponds to the deviation of a second manual segmentation.

Acknowledgement. This work was supported by research grants from the German Federal Ministry of Education and Research (grant number 01EZ0401) as part of the cooperation project VICORA – Virtual Institute for Computer Assistance in Clinical Radiology – and from Siemens Healthcare, Computed Tomography, Forchheim, Germany. We thank Holger Bourquain and Benjamin Geisler for their radiological advice.

References

[1] Lars Bornemann, Volker Dicken, Jan-Martin Kuhnigk, Dag Wormanns, Hoen-Oh Shin, Hans-Christian Bauknecht, Volker Diehl, Michael Fabel, Stefan Meier, Oliver Kress, Stefan Krass, and Heinz-Otto Peitgen. OncoTREAT: a software assistant for cancer therapy monitoring. *Int. J. CARS*, 1(5):231–242, 2007. 2

[2] Marcin Ciecholewski and Marek R. Ogiela. Automatic segmentation of single and multiple neoplastic hepatic lesions in CT images. In *Proc. IWINAC*, pages 63–71, 2007. 2

[3] Marie-Pierre Jolly and Leo Grady. 3d general lesion segmentation in CT. In *Proc. ISBI*, pages 796–799, 2008. 2

[4] Jan-Martin Kuhnigk, Volker Dicken, Lars Bornemann, Annemarie Bakai, Dag Wormanns, Stefan Krass, and Heinz-Otto Peitgen. Morphological segmentation and partial volume analysis for volumetry of solid pulmonary lesions in thoracic CT scans. *IEEE Trans. Med. Imag.*, 25(4):417–434, 2006. 3.1, 3.2

[5] Kang Li and Marie-Pierre Jolly. Simultaneous detection of multiple elastic surfaces with application to tumor segmentation in CT images. In *Proc. SPIE*, page in print, 2008. 2

[6] Yuanzhong Li, Shoji Hara, and Katsuo Shimura. A machine learning approach for locating boundaries of liver tumors in ct images. In *Proc. ICPR*, pages 400–403, 2006. 2

[7] Rui Lu, Pina Marziliano, and Choon Hua Thng. Liver tumor volume estimation by semi-automatic segmentation method. In *Proc. IEEE EMBS*, pages 3296–3299, 2005. 2

[8] Seung-Jin Park, Kyung-Sik Seo, and Jong-An Park. Automatic hepatic tumor segmentation using statistical optimal threshold. In *Proc. ICCS*, pages 934–940, 2005. 2