



HAL
open science

Supervoxel based 3D diseased lung segmentation

Leo Stepien, Younes Belkouchi, Jean-Luc Dupire, Nicolas Billet, Lama Dawi,
Felix Wirth, Nathalie Lassau, Hugues Talbot

► **To cite this version:**

Leo Stepien, Younes Belkouchi, Jean-Luc Dupire, Nicolas Billet, Lama Dawi, et al.. Supervoxel based 3D diseased lung segmentation. 2024. hal-04418158

HAL Id: hal-04418158

<https://hal.science/hal-04418158>

Preprint submitted on 25 Jan 2024

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.

Supervoxel based 3D diseased lung segmentation

Leo Stepien¹, Younes Belkouchi^{1,2}, Jean-Luc Dupire¹, Nicolas Billet³, Lama Dawi³, Felix Wirth², Nathalie Lassau^{2,3}, and Hugues Talbot¹

¹ CentraleSupélec, Université Paris-Saclay, Inria, 91190 Gif-Sur-Yvette

² BIOMAPS, Université Paris-Saclay, 94800 Villejuif

³ Gustave Roussy, 94800 Villejuif

Abstract. Computer-Aided Diagnosis relies on accurate tumor volume and heterogeneity assessment through CT-scans. Precise lesion segmentation is essential for patient diagnosis, therefore the development of automatic tools for lesion segmentation is needed. To improve lung nodule segmentation performance, lung segmentation masks serve as valuable priors, narrowing the focus to lung regions. Current methods suffer from the exclusion of pathological areas, especially in oncology patients, since tumor tissue differ in voxel density from other structures in the lung. Consequently, ensuring accurate lung segmentation encompassing all lesions is crucial. We developed a method based on supervoxels to fully segment the lung while encompassing nodules if present using a propagation algorithm based on geometrical properties. We compared our method to a morphology based method and neural networks trained to segment the lungs. Our method had the best performance in the inclusion of lung lesions, while retaining an adequate level of precision.

Keywords: Lung segmentation · Medical Imaging · Supervoxel · Mathematical morphology.

1 Introduction

Lung cancer, one of the deadliest forms of cancer, requires early detection to enhance the chances of successful treatment and recovery. This type of cancer has the capacity to evolve quite rapidly, hence growing in size fast. It often evolves into multiple small tumors covering large areas of the lungs, which makes disease quantification more difficult since evaluation methods rely on tumors measurements. Currently, radiologists use the RECIST criteria to evaluate patient's follow-up to treatment. This criterion relies on manual measurements of the diameters of maximum 5 lesions, 2 per metastatic site [2]. One of the main limitations of this tool is its selection process, which might underestimate the true tumor burden, especially in metastatic patients [11]. From a practical perspective, RECIST is an adequate tool to assess patient response to treatment, especially in routine practice where time is an issue. Manually delineating all tumors is a time-consuming task, but offers better prognosis for patient's response. For this purpose, segmentation algorithms are of interest, as they might provide more information and reduce radiologist annotation time, hence leading towards

more precise medicine. Modern diagnostic procedures often involve Computer-Aided Diagnosis (CAD) systems, which play a crucial role in assisting healthcare professionals. For lung cancer, these CAD systems perform tumor segmentation on lung images [4], necessitating an initial pre-processing step to segment the lungs within the image and reduce the search space.

Our focus lies specifically on lung segmentation in Computed Tomography (CT) scans, as the accuracy of CAD systems heavily relies on the accurate delineation of lung structures [1]. While numerous lung segmentation algorithms that rely on morphological methods or deep learning have been developed, they often encounter challenges when dealing with scans containing tumors, since they can be very similar to vessels or organs in the mediastinum and differ a lot from lung tissue. Failure to accurately segment the lungs in such cases prevents CAD systems from detecting these tumors. In response, we propose an efficient method for integrating lung nodules into the lung segmentation mask, which combines the strengths of morphological and graph-based image segmentation techniques.

2 Related Works

2.1 Lung segmentation

Over the years, various morphological techniques, including active contours and region growing, along with thresholding methods, have been employed for lung segmentation on CT-scans. However, these methods often struggle to deliver precise results when dealing with diseased lungs. Huidrom, R. et al. [4] have specifically addressed the issue of pulmonary nodules that are closely connected to the chest wall and pleural surface. Traditional morphological methods tend to overlook such nodules, resulting in incomplete lung segmentation. In response, Shanhui Sun et al. [9] proposed a robust active shape model approach. Nonetheless, they acknowledge limitations when dealing with 'out-of-the-norm' cases encountered in real-life scenarios, such as unique lung shapes or those afflicted by numerous pathological regions. Deep learning methods, such as the U-Net model, have been explored by researchers like Bizopoulos, P. et al. [1] for lung segmentation. However, it's important to note that even these advanced techniques face significant challenges in accurately segmenting the lungs, particularly in difficult cases. Moreover, they require the availability of manual annotations for training as is the case for the Scancovia model [6] or TotalSegmentator [10]. Scancovia's design for Covid-19 severity prediction on CT-scans includes a 3D lung segmentation U-Net model trained on an in-house dataset. TotalSegmentator, based on a nnU-net, is capable of segmenting 117 anatomical structures on CT-scans. While both achieve high performance in lung segmentation, they were trained on datasets from various origins, which may constitute a problem when dealing with lungs affected by specific pathologies such as cancer. Since we wish to develop a method of lung segmentation that does not require previous annotations, we will be focusing on methods based on mathematical morphology or unsupervised learning.

Özsavaş, E.E et al. [7] proposed a method for accurate lung segmentation and inclusion of pathological areas based on prior knowledge of CT-scan’s physics, which can be summed up as follows:

Step 1. Rough lung segmentation To initiate the morphological segmentation process, the method begins by identifying the body region within the 3D scan. This is achieved by isolating the largest high-density structure, between -175 and 750 Hounsfield Unit (HU) on each slice, removing all the surrounding objects such as the plate. Subsequently, a preliminary segmentation of the lung is established through thresholding by keeping values with density below -300 HU in the body region, allowing for the initial delineation of the lung structure.

Step 2. Trachea segmentation and lung mask refinement To exclude the trachea from the previously obtained mask, a process is employed to identify the air-filled regions within the image. These regions, characterized by very low density, are carefully evaluated using a gradually increasing threshold. This incremental thresholding approach is applied to prevent any overlapping with the lung structures, ensuring that the trachea is accurately delineated and removed from the mask. A hole-filling algorithm is applied to include the vessels, nodules, tumors or any other high density pathology that are inside the remaining lung fields.

Step 3. Intestines removal Since intestines may have a density greater than -300 HU, they ought to be removed. First, connectivity is checked to identify the objects within the rough lung segmentation. Then, small regions in the rough lung segmentation are removed if lower than 200mm^2 .

Step 4. Separation of the lungs The right and left lung are separated by a high density junction that divides both structures. The algorithm is described thoroughly in the original paper.

Step 5. Inclusion of pathological areas At this step, only the tumors that are in relation with the borders of the lungs were excluded from the lung mask. To remedy this problem, a three-step approach was established. Firstly, for each voxel coordinates (x, y, z) contained in the rough lung mask, we identify z_{\min} and z_{\max} for which (x, y, z) is included in the segment $(x, y, z_{\min}), (x, y, z_{\max})$. All voxels in the segment between (x, y, z_{\min}) and (x, y, z_{\max}) (along the z axis) are then labeled as lung on a new mask, that we denote as the *extended mask*. Secondly, a propagation algorithm is employed to expand the lung mask into regions belonging to the extended mask. This is achieved by examining the surrounding $7 \times 7 \times 7$ voxel neighborhood of each voxel in the border of the lung mask. If more than half of the voxels’ neighborhood are part of the existing extended mask, the voxel in question is added to the lung mask. This process

continues until no more qualifying voxels can be identified. Finally, the updated segmentation mask is refined by taking the intersection with the extended mask, resulting in a comprehensive and accurate representation of the lung structure, including the pathological areas.

2.2 Supervoxels

Supervoxels and superpixels algorithms group voxels/pixels into perceptually meaningful regions. The image information are now represented by these supervoxels/superpixels (instead of voxels/pixels), providing a convenient unit to process local features on an image. In the domain of medical imaging, their usage has been progressively growing, particularly in segmentation tasks, as they facilitate the formation of coherent regions. Notably, Hansen, S. et al. [3] have leveraged supervoxels and superpixels as a preliminary step in the segmentation process to enhance the accuracy of abdominal organ segmentation and cardiac segmentation. By employing these regions as sub-segmentation units, they have achieved improved results in these critical medical imaging applications. One notable approach within this framework is SLIC, or Simple Linear Iterative Clustering, introduced by Achanta, R. et al. [8]. The objective of SLIC is to clusterize pixels based on their color similarity and proximity in the image space. SLIC operates by first over-segmenting an image into compact, roughly equally-sized regions and then adjusts these regions to gather similar intensity regions together. The key features of SLIC include its simplicity and efficiency, making it well-suited for real-time or large-scale image processing tasks. It also allows users to control the compactness of the superpixels, providing a balance between more cubed shaped clusters, or more irregular supervoxels.

3 Method

We introduce a novel approach that capitalizes on the rich local information offered by supervoxel segmentation to enhance the segmentation of lungs in CT-scans. Drawing inspiration from Özsavaş, E.E et al.’s method, we adopt supervoxels as a superior alternative to scrutinizing individual voxel surroundings. This adaptation significantly improves efficiency, as we now analyze contiguous groups of voxels that exhibit similarities.

3.1 Supervoxel segmentation

To perform the supervoxel segmentation of CT-scans, we employ an efficient GPU-based implementation of the SLIC algorithm, known as `cuda-slic`⁴. Firstly, the 3D scan is clipped to the range -1024 to 400. Voxels of the lungs and surrounding organs all belong to this range. By clipping high density objects, we make sure that those areas are not segmented into multiple supervoxels, as they

⁴ <https://github.com/rosalindfranklininstitute/cuda-slic>

are not relevant to lung segmentation. Secondly, similarly to Özsvaş, E.E et al.’s **Step 1** (Section 2.1), we extract the body region mask. Thirdly, we feed the cuda-slic algorithm the clipped 3D scan multiplied by the body region mask to obtain the most supervoxels in the body region. Lastly, all supervoxel outside the body region are discarded.

Two key parameters are configured to guide the supervoxel generation process:

- **Number of Supervoxels** : We ensure that the algorithm generates supervoxels with a size around $7 \times 7 \times 7$ voxels. This is roughly estimated by dividing the total number of voxels in the 3D CT scan by the estimated size of supervoxels.
- **Compactness** : The compactness parameter is set to 100, which was determined visually to obtain regular shaped supervoxels, while delineating organ boundaries.

The output of the supervoxel algorithm is a 3D image where all voxels are assigned a label corresponding to the supervoxel they belong to. From this we generate two additional outputs, a 3D scan where all voxels are assigned their supervoxel’s median density. We refer to this scan as the *Supervoxel CT-scan*. Moreover, we construct a Region Adjacency Graph (RAG) to capture the spatial relationships between supervoxels.

3.2 Rough lung segmentation and trachea removal

To obtain an initial lung segmentation, we use the methodology in [7] working with the Supervoxel CT-scan, instead of the original scan. We follow their process, performing the thresholding step to extract the rough lung segmentation, and intestines removal as outlined in their paper.

We developed a new method for the segmentation of the trachea. We first select an upper slice ($\sim 35\text{mm}$ from the top) to determine a trachea seed point. We perform a 2D connected component and select the object closest to the center of the slice. Its center coordinates will represent the seed. Starting from a subset of the scan including the upper slices, we incrementally include 10 more slices and perform a 3D connected component. We continue including slices until the object containing the seed doubles in size or does not change volume upon inclusion, upon which we consider the previous objects voxels as trachea. This mask is dilated to include airway walls and then subtracted from the lung mask.

Finally, the right and left lung are separated using the same algorithm in Özsvaş, E.E et al.’s paper (Section 2.1).

3.3 Inclusion of pathological areas

For this step, in a similar fashion we construct an extended mask as in **Step 5** of the previous method (Section 2.1). Two new masks were created after this step: the *broad mask* containing all supervoxels that have at least one voxel in

the extended mask; and the *narrow mask* including supervoxels whose centroids are within the extended mask. Initially, all voxels contained in the narrow mask are labeled as lung.

Next, we proceed to expand the narrow mask by including adjacent supervoxels when the following criterion is met: if a defined percentage of the adjacent supervoxel’s surface is in contact with lung supervoxels, then it is considered lung. The rationale behind this procedure is that lung nodules will be enveloped by lung tissue or other lung nodules supervoxels. To prevent unintended propagation into the in-between region of the lungs (heart, etc.), while allowing for the inclusion of nodules attached to the external bounds of the lung, we employ a dynamic acceptance threshold strategy. The acceptance threshold for supervoxel candidates is initially set to 0.5 for the supervoxel closest to the center of the scan, and linearly decreases to 0.45 for those located further away. This strategy ensures controlled propagation within the lung region, since the lungs wraps more around the organs situated in the middle part of the body, which is not the case for the exterior of the lung.

Our propagation process is designed to ensure controlled expansion from the narrow mask into the pathological areas, while avoiding overflow beyond the broader mask. To achieve this, we follow these steps:

- **Initial Neighbor Consideration** : We begin by considering all neighboring supervoxels of those within the narrow mask that are included in the broad mask.
- **Lung Supervoxel Check** : Among these neighboring supervoxels, we assess whether they meet the criteria for being lung supervoxels. When we identify supervoxels that qualify as lungs, we add them to the narrow mask. This addition extends the segmentation further into the lung region.
- **Propagation** : We then repeat the process by considering the neighbors of the newly added lung supervoxels that are still within the broad mask. The propagation process continues until we no longer identify any new lung supervoxels within the broad mask.

This process is showed on Figure 1 where the supervoxels appear on one axial slice.




Fig. 1. Propagation process projected on one slice for 3 iterations. in blue the segmented lung supervoxels, in green the not selected lung candidates, and in red the added lung supervoxels at each iteration. The propagation process is performed in 3D.




Fig. 2. Example of lung segmentation obtained through thresholding, and the lung segmentation of our algorithm. Left: Lung mask obtained through thresholding. Right: Our algorithm's mask. In purple: Lung segmentation mask. In yellow: tumors not included in the mask. In orange: tumors included in the mask.

Algorithm 1 Get Supervoxel Lung Segmentation

Input: S ▷ 3D CT scan data
Output: M_{lung} ▷ Lung segmentation mask

- 1: $S_{super} \leftarrow \text{SLIC}(S)$ ▷ Using cuda_slic
- 2: $S_{super}^{ct} \leftarrow \text{Supervoxel CT scan}$ ▷ Median intensity of each supervoxel
- 3: $G_{RAG} \leftarrow \text{RAG}(S_{super})$ ▷ Section 3.1
- 4: $M_{rough} \leftarrow \text{Rough_segmentation}(S_{super}^{ct})$ ▷ Section 3.2
- 5: $M_{rough} \leftarrow M_{rough} - \text{Intestines}(M_{rough}) - \text{Trachea}(M_{rough})$
- 6: $[M_{rough}^{left} \ M_{rough}^{right}] \leftarrow \text{Separate_lungs}(M_{rough})$ ▷ Section 2.1
- 7:
- 8: ▷ Begin the propagation
- 9: $M_{lung} \leftarrow$ empty array of size S
- 10: **for** M_{rough}^{side} in $[M_{rough}^{left} \ M_{rough}^{right}]$ **do**
- 11: $M_{extended}^{side} \leftarrow \text{Extend_mask}(M_{rough}^{side})$ ▷ See Section 2.1
- 12: $[M_{narrow}^{side} \ M_{broad}^{side}] \leftarrow$ Get the narrow and broad mask from $M_{extended}^{side}$
- 13: candidates $\leftarrow \text{Get_candidates}(M_{narrow}^{side}, G_{RAG}, M_{broad}^{side})^a$
- 14: **while** candidates is not empty **do**
- 15: $L_{new} \leftarrow \{\}$ ▷ New lung supervoxels
- 16: **for** c in candidates **do**
- 17: **if** $\frac{\text{Area}(c \text{ in contact with } M_{narrow}^{side})}{\text{Area}(c)} > \text{thresh}$ **then** ▷ See 3.3
- 18: $L_{new} \leftarrow L_{new} \cup \{c\}$
- 19: **end if**
- 20: **end for**
- 21: $M_{narrow}^{side} \leftarrow M_{narrow}^{side} \cup L_{new}$
- 22: candidates $\leftarrow \text{Get_candidates}(M_{narrow}^{side}, G_{RAG}, M_{broad}^{side})$
- 23: **end while**
- 24: $M_{lung} \leftarrow M_{lung} \cup M_{narrow}^{side}$

The difference between the lung segmentation mask obtained after thresholding selection of lung tissue and the mask obtained after the propagation step is displayed in Figure 2. A majority of tumors that were not included in the first mask are now included with our method, and the surrounding organs were removed. Algorithm 1 sums up our method.

4 Experiments

We designed three experiments to demonstrate the efficiency and effectiveness of our method in comparison to the Özsvağ, E.E et al.’s method [7]. These experiments encompass evaluating the method’s ability to handle scans of non-sick patients without overflowing, and verifying that the segmentation results across a large sample of scans includes tumor regions. For this purpose, we analyzed the results of segmentation of 4 methods: ours (superlung), Özsvağ, E.E et al., Scancovia and TotalSegmentator. Three datasets were studied: Lung dataset, Tumors_{2D} and Tumors_{3D}.

4.1 Lung dataset

We assessed our method using 10 CT-scans from Nanjing Drum Tower Hospital [5], where both lungs were meticulously annotated by two radiologists and validated by an experienced radiologist. Subsequently, we computed four key metrics using the lung annotations: Intersection over Union (IOU), Precision, Recall, and DICE score. These metrics provided valuable insights into our method’s capability to avoid overflowing when propagating the segmentation mask into nodules.

Since lung voxels are included using a propagation process, we might expect the output lung mask to overflow in non-lung regions. Utilizing fully annotated lungs provides a reliable means to gauge the accuracy of our method in correctly segmenting the lung. The mean scores for each metric are presented in Table 1.

Table 1. Lung segmentation metrics on the Lung dataset. The metrics represent the mean (\pm std) over all scans.

Method	IoU	Precision	Recall	DICE
Superlung	0.87 (\pm 0.03)	0.89 (\pm 0.03)	0.98 (\pm 0.02)	0.93 (\pm 0.02)
Özsvağ, E.E et al.	0.89 (\pm 0.07)	0.94 (\pm 0.02)	0.94 (\pm 0.08)	0.94 (\pm 0.04)
Scancovia	0.95 (\pm 0.01)	0.96 (\pm 0.01)	0.99 (\pm 0.00)	0.98 (\pm 0.01)
TotalSegmentator	0.96 (\pm 0.01)	0.98 (\pm 0.00)	0.98 (\pm 0.00)	0.98 (\pm 0.00)

From these results, we can observe that the lung masks computed by our method overflow to areas close to the lung. This is manifested by a portion of false positive voxels included in the segmentation mask. Specifically, the average precision for our mask stands at 0.89, in contrast to 0.96 and 0.98 achieved by Scancovia and TotalSegmentator, respectively. However, our method does not

seem to suffer from false negative voxels, as its average recall stands at 0.98, on par with the U-Net model’s performance. Overall, all the methods seem to perform well when it comes to lung segmentation in non-cancer patients, with both neural network exhibiting an outstanding ~ 0.98 Dice coefficient.

We identify two primary causes for this overflow phenomenon. Firstly, an insufficiently accurate segmentation of the trachea results in the mask extending into regions between the lungs and the trachea. Secondly, propagation from supervoxels located between lung supervoxels vertically, as these regions are not always tumors. These problematic scenarios are exemplified in Figure 3, featuring one of the fully annotated lungs. The image on the right showcases a lung segmentation that extends into a region featuring tissue characteristics resembling tumors, which may as well be tumors themselves. This type of error represents a trade-off made to ensure the inclusion of as many tumors as possible, as will be seen in the next section (Section 4.2-4.3). The slight overflow into the rib area can be attributed to the acceptance threshold for new lung supervoxels, which might be set too low for this particular scan. This highlights the potential necessity of customizing the acceptance threshold for each scan to optimize segmentation results.




Fig. 3. Example of a segmentation using our algorithm. On the left, the trachea was not thoroughly removed, which may impact the inclusion of non-pathological areas. As a consequence, some of the heart (right image) was included in the lung mask.

4.2 Tumors_{2D}

The dataset contains 622 contrast-injected thoracic CT-scans, featuring a total of 15386 lung lesions. 2D fine contour annotations of each lung lesion on the axial slice with the largest diameter is provided. These scans originated from Gustave Roussy (GR) and were annotated by 2 radiologists, with subsequent review by 1 expert radiologist.

To assess the quality of the segmentation masks produced by these methods, we quantified the degree of overlap between these masks and the associated tumors mask as described by Equation (1). This evaluation allowed us to gauge how effectively the lesions were integrated into the lung masks.

$$\text{overlap} = \frac{|\text{Lung mask} \cap \text{Tumors mask}|}{|\text{Tumors mask}|} \quad (1)$$

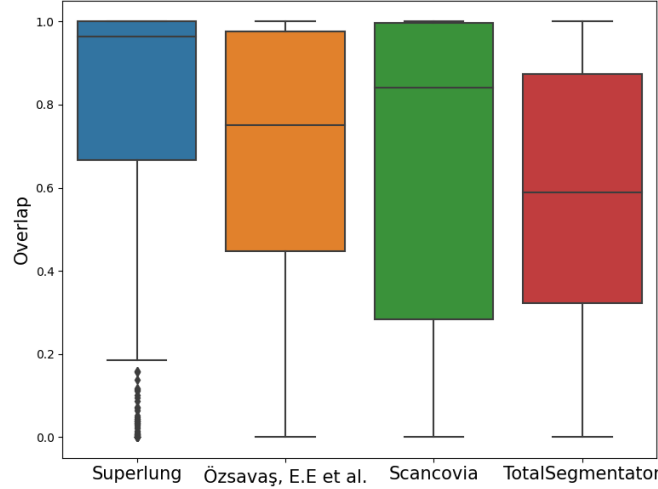


Fig. 4. Distribution of the overlap scores with the 622 scans from GR. The overlap scores reflect the percentage of included tumor masks in the lung segmentation.

We conducted an evaluation of the overlap between the segmentation masks generated by all methods with the annotated lesions. As evidenced by the distribution of overlap scores in Figure 4, Superlung manages to include more tumors in the lung segmentation mask. Furthermore, our method achieves great results consistently, as the first quartile corresponds to a 67% overlap of the annotations. Some difficult cases include patients where the lung lesion was quite large, covering most of the lung, or juxtapleural nodules that are big enough to not be integrated in the extended mask.

4.3 Tumors_{3D}

3D annotations for all nodules present in five scans from Tumors_{2D} were analyzed, totaling 112 lesions. The same method of evaluation as before was used: the lung mask was evaluated depending on the percentage of integrated nodules.

The overlap scores obtained by each method on each scan with fully annotated lesions are summarized in Table 2.

Both TotalSegmentator and Scancovia fall short in detecting as many tumors as Özsvaşı, E.E et al. and our method. In Figure 5, we display an example of segmented lungs of Scan 4 on one axial slice for each method. On this scan TotalSegmentator and Scancovia completely overlook the tumor closest to the center of the scan (missed parts in red), it underlines their weaknesses when it

Table 2. Summary of the overlap score on 5 scans with fully 3D annotated lung lesions. The overlap score reflects the percentage of lesion mask included in the obtained lung mask. The number between parentheses are the number of nodules in the scan.

Method		Scan 1 (3)	Scan 2 (5)	Scan 3 (5)	Scan 4 (98)	Scan 5 (1)
Overlap	Superlung	1.0	1.0	0.34	0.93	0.66
	Özsavaş, E.E et al.	1.0	1.0	0.28	0.83	0.53
	Scancovia	1.0	1.0	0.05	0.82	0.14
	TotalSegmentator	0.31	0.98	0.05	0.71	0.86

comes to out of the norm lesions that result in unusual lung shape. This is most probably due to the lack of cancer patients in the training of the U-Net models.

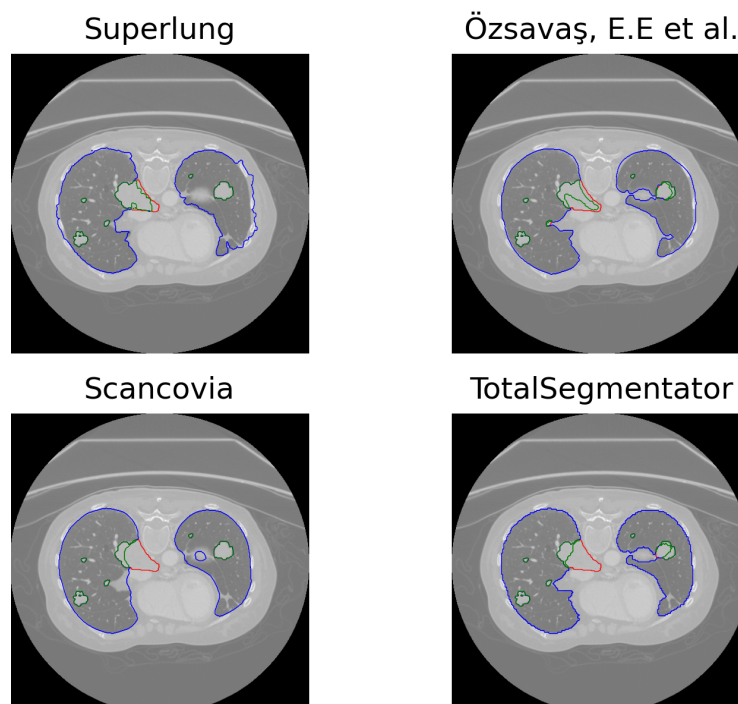


Fig. 5. Example of lung segmentation on an axial slice of Scan 4 of Tumors_{3D} dataset. In blue: the segmented lung. In green: tumors included by the lung segmentation mask. In red: tumor not included in the lung segmentation. A perfect segmentation would contain no red.

5 Conclusion

Segmenting the lung region is a crucial step in CAD, especially when the primary goal is to identify and segment lung lesions. Therefore, it becomes essential for the lung segmentation mask to effectively encompass all lesions. Our method does encounter challenges related to the generalization of certain hyperparameters, including the number and compactness of supervoxels, and the acceptance threshold. These challenges occasionally result in minor overflow issues within the segmentation mask. While our method overflows in the non-lung area, resulting in a degraded precision compared to other methods in non-cancer patients, it has demonstrated its capacity in tumor inclusion in the lung mask compared to other approaches. This accomplishment is attributed to our effective utilization of spatial and structural information contained in supervoxels.

References

1. Bizopoulos, P. et al.: Comprehensive Comparison of Deep Learning Models for Lung and COVID-19 Lesion Segmentation in CT scans, <http://arxiv.org/abs/2009.06412>, (2022).
2. Eisenhauer, E.A. et al.: New response evaluation criteria in solid tumours: Revised RECIST guideline (version 1.1). *European Journal of Cancer*. 45, 2, 228–247 (2009). <https://doi.org/10.1016/j.ejca.2008.10.026>.
3. Hansen, S. et al.: Anomaly detection-inspired few-shot medical image segmentation through self-supervision with supervoxels. *Medical Image Analysis*. 78, 102385 (2022). <https://doi.org/10.1016/j.media.2022.102385>.
4. Huidrom, R. et al.: Automated Lung Segmentation on Computed Tomography Image for the Diagnosis of Lung Cancer. *CyS*. 22, 3, (2018). <https://doi.org/10.13053/cys-22-3-2526>.
5. Jun, M. et al.: COVID-19 CT Lung and Infection Segmentation Dataset, <https://zenodo.org/record/3757476>, (2020). <https://doi.org/10.5281/ZENODO.3757476>.
6. Lassau, N. et al.: Integrating deep learning CT-scan model, biological and clinical variables to predict severity of COVID-19 patients. *Nat Commun*. 12, 1, 634 (2021). <https://doi.org/10.1038/s41467-020-20657-4>.
7. Özsavaş, E.E. et al.: Automatic Segmentation of Anatomical Structures from CT Scans of Thorax for RTP. *Computational and Mathematical Methods in Medicine*. 2014, 1–14 (2014). <https://doi.org/10.1155/2014/472890>.
8. Achanta, R. et al.: SLIC Superpixels, EPFL Technical Report no. 149300, June 2010.
9. Shanhui Sun et al.: Automated 3-D Segmentation of Lungs With Lung Cancer in CT Data Using a Novel Robust Active Shape Model Approach. *IEEE Trans. Med. Imaging*. 31, 2, 449–460 (2012). <https://doi.org/10.1109/TMI.2011.2171357>.
10. Wasserthal, J. et al.: TotalSegmentator: robust segmentation of 104 anatomical structures in CT images. (2022). <https://doi.org/10.48550/ARXIV.2208.05868>.
11. Zimmermann, M., et al: (2021). CT-based whole-body tumor volumetry versus RECIST 1.1: Feasibility and implications for inter-reader variability. *European Journal of Radiology*, 135, 109514. <https://doi.org/10.1016/j.ejrad.2020.109514>