



# Fully automated lung segmentation from chest radiographs using SLICO superpixels

Bilal Yassine<sup>1</sup> · Paul Taylor<sup>1</sup> · Al Story<sup>2</sup>

Received: 14 January 2018 / Revised: 26 February 2018 / Accepted: 2 March 2018 / Published online: 20 March 2018  
© The Author(s) 2018

## Abstract

This project aims to create a computer-aided diagnosis (CAD) system that can be used to identify tuberculosis (TB) from chest radiographs (CXRs) and, in particular, to observe the progress of the disease where patients have had multiple images over a period of time. Such a CAD tool, if sufficiently automated could run in the background checking every CXR taken, regardless of whether the patient is a suspected carrier of TB. This paper outlines the first phase of the project: segmenting the lung region from a CXR. This is a challenge because of the variation in the appearance of the lung in different patients and even in images of the same patient.

**Keywords** Tuberculosis · Radiographs · Segmentation · SLICO · CAD

## 1 Introduction

TB remains a global health challenge. Estimates of the number of people infected with TB are likely to be an underestimate since they are based on calculations from reported cases, notification data and expert opinion, however in 2013 the World Health Organisation (WHO) estimated that there were 9 million people with TB, including an estimated 1.2 million people with TB and HIV. TB is believed to be responsible for around 1.5 million deaths a year [1]. In Africa, the size of the TB problem is technically considered undefined due to the poor infrastructure for case detection, recording and reporting [2].

Multi-drug-resistant tuberculosis (MDR-TB) and extensively drug-resistant TB (XDR-TB) are now major threats to health in Europe, Asia and southern Africa. The

number of new cases of MDR-TB caused by *Mycobacterium tuberculosis* (M.tb) strains resistant to rifampicin and isoniazid is increasing. An increase in numbers is also seen in strains of TB resistant to Rifampicin, Isoniazid, plus any fluoroquinolone and at least one of the three injectable second-line TB drugs: Amikacin, Kanamycin, Capreomycin [2]. In 2013 there were an estimated 480,000 new cases of MDR-TB discovered around the world. This number is also likely to be an underestimate [2].

The development of an effective vaccine against TB is a very difficult challenge. Avenues for research and development include experimental medicine, biomarker discovery, and proof of concept studies that aim to streamline vaccine development and maximise the probability of success in late-stage trials [3].

The pathophysiology of TB is very complicated, the disease can have various effects on the lung and other organs. The morphological patterns that are observable on a CXR vary significantly. This variation depends on factors such as age, ethnicity and HIV infection [4]. Appearance also varies as the disease progresses. TB, lesions are small in the early stages and change their morphology over time.

Public Health England's issues the following guidelines for the management of patients with active pulmonary TB. A patient with any of the following findings on a CXR must submit sputum specimens for examination [5].

1. Infiltrate or consolidation.

---

✉ Bilal Yassine  
bilal.yassine.14@ucl.ac.uk

Paul Taylor  
p.taylor@ucl.ac.uk

Al Story  
al.story@nhs.net

<sup>1</sup> Institute of Health Informatics, CHIME, 222 Euston Road, London NW1 2DA, UK

<sup>2</sup> University College Hospitals NHS Foundation Trust, 5th Floor East, 250 Euston Rd, London NW1 2PG, UK

2. Any cavitory lesion.
3. Nodule with poorly defined margins.
4. Pleural effusion.
5. Hilar or mediastinal lymphadenopathy (bihilar lymphadenopathy).
6. Linear, interstitial disease (in children only).
7. Other—any other finding suggestive of active TB, such as miliary TB. Miliary findings are nodules of millet size (1–2 mm) distributed throughout the parenchyma.

In a recent study, it was estimated that the prevalence of tuberculosis in hard to reach groups in London was 788 per 100,000 in homeless people, 354 per 100,000 in people with problematic drug use, and 208 per 100,000 in prisoners. In comparison, the overall prevalence of tuberculosis in London was 27 per 100,000 people. Although only 17% of tuberculosis cases in London are in hard to reach areas, they make up nearly 38% of non-treatment cases, 44% of lost cases, and 30% of all highly infectious cases. Tuberculosis controls are therefore needed in targeted interventions to address transmission within these groups. In April 2005, the English Department of Health provided funding to set up a mobile radiography unit that could actively screen for tuberculosis disease in London's vulnerable populations. The service, known as Find and Treat visits locations where high-risk groups can be found, including drug treatment services and hostels or day centres for homeless and impoverished people. All individuals are screened on a voluntary basis regardless of their current symptom status [6].

This project uses over 90,000 images recorded by Find and Treat between 2005 and 2015. All the images used in this project were obtained by the same machine.


The first phase of the project involves the segmenting the lungs and then dividing them up into regions of interest. The approach taken here uses SLICO superpixels. Superpixel algorithms group pixels into clusters whose boundaries replace the conventional pixel grid [7]. Since clusters are groups of similar pixels, they can be treated as single entities reducing the computational cost of processing the image post-segmentation [8, 9]. Many computer vision algorithms now use superpixels as a basic building block due to their effectiveness at multiclass object segmentation [10], depth estimation [11], body model estimation [12], and object localization [8].

There are many different techniques to generate superpixels, each has its own advantages and disadvantages that may be more useful in certain scenarios. Simple linear iterative clustering (SLIC) by Achanta et al. [7] is a novel method for generating superpixels. The benefits of SLIC are that it is faster than existing methods, more memory efficient has extremely accurate boundary conformity, and improves the performance of segmentation algorithms. In

this project, we used a modified version of SLIC called SLICO. The difference is in the input parameters: SLIC requires the user to input both the number of superpixels to create and the compactness factor, SLICO only needs the number of superpixels. Compactness is a measure of shape calculated as a ratio of the perimeter to the area. In the case of SLIC, this means that if there are regions in the image with a smooth texture then smooth regular sized superpixels are produced. If the region texture is coarse then the superpixels will be very irregular in shape. This issue arises from the dependence on a compactness factor. SLICO only requires the number of superpixels required. It adaptively changes the compactness factor depending on the texture of the region, this creates regularly shaped superpixels regardless of the texture. This improvement is claimed to have a very little impact on performance, making it an ideal solution for complex segmentations. Figure 1 shows a comparison of performance between SLIC and SLICO algorithms.

## 2 Proposed segmentation method

The proposed method for segmenting the lungs involves at first making a rough segmentation using basic image processing techniques. The first step involves some pre-processing. The CXRs in our dataset are 16-bit images giving a very large range of available grey values. We reduced the number of grey values from  $2^{16}$  to 64. This makes the edges more prominent, eliminating the gradual changes in gradient typically found in the image. This binning process is done using histogram equalization: an approximately equal number of pixels is mapped to each the bin. This step results in a huge loss of detail that is not required for segmentation. The next step is to adjust the contrast of the




**Fig. 1** Images segmented using SLIC and SLICO into superpixels of approximate size 64, 256, and 1024 pixels. A bone suppressed image (BSI) was also applied to both sets of images (all examples are of the same CXR)

image. This removes even more of the detail in the image by mapping the intensity values in the grayscale image to new values so that 1% of the data is saturated at low and high intensities (this value was determined through trial and error). This increases the contrast of the output image. It makes the inside of the lung almost entirely white and shows much clearer edges in areas where the lung is obscured by other tissue. However, this is only the first step and subsequent work is needed to drastically improve the result. Each pre-processing step aims to show a better definition of the Lung edges. The following step does this through an algorithm that creates a small user-defined box (In testing a 9 by 9 box performed well). The values of the pixels that make up the box are analysed and a decision is made about and the box in order to decide if it lies on an edge within the image or not. It then moves on to the adjacent area and repeats the process, scanning the entire image as it moves. The analysis within the box begins with the centre pixel which is taken as the reference point and all pixels around it are investigated. It is known the bones in the image are dark and the inside of the Lungs are lighter in appearance (a colour map is used to show the bones as white when observed by humans Fig. 2).

This step is followed by a SLICO superpixel segmentation to provide a more accurate delineation of the lung boundaries in the CXR. The superpixel algorithm is run on the original image and overlaid on the processed image created in the previous step. The accuracy of the SLICO segmentation is dependent on the number of superpixels used.

The process of finding the lungs begins by randomly choosing superpixels a third of the way in from the image left and right edges, and halfway down from the top. Unless the CXR quality is extremely poor and too bad for even human reading this guarantees the first superpixel will




**Fig. 2** CXR before normal correct colour map is applied and blown up region showing the shallow gradient on the lung edge (Color figure online)

always be inside the lungs. The surrounding superpixels are then investigated to see if they overlap with the pixels representing the lung edge in the pre-processed image; this process is repeated until the edges are found (Fig. 3).

### 3 Evaluation

To assess the accuracy of the segmentation algorithm, the generated segmentations were compared to a gold standard of hand segmentations performed by two radiographers and a doctor. All three work for Find and Treat and have a wealth of experience reading CXR's and diagnosing TB. 40 images were used with each of the three human readers segmenting all 40. The images were segmented using a Surface Pro Tablet and digital pen. The software used was MicroDicom which is an application for primary processing and preservation of medical images in the DICOM format. The hand segmentations were done by the participant drawing an outline with the pen tool in the software around the lung edge. The images were then imported into a photo editing program to turn it into a binary image for comparison, using selection tools to mark where the hand segmentation was made.

The images segmented using the algorithm are processed in a similar fashion. The algorithm identifies the superpixels that make up the lung edges in the images and therefore can be turned into binary images quite easily. The binary images being compared were overlaid to show how



**Fig. 3** Example of superpixel segmentation overlaid on pixel processed lung image (Color figure online)

much overlap there was as well as the areas that didn't overlap. The comparison between the images was made using DICE similarity scores calculated for the algorithm compared to each of the three hand segmentations. The doctor's segmentation was also compared against each of the radiographers and the radiographers were compared with each other. Due to variances in the hand segmentations, the results were calculated twice. Firstly, with each hand segmentation individually against the algorithm and secondly with a combination of the three hand segmentations against the algorithm.

## 4 Results


Table 1 shows the mean DICE scores for all 40 images on each of six pairwise comparisons and one comparison between the algorithm and the pooled result. A similarity score closer to 1.0 between manual and automatic indicates that the result was similar. Figures 4, 5, 6 and 7 display the overlap of the binarized segmentations comparing the algorithm with each of the manuals separately. The figures are all laid out in the same order: top left is the doctor versus the algorithm, bottom left is radiographer 1 versus the algorithm, the top right is radiographer 2 versus the algorithm and the bottom right of each of the figures shows how the algorithm compares with the pooled manual segmentations. In each of the figures, the white represents the automatic segmentation with the gold standard shown in grey.

## 5 Discussion

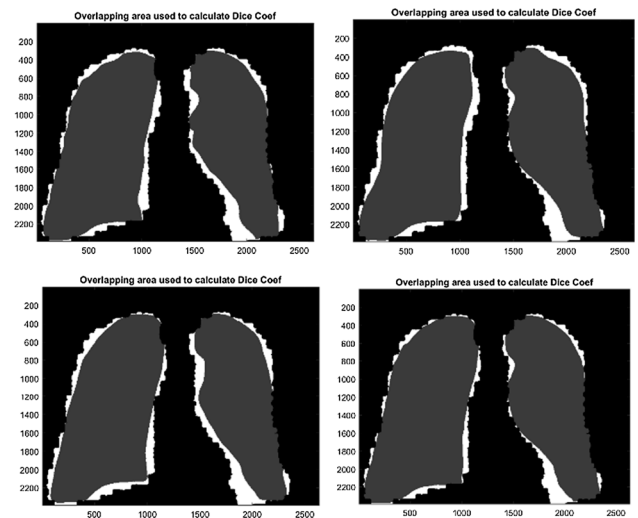
All the segmentations had an acceptable range of variation from each other. The similarity scores showed that the doctor and second radiographer were almost identical in terms of the similarity between their segmentations and that of the algorithm, the first radiographer fell behind only slightly. The largest difference between any two scores was

**Table 1** The mean result of DICE similarity score

|                           |       |
|---------------------------|-------|
| Doc versus algorithm      | 0.953 |
| Rad1 versus algorithm     | 0.957 |
| Rad2 versus algorithm     | 0.952 |
| Combined versus algorithm | 0.901 |
| Doc versus Rad1           | 0.882 |
| Doc versus Rad2           | 0.876 |
| Rad2 versus Rad3          | 0.877 |



**Fig. 4** Patient 011621\_20061017101256A



**Fig. 5** Patient 006051\_20050711100551A

not more than 0.07, the lowest score was 0.76 between radiographer 1 and the algorithm. This image had very small lungs most likely due to the patient not taking the correct breath, however, this didn't affect the clarity of the lung edges. The overlap area image in Fig. 8 shows that the algorithm seems to have over segmented the Lung, including part of the heart and more of the rib cage wall. This error is most likely due to the parameters used by the algorithm in the preprocessing phase. Inspection of the DICE scores across the test set suggests that this is an isolated problem.






Fig. 6 Patient 012696\_20061109113616A




Fig. 7 Patient 007714\_20050811093815A

## 6 Conclusion

The DICE scores show that the performance of the algorithm is very similar to that of the manual segmentations. The result of combining the manual segmentation produced a closer result to the automated segmentation than any of the comparisons between individual manual segmentations. The variation between human readers seems higher than between the algorithm and the human readers, which suggests that the algorithm is accurate. Figures 4, 5, 6, 7 and 8 show some typical examples of segmentations. The grey region is the overlap between manual and automated segmentations. The white, non-overlapping, region mostly comes from the automated segmentation, which is appropriate for this application because the role of the

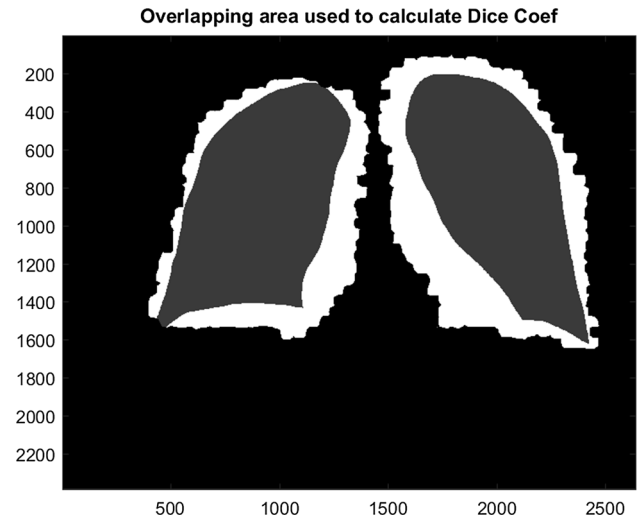


Fig. 8 Example of over segmentation by the algorithm

segmentation step is to define a region of interest to be processed in subsequent steps looking for abnormalities. Over-segmenting slightly is therefore preferable to under-segmenting.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.


## References

- Zumla, A., Maeurer, M., Marais, B., et al. (2015). Commemorating world tuberculosis day 2015. *International Journal of Infectious Diseases*, 32, 1–4. <https://doi.org/10.1016/j.ijid.2015.01.009>.
- Zumla, A., Petersen, E., Nyirenda, T., & Chakaya, J. (2015). Tackling the tuberculosis epidemic in sub-Saharan Africa—Unique opportunities arising from the second European developing countries clinical trials partnership (EDCTP) programme 2015–2024. *International Journal of Infectious Diseases*, 32, 46–49. <https://doi.org/10.1016/j.ijid.2014.12.039>.
- da Costa, C., Walker, B., & Bonavia, A. (2015). Tuberculosis vaccines—State of the art, and novel approaches to vaccine development. *International Journal of Infectious Diseases*, 32, 5–12. <https://doi.org/10.1016/j.ijid.2014.11.026>.
- Salazar, G. E., Schmitz, T. L., Cama, R., et al. (2001). Pulmonary tuberculosis in children in a developing country. *Pediatrics*, 108, 448–453. <https://doi.org/10.1542/peds.108.2.448>.
- Public Health England. (2014). Tuberculosis (TB) and other mycobacterial diseases: Diagnosis, screening, management and data—GOV.UK. <https://www.gov.uk/government/collections/tuberculosis-and-other-mycobacterial-diseases-diagnosis-screening-management-and-data#diagnosis>. Accessed Dec 3, 2017.
- Jit, M., Stagg, H. R., Aldridge, R. W., et al. (2011). Dedicated outreach service for hard to reach patients with tuberculosis in

- London: Observational study and economic evaluation. *BMJ*, 343, d5376–d5376. <https://doi.org/10.1136/bmj.d5376>.
7. Achanta, R., Shaji, A., Smith, K., et al. (2012). SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34, 2274–2281. <https://doi.org/10.1109/TPAMI.2012.120>.
  8. Fulkerson, B., Vedaldi, A., & Soatto, S. (2009). Class segmentation and object localization with superpixel neighborhoods. In *Comput Vision, IEEE 12th International Conference on 2009* (pp. 670–677). <https://doi.org/10.1109/iccv.2009.5459175>.
  9. Vincent, L., & Soille, P. (1991). Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13, 583–598.
  10. Li, Y., Sun, J., Tang, C.-K., & Shum, H.-Y. (2004). Lazy snapping. *ACM Transactions on Graphics*, 23, 303. <https://doi.org/10.1145/1015706.1015719>.
  11. Zitnick, C. L., & Kang, S. B. (2007). Stereo for image-based rendering using image over-segmentation. *International Journal of Computer Vision*, 75, 49–65. <https://doi.org/10.1007/s11263-006-0018-8>.
  12. Mori, G., & Va, C. (2005). Guiding model search using segmentation 2. Related work. *Computer (Long Beach Calif)*, 2, 1417–1423. <https://doi.org/10.1109/ICCV.2005.112>.



**Bilal Yassine** received his B.E. in Mechanical and Medical Engineering at the University of Hull, he then went to the University of Reading for his M.Sc. in Digital Signal Processing and Communications. He is currently working on his Ph.D. at University College London in Medical Image Processing.



**Paul Taylor** his main research interest is in clinical decision support systems which includes understanding the information needs of clinicians, designing tools to present relevant clinical information and assessing the effectiveness of technology designed to help clinicians make decisions. The main focus of his research has been in medical imaging with two major projects involving (a) designing and building a system for the differential diagnosis of mammographic calcifications and (b) assessing the potential of a commercial system designed to alert radiologists to potential abnormalities. This research is developing in a new direction and he is now working on the development of a system to help train radiologists seeking to specialize in screening mammography. He hopes to extend these ideas to other areas of radiology.



**Al Story** is the founder and Clinical Lead of the pan-London Find&Treat Service. His core expertise is in tackling communicable diseases among vulnerable populations. Research interests include outreach, point of care diagnostics, computer aided diagnostics, Peer interventions and using mobile technologies to promote engagement and treatment continuity. He is an original member of the Faculty for Homeless and Inclusion Health and the Pathway Team and a senior lecturer with the Farr Institute at University College London.