

Adaptive Local Multi-Atlas Segmentation: Application to Heart Segmentation in Chest CT Scans

Eva M. van Rikxoort, Ivana Isgum, Marius Staring, Stefan Klein and Bram van Ginneken
Image Sciences Institute, University Medical Center Utrecht
Heidelberglaan 100, 3584 CX Utrecht, The Netherlands
eva@isi.uu.nl

ABSTRACT

Atlas-based segmentation is a popular generic technique for automated delineation of structures in volumetric data sets. Several studies have shown that multi-atlas based segmentation methods outperform schemes that use only a single atlas, but running multiple registrations on large volumetric data is too time-consuming for routine clinical use. We propose a generally applicable adaptive local multi-atlas segmentation method (ALMAS) that locally decides how many and which atlases are needed to segment a target image. Only the selected parts of atlases are registered. The method is iterative and automatically stops when no further improvement is expected. ALMAS was applied to segmentation of the heart on chest CT scans and compared to three existing atlas-based methods. It performed significantly better than single-atlas methods and as good as multi-atlas methods at a much lower computational cost.

1. INTRODUCTION

Atlas-based segmentation uses registration to achieve segmentation and this has proven to be a powerful concept.¹ Atlas-based methods start by registering an atlas image with a target image that is to be segmented. To obtain a segmentation of the target image, a manual labelling of the atlas is transformed using the mapping determined during the registration. This process is called label propagation. Major advantages of this segmentation approach, especially for 3D applications, are simplicity (only a registration framework and a number of pre-segmented data sets are required, no need for landmarking and complex training procedures), general applicability (a wide range of segmentation tasks can be solved by this method) and robustness (fairly stable volumetric registration methods exist).

A critical assumption of atlas-based segmentation is that it is possible to find a deformation to the atlas that aligns it with the target such that the objects of interest line up perfectly. However, insufficient similarity between the atlas and the target image often results in local mismatches, which in turn leads to segmentation errors.²

This fundamental problem of single atlas-based segmentation has been recognized and an obvious solution has been attempted: instead of a single atlas, multiple atlases are registered and the propagated labels are fused, for example by averaging. Local errors of a single atlas can now be corrected, as long as the majority of atlases register correctly. Several studies have shown that multi-atlas segmentation outperforms methods that use a single atlas.^{1,3,4}

The main drawback of multi-atlas registration methods is computational complexity. Since atlases and target images are taken from different patients, an affine transformation is not sufficient and elastic registration is needed. Elastic registration of large data sets can be computationally expensive. In practice a large set of atlases is needed to represent the variety of data encountered in a real application, which makes the method extremely time-consuming and therefore not suitable for routine clinical use. Another drawback of multi-atlas segmentation is that if the target image shows substantial local deviations from the majority of atlases (and most images will show such deviations at some locations), the method will still produce errors, i.e. the majority is not always right.

In this work, we take the concept of multi-atlas registration one step further to address these two problems. Multiple atlas registrations are only needed when a single atlas registration has failed in a region around the

edge of the object to be segmented. In that case, additional registrations can be beneficial, especially if scans that are *locally similar* to the target image are used. We propose a framework that starts with a single atlas registration and subsequently locally inspects the likelihood that the segmentation is correct. If not, blocks from additional atlases are selected that are expected to be locally similar to the target and thus will improve the segmentation. The process is iterative and uses different atlas blocks for different locations. Updating stops automatically when no further improvement for a block is expected. We have named our method adaptive local multi-atlas segmentation (ALMAS). The concept is illustrated in Fig. 1.

We use ALMAS to segment the heart on non-contrast volumetric chest computed tomography (CT) scans. This task is important for the detection of cardiac abnormalities. ALMAS is compared to three existing atlas-based segmentation methods.

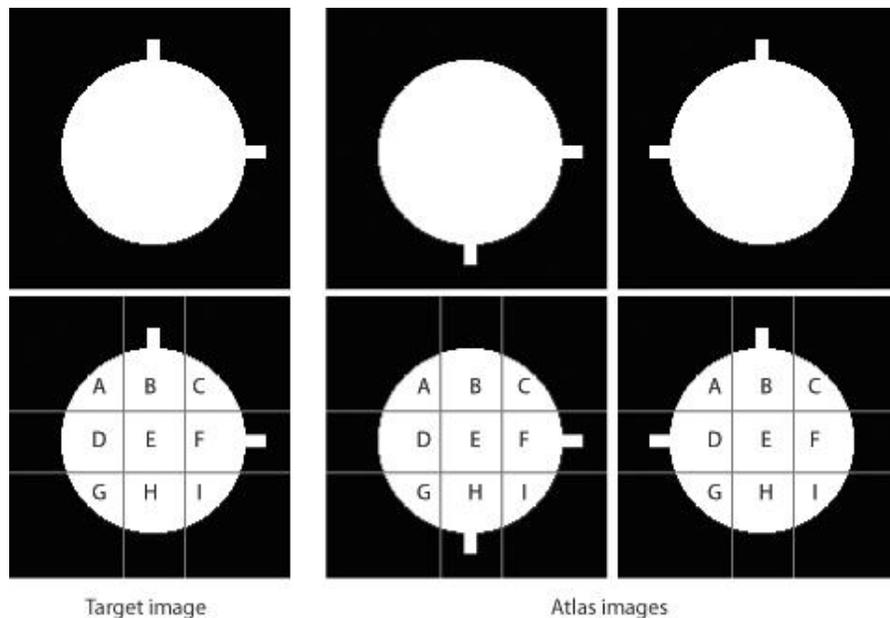


Figure 1. Illustration of the concept of local atlas segmentation. Imagine the target image is the image on the left and the atlas images are the two images on the right. Although the target image can be composed with parts of the different atlases, it is difficult to transform either of the complete atlases to the target image. We propose to divide the images into a number of blocks and register separate blocks. In this example, block D and F from the left atlas will be used and block B and H from the right atlas. For the other blocks either of them can be used. This way the final registration result for the target image will be composed of parts of different atlases.

2. METHOD

2.1. Registration

The method we propose is generally applicable and any registration method can be plugged in. The particular method used in all methods in this paper formulates the registration problem as an optimization problem in which the similarity between the target and atlas image is maximized. For the cost function negative mutual information was used, following the implementation in.⁵ The transformation is initialized with an affine transformation followed by a non-rigid transformation modelled by a B-spline. The employed registration framework is largely based on⁶ and.⁷ For the optimization of the cost function, an iterative stochastic gradient optimizer is used. To avoid local minima, a multi-resolution strategy was taken. Parameter settings are provided in the next section.

2.2. Adaptive local multi-atlas segmentation (ALMAS)

During the registration, a coordinate transformation is determined which transforms the atlas image to the target image. In the ideal case the transformed atlas image would be equal to the target image and the difference between them would be a zero image. In reality, registration does not perfectly align the two images and local misalignment occurs. Based on this observation, we propose to decide which atlas image is locally most similar to the target image and only register part of that atlas.

The ALMAS method is an iterative method which starts by registering one atlas (referred to as the reference atlas) to the target image, as in single atlas segmentation. Next, in each iteration, a part of the target image (we use rectangular blocks here) that has not been registered correctly yet is automatically detected and that part of the segmentation of the target image is updated using a corresponding part of the most promising atlas. The updating automatically stops when no improvement is expected anymore for a block.

This strategy was implemented as follows. ALMAS starts with a training phase that only needs to be performed once. In this training phase, from the total set of N atlases A_1, A_2, \dots, A_N with corresponding manual segmentations S_1, S_2, \dots, S_N one reference atlas A_r with its corresponding manual segmentation S_r is chosen randomly. A_r is registered to all $N - 1$ remaining atlases and its labels propagated. The propagated labels from A_r are used to define a number of blocks on the atlases. For this paper, eight blocks in the heart were defined by dividing the volume of the resulting segmentation in each direction. Each block contained exactly $\frac{1}{8}$ volume of the heart and was guaranteed to contain the border of the heart. For each block j in every atlas A_i , the mean absolute difference d_{ij} with the corresponding block in the registered A_r in a region within 10 voxels around the border of the propagated labels is computed.

The segmentation of an unseen target image starts by registering A_r and propagating its labels. Based on the propagated labels the same blocks are defined in the target image as for all atlases and d_{tj} to A_r is calculated for each block j . The probabilistic segmentation of the target image (S_t) is initialized by the propagated labels of A_r . Next, an iterative process is started that updates different blocks of the segmentation. To determine which block of the target image needs to be updated, the percentage of voxels that is likely to switch segmentation label if another set of propagated labels for that block were available is computed. We define these voxels as those where $n/2$ of n propagated labels agree (if n is odd we floor $n/2$). The percentage of those voxels in a block is indicated with T_j . The block with the largest T_j is chosen to be updated. After the initial registration of A_r , T_j is not yet defined and initialized to 1 so that all blocks are updated at least once. To select the atlas that is most promising to improve the segmentation, d values calculated for both the target block and all atlas blocks are employed. A similar d for two blocks is assumed to indicate that the registration of A_r failed in a similar way for the two blocks. Therefore, the atlas block with the most similar d is selected. This block is registered to the target image and its labels propagated. Since the definition of the blocks as used in this paper is not guaranteed to contain exactly the same anatomical part, the blocks to be registered are defined with overlapping borders. S_t is updated with the propagated labels from the selected atlas block, using averaging. Finally, a stopping criterion is defined: Only blocks with $T > T_{min}$ are updated.

To summarize, a stepwise description of ALMAS is given:

1. Register A_r to the target image
2. Initialize S_t to the propagated labels of A_r
3. Divide the target image into blocks and calculate d to A_r for each block, initialize T_j for each block j to 1
4. Take the block with the largest $T_j > T_{min}$ for which there are still unused atlas blocks available, if any
 - (a) Select the block from the atlases with the most similar d_j that has not been used before and register it to target block j .
 - (b) Propagate the labels of the atlas block, update S_t , update T_j
5. The final segmentation S_t is the result of thresholding.

2.3. Previously proposed atlas-based segmentation methods

Three existing methods were implemented for comparison.

1. **Atlas-based segmentation with a single best atlas (SBA).** When segmentation is performed using a single atlas, a criterion needs to be defined to select the atlas. We selected one atlas from the total set judged to be representative by visual inspection. This atlas is registered to the target image followed by the propagation of the label image to obtain the segmentation. A similar approach was used in.³
2. **Average-shape atlas-based segmentation (ASA).** A single atlas A_j is randomly chosen from the set of N atlases. The remaining $N - 1$ atlases are registered to A_j and their labels propagated. The gray values as well as the propagated labels are averaged to obtain the average shape atlas and its segmentation. Segmentation of the target image is computed by registering the single average-shape atlas to the target image, and propagating its segmentation. This approach was used in.⁸
3. **Multi-atlas segmentation with averaging as decision fusion (ADF).** All training atlases A_1, A_2, \dots, A_N are registered to the target image, followed by propagation of the labels. The probability of a label in the segmentation of the target image is determined by averaging the propagated labels, followed by thresholding. This method was used in e.g..^{3,4}

3. EXPERIMENTS

Twenty-nine CT scans of the thorax were randomly taken from a lung cancer screening trial with low dose CT (30 mAs at 120 kV for patients weighing ≤ 80 kg and 30 mAs at 140 kV for those weighing over 80 kg). Data was acquired in spiral mode with 16×0.75 mm collimation. No contrast material was injected. Axial images of 1.0 mm thickness at 0.7 mm increment were reconstructed. All scans were reconstructed with a 512×512 matrix, yielding an axial resolution in between 0.6 and 0.8 mm. The scans were randomly divided into two sets, a set of 15 atlases and a set of 14 target images. The heart was segmented by two medical students who were trained and supervised by a radiologist. The 15 atlases were segmented by one student, and the target images by both of them. The results of the observer who segmented all images were used as the reference standard. Typically, manual segmentation of the heart took about 90 minutes.

All 29 images were down-sampled by a factor two in each dimension in order to reduce computer memory requirements and computation time of registration. Registration parameters were determined in a set of pilot experiments. For the affine registration 4 resolutions were used, in each of which 512 iterations of the stochastic gradient descent optimizer were performed. For the nonrigid B-spline registration 5 resolutions were used. The B-spline grid spacing used in these resolutions was 64, 64, 32, 16, and 8 voxels respectively. The optimizer performed 256 iterations in each resolution. For both affine and nonrigid registration 32 histogram bins were used. With these settings, one registration takes about ten minutes on a standard high-end PC.

ALMAS, SBA, ASA, and ADF were applied to the 14 target images using the manual segmentations of the 15 atlases. For ALMAS the training phase was computed using the 15 atlases. Based on pilot experiments, the threshold T_{min} was set to 8%. To obtain segmentation of the images in their original size, the segmentation results were super-sampled to the original resolution.

4. RESULTS

ALMAS performed on average 3.0 updates per block (ranging from 1.5 to 5.4 average per block), requiring around 30 minutes for segmenting a complete scan. Figure 2 shows the segmentation results in two slices from two different subjects. To quantify the performance of the different methods the volumetric overlap is calculated, defined as the number of voxels in the intersection of automatic segmentation and reference divided by the number of voxels in the union. Table 1(a) gives for each method, and for the second observer, the average overlap on the target images and standard deviation. Since on average three updates per block were used for ALMAS, the

Table 1. Results for the different atlas-based segmentation methods.

(a) Segmentation results: mean overlap and standard deviation (SD) between the reference standard and the automatic methods.

method	Mean	SD
2nd observer	0.8791	0.0222
ALMAS	0.8733	0.0363
ADF 15 atlases	0.8744	0.0375
ADF 4 atlases	0.8566	0.0351
ASA	0.8520	0.0426
SBA	0.8318	0.0468

(b) p -values for the significance of the differences in overlap for the various methods. The p -values are from a two-tailed paired t -test.

methods	p -value
ALMAS vs. 2nd observer	$p = 0.62$
ALMAS vs. ADF 15 atlases	$p = 0.60$
ALMAS vs. ADF 4 atlases	$p < 0.01$
ALMAS vs. ASA	$p < 0.01$
ALMAS vs. SBA	$p < 0.01$

computation time is slightly less than that of ADF using 4 atlases. Therefore ALMAS was also compared to ADF using 4 atlases.

To compare the performance of ALMAS to the other methods two-tailed paired t -tests were performed. Table 1(b) lists the results. There is no significant difference between ALMAS and the second observer. In addition ALMAS is comparable to ADF using 15 complete atlas scans. When only 4 atlas scans were used for ADF, ALMAS performed significantly better. ALMAS is also significantly better than ASA and SBA.

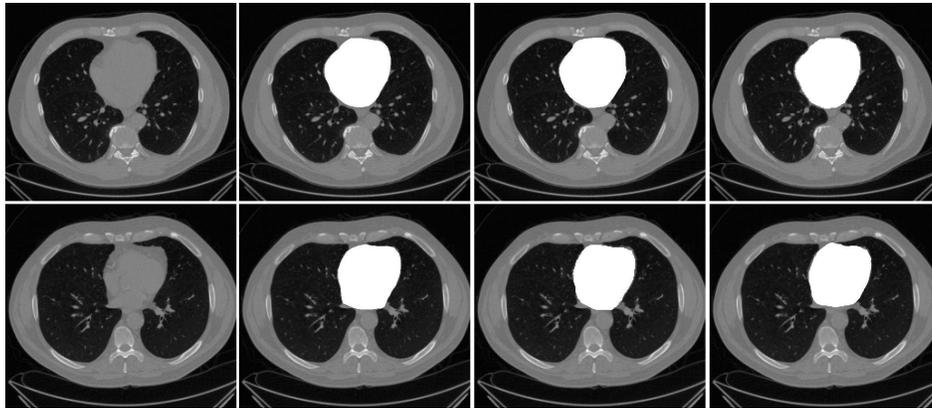


Figure 2. Example of the results of automatic segmentation of the heart for two slices of two different scans. The first column shows the original slice. In the second column the reference standard is shown. The third column shows the result of the second observer and finally the last column shows the results of ALMAS.

5. DISCUSSION

In this paper a fully automatic segmentation by registration method was presented (ALMAS) which locally decides which atlas registrations to perform. The method was applied to segmentation of the heart on non-contrast volumetric CT data. The results show that the ALMAS approach is feasible and able to perform as well as multi-atlas methods which are computationally more expensive. We hypothesize that ALMAS will prove especially useful for tasks where anatomical variations or pathologies are present. For those tasks, many more atlases than were used here might be needed to contain similar atlases for all locations in the target images. Not only is registering a very large amount of atlases using the ADF approach not feasible in a reasonable time, it will also not produce a correct result when only a few out of the complete set of atlases contain the correct

anatomy. ALMAS may be able to select the correct atlases locally and only register those, which leads to better performance and lower computational time.

A critical point of ALMAS is how to select the block in the target image to be updated and how to select the most promising atlas. The criteria used in this paper are fairly simple, yet seem to perform well already. But although the d value may be a valid measure for the success of a registration, two blocks having a similar d to the reference image are not necessarily similar. If the method is scaled up to using many more atlases, investigation of effective, more sophisticated selection criteria is a challenging research direction.

Our implementation for the application of the segmentation of the heart uses a straightforward division of the heart into a number of blocks. However the framework allows any division into blocks, so for other applications blocks can for example be defined based on anatomical positions. When very small blocks are defined, an affine registration might be sufficient which will further speed up the algorithm.

Note that in the current implementation, the processing of blocks is independent, and can be performed in parallel. It may also be possible, especially if many more small blocks are used in the process, to infer which blocks need updates and to infer which atlases are most promising to be used in an update from results on neighbouring blocks. Updating of the segmentation could also be based on other decision fusion rules than averaging. So where this study shows the potential and a proof of principle of ALMAS, there are many avenues for further research that present themselves.

In summary, a new adaptive local multi-atlas based segmentation method was proposed and shown, in an experiment on heart segmentation in CT data, to perform better than single atlas methods and as good as an independent human observer. The proposed approach is expected to be applicable to other tasks, including segmentation of data that contains anatomical variations or pathology.

REFERENCES

1. T. Rohlfing, R. Brandt, R. Menzel, D. B. Russakoff, and C. R. Maurer, Jr., "Quo vadis, atlas-based segmentation?," in *The Handbook of Medical Image Analysis – Volume III: Registration Models*, J. Suri, D. L. Wilson, and S. Laxminarayan, eds., ch. 11, pp. 435–486, Kluwer Academic / Plenum Publishers, New York, NY, Aug. 2005.
2. W. Crum, L. Griffin, D. Hill, and D. Hawkes, "Zen and the art of medical image registration: correspondence, homology, and quality," *NeuroImage* **20**, pp. 1425–1437, 2003.
3. T. Rohlfing, R. Brandt, R. Menzel, and C. R. Maurer, Jr., "Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains," *NeuroImage* **21**(4), pp. 1428–1442, 2004.
4. R. A. Heckemann, J. V. Hajnal, P. Aljabar, D. Rueckert, and A. Hammers, "Automatic anatomical brain MRI segmentation combining label propagation and decision fusion," *NeuroImage* **33**(1), pp. 115–26, 2006.
5. P. Thévenaz and M. Unser, "Optimization of mutual information for multiresolution image registration," *IEEE Transactions on Image Processing* **9**, pp. 2083 – 2099, December 2000.
6. D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: Application to breast MR images," *IEEE Transactions on Medical Imaging* **18**, pp. 712 – 721, August 1999.
7. D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank, "PET-CT image registration in the chest using free-form deformations," *IEEE Transactions on Medical Imaging* **22**, pp. 120 – 128, January 2003.
8. I. Sluimer, M. Prokop, and B. van Ginneken, "Towards automated segmentation of the pathological lung in CT," *IEEE Transactions on Medical Imaging* **24**(8), pp. 1025–1038, 2005.