

A Fully Automated Approach to Segmentation and Registration of 3D Medical Image Data for Pulmonary Diagnosis

Alvin Ihsani¹, Jan Modersitzki², and Troy Farncombe³

¹Computing and Software, McMaster University, Ontario, Canada

²Institute for Mathematics, University of Lübeck, Lübeck, Germany

³Radiology, McMaster University, Ontario, Canada

ihsania@mcmaster.ca

Abstract. Molecular imaging is an important tool that has found widespread use in the diagnosis and observation of various diseases and has more recently been used in areas such as drug development in order to facilitate the observation and analysis of newly developed drugs. The amounts of data in drug development experiments may be very large due to the involvement of both spatial and temporal information of medical images. Imaging techniques can facilitate the analysis of this data by automating information extraction and providing meaningful results. We propose a fully automated approach to pulmonary diagnosis in the context of drug development experiments using image segmentation and registration techniques. In particular, we propose a modification of the Chan-Vese approach for a stable segmentation of the lungs which then serves as a starting point in a spatial alignment process. Our results demonstrate the feasibility and potential of the proposed approach.

Keywords: image segmentation, image registration, automated diagnosis

1 Introduction

Medical imaging is a relatively new technology that has found widespread use in the observation and diagnosis of various diseases and symptoms. More recently, this technology has penetrated such areas as drug development in order to facilitate the observation of the effects of newly developed drugs on various subjects.

Drug development experiments employing medical imaging techniques (such as CT and SPECT) may produce large amounts of data since spatial and temporal information spanning over a few weeks or months (e.g. 8 months) may be utilized to derive an informed conclusion of the effectiveness of the drug and its short and long term effects. The scans may be taken in regular intervals over the duration of the treatment and therefore corrections may need to be made to the data due to

- the position of the subject changing in consecutive scans,
- different machine artifacts over the period of treatment,
- artifacts introduced by motion such as breathing, heart-beats or other involuntary motions and
- physiological differences in the subject over the period of treatment.

When the amount of data in such experiments is large the automatic correction of the aforementioned artifacts becomes indispensable in order to minimize the time and resources needed to process the information and provide meaningful results.

This paper focuses on the automated correction of motion artifacts introduced by the placement of the subject in the scanning machine in consecutive scans.

1.1 Problem Description

The objective of the project presented in this paper is to facilitate the observation of the effects of a newly developed drug on the lungs of a subject over the period of treatment.

The approach taken by the McMaster Pre-clinical Imaging Facility to observe the effects of the treatment on the lungs is by tracking the development of the air flow and blood circulation in the lungs using SPECT (Single Photon Emission Computed Tomography) and CT (Computed Tomography) data acquired using a combined SPECT–CT scanner. This approach is commonly taken when looking for such diseases as cancer or pulmonary embolism [5,9,10].




Fig. 1. Visualization of the diffusion of the perfusion contrast agent (right) showing blood-flow, ventilation contrast agent (left) showing air-flow and fused data (centre) showing potentially unhealthy regions in the lungs.

In order to track air flow in the lungs a ventilation contrast agent (Tc-99m Technegas) is inhaled by the subject and its diffusion is tracked using a SPECT scanner. Similarly, a perfusion contrast agent (Tc-99m labeled macro-aggregate albumin) is injected in the bloodstream and its diffusion is tracked using the same scanning machine [5,9]. The air and blood diffusion data, “fused” together, provide information about healthy and unhealthy regions in the lungs [5,9,10]

as shown in Figure 1. As an example, cancer can be identified by regions in the lungs where the blood flow is high but air flow is poor.

The fusion of the ventilation and perfusion SPECT data and its interpretation is outside the scope of this paper. The problem treated here deals with the automatic (or unsupervised) alignment of the ventilation and perfusion SPECT data, since these are taken more than two hours apart, so that the data is ready for fusion and interpretation.

SPECT data provides only information about the diffusion of a contrast agent in the lungs and this may not at all be indicative of the shape of the lungs. Also, the structural differences between ventilation and perfusion SPECT may be large (see red and blue areas in Figure 1). This complicates the task of aligning the SPECT data meaningfully.

Using a combined SPECT-CT scanner, CT and SPECT data can be acquired almost concurrently (two minutes apart) and output a pre-aligned pair of ventilation (or perfusion) SPECT and CT. Therefore, to obtain the desired alignment between the two SPECT data, the corresponding pre-aligned CT data are used since they contain consistent information about the shape of the lungs and exhibit more structural consistency over different scans.

Since SPECT data contains information about the lungs then the CT data pair must be aligned using information from the chest region (i.e. lungs, heart, ribs and other proximal tissues), which in this paper are referred to as the *target region*. The intent is to remove far away objects or organs (head, limbs, muzzle holder etc.) that might obstruct the proper alignment of the chest between the two CT data. The *target region* is found using a novel segmentation technique which is a central step in the correct alignment of the ventilation and perfusion CT data.

1.2 A High-Level Description of the Alignment Process

At a high-level the alignment process designed in this project aims to automatically align the perfusion and ventilation SPECT by finding a transformation which aligns the corresponding perfusion and ventilation CT using the the *target region*. The alignment process is broken down into the following steps.

1. Find an approximate region of the lungs in the perfusion CT data using the perfusion SPECT data. Since the location of the lungs in the perfusion CT data is not known the perfusion SPECT can provide some “hints” due to the fact that it only contains information about the blood-flow inside the lungs.
2. Use a segmentation technique, which uses the region found in the previous step as an initial guess, to extract the lungs from the perfusion CT and lead to the *target region*.
3. Finally, a rigid model multi-level image registration method is used to align the *target region* in the perfusion CT to the ventilation CT. The resulting transformation is used to align the perfusion and ventilation SPECT data.

Figure 2 shows a sketch of the aforementioned steps.




Fig. 2. Diagram of the perfusion/ventilation alignment process.

2 A Detailed View of the Methods Used in the Alignment Process

This section describes the reasons behind each step of the alignment process as well as some of the mathematical details of the segmentation and registration methods.

2.1 Ellipsoid Fitting: Finding the Initial Region

In the first step of the alignment process, the aim is to find an approximate location where the lungs are situated in the perfusion CT scan. Because the region where the lungs may be situated is narrowed down the initial guess for the segmentation step is *stabilized*. Initially, the alignment process “does not know” where the lungs are in the perfusion CT, but because the corresponding perfusion SPECT contains only information from inside and closely around the lungs the SPECT can be used to indicate an initial region of the lungs in the CT as shown in Figure 3.

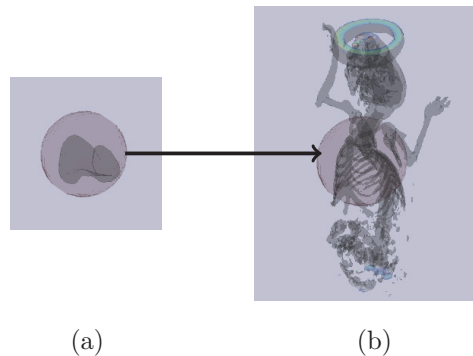


Fig. 3. A typical result of the initial region (a) on the perfusion SPECT and (b) on the perfusion CT.

In this step, the initial region is described by a smooth ellipsoid since it is an easy shape to describe mathematically. The shape of the initial region is not of importance since it is modified in the segmentation step.

While a simple thresholding of the perfusion CT data may seem sufficient at first this complicates the problem of finding a region of interest that does not include structures such as the arms and neck. For instance, a simple thresholding (e.g. at -300 HU) would include “fatty” structures in the arms of the mouse in addition to the lungs rendering the initial guess for the *target region* too large and making the result of the following segmentation step “less reliable”. In addition, it is a challenging task to find an appropriate thresholding value and a considerable amount of time needs to be spent analyzing each sample of data. In order to avoid these challenges, the ellipsoidal region bases its initial guess on the perfusion SPECT data which contains only lung information however small this may be.

The ellipsoid used in this project is three-dimensional and is described by the function

$$W(\omega, x) = \begin{cases} 1 & R < 0.5 \\ -\frac{\cos(2\pi R)}{2} + \frac{1}{2} & 0.5 \leq R \leq 1 \\ 0 & R > 1 \end{cases} \quad (1)$$

with

$$R = \sqrt{\left(\frac{x_1 - \omega_4}{\omega_1}\right)^2 + \left(\frac{x_2 - \omega_5}{\omega_2}\right)^2 + \left(\frac{x_3 - \omega_6}{\omega_3}\right)^2}. \quad (2)$$

This defines an ellipsoid parameterized by ω whose “density” increases smoothly towards the centre with $0 \leq W(\omega, x) \leq 1$.

The aim is to find the smallest ellipsoid enclosing a thresholded version of the perfusion SPECT data as shown in Figure 3(a). Considering that the highest concentration of contrast agent is exhibited in the inside of the lungs, the threshold value is chosen so as to remove small nonzero values outside the lungs (or “miscounts”).

Minimizing the objective function

$$F(\omega) = \|\omega_1^2 + \omega_2^2 + \omega_3^2\|_2 + \alpha \int_{\Omega} u(x)(1 - W(\omega, x))dx \quad (3)$$

the smallest ellipsoid enclosing the thresholded perfusion SPECT data

$$u(x) = \begin{cases} 1 & g(x) \geq \text{threshold} \\ 0 & g(x) < \text{threshold} \end{cases} \quad (4)$$

is found where $g : \Omega \rightarrow \mathbb{R}$ is the image and $\Omega \subset \mathbb{R}^d$ is the image domain with $d = 3$. The parameter $\alpha > 0$ is chosen empirically. In equation (3), the first term aims to minimize the radii of the ellipsoid while the second term aims to include as much data as possible inside the ellipsoid. The minimizer of equation (3) is found using the Gauss-Newton method available in FAIR [4].

2.2 Lung Segmentation: Refining the Initial Region

The initial region obtained in the step described in the previous section may be so large that it may enclose organs or objects of the perfusion CT data which might obstruct proper alignment of the chest. Conversely, the initial region may be so small (due to poor circulation in the lungs) that the information it contains may be insufficient to properly align the perfusion and ventilation CT. For these reasons, the initial region must be refined to a shape that is comparable to the size of the lungs as they appear in the perfusion CT (see Figure 4). The extraction of the lungs (or *target region*) can be performed using a segmentation method.

While a large variety of segmentation methods exist, including region growing, histogram-based, model-based and edge detection methods [1,6], a segmentation method based on the Chan-Vese model [13] was chosen. All of the aforementioned methods may be used for lung segmentation, however, the choice for

this project was based on the fact that a “stable” initial guess for the *target region* is obtained from the perfusion SPECT data in the step described in the previous section (i.e. a guess which lies either mostly inside the lungs or very proximally to their surface) and also on the requirement of avoiding the supervision of the segmentation and registration steps. Region growing methods may also have provided the desired result, however, this has not been tested.

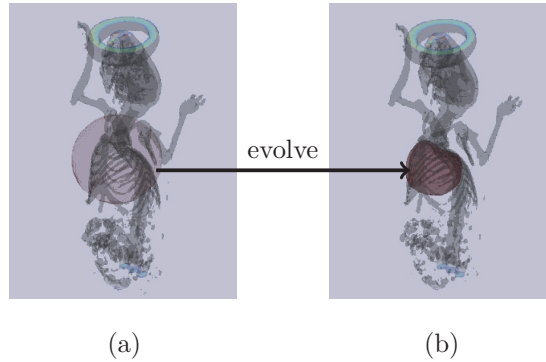


Fig. 4. The evolution from (a) the initial region to (b) the *target region* closely enclosing the lungs in the perfusion CT data.

A Novel Segmentation Method The Chan-Vese method [13] seemed to better fit the needs of this project since the image needs to be partitioned into only two regions (the lungs and the “rest”) and therefore a simple two phase piecewise constant model is sufficient. However, the Chan-Vese approach [13] did not provide the desired result (see Figure 5) and the energy functional was modified to

$$E(\phi) = \mu \int_{\Omega} |\nabla H(\phi(x))| dx + \nu \left| \mathbf{V} - \int_{\Omega} H(\phi(x)) dx \right| + \lambda \int_{\Omega} \|g(x) - c\|^2 H(\phi(x)) dx. \quad (5)$$

Here, $\mu, \nu, \lambda > 0$ are scalars which are chosen empirically, $\Omega \subset \mathbb{R}^d$ is the image domain with $d = 3$, $g : \Omega \rightarrow \mathbb{R}$ is the image, $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ is the level-set function [11], H is a smooth approximation to the Heaviside function as proposed in [13,14], \mathbf{V} is a constant representing a rough estimate of the volume of the lungs and c is a constant representing an estimated average intensity value (in this case in Hounsfield Units) inside the lungs in the CT data.

In equation (5), the first term aims to minimize the segmenting surface defined as $\Gamma = \{\mathbf{x} | \phi(\mathbf{x}) = 0\}$ [7,8,13], the second term asks that the volume inside Γ be as close as possible to the estimated volume \mathbf{V} and the third term asks that the region of the perfusion CT data g inside Γ be as close as possible to the estimated average intensity value c . The functional in (5) is minimized using a variational approach and the same discretization scheme proposed in [13].




Fig. 5. Coronal slices of 3D data showing (a) a typical initial guess and (b) a typical result of the Chan-Vese segmentation method with segmented area shown in red.

The main difference between our functional and that proposed by Chan and Vese in [13] is that our functional is only affected by intensity values of the data inside the segmenting surface while the functional proposed in [13] is affected by values both inside and outside the segmenting surface.

As a concluding step in the refinement of the initial region, an ellipsoid is fit on the segmented region using the same approach described in Section 2.1. The motivation behind this additional step is to stabilize the registration. The less information is used from the perfusion CT the more sensitive to noise the registration becomes and also other regions in the ventilation CT (e.g. fatty tissues in the arms) may match the intensity values in the lungs. The ellipsoid enclosing the segmented region is a simple way to include closely surrounding tissues, specifically the ribs, which may help in providing more information for the rigid registration due to their relatively rigid structure. This final ellipsoid will be referred to as the *target region*.

2.3 Registration: Aligning the Ventilation and Perfusion Data

In the last step of the alignment process, the aim is to use a registration method to find a rigid transformation that aligns the perfusion and ventilation CT by using the *target region*. For the purposes of this project (in small animal imaging) a rigid registration model is sufficient and was also recommended by the staff at the McMaster Pre-clinical Imaging Facility, however, for larger animals and humans an affine or elastic registration model may have to be used.

In particular, the objective functional

$$\mathcal{J}[y](\nu) = \frac{1}{2} \int_{\Omega} (\mathcal{T}(y(\nu, x)) - \mathcal{R}(x))^2 dx \quad (6)$$

was minimized in order to find a rigid transformation y parameterized by ν (allowing only for rotation and translation) that increases the similarity between the *target region* weighted perfusion CT data, \mathcal{T} , and the ventilation CT data, \mathcal{R} . The reason for the spatial filtering is to remove unwanted organs or objects (head, legs, muzzle holder etc.) in the perfusion CT data that might obstruct registration of the chest area between perfusion and ventilation CT.

The registration of the CT data is multi-level [2,3], where the two images have been registered at multiple gradually increasing resolutions. Multi-level image registration is an effective strategy to avoid local minima [2]. Functional (6) was minimized using the Gauss-Newton method available in FAIR [4].

Finally, after obtaining the optimal parameters ν , the rigid transformation is applied to the perfusion SPECT for alignment to the ventilation SPECT.

3 Testing the Alignment Process

Materials The materials provided by our clinical partners for this project consist of SPECT and CT samples of a mouse taken over a period of six weeks for a total of 24 samples. For each week two pairs of pre-aligned data were provided, namely ventilation SPECT and CT and perfusion SPECT and CT at a resolution of $256 \times 256 \times 256$ voxels. The size of each sample is $(58.88 \times 58.88 \times 58.88)mm^3$.

Experiment Setting To test the alignment process the ground truth was simulated on real-life data using a set of 60 randomly generated transformations with mean zero for both rotation and translation and standard deviation $\sigma_{rotation} = 0.3rad$ and $\sigma_{translation} = 7mm$. The simulated transformations are roughly 20 times larger than those found in real-life data. The results of the segmentation and alignment process for the real-life data were also visually inspected for correctness.

The perfusion SPECT data, after some observation, was thresholded at intensity value above 300 counts. The parameter α in the ellipsoid fitting objective function (3) was set to $\alpha = 1$.

The parameters for the minimization of the novel segmentation functional (5) were found empirically and set to $\mu = 0.5 \cdot 255^2$, $\nu = 10^5$, $\lambda = 1$ with a fixed time-step $\tau = 0.046$. The segmentation was performed at $16 \times 16 \times 16$ voxels.

The resolutions used for the multi-level registration are $2^i \times 2^i \times 2^i$ voxels with $i = \{3, 4, 5, 6\}$.

Alignment Process Results The results obtained for the simulated set of transformations are summarized in Table 1 in terms of the absolute error between the simulated transformation and the result of the alignment process showing both the mean and standard deviation for translation and rotation.

In order to get a better idea of the “reliability” of the alignment process it is also important to look at the *target regions* obtained for each week. Table 2 summarizes these results.

The alignment process takes on average 5 iterations (0.5 seconds) for the ellipsoid fitting steps, 15 iterations (20 seconds) for segmentation and 3–6 iterations (3 minutes and 20 seconds) for the multilevel registration step. The machine used is a standard PC with 2 GB of RAM and a 2 GHz AMD Turion X2 processor running in Matlab.

	Mean		
Rotation (rad)	$-2.0447 \cdot 10^{-5}$	$-3.5597 \cdot 10^{-5}$	$-2.2869 \cdot 10^{-5}$
Translation (mm)	$-7.9440 \cdot 10^{-5}$	$-2.4414 \cdot 10^{-4}$	$-1.2205 \cdot 10^{-3}$

	Standard Deviation		
Rotation (rad)	$2.7496 \cdot 10^{-5}$	$3.3461 \cdot 10^{-5}$	$3.4985 \cdot 10^{-5}$
Translation (mm)	$1.2506 \cdot 10^{-3}$	$1.3136 \cdot 10^{-3}$	$1.4992 \cdot 10^{-3}$

Table 1. Absolute error between the simulated transformations and alignment process results summarized in terms of mean and standard deviation.

Radii (mm)	Week					
	1	2	3	4	5	6
ω_1	6.4212	4.8325	4.1047	4.2579	2.4686	2.0770
ω_2	7.1103	5.5471	6.1601	4.4618	3.5114	2.0770
ω_3	5.8149	3.5433	3.4566	3.0731	2.1633	0.7248

Table 2. The target region radii may vary for each week depending on the perfusion SPECT thresholding value (here ≥ 300 counts).

4 Discussion and Conclusions

The novel segmentation method proposed in this paper has shown to be effective for the purposes of lung segmentation and has provided the desired results. The objective functional (5) is not convex since the result of the segmentation depends on the starting guess, however, by stabilizing the starting guess using the ellipsoid fitting method described in Section 2.1 the result of the segmentation has shown to be very consistent (see Figure 6).

Furthermore, the overall alignment process has shown to be reliable since the absolute error in the resulting transformation was in the order of 10^{-3} for both rotation (rad) and translation (mm) given that the target region changes considerably in each week (see Table 2) and the smallest cell size used for registration is $\approx 0.92mm$ at the edge. The down-sampling of the data for both segmentation and registration is done to decrease computation time but also to test the sensitivity of the overall alignment process against varying resolutions. This has shown to work quite well for very “rough” samples of data as presented here as well as for finer data without changing parameters.

The novel segmentation method proposed in this paper achieves the desired result using only two additional pieces of information, namely an estimated lung volume and average intensity value in the lungs. While it is hard to estimate these two parameters accurately, from the tests carried out so far it seems the segmentation method will provide the desired result even for very rough estimates. For instance, in this project the lung volume was estimated as one-sixteenth of the total volume of the data which is a very rough estimate. In reality, the lung volume may be estimated using statistical data based on species, age, weight and other factors. In addition, the average Hounsfield unit values used for the lungs

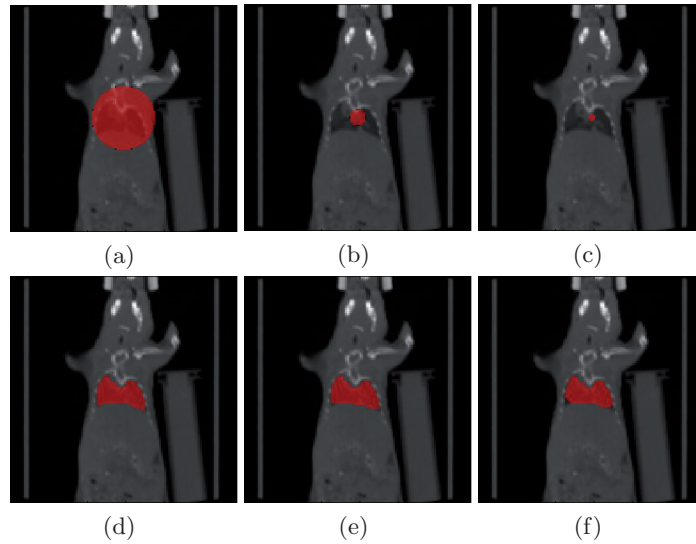


Fig. 6. Typical initial guesses 6(a)–6(c) and their respective results 6(d)–6(f) shown in coronal slices of 3D real-life data. The area in red represents the region inside the segmenting surface.

were set between -100 HU and -400 HU but the segmentation method still provides the desired result even though the actual average value is -150 HU. The segmentation parameters μ , ν and λ in equation (5) remain unchanged for the samples of data presented here. In fact the parameter μ has not been changed for other (synthetic) data sets as well, however, the parameters β and λ may need to be tweaked very slightly from the values presented above to obtain the desired result.

One limitation of the alignment process is the thresholding of the perfusion SPECT data for ellipsoid fitting. In the future this step will be substituted by the Chan-Vese segmentation model which, after some preliminary tests, has shown to provide much better results without the need of estimating a threshold value. Furthermore, the resulting level-set function from the Chan-Vese segmentation method can be used as the initial guess for the novel segmentation functional (5).

Other considerations for the improvement of the alignment process include improving the numerical optimization scheme of the segmentation step, performing multi-scale segmentation of the perfusion CT data and using non-rigid registration models.

References

1. D. L. Pham and C. Xu and J. L. Prince. A Survey of Current Methods in Medical Image Segmentation. In *Annual Review of Biomedical Engineering*, volume 2, pages 315–338. 2000.
2. E. Haber and J. Modersitzki. A Multilevel Method for Image Registration. *SIAM J. Sci. Comput.*, 27(5):1594–1607, 2006.

3. J. Modersitzki. *Numerical Methods for Image Registration*. Oxford University Press, 2004.
4. J. Modersitzki. *FAIR: Flexible Algorithms for Image Registration*. SIAM, 2009.
5. J. Palmer and U. Bitzén and B. Jonson and M. Bajc. Comprehensive Ventilation/Perfusion SPECT. *The Journal of Nuclear Medicine*, 42:1288–1294, 2001.
6. J. S. Suri and D. Wilson and S. Laxminarayan. *Handbook of Biomedical Image Analysis: Volume 1: Segmentation Models Part A (Topics in Biomedical Engineering International Book Series)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.
7. L. Ambrosio and V. Tortorelli. Approximation of Functionals Depending on Jumps by Elliptic Functionals via Γ -Convergence. volume 28, n.3, pages 213–221, 1998.
8. L. C. Evans and R. F. Gariepy. *Measure Theory and Fine Properties of Functions*. 2000.
9. M. Bajc and B. Jonson. Ventilation/perfusion SPECT an essential but underrated method for diagnosis of pulmonary embolism and other diseases. *European Journal of Nuclear Medicine and Molecular Imaging*, 2009.
10. O. Schillaci. Hybrid SPECT/CT: a new era for SPECT imaging? *European Journal of Nuclear Medical Molecular Imaging*, 32(5):521–524, 2005.
11. S. Osher and J. A. Sethian. Fronts Propagating with Curvature Dependent Speed: Algorithms Based on Hamilton-Jacobi Formulations. *Journal of Computational Physics*, 79:12–49, 1988.
12. T. F. Chan and L. A. Vese. A level set algorithm for minimizing the Mumford-Shah functional in image processing. In *VLSM '01: Proceedings of the IEEE Workshop on Variational and Level Set Methods (VLSM'01)*, page 161, Washington, DC, USA, 2001. IEEE Computer Society.
13. T. F. Chan and L. A. Vese. Active Contours Without Edges. *IEEE Transactions on Image Processing*, 10(2):, 2001.
14. Zhao, Hong-Kai and Chan, T. and Merriman, B. and Osher, S. A variational level set approach to multiphase motion. *J. Comput. Phys.*, 127(1):179–195, 1996.