

Volumetric Quantification of Metastatic Burden from SPECT/CT Images

Hanna Carlsson

November 12, 2012

Abstract

Bone scintigraphy images are used to investigate the presence of metastases in patients with prostate cancer. An analysis of these images indicates the proportion of cancer in the bones; the Bone Scan Index (BSI). This value is an important prognostic factor used to predict the future for the patients. It is common to extend the examination with a combination of SPECT and CT images to be able to study possible metastases in more detail. The aim of this thesis is to integrate the information from SPECT and CT images to get a more accurate calculation of the BSI.

The strategy has been to use the analysis of the bone scintigraphy images as an initialization of a segmentation of metastases in the SPECT images. The CT images are used to produce a simple segmentation of the bones.

The problem is divided into two main parts; image registration and segmentation of metastases. The image registration is needed to align the coordinate system of the bone scintigraphy images and the SPECT images. The Morphon method has been chosen and the results are good; twelve of fifteen tested registrations are classified as successful. A combination of a more robust method to find start guesses and a more generous transformation model would probably improve the results even further. Seeded region growing has been chosen as the segmentation algorithm. An implementation with an automatic termination criterion has been created and the results seem promising.

The conclusions are that the Morphon method works fine for registration of the images and that the strategy of initializing a segmentation gives good results. In this work, the size of the segmentation was not used to update the BSI but the results can still be useful. Another goal was to facilitate navigation between different types of images by aligning them to one coordinate system. A user interface has been created to reduce the amount of time spent by doctors and biomedical scientists to navigate through the images searching for possible metastases.

Preface

This project is a master thesis in engineering physics at the Faculty of Engineering at Lund University. The thesis has been carried out at EXINI Diagnostics AB and I would like to thank everyone at the company for giving me such a welcoming feel, for help and for interesting conversations. Special thanks go to my supervisors; Karl Sjöstrand and Lars Edenbrandt at EXINI Diagnostics and Kalle Åström and Niels Christian Overgaard at the Centre for Mathematical Sciences at Lund University. Your thoughts and comments have been extremely valuable for me. I would also like to thank Anton Sörensen, a fellow master thesis student at EXINI from Chalmers University of Technology, for many meaningful discussions. It has been a pleasure getting to know you all!

Contents

1	Introduction	1
1.1	Aim of the thesis	1
1.1.1	Milestones	2
1.2	Related work	3
2	Background	4
2.1	Prostate cancer	4
2.2	Scintigraphy	4
2.2.1	SPECT	5
2.3	CT	6
3	Theory	7
3.1	Image registration using the Morphon method	7
3.1.1	Displacement field estimation	8
3.1.2	Regularisation	14
3.1.3	Deformation field accumulation	17
3.1.4	Deformation of image	18
3.1.5	Registration example of 2D image	19
3.2	Principal component analysis	20
3.3	Segmentation using region growing methods	22
3.3.1	Finding or choosing seed points	22
3.3.2	Increasing the region	23
3.3.3	Termination	23
4	Methods	25
4.1	Image registration	25
4.1.1	Test of the Morphon method implementation	26
4.2	Segmentation	28
5	Results	29
5.1	Image registration	29
5.2	Segmentation	33
6	Discussion	35
6.1	Image registration	35
6.1.1	The accumulation step	36
6.2	Segmentation	38
6.3	Evaluation of milestones	38

6.4	Conclusions	41
6.5	Future work	41
7	Bibliography	42
A	Image registration results	44
B	Segmentation results	52

Chapter 1

Introduction

1.1 Aim of the thesis

Prostate cancer is one of the most common forms of cancer and there are many patients in need of good care, where a clear diagnosis can be determined and the progress of the disease can be monitored and limited. Prostate cancer originates from a tumour in the prostate. A common result of the disease is that the tumour eventually spreads to the bones, often to the pelvic area or the region of the lower back. These metastases can be seen as spots with high intensity in a bone scintigraphy image. When a bone scintigraphy examination is performed, the doctor could choose to extend the investigation with SPECT images and/or CT images. Thanks to such three-dimensional images it is possible to get a more thorough view of the metastases and their locations in the body.

SPECT and CT images are useful tools to get a visual complement to the planar scintigraphy images, but they have not yet been included in a fully automatic and quantitative analysis of how far the cancer has spread. The Bone Scan Index (BSI) is an estimate of the proportion of the skeleton that has been affected by metastases. This value is calculated from analysis of the bone scintigraphy images and general knowledge of the thickness of different bones in the body. The idea with this thesis is to investigate the possibilities to include the three-dimensional information from the tomography in this analysis to get an improved BSI.

Another goal is to simplify the overview of the combination of bone scintigraphy images with tomography images like SPECT and CT. Computer programs visualize the three-dimensional images with three cross-sectional images and a click somewhere in one of those images updates the view with three new images. However, the connection with the scintigraphy images is missing and the person investigating the result of the examination has to scroll through the layers of images to find what he or she believes is the corresponding object of interest. A correlation between the different types of images could make it possible to identify a location of interest in one of the planar images and directly find the three-dimensional correspondence (or the other way around). This would facilitate the visual work for doctors handling image combinations like the ones described in this section.

Today, doctors and scientist seek more independent, quantitative evaluation

methods that could be used for diagnosis of various types of diseases. It is desirable both to automate and to improve the quality of the examination process, and to be able to ensure that a patient gets the same prognosis independent of the doctors involved. The already existing computer aided system for finding out how far the cancer has spread is one step in this development and the aim of this thesis is to improve the result of the analysis even further.

1.1.1 Milestones

Some milestones were set to reach the aim of the thesis. These milestones can be seen below and will be followed up in Section 6.3. The two main parts of the thesis have been considered to be image registration between scintigraphy images and a corresponding SPECT volume and segmentation of possible metastases. The explanation of the work will follow this partition.

1. *Two-dimensional projections of SPECT images*
To find the connection between different types of images used in the examination of prostate cancer patients, a comparison is needed. To save computational time it would be desirable to compare bone scintigraphy images and SPECT images in two dimensions. The first milestone is therefore to create some 2D projection of the SPECT volume. The relationship between SPECT and CT is saved in the DICOM file format. In other words, the connection between the bone scintigraphy images with the SPECT volume would be enough to align all types of images to each other.
2. *Image registration of SPECT and bone scintigraphy images*
Image registration is the key to connect images. This step is needed to reach the goal of parallel navigation through the cross-sectional images and the bone scintigraphy images. It is also essential if the analysis already present on the bone scintigraphy images shall be used as initialization for the process of finding hotspots in three dimensions.
3. *Validation of the result*
A validation of the image registration is needed to find out whether the connection found gives an accurate result. Most importantly it is necessary to test the registration on a wide range of image pairs to try to make sure that the method works fine in all possible conditions.
4. *User interface for navigation through the images*
To create the easy overview of images described in the aim of this project, some kind of user friendly interface is needed.
5. *Segmentation of metastases in three dimensions*
It seems reasonable to make use of the existing information from analysis of the two-dimensional images. After the image registration has been completed, the preparations for initializing the three-dimensional segmentation of metastases are done. The segmentation is required to find out the size of possible metastases.
6. *Distribution of metastases*
A good measurement of the size of metastases is essential since this would

be the foundation for the updated value of the BSI. The size will be given in number of pixels and the result has to be transformed into some unit that could be compared to the BSI from today.

7. *Three-dimensional visualisation of metastases*

The ideal result is to find a good segmentation of possible metastases from the combination of bone scintigraphy images and a SPECT volume. A segmentation of the bones from the CT image could then be used to view the result in a three-dimensional figure showing the bone of the patient in some transparent colour and possible metastases as distinct hotspots.

1.2 Related work

The Bone Scan Index was first mentioned in [1]. Here, a method for determining the BSI was described and a correlation between the BSI and the progress of the disease was found. The results are consistent with the indications found in [2], where the BSI value was investigated on a larger data set. The method of determining the BSI is based on a visual inspection of bone scintigraphy images. A semi-automated segmentation method, where a user identifies metastatic regions and inserts seed points at those locations, were tried to investigate the fraction of cancer in the skeleton. A region growing method uses the seed points to calculate a segmentation of the metastases. The correlation between the fraction found and the BSI values from visual inspection was high, but the results were consistently lower than the conventional visual estimates. The benefits of a semi-automated method are the reproducibility and the possibility to detect small variations in follow-up scans. The dependency of the user was eliminated in [3], where image-processing techniques and artificial neural networks are used to interpret the scintigraphy images. A completely automated computer-assisted diagnosis system was created showing good results. The method is considered to have significant potential and to be of clinical value as a decision support tool. This theory is the foundation for the analysis of bone scintigraphy images and the results that will be used in this project.

There has been no related work found on the task of estimating the BSI value from SPECT/CT images, but the problem of searching for segmentations of metastases can be found. A similar problem to the one in this thesis is the one given in [4]. A registration between PET and CT images is estimated from cross-correlation and candidate points from the PET image are fed to a seeded region growing algorithm to find the boundary of lesions in the CT image. Another method of finding a segmentation of metastases in the bones has been proposed in [5]. Here an iterative thresholding method has been implemented for finding tumours in SPECT images.

Even if the specific aim of this thesis is new, the subproblems of image registration and segmentation appears frequently in different kinds of published material. An overview of different image registration methods can be found in [6] and a description of a variety of available methods in medical image segmentation are given in [7].

Chapter 2

Background

2.1 Prostate cancer

Prostate cancer is the most common form of cancer in Sweden today. Nearly 10 000 men were diagnosed in 2010 and around 75 000 are living with the disease. The number of affected persons has been increasing for decades and the causes are not yet fully understood. [8]

As all forms of cancer, prostate cancer involves an unregulated growth of cells forming malignant tumours. In an early stage this might not cause any symptoms at all but as the tumour reaches a larger size it might generate urinary problems and/or erection problems. Sometimes symptoms are not present before the cancer has spread outside the prostate. The most common ones are then frequent pain in the lower back, hips, or upper thighs due to metastases in the bones or in the lymph nodes. [8, 9]

It is possible to detect prostate cancer even if there are no symptoms to be seen. A blood sample could be checked for prostate-specific antigen (PSA). If the levels of PSA in the blood exceed some given limit, further examination is executed. Since prostate cancer often is a slowly growing form of cancer and that diagnosis among men under 45 is rare, few people live a shorter life because of their disease. There are though many courses of the disease and a general prognosis is hard to give. Desirable is to identify the right type of treatment for each patient and to slow down the progress as much as possible. [8, 9]

2.2 Scintigraphy

Scintigraphy is a form of functional scanning test for diagnostic purposes in nuclear medicine. When it comes to prostate cancer, scintigraphy is an effective tool to find out whether the cancer has spread to the bones in the form of metastases. The examination could be used to follow the progress of the disease over time or the response to some kind of treatment. [8]

Before the examination starts, a radioactive isotope is injected into the blood and the isotope spreads in the body. Because of the cancer cells' unregulated growth, bone close to metastases will absorb more of this radioactive material than ordinary growing cells. When a gamma camera then detects the photons

from the decay, such abnormalities in the bone will be visible in the image as spots with higher intensity. [10]

Bone scintigraphy is the type of scintigraphy that has been used to acquire the two-dimensional images in this thesis. An example can be seen in Figure 2.1.




Figure 2.1: Bone scintigraphy images of a patient, anterior to the left and posterior to the right.

2.2.1 SPECT


The single-photon emission computed tomography (SPECT) technique is another type of scintigraphy. An injected radioactive material is used for detection with a gamma camera as described above. The main difference from the bone scintigraphy examination is that the camera is moved around the body to create images from multiple angles. These images are then put together by a tomographic reconstruction algorithm to create volumetric cross-sectional slices through the patient. This means that SPECT presents three-dimensional images of the part of the body that is examined instead of planar images. [11]

It is common to perform both a bone scintigraphy and a SPECT examination at the same time on prostate cancer patients. The planar bone scintigraphy is used to detect suspected metastases and the three-dimensional SPECT image to study these areas in more detail. Typically the bone scintigraphy first generates two planar images showing the whole body of the patient; one anterior (from the front) and one posterior (from the back). Then a doctor or a biomedical scientist determines the need of an additional SPECT image. If the choice is to continue, the interesting part of the patient is examined again using either the SPECT technique or both SPECT and Computed Tomography (CT). The primary use of an additional CT image is to get a more detailed view of the location of a suspected metastasis.


2.3 CT

Computed tomography (CT) is used in a wide range of areas. In medicine, CT images provide structural information about different parts of the body. X-ray beams are sent through a patient from different angles. As with SPECT, the resulting images are processed to generate a three-dimensional volume, but in this case the different intensities correspond to various abilities to block the X-ray beam in different parts of the body. CT images have a high contrast and high resolution and it is possible to distinguish tissues that differ in physical density by less than 1%. These characteristics make it possible to, for example, distinguish the bones in an image. For a cancer patient a CT examination could be useful to detect or confirm tumours, provide information of size and location, guide a biopsy or to determine whether the disease is responding to treatment or spreading further. [12, 13]

Figure 2.2 shows one cross sectional image from a SPECT volume and one from CT. SPECT generates images indicating the function of some interesting part of the body while CT images give a detailed view of the anatomical structure of the area.



(a) SPECT



(b) CT

Figure 2.2: Cross-sectional images from SPECT and CT. Notice the difference in type of information; SPECT images provide functional information, while CT images give anatomical information. The combination can be used to find and locate metastases in the examination of a cancer patient.

Chapter 3

Theory

3.1 Image registration using the Morphon method

Image registration is the process of transforming different sets of data into one common coordinate system [14]. The goal is to find the deformation that will transform a source data set to match a target data set. The need for image registration in the medical area is frequent. Data sets from an examination of a patient could be acquired at different times or by different image modalities. The combination of a bone scintigraphy with SPECT and/or CT is common, and an example of where image registration is a useful tool to be able to align sets of data to each other.

The Morphon method, first presented in [15] and described further in for example [16], is one technique to perform image registration. The theory is applicable to problems in different dimensions; one, two or three, and allows non-rigid transformations. The two images that are to be registered are assumed to have parts with similar information. Comparisons between source and target focus on structure, which makes the procedure independent of global intensity variation. This is often necessary when working with images from different medical image modalities. Figure 3.1 shows a bone scintigraphy as source image and a 2D projection of a SPECT image as target. The deformed source and the target image should look similar and share coordinate system after the registration process is completed.

The procedure to find the transformation from source to target is an iterative process consisting of a number of steps. The algorithm starts on a coarse resolution scale and proceeds to finer scales with the following steps performed on each scale.

- *Displacement field estimation*

Here, a displacement field between the source and the target images is sought. The field has a value for every pixel representing a movement from source towards target. The purpose of this field is to deform the source so that the similarity to the target is increased.

- *Regularisation*

This part defines the transformation model of the problem, specifying which types of deformations that should be allowed. Regularisation is




Figure 3.1: Bone scintigraphy as source (left) and SPECT as target (right). The source is created from the sum of the anterior and the posterior scintigraphy images and the target from the sum of all slices in the SPECT volume from an anterior view. The aim is to find an image registration from source to target.

needed to create a deformation field with smooth changes. Without this step, the source image is likely to be disrupted due to locally divergent displacement estimates.

- *Deformation field accumulation*
The displacement field from each iteration is added to an accumulated deformation field to avoid the smoothing effect from repeated interpolations of the source image.
- *Deformation of image*
In the end of each iteration, the original source image is transformed according to the accumulated deformation field, and the result is used as source in the next displacement field estimation.

After a number of iterations on each resolution scale, the source reaches the target and the algorithm terminates. The transformation needed to match the source image with the target image is found and the image registration is complete.

3.1.1 Displacement field estimation

The local displacements are estimated by computing the optical flow between source and target. Optical flow is the pattern of apparent motion of objects, surfaces, and edges in a visual scene [17]. In this case it means trying to find the displacement of structures like lines and edges in the target image and the source image. There are various ways to do this but the method used in the Morphon algorithm focuses on the difference in local image phase.

To compute the image phase, a number of quadrature filters are used. The idea with these complex filters is that the real part of the filters can detect lines and the imaginary part can detect edges. The size of the filter corresponds to how large or small differences the filter can trace between source and target. This specific property is the reason for rescaling the images during the algorithm. Instead of using filters of bigger size, which would require a longer computational time, the images are subsampled to a more coarse scale and larger differences can still be found.

A one-dimensional example is introduced to illustrate the different steps in the Morphon method. In this example the source is a simple signal of ones and zeros with one line and one edge. The target has been created moving the source signal a short distance to the right. This means that there are identical parts of source and target in this specific case. Source and target can be seen in figure 3.2.

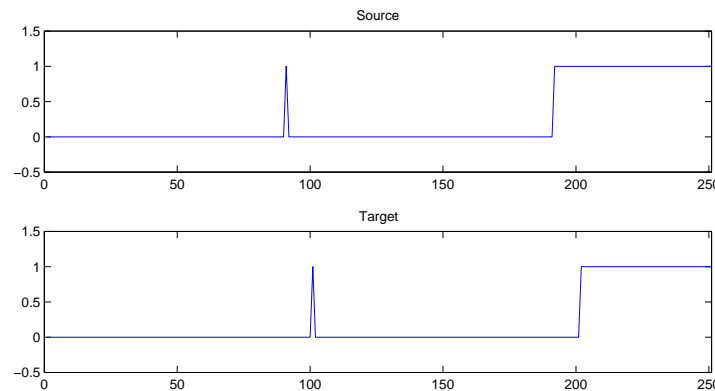


Figure 3.2: The figure shows two one-dimensional signals (source above and target below), used to illustrate image registration.

Filtering of a one-dimensional signal can be done with a lognormal filter. A lognormal filter in frequency and spatial domain can be seen in Figure 3.3. The real part of the lognormal filter in the spatial domain shows a line filter and the imaginary part an edge filter.




Figure 3.3: Lognormal filter in frequency (left) and spatial domain (right). The filter to the left is given by $\hat{f} = 0$ for $x \leq 0$ and $\hat{f} = e^{-\frac{4}{1.82 \ln 2} \ln^2(x/0.15)}$ for $x > 0$.

Using the complex filter from Figure 3.3 on a one-dimensional signal gives a response with a phase that indicates the type of event found and a magnitude that increases when the signal neighbourhood fits the filter well. The type of event found could be a positive line, a negative line, a positive edge, a negative edge or something in between.

Figure 3.4 illustrates the connection between the phase and the type of event found. The colour of each event is incorporated in a purpose specific colour scale that has been created to visualize the filter responses of images. The colour map is based on the principle of the HSV colour scale (*hue, saturation and value*) [18]. Red and blue are associated with the real part of the filter output, while magenta and green represents the imaginary part. This means that blue and red indicates positive and negative lines respectively, and magenta and green represents positive and negative edges. Unlike with the HSV colour map the colours fade to black in between the angles to intensify pure lines and edges.




Figure 3.4: Connection between phase and event type.

Filtering of a signal h will generate outputs of the following form for every position x of the signal

$$q = (f * h)(x) = A(x)e^{i\theta(x)}, \quad (3.1)$$

where f is the quadrature filter, $A(x)$ the magnitude of the filter output and $\theta(x)$ is the phase of the filter output. The result of filtering the one-dimensional example in Figure 3.2 can be seen in Figure 3.5.

It is possible to confirm that the local phase corresponding to the edge actually is zero (real part) and the phase corresponding to the line is $\pi/2$ (imaginary part). It is also clear that the certainties have their largest values where these structures are located, apart from some edge effect that can be seen to the right. This edge effect can be eliminated by manually setting the certainties at a distance of half the filter size from the edge to zero.

Combining the filter responses from source and target, by forming the complex valued product of the filter output from the source signal and the complex conjugate of the output from the target signal, will give a new output as follows

$$q_S q_T^* = A_S(x)A_T(x)e^{i\Delta\theta(x)}. \quad (3.2)$$

Here q_S and q_T denotes the filter output from the source and target respectively. The complex conjugate of q_T is denoted q_T^* . The magnitude of the filter responses from source and target respectively is represented by $A_S(x)$ and $A_T(x)$, while $\Delta\theta(x) = \theta_S(x) - \theta_T(x)$ is the local phase difference. The product in Equation 3.2 makes it possible to identify the phase difference between source and target and a product of the outputs magnitude. The phase difference will turn out to be proportional to the sought displacement estimates d for every

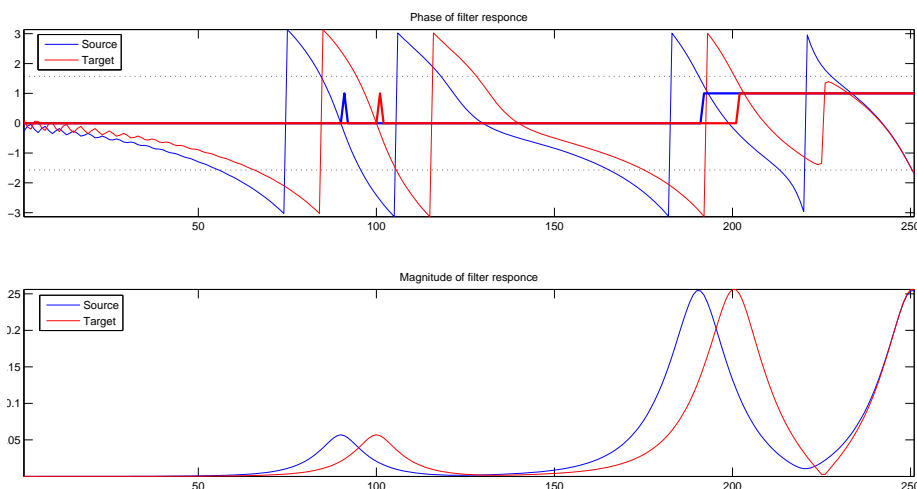


Figure 3.5: Above, phase of filter responses from source and target together with the signals. Below, corresponding certainty measurements. The markings show $-\pi/2$ and $\pi/2$. A positive line should give the phase zero, and a positive edge the phase $\pi/2$. The result can be verified in the upper part of the figure.

filter and position of the signal, and the combined magnitude $A_S(x)A_T(x)$ will correspond to an estimation of how certain each of those local displacement estimates are.

Some requirements on source and target are needed to make the local phase differences useful as displacement estimates. First, a perfect registered pair of signals is assumed to have the same filter responses and therefore no phase differences. If the signals are close but not equal a small difference in phase must occur. The larger spatial difference between source and target, the larger this phase difference becomes. Eventually, this phase difference will be affected by a shift from $-\pi$ to π (see Figure 3.5) and the difference is not certain to increase anymore. Until this point, the local phase difference must be proportional to the spatial difference between source and target

$$d \propto \Delta\theta(x). \quad (3.3)$$

The user of the Morphon method has to make sure that the differences between source and target are small enough, for the relation in Equation 3.3 to hold. Initially some start guess could be needed. Further on, it is the size of the filters that affects the spatial differences allowed. This is managed by rescaling the signal (subsampling of an image), as discussed earlier in this section.

For the given one-dimensional example, the phase differences calculated from the filter responses from source and target and the combined certainty measurement can be seen in Figure 3.6.

A simple two-dimensional image pair is introduced as an example in the explanation of the 2D version of the Morphon method, see Figure 3.7.

For images the idea of the lognormal filter is incorporated in quadrature filters. The one-dimensional line and edge detection filters are placed in a specific direction. The structure of such a filter can be seen in Figure 3.8.




Figure 3.6: The figure shows phase differences (above) and total certainty estimates (below). Notice the connection between the certainty measurements of higher value and the locations of the events found (apart from the edge effect to the right). The phase differences are not directly linked to a specific spatial difference, but if the spatial differences between source and target are small enough there is a proportionality to the spatial movement needed to bring the images closer to one another. This means that it is the sign of the phase differences that gives the direction of the desired movement.




Figure 3.7: The figure shows two images used to illustrate the two-dimensional image registration with the Morphon method. The source image consists of two lines with different orientation and one rectangle. The target is translated three pixels to the right from the source. The small difference is chosen to eliminate the need of a start guess.




Figure 3.8: Example of a quadrature filter in the spatial domain. Note the one-dimensional line and edge detection filters along the direction of the filter.

The filters are applied in four different directions to be able to detect the structure needed; from the left, diagonally from above/left, from above and diagonally from above/right. The procedure described above will accordingly give one displacement estimate per pixel and filter direction, or in other words; after applying the filter set to an image there will be as many displacement field estimates as the number of filters in the set. The result for the source and target in Figure 3.7 is visualized in Figure 3.9 and Figure 3.10. Figure 3.9 shows the phase of the filter responses for the source on the upper row and for the target on the bottom row. The filter directions vary in the mentioned order from left to right. The certainty measurements corresponding to the phase differences are given in Figure 3.10. Studying both figures together makes it possible to find areas with higher certainty and identify the type of event at that location.




Figure 3.9: The phase from all filter responses for source (top) and target (bottom row). The colour corresponds to the type of event found.




Figure 3.10: Magnitudes, or certainty measurements, from all filter responses for source (top) and target (bottom row). Brighter areas have higher certainties.

The four estimated displacement fields need to be combined into one field somehow. Solving the following weighted least squares problem for each iteration k will give the field d_k that in each direction \hat{n}_i of filter i yields the result closest to the computed displacement estimates $d_{k,i}$

$$\min_{d_k} \sum_i [c_{k,i}(\hat{n}_i^T d_k - d_{k,i})]^2. \quad (3.4)$$

The certainty measurements $c_{k,i}$, derived from the magnitude of Equation 3.2, are used as weights and the equation is minimized with respect to d_k . The solution provides one component of the movement in x and one in y for each pixel. The result is the total estimated displacement field for the current iteration and the information needed to shift every pixel towards the target is found.

The certainty measurements are combined into one by addition of the estimated certainty fields for each filter direction

$$c_k = \sum_i c_{k,i}. \quad (3.5)$$

3.1.2 Regularisation

The computed local displacement estimates can be very different across the image, and using these displacements without regularisation might tear the image apart. Figure 3.11 shows an example of a deformed image with and without regularisation. Figure 3.12 shows the regularisation of the displacement field from the first iteration on a subsampled version of the images in Figure 3.11.




Figure 3.11: Result after registration without regularisation to the left, and regularisation allowing shift in scale and translation to the right.




Figure 3.12: To the left, an estimated displacement field before regularisation. To the right, the same displacement field after regularisation allowing a shift in scale and translation.

It is the regularisation step that defines which type of transformation that is desirable going from source to target. It can be anything from a simple translation to a non-rigid transformation. Deciding what is allowed determines the degree of freedom for the deformation. The calculations cannot be described in general since the transformation model has to be chosen to match the specific problem. Four different models will be described below; translation, translation and scaling, affine transformation and elastic transformation.

3.1.2.1 Translation

The first example will follow up the one-dimensional source and target signal from Figure 3.2. The computed phase differences and the corresponding certainties will together give the sought values proportional to the displacement between source and target. In this case the only transformation allowed is translation of the whole signal. This reduces the regularisation step to a calculation of the average of the estimated displacements with the certainties used as weights. The formula for a weighted mean is

$$\hat{x} = \frac{w_1x_1 + w_2x_2 + \dots + w_nx_n}{w_1 + w_2 + \dots + w_n}, \quad (3.6)$$

where $\{x_1, x_2, \dots, x_n\}$ in this example is the data set of displacement estimates d_k and $\{w_1, w_2, \dots, w_n\}$ the certainty weights c_k . The result is the regularised displacement estimate for each location. Here, all locations share the same value. The same procedure is carried out for the two-dimensional example given, since the desired transformation was translation also in this case.

3.1.2.2 Scale and translate

For reasons explained in Chapter 4.1 the chosen model for aligning bone scintigraphy images with SPECT volumes is a difference in scale and a translation along the x- and y-axis between source and target. The aim of the regularisation is then to replace the calculated deformation field with a new one, representing scaling and translating, as close to the estimated field as possible. The deformation given by the Morphon method indicates what pixel value in the source to place where in the target to create a deformed version of the source. Therefore the model of the transformation is to first scale and then translate the target image to reach the source. The corresponding formula is

$$s \cdot d_T + e_x \cdot t_x + e_y \cdot t_y = d_S, \quad (3.7)$$

where s is the scale, t_x and t_y the translation in x and y , e_x and e_y are unit vectors and d_T and d_S the pixel locations in target and source respectively. The pixel locations are represented with column stacked vectors. Note that the local displacement estimates $d = d_S - d_T$.

The regularised field is the one minimizing the difference to the deformation field calculated earlier, that allows all sorts of deformations. The sought field can be found solving the following least squares problem

$$\min_{s, t_x, t_y} \|d_S - (s \cdot d_T + e_x \cdot t_x + e_y \cdot t_y)\|^2, \quad (3.8)$$

with the same notations as before.

Rearranging the expression from Equation 3.8 into matrix form makes it possible to use the standard solution of a linear least squares problem

$$\min_x \|d_S - Ax\|^2. \quad (3.9)$$

Here, A becomes a matrix with three columns. The first column consists of the pixel locations in the target image d_T , the second of a column with ones on the elements corresponding to the x-direction and zeros on the ones representing the y-direction in d_T . The third column has ones on the elements corresponding to y and zeros otherwise. The vector x is a column vector with elements s , t_x and t_y .

It is desirable to weight the local differences with the calculated certainty measurements. The scale and translation representing the new regularised field is then given by

$$\hat{x} = [s \quad t_x \quad t_y]^T = (A^T W A)^{-1} A^T W d_S, \quad (3.10)$$

where the weight matrix W has the certainty measurements c_k on the diagonal and zeros elsewhere.

The result needs to be transformed to a deformation field that can be accumulated in the next step of the algorithm. One simple way to do this is to use the pixel locations for the target image and the calculated scale and translations to transform the locations to locations belonging to the deformed source, according to

$$\text{vec}(d_k^{\text{reg}}) = s \cdot d_T + e_x \cdot t_x + e_y \cdot t_y - d_T, \quad (3.11)$$

where $\text{vec}(d_k^{\text{reg}})$ denotes the vectorization of the regularised displacement field (the column stacked version). The difference between those locations represents the values in the regularised deformation field. The last step is to turn this column stacked vector back into its matrix form (the same size as the target image) and the result of the regularisation is found.

3.1.2.3 Affine transformation

The assumption of always having a source and target according to the transformation model above might be wrong. The theory for regularising according to an affine model, or a non-rigid model as will be described next, will be given to be able to improve the registration when needed.

An affine transformation could include any type of transformation that preserves parallel lines in an image. This transformation is given from a 2×2 matrix plus a vector representing the translation.

The procedure of finding the affine regularisation is similar to the one described above. Here the coordinates of the source image and the target image are arranged to matrices with the x-coordinates in the first column and the y-coordinates in the second column. The matrix representing the coordinates in the target image is extended with a third column consisting of ones. This matrix is denoted D_T and the one including the source coordinates D_S . The extra ones makes it possible to handle the 2×2 matrix and the translation at the same time in a 3×2 matrix X with elements as

$$X = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ t_x & t_y \end{bmatrix}. \quad (3.12)$$

The minimization problem for the affine transformation becomes

$$\min_X \|D_S - D_T X\|^2, \quad (3.13)$$

and the solution, including the certainty measurements as weights,

$$\hat{X} = (D_T^T W D_T)^{-1} D_T^T W D_S. \quad (3.14)$$

The transformation from the resulting \hat{X} to a displacement field is done similarly as in the previously described transformation model.

3.1.2.4 Non-rigid transformation

A non-rigid transformation model will allow all kinds of movement to reach target from source. To prevent the image from tearing apart, as shown in Figure 3.11, a smoothing of the estimated displacement field is introduced. The regularisation is performed using normalized averaging with a Gaussian kernel g

$$d_k^{reg} = \frac{(c_k d_k) * g}{c_k * g}, \quad (3.15)$$

using the same notations as before. By using this method the certainty measurements c_k can be included to strengthen more reliable estimates. The size of the Gaussian kernel represents how much variety that should be allowed over the field; a small kernel will give a more elastic field and a large kernel will generate a stiffer displacement field.

3.1.3 Deformation field accumulation

The displacement field found in each iteration is added to an accumulated deformation field. This field is needed to transform the original source. It is important to accumulate and always transform the original image since an interpolation is performed in every transformation, where a smoothing effect is introduced. This cannot be prevented, but always using the original image or signal makes this effect appear only once. Repeated transformations would eventually make the source completely smooth and the information in the image would be lost. Figure 3.13 shows the smoothing effect of five repeated transformations on a one-dimensional signal.

The accumulated displacement field d_a is updated with the current estimated displacement field d_k

$$d'_a = d_a + d_k. \quad (3.16)$$

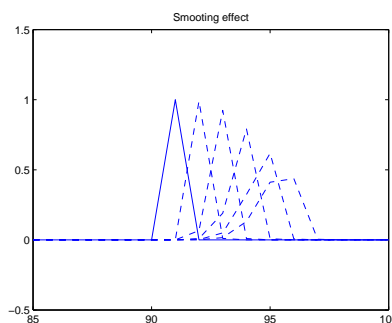


Figure 3.13: Figure illustrating the smoothing effect that is introduced on a line if the deformation field is not accumulated as described above.

The method suggested, uses the corresponding certainty measurements c_a and c_k as weights according to

$$d'_a = \frac{c_a d_a + c_k (d_a + d_k)}{c_a + c_k}. \quad (3.17)$$

The accumulated certainty field c_a is then updated with the current certainty estimates c_k , weighted with themselves. The certainty values are thus used as certainties of themselves

$$c'_a = \frac{c_a^2 + c_k^2}{c_a + c_k}. \quad (3.18)$$

For reasons discussed in Section 6.1.1, the proposed method was rejected in favour of accumulating the displacement field without certainties as weights.

3.1.4 Deformation of image

At the end of an iteration the accumulated deformation field is used to deform the original source image. The deformed image is found from linear interpolation of the source and is set to be the new image to be filtered and compared to the target. After termination of the Morphon algorithm, the last deformed source is the solution to the registration problem. This image should now look similar to the target.

The resulting registration from the one-dimensional example in Figure 3.2 can be seen in Figure 3.14.




Figure 3.14: The source signal, the deformed source signal after termination and the target signal. Notice that the source signal has been moved to the right and that it now looks similar to the target. The difference in height that can be seen for the line is due to interpolation artefacts.

The corresponding result for the two-dimensional example in Figure 3.7 is given in Figure 3.15. Small differences between the deformed source and the target indicate a good registration result.

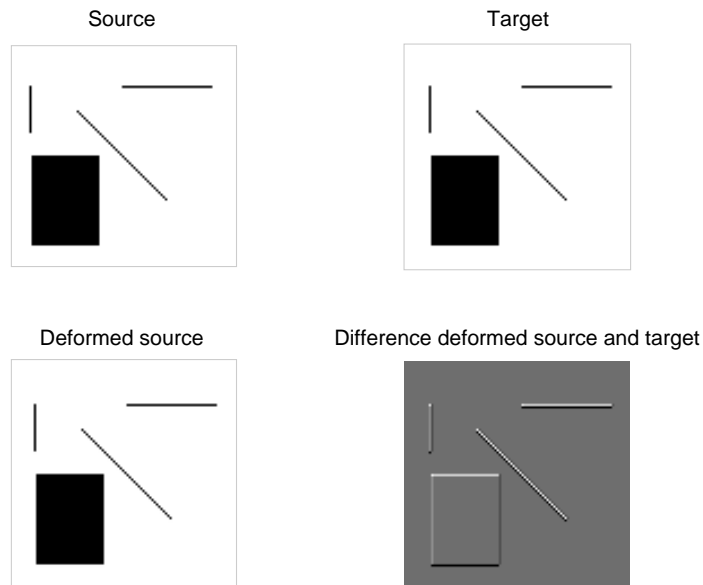


Figure 3.15: The result after image registration. On the upper row source and target can be seen, whilst the bottom row shows the deformed source and a comparison between the deformed source and the target. The differences are small and the registration successful.

3.1.5 Registration example of 2D image

There is no need for source and target to have the same size or scale, as in the simple examples above. A final example is included to illustrate the result of the Morphon method on a real image pair. Figure 3.16 shows the original source and target images.

A start guess ensures that the difference between source and target is small enough for the phase differences to be proportional to the spatial differences between the images. The algorithm starts on a coarse resolution scale to pick up larger differences at start. The images are filtered with the quadrature filters and the phase differences and certainty estimates are computed. The outcome is used to increase the similarity of source and target and the source image is moved towards the target. The iteration continues a number of steps until the registration is acceptable at the current resolution scale. The procedure continues at finer scales until the image registration is complete. The result for this example can be seen in Figure 3.17.




Figure 3.16: Example of source and target with different size, scale and translation. The target image is smaller and it covers only a part of the source. This is similar to the registration of bone scintigraphy images and a SPECT volume.



Figure 3.17: The result after registration of the example in Figure 3.16. The differences are small, besides those with origin in information only available in the source, and the result is good.

3.2 Principal component analysis

Principal component analysis (PCA), explained in [19], is used as a part of testing the result of the image registration. The test will be explained further in Section 4.1.1.

PCA is a mathematical method of transforming a set of observations of possibly correlated variables to a set of uncorrelated variables called principal components. An orthogonal linear transformation gives a new coordinate system such that the largest variance by any projection of the data comes to lie on the first coordinate, or principal component. The second component corresponds to the second greatest variance, and so on. This property makes PCA useful for reducing the dimensionality in a problem. In this application PCA is needed to illustrate the variance of a Gaussian set of data. The purpose is to find an ellipse that encloses most of the data points and to use this ellipse to evaluate the quality of an image registration using a number of landmarks.

The original two-dimensional coordinate system has axes denoted x_1 and x_2 , while the principal axes of the transformed coordinate system is given the names l_1 and l_2 . The matrix $L = [l_1 \ l_2]$ is called the loading matrix. The data points, represented in X , are projected onto the principal axes

$$S = X \cdot L = [s_1 = Xl_1 \quad s_2 = Xl_2]. \quad (3.19)$$

S is called the scores matrix and consists of the principal components s_1 and

s_2 . The variance σ_1^2 along the first principal axis l_1 , that is to be maximized, is

$$\sigma_1^2 = \frac{1}{n} s_1^T s_1 = \frac{1}{n} l_1^T X^T X l_1 = l_1^T \left(\frac{1}{n} X^T X \right) l_1 = l_1^T \Sigma_X l_1. \quad (3.20)$$

The problem becomes

$$\max_{l_1} l_1^T \Sigma_X l_1 \quad \text{subject to} \quad l_1^T l_1 = 1. \quad (3.21)$$

Equation 3.21 is solved using the technique of Lagrange multipliers. If there is a maximum $f(x_0)$ to the original constrained problem $f(x)$, then there exists a stationary point for the Lagrange function $L(x, \lambda)$ at (x_0, λ_0) . The Lagrange function for this problem is defined by

$$L_1 : \max_{l_1} l_1^T \Sigma_X l_1 - \lambda_1 (l_1^T l_1 - 1), \quad (3.22)$$

where λ_1 is a Lagrange multiplier. Differentiation with respect to l_1 gives

$$\frac{\partial L_1}{\partial l_1} = 2\Sigma_X l_1 - 2\lambda_1 l_1 = 0 \quad \Leftrightarrow \quad \Sigma_X l_1 = \lambda_1 l_1, \quad (3.23)$$

which is an eigenvalue problem. Solving Equation 3.23 will give eigenvectors l_j with corresponding eigenvalues λ_j for $j = 1, 2$. The variance along l_1 is now

$$\sigma_1^2 = l_1^T \Sigma_X l_1 = \lambda_1 l_1^T l_1 = \lambda_1. \quad (3.24)$$

The conclusion is that the first principal axis is given by the eigenvector corresponding to the largest eigenvalue.

The second principal axis is orthogonal to the first. Besides this, the same conditions apply as in finding l_1 . The Lagrange function becomes

$$L_2 : \max_{l_2} l_2^T \Sigma_X l_2 - \lambda_2 (l_2^T l_2 - 1) - \mu (l_2^T l_1), \quad (3.25)$$

where λ_2 and μ are Lagrange multipliers. Differentiation gives

$$\Sigma_X l_2 - \lambda_2 l_2 - \mu l_1 = 0. \quad (3.26)$$

Multiplying with l_1^T from the left and using Equation 3.23 and the orthogonality $l_1^T l_2 = 0$ reduces the relation to

$$l_1^T \Sigma_X l_2 - \lambda_2 l_1^T l_2 - \mu l_1^T l_1 = \mu l_1^T l_1 = 0. \quad (3.27)$$

Since $l_1^T l_1 = 1$, μ has to be zero and the part remaining from Equation 3.26 is the eigenvalue problem from Equation 3.23. This means that the second principal axis is given by the eigenvector l_2 corresponding to the second largest eigenvalue λ_2 . Similarly, the variance σ_2^2 along the axis l_2 is λ_2 according to

$$\sigma_2^2 = l_2^T \Sigma_X l_2 = \lambda_2 l_2^T l_2 = \lambda_2. \quad (3.28)$$

The result can be used to create an ellipse capturing most of the variance among the data points. An example is given in Figure 3.18.




Figure 3.18: PCA of an example data set. The figure shows an ellipse that extends 2.5 standard deviations along each principal axis, which captures 98.8% of the total variance. The unit of the axes is number of pixels, the original coordinate axes are shown with dashed arrows and the calculated principal axes with solid arrows. The different symbols and colours of the points are associated with origin in different images used to collect the data.

3.3 Segmentation using region growing methods

In image processing, segmentation means the process of partitioning an image according to some meaningful relationship among sets of pixels. The purpose of computing segments could typically be to find an object or area of interest for further analysis in some problem. Pixels, or voxels for three-dimensional images, that are similar with respect to some characteristic property like intensity, texture or colour are set to belong to the same region. Adjacent regions are in the same manner significantly different to these pixels according to the same properties. [20]

The applications of segmentation are widely spread and there are many different methods available depending on the character of each problem. In this thesis a segmentation in two dimensions is known and the goal is to find the corresponding segmentation in three dimensions. A method that can use this prior knowledge is preferable to limit the search space. A seeded region growing method fulfils these requirements.

Seeded region growing, presented in [21], consists mainly of three steps, which will be described in the following sections; *finding or choosing seed points*, *increasing of the region* and *termination of the algorithm*.

3.3.1 Finding or choosing seed points

In seeded region growing, the segmentation starts from one pixel or voxel in the image that is known to belong to the wanted segment. For this specific problem there is a two-dimensional segmentation available that can be used for initialization of the three-dimensional segmentation. The segmentation comes

with information like size, centroid point and a bounding box for each segment. The result of the image registration is used to transform the bounding box to the coordinate system of the volume. The voxel with the highest intensity in the three-dimensional generalisation of the bounding box is chosen as a seed point for the region growing algorithm.

3.3.2 Increasing the region

From the initial seed point, one pixel or voxel at a time is added to the region until some criterion terminates the iterations. The pixel or voxel that is added must belong to the neighbouring structure of the region. This makes the segment cohesive. Common neighbouring structures for two-dimensional images are the 4-connected neighbourhood and the 8-connected neighbourhood, which can be seen in Figure 3.19. The theory is easy to extend from two to three dimensions. To reduce the amount of computations needed in three dimensions, the generalisation of the 4-connected area to a 6-connected area is chosen.

The property of the selection of pixels or voxels in the neighbourhood of the current region is compared to the region according to some criterion. A common rule that is used in this application is to add the pixel or voxel with the intensity closest to the average intensity of the current region.




Figure 3.19: Figure showing the two-dimensional 4-connected and 8-connected neighbourhood. A three-dimensional generalisation of the 4-connected neighbourhood to a 6-connected neighbourhood is used in this implementation of the seeded region growing method.

3.3.3 Termination

The most difficult in the segmentation process with region growing methods is to know when to terminate the algorithm. The adding of pixels or voxels continues until some stopping criterion is fulfilled, and the quality of the segmentation is strongly connected to the ability to stop at the right moment.

To create a more robust method an automated stopping criterion, inspired by [22], has been implemented. The criterion has three parts; the first is a maximum size allowed for each segment, the second is an area that the segmentation cannot leave and the third is a maximum difference allowed in intensity variations between the region and the potential pixel/voxel to add. The size is limited by the dimensions of the bounding box mentioned earlier. The maximum volume is calculated assuming that the third dimension is no larger than the highest value of the two known dimensions. The bounding box is also used to define what area the segmentation should stay inside. An area with five pixels

margin from the bounding box is chosen. The allowed intensity difference is increased gradually from 1% of the intensity of the seed point, with steps of 0.5%, and each segmentation found is considered to be the result. The termination happens either due to the fact that the current segmentation is considered too large or that the segmentation continues outside the region of interest. The segmentation belonging to the previous intensity threshold is chosen as the result of the segmentation.

The method of finding the right segmentation is extended with an extra criterion, revealing indications of an unwanted explosion in the segmentation process. An explosion in a region growing method is when the region suddenly grows with a much higher rate than earlier. Ideally this would stop the iterations just after the correct segmentation is found. However, this rapid change in growing rate might happen both before and after this point. To avoid under-segmentation the first threshold for the intensity is given a very small value. If an explosion has occurred during the algorithm, the segmentation previous to the last exploded one is chosen as the segmentation of interest.

The criterion for an explosion is used on each resulting segmentation before the threshold value of intensity is increased and the algorithm continues. The size of the current region is compared to the previous region and the difference between those is divided by the change in intensity variation allowed. This corresponds to finding the slope of an equation showing the region size as a function of the intensity threshold. Figure 3.20 shows an example of such a diagram. It is not obvious what is an explosion and what is not. If the slope has a value larger than five, an explosion is regarded to have happened.

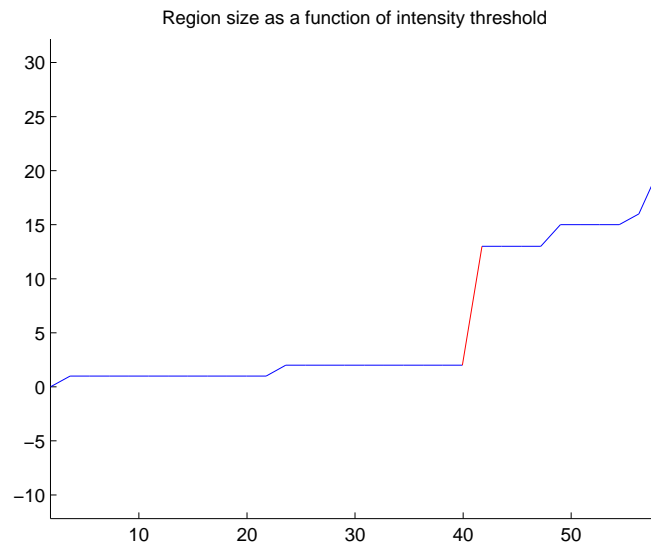


Figure 3.20: Figure showing the region size as a function of the intensity difference allowed for a voxel to be included in the region. The red colour marks a threshold change that gives an explosion.

Chapter 4

Methods

This section will describe how the theory from above is applied to the specific problem in this thesis. The assumptions made will be motivated and the methods used to evaluate the results will be discussed here.

4.1 Image registration

Image registration is needed to find the connection between a set of bone scintigraphy images and their corresponding SPECT volume. The primary goal is to be able to identify hotspots of some kind in one of the planar bone scintigraphy images and directly know where in the tomography image the hotspot is located. A hotspot could be a metastasis but there could also be other areas that show high intensity in the images.

The first problem to solve is how to register a two-dimensional image and a three-dimensional image. The Morphon method could handle a pair of two-dimensional images or a pair of three-dimensional images but in this case there is one of each. To save computational time and effort it would be desirable to perform the calculations in 2D. The next question is whether it is needed to register one anterior and one posterior version of the SPECT image, or if one of them would give enough information to solve the problem.

Initially the registration was accomplished in two versions. To do this the SPECT volume was projected to two dimensions using a quadratic weight to simulate the anterior and the posterior image. In this way, the part of the body that was closest to the camera contributes more to the image. Some experimentation showed that the result was equivalent to register one image pair and flip the result when needed. For this reason the image registration is performed with a source that consists of the sum of the anterior bone scintigraphy and a flipped version of the posterior. The target is set to be the sum of all corresponding slices in the SPECT image. Figure 3.1 in Section 3.1 shows a setup of one of those image pairs.

To ensure that the phase differences and the spatial differences between source and target are proportional, a start guess is needed to initialize the Morphon method. The strategy to find this guess has been to let the source image slide over the target image from top to bottom and for every step save the correlation between source and target. The guess that leads to the highest

correlation is chosen as the guess to initialize the algorithm.

To be able to perform the regularisation step, a suitable transformation model has to be chosen. It is known that source and target are images of the same patient and that the images are acquired at the same examination. Since there is some time between the acquisition of the bone scintigraphy images and the SPECT image, it is possible that the patient has moved. If the body is not in the same exact position there could be need of a more complex transformation model. In most cases the time period between bone scintigraphy and SPECT is short and this is not a problem. According to this there should only be a shift in scale due to different resolution of the images and a translation from source to target. This assumption has been discussed with a doctor and a large amount of image pairs have been studied. It happens that the model is too tight and this is something a potential user or a future developer has to bear in mind. It is simple to extend the code to handle affine or elastic transformations but with a higher degree of freedom more effects could interfere with finding a good registration.

4.1.1 Test of the Morphon method implementation

The results need to be tested in some way. A number of patients have been chosen from an archive so that the images are acquired in a relatively even distributed period of time. The images have also been chosen so that different parts of the patients' bodies are visible in the matching tomography. The idea is to represent a couple of different types of images with different resolution, scale and translation. In each image pair, some landmarks are marked manually. Since different parts of the body are shown, it is hard to find locations for the landmarks that are equivalent in all image pairs. The important criterion for the landmarks is that they most certainly point to the same position in both source and target. After image registration, those landmarks could be transformed from source to target according to the resulting deformation field (or the other way around). The deformed landmarks could then be compared to the manually marked points, both visually and with a calculated distance. Figure 4.1 shows an example of what this could look like. The green dots belong to manually marked landmarks and the red dots to transformed landmarks.

The quality of the landmarks, and in extension the result of the registration, is affected by the ability to find the exact corresponding points in source and target of the person that are supposed to define the landmarks. In other words, even if the registration provides the perfect deformation field to match source and target, there would most likely be a difference between the predefined landmarks and the deformed ones. To be able to say something about how large this difference could be expected to be, another test is created. Here landmarks are given in a number of source images and the user is asked to point out the corresponding location in the target image. This is repeated with some time in between. Figure 4.2 shows a screenshot from one step in this test. The process will lead to clusters of points with some deviation from a centre. This deviation could be used as a measurement of the variability and is calculated using PCA, as explained in Section 3.2. The result is shown as ellipses around the pre-defined landmarks in the target images, indicating their spatial variation. A deviation from a deformed source point to the corresponding target landmark within this ellipse is considered acceptable.




Figure 4.1: The figure shows bone scintigraphy images as the source (left) and SPECT images as the target (right). Landmarks are defined in both images before registration, and can be seen as green dots. The landmarks in the source image are transformed according to the result of the registration, and can be seen as red dots in the target image. The transformed landmarks are compared to the pre-defined ones in the target image, and the result is used to evaluate the quality of the image registration.




Figure 4.2: The figure shows a screenshot from testing the variability for a person to find the same corresponding point in the target image a repeated number of times. One point at a time is given in the source image to the left and the user clicks at the point in the image to the right that he or she believes is viewing the same location. The test is repeated and the gathered data is analyzed using principal component analysis.

4.2 Segmentation

One of the goals in this thesis is to find a value of how large the metastases are that are found in the two-dimensional analysis of the bone scintigraphy images. The idea is then to, if possible, use this value to improve the Bone Scan Index. The problem is somewhat simplified since the question of what is a metastasis or not is completely left to the two-dimensional analysis. The goal here is to find the three-dimensional segmentation that best describes the existing 2D segmentation.

The methods described in Section 3.3 were implemented and the result visualized through cross-sectional images in the same way as the original SPECT images. First, a successful image registration is necessary to be able to find good seed points for the region growing method. Next, the script might find a segmentation that looks similar to the two-dimensional version of the segmentation or not. The resulting segmentations are compared to the two-dimensional analyzes available to evaluate the quality of the result. Figure 4.3 shows an example of how an existing two-dimensional analysis can look like. The blue colour marks hotspots while the red colour marks hotspots that have been classified as metastases. To simplify the comparison between a result like the one in Figure 4.3, the 2D sum of the calculated three-dimensional segmentation is created (from an anterior view).




Figure 4.3: The figure shows an analysis made on the bone scintigraphy images of an example patient. Hotspots (areas with significantly higher intensity than the surrounding pixels) can be seen in blue, while the hotspots that have been classified as metastases are red. It is the red hotspots that are the foundation to initializing the three-dimensional region growing. The segmentation result is compared to views corresponding to the one in this figure.

Chapter 5

Results

Here, the results of the image registration and the segmentation are presented. A few comments will be given about the results, with a more thorough discussion in the next chapter. All implementations have been carried out in Matlab.

5.1 Image registration

The implementation of the Morphon method has been tested on fifteen patients with scintigraphy images and SPECT images. The results of the image registrations are compared using eight landmarks in each image pair. At least six of eight deformed landmarks must be close to the pre-defined landmarks for the registration to be successful.

The distance allowed between a pre-defined landmark and a deformed landmark in the target image, is based on the test described in Section 4.1.1. An ellipse extending 3 standard deviations along each principal axes will capture 99.7% of the total variance of the points from the variability test. The problem is to know where the centre of the ellipse should be located. The assumptions in the test means that the ellipse should have a centre in the unknown correct point and that the point chosen by the person defining the landmarks should be located somewhere inside the ellipse. If a worst case scenario is considered and the point chosen is located at the boundary, the criterion for a good match becomes very narrow as this point is set to be the approximation of the centre of the ellipse. An illustration of the problem is given in Figure 5.1. To make sure that points inside the correct ellipse is included, the ellipse extending 6 standard deviations is chosen.

The test data set consists of bone scintigraphy images of size 1024×256 and SPECT images with slices of size 128×128 . The number of slices differ between 76 and 127. If there are slices with no information, the SPECT image is cropped before the registration. The initial guess for the scale is two and for the translation sideways it is zero. Vertically, the initial guess is found comparing the correlations found letting the target image slide over the source image. These initial guesses are used to create an initial field for the algorithm. This implementation of the Morphon method rescales the images three times and the number of iterations on each scale is seventeen. The size of the filter has been nine and a proportionality constant of one has been used for the relation





Figure 5.1: A justification of the choice of ellipses around landmarks. The black star is assumed to be located at the spot where the landmark should have been placed and the arrows represent the original coordinate axes and the calculated principal axes. The red star is an example of where the user could have clicked to set the landmark, and since the correct point is unknown the center of the ellipse end up here. To make sure that the points inside the ellipse around the correct point is considered right, the ellipse drawn is enlarged with a factor two. This ellipse (dashed, red line) is the one used in evaluating the result of the image registration.

between the phase differences and the spatial differences. The results of the registration are given in Table 5.1. The corresponding images can be seen in Appendix A.


It is not straight forward to choose a criterion for a successful registration and the result needs to be studied further. The registrations that have failed are presented in Figure 5.2. Patient number four is not far from being considered okay and the failure could be due to badly chosen landmarks. The registration of patient number nine seems to be shifted and the result could be due to too few iterations or perhaps changes in the body that makes it difficult to find a perfect registration. An example of this could be that more of the injected radioactive material has reached the bladder during the acquisition of the SPECT image than at the time for the bone scintigraphy. Since the bladder has such a high intensity in the image this could lead to the fact that edges and lines move in a way that the chosen transformation model cannot reflect. The registration of patient ten is far from close, and the results looks like an example of when the phase differences between source and target are too large to be proportional to the spatial differences. Visualizing the results step by step throughout the iterations, shows that the initial guess for registering patient number ten is not the best possible choice. A manual correction improves the result, see Figure 5.3. The landmarks are still not close enough to classify the registration as successful, but now the differences are more likely due to the fact that the image has a poor resolution. This complicates both the registration and the process of defining landmarks.

Target




(a) Registration result of patient number four.

Target



(b) Registration result of patient number nine.

Target



(c) Registration result of patient number ten.

Figure 5.2: The figure shows the registrations classified as unacceptable, according to the criterion that at least six of the eight transformed landmarks should be close to the landmarks in the target image.

Patient	Correct landmarks	Incorrect landmarks	Result
1	7	1	Successful
2	8	0	Successful
3	8	0	Successful
4	5	3	Unacceptable
5	6	2	Successful
6	7	1	Successful
7	7	1	Successful
8	7	1	Successful
9	1	7	Unacceptable
10	1	7	Unacceptable
11	6	2	Successful
12	6	2	Successful
13	6	2	Successful
14	7	1	Successful
15	7	1	Successful

Table 5.1: The table shows a summary of the result of testing the implementation of the Morphon method. The criterion for a successful image registration is that at least six of eight deformed landmarks are close to the corresponding pre-defined landmarks. The definition of close in this application is given by the large ellipse in Figure 5.1.





Figure 5.3: The registration result after adjusting the initial guess for patient number ten. The result is improved but still not classified as a successful registration. The problem could be due to a poor resolution of the image, which complicates both the registration procedure and the task of defining landmarks.

5.2 Segmentation

The segmentation part has been tested on the images of sixteen patients, different from those in the registration part. The goal is, as mentioned, to achieve a segmentation of potential metastases as close to the result of analyzing the two-dimensional bone scintigraphy images as possible. There should be more information available in a SPECT image, and therefore it is possible that the script might find a better segmentation than the 2D analysis even though it might differ from the 2D result. For this reason, the results of the segmentation are not evaluated as successful or unacceptable. A couple of examples will be given below as a proof of concept, while the result for each of the sixteen patients can be seen in Appendix B. It is worth noting that patients can have hotspots at a location that is not covered by the SPECT image. The script for finding a segmentation ignores hotspots that are not completely inside the SPECT volume, which means that there are results with no segmentation at all. The reason for including these patients is to investigate whether the implementation of the Morphon method produces good registrations, which is necessary to be able to initialize the region growing method.

The choice of parameters for the region growing method has been included in the description of the segmentation method in Section 3.3. The segmentation results of patient number seven and eight have been chosen to represent how the theory can be used. Figure 5.4 shows a possible metastasis in the spine. The corresponding three-dimensional segmentation seems to have reached a similar estimation. Figure 5.5 shows a patient with hotspots in the pelvic area and in the ribs. Also here, the segmentation result reflects the two-dimensional estimation in a good way.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure 5.4: Result for patient number seven. A number of hotspots have been found in the spine, and the segmentation looks similar. The result is therefore correct.



(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure 5.5: Segmentation results for patient number eight. The two-dimensional analysis has found hotspots in the pelvis area and in the ribs. The three-dimensional segmentation result looks similar to the two-dimensional analysis, and the outcome is considered successful.

Chapter 6

Discussion

The discussion will start with comments on the registration and the segmentation results. Next, a follow-up of the milestones from Section 1.1.1 will be given. Some conclusions will summarize the result of this project and at last, some thoughts on future work will be given.

6.1 Image registration

The image registration result indicates that the Morphon method is a good algorithm to register bone scintigraphy images and SPECT images. There are some problems though; all registrations are not successful. The method for finding a good start guess for the algorithm was not robust enough. This is not a part of the registration method, but the outcome is essential for the Morphon method to work. A start guess too far from the target becomes a problem in the deformation field estimation. A distance too large means that the proportionality between the phase differences and the spatial differences cannot be guaranteed. A larger filter or rescaling the images even further might help compensating a bigger spatial difference.

Another factor that affects the result is the choice of parameters. It has been desirable to have the same choice for all images. Experimenting with this indicates that a change in parameters can improve some results and make other worse. This suggests that testing the registrations one by one could give better results. If there is no method for finding optimal parameters, this could be a problem when implementing the algorithm in a program. For this reason it is preferable to stick to parameters common for all images. What choice is the best could be questioned, but the parameters in this implementation seem okay.

Working with the images has made it obvious that the quality of the images are quite different. Some images are harder, both to identify good locations for landmarks in and to define those landmarks precisely. Looking at the data set from the PCA analysis in Figure 3.18, tells that an ellipse specific for each image pair would vary a lot in size. Many of the points close to the border of the ellipse seem to originate from the same image. This means that the test method might be too tough for some image combinations and too kind for others. Furthermore, the size of the ellipse is just a guess that seemed to perform well and many of the landmarks are close to the edge. A small difference in this choice might

affect the result significantly. Note that a result close to successful might be good enough to initialize a segmentation.

The result could possibly be improved by redefining the landmarks. There are landmarks that in a close-up view obviously are placed further away from the correct position than the ellipse allows. This has probably happened because a location in the body has been mistaken for another part nearby. A correct registration would still classify the misplaced point as far from its corresponding landmark.

It is important to capture all possible variations in the transformation model chosen. There are changes appearing from the acquisition of the bone scintigraphy images to the SPECT image that cannot be captured here. An example of this is when the patient has moved (patient number twelve in the segmentation test has moved his head, see Appendix A). Another is when internal changes in the body affect the result. This could happen when a larger amount of radioactive material is gathered in the bladder at the time for the SPECT image than at the time for the bone scintigraphy investigation. An affine or elastic transformation model is needed to capture such changes.

6.1.1 The accumulation step

As mentioned briefly in Section 3.1.3 the proposed method for handling the accumulation of the estimated displacement field and the corresponding certainty measurements was questioned. One problem with accumulating the displacement field according to Equation 3.17 after the regularisation step, is that the accumulated field will deviate from the chosen transformation model. It is suggested that the accumulation should take place either before or after the regularisation, but accumulating before regularisation would result in the fact that the first field is regularised a repeated number of times, while newer fields are regularised fewer times. This might be better than the first approach, but the feeling of some missing part kept coming back during the experimentations.

Implementing the example with no deformations allowed besides translation of the whole image, made it feel intuitive to add a regularisation of the certainty field. If the values in the deformation field are changed, the corresponding certainties should change in a similar way. For example, if a weighted mean is used to find the translation the corresponding certainties are approximated to a weighted mean using themselves as weights. If an elastic model is used, the same averaging method could be used to the certainty field. This procedure is however hard to generalize to an affine transformation model or a model like the one chosen for the registration problem in this project.

The procedure of adding a regularisation of the certainties was changed to the simplified accumulation described in Section 3.1.3. The results were still good after a few changes in parameters; better than with the suggested accumulation applied after the regularisation. What was surprising, was the fact that this change in the Morphon method was tried on EXINI's existing registration of an atlas to bone scintigraphy images and the results was improved. Figure 6.1 shows three examples of improvements. Studying a large set of images, noting that the results never seem to get worse with this change, leads to the conclusion that the Morphon method could be improved and the parts including regularisation and accumulation is something that would be interesting to investigate further in the future.




Figure 6.1: Registrations of an atlas to bone scintigraphy images, using the suggested accumulation to the left and the simplified accumulation to the right. Improvements are visible primarily in registration of the arms.

6.2 Segmentation

First the conclusion is made that the image registration provides results good enough to initialize the segmentation for all patients. The segmentation results are harder to evaluate in a general manner. The presented result works as a proof of concept and the idea of analyzing bone scintigraphy images and use the result to start a three-dimensional segmentation with a seeded region growing method works. In classifying a three-dimensional segmentation as good or bad, the hardest part is to know whether the segmentation in 2D or the one in 3D is the correct one or if perhaps both are incorrect. In this project, the aim was to find a three-dimensional segmentation as similar to the two-dimensional segmentation as possible. In other words, the segmentation could differ and be considered bad even though it might reflect the truth in a better way.

The most uncertain parameter that affects the result of the region growing algorithm, is the part that determines what is an explosion and not. There might be room for improvement here, even if the choice made seem to reflect the two-dimensional analyses in most cases. It seems as though it is harder to find good segmentations of larger hotspots. Over all, a lot of assumptions are made in implementing the stopping criterion for the region growing method and every one of them contributes to the result. It is likely that the choices have to be optimized for the application of each specific problem.

6.3 Evaluation of milestones

Here a follow-up of the milestones from Section 1.1.1 is given. Six out of seven milestones have been reached.

1. *Two-dimensional projections of SPECT images*
It turned out to be enough to perform one registration with the sum of the anterior and a flipped version of the posterior bone scintigraphy image as source and a sum of all slices of the SPECT volume from an anterior view as target. It should be possible to use this strategy for other 2D-3D image registration problems.
2. *Image registration of SPECT and bone scintigraphy images*
With the right choice of parameters and a good start guess, the Morphon method produces image registrations of bone scintigraphy images and SPECT images in a good way.
3. *Validation of the result*
Testing the image registration results using landmarks indicates a need of an affine or elastic transformation model in some cases. It also happens that the method for finding a start guess fails in giving the best possible option. Twelve of fifteen tested registrations are classified as successful and at least one of the failures could be due to badly chosen landmarks. The sixteen registrations performed in the segmentation part are all good enough to initialize the region growing algorithm.
4. *User interface for navigation through the images*
A user interface has been created using Matlab, where it is possible to navigate through both bone scintigraphy images and SPECT images at the




Figure 6.2: The figure shows a screenshot from the user interface created to simplify the overview of bone scintigraphy images with SPECT images. The first image from the left shows the anterior bone scintigraphy image and the second the posterior image. The next three shows the SPECT volume from different angles. A click in one of the first two images indicates where the chosen point is located in the SPECT image, but since the depth cannot be determined a line is drawn at which the point should be found. A click in one of the others marks the spot chosen in all images. The sliders under the images adjust the display range of gray levels.

same time. The result has been viewed to a couple of doctors and biomedical scientists who have contributed with data, and the user interface has gained a lot of positive response. A screenshot is given in Figure 6.2.

5. *Segmentation of metastases in three dimensions*


The seeded region growing method works to find three-dimensional segmentations of the hotspots found in 2D. It might be necessary to evaluate the result in a more general way, but the result looks promising.

6. *Distribution of metastases*


This milestone has not been reached. A size of the metastases has been calculated in number of voxels, but it needs to be linked to the proportion of the bones to be able to compare the result with the BSI of today. If the CT image would have covered the whole body and not just a small part, a segmentation of the bones could contribute to this. A proportion of the amount of metastases in the bones for the part covered by SPECT and CT would unfortunately not say anything about the total BSI.

7. *Three-dimensional visualisation of metastases*

The analysis to find hotspots is carried out on the bone scintigraphy images, and the result is used to initialize a segmentation of possible metastases in the SPECT image. The segmentation is transformed to the coordinate system of the CT image, where a segmentation of the bones is made from a threshold value. The visualisation can be seen in Figure 6.3.



(a) Analysis from EXINI Bone.



(b) 3D visualisation of metastases.

Figure 6.3: The figure below shows a screenshot from a rotatable three-dimensional visualisation of the bones of a patient, with possible metastases plotted in red. The image registration is used to initialize the seeded region growing method from the two-dimensional analysis above, and the resulting segmentation is transformed to the coordinate system of the corresponding CT image. The segmentation of the bones is created using an intensity threshold on the CT image, requiring a connected area.

6.4 Conclusions

The methods chosen works for solving the tasks proposed in this thesis. The Morphon method is reliable if the proportionality between phase differences and spatial differences (see Section 3.1.1) can be ensured and if the right transformation model is chosen. Benefits are the independency of intensity values and that it is very easy to change the regularisation part to handle different types of transformation models.

The need for a segmentation method that could make use of the existing analyses in 2D was met with the seeded region growing method. The algorithm generates good segmentations for the patients tested here. It is however not obvious what is a good and a bad segmentation and further testing might be needed if the result is to be implemented in a commercial product.

Even though the image registration part and the segmentation part have been successful, one part is missing to achieve the aim of the thesis. Unfortunately the voxel size of the segmented metastases has not been transformed to some unit comparable to the BSI value. I believe that it is possible to calculate an updated BSI closer to the truth, but I am not convinced that this value would affect the conclusions drawn from the existing BSI.

6.5 Future work

All methods used could probably be refined further by choosing even better sets of parameters and by optimizing the code with respect to computational time. To reach the goal of this thesis it is necessary to investigate the connection between voxel size of a hotspot and the size in mm^3 or at least to say something about the ratio between the size of the metastases and the size of the bones.

Another possible approach to include the information from SPECT/CT images in the estimation of the BSI value could be to perform the analysis of suspected metastases in the SPECT image. This strategy would require more computational effort but a possible gain could be the ability to find metastases that can't be seen in the scintigraphy images.

It would be interesting to analyze the accumulation step and the effect of the certainty measurements in the Morphon method further. There seems to be room for improvement of the method, and an investigation of this could be useful for applications in many different areas.

Hopefully the user interface for viewing scintigraphy images with corresponding SPECT images can be implemented as improvements in future software versions. The same applies for the registration improvements which may emerge from changing the accumulation step of the Morphon method.

Chapter 7

Bibliography

- [1] Y. Erdi, J. Humm, M. Imbriaco, H. Yeung, and S. Larson, “Quantitative bone metastases analysis based on image segmentation,” *The Journal of nuclear medicine*, vol. 38, no. 9, pp. 1401–1406, 1997.
- [2] M. Imbriaco, S. Larson, H. Yeung, O. Mawlawi, Y. Erdi, E. Venkatraman, and H. Scher, “A new parameter for measuring metastatic bone involvement by prostate cancer: the bone scan index,” *Clinical cancer research*, vol. 4, no. 7, pp. 1765–1772, 1998.
- [3] M. Sadik, I. Hamadeh, P. Nordblom, M. Suurkula, P. Höglund, M. Ohlsson, and L. Edenbrandt, “Computer-assisted interpretation of planar whole-body bone scans,” *Journal of Nuclear Medicine*, vol. 49, no. 12, pp. 1958–1965, 2008.
- [4] I. Avazpour, R. Roslan, P. Bayat, M. Saripan, A. Nordin, and R. Abdullah, “Segmenting ct images of bronchogenic carcinoma with bone metastases using pet intensity markers approach,” *Radiology and Oncology*, vol. 43, no. 3, pp. 180–186, 2009.
- [5] M. Pacilio, C. Basile, S. Shcherbinin, F. Caselli, G. Ventroni, D. Aragno, L. Mango, and E. Santini, “An innovative iterative thresholding algorithm for tumour segmentation and volumetric quantification on spect images: Monte carlo-based methodology and validation,” *Medical physics*, vol. 38, p. 3050, 2011.
- [6] B. Zitova and J. Flusser, “Image registration methods: a survey,” *Image and vision computing*, vol. 21, no. 11, pp. 977–1000, 2003.
- [7] D. Pham, C. Xu, and J. Prince, “Current methods in medical image segmentation 1,” *Annual review of biomedical engineering*, vol. 2, no. 1, pp. 315–337, 2000.
- [8] A. Nystrand, “Cancerfonden, Prostatacancer.” <http://www.cancerfonden.se/sv/cancer/Cancersjukdomar/Prostatacancer/>. 2012-04-27.
- [9] “National Cancer Institute, What You Need To Know About Prostate Cancer.” <http://www.cancer.gov/cancertopics/wyntk/prostate>. 2012-04-27.

- [10] “Bone scintigraphy.” http://en.wikipedia.org/wiki/Bone_scintigraphy. 2012-04-27.
- [11] “Single-photon emission computed tomography.” http://en.wikipedia.org/wiki/Single_photon_emission_computed_tomography. 2012-04-27.
- [12] “X-ray computed tomography.” http://en.wikipedia.org/wiki/X-ray_computed_tomography. 2012-04-27.
- [13] “National Cancer Institute, Computed Tomography (CT): Questions and Answers.” <http://www.cancer.gov/cancertopics/factsheet/detection/CT>. 2012-04-27.
- [14] “Image registration.” http://en.wikipedia.org/wiki/Image_registration. 2012-05-08.
- [15] H. Knutsson and M. Andersson, “Morphons: Segmentation using elastic canvas and paint on priors,” in *Image Processing, 2005. ICIP 2005. IEEE International Conference on*, vol. 2, pp. II–1226, IEEE, 2005.
- [16] J. Pettersson, “Automatic generation of patient specific models for hip surgery simulation,” 2006.
- [17] “Optical flow.” http://en.wikipedia.org/wiki/Optical_flow. 2012-06-01.
- [18] “HSL and HSV.” http://en.wikipedia.org/wiki/HSL_and_HSV. 2012-08-22.
- [19] I. Jolliffe, *Principal component analysis*. Wiley Online Library, 2005.
- [20] “Segmentation (image processing).” [http://en.wikipedia.org/wiki/Segmentation_\(image_processing\)](http://en.wikipedia.org/wiki/Segmentation_(image_processing)). 2012-06-07.
- [21] R. Adams and L. Bischof, “Seeded region growing,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 16, no. 6, pp. 641–647, 1994.
- [22] J. Wu, S. Poehlman, M. Noseworthy, and M. Kamath, “Texture feature based automated seeded region growing in abdominal mri segmentation,” in *BioMedical Engineering and Informatics, 2008. BMEI 2008. International Conference on*, vol. 2, pp. 263–267, IEEE, 2008.

Appendix A

Image registration results

This appendix contains all image registration results.




Figure A.1: Registration result for patient number one. Seven of eight landmarks are close.




Figure A.2: Registration result for patient number two. Eight of eight landmarks are close.




Figure A.3: Registration result for patient number three. Eight of eight landmarks are close.




Figure A.4: Registration result for patient number four. Five of eight landmarks are close.




Figure A.5: Registration result for patient number five. Six of eight landmarks are close.




Figure A.6: Registration result for patient number six. Seven of eight landmarks are close.




Figure A.7: Registration result for patient number seven. Seven of eight landmarks are close.




Figure A.8: Registration result for patient number eight. Seven of eight landmarks are close.




Figure A.9: Registration result for patient number nine. One of eight landmarks is close.




Figure A.10: Registration result for patient number ten. One of eight landmarks is close.




Figure A.11: Registration result for patient number eleven. Six of eight landmarks are close.




Figure A.12: Registration result for patient number twelve. Six of eight landmarks are close.




Figure A.13: Registration result for patient number thirteen. Six of eight landmarks are close.




Figure A.14: Registration result for patient number fourteen. Seven of eight landmarks are close.





Figure A.15: Registration result for patient number fifteen. Seven of eight landmarks are close.

Appendix B

Segmentation results

This appendix contains all segmentation results. Note that it is not the same patients as in the test of the image registration part.




(a) Analysis from EXINI Bone.




(b) Segmentation from SPECT image.

Figure B.1: Segmentation result for patient number one.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.2: Segmentation result for patient number two.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.3: Segmentation result for patient number three.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.4: Segmentation result for patient number four.




(a) Analysis from EXINI Bone.




(b) Segmentation from SPECT image.

Figure B.5: Segmentation result for patient number five.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.6: Segmentation result for patient number six.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.7: Segmentation result for patient number seven.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.8: Segmentation result for patient number eight.




(a) Analysis from EXINI Bone.




(b) Segmentation from SPECT image.

Figure B.9: Segmentation result for patient number nine.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.10: Segmentation result for patient number ten.




(a) Analysis from EXINI Bone.




(b) Segmentation from SPECT image.

Figure B.11: Segmentation result for patient number eleven.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.12: Segmentation result for patient number twelve.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.13: Segmentation result for patient number thirteen.




(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.14: Segmentation result for patient number fourteen.




(a) Analysis from EXINI Bone.




(b) Segmentation from SPECT image.

Figure B.15: Segmentation result for patient number fifteen.



(a) Analysis from EXINI Bone.



(b) Segmentation from SPECT image.

Figure B.16: Segmentation result for patient number sixteen.