



LUND
UNIVERSITY

Master of Science dissertation:

Automatic segmentation of SPECT/CT-images using a deformable surface model based on Fourier descriptors with application in radionuclide therapy


Johan Gustafsson

Supervisors:

Katarina Sjögren Gleisner

Michael Ljungberg

Department of Medical Radiation Physics, Clinical Sciences, Lund
Lund University, Sweden, 2011



Abstract

Introduction: The purpose of this study was to develop a technique for automatic segmentation based on Fourier descriptors for separating smooth objects from the background in SPECT-images. The aim of the segmentation method was to be able to automatically outline kidneys in SPECT-images of patients undergoing radionuclide therapy with ^{177}Lu -DOTATATE. The potential advantage of combining SPECT- and CT-images for the task of segmentation using the same technique was also explored.

Material and methods: A method for segmentation of SPECT- and SPECT/CT-images was developed and implemented in IDL. The process included a first crude estimation of the organ position and size in a graphical user interface, producing an ellipsoid for initialisation of the coming optimisation procedure. The ellipsoid was used as input to an optimisation process where a surface, described by Fourier descriptors of successively increasing order, was adapted to a boundary measure based on the grey level gradient of the image. In order to avoid uncontrolled growth of the surface, the objective function was normalised to the total surface area raised to an exponent, termed the normalisation exponent. The surface description of the object could finally be transformed into a volume description of the object, *i.e.* a labelling of the voxels encompassed by the surface, by using a method based on the relative directions of the surface normal and the directed line segment joining a voxel coordinate and its closest point on the surface. The performance of the method was evaluated in different test images. Single organ images based on the left kidney of the XCAT anthropomorphic computer phantom were used in a first step, with different contrast to noise ratios and spatial resolutions. The degree of complexity was then increased by using test images designed to imitate the biodistribution of ^{177}Lu -DOTATATE in patients undergoing radionuclide therapy, again using the XCAT phantom. The same distribution was also used to create Monte Carlo simulated SPECT-images, to also model the actual camera characteristics. Finally the segmentation was performed in a set of real clinical images. In the simulated images the performance of the method was quantified both in terms of the distance between the estimated surface and the true surface, and in terms of a volume measure. In the single organ images evaluation was performed as a function of the number of Fourier orders used in the segmentation, for varying contrast to noise ratios and varying spatial resolution. Also investigated was the behaviour with respect to the normalisation exponent. The sensitivity of the method to a heterogeneous background distribution was performed in the simulated images mimicking the biodistribution of ^{177}Lu -DOTATATE. Finally the operator dependence, *i.e.* the initialisation dependence, on the final segmentation result was tested in the set of clinical images.

Result: The segmentation method was generally found to be resistant to noise. For a contrast to noise ratio of approximately unity combined with a high enough spatial resolution, the mean distance between the segmented surface and a reference surface was below one voxel side. With a poorer spatial resolution the result deteriorated for contrast to noise ratios of approximately one to two. When a background was added to the images, imitating a clinical activity distribution, the segmentation method showed some sensitivity to nearby high signal objects, for instance for the spleen which in some patients is in close to the left kidney. The operator dependence of the method when applied to clinical images was found to be low. No obvious advantage was observed when CT- images were combined with SPECT-images as the basis for the segmentation.

Discussion: The segmentation method gives a good agreement with the original volume segmented and has the potential of being useful as a tool for clinical SPECT-images. It is insensitive to noise, and is particularly suited for smooth images, such as those obtained with SPECT imaging. The inclusion of CT-images was found not to contribute to the precision of the method. A possible explanation is that the kidneys are in themselves smooth objects and no basic information of the object characteristics is added by the CT-images. Another explanation may be that the weighting used of the SPECT- and CT-based information was suboptimal. In cases of nearby high signal objects close to the object of interest, the method showed a tendency of incorporating the nearby object. One way to handle this problem could be to incorporate *a priori* information into the segmentation model, of interest for future work.

Author: **Johan Gustafsson**

Supervisors: **Katarina Sjögren Gleisner** and **Michael Ljungberg**

Master of Science thesis in medical radiation physics 30 hp, spring 2011

Department of Medical Radiation Physics, Lund University

Table of contents

1.	Introduction	1
2.	Background	1
2.1	SPECT-images	1
2.2	Radionuclide therapy with ¹⁷⁷ Lu-DOTATATE	1
2.3	Automatic segmentation of medical images	2
3.	Theory	2
3.1	Fourier descriptors	3
3.2	The downhill simplex method.....	6
3.3	Golden section search	7
3.4	The Otsu method.....	8
3.5	The Hotelling transform	8
4.	Materials and methods.....	9
4.1	Materials	9
4.2	Methods	12
4.3	Evaluation.....	22
5.	Results	28
5.1	Single organ images.....	28
5.2	Evaluation in activity distributions mimicking patients	36
5.3	Monte Carlo simulated SPECT-images	40
5.4	Operator dependence in the segmentation of clinical SPECT-images.....	44
6.	Discussion	45
7.	Conclusions	53
8.	Acknowledgements	53
9.	Appendix	53
10.	References	56

1. Introduction

Traditionally, segmentation of medical images is made by hand. Manual segmentation has two major disadvantages; it is time consuming and subjective. Therefore an automatic method for segmentation of images can be useful. It is hard to incorporate the full knowledge and skill of an experienced human observer into an automatic method. Still, the reproducibility and objectivity of an automatic method could in some situations prove to be advantageous. Furthermore, an automatic method can be designed to work with a full three dimensional volume rather than the classical human approach, where a volume is considered slice by slice.

Commonly used automatic segmentation methods are different types of thresholding, where a voxel is classified to belong to a certain object based on the voxel value. The cut off value for an object can be found by studying the image histogram [1]. Such an approach is unfortunately sensitive for the actual choice of cut off level, *i.e.* the threshold, and is therefore unsuitable for images exhibiting a low signal to noise ratio and poor spatial resolution. An example of an image showing these poor image characteristics would be SPECT¹-images in nuclear medicine.

One reason for wishing to do segmentation based on nuclear medicine images comes from the needs of radionuclide therapy, where the mean absorbed dose to an organ is quantified from images of the activity distribution. A correct segmentation of the image is therefore important for the correct assessment of absorbed dose. In this work, a method for automatic image segmentation of SPECT/CT-images based on Fourier descriptors is presented and investigated, with the aim to use the method for the delineation of kidneys in patients undergoing radionuclide therapy with the pharmaceutical ¹⁷⁷Lu-DOTATATE.

Fourier descriptors is a mathematical tool for describing curves and surfaces using Fourier series. Since most of the information about a function is contained in its first Fourier coefficients, a smoothly appearing surface can be accurately described using a limited number of coefficients, thereby reducing the number of parameters needed in the computations. A further advantage of using Fourier descriptors is that, though truncated on the Fourier side, the surface description becomes continuous in spatial space, permitting an arbitrary fine discretisation of the surface at realisation.

Obtaining a surface representing the object could in some instances be of interest in itself, but for the application of using the segmentation results for activity quantification in SPECT-images, the surface description of the organ must be converted into a description based on the voxels enclosed by the surface. This conversion should preferably take advantage of the continuous surfaces that Fourier descriptors produce.

2. Background

2.1 SPECT-images

The gamma camera, invented by Anger in the 1950 [2], is a device for imaging the spatial distribution of a radionuclide in an object. The technique can be used for obtaining planar images as well as for tomographic imaging, where a number of projections in different angles of the object are collected and a three dimensional volume is obtained through either filtered back projection or an iterative method such as MLEM² or OSEM³. The iterative methods may also incorporate compensations for attenuation, scattering and collimator response as a part of the reconstruction. The tomographic use of the camera is often referred to as Single Photon Emission Computed Tomography, or SPECT.

2.2 Radionuclide therapy with ¹⁷⁷Lu-DOTATATE

Radionuclide therapy with ¹⁷⁷Lu-DOTATATE is a method for treatment of patients with neuroendocrine tumours expressing somatostatin receptors. The radioactive nuclide ¹⁷⁷Lu has a physical half-life of 6.7 days and decays emitting beta particles with a maximum energy of 149 keV. It also emits several gamma photons with discrete energies, where one of the energies is at 208 keV. The dose limiting organ for the therapy is considered to be the kidneys since they retain the radiopharmaceutical [3]. The gamma emission of ¹⁷⁷Lu can be used for quantification of the activity in the kidneys and evaluation of the absorbed dose.

¹Single Photon Emission Computed Tomography

²Maximum Likelihood Expectation Maximisation

³Ordered Subset Expectation Maximisation

The quantification can be based on gamma camera planar images as well as SPECT-images or a combination of the two techniques [3].

2.3 Automatic segmentation of medical images

The general goal of segmentation of an images is to divide the image into two or more regions such as an object and a background [1]. This task can be done either manually or by some automatic or semiautomatic method. For most automatic methods, the segmentation is normally confined to rely on a single image, but it is also possible to use several images or images recorded from different modalities for the segmentation task, so called multispectral or multimodal segmentation, for example by using MR images with different contrast weighting [1].

Several methods, of varying complexity, have been described for automatic segmentation of medical images. A review by Pham *et al.* [4] included methods such as thresholding, region growing, classifiers, clustering, Markov random field models, artificial neural networks, atlas guided approaches and deformable models. As stated by Rogowska [1], no single segmentation technique can be used for all purposes. Because of the variation in image data and variation in purpose of the segmentation, different techniques are optimal in different situations.

Deformable models, one of the techniques mentioned above, can be implemented for both two and three dimensional images. It is a segmentation method where a curve or a surface is attracted to some image characteristics captured in an objective function derived from the image. Also, some internal forces of the model favouring for example smoothness [4, 5] can be incorporated, and possibly a statistically based model for *a priori* information about the object [6-8]. One of them most well-known example of a deformable model, meant for segmentation of two dimensional images, is so called snakes introduced by Kass *et al.* [9], given their name from their movement as the structure tried to minimise its associated objective function.

The main advantage of deformable models, as described by Boudraa and Zaidi [10], is that they can be designed to always form closed contours or surfaces, and due to the incorporation of smoothness constraints they can be relatively resistant to image noise. Among the disadvantages is the need for user interaction when placing the initial contour fairly close to the final result and determining the right parameters for the segmentation.

Automatic segmentation of nuclear medicine images

Automatic methods for segmentation of nuclear medicine images are often limited by the typically poor spatial resolution and high noise levels, and in the literature, the number of reports in this field are relatively few. Applications of image segmentation in nuclear medicine mentioned by Boudraa and Zaidi [10] included *e.g.* organ volume assessment, determination of parameters such as the left ventricular volume, dynamic functional imaging and imaging in connection with therapy. Techniques that have been used for segmentation of nuclear medicine images, also as listed by Boudraa and Zaidi, included for example thresholding, region growing, clustering, Markov random fields and deformable models. King *et al.* [11] made an investigation of segmentation using thresholding based on the maximum count in the object, and edge detection using gradient operators. Deformable models have for example been utilised in nuclear medicine for semi-automatic segmentation of the cortex in PET brain imaging as reported by Mykkänen *et al.* [12]. In this study the automatic method was found to have several advantages compared to manual delineation for the purpose of quantification.

3. Theory

The segmentation method used in this work includes a series of operations, of varying theoretical complexity. In this section some of the more complex theoretical concepts used within the segmentation method are described. In order to put the different theoretical concepts into the context of the segmentation method, a brief summary is given in Table 1.

Table 1: The use of theoretical concepts and methods for different aspects of this work

Aspect	Procedures	Methods and concepts used
1. Initialisation	Identification of the organ to be segmented, estimation of ellipsoid parameters that approximate the organ and creation of grey level gradient images	<ul style="list-style-type: none"> • The Otsu method for automatic selection of threshold • The Hotelling transform for alignment of the object the with external coordinate system • Downhill simplex numerical optimisation method
2. Optimisation	Optimisation of the surface to image data	<ul style="list-style-type: none"> • Fourier descriptors for surface description • Downhill simplex numerical optimisation method
3. Surface to volume conversion	Conversion from closed surface to contained volume	<ul style="list-style-type: none"> • Golden section search
4. Evaluation	Evaluation of the segmentation method	<ul style="list-style-type: none"> • Golden section search

3.1 Fourier descriptors

Fourier descriptors in two dimensions

The Cartesian coordinates, x and y , of a planar curve in two dimensions can be represented by two functions of some parameter t , $x(t)$ and $y(t)$, or expressed as a single vector valued function [13],

$$\mathbf{R}(t) = (x(t), y(t)), t \in [0, 2\pi]. \quad (1)$$

For a closed curve, the functions $x(t)$ and $y(t)$ are obviously periodic. It is then natural to describe the curve using the Fourier coefficients of its coordinate functions.

Fourier descriptors have previously been used for describing deformable curve models for segmentation tasks in two dimensions, as described by Staib and Duncan [14] who also incorporated *a priori* information into the model by biasing the objective function used for adapting the curve to the object of interest towards the average shape of the object. In their work the technique was described not only for closed curves, but also for open curves. For open curves, the Fourier parameterisation was restricted to a cosine series and the curve traced in one direction in the first part of the period and then in the other direction during the second part of the period. In this way, discontinuities were avoided in the curve description.

If possible, discontinuities are to be avoided since these require the inclusion of high order coefficient for a good approximation. In practice, a truncated description must always be used; *i.e.* the Fourier series cannot have an infinite number of terms, and the fewer coefficients that are needed, the easier the curve description. In theory, the number of Fourier coefficients needed in order to approximate a curve to a given accuracy can be predicted theoretically [15], but, as shown by Kuhl and Giardina [16], the mathematical expression, which was described as a maximum error, when evaluated for a range of shapes overestimated the actual error by at least a factor two.

A version of snakes based on Fourier descriptors has been described by Székely *et al.* [17] who also used the technique for segmentation of brain structures in MRI data. The authors recognised that the natural truncation of the Fourier series to a finite number of Fourier coefficients imposed a regularisation of the obtained contour, but pointed out that even surfaces of low order were capable of producing often undesired effects such as sharp corners and self-crossing. In order to avoid the first effect, the authors incorporated an internal energy of the snake in the objective function making high derivatives of the functions describing the curve unfavourable for the optimisation process. The authors also proposed a way to incorporate *a priori* knowledge into the segmentation model in a way that differed from the one proposed by Staib and Duncan. Instead of biasing the objective function, the solutions were restricted to the most prominent modes of variation obtained from a principal component analysis of the results in the segmentations of a training set of images.

The meaning of t

In the article by Staib and Duncan [14] t in equation (1) was suggested to correspond to arclength, s , by setting

$$t(s) = \frac{2\pi s}{S}, \quad (2)$$

where S is the total length of the curve. As pointed out by the authors, the interpretation as arclength generally does not hold exactly for curves not described by infinite Fourier series except for some special cases such as circles, with reference to Persoon and Fu [18]. The latter pointed out that other parameterisations than arclength correspond to a nonhomogeneous parameterisation of the curve and that such a parameterisation could be more efficient in terms of describing a curve to a predefined accuracy with a low number of Fourier coefficients. The complication of interpreting t as arclength is realised by noting that

$$\frac{ds}{dt} = \sqrt{\left(\frac{dx}{dt}\right)^2 + \left(\frac{dy}{dt}\right)^2} \quad (3)$$

is not constant for a general finite Fourier series. For the special case of a series describing a circle of radius r and centre point at (a_0, b_0) , *i.e.*

$$\begin{cases} x = a_0 + r \cos t \\ y = b_0 + r \sin t \end{cases}, \quad (4)$$

equation (3) does reduce to a constant, which is easily proven by applying the Pythagorean trigonometric identity. By a similar argument, it is also realised that equation (3) will not return a constant value when given for example a canonical ellipse, *i.e.* equation (4) with r replaced by two different coefficients for x and y respectively, as input.

In conclusion, it is not clear what meaning to give to the parameter running along the curve, which is equally true for parameters spanning a surface. It is the experience in this work that an exact interpretation of the parameter is not crucial for the implementation of Fourier descriptors as the basis for deformable models. To think of the parameter as just a way to move between different points at the curve, without considering its geometrical meaning, is sufficient for the application.

Fourier descriptors in three dimensions

The Cartesian coordinates of a surface in three dimensions can be represented by three functions of two parameters, $x(u, v)$, $y(u, v)$ and $z(u, v)$, or, expressed as a single vector valued function,

$$\mathbf{R}(u, v) = (x(u, v), y(u, v), z(u, v)). \quad (5)$$

The parameters u and v can be constructed to lie in the interval $[0, 2\pi)$. The coordinate functions can be written as Fourier series according to [13, 19]

$$\begin{aligned} x(u, v) = & a_{x,0,0} + 2 \sum_{m=1}^{\infty} (a_{x,m,0} \cos mu + b_{x,m,0} \sin mu) + 2 \sum_{l=1}^{\infty} (a_{x,0,l} \cos lv + c_{x,0,l} \sin lv) + \\ & + 4 \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} (a_{x,m,l} \cos mu \cos lv + b_{x,m,l} \sin mu \cos lv + c_{x,m,l} \cos mu \sin lv + d_{x,m,l} \sin mu \sin lv) \end{aligned} \quad (6)$$

$$\begin{aligned}
y(u, v) = & a_{y,0,0} + 2 \sum_{m=1}^{\infty} (a_{y,m,0} \cos mu + b_{x,m,0} \sin mu) + 2 \sum_{l=1}^{\infty} (a_{y,0,l} \cos lv + c_{y,0,l} \sin lv) + \\
& + 4 \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} (a_{y,m,l} \cos mu \cos lv + b_{y,m,l} \sin mu \cos lv + c_{y,m,l} \cos mu \sin lv + d_{y,m,l} \sin mu \sin lv)
\end{aligned} \tag{7}$$

$$\begin{aligned}
z(u, v) = & a_{z,0,0} + 2 \sum_{m=1}^{\infty} (a_{z,m,0} \cos mu + b_{x,m,0} \sin mu) + 2 \sum_{l=1}^{\infty} (a_{z,0,l} \cos lv + c_{z,0,l} \sin lv) + \\
& + 4 \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} (a_{z,m,l} \cos mu \cos lv + b_{z,m,l} \sin mu \cos lv + c_{z,m,l} \cos mu \sin lv + d_{z,m,l} \sin mu \sin lv),
\end{aligned} \tag{8}$$

where $a_{x,m,l}$, $b_{x,m,l}$, $c_{x,m,l}$ and $d_{x,m,l}$ and the corresponding coefficients in the y- and z-functions are scalar valued parameters. There is a direct link between this real version of two dimensional Fourier series and the, perhaps more familiar, version that uses complex basis functions. The mathematical connection is described in the appendix. In the appendix the factors 2 and 4 that appear in equations (6) to (8) are also explained as a consequence of the conversion between complex and real Fourier series.

The underlying assumption of using equation (6) to (8) is that the coordinate functions are periodic with period 2π in both parameters. In general this will describe a torus, but these tori may be collapsed in certain aspects that allow other shapes to be formed. With certain restrictions placed on the Fourier coefficients, open, tubular and closed surfaces can be produced [19].

Closed surfaces can be described by the equations [13, 19, 20]



$$x(u, v) = a_{x,0,0} + 2a_{x,0,1} \cos v + 2 \sum_{l=1}^{\infty} c_{x,0,l} \sin lv + 4 \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} (c_{x,m,l} \cos mu \sin lv + d_{x,m,l} \sin mu \sin lv) \tag{9}$$

$$y(u, v) = a_{y,0,0} + 2a_{y,0,1} \cos v + 2 \sum_{l=1}^{\infty} c_{y,0,l} \sin lv + 4 \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} (c_{y,m,l} \cos mu \sin lv + d_{y,m,l} \sin mu \sin lv) \tag{10}$$

$$z(u, v) = a_{z,0,0} + 2a_{z,0,1} \cos v + 2 \sum_{l=1}^{\infty} c_{z,0,l} \sin lv + 4 \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} (c_{z,m,l} \cos mu \sin lv + d_{z,m,l} \sin mu \sin lv), \tag{11}$$

where $u \in [0, 2\pi)$ and $v \in [0, \pi]$.

In equation (9) to (11) the surface is described using both sine and cosine functions in the u-parameter, while the surface is described using sine series in the v-parameter except for the second terms in the formulas. The v-parameter is further restricted to the interval $[0, \pi]$ while the u-parameter takes values in the full interval $[0, 2\pi)$. This difference can be understood in terms of that the coordinate functions are basically described by sine series in the v-direction and are thus antisymmetric with respect to v . The incorporation of the $\cos v$ terms in the Fourier series destroys the symmetry, but still values of v outside the interval $[0, \pi]$ will not form a part of the closed surface, but rather a ‘mirror’ surface, as illustrated in Figure 1 and Figure 2.

Figure 1: Surface with ν lying in the interval $[0, \pi]$ Figure 2: Surface with ν lying in the interval $[0, 2\pi)$. The surface shown in Figure 1 is accompanied by a surface described by ν in the interval $(\pi, 2\pi)$ seen as the 'outgrowth' in the image.

Fourier descriptors for image segmentation in three dimensions

Fourier descriptors have previously been used for segmentation based on deformable surface models in three dimensions by Staib and Duncan [19], Gao *et al.* [21] and Floreby *et al.* [20, 22], the later aiming to use the technique in nuclear medicine images. In the article by Staib and Duncan, a method for biasing the surface towards its average shape was proposed, thereby incorporating *a priori* knowledge into the segmentation in the same ways as was done for the two-dimensional case [14]. As pointed out by Staib and Duncan the parametric description of an object can be accurate with relatively few parameters and does in a natural way consider the surface as a whole, as opposed to surface descriptions using a decomposition of the object.

Criticism against the use of Fourier descriptors for segmentation tasks has been expressed by Brechbühler *et al.* [23]. The objections involved the fact that the two parameters spanning the surface have different characteristics and that the parameterisation of the surface does not become homogenous. Special problems occur at the poles of the objects, where the surface description collapses to a single point. Brechbühler *et al.* instead proposed another surface parameterisation that involved a mapping of the object to be segmented onto a space with spherical topology, thereby circumventing some of their concerns.

3.2 The downhill simplex method

The downhill simplex method is used for optimisation both within the initialisation and boundary optimisation steps of the segmentation method. A simplex in N dimensions is a set of $N+1$ vertices that must not lie in the same hyper plane [13]. The downhill simplex method, or Nelder-Mead simplex-reflection method or amoeba method as it is also called, is a method for numerical optimisation that does not need evaluation of derivatives of the function to be optimised. Instead a simplex is spanned according some criterion and is then evolved based on the function values in its vertices [13, 24]. In a single iteration of the method, the value of the function is first evaluated in the $N+1$ vertices and sorted. Then the algorithm tries to modify the position of the poorest vertex in a series of reflections, expansions or contractions along the line joining the poorest vertex with the centroid of the other N vertices. If the poorest vertex cannot advance in the ordering of vertices through these operations, then the whole simplex is shrunk towards the best vertex [24]. When a new and better set of coordinates for the poorest vertex has been found, or the simplex has been shrunk, the process is repeated until a convergence criterion has been reached.

The downhill simplex method has been used as the optimisation method for adapting Fourier descriptor surfaces to objects in nuclear medicine images by Floreby *et al.* [20, 22], who also compared the method with conjugate gradient optimisation, another method for numerical optimisation but that unlike the downhill simplex method is based on evaluation of the function gradient, and to a hybrid scheme of the two methods. The downhill simplex method was in this study found to be superior to pure conjugate gradient optimisation. The hybrid optimisation scheme was found to give comparable results to the downhill simplex method, but could also use the faster convergence of the conjugate gradient method, as compared to the downhill simplex method, for later stages of the optimisation process where many Fourier coefficients were used for the surface description, and the optimisation problem thus had increased in computational complexity compared to the initial stages.

3.3 Golden section search

Golden section search is used both in the method for conversion of the segmented boundary to a volume description, and within the evaluation of the segmentation method. It is a method used for gradient free numerical optimisation in one or several dimensions of a function with a single local optimum in the interval of interest. The explanation of the method given below is adapted from Chang [25].

For the case of a minimisation of a one dimensional function, suppose $f(x)$ to be a function with a unique minimum in the interval $[a, b]$. In golden section search, the interval $[a, b]$ is first evaluated in two points according to

$$x_1 = a + (1-r)(b-a) \quad (12)$$

$$x_2 = a + r(b-a), \quad (13)$$

where

$$r = \frac{\sqrt{5}-1}{2} \quad (14)$$

is the inverse of the golden ratio. If in the evaluation of these two points it is found that $f(x_1) \geq f(x_2)$, then the minimum must lie in the interval $[x_1, b] = [a + (1-r)(b-a), b]$. In the next step, the function is to be evaluated accordingly in the interval $[x_1, b]$, *i.e.*

$$\begin{aligned} x_1^{\text{new}} &= a + (1-r)(b-a) + (1-r)(b-a - (1-r)(b-a)) = a + (1-r + (1-r)(1-1+r))(b-a) = \\ &= a + (1-r^2)(b-a) \end{aligned} \quad (15)$$

and

$$\begin{aligned} x_2^{\text{new}} &= a + (1-r)(b-a) + r(b-a - (1-r)(b-a)) = a + (1-r + r(1-1+r))(b-a) \\ &= a + (1-r+r^2)(b-a), \end{aligned} \quad (16)$$

where x_1^{new} and x_2^{new} are the points splitting the interval $[x_1, b]$ according to the rule in equations (12) and (13). According to the definition of r as the inverse of the golden ratio it is true that $1-r = r \cdot r = r^2$ and thus

$$x_1^{\text{new}} = a + (1-r^2)(b-a) = a + r(b-a) = x_2. \quad (17)$$

Therefore, except for the first step of the optimum search, only one function evaluation has to be performed for each step, *i.e.* the number of function evaluations are almost halved compared to an arbitrary choice of the points splitting the interval. The same is true by a similar argument had it been found that $f(x_2) > f(x_1)$.

In the case of a function of two variables, $f(x,y)$, that has a unique minimum in the interval $[a, b]$ and $[c, d]$ for x and y respectively. Analogously to the one dimensional case, the intervals are divided according to

$$x_1 = a + (1-r)(b-a) \quad (18)$$

$$x_2 = a + r(b-a) \quad (19)$$

$$y_1 = c + (1-r)(d-c) \quad (20)$$

$$y_2 = c + r(d-c). \quad (21)$$

The first step in the search for the optimum point requires four function evaluations to be performed. For later steps in the process only three function evaluations will be necessary by the same sort of arguments as for the one dimensional case.

3.4 The Otsu method

The Otsu method is an automatic method for finding a global threshold value for segmentation of an image containing an object with a certain distribution of voxel values and a background with another distribution of voxel values [10, 26]. The method has been used by Mortelmans *et al.* [26] for segmentation of SPECT-volumes. It is based on the principle of maximising the between class variance of the image histogram by adjusting the separation threshold between object and background, which, for a threshold k , can be calculated as

$$\sigma_B^2 = \omega_0(\mu_0 - \mu_T)^2 + \omega_1(\mu_1 - \mu_T)^2, \quad (22)$$

where, if p_i corresponds to the probability of a voxel having the i th grey level of a total of L grey levels,

$$\omega_0 = \sum_{i=1}^k p_i \quad (23)$$

$$\omega_1 = \sum_{i=k+1}^L p_i \quad (24)$$

$$\mu_0 = \sum_{i=1}^k \frac{ip_i}{\omega_0} \quad (25)$$

$$\mu_1 = \sum_{i=k+1}^L \frac{ip_i}{\omega_1} \quad (26)$$

$$\mu_T = \sum_{i=1}^L ip_i. \quad (27)$$

The expressions of equation (23) to (27) are to be interpreted as: (23) the cumulative probability for a voxel to have a value causing it to be classified as a part of the background according to the current threshold selection, (24) the cumulative probability for a voxel to have a value causing it to be classified as a part of the object according to the current threshold selection, (25) the mean value for voxels classified as background, (26) the mean value for voxels classified as object, and (27) the mean value for all voxels respectively.

In the method by Mortelmans *et al.*, all possible thresholds were evaluated with respect to equation (22), and the one maximising the result was chosen as the threshold for segmentation of the image.

3.5 The Hotelling transform

The Hotelling transform reorients an image object so its principal axes of variation, assumed to be orthogonal, are aligned with the x-, y-, and z-axes of the coordinate system. The description of the method follows the one given by Gonzalez and Woods [27].

In the context of the Hotelling transform, the coordinates of a three dimensional binary image are considered as a set of M three dimensional random vectors. The expectation value of the vector distribution is defined as the expectation value of the vector coordinates. The expectation value, $\mu_{\mathbf{x}}$, can be estimated as the mean value, $\mathbf{m}_{\mathbf{x}}$, according to

$$\mathbf{m}_x = \frac{1}{M} \sum_{k=1}^M \mathbf{x}_k. \quad (28)$$

The covariance matrix of the \mathbf{x} -vectors, \mathbf{C}_x , *i.e.*

$$\mathbf{C}_x = E\left\{(\mathbf{x} - \boldsymbol{\mu}_x)(\mathbf{x} - \boldsymbol{\mu}_x)^T\right\}, \quad (29)$$

where $E(\cdot)$ denotes the expectation value and T denotes the transpose of a matrix, can be estimated as

$$\tilde{\mathbf{C}}_x = \frac{1}{M-1} \sum_{k=1}^M (\mathbf{x}_k \mathbf{x}_k^T - \mathbf{m}_x \mathbf{m}_x^T). \quad (30)$$

The Hotelling transform is the transformation of the vector \mathbf{x} representing the original image object coordinates to new coordinates \mathbf{y} according to

$$\mathbf{y} = \mathbf{T}(\mathbf{x} - \mathbf{m}_x), \quad (31)$$

where \mathbf{T} is the matrix whose rows are the normalised eigenvectors of the covariance matrix of the \mathbf{x} -vectors. After transformation of an object with equation (31) the principal axes of variation of the object are aligned with the axes of the coordinate system.

Since \mathbf{T} by definition is an orthogonal matrix its inverse is given by its transpose. Therefore, with $\mathbf{M}=\mathbf{T}^T$, the inverse of equation (31) is

$$\mathbf{x} = \mathbf{M}\mathbf{y} + \mathbf{m}_x. \quad (32)$$

4. Materials and methods

In this section the method for segmenting SPECT- and SPECT/CT-images is presented together with methods for evaluation of the segmentation result. The first part describes the image material, the second part describes the segmentation method, and the third part describes the evaluation of the method.

4.1 Materials

Patient images

Five SPECT-series of patients undergoing therapy with ^{177}Lu -DOTATATE were obtained. For each case, SPECT-studies were performed at 0.5 h, 24 h, 96 h and 168 h post injection of the pharmaceutical. The images were obtained with a Discovery VH SPECT system with a Hawkeye CT unit. Acquisition was performed using a medium energy collimator for the 208 keV photons of ^{177}Lu , and was made in 60 projections in full rotation mode. Projection images were reconstructed into a tomographic volume in a 128x128x128 matrix with 4.02 mm voxel size, using OSEM with eight iterations and six angles per subset, employing attenuation correction, scatter correction and collimator response function compensation in the reconstruction. The five image series were obtained from three different patients, *i.e.* two of the patients were imaged twice in different therapy sessions.

Activity distribution images

The biodistribution of ^{177}Lu -DOTATATE was estimated from literature for blood and urine [28-30] and from clinical patient SPECT-images for liver, spleen, renal cortex, renal medulla and background at approximately 0.5 h, 24 h, 96 h and 168 h post injection. The activity distribution was partly obtained from planar images from 25 therapies and partly from the five cases of SPECT-series described above. As an example of when the SPECT-images were used for finding typical activity concentrations were the relative distribution of activity between the renal cortex and the renal medulla. The quotient between these two compartments could then be used at the activities obtained for the kidney as a whole from planar images.

The relative organ concentrations were transcribed to the male version of the XCAT phantom [31], and 128x128x128 images created with 4x4x4 mm³ voxels. The spatial resolution of a SPECT-system was simulated by convolution with a Gaussian kernel for two values of the FWHM, 12 and 16 mm. Independent Gaussian noise was added to the images with a signal to noise ratios in different compartments corresponding to those estimated from the patient SPECT-images in the renal cortex, the renal medulla, the liver, the spleen and an area representing the background.

Parallel to the activity images, a density map was also created using the XCAT computer phantom. A mass density map was created by annotating values obtained from ICRU 44 [32] to the different organs. A CT-like image was then created by mapping the density values to CT-numbers according to the relation

$$H = 0.96799 \cdot \rho + 15.16571 \quad (33)$$

for densities below 1.070 kg/dm³ and

$$H = 1.36883 \cdot \rho - 402.86418 \quad (34)$$

for densities above 1.070 kg/dm³, where ρ is the tissue density in gram per cubic decimetre and H is the Hounsfield number plus a value of 1 000. After remapping, independent normal distributed noise was added to the image with a standard deviation of 14. The standard deviation was chosen by estimation of the noise in the CT-images obtained together with the patient SPECT-images. The background density of the body was set to 0.95 g/cm³ and the kidneys to a density of 1.05 g/cm³, resulting in CT-numbers of 935 and 1 032 respectively. The reason for not setting the background to unit density, *i.e.* to water density, was that it resulted in a poorer contrast for the kidneys, as compared to the patient images at visual inspection. The reason for the good contrast in the patient CT-images was thought to be due to adipose tissue.

An example of a simulated SPECT-image together with a corresponding patient example is shown in Figure 3 and Figure 4. An example of a CT image and a corresponding clinical case is shown in Figure 5 and Figure 6.




Figure 3: Coronal slice of a XCAT phantom simulating a SPECT image at 30 minutes post injection with ¹⁷⁷Lu-DOTATATE




Figure 4: Coronal slice of a patient image corresponding to the simulated image in Figure 3.




Figure 5: Coronal slice of a simulated CT-image




Figure 6: Coronal slice of patient CT-image corresponding to the simulated example in Figure 5

Monte Carlo simulated SPECT-images

Four SPECT studies were created using the SIMIND Monte Carlo program [33]. The phantoms in the simulations were basically the same as the ones used in the previous section, but with a finer voxel discretisation of $2 \times 2 \times 2 \text{ mm}^3$. An example of a simulated result showing a simulated male patient approximately 0.5 h post injection is shown in Figure 7.

Sixty noise free projections were simulated and the image series was then rescaled to different levels of total number of counts. The number of counts to scale the results was estimated from a clinical case for each point in time and decided to correspond to 20 million counts, eight million counts, three million counts and two million counts for points in time 0.5 h, 24 h, 96 h and 168 h post injection respectively. The projection voxel values were then replaced by random numbers drawn from Poisson distributions with the original voxels values as their means. Reconstruction was performed in $128 \times 128 \times 128$ matrices with $4 \times 4 \times 4 \text{ mm}^3$ voxels using OSEM with eight iterations and six angles per subset employing compensation for attenuation and scatter. In order to manage noise, the reconstructed images were post processed with a three dimensional Gaussian filter with an FWHM of two voxel sides. Reconstruction using collimator response function compensation was also tested, but the images reconstructed with this compensation were found to produce systematically too small volumes of the kidneys when the images were segmented independent of the parameters used for the segmentation procedure. Since the reason for this behaviour was unknown and a similar behaviour was not observed for the images reconstructed without collimator response function compensation, the use of this compensation was not considered further.

A CT-volume similar to the one used in the previous section was created parallel to the SPECT-volumes. The only difference was that this density volume had a voxel size of $2 \times 2 \times 2 \text{ mm}^3$, which was resampled to $4 \times 4 \times 4 \text{ mm}^3$ prior to segmentation.




Figure 7: Example of a Monte Carlo simulated male patient approximately 30 min after injection of ^{177}Lu -DOTATATE. The projections had a total number of 20 million counts. The volume was post processed with a Gaussian filter in order to suppress image noise.

Single organ images

A binary volume was created based on the left kidney of the XCAT phantom [31] including the renal medulla and renal cortex. In the XCAT phantom, blood vessels can be excluded from the organ description. This was done in order to get a solid volume for comparison, but even without blood vessels the organ was deemed too irregular for being a good test object for the performance of a segmentation method aimed for outlining smooth image objects. The problem was due to the renal pelvis, whose inclusion was deemed to form an unrealistic outgrowth but whose exclusion formed a hole in the organ. In order to assure the test volume to be solid and representative of how objects appeared in clinical images, a surface was segmented to the original volume with the method described below using six Fourier orders. The created surface was then used as the basis for the binary test volume. The test volume was created using a $128 \times 128 \times 128$ matrix with $4 \times 4 \times 4 \text{ mm}^3$ voxels.

Test images were created by blurring the test volume using a $9 \times 9 \times 9$ Gaussian kernel with varying FWHM. After blurring, independent Gaussian noise with zero mean was added to the image. The standard deviation of the noise was varied thereby creating images with different contrast to noise ratios (CNR). Examples of the appearance of the binary object in the final images are shown in Figure 8.

The contrast of the image was defined as the absolute difference between the object and the background before blurring. In practice, the object always had a value of unity and the background always had a value of zero, thereby forming a contrast of unity according to the stated definition. The rationale for using this

definition of contrast was that the segmentation method described below was dependent on the gradient of the original images. Data directly used for segmentation was thus only due to the signal difference between object and background and thereby not affected by any constant background added to the image. The contrast to noise ratio was then defined as the quotient of the contrast and the standard deviation of the noise distribution.

As a complement to the kidney based single organ images, a binary version of the XCAT phantom spleen was created also excluding blood vessels. This organ was deemed more complex due to its irregular shape but still smooth enough to be used as a realistic test object for the segmentation method. The intermediate step of fitting a surface to the original organ and then base the test object on the result from that segmentation was thus not necessary. Just as for the kidney based test object, the binary spleen was blurred with Gaussian kernels of varying FWHM and normal distributed noise was added in order create test images of varying spatial resolution and contrast to noise ratios.



Figure 8: The appearance of the kidney based binary object in images with 4 mm spatial resolution (upper row) and 12 mm spatial resolution (lower row) for CNR ranging from eight (left) two one (right) in order eight, six, four, two and one.

4.2 Methods

Surface description

A closed surface laying the foundation for a deformable surface model was described using equations (9) to (11). This section describes the practical implementation of the surface description for the coming optimisation of the model to an image object. For an exact description of a surface, generally an infinite number of coefficients are necessary in the Fourier expansions. For practical purposes, only a finite number can be used, *i.e.* instead of infinity, the upper limits of the double sums in equations (9) to (11) are replaced by K_1 and K_2 for sums over m and l respectively. There is in principle nothing that prevents K_1 and K_2 to be different numbers, but in this work K_1 and K_2 always had a common value of $K_1=K_2=K$. This number will be referred to as the order of the surface or the number of Fourier orders used.

Organising the Fourier coefficients

The Fourier coefficients in equations (9) to (11) were organised in three matrices, **A**, **C** and **D**, for each function, x , y and z for the a -, c - and d -coefficients respectively with the m -indices organised by row and the l -indices organised by column, *i.e.* there were nine matrices in total. For a surface description of orders K_1 and K_2 these matrices were $K_1 \times K_2$ matrices. The matrices were indexed from zero to K_1-1 and zero K_2-1 in rows and columns respectively. These matrices will be referred to as Fourier coefficients matrices.

It is realised that some of the elements in the Fourier coefficient matrices do not represent coefficients included in equation (9) to (11), for example the equations do not include any other coefficients contained in the **A** matrix than $a_{0,0}$ and $a_{0,1}$. In order to enable these other elements to be constantly zero and in order to simplify the optimisation of the Fourier coefficients, the elements were reorganised into a vector

\mathbf{p} with $3(2 + (2K_1 - 1) \cdot (K_2 - 1))$ elements [13] prior to the optimisation process. All coefficients not present in equations (9) to (11) were excluded from \mathbf{p} . The Fourier coefficient matrices were transformed to \mathbf{p} and vice versa using a system of binary matrices parallel to the Fourier coefficient matrices, where all elements allowed to vary were marked by 1 and all elements that were to be constant zero were marked by 0.

Initialisation

A first approximation of the surface of an organ was made using an ellipsoid, which can be described by equation (9) to (11) with $K=2$. In order to describe an ellipsoid, a centre point, three semi axes and an orientation in three dimensions for the object had to be determined. These parameters were obtained using 1) a presegmentation algorithm based on the Otsu method, followed by 2) a reorientation of the object using the Hotelling transform and finally, 3) an optimisation process where an ellipsoid was fitted to the binary reoriented data.

Presegmentation

The first step for finding an approximate ellipsoid was to get a binary image of the organ, *i.e.* to do a presegmentation of the organ. This presegmentation was based on thresholding using the Otsu method described above. A rough first outline of the organ was made by hand in a graphical user interface, producing a volume of interest that was assumed to contain only the organ and background. The primary purpose of the manual delineation was to exclude nearby objects from the coming analysis in the initialisation process, thus only a crude outline of the organ was required, as shown in Figure 9. The interface was later also equipped with a function for nearest neighbour interpolation in the axial direction for intermediate slices where the organ was not outlined. Outlining around 5 slices was typically sufficient, requiring about one minute for the operator.

The histogram of the marked volume was created, excluding zeros, and the threshold decided based on the maximisation of equation (22). The histogram of a volume of interest and the threshold obtained by the Otsu method is shown in Figure 10. In order to avoid getting a perforated volume due to image noise and inhomogeneous uptake of the radiopharmaceutical, the volume was post processed using a morphological closing operation [34].




Figure 9: Rough delineation of the left kidney as a part of the presegmentation process




Figure 10: Example of a histogram with the obtained threshold marked with a cross (x)

Finding the centre point and orientation of the organ

The binary organ obtained from the presegmentation process was reoriented by use of the Hotelling transform. By applying the Hotelling transform a new version of the binary image was created, where the principal axes of variation were aligned with the x-, y- and z-axes of the image coordinate system. Since voxel coordinates can only be assigned to integer values, the result from equation (31) had to be rounded off before used. This resulted in round off errors causing small holes in the object. In order to close these holes, another morphological closing operation was performed.

The Hotelling transform requires the mean in each dimension to be subtracted from the data set. These mean values were taken as coordinates for the centre point for the initial ellipsoid. The matrix of eigenvectors to the covariance matrix of the object coordinates describes the mapping of the object to its canonical position, so the inverse, \mathbf{M} in equation (32), of the matrix described the orientation of the ellipsoid in the original image.

Finding the length of the semi axes

The semi axes of the initialising ellipsoid was found by an optimisation process based on the binary image of the object with its principal axes of variation aligned with the axes of the coordinate system, as obtained from the Hotelling transform. An ellipsoid aligned with the axes of the coordinate system can be described as the set of points fulfilling the condition

$$\left(\frac{x-x_c}{x_1}\right)^2 + \left(\frac{y-y_c}{y_1}\right)^2 + \left(\frac{z-z_c}{z_1}\right)^2 = 1, \quad (35)$$

where (x_c, y_c, z_c) is the centre point of the ellipsoid and x_1, y_1 and z_1 are the semi axes of the ellipsoid in the x-, y- and z-direction, respectively. The interior of the ellipsoid, including its surface, is thus described by the condition

$$\left(\frac{x-x_c}{x_1}\right)^2 + \left(\frac{y-y_c}{y_1}\right)^2 + \left(\frac{z-z_c}{z_1}\right)^2 \leq 1. \quad (36)$$

The voxel coordinates containing the reoriented binary organ was tested against the inequality in (36) with variable values of the semi axes in an optimisation process. This optimisation was formulated to minimise the number of voxels not belonging to the binary organ inside the ellipsoid, and, at the same time, minimising the number of voxels belonging to the organ but lying outside the ellipsoid. If a voxel found to lie inside the ellipsoid had a value of zero, then a penalty of plus one was given to an objective function. Similarly, if a voxel with the value of one was found to lie outside the ellipsoid, again a penalty of plus one was given to the objective function.

A first guess of the length of the semi axes was made as half the length between the extreme values of the voxel coordinates in the reoriented binary image. The lengths were then optimised with respect to the objective function using the amoeba procedure in IDL [35].

Describing an ellipsoid using Fourier descriptors

An ellipsoid with its semi axes of lengths x_1, y_1 and z_1 aligned with the x-, y-axes and z-axes of the image can be described by the equations

$$\begin{cases} x(u, v) = x_c + x_1 \cos u \sin v \\ y(u, v) = y_c + y_1 \sin u \sin v \\ z(u, v) = z_c + z_1 \cos v \end{cases}, \quad (37)$$

where $u \in [0, 2\pi)$ and $v \in [0, \pi)$. It is thus realised that in terms of the coefficients in equation (9) to (11) an ellipsoid of this form is described by $a_{x,0,0} = x_c$, $a_{x,0,1} = z_1/2$, $a_{y,0,0} = y_c$, $a_{y,1,1} = y_1/4$, $a_{z,0,0} = z_c$ and $c_{z,1,1} = z_1/4$. So, in terms of Fourier coefficient matrices with $K=2$, the following description could be used:

$$\mathbf{A}_x = \begin{bmatrix} x_c & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{A}_y = \begin{bmatrix} y_c & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{A}_z = \begin{bmatrix} y_c & \frac{z_1}{2} \\ 0 & 0 \end{bmatrix} \quad (38)$$

$$\mathbf{C}_x = \begin{bmatrix} 0 & 0 \\ 0 & \frac{x_1}{4} \end{bmatrix}, \quad \mathbf{C}_y = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{C}_z = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \quad (39)$$

$$\mathbf{D}_x = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{D}_y = \begin{bmatrix} 0 & 0 \\ 0 & \frac{y_1}{4} \end{bmatrix}, \quad \mathbf{D}_z = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}. \quad (40)$$

The above description is only capable of describing an ellipsoid aligned with the coordinate axes. In order to get an arbitrary ellipsoid, the description must be reoriented. Simply applying the transformation to the (x,y,z) -vector describing the ellipsoid would not be correct, since the rotation is not to be applied with respect to the origin, but with respect to (x_c, y_c, z_c) . The rotation should instead be applied according to equation (35). In practice the rotation was applied to the Fourier coefficient matrices according to

$$\mathbf{A}'_x = \begin{bmatrix} x_c & 0 \\ 0 & 0 \end{bmatrix} + \mathbf{M}(0,2) \cdot \begin{bmatrix} 0 & \frac{z_1}{2} \\ 0 & 0 \end{bmatrix} \quad (41)$$

$$\mathbf{A}'_y = \begin{bmatrix} y_c & 0 \\ 0 & 0 \end{bmatrix} + \mathbf{M}(1,2) \cdot \begin{bmatrix} 0 & \frac{z_1}{2} \\ 0 & 0 \end{bmatrix} \quad (42)$$

$$\mathbf{A}'_z = \begin{bmatrix} z_c & 0 \\ 0 & 0 \end{bmatrix} + \mathbf{M}(2,2) \cdot \begin{bmatrix} 0 & \frac{z_1}{2} \\ 0 & 0 \end{bmatrix} \quad (43)$$

$$\mathbf{C}'_x = \mathbf{M}(0,0) \cdot \mathbf{C}_x + \mathbf{M}(0,1) \cdot \mathbf{C}_y + \mathbf{M}(0,2) \cdot \mathbf{C}_z \quad (44)$$

$$\mathbf{C}'_y = \mathbf{M}(1,0) \cdot \mathbf{C}_x + \mathbf{M}(1,1) \cdot \mathbf{C}_y + \mathbf{M}(1,2) \cdot \mathbf{C}_z \quad (45)$$

$$\mathbf{C}'_z = \mathbf{M}(2,0) \cdot \mathbf{C}_x + \mathbf{M}(2,1) \cdot \mathbf{C}_y + \mathbf{M}(2,2) \cdot \mathbf{C}_z \quad (46)$$

$$\mathbf{D}'_x = \mathbf{M}(0,0) \cdot \mathbf{D}_x + \mathbf{M}(0,1) \cdot \mathbf{D}_y + \mathbf{M}(0,2) \cdot \mathbf{D}_z \quad (47)$$

$$\mathbf{D}'_y = \mathbf{M}(1,0) \cdot \mathbf{D}_x + \mathbf{M}(1,1) \cdot \mathbf{D}_y + \mathbf{M}(1,2) \cdot \mathbf{D}_z \quad (48)$$

$$\mathbf{D}'_z = \mathbf{M}(2,0) \cdot \mathbf{D}_x + \mathbf{M}(2,1) \cdot \mathbf{D}_y + \mathbf{M}(2,2) \cdot \mathbf{D}_z, \quad (49)$$

where the primed matrices are the matrices after rotation and $\mathbf{M}(\cdot, \cdot)$ is an element of the rotation matrix in equation (32) indexed by row (first argument) and column (second argument). Note how the \mathbf{A} -matrices had to be specially treated since they contained information about the centre point of the ellipsoid. An example of a coronal slice with the right kidney marked by the cross section of its corresponding ellipsoid is shown in Figure 11.




Figure 11: Ellipsoid marking the right kidney in a coronal slice. The ellipsoid was later used as a first guess for an optimisation scheme where successively higher order surfaces were adapted to the organ.

Computing the image gradient

The image gradient was computed using three 3x3x3 kernels, shown in Figure 12. Depending on how this kernel was oriented with respect to the image coordinate system, it gave the partial derivatives in the x-, y- and z-directions. The image volume was convolved with these three differently oriented kernels and thus three volumes, representing the derivatives in the x-direction, y-direction and z-direction were obtained.




Figure 12: The kernel used for computing the grey level gradient. Depending on how the faces with 1 and -1 are oriented, the image gradient in the x-, y or z-direction can be calculated

In order to reduce the amount of data needed to be handled in the coming optimisation, the image gradient was only computed in a cut out of the original image. The cuboid cut out was based on the extreme values of the initialising ellipsoid in the x-, y, and z-direction plus a margin. The global position of the a surface could later be placed correctly in the image by relating the coefficients $a_{x,0,0}$, $a_{y,0,0}$ and $a_{z,0,0}$ to the position of the cut out.

Optimisation based on a single image volume

Optimisation was performed on the elements in the vector \mathbf{p} with respect to an objective function, E , further described below, using the downhill simplex method.

For segmentation of an image based on a single image volume, *i.e.* a SPECT-volume without a corresponding CT-volume, optimisation was performed in a successive manner, where the number of Fourier orders was gradually increased for each step in the optimisation. First a surface with $K=2$ was adapted using the ellipsoid from the initialisation process as input. Then all Fourier coefficient matrices were expanded with one row and one column representing $m=3$ and $l=3$ filled with zeros. These matrices were then used as input for the next surface with $K=3$. The process was repeated until a predefined order had been reached. A flowchart showing the optimisation procedure is shown in Figure 13.

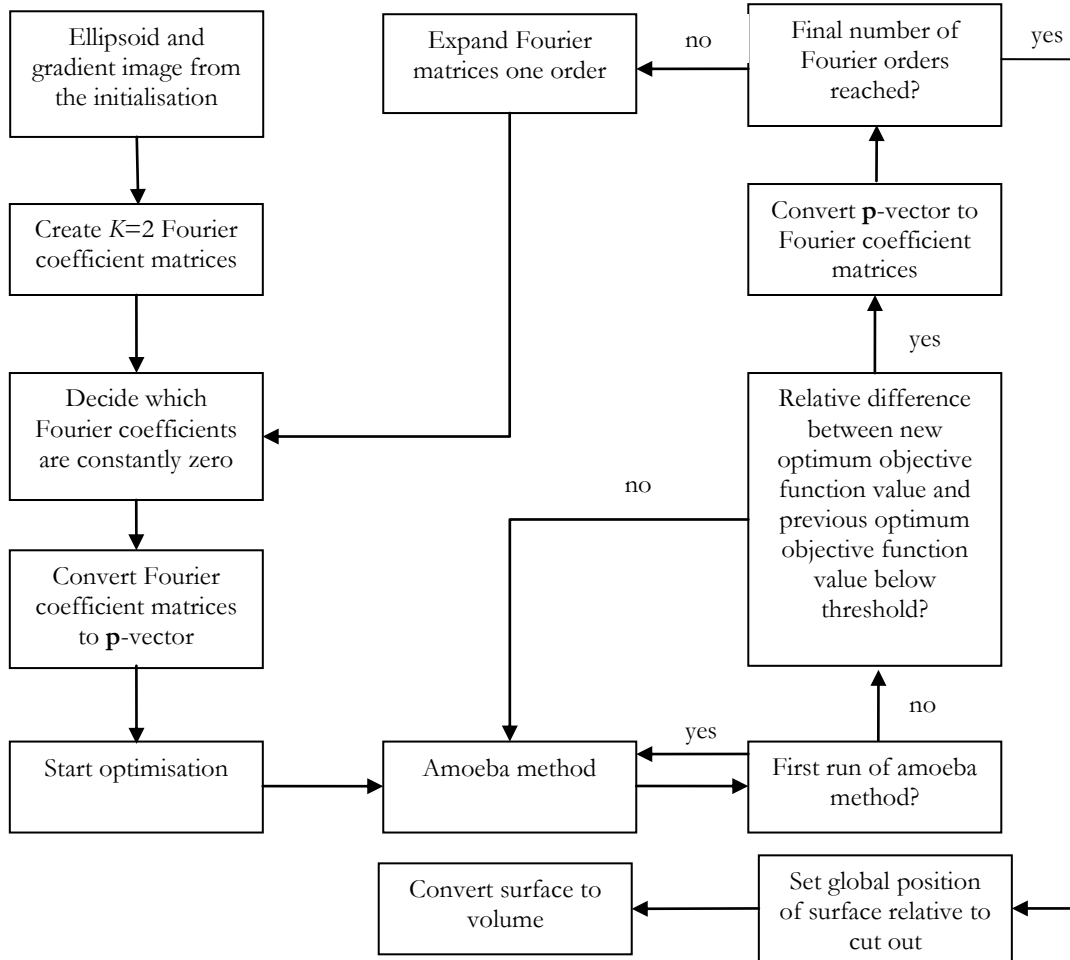


Figure 13: Flowchart showing the general idea of the optimisation procedure based on a single image volume

The objective function

As objective function for the optimisation of a surface to an image object Staib and Duncan [19] proposed a function for maximisation of the form

$$M = \iint_S b[x(u,v), y(u,v), z(u,v)] dA, \quad (50)$$

where b is a boundary measure integrated over the surface S . In the original version of Staib and Duncan this was derived using a probabilistic approach and also involving the use of *a priori* information about the object of interest. As a boundary measure the authors proposed the image grey level gradient. The practical evaluation of the surface integral was done using numerical integration.

For the purpose of marking kidneys in abdominal CT-images Gao *et al.* [21] evaluated the performance of three different objective functions, one using the form of Staib and Duncan using the scalar product between the grey level gradient and the surface normal, one using the same boundary measure but also involving a normalisation of the objective function to the total area of the surface and one objective function using the grey level gradient magnitude as boundary measure, still normalising the function to the total surface area. For their purpose, the authors found the use of the objective function involving the product of the grey level gradient and normalisation to the total surface area to perform the best. Staib and Duncan also discussed the use of normalisation against the total surface area as a way to avoid a tubular surface to grow in the ends despite a weak support. The authors suggested a normalisation against the surface area raised to one or raised to $\frac{1}{2}$, depending on how strong punishment was needed in order to prevent the surface growing in an uncontrolled fashion.

For the propose of nuclear medicine image segmentation Floreby *et al.* [20] used an objective function, E , for minimisation similar to Gao *et al.*:

$$E = \frac{\sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} [\nabla I \cdot \mathbf{n}(u_i, v_j)] \delta(u_i, v_j, \nabla I)}{\sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} |\mathbf{n}(u_i, v_j)|}, \quad (51)$$

where

$$\delta = \begin{cases} 1, & \nabla I \cdot \mathbf{n} < 0 \\ 0, & \nabla I \cdot \mathbf{n} \geq 0 \end{cases}. \quad (52)$$

In equation (51) the surface has been discretised into N_1 and N_2 steps for u and v respectively. This discretisation was performed in even steps with respect to u and v .

The relationship between equation (51) and equation (50) is realised considering the expression for the infinitesimal surface element

$$dA = \left| \frac{\partial \mathbf{R}}{\partial v} \times \frac{\partial \mathbf{R}}{\partial u} \right| du dv \quad (53)$$

and that the normal vector of the surface is given by

$$\mathbf{n}(u, v) = \frac{\partial \mathbf{R}}{\partial v} \times \frac{\partial \mathbf{R}}{\partial u}. \quad (54)$$

The magnitude of the surface element is thus contained in the scalar product of the surface normal and the image gradient. The approximation of du and dv with finite elements results in multiplication of each term in the sums with a constant factor if the surface is homogenously sampled with respect to u and v . This factor is cancelled in the quotient. Even without the denominator, this factor could be omitted since multiplication with a positive constant does not affect the optimisation.

Whether the normal of equation (54) points outwards from the surface or inwards depends on the relative directions of u and v parameters spanning the surface. In practice it also depends on the orientation of the coordinate system. This only affects the sign of the vector product and thus the sign for when to assign δ a value of 0 or 1.

In this work, an objective function with a variable exponent in the normalisation was introduced, giving the objective function, E_{SPECT} , according to

$$E_{\text{SPECT}} = \frac{\sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} [\nabla I \cdot \mathbf{n}(u_i, v_j)] \delta(u_i, v_j, \nabla I)}{\left(\sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} |\mathbf{n}(u_i, v_j)| \right)^a}, \quad (55)$$

where a became a parameter to be experimentally determined. The parameter a will be referred to as the normalisation exponent.

The partial derivatives of equation (55) needed for the evaluation of equation (51) were calculated analytically as [13]:

$$\frac{\partial x}{\partial u} = 4 \cdot \sum_{m=1}^{K_1-1} \sum_{l=1}^{K_2-1} m \cdot (-c_{x,m,l} \sin mu \sin lv + d_{x,m,l} \cos mu \sin lv) \quad (56)$$

$$\frac{\partial x}{\partial v} = -2a_{x,0,1} \sin v + \sum_{l=1}^{K_2-1} c_{x,0,1} \cos lv + 4 \cdot \sum_{m=1}^{K_1-1} \sum_{l=1}^{K_2-1} l \cdot (c_{x,m,l} \cos mu \cos lv + d_{x,m,l} \sin mu \cos lv) \quad (57)$$

$$\frac{\partial y}{\partial u} = 4 \cdot \sum_{m=1}^{K_1-1} \sum_{l=1}^{K_2-1} m \cdot (-c_{y,m,l} \sin mu \sin lv + d_{y,m,l} \cos mu \sin lv) \quad (58)$$

$$\frac{\partial y}{\partial v} = -2a_{y,0,1} \sin v + \sum_{l=1}^{K_2-1} c_{y,0,1} \cos lv + 4 \cdot \sum_{m=1}^{K_1-1} \sum_{l=1}^{K_2-1} l \cdot (c_{y,m,l} \cos mu \cos lv + d_{y,m,l} \sin mu \cos lv) \quad (59)$$

$$\frac{\partial z}{\partial u} = 4 \cdot \sum_{m=1}^{K_1-1} \sum_{l=1}^{K_2-1} m \cdot (-c_{z,m,l} \sin mu \sin lv + d_{z,m,l} \cos mu \sin lv) \quad (60)$$

$$\frac{\partial z}{\partial v} = -2a_{z,0,1} \sin v + \sum_{l=1}^{K_2-1} c_{z,0,1} \cos lv + 4 \cdot \sum_{m=1}^{K_1-1} \sum_{l=1}^{K_2-1} l \cdot (c_{z,m,l} \cos mu \cos lv + d_{z,m,l} \sin mu \cos lv). \quad (61)$$

Optimisation of the Fourier coefficients using the amoeba procedure

The amoeba procedure is a program in IDL for optimisation using the downhill simplex method [35] previously described in this work. In this procedure a first guess of the minimum of a function is given, after which the other vertices of the simplex is constructed by adding base vectors to this initial guess multiplied by a number giving the characteristic length of the problem. This scale parameter may be different for different dimensions, but for this work only one single value was used for all dimensions. One round of optimisation was finished once a tolerance, quantified as the relative difference between the function values was reached. In order to avoid the downhill simplex method to get stuck in poor local minima, the optimisation procedure was restarted several times for each step, *i.e.* one value of K , until the relative difference between two consecutive rounds of optimisation fell below a defined tolerance level, using the result from the former round as initial guess for the next one, as suggested by Floreby *et al.* [20]. Moreover, the spanning of the initial simplex alternated between adding base vectors in the positive and negative direction in order to further prevent the optimisation to get stuck without having found the true optimum solution.

Combination of SPECT- and CT-images

In order to combine the information from two coregistered images in the forms of one SPECT-volume and one CT-volume, a modified version of the objective function in equation (55) was used. The background to this modified objective function is given below. The derivation is analogue to the one given by Staib and Duncan [14, 19] on their method of how to incorporate *a priori* information into the segmentation scheme.

Assume there are two boundary measures of a given object, named b_{ect} and b_{ct} . These boundary measures are a realisation of a stochastic process ruled by a certain probability distribution based on a true boundary measure. Call the events of getting these realisations B_{ect} and B_{ct} . Assume the true boundary measure for a particular object to be an event T , then according to Bayes' theorem

$$P(T|B_{\text{ect}} \cap B_{\text{ct}}) = \frac{P(B_{\text{ect}} \cap B_{\text{ct}}|T) \cdot P(T)}{P(B_{\text{ect}} \cap B_{\text{ct}})}. \quad (62)$$

In equation (62) $P(B_{\text{ect}} \cap B_{\text{ct}})$ is the probability for getting the image data, which will be equal for all surfaces adapted to the organ and the factor can thus be ignored in terms of the optimisation process. Further all organ shapes are considered equally probable, leading $P(T)$ to be constant. Also, B_{ect} and B_{ct} are assumed to be independent. Then

$$\mathbb{P}(\mathbb{T}|\mathbb{B}_{\text{ect}} \cap \mathbb{B}_{\text{ct}}) \propto \mathbb{P}(\mathbb{B}_{\text{ect}}|\mathbb{T}) \cdot \mathbb{P}(\mathbb{B}_{\text{ct}}|\mathbb{T}). \quad (63)$$

Assume b_{ect} and b_{ct} to be noise corrupted versions of the true boundary measures t_{ect} and t_{ct} , *i.e.*

$$b = t + n, \quad (64)$$

where n is a term describing the noise. Considering the expression $\mathbb{P}(\mathbb{B}|\mathbb{T})$ is then equivalent to considering the expression $\mathbb{P}(n = b - t)$. Assume the noise in different image voxels to be independent and normally distributed with standard deviations σ_{ect} and σ_{ct} respectively. In that case equation (63) can be written as

$$\mathbb{P}(\mathbb{T}|\mathbb{B}_{\text{ect}} \cap \mathbb{B}_{\text{ct}}) \propto \prod_{\mathbb{V}} \frac{1}{\sqrt{2\pi}\sigma_{\text{ect}}} \exp\left(-\frac{(b_{\text{ect}} - t_{\text{ect}})^2}{2\sigma_{\text{ect}}^2}\right) \cdot \prod_{\mathbb{V}} \frac{1}{\sqrt{2\pi}\sigma_{\text{ct}}} \exp\left(-\frac{(b_{\text{ct}} - t_{\text{ct}})^2}{2\sigma_{\text{ct}}^2}\right), \quad (65)$$

where \mathbb{V} is the imaged volume.

The logarithm of this expression yields

$$\sum_{\mathbb{V}} \ln \frac{1}{\sqrt{2\pi}\sigma_{\text{ect}}} - \sum_{\mathbb{V}} \frac{(b_{\text{ect}} - t_{\text{ect}})^2}{2\sigma_{\text{ect}}^2} + \sum_{\mathbb{V}} \ln \frac{1}{\sqrt{2\pi}\sigma_{\text{ct}}} - \sum_{\mathbb{V}} \frac{(b_{\text{ct}} - t_{\text{ct}})^2}{2\sigma_{\text{ct}}^2}. \quad (66)$$

If constant terms are removed in the algebraic expression (66), then an expression

$$\begin{aligned} & - \sum_{\mathbb{V}} \frac{(b_{\text{ect}} - t_{\text{ect}})^2}{2\sigma_{\text{ect}}^2} - \sum_{\mathbb{V}} \frac{(b_{\text{ct}} - t_{\text{ct}})^2}{2\sigma_{\text{ct}}^2} = \\ & = - \frac{1}{2\sigma_{\text{ect}}^2} \left\{ \sum_{\mathbb{V}} b_{\text{ect}}^2 + \sum_{\mathbb{V}} [-2b_{\text{ect}}t_{\text{ect}} + t_{\text{ect}}^2] \right\} - \frac{1}{2\sigma_{\text{ct}}^2} \left\{ \sum_{\mathbb{V}} b_{\text{ct}}^2 + \sum_{\mathbb{V}} [-2b_{\text{ct}}t_{\text{ct}} + t_{\text{ct}}^2] \right\} \end{aligned} \quad (67)$$

is obtained.

By further removing terms not affected by t , then

$$- \frac{1}{2\sigma_{\text{ect}}^2} \sum_{\mathbb{V}} [-2b_{\text{ect}}t_{\text{ect}} + t_{\text{ect}}^2] - \frac{1}{2\sigma_{\text{ct}}^2} \sum_{\mathbb{V}} [-2b_{\text{ct}}t_{\text{ct}} + t_{\text{ct}}^2]. \quad (68)$$

Assume the templates only to have support along the surface, \mathbb{S} . This change the set to be summed over from a volume to a surface and expression (68) is thereby transformed to

$$- \frac{1}{2\sigma_{\text{ect}}^2} \sum_{\mathbb{S}} [-2b_{\text{ect}}t_{\text{ect}} + t_{\text{ect}}^2] - \frac{1}{2\sigma_{\text{ct}}^2} \sum_{\mathbb{S}} [-2b_{\text{ct}}t_{\text{ct}} + t_{\text{ct}}^2]. \quad (69)$$

Expression (69) can be further simplified by assuming the boundary support to be constant for the SPECT-image and the CT-image with values k_{ect} and k_{ct} respectively:

$$- \frac{1}{2\sigma_{\text{ect}}^2} \sum_{\mathbb{S}} [-2k_{\text{ect}}t_{\text{ect}} + k_{\text{ect}}^2] - \frac{1}{2\sigma_{\text{ct}}^2} \sum_{\mathbb{S}} [-2k_{\text{ct}}t_{\text{ct}} + k_{\text{ct}}^2]. \quad (70)$$

If the total surface area is assumed constant in the optimisation process, then the constant terms in expression (70) can be removed. An expression of the form

$$\frac{k_{\text{ect}}}{\sigma_{\text{ect}}^2} \sum_S t_{\text{ect}} + \frac{k_{\text{ct}}}{\sigma_{\text{ct}}^2} \sum_S t_{\text{ct}} \quad (71)$$

is thus obtained.

Just like for the segmentation of a single SPECT-image with an objective function based on equation (55), it was deemed advantageous to normalise against the total surface area raised to an exponent and to choose the boundary measure as the scalar product between the image gradient and the surface normal. The final objective function for combined SPECT/CT-images becomes

$$E_{\text{SPECT/CT}} = \frac{\frac{k_{\text{ect}}}{\sigma_{\text{ect}}^2} \sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} \nabla I_{\text{ect}} \cdot \mathbf{n}(u_i, v_j) \delta_1(u_i, v_j, \nabla I_{\text{ect}}) + \frac{k_{\text{ct}}}{\sigma_{\text{ct}}^2} \sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} \nabla I_{\text{ct}} \cdot \mathbf{n}(u_i, v_j) \delta_2(u_i, v_j, \nabla I_{\text{ct}})}{\left(\sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} |\mathbf{n}(u_i, v_j)| \right)^\alpha}. \quad (72)$$

Estimating k and σ^2

The value of k in equation (72) was estimated as a weighted average of the response for a surface in the image to be segmented, *i.e.* applying equation (55) with $a = 1$. Simultaneously, the variance of the boundary measure, σ^2 , could be estimated using a weighted variance, s^2 , as given by the NIST DATAPLOT Reference Manual [36], of the boundary measure along the surface according to

$$s^2 = \frac{N' \sum_{i=1}^N w_i (E_i - k)^2}{(N' - 1) \sum_{i=1}^N w_i}, \quad (73)$$

where N' is the number of elements having nonzero weights, E_i is the scalar product between the grey level gradient and the surface unit normal in the i th point along the surface, k is the mean boundary support calculated as described above and w_i is the weight given to the i th sample. For the purpose of estimating the variance along the surface, the weight in the i th point is proportional to the area represented by the i th point of the surface.

The practical estimation of the boundary response and the variance of the boundary measure were performed using a surface with a low Fourier order. Optimisation was first performed on the pure SPECT-image and then on the pure CT-image using the SPECT-image results as a first approximation. The segmentation algorithm was restarted when the image basis was changed, so in a first instance only the lowest order Fourier coefficients of the SPECT-result was used as an initial estimate for the CT-result. This was done in order for the optimisation procedure not to produce artefacts at the surface. The parameters k and s^2 were calculated along these surfaces for the SPECT- and CT-cases respectively. After that, the optimisation procedure was restarted using the weighted objective function with the CT-result as first approximation and run until a predefined order of the surface was obtained just like when the segmentation was based on a single image volume, *i.e.* based on the SPECT-image only.

Surface to volume conversion

The classification of voxels contained in the segmented closed surface was done in a way similar to the one indicated by Floreby *et al.* [20]. A discrete approximation of the segmented surface was done by evaluating a number of points on the surface evenly distributed in u and v . A cuboid containing the volume of interest was then created based on the maximum and minimum values of the x -, y - and z -coordinate values plus a margin.

For all voxel coordinates in the cuboid, the closest point at the surface to the voxel was found using golden section search with respect to u and v . The solution for the u -parameter was unrestricted in this procedure, while the v -parameter was restricted to stay in the interval $[0, \pi]$. The restriction on v was necessary because the coordinate functions were periodic with period 2π but only half that interval was used for describing the organ of interest. It is true that the coordinate functions were periodic also with

respect to u , but here the entire interval is used for describing a closed surface, *i.e.* any value of u would still describe a point on the surface and hard restrictions for the parameter were thus not necessary.

In the next step, the directed line segment between each voxel coordinate and the corresponding closest points on the surface was formed. The scalar product was then calculated between the vector associated with the directed line segment and the surface unit normal in the closest point of the surface. A negative value of this product led to a classification of the voxel as being outside the object, while a positive value, led to a classification of the voxel as inside the object.

The procedure described above could not be directly applied to the poles of the object, *i.e.* for v equal to zero or π . This was because the surface normal is zero by definition in those two points, making every attempt to use the surface normal impossible, and even in a neighbourhood of the poles, the calculations were found to be unreliable due to small values. In order to avoid that situation, all v -values associated with normal lengths below a threshold were moved in the v -direction until the surface normal was larger than the threshold value. This movement was done in the positive v -direction if v was close to 0 and in the negative v -direction if v was close to π , this in order to assure the v -parameter to lie in the interval $[0, \pi]$.

4.3 Evaluation

Methods for evaluation of the segmentation result

The segmentation method was evaluated with respect to two different aspects, the mean distance between the found surface and the true object surface, and the correspondence of the voxels contained by the surface and the voxels building up the original object.

Surface based evaluation

As proposed by Staib and Duncan [19] and Floreby and Floreby *et al.* [13, 20, 22] an error measure, ε , was used of the form

$$\varepsilon = \frac{\sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} \left| \mathbf{R}_p(u', v') - \mathbf{R}_{\hat{p}}(u_i, v_j) \right| \cdot \left| \mathbf{n}(\hat{p}, u_i, v_j) \right|}{\sum_{i=0}^{N_1-1} \sum_{j=0}^{N_2-1} \left| \mathbf{n}(\hat{p}, u_i, v_j) \right|}, \quad (74)$$

where $\mathbf{R}_{\hat{p}}(u_i, v_j)$ is the coordinates of point (u_i, v_j) at the surface obtained from the segmentation of the image and $\mathbf{R}_p(u', v')$ is the closest point to $\mathbf{R}_{\hat{p}}(u_i, v_j)$ at the true surface of the object.

In practice the true surface of the object was replaced by a reference surface that was obtained by optimisation under idealised conditions, as further described below. The closest distance between the estimated surface and the reference surface was determined using golden section search in a similar way as was done for the conversion from a closed surface to volume. The error presented in equation (74) was given the name ‘mean surface distance’.

As an alternative measure the median surface distance was used, *i.e.* the distance that half the obtained surface fell below. In connection with the median, the tenth and 90:th percentiles were also calculated, thereby giving an indication of the spread of the error. The measures were given the names ‘median surface distance’, ‘tenth percentile surface distance’ and ‘90:th percentile surface distance’ respectively.




Figure 14: Comparison between an obtained surface (wire mesh) and reference surface (solid), which were compared in terms of different distance measures, *i.e.* the mean, median, tenth percentile and 90:th percentile surface distances.

Volume based evaluation

Evaluation of the volume enclosed by the segmented surface was evaluated by classifying voxels as ‘correct’ if they were included in both the original volume and the volume obtained in the segmentation process, ‘falsely included’ if they were included in the volume obtained in the segmentation process but not in the original volume, and ‘missed’ if they were included in the original volume but not in the volume obtained in the segmentation process. As a single measure of the error, the fraction with respect to the original volume of the sum of falsely included and missed voxels was used. This fraction was given the name ‘fraction of misclassified voxels’ and could mathematically be described as

$$\text{FMV} = \frac{|(S \cap T) \cup (T \cap S)|}{|T|}, \quad (75)$$

where FMV is the fraction of misclassified voxels, T is the set of voxels belonging to the original object, S is the set of voxels classified as belonging to the object from segmentation and $|\cdot|$ signifies the cardinality of a set.

In order to investigate the volume conservation of the segmentation method, a measure based on the fraction of the fraction of misclassified voxels due to missed voxels was defined. This measure was given the name ‘partial fraction of missed voxels’ and could mathematically be described as

$$\text{PFMV} = \frac{|T \cap S|}{|(S \cap T) \cup (T \cap S)|}, \quad (76)$$

provided that $|(S \cap T) \cup (T \cap S)| \neq 0$.

Ideally, the partial fraction of missed voxels is $\frac{1}{2}$, since in that case the number of missed voxels is equal to the number of falsely included voxels and thus the total volume of the object is preserved.




Figure 15: Example from the evaluation with respect to volume, with matching voxels shown in green, voxels included in the volume obtained in the segmentation process but not in the original volume shown in red and voxels included in the original volume but not in the volume obtained in the segmentation process shown in blue. From left to right, a transversal, a coronal and a sagittal view is presented respectively.

As a comparison for the volume error the kidney based binary test object was translated in the x-, y- and z-direction and the obtained mismatch quantified. A translation in the x-direction corresponded to a fraction of misclassified voxels of 0.22, a translation in the y-direction a fraction of 0.23 and a translation

in the z-direction a fraction of 0.17. A combined mismatch of one voxel side in all directions produced a fraction of misclassified voxels of 0.35.

Sensitivity analysis: Factors believed to affect the performance of the segmentation method

Factors that were hypothesised to affect the performance of the segmentation method were identified and listed as:

- **The contrast to noise ratio of the image**
The contrast between organ and background determines the signal in the gradient images. It is thus the contrast to noise ratio of the original image that is likely to be of importance in terms of image noise.
- **The spatial resolution of the image**
A poorer spatial resolution makes the organ edges less well defined. Since it smears out the edge it also affects the signal to noise ratio of the gradient images. A poor spatial resolution helps the surface to be attracted to the organ from large distances and can therefore also be an advantage.
- **The organ shape**
The organ shape affects the strategy that must be used for the segmentation in terms of the choice of parameters. The segmentation method was not believed to work well for complex shapes.
- **The accuracy of the initialisation**
An accurate initialisation makes the optimisation task simpler and is thus believed to affect the rate of successful segmentations as well as the reproducibility of acceptable segmentations.
- **The constraints given to the optimisation**
If the optimisation procedure is given too hard constraints in relation the precision of the initialisation, it risks missing the optimum solution. If the optimisation is given too loose constraint, the risk of getting stuck in poor local minima is increased.
- **The number of Fourier orders used**
As more degrees of freedom are added to the surface, the better the surface can form itself to the organ edge. The number of Fourier orders must be limited due to time considerations. The use of high Fourier orders also risks making the surface more susceptible to noise.
- **Nearby high signal objects in the image**
Nearby high signal objects in the image may affect the optimisation step in the segmentation as they risk attracting the surface.
- **The combination of SPECT- and CT-images**
Compared to SPECT-images, CT-images have much higher spatial resolution. By combining SPECT- and CT-images, the performance of the segmentation algorithm was believed to improve since the organ edge would be better defined.
- **The matrix voxel size**
The matrix voxel size determined the importance of the round off error due to the discrete map when the objective function is evaluated. When a crisp set of voxels is determined in the surface to volume conversion, the voxel size will determine the importance of effects due to the discretisation of the volume, *e.g.* the effect introduced at the boundary when a voxel is partly in the background and partly in the object.

All of the listed aspects above were investigated in the evaluation except for the matrix voxel size that was held constant at four millimetres for all tests. The effects of spatial resolution and noise were investigated in single organ images as well as the effects of organ shape, the constraints given to the optimisation and the number of Fourier orders used. The effect of nearby high signal objects could be observed in simulated SPECT-images, as could the result of combining SPECT- and CT-images. The dependence of the method to the initialisation process and the operator intervariability to the segmentation result were studied in the set of clinical SPECT-images described in section 4.1.

Evaluation in single organ images

A first test of the segmentation method was performed on the single organ images based on the XCAT phantom left kidney described in section 4.1.

Creating a reference surface

A reference surface was created to the test volume by applying the segmentation method under idealised conditions, *i.e.* the original binary object volume was used with a surface adapted to the object employing Fourier orders up to $K=8$. The segmentation was performed using images with increasingly finer spatial resolution. The $K=2$ surface was created using a volume blurred with a three dimensional Gaussian kernel with 1 unit of length FWHM (one unit of length corresponded to 4 mm). The spatial resolution of the volume was then improved to 0.5 units of length, which was used up to $K=4$. After $K=4$, the resolution was further improved to 0.1 units of length, and after $K=6$ it was increased to 0.05 units of length up to $K=8$. The reference surface is shown in Figure 16. The reason for using a scheme with increasingly higher spatial resolution was to guide the surface to the boundary object in a coarse-to-fine manner.




Figure 16: Reference surface with Fourier orders up to $K=8$

Tests performed

In a first step, all segmentations were done with normalisation exponent $a=1$, since this implicitly was the exponent used by Floreby *et al.* [20, 22] and Gao *et al.* [21]. The test image was segmented up to Fourier order $K=6$ with 6 different resolutions ranging from 0.5 voxel sizes to 5 voxel sizes corresponding to 2 mm to 20 mm spatial resolution. In order to get two realisations of the segmentations, the method was applied twice with the initialising ellipsoid translated 2 mm in the x -, y - and z -direction between the two. This translation was thought not to affect the segmentation performance as a whole but still make it possible for the optimisation to take another path through the space of Fourier coefficients when finding the optimum solution. The mean surface distance and the fraction of misclassified voxels were calculated as the mean value of the two segmentation realisation. The measures were evaluated as a function of the number of Fourier orders used.

In the next step independent Gaussian noise was added to the images. A resolution of 12 mm was used and the CNR of the images as defined above was four. The mean surface distance and the fraction of misclassified voxels were evaluated for six noise realisations. This was done in order to get an idea of how much the result could vary due to the noise realisation being different on otherwise identical images with identical initialisations.

In order to examine the behaviour of the segmentation process performance under varying conditions of spatial resolution and CNR, the mean surface distance and fraction of misclassified voxels were determined for spatial resolutions between 4 mm and 16 mm and CNR ranging from eight to 1/2. The two measures were evaluated as a function of the number of Fourier orders used up to $K=6$. Two noise realisations were used for each resolution and CNR. The mean of the two results were used for the analysis.

The influence of the scale parameter when spanning the simplexes in the optimisation process was tested by varying the scale parameter in the first step ($K=2$) in the segmentation process between one and six in steps of unity. The image used had a spatial resolution of 12 mm and CNR=4. Again two noise realisations were employed and the analysis was based on the mean result of the two segmentations. The experiment was also repeated for the second scale parameter ($K=3$) in the segmentation process, which was varied between 0.1 and three.

In the next step, the normalisation exponent was varied between 0.1 and one for images of spatial resolution between 4 mm and 16 mm and CNR from one to eight. Surfaces were used up to order $K=5$ and the fraction of misclassified voxels were studied as a function of the normalisation exponent. Also, the partial fraction of missed voxels was studied as a function of the normalisation exponent.

In order to test the effect of a more complex organ shape, single organ images based on the XCAT spleen were used. A normalisation exponent of 0.6 was used for the segmentation. It was recognised that the optimisation algorithm needed looser constraints in order to segment images with objects having more complex shapes, and thus the scale parameter was increased for later steps in the optimisation scheme for this particular investigation compared to the values used for the kidney images. A surface was adapted to the organ up to Fourier order $K=6$ in an image with 4 mm spatial resolution and one image with spatial resolution 12 mm and contrast to noise ratios ranging from two to eight. The results were evaluated with respect to the fraction of misclassified voxels only. An example of a surface adapted to the spleen object is shown in Figure 17.

The initialising ellipsoids for all segmentations of single organ images, except for the one based on the XCAT phantom spleen, were identical, thus excluding the initialisation as a part of the variation for different image conditions. The ellipsoid used was based on an initialisation, excluding the presegmentation step, of the original object used for creating the images, *i.e.* no effects of spatial resolution or noise were present. For the spleen based images a corresponding ellipsoid based on that object was used for initialisation.




Figure 17: Example of a surface marking the spleen object

Evaluation in activity distributions mimicking patients

After investigation of the method on single organ images, the method was tested on images of activity distributions mimicking patients described in section 4.1.

Creation of reference surface and reference volume

Two reference surfaces with $K=8$ and reference volumes for left and right kidney were developed in a similar way to what was done for the objects in single organ images. For the same reasons as why the pure XCAT phantom kidneys were not deemed a valid test object in the context of single organ images, they could not be used unprocessed for evaluation, this even more since arteries and veins created holes in this volume. In order to circumvent that problem an intermediate step was introduced by using the contained volume of the reference surface for comparison when calculating the volume based measures. A comparison between the reference volume for the left kidney and the original XCAT kidneys is shown in Figure 18 and Figure 19.



Figure 18: The original kidneys from the xcat phantom



Figure 19: The reference volume obtained after segmentation of Figure 18.

Tests performed

The images were segmented marking the left and right kidney using surfaces up to Fourier order $K=5$ for all images with the normalisation exponent varying between 0.1 and 0.9. The fraction of misclassified voxels and the partial fraction of missed voxels were calculated as a function of the normalisation exponent. The experiment was repeated also including CT-images in the segmentation basis for the image mimicking the activity distribution at 0.5 hours post injection. For normalisation exponents 0.9, the combined segmentation was tested also for later points in time and studied as a function of the number of Fourier orders used.

As opposed to the single organ images, the initialisation of the segmentation was based on the image to be segmented. Thus, the initialisation was a part of the result for these investigations.

Evaluation in Monte Carlo simulated SPECT-images

In order to further mimic a patient SPECT-study, including the characteristics of the gamma camera, the segmentation method was evaluated on the Monte Carlo simulated SPECT-images described in section 4.1.

Creation of reference surface and reference volume

A reference surface and a reference volume were created for the left and right kidney of the phantom used as the basis for the SIMIND simulation of SPECT-images. Since the phantom used had $2 \times 2 \times 2 \text{ mm}^3$ voxels, but the SPECT-volumes had $4 \times 4 \times 4 \text{ mm}^3$ voxel size, the organ had to be resampled before the construction of the reference surface. The resampling was performed with trilinear interpolation.

Tests performed

The images were segmented by marking left and right kidney using surfaces up to Fourier order $K=5$ and the mean surface distances and fractions of misclassified voxels studied as a function of the normalisation exponent for values 0.6, 0.7 and 0.9. This was done for all time point images using both SPECT-images solely and combined SPECT- and CT-images. Further, a combined segmentation was performed where the CT-images were given a factor four higher importance in the weighting of the SPECT- and CT-images than was the case for the original segmentation scheme.

For the first time point image the result was also studied as a function of the number of Fourier orders used for normalisation exponent 0.7. For later images the same study was done for normalisation exponent 0.9.

Evaluation in clinical images

Since it was recognised that the initialisation of the optimisation process introduced a subjective step in the segmentation process, the effect of operator dependency was tested for the set of clinical patient images described in section 4.1. The principles for the initialisation process described above was incorporated in a graphical user interface and used by four different operators on the left and right kidney. The users were instructed to make a rough delineation extent of the kidneys including the whole organ but excluding other bright object in the images. The different initialisations were used as input for segmentation of the SPECT-images using surfaces up to Fourier order $K=5$.

The operator segmentations were cleared from obviously erroneously segmented images with respect to the marking of left and right kidney, and for the cases where the definition of the organs from all operators passed this sorting, the intervariability was investigated. The intervariability was tested as the number of voxels included by one, two, three and four segmentations respectively as a fraction of the union all voxels included by any of the segmentations. The reason for excluding segmentations with erroneously marked volumes was that the question of failed segmentations was partly different from the question of operator dependence. The question of whether or not the segmentation was successful was a more profound question that needed to be answered before the reproducibility was investigated. It should be noted that the sorting of segmentations was not intended to distinguish clinically useful from clinically useless results. The clinical use of the method has not been systematically investigated in this work.

5. Results

5.1 Single organ images

The results from the investigation of the segmentation method based on single organ images are divided into the effects of resolution and number of Fourier orders, the uncertainty due to different noise realisations, the effect of image noise, the effect of the scale parameters in the optimisation scheme, the effect of the value of the normalisation exponent and the effect of organ shape.

Varying spatial resolution and number of Fourier orders

The mean surface distance and fraction of misclassified voxels as a function of the number of Fourier orders up to $K=6$ for noise free images of spatial resolution 2 mm, 4 mm, 8 mm, 12 mm, 16 mm and 20 mm is shown in Figure 20. In Figure 20 is also shown the median surface distance together with the tenth percentile surface distance and 90:th percentile surface distance for the test image with 12 mm spatial resolution, indicating the spread in distance between the obtained surface and the reference surface. All images were segmented using a normalisation exponent of unity. The error measures are generally seen to decrease as more Fourier orders are used and also get lower for better spatial resolutions compared to poorer spatial resolutions.

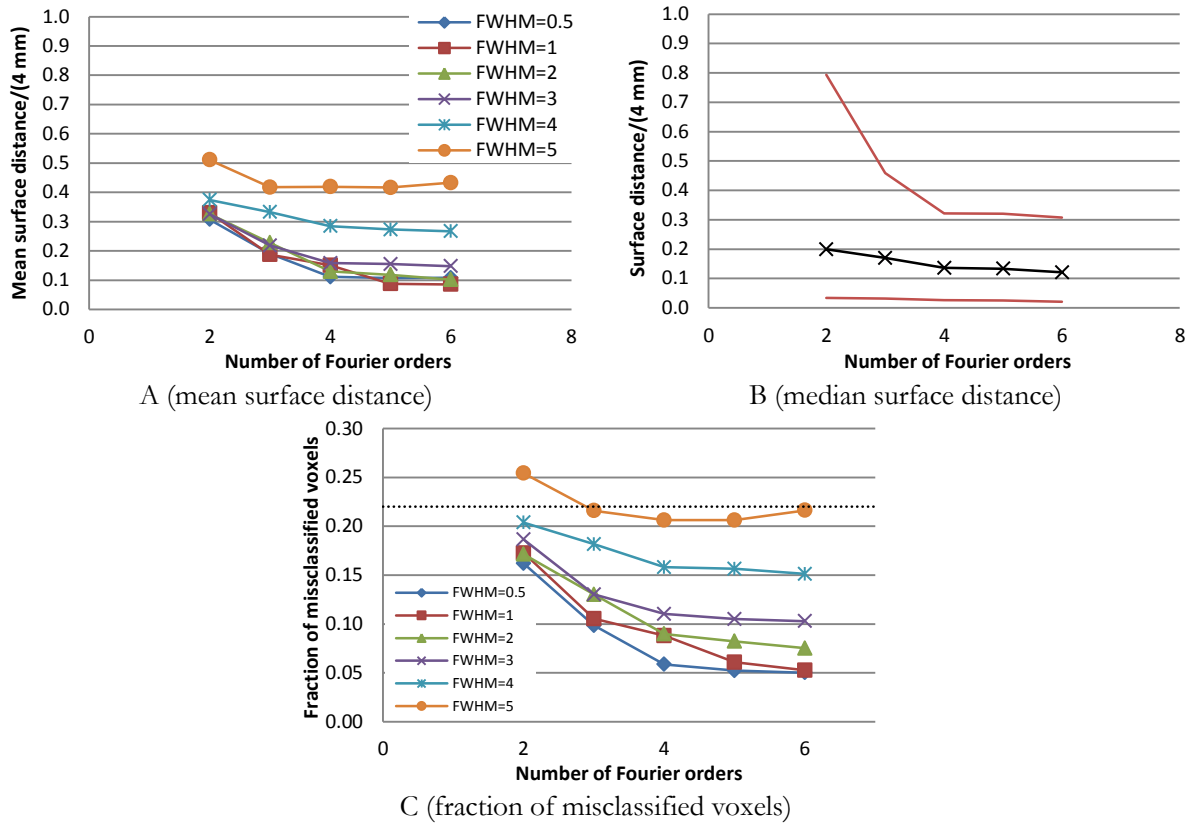


Figure 20: The behaviour of the segmentation method when applied to noise free single organ images. In A the behaviour with respect to the mean surface distance as a function of the number of Fourier orders for different spatial resolutions ranging from 2 mm to 20 mm is shown. In the diagram the resolution is specified as units of length, where one unit of length corresponds to 4 mm. In B the median surface distance (black line with crosses) is shown together with the tenth and 90:th percentile surface distances (red lines with no marks) for the case of FWHM=3. In C the corresponding results to A is shown in terms of fraction of misclassified voxels. The dotted line in this diagram indicates the effect of translating the original object one voxel in the x-direction. The mean surface distance as well as the fraction of misclassified voxels is seen to increase as the spatial resolution gets poorer. For the case of 12 mm spatial resolution, the distance between the obtained surface and the reference surface is, with respect to the 90:th percentile surface distance, seen to be less than one voxel side. With respect to the fraction of misclassified voxels the result obtained by the segmentation method is for most cases below the result obtained by deliberately translate the original object one voxel in the x-direction. All results were obtained as the average of two segmentations where the initialising ellipsoid has been slightly changed permitting different paths for the optimisation algorithm.

Different noise realisations

The mean surface distance and fraction of misclassified voxels as a function of the number of Fourier orders for six different noise realisations in single organ images segmented using surfaces up to $K=6$ for 12 mm spatial resolution and $CNR=4$ is shown in Figure 21. The segmentations were performed with the normalisation exponent set to unity. The error measures are observed to demonstrate a spreading due to the different noise realisations.

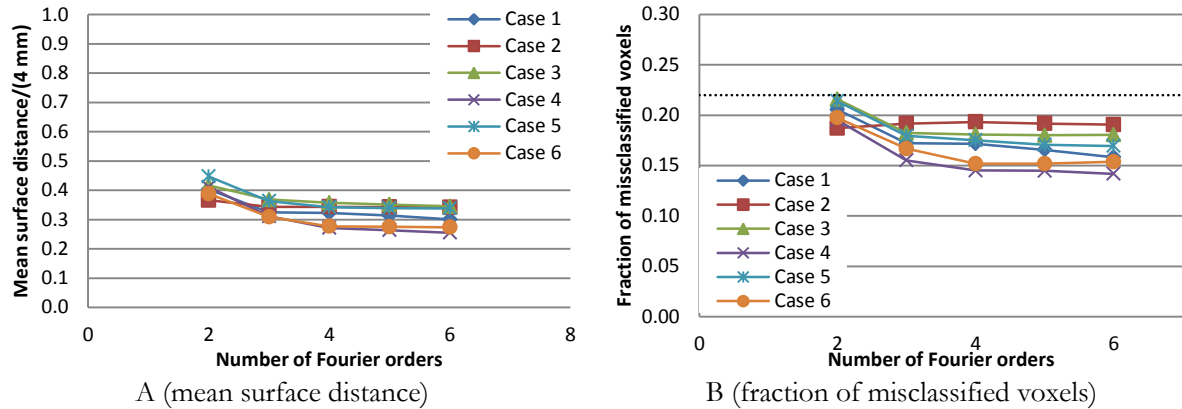


Figure 21: The effect of different noise realisations for the performance of the segmentation method on single organ images with 12 mm spatial resolution and a contrast to noise ratio of four. In A the result is presented with respect to the mean surface distance and in B the result is presented with respect to the fraction of misclassified voxels, where the dotted line indicates the effect of deliberately translating the original object one voxel in the x-direction. For the six cases investigated, the mean surface distance varied about 0.1 units of length and the fraction of misclassified voxels varied about 0.05.

Varying spatial resolution, contrast to noise ratio and number of Fourier orders

The mean surface distance and fraction of misclassified voxels for segmentations of images of varying spatial resolution between 4 mm and 16 mm and CNR between eight and $\frac{1}{2}$ as a function of the number of Fourier orders up to $K=6$ is shown in Figure 22. In Figure 22 is also shown the median surface distance together with the tenth percentile surface distance and 90:th percentile surface distance for the case of CNR=4 and spatial resolution 12 mm, indicating the spread in distance between the obtained surface and the reference surface. All segmentations were performed on two images with different noise realisations and the result was obtained as the average of these results. All segmentations were performed with the normalisation exponent set to unity. The error measures are seen to increase when the segmentation method is applied to images with poor spatial resolutions and low contrast to noise ratio compared to when applying it to images with good spatial resolution and high contrast to noise ratio.

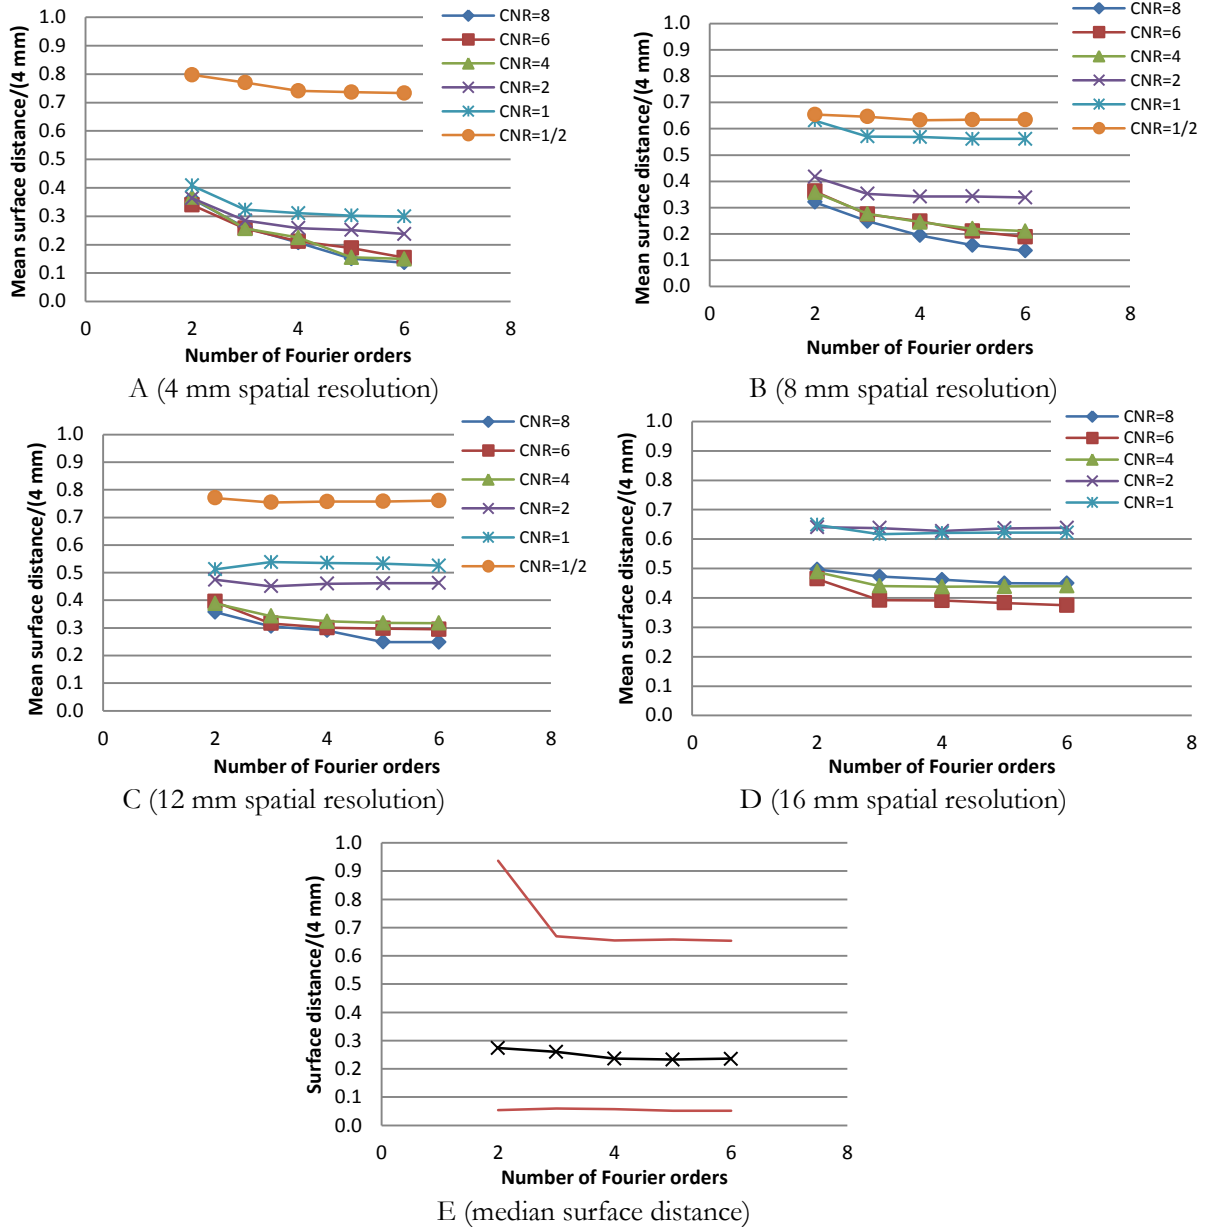


Figure 22: The performance of the segmentation method with respect to the mean surface distance for varying contrast to noise ratios. In E the median surface distance (black line with crosses) is shown together with the 10:th and 90:th percentile surface distances (red lines with no marks) for the case of 12 mm spatial resolution and a contrast to noise ratio of four, indicating the spread in distance for the surface. The performance of the segmentation method is seen to deteriorate as the contrast to noise ratio gets lower and the spatial resolution gets poorer. All results were obtained as the average when applying the method to two different noise realisations. For the 16 mm spatial resolution case (D), the CNR=1/2 case showed a mean surface distance above 4 mm for all number of Fourier orders, and is thus not shown in the diagram.

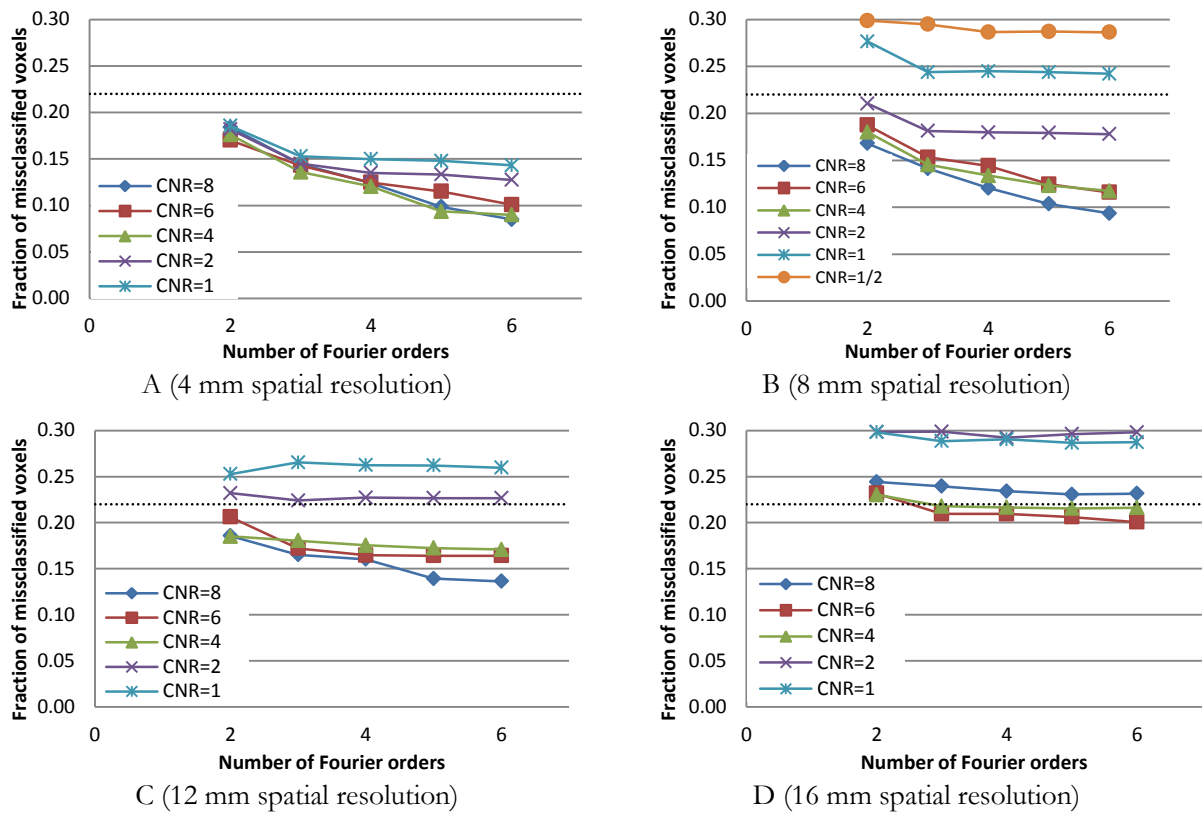


Figure 23: The performance of the segmentation method with respect to the fraction of misclassified voxels for varying contrast to noise ratios. Indicated as the dotted line, is also the effect of deliberately translating the original object one voxel in the x-direction. The performance of the segmentation method is seen to deteriorate as the contrast to noise ratio gets lower and the spatial resolution gets poorer in agreement to what was demonstrated in Figure 22. All results were obtained as the average when applying the method to two different noise realisations. For all cases except for the 8 mm spatial resolution case (B), the CNR=1/2 case showed a fraction of misclassified voxels above 0.30 for all number of Fourier orders, and is thus not shown in the diagrams.

Varying scale in the first two steps of the segmentation process

The mean surface distance and fraction of misclassified voxels as a function of the scale parameter in the first two steps in the optimisation process ($K=2$ and $K=3$) is shown in Figure 24. All segmentations were performed with the normalisation exponent set to unity. It is seen that too low a choice of the scale parameters deteriorates the performance of the segmentation method with respect to the error measures.

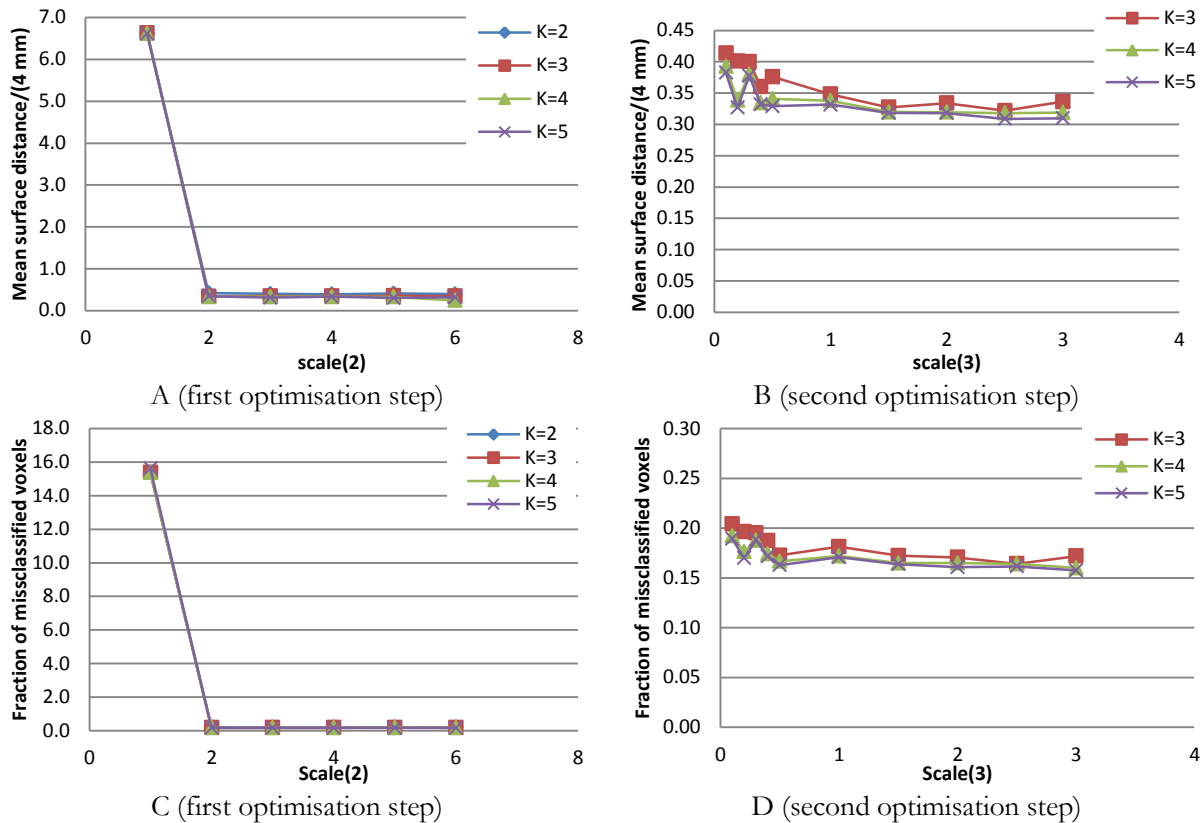


Figure 24: The performance of the segmentation method with respect to the mean surface distance (A and B) and the fraction of misclassified voxels (C and D) when changing the scale parameter in the first optimisation step (A and C) and the second optimisation step (B and D) for single organ images with 12 mm spatial resolution and CNR=4. Using too small scale parameters is shown to deteriorate the performance of the method. All results were obtained as the average when applying the method to two images with different noise realisations.

Varying normalisation exponent

The fraction of misclassified voxels as a function of the normalisation exponent for segmentations of images with CNR varying from one to eight is shown in Figure 25 for spatial resolutions of 4 mm, 8 mm, 12 mm and 16 mm using surfaces of Fourier order $K=5$ for the segmentations. The partial fraction of missed voxels is shown in Figure 26. The volume of the object of interest is preserved by the segmentation process if this latter measure is $1/2$. Also shown in Figure 27 are the median surface distance and the tenth and 90:th percentile surface distances for 12 mm spatial resolution, CNR=4 and normalisation exponent 0.7 as a function of the number of Fourier orders used. This can be compared to the corresponding diagram showing the result for the normalisation exponent set to unity in Figure 22.

The need of tuning the normalisation exponent in order not to bias the segmentation method towards too large or too small organ volumes is seen, as is the tendency for the fraction of misclassified voxels to have a minimum in the range where the total volume is approximately preserved.

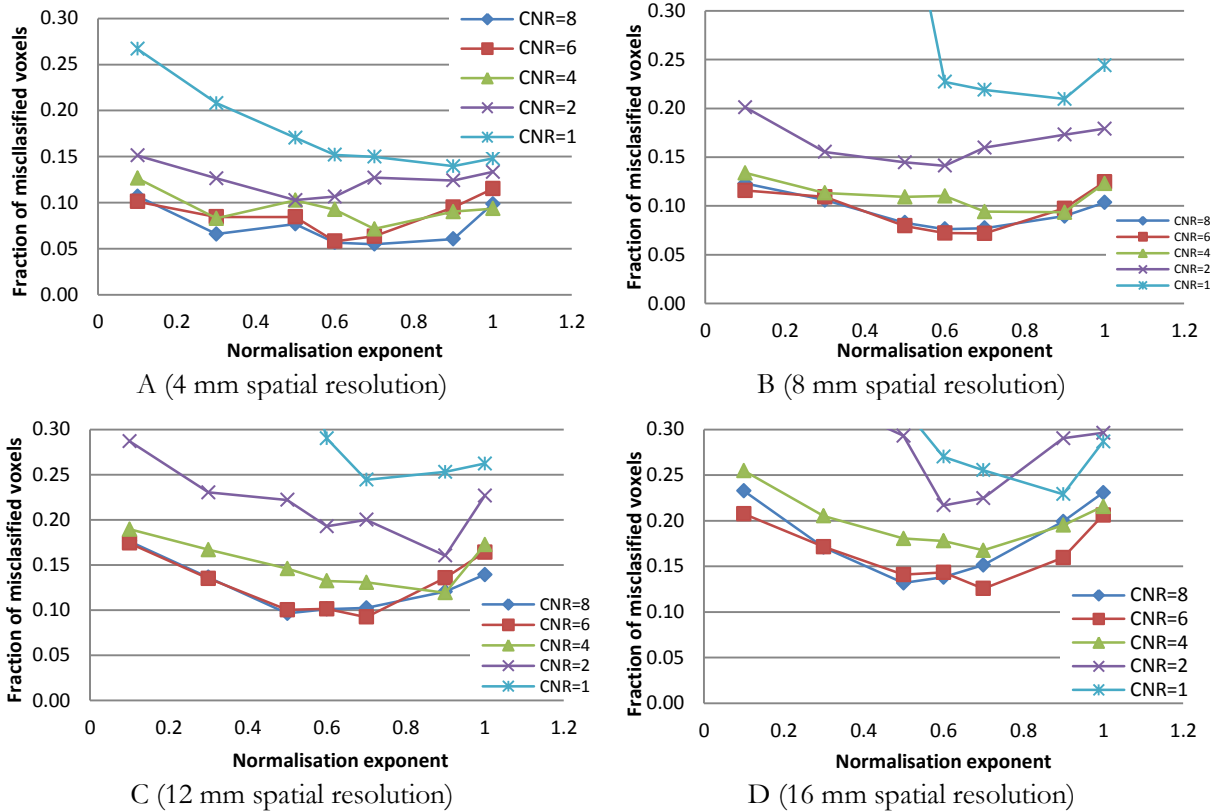


Figure 25: The performance of the segmentation method with respect to the fraction of misclassified voxels for different spatial resolutions and varying contrast to noise ratios when using up to five Fourier orders. The value of the normalisation exponent that causes the best performance of the method seems to lie in the region where the total volume of the object under consideration is preserved according to Figure 26. All results were obtained as the average when applying the method to two images with different noise realisations.

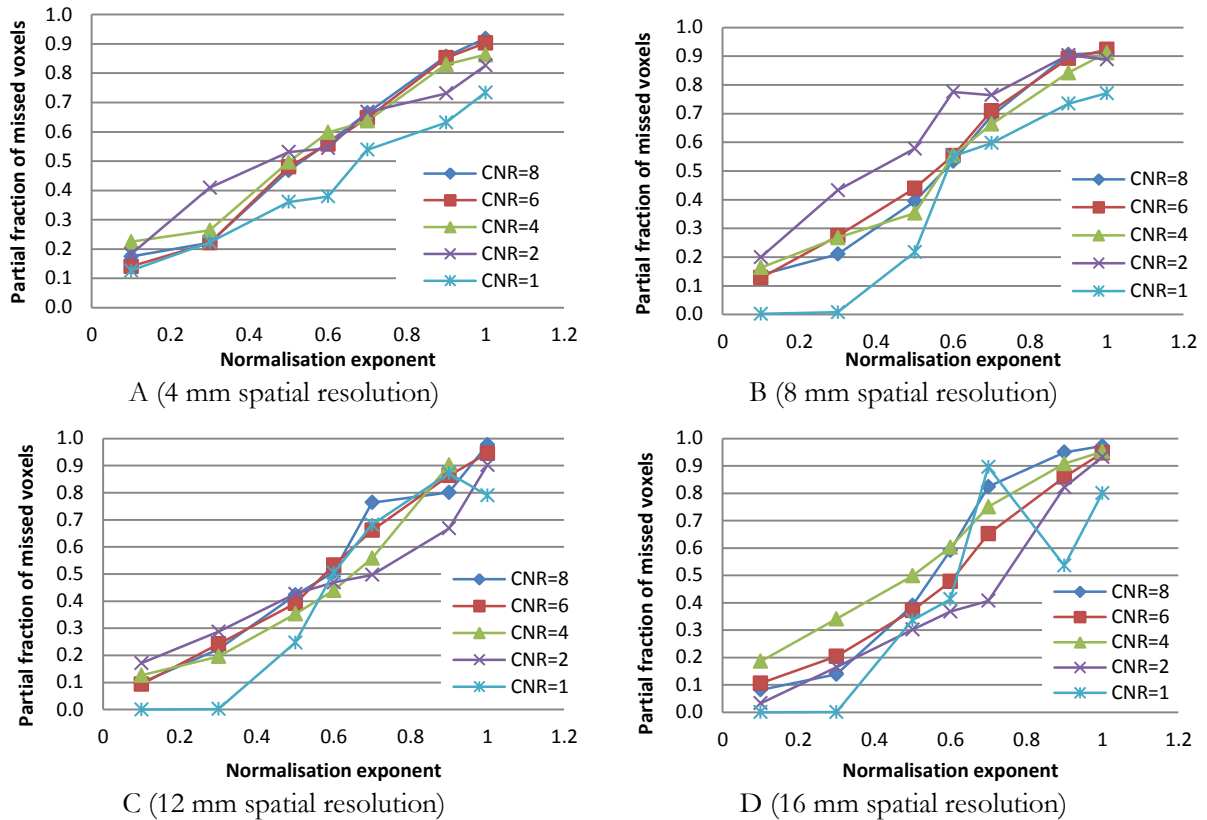


Figure 26: The partial fraction of missed voxels for different spatial resolutions and contrast to noise ratios as a function of the normalisation exponent. The segmentation method has preserved the volume of the object under consideration if the partial fraction of missed voxels equals $\frac{1}{2}$. For the investigated cases, this is the case for normalisation exponents of approximately 0.5 to 0.7. For lower values, the volume is overestimated, while for higher values the volume is underestimated.

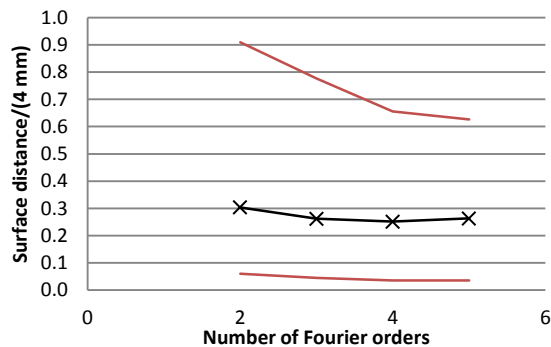


Figure 27: The median surface distance and the tenth and 90:th percentile surface distances as a function of the number of Fourier orders used for segmentations performed on single organ images with 12 mm spatial resolution and CNR=4 with the normalisation exponent set to 0.7. The median surface distance is shown as the black line with crosses, while the tenth and 90:th percentile surface distances are shown in red with no marks. The figure indicates the spread in the surface distance when not using a normalisation exponent of unity. All points were obtained as the mean result from two different noise realisations.

Organ shape

The fraction of misclassified voxels as a function of the number of Fourier orders used for the spleen in images with 4 mm and 12 mm spatial resolution and contrast to noise ration ranging from two to eight is shown in Figure 28. The spleen shape is seen to produce higher error measures than was the case for the previously used kidney-like object, especially for low number of Fourier orders.




Figure 28: The performance of the segmentation method when applied to a spleen-like object image with two different spatial resolutions and varying contrast to noise ratios. Compared to the kidney-like object presented earlier, the fraction of misclassified voxels is generally higher for the spleen-like object.

5.2 Evaluation in activity distributions mimicking patients

The fraction of misclassified voxels as a function of the normalisation exponent for segmentations adapting a surface up to Fourier order $K=5$ to the left and right kidneys in images with an imitated distribution of ^{177}Lu -DOTATATE for a patient undergoing radionuclide therapy is shown in Figure 29, depicting the situation for 0.5 h, 24 h, 69 h and 168 h post injection. The segmentation of the image showing the first point in time was performed for one segmentation scheme using only the simulated SPECT-images and one segmentation scheme also using the simulated CT-images. For the other points in time, generally only the pure SPECT-images were used. The graphs for the partial fraction of missed voxels are shown in Figure 30. Just as for single organ images, the value of the normalisation exponent is seen to affect the behaviour of the segmentation procedure, but the value for preservation of organ volume and minimum error measures tends to be higher.

The mean surface distance and median surface distance together with the tenth and 90:th percentile surface distances for the right kidney as a function of the number of Fourier orders used in different time point images is shown in Figure 31. The corresponding fractions of misclassified voxels to the mean surface distances are shown in Figure 32. The result for left kidney is not shown when studying the dependency on the number of Fourier orders used. This is because the segmentations frequently failed since also the spleen was included as a part of the left kidney. For the first point in time, the segmentations were performed using a normalisation exponent of 0.7, while for the later points in time a normalisation exponent of 0.9 was used when studying the result as a function of the number of Fourier orders. In Figure 31 and Figure 32 the result of including CT-images is shown for all time point images together with the result of using the pure SPECT-volumes. No major difference is observed between using pure SPECT-images and combined SPECT/CT-images.

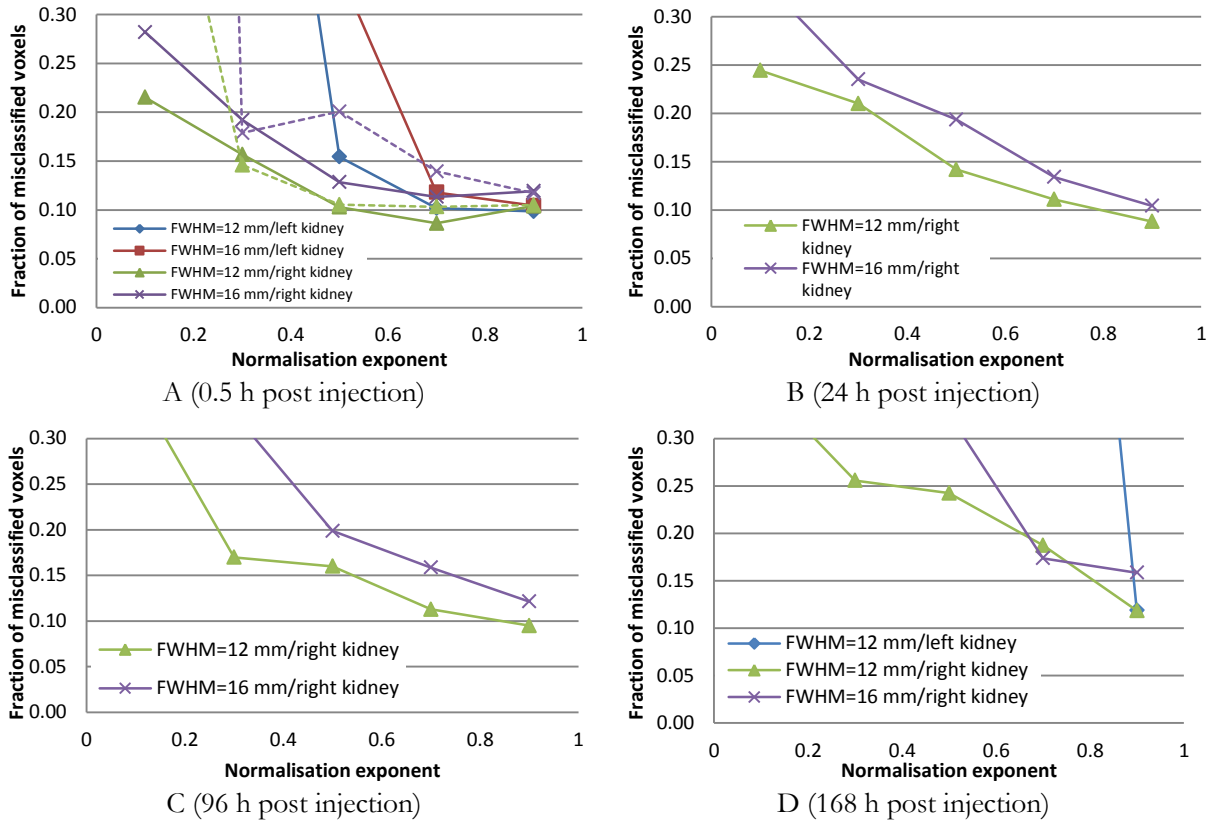


Figure 29: The fraction of misclassified voxels when applying the segmentation method to left and right kidney in images where the activity distribution for patients undergoing therapy with ^{177}Lu -DOTATATE using up to five Fourier orders as a function of the normalisation exponent. The result is displayed for 0.5 h post injection (A), 24 h post injection (B), 96 h post injection (C) and 168 h post injection (D). For the cases where the result for the left kidney is not displayed, the segmentation has failed due to inclusion of the spleen. In A the result when combining SPECT- and CT-images is shown as the broken lines. When compared to the corresponding plots for single organ images, the lowest fraction of misclassified voxels tends to be for a higher normalisation exponent when a clinical activity distribution is used.

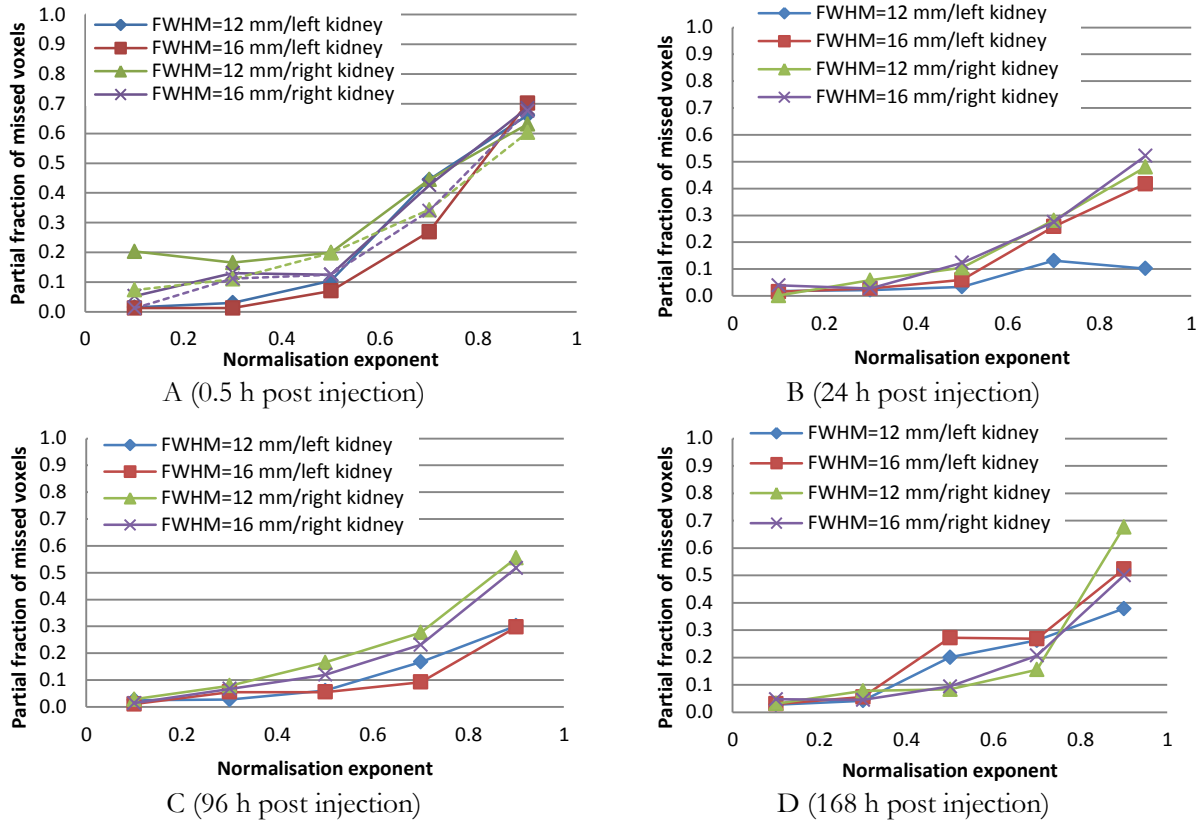


Figure 30: The partial fraction of missed voxels as a function of the normalisation exponent when applying the segmentation method to activity distributions imitating a patient undergoing therapy with ¹⁷⁷Lu- DOTATATE using up to five Fourier orders as a function of the normalisation exponent. The result is displayed for 0.5 h post injection (A), 24 h post injection (B), 96 h post injection (C) and 168 h post injection (D). In A the result when combining SPECT- and CT-images is shown as the broken lines. The segmentation method has preserved the total volume of the object under consideration if the partial fraction of misclassified voxels is 1/2. When compared to the corresponding plots for single organ images, this tends to happen for higher values of the normalisation exponent when using clinical activity distributions.

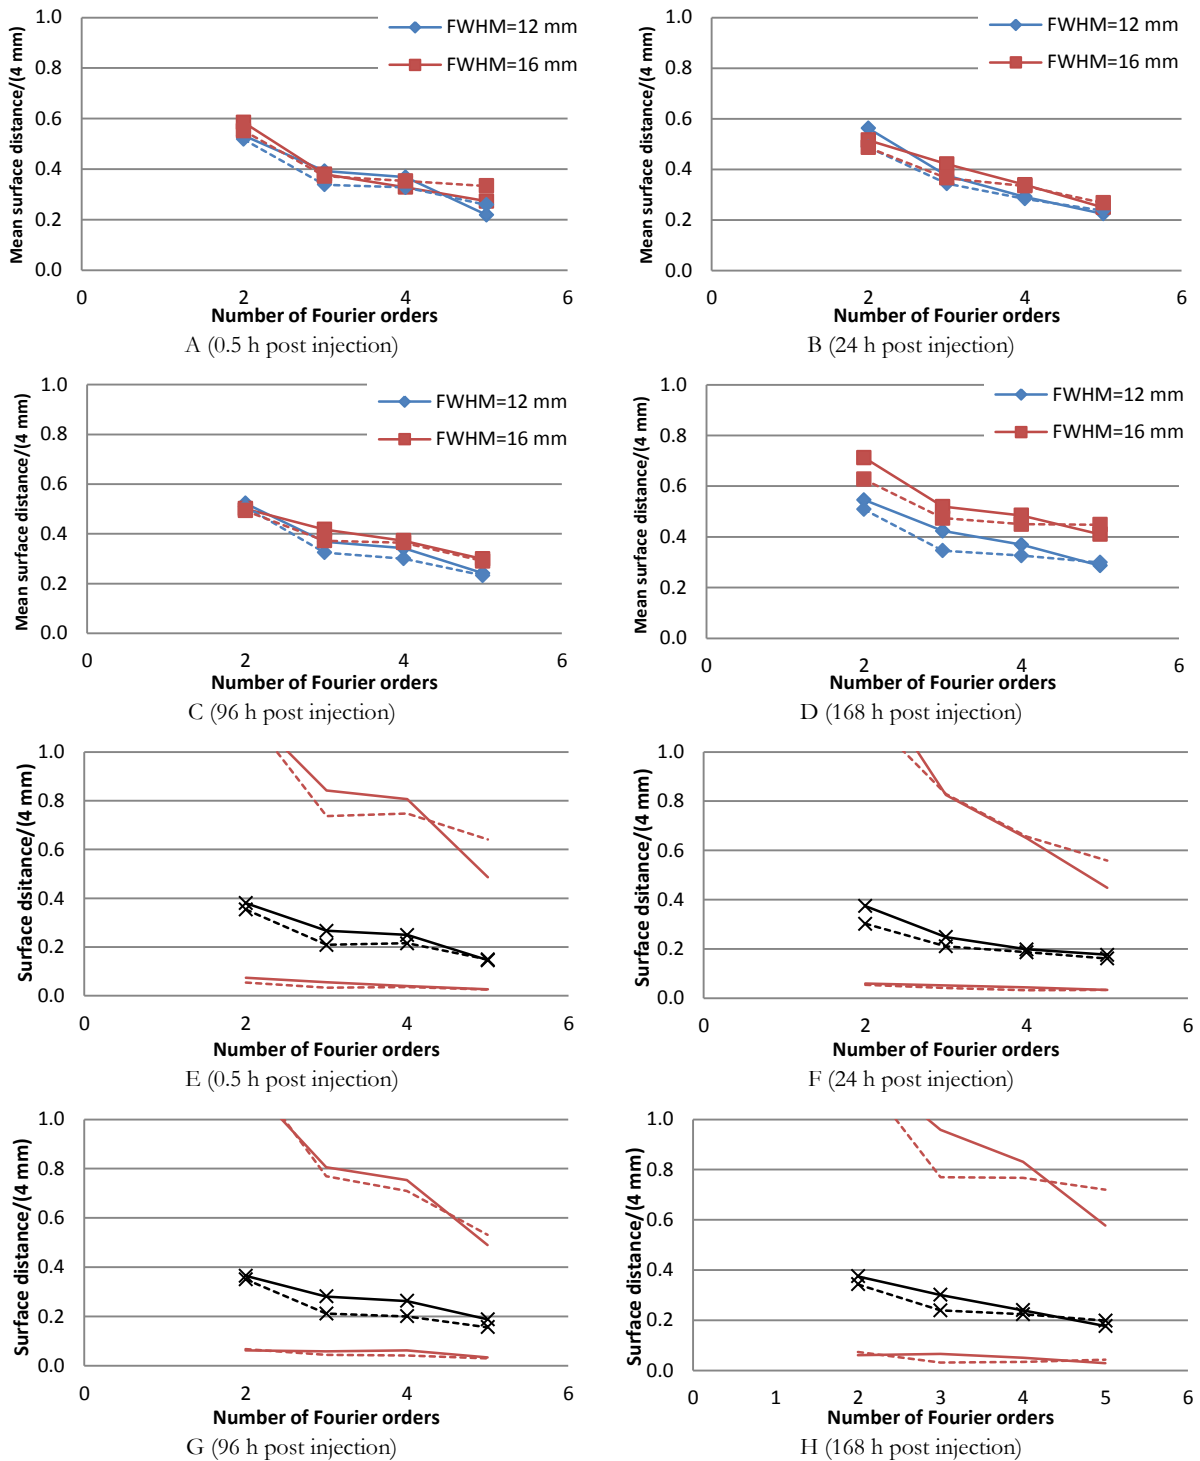


Figure 31: The mean surface distance as a function of the number of Fourier orders used when applying the segmentation method to images imitating the activity distribution in a patient undergoing therapy with ^{177}Lu -DOTATATE for right kidney. For the left kidney, the segmentation frequently failed due to the spleen being included in the volume marking the object under consideration. In A to D the result is displayed for 0.5 h post injection, 24 h post injection, 96 h post injection and 168 h post injection respectively. In E to H the median surface distance (back line with crosses) is shown together with the tenth percentile surface distances and 90:th percentile surface distances (red lines with no marks). Solid lines indicate the result when the pure SPECT-image was considered while broken lines indicate the result when the combined SPECT/CT-images were considered. For the case 0.5 h post injection (A and E) a normalisation exponent of 0.7 was used while for the other cases a normalisation exponent of 0.9 was used. Just as for single organ images, the mean surface distance decreases as more Fourier orders are used, with a spread of the distance for these particular cases indicated by the median surface distance and the tenth and 90:th percentiles. No major difference is observed between the results obtained when using pure SPECT-images and using combined SPECT/CT-images.

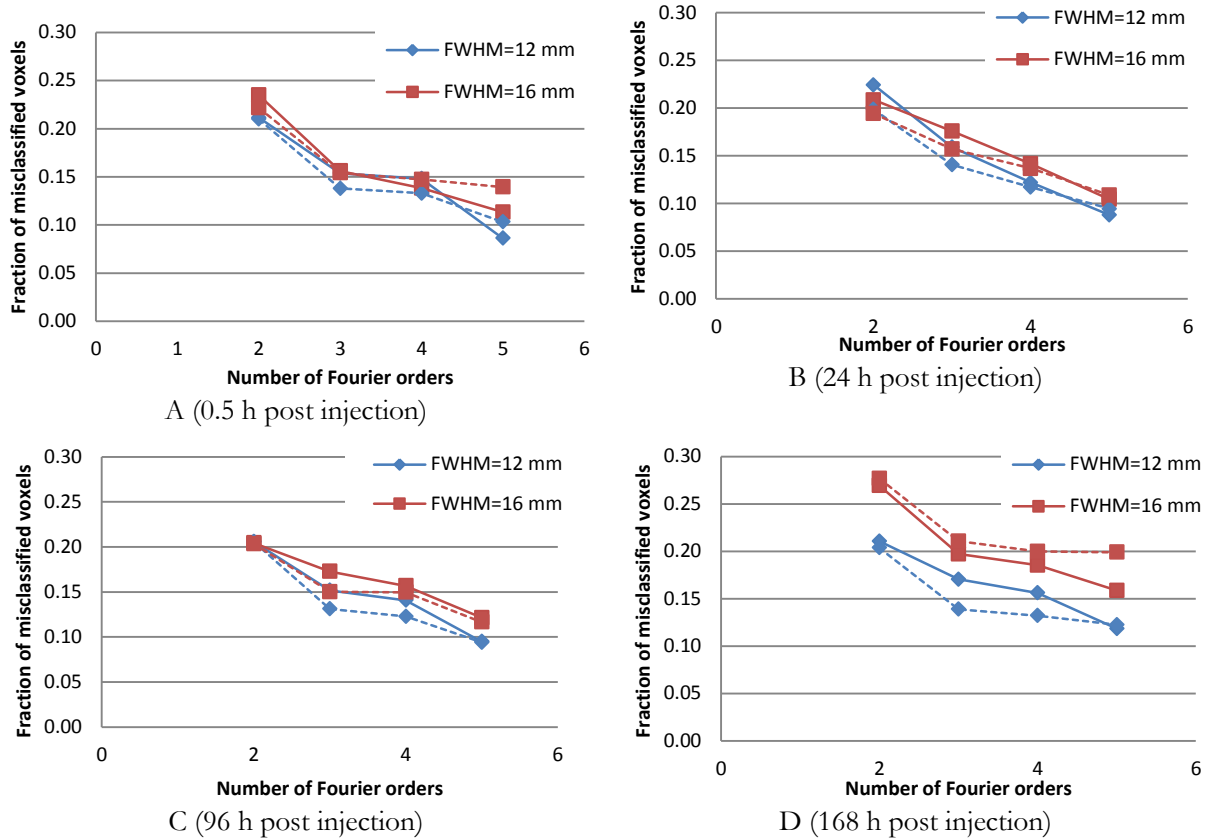


Figure 32: The corresponding results to Figure 31, but with respect to the fraction of misclassified voxels. Results are shown for right kidney only. When marking the left kidney, the segmentation frequently failed due to inclusion of the spleen. For the case 0.5 h post injection (A) a normalisation exponent of 0.7 was used, while for the other cases a normalisation exponent of 0.9 was used. Just as for single organ images, the fraction of misclassified voxels decreases as more Fourier orders are used. No major difference is observed between the results obtained when using pure SPECT-images and using combined SPECT/CT-images.

5.3 Monte Carlo simulated SPECT-images

The fraction of misclassified voxels as a function of the normalisation exponent, with values between 0.6 and 0.9, for SPECT-images obtained from SIMIND imitating SPECT-studies of a patient undergoing therapy with ¹⁷⁷Lu-DOTATATE at time points 0.5 h, 24 h, 69 h and 168 h post injection is shown in Figure 33. The corresponding partial fractions of missed voxels are shown in Figure 34. For normalisation exponents 0.7 and 0.9 for points in time 0.5 h and 24 h, 96 h and 168 h post injection respectively the mean surface distance for left and right kidney and the median surface distance together with the tenth and 90:th percentile surface distances for right kidney is shown in Figure 35 as a function of the number of Fourier orders used. The reason for not showing the result of the left kidney with respect to median surface distance was that segmentations marking that organ frequently failed due to incorporation of the spleen. The fractions of misclassified voxels for left and right kidney for the same normalisation exponents are shown as a function of the number of Fourier orders used in Figure 36.

Three weightings of SPECT- and CT-images are presented; one using the pure SPECT-image, one using the combined SPECT/CT-image and one using the combined SPECT/CT image where the CT-image has been given a factor four higher weight than for the original scheme. The results for the Monte Carlo simulated images are in agreement to what was observed for the SPECT-model with result presented in 5.2. No major difference is seen between using a scheme with higher weight given to the CT-image and the original scheme.

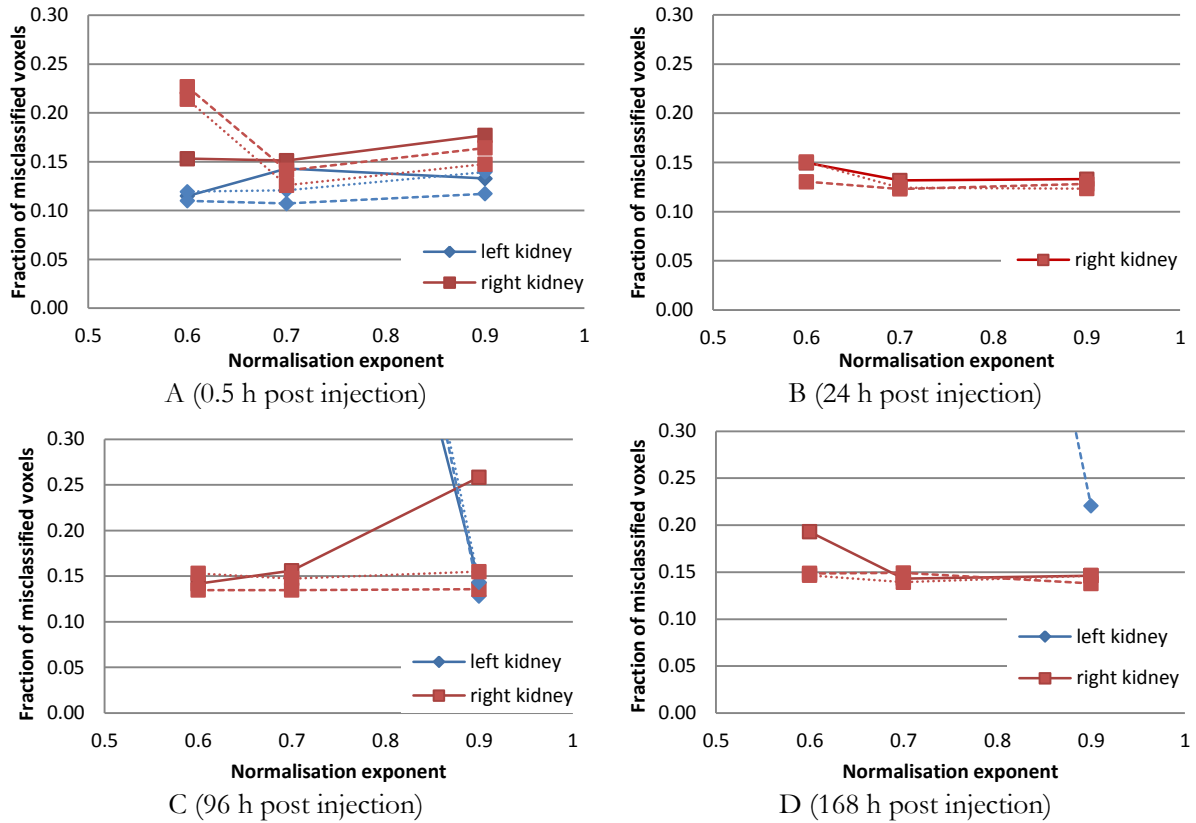


Figure 33: The fraction of misclassified voxels when applying the segmentation method to left and right kidney in images where the activity distribution for patients undergoing therapy with ^{177}Lu -DOTATATE using up to five Fourier orders as a function of the normalisation exponent. The result is displayed for 0.5 h post injection (A), 24 h post injection (B), 96 h post injection and 168 h post injection. For the cases where the result for the left kidney is not displayed, the segmentation has failed due to inclusion of the spleen. The results when segmenting the pure SPECT-images are shown as solid lines, the results when combining SPECT- and CT-images are shown as broken lines and the result when using combines SPECT/CT-images where the CT-images has been weighted a factor four higher compared to the standard method are shown as dotted lines. The minimum fractions of misclassified voxels appear for normalisation exponents 0.7 or 0.9. For most cases no major difference is observed between the different methods with respect to combing SPECT- and CT-information.

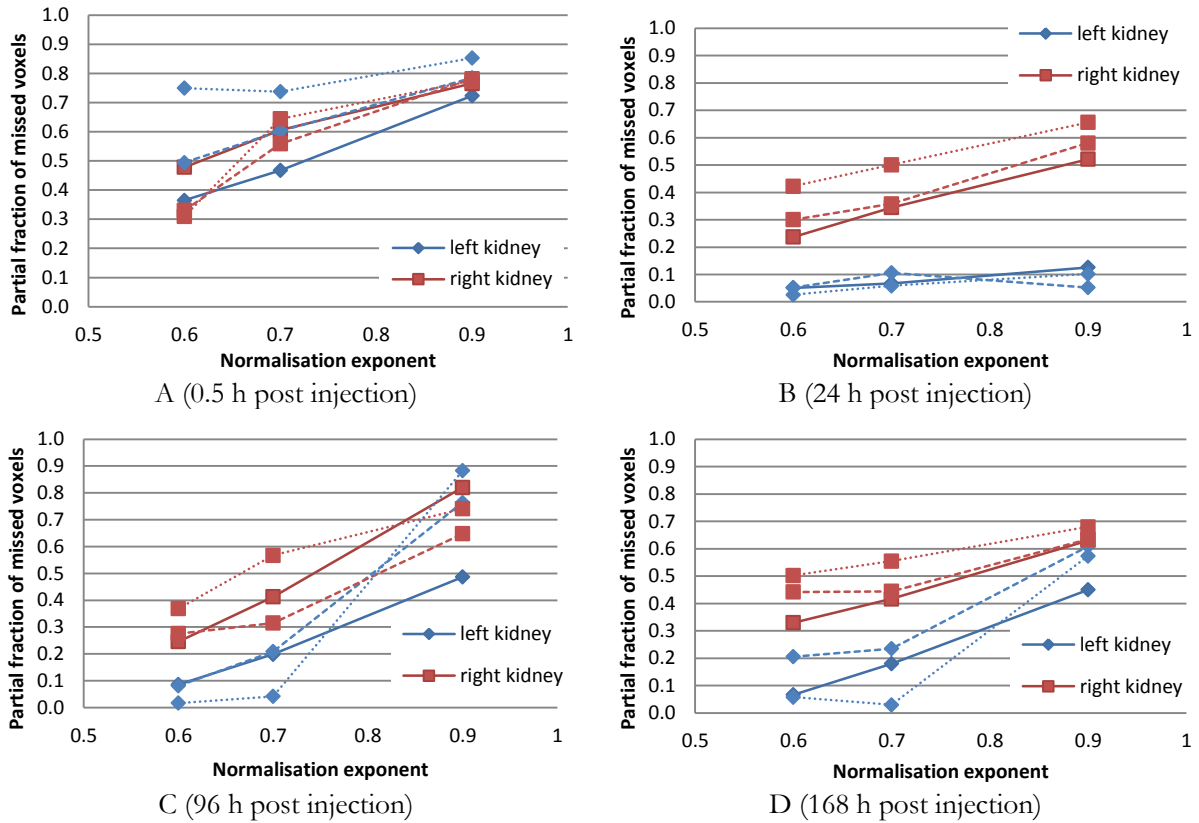


Figure 34: The partial fraction of missed voxels as a function of the normalisation exponent when applying the segmentation method to activity distributions imitating a patient undergoing therapy with ^{177}Lu -DOTATATE using up to five Fourier orders as a function of the normalisation exponent. The result is displayed for 0.5 h post injection (A), 24 h post injection (B), 96 h post injection and 168 h post injection. The results when segmenting the pure SPECT-images are shown as solid lines, the results when combining SPECT- and CT-images is shown as the broken lines and the result when using combines SPECT/CT-images where the CT-images has been weighted a factor four higher compared to the standard method are shown as the dotted lines. The segmentation method has preserved the total volume of the object under consideration if the partial fraction if misclassified voxels is $\frac{1}{2}$. When considering the result for left kidney in for example B, it must be remembered that these segmentations frequently failed due to inclusion of the spleen, causing a overestimation of the organ volume for reasons not directly connected to the normalisation exponent. Apart from that, the normalisation exponent preserving the organ volume seem to be between 0.7 and 0.9 for most cases, in approximate agreement to the normalisation exponent causing the minimum fraction of misclassified voxels in Figure 33.

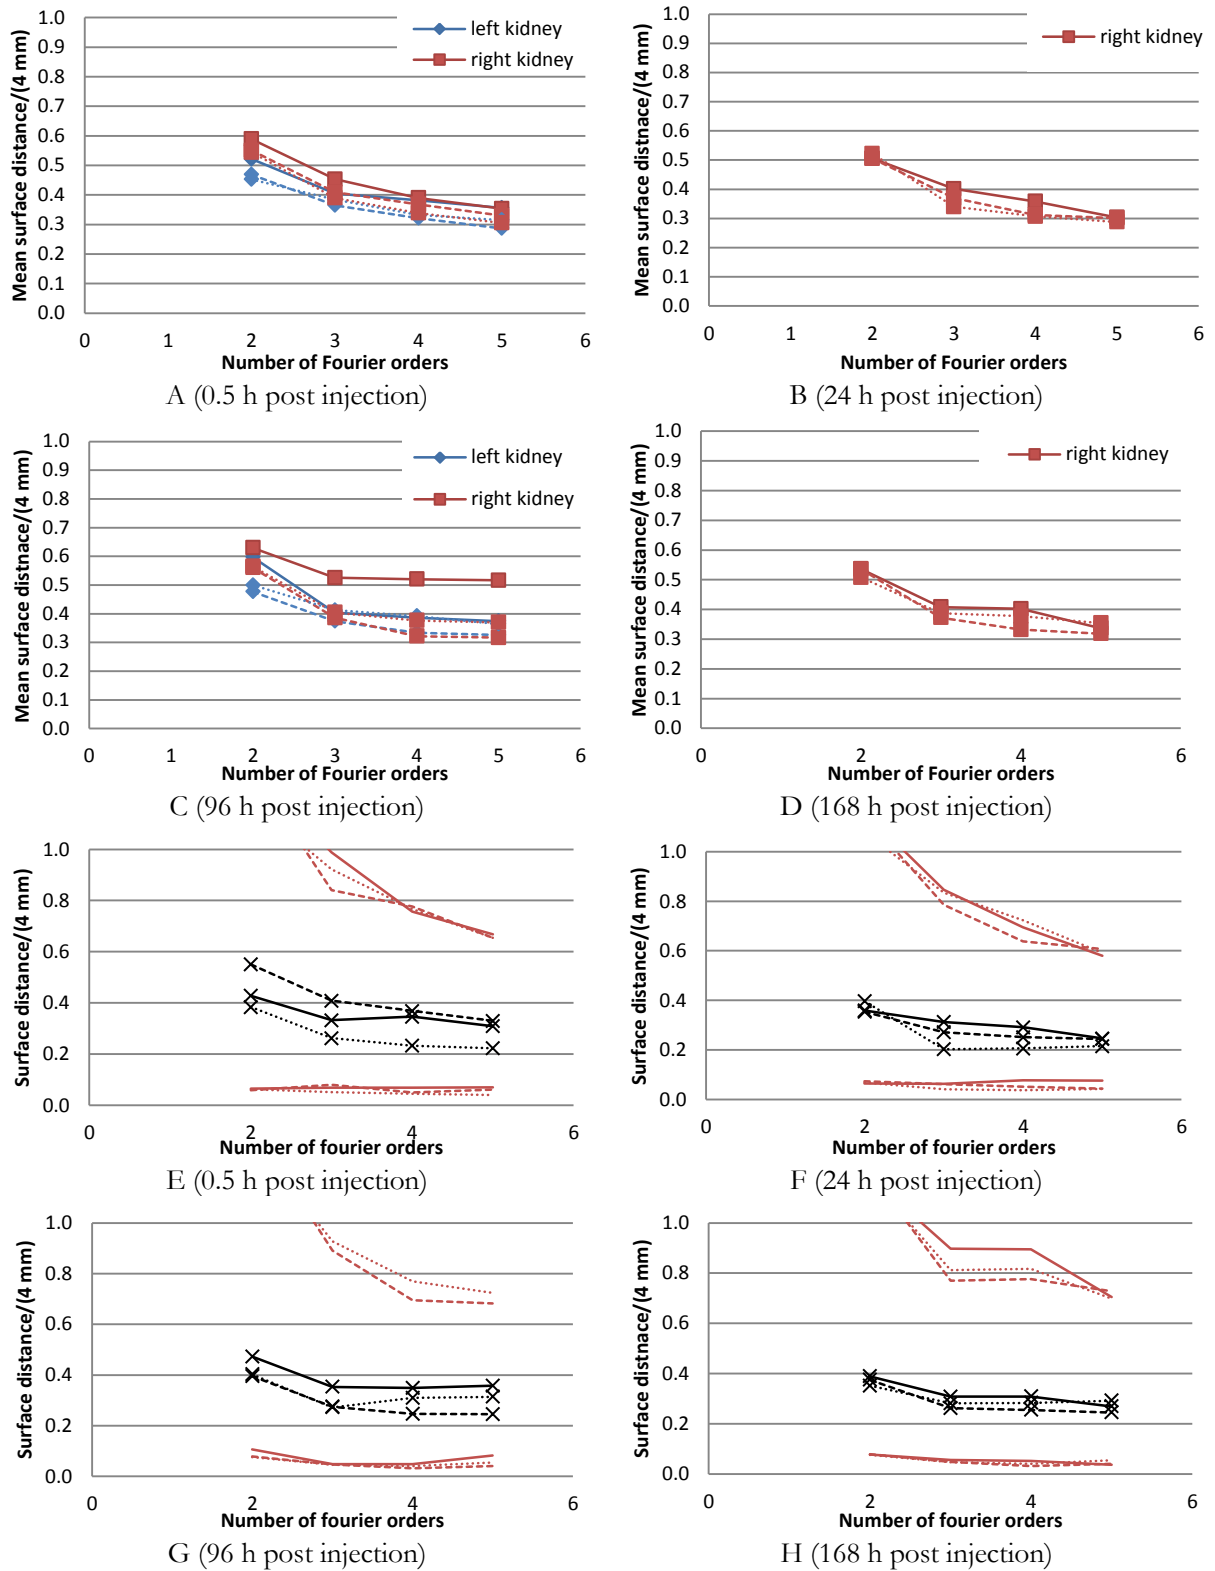


Figure 35: The mean surface distance as a function of the number of Fourier orders used when applying the segmentation method to Monte Carlo simulated SPECT-images imitating a patient undergoing therapy with ¹⁷⁷Lu-DOTATATE. For the left kidney, the segmentation frequently failed due to the spleen being included in the volume marking the object under consideration, and for B and D it is not shown in the diagram. In A to D the result is displayed for 0.5 h post injection, 24 h post injection, 96 h post injection and 168 h post injection respectively. In E to H the median surface distance (back line with crosses) is shown together with the tenth percentile surface distances and 90:th percentile surface distances (red lines with no marks) for the right kidney. Solid lines indicate the result when the pure SPECT- image was considered, while broken lines indicate the result when the combined SPECT/CT-images were considered and the result when using combines SPECT/CT-images where the CT-images has been weighted a factor four higher compared to the standard method are shown as the dotted lines. For the case 0.5 h post injection (A and E) a normalisation exponent of 0.7 was used while for the other cases a normalisation exponent of 0.9 was used. The result is similar to the simpler SPECT-model shown in Figure 31.

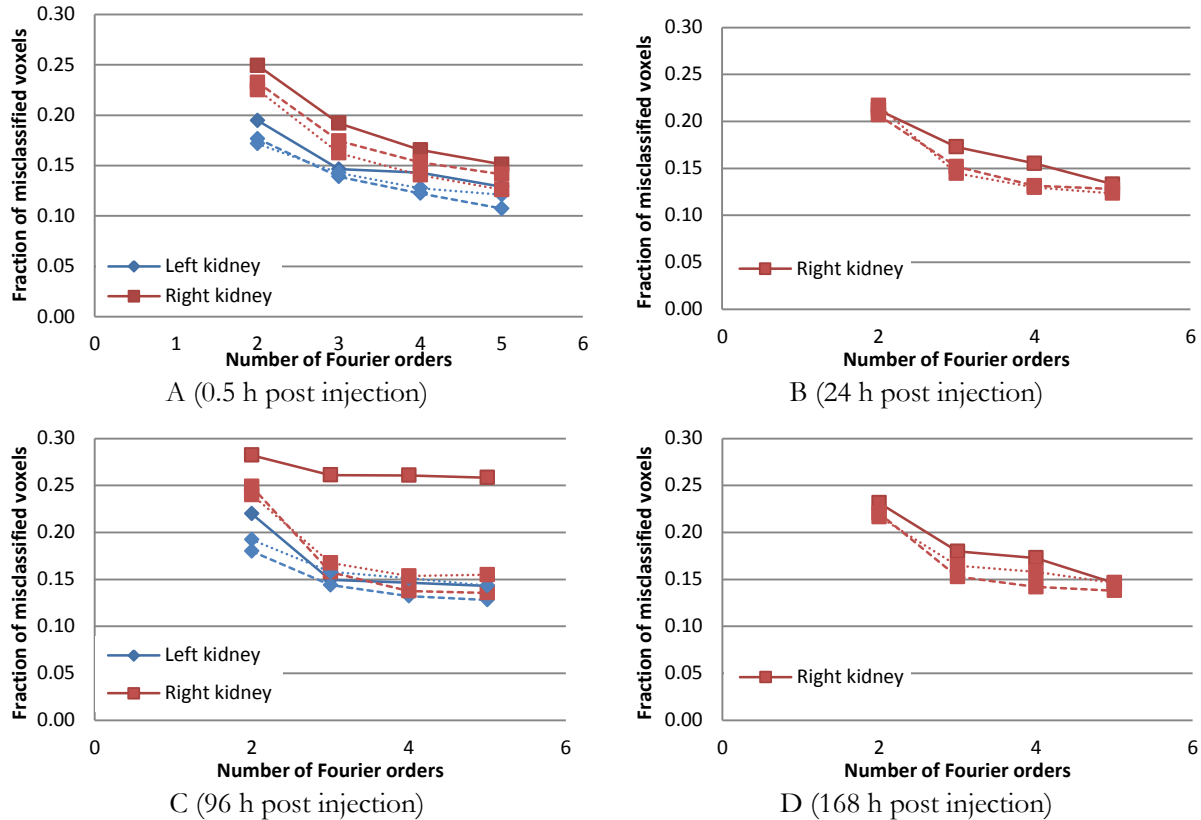


Figure 36: The corresponding results to Figure 31, but with respect to the fraction of misclassified voxels. Results are shown for right kidney. When marking the left kidney, the segmentation frequently failed due to inclusion of the spleen. For the case 0.5 h post injection (A) a normalisation exponent of 0.7 was used, while for the other cases a normalisation exponent of 0.9 was used. Just as for single organ images, the fraction of misclassified voxels decreases as more Fourier orders are used. No major difference is observed between the results obtained when using pure SPECT-images and using combined SPECT/CT-images.

5.4 Operator dependence in the segmentation of clinical SPECT-images

Out of 160 marked kidneys of 40 cases and four operators in clinical SPECT-images, 105 were found valid for future analysis on operator dependence, which corresponded to 66 %. Out of the not accepted segmentations, seven were believed only to require minor modifications in order to become acceptable and two could not be decided upon whether they were acceptable or not. In total there were 55 segmentations that were not accepted, whereof 43 showed patients' left kidneys (78 %). For each patient, a maximum of 32 segmentations could be accepted. The distribution between the five patients was 21, 26, 20, 31 and 7 accepted segmentations respectively.

Since the segmentations had to be valid for all operators in order for it to be included in the further analysis, the operator dependence could be investigated in 22 objects. Out of the 18 objects where not all segmentations were found acceptable, the segmentations were found unacceptable for all operators in ten cases. Examples of unacceptable segmentations for further analysis are seen in Figure 37 to Figure 40. The fraction of all voxels included by four, three, two and one operators as a mean plus/minus one standard deviation is shown in Table 2. Note that these numbers include both right and left kidneys and segmentations based on images from all four time point images investigated and the result is thus not based on completely equivalent images.




Figure 37: Segmentation marking the left kidney where the spleen has been included




Figure 38: Segmentation aiming to mark the right kidney where just a hotspot in the renal pelvis has been included by the algorithm. The image is heavily gamma corrected to suppress hotspots.




Figure 39: Segmentation aiming to mark the right kidney. The obtained surface was heavily distorted and the conversion to volume could not produce a solid volume.




Figure 40: Segmentation marking the left kidney where a small part of the spleen has been included in the region. This segmentation was thought only to require minor adjustments in order to be acceptable.

Table 2: The fraction of all included voxels by four (all), three, two and one operators plus/minus one standard deviation of the 22 cases where all segmentations were considered valid

Included by number of operators	Fraction of all included voxels
4	0.89 ± 0.03
3	0.04 ± 0.01
2	0.03 ± 0.01
1	0.04 ± 0.02

6. Discussion

In this work a method for segmentation of images based on a procedure where a surface described by Fourier descriptors was optimised to an object boundary was investigated. The aim was to segment SPECT/CT-images by marking kidneys in patients undergoing radionuclide therapy with ^{177}Lu -DOTATATE, where segmentation is necessary for the assessment of organ mean absorbed dose. The performance of the method has been tested on images showing objects of varying complexity. Special attention has been drawn to the conservation of volume by the segmentation algorithm, since it was found that there exists a parameter that must be tuned in order to make the objective function of the optimisation as neutral as possible to the object size, *i.e.* there should be no advantage for the optimisation procedure to expand or shrink the surface if there is no support for this in image data. The parameter to be tuned was given the name normalisation exponent.

It was also tested to combine the information from the SPECT-images with the information from the coregistered CT-images. A method was developed to weight together the information from the two modalities based on the boundary response and the image noise.

The segmentation performance was evaluated in two respects: 1) based on an error measure describing the distance between the segmented surface and a reference surface, and 2) based on the voxels enclosed by the surface. The surface comparison and the volume comparison reflected different aspects of the algorithm performance, and together enabled a thorough investigation of the sensitivity of the segmentation method to different sources of errors.

The segmentation algorithm was first tested on a set of single organ images and then on images imitating patient SPECT-studies. For the latter investigation, two SPECT-image models were used, one analytically based simulation performed directly on a three dimensional image volume by blurring and adding independent Gaussian noise, and one Monte Carlo based simulation of gamma camera projections that were reconstructed into a three dimensional tomographic volume. In both simulations an anthropomorphic phantom with a biodistribution of ^{177}Lu -DOTATATE similar to the average patient case was used.

The method was found to be resistant to image noise, with acceptable performance down to a contrast to noise ratio of approximately two and even lower for images of good spatial resolution. The method was also found to be able to delineate the kidneys in the simulated SPECT-images, but it encountered substantial problems with the spleen, since this organ also took up the radiopharmaceutical and the poor resolution in SPECT-images caused the nearby organ to be considered a part of the kidney by the

algorithm. This problem was probably even more pronounced due to the fact that in some of the simulated images the spleen had a higher uptake of the radiopharmaceutical than the kidneys.

No obvious advantage was seen when CT-images were combined with the SPECT-images. This could be due to the fact that the parameters used to weight together SPECT- and CT-turned out to be hard to estimate from the images, causing the CT-images to likely be underweighted compared to the SPECT-images. An alternative explanation could be that the kidneys are smooth objects in themselves and therefore the better spatial resolution of the CT-images did not contribute any further information to the segmentation basis.

The mean surface distance as a measure of the performance of the segmentation procedure

The accuracy of the mean surface distance as a measure of the segmentation performance has several aspects that motivate discussion. The first point to consider is the accuracy of the calculation of the mean distance between the segmented surface and the reference surface. The aim of evaluating equation (74) is to find the least distance between the segmented surface and the reference surface for a set of points distributed over the segmented surface and then weight that distance with respect to the area element each point represented. The finding of the least distance was done using golden section search. In this search it cannot be guaranteed that the optimisation method finds the truly shortest distance, since a unique minimum cannot be guaranteed, and the evaluation of equation (74) is likely to overestimate the mean surface distance in this perspective. When testing the evaluation algorithm by comparing a surface to itself, the mean surface distance was calculated to be 0.001 voxel sides and the maximum surface distance was calculated to be 0.09 voxel sides. The introduction of this error was not considered important for the purpose of evaluating the method, since the mean surface distance was typically much larger than 0.001 voxel sides.

On the other hand, there is an effect in the least distance estimation that counteracts the effect of the overestimation introduced by the numerical optimisation procedure. This effect is due to the search only being for the minimum distance and not necessarily for the distance between corresponding points. Ideally, the error measure should be evaluated for corresponding surface points. For fairly accurate segmentations, the distance between corresponding points and the least distance are likely to be approximately equal, but this can never be guaranteed. Even if this is an effect that does not affect the mean surface distance as defined by equation (74), this effect makes the mean surface distance smaller than what it should be according to a more intuitive definition, and for very dissimilar surfaces the effect could be substantial.

The same remarks are true for the median, tenth and 90:th percentile surface distances. In these estimations a round-off error also arises from that the segmented surface was only evaluated in a finite number of points, *i.e.* it was in general not possible to find a distance that fell below by *exactly* half of the surface. In this situation, the evaluation program was constructed to overestimate the distance, *i.e.* to assure that *at least* half, ten per cent or 90 per cent of the surface fell below the nominal distance. Typically the part of the surface exceeding the nominal fraction was less than one per cent.

Another aspect of the uncertainty of the mean surface distances is the reproducibility of a segmentation applied at equivalent images, in this aspect defined as two images with identical objects to be segmented, identical spatial resolution and identical CNR, but two different noise realisations. This aspect of the segmentation procedure is shown in Figure 21. In this diagram it is seen that the mean surface distance could vary almost 0.1 units of length for 12 mm resolution and CNR=4 just by alternating the noise realisation six times. In order to some extent avoid the variability introduced by that aspect all other results for the single organ images were based on the mean of two noise realisations. For single organ images with no added noise two realisations was obtained by changing the initialising ellipsoid slightly and thereby obtaining a similar effect. For the simulated SPECT-images one single noise realisation was used, making these more susceptible to the effect in question.

It must also be realised that an uncertainty is introduced in the results of the mean surface distance since the reference surface is not guaranteed to exactly follow the segmented volume. Therefore it is hard to interpret the absolute value of the mean surface distance as a measure of the performance of the segmentation procedure. The reference surface was manufactured under idealised conditions, since the spatial resolution was better and the image was free of noise. Therefore a relative comparison could still be done, *i.e.* a sensitivity analysis by this approach is justified.

The fraction of misclassified voxels as a test of the performance of the segmentation procedure

The fraction of misclassified voxels as a measure of the accuracy of the segmentation is not susceptible to many of the concerns expressed for the mean surface distance above. Since it is a simple comparison of voxels, no bias is introduced due to an optimisation procedure. It is however not a pure measure of the main segmentation scheme, since it also requires the step of converting the surface description into a volume description. The conversion from surface to volume was not a trivial task, and the method used was found to have several limitations. Since the classification of voxels was dependent on the direction of the surface normal, the method was found to be sensitive to surface artefacts such as folds, which in some cases could turn the surface ‘inside out’. Partly, this problem could be overcome using post processing trying to make the final volume to be solid, *i.e.* without holes, and connected, *i.e.* not divided into two or more parts.

Another approach for the volume conversion was also tested, based on the surface to be approximated by triangles prior to the interior being filled and the included voxels obtained. This approach had the appeal of being fairly simple compared to the method based on the direction of the surface normal. Unfortunately, it was hard to make it produce a good volume agreement. The reason for this was thought to be threefold. The first problem was whether the surface itself should be counted as a part of the volume or not. This problem was solved using the normal based approach, since it truly considered the surface as being infinitely thin. When finding a voxel coordinate exactly at the surface, or at least close enough to the surface in view of the numerical precision, a choice has to be made whether to count the voxel as a part of the object or a part of the background. This problem was much more emphasised if the alternative method was used.

The second problem, related to the first one, was due to that the method caused round off errors at the surface leading to the surface not to be just one voxel thick. In order to reduce this error the calculation was performed based on an enlarged version of the surface, but the magnification factor was limited by computer memory requirements. The third problem, related to the second one, was to decide how densely to discretise the surface in the triangle approximation. This discretisation had to be rather dense if a large magnification factor was used. Too sparse a discretisation caused the volume not to give a true estimation of the true enclosed volume. One way to solve the third problem would have been to implement a method to automatically decide the density needed for a certain error tolerance. If this last problem was solved, the method could have been interesting at least as a complement to the one that was finally used, but overall, the method based on the surface normal was found to take better advantage of the analytical surface description given by Fourier descriptors.

The variability of the performance of the segmentation method due to different noise realisation is still an issue for the fraction of misclassified voxels, just as it was an issue for the mean surface distance, since that uncertainty is a part of the segmentation method itself rather than a concern of the error measure used. In Figure 21 the fraction of misclassified voxels ranges from 0.11 to 0.18 for segmentations of equivalent images.

For the single organ images, the comparison between the segmented volume and the object was direct and therefore no uncertainty due an intermediate step has to be taken into account as for the use of reference surfaces. For reasons previously discussed for the simulated SPECT-images, an intermediate step had to be taken via the generation of a reference surface by segmentation under idealised conditions. This step was motivated by the geometry in the XCAT phantom, including structures in the kidneys such as blood vessels and renal pelvis and medulla. Methodologically, a choice had thus to be made concerning the aims of the segmentation. For example, the purpose of the segmentation was not to be able to separate the kidney from blood vessels in the kidney, thus the blood vessels had to be removed from the reference volume. The same reasoning had to be done with respect to the renal pelvis, since the method was not intended to be able to separate the kidney medulla and the renal pelvis. Neither was it intended to be able to include the outgrowth existing in the renal pelvis in the XCAT phantom. Using the intermediate step of a reference volume was a way to manage these considerations, and a sensitivity analysis based on this approach is considered justified.

Mean surface distance compared to fraction of misclassified voxels

The use of mean surface distance for evaluation of the segmentation process has several drawbacks, as partly described above, not shared by the use of fraction of misclassified voxels. The latter also has the advantage of making it possible to say whether the segmentation systematically overestimates or

underestimates the volume to be segmented by applying the measure that was called partial fraction of missed voxels. The major disadvantage of using the volume based method is that it is not a pure measure of the segmentation method based on surfaces but is also susceptible to errors in the conversion from a surface to a volume description.

One advantage of the surface based error measure is that it can give information about the distribution of the error if the version based in the median and tenth and 90:th percentile is used. This information is not easily extracted in numerical form from the fraction of misclassified voxels since this measure basically just says whether a voxel is correctly classified, not by which amount a certain voxel is misjudged.

In conclusion the surface based and volume based measures show two complementary aspects of the segmentation performance. The volume based measure is on the whole a better one for evaluation of the segmentation procedure, and it can give more information than just a comparison between the segmented surface and a reference surface is able to give, for example if the volume of the organ is underestimated or overestimated. Which measure to use is finally due to the application of the segmentation method. If the aim is to classify voxels, as was the objective for this work, the volume based measure is to be preferred. If on the other hand the primary interest is on the surfaces themselves, then a surface based error measure is preferable. For this work, the combination of the two evaluation measures was found very useful, reflecting different aspects of the segmentation method.

Performance of the segmentation method with respect to noise and spatial resolution

Generally, the performance of the segmentation algorithm is negatively affected by the non-ideal resolution of images and by image noise. This is illustrated by the investigation in single organ images. In the first instance of that investigation, where the normalisation exponent was set to unity, the mean surface distance increased from approximately 0.4 mm to approximately 1.6 mm between the images with best resolution, 4 mm, and the image with the worst resolution, 20 mm, as seen in Figure 20. In terms of the fraction of misclassified voxels, this mean surface distance corresponded to a fraction of misclassified voxels from approximately 0.05 to more than 0.20. The same pattern goes for CNR, where high CNR generally gives a better correspondence than a low CNR. Interplay is also evident between the resolution and CNR, since the segmentation algorithm gives fairly correct segmentations even for CNR as low as unity for a spatial resolution of 4 mm (Figure 22 A and Figure 23 A), while the method seems to get into troubles if that CNR is applied to poorer spatial resolutions. The effect of an inferior spatial resolution is likely to be twofold since there is one direct effect of the blurring that makes the edge of the object less clearly defined and one indirect effect that decrease the signal in the gradient images and thereby their signal to noise ratio. The effect of image resolution and noise is evident also when the normalisation exponent is changed, as seen in *e.g.* Figure 25.

It should be remembered that the tests of the performance of the method with respect to noise and spatial resolution in single organ images was performed with an initialisation based on perfect images. So, what has actually been shown is that provided that a good enough initialisation the method is capable of performing acceptable segmentations of an image. In practice, also the initialisation will deteriorate as the noise level increases and the spatial resolution of the image decreases, which will affect the segmentation negatively.

The number of Fourier orders used

As is evident for most cases of successful segmentations, the mean surface distance and the fraction of misclassified voxels decreases as the number of Fourier orders are increased. Often there is a quick decrease between $K=2$ and $K=3$ and then the rate of the improvement is decreased, even if there are examples of a steady decrease with Figure 32 B as an obvious example. For the single organ images it must be remembered that this test object was based on a segmentation using $K=6$, limiting the number of Fourier orders likely to be present in the object itself and thus $K=6$ should be enough for the surface to form itself to the object.

How many Fourier orders to be used in clinical images is not obvious. In this work the number was finally limited to $K=5$. This was much due to time considerations, since a normal segmentation then took between five and ten minutes, which was thought to be acceptable, while the decrease in error was considered to be acceptable small for the current typed of objects when higher orders were. A much faster segmentation, maybe two or three minutes, could be achieved if fewer Fourier orders were used.

In terms of time spent, the optimisation method plays a key role, and the downhill simplex method is a slow way of doing numerical optimisation. As described by Floreby *et al.* [20] other optimisation methods

could be favourable in this perspective. It is also worth mentioning that the optimisation method is likely to be one of the key components for the method to be successful. Alternatives to the amoeba method and their performance with respect to the error measures have not been investigated in this work.

Constraints of the optimisation procedure

The amoeba method for numerical optimisation must be given a scale when the initial simplex is spanned. The scale parameter is not a hard constraint to the optimisation but rather a way to give some general size of the uncertainty of the initial guess on where the optimal solution lies. To give too small a scale of the problem makes it hard, but not impossible, for the method to find the correct solution. This is illustrated in Figure 24 for the scale(2) and scale(3) parameters, where the segmentations fails for scale(2)=1, while values larger than that are all equivalent in their performance. The sharp increase in the error when lowering the scale(2) parameter is truly remarkable. There are indications that this result was partly due to miscommunication between the initialisation procedure and the optimisation procedure. The communication problem thought to exist made it necessary to have a large enough scale(2) parameter in order to correct for the introduced error. It is believed that this error did not substantially influence the segmentation result as this first scale parameter was just large enough, motivated by the steady result for scale(2) larger than one in Figure 24 A and C.

The same trend is seen for the scale(3) parameter in Figure 24 B and D. The effect of the scale(3) parameter is less marked than when the scale(2) parameter is changed, which also should be the case since the $K=3$ step in one perspective is a fine tuning of the result from the $K=2$ step and since the potential communication error had already been corrected for by the $K=2$ step. If the different curves of Figure 24 B and Figure 24 D are considered, the increase in the error for low values of scale(3) is more pronounced for the $K=3$ curve, while the segmentation seems to be able to partly recover for higher values of K .

It should be noted that even if no negative effect was seen when the scale parameters were increased in the single organ images, this effect is likely to exist in more complex images, since a large interval is likely to contain poor local minima. Therefore, it is probably advisable to keep the scale parameters as small as possible while still assuring they are large enough to cover the optimal solution. This is considered a subject where further optimisation of the method is required.

The distribution of the error

A few examples are given in the results showing the median of the distance between the segmented surface and the reference surface together with the tenth and 90:th percentile of the distance. The median surface distance gives roughly the same sort of information as the mean surface distance, but when combined with the tenth and 90:th percentile surface distances it also gives information about the distribution of the error. The tenth and 90:th percentile surface distances have a straightforward interpretation and were therefore chosen in favour of the standard deviation as an indication of the spread of the surface error. From the shown examples it is evident that more than 90 % of the surface is less than one voxel side from the reference surface for the shown examples. The reference to one voxel side is interesting since effects of the discretisation of the images should be evident for smaller distances.

The normalisation exponent

In this work a variable a called the normalisation exponent has been introduced as a parameter in the objective function. For most other works, this parameter has been implicitly set to either one or zero. The choice of zero can be motivated based on the derivation by Staib and Duncan [14, 19] giving the objective function a maximum likelihood interpretation. Their derivation was however based on the assumption of the surface not changing its area in the optimisation. That assumption is probably justified if only minor corrections are to be done to the surface during the optimisation, *i.e.* the initialisation is close to the final solution. If only a rough first estimation is available and the parameter space has to be evaluated in a larger interval, then the assumption of a constant area fails. Since a larger area will increase the value of the objective function, the optimisation risks going out of bounds. This is the rationale for including a normalisation of the objective function to the surface area. Unfortunately, such a normalisation destroys the maximum likelihood interpretation of the objective function. Instead of a normalisation, for a maximum likelihood interpretation to hold the true boundary response, k , should be subtracted from twice the measured boundary response for every term in the summation, *i. e.*

$$E = \sum_s (2b(u,v) - k) \Delta A(u,v). \quad (77)$$

The true boundary response is normally not known, and it is hard to estimate from images, as will be discussed later in the context of combining SPECT- and CT-images, making this objective function impractical.

Setting the normalisation exponent to a value of one will cause the objective function to have a nice interpretation as the mean response along the surface. As demonstrated in for example Figure 26 the use of this function will cause the segmentation to be biased toward too a small organ volume. A possible explanation for this phenomenon could be that the optimisation finds it favourable to shrink the surface in areas where the surface response is below average, and a smaller surface generally means a smaller volume, at least for relatively simple shapes. It is true by the same arguments that the optimisation algorithm will find it advantageous to increase the surface area in spots where the surface response is above average. A strong boundary response indicates however that the true boundary location is found, and it will be hard for the optimisation algorithm to find a way to increase the area and still stay on the boundary, making the tendency for expansion much less of a problem in terms of biasing the organ size.

As understood from the discussion above, the area raised to one is not obvious as a choice for normalisation and another exponent could be chosen. Since the division can never replace the subtraction in equation (77) for all cases, it is probably not feasible to make the objective function completely area neutral, but at least a strong bias could be avoided by a good choice of the exponent as also shown in Figure 26, indicating the best choice to be about 0.5 to 0.9 for most images. In Figure 25 it is possibly indicated that this choice for exponent also improves the overall performance of the segmentation algorithm, at least for high to moderate CNR. It has not been investigated in this study to what extent this result can be generalised to other types of objects, for example the liver.

When a distribution of a radiopharmaceutical is considered, there is a tendency for later points in time to gain from increasing the normalisation exponent compared to the first point in time. It should be remembered that these segmentations are just done once for each organ and point in time, so the effect of varying result for different noise realisations makes it hard to draw any definitive conclusions. The reason for such an effect is unknown. Possible explanations could be the different distributions of the radiopharmaceutical, where the first point in time is rather different compared to the later ones with for example much more activity being in the blood pool, or it could be due to different noise levels. Even if the exact cause is unknown, this illustrates that tuning the normalisation exponent is not trivial.

Effects of organ shape

As seen in Figure 28, the fraction of misclassified voxels increases when a more complex organ such as the spleen is the organ under consideration compared to the kidney based shapes. This result is natural since the shape of the spleen is much more demanding for the optimisation procedure. Also natural is the poor result of low order surfaces in these shapes since a surface of low order simply is incapable of correctly representing the organ in question. It should also be remembered that the initialisation works less well for less ellipsoid like organs, further worsening the result of early stages in the optimisation and making the task of optimisation more difficult as a whole. The strategy of the optimisation had to be changed when the spleen was considered instead of the kidneys, therefore it is possible that the spleen results would be better had a more thorough tuning of the segmentation parameters, such as the scale parameters, been performed.

Perspectives on the fraction of misclassified voxels

To put the obtained fractions of misclassified voxels into perspective, the results can be compared to a known introduced error introduced to two otherwise identical volumes. As such a known introduced error, the binary volume based on the XCAT phantom left kidney was translated in one pixel in all directions. For the left kidney, a translation in the x-direction was equivalent to a fraction of misclassified voxels of 0.22, a translation in the y-direction a fraction of 0.23 and a translation in the z-direction a fraction of 0.17. A combined mismatch of one voxel side in all direction produced a fraction of misclassified voxels of 0.35.

These errors illustrate the fact that the kidneys are small organs. The fraction of misclassified voxels would probably be smaller if a larger structure, such as the liver, had been considered. It must also be remembered that that the fraction of misclassified voxels is composed of two types of errors, falsely

included voxels and missed voxels. As clearly illustrated by the experiments where the normalisation exponent was varied, the relative proportions of these two components can vary substantially. There is a tendency that the overall result of the segmentation is improved if these components are approximately equal, but regardless of that, a segmentation where the total volume of the object of interest is preserved must be considered better than segmentations where the volume is systematically overestimated or underestimated.

Disturbing objects and the possible incorporation of a priori knowledge

When considering the marking of the left kidney for the later times in the test images based on the biodistribution of ^{177}Lu -DOTATATE, it becomes obvious that high signal objects close to the object under consideration is a potential problem for the segmentation algorithm. It must be recognised that the left kidney and the spleen are very hard to clearly separate in images of poor resolution and high noise levels. This also goes for a human observer, even if a human has the capacity of judging, at least roughly, what part of the image is the kidney and what part of the image is the spleen.

The problem experienced for the automatic segmentation algorithm is thought to be an effect from two different sources. First, the spleen actually has higher signal than the kidneys for the later points in time. Since a higher signal result in a stronger boundary support, it is natural for the optimisation routine to pick out this object rather than just the kidney. In the current state of the segmentation algorithm there is little that can be done about this. Of course a good initialisation will make an erroneous segmentation less likely, especially since a more accurate initialisation would allow a more constrained search for the optimal fit of the surface to image data, but there is little hope this will substantially change the overall outcome substantially. The second reason why the segmentation algorithm makes such remarkable mistakes is that it is blind with respect to the organ shape and size and only judge the goodness of a certain surface based on image data.

One way to improve the robustness of the method, even if image data alone does not support the correct segmentation of the object, would be the incorporation of *a priori* knowledge of the organ appearance into the surface model. There are numerous ways to incorporate such information into deformable models and several reviews have been published on the subject [6, 7]. For the purpose of Fourier descriptors, Staib and Duncan [14, 19] proposed a method based on assigning a probability distribution to each Fourier coefficient, thereby introducing a bias of the organ shape towards its average value. It is uncertain whether such a probability distribution would be feasible for the higher order Fourier coefficients based on the image material considered in this work, but it is possible that a bias would be helpful for the coefficients associated with the overall size of the object, *i.e.* the coefficients associated with the initialising ellipsoid.

The combination of SPECT- and CT-images

In this study no beneficial effects have been found by combining SPECT-images with its corresponding CT-images in terms of obtaining a good segmentation. This conclusion is not based on a statistical evaluation of segmentations based on pure SPECT-images versus combined SPECT/CT-images, but no systematic difference between the two groups is obvious if the graphs of the evaluation of simulated SPECT-images are considered. A more rigorous investigation of the possible benefits of the inclusion of CT-images is therefore likely to be time consuming and unlikely to demonstrate any major advantage.

The reason for the combination of SPECT- and CT-images not to demonstrate any advantage for the segmentation result compared to pure SPECT-images can be due to several reasons. First, the weighting of SPECT and CT-images, though based on a theoretical ground, in practice turned out to be very difficult. The estimation of the boundary response and the variance of the gradient images were based on surfaces of low order. These surfaces were not likely to fit perfectly to the object boundary, thereby decreasing the measured average boundary response and increasing the variance. This effect is likely to be more severe on the relatively high resolution CT-images compared to the low resolution SPECT-images. This effect results in systematic underweighting of the CT-images compared to the SPECT-images. This was the reason for increasing the weight of the CT-images, as shown in Figure 33 to Figure 36, but it resulted in no obvious improvement.

Another possible reason for the not found advantage of incorporating CT-images could be that the kidneys are relatively smooth objects. It is possible that the CT-images only contribute to a better segmentation result if higher orders are used. It is also possible that such high Fourier coefficients are not supported by the kidneys as an organ, since not only is the kidney appearance in SPECT-images smooth,

but also the kidneys in themselves are smooth organs. It should also be noted that the CT-images that were tried to be simulated was in the form of a low dose CT and thus not of a diagnostic quality. It is possible that the better general image quality of diagnostic CT-images would greatly increase the contribution of the inclusion of CT-images to the segmentation result.

Unfortunately, even if a CT of diagnostic quality were included in the basis for segmentation of the image, there are a number of effects that have not been considered. In practice the SPECT- and CT-images are often not perfectly coregistered due to for example patient movement. This will obviously destroy many of the potential advantages of including CT-images.

The use of the method on clinical images

The use of the method on clinical images partly gives another result than its use on simulated images. For example, many more segmentations of images by marking the left kidney for times later than 0.5 h post injection were successful than for the corresponding situation in simulated images. This illustrates that the differences between patients, often seen as a problem for automatic segmentation algorithms, also can be an advantage. In one of the patient cases there was only little uptake of the radiopharmaceutical in the spleen, and thus the interference between the spleen and left kidney always present in the simulated images was not a problem. On the other hand some deleterious effects, such as high concentration of the radiopharmaceutical in the renal pelvis as seen in Figure 38, were not considered in the simulated images.

The practical use of the method is today hindered by the high percentage of the images where the segmentation method fails. In the clinical examples used in this work, the segmentation method was judged to deliver valid results in 66 % of the cases. It must again be stressed that this figure is not based on the clinical usefulness of the method, but only reflects gross mistakes made by the algorithm. If the obtained segmentations actually are useful for clinically relevant questions remains a question to be answered. It is perhaps possible to further reduce the number of invalid segmentations by tuning the parameters used for the segmentation, but in order to substantially increase the rate of success for the method, *a priori* information is likely to have to be incorporated in the model as previously discussed.

Another aspect of the practical usefulness of the method is the degree of precision of successful segmentations and to what extent they depend on the operator doing the initialisation of the process. As shown in Table 2 this operator dependence is low on average. Accordingly, the problem of the segmentation algorithm does not appear to be the error of an acceptable segmentation, but the frequency by which the method fails to correctly segment the image at all.

It could be argued that the high figure of 89 % of voxels included by all operators in Table 2 is affected by that the objects that were successfully marked based on the initialisation by just some of the operators could have been harder examples than the ones where all operators were successful, and thus would affect the result downwards if included in the analysis. This could of course be the case, but if the three cases where all but one segmentation were successful are studied, excluding the failed segmentation, again it turns out that about 90 % of the voxels included by any operator are included by all operators. The same is true for the two cases where just two segmentations were correct. It is unfortunately not straightforward to compare results of fractions included by all operators if just two operators are considered with results based on four operators, since a voxel is more likely to be included by all operators if just two operators are considered than if four operators are considered.

Processing the images prior to segmentation

A possibility, whose effect on the outcome has not been investigated in this work, is to filter the image before segmentation in order to suppress noise. For images with low contrast to noise ratio it is possible that the loss of spatial resolution due to the filtering is more than compensated by the better noise characteristics. Implicitly, a smoothing of the images is already present in the method in its current state. This since the calculation of the image gradient is performed using 3x3x3 kernels and the approximation of derivatives thus averaged over nine differences.

Alternative applications - open and tubular surfaces

This work has been focused on the use of Fourier descriptors for describing closed surfaces. As mentioned in section 3.1, surfaces described by Fourier descriptors do not necessarily have to be closed. On the contrary, the surface description will in many perspectives be simpler for open or tubular surfaces. The closed surface description is important because many structures relevant in medical imaging are best

described by a closed surface. For segmentation of structures such as for example the aorta, the tubular surface description would be the proper choice. This is a field open for further investigations.

7. Conclusions

The proposed method for segmentation of smoothly appearing objects has been demonstrated to have a high resistance to image noise and is able to segment even noisy images to a fairly high degree of accuracy. The performance of the methods deteriorates when the noise levels are high (typically CNR less than two) and the image spatial resolution poor, but the performance is acceptable for clinically relevant noise levels and resolutions. The main drawback of the method seems to be its tendency also to incorporate other high signal objects near the organ of interest, causing the procedure to fail with its aim for such situations.

There appears to exist a parameter, in this work called the normalisation exponent, that must be tuned in order for the segmentation to be volume neutral, *i.e.* not to prefer large or small volumes. A general optimum value for this parameter is hard to find and is likely to be dependent on the image characteristics. Also for each application of the method an adaption of the optimisation scheme has to be performed. Once that is done, the operator dependence of the method is low.

The proposed segmentation method has been proven to have some advantages that make it interesting for further investigations. For the method to be practically useful, it must be applied to images showing smooth shapes. Before a widespread practical use of the method, the rate of successful segmentations must also increase from its current value of approximately two out of three.

8. Acknowledgements

I would like to thank my supervisors Katarina Sjögren Gleisner and Michael Ljungberg for help and support. Special thanks are also directed to Gustav Brolin for cooperation and fruitful discussions, Karin Wingårdh, for help with testing the initialisation process and discussions on the result, and André Ahlgren also for help with testing the initialisation procedure.

9. Appendix

In this appendix the direct equivalence of the complex version of Fourier series and the real version of Fourier series used in this work is shown.

Assume f to be a function of two variables u and v . Use the definition

$$\hat{f}_{m,l} = \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) e^{-i(mu+lv)} du dv \text{ for the Fourier coefficients.}$$

The function is then described by a Fourier series according to

$$\begin{aligned} f(u, v) &= \sum_{m=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} \hat{f}_{m,l} e^{i(mu+lv)} = \hat{f}_{0,0} + \sum_{m=1}^{\infty} \left\{ \hat{f}_{m,0} e^{im u} + \hat{f}_{-m,0} e^{-im u} \right\} + \sum_{l=1}^{\infty} \left\{ \hat{f}_{0,l} e^{il v} + \hat{f}_{0,-l} e^{-il v} \right\} + \\ &+ \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} \left\{ \hat{f}_{m,l} e^{i(mu+lv)} + \hat{f}_{-m,l} e^{i(-mu+lv)} + \hat{f}_{m,-l} e^{i(mu-lv)} + \hat{f}_{-m,-l} e^{i(-mu-lv)} \right\} = \hat{f}_{0,0} + \sum_{m=1}^{\infty} \left\{ \hat{f}_{m,0} (\cos mu + i \sin mu) + \right. \\ &+ \left. \hat{f}_{-m,0} (\cos mu - i \sin mu) \right\} + \sum_{m=1}^{\infty} \left\{ \hat{f}_{0,l} (\cos lv + i \sin lv) + \hat{f}_{0,-l} (\cos lv - i \sin lv) \right\} + \\ &+ \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} \left\{ \hat{f}_{m,l} (\cos[mu+lv] + i \sin[mu+lv]) + \hat{f}_{-m,l} (\cos[-mu+lv] + i \sin[-mu+lv]) + \right. \\ &+ \left. \hat{f}_{m,-l} (\cos[mu-lv] + i \sin[mu-lv]) + \hat{f}_{-m,-l} (\cos[mu+lv] - i \sin[mu+lv]) \right\} = \hat{f}_{0,0} + \sum_{m=1}^{\infty} \left\{ (\hat{f}_{m,0} + \hat{f}_{-m,0}) \cos mu + \right. \\ &+ \left. i (\hat{f}_{m,0} - \hat{f}_{-m,0}) \sin mu \right\} + \sum_{l=1}^{\infty} \left\{ (\hat{f}_{0,l} + \hat{f}_{0,-l}) \cos mu + i (\hat{f}_{m,0} - \hat{f}_{-m,0}) \sin mu \right\} + \end{aligned}$$

$$\begin{aligned}
& + \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} \left\{ \hat{f}_{m,l} (\cos mu \cos lw - \sin mu \sin lw + i \sin mu \cos lw + i \cos mu \sin lw) + \right. \\
& + \hat{f}_{-m,l} (\cos mu \cos lw + \sin mu \sin lw - i \sin mu \cos lw + i \cos mu \sin lw) + \\
& + \hat{f}_{m,-l} (\cos mu \cos lw + \sin mu \sin lw + i \sin mu \cos lw - i \cos mu \sin lw) + \\
& + \left. \hat{f}_{-m,-l} (\cos mu \cos lw - \sin mu \sin lw - i \sin mu \cos lw - i \cos mu \sin lw) \right\} = \hat{f}_{0,0} + \sum_{m=1}^{\infty} \left\{ (\hat{f}_{m,0} + \hat{f}_{-m,0}) \cos mu + \right. \\
& + i (\hat{f}_{m,0} - \hat{f}_{-m,0}) \sin mu \left. \right\} + \sum_{l=1}^{\infty} \left\{ (\hat{f}_{0,l} + \hat{f}_{0,-l}) \cos mu + i (\hat{f}_{m,0} - \hat{f}_{-m,0}) \sin mu \right\} + \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} \left\{ (\hat{f}_{m,l} + \hat{f}_{-m,l} + \right. \\
& \hat{f}_{m,-l} + \hat{f}_{-m,-l}) \cos mu \cos lw + \left. (-\hat{f}_{m,l} + \hat{f}_{-m,l} + \hat{f}_{m,-l} - \hat{f}_{-m,-l}) \sin mu \sin lw + \right. \\
& + i (\hat{f}_{m,l} - \hat{f}_{-m,l} + \hat{f}_{m,-l} - \hat{f}_{-m,-l}) \sin mu \cos lw + \left. i (\hat{f}_{m,l} + \hat{f}_{-m,l} - \hat{f}_{m,-l} - \hat{f}_{-m,-l}) \cos mu \sin lw \right\}.
\end{aligned}$$

Define:

$$\begin{aligned}
a_{m,l} &= \frac{1}{4} (\hat{f}_{m,l} + \hat{f}_{-m,l} + \hat{f}_{m,-l} + \hat{f}_{-m,-l}) \\
b_{m,l} &= \frac{1}{4} i (\hat{f}_{m,l} - \hat{f}_{-m,l} + \hat{f}_{m,-l} - \hat{f}_{-m,-l}) \\
c_{m,l} &= \frac{1}{4} i (\hat{f}_{m,l} + \hat{f}_{-m,l} - \hat{f}_{m,-l} - \hat{f}_{-m,-l}) \\
d_{m,l} &= \frac{1}{4} (-\hat{f}_{m,l} + \hat{f}_{-m,l} + \hat{f}_{m,-l} - \hat{f}_{-m,-l}).
\end{aligned}$$

Note that the definition of $a_{m,l}$ and $b_{m,l}$ and $c_{m,l}$ means that:

$$\begin{aligned}
a_{m,0} &= \frac{1}{4} (2\hat{f}_{m,0} + 2\hat{f}_{-m,0}) = \frac{1}{2} (\hat{f}_{m,0} + \hat{f}_{-m,0}) \\
a_{0,l} &= \frac{1}{4} (2\hat{f}_{0,l} + 2\hat{f}_{0,-l}) = \frac{1}{2} (\hat{f}_{0,l} + \hat{f}_{0,-l}) \\
a_{0,0} &= \hat{f}_{0,0} \\
b_{m,0} &= \frac{1}{4} i (2\hat{f}_{m,0} - 2\hat{f}_{-m,0}) = \frac{1}{2} i (\hat{f}_{m,0} - \hat{f}_{-m,0}) \\
c_{0,l} &= \frac{1}{4} i (2\hat{f}_{0,l} - 2\hat{f}_{0,-l}) = \frac{1}{2} i (\hat{f}_{0,l} - \hat{f}_{0,-l}).
\end{aligned}$$

Thus,

$$\begin{aligned}
f(u,v) &= a_{0,0} + 2 \sum_{m=1}^{\infty} \left\{ a_{m,0} \cos mu + b_{m,0} \sin mu \right\} + 2 \sum_{l=1}^{\infty} \left\{ a_{0,l} \cos lw + c_{0,l} \sin lw \right\} + \\
& + 4 \sum_{m=1}^{\infty} \sum_{l=1}^{\infty} \left\{ a_{m,l} \cos mu \cos lw + b_{m,l} \sin mu \cos lw + c_{m,l} \cos mu \sin lw + d_{m,l} \sin mu \sin lw \right\}.
\end{aligned}$$

This is the function representation used for the coordinate functions in equations (6) to (8). Further:

$$\begin{aligned}
a_{m,l} &= \frac{1}{4} \cdot \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \left\{ e^{i(mu+lv)} + e^{i(-mu+lv)} + e^{i(mu-lv)} + e^{-i(mu+lv)} \right\}^* dudv = \\
&= \frac{1}{16\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \{ 2 \cos(mu + lv) + 2 \cos(mu - lv) \}^* dudv = \\
&= \frac{1}{8\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \{ \cos mu \cos lv - \sin mu \sin lv + \cos mu \cos lv + \sin mu \sin lv \} dudv = \\
&= \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \cos mu \cos lv dudv
\end{aligned}$$

$$\begin{aligned}
b_{m,l} &= \frac{i}{4} \cdot \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \left\{ e^{i(mu+lv)} - e^{i(-mu+lv)} + e^{i(mu-lv)} - e^{-i(mu+lv)} \right\}^* dudv = \\
&= \frac{i}{16\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \{ 2i \sin(mu + lv) + 2i \sin(mu - lv) \}^* dudv = \\
&= \frac{-i \cdot i}{8\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \{ \sin mu \cos lv + \cos mu \sin lv + \sin mu \cos lv - \cos mu \sin lv \} dudv = \\
&= \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \sin mu \cos lv dudv
\end{aligned}$$

$$\begin{aligned}
c_{m,l} &= \frac{i}{4} \cdot \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \left\{ e^{i(mu+lv)} + e^{i(-mu+lv)} - e^{i(mu-lv)} - e^{-i(mu+lv)} \right\}^* dudv = \\
&= \frac{i}{16\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \{ 2i \sin(mu + lv) + 2i \sin(-mu + lv) \}^* dudv = \\
&= \frac{-i \cdot i}{8\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \{ \sin mu \cos lv + \cos mu \sin lv - \sin mu \cos lv + \cos mu \sin lv \} dudv = \\
&= \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \cos mu \sin lv dudv
\end{aligned}$$

$$\begin{aligned}
d_{m,l} &= \frac{1}{4} \cdot \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \left\{ e^{i(mu+lv)} + e^{i(-mu+lv)} - e^{i(mu-lv)} - e^{-i(mu+lv)} \right\}^* dudv = \\
&= \frac{1}{16\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \left\{ -2\cos(mu+lv) + 2\cos(mu-lv) \right\}^* dudv = \\
&= \frac{1}{8\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \left\{ -\cos mu \cos lv + \sin mu \sin lv + \cos mu \cos lv + \sin mu \sin lv \right\} dudv = \\
&= \frac{1}{4\pi^2} \iint_{\substack{u \in [0, 2\pi) \\ v \in [0, 2\pi)}} f(u, v) \sin mu \sin lv dudv.
\end{aligned}$$

With this definition of $a_{m,l}$, $b_{m,l}$, $c_{m,l}$ and $d_{m,l}$ the definition of scalar product becomes consistent due to the inclusion of the factors $\frac{1}{4}$ in the coefficient definitions, which are responsible for the factors 2 and 4 in the Fourier series. Another definition of scalar product would have caused other factors to appear. By defining scalar product differently depending on the value of m and l all constant factors could have been avoided. For theoretical reasons that would not be satisfactory, but for the task of optimising the surface to an image object there is no difference in optimising a coefficient or optimising a coefficient times a factor.

10. References

1. Rogowska, J., *Overview and Fundamentals of Medical Image Segmentation*, in *Handbook of Medical Imaging*, I. Bankman, Editor. 2000, Elsevier. p. 69-79.
2. Anger, H.O., *Scintillation Camera*. Review of scientific instruments., 1958. **29**(1): p. 27-33.
3. Garkavij, M., et al., ^{177}Lu -[DOTA0,Tyr3] octreotate therapy in patients with disseminated neuroendocrine tumors: Analysis of dosimetry with impact on future therapeutic strategy. *Cancer*, 2010. **116**(S4): p. 1084-1092.
4. Pham, D.L., C. Xu, and J.L. Prince, *Current Methods in Medical Image Segmentation*. Annual Review of Biomedical Engineering, 2000. **2**: p. 315-337.
5. McInerney, T. and D. Terzopoulos, *Deformable Models*, in *Handbook of Medical Imaging*, I. Bankman, Editor. 2000, Elsevier. p. 127-145.
6. Heimann, T. and H.P. Meinzer, *Statistical shape models for 3D medical image segmentation: A review*. Medical Image Analysis, 2009. **13**(4): p. 543-563.
7. Neumann, A. and C. Lorenz, *Statistical shape model based segmentation of medical images*. Computerized Medical Imaging and Graphics, 1998. **22**(2): p. 133-143.
8. Staib, L.H., et al., *Shape Constraints in Deformable Models*, in *Handbook of Medical Imaging*, I. Bankman, Editor. 2000, Elsevier. p. 147-157.
9. Kass, M., A. Witkin, and D. Terzopoulos, *Snakes: Active Contour Models*. International Journal of Computer Vision, 1988. **1**: p. 321-331.
10. Boudraa, A.O. and H. Zaidi, *Image Segmentation Techniques in Nuclear Medicine Imaging*, in *Quantitative Analysis in Nuclear Medicine Imaging* H. Zaidi, Editor. 2006, Springer. p. 308-357.
11. King, M.A., D.T. Long, and A.B. Brill, *SPECT volume quantitation: Influence of spatial resolution, source size and shape, and voxel size* Medical Physics, 1991. **18**(5): p. 1016-1024.
12. Mykkanen, J.M., M. Juhola, and U. Ruotsalainen, *Extracting VOIs from brain PET images*. International Journal of Medical Informatics, 2000. **58-59**: p. 51-57.
13. Floreby, L., *Closed Parametric Boundary Finding in Medical Images*, in *Image Analysis and Enhancement with Applications in Medicine*. 1998, Lund University: Lund. p. 205-254.
14. Staib, L.H. and J.S. Duncan, *Boundary Finding with Parametrically Deformable Models*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1992. **14**: p. 1061-1075.
15. Giardina, C.R. and F.P. Kuhl, *Accuracy of curve approximation by harmonically related vectors with elliptical loci*. Computer Graphics and Image Processing, 1977. **6**(3): p. 277-285.
16. Kuhl, F.P. and C.R. Giardina, *Elliptic Fourier features of a closed contour*. Computer Graphics and Image Processing, 1982. **18**(3): p. 236-258.

17. Székely, G., et al., *Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models*. Medical Image Analysis, 1996. **1**: p. 19-34.
18. Persoon, E. and K.-S. Fu, *Shape Discrimination Using Fourier Descriptors*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1986. **8**: p. 388-397.
19. Staib, L.H. and J.S. Duncan, *Model-Based Deformable Surface Finding for Medical Images*. IEEE Transactions on Medical Imaging, 1996. **15**: p. 720-731.
20. Floreby, L., L. Sörnmo, and K. Sjögreen. *Boundary Finding Using Fourier Surfaces of Increasing Order*. in *Fourteenth International Conference on Pattern recognition*. 1998: IEEE Computer Society.
21. Gao, L., D. Heath, and E. Fishman, *Abdominal Image Segmentation Using Three-dimensional Deformable Models*. Investigative Radiology, 1998. **33**: p. 348-355.
22. Floreby, L., et al. *Deformable Fourier Surfaces for Volume Segmentation in SPECT*. in *Fourteenth International Conference on Pattern Recognition*. 1998: IEEE Computer Society.
23. Brechbühler, C., G. Greig, and O. Kübler, *Parametrization of Closed Surfaces for 3-D Shape Description*. Computer Vision and Image Understanding, 1995. **61**: p. 154-170.
24. Nocedal, J. and S.J. Wright, *Nelder-Mead method*, in *Numerical Optimization*. 2006, Springer. p. 238-240.
25. Chang, Y.-C., *N-Dimension Golden Section Search: Its Variants and Limitations*. 2009 2nd International Conference on Biomedical Engineering and Informatics, 2009: p. 1-6.
26. Mortelmans, L., et al., *A new thresholding method for volume determination by SPECT*. European Journal of Nuclear Medicine, 1986. **12**: p. 284-290.
27. Gonzalez, R.C. and R.E. Woods, *the Hotelling transform*, in *Digital Image Processing*. 1993, Addison-Wesley. p. 148-156.
28. Esser, J.P., et al., *Comparison of [177Lu-DOTA0,Tyr3]octreotate and [177Lu-DOTA0,Tyr3]octreotide: which peptide is preferable for PRRT?* European journal of nuclear medicine and molecular imaging, 2006. **33**(11): p. 1346-1351.
29. Forrer, F., et al., *A comparison of 111In-DOTATOC and 111In-DOTATATE: biodistribution and dosimetry in the same patients with metastatic neuroendocrine tumours*. European journal of nuclear medicine and molecular imaging, 2004. **31**(9): p. 1257-1262.
30. Kwekkeboom, D.J., et al., *[177Lu-DOTA0,Tyr3]octreotate: comparison with [111In-DTPA0]octreotide in patients*. European Journal of Nuclear Medicine, 2001. **28**(9): p. 1319-1325.
31. Segars, W.P., et al., *4D XCAT phantom for multimodality imaging research*. Medical Physics, 2010. **37**: p. 4902-4915.
32. ICRU, *ICRU Report 44: Tissue Substitutes in Radiation Dosimetry and Measurements*. 1989, International Commission on Radiation Units and Measurements: Bethesda, MD.
33. Ljungberg, M. and S.E. Strand, *A Monte Carlo program for the simulation of scintillation camera characteristics*. Computer Methods and Programs in Biomedicine, 1989. **29**(4): p. 257-272.
34. Gonzalez, R.C. and R.E. Woods, *Morphology*, in *Digital Image Processing*. 1993, Addison-Wesley. p. 518-560.
35. *IDL Version 6.4*. 2007, ITT Visual Information Solutions: Boulder.
36. NIST, *Weighted variance*, in *DATAPLOT Reference Manual*. 1996. p. 68.