Competing level sets for segmentation of the thalamus

Cecilia Richero Wilson

22nd June 2004

Contents

List of Figures

Abstract

Alterations of the thalamus are implicated in a great number of brain diseases such as Parkinson's disease or schizophrenia, among others. Some of those diseases must be surgically treated, others need deep-brain electrical stimulation or similar treatments. A need arouses to be able to clearly see the thalamic region and its cytoarchitecture.

Diffusion Tensor Magnetic Resonance Imaging provides a non-invasive technique to observe molecular diffusion in brain tissues. Making use of this technique, we have developed a method to segment the thalamus as a whole by means of active contours. Several brain maps can be obtained from DT-MRI, such as the Fractional Anisotropy map or the Diffusion Coefficients. Regarding such maps, statistical data is extracted for the different regions of the brain, and the probability of a voxel being in a certain region can be determined. Then, the front propagation, implemented with level sets methods, is based on those statistics and a curvature term which helps provide smoothness. Coupling surfaces are used to assure interaction between the different curves and help define each region's boundaries.

The results show that the method can be used to segment the thalamic region together with the subthalamus. This will be useful in presurgical planning of the mentioned diseases and future generation of thalamic and subthalamic nuclei atlases.

Chapter 1

Introduction

The thalamus can be considered the central relay station for nerve impulses in the brain. It acts as the gatekeeper for anything that wants to get up to the cortex. Axons from every sensory system (except olfaction) synapse here as the last site before the information reaches the cerebral cortex. The thalamus sends to it the information received from diverse brain regions.

Anatomically, it is a large, dual lobed mass made up mainly of grey matter cells, located in the center of the brain. Each lobe measures approximately 4 centimetres. It is part of the diencephalon and we find it deep inside the cerebral hemispheres and next to the ventricles. The third ventricle shares its lateral walls with the thalamus, as it separates the two thalamic bodies. These are connected by a piece of thalamic white matter tissue called *massa intermedia*. This is not the only white matter in the thalamus, since several fibers pass through and around it. The thalamus, situated at the top of the brainstem, superior to the hypothalamus, communicates sensory, motor and associative brain regions.

Some parts of the thalamus play a major role in the regulation of consciousness, alertness, arousal, and possibly attention, which partially explains why the thalamus is considered to be part of the limbic system.

The thalamic cytoarchitecture is divided into different clusters, a heterogeneous group of nuclei, each with a specific function. The thalamic nuclei have traditionally been studied with histological methods and their number varies depending on the method used. However most studies identify 14 major nuclei, some of them being subdivided.

The importance of generating an exact map of the thalamus comes from the fact that thalamic changes are involved in a large number of diseases, such as schizophrenia, Parkinson's disease or multiple sclerosis, among others. The thalamic atlases are used to target the pertinent nucleus in presurgical planning of these diseases. However there is a high inter-subject variability in the location and size of the thalamic nuclei and therefore generic thalamic atlases may be highly inaccurate.

Resolving thalamic nuclei by noninvasive imaging would be a great step

forward as it would enable, among other things, more accurate neurosurgical planning for the diseases mentioned above, giving the possibility of generating a personal thalamic atlas for each patient. Unfortunately current imaging methods such as CT[1] and conventional Magnetic Resonance Imaging do not provide the necessary image contrast to differentiate the nuclei so radiological identification of individual thalamic nuclei is not currently possible.

M.R. Wiegell *et al.* [18] have shown how DT-MRI can differentiate the principal thalamic nuclei, noninvasively, basing on the characteristic fiber orientation, which is assumed to stay the same all along one certain nucleus and varies from one nucleus to another.

The present work aims to be a first step in the process of resolving the thalamic nuclei. It's objective is to accomplish the segmentation of the thalamus as a whole by the process through which the nuclei will be later segmented.

The recovery of shapes of the human body is more difficult compared to other imaging fields. This is mainly due to factors as the large variability in shapes, complexity of medical structures and several kinds of artifacts. However, we will try to overcome the mentioned complications and propose a method based on the class of deformable models known as *level sets* and *geodesic active contours.*

First we will present the concept of Diffusion Tensor and how diffusion is present in the different parts of the brain. We will also introduce the segmentation theory of Geodesic Active Contours and the front propagation implemented with Level Set Method and Geometric Flows. That is done in chapter two.

In the third chapter the initial objectives of this work are presented, along with the method developed. We will show how to use the statistics in the different images in order to obtain the speed which will help propagate the surface. And how the Level Sets Method should be applied to this case.

The fourth chapter is dedicated to explaining the actual implementation. We present our signed distance function, explain how the probabilities for the different regions are calculated and show how the coupling forces are implemented for this case. Details are also given on the data images used and the treatment applied to some of them before being used.

The results are presented in chapter 5, both for synthetic and for real images.

The last chapter is dedicated to the conclusions that can be extracted from this work and possible future work on the subject.

[1] CT: Computerized Tomography

Chapter 2

Background Theory

In this chapter we will introduce the concept of Diffusion Tensor and the mathematics in it. We will also explain how diffusion is present in the different parts of the brain, how it can help classify them and how DT-MRI[1] can be used. The theories of Geodesic Active Contours and Level Sets are introduced as they will be used further on.

2.1 Diffusion Tensor Imaging

The diffusion tensor (DT) provides information about the intensity of the water diffusion in any given direction, at a certain point, when applying a magnetic field [7]. Water diffusion is described by what is called Brownian motion. The Brownian motion is an irregular motion exhibited by minute particles of matter when suspended in a fluid. This effect, independent of all external factors, is ascribed to the thermal motion of the molecules of the fluid.

Diffusion in tissue is dependent on the tissue's characteristics. Because DT becomes highly anisotropic in areas of compact nerve fiber organization, it provides an indirect way of fiber tract identification and tissue classification.

2.1.1 Diffusion in the brain

Water diffusion in the brain is highly affected by the tissues' cellular organization. In white matter, the values for diffusion can be extremely variable, but its direction is mainly the same as the fibers. This means that it is highly anisotropic.

On the other hand, in regions with fluid, like CSF[2] filled ventricles, diffusion is mainly spherical.

The values of diffusion in grey matter stand in the middle, as this region shows nearly no anisotropy but diffusion is not spherical either. Moreover,

[1] Diffusion Tensor Magnetic Resonance Imaging
[2] CSF: Cerebro-Spinal Fluid

some regions made of grey matter are difficult to classify since they also have white matter or other kinds of tissue inside them, like fibers going through them. That is the case of the thalamic region, which concerns us.

The value and type of diffusion directly depends on the relative orientation of the fibers and the gradient direction [6]. For example, myelin sheath is one of the main components restricting water mobility: diffusion is dependent on the orientation of myelin fiber tract. The measured DT in it becomes highly anisotropic and orientated in the same way as fiber organization. When measuring molecular displacements perpendicular to the direction of the myelin fiber tracts, diffusion is significantly decreased. On the other hand, the diffusion measured parallel to myelin fiber direction can be about three times larger than the one measured in the perpendicular direction. Hence, diffusion helps identify fiber tracts and distinguish different types of tissue.

Density of fibers, degree of myelinisation, presence of fluid and the directional similarity of fibers in a voxel are some of the factors which would affect the shape of the apparent DT in different parts of the brain.

2.1.2 DT-MRI

Through DT-Magnetic Resonance Imaging (DT-MRI), the averaged self-diffusion tensor (or effective diffusion tensor) in each voxel of a 3D image can be measured. While conventional MRI cannot provide information on fiber orientation, DT-MRI allows to see the orientation of the fibers, and can therefore be used to make a classification of the different tissues depending on the type of diffusion in them. It is therefore, a non-invasive, indirect way of fiber tract and tissue classification.

Diffusion can be measured by diffusion weighted MRI along at least six independent axis. That means we require six different diffusion weighted images. A normalizing image without diffusion weighting is also needed.

The diffusion tensor in each voxel is obtained to create a 3D field of diffusion tensors. Anisotropic diffusion can also be measured. Modeled by an anisotropic Gaussian that can be parametrized by the diffusion tensor in each voxel [3] it can provide a 3D map of Fractional Anisotropy values.

2.1.3 Mathematics in The Diffusion Tensor

Mathematically, the DT is a symmetric, second order tensor, represented by a 3x3 semi-positive definite matrix. It represents the diffusion in each direction of a voxel and enables us to obtain the value for the diffusion in any direction, as follows:

$$D(\underline{\hat{x}}) = \underline{\hat{x}}^T D \underline{\hat{x}}, \tag{2.1}$$

where $\underline{\hat{x}}$ is the unit vector representing the direction for which we want to know the diffusion.

$D(\underline{\hat{x}})$ is proportional to the rate of transfer of water molecules across a plane perpendicular to $\underline{\hat{x}}$.

We can also calculate the mean diffusivity, which is:

$$MeanD = \frac{Trace(D)}{3} = \frac{D_{1,1} + D_{2,2} + D_{3,3}}{3} \tag{2.2}$$

The DT can be interpreted by its eigenvectors and eigenvalues, obtained by diagonalizing the DT matrix. Each eigenvalue is a quantitative measure of the diffusion in its corresponding direction. The eigenvector corresponding to the highest eigenvalue describes the direction of the principal diffusion.

In 3D, the DT can be thought of as an ellipsoid with the principal axes having a length proportional to the eigenvalues (to $\sqrt{\lambda}$) and coinciding with the directions of the eigenvectors.

Having calculated the diffusion tensor's eigenvalues and eigenvectors, the type of diffusion can be classified after the geometry of the diffusion ellipsoid as follows:

1) Linear, where: $\lambda_1 \gg \lambda_2 \simeq \lambda_3$. Diffusion is present only along one direction. Therefore, it is anisotropic diffusion.

2) Planar, where: $\lambda_1 \simeq \lambda_2 \gg \lambda_3$. Diffusion of this type is restricted to planes.

3) Spherical, where: $\lambda_1 \simeq \lambda_2 \simeq \lambda_3$. This means diffusion is similar in all directions, it is therefore, isotropic.

Through the eigenvalues we can also calculate the diffusion coefficients, which give a measure of how close the diffusion tensor is to the generic cases. The coefficients are:

- Linear case:

$$c_l = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3} \tag{2.3}$$

- Planar case:

$$c_p = \frac{2(\lambda_2 - \lambda_3)}{\lambda_1 + \lambda_2 + \lambda_3} \tag{2.4}$$

- Spherical case:

$$c_s = \frac{3\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3} \tag{2.5}$$

There are several measures of diffusion anisotropy. They measure the directional bias of a DT. They have higher values for more anisotropic, less spherical, diffusion tensors.

One of the most popular of this measures is the Fractional Anisotropy. The Fractional Anisotropy (FA) index measures the fraction of the magnitude of the diffusion that can be ascribed to anisotropic diffusion [6]. It is quantitative and dimensionless, and is calculated as follows:

$$FA = \sqrt{\frac{3}{2}} \sqrt{\frac{(\lambda_1 - \langle\lambda\rangle)^2 + (\lambda_2 - \langle\lambda\rangle)^2 + (\lambda_3 - \langle\lambda\rangle)^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} \tag{2.6}$$

where $\langle\lambda\rangle$ is the average of the eigenvalues.

2.2 Geodesic Active Contours (*Snakes*)

The geodesic active contours approach for object segmentation found in [5] allows to connect the classical "snakes" approach, based on energy minimization, with geometric active contours, based on the theory of curve evolution.

The classical "snakes" approach is based on deforming an initial contour C_0 towards the boundary of the object. It has been used so far for edge and curve detection, segmentation and shape modelling, among others.

In order to get the deformation of the curve we must minimize a functional designed so that its minimum is achieved on the boundaries of the object. This is the energy functional and it is basically composed of two terms. One of them controls the smoothness of the curve whereas the other attracts the curve towards the boundary. The drawback of this model is that it is not able to handle topology changes of the evolving contour, so the topology of the final curve will be the same as the one of C_0. This is a problem when the number of objects to be simultaneously detected is unknown.

In the geometric active contours model, based on curve evolution and geometric flows, the curve propagates by means of a speed. This speed is also composed of two terms. One of them is related to the regularity of the curve. The other term shrinks or expands the curve towards the boundary. The model is given by a geometric flow (in PDE form) and it is not dependent on energy minimization. If implemented using the level-sets algorithm [12], explained in Section 2.3, it allows the handling of topology changes. Hence, it can be used for detection of several objects simultaneously without knowing their exact number a priori.

Caselles, Kimmel and Sapiro [5] presented a way of object boundaries detection based on active contours evolving in time according to intrinsic geometric measures of the image. The evolving contours are designed to split and merge, so detection of several objects, and both interior and exterior boundaries, can be possible. They also proved a particular case of the classical energy snakes model to be equivalent to finding a geodesic curve in a Riemannian space. This means that, under certain conditions, object boundary detection can be considered equivalent to finding a curve of minimal weighted length in a new metric.

Assuming that this geodesic active contour is represented as the zero level-set of a 3D function, the geodesic curve computation is reduced to a geometric flow similar to the one obtained by curve evolution approaches. The main difference is that this flow includes a new speed term which is based on image information, that improves the previous models. This way, geodesic active contours allow to track boundaries with high variation in their gradient and give a unique and stable solution.

To get to this solution, we first have to show the connection between energy based active contours (*snakes*) and the computation of geodesics or minimal distance curves in a Riemannian space.

We will start describing the classical energy based snakes.

Let $C(q) : [0, 1] \to R^2$ be a parametrized planar curve and $I : [0, a] \times [0, b] \to R^+$ be a given image in which we want to detect the object boundaries. The

classical snakes approach (Kass et al., 1988) associates the curve C with an energy $E(C)$, composed by two terms for internal energy and another term for external energy:

$$E(C) = \alpha \int |C'(q)|^2 \, dq + \beta \int |C''(q)|^2 \, dq - \lambda \int |\nabla I(C(q))| \, dq \tag{2.7}$$

The first two terms are the ones for internal energy and control the smoothness of the contours to be detected. The third term is the one corresponding to external energy and it is responsible for attracting the contour towards the object in the image (this term is dependent on the image). However, it has been previously proven that curve smoothing can also be obtained with $\beta = 0$. So that will reduce (2.7) to:

$$E(C) = \alpha \int |C'(q)|^2 \, dq - \lambda \int |\nabla I(C(q))| \, dq \tag{2.8}$$

To perform segmentation, we need to find the curve C that minimizes the energy $E(C)$.The problem is that, having more than one object in the image, it is not possible to detect all of them. In other words, the classical snakes approach cannot directly deal with topology changes.

The solution to this problem is given by a geodesic curve in a Riemannian space induced from the image I. A geodesic curve is a minimal weighted distance path between given points. Following Caselles, Kimmel and Sapiro's studies, we find that the problem of minimizing (2.8) can be transformed into a problem of geodesic computation in a Riemannian space, which is minimizing:

$$\int g\left(|\nabla I\left(C(q)\right)|\right) |C'(q)| \, dq \tag{2.9}$$

In a different Riemannian space, this gives a new length definition:

$$L_R = \int g\left(|\nabla I\left(C(q)\right)|\right) |C'(q)| \, dq \tag{2.10}$$

We can therefore see that, when trying to detect an object, we are interested not only in finding the path of minimal classical length, but the one that minimizes a new length definition which bears in consideration the image characteristics.

The minimization of (2.10) is done by the steepest descent method, so we will search for the gradient descent direction in the equation. Calculating Euler-Lagrange of this, we find the curve evolution equation:

$$\frac{\partial C(t)}{\partial t} = g(I) k \vec{N} - \left(\nabla g \cdot \vec{N}\right) \vec{N}, \tag{2.11}$$

being C_t the flow, k the Euclidean curvature and $\vec{N}$ the unit inward normal. This equation is represented using the level-sets approach and shows how each point in the active contour C should evolve to minimize the length L_R. The detected object will then be given by the steady state solution of (2.11), which is $C_t = 0$.

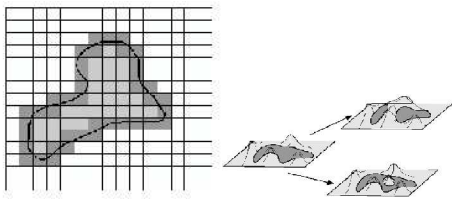

Figure 2.1: A curve can be thought of as a level set of some discretely sampled 2D function or as a level set of a surface

2.3 Level Sets and Geometric Flows

The Level Set Method Let's assume we have an image in which an interface separates one region from another, and there is a speed for each point of the interface. The level set method can be used to know how to evolve the surface.

Introduced by Osher and Sethian, the level set approach takes the original curve and builds it into a surface, called the level set function.

Given a function $u : [0, a] \times [0, b] \rightarrow \Re$. Assuming that C is a level-set of that function, then C coincides with the set of points $u =$constant. Therefore, u is an implicit representation of the curve C and has one more dimension.

This function accepts any point in the plane as an input, and hands out its distance. Instead of moving the initial front, which could result in strange shapes, we calculate the position in the surface of any given point. The built surface intersects the x-y plane exactly at the location of the initial curve. The initial curve is called the zero level set, because it contains all the points with distance zero.

One of the main advantages of the level sets algorithm is that this representation is not dependent on topology since different topologies of the zero level-set do not imply different topologies of u. So it easily handles topology changes of the given surface.

The level set methods are very useful in problems where the speed of the curve can be either positive or negative, which means that the front can either propagate or shrink.

If we take the evolution of the planar curve according to:

$$C_t = \beta \vec{N}$$

for any given function β, then the function u will deform as: [11]

$$u_t = \beta \left| \nabla u \right| .$$

By considering C as the zero level-set of the function u, all topological changes of $C(t)$ are automatically handled and an accurate and stable solution can be achieved.

Even though some parts of the level set theory in two dimensions, such as the property of a curve shrinking to a point under a curvature flow, cannot be

extended to the 3D case, the greater part of it remains valid and gives good results for segmentation of 3D objects. So the level set approach will take the original curve in 2D and build it into a 3D surface, which will be the level set function.

A curvature flow is a curve or surface that evolves at each point along the normal with a speed depending on the curvature at that point. This evolution leads to a smoothing of the curves or surfaces, which helps to eliminate the effects of noise. When trying to segment an image, the evolution of the curve or surface, if geometrical flows are used, will be dependent on external properties, determined by the image features. An example is segmentation of grey scale images using a classical speed function based on the gradient of the images. The speed tends to zero when the surface approaches an edge [5].

We will describe briefly here how geometric flows work and their formulation as level sets. A curvature flow for a 3D closed surface can be described as:

$$\frac{\partial S}{\partial t} = G\vec{N} \tag{2.12}$$

$\partial S/\partial t$ is the derivative of the surface through time, and represents the curvature flow. G is an intrinsic speed dependent on the curvature of the surface, S is the surface and $\vec{N}$ is the normal to the surface.

This time dependent PDE can be solved using the level set method.

In our case, the function of higher dimension is the signed distance function, $\phi(t)$, of the evolving surface. Being so, the evolution of the zero level set coincides with the evolution of $S(t)$.

What we have is then:

$$\phi(S(t)) = constant. \tag{2.13}$$

And the evolution of the signed distance function is therefore:

$$\frac{d\phi}{dt} = \phi_t + S_t \nabla\phi = 0, \tag{2.14}$$

where $S_t = \frac{dS}{dt}$. If we insert equation (2.12) in (2.14) and consider:

$$\vec{N} = \frac{\nabla\phi}{|\nabla\phi|},$$

we obtain the following flow:

$$\phi_t = |\nabla\phi| G = |\nabla\phi| \, div\left(\frac{\nabla\phi}{|\nabla\phi|}\right) \tag{2.15}$$

This is the level set formulation of the mean curvature flow.

An extension of the mean curvature flow is the geodesic active contours [5]. These are used for image segmentation by trying to find an harmonic map which

minimizes a functional. In the 3D case, the harmonic map is a minimal surface and the functional to minimize is defined as a weighted area:

$$A = \int \int g(I) da,$$

where *g(I)* is an image dependent function. In image segmentation, the aim is to locate the object boundaries, so this function is usually defined as one which approaches zero as the image gradient increases. For example, a commonly used function is $g(I) = 1/(1 + |\nabla I|)$.

To minimize the weighted area, the procedure described above is followed. In this case, it leads to the following minimizing flow:

$$\frac{\partial S}{\partial t} = \left(gG - \nabla g \cdot \vec{N} \right) \vec{N} \tag{2.16}$$

which is equivalent to equation (2.11).

On the right side of the equation, the first term is the mean curvature flow multiplied by the image dependent function. This term will stop the front propagation at edges. The second term forms an attraction valley that pushes the front towards the edge when the first is close enough to the second. If a balloon force is added to the propagation speed, this is especially useful. The balloon force leads to a constant speed and is added to increase the convergence speed. This can sometimes have the drawback of the front being pushed too far and pass the edge. It is in this situation where the second term starts acting to push the front back to the edge.

The level set implementation for the above flow will now be:

$$\phi_t = g(I) \left| \nabla \phi \right| div \left(\frac{\nabla \phi}{|\nabla \phi|} \right) + \nabla g(I) \cdot \nabla \phi. \tag{2.17}$$

2.4 Geodesic Active Regions and their Statistics

Paragios presented a unified approach for image segmentation [13] which incorporates boundary and region information under a curve-based minimization framework, exploited directly from the Geodesic Active Region model.

It consists on segmenting an image into different regions by calculating the probability of every intensity value in the image of being in each region.

On a first stage, an analysis of the image histogram (where the global intensity properties of the image are reflected) is made. This analysis is based on the Minimum Description Length criterion and the Maximum Likelihood Principle. It associates a Gaussian component, and therefore a Gaussian probability, to each region of the image.

Let I be the input image and *H(I)* its observed density function (histogram). Considering a partition of the image into N non-overlapping regions, let $\{R_i : i \in [1, N]\}$ be the regions and $\{\partial R_i : i \in [1, N]\}$ be the region boundaries.

The key hypothesis that is made to perform segmentation relies on the fact that the image is composed of homogeneous regions. Hence, the intensity properties of a given region can be determined using a Gaussian distribution.

Let $p(.)$ be the probability density function with respect to the intensity space of the image. Assuming that this probability density function is homogeneous, then an intensity value x is derived by selecting a component k with *a priori* probability P_k and then selecting this value according to the distribution of this element $p_k()$. This hypothesis leads to a mixture model of Gaussian elements:

$$p(x) = \sum P_k p_k(x) \tag{2.18}$$

$$p_k(x) = \frac{1}{\sqrt{2\Pi}\sigma_k} e^{-\frac{(x-\mu_k)^2}{2\sigma_k^2}} \tag{2.19}$$

where k corresponds to one of the regions, R_i.

According to the Geodesic Active Region model, the best segmentation map is determined using a set of regular curves, where each curve:

1) Is attracted by the boundaries of a specific region,

2) Defines an interior region with maximum joint segmentation probability given the observed intensities.

This map is obtained by minimizing an Energy function using a gradient descent method. A system of Euler-Lagrange motion equations with respect to different curves, as in Section 2.2, is obtained (one equation for each region).

The system is given by:

$$\begin{array}{l} \forall i \in [1, N], \\ \frac{\partial}{\partial t}\partial R_i = \alpha \left[g\left(p_i\left(I\left(\partial R_i\right)\right), \sigma_R\right) - g\left(p_{k_i}\left(I\left(\partial R_i\right)\right), \sigma_R\right)\right] N_i\left(\partial R_i\right) + \\ (1-\alpha)\left(g\left(p_{B,i}\left(\partial R_i\right), \sigma_B\right) K_i\left(\partial R_i\right) + \nabla g\left(p_{B,i}\left(\partial R_i\right), \sigma_B\right) N_i\left(\partial R_i\right)\right) N_i\left(\partial R_i\right) \end{array} \tag{2.20}$$

where K_i is the Euclidean curvature (respect to the normal) with respect to the curve ∂R_i, g is an image dependent function and $p_i(x)$are the probabilities in equation (2.19).

Each equation is composed of a Region-based force (which is the first term) and a Boundary-based force (second and third terms), both acting in the direction of the inward normal.

The region-based force aims at moving the curve towards the direction that maximizes the *a posteriori* segmentation probability. The boundary-based force aims at shrinking the curve towards the region boundaries constrained by the curvature.

The whole system of equations relies on a multi-phase curve propagation since several curves are propagated simultaneously, one for each region. The interaction between the different curves is obtained thanks to the region-based term.

Chapter 3

Methodology

The following pages are dedicated to explaining the purpose of this work, the approach followed and the method developed. We will show how the theory presented in chapter two is used and how we obtain the equations which will determine the evolution of our curves.

3.1 Objectives

The principal aim of this paper is to find a way of segmenting the thalamus as a whole from the rest of the brain, using coupling level sets. The segmentation will be based on the diffusion properties of the different parts of the brain. This means that the different values and types of diffusion in different areas of the brain will be considered.

In order to do this, we propose a method which consists on the propagation of several surfaces, each in a different part of the brain and on different images but still coupled. The propagation of the surfaces responds to the theories of level-sets and geodesic active contours.

3.2 Initial Approach

The starting point of this work was an implementation of image segmentation with a single surface evolving in the image. The evolution of the surface was based on the region-based term, calculated from a tensor similarity measure, the Tensor Scalar Product. The data images were also diffusion tensor MRI, and to calculate the speed of each voxel, the similarity was looked for in neighbour voxels.

The speed term being proportional to a TSP means that the more similar the tensors are in two neighbour voxels, the higher the speed with which the surface evolves. On the other hand, the more similar two tensors are, more probable it is that they lie in the same region. This is where the connection between

the similarity measure and calculating the speed term from the probabilities is found.

3.3 Geodesic Active Contours

From the point of view of image segmentation, in order to detect the object boundaries in the image, we will use the geodesic active contours theory together with level sets implementation. Following the theory explained in Section 2.2 we will initialize u to be the signed distance function. Solving the geodesic problem is then equivalent to searching for the steady state solution ($\frac{\partial u}{\partial t} = 0$) of the following evolution equation: (using the curve evolution equation (2.11) together with u)

$$\frac{\partial u}{\partial t} = g(I)\,|\nabla u|\,k + \nabla g(I) \cdot \nabla u \tag{3.1}$$

This is the main part of the active contour model we are using. Then, considering boundary detection, the model is developed further and another term is added. The resulting model proposed in [5] is:

$$\frac{\partial u}{\partial t} = g(I)(c+k)\,|\nabla u| + \nabla g(I) \cdot \nabla u$$

So the curve evolution equation in (2.11) will now be:

$$C_t = g(I)(c+k)\vec{N} - \left(\nabla g(I) \cdot \vec{N}\right)\vec{N}. \tag{3.2}$$

In this equation:

i) c is a positive real constant which, together with $\vec{N}$, gives a constant velocity. This term introduces a balloon force that pushes the curve inwards (or outward) with a constant speed.

ii) k is the curvature. The product $k\vec{N}$ is the regularization curvature flow, which provides smoothness.

iii) *g(I)* is a stopping function which gives the external image dependent force. Its aim is to stop the evolving curve when it arrives to the objects boundaries. For that reason, it should approximate to zero when the curve is getting close to an edge.

iv) $\nabla g(I) \cdot \vec{N}$ pushes the curve towards the boundaries of the objects, as $\nabla g(I)$points towards the middle of the boundaries. It is necessary to restrict the g values so that the propagating curve is guaranteed to stop, because real images don't contain ideal edges.

The term corresponding to $(c+k)$ acts as the internal force in the classical energy based snakes model.

In the present work, we will use a model partially inspired by the one explained above.

Equation (3.2) can be seen as a speed evolving in the direction of the inward normal, $\vec{N}$. Then, we have:

$$C_t = \mathcal{F}\vec{N} \tag{3.3}$$

as the evolving curve.

In the model we propose, the speed will consist of two terms. One corresponding to the curvature and the other one, region-based, dependent on the image. So, in (3.5), the speed will be:

$$\mathcal{F} = v + k, \tag{3.4}$$

where k is the curvature term and v is the region-based term.

We will then need a force to stop the surfaces evolving when they come across each other, the coupling force. When the curves meet next to a boundary, they should push each other so both of them end up defining the boundary. Then, our evolving level sets will be:

$$C_t = (\mathcal{F} + \mathcal{H})\vec{N}, \tag{3.5}$$

being H the coupling force.

3.4 Level Sets Implementation

Implementing the motion equations, which define the evolution of the curves, with the level sets method has many advantages. The level sets theory provides stable numerical approximations for the normal and the curvature, ability of handling topology changes in the curvature (splitting, merging) and stable implementation schemes.

Using this implementation and considering the signed distance function ϕ as our function of higher dimension, the system of equations in (2.20), which represents the evolution of the different curves, is transformed into:

$$\begin{array}{l} \forall i \in [1, N], \\ \frac{\partial}{\partial t}\phi_i(s) = \alpha\left[g\left(p_i\left(I(s)\right), \sigma_R\right) - g\left(p_{k_i}\left(I(s)\right), \sigma_R\right)\right]\left|\nabla\phi_i(s)\right| + \\ (1-\alpha)\left(g\left(p_{B,i}(s), \sigma_B\right) K_i(s)\left|\nabla\phi_i(s)\right| + \nabla g\left(p_{B,i}(s), \sigma_B\right) \cdot \nabla\phi_i(s)\right) \end{array} \tag{3.6}$$

The first term is the region-based force and the two following ones correspond to the boundary-based force, being K_i the Euclidean curvature respect to the normal.

This system represents the coupled surfaces evolution. However, with our speed being (3.4) and introducing the coupling forces, this system will be modified. In the next section we explain how.

3.5 Image Segmentation

3.5.1 Regions' Statistics

As seen before, our segmentation of the brain images is based on associating a Gaussian component to every region in the image and then finding the probabilities of one certain intensity value in a voxel, belonging to a certain region.

That way, we will find the region-based force for every curve. This term, which is actually the propagation speed, aims at propagating the curve towards the most probable region and will also provide the interaction between different curves.

Following [13], let's consider the propagation of the curve ∂R_i (region R_i boundaries). The region-based term will be proportional to:

$$-\alpha \log \left[\frac{p_i\left(I\left(s\right)\right)}{p_j\left(I\left(s\right)\right)} \right] \tag{3.7}$$

where $p_i\left(I\left(s\right)\right)$ is the intensity probability density function followed by region R_i. It is therefore the Gaussian distribution that determines the probability of the intensity value corresponding to voxel s, of being in region R_i.

Using this function we obtain that when voxel s in effect, belongs to region R_i, then:

$$p_i\left(I\left(s\right)\right) > p_j\left(I\left(s\right)\right) \Rightarrow \frac{p_i\left(I\left(s\right)\right)}{p_j\left(I\left(s\right)\right)} > 1 \Rightarrow -\alpha \log \left[\frac{p_i\left(I\left(s\right)\right)}{p_j\left(I\left(s\right)\right)} \right] < 0$$

and so, the speed term is negative and aims at expanding the curve.

If, on the contrary, the voxel s doesn't belong to R_i, then:

$$p_i\left(I\left(s\right)\right) < p_j\left(I\left(s\right)\right) \Rightarrow \frac{p_i\left(I\left(s\right)\right)}{p_j\left(I\left(s\right)\right)} < 1 \Rightarrow -\alpha \log \left[\frac{p_i\left(I\left(s\right)\right)}{p_j\left(I\left(s\right)\right)} \right] > 0$$

and therefore the speed is positive and the force is applied to shrink the curve, so that voxel s can be attributed to another region.

In our case, the intensity values correspond to the fractional anisotropy values or some of the diffusion coefficients values in the voxels.

3.5.2 Coupling Forces

When propagating several curves, the overlapping between some of them is almost inevitable. When that occurs, it means that a voxel has been initially attributed to two different regions. This is an undesired situation and a constraint to solve or avoid it from the beginning has to be applied. This is done by adding an artificial force in the direction of the normal, to the corresponding level set motion equations. The force will penalize voxels which have been attributed to more than one region. And if necessary, also those voxels which haven't been labeled yet (which don't yet belong to a region).

Inspired by Paragios and Deriche [14], our coupling forces for a given voxel s and a region i, are given by:

$$\sum_{j \in [1,\mathbb{N}]} H_i\left(i, \phi_j\left(s\right)\right) \left|\nabla \phi_i\left(s\right)\right| \tag{3.8}$$

where the function $H_i\left(,\phi\left(\right)\right)$ is:

$$H_i\left(m, \phi_n\left(s\right)\right) = \begin{cases} 0, & if\ m = i \\ -sign\left(\phi_j\left(s\right)\right), & if\ m \neq i \end{cases}$$

Analyzing the force added, we see it can have two possible effects. The first one is expanding the corresponding curve. If the voxel hasn't yet been attributed to any region, the new force is negative and helps region R_i to expand, occupying the given voxel. The other effect is the shrinking of the curve. If the voxel has been attributed both to region R_i and to some other region R_k then the force will be positive and aim at shrinking region R_i so that overlapping is avoided and voxel s ends up belonging to just one region.

However, the function presented above, which defines the coupling force, presents some problems. To start with, it penalizes the non-attributed voxels in the same way as the ones attributed to more than one region. Another problem is that the function is not continuous and that may create stability problems during the level sets evolution.

To achieve a more suitable coupling function, we should consider two more properties for it. The first one is the fact that, when a voxel is already attributed to a region j and it is far away from that region's boundaries, the evolution of the level set $\phi_i\left()\right.$ (being R_i another region) should be discouraged to include that voxel in R_i. The second one is that a certain overlapping between two neighbour regions should be tolerated for voxels inside a region which stand very close to its boundaries.

For that purposes, we define the following function:

$$H_a\left(x\right) = -\begin{cases} +1, & if\ x > a \\ -1 & if\ x < -a \\ \frac{1}{\tan(1)} \tan\left(x/a\right), & if\ \left|x\right| \leq a \end{cases}$$

which will be the basis of the coupling force, defined as:

$$H_i\left(j, \phi_j\left(s\right)\right) = \begin{cases} 0, & if\ j = i \\ H_a\left(\phi_j\left(s\right)\right), & if\ j \neq i\ and\ \phi_j\left(s\right) \leq 0 \\ \frac{1}{N-1} H_a\left(\phi_j\left(s\right)\right), & if\ j \neq i\ and\ \left[\cap_{\{k=1,k\neq i\}}^{N} \phi_k\left(s\right) > 0\right] \end{cases}$$

The result of introducing this coupling force is that, considering a voxel s and a level set function $\phi_i\left()\right.$, if the voxel is already belonging to another region R_j, then $\phi_j\left()\right. \leq 0$, and the coupling force will be positive, so it will have a shrinking effect, proportional to the distance of the voxel to the boundaries of R_j:

$$H_i\left(j,\phi_j\left(s\right)\right)=\left\{\begin{array}{ll}+1, & if\ x<-a\\ -\frac{1}{\tan(1)}\tan\left(x/a\right), & if\ \ |x|\leq a\end{array}\right.$$

On the other hand, if the voxel is not attributed to any region we will find:

$$\phi_j\left(s\right)\geq 0\Rightarrow H_i\left(j,\phi_j\left(s\right)\right)=\frac{1}{N-1}H_a\left(\phi_j\left(s\right)\right)$$

and therefore:

$$H_i\left(j,\phi_j\left(s\right)\right)=\frac{1}{N-1}\left\{\begin{array}{ll}+1, & if\ x>a\\ -\frac{1}{\tan(1)}\tan\left(x/a\right), & if\ \ |x|\leq a\end{array}\right.$$

This force also allows the overlapping when the curves are on top of the real region boundaries.

3.5.3 Motion Equations

The system of equations representing the evolution of the coupled surfaces in (3.6) has now been modified due to our concept of the speed and to the coupling forces that we want to introduce.

At this stage, the system of level set motion equations is given by:
$\forall i\in[1,N]$,

$$\begin{array}{l}\frac{\partial\phi_i(s)}{\partial t}=\beta\sum_{j\in[1,\mathbb{N}]}H_i\left(i,\phi_j\left(s\right)\right)\left|\nabla\phi_i\left(s\right)\right|-\\ \gamma[\log\left(\frac{p_i(I(s))}{p_{k_i}(I(s))}\right)]\left|\nabla\phi_i\left(s\right)\right|+\\ \delta K_i\left(s\right)\left|\nabla\phi_i\left(s\right)\right|\end{array}\qquad(3.9)$$

where β, γ, δ are positive constants and $\beta+\gamma+\delta=1$. The first term is the coupling force. Looking at equation (3.6) we can see it corresponds to the boundary term which pushes the curve towards the edges when necessary. The second one is the region-based term (v in (3.4)) and the third one corresponds to the function which calculates the curvature in 3D (k in (3.4)).

Chapter 4

Implementation

In this chapter we will explain how the whole procedure to segment the thalamus was implemented. Our signed distance function is presented, along with explanations on how we calculate the probabilities for every region and all the forces involved in the evolution of the curves. Details are given on the data images used, and the treatment applied to some of them when necessary.

The method for the segmentation of the thalamus has been implemented in Matlab 6.1.

4.1 Algorithm

When starting to implement the explained method, we first developed the following algorithm, on which the implementation should be based.

1. The first thing to do is choose in which images (fractional anisotropy, or the different diffusion coefficients) are we going to look for each part of the brain and into how many regions do we want to divide each image. We then calculate the probabilities for each region.

2. Once in the program, we define the initial surfaces from where the propagation will start. They will evolve at a certain speed until they find the boundaries of the different regions.

3. Then, for every surface:

- The signed distance function is calculated (that is our zero level set) and a feature map is created.

- The curvature is then calculated in the whole feature map.

- We define a band around the contour of the surface to evolve which will comprise all the voxels around the contour where phi is smaller than 3. (It is just to give it a band so that it doesn't take a lot of time to compute the whole map unnecessarily).

- The speed is calculated for every voxel in the band:
  - First, we obtain the value in side that voxel.
  - For that value, the probabilities for being in each of the different regions in which the image is divided are calculated.
  - The speed is obtained through a relation between the two higher probabilities (most probable regions containing that value).
- We calculate a coupling force between every surface and all the others, which will help them stop each other at the edges.
- Then the propagating curve is recalculated.
- After some iterations, the signed distance function is reinitialized.
- Back to third step.

4.2 The signed distance function

When segmenting images by propagating curves with the level sets method, the surface introduced evolves as the zero level set of the signed distance function (SDF). If the signed distance function is not correctly chosen, irregularities can appear when the non-zero level sets evolve, which can deform the SDF until it is no longer a signed distance function.

Along with that, we should also consider that the calculation of the normals depends directly on the SDF.

The fast marching method is used to implement the solving of the equation:

$$\phi_t = sign(\phi_0)(1 - |\nabla\phi|) \tag{4.1}$$

The reinitialization program, implemented in C and compiled with the mex-library, so it can be called from Matlab, follows the procedure explained herein:

1. First, it finds the zero level set.
2. Then, all voxels surrounding the zero level set are marked as the next level to update. (The voxels lying inside the zero level set will not be updated so that the exact position of this level will not be changed.)
3. Start front propagation. Values in voxels that were marked in the previous stage are updated using the PDE. They are then marked as updated.
4. Continue the front propagation until all voxels have been updated, except for the zero level set.

This implementation is very fast compared to other methods, but it is not 100% exact since the update of some voxels can depend on others which may have not yet been updated. However, it has been chosen because the error is very small and the high speed is worth it.

(For details on the program, refer to [8].)

4.3 Calculating the probabilities and the speed term

In order to calculate the region-based term as in (3.7), we will first have to calculate the density probability functions which determine, for each value, the probability of being in a certain region.

The first step is determining into how many regions do we want to divide the image. Once that is clear, we proceed to associate a Gaussian component to each region.

A Matlab function has been implemented to calculate all the necessary parameters. It follows the procedure explained herein:

1. The histogram of the 3D image is calculated. With this, we are binning all the possible values in the image into a more treatable number of values. This is necessary since in one single image we can have as many different values as voxels, so it is a very big number to deal with (or can be dealt with at a very high computational cost). This number will be determined, in this case, by a vector that goes from zero to the maximum value in the image, and can have as many elements as we want. When generating the histogram we will also obtain the number of voxels in the image, with each value.

2. The probability of each region in the image is computed by adding up all the voxels in the region and dividing them by the total of voxels in the image.

3. Then a conditioned probability is calculated for each of the possible values in the image of being in each of the regions. The reason why the calculated probability is conditioned is that when calculating the speed, we start from the supposition that the given voxel belongs to a certain region. This probability is calculated with the total number of voxels with the same value, inside each region.

4. The mean and variance are then calculated, with the usual formulas:

$$\mu = \sum_i x_i p(x_i), \quad \sigma^2 = \sum_i p(x_i)(x_i - \mu)^2$$

5. The Gaussian distribution is computed for every region, from the mean and variance.

6. Finally every Gaussian will be multiplied by the total probability of the corresponding region in the image.

With this procedure, we obtain the Gaussian probability function for every region and we can therefore calculate the probability for value in a certain voxel of belonging to a certain region.

With those probabilities we will be able to obtain the speed term as in (3.7).

This procedure is repeated for all the images used.

Initially, to start from simple to complicated, we divided the images into only two regions. So the values in every voxel were either in or outside one region. The speed was then calculated as explained above, having just one hypothesis where every value had a probability of being in region one and a probability of being in region two. But that doesn't give us the parts we are looking for.

When more regions are introduced, the problem becomes more complicated, because before calculating the speed, we have to decide the probabilities of which two regions are the correct ones to use. Assuming the evolution of ∂R_i (region R_i's boundaries) the first probability will clearly be p_i. The other probability will be the one corresponding to the most close neighbour region, which will be:

$$max\{p_j,\ \forall j \neq i\}$$

4.4 Coupling the level sets

Basing on the theory for coupling forces explained in Section 3.5.2, we have implemented a coupling force which will help a curve representing the evolution of region R_i's boundaries to expand when finding a voxel that certainly belongs to R_i. And will force it to shrink when the voxel found belongs to another region. What we want to accomplish with this force is avoiding the overlapping of surfaces evolving in different regions and, along with that, helping the curves to find exactly the boundaries of the region they are looking for.

The coupling force is implemented as follows:

$$H_i\left(j,\phi_j\left(s\right)\right) = -\begin{cases} 0, & if\ j = i \\ sign\left(\phi_j\left(s\right)\right), & if\ j \neq i\ and\ \phi_j\left(s\right) < -a \\ \frac{1}{\tan(1)}\tan\left(\frac{\phi_j(s)}{a}\right), & if\ j \neq i\ and\ -a \leq \phi_j\left(s\right) < 0 \\ \frac{1}{N-1}\frac{1}{\tan(1)}\tan\left(\frac{\phi_j(s)}{a}\right), & if\ j \neq i\ and\ 0 < \phi_j\left(s\right) \leq a \\ \frac{1}{N-1}sign\left(\phi_j\left(s\right)\right) & if\ j \neq i\ and\ \phi_j\left(s\right) > a \end{cases} \tag{4.2}$$

Having experimented, we found that the best value for a in this case was $a = 0.5$.

4.5 Weighting factors

As seen in the theory explained before, the speed with which our propagating curve evolves is based on what corresponds to the internal force in the classical snakes model. That is, the combination of two terms. One of them represents the constant speed with which the curve is pushed towards the contours. The other one represents the curvature and is a regularization term. It is important to weight these two terms properly, as the correct propagation of the surface will depend on them.

The regularization effect is necessary, but the appropriate weighting terms must be found. When evolving the curve, for example, along a fiber, we want the curve to evolve regularly in the normal direction and the curvature shouldn't inhibit the propagation. On the other hand, we need the regularization for those parts where the propagation is slower, due to irregularities in the way.

Our geometric flow will therefore be:

$$\frac{\partial S}{\partial t} = (\gamma v + \delta k)\vec{N}, \tag{4.3}$$

When including the coupling forces in the equation, another weighting factor will be needed, since every force has to have the appropriate influence in the evolution of the curve. Therefore, the flow is actually:

$$\frac{\partial S}{\partial t} = (\gamma v + \delta k + \beta H)\vec{N}, \tag{4.4}$$

as stated in Section 3.5.3. Following the theory explained, the weighting factors should add one.

Different combinations of weighting factors will provide different results, which are shown in chapter 5. For example if β is not high enough compared to the others, the influence of the coupling forces will not be enough. In that case, we can find that two surfaces evolving, looking for the same region, may not merge when they find each other, or if looking for different regions, they may go over each other. On the other hand, if this term is too high surfaces may not evolve enough to reach the boundaries of the region. The same happens with the other two factors. If the curvature term has too much influence, the function may result in curve close to the one we are looking for, but much more rounded, while if it is too low, the resulting shape may be too angular.

One of the things we try to do, is consider every region in which we divide an image, as a whole. This means that we may find a voxel inside a certain region, whose probabilities would say it belongs to a different region, but if it is a single, left alone voxel, we want the regularization term to help omit that voxel so it doesn't break the regularity. γ, the term which determines the influence of the region-based term, when being too high with regard to the other two, may inhibit the function of the coupling forces or the regularization term. This can be disadvantageous since then, those isolated voxels we were describing may be had in consideration, causing for example the starting of small new surfaces (due to the property which enables to handle topology changes) where they shouldn't.

4.6 Stability Condition

To assure stability, the time-step for the front propagation is calculated according to the CFL condition, as in [8]. This criteria states that when looking for stability in a numerical scheme, the domain of dependence for each point should comprise the domain of dependence of the PDE itself.

The time-step obtained following the above is:

$$\triangle t \leq \frac{1}{max\left(\gamma v + \delta k \overrightarrow{N} + \beta H\right)}$$

In order to keep our signed distance function as what it is, and not having to deal with strange behaviour of it, a very small time step has been chosen for this case. We have observed that $\triangle t = 0.1/max\left(\gamma v + \delta k \overrightarrow{N} + \beta H\right)$ is the best choice. The time-step being so small is a consequence of the way of fixing the level set in the reinitialization algorithm.

4.7 Surface initialization

The idea is to initialize surfaces in several parts of the brain. It would be very inaccurate to initialize just one surface inside the thalamus because its boundaries are not very clear. However, if we have various different surfaces propagating, due to the coupling forces, the surfaces will force each other to define the exact borders of the different parts. Every surface will evolve with its own speed. The best way of choosing the parts where surfaces will be initialized is to look for adjoining parts. This way, they will help define each other.

In our case, the best parts to look for, as will be later explained, are the thalamus itself, the ventricles, and the fibers which, in the axial slices, can be seen lying next to the thalamus. The method was tested initializing 3, 4 and 6 different surfaces in the same and in different images.

As explained earlier on, the method used can handle topology changes. When initializing three different surfaces, each one of them looks for one of the parts mentioned above. Once found the corresponding area, the surface jumps onto the next area with the same characteristics. Hence, if we initialize a surface on the left lobe of the thalamus, once completed, the level sets start another surface on the right lobe, as it considers it the same type if region.

When six surfaces are initialized, two of them look one for the left and one for the right thalamic lobes. Two of them look for the left and right side fibers, a fifth one looks for fibers behind the thalamus, which are next to the ventricles, and the last one looks for ventricles.

4.8 Data Images

The images used in this project are diffusion tensor magnetic resonance images obtained from healthy volunteers.

29 axial slices in a 128 by 128 matrix covering the center of the brain were acquired. Diffusion weighting was performed along 6 independent axes and a normalizing image without diffusion was also required. The resulting images were those of the diffusion tensor.

Brain Compartment	Diffusion coefficient ($\cdot 10^{-3} mm^2/seg.$)
CSF	2.94 ± 0.05
Gray Matter	0.76 ± 0.03
White Matter:	
Corpus Callosum	0.22 ± 0.22
Axial Fibers	1.07 ± 0.06
Transverse fibers	0.64 ± 0.05

Table 4.1: Diffusion coefficients of water in several human brain compartments.

As seen in the theory explained in Section 2.1.3, the diffusion tensor for each voxel is a 3x3 matrix. Therefore, we cannot see it as a value, so we must calculate the mean diffusivity or some other scalar measure.

With the diffusion tensor images, using the diffusion tensor theory, we can also calculate the diffusion coefficients and therefore obtain the brain maps for c_l, c_p, c_s and FA (equations (2.3) to (2.6)).

It is from the values in these images that we calculate the different probabilities for the regions. The difficulty lies on the fact that not all the regions are clear in all the images. And, above all, the main problem is that the values and type of diffusion in the thalamus are not clearly defined.

We will now explain the procedure through which we differentiate the regions to obtain the probabilities and further segment the thalamus.

The first thing to do is determine in which image we will best find each region. In order to do that, let's evaluate the information we have on diffusion coefficients in the brain.

Table (4.1) [15] resumes the diffusion coefficients of water in several human brain compartments.

Analyzing the information on diffusion and images obtained, we find that each one of them shows most clearly the components detailed herein:

- Mean Diffusivity: As can be deduced from the table above, the most easily distinguishable component when using the mean diffusion images is the CSF. Because the diffusion values in this component are much higher than the rest, we'll find a visible difference. This fluid is what fills the ventricles, so that will be the part that we will most clearly see. In fact, looking at the mean diffusivity, we can see that the images can only be clearly separated into two regions: the ventricles and the rest of the brain.

- Fractional Anisotropy: The parts better distinguished in the FA images are mainly the fibers found next to the thalamus. Fibers are part of white matter, where diffusion is extremely variable. However, inside the fibers diffusion tends to be highly anisotropic as it is mainly in the direction of the fiber tracts, while grey matter and fluid don't exhibit high anisotropy. So the values for the fibers in the FA images are much higher than the rest and therefore, easier to differentiate.

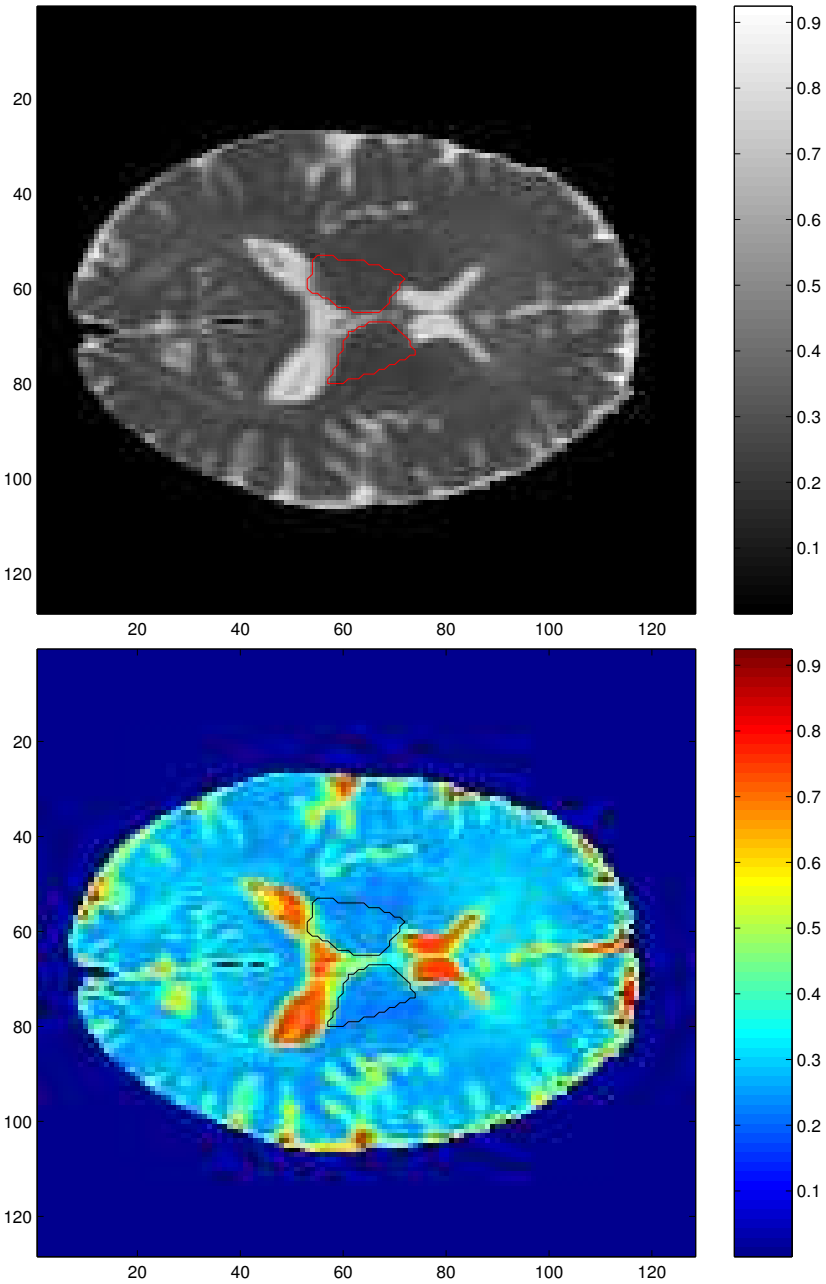

Figure 4.1: Normalized Mean Diffusivity maps of a DT-MRI, axial slices. Same cut (slice #20) in black and white, and colours. The thalamus is the area inside the red (black and white image) and blue (colour image) contours, in the center of both images.

- Linear diffusion (c_l): The only part which can be clearly differentiated in these images are the fibers, as they are highly anisotropic, and the values for c_l are higher than in the rest of the brain. However, for that matter, FA images are clearer, so we discard c_l ones in favour of FA.

- Planar diffusion (c_p): There are intermediate values of planar diffusion all over the brain, so these images are not very useful for the purpose of this study.

- Spherical diffusion (c_s): The highest values of spherical diffusion coefficients are found inside the ventricles, since there, diffusion is totally spherical. The next values in the scale are in the thalamus, since it is made of grey matter and the diffusion there is not highly spherical but not highly anisotropic either.

Having said all the above, it seems clear that the appropriate images to look for the thalamus are both the FA and the spherical diffusion maps. In both of these two maps, the values in the thalamic region will stand in the middle. In the fractional anisotropy they will be closer to the lower values, which in this case will correspond to the ventricles. And in the spherical diffusion coefficients map, the values in the thalamic region will also be closer to the ventricle ones, but in this case, these will correspond to the highest values. Obviously, when a region has a high anisotropy, the values for the spherical diffusion are low, and vice versa. These characteristics can be appreciated in Figures (4.2) and (4.3), where maps of the fractional Anisotropy and Spherical Diffusion are shown. The thalamus is shown by a manually drawn contour in red (blue for the spherical diffusion in colours). These contours were handled by a neurosurgeon who has segmented the thalamus by hand.

Inasmuch as its situation in the brain is concerned, the thalamus is surrounded by ventricles (lateral ventricles on the sides, and third ventricle going through the middle of the thalamus) and white matter (mainly the corpus callosum going over it).

This is useful information for the segmentation because, as explained above, there are images where those parts appear clearer than the thalamus. Then, we can look for them and in finding those other parts, we will also be defining the thalamus' limits. That is the sense of using coupling level sets for this application.

A surface can be initialized in a region different from the thalamus, but which is next to it. Due to the coupling forces, when that surface reaches the one evolving inside the thalamic region, they will stop each other and adapt until they find the boundaries of both regions.

From the images obtained from the magnetic resonances we also concluded that the thalamus could be found, in this case, somewhere between the 11th and the 25th axial slices (out of the 29 initial ones). Using this information we were able to zoom the images in order to better see the area of interest.

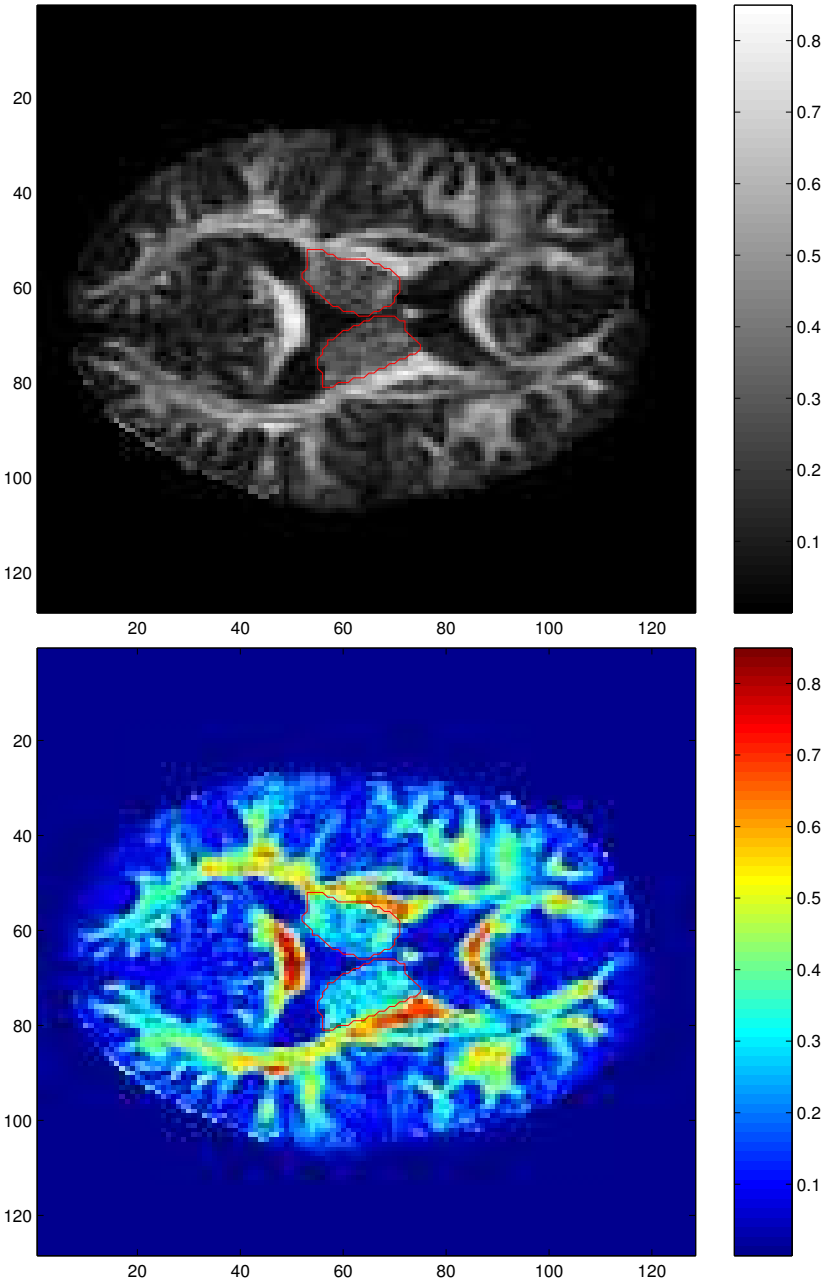

Figure 4.2: Fractional Anisotropy maps of a DT-MRI, axial slices. Same cut (slice #19) in black and white, and colours. The thalamus is the area inside the red contours, in the center of both images.

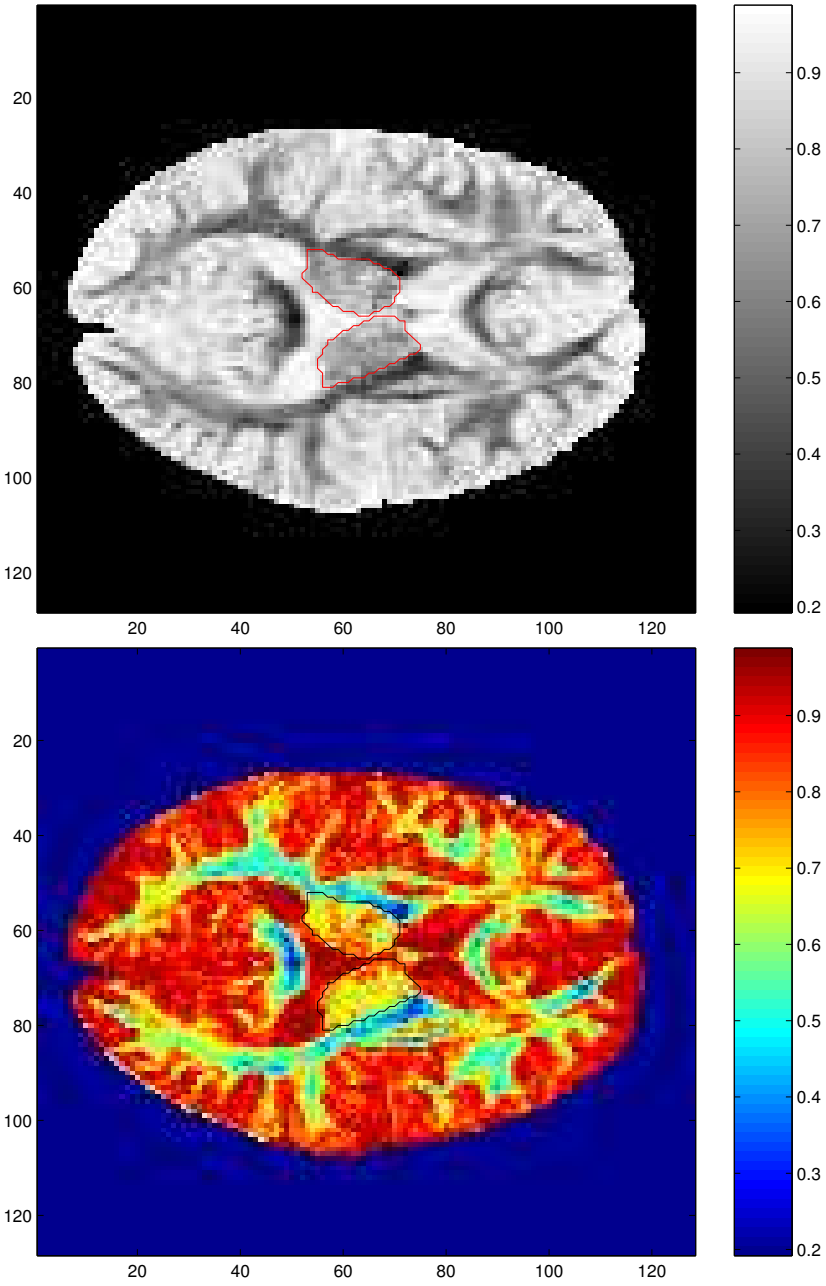

Figure 4.3: Spherical diffusion coefficient (c_s) of a DT-MRI, axial slices. Same cut (slice #19) in black and white, and colours. The thalamus is the area inside the small contours, in the center of both images.

4.9 Image Treatment

Given we are working with real brain images and the purpose of this project is to obtain an application suitable for medical surgery, the idea is to accomplish our aim using directly the real data, without modifying the images before making use of them. However, we still want to obtain the best possible results, and to do that, there are some problems we can solve. Therefore, as a part of the process, some treatment can be done to the real images if we are sure not to lose any information.

First, to reduce noise and smooth some edges, the images used as data can be pre-processed. For this case, we can convolve them with a Gaussian filter before being used. A rotationally symmetric Gaussian low pass filter of size 3x3 and with a standard deviation sigma of $\sigma = 0.5$ will do.

The first difficulty with which we have to deal is the fact that the thalamus is, in some images, very difficult to distinguish from other parts of the brain when using diffusion measures. Its boundaries are not clear.

We can try to highlight the differences between the values in neighbour regions. A very simple way to do it is using the square root or second power of the values in the data images. One or the other solution depending on if we want to differentiate the higher or the lower values in the image.

For example, in the Fractional Anisotropy maps we will try to segment the axial fibers. Inside this region we find the highest FA values (Section 4.8), which go approximately from 0.55 to 0.9 in a scale from 0 to 1. That makes them very difficult to group in just one region. The difference between a value in that region and one in a neighbour region could be smaller than between two values inside the fibers. If we calculate the square root of the fractional anisotropy, we find that the thresholds separating the gaussians which correspond to the three different regions are nearly the same as for the fractional anisotropy image (Figure 4.4). However, because the values are between 0 and 1, when computing their square root, they augment. So now, some values for which it wasn't clear if they belonged to the second or the third region, are now confined clearly to the third region. In Figure 4.4 we can see an amplified image of the square root of the fractional anisotropy in the thalamic region and neighbour zones. The gaussians are also presented, for three regions in the fractional anisotropy map (blue plot), and its square root (red plot).

It would also be helpful to remove from the images the parts which we know for sure are not the thalamus. For this purpose we will use, as in the rest of the segmentation process, the probabilities for the different regions.

The first and easiest part to remove, as seen in Section 4.8, are the ventricles. In the images of mean diffusivity we can clearly distinguish this part from the rest of the brain as the values for mean diffusivity are much higher. We also know that the thalamus is found next to the ventricles, and therefore, in the axial cuts, they share their contours. Hence, we will use the mean diffusivity image to make a mask where the ventricles are segmented, removed from the image. The thalamus will be contained in the remaining zone and some parts of its boundaries will be already defined.

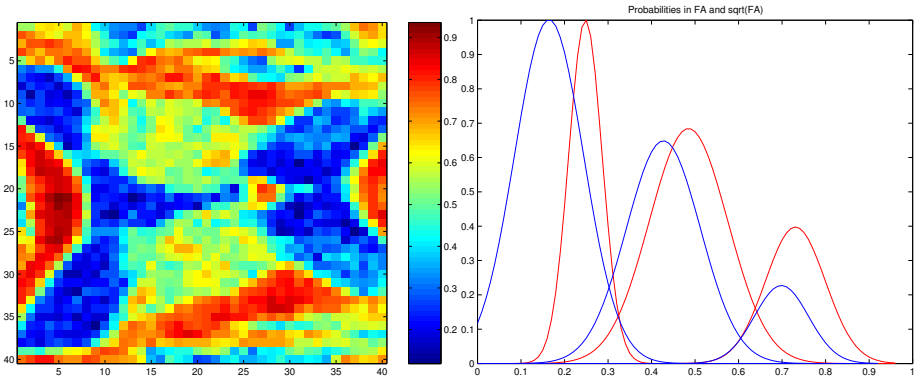

(a) (b)

Figure 4.4: (a) Amplified image of the square root of the fractional anisotropy in the thalamic region and neighbour zones (axial slice #20). (b) Probabilities for three regions in the fractional anisotropy map (blue), and its square root (red).

To do that, we calculate the probabilities for the mean diffusivity images, separating two different regions. With those we can make a probability map for each region, where the probability of each voxel being in that region is represented. Taking the probability map for the region in which we are interested (the part which is not ventricle) we find that there are very low values inside the ventricle part, and very high values for the rest. A mask can then be extracted by eliminating the parts with probability under a certain threshold. Having tested, a threshold of 0.01 was found to fit the needs of this case.

The mask is then multiplied by any of the images we want to use, and that way, removes the unwanted parts (in this case, the ventricles) according to the most probable regions. The result applied to the spherical diffusion coefficient image can be seen in Figure 4.5.

Another of the problems we find arouses from the fact that there are other areas, near the thalamus, which are also made of grey matter. Inside those areas, values and diffusion properties will be quite similar to the ones in the thalamic region. We stand in front of one of the drawbacks of using a method which can deal with topology changes. Because surfaces can split and merge if needed, the initialization surface or surfaces looking for the thalamus, once found it, will probably jump onto some other nearby grey matter zone.

For that reason, it seems a good idea to restrict the zone in which we look for the thalamus. Given we know approximately where in the brain is the thalamus found, what we can do to restrict that zone is make a map which highlights the voxels standing inside that zone or its surroundings, and attenuates the

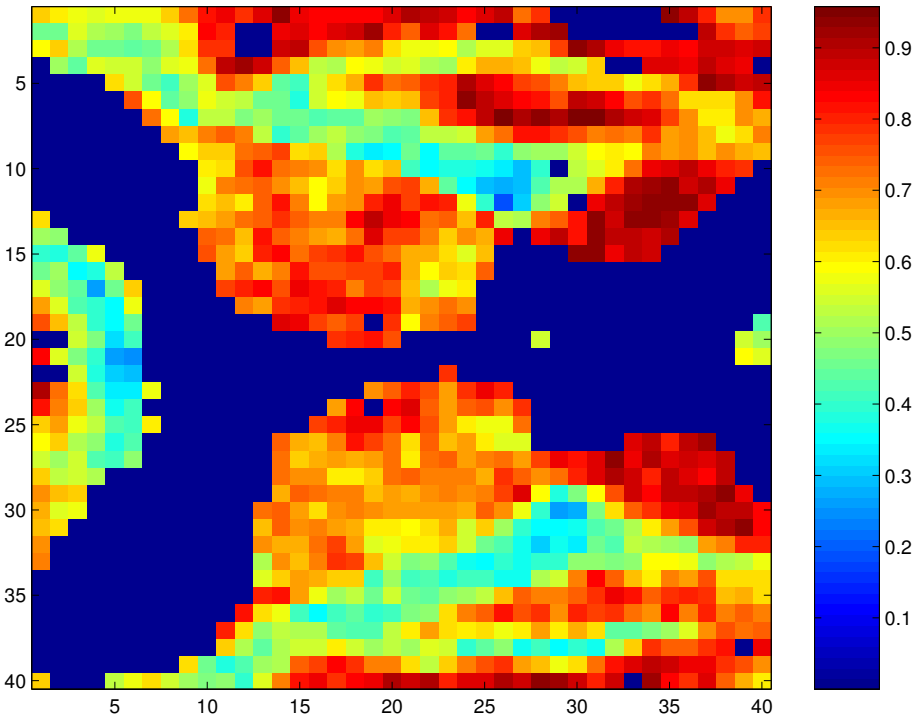

Figure 4.5: Amplified image of spherical diffusion in axial slice #19, having removed the ventricles with a probability mask. The thalamic lobes are the two semi-rounded volumes in the center of the image, with values between0.6 and 0.7.

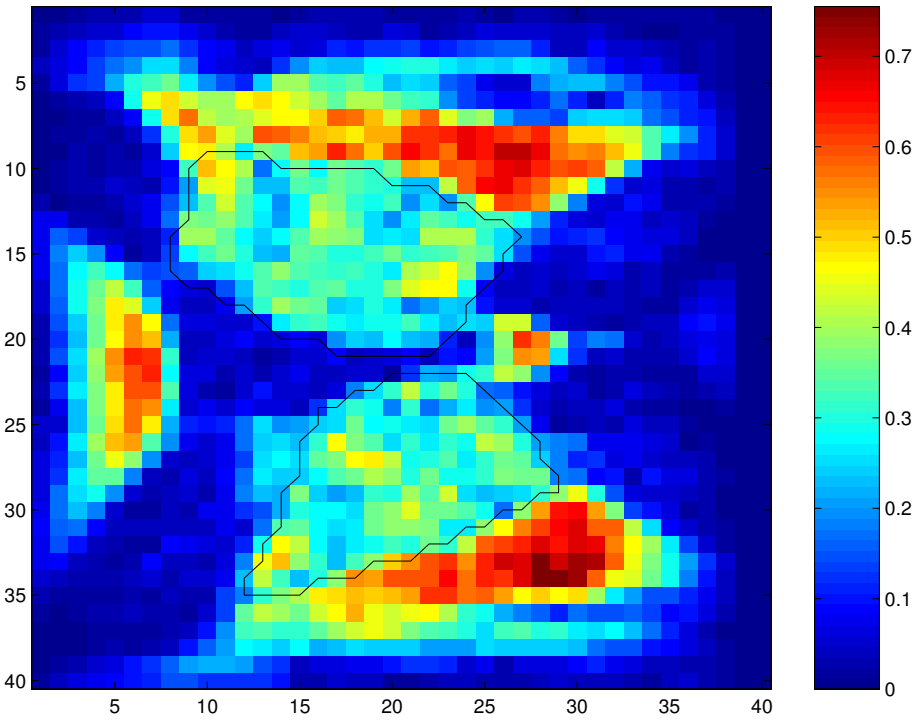

Figure 4.6: Amplified image of Fractional Anisotropy multiplied by a Gaussian centered in the thalamic region. Axial slice #20. The thalamus is shown by the black contours drawn by a neurosurgeon.

ones outside. With that purpose, we will use a pulse (centered in the middle of the image) convolved with a Gaussian filter, and that will be multiplied by the image we want to use to find the thalamus. The pulse is a 40x40 matrix with a rectangle of 35 by 30 ones in the center of the matrix, and zeros in the rest. The filter is a normalized 8x8 Gaussian, with standard deviation sigma of $\sigma = 15$. We can see the result applied to the fractional anisotropy image in Figure 4.6.

Chapter 5

Results

All along the present chapter we are going to show the segmentation results obtained with the method developed in this work. We will first present the results obtained when trying the method on synthetic images (Section 5.1). Once the method was found to work for those, we went on to try it on real brain images, results of which are shown in Section 5.2.

The shapes for which we looked for in the synthetic images and the values we were dealing with, were different to those in the real images. Therefore, the parameters needed are not the same for ones and for the others, and even if they were correctly chosen for the synthetic images, they would have to be readjusted for the real ones. For that reason, experimentation on real images started before having accurately found the exact parameters for the synthetic fields.

When testing the method with real brain data, first we used directly the fractional anisotropy and diffusion coefficients images. Then some of them were processed, to see if better results were obtained.

We will now proceed to show the results.

5.1 Synthetic Images

On a first stage, the method was tested on synthetic images, created from synthetic tensor fields. The images consist of two regions, one with isotropic and another with anisotropic tensor values, taken from real data of DT-MRI from a healthy human brain. The isotropic region is used as a background for the anisotropic one, which represents fibers in the brain. A non-identical tensor field is closer to the real case, so uniformly distributed random noise was added to the isotropic tensors. The anisotropic tensors were rotated to obtain the desired direction and an approximation of Rician noise was added. After that, the tensor images were recreated. Then, the diffusion coefficients were calculated from the image, following equations (2.3), (2.4) and (2.5), as would have been done for real images. The synthetic images were created to have an isotropic

background and an anisotropic region representing a fiber. Therefore, the two regions will be clearly differentiated in the linear diffusion coefficient and the spherical diffusion coefficient images. Let's take, for an example, the linear diffusion coefficient one. The image can be seen in Figure 5.1 (a).

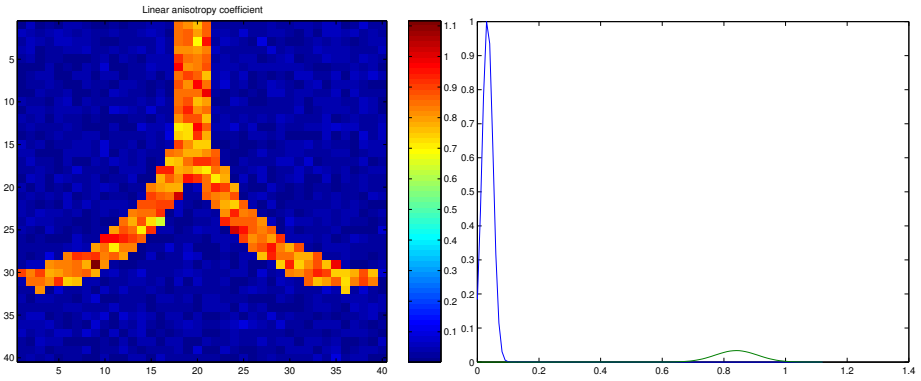

(a) Synthetic Linear Diffusion Coefficient (b) Probabilities for two regions

Figure 5.1: Map of synthetic linear diffusion coefficient and probabilities when separated in two regions.

The method was then tested on this synthetic image, to distinguish two different regions.

First, the probabilities were calculated for the two regions, using the function *probab_whole*. The obtained probabilities are shown in Figure 5.1 (b), being the means 0.0329 and 0.8404, and the variance 0.0178 and 0.0686 (respectively, for the first and second region). It can be clearly seen that for this case two very differentiated regions are obtained.

Then, to start the segmentation, three surfaces were initialized in the image. A small one inside the fibers and two slightly bigger ones in the isotropic region, all of them rectangular shaped.

The surfaces evolved normally and an acceptable segmentation was obtained, finding the whole synthetic fiber tract in 500 iterations. Figure (5.2) shows one of the results. The contour of the ϕ function once having found the fiber tract can be seen (shown by a thin line in magenta). It is visibly not the exact contour of the synthetic fiber. However, the adjustment is a matter of finding the right weighting factors. Once at this point, we decided to start trying the method in the real brain images, since finding the right parameters for the synthetic images wouldn't assure having them for the real images.

Several simulations were made to see the dependency on the placement and size of the initial surface, and no significant difference was observed between them.

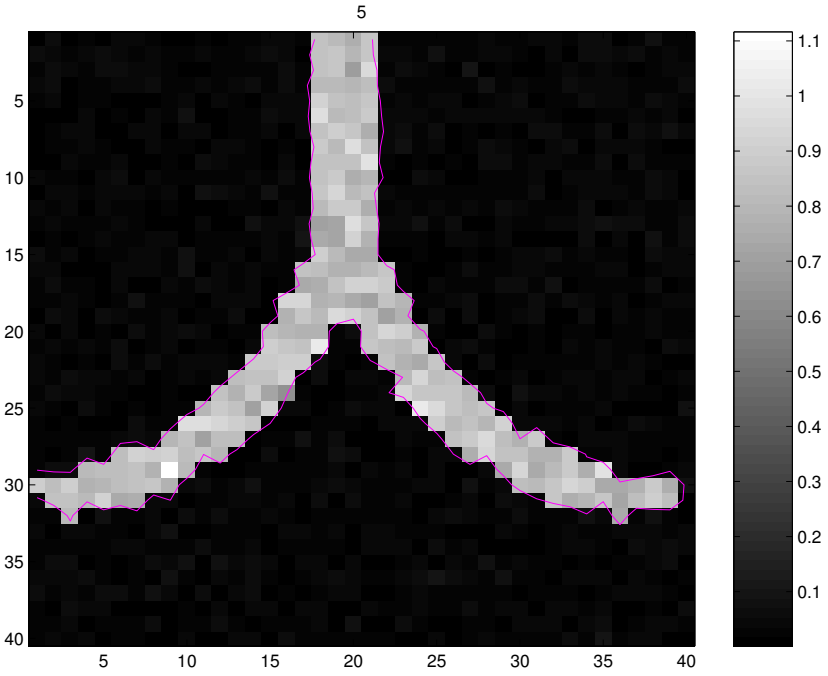

Figure 5.2: Image shows the contour found by the surface initialized inside the synthetic fiber tract (contour in magenta).

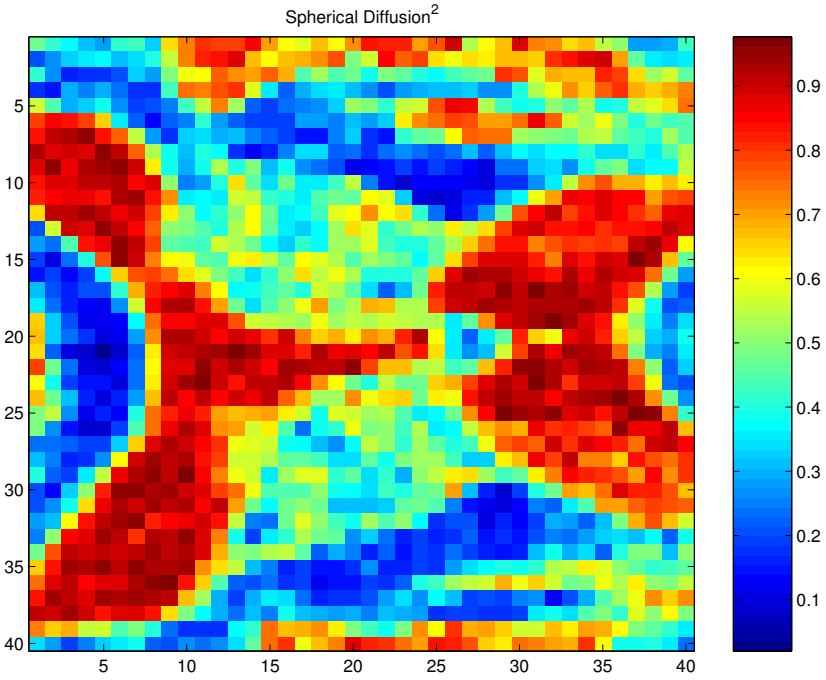

Figure 5.3: Spherical diffusion raised to the 2nd power of an amplified image of the thalamic region (axial slice #20). The zone with the higher values (approx. from 0.7) are the ventricles, the lower values (up to 0.3) correspond to the fibers, and the middle values (zone in light blue) represent the thalamus.

5.2 Real DT-MRI

Having tested the method with synthetic images, we went on to try it on real human brain images. The data images used, as explained in Section 4.8, are obtained from DT-MRI. Before doing anything else, we zoomed the images in the axial slices so the zone in which we want to work remained centered and could be seen better. The resulting amplified images were 40x40x15. So, from this point on, we worked with 15 axial slices of 40x40 voxel matrices obtained from the brain of one subject.

Except for the mean diffusivity images, where only two zones can be differentiated (the ventricles and the rest of the brain), in all the other images we can visibly distinguish at least three zones with different values. The three main regions correspond to the thalamus, the ventricles, and the axial fibers, as can be seen in Figure 5.3.

The image also shows some other parts, surrounding that group of elements, which have got similar values to the ones in the thalamus, as they are also composed of grey matter. So, inevitably, when doing any kind of separation

between regions, those areas will be included within the thalamic one.

We must say that in all simulations, carried out both on the data images and on the treated images, we looked for the thalamus, the fibers which surround it, and the ventricles next to it. So at least one surface was initialized inside each one of these regions. This way, the coupling forces can have their effect on the evolving surfaces. Not all the contours found are shown in all of the images, so the reader can see more clearly the ones that do appear.

First we carried out the segmentation on the data images, without giving them any previous treatment. We looked for the thalamic region in the fractional anisotropy image. To do so, we divided it into three different regions, since it seemed the most logical thing to do as there are three different types of tissue we are looking for. Because the values for the fractional anisotropy in the thalamic region are between the ones in the fibers and the ones in the ventricles, the second region corresponded to the thalamus. The fibers and ventricles were searched in the spherical diffusion images. These were separated into four regions because we observed that this way, the values inside the corresponding gaussians fitted more accurately the ones in the real areas.

The target was then finding the appropriate weighting factors. The weighting factors, β, γ, δ, determine how much influence does each force has on the evolution of the curve (Section 4.5). The γ corresponds to the region-based term, that is, the speed dependent on the probabilities. The β represents the influence of the coupling force, and the δ corresponds to the curvature term. Several combinations of these factors were experimented in different simulations, following the theory explained in Section 4.5 and trying to find the best set of factors by experimenting and regarding the results. We observed that some good results were obtained for: $\gamma = 0.03$, $\delta = 0.8$, $\beta = 0.17$. We also observed that the thalamus was entirely segmented in 500 iterations although, depending on the proportions of the weighting factors, the other regions were not always completely segmented and required 100 or 200 more iterations. The best frequency to reinitialize the signed distance function was found to be every two iterations. If this number was increased, small unwanted surfaces started to develop in different areas. If we put this number to one, not a great difference is observed, but computational time increases.

Another important thing we noticed was that in the segmentation, even though being able to distinguish the shape of the thalamus in the middle axial slices, in the bottom slices the shape wasn't correct. The resulting curve at the end (bottom) of the thalamic region was much bigger than the one drawn by the neurosurgeon. This, we observed later, occurs in all segmentations and is due to the fact that under the thalamus we find the subthalamic region, whose diffusion values are very close to the thalamic ones.

We proceeded then, with the same input parameters mentioned above, to experiment segmenting the thalamic region in the spherical diffusion image, dividing it into four different regions as before, but now we would look for the third region. The fractional anisotropy image was then used to segment the fibers and ventricles, assuming again a division of the image in three regions.

So, being the weighting factors: $\gamma = 0.03$, $\delta = 0.8$, $\beta = 0.17$, with 500

iterations and reinitializing the signed distance function every 2 iterations, we obtained the results in Figure 5.4 (a). We can see an axial slice with a contour showing the result of the function which evolved inside the thalamic region (contours in blue), and another contour (in green) which shows the ventricles. The surface inside the thalamic region hasn't grown enough, probably due to a low β. In (b), the same slice using different weighting factors is shown. The surface looking for the thalamus has, in this case, grown too much and gone over the fiber region. This may be due to the fact that γ is higher, and therefore small surfaces are initialized on the outer side of the fibers (opposite to the thalamus), quickly merging with the ones inside the thalamic region because β is also higher. Weighting factors are for this case: $\gamma = 0.05$, $\delta = 0.7$, $\beta = 0.25$.

(NOTE: In all images, from now on, the white contours are the handmade segmentations handled by the neurosurgeon.)

Considering that the shape of thalamus tends to be more spherical, whereas the fibers, for instance, have got a much linear shape, we thought it would be helpful to have a higher curvature factor for the surface searching the thalamus, than for the others. So, several simulations were experimented, doubling the curvature factor only for the surface evolving in the thalamic region. Better results were obtained. Figure 5.5 shows an example.

Later, simulations were made on treated images. In the first place, we looked for the thalamic region in the spherical diffusion raised to the 2nd power, considering three different regions in the image. The fibers and ventricles were searched in the square root of the fractional anisotropy, which was also divided into three regions. Under the same conditions (weighting factors, iterations, frequency of reinitialization of the signed distance function) as with the non-processed data images, slightly better results were obtained.

For example, for the weighting factors: $\gamma = 0.03$, $\delta = 0.8$, $\beta = 0.17$ some good results were obtained for 500 iterations. In Figure 5.6 we can see an axial slice where the thalamus and the fibers (contours in blue) have been segmented with those (a). The bottom image (b) represents the same slice of segmentation using the same maps, but with inadequate weighting factors. The surfaces inside the thalamus haven't grown enough, probably due to a bad proportion of coupling force to region-based term.

In Figure 5.7 we have the same segmentation as in 5.6 (a), but showing the thalamus (contours in blue) and the ventricles (contours in green), in another axial slice. Underneath it, the same slice, segmented with the inadequate set of weighting factors. As seen in Figure 5.6, the contours found in the second segmentation are not as close to the handmade segmentation as the ones in the first segmentation.

Image treatment previous to the segmentation has also been tested on several simulations. Segmentation was tried using a probability map as a mask on the fractional anisotropy image, as explained in Section 4.9. Results came out better than the ones shown previously, since the parts outside the thalamus which have similar FA values are no longer there. So the surfaces evolving inside the thalamic region, once found it, don't go any further. Results can be seen in Figure 5.8 for weighting factors: $\gamma = 0.03$, $\delta = 0.8$, $\beta = 0.17$.

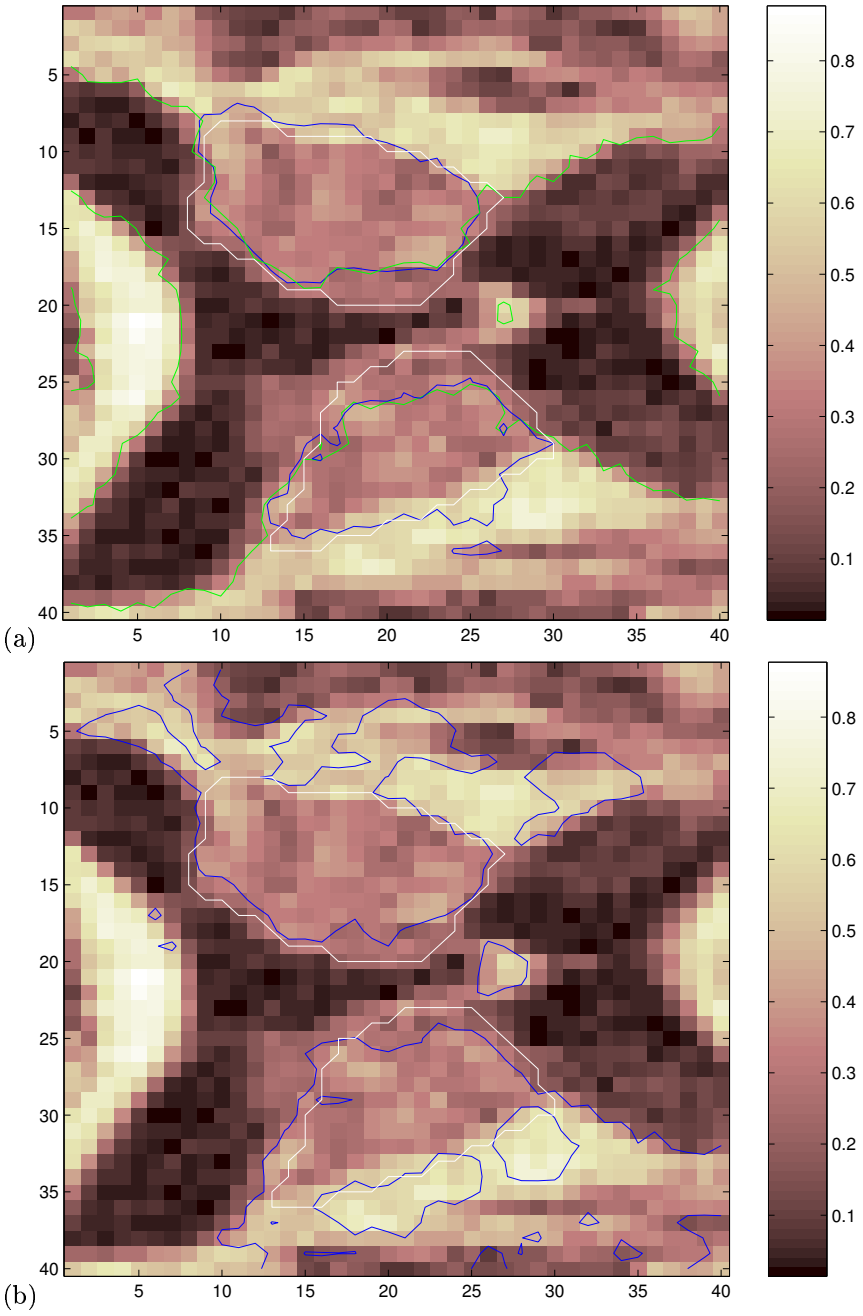

Figure 5.4: (a) Automatic segmentation of thalamus (in blue), ventricles (in green) and fibers (not shown on the image) with weighting factors: $\gamma = 0.03$, $\delta = 0.8$, $\beta = 0.17$, together with handmade segmentation (in white). Axial slice #20. (b) Same axial slice showing segmentation of thalamus (in blue) with weighting factors: $\gamma = 0.05$, $\delta = 0.7$, $\beta = 0.25$, together with handmade segmentation (in white).

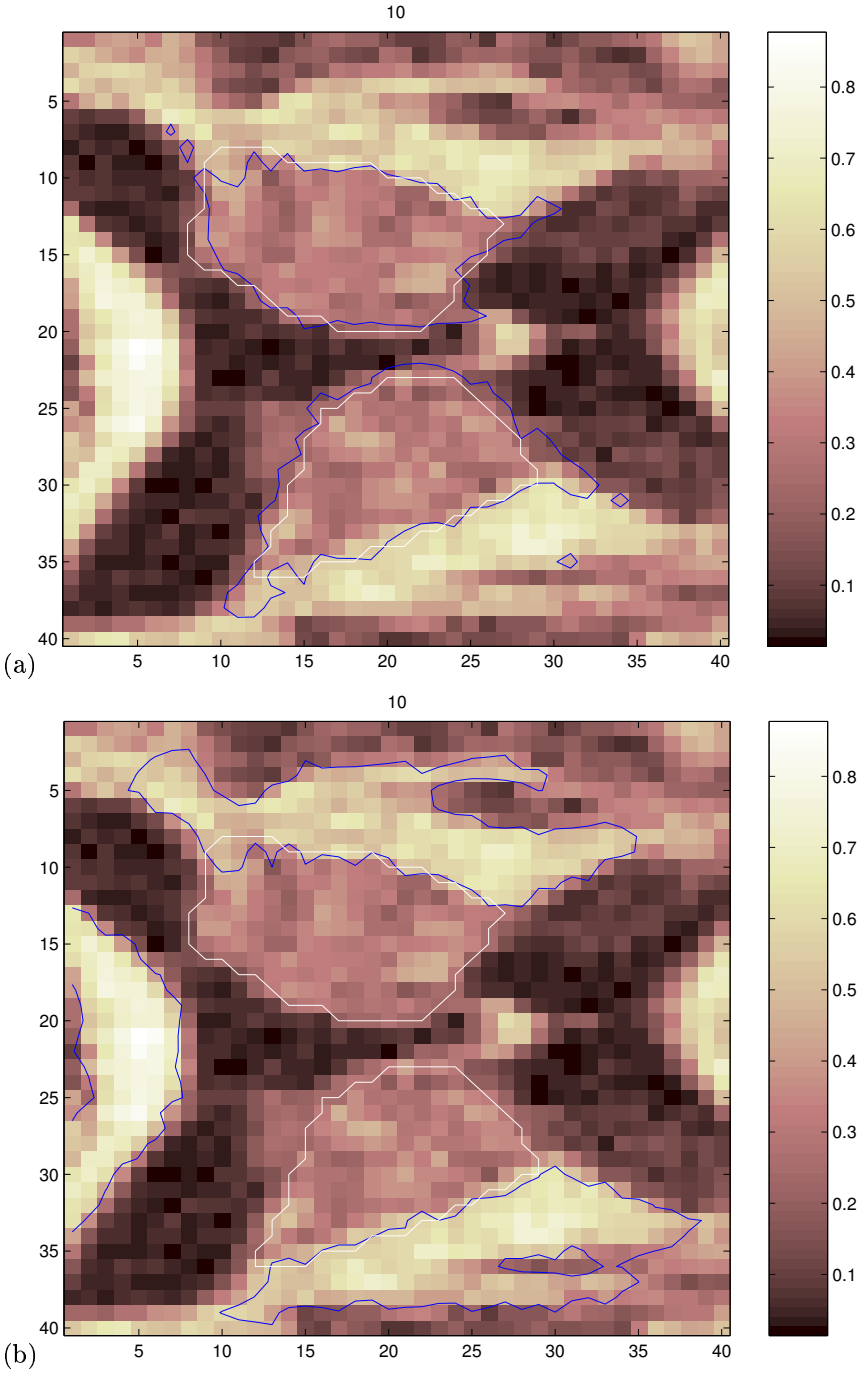

Figure 5.5: (a) Automatic segmentation of thalamus shown by blue contours, having doubled δ for the surface evolving inside the thalamic region. Hand-made segmentation in white. Axial slice #20. (b) Fibers found in the same segmentation.

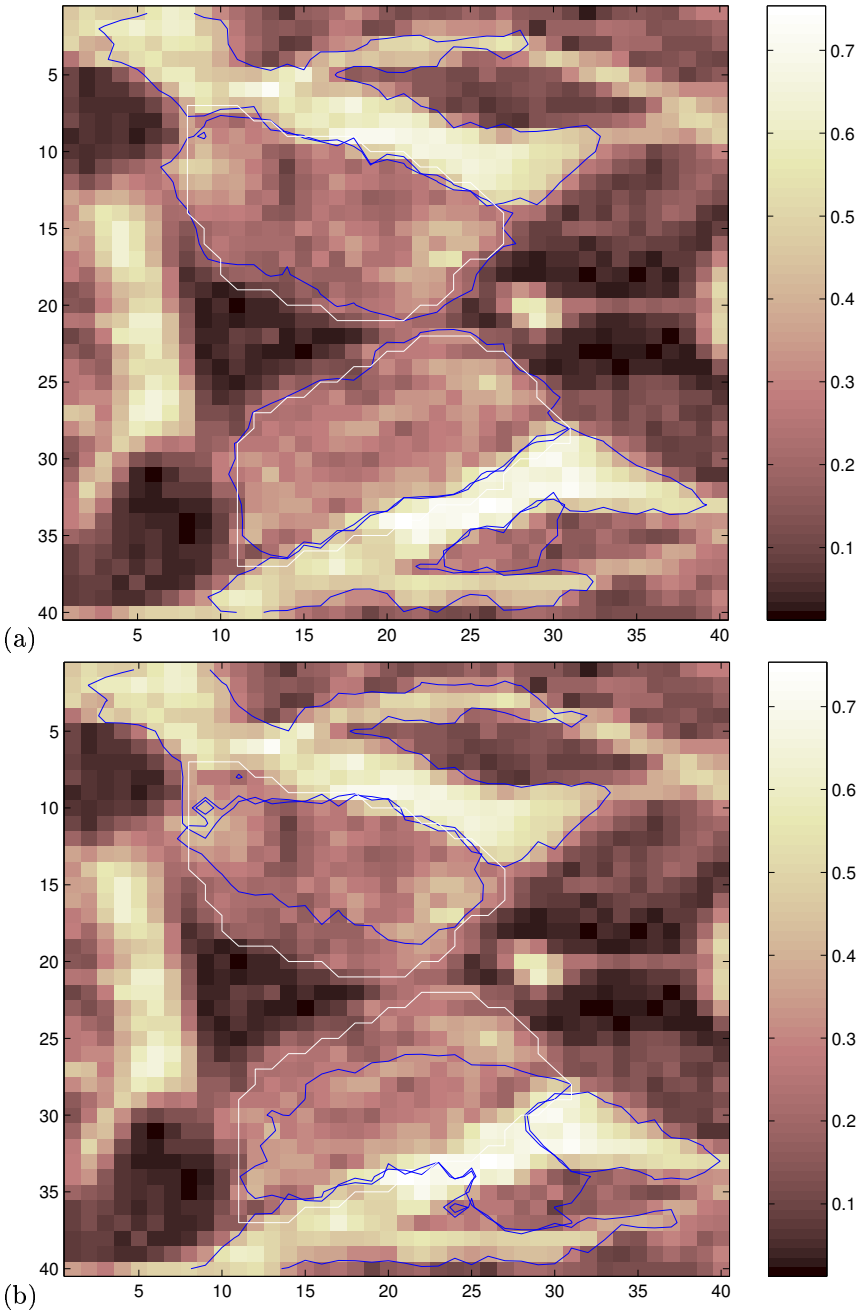

Figure 5.6: (a) Automatic segmentation of thalamus performed on spherical diffusion raised to the 2nd power. Axial fibers segmented in square root of fractional anisotropy. The handmade segmentation is in white. Axial slice #18. (b) Same axial slice of automatic segmentation on same images, with inadequate weighting factors, together with handmade segmentation (in white).

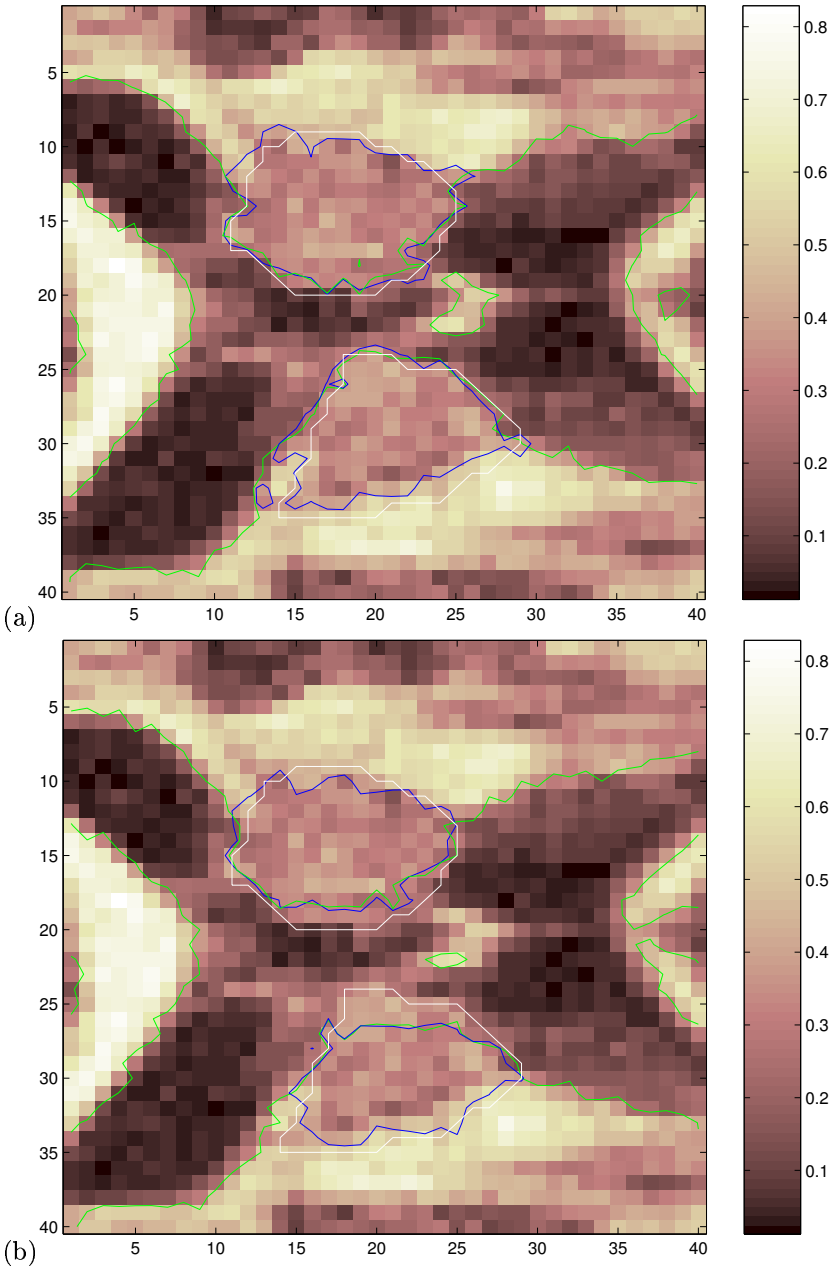

Figure 5.7: (a) Automatic segmentation of thalamus performed on spherical diffusion raised to the 2nd power (blue contours). Ventricles segmented in square root of fractional anisotropy (green contours). The handmade segmentation is in white. Axial slice #21. (b) Same axial slice of another automatic segmentation on same images. Contours found for the thalamus (in blue) and ventricles (in green) with inappropriate weighting factors, together with handmade segmentation (in white).

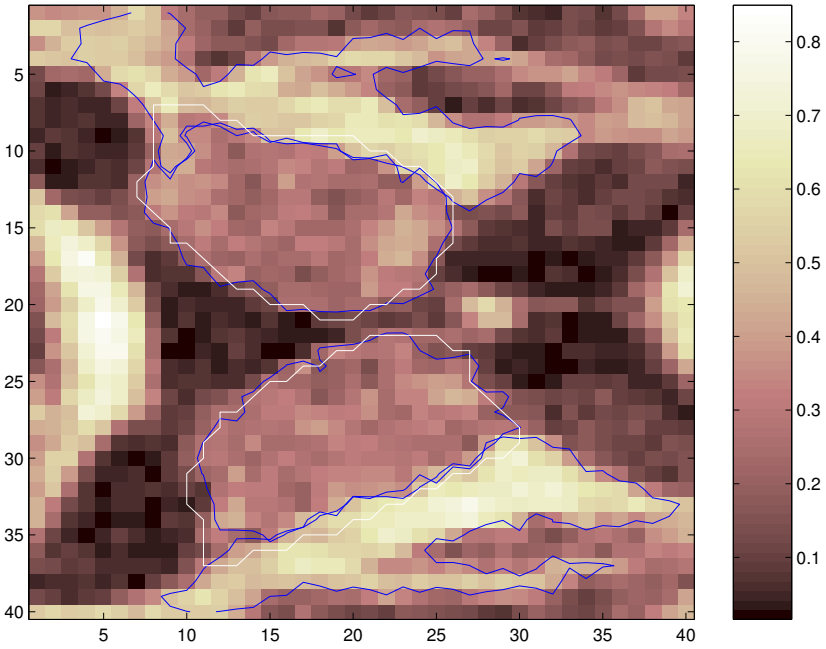

Figure 5.8: Automatic segmentation of thalamus and axial fibers (contours in blue) in fractional anisotropy image, treated with a probability map mask. Handmade segmentation of the thalamic region in white. Axial slice #19.

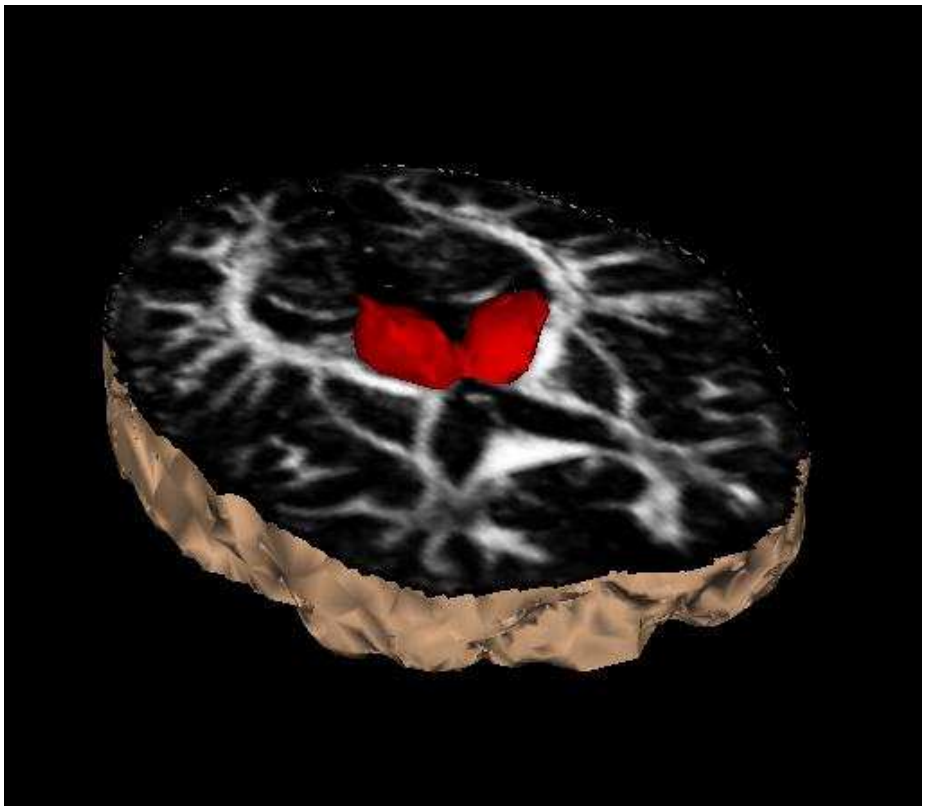

Figure 5.9: 3D segmentation of thalamic and subthalamic regions.

We can see that no new surfaces have been initialized on the outer side of the fibers.

All results shown in this chapter belong to DT-MRI of the brain of the same subject in order to keep a reference when evaluating them. However, simulations have been done on DT-MRI from the brains of other subjects to verify the method. Similar results were obtained.

To conclude this chapter, we proceed to show the 3D representation of one of the segmentations (Figure 5.9). The whole thalamic region, including the subthalamus, has been segmented. The axial, coronal and sagittal slices of the same segmentation are shown in Figure 5.10.

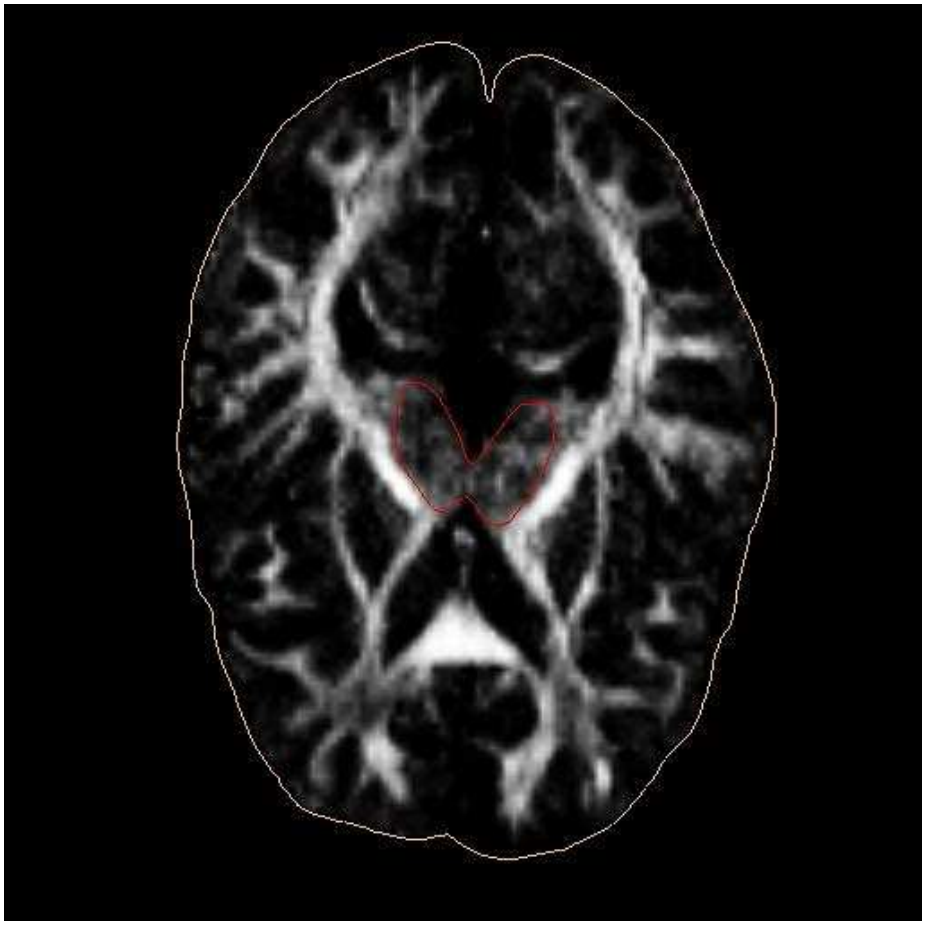

(a) Axial

(b) Coronal (c) Sagittal

Figure 5.10: Axial, coronal and sagittal slices of segmentation in Figure 5.9.

Chapter 6

Conclusions and Future Work

6.1 Conclusions

In the present work a new method has been introduced to segment the thalamus as a whole. It uses the statistics obtained from the water diffusion characteristics in the different areas of the brain.

As it has been shown in Section 5.2, the segmentation results are sensitive to parameters' choice. The function obtained may vary considerably depending on several elements.

The weighting factors, for instance, which determine how much influence does each force have in the evolution of the surfaces play a very important role. If they are not correctly chosen, the curve obtained could result very different from the contour we are looking for.

The kind of measures used have their influence too. It won't be the same trying to segment the thalamus on an image where the voxels reflect the fractional anisotropy values than on one which shows the linear diffusion coefficients. Moreover, the number of gaussians that are fitted for each image, that is, the number of regions in which each image is divided, can make a difference.

The segmentation will also depend on whether we use the direct values or we treat them before, for example, using a probability map mask.

Once having found the right parameters, the results obtained using the method presented were found to be highly accurate in some parts of the brain, mainly in the center of the thalamus. The middle axial slices of the data images used show very clearly the thalamic shape, and therefore the resulting curve of our function stands very close to it. However, it is not like that for the lower axial slices, where the thalamus gets confused with the subthalamus and it is difficult to distinguish. This is possibly due to the fact that some of the matter which composes the thalamus is also present in the subthalamus. Therefore the diffusion values and coefficients in both areas are very similar and when separating the images into regions, there is no distinction between the thalamic and the subthalamic zones.

It is true that the current segmentation results and accuracy were also limited by the image resolution. With higher resolution, the difference between the values in the thalamus and the regions standing next to it could be clearer.

Due to the reasons stated above, we could conclude that the chosen method is not yet ready to segment the thalamus on its own, since it doesn't yet give the high level of accuracy needed for neurosurgery purposes. Nevertheless it would be suitable to segment the thalamus together with the subthalamus, or even the whole diencephalon, and with some improvements it will be able to reach its purpose.

6.2 Future Work

On a first stage, there are several things that could be done to improve the present method. For a start, images with more resolution would help accomplish a better segmentation. Incorporation in the algorithm of an automatic procedure for determining the number of regions in every image would also be interesting, as well as detecting automatically which region should be searched by every initialization surface. A condition could be introduced to restrict the new surfaces which start in a new region after having found the area of interest, for example by volume. If we are looking for a region which is not smaller than a certain volume, we could ban new surfaces starting in regions smaller than that volume.

Finding a more accurate way of obtaining statistical data on the images could also give better results. Experimentation is being done to test the method using probabilities calculated using a finite Gaussian mixture model based on the Expectation Maximization Algorithm (EM), instead of the ones presented in Section 4.3.

Another thing that could be done to improve this work is incorporate in the algorithm a method to calculate the error in the segmentation once it is finished. Up to the moment, the results have been evaluated by visual comparison with the contours drawn by a neurosurgeon.

In addition to all that, we should bear in mind the initial idea of this project, which was to segment the different thalamic nuclei individually.

Even though the presented method is not appropriate to segment the thalamus on its own, when coming to resolve the thalamic nuclei inside the thalamus, the method can be used. The subthalamus is also composed in part by grey matter and divided into sub-thalamic nuclei. When looking for thalamic nuclei, the importance lies on distinguishing small nuclei on their own, independently of them belonging to the thalamus or the subthalamus, so resolving these two parts as a whole instead of separately would be an interesting starting point.

As far as segmentation of the individual thalamic nuclei is concerned, the method introduced in the present work wouldn't provide acceptable results. A more suitable method should be employed. Progress is being done on developing a method based on tensor similarity measures. Using a measure which is sensitive to differences in size, shape and orientation of the diffusion tensors in

the different voxels, the different nuclei could be distinguished. Then coupling level sets could be used to evolve the surfaces in different regions.

Finally, I believe the present method should be compared to other segmentation methods in order to find the most accurate way of accomplishing our purpose. Combining it with some other method could help to overtake the drawbacks mentioned above.

The improvement of the segmentation of the thalamus with DT-MRI and, further on, of its nuclei, would allow non-invasive morphological analysis and would be very helpful in the detection and further treatment of Parkinson's disease, schizophrenia and other illnesses which manifest in changes of the thalamic structure.

No matter how small the step is, when collaborating with medicine, it is always a big progress.

Bibliography

[1] D. Alexander. "Note Indices of Shape and Similarity for Diffusion Tensors". Department of Computer Science, UCL (University College London). October, 2000.

[2] P.J. Basser, D.K. Jones. "Diffusion-tensor MRI: theory, experimental design and data analysis - a technical review". NMR in Biomedicine, 15, 456-467. 2002.

[3] P.J. Basser, J. Mattinello, D. Le Bihan, "MR Diffusion Tensor Spectography and Imaging". Biophys. J. 66, 259-267, 1994.

[4] P.J. Basser, C. Pierpaoli. "A simplified method to measure the diffusion tensor from seven MR images". Magnetic Resonance Med, Vol. 39, 928-934, 1998.

[5] Vincent Caselles, Ron Kimmel, Guillermo Sapiro. "Geodesic Active Contours". International Journal of Computer Vision, 22(1): 61-79, 1997

[6] Patric Hagmann. "Statistical mapping of brain connectivity from diffusion tensor data". Technical Report. Swiss Federal Institute of Technology (EPFL), 2001.

[7] L. Jonasson, P. Hagmann, X. Bresson, O. Cuisenaire, R. Meuli, J.-P. Thiran. "White Matter Mapping in DT-MRI using Geometric Flows". 9th International Workshop on Computer Aided Systems Theory, Eurocast 2003, 2003.

[8] Lisa Jonasson, Xavier Bresson, Patric Hagmann, Olivier Cuisenaire, Jean-Philippe Thiran. "White Matter Mapping in DT-MRI using Geometric Flows". Technical Report. Swiss Federal Institute of Technology (EPFL), 2003

[9] Satyanad Kichenassamy, Arun Kumar, Peter Olver, Allen Tannenbaum, Anthony Yezzi Jr. "Gradient Flows and Geometric Active Contour Models". In International Conference on Computer Vision, pages 810-815, 1995.

[10] D. LeBihan, J.F. Mangin, C. Poupon, C.A. Clark, S. Pappata, N. Molko, H. Chabriat. "Diffusion tensor imaging: concepts and applications". Journal Magnetic Resonance Imaging, Vol. 13, No. 4, 534-546. 2001.

[11] Liana M. Lorigo, Olivier Faugueras, W.E.L. Grimson, Renaud Keriven, Ron Kikinis, Carl-Frederik Westin. "Co-Dimension 2 Geodesic Active Contours for MRA Segmentation". IPMI 1999, 126-139. 1999.

[12] S. Osher, J. Sethian. "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations". Journal of Computational Physics 79, 12-49. 1988

[13] Nikos K. Paragios. "Geodesic Active Regions and Level Set methods". Master Thesis (Chapter 3). January 2000.

[14] Nikos Paragios, Rachid Deriche. "Coupled Geodesic Active Region for Image Segmentation: A Level Set Approach". Technical report. INRIA. October 1999.

[15] Claudio Pollo. "Statistical Mapping of the Subthalamic Region from Diffusion Tensor Magnetic Resonance Imaging". Postgraduate Diploma in Biomedical Engineering. Swiss Federal Institute of Technology (EPFL), September 2003.

[16] J.A. Sethian. "Level set methods and fast marching methods: Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science". Cambridge University Press. 1999.

[17] Jasjit Suri, Kecheng Liu, Sameer Singh, Swamy Laxminarayan, Xiaolan Zeng, Laura Reden. "Shape Recovery Algorithms Using Level Sets in 2D/3D Medical Imagery: A State-of-the-Art Review". IEEE Transactions on Information Technology in Biomedicine, Vol. 6, No. 1, March 2002.

[18] Mette R. Wiegell, David S. Tuch, Henrik B. W. Larsson, Van J. Wedeen. "Automatic segmentation of thalamic nuclei from diffusion tensor magnetic resonance imaging". Neuroimage, 19: 391-401, 2003.