Uncertainty-driven Non-parametric Knowledge-based Segmentation: The Corpus Callosum Case

Maxime Taron[2], Nikos Paragios[2], and Marie-Pierre Jolly[1]

[1] Imaging & Visualization Department,
Siemens Corporate Research, Princeton, NJ, USA
755 College Road East, Princeton, NJ 08540, USA

[2] CERTIS - Ecole Nationale des Ponts et Chaussees,
Champs-sur-Marne, France
6 - 8 Avenue Blaise Pascal, 77455 Champs-sur-Marne, France

Abstract. In this paper we propose a novel variational technique for the knowledge based segmentation of two dimensional objects. One of the elements of our approach is the use of higher order implicit polynomials to represent shapes. The most important contribution is the estimation of uncertainties on the registered shapes, which can be used with a variable bandwidth kernel-based non-parametric density estimation process to model prior knowledge about the object of interest. Such a non-linear model with uncertainty measures is integrated with an adaptive visual-driven data term that aims to separate the object of interest from the background. Promising results obtained for the segmentation of the corpus callosum in MR mid-sagittal brain slices demonstrate the potential of such a framework.

1 Introduction

Over the last decade, shape-based segmentation methods have become more and more common. First introduced in 1995, active shape models (ASM) and active appearance models (AAM) [3] have been very popular tools for the segmentation of anatomical structures in medical images [2, 5, 1, 10]. More recently, principal component analysis (PCA) has also been applied to distance transforms for an implicit representation of shapes [9]. Shape-based segmentation is usually equivalent to recovering a geometric structure which is both highly probable in the model space and well aligned with strong features in the image. The advantage of the shape based methods over classical deformable templates [11] is that they allow the deformation process to be constrained to remain within the space of allowable shapes. These methods have proven to be a good compromise between complexity and shape generalization. However, since modeling is performed after registration, errors in the registration can be propagated into the model space. Furthermore, the assumption of Gaussian shape models might be a little restrictive.

In this paper, shapes are represented implicitly using the distance transform. To generate a model of the structure of interest, we register shape examples using a spline based free form deformation. The main contribution of this paper is the derivation of

a measure representing the uncertainty of the registration at the zero iso-surface. After dimensionality reduction, these measures are combined with a variable bandwidth kernel-based approach to derive a density function that models the family of shapes under consideration. Given a new image, the segmentation process is expressed in a variational level set framework [14] where the energy function makes use of the uncertainties of the registration between the deformed shape which aligns to the image features and the model.

We apply our novel modeling and segmentation technique to the case of the corpus callosum. The corpus callosum is a thick bundle of nerve fibers that connect the left and right hemispheres in the brain. It is believed to be responsible for balancing the load of learning tasks across each hemisphere, making each specialized in certain tasks. While not learning, it is responsible for routing most of the communication between the two hemispheres. This is the reason why a surgical procedure has been developed to cut the corpus callosum in patients with severe epilepsy for which drug treatment is ineffective. In addition, several studies indicate that the size and shape of the corpus callosum is related to various types of brain dysfunction such as dyslexia [4] or schizophrenia [6]. Therefore, neurologists are interested in looking at the corpus callosum and analyzing its shape. Magnetic resonance imaging (MRI) is a safe and non-invasive tool to image the corpus callosum. Since manual delineation can be very time consuming, we demonstrate how our algorithm can be used to segment the corpus callosum on mid-sagittal MR slices.

The remainder of this paper is organized as follows. In Section 2, we introduce registration with uncertainties and probabilistic modeling to describe the corpus callosum structure. The segmentation component combining data and shape terms is described in Section 3. Experimental results are presented in Section 4. Finally, conclusions and future directions are discussed in Section 5.

2 Shape Representation through Implicit Polynomials

Let us consider a training set $\{C_1, C_2, ..., C_{\mathrm{N}}\}$ of shapes representing the structure of interest. The model building task consists of recovering a probabilistic representation of this set. In order to remove all the pose variation from the training set, all shapes have to be registered to a common pose with respect to an affine transformation. Then a reference model $C_{\mathcal{M}}$ is locally registered to every sample of the training set C_i using implicit polynomials. We will first describe the registration process and the calculation of uncertainties on the registered model. The uncertainty measures represent the allowable range of variations in the deformations of the model that still match C_i. Then we describe the way these uncertainties are used in the estimation of probability density function of the deformations.

2.1 Registration through implicit polynomials

In the classical ASM the initial step is used to recover explicit correspondence between the discretized contour of the model shape and the training examples. In the present

framework, the model shape is non rigidly registered to every sample from the training, and the statistical shape model is actually built on the parameters of the recovered transformation.

Shapes C_i are represented in an implicit fashion using the Euclidean distance transform [9, 15] . In the 2D case, we consider the function defined on the image domain Ω :

$$\phi_{C_i}(\mathbf{x}) = \begin{cases} 0, & \mathbf{x} \in C_i \\ +\mathcal{D}(\mathbf{x}, C_i), & \mathbf{x} \in \mathcal{R}_{C_i} \\ -\mathcal{D}(\mathbf{x}, C_i), & \mathbf{x} \notin \mathcal{R}_{C_i} \end{cases}$$

where $\mathcal{R}_{C_i}$ is the region enclosed by C_i. Such a space is invariant to translation, rotation and can also be modified to account for scale variations. This representation has already been used along with simple criteria like sum of squared differences to address similarity registration [15] or mutual information for affine transformations [7].

The retained framework for density estimation does not put any constraint on the reference model used for registration. In practice we choose a shape characteristic of the object to segment. Without loss of generality, we can choose for $C_{\mathcal{M}}$ a smoothed version of C_1. All contours of the training set are now registered to $C_{\mathcal{M}}$ with respect to an affine transform and from now on, we will denote $\{C_1, C_2, ..., C_{\mathrm{N}}\}$ as the globally registered training set.

Local registration is crucial to model building. To this end one would like to recover an invertible transformation (diffeomorphism) $\mathcal{L}_{\mathbf{\Theta}_i}$ parameterized by a vector $\mathbf{\Theta}_i$ that creates a one to one mapping between each contour of the training set C_i and the model $C_{\mathcal{M}}$:

$$\mathcal{L}_{\mathbf{\Theta}_i} : \mathcal{R}^2 \rightarrow \mathcal{R}^2 \text{ and } \mathcal{L}_{\mathbf{\Theta}_{\mathrm{i}}}(\mathrm{C}_{\mathcal{M}}) \approx \mathrm{C_i}$$

When $\mathcal{L}_{\mathbf{\Theta}}$ is chosen as a 2D polynomial with coefficients $\mathbf{\Theta}$ in an appropriate basis, the expression $\phi \circ \mathcal{L}_{\mathbf{\Theta}}$ inherits the invariance properties of implicit polynomials, i.e. linear transformations applied to $\mathbf{\Theta}$ are related to linear transformations applied to the data space. In the present paper, we used a simple polynomial warping technique to address the demand of local registration: the free form deformations method (FFD) [16]. The essence of FFD is to deform an object by manipulating a regular control lattice overlaid on its embedding space. We use a cubic B-spline FFD to model the local transformation $\mathcal{L}$. Consider the $M \times N$ square lattice of points, $[\{\mathbf{P}^0_{m,n}\}; (m, n) \in [1; M] \times [1; N]]$. In this case the vector of parameters $\mathbf{\Theta}$ defining the transformation $\mathcal{L}$ is the displacement coordinates of the control lattice. $\mathbf{\Theta}$ has size $2MN$:

$$\mathbf{\Theta} = \{\delta\mathbf{P}^x_{m,n}, \delta\mathbf{P}^y_{m,n}\}; (m, n) \in [1; M] \times [1; N]$$

The motion of a pixel $\mathbf{x}$ given the deformation of the control lattice, is defined in terms of a tensor product of Cubic B-splines [17]. As FFD is linear in the parameter $\mathbf{\Theta} = \delta\mathbf{P}$, it can be expressed in a compact form by introducing $\mathcal{X}(\mathbf{x})$ a $[2 \times 2MN]$ matrix:

$$\mathcal{L}(\mathbf{\Theta}; \mathbf{x}) = \sum\sum B_i(u)B_j(v)(\mathbf{P}^0_{i,j} + \delta\mathbf{P}_{i,j}) = \mathbf{x} + \mathcal{X}(\mathbf{x})\mathbf{\Theta}$$

where (u, v) are the coordinates of $\mathbf{x}$, and (B_i, B_j) the cubic B-spline basis functions.

Local registration now is equivalent to finding the best lattice configuration such that the overlaid structures coincide. Since structures correspond to distance transforms

of globally aligned shapes, the sum of squared differences (SSD) can be considered as the data-driven term to recover the deformation field $\mathcal{L}(\mathbf{\Theta};\mathbf{x})$ between the element C_i of the training set and the model $C_{\mathcal{M}}$ (corresponding respectively to the distance transform ϕ_i and $\phi_{\mathcal{M}}$)

$$E_{\text{data}}(\mathbf{\Theta}) = \iint_{\Omega} \chi_\alpha(\phi_i(\mathbf{x}))\left[\phi_i(\mathcal{L}(\mathbf{\Theta};\mathbf{x})) - \phi_{\mathcal{M}}(\mathbf{x})\right]^2 d\mathbf{x} \tag{1}$$

with $\chi_\alpha(\phi_i(\mathbf{x}))$ being an indicator function that defines a band of width α around the contour. In order to further preserve the regularity of the recovered registration, one can consider an additional smoothness term on the deformation field $\delta\mathcal{L}$. We consider a computationally efficient smoothness term :

$$E_{\text{smooth}}(\mathbf{\Theta}) = \iint_{\Omega} \left(\left|\mathcal{L}_{xx}(\mathbf{\Theta};\mathbf{x})\right|^2 + 2\left|\mathcal{L}_{xy}(\mathbf{\Theta};\mathbf{x})\right|^2 + \left|\mathcal{L}_{yy}(\mathbf{\Theta};\mathbf{x})\right|^2 \right) d\mathbf{x}.$$

The data-driven term and the smoothness constraint component can now be integrated to recover the local deformation component through the calculus of variations. We denote as $\mathbf{\Theta}_i$ the reached minimum.

However, one can claim that the local deformation field is not sufficient to characterize the registration between two shapes. Data is often corrupted by noise so that the registration retrieved using a deformable model may be imprecise. Therefore, recovering uncertainty measurements [8] that do allow the characterization of an allowable range of variation for the registration process is an eminent condition of accurate shape modeling.

2.2 Uncertainty estimation on registered shapes

We aim to recover uncertainties on the vector $\mathbf{\Theta}$ in the form of a $[2MN \times 2MN]$ covariance matrix by adapting a method initially introduced in [18]. We are considering the quality of the local registration on shapes, that is the zero levelset of the distance transform. Therefore, E_{data} is formulated in the limit case where α the size of the limited band around the model shape tends to 0. The data term of the energy function (1) can now be expressed as:

$$E_{\text{data}}(\mathbf{\Theta}) = \oint_{\mathcal{C}_{\mathcal{M}}} \phi_i^2(\mathcal{L}(\mathbf{\Theta};\mathbf{x}))d\mathbf{x} = \oint_{\mathcal{C}_{\mathcal{M}}} \phi_i^2(\mathbf{x}')d\mathbf{x},$$

where we denote $\mathbf{x}' = \mathcal{L}(\mathbf{\Theta}_i;\mathbf{x})$. Let us consider $\mathbf{q}$ to be the closest point from $\mathbf{x}'$ located on C_i. As ϕ_i is assumed to be a Euclidean distance transform, it also satisfies the condition $\|\nabla\phi_i(\mathbf{x}')\| = 1$. Therefore one can express the values of ϕ_i at the first order in the neighborhood of $\mathbf{x}'$ in the following manner :

$$\begin{aligned}\phi_i(\mathbf{x}' + \delta\mathbf{x}') &= \phi_i(\mathbf{x}') + \delta\mathbf{x}' \cdot \nabla\phi_i(\mathbf{x}') + \circ(\delta\mathbf{x}') \\ &= (\mathbf{x}' + \delta\mathbf{x}' - \mathbf{q}) \cdot \nabla\phi_i(\mathbf{x}') + \circ(\delta\mathbf{x}')\end{aligned}$$

This local expression of ϕ_i with a dot product reflects the condition that a point to curve distance was adopted. Under the assumption that E_{data} is small when reaching the

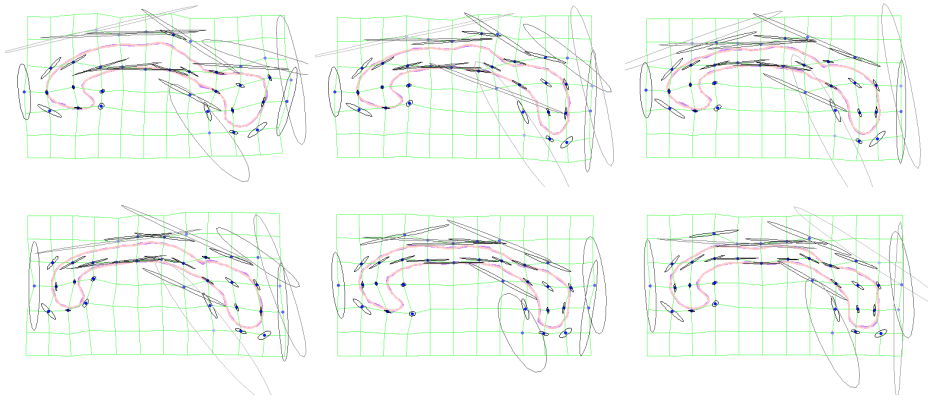

Fig. 1. *Implicit higher order polynomials and registration of corpus callosum with uncertainty estimates (this figure should be seen in color).*

optimum, we can write the classical second order approximation of quadratic energy in the form:

$$E_{\text{data}}(\mathbf{\Theta}) = \oint_{\mathcal{C}_\mathcal{M}} \left[(\mathbf{x}' - \mathbf{q}) \cdot \nabla\phi_i(\mathbf{x}')\right]^2 = \oint_{\mathcal{C}_\mathcal{M}} \left[(\mathbf{x} + \mathcal{X}(\mathbf{x})\mathbf{\Theta} - \mathbf{q}) \cdot \nabla\phi_i(\mathbf{x}')\right]^2$$

Localizing the global minimum of an objective function E is equivalent to finding the major mode of a random variable with density $\exp(-E/\beta)$. The coefficient β corresponds to the allowable variation in the energy value around the minimum. In the present case of a quadratic energy (and therefore Gaussian random variable), the covariance and the Hessian of the energy are directly related by $\Sigma_{\mathbf{\Theta}_i}^{-1} = H_{\mathbf{\Theta}_i}/\beta$. This leads to the following expression for the covariance :

$$\Sigma_{\mathbf{\Theta}_i}^{-1} = \frac{1}{\beta} \oint_{\mathcal{C}_\mathcal{M}} \mathcal{X}(\mathbf{x})^T . \nabla\phi_i(\mathbf{x}') . \nabla\phi_i(\mathbf{x}')^T . \mathcal{X}(\mathbf{x}) d\mathbf{x}$$

In the most general case one can claim that the matrix $H_{\mathbf{\Theta}}$ is not invertible because the registration problem is under-constrained. Then, additional constraints have to be introduced towards the estimation of the covariance matrix of $\mathbf{\Theta}_i$ through the use of an arbitrarily small positive parameter γ :

$$E(\mathbf{\Theta}) = \oint_{\mathcal{C}_\mathcal{M}} \left[(\mathbf{x} + \mathcal{X}(\mathbf{x})\mathbf{\Theta} - \mathbf{q}) \cdot \nabla\phi_i(\mathbf{x}')\right]^2 d\mathbf{x} + \gamma\, \mathbf{\Theta}^T \mathbf{\Theta}$$

This leads to the covariance matrix for the parameter estimate :

$$\Sigma_{\mathbf{\Theta}_i} = \beta \left(\oint_{\mathcal{C}_\mathcal{M}} \mathcal{X}(\mathbf{x})^T . \nabla\phi_i(\mathbf{x}') \nabla\phi_i(\mathbf{x}')^T \mathcal{X}(\mathbf{x}) d\mathbf{x} + \gamma \mathbf{I} \right)^{-1} \tag{2}$$

2.3 Hybrid kernel based density function and kernel selection

Now that all shapes of the training set have been aligned, standard statistical techniques like PCA or ICA could be applied to recover linear Gaussian models. But in the most general case shapes that refer to objects of particular interest vary non-linearly and therefore the assumption of simple parametric models likes Gaussian is rather unrealistic. Therefore within our approach we propose a non-parametric form of the probability density function.

Let $\{\mathbf{\Theta}_1 ... \mathbf{\Theta}_N\}$ be the N vectors of parameters associated with the registration of the N sample of the training set. Considering that this set of vectors is a random sample drawn from the density function f describing the shapes, the fixed bandwidth kernel density estimator consists of:

$$\hat{f}(\mathbf{\Theta}) = \frac{1}{\mathrm{N}} \sum_{i=1}^{\mathrm{N}} \frac{1}{\|\mathbf{H}\|^{1/2}} \mathrm{K}\left(\mathbf{H}^{-1/2}(\mathbf{\Theta} - \mathbf{\Theta}_i)\right)$$

where $\mathbf{H}$ is a symmetric definite positive (bandwidth matrix) and K denote the centered Gaussian kernel with identity covariance. Fixed bandwidth approaches often produce under-smoothing in areas with sparse observations and over-smoothing in the opposite case.

Kernels of variable bandwidth can be used to encode such a condition and provide a structured way for utilizing the variable uncertainties associated with the sample points. In the literature, kernel density estimation methods that do rely on varying bandwidths are generally referred to as adaptive kernels. Density estimation is performed with kernels whose bandwidth adapts to the sparseness of the data [19].

In the present case, the vectors $\{\mathbf{\Theta}_i\}$ come along with associated uncertainties $\{\Sigma_i\}$. Furthermore, the point $\mathbf{\Theta}$ where the density function is evaluated corresponds to a deformed model, and therefore is also associated to a measure of uncertainty Σ. In order to account for the uncertainty estimates both on the sample points themselves as well as on the estimation point, we adopt a hybrid estimator [12].

$$\begin{aligned}\hat{f}_H(\mathbf{\Theta}, \Sigma) &= \frac{1}{N} \sum_{i=1}^{N} \mathcal{K}(\mathbf{\Theta}, \Sigma, \mathbf{\Theta}_i, \Sigma_i) \\ &= \frac{1}{N} \sum_{i=1}^{N} \frac{1}{\|\mathbf{H}(\Sigma_{\mathbf{\Theta}}, \Sigma_{\mathbf{\Theta}_i})\|^{1/2}} \mathrm{K}(\mathbf{H}(\Sigma_{\mathbf{\Theta}}, \Sigma_{\mathbf{\Theta}_i})^{-1/2}(\mathbf{\Theta} - \mathbf{\Theta}_i)\end{aligned}$$

where we choose for the bandwidth function: $\mathbf{H}(\Sigma_{\mathbf{\Theta}}, \Sigma_{\mathbf{\Theta}_i}) = \Sigma_{\mathbf{\Theta}} + \Sigma_{\mathbf{\Theta}_i}$ as proposed in [12]. Using this estimator, the density decreases more slowly in directions of large uncertainties when compared to the other directions.

This metric can now be used to assess the probability of a new sample being part of the training set and account for the non-parametric form of the observed density. However, the computation is time consuming because it leads to the calculation of large matrix inverses. Since the cost is linear in the number of samples in the training set, there is an eminent need to decrease its cardinality by selecting the most representative kernels.

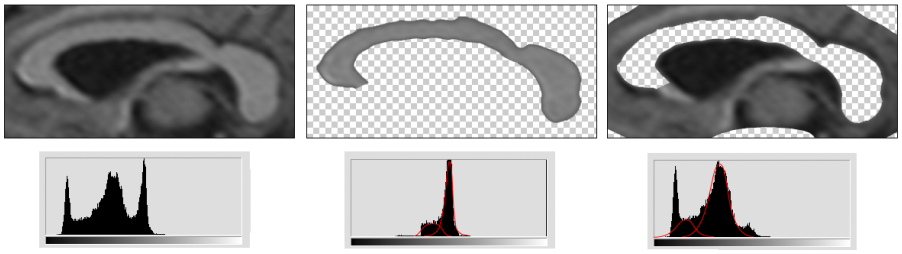

Fig. 2. *Histograms of the corpus callosum and the background area. The of a use gaussian mixture to model the corpus callosum and background intensity distribution in MR is appropriate (this figure should be seen in color).*

The maximum likelihood criterion expresses the quality of approximation from the model to the data. We use a recursive sub-optimal algorithm to select kernels and therefore build a compact model that maximizes the likelihood of the whole training set.

Consider a set $\mathcal{Z}_{\mathrm{K}} = \{X_1, X_2, \ldots, X_{\mathrm{K}}\}$ of K kernels extracted from the training set with mean and uncertainties estimates $\{X_i = (\mathbf{\Theta}_i \Sigma_i)\}_{i=1}^{K}$. The log likelihood of the entire training set according to this model is:

$$C_{\mathrm{K}} = \sum_{i=1}^{N} \log \left(\frac{1}{K} \sum_{(\mathbf{\Theta}_j, \sigma_j) \in \mathcal{Z}_{\mathrm{K}}} \mathcal{K}(\mathbf{\Theta}_j, \Sigma_i, \mathbf{\Theta}_i, \Sigma_i) \right)$$

A new kernel X_{K+1} is extracted from the training set as the one maximizing the quantity $C_{\mathrm{K}+1}$ associated with $\mathcal{Z}_{\mathrm{K}+1} = \mathcal{Z}_{\mathrm{K}} \bigcup X_{K+1}$. The same kernel may be chosen several times in order to preserve a increasing sequence C_K. Consequently the selected kernels X_i in $\mathcal{Z}_{\mathrm{K}}$ are also associated with a weight factor w_i. Once such a selection has been completed, the hybrid estimator is evaluated over $\mathcal{Z}_{\mathrm{K}}$:

$$\hat{f}_H(\mathbf{\Theta}, \Sigma) = \frac{1}{N} \sum_{(\mathbf{\Theta}_i, \sigma_i, w_i) \in \mathcal{Z}_{\mathrm{K}}} w_i \mathcal{K}(\mathbf{\Theta}, \Sigma, \mathbf{\Theta}_i, \Sigma_i) \tag{3}$$

3 Shape based Segmentation applied to the Corpus Callosum

Let us consider an image $\mathcal{I}$ where the corpus callosum structure is present and is to be recovered. Recall that we now have a model of the corpus callosum: a shape that can be transformed using an affine transformation and a FFD, and a measure of how well the deformed shape belongs to the family of trained shapes.

Let $\phi_{\mathcal{M}}$ be the distance transform of the reference model. Segmentation consists of globally and locally deforming $\phi_{\mathcal{M}}$ towards delineating the corpus callosum in $\mathcal{I}$. Let $\mathcal{A}$ be an affine transformation of the model and $\mathcal{L}(\mathbf{\Theta})$ its local deformation using FFD as previously introduced.

For now, we assume that the visual properties of the corpus callosum $\pi_{\mathrm{cor}}()$ as well as the ones of the local surrounding area $\pi_{\mathrm{bck}}()$ are known. Then segmentation of the

corpus callosum is equivalent to the minimization of the following energy with respect to the parameters $\mathbf{\Theta}$ and $\mathcal{A}$:

$$E_{\text{image}}(\mathcal{A}, \mathbf{\Theta}) = -\iint_{\mathcal{R}_{\mathcal{M}}} \log\left[\pi_{\text{cor}}\left(\mathcal{I}\left(\mathcal{A}(\mathcal{L}(\mathbf{\Theta}; \mathbf{x}))\right)\right)\right] d\mathbf{x} - \iint_{\Omega - \mathcal{R}_{\mathcal{M}}} \log\left[\pi_{\text{bkg}}\left(\mathcal{I}\left(\mathcal{A}(\mathcal{L}(\mathbf{\Theta}; \mathbf{x}))\right)\right)\right] d\mathbf{x}$$

where $\mathcal{R}_{\mathcal{M}}$ denotes the inside of $C_{\mathcal{M}}$. However, the direct calculation of variations involves image gradient and often converges to erroneous solutions due to the discretization of the model domain. In that case, we change the integration domain to the image by implicitly introducing the inverse transformation (see Appendix). A bimodal partition in the image space is now to be recovered. The definition of this domain $\mathcal{R}_{\text{cor}}$ depends upon the parameters of the transformation $[\mathcal{A}, \mathbf{\Theta}]$ as :

$$\mathcal{R}_{\text{cor}} = \mathcal{A}(\mathcal{L}(\mathbf{\Theta}, \mathcal{R}_{\mathcal{M}})) \text{ and } \mathbf{y} = \mathcal{A}(\mathcal{L}(\mathbf{\Theta}, \mathbf{x}))$$

The actual image term of the energy to be minimized then becomes:

$$E_{\text{image}}(\mathcal{A}, \mathbf{\Theta}) = -\iint_{\mathcal{R}_{\text{cor}}} \log\left[\pi_{\text{cor}}\left(\mathcal{I}\left(\mathbf{y}\right)\right)\right] d\mathbf{y} - \iint_{\Omega - \mathcal{R}_{\text{cor}}} \log\left[\pi_{\text{bkg}}\left(\mathcal{I}\left(\mathbf{y}\right)\right)\right] d\mathbf{y} \tag{4}$$

where statistical independence is considered at the pixel as well as hypotheses level. In practice the distributions of the corpus callosum as well as the ones of the surrounding region $[\pi_{\text{cor}}, \pi_{\text{bkg}}]$ can be recovered in an incremental fashion using the Mumford-Shah principle [13]. In the present case, each distribution is estimated by fitting a mixture of Gaussians to the image histogram using an Expectation-Maximization algorithm (Fig. 2).

The shape based energy term, making use of the non parametric framework introduced earlier is also locally influenced by a covariance matrix of uncertainty calculated on the transformed model. This covariance matrix is computed in a fashion similar to (2) with the difference that it may only account for the linear structure of the transformed model and therefore allow variations of $\mathbf{\Theta}$ that creates tangential displacements of the contour:

$$\Sigma_{\mathbf{\Theta}}^{-1} = \frac{1}{\beta} \oint_{\mathcal{C}_{\mathcal{M}}} \mathcal{X}(\mathbf{x})^T \nabla\tilde{\phi}_{\mathcal{M}}(\mathbf{x}') \nabla\tilde{\phi}_{\mathcal{M}}(\mathbf{x}')^T \mathcal{X}(\mathbf{x}) d\mathbf{x}$$

where $\tilde{\phi}_{\mathcal{M}}$ is the transformation of $\phi_{\mathcal{M}}$ under the deformation $\mathcal{A}(\mathcal{L}(\mathbf{\Theta}))$. Direct computation leads to:

$$\nabla\tilde{\phi}_{\mathcal{M}}(\mathbf{x}') = \text{com}\left[\frac{d}{d\mathbf{x}}(\mathcal{L}(\mathbf{\Theta}, \mathbf{x}))\right] . \nabla\phi_{\mathcal{M}}(\mathbf{x})$$

where 'com' denotes the matrix of cofactors. Then we introduce the shape based energy term using the same notations as in (3) as:

$$E_{\text{shape}}(\mathbf{\Theta}, \Sigma_{\mathbf{\Theta}}) = -\log(\hat{f}_H(\mathbf{\Theta}, \Sigma))$$

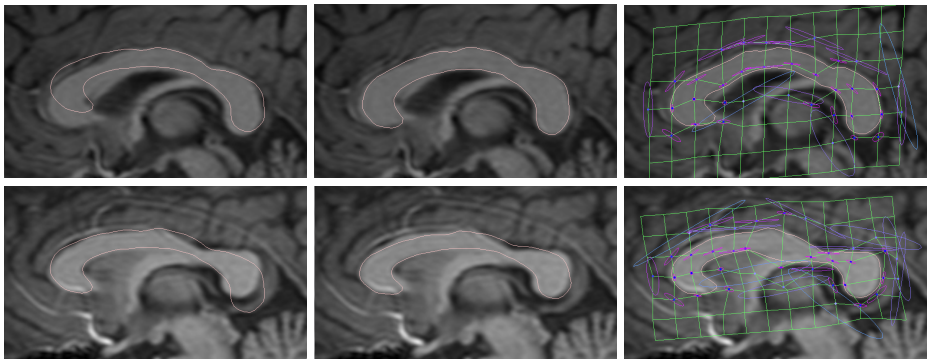

(a) (b) (c)

Fig. 3. *Segmentation with uncertainties estimates of the corpus callosum; (a) Automatic rough positioning of the model, (b) segmentation through affine transformation of the model (c) segmentation using the local deformation of the FFD grid and uncertainties estimates on the registration/segmentation process (this figure should be seen in color).*

The global energy is minimized with respect to the parameters of $\mathcal{A}$ and $\boldsymbol{\Theta}$ through the computation of variations on $E = E_{\text{image}} + E_{\text{shape}}$ and implemented using a standard gradient descent.

4 Experimental Results

We have applied our method to the segmentation of the corpus callosum in MR mid-sagittal brain slices.

The first step was to build a model of the corpus callosum. Minimization of the registration energy is performed using gradient descent. In parallel, we successively refine the size of the band α around the contour (from .3 to .05 times the size of the shape), while we increase the complexity of the diffeomorphism (from an affine transformation to an FFD with a regular $[7 \times 12]$ lattice).

Fig. 1 shows examples of FFD deformations along with uncertainty ellipses. These ellipses are the representation of the 2D conic obtained when projecting the covariance matrix $\Sigma_{\boldsymbol{\Theta}}$ (of size 168×168) on the control points. It therefore does not allow us to represent the correlations between control points.

The segmentation process is initialized by positioning the initial contour according to the method proposed in [10]. Energy minimization is performed through gradient descent, while the PDF π_{cor} and π_{bkg} are estimated by mixtures of Gaussians. Fig. 2 shows the histogram of a typical image of the corpus callosum. The figure illustrates how well mixtures of two Gaussian distributions can represent the individual histograms for the corpus callosum and the background, respectively. Segmentation results are presented in (Fig. 3 and Fig. 4) along with the associated uncertainties. In Fig. 3, we demonstrate the individual steps of the segmentation process: the left most image shows the automatic initialization of the contour, the middle image shows the contour after the affine transformation has been recovered, and the right image shows the local deformations. Fig. 4 shows additional results and illustrates that our method can handle a wide variety

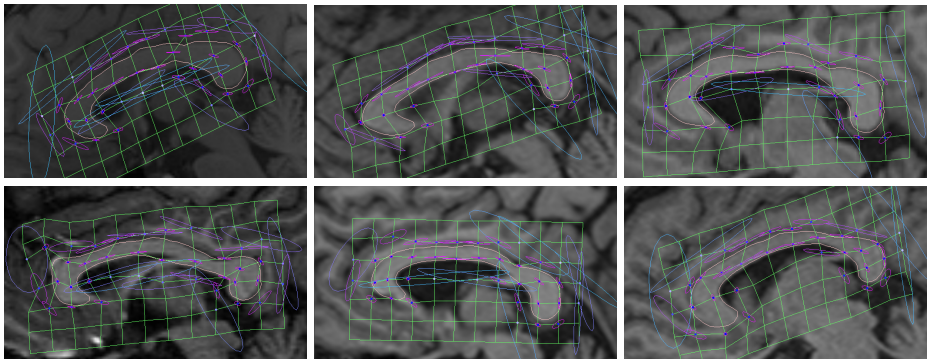

Fig. 4. *Additional segmentation results with uncertainty measures.*

of shapes for the corpus callosum as well as large variations in image contrast. It can be seen that the results in the bottom left image is not perfect. In general, failures may be due to the fact that the shape constraint is not strong enough and the contrast in the image dominates the deformation. Also, it might be that the shape of this particular corpus callosum cannot be captured with the current PDF because it has been reduced to only 10 kernels.

5 Conclusions

In this paper we have introduced a novel method to account for prior knowledge in the segmentation process using non-parametric variable bandwidth kernels that are able to account for errors in the registration and the segmentation process. We have shown that the method can generate a very good model of the object of interest and produce very good segmentation results.

However the method of kernel selection presented in Section 3 has shown some limitation in practice. Therefore there is a strong need to build more efficient and compact estimators of the shape variation PDF which account for these uncertainty measures. It is also important to note that this method can be extended to higher dimensions. Building models in 3D and segmenting objects of large variability is the next step of our research work.

The covariance matrices of uncertainty $\Sigma_{\boldsymbol{\Theta}}$ are very sparse. Indeed, while using regular FFD, the influence of every grid point is local and therefore many cross correlation coefficients are null. Different types of B-spline deformations using an irregular positioning of control points (but dependent on the model) will be tried to address this issue and therefore reduce the dimensionality of the problem.

Last, but not least, introduction of uncertainties directly measured in the image as part of the segmentation process will provide local measures of confidence and could be considered as a major breakthrough in the area of knowledge-based object extraction.

Appendix

In this section we give some further exploration of the calculus of the derivative on the energy term E_{image}. We need first to introduce the Heaviside distribution which we note H and the inverse diffeomorphism of $\mathcal{A} \circ \mathcal{L}(\mathbf{\Theta})$ which we note $\mathcal{G}(\mathbf{\Theta})$. This diffeomorphism therefore verifies:

$$\mathcal{A}(\mathcal{L}(\mathbf{\Theta}, \mathcal{G}(\mathbf{\Theta}, \mathbf{y}))) = \mathbf{y} \tag{5}$$

For simpler notation purpose we also pose:

$$D(\mathbf{x}, \mathbf{y}) = -H(\phi_{\mathcal{M}}(\mathbf{x}))\log(\pi_{\text{cor}}(\mathcal{I}(\mathbf{y}))) - (1 - H(\phi_{\mathcal{M}}(\mathbf{x})))\log(\pi_{\text{bkg}}(\mathcal{I}(\mathbf{y})))$$

Then the image term of the energy (eq. 4) can be rewritten as:

$$E_{\text{image}}(\mathbf{\Theta}) = \int_{\Omega} D(\mathcal{G}(\mathbf{\Theta}, \mathbf{y}), \mathbf{y}) d\mathbf{y}$$

When differentiating Eq. (5) with respect to $\mathbf{\Theta}$ and substituting the expression obtained for $d\mathcal{G}/d\mathbf{\Theta}$ into the expression of $dE_{\text{image}}(\mathbf{\Theta})/d\mathbf{\Theta}$, we get the following:

$$\frac{dE_{\text{image}}(\mathbf{\Theta})}{d\mathbf{\Theta}} = \\ -\int_{\Omega} \frac{\partial D}{\partial \mathbf{x}^T}(\mathcal{G}(\mathbf{\Theta}, \mathbf{y}), \mathbf{y}) \left[\frac{\partial(\mathcal{A} \circ \mathcal{L})}{\partial \mathbf{x}^T}(\mathcal{G}(\mathbf{\Theta}, \mathbf{y}), \mathbf{\Theta})\right]^{-1} \frac{\partial(\mathcal{A} \circ \mathcal{L})}{\partial \mathbf{\Theta}^T}(\mathcal{G}(\mathbf{\Theta}, \mathbf{y}), \mathbf{\Theta}) d\mathbf{y}$$

Now changing the integration variable according to the diffeomorphism $\mathbf{x} = \mathcal{G}(\mathbf{\Theta}, \mathbf{y})$

$$\frac{dE_{\text{image}}(\mathbf{\Theta})}{d\mathbf{\Theta}} = -\int_{\Omega} \frac{\partial D}{\partial \mathbf{x}^T}(\mathbf{x}, \mathcal{A}(\mathcal{L}(\mathbf{\Theta}, \mathbf{x})))\text{com}\left(\frac{\partial(\mathcal{A} \circ \mathcal{L})}{\partial \mathbf{x}^T}(\mathbf{x}, \mathbf{\Theta})\right)^T \frac{\partial(\mathcal{A} \circ \mathcal{L})}{\partial \mathbf{\Theta}^T}(\mathbf{x}, \mathbf{\Theta}) d\mathbf{x}$$

where 'com' denotes the matrix of cofactors. When calculating explicitly the partial derivative of D with respect to its first variable, this integral further simplifies into a curve integral along the reference model:

$$\frac{dE_{\text{image}}(\mathbf{\Theta})}{d\mathbf{\Theta}} = \\ -\oint_{C_{\mathcal{M}}} \tilde{D}(\mathcal{A}(\mathcal{L}(\mathbf{\Theta}, \mathbf{x}))) \left[\text{com}\left(\frac{\partial(\mathcal{A} \circ \mathcal{L})}{\partial \mathbf{x}^T}(\mathbf{x}, \mathbf{\Theta})\right) . \nabla \phi_{\mathcal{M}}(\mathbf{x})\right]^T \frac{\partial(\mathcal{A} \circ \mathcal{L})}{\partial \mathbf{\Theta}^T}(\mathbf{x}, \mathbf{\Theta}) d\mathbf{x}$$

with $\tilde{D}$ defined as:

$$\tilde{D}(\mathbf{y}) = -\log(\pi_{\text{cor}}(\mathcal{I}(\mathbf{y}))) + \log(\pi_{\text{bkg}}(\mathcal{I}(\mathbf{y})))$$

This expression of the derivative refers only to the contour in the model space. Therefore there is no need to parse the entire image domain at every iteration of the gradient descent used in our implementation. Instead, we only scan the model contour at every iterations. Parsing of the images is only necessary when we reevaluate the parameters of the gaussian mixtures for π_{cor} and π_{bkg} (every 20 iteration).

References

1. J. G. Bosch, S. C. Mitchell, B.P.F. Lelieveldt, F. Nikland, O. Kamp, M. Sonka, and J. H. Reiber. Automatic segmentation of echocardiographic sequences by active appearance motion models. *IEEE Trans. Medical Imaging*, 21(11):1374–1383, 2002.
2. T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):255–266, 1994.
3. T. F. Cootes and C. J. Taylor. Statistical models of appearance for computer vision. *Technical Report, University of Manchester*, 2004.
4. J. Duncan, A. Wang, A. Amini, R. Greene, L. Kier, J. Gore, J. Holahan, S. Shaywitz, J. Fletcher, R. Bronen, and B. Shaywitz. A MRI-based study of the corpus callosum in dyslexic and normal children. *Neurology*, 1996.
5. N. Duta and M. Sonka. Segmentation and interpretation of MR brain images: An improved active shape model. *IEEE Trans. Medical Imaging*, 17(6):1049–1062, 1998.
6. M. Frumin, P. Golland, R. Kikinis, Y. Hirayasu, D. F. Salisbury, J. Hennen, C. C. Dickey, M. Anderson, F. A. Jolesz, W. E. L. Grimson, R. W. McCarley, and M. R. Shenton. Shape differences in the corpus callosum in first-episode schizophrenia and first-episode psychotic affective disorder. *American Journal of Psychiatry*, 159:866–868, 2002.
7. X. Huang, N. Paragios, and D. Metaxas. Registration of Structures in Arbitrary Dimensions: Implicit Representations, Mutual Information & Free-Form Deformations. Technical Report DCS-TR-0520, Division of Computer & Information Science, Rutgers University, 2003.
8. K. Kanatani. Uncertainty modeling and model selection for geometric inference. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(10):1307–1319, 2004.
9. M. Leventon, E. Grimson, and O. Faugeras. Statistical Shape Influence in Geodesic Active Controus. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages I:316–322, 2000.
10. A. Lundervold, N. Duta, T. Taxt, and A. Jain. Model-guided segmentation of corpus callosum in MR images. In *CVPR*, pages 1231–1238, 1999.
11. T. McInerney, G. Hamarneh, M. Shenton, and D. Terzopoulos. Deformable organisms for automatic medical image analysis. *Medical Image Analysis*, 6:251–266, 2002.
12. A. Mittal and N. Paragios. Motion-based background substraction using adaptive kernel density estimation. In *Computer Vision and Pattern Recognition*, volume 2, pages 302–309, 2004.
13. D. Mumford and J. Shah. Boundary detection by minimizing functionals. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 22–26, 1985.
14. S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed : Algorithms based on the Hamilton-Jacobi formulation. *Journal of Computational Physics*, 79:12–49, 1988.
15. N. Paragios, M. Rousson, and V. Ramesh. Matching Distance Functions: A Shape-to-Area Variational Approach for Global-to-Local Registration. In *European Conference on Computer Vision*, pages II:775–790, 2002.
16. D. Rueckert, L.I. Sonoda, C. Hayes, D. Hill, M. Leach, and D. Hawkes. Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18:712–721, 1999.
17. T. Sederberg and S. Parry. Free-form deformation of solid geometric models. *Proceedings SIGGRAPH '86*, 20:151–160, 1986.
18. C. Stewart, C.-L. Tsai, and B. Roysam. The dual bootstrap iterative closest point algorithm with application to retinal image registration. *IEEE Trans. Med. Img.*, 22:1379–1394, 2003.
19. M. Wand and M. Jones. *Kernel Smoothing*. Chapman & Hall, 1995.