

From Uncertainties to Statistical Model Building and Segmentation of the Left Ventricle

Maxime TARON Nikos PARAGIOS
Ecole Centrale Paris, MAS
92295 Chatenay-Malabry, France
{maxime.taron;nikos.paragios}@ecp.fr

Marie-Pierre Jolly
Siemens Corporate Research
08540 Princeton, USA
marie-pierre.jolly@siemens.com

Abstract

Reliable segmentation of the left ventricle is a long sought objective in medical imaging for automatic retrieval of anatomical and pathological measurements and detection of malfunctions. In this paper, we propose a novel model-constrained approach to address this task. The method is based on an implicit representation of the shape model used in a shape registration framework with a Thin Plate Spline transform to retrieve possible deformations. The main innovation of our approach resides in the use of uncertainties defined on the registered shape to augment the training set and improve the robustness of the statistical deformable model. We use ICA to reduce the dimensionality of the space of deformations and provide a good separation of the different deformable parts of the heart. Furthermore the estimation of uncertainties is also introduced in the segmentation process which is addressed in a variational framework where prior knowledge and visual support are considered. The method lead to very promising qualitative and quantitative experimental results in CT.

1. Introduction

Segmentation is one of the most critical problems in medical image analysis. The extraction of volumetric and geometric information about organs from images can be a valuable tool in computer-aided diagnosis. In the literature, both model-free and model-based methods have been considered.

Model-free methods do not make any prior assumptions on the geometric or iconic properties of the structure of interest. Therefore these algorithms aim to highlight homogeneity directly on the grey levels or through the use of filters capable of separating textures. Methods based on clustering, snakes, and MRFs are the most prominent in this direction.

Medical imaging is one area where domain knowl-

edge is available and can be considered to impose constraints and improve the solution obtained from images. Deformable templates [10], shape-driven snakes, shape-constrained level sets deformable contours, active shape and appearance models [1] as well as models based on non-parametric densities are the most prominent approaches to knowledge-based segmentation of the left ventricle. These methods make the explicit assumption that a training set is available. Once a model (statistical representation of the samples) has been built, the goal of segmentation is to recover a solution that is both supported by the observed image and present in the learned statistical model.

Model-based segmentation of the left ventricle of the heart has been investigated quite well in the past. In [8] a method that combines a shortest path algorithm and shape matching is present. In [13] a more specific prior model was considered in the space of level set functions. Deformable surfaces as well as active shape and appearance models [12] and their robust variants have proven to be quite efficient prior models. An overview of a large set of statistical analysis methods applied to medical models was recently presented in [9].

The most important criticism of these methods is the fact that model construction and segmentation are considered separately. Recently, attempts were made to retain more information from every sample of the training set than a sheer point of a high dimensional vector space. For instance by adding locality information on the shape [11] or constraining deformations to legal displacements using m-reps [5].

In the context of this paper, we aim to address this fundamental limitation of existing methods in a different and computationally efficient manner. To this end, a prior deformation model using Thin Plate Splines (TPS) is constructed that consists of a set of control points attached to a distance function. The main innovation of the method refers to the estimation of registration uncertainties that are propagated to the model construction process. Then, an Independent Component Analysis (ICA) is used to capture the variations in the training set where the registration uncertainties are

explicitly transmitted to the statistical model. Segmentation is then performed using a multi-phase propagation with cost terms attached to the TPS implicit representation. Uncertainties are also computed during the segmentation to assess the quality of the results.

The remainder of the paper is organized as follows. In section 2 we discuss registration with uncertainties and TPS, while in section 3 we propose a modified ICA that encodes registration uncertainties. Section 4 is dedicated to the segmentation. Experimental results are discussed in Section 5. Conclusions are given in Section 6.

2. Shape Registration through Thin Plate Spline Transformation

Smoothness and topology preservation are desirable properties in registration. Our model fulfills both by applying TPS deformations whose partial derivatives are all continuous.

In the present framework, a 3D shape \mathcal{S} representing the left ventricle myocardium is embedded in a higher dimensional space through the use of Euclidean distance transform. The distance transform of the shape \mathcal{S} is denoted $\phi_{\mathcal{S}}$, it is assumed negative inside (muscle) and positive outside (blood pool and external parts) (Fig. 1).

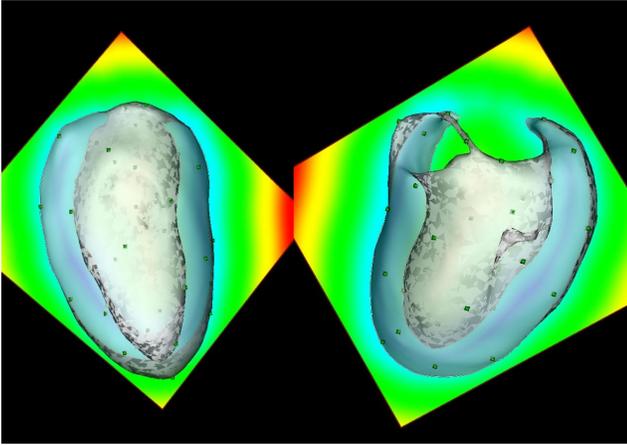


Figure 1. The average shape model of the left ventricle with its 90 control points and signed distance function. Used as a reference for registration and segmentation.

Such a space is invariant to translation T and rotation R and can also be modified to account for scale variations s :

$$\mathcal{S}_2 = (s.R + T) \circ \mathcal{S}_1 \Rightarrow s.\phi_{\mathcal{S}_1} = \phi_{\mathcal{S}_2} \circ (s.R + T). \quad (1)$$

In the most general case an explicit relation between the distance function of the source and the target does not exist.

Now consider a smooth diffeomorphism \mathcal{L} defined on the domain Ω and depending upon a vector of parameters

$\Theta \in \mathbb{R}^n$. Standard point-based registration consists of finding \mathcal{L} applied to the source shape \mathcal{S} which minimizes the integral defined on \mathcal{S} of some metric error between the transformed source and the target \mathcal{T} . In order to prevent the minimization process of such energy to fall into local minima, one can extend registration within a band including numerous isosurfaces of the source distance transform. Therefore, a robust quadratic registration energy is proposed:

$$E_{\alpha}(\mathcal{L}(\Theta)) = \int_{\Omega} \mathbb{1}_{\alpha}(\phi_{\mathcal{S}}(\mathbf{x})) (\phi_{\mathcal{S}}(\mathbf{x}) - \phi_{\mathcal{T}}(\mathcal{L}(\Theta, \mathbf{x})))^2 d\mathbf{x} \quad (2)$$

where $\mathbb{1}_{\alpha}$ is the indicator function of the segment $[-\alpha/2, \alpha/2]$. This approach is a generalisation of [6] where the registration energy was defined considering the entire domain Ω .

This energy is minimized through the calculus of variations. Within such a process the selection of the parameter α is crucial, since to some extent, it refers to the scale of the shapes to be registered. Therefore minimization is performed for a decreasing set of α values while the number of degrees of freedom of the transformation \mathcal{L} also increases. The considered transformations do not preserve the distance transform constraint, therefore a correction process is applied to the source implicit representation when refining the narrow band [15].

The objective function (2) is used to address global registration as well as local deformations.

When designing a shape model it is desirable to separate global motion from deformation and therefore build a statistical model invariant under a certain class of global transformations [14]. Within our approach, shapes are initially registered with respect to a similarity transformation to estimate motion.

We use a Thin Plate Spline (TPS) transformation [2] to address local deformations using an increasing set of control points. Consider N control points $\Theta = \{P_i\}_{i=1}^N$ located on the surface of the source shape \mathcal{S} . TPS is defined as the transformation \mathcal{L} minimizing the bending energy under hard constraints on the displacement of control points. The bending energy is the Frobenius norm of the Hessian matrix of the transformation:

$$E_f = \int_{\Omega} \|H_{\mathcal{L}}(\mathbf{x})\|_F^2 d\mathbf{x}, \quad \text{under constraint } \mathcal{L}(P_i) = P'_i, \\ \text{where } \|H_{\mathcal{L}}(\mathbf{x})\|_F^2 = \text{tr}(H_{\mathcal{L}}(\mathbf{x})H_{\mathcal{L}}(\mathbf{x})^T), \quad (3)$$

where $H_{\mathcal{L}}(\mathbf{x})$ is the Hessian matrix of $\mathcal{L}(\Theta, \mathbf{x})$ and P'_i is the new position of the control point P_i .

The minimum of this functional verifies the biharmonic equation and the solution may be written with the form

$$\mathbf{x}' = \mathcal{L}(\mathbf{x}, A, T, V_i) = A.\mathbf{x} + T + \sum_{i=1}^N V_i U(\|P_i - \mathbf{x}\|), \quad (4)$$

where $A \cdot \mathbf{x} + T$ represent the affine part of the transformation and the set of vectors $\{V_i\}_{i=1}^n$ the weights of the non affine warping. $U(r)$ is a radial basis function, solution of the biharmonic equation, that is in 3 dimensions $U(r) = -|r|^4$.

The parameters of the transformation (V, A, T) , can be expressed as a linear combination of the displacement of the control points. Therefore in the next section we will refer to the transformation with the formula:

$$\mathcal{L}(\Theta, \mathbf{x}) = \mathbf{x} + \mathcal{X}(\mathbf{x}) \cdot \Theta, \quad (5)$$

with $\mathcal{X}(\mathbf{x})$ a $(3 \times 3N)$ matrix.

All computations in the article make use of discrete sets of points, therefore all non linear computations representing the matrix $\mathcal{X}(\mathbf{x})$ are performed at initialization and subsequent transformations of the sets of points are computed in linear time.

To prevent non invertible transformations to appear during the minimization process, a regularisation term with decreasing influence was added according to [2]. This additional term is actually the bending energy defining the TPS, this term is smoothing out the local part of transformation and can be expressed directly as a quadratic form on the position of the control points.

2.1. Similarity-Invariant TPS

The registration process separates the global transformation from the deformations in the sense that the initially estimated similarity transformation is subtracted from the target mesh. However the subsequently estimated TPS transform also contains an affine transform. We use the general form of the TPS to extract and remove the remaining similarity closest to this affine. We choose constrain the global transformation to a similarity rather than an affine to prevent the shape from anisotropic scaling and shearing.

The extraction of the similarity can be implemented by determining the projection of the affine part of the TPS (eq. 4) in the manifold of similarity transformations.

This projection is obtained by extracting the affine part A of the TPS and minimizing the Froebenius norm $\|A - sR\|_F$ with respect to a rotation R and a scaling factor s . This minimization problem has a closed form solution [3]. Consider the SVD decomposition of the affine matrix $A = V \cdot D \cdot W^T$, the closest similarity is:

$$S = s \cdot R \text{ with } R = V \cdot W^T \text{ and } s = \text{trace}(D)/3 \quad (6)$$

In order to obtain an efficient approximation of the similarity transform, it is necessary to locate the center of rotation at the center point of the initial shape. Finally this transformation is subtracted from both the target shape and the current transformation to preserve the similarity invariant warp of shapes.

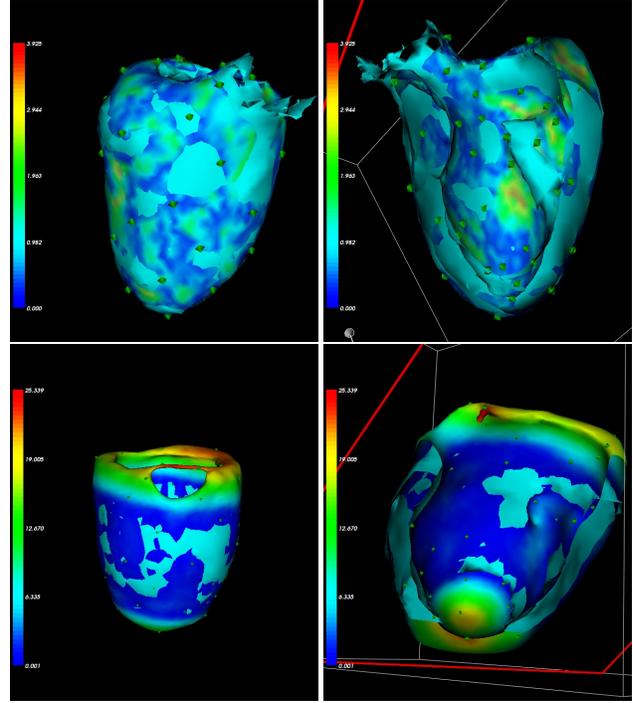


Figure 2. Left side: Registration of the shape model to segmented data extracted from CT volumes. Right side: Registration obtained from MRI data with missing apex and valves. Registration results are obtained with 90 TPS control points located on both endocardium and epicardium. Left: View of the epicardium; Right: Clipped view showing the endocardium. (Colormap errors measures are expressed in mm).

The output of the process is a deformation field that superposes the deformed source shape and the target shape (see Fig. 2). In order to assess the quality of the solution and how its support varies spatially, we can determine some measures of uncertainties which relate the obtained solution with the observable data support.

2.2. Uncertainty estimation on registered shapes

We aim to recover uncertainties on the vector Θ in the form of a $[3N \times 3N]$ covariance matrix by adapting a method initially introduced in [15] to 3D with thin plate spline. This approach is based on an approximation of the Hessian of the registration energy (2) when reaching a local minimum. These uncertainties should only account for the local variability of the shape, so that we only consider the registration energy in a very thin narrow band, surrounding the model shape. We use the first order approximation of the energy in the same way as the Gauss Newton method applied to quadratic energies.

One notices that localizing the global minimum of the objective function E is equivalent to finding the major mode of a random variable with density $\exp(-E/\beta)$. The coef-

ficient β corresponds to the allowable variation in the energy value around the minimum. In the present case of a quadratic energy, the covariance and the Hessian of the energy are directly related by $\Sigma_{\Theta}^{-1} = H_{\Theta}/\beta$. The choice of TPS transformation ensures the Hessian matrix to be invertible. This leads to the following expression for the covariance :

$$\Sigma_{\Theta}^{-1} = \frac{1}{\beta} \oint_{\partial S} \mathcal{X}(\mathbf{x})^T \cdot \nabla \phi_T(\mathbf{x}') \cdot \nabla \phi(\mathbf{x}')^T \cdot \mathcal{X}(\mathbf{x}) d\mathbf{x} \quad (7)$$

where $\mathbf{x}' = \mathcal{L}(\Theta, \mathbf{x})$ represents the position of a transformed point. Fig. 3 shows a projection of the uncertainties on two training shapes.

An important use for the uncertainty is to overcome the limited size of the training set available on myocardial data: these are used to generate new samples and will be developed the Section 3.

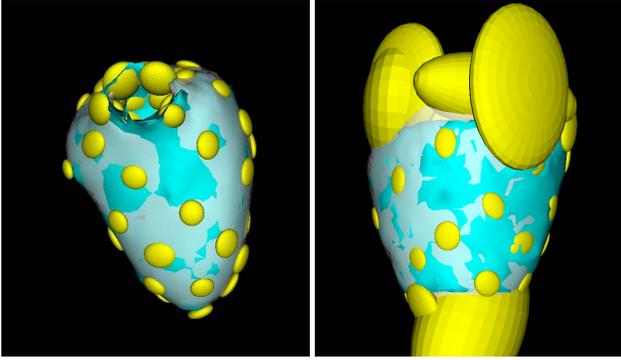


Figure 3. Projection of the uncertainty covariance matrix Σ_{Θ} on the TPS control points. Uncertainty is represented as a set of ellipsoids which do not account for correlation between control points. Left: CT registration; Right: MRI registration showing very large uncertainties in areas with missing data.

2.3. Building the Training Set

In the following we consider the source shape representing the left ventricle \mathcal{S} as a unique model used for both the registration and segmentation process. In order to perform high quality registration and efficient dimensionality reduction, it is desirable for \mathcal{S} to be smooth and represent an "average shape" in the space of deformations [1].

We dispose of a training set of N semi-automatically segmented ventricles without prior using a graph-based algorithm [4]. Our training set is also composed of manually segmented samples acquired from MRI. All the MRI datasets were missing slices at the valves and the apex, thus resulting in incomplete training sets. One can see from Fig. (2,3) that the areas where data is absent show high uncertainty on the control points. Similarity invariant registration with computation of uncertainty was performed on both

training sets. In order to estimate the average deformation $\bar{\Theta}$ over the whole training set, we use the Mahalanobis distance with the estimated uncertainty for each sample:

$$\bar{\Theta} = \arg \min_{\Theta} \left(\sum_i (\Theta - \Theta_i)^T \Sigma_{\Theta_i}^{-1} (\Theta - \Theta_i) \right) \quad (8)$$

3. Dimension reduction using ICA

Conventional techniques of shape modeling assume a *reference average shape* (\mathcal{M}) and then build statistical densities on the space of deformations [1] between the reference shape and the training examples.

Independent component analysis (ICA) is a statistical/computational technique that can determine the statistical nature for sets of random variables, measurements, or signals. This technique proposes a linear decomposition of a multidimensional model, just as Principal Component Analysis (that computes directions of largest variance) and Factor Analysis (that explain a statistical measurements as a linear combination of a certain number of factors).

ICA relies on a different model. Like PCA, it requires multidimensional random variables, but retrieves independent directions. This means that the projections of the observed data along these directions are independent scalar random variables [7]. On the other hand factor analysis uses a set of candidate factors to retrieve the weights that generate the one dimensional observed data, while ICA directly uses multivariate data to estimate both the factor and the weights.

In brief ICA relies on a generative model that assumes that the data can be generated as a linear combination of random variables. These latent variables are assumed non-Gaussian and mutually independent, they are called the 'independent components' of the observed data. Regarding our application to modeling of shape deformations, this means that any sample shape can be explained as a linear combination of specific independent deformations. The simplest proof of independence is found in the fact that it affects different parts of the shape under study:

$$\Theta = \mathbf{A} \cdot \mathbf{s}, \quad (9)$$

where Θ is the observed data vector, \mathbf{s} a m -dimensional random vector with independent components, and \mathbf{A} a constant mixing matrix. Several approaches were developed to robustly compute statistical independence of random variables and estimate the values of the de-mixing matrix for a fixed number of independent components:

$$\mathbf{s} = \mathbf{W} \cdot \Theta = \mathbf{A}^+ \cdot \Theta. \quad (10)$$

ICA was previously used in shape modeling [16] where the authors proposed a method to order and select the "most relevant" independent components. Their method represents

shapes using the point distribution models in a very high dimensional space. Our representation of deformations is in much lower dimensional space. Thus, our application does not require classification of independent components, we therefore directly compute m components using the *Fast-ICA* algorithm. The components of \mathbf{s} being independent, their distributions can be expressed as the product of 1-dimensional random variables densities:

$$p(\mathbf{s}) = \prod_{i=1}^m p_i(s_i) \quad (11)$$

where a Gaussian mixture model is fitted with EM algorithm to estimate the actual density of each components :

$$p_i(s_i) = \sum_{j=1}^{M_i} \alpha_j \mathcal{N}(\mu_{i,j} - s_i, \sigma_{i,j}). \quad (12)$$

In practice, the choice $M_i = 2$ is sufficient to enclose the variability of the training set. Fig. 4 shows the training samples projected on the first 3 components. This method can produce a compact statistical model given a training set where all samples are brought to the same pose. However, it does not account for the uncertainties determined during the registration process. In order to include the registration uncertainties in the model building, we generate additional data prior to ICA. Our approach consists of sampling each measurement on the training set using the normal law centered on the registration results with uncertainty being the covariance matrix $\mathcal{N}(\Theta_i, \Sigma_{\Theta_i})$. For each example the same number of samples is drawn using the corresponding densities. Such a process will lead to an augmentation of the training set where samples of low uncertainty will have more influence on the model to be recovered. On the other hand, for registration results showing high values of the covariance matrix determinant, the drawn samples will be dispersed and therefore will have limited influence in the model construction. The augmented set of variables can now be used within the Independent Component Analysis framework to produce a multi-component density of deformations.

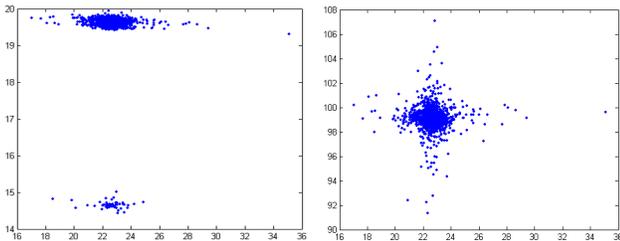


Figure 4. Density plot of independent components $\mathbf{s} = \mathbf{W} \cdot \Theta$, independence and gaussian mixture estimation of components appears relevant.

4. Shape Inference from Images

Computer-aided diagnosis is a growing application domain of medical image analysis. Segmentation of cardiac structures is used to assist physicians in various stages of treatment of cardiovascular diseases. In particular, measuring the blood volume, wall motion and wall thickening properties of the left ventricle (LV) is of great importance because the LV pumps oxygenated blood out to distant tissue in the entire body and malfunction leads to severe pathology. Automatic detection, segmentation, and tracking of the LV are powerful tools to assist physicians in calculating these measurements efficiently.

4.1. Segmentation of the Left Ventricle

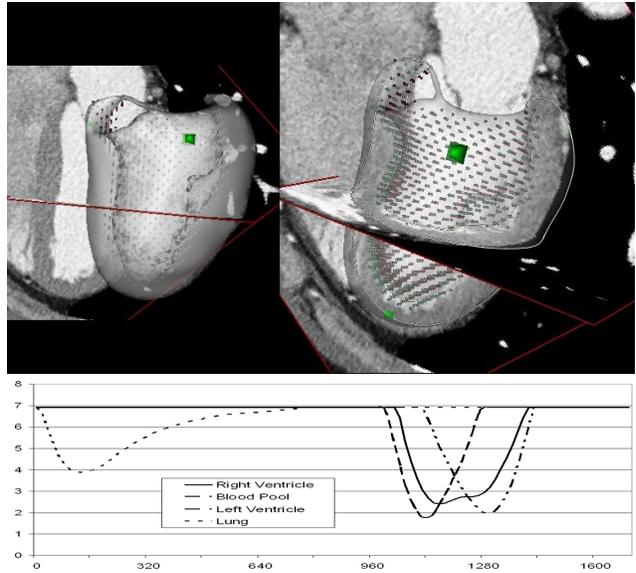


Figure 5. Left: the rough segmentation of the left ventricle, simply using the endocardium and minimized with respect to a similarity transform. Right: the histograms of the greylevels of different areas of the myocardium.

Let \mathcal{I} represent a 3D CT-scan of the heart. We take advantage of the Hounsfield units and the characteristic gray values of the blood and muscle to propose a region based segmentation strategy [8]. It is natural to separate the segmentation into 4 classes: Blood Pool (BP), Myocardium (MC), Lung (L) and Right Ventricle (RV). The first three classes have a single monochromatic component (see Fig. 5) while the right ventricle actually shows a more complex mixture. We assume that the visual properties of these classes were learned from training samples and are known $\{\pi_{BP}(), \pi_{MC}(), \pi_L() \text{ and } \pi_{RV}()\}$. Then segmentation of the ventricle is equivalent to the minimization of the following energy with respect to the parameters Θ under the assumption that pixels of the same class are statistically independent.

dent, that is often considered in computer vision and medical imaging:

$$E_{\text{image}}(\Theta) = \sum_{\Omega_i=\{\text{BP,MC,L,RV}\}} \int_{\mathcal{R}_{\Omega_i}} -L_{\Omega_i}(\mathcal{L}(\Theta, \mathbf{x}))d\mathbf{x} \quad (13)$$

where L_{Ω_i} refers to the log likelihood of a pixel intensity ($L_{\Omega_i}() = \log[\pi_{\Omega_i}(\mathcal{I}())]$). Direct calculus of variations with respect to Θ involves image gradient, derivation of densities ($\pi_{\text{BP}}, \pi_{\text{MC}}, \pi_{\text{L}}, \pi_{\text{RV}}$), and often converges to erroneous solutions. In our case, we change the integration domain to the image by implicitly introducing the inverse transformation as was introduced in [15]. In the case of the heart, the 4 phase segmentation takes the form of 3 integrals on the interfaces: the endocardium [Blood Pool/Myocardium] and epicardium [Myocardium/Right Ventricle] and [Myocardium/Lung], as the other interfaces are not relevant or do not exist.

$$\frac{dE_{\text{image}}(\Theta)}{d\Theta} = \sum_{i,j} - \int_{\partial C_{\Omega_i/\Omega_j}} L_{\Omega_i/\Omega_j}(\mathcal{L}(\Theta, \mathbf{x})) \nabla \phi_{S'}(\mathbf{x}')^T \frac{\partial \mathcal{L}}{\partial \Theta^T}(\Theta, \mathbf{x})d\mathbf{x} \quad (14)$$

with $L_{\Omega_i/\Omega_j} = L_{\Omega_i} - L_{\Omega_j}$ being the difference of the log likelihood of two regions sharing an interface. In this equation, we use the following interfaces defined with $(\Omega_i, \Omega_j) = \{(\text{BP}, \text{MC}), (\text{RV}, \text{MC}), (\text{L}, \text{MC})\}$. $\nabla \phi_{S'}(\mathbf{x}')$ is the gradient of the transformed shape, evaluated at the transformed position $\mathbf{x}' = \mathcal{L}(\Theta, \mathbf{x})$. In practice, this can be explicitly computed as a function of the initial source shape gradient and the transformation:

$$\nabla \phi_{S'}(\mathbf{x}') = \text{com} \left(\frac{\partial \mathcal{L}}{\partial \mathbf{x}^T}(\mathbf{x}, \Theta) \right) \cdot \nabla \phi_S(\mathbf{x}), \quad (15)$$

where ϕ_S is the signed distance transform of the original left ventricle, negative on the muscle (see Fig. 1), and $[\text{com}(\frac{\partial \mathcal{L}}{\partial \mathbf{x}^T}(\Theta, \mathbf{x}))]$ denotes the matrix of cofactors of the Jacobian of the transformation.

We use standard estimates from the training set for the densities $\pi_{\text{BP}}, \pi_{\text{MC}}, \pi_{\text{L}}, \pi_{\text{RV}}$. Initially, a rough segmentation is performed by aligning the model with respect to a similarity transform and considering only the myocardium and blood pool regions, taking advantage of the high contrast between blood and muscle. Subsequently the 4 class segmentation is used to adjust the position of the epicardium. The gray level distributions of the 4 classes are recovered in an incremental fashion using the Mumford-Shah principle.

The shape based energy term is defined on the entire space of deformation and makes use of the prior density

function (eq. 11). However, recall that only the local deformation component of the transformation was learned, so we need to remove the similarity part of the transformation from the current estimate of Θ before projecting onto the space of independent components. The current similarity carried by the TPS is extracted using the same approach as presented in subsection (2.1). The inverse of this transformation is subsequently applied to the control points to remove the pose parameters associated with the current segmentation and lead to the similarity invariant parameter vector Θ' .

The shape prior energy term is equivalent to the maximization of the log likelihood of the current vector of parameter Θ .

$$E_{\text{shape}}(\Theta) = - \sum_{i=1}^m \log(p_i([\mathbf{W} \cdot \Theta']_i)), \quad (16)$$

where $\mathbf{W} \cdot \Theta'$ refers to the projection of the current shape on the independent components. p_i refers to the density estimated in (eq. 11) and (eq. 12) for the i^{th} ICA component.

The global energy is minimized with respect to the parameters Θ through the calculus of variations on the blended objective function $E = E_{\text{image}} + \alpha E_{\text{shape}}$ and implemented using standard gradient descent. This process is demonstrated in Fig. (6,5) with the initial, intermediate and final stages.

4.2. Segmentation Uncertainties

Medical imaging is an area where errors could be critical. Furthermore, computer-aided diagnosis should not aim at replacing humans but should be a tool that helps the physicians to make more appropriate decisions. Therefore methods being able to determine the quality of the obtained solution are required. As stated earlier, this can be done using the uncertainties of the process that relate the Hessian matrix of the objective function to the variance of the distribution describing the space of solutions.

To validate the segmentation results in a quantitative manner we propose to estimate uncertainties on the deformation, expressed on the position of TPS control points. These uncertainties are actually based on the structure matrix of the image near the shape boundary and projected on the vectors of parameters Θ . In practice, these uncertainties are expressed in the same fashion as (eq. 7) where the local shape information part $[\nabla \phi_T(\mathbf{x}') \cdot \nabla \phi_T(\mathbf{x}')^T]$ is replaced with the following symmetric matrices:

$$\Sigma_i^+ = \frac{\|\mathbf{a}_i\|}{\|\mathbf{b}_i\|} \cdot \mathbf{b}_i \cdot \mathbf{b}_i^T + \frac{\|\mathbf{b}_i\|}{\|\mathbf{a}_i\|} \cdot \mathbf{a}_i \cdot \mathbf{a}_i^T \quad (17)$$

with $\begin{cases} \mathbf{a}_i = \nabla(L_{\Omega_i/\Omega_j}(\mathbf{y}_i)) \\ \mathbf{b}_i = \nabla(\phi_{S'}(\mathbf{y}_i)) \end{cases}$

where integration is performed along the interfaces considered for segmentation. This quantity represents a fair

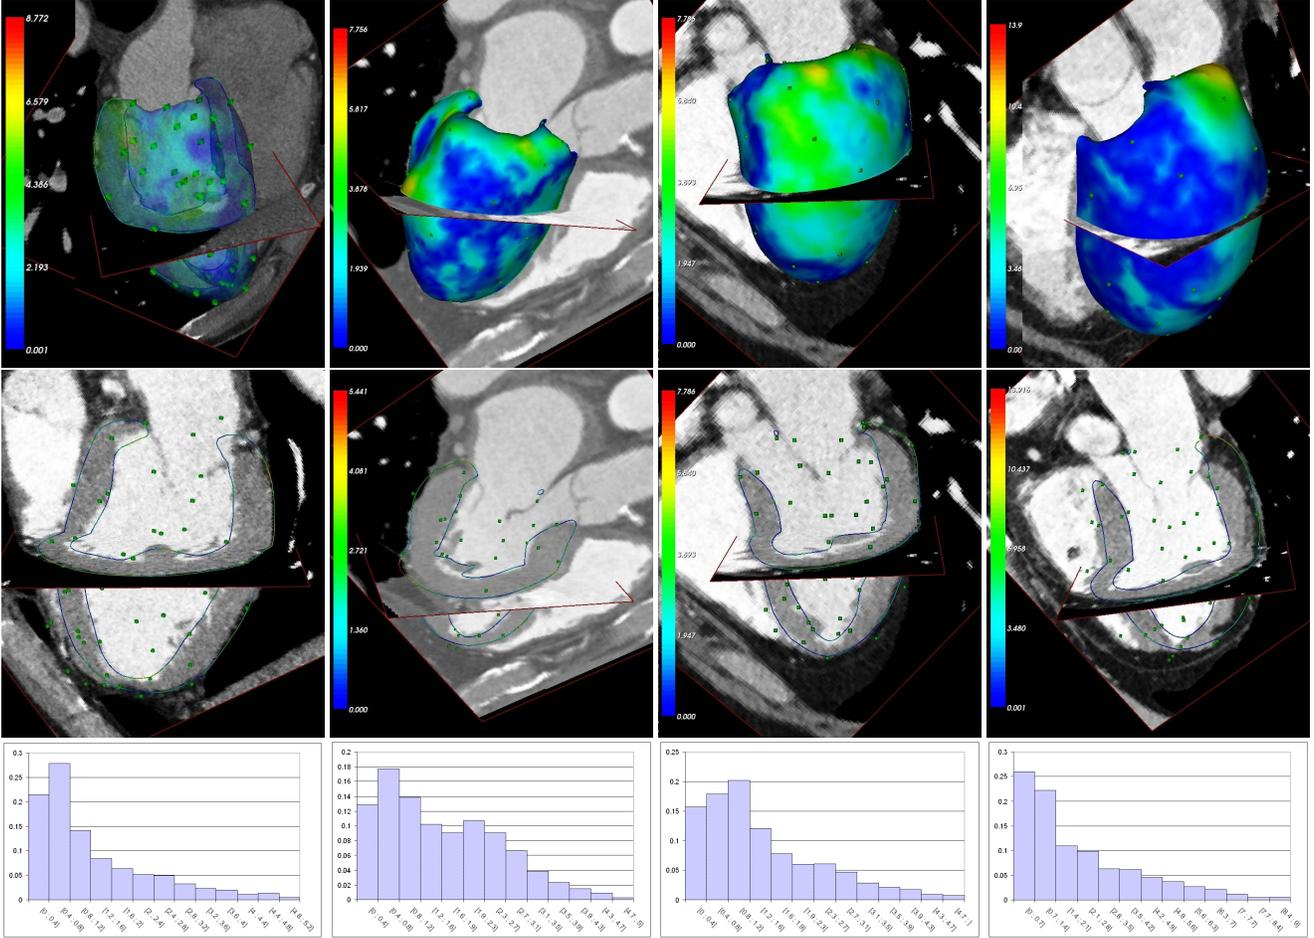


Figure 6. Results of the segmentation process. Left: Segmentation of the myocardium displaying the distance to the groundtruth along with a colormap considering four different sample cases outside of the learning database. All results are displayed for a TPS deformation using the ICA shape model with 90 control points. Middle column: The intersection of the segmentation with the data, papillary muscles are correctly segmented. Right: Distribution of errors displayed as histograms. All numerical values are expressed in millimeters.

approximation of the Hessian of the image energy term (eq. 13) where $\nabla \phi_{S'}(\mathbf{x}')$ contributes to the uncertainty along the shape surface and $\nabla(L_{\Omega_i/\Omega_j}(\mathbf{x}'))$ represents the quality of region separation on the interface (edge information). If the direction of these two vectors match exactly, then the uncertainty is concentrated in the contour direction. If these directions differ, uncertainties are also present in the direction normal to the contour. This contribution is integrated along the segmented interface and projected on the parameter space (eq. 7) to define the segmentation uncertainty information.

5. Results

In order to validate our approach, we have used segmented CT and MRI volumes. This set includes both healthy and pathological hearts. The model was constructed using a set of 25 diastolic CT samples which were semi-

automatically segmented using the method presented in [4] and 18 diastolic MRI samples which were manually segmented. Large errors were induced in the training set (erroneous samples) which aimed to demonstrate the importance of uncertainties estimation and propagation to the model construction (Fig. 3). Our thin plate spline model contains 90 control points, located on the endocardium and epicardium. Using the uncertainty estimate on the registration, we generate 300 random samples for each registration, for a total of 12900 shapes for the entire training set. Then using the *FastICA* algorithm, we extract 15 independent components. Each component distribution is modeled using a mixture of 2 Gaussians with an Expectation-Maximization algorithm.

Energy and gradients are computed using discrete points sets. Therefore the integral in (eq. 14) is discretized using a narrow band surrounding the ventricle surface. This improves stability and favors convergence of the gradient

descent.

Subsequently, segmentation is performed in 3 successive steps using gradient descent on the global energy $E = E_{image} + \alpha E_{shape}$. First, image based ($\alpha = 0$) affine segmentation of the endocardium, then image based segmentation of the entire left ventricle with 60 control points, then segmentation with prior of the entire ventricle with 90 control points. Fig. 6 shows that we obtain very promising results with this method. For the first 3 examples, the average error is around 1.5 mm and the maximum error is less than 5mm. The errors are larger for the last example because this patient has a thick layer of fat around the heart and the method mistakenly combines the fat with the myocardium. In the future we are planning to broaden the lung class with a fat class.

6. Conclusion

In this paper we have proposed a novel method of image segmentation. The method exploits recent results on shape registration with uncertainties to build a compact statistical model. In these problems, a quantitative and qualitative solution is obtained where the final output encodes the spatially varying support of geometric (surface) and photometric features (image).

Improving the modeling as well as the inference aspect of our approach involves a number of potential future directions. Future work will make the shape learning framework available in the coarse to fine approach used for the TPS transformation, that means also learning deformation for lower degrees of freedom in the TPS transformation. The algorithm does not suffer from any limitation regarding the number of classes to be segmented, the next step will be to amend the model including right ventricle atria and aorta. Furthermore, the introduction of more advanced photometric properties, like texture recognition, could also improve the segmentation method.

References

- [1] T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61:38–59, 1995.
- [2] G. Donato and S. Belongie. Approximate thin plate spline mappings. In *European Conference on Computer Vision*, pages 21–31, 2002.
- [3] G. Golub and C. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1989.
- [4] L. Grady, Y. Sun, and J. Williams. Three interactive graph-based segmentation methods applied to cardiovascular imaging. In O. F. Nikos Paragios, Yunmei Chen, editor, *Handbook of Mathematical Models in Computer Vision*, pages 453–469. Springer, 2006.
- [5] Q. Han, D. Merck, J. Levy, C. Villarruel, J. N. Damon, E. L. Chaney, and S. M. Pizer. Geometrically proper models in statistical training. In *IPMI*, pages 751–762, 2007.
- [6] X. Huang, N. Paragios, and D. N. Metaxas. Shape registration in implicit spaces using information theory and free form deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(8):1303–1318, 2006.
- [7] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. John Wiley & Sons, 2001.
- [8] M.-P. Jolly. Automatic segmentation of the left ventricle in cardiac mr and ct images. *International Journal of Computer Vision*, 70(2):151–163, 2006.
- [9] H. Krim and A. Yezzi. *Statistics and Analysis of Shapes (Modeling and Simulation in Science, Engineering and Technology)*. Birkhauser, 2005.
- [10] P. Lipson, A. Yuille, D. Okeefe, J. Cavanaugh, J. Taaffe, and D. Rosenthal. Deformable Templates for feature extraction from Medical Images. In *European Conference on Computer Vision*, pages 413–417, 1990.
- [11] M. Loog. Localized maximum entropy shape modelling. In *IPMI*, pages 619–629, 2007.
- [12] S. Mitchell, J. Bosch, B. Lelieveldt, R. van der Geest, J. Reiber, and M. Sonka. 3-d active appearance models: Segmentation of cardiac MR and ultrasound images. *IEEE Trans. Med. Img.*, 21:1167–1178, 2003.
- [13] M. Rousson and N. Paragios. Shape Priors for Level Set Representations. In *European Conference on Computer Vision*, pages II:78–93, Copenhagen, Denmark, 2002.
- [14] S. Soatto and A. J. Yezzi. DEFORMATION: Deforming motion, shape average and the joint registration and segmentation of images. In *ECCV (3)*, pages 32–57, 2002.
- [15] M. Taron, N. Paragios, and M.-P. Jolly. Uncertainty-driven non-parametric knowledge-based segmentation: The corpus callosum case. In *VLSM*, pages 198–207, 2005.
- [16] M. Uzumcu, A. Frangi, J. Reiber, and B. Lelieveldt. Independent component analysis in statistical shape models. In *SPIE Medical Imaging*, volume 5032, pages 375–383, 2003.