PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE

SCHOOL OF ENGINEERING

PRIOR KNOWLEDGE FOR LEVEL SET SEGMENTATION

CRISTOBAL IGNACIO ARRIETA PELLEGRIN

Thesis submitted to the Office of Graduate Studies in partial fulfillment of the requirements for the Degree of Doctor in Engineering Sciences

Advisor:

CRISTIAN TEJOS

Santiago de Chile, October 2017

PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE

SCHOOL OF ENGINEERING

PRIOR KNOWLEDGE FOR LEVEL SET SEGMENTATION

CRISTOBAL IGNACIO ARRIETA PELLEGRIN

Members of the Committee:

CRISTIAN TEJOS

DOMINGO MERY

PABLO IRARRAZAVAL

JOAQUIN MURA

RENE VIDAL

JORGE VASQUEZ

Thesis submitted to the Office of Graduate Studies in partial fulfillment of the requirements for the Degree of Doctor in Engineering Sciences

Santiago de Chile, October 2017

Para las Pequitas

ACKNOWLEDGEMENTS

I would like to thank the Grant funding: CONICYT FONDECYT/Regular 1130887, CONICYT FONDECYT/Regular 1161448, CONICYT FONDEF/I Concurso IDeA en dos etapas ID15|10284, CONICYT - PIA - Anillo ACT1416, and also CONICYT Beca de Doctorado Nacional.

I also would like to thank my parents, Victor Hugo and Lorena, my brothers Daniel and Gabriel, for their unconditional support, and all my family, specially my grandmother Mane.

I would like to thank professors Pablo Irarrazaval, Sergio Uribe, Marcelo Andia and Joaquin Mura, for their patient and support. This thesis is undoubtedly better due to their contributions.

I also would like to specially thank my advisor, Cristian Tejos, for believing in me. I will always remember his detailed and insightful observations.

I would like to thank professor Sing-Long for the inspiration to my work and for his friendship. I also like to thank my colleges at Biomedical Imaging Center. All those guys turn a workspace into a home.

I would like to thank my old friends Felipe Araya, Ivan Parra, Gonzalo Aguirre, Diego Diaz, Daniel Aguirre, Camila Muoz, Ignacio Gomez, Gabriel Acua and Alejandra Olave. They added the greatest poetry, art and music to these PhD years.

Finally, I would like to dedicate this thesis to my wife, Constanza Gomez (las Pequitas), because at her side, all this effort was worth it, and all these years were perfectly knit for getting us together, as a great team. She showed me the deepest love that I could ever imagine: the love of sharing simple things in life, which she magically makes them legendary.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Level set segmentation has been successfully used in several image applications. However, it performs poorly when applied to severely corrupted images or when the object boundaries are blurred or occluded. Shape prior knowledge adds information inferred from a database that allows one to compensate for poor data quality. This prior knowledge can be included into the level set segmentation in two ways: weak prior knowledge, which makes simple geometric assumptions about the shape and usually implies a rigid topology; and strong prior knowledge, which consists of adding a regularization term that penalizes shapes that differ from a database and it is normally topologically flexible.

In this thesis we present two works that involve prior knowledge: the first one corresponds to the application of a weak prior knowledge approach to cardiac segmentation and the second one is a new method to incorporate strong prior knowledge.

For the weak prior knowledge approach, we propose a method based on level sets with preserved topology that allows simultaneous, fast and accurate segmentation of the left and right ventricles. We compared our segmentation results of the left and right ventricles with those obtained with clinically validated software (Viewforum, Philips, Best and Segment, Medviso, Lund) using two-tailored paired t-test, Pearsons correlation, Bland-Altman plots of standard functional indexes and voxel-by-voxel analysis with Dice. Two-tailored paired t-test showed no significant difference between our method and gold standards ($P<0.05$), Pearsons correlation showed a high correlation of our measurement with gold standards

(over 0.98), Dice showed an average agreement of at least 0.90 and Bland Altman analysis showed that our method has a good agreement with the gold standard segmentation.

For the strong prior knowledge application, one of the challenges is to define a pose invariant regularization term, that is a regularizer that does not change when subjected to translations, rotations and scaling. Previous works have accounted for this by coupling the curve evolution to a registration problem through an optimization procedure. This approach is slow and its results depend on how this optimization is implemented. To overcome this issue, Cremers et al. (2006) introduced an intrinsic alignment, which normalizes each shape to a common coordinate system, avoiding the registration process. Nevertheless, their proposed solution considered only scaling and translation but not rotation, which is critical in several image applications. We propose an extension to Cremers' work, that considers intrinsic scaling, translation and rotation. Our regularization term is based on the eigenvalues and eigenvectors of the covariance matrix of each training shape, and this eigendecomposition dependency leads to a new set of evolution equations. We tested our regularizer, combined with a Chan-Vese functional, in 2D and 3D synthetic and medical images, demonstrating the effectiveness of using shape priors with intrinsic scaling, translation and rotation alignment in different segmentation problems.

Keywords: segmentation, level sets, prior knowledge, preserved topology, pose invariant, intrinsic alignment.

RESUMEN

La segmentación usando Level Sets ha sido utilizada exitosamente en varios tipos de imágenes. Sin embargo, ésta falla cuando las imágenes tienen mucho ruido, bordes difusos u oclusión. El uso de información a priori inferida a partir de bases de datos de entrenamiento permite compensar la baja calidad de las imgenes. La informaión a priori, puede ser débil o fuerte. La información a priori débil supone geometrías simples de la forma buscada y fuerza una topología rígida. La información a priori fuerte agrega un término de regularización que penaliza formas distintas a las del conjunto de entrenamiento y permite una topología flexible.

Esta tesis presenta dos trabajos relacionados con conocimiento a priori, el primero es una aplicación de conocimiento a priori débil en segmentación cardíaca, y el segundo propone un nuevo método de conocimiento a priori fuerte.

Para el conocimiento a priori débil, proponemos un método basado en level sets que preserva la topología y segmenta simultánea, rápida y precisamente los ventrículos izquierdo y derecho del corazón. Los resultados de este método fueron comparados con otros obtenidos utilizando *softwares* clínicamente validados (Viewforum, Philips, Best and Segment, Medviso, Lund), usando *two-tailored paired t-test*, correlación de Pearson, gráficos de Bland-Altman de índices funcionales y comparación voxel a voxel usando el índice Dice. El *two-tailored paired t-test* mostró que no hay diferencia significativa entre nuestro método y el *gold standard* ($P<0.05$), Pearson mostró alta correlación entre nuestro método

y el *gold standard* (sobre 0.98), el índice Dice mostró alta intersección entre nuetro método y el *gold standard* (igual o mayor al 90%) y los gráficos de Bland-Altman mostraron alto grado de conscenso y bajo sesgo entre nuestro método y el *gold standard*.

Para el conocimiento a priori fuerte, definir un regularizador invariante a la traslación, rotación y escalmiento de las formas presentes en la base de datos es un gran desafío. Trabajos anteriores lograron esta invarianza acoplando procesos de registro y segmentación. Esta solución es lenta y presenta resultados variables, según cómo se realice el registro. Frente a este problema, Cremers et al. (2006) propusieron un alineamiento intrínseco, mediante el cual las formas de la base de datos son normalizadas y llevadas a un sistema de coordenadas común que permite compararlas. Desgraciadamente, esta solución sólo considera invarianza a la translación y el escalamiento, sin considerar la rotación, crucial en imágenes médicas. Nuestro trabajo es una extensión del trabajo de Cremers, considerando alineamiento intrínseco invariante a la translación, escalamiento y rotación. El nuevo regularizador considera los vectores y valores propios de la matriz de covarianza de las formas de entrenamiento, produciendo un nuevo conjunto de ecuaciones de evolución. Probamos el nuevo regularizador combinado con el algoritmo de Chan-Vese en imágenes sintéticas e imágenes médicas, en 2D y 3D, mostrando resultados efectivos, precisos y destacando la importancia de considerar rotaciones en el proceso de alineamiento intrínseco.

Palabras claves: segmentación, level sets, conocimiento previo, preservación de la topología, invarianza a la posición, alineamiento intrínseco.

1. INTRODUCTION

Image segmentation based on level set algorithms has been successfully used in several applications. It was introduced by Caselles et al. (Caselles, Catté, Coll, & Dibos, 1993) as geometric active contours. This method moves a smooth curve in the image to enclose an object of interest, and isolates it from the background. The curve is implicitly defined as the zero level set of a signed distance function $\phi(\vec{x})$, and it is driven following two criteria: the curve has to be smooth and moves according to its mean curvature; and it is stopped by a forced that depends on a desired image feature (e.g., object's edges). This is done through an iterative process that depends on a parameter t, usually identified with time. Hence, in this first formulation, the curve evolves proportionally to its mean curvature, which is multiplied by a term that sets the velocity to zero when the curve passes through an edge within the image. Remarkably, level sets are the first active contour formulation that is topologically flexible, i.e., the number of zero-level set curves (and thus the number of segmented objects) could change during the deformation process.

This first formulation was proved to be numerically unstable and it worked properly only under special conditions. The Geodesic Active Contours (Caselles, Kimmel, & Sapiro, 1997) solved some of these issues by formulating the same edge-based problem as an energy minimization process, in which the energy is differentiated with respect to $\phi(\vec{x})$ in order to find the first order condition. The resulting expression, called evolution equation, drives the deformation of the curve implicitly embedded by $\phi(\vec{x})$. Subsequently, Chan and Vese (Chan & Vese, 2001) formulated a region-based minimization problem

which groups pixels with similar intensity into the regions inside or outside the curve. The energy to be minimized has two terms, one that accounts for the homogeneity of the intensity of each one of the two regions, and a second one that acts as a regularizing term enforcing the curve to be smooth.

These techniques led to a big family of algorithms that have made an important contribution to many kinds of segmentation problems. Particularly, they found an important field of application in medical images due to their ability to deal with some commonly encountered image distortions (Niessen, ter Haar Romeny, & Viergever, 1998; Suri et al., 2002; He et al., 2008; Q. Liu, Jiang, Bai, & Yang, 2016; Ivanovska et al., 2016; L. Liu, Zhang, Wu, Li, & Shang, 2013; Dong, Chen, & Wang, 2013).

In general, medical images (e.g., magnetic resonance (MR) imaging, computerized tomography, ultrasound), microscopy images, security cameras videos or satellite images, are highly corrupted by noise and acquisition artifacts. Furthermore, sometimes the available visual information is not enough to correctly delimit the structures of interest. In those cases, prior knowledge needs to be used in order to get better segmentation results (Tejos, Irarrazaval, & Cárdenas-Blanco, 2009).

Level set segmentation has used two kinds of prior knowledge: weak or strong, depending on the geometric assumptions made about the shape (Petitjean & Dacher, 2011). An example of weak prior knowledge is to fix the number of objects allowed to be segmented. The idea is to keep the same number of connected objects throughout the entire

curve evolution. This is done by allowing the image pixels to enter or leave the area inside a curve, only if they do not change the topology of the curve (Han, Xu, & Prince, 2003). This process is repeated at the end of each iteration. Another example of weak prior knowledge is to enforce the shape of the segmenting curve to adopt a simple geometric shape (e.g. an ellipse). This can be done by fitting the segmenting curve to the chosen shape at the end of every iteration (Pluempitiwiriyawej, Moura, Wu, & Ho, 2005).

Despite being easy to implement, these two approaches have important limitations. Imposing a rigid topology is not always desirable, since it might add excessive restrictions to the final result. Additionally, restricting the topology does not guarantee a convergence to the correct result in case of corrupted images. Alternatively, simple geometric shapes might not always be a reasonable constraint, especially for complex or variable shapes.

Strong prior knowledge tries to overcome all these issues. The strategy behind strong prior knowledge is to learn a shape and its valid variations from a training data set. This approach usually adds a regularization term to the minimization process, which penalizes curves that are too different from the shapes observed in the training set. This regularization term compares two shapes using a similarity distance or metric.

An important work that uses this approaches was proposed by Tsai et al. (2003). They generated a template shape from a training set using principal component analysis. At each iteration, the principal components of the template are compared with the evolving curve. Since the concept of shape is invariant to translations, scaling and rotations, the template and the evolving object must be aligned before comparing them through the

regularization term. Tsai et al. solved this issue by alternating the level set segmentation with a registration process. Despite reporting major improvements in the segmentation of complex structures, the processing time of Tsai's algorithm increases considerably and the final result varies depending on how the registration parameters are optimized. Moreover, the number of variables increases since new parameters need to be estimated: the template pose parameters and the weight of each principal component.

In summary, in level set segmentation shape prior knowledge can be weak or strong. Whereas weak prior knowledge is easy to implement, demands less computational resources and runs considerably faster than strong prior knowledge, the latter dramatically improves segmentation accuracy, it can accommodate more complex and variable shapes, and is robust to severely corrupted images, but this comes at a higher computational cost and algorithm complexity (e.g., number of parameters). This is the key trade-off that needs to be analyzed in any problem that requires shape prior knowledge.

Particularly, medical image segmentation proposes a continuous challenge of developing new and more robust methods to get more accurate results. To achieve this, shape anatomy information can be included as prior knowledge. It is important to choose a good balance between method complexity and the degree of automation and accuracy, given the characteristics of the searched shape.

In concrete, this thesis presents two ways of incorporating prior knowledge into segmentation problems using level sets. The first one corresponds to an important task in medical applications, the segmentation of left and right ventricles. These ventricles are

segmented in different points of the cardiac cycle in order to measure clinically relevant variables, such as the blood volume ejected by the heart. This is usually done through a manual segmentation performed by an expert, resulting in a tedious and time-consuming process. Even though there are methods that have automatized this task, they required training data and they segment both ventricles in different steps, thereby increasing the processing time. Here we propose a level set segmentation method combined with weak prior knowledge that allows for simultaneously segmenting the left and right ventricles in magnetic resonance images. This application shows that weak prior knowledge can be very useful and presents several advantages in terms of accuracy and reduction of human intervention.

The second development involves strong prior knowledge. Our development is an extension of the method proposed by Cremers et al. (2006). Cremers' method corresponds to a strong prior knowledge approach based on level sets that can learn the shape of objects of interest from training examples. Since the concept of "shape" is invariant to pose changes, Cremers et al. integrated an intrinsic alignment procedure into their level set formulation. However, the procedure only accounts for translation and scale invariance. In our approach we include translation, scale and also rotation invariance. This implies a new formulation and a new set of equations that drive the segmentation process. This improvement leads to accurate results and paves the way for a wide range of medical segmentation applications.

The remainder of this thesis is structured as follows: Section 1.1 includes objectives and hypothesis. Chapter 2 review shape prior knowledge in the context of a wide family of segmenting curves, called active contours. Chapter 3 presents the first application related to weak prior knowledge and Chapter 4 presents the new strong prior knowledge method.

1.1. Objectives and hypothesis

The objective of this thesis is to develop strategies for medical segmentation problems, which can deal with all those challenging issues. In particular, we choose the level set segmentation framework because it can deal with noise and low resolution, which are common factors in almost all applications. Furthermore, since the visual information available within medical images may not be enough to perform an accurate segmentation, we need to add shape prior knowledge to improve level set segmentation results. We propose two different approaches: (1) to solve a specific, but clinically relevant problem, combining existing techniques and (2) to develop a more general and novel approach to be used in a wider range of medical segmentation applications. In this context, as specific goals, we aim to explore and evaluate, weak and strong prior knowledge methods. Our hypothesis has two-fold. First, the use of a segmentation method with topology preservation as a weak prior knowledge, would allow the simultaneous segmentation of the right and left ventricles of the heart. Second, the use of strong prior knowledge based on shape distribution and intrinsic alignment would solve a wide range of segmentation problems, including poorly defined edges, noise and other commonly encountered artifacts within medical imaging.

2. ACTIVE CONTOURS AND PRIOR SHAPE KNOWLEDGE

2.1. Parametric active contours

Active contours were introduced by Kass et al. (1988) and consist on a segmentation technique in which a curve moves driven by two forces, (1) an internal force or a smooth term, and (2) an external force, i.e., a force that depends on the image. Initially, the proposed external force was an edge-detector, i.e., an attractor that depends on the gradient of the image. Let define the parametrized curve $\mathcal{C}(s) : [0, 1] \to \mathbb{R}^2$, and let $u_0 : \Omega \subset \mathbb{R}^2 \to \mathbb{R}$ be an image, the parametric active contour problem was formulated as follows,

$$\underset{C(s)}{\operatorname{argmin}} E_{snake} = \underset{C(s)}{\operatorname{argmin}} \frac{a_1}{2} \int_0^1 |\mathcal{C}'(s)|^2 \mathrm{d}s + \frac{a_2}{2} \int_0^1 |\mathcal{C}''(s)|^2 \mathrm{d}s - \int_0^1 |\nabla u_0(\mathcal{C}(s))|^2 \mathrm{d}s, \tag{2.1}$$

where $a_1/2$ and $a_1/2$ are positive scaling parameters. The first two terms in 2.1 control the smoothness of the curve and the third one attracts the curve to the high intensity image gradients, i.e., the edges of the image.

Although this approach presents several drawbacks that we will discussed later, the core of this variational approach is kept intact across all active contours techniques. The idea is to minimize an energy functional such as (2.1), which is equivalent to move a curve towards edges. This is done in a gradient descent direction of that functional. This idea

leads us to write an evolution equation as follows,

$$\frac{\mathrm{d}E_{snake}(\mathcal{C}(s))}{\mathrm{d}t} = -\frac{\mathrm{d}E_{snake}(\mathcal{C}(s))}{\mathrm{d}\mathcal{C}(s)} = \nabla|\nabla F(u,v)|^2 + a_1\mathcal{C}''(s) - a_2\mathcal{C}''''(s) \qquad (2.2)$$

The main drawbacks of this algorithm are two: (1) Due to the presence of noise, particularly critical in medical images, the edge-detector approach is not robust and solutions erroneously converge to local minima. Thus, more sophisticated solutions were proposed, such as Balloons (Cohen, 1991; Cohen & Cohen, 1993), which adds a preferential direction of curve motion, gradient vector flow (Santarelli, Positano, Michelassi, Lombardi, & Landini, 2003; Xu & Prince, 1998b, 1998a; Paragios, Mellina-Gottardo, & Ramesh, 2001; Cheng & Foo, 2006; El-Berbari et al., 2007; Pieciak, 2012), which takes into account a more global behavior of the image gradient instead of only its local behavior, and region-based approach (Cremers, Tischhäuser, Weickert, & Schnör, 2002; S. C. Zhu & Yuille, 1996; Chakraborty, Staib, & Duncan, 1996; Ronfard, 1994), which finally replace the edge-detection by another energy term that enforces the inside and outside the curve to have homogenous intensities, respectively. (2) Due to the nature of the curve parametrization, Kass et al. formulation does not allow topological changes of the curve, i.e., no contour splitting or merging are possible. To overcome this McInerney and Terzopoulos (2000) developed a set of ad-hoc rules to discriminate whereas the curve need to be split or merged-. Alternative solution introduced geodesic or level set-based active contours (Caselles et al., 1993, 1997; Chan & Vese, 2001).

2.2. Level set framework

Level set segmentation is an active contour-based algorithm that has been extensively used in many image applications. Osher and Sethian (Osher & Sethian, 1988) were the first ones to introduce level sets, in order to model front propagation and surface motion problems. Level set methods evolve a smooth curve driven by image features. The curve $\mathcal{C}$ lying in the image domain $\Omega \subset \mathbb{R}^n$ does not depend on a set of parameters or control points, but is implicitly defined using a signed distance function (also called embedding function) $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$, such that

$$\mathcal{C} = \{\vec{x} \in \Omega : \phi(\vec{x}) = 0\}, \tag{2.3}$$

$$interior(\mathcal{C}) = \{\vec{x} \in \Omega : \phi(\vec{x}) > 0\}, \tag{2.4}$$

$$exterior(\mathcal{C}) = \{\vec{x} \in \Omega : \phi(\vec{x}) < 0\}. \tag{2.5}$$

The definition of $\phi(\vec{x})$ as a signed distance function has interesting implications. Firstly, since $\phi(\vec{x})$ is a level set function, we know that $\frac{\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|} = \hat{n}$, which means that the gradient of ϕ at every point is normal to the level curves or sets, and importantly for the zero level curve, we can compute its normal direction by taking $\nabla\phi(\vec{x})$. That is a relevant feature since the curve $\phi(\vec{x}) = 0$ is moved by perturbing its normal. Secondly, $\phi(\vec{x})$ is signed, which means that it is trivial to know weather a point is inside or outside of it. Thirdly, $\phi(\vec{x})$ is a distance function which means that globally $|\nabla\phi(\vec{x})| = 1$, but this condition is

not satisfied locally, for example near the zero level curve. Thus, particular care need to be taken for assuming that and in several cases it is not recommended to assume it.

Let $\vec{V}(\vec{x},t)$ a 2D velocity field. In order to move our curve of interest $\phi(\vec{x})=0$, we can use the convection equation,

$$\frac{\partial}{\partial t}\phi(\vec{x})+\vec{V}(\vec{x},t)\cdot\nabla\phi(\vec{x})=0 \tag{2.6}$$

which is basically the projection of the velocity $\vec{V}(\vec{x},t)$ in the direction normal to the zero level curve. If we separate the velocity in tangential and normal components with respect to the zero level curve of $\phi(\vec{x})$, $\vec{V}(\vec{x},t)=V_t(\vec{x},t)\hat{t}+V_n(\vec{x},t)\hat{n}$, and we know that $\hat{t}\cdot\nabla\phi(\vec{x})=0$, then we rewrite 2.6,

$$\begin{aligned}
&\frac{\partial}{\partial t}\phi(\vec{x})+(V_t(\vec{x},t)\hat{t}+V_n(\vec{x},t)\hat{n})\cdot\nabla\phi(\vec{x})=\frac{\partial}{\partial t}\phi(\vec{x})+V_n(\vec{x},t)\frac{\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\cdot\nabla\phi(\vec{x})\\
&\frac{\partial}{\partial t}\phi(\vec{x})+V_n(\vec{x},t)|\nabla\phi(\vec{x})|=0
\end{aligned} \tag{2.7}$$

which is called the level set equation. To complete the evolution, we need to define an initial condition $\phi(\vec{x},t=0)=\phi_0(\vec{x})$.

A special case is the mean curvature motion. In this case the velocity is proportional to the mean curvature of the zero level curve, which is $V_n(\vec{x},t)=-\mathrm{div}\left(\frac{\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\right)$, and

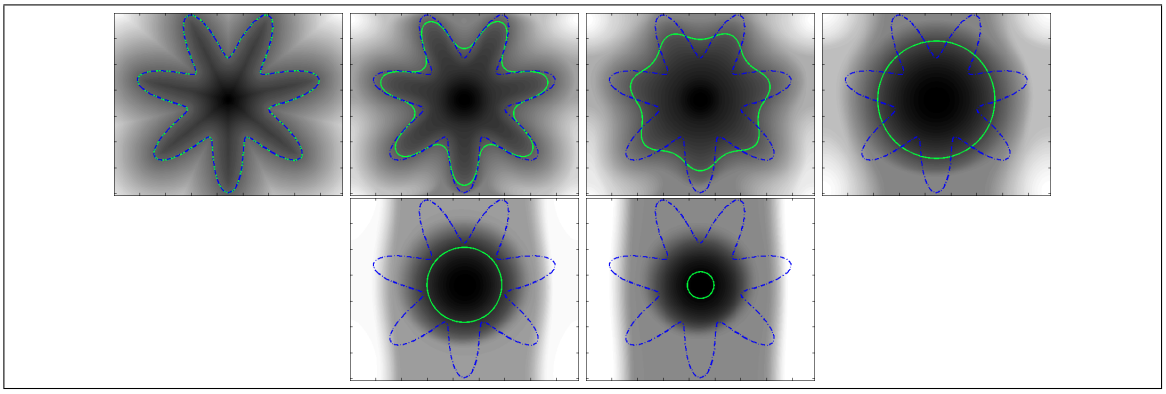

Figure 2.1. Mean curvature flow evolution process (from left to right) of a star-like shape. The blue dashed line shows the initial curve $\phi_0(\vec{x})$ and the green curve shows the evolution of the zero level curve. The zero level curve moves proportionally to its mean curvature, so firstly it is smoothed until it reaches a circular shape and then, it starts to shrink.

the evolution equation is given by 2.8,

$$\begin{cases} \dfrac{\partial}{\partial t}\phi(\vec{x}) = \operatorname{div}\left(\dfrac{\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\right)|\nabla\phi(\vec{x})| \\ \\ \phi(\vec{x},0) = \phi_0(\vec{x}) \end{cases} \tag{2.8}$$

In practice, this equation is very important since it smooths the zero level curve until it converges into a circle and then it starts to shrink until it collapses to a point and finally disappear (Fig. 2.1).

However, in order to apply this framework for segmenting images, we need to add forces that depend on the image. Therefore, the mean curvature motion can be combined with an edge detector $g(|\nabla u_0|)$, and a balloon force ν, leading to the so-called geometric

active contours (Caselles et al., 1993),

$$\begin{cases} \dfrac{\partial}{\partial t}\phi(\vec{x}) = g(|\nabla u_0|)\mathrm{div}\left(\dfrac{\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|} + \nu\right)|\nabla\phi(\vec{x})| \\ \\ \phi(\vec{x},0) = \phi_0(\vec{x}) \end{cases} \tag{2.9}$$

This extra term allows the curve to stop in a desired feature (i.e., an edge) avoiding it to collapse and disappear as in the traditional mean curvature approach.

Although this approach was original and innovative, it shows to be unstable, especially in the presence of noise. A big jump in this field was the Chan and Vese (Chan & Vese, 2001) level set segmentation algorithm, which moves away from edge detectors and introduces a region-based energy functional. The proposed functional is,

$$\begin{aligned} &\min_{\phi} E_{CV}(\bar{u}_+, \bar{u}_-, \phi) = \min_{\phi} \mu \int_{\Omega} \delta(\phi(\vec{x}))\,|\nabla\phi(\vec{x})|\,\mathrm{d}\vec{x} + \\ &\lambda_1 \int_{\Omega} (u_0(\vec{x}) - \bar{u}_+)^2 H(\phi(\vec{x}))\mathrm{d}\vec{x} + \lambda_2 \int_{\Omega} (u_0(\vec{x}) - \bar{u}_-)^2 (1 - H(\phi(\vec{x})))\mathrm{d}\vec{x}, \end{aligned} \tag{2.10}$$

where the region inside the curve is a differentiable approximation of the Heaviside function $H(\phi(\vec{x}))$, and the region outside the curve $1 - H(\phi(\vec{x}))$, $\delta(\phi)$ is the derivative of the Heaviside (the Dirac Delta); $u_0(\vec{x})$ is the image to be segmented, $\bar{u}_+$ and $\bar{u}_-$ are the average image intensity inside and outside the curve, respectively, defined as follows,

$$\begin{aligned} \bar{u}_+ &= \frac{\int_{\Omega} u_0(\vec{x}) H(\phi(\vec{x}))\mathrm{d}\vec{x}}{\int_{\Omega} H(\phi(\vec{x}))\mathrm{d}\vec{x}} \\ \bar{u}_- &= \frac{\int_{\Omega} u_0(\vec{x})(1 - H(\phi(\vec{x})))\mathrm{d}\vec{x}}{\int_{\Omega}(1 - H(\phi(\vec{x})))\mathrm{d}\vec{x}}, \end{aligned} \tag{2.11}$$

and μ, λ_1 and λ_2 are weighting terms that need to be specified by the user.

The first term of (2.10) represents the curve length, and its minimization enforces the resulting curve to be smooth. The second and third terms of this equation penalize the variance of the image intensities inside and outside the curve, producing two disjoint and homogeneous regions.

The minimization of this energy functional is solved following the curve evolution equation,

$$\frac{\partial \phi(\vec{x})}{\partial t} = -\frac{\partial E_{CV}(\phi)}{\partial \phi} = 0. \tag{2.12}$$

This derivative is calculated using the Gâteaux derivative of ∂E_{CV} in the $\tilde{\phi}(\vec{x})$ direction,

$$\left.\frac{\partial E_{CV}(\phi)}{\partial \phi}\right|_{\tilde{\phi}} = \lim_{\varepsilon \to 0} \frac{1}{\varepsilon}(E_{CV}(\phi + \varepsilon\tilde{\phi}) - E_{CV}(\phi)) = \left.\frac{\mathrm{d}}{\mathrm{d}\varepsilon} E_{CV}(\phi + \varepsilon\tilde{\phi})\right|_{\varepsilon=0}, \tag{2.13}$$

and then, we can calculate our derivative of interest as,

$$\left.\frac{\partial E_{CV}(\phi)}{\partial \phi}\right|_{\tilde{\phi}} = \int_{\Omega} \tilde{\phi}(\vec{x}) \frac{\partial E_{CV}(\phi)}{\partial \phi} \mathrm{d}\vec{x}. \tag{2.14}$$

Using (2.13) in (2.10), we have,

$$\begin{aligned}\left.\frac{\partial E_{CV}(\phi)}{\partial \phi}\right|_{\tilde{\phi}} =& \frac{\mathrm{d}}{\mathrm{d}\varepsilon}\Bigg(\mu \int_{\Omega} \delta(\phi(\vec{x})+\varepsilon\tilde{\phi}(\vec{x}))\left|\nabla(\phi(\vec{x})+\varepsilon\tilde{\phi}(\vec{x}))\right| \mathrm{d}\vec{x} \\ &+ \lambda_1 \int_{\Omega} (u_0(\vec{x})-\bar{u}_+)^2 H(\phi(\vec{x})+\varepsilon\tilde{\phi}(\vec{x}))\mathrm{d}\vec{x} \\ &+ \lambda_2 \int_{\Omega} (u_0(\vec{x})-\bar{u}_-)^2 (1-H(\phi(\phi(\vec{x})+\varepsilon\tilde{\phi}(\vec{x}))))\mathrm{d}\vec{x}\Bigg)\Bigg|_{\varepsilon=0}, \end{aligned} \tag{2.15}$$

being the derivative of the last two terms straightforwards. Isolating the derivative respect to ε we get the first term as,

$$\begin{aligned}&\frac{\mathrm{d}}{\mathrm{d}\varepsilon}\Big(\delta(\phi(\vec{x})+\varepsilon\tilde{\phi}(\vec{x}))|\nabla(\phi(\vec{x})+\varepsilon\tilde{\phi}(\vec{x}))|\Big)\Big|_{\varepsilon=0} = \Big(\delta'(\phi(\vec{x}))|\nabla(\phi(\vec{x})+\varepsilon\tilde{\phi}(\vec{x}))| \\ &+ \delta(\phi(\vec{x}))\frac{\mathrm{d}}{\mathrm{d}\varepsilon}|\nabla(\phi(\vec{x})+\varepsilon\tilde{\phi}(\vec{x}))|\Big)\Big|_{\varepsilon=0}. \end{aligned} \tag{2.16}$$

For the last term of 2.16, without lose of generality, we assume $\phi(\vec{x}) = \phi(x,y) \in \mathbb{R}^2$, and using the abbreviate notation for $\dfrac{\partial\phi(x,y)}{\partial x} = \phi_x$ and $\dfrac{\partial\phi(x,y)}{\partial y} = \phi_y$, then,

$$\begin{aligned}&\frac{\mathrm{d}}{\mathrm{d}\varepsilon}|\nabla(\phi(x,y)+\varepsilon\tilde{\phi}(x,y)|\Big|_{\varepsilon=0} = \frac{\mathrm{d}}{\mathrm{d}\varepsilon}|(\phi_x,\phi_y)+\varepsilon(\tilde{\phi}_x,\tilde{\phi}_y)|\Big|_{\varepsilon=0} \\ &= \frac{\mathrm{d}}{\mathrm{d}\varepsilon}\sqrt{(\phi_x-\varepsilon\tilde{\phi}_x)^2+(\phi_y-\varepsilon\tilde{\phi}_y)^2}\Big|_{\varepsilon=0} \approx \frac{\mathrm{d}}{\mathrm{d}\varepsilon}\sqrt{\phi_x^2-2\varepsilon\tilde{\phi}_x\phi_x+\phi_y^2-2\varepsilon\tilde{\phi}_y\phi_y}\Big|_{\varepsilon=0} + \mathcal{O}(\varepsilon^2) \\ &- \frac{-2\tilde{\phi}_x\phi_x-2\tilde{\phi}_y\phi_y}{2\sqrt{\phi_x^2-2\varepsilon\tilde{\phi}_x\phi_x+\phi_y^2-2\varepsilon\tilde{\phi}_y\phi_y}}\Big|_{\varepsilon=0} = \frac{\tilde{\phi}_x\phi_x+\tilde{\phi}_y\phi_y}{\sqrt{\phi_x^2+\phi_y^2}} = \frac{\nabla\phi(\vec{x})\cdot\nabla\tilde{\phi}(\vec{x})}{|\nabla\phi(\vec{x})|}. \end{aligned} \tag{2.17}$$

Replacing (2.16) and (2.17) in (2.15), we obtain,

$$\left.\frac{\partial E_{CV}(\phi)}{\partial \phi}\right|_{\tilde{\phi}} = \Bigg(\mu \int_\Omega \delta'(\phi(\vec{x})\tilde{\phi}(\vec{x})|\nabla\phi(\vec{x})| + \delta(\phi(\vec{x}))\frac{\nabla\phi(\vec{x})\cdot\nabla\tilde{\phi}(\vec{x})}{|\nabla\phi(\vec{x})|}\mathrm{d}\vec{x}$$
$$+ \lambda_1 \int_\Omega |u_0(\vec{x}) - \bar{u}_+|^2\,\delta(\phi(\vec{x}))\tilde{\phi}(\vec{x})\mathrm{d}\vec{x} + \lambda_2 \int_\Omega (u_0(\vec{x}) - \bar{u}_-)^2\delta(\phi(\vec{x}))\tilde{\phi}(\vec{x})\mathrm{d}\vec{x}\Bigg). \quad (2.18)$$

We need to transform the $\nabla\tilde{\phi}(\vec{x})$ into something that depends explicitly on $\tilde{\phi}(\vec{x})$ to complete the form of (2.14). Using the Green's first identity we get,

$$\int_\Omega \delta(\phi(\vec{x}))\frac{\nabla\phi(\vec{x})\cdot\nabla\tilde{\phi}(\vec{x})}{|\nabla\phi(\vec{x})|}\mathrm{d}\vec{x} = \oint_{\partial\Omega} \tilde{\phi}(\vec{x})\delta(\phi(\vec{x}))\frac{\nabla\phi(\vec{x})\cdot\hat{n}}{|\nabla\phi(\vec{x})|}\mathrm{d}s$$
$$- \int_\Omega \tilde{\phi}(\vec{x})\nabla\cdot\frac{\delta(\phi(\vec{x}))\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\mathrm{d}\vec{x} = \oint_{\partial\Omega} \tilde{\phi}(\vec{x})\delta(\phi(\vec{x}))\frac{\nabla\phi(\vec{x})\cdot\hat{n}}{|\nabla\phi(\vec{x})|}\mathrm{d}s$$
$$- \int_\Omega \tilde{\phi}(\vec{x})\delta'(\phi(\vec{x}))\frac{\nabla\phi(\vec{x})\cdot\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\mathrm{d}\vec{x} - \int_\Omega \delta(\phi(\vec{x}))\tilde{\phi}(\vec{x})\frac{\nabla\cdot\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\mathrm{d}\vec{x}$$
$$= \oint_{\partial\Omega} \tilde{\phi}(\vec{x})\delta(\phi(\vec{x}))\frac{\nabla\phi(\vec{x})\cdot\hat{n}}{|\nabla\phi(\vec{x})|}\mathrm{d}s - \int_\Omega \tilde{\phi}(\vec{x})\delta'(\phi(\vec{x}))|\nabla\phi(\vec{x})|\mathrm{d}\vec{x}$$
$$- \int_\Omega \tilde{\phi}(\vec{x})\delta(\phi(\vec{x}))\frac{\nabla\cdot\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\mathrm{d}\vec{x}. \quad (2.19)$$

For the line integral of (2.19), we can impose the boundary condition $\delta(\phi(\vec{x}))\nabla\phi(\vec{x})\cdot\hat{n} = \delta(\phi(\vec{x}))\frac{\partial\phi(\vec{x})}{\partial\hat{n}} = 0$, on $\partial\Omega$. Finally, replacing the rest of (2.19) in (2.18),

$$\left.\frac{\partial E_{CV}(\phi)}{\partial \phi}\right|_{\tilde{\phi}} = \mu \int_\Omega \tilde{\phi}(\vec{x})\delta'(\phi(\vec{x}))|\nabla\phi(\vec{x})|\mathrm{d}\vec{x} - \mu \int_\Omega \tilde{\phi}(\vec{x})\delta'(\phi(\vec{x}))|\nabla\phi(\vec{x})|\mathrm{d}\vec{x}$$
$$- \mu \int_\Omega \tilde{\phi}(\vec{x})\delta(\phi(\vec{x}))\frac{\nabla\cdot\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\mathrm{d}\vec{x} + \lambda_1 \int_\Omega (u_0(\vec{x}) - \bar{u}_+)^2\delta(\phi(\vec{x}))\tilde{\phi}(\vec{x})\mathrm{d}\vec{x}$$
$$- \lambda_2 \int_\Omega (u_0(\vec{x}) - \bar{u}_-)^2\delta(\phi(\vec{x}))\tilde{\phi}(\vec{x})\mathrm{d}\vec{x}, \quad (2.20)$$

and using (2.14), we can write the partial derivative of $E_{CV}(\phi(\vec{x}))$ as,

$$\frac{\partial E_{CV}(\phi)}{\partial \phi} = \delta(\phi(\vec{x}))\Big[-\mu \text{div}\left(\frac{\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\right) + \lambda_1(u_0(\vec{x}) - \bar{u}_+)^2 \quad -\lambda_2(u_0(\vec{x}) - \bar{u}_-)^2\Big], \tag{2.21}$$

Finally, using (2.12), we can write the curve evolution equation as follows,

$$\begin{cases} \dfrac{\partial\phi(\vec{x})}{\partial t} = \delta(\phi(\vec{x}))\left[\mu \ \text{div}\left(\dfrac{\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\right) - \lambda_1(u_0(\vec{x}) - \bar{u}_+)^2 + \lambda_2(u_0(\vec{x}) - \bar{u}_-)^2\right] \\ \delta(\phi(\vec{x}))\nabla\phi(\vec{x}) \cdot \hat{n} = \delta(\phi(\vec{x}))\dfrac{\partial\phi(\vec{x})}{\partial\hat{n}} = 0, \text{on } \partial\Omega \\ \phi(\vec{x}, 0) = \phi_0(\vec{x}) \end{cases} \tag{2.22}$$

All the evolution equation is multiplied by $\delta(\phi(\vec{x}))$ (i.e. the derivative of a differentiable version of the Heaviside function), which corresponds to the zero level of $\phi(\vec{x})$. The term $\text{div}\left(\dfrac{\nabla\phi(\vec{x})}{|\nabla\phi(\vec{x})|}\right)$ corresponds to the mean curvature, so the curve has to be smooth as in 2.8. The terms $(u_0(\vec{x}) - \bar{u}_+)^2$ and $(u_0(\vec{x}) - \bar{u}_-)^2$ force the curve to move so that it encloses only those pixels whose intensities are similar to $\bar{u}_+$ and it excludes those intensities similar to $\bar{u}_-$. This model, also known as active contours without edges, can be understood as the piecewise constant version of the Mumford-Shah functional (Mumford & Shah, 1989), and it has been extensible applied in several applications (see chapter 3), as well as alternative adaptations (Daněk, Matula, Maška, & Kozubek, 2012; Duan, Pa, Yin, Wei, & Wang, 2014), including a convex version of the energy functional (Brown, Chan, & Bresso, 2012).

An interesting feature of this formulation is that it is independent of the image dimensions. The only change that needs to be done to extend the model from 2D to higher dimensions is to re-define the differential operator ∇.

The evolution equation (2.22)is solved in practice as an iterative process. This produces $\phi(\vec{x})$ to move away from being a distance function and it tends to become flat. A typical solution is to reinitialize $\phi(\vec{x})$ so that to make it becoming a signed distance function satisfying the global condition $|\nabla\phi(\vec{x})| = 1$. The standard approach for the reinitialization problem is to solve the Hamilton-Jacobi equation (Duan, Pa, Yin, Wei, & Wang, 1994),

$$\begin{cases} \dfrac{\partial\psi(\vec{x})}{\partial t} = \text{sign}(\phi(\vec{x}))(1 - |\nabla\psi(\vec{x})|) \\ \\ \psi(\vec{x}, t = 0) = \phi(\vec{x}) \end{cases} \tag{2.23}$$

This equation move the level curves without changing the sign of $\phi(\vec{x})$, until it reaches the steady state, i.e., satisfying $|\nabla\psi(\vec{x})| = 1$. Unfortunately, the numerical errors produce that in practice the zero level curve suffers small perturbations. Therefore, attention needs to be taken into the discretization scheme, since oscillatory behavior of the process could lead to unstable solutions (Osher & Fedkiw, 2003).

This brief active contour review brings the background to understand level set basics. To move forward, we are going to introduce the concept prior shape knowledge, i.e., and how to add information about known shapes into the evolution process.

2.3. Prior knowledge in active contours

In many image applications the visual information is not enough to segment correctly a particular structure. In medical images this issue is critical, since we normally find high level of noise, acquisition artifacts (e.g., motion), reconstruction artifacts (e.g., aliasing, ghosting, Gibbs rings) and partial volume effects (e.g., smooth and blurred edges). All those issues motivated early developments of shape prior knowledge with parametric active contours. For example, Pentland and Sclaroff (1991) adds mechanical restrictions to the moving curve based on nodes connected by springs with certain mass and elasticity.

Cootes et al. (1995) developed the so-called active shape models. This methods requires shapes defined by corresponding landmarks as a training set. The algorithm constrains the parametric curve to shape variations represented by a linear combination of principal components of the landmarks covariance matrix. Since in many cases shape variations cannot be represented as linear combinations of principal components, alternative non-linear models were developed (Cremers, Kohlberger, & Schnörr, 2003; Romdhani, Gong, & Psarrou, 1999; Twining & Taylor, 2001).

Cremers et al. (1996) proposed a 2D linear methods based on shape priors given by non-correspondening landmarks of a curve parametrized as a B-spline. Tejos et al. (2009) extended this approach for 3D using simplex meshes as surface parametrization.

Prior knowledge was later introduced in level set-based methods. For instance, Han et al. (2003) proposed a preserving topology level set, which keeps constant the number of

initial curves. In certain way, this formulation emulates the behavior of parametric active contours, since no splitting nor merging is allowed, but with the advantage of naturally evolving several curves simultaneously, which allows to segment two or more disjoint objects in the same evolution process. Pluempitiwiriyawej et al. (Pluempitiwiriyawej et al., 2005) fit an elliptical shape to the curve during the evolution process. Those approaches can be referred as weak prior knowledge, since simple geometric assumptions are made and they do not need any training set.

Approaches that extract information of a set of training shapes (i.e., strong prior knowledge) have been developed using different statistical models. For instance, Leventon et al. (Leventon, Grimson, & Faugeras, 2000) used principal component analysis (PCA) to consider an average shape and its variations in the implicit function space, combined with a geodesic level set formulation (Caselles et al., 1997). A Gaussian distribution of shapes knowledge is assumed. Here, an extra optimization step (registration process) over pose parameters (scaling, translation and rotation) was introduce, which is needed to compare the shape embedded by $\phi(\vec{x})$ with the PCA model. Tsai et al. (2003) used a similar model, including optimization of eigenmodes weights of a PCA model. This adds new variables to the registration, and the results vary depending on the order in which all those variables are optimized. Chen et al. (2002) introduced an average shape using a parametric model and they evolved the contour projecting it onto the shape space using a registration process that considers scaling, translation and rotation. Rousson and Paragios (2008) presented an approach that generates, in a first step, an optimal model from the training data and also a

confidence map of this model with respect to the evolving curve. In a second step, it incorporates the prior shape knowledge into a region-based algorithm through a registration process.

Cremers et al. (2006) proposed two new ideas: (1) an intrinsic alignment of shapes so that to avoid a registration process, and (2) integrating the information of several shapes using a kernel density estimation, which assumes an arbitrary distribution of the shape (avoiding the generation of models, average shape and PCA). This approach has been successfully applied in medical image problems (S. Chen & Radke, 2009; Aslan et al., 2011; Yeo, Xie, Sazonov, & Nithiarasu, 2011). Although this approach overcomes several issues, it only accounts for scaling and translation, leaving out the rotation.

These works represent the core of shape prior knowledge applied to segmentation in the active contours context. Even today, most of the works are influenced by these approaches, including convex formulation. There is a related field in level sets called shape optimization. This field also produced many advances regarding shapes, but it is neither trivial nor clear how to apply that kind of knowledge to the segmentation task. Therefore, shape optimization is out of the scope of this thesis.

3. SIMULTANEOUS LEFT AND RIGHT VENTRICLE SEGMENTATION USING TOPOLOGY PRESERVING LEVEL SETS

This chapter shows an application of weak prior knowledge for cardiac segmentation. Particularly, the goal of this method is to segment simultaneously, accurately, fast and with high level of automation, the left and right ventricles of the heart. The first section (3.1) introduces the cardiac segmentation problem, the challenges involved, and a review of different approaches that have used Active Contours. Section 3.2 presents the methods related to our weak prior knowledge algorithm and also the experiments and statistics tools to measure the performance and validate the proposed method. Section 3.3 shows the results and Section 3.4 discusses these results and analyzes the advantages and disadvantages of the proposed method.

3.1. Introduction to cardiac segmentation

The assessment of cardiac function is highly relevant for several cardiovascular diseases. Cardiac performance is typically evaluated using Cardiovascular Magnetic Resonance (CMR) imaging since it has been shown to be accurate and reproducible for normal and abnormal hearts (Bellenger et al., 2000). This technique allows for computing accurately the Left Ventricle (LV) volume and it has also shown to be a reliable tool to measure the Right Ventricle (RV) volume (Petitjean & Dacher, 2011). 2D multi-slice cine images of the heart, using balanced Steady-State Free Precession (SSFP) sequences, are the most widely used CMR technique for measuring ventricular volumes (Carr, Simonetti, Bundy, & et al., 2001; Thiele, Nagel, Paetsch, & et al., 2001; Grothues et al., 2002).

Once the cine images are acquired, they need to be processed by an expert. The expert has to identify the end-systolic and end-diastolic frames from the cine sequence, and then, perform a manual image segmentation in order to compute the ventricular volumes, which are usually calculated using a method based on the Simpsons rule (Ino, Benson, Mikalian, Freedom, & Rowe, 1989). In the case of CMR, this is defined as the sum of the cross-sectional areas of each slice considering the slice thickness and the spacing between slices (Souto, Dias, et al., 2013). This is currently the most accurate and robust method to measure the ventricular volumes (Jahnke et al., 2011). Unfortunately, manual segmentation is a tedious and labor-intensive process and has high inter- and intra-observer variability. Promising alternatives have started to arise using different physical constraints and basis functions to model the cardiac volume (Tavakoli & Amini, 2013).

A wide variety of image segmentation techniques exists in order to automate and speed up this process (Petitjean & Dacher, 2011; W. Zhu, Ha Kang, & Biros, 2013; Santarelli et al., 2003; Xu & Prince, 1998b, 1998a; Paragios et al., 2001; Cheng & Foo, 2006; El-Berbari et al., 2007; Pieciak, 2012; Tavakoli & Amini, 2013; Mitchell et al., 2001; Mahapatra & Buhmann, 2012, 2013; Montillo, Metaxas, & Axel, 2003; Sardanelli, Quarenghi, Di Leo, Boccaccini, & Schiavi, 2008; Wang et al., 2014; Souto, Masip, et al., 2013). Those techniques involve methods based on pixel classification and active contours, with different degrees of prior knowledge. Commonly used pixel classification techniques are region growing, clustering and atlas-based methods (Petitjean & Dacher, 2011). Region growing and clustering implementations are simple and fast. However, those techniques

have several drawbacks: both of them are very sensitive to noise and therefore, they produce corrupted results; region growing methods depend on a parameter which is hard to determine for images with different contrast and intensities as it usually occurs in CMR images; and clustering-based methods do not consider spatial information and thus, they require an extra manual process to select the correct regions among those several structures that are commonly segmented. Atlas based methods consist in matching the image being segmented with a template generated from a big training data set created by experts. The main drawbacks of atlas-based methods are that they need a large training set for generating one template that tends to fail with severely abnormal shapes. This is especially important in patients with congenital heart diseases (CHD), whose the anatomy is not well represented by standard templates.

Active Contour (AC) techniques have been extensively used in cardiac segmentation. They consist in an iterative process in which a curve is deformed in order to capture a specific feature (typically the edges) of an object of interest within an image. These techniques can be classified as explicit or implicit. Explicit (or parametric) AC (Kass et al., 1988) have the advantage of preserving the topology, i.e. keeping constant the number of contours defined during the initialization process, avoiding any merge or split of different structures. However, standard explicit AC can handle the deformation of only one curve, and thus multiple boundaries are difficult to segment simultaneously (McInerney & Terzopoulos, 2000).

As we introduced before, level set based algorithms (Caselles et al., 1997) are an elegant solution for the simultaneous segmentation of multiple structures. Several authors have used them in cardiac imaging segmentation problems. They firstly tried the edge-based level sets with promising results (Gupta et al., 1993; Geiger, Gupta, Costa, & Vlontzos, 1995; W. Zhu et al., 2013) and then, other authors chose slightly different approaches using AC based on Gradient Vector Flow or modifications of it (Santarelli et al., 2003; Xu & Prince, 1998b, 1998a; Paragios et al., 2001; Cheng & Foo, 2006; El-Berbari et al., 2007; Pieciak, 2012). Those techniques have been applied to the LV with relative success, but they have not solved the simultaneous segmentation of the right and left ventricles.

An interesting level set based approach is the Stochastic Active Contour Scheme (STACS) (Pluempitiwiriyawej et al., 2005), which is a special formulation of cardiac image segmentation that combines edge-based, region-based and weak prior knowledge energies. This technique allows for segmenting LV and RV, but through separated consecutive processes. The method tries to approximate the segmentation result to an elliptical shape, which may not always yield good results, particularly in patients with CHD. Comprehensive reviews of cardiac image segmentation methods can be found in (Petitjean & Dacher, 2011; Tavakoli & Amini, 2013).

Despite the wide variety of approaches for segmenting automatically LV volumes, there is still no consensus about which is the most accurate one. Additionally, the problem of segmenting automatically RV volumes still remains unsolved (Petitjean & Dacher, 2011), especially in patients with CHDs.

LV and RV segmentation has started to be explored with more attention (Mitchell et al., 2001; Mahapatra & Buhmann, 2012, 2013; Montillo et al., 2003; Sardanelli et al., 2008; Wang et al., 2014; Souto, Masip, et al., 2013; Grosgeorge, Petitjean, Caudron, Fares, & Dacher, 2011), even as simultaneous processes. For example Mahapatra et al. (Mahapatra & Buhmann, 2012) and Mahapatra (Mahapatra & Buhmann, 2013) use mutual context information of both ventricles iteratively, fixing one of them and then optimizing the shape of the other. Grosgeorge et al. (Grosgeorge et al., 2011) used the region-based level set (Chan & Vese, 2001) in order to simultaneously segment the left and right ventricles. The Chan-Vese algorithm produced the segmentation of several structures of high intensity with different unconnected curves. For this reason, they proposed a second step in which the final segmentation only keeps the two largest components, assuming that those components corresponded to the ventricles. Although they reported interesting results, the authors did not evaluate the standard functional indexes. Additionally, they reported problems on the apical zone of the heart, due to the poor contrast obtained at the septum.

The Chan-Vese algorithm is a good option for segmenting structures with well-defined image contrasts and it has been previously used, showing interesting and promising results (Grosgeorge et al., 2011), but it tends to fail when those contrasts decrease. This problem commonly happens at apical slices, but it is even more serious in abnormal hearts, causing region-based algorithms to fail. For example, in patients with repaired tetralogy of Fallot, a severe hypertrophied right ventricle can produce MR images in which the septum can be barely distinguished from the ventricles. The resulting Chan-Vese segmentation can therefore consider LV and RV as a single structure (Fig. 3.1 (a)). This effect makes

simple rules such as that proposed by Grosgeorge et al. (Grosgeorge et al., 2011) (i.e. selecting the two largest components) fail and important human intervention is therefore needed to correct the segmentation. Similarly, in the presence of ventricular septal defects, the separation between ventricles can be missed for several slices causing the Chan-Vese algorithm to fail (Fig. 3.1 (b)).

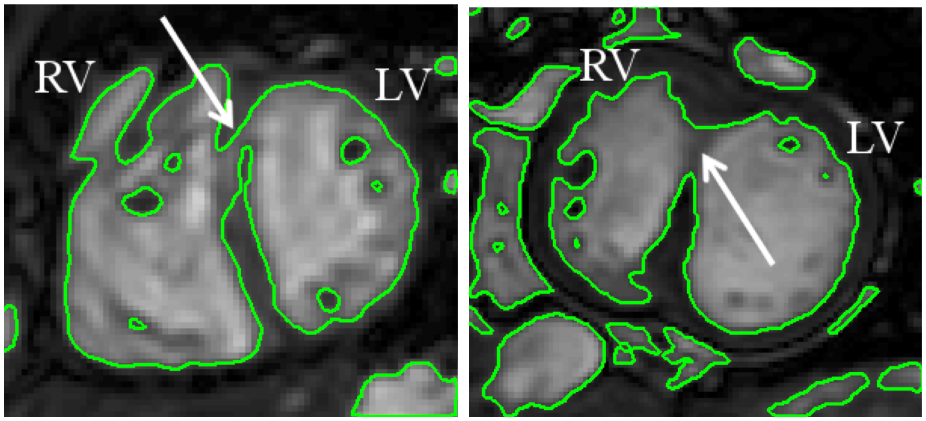

Figure 3.1. Segmentation result using Chan-Vese algorithm. (a) Due to the presence of tetralogy of Fallot, both ventricles are erroneously segmented as a single structure. (b) This patient presents an interventricular connection near the apex, which causes both ventricles to be erroneously segmented as the same structure in several slices.

The new approach proposed in this work includes additional information to the Chan-Vese segmentation technique. We combine the Chan-Vese level set framework with a topology-preserving algorithm (Han et al., 2003) to set a fixed number of structures during the entire segmentation process. This approach provides two important contributions: (1) it allows us to segment simultaneously the left and right ventricles considering them as two unconnected structures; and (2) it allows us to successfully segment the left and right

ventricles even in the presence of several abnormal anatomies, which is particularly useful for different types of CHDs. The resulting algorithm is fast, does not need any training data, and has a reasonable degree of automation since the amount of human interaction needed is minimal. We tested our approach on a database of 6 healthy volunteers and 35 CHD patients (1166 images) that emphasizes all the issues that make the segmentation task highly challenging, and we evaluated our method using clinically relevant measurements and additional metrics obtaining accurate results.

3.2. Material and methods

This section is organized as follows: Section 3.2.1 describes the segmentation framework, Section 3.2.2 details the relevant implementation steps in order to integrate the different processes of the segmentation framework, and Section 3.2.3 describes the different experiments done in order to test the performance of our proposed method.

3.2.1. Segmentation Methods: Weak prior knowledge algorithm

The Chan-Vese algorithm is topologically flexible, i.e. curves are allowed to split or merge during their evolution. This means that the number of curves might vary during the segmentation process. Therefore, the final result might include regions that do not belong to the structures of interest as seen in Fig. 3.2 (a). In order to constrain the topological changes of the curves, we included an additional weak prior knowledge algorithm that preserves the topology (Han et al., 2003). The idea is to start with two initial curves, as

in Fig. 3.2, and to evolve them while maintaining a fixed the number of curves. This algorithm is applied at the end of each Chan-Vese iteration and it tracks all the changes with respect to the previous iteration, allowing only those changes that keep the topology intact. In this extra step, pixels that entered or left the area enclosed by any curve are analyzed (i.e., those pixels that change their sign of ϕ from one iteration to the next), and the curve is allowed to move only when those pixels do not change the topology of the enclosed area. This process allows the simultaneous segmentation of the right and left ventricles, avoiding any erroneously curve split.

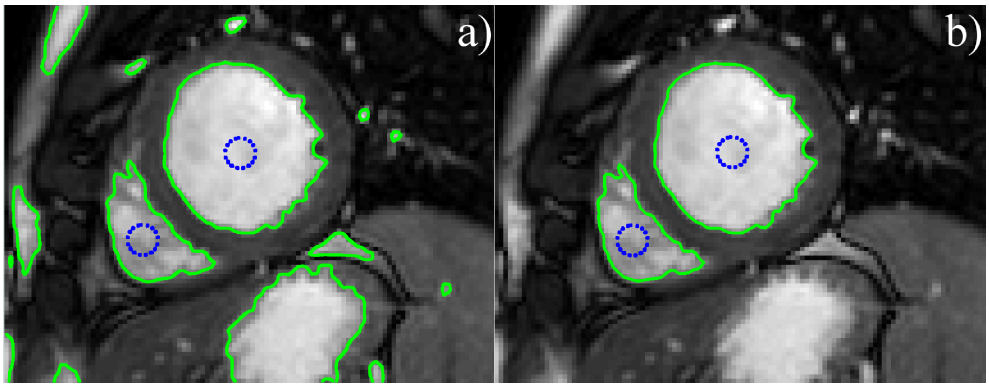

Figure 3.2. Comparison between Chan-Vese (a) without and (b) with the preserving-topology algorithm. If the topology is preserved, additional structures are discarded and only the ventricles are extracted. The blue dotted lines correspond to the initial curves and green lines correspond to the final curve.

The topology is preserved by checking a property called simple point. A simple point is a pixel that can be added or deleted from a digital object without changing the number of connected objects. The simple point property is locally calculated for every candidate to enter or leave the curve, in a neighborhood using different but complementary connectivities for foreground and background. This complementary connectivity pair is needed

in order to satisfy the Jordan curve theorem for digital topology, i.e. a simple closed curve divides the domain into two disjoint regions. In our case, we use 8-connectivity for the foreground and 4-connectivity for the background.

For example, if the center pixel (gray) in Fig. 3.3 (a) is added to the curve, that neighborhood still contains one connected object, so the topology does not change. Therefore, the curve could move and incorporate such gray pixel, since the image preserves its topology. In contrast, Fig. 3.3 (b) shows a neighborhood that contains two different objects (in white). If the center pixel (gray) is added, that neighborhood will contain only one connected object, because the center pixel will merge those two initial structures, causing a topology change. Therefore, curves could not move to incorporate such gray pixel, since the image changes its topology. A more detailed explanation of the algorithm can be found in appendix A.

3.2.2. Integrating Chan-Vese algorithm with the topology preservation step and additional implementation details

Integrating Chan-Vese with the topology preservation step is not straightforward since it heavily depends on implementations issues. These implementation details are not commonly discussed in the literature, despite of critically affecting the segmentation results. The designed integration strategy is now discussed. Firstly, since the level set curve may move significantly on each iteration, the simple point property may need to be verified for several pixels in order to decide whether or not they could get in or out from the region

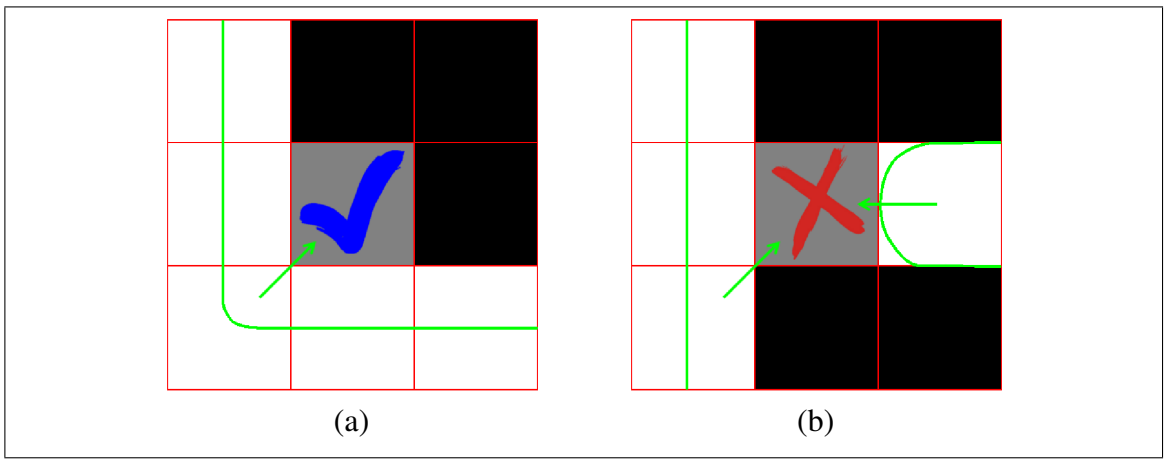

Figure 3.3. Topology preserving property: simple point. (a) Example of preserved topology. If the curve (green) is moved to include the center pixel (gray) to the object (white), the topology does not change, i.e., the number of connected objects in the neighborhood keeps constant. (b) Example of a topology that is not preserved. If the curve (green) is moved to include the center pixel (gray) to the object (white), the topology changes because the number of connected objects in the neighborhood does not keep constant. Initially the neighborhood contains two different structures. After include the center pixel to the curve, the two initial structures are merged, resulting in only one object.

enclosed by the curve. Therefore, one needs to define a consistent order to analyze those pixels. We checked first those pixels close to the segmenting curve since this allows accepting or rejecting several pixels per iteration, making the algorithm more efficient. This also ensures that the topology evaluation is performed continuously and without artificial gaps that could be erroneously considered as topological changes. Since ϕ is a signed distance function, we order the candidates according to their value of ϕ.

Secondly, after a few iterations ϕ moves away from being a distance function, producing the curve to stop in local minima. To avoid this, ϕ must be reinitialized periodically. There are a few alternatives for the reinitialization method; in our case we reinitialize our

distance function through the Hamilton-Jacobi equation (Sussman, Smereka, & Osher, 1994). The reinitialization process requires particular attention, because some smoothing or numerical approximations may cause topology changes. Including the topology preservation step into the reinitialization is not an option since the later may become unstable. A simple solution to avoid topological changes during reinitialization is to reject any sign change in ϕ. If a pixel changed its sign we kept the distance value but we switched the sign.

To stabilize the parameter setting, image intensities were re-scaled from 0 to 255. Chan-Vese parameters (Eq. 2.10) were $\mu = 0.02\times255^2$, $\lambda_1 = 1$, $\lambda_2 = 1$, $\Delta t = 5\times10^{-4}$, a maximum of 150 iterations, and reinitialization every 30 iterations. The stopping condition is that the area does not change for 15 iterations. After that, we performed the segmentation process once again modifying only $\mu = 0.01\times255^2$, $\lambda_1 = 0.8$, and $\lambda_2 = 2$, and a maximum of 50 iterations. This second process helped us to improve segmentation details especially at low-contrast regions of the right ventricle and apical slices.

For the initialization of the next slice, we define two circles of radius 3 pixels, centered on the centroids of each segmented ventricle of the previous slice.

The image intensity in Multi-slice cardiac MR tends to significantly change from one slice to the other. If that happens, segmentation parameters need to be adjusted for each slice. Alternatively, it is possible to rescale the intensities of each slice. We considered a 0 to 255 intensity rescaling for each slice using as reference the maximum of the area inside the initial curve (this assumes that the intialization will always be inside the ventricles).

Considering all these implementation details, the total computational time may vary especially for the first iterations. However, the total processing time per slice is reasonable small, i.e., in general less than a second and reaching a maximum value of three seconds for critical cases.

3.2.3. Experiments

We used 2D multi-slice short-axis cine SSFP MRI, acquired with a 1.5T scanner (Achieva; Philips Medical Systems) of 41 patients (average age 20.41 ± 12.12, 15 women) with different heart related diseases. Table 3.1 shows the patient diagnosis and highlights that our database is especially focused on repaired Tetralogy of Fallot, a very challenging CHD due to changes in the shape of the ventricles after the repairing surgery. We acquired 30 frames per cardiac cycle; 10 to 14 slices to cover the entire heart; TE and TR were 1.5 ms and 3.2 ms, respectively. The image size was 224×224 pixels, with pixel size of approximately 1.74 mm$\times$1.74 mm, and slice thickness of about 8 mm.

For our proposed method, the user is required to run a few iterations over one data set to calibrate the parameters of the algorithm. After that, these parameters are fixed for all the data. The user also has to choose the End Diastolic (ED) and End Systolic (ES) frames, and the basal and apical slices of each ventricle. The algorithm is initialized by placing in one slice a small circle with a predefined radius (3 pixels) in each ventricle (blue dotted lines in Fig. 3.2). The contours were automatically propagated through all slices. A few final manual corrections at the basal slices were typically needed, because

Disease	Number of patients
Repaired tetralogy of Fallot	18
Dextro-transposition of the great arteries	4
Hypertrophic cardiomyopathy	2
Inter-ventricular communication	2
Inter-auricular communication	1
Aortic stenosis	1
Tricuspid regurgitation	1
Alagille syndrome	1
Definite small vessel	1
Marfan syndrome	1
Ehler Danlos syndrome	1
Secundum atrial septal defect	1
Arterial trunk type I	1
Volunteers	6

Table 3.1. Patients summary

of the poor contrast produced by the mitral and tricuspid valves. Considering these two cardiac frames, the number of processed images was 1166, which is a significant number considering that our database is specialized in CHDs. We evaluated the algorithm through two experiments. In a first experiment we compared the standard functional indexes End Systolic Volume (ESV) and End Diastolic Volume (EDV) for the LV and RV obtained with our method against those obtained from segmentations done by experts using the ViewForum software (Philips Healthcare, Best, The Netherlands) and manual editions. To robustly verify that the difference between those two sets of measurements is not statistically significant, we used a two-tailored paired t-test (Santarelli et al., 2003; Souto, Masip, et al., 2013), Pearson correlation (Souto, Masip, et al., 2013) and Bland-Altman statistical plots (Santarelli et al., 2003; Mitchell et al., 2001; Sardanelli et al., 2008; Wang et al., 2014; Souto, Masip, et al., 2013; Martin Bland & Altman, 1986) for this analysis.

Since erroneously overestimated regions could cancel out erroneously underestimated regions, standard indexes can underestimate the overall error. To avoid this problem, in a second experiment we took a random subsample of 15 patients and we calculated the voxel-by-voxel correspondence of our segmentation results compared with those obtained by an expert using the freely available software Segment version 1.9 R3061 (http://segment.heiberg.se) (Heiberg et al., 2010) and manual editions. The correspondence was done using the index (Mahapatra & Buhmann, 2013; Montillo et al., 2003; Grosgeorge et al., 2011; Dice, 1945),

$$\text{Dice} = 2\frac{|A \cap B|}{|A| + |B|}, \tag{3.1}$$

where A and B represent the segmented volumes obtained with our method and the Segment software, respectively. The numerator of this expression represents the intersection of both volumes, and the denominator represents the sum of the volumes. This index corresponds to a normalized metric that gives a value of 1 when both volumes are identical, and penalizes every underestimated or overestimated voxel decreasing the output value down to 0 when there is no match between the two volumes.

3.3. Results

Fig. 3.4 shows the results of the complete segmentation process using our proposed method for one data set. For ease of visualization, we show 2D slices at the top row,

for systole (left) and diastole (right). The bottom row shows the 3D reconstructions after applying the process over the entire heart, for systole (left) and diastole (right). The simultaneous segmentation of both ventricles takes approximately 15 seconds per cardiac phase. Additionally, minor manual corrections are sometimes needed typically at the basal slices, in which it is not clear the precise location of the valves. Fig. 3.5 shows the same slice segmented with the Viewforum software (gold standard).

Our method was particularly useful in the presence of CHD, due to its ability to segment the ventricles even with unusual anatomies. Fig. 3.6 and 3.7 show the segmentation results of the two critical cases presented in the Introduction. In our case, we were able to successfully segment both ventricles simultaneously using our topology-preserving Chan-Vese algorithm.

Table 3.2 shows the p-value of a two-tailed paired t-test and the Pearsons correlation between our method and the gold standard generated using ViewForum for the 41 subjects. In all cases p-values are greater than 0.05, which means that we cannot reject the null hypothesis, i.e. that our measurements are not significantly different from the gold standard (with $P<0.05$). Pearson correlation shows a high correlation between our segmentation and the gold standard.

Bland-Altman plots (Fig. 3.8) show the statistical comparison for ESV and EDV of the LV and RV obtained with our method and the manually assisted segmentation using ViewForum. As can be observed, our method shows almost no bias and except from a few

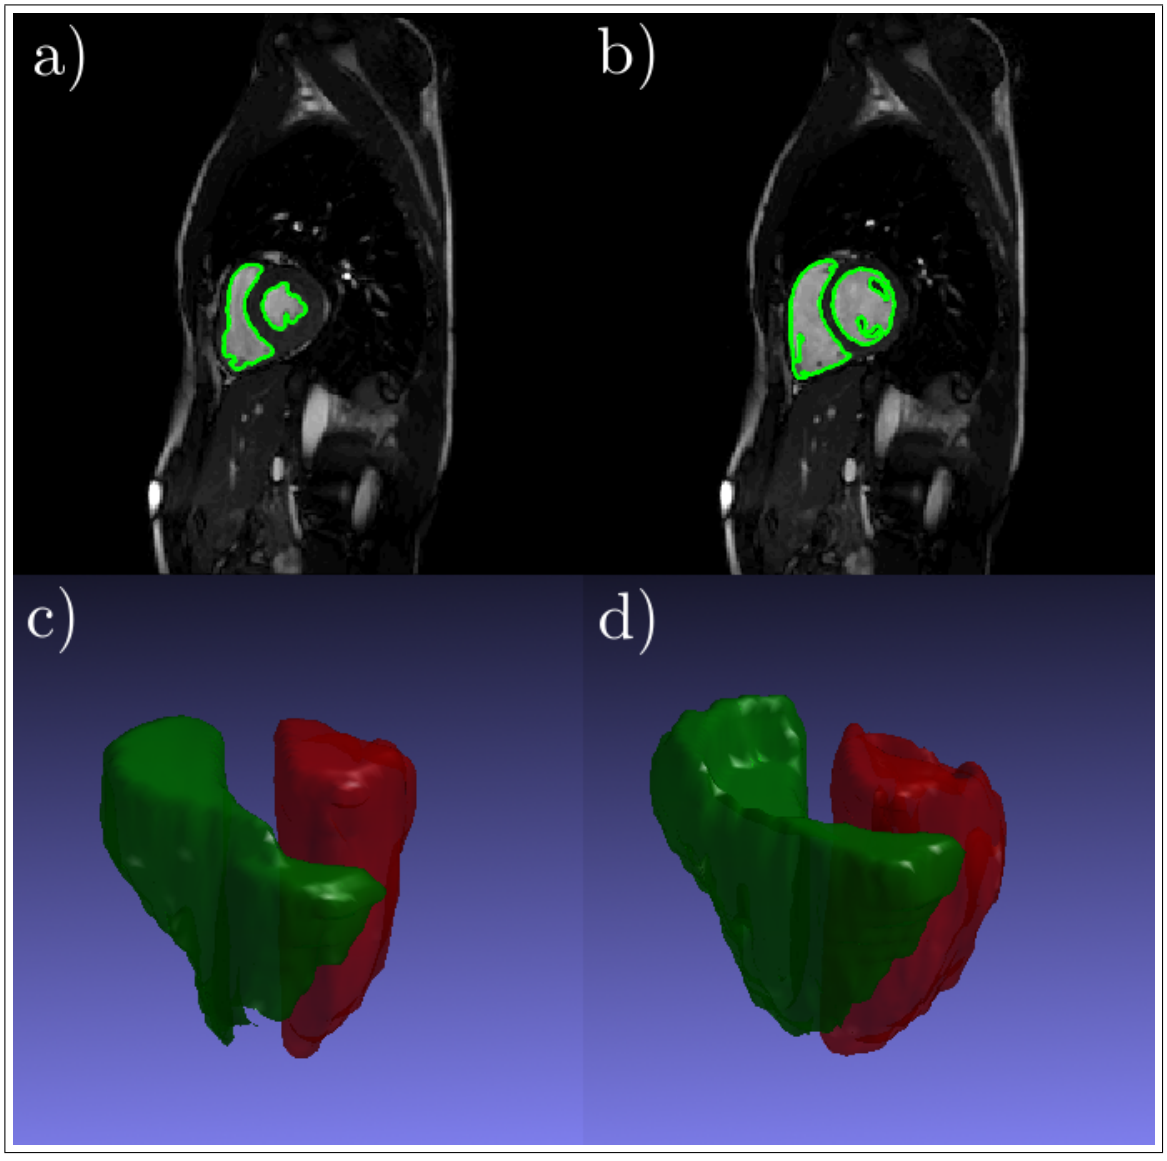

Figure 3.4. Segmentation results using Chan-Vese with preserved topology. (a) and (b) show the same slice for end of systole and for end of diastole, respectively. (c) and (d) show the segmented volumes for systole and diastole, respectively.

outliers, all measurements were within the confidence range (± 1.96 times the standard deviation).

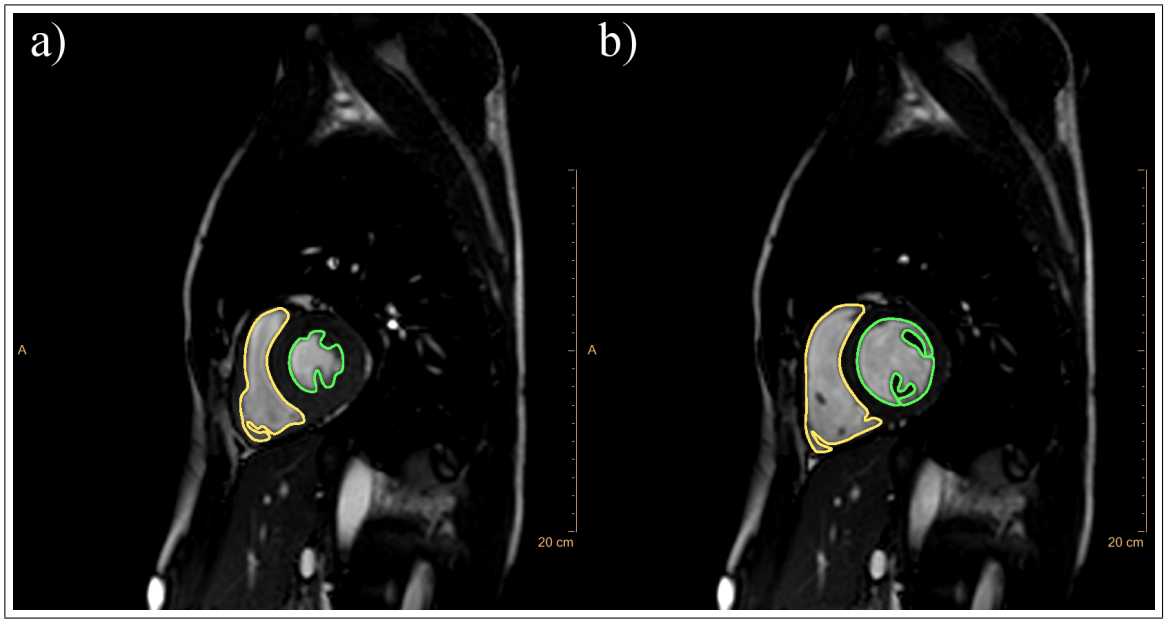

Figure 3.5. Gold standard segmentation using ViewForum. (a) and (b) show the same slice for end of systole and for end of diastole, respectively

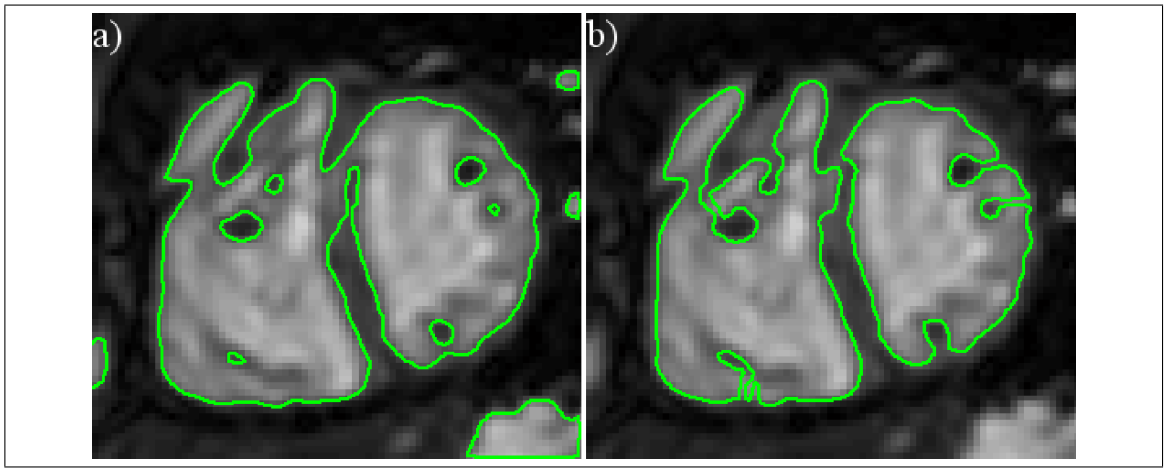

Figure 3.6. Segmentation results in the presence of tetralogy of Fallot (same case of Fig. 3.1 (a)). (a) Segmentation result applying Chan-Vese. (b) Segmentation result applying topology-preserving Chan-Vese. Topology preserving Chan-Vese forces the ventricles to be segmented as disjoint structures.

Fig. 3.9 shows the Dice index for each patient, for LV (a) and RV (b), and for end systole and end diastole frames. Table 3.3 shows the average Dice. It is accepted that a

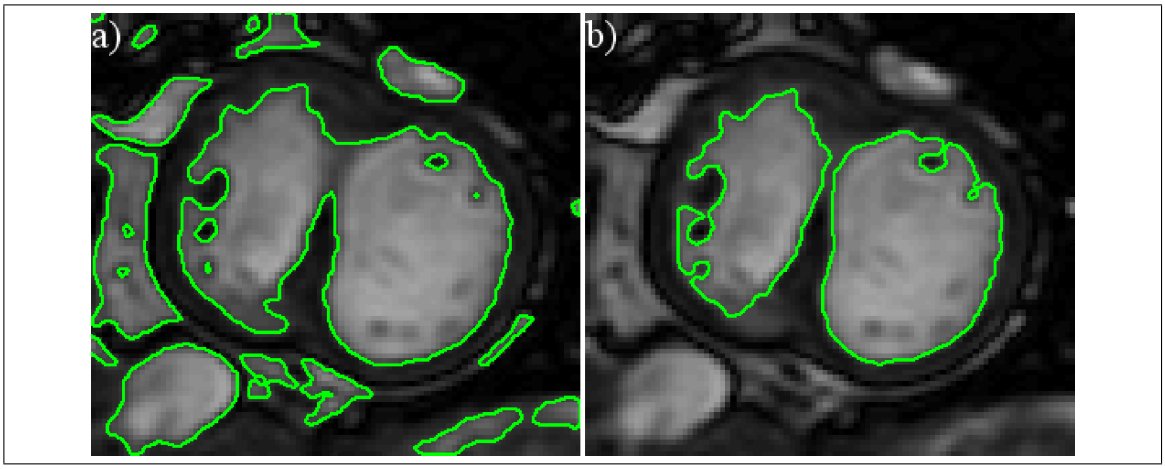

Figure 3.7. Segmentation results in the presence of interventricular connection (same case of Fig. 3.1 (b)). (a) Segmentation result applying Chan-Vese. (b) Segmentation result applying topology-preserving Chan-Vese. Topology preserving Chan-Vese forces that the ventricles to be segmented as disjoint structures.

Volume	p-value of t-test (paired two tail)	Pearson correlation
LV EDV	0.951	0.994
LV ESV	0.212	0.988
RV EDV	0.053	0.995
RV ESV	0.245	0.995

Table 3.2. P-value and Pearson correlation for each ventricle and cardiac phase considering the Viewforum gold standard.

Dice ≥ 0.70 is a reasonable result (Pluempitiwiriyawej et al., 2005). In our case, all Dice indexes were over 0.80 and the average Dice was always greater than or equal to 0.90. It also shows more accurate results for the diastolic phase.

Dice	ESV	EDV
LV	0.92 ± 0.021	0.94 ± 0.018
RV	0.90 ± 0.027	0.92 ± 0.014

Table 3.3. Average Dice for 15 patients, comparing our algorithm with software Segment.

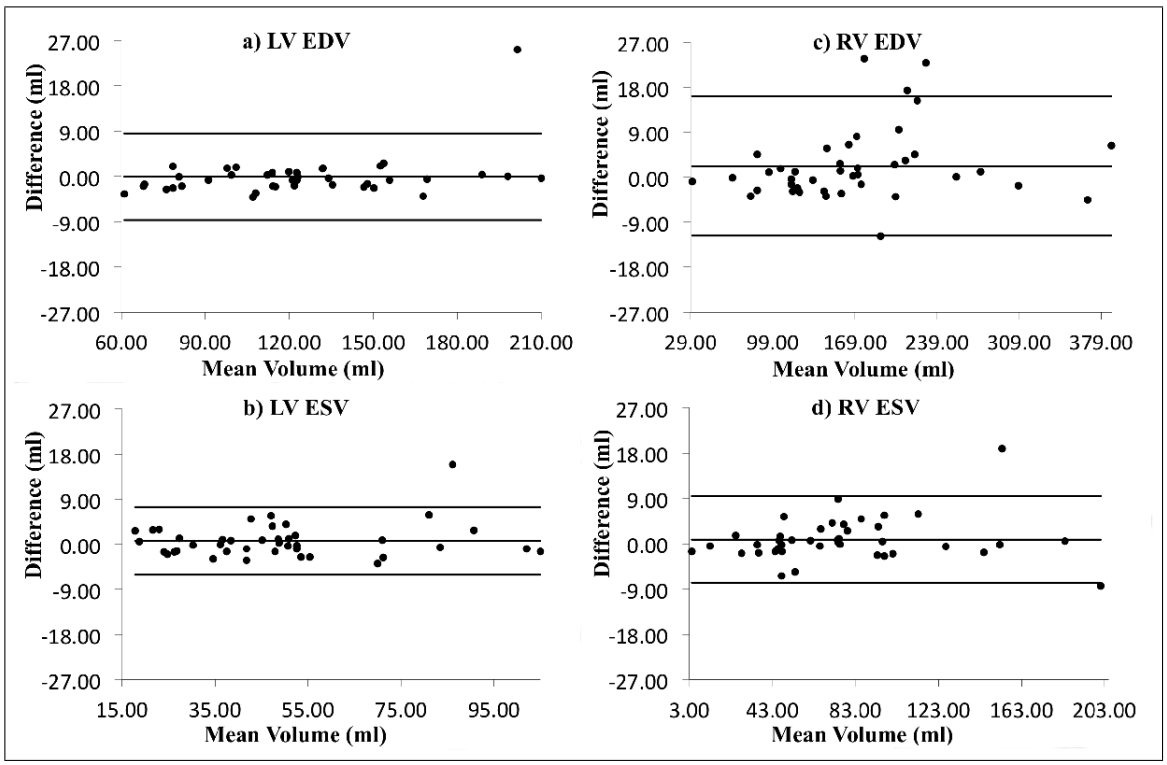

Figure 3.8. Bland-Altman analysis. (a) Left ventricle end diastolic volume; (b) Left ventricle end systolic volume; (c) Right ventricle end diastolic volume; (d) Right ventricle end systolic volume.

3.4. Discussion

The proposed topology preserving level set algorithm enables us to segment simultaneously left and right ventricles. Our method has a reasonable degree of automation without any need for training data sets. Results show that it can achieve accurate segmentations of the left and right ventricles, even in the presence of different CHDs. This represents an important challenge in terms of large variations on the size, shape and location of the ventricular structures. Even though it is still a matter of contention whether trabeculations and papillary muscles should be considered as part of the ventricular cavity, in our approach we excluded those structures following the approach suggested by Jacquier et al. (2010).

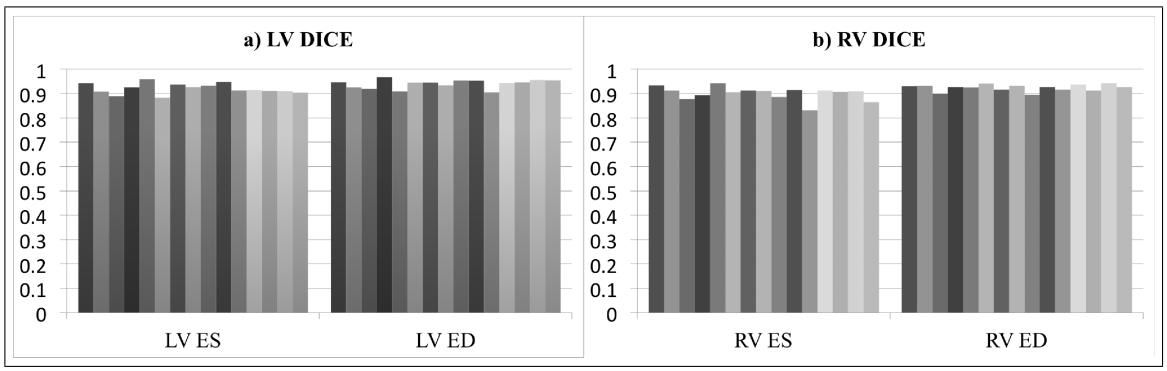

Figure 3.9. Dice for 15 patients, comparing our algorithm with software Segment. (a) Left ventricle for the end of systole (the group on the left) and end of diastole (the group on the right), and (b) right ventricle results for end of systole (the group on the left) and end of diastole (the group on the right).

Simultaneous ventricle segmentation is done fast (about 15 seconds per cardiac phase). Importantly, the computational cost of segmenting additional structures with the same intensity and within the same slices is minimal; thus, exactly the same algorithm could also be used to segment more objects without increasing significantly the processing time.

Manual corrections are sometimes needed at regions with low contrast. This typically happens at the basal slices, in which valves are not clearly shown.

Two-tailed paired t-test is used to evaluate if there are significant differences between our method and the gold standard. The null hypothesis states that there are no differences between those measurements. Since all p-values are greater than 0.05, the null hypothesis cannot be rejected and thus we can conclude that the results obtained with both methods are not significantly different with $P<0.05$. Additionally, Pearsonn tests showed a high correlation between our proposed method and the gold standard.

Bland Altman analysis showed a high level of accuracy between our measures and the chosen gold standard, being almost all data sets inside the acceptable range of 1.96 times the standard deviation. Bland-Altman also showed a minimal bias on the results. Only the RV EDV showed a small bias of about 2.22 ml with a slightly larger standard deviation. That makes explicit the large variations that the RV shows in terms of the ventricular shape (Sardanelli et al., 2008), and the number, size and position of trabeculations, especially in patients with CHDs. Despite that, the bias is small and in clinical practice 2 ml may not be a critical volumetric variability.

Dice indexes showed an accuracy greater or equal to 0.90. As expected, the left ventricle showed a greater Dice than the right one, especially in the diastolic phase, which confirms Bland Altman results. Indeed, trabeculation exclusions are one of the main variation sources, due to the poor contrast shown at some boundaries.

Our method worked well for all but basal slices. The same behavior was observed in the results obtained using the gold standard and can be explained by the poor definition of the valves. In this work we performed corrections manually; otherwise our algorithm would be almost entirely automated. This is an important advantage compared with the Segment software, since the latter failed almost everywhere in the RV.

4. LEVEL SET SEGMENTATION WITH SHAPE PRIOR KNOWLEDGE USING INTRINSIC ROTATION, TRANSLATION AND SCALING ALIGNMENT.

In this chapter we present an algorithm that combines level sets with strong prior knowledge. This new approach extends the intrinsic alignment method proposed by (Cremers et al., 2006). This last method proposes a prior knowledge regularization term which is invariant under translation and scaling. We added rotation invariance since it is crucial, especially in medical image segmentation, as is demonstrated in this chapter. This chapter is structured as follows: Section 4.1 introduces strong prior knowledge in the context of level sets. Section 4.2 defines the general mathematical framework of level sets with a strong prior knowledge term, with particular attention to Cremers' approach (2006). Section 4.3 describes our contribution, which is to generalize Cremers method, so that it can deal with arbitrary rotations in the training database. Sections 4.4 and 4.5 present the results of some experiments demonstrating the benefits of considering rotations into the intrinsic alignment process.

4.1. Level set segmentation with shape prior knowledge using intrinsic rotation, translation and scaling alignment.

Region-based level set algorithm, such as Chan-Vese, showed interesting results for low SNR images or objects with blurred edges within the image. Nevertheless, for more complex distortions (e.g. highly corrupted images or when images show partially occluded objects), additional information is required to perform a correct segmentation. A common approach to overcome this issue has been the introduction of strong shape prior knowledge

(Pluempitiwiriyawej et al., 2005; Tsai et al., 2003; Cremers et al., 2006). This is normally done by adding an additional regularizing term to evolve the curve, considering now a set of known shapes.

As explained in section 1, strong prior knowledge adds a regularization term that penalizes shapes that differ from the training database. Thus, the regularization term needs to measure the distance between two shapes. For this comparison shapes have to be aligned into a common reference frame, and hence, a registration process is needed. This additional process not only increases the computing time and but also, in some cases, generates variable results depending on how this registration is solved.

In order to avoid a registration process, Cremers et al. (2006) proposed a level set algorithm based on shape prior knowledge using intrinsic alignment, i.e. each training shape is normalized depending only on its own features. Thus, all shapes can be easily aligned into a common coordinate system, and subsequently compared directly in that space. However, Cremers et al. solved this intrinsic alignment only for scaling and translation, but not for rotation.

Rotations are critical in several image applications, especially in medical images, due to inter-subject variability or due to changes in patient positioning with respect to the scanner. In this work, we propose an extended framework for the intrinsic alignment developed by Cremers et al. (2006), which now includes rotation, translation and scaling. We used the first moment for translations, as in Cremers' approach, but we used eigenvalues and eigenvectors of the covariance matrix of each training data, in order to find the scaling

and rotation alignment parameters. We derived a new set of evolution equations to include a new regularization term, which now depends on the derivatives of the matrices of eigenfactors.

In the context of medical imaging there is a critical constrain. In general, the size of training sets is small, since it is not trivial to use clinical data, and the gold standard generation is an extremely tedious and time consuming task (Golland & Fischl, 2003; Mazurowskia et al., 2008; Masutani, Nemoto, Nomura, & Hayashi, 2012; Tajbakhsh et al., 2016). These issues turn the use of several state-of-art machine learning algorithms into a really challenging problem as they require the development of complex strategies to validate the results (Mazurowskia et al., 2008; Masutani et al., 2012). The scenario is even worst in the presence of diseases or pathologies, since it is not clear how to generate a representative training set. For these reasons, we assumed a limited number of training samples and thus, our method must deal with the problem of having small training data set. As a final requirement for this approach, is to deal with several medical imaging applications. This means to work with 2D and 3D images.

In this chapter, we propose a new shape regularization term for level set segmentation with strong shape prior knowledge. This is based on an intrinsic alignment approach, which is invariant to translation, scaling and rotation. Our regularizer can be used with any other level set energy term, such us geometric, geodesic or Chan-Vese energies.

We evaluated our algorithm using synthetic and real images in two and three dimensions. We show how our method, combined with the Chan-Vese functional, improves the

results and produces accurate segmentation of 2D and 3D shapes in comparison with (1) Chan-Vese approach (with no prior) and (2) the original approach proposed by Cremers et al. (combined with Chan-Vese functional).

4.2. Shape prior knowledge

Strong prior knowledge learns the shape of an object from a training set and includes this information into the evolution equation through a regularization term. For that purpose, the functional now considers two energy terms: one that depends on the image itself, and a regularizer that depends on the prior shape information. The new minimization problem becomes,

$$\inf_{\phi} \left(E_{image}(\phi) + \beta E_{shape}(\phi, \hat{\phi}) \right) \tag{4.1}$$

where E_{image} is an energy functional that only considers image features (such as the Chan-Vese functional), E_{shape} favors the segmentation of shapes (defined by the signed distance function ϕ) similar to the prior $\hat{\phi}$ and β is a weighting parameter. In order to compare two different shapes, this new regularizer needs to measure the similarity between ϕ and $\hat{\phi}$, and thus it depends on a similarity metric. A discussion on similarity metrics can be found in (Cremers et al., 2006). This regularizing term also needs to be invariant to pose changes of the shapes (such as translation, scaling and rotation). In other words, shapes must be aligned into a common and corresponding coordinate system before evaluating the similarity metric.

A common approach for achieving shape alignment is to solve a registration subproblem. This implies that the new regularization term considers a coordinate transformation, which typically includes translation, scaling and rotation as unknown parameters. Hence, the regularizer is optimized alternating between the segmentation based on image features and the registration process for updating the transformation parameters (Y. Chen et al., 2002; Leventon et al., 2000; Rousson & Paragios, 2002, 2008; Tsai et al., 2003). The registration process not only increases the computation time, but also produces variable segmentation results depending on the order in which the pose parameters are optimized, due to the existence of several local minima for the registration problem.

To avoid this registration process, Cremers et al. (2006) proposed an intrinsic alignment procedure. This intrinsic alignment normalizes each training level set function and put all of them into a common space. Importantly, this normalization is computed using only intrinsic parameters, i.e., parameters obtained from each training data without taking into account other shapes, or any statistical information such as the mean or the standard deviation of the whole database. Considering the signed distance function $\phi(\vec{x})$ with $\vec{x} = (x_1, x_2, ..., x_i, ..., x_n)$, its first and second moments can be expressed as

$$\mu_{x_i}(\phi) = \int_\Omega x_i \, h(\phi(\vec{x})) \, \mathrm{d}\vec{x} \tag{4.2}$$

$$\sigma^2_{x_i}(\phi) = \int_\Omega (x_i - \mu_{x_i}(\phi(\vec{x})))^2 \, h(\phi(\vec{x})) \, \mathrm{d}\vec{x} \tag{4.3}$$

with

$$h(\phi(\vec{x})) = \frac{H(\phi(\vec{x}))}{\int_\Omega H(\phi(\vec{x})) \, \mathrm{d}\vec{x}} \tag{4.4}$$

where $H(\phi(\vec{x}))$ is the Heaviside function applied to ϕ. Considering the normalized level set function $\hat{\phi}$ (i.e., it is already intrinsically aligned), the regularizer $d^2(\phi, \hat{\phi})$ is defined in the normalized (intrinsic) space as

$$d^2(\phi, \hat{\phi}) = \int_\Omega (H(\phi(\sigma_{\vec{x}}(\phi)\vec{x} + M(\phi))) - H(\hat{\phi}(\vec{x})))^2 \mathrm{d}\vec{x}, \tag{4.5}$$

being $\sigma_{\vec{x}}(\phi) = \mathrm{diag}(\sigma_{x_1}(\phi), \sigma_{x_2}(\phi), ..., \sigma_{x_n}(\phi))$ and $M = [\mu_{x_1}(\phi), \mu_{x_2}(\phi), ..., \mu_{x_n}(\phi)]^T$.

As demonstrated in (Cremers et al., 2006), this metric is invariant to translation and scaling, and can be effectively used as E_{shape}. This was demonstrated, without loss of generality, considering the 1D case and supposing that the function ϕ is already centered and scaled so that $\mu_x(\phi) = 0$ and $\sigma_x(\phi) = 1$. If we apply an arbitrary translation and scaling to ϕ we will have a new ϕ^* that embeds the same shape scaled by σ^* and shifted by μ^* and,

$$H(\phi^*(x)) = H\left(\phi\left(\frac{x - \mu^*}{\sigma^*}\right)\right). \tag{4.6}$$

Evaluating the intrinsic alignment parameters for ϕ^* we have,

$$\mu_x(\phi^*) = \frac{\int_\Omega x\, H(\phi^*(x))\mathrm{d}x}{\int_\Omega H(\phi^*(x))\,\mathrm{d}x} = \frac{\int_\Omega x\, H\Big(\phi(\frac{x-\mu^*}{\sigma^*})\Big)\mathrm{d}x}{\int_\Omega H\Big(\phi(\frac{x-\mu^*}{\sigma^*})\Big)\,\mathrm{d}x} = \frac{\int_\Omega(\sigma^* u + \mu^*)\, H(\phi(u))\sigma^*\,\mathrm{d}u}{\int_\Omega H(\phi(u))\sigma^*\,\mathrm{d}u}$$

$$= \sigma^* \frac{\int_\Omega u H(\phi(u))\,\mathrm{d}u}{\int_\Omega H(\phi(u))\,\mathrm{d}u} + \mu^* \frac{\int_\Omega H(\phi(u))\,\mathrm{d}u}{\int_\Omega H(\phi(u))\,\mathrm{d}u} = \sigma^*\mu_x(\phi) + \mu^* = \mu^*,$$

and,

$$\sigma_x(\phi^*) = \left(\frac{\int_\Omega (x-\mu_x(\phi^*(x)))^2\, H(\phi^*(x))\,\mathrm{d}x}{\int_\Omega H(\phi^*(x))\,\mathrm{d}x}\right)^{\frac{1}{2}} = \left(\frac{\int_\Omega (x-\mu^*)^2\, H\big(\phi(\frac{x-\mu^*}{\sigma^*})\big)\,\mathrm{d}x}{\int_\Omega H\big(\phi(\frac{x-\mu^*}{\sigma^*})\big)\,\mathrm{d}x}\right)^{\frac{1}{2}}$$

$$= \left(\frac{\int_\Omega (\sigma^* u)^2\, H\big(\phi(u)\big)\,\mathrm{d}x}{\int_\Omega H\big(\phi(u)\big)\,\mathrm{d}u}\right)^{\frac{1}{2}} = \sigma^* \left(\frac{\int_\Omega u^2\, H\big(\phi(u)\big)\,\mathrm{d}u}{\int_\Omega H\big(\phi(u)\big)\,\mathrm{d}u}\right)^{\frac{1}{2}} = \sigma^*\sigma_x(\phi) = \sigma^*.$$

Evaluating the regularizer for $\phi^*(x)$ we get,

$$d^2(\phi^*, \hat{\phi}) = \int_\Omega (H(\phi^*(\sigma_x(\phi^*)x + \mu_x(\phi^*))) - H(\hat{\phi}(x)))^2\mathrm{d}x$$

$$= \int_\Omega (H(\phi^*(\sigma^* x + \mu^*)) - H(\hat{\phi}(x)))^2\mathrm{d}x$$

$$= \int_\Omega \left(H\left(\phi\big(\sigma^*\big(\frac{x-\mu^*}{\sigma^*}\big) + \mu^*\big)\right) - H(\hat{\phi}(x))\right)^2 = \int_\Omega (H(\phi(x)) - H(\hat{\phi}(x)))^2$$

$$= d^2(\phi, \hat{\phi}),$$

which shows the invariance of this regularizer under translation and scaling. The next step is to find the evolution equation that generates this regularizer. In other words, we need to compute the derivative of the regularizer with respect to ϕ, $\frac{\partial}{\partial\phi}d^2(\phi,\hat{\phi})$. For simplicity, we consider the 1D case again. We can write the Gâteaux derivative of E_{shape} in the $\tilde{\phi}$ direction, in the same way as in (2.13),

$$\left.\frac{\partial E_{shape}(\phi,\hat{\phi})}{\partial\phi}\right|_{\tilde{\phi}} = \lim_{\varepsilon\to 0}\frac{1}{\varepsilon}(E_{shape}(\phi+\varepsilon\tilde{\phi},\hat{\phi}) - E_{shape}(\phi,\hat{\phi})) = \left.\frac{\mathrm{d}}{\mathrm{d}\varepsilon}E_{shape}(\phi+\varepsilon\tilde{\phi},\hat{\phi})\right|_{\varepsilon=0}. \tag{4.7}$$

Considering translation and scaling, we write the term $E_{shape}(\phi+\varepsilon\tilde{\phi},\hat{\phi})$ in (4.7) as follows,

$$E_{shape}(\phi+\varepsilon\tilde{\phi},\hat{\phi}) = \int_\Omega (H(\phi+\varepsilon\tilde{\phi})(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi})) - H(\hat{\phi}(x)))^2\mathrm{d}x. \tag{4.8}$$

Hence, in the first place we need to calculate $\mu_x(\phi+\varepsilon\tilde{\phi})$ and $\sigma_x(\phi+\varepsilon\tilde{\phi})$. For these two terms we will have,

$$\mu_x(\phi+\varepsilon\tilde{\phi}) = \int_\Omega x\, h((\phi+\varepsilon\tilde{\phi})(x))\,\mathrm{d}x, \tag{4.9}$$

$$\sigma_x(\phi+\varepsilon\tilde{\phi}) = \left(\int_\Omega (x-\mu_x((\phi+\varepsilon\tilde{\phi})(x)))^2\, h((\phi+\varepsilon\tilde{\phi})(x))\,\mathrm{d}x\right)^{\frac{1}{2}}, \tag{4.10}$$

therefore, we need to calculate the $h((\phi+\varepsilon\tilde{\phi})(x))$ in both cases. The Taylor series expansion of this term can be write as follows,

$$h((\phi+\varepsilon\tilde{\phi})(x)) = \frac{H((\phi+\varepsilon\tilde{\phi})(x))}{\int_\Omega H((\phi+\varepsilon\tilde{\phi})(x))\,\mathrm{d}x} = \frac{H(\phi(x))+\varepsilon\tilde{\phi}(x)H'(\phi(x))}{\int_\Omega H(\phi(x))+\varepsilon\tilde{\phi}(x)H'(\phi(x))\,\mathrm{d}x} + \mathcal{O}(\varepsilon^2)$$
$$= \frac{H(\phi(x))+\varepsilon\tilde{\phi}(x)\delta(\phi(x))}{\int_\Omega H(\phi(x))+\varepsilon\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x} + \mathcal{O}(\varepsilon^2). \tag{4.11}$$

We need to rewrite the term $\frac{1}{\int_\Omega H(\phi(x))+\varepsilon\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}$. For a small ε, we know the following Taylor series expansion,

$$\frac{1}{1+\varepsilon} = 1-\varepsilon+\mathcal{O}(\varepsilon^2),$$

and thus,

$$\frac{1}{a+b\varepsilon} = \frac{1}{a}\left(\frac{1}{1+\frac{b}{a}\varepsilon}\right) = \frac{1}{a}\left(1-\frac{b}{a}\varepsilon+\mathcal{O}(\varepsilon^2)\right). \tag{4.12}$$

Finally, using (4.12) we can write,

$$\frac{1}{\int_\Omega H(\phi(x))+\varepsilon\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x} = \frac{1}{\int_\Omega H(\phi(x))\mathrm{d}x}\left(\frac{1}{1+\frac{\int_\Omega \varepsilon\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x}}\right)$$
$$= \frac{1}{\int_\Omega H(\phi(x))\mathrm{d}x}\left(1-\frac{\varepsilon\int_\Omega \tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x}\right). \tag{4.13}$$

Replacing (4.11) and (4.13) in (4.9), we obtain the final expression for $\mu_x(\phi+\varepsilon\tilde{\phi})$,

$$
\begin{aligned}
\mu_x(\phi+\varepsilon\tilde{\phi}) \approx & \int_\Omega x\, H(\phi(x)) + \varepsilon\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x \left[\frac{1}{\int_\Omega H(\phi(x))\mathrm{d}x}\left(1-\frac{\varepsilon\int_\Omega \tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x}\right)\right] \\
=& \frac{\int_\Omega x\, H(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} + \frac{\varepsilon\int_\Omega x\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} \\
& - \frac{\varepsilon\int_\Omega \tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x}\frac{\int_\Omega x\, H(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} + \mathcal{O}(\varepsilon^2) \\
=& \mu_x(\phi(x)) + \frac{\varepsilon}{\int_\Omega H(\phi(x))\mathrm{d}x}\int_\Omega (x-\mu_x(\phi(x)))\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x + \mathcal{O}(\varepsilon^2).
\end{aligned}
\tag{4.14}
$$

In the same way, we can rewrite $\sigma_x^2(\phi+\varepsilon\tilde{\phi})$ in (4.10) using (4.13) and (4.14). For simplicity we calculate the expression powered by 2 instead of $\sigma_x(\phi+\varepsilon\tilde{\phi})$), resulting,

$$\begin{aligned}
\sigma_x^2(\phi+\varepsilon\tilde{\phi}) &\approx \int_\Omega (x-\mu_x((\phi+\varepsilon\tilde{\phi})(x)))^2\,(H(\phi(x))+\varepsilon\tilde{\phi}(x)\delta(\phi(x)))\,\mathrm{d}x\Bigg[\frac{1}{\int_\Omega H(\phi(x))\mathrm{d}x} \\
&\quad \left(1-\frac{\varepsilon\int_\Omega \tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x}\right)\Bigg] \\
&\approx \int_\Omega \left(x-\mu_x(\phi(x))+\frac{\varepsilon}{\int_\Omega H(\phi(x))\mathrm{d}x}\int_\Omega (x-\mu_x(\phi(x)))\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x\right)^2 \\
&\quad \left[\frac{H(\phi(x))}{\int_\Omega H(\phi(x))\mathrm{d}x}-\frac{\varepsilon H(\phi(x))\int_\Omega \tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\left(\int_\Omega H(\phi(x))\mathrm{d}x\right)^2}+\frac{\varepsilon\tilde{\phi}(x)\delta(\phi(x)))}{\int_\Omega H(\phi(x))\mathrm{d}x}\right]+\mathcal{O}(\varepsilon^2) \\
&\approx \int_\Omega \Bigg((x-\mu_x(\phi(x)))^2\frac{H(\phi(x))}{\int_\Omega H(\phi(x))\mathrm{d}x}+\varepsilon\frac{(x-\mu_x(\phi(x)))^2}{\int_\Omega H(\phi(x))\mathrm{d}x}\tilde{\phi}(x)\delta(\phi(x)) \\
&\quad -\varepsilon\frac{(x-\mu_x(\phi(x)))^2H(\phi(x))}{\int_\Omega H(\phi(x))\mathrm{d}x}\frac{\int_\Omega \tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} \\
&\quad -2\varepsilon\frac{(x-\mu_x(\phi(x)))H(\phi(x))}{(\int_\Omega H(\phi(x))\mathrm{d}x)^2}\int_\Omega (x-\mu_x(\phi(x)))\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x\Bigg)\,\mathrm{d}x \\
&\quad +\mathcal{O}(\varepsilon^2). \qquad (4.15)
\end{aligned}$$

Now, the last term of this expression is zero, since,

$$\begin{aligned}
\frac{\int_\Omega (x-\mu_x(\phi(x)))H(\phi(x))\mathrm{d}x}{(\int_\Omega H(\phi(x))\mathrm{d}x)^2} &= \frac{\int_\Omega xH(\phi(x))\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x}-\mu_x(\phi(x))\frac{\int_\Omega H(\phi(x))\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} \\
&= \mu_x(\phi(x))-\mu_x(\phi(x))=0,
\end{aligned}$$

and thus, (4.15) leads to,

$$\begin{aligned}\sigma_x^2(\phi+\varepsilon\tilde{\phi}) \approx & \frac{\int_\Omega (x-\mu_x(\phi(x)))^2 H(\phi(x))\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} + \varepsilon \frac{\int_\Omega (x-\mu_x(\phi(x)))^2 \tilde{\phi}(x)\delta(\phi(x))\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} \\ & - \varepsilon \frac{\int_\Omega (x-\mu_x(\phi(x)))^2 H(\phi(x))\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} \frac{\int_\Omega \tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x}{\int_\Omega H(\phi(x))\mathrm{d}x} + \mathcal{O}(\varepsilon^2) \\ = & \sigma_x^2(\phi(x)) + \frac{\varepsilon}{\int_\Omega H(\phi(x))\mathrm{d}x} \int_\Omega \big((x-\mu_x(\phi(x)))^2 - \sigma_x^2(\phi(x))\big)\tilde{\phi}(x)\delta(\phi(x))\mathrm{d}x + \mathcal{O}(\varepsilon^2).\end{aligned} \tag{4.16}$$

Using (4.16) we can easily obtain an expression for $\sigma_x(\phi+\varepsilon\tilde{\phi})$, since the Gâteaux derivative of $\sigma_x^2(\phi+\varepsilon\tilde{\phi})$ in the $\tilde{\phi}$ direction can be written as follows,

$$\begin{aligned}\frac{\partial}{\partial\phi}\sigma_x^2(\phi)\bigg|_{\tilde{\phi}} &= \lim_{\varepsilon\to 0} \frac{\sigma_x^2(\phi+\varepsilon\tilde{\phi}) - \sigma_x^2(\phi)}{\varepsilon} \\ &= \frac{1}{\int_\Omega H(\phi(x))\mathrm{d}x} \int_\Omega \big((x-\mu_x(\phi(x)))^2 - \sigma_x^2(\phi(x))\big)\tilde{\phi}(x)\delta(\phi(x))\mathrm{d}x,\end{aligned}$$

and because we also know that,

$$\frac{\partial}{\partial\phi}\sigma_x(\phi)\bigg|_{\tilde{\phi}} = \frac{\frac{\partial}{\partial\phi}\sigma_x^2(\phi)\Big|_{\tilde{\phi}}}{2\sigma_x(\phi)},$$

we can write $\sigma_x(\phi+\varepsilon\tilde{\phi})$ as follows,

$$\sigma_x(\phi+\varepsilon\tilde{\phi}) \approx \sigma_x(\phi(x)) + \frac{\varepsilon}{2\sigma_x(\phi(x))\int_\Omega H(\phi(x))\mathrm{d}x}\int_\Omega \big((x-\mu_x(\phi(x)))^2-\sigma_x^2(\phi(x))\big)\tilde{\phi}(x)\delta(\phi(x))\mathrm{d}x+\mathcal{O}(\varepsilon^2). \tag{4.17}$$

Back to (4.8), we are now able to calculate,

$$\begin{aligned}
\frac{\partial E_{shape}(\phi,\hat{\phi})}{\partial\phi}\bigg|_{\tilde{\phi}} =& \frac{\mathrm{d}}{\mathrm{d}\varepsilon}E_{shape}(\phi+\varepsilon\tilde{\phi},\hat{\phi})\bigg|_{\varepsilon=0}\\
=& \frac{\mathrm{d}}{\mathrm{d}\varepsilon}\int_\Omega (H(\phi+\varepsilon\tilde{\phi})(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi}))-H(\hat{\phi}(x)))^2\mathrm{d}x\bigg|_{\varepsilon=0}\\
=& 2\int_\Omega (H(\phi+\varepsilon\tilde{\phi})(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi}))-H(\hat{\phi}(x)))\\
& \frac{\mathrm{d}}{\mathrm{d}\varepsilon}H\Big(\phi\big(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi})\big)+\varepsilon\tilde{\phi}\big(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi})\big)\Big)\bigg|_{\varepsilon=0}\mathrm{d}x\\
=& 2\int_\Omega (H(\phi+\varepsilon\tilde{\phi})(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi}))-H(\hat{\phi}(x)))\\
& \delta\big(\phi(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi}))\big)\Big(\big(\nabla\phi(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi}))\big)^T\\
& \frac{\mathrm{d}}{\mathrm{d}\varepsilon}\big(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi})\big)+\tilde{\phi}\big(\sigma_x(\phi+\varepsilon\tilde{\phi})x+\mu_x(\phi+\varepsilon\tilde{\phi})\big)\Big)\bigg|_{\varepsilon=0}\mathrm{d}x,
\end{aligned} \tag{4.18}$$

and replacing (4.17) and (4.14) in (4.18) and evaluating for $\varepsilon = 0$, the new Gâteaux derivative can be written as follows,

$$\frac{\partial E_{shape}(\phi,\hat{\phi})}{\partial \phi}\bigg|_{\tilde{\phi}} = 2\int_{\Omega}\big(H(\phi(\sigma_x(\phi(x))x+\mu_x(\phi(x))))-H(\hat{\phi}(x))\big)\delta\big(\phi(\sigma_x(\phi(x))x+\mu_x(\phi(x)))\big)$$
$$\Bigg[\tilde{\phi}\big(\sigma_x(\phi(x))x+\mu_x(\phi(x))\big)+\big(\nabla\phi(\sigma_x(\phi(x))x+\mu_x(\phi(x)))\big)^T$$
$$\Bigg(\frac{x}{2\sigma_x(\phi(x))\int_{\Omega}H(\phi(x))\mathrm{d}x}\int_{\Omega}\big((x-\mu_x(\phi(x)))^2-\sigma_x^2(\phi(x))\big)\tilde{\phi}(x)\delta(\phi(x))\mathrm{d}x$$
$$+\frac{1}{\int_{\Omega}H(\phi(x))\mathrm{d}x}\int_{\Omega}(x-\mu_x(\phi(x)))\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x\Bigg)\Bigg]\mathrm{d}x. \tag{4.19}$$

As we can see in (4.19), we have the direction of the Gâteaux derivative evaluated in the coordinate transformation $\tilde{\phi}\big(\sigma_x(\phi(x))x+\mu_x(\phi(x))\big)$, instead of $\tilde{\phi}\big(x\big)$ which makes difficult to find the total derivative, i.e.,

$$\frac{\partial E_{shape}(\phi,\hat{\phi})}{\partial \phi}\bigg|_{\tilde{\phi}} = \int_{\Omega}\tilde{\phi}(x)\frac{\partial E_{shape}(\phi,\hat{\phi})}{\partial \phi}\mathrm{d}x, \tag{4.20}$$

and thus, we need to perform the change of variables $y=\sigma_x(\phi(x))x+\mu_x(\phi(x))$, turning (4.19) into,

$$\frac{\partial E_{shape}(\phi,\hat{\phi})}{\partial \phi}\bigg|_{\tilde{\phi}} = 2\int_{\Omega}\big(H(\phi(y)) - H(\hat{\phi}(\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}))\big)\delta\big(\phi(y)\big)$$
$$\Bigg[\tilde{\phi}(y) + \big(\nabla\phi(y)\big)^T$$
$$\Bigg(\frac{\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}}{2\sigma_x(\phi(x))\int_{\Omega}H(\phi(x))\mathrm{d}x}\int_{\Omega}\big((x-\mu_x(\phi(x)))^2 - \sigma_x^2(\phi(x))\big)\tilde{\phi}(x)\delta(\phi(x))\mathrm{d}x$$
$$+\frac{1}{\int_{\Omega}H(\phi(x))\mathrm{d}x}\int_{\Omega}(x-\mu_x(\phi(x)))\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x\Bigg)\Bigg]\frac{1}{\sigma_x(\phi(x))}\mathrm{d}y. \tag{4.21}$$

Separating the integrals and making explicit the linear dependency of each of the three term to $\tilde{\phi}(x)$,

$$\frac{\partial E_{shape}(\phi,\hat{\phi})}{\partial \phi}\bigg|_{\tilde{\phi}} = 2\int_{\Omega}\big(H(\phi(y)) - H(\hat{\phi}\bigg(\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}\bigg))\big)\delta\big(\phi(y)\big)\tilde{\phi}(y)\frac{1}{\sigma_x(\phi(x))}\mathrm{d}y$$
$$+2\int_{\Omega}\big(H(\phi(y)) - H(\hat{\phi}\bigg(\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}\bigg))\big)\delta\big(\phi(y)\big)\big(\nabla\phi(y)\big)^T\bigg(\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}\bigg)\frac{1}{\sigma_x(\phi(x))}\mathrm{d}y$$
$$\Bigg(\frac{1}{2\sigma_x(\phi(x))\int_{\Omega}H(\phi(x))\mathrm{d}x}\int_{\Omega}\big((x-\mu_x(\phi(x)))^2 - \sigma_x^2(\phi(x))\big)\tilde{\phi}(x)\delta(\phi(x))\mathrm{d}x\Bigg)$$
$$+2\int_{\Omega}\big(H(\phi(y)) - H(\hat{\phi}\bigg(\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}\bigg))\big)\delta\big(\phi(y)\big)\big(\nabla\phi(y)\big)^T\frac{1}{\sigma_x(\phi(x))}\mathrm{d}y$$
$$\Bigg(\frac{1}{\int_{\Omega}H(\phi(x))\mathrm{d}x}\int_{\Omega}(x-\mu_x(\phi(x)))\tilde{\phi}(x)\delta(\phi(x))\,\mathrm{d}x\Bigg). \tag{4.22}$$

Note that we have assumed, $\sigma_y(\phi) = \sigma_x(\phi)$ and $\mu_y(\phi) = \mu_x(\phi)$, because they only differ in the integration variable. Due to (4.20), we can finally write,

$$\frac{\partial E_{shape}(\phi,\hat{\phi})}{\partial \phi} = \delta\big(\phi(x)\big)\Bigg[\frac{2}{\sigma_x(\phi(x))}\big(H(\phi(x)) - H(\hat{\phi}\bigg(\frac{x-\mu_x(\phi)}{\sigma_x(\phi)}\bigg))\big)$$
$$+ \int_\Omega \big(H(\phi(y)) - H(\hat{\phi}\bigg(\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}\bigg))\delta\big(\phi(y)\big)\big(\nabla\phi(y)\big)^T\bigg(\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}\bigg)\mathrm{d}y$$
$$\bigg(\frac{\big(\frac{x-\mu_x(\phi(x))}{\sigma_x(\phi(x))}\big)^2 - 1}{\int_\Omega H(\phi(x))\mathrm{d}x}\mathrm{d}x\bigg) + 2\int_\Omega \big(H(\phi(y)) - H(\hat{\phi}\bigg(\frac{y-\mu_x(\phi(x))}{\sigma_x(\phi(x))}\bigg))\delta\big(\phi(y)\big)\big(\nabla\phi(y)\big)^T\mathrm{d}y$$
$$\bigg(\frac{\frac{x-\mu_x(\phi(x))}{\sigma_x(\phi(x))}}{\int_\Omega H(\phi(x))\mathrm{d}x}\bigg)\Bigg] \tag{4.23}$$

This is valid for one training shape. In order to consider several training shapes within E_{shape}, Cremers et al. combined this metric with a kernel density estimation and derived the evolution equation to minimize the new $E_{shape}(\phi)$, which, for N different shapes, consists on the term,

$$\frac{\partial E_{shape}(\phi)}{\partial \phi} = \frac{\sum_{i=1}^N \alpha_i \frac{\partial}{\partial\phi} d^2(\phi,\phi_i)}{2\eta^2 \sum_{i=1}^N \alpha_i}, \tag{4.24}$$

with $\alpha_i = \exp\bigg(-\frac{1}{2\eta^2}d^2(\phi,\phi_i)\bigg)$, $\eta^2 = \frac{1}{N}\sum_{i=1}^N \min_{i\neq j} d^2(\phi_i,\phi_j)$ and $\frac{\partial}{\partial\phi}d^2(\phi,\phi_i)$ given in (4.23).

They demonstrated the ability of their algorithm to successfully segment pedestrians from digital images under normal and occlusion conditions, as well as left ventricles from ultrasound images.

Unfortunately, several image segmentation problems involve rotations. In medical image segmentation rotation is critical since there is a high variability in patient positioning with respect to the scanner, and also a high inter-subject anatomic variability. As shown subsections 4.4 and 4.5, not considering rotations within the intrinsic alignment induces severe errors.

4.3. Our work: Intrinsic rotation

The main contribution of this work is the introduction of rotations into the alignment approach proposed by Cremers et al. (2006). The challenge is to align shapes using intrinsic rotation parameters that depend only on the analyzed shape itself. We propose to use eigenvectors and eigenvalues of the shape covariance matrix for intrinsic rotation and scaling, respectively. This can be thought as aligning the shapes with respect to their principal axes. We use the first moment (or the center of mass) to define the intrinsic translation.

For simplicity, we discuss the 2D case, as the method can be extended with minor modifications to 3D. Let $\vec{x} = (x, y)$ be an arbitrary coordinate in $\mathbb{R}^2$. Let $\phi(\vec{x}) : \mathbb{R}^2 \rightarrow \mathbb{R}$ be the signed distance function that embeds a shape of interest. As in 4.2 and 4.25, we define the first moments for x and for y as,

$$\begin{aligned} \mu_x(\phi(\vec{x})) &= \int_\Omega x\, h(\phi(\vec{x}))\, \mathrm{d}\vec{x} \\ \mu_y(\phi(\vec{x})) &= \int_\Omega y\, h(\phi(\vec{x}))\, \mathrm{d}\vec{x} \end{aligned} \tag{4.25}$$

As in 4.3 and 4.26, we define the second moments (or variances) of a shape for x and for y as,

$$\sigma_x^2(\phi(\vec{x})) = \int_\Omega (x - \mu_x(\phi(\vec{x})))^2 \, h(\phi(\vec{x})) \, \mathrm{d}\vec{x}$$
$$\sigma_y^2(\phi(\vec{x})) = \int_\Omega (y - \mu_y(\phi(\vec{x})))^2 \, h(\phi(\vec{x})) \, \mathrm{d}\vec{x}. \tag{4.26}$$

The anti-diagonal term of the covariance matrix can be write as in (4.27).

$$\sigma_{xy}(\phi(\vec{x})) = \int_\Omega (x - \mu_x(\phi(\vec{x})))(y - \mu_y(\phi(\vec{x}))) \, h(\phi(\vec{x})) \, \mathrm{d}\vec{x}, \tag{4.27}$$

with $h(\phi)$ defined by (4.4). Since the covariance matrix $\Sigma(\phi)$ is real valued and symmetric, the eigendecomposition is defined as

$$\Sigma(\phi) = \begin{bmatrix} \sigma_x^2(\phi(\vec{x})) & \sigma_{xy}(\phi(\vec{x})) \\ \sigma_{xy}(\phi(\vec{x})) & \sigma_y^2(\phi(\vec{x})) \end{bmatrix} = Q \Lambda Q^T \tag{4.28}$$

where Q and Λ are the matrices containing the eigenvectors and eigenvalues of Σ, respectively. We redefine Cremers' metric (4.5), including now intrinsic translation, scaling and rotation as

$$d^2(\phi, \hat{\phi}) = \int_\Omega (H(\phi(Q\Lambda^{\frac{1}{2}}\vec{x} + M) - H(\hat{\phi}(\vec{x})))^2 \mathrm{d}\vec{x} \tag{4.29}$$

with M, the intrinsic translation, defined as

$$M(\phi) = \begin{bmatrix} \mu_x(\phi(\vec{x})) \\ \mu_y(\phi(\vec{x})) \end{bmatrix} \tag{4.30}$$

and Q and $\Lambda^{\frac{1}{2}}$ as the intrinsic rotation and intrinsic scaling, respectively.

After replacing E_{shape} by $d^2(\phi, \hat{\phi}))$ in (4.1), the new energy functional becomes

$$\inf_{\phi} \left(E_{image}(\phi) + \beta E_{shape}(\phi, \hat{\phi})\right) = \inf_{\phi} \left(E_{image}(\phi) + \beta d^2(\phi, \hat{\phi})\right) \tag{4.31}$$

with d^2 defined as in (4.29) and $\hat{\phi}$ being a shape prior.

To compute the evolution equation of our functional we differentiated it with respect to ϕ and applied the curve evolution equation defined in (2.22), leading to the following expression:

$$\frac{\partial \phi}{\partial t} = -\frac{\partial E_{image}}{\partial \phi} - \beta \frac{\partial d^2(\phi, \hat{\phi})}{\partial \phi} = 0. \tag{4.32}$$

In our experiments, we used the Chan-Vese functional (Chan & Vese, 2001) for E_{image}. We compute the derivative of $d^2(\phi, \hat{\phi})$ with respect to ϕ using Gâteaux derivatives as in (4.22). The Gâteaux derivative of our shape energy in the direction $\tilde{\phi}$ is,

$$\left.\frac{\partial E_{shape}(\phi, \hat{\phi})}{\partial \phi}\right|_{\tilde{\phi}} = 2 \int_{\Omega} (H(\phi(Q\Lambda^{1/2}\vec{u} + M)) - H(\hat{\phi}(\vec{u})))\delta(\phi(Q\Lambda^{1/2}\vec{u} + M))$$
$$\left[\tilde{\phi}(Q\Lambda^{1/2}\vec{u} + M) + \left(\nabla\phi(Q\Lambda^{1/2}\vec{u} + M)\right)^T \left(\left.\frac{\partial Q}{\partial \phi}\right|_{\tilde{\phi}} \Lambda^{1/2}\vec{u} + \frac{1}{2} Q\Lambda^{-1/2} \left.\frac{\partial \Lambda}{\partial \phi}\right|_{\tilde{\phi}} \vec{u} + \left.\frac{\partial M}{\partial \phi}\right|_{\tilde{\phi}}\right)\right] \mathrm{d}\vec{u} \tag{4.33}$$

being $\left.\frac{\partial Q}{\partial \phi}\right|_{\tilde{\phi}}, \left.\frac{\partial \Lambda}{\partial \phi}\right|_{\tilde{\phi}}$ and $\left.\frac{\partial M}{\partial \phi}\right|_{\tilde{\phi}}$ the Gâteaux derivatives of Q, Λ and M with respect to ϕ in the $\tilde{\phi}$ direction. In order to have the direction $\tilde{\phi}$ depending on the space variable instead of a transformation of it, we need to perform a change of variable. We set $\vec{u} =$

$\Lambda^{-1/2}Q^T(\vec{x}-M)$ in (4.33) and rewriting, we obtain

$$\frac{\partial E_{shape}(\phi)}{\partial \phi}\bigg|_{\tilde{\phi}} = 2\int_{\Omega}[H(\phi(\vec{x})) - H(\hat{\phi}(\Lambda^{-1/2}Q^T(\vec{x}-M)))]\delta(\phi(\vec{x}))\Bigg[\tilde{\phi}(\vec{x}) + \big(\nabla\phi(\vec{x})\big)^T\left(\frac{\partial Q}{\partial \phi}\bigg|_{\tilde{\phi}} Q^T(\vec{x}-M) + \frac{1}{2}Q\Lambda^{-1/2}\frac{\partial \Lambda}{\partial \phi}\bigg|_{\tilde{\phi}}\Lambda^{-1/2}Q^T(\vec{x}-M) + \frac{\partial M}{\partial \phi}\bigg|_{\tilde{\phi}}\right)\Bigg]|\det(\Lambda^{-1/2})|\,\mathrm{d}\vec{x} \tag{4.34}$$

Particular attention needs to be taken to the derivatives of the eigendecomposition Q and Λ, since we only know the explicit dependence between Σ and ϕ. Thus, we could differentiate those matrices using the chain rule dependence of the Jacobian of the eigendecomposition

$$\begin{aligned}\frac{\partial Q}{\partial \phi}\bigg|_{\tilde{\phi}} &= \frac{\partial Q}{\partial \Sigma}\frac{\partial \Sigma}{\partial \phi}\bigg|_{\tilde{\phi}}, \\ \frac{\partial \Lambda}{\partial \phi}\bigg|_{\tilde{\phi}} &= \frac{\partial \Lambda}{\partial \Sigma}\frac{\partial \Sigma}{\partial \phi}\bigg|_{\tilde{\phi}}.\end{aligned} \tag{4.35}$$

Note that $\partial Q/\partial\phi$ and $\partial\Lambda/\partial\phi$ are interpreted element-wise, i.e., each element is the derivative of the corresponding element of Q and Λ with respect to each entry of Σ. To obtain such derivatives we adapted the approach used in (Candes, Sing-Long, & Trzasko, 2013) for singular value decomposition, to the eigendecomposition problem. We combined (4.28) with the fact that Q is orthonormal, obtaining

$$\begin{cases}\Sigma(\phi) = Q\Lambda Q^T \\ \\ \quad I = Q^TQ\end{cases} \tag{4.36}$$

Taking the derivative of both equations with respect of each element s_{ij} of the matrix Σ and multiplying the first equation by $Q^T(\cdot)Q$, we obtained the linear system

$$\begin{cases} Q^T \dfrac{\partial \Sigma}{\partial s_{ij}} Q = Q^T \dfrac{\partial Q}{\partial s_{ij}} \Lambda + \dfrac{\partial \Lambda}{\partial s_{ij}} + \Lambda \dfrac{\partial Q^T}{\partial s_{ij}} Q \\ 0 = \dfrac{\partial Q^T}{\partial s_{ij}} Q + Q^T \dfrac{\partial Q}{\partial s_{ij}} \end{cases}$$

The second equation shows that $Q^T \dfrac{\partial Q}{\partial s_{ij}}$ is antisymmetric, and thus, defining

$$\Gamma_{ij} := Q^T \frac{\partial Q}{\partial s_{ij}} \tag{4.37}$$

the linear system could be written as

$$\begin{cases} Q^T \dfrac{\partial \Sigma}{\partial s_{ij}} Q = \Gamma_{ij} \Lambda + \dfrac{\partial \Lambda}{\partial s_{ij}} + \Lambda \Gamma_{ij}^T \\ \Gamma_{ij} = -\Gamma_{ij}^T \end{cases} \tag{4.38}$$

In 2D, for the derivative of the eigenvector matrix $\frac{\partial Q}{\partial s_{ij}}$ with respect of the elements of Σ, we have,

$$\frac{\partial Q}{\partial s_{ij}} = Q\Gamma_{ij} = \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix} \begin{bmatrix} 0 & \gamma^{ij} \\ -\gamma^{ij} & 0 \end{bmatrix}, \tag{4.39}$$

where

$$\gamma^{ij} := \frac{q_{i1} q_{j2}}{\lambda_2 - \lambda_1} \tag{4.40}$$

being q_{ij} the ij-th element of the Q matrix, λ_1 and λ_2 the eigenvalues and the supra-indexes of γ^{ij} depend on the correspondent element s_{ij}. For the derivative of the eigenvalues with

respect to the elements of Σ, we obtain

$$\frac{\partial \lambda_k^{ij}}{\partial s_{ij}} = q_{ik}q_{jk} \quad \text{for } k = 1, 2 \tag{4.41}$$

and therefore, we can define the matrix $\frac{\partial \Lambda}{\partial s_{ij}}$ as,

$$\frac{\partial \Lambda}{\partial s_{ij}} = L_{ij} = \begin{bmatrix} q_{i1}q_{j1} & 0 \\ 0 & q_{i2}q_{j2} \end{bmatrix}. \tag{4.42}$$

Now, we can write (4.35) as follows,

$$\begin{aligned}
\left.\frac{\partial Q}{\partial \phi}\right|_{\tilde{\phi}} =& \sum_i \sum_j Q\Gamma_{ij} \left.\frac{\partial s_{ij}}{\partial \phi}\right|_{\tilde{\phi}} \\
=& Q\left(\Gamma_{11}\left.\frac{\partial \sigma_x^2(\phi(\vec{x}))}{\partial \phi}\right|_{\tilde{\phi}} + (\Gamma_{12}+\Gamma_{21})\left.\frac{\partial \sigma_{xy}(\phi(\vec{x}))}{\partial \phi}\right|_{\tilde{\phi}} + \Gamma_{22}\left.\frac{\partial \sigma_y^2(\phi(\vec{x}))}{\partial \phi}\right|_{\tilde{\phi}}\right), \\
\left.\frac{\partial \Lambda}{\partial \phi}\right|_{\tilde{\phi}} =& \sum_i \sum_j L_{ij} \left.\frac{\partial s_{ij}}{\partial \phi}\right|_{\tilde{\phi}} \\
=& L_{11}\left.\frac{\partial \sigma_x^2(\phi(\vec{x}))}{\partial \phi}\right|_{\tilde{\phi}} + (L_{12}+L_{21})\left.\frac{\partial \sigma_{xy}(\phi(\vec{x}))}{\partial \phi}\right|_{\tilde{\phi}} + L_{22}\left.\frac{\partial \sigma_y^2(\phi(\vec{x}))}{\partial \phi}\right|_{\tilde{\phi}}.
\end{aligned} \tag{4.43}$$

The terms $\left.\frac{\partial \sigma_x^2(\phi)}{\partial \phi}\right|_{\tilde{\phi}}$ and $\left.\frac{\partial \sigma_y^2(\phi)}{\partial \phi}\right|_{\tilde{\phi}}$ were calculated in (4.17). In the same way, we can calculate $\left.\frac{\partial \sigma_{xy}(\phi)}{\partial \phi}\right|_{\tilde{\phi}}$ as follows,

$$\sigma_{xy}(\phi+\varepsilon\tilde{\phi}) \approx \sigma_{xy}(\phi) + \frac{\varepsilon}{\int_\Omega H(\phi(\vec{x}))\mathrm{d}\vec{x}} \int_\Omega \big((x-\mu_x(\phi(\vec{x})))(y-\mu_y(\phi(\vec{x}))) - \sigma_{xy}^2(\phi(\vec{x}))\big)$$

$$\tilde{\phi}(\vec{x})\delta(\phi(\vec{x}))\mathrm{d}\vec{x} + \mathcal{O}(\varepsilon^2), \tag{4.44}$$

and

$$\left.\frac{\partial\sigma_{xy}(\phi)}{\partial\phi}\right|_{\tilde{\phi}} = \frac{1}{\int_\Omega H(\phi(\vec{x}))\mathrm{d}\vec{x}} \int_\Omega \big((x-\mu_x(\phi(\vec{x})))(y-\mu_y(\phi(\vec{x}))) - \sigma_{xy}^2(\phi(\vec{x}))\big)\tilde{\phi}(\vec{x})\delta(\phi(\vec{x}))\mathrm{d}\vec{x} \tag{4.45}$$

For a detailed explanation of those derivatives calculation, including the linear system solution, see Appendix B. Replacing those derivatives into (4.34), allows to write the final evolution equation as

$$\frac{\partial E_{shape}(\phi)}{\partial\phi} = 2\delta(\phi)D(\phi,\hat{\phi})|\det(\Lambda^{-\frac{1}{2}})| + \frac{\delta(\phi)}{\int_\Omega H\phi\mathrm{d}\vec{x}}\Bigg(S_{xx}R_1 + S_{xy}R_2 + S_{yy}R_3$$

$$+ S_{xx}S_1 + S_{xy}S_2 + S_{yy}S_3 + T_1(\vec{x}-M)\Bigg) \tag{4.46}$$

where

$$
\begin{aligned}
D(\phi,\hat{\phi}) =& H(\phi(\vec{x}) - H(\hat{\phi}(\Lambda^{-\frac{1}{2}}Q^T(\vec{x}-M))) \\
S_{xx} =& (x-\mu_x)^2 - \sigma_x^2 \\
S_{yy} =& (y-\mu_y)^2 - \sigma_y^2 \\
S_{xy} =& (x-\mu_x)(y-\mu_y) - \sigma_{xy} \\
R_1 =& \int_\Omega 2A^TQ\Gamma_{11}B\mathrm{d}\vec{x} \\
R_2 =& \int_\Omega 2A^TQ(\Gamma_{12}+\Gamma_{21})B\mathrm{d}\vec{x} \\
R_3 =& \int_\Omega 2A^TQ\Gamma_{22}B\mathrm{d}\vec{x} \\
S_1 =& \int_\Omega A^TQ\Lambda^{-\frac{1}{2}}L_{11}\Lambda^{-\frac{1}{2}}B\mathrm{d}\vec{x} \\
S_2 =& \int_\Omega A^TQ\Lambda^{-\frac{1}{2}}(L_{12}+L_{21})\Lambda^{-\frac{1}{2}}B\mathrm{d}\vec{x} \\
S_3 =& \int_\Omega A^TQ\Lambda^{-\frac{1}{2}}L_{22}\Lambda^{-\frac{1}{2}}B\mathrm{d}\vec{x}
\end{aligned}
$$

$$
\begin{aligned}
T_1 =& \int_\Omega 2D(\phi,\hat{\phi})\delta(\phi)(\nabla\phi)^T|\det(\Lambda^{-\frac{1}{2}})|\mathrm{d}\vec{x} \\
A =& D(\phi,\hat{\phi})\delta(\phi)\nabla\phi \\
B =& Q^T(\vec{x}-M)|\det(\Lambda^{-\frac{1}{2}})|
\end{aligned}
$$

The terms R_1, R_2, and R_3 are related to the rotation matrix and are multiplied by the terms that minimize the variances and covariance of the shapes considering the three axes

of symmetry, S_{xx}, S_{yy} and S_{xy}, and they also consider the difference of the scaling factors through the matrices Γ_{ij}. The terms S_1, S_2, and S_3 are related to the scaling factors and they consider the same three axes. Interestingly they do not depend on the scaling factors themselves, but only on the rotation matrix. Since the σ_x, σ_{xy}, and σ_y variables account for scaling, matrices S_1, S_2 and S_3 rotate the axes in order to generate the common space. Finally, the term T_1 comes from the translation. Note that the term $|\det(\Lambda^{-\frac{1}{2}})|$ is to account for the change of variable that was needed to compute the total differentiation using Gâteaux derivative.

Our approach can be easily extended to a higher dimensional space. For example, in 3D the eigenvalue decomposition matrices (Q and Λ) are 3×3, so the indexes in (4.39), (4.40), (4.41), and (4.42) become $i = 1...3, j = 1...3, k = 1...3$. Additionally, a new gradient operator needs to be defined and its proper discretization.

This approach is only valid for comparing two signed distance functions. For considering two or more training shapes, we used, as in (Cremers et al., 2006), a kernel density estimation, which does not assume any particular distribution of the data and compares every shape of the training set with the current ϕ, weighting them depending on the similarity of the shapes after the intrinsic alignment process. The final expression is defined in (4.24) substituting $\frac{\partial}{\partial \phi} d^2(\phi(\vec{x}))$ by the new regularization derivative in 4.46.

4.3.1. Implementation details

We implemented our algorithm in MATLAB. For the level set framework we used the differentiable approximation of the Heaviside function,

$$H(\phi,\epsilon)=\begin{cases} 1 \text{ if } \phi>\epsilon \\ 0 \text{ if } \phi<-\epsilon \\ \frac{1}{2}\left[1+\frac{\phi}{\epsilon}+\frac{1}{\pi}\sin\left(\frac{\pi\phi}{\epsilon}\right)\right] \text{ if } |\phi|\leq\epsilon, \end{cases} \tag{4.47}$$

and the discrete Dirac function corresponding to the derivative of the (4.47),

$$\delta(\phi,\epsilon)=\begin{cases} \frac{1}{2\epsilon}\left[1+\cos\left(\frac{\pi\phi}{\epsilon}\right)\right] \text{ if } |\phi|\leq\epsilon \\ 0 \text{ otherwise} \end{cases}$$

We used $\epsilon=5$. This Dirac Delta approximation has values only in the interval $[-\epsilon,\epsilon]$, in contrast with other approximations which have values different from zero in the whole domain. Therefore, this Dirac Delta approximation allows us to apply the shape energy locally, i.e., only in a neighborhood to the curve. Even though this option restricts the curve evolution, the results obtained are more stable for the shape regularization, compared with those obtained with non-compact support approximations.

After we calculated the coordinate transformation from the normalized space to the space of ϕ, we obtained the aligned shape prior using a bilinear or trilinear interpolation for 2D or 3D, respectively.

For training purposes, we estimated ϕ from silhouette images assuring that they correspond to signed distance functions. We then stored a cropped version of those functions (cropping at the -10 level curve).

4.4. Experiments and results in 2D

We show five 2D experiments to test the performance of our algorithm under different conditions. In our first 2D experiment, we created a synthetic image (128×128 pixels) with intensity values in $[0, 1]$, and with an 80×30-pixel rectangle shape occluded by an ellipse of semi-axes $(25, 15)$ pixels. The figure was rotated 50° with respect to the center of the image. We blurred (Gaussian filter of 10×10 pixels and standard deviation 2 pixels) and then we added noise (with mean 0 and standard deviation 0.1) to the image. The training data consisted of five rectangles of different sizes and positions.

In our second 2D experiment, we used a horse silhouette with the same noise, blurring and rotation as in the previous experiment, but now the head of the horse was occluded by an ellipse. Shape priors were a set of five non-rotated horse silhouettes (Fig. 4.1). All horse silhouettes were taken from the Weizzman Horse Database (Borenstein, Sharon, & Ullman, 2004). The initialization was another horse silhouette which is not part of the training shapes, manually placed close to the desired figure.

In our third 2D experiment, we used a real picture of a horse (Borenstein et al., 2004) with a complex background, rotated 30°. We trained our algorithm with the same set of non-rotated horse silhouettes as in the previous experiment.

In our fourth 2D experiment, we used a 2D cardiac Magnetic Resonance (MR) image of a right ventricle with an abnormal anatomy. We used five similar 2D cardiac MR images of volunteers (without any pathology) as shape priors. The selected images correspond approximately to the same slice of the heart and the ground truth was generated by an expert.

In our fifth 2D experiment, we segmented the Corpus Callosum (CC) from MR images taken from the public data base OASIS (Marcus et al., 2007). Shape priors and ground truth are also available online (Ardekani, Bachman, Figarsky, & Sidtis, 2014). We randomly chose 100 training shapes of healthy subjects (Fig. 4.3) and we tested our method over 60 randomly chosen images of subjects with Alzheimer's disease. We assessed the performance of our regularizer measuring the average Dice index, and false positives and false negatives expressed as percentage of the ground truth area.

In all experiments we compared our approach, which considers the Chan-Vese energy functional combined with intrinsic translation, scaling and rotation, against the Chan-Vese energy functional (without any prior), and also against the same functional combined with the Cremers' regularizer, which considers only intrinsic translation and scaling.

Results with synthetic images of a rectangle (Fig. 4.4) show the effectiveness of including the intrinsic rotation alignment under the presence of noise and blurring. As expected, our regularizer fitted the rectangular shape better than the alternative methods. Fig. 4.5 (a) shows the evolution of the Chan-Vese energy, which has some ripples and oscillations produced by the re-initialization of the level set function every 11 iterations. Fig.

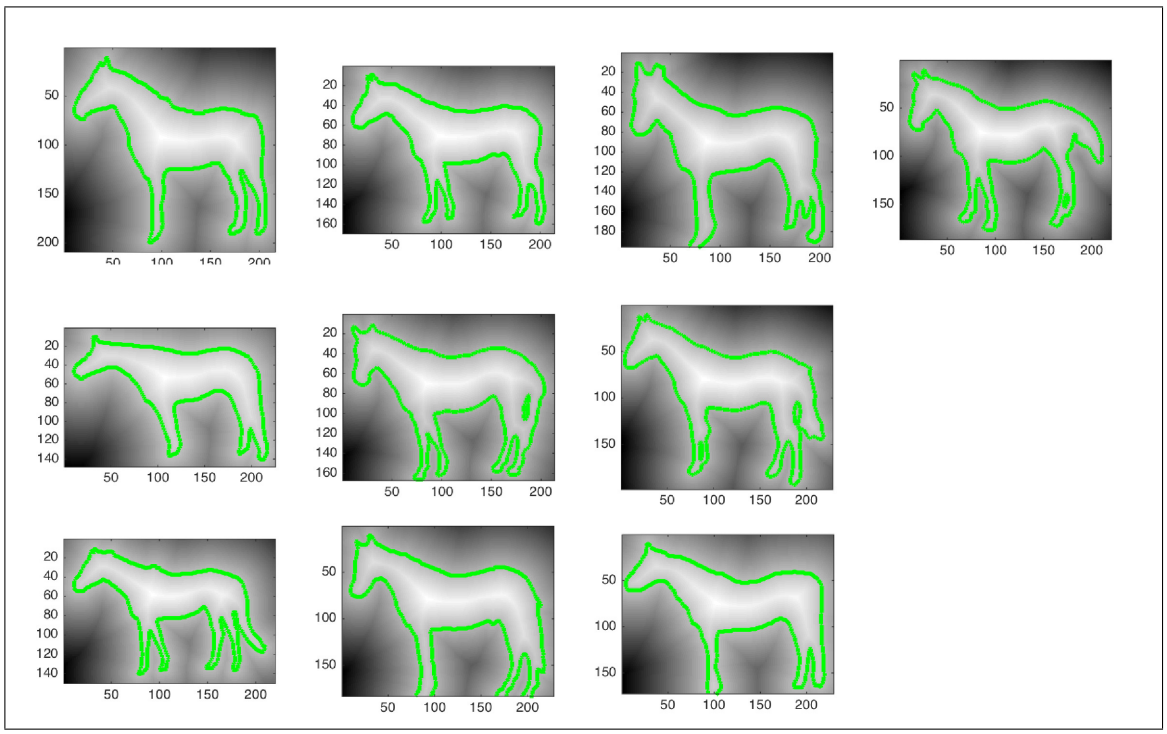

Figure 4.1. Horse shape priors. Figures show the level set function and the green curves correspond to the zero level curves.

4.5 (b) shows how the shape regularization term decreases, starting from a high value, due to the elliptic initialization, and then imposing the rectangular shape, while Chan-Vese energy starts to stabilize. We used the stabilization of both energies as a stopping criteria.

The experiment with a horse silhouette (Fig. 4.6) shows that our algorithm can successfully deal with noise and blurring, considering a more complex shape, even in the presence of occlusion and missing visual data. Our algorithm achieves a good consensus between Chan-Vese energy and the shape energy, completing the unknown information, in this case the head of the horse.

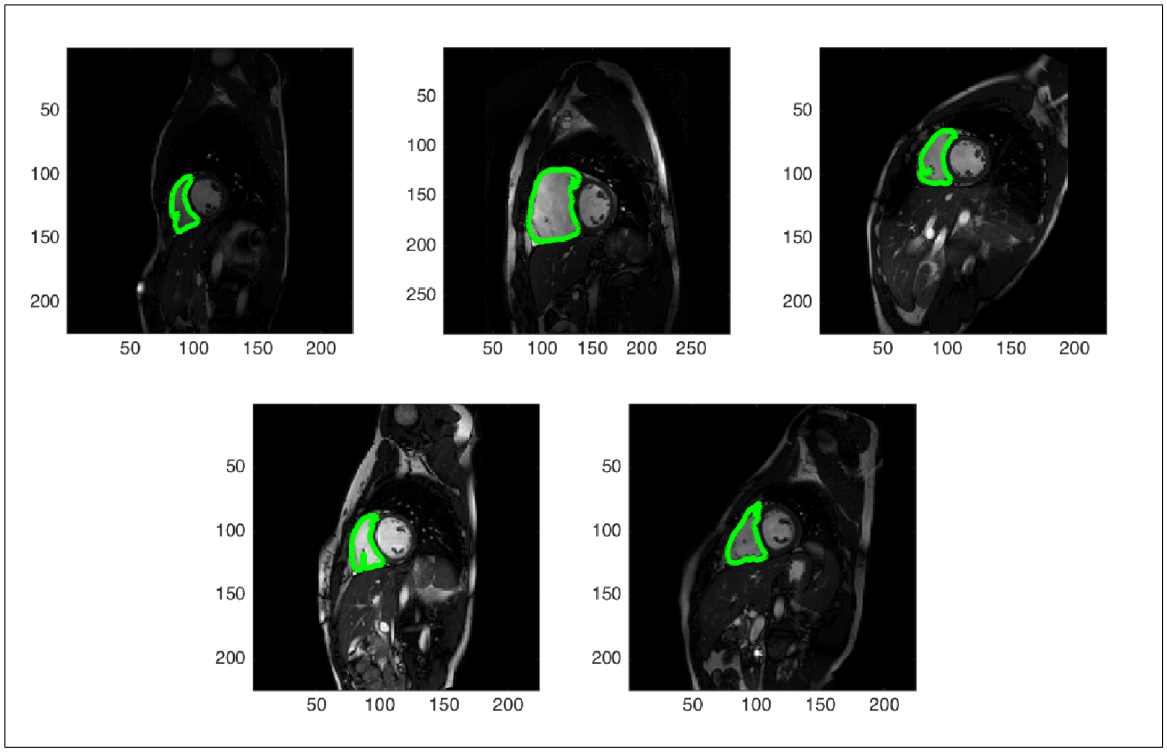

Figure 4.2. Right ventricle shape priors. Figures show MRI images of the heart and the green curves correspond to the zero level curves.

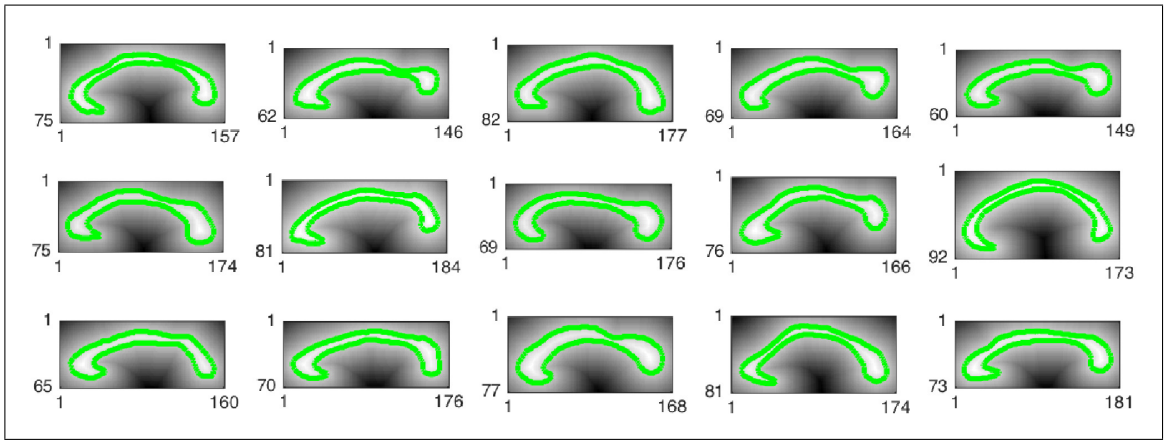

Figure 4.3. 15 corpus callosum shape priors from the 100 shapes of the entire database. Figures show the level set function and the green curves correspond to the zero level curves.

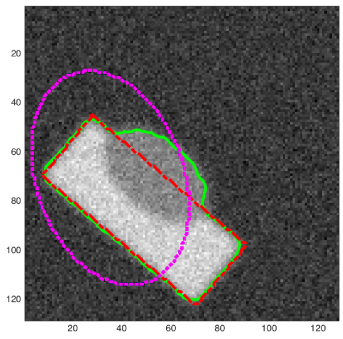

(a) Chan-Vese.

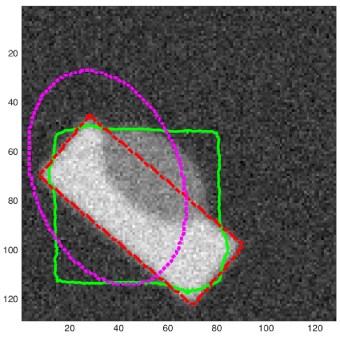

(b) Cremers' metric.

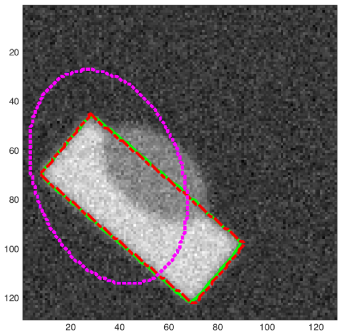

(c) Our metric.

Figure 4.4. Comparing the segmentation of a rectangle using (a) Chan-Vese (CV) without priors, (b) Cremers' approach (CV + rectangular priors with intrinsic scaling and translation) and (c) Our approach (CV + rectangular priors with intrinsic scaling, translation and rotation). Solid green curve shows the segmentation, dotted magenta curve shows the initialization and dashed red curve shows the ground truth.

The experiment with the image of a real horse (Fig. 4.7) shows that our algorithm can deal with background clutter. Additionally, it shows how our regularizer manages the shape variability of the prior set. The head and the body of the horse have low shape variability, hence the force is strong enough to impose the shape regularization over the Chan-Vese energy and a consensus can be easily achieved. On the other hand, the size

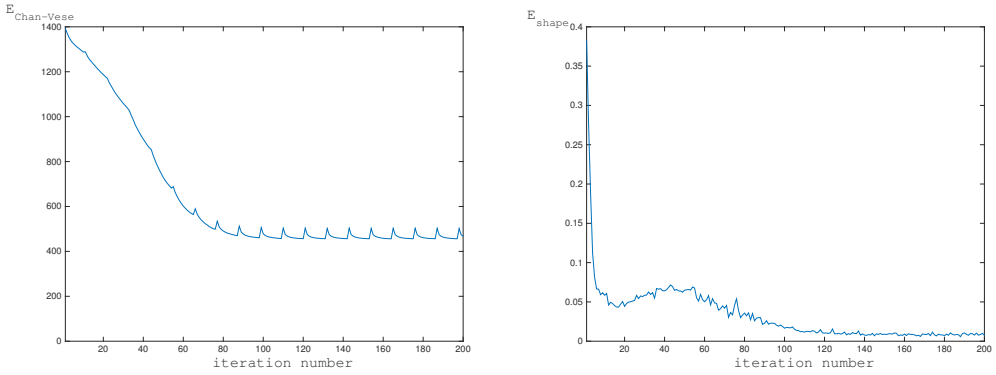

(a) Chan-Vese energy.

(b) Shape energy.

Figure 4.5. Energies of the minimized functional for the segmentation of a rectangle. (a) Chan-Vese energy, (b) Shape energy. The ripples of the Chan-Vese energy are produced by the re-initialization process which was set every 11 iterations. The shape energy is progressively minimized starting from a large value since the initialization and then that value decreases since the shape progressively becomes a rectangle.

and the exact position of the horses' legs have high variability in the shape prior set, hence it is more difficult to get a good balance between the shape regularization term and the Chan-Vese energy. Since the amount of background clutter was larger behind the head of the horse than in with the rest of the image, we tended to increase the weight of the shape regularization to get better shape definition in that region. This produces over-training of the legs. In practice, this means the curve evolution resembles closely the shape of the legs of the training set, and as a consequence, the image intensity in that region is mostly ignored.

Right ventricle experiment (Figs. 4.8, 4.9 and 4.10) showed that our algorithm can be applied in medical images, imposing a specific shape but allowing some variability, as we can see the trabeculation and at the boundary of the ventricle. Additionally, it

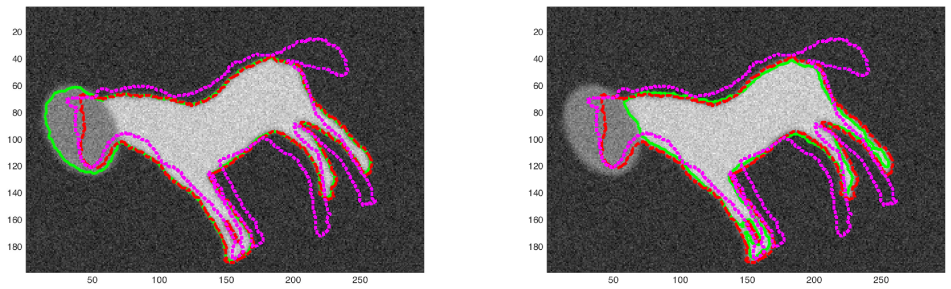

(a) Chan-Vese.

(b) Cremers' metric.

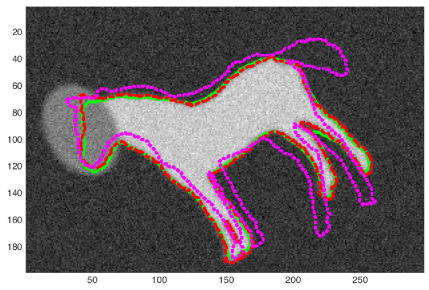

(c) Our metric.

Figure 4.6. Comparing the segmentation of a horse silhouette using (a) Chan-Vese (CV) without priors, (b) Cremers' approach (CV + six horse silhouettes priors with intrinsic scaling and translation) and (c) Our approach (CV + six horse silhouettes priors with intrinsic scaling, translation and rotation). Solid green curve shows the segmentation, dotted magenta curve shows the initialization and dashed red curve shows the ground truth.

can reject the anomaly of the pulmonary artery. This result also shows the importance of including intrinsic rotation, since although the images were acquired in a consistent anatomical orientation (so-called short axis), the result of our method is closer the ground truth than the Cremers' approach.

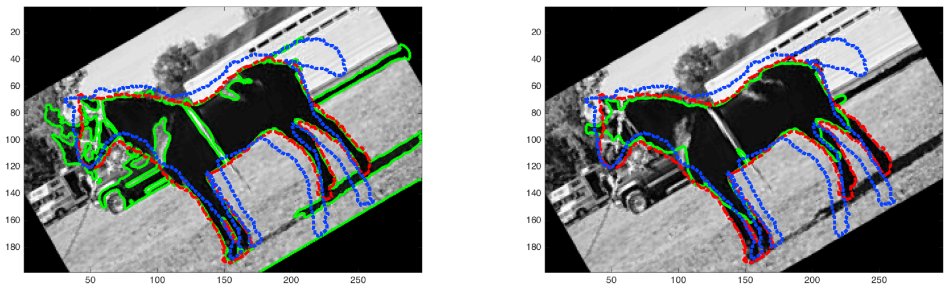

(a) Chan-Vese.

(b) Cremers' metric.

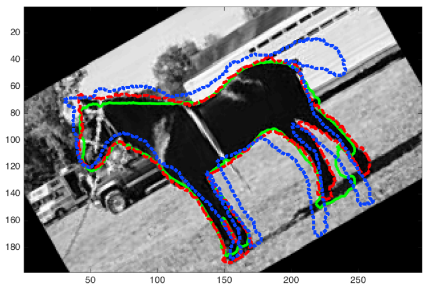

(c) Our metric.

Figure 4.7. Comparing segmentation of a real image of a horse using (a) Chan-Vese (CV) without priors, (b) Cremers' approach (CV + six horse silhouettes priors with intrinsic scaling and translation) and (c) Our approach (CV + six horse silhouettes priors with intrinsic scaling, translation and rotation). Solid green curve shows the segmentation, dotted blue curve shows the initialization and dashed red curve shows the ground truth.

The corpus callosum experiment (Figs. 4.11, 4.12, 4.13) confirms that our algorithm produces accurate results. Importantly, the database was built after registering the images and locating the mid-sagittal slice of the brain. Hence, the images were supposed

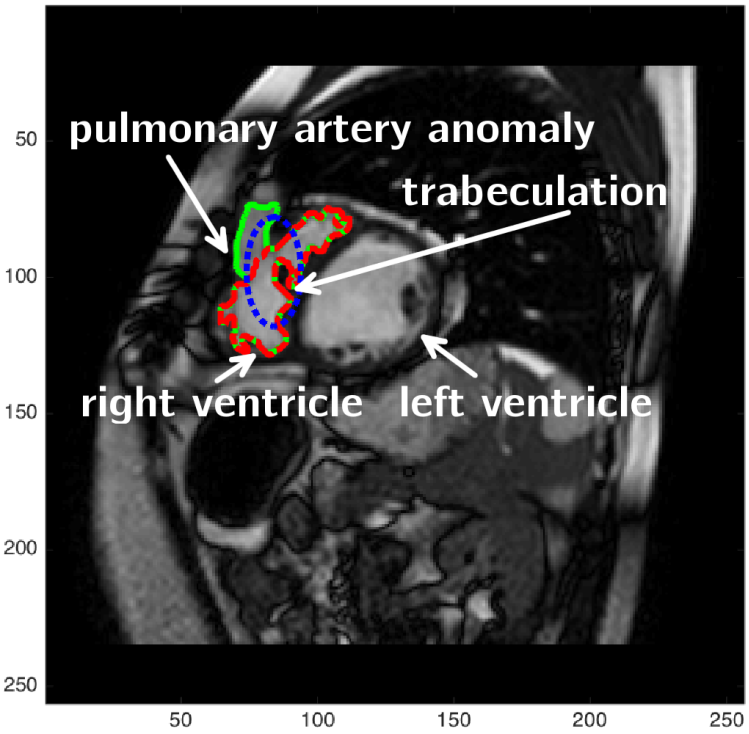

Figure 4.8. Comparing segmentation of a Right Ventricles (RV) using Chan-Vese algorithm without priors. Solid green curve shows the segmentation, dotted blue curve shows the initialization and dashed red curve shows the ground truth.

to be aligned. Despite the existence of that process, differences between corpus callosum's anatomies might include rotation variations and thus, affect the segmentation results. Chan-Vese algorithm without any prior gave an average Dice index of 0.756, which is in principle reasonable, but false positives reported an average of 68% of the ground truth area (first row of Table 4.1). This means that despite the intersection between Chan-Vese segmentation and the ground truth was acceptable, this algorithm segmented a big amount of pixels which are not part of the ground truth. Introducing Cremers' regularized

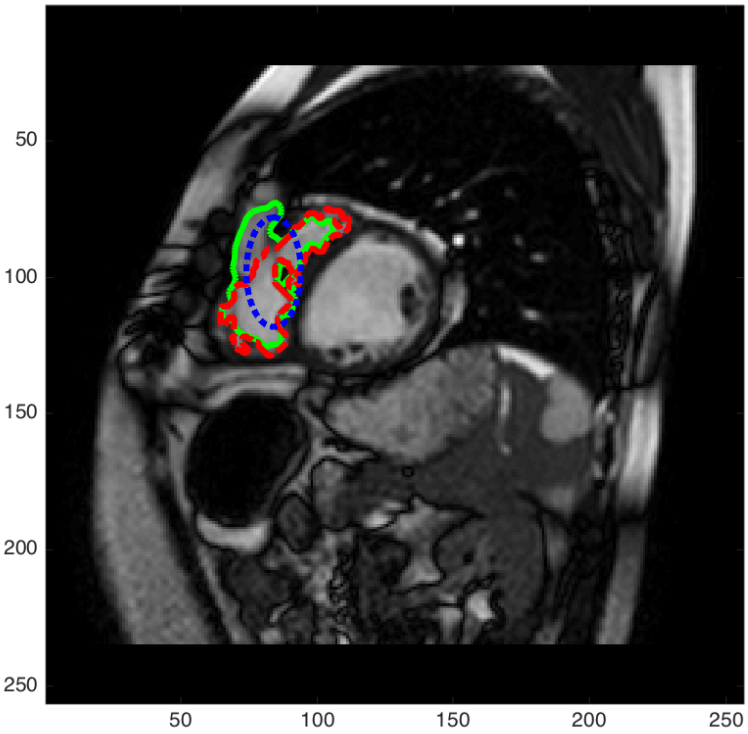

Figure 4.9. Comparing segmentation of a Right Ventricles (RV) using Cremers' approach (Chan-Vese + five RV priors with intrinsic scaling and translation). Solid green curve shows the segmentation, dotted blue curve shows the initialization and dashed red curve shows the ground truth.

incremented the average Dice index to 0.859 and decreased the false positives to 24.0%. Unfortunately, Cremers' method increased the false negatives from 0.5% to 4.7%, (second row of Table 4.1). Finally, our regularizer increased the average Dice index to 0.919 and also decreased the false positives to 14% and false negatives to 3.1%, clearly presenting the best performance of the three methods (third row of Table 4.1).

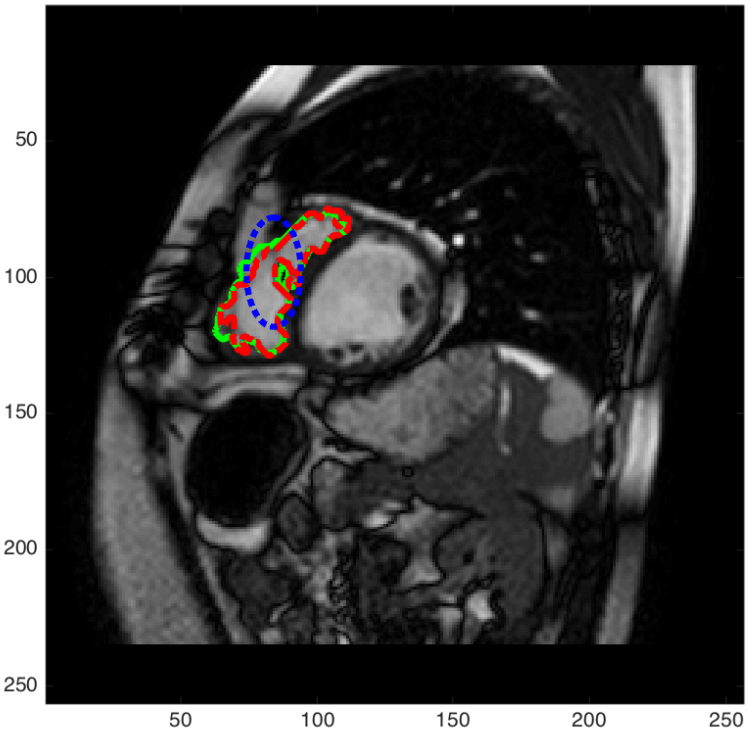

Figure 4.10. Comparing segmentation of a Right Ventricles (RV) using our approach (Chan-Vese + five RV priors with intrinsic scaling, translation and rotation). Solid green curve shows the segmentation, dotted blue curve shows the initialization and dashed red curve shows the ground truth.

Corpus callosum	Dice	False positives[%]	False negatives [%]
Chan-Vese	0.765 ± 0.120	68.2 ± 49.4	0.5 ± 0.6
Cremers' regularizer	0.859 ± 0.162	24.0 ± 9.2	4.7 ± 17.7
Our regularizer	0.919 ± 0.020	14.0 ± 5.1	3.1 ± 2.0

Table 4.1. Performance for the corpus callosum test data. Average Dice, false positives (expressed in % of the ground truth area) and false negatives (expressed in % of the ground truth area) of Chan-Vese algorithm with no prior, Cremers' regularizer and our regularizer. The training set consisted on 100 corpus callosum normal shapes, and the test set consisted on 60 corpus callosum pathological shapes (Alzheimer's disease).

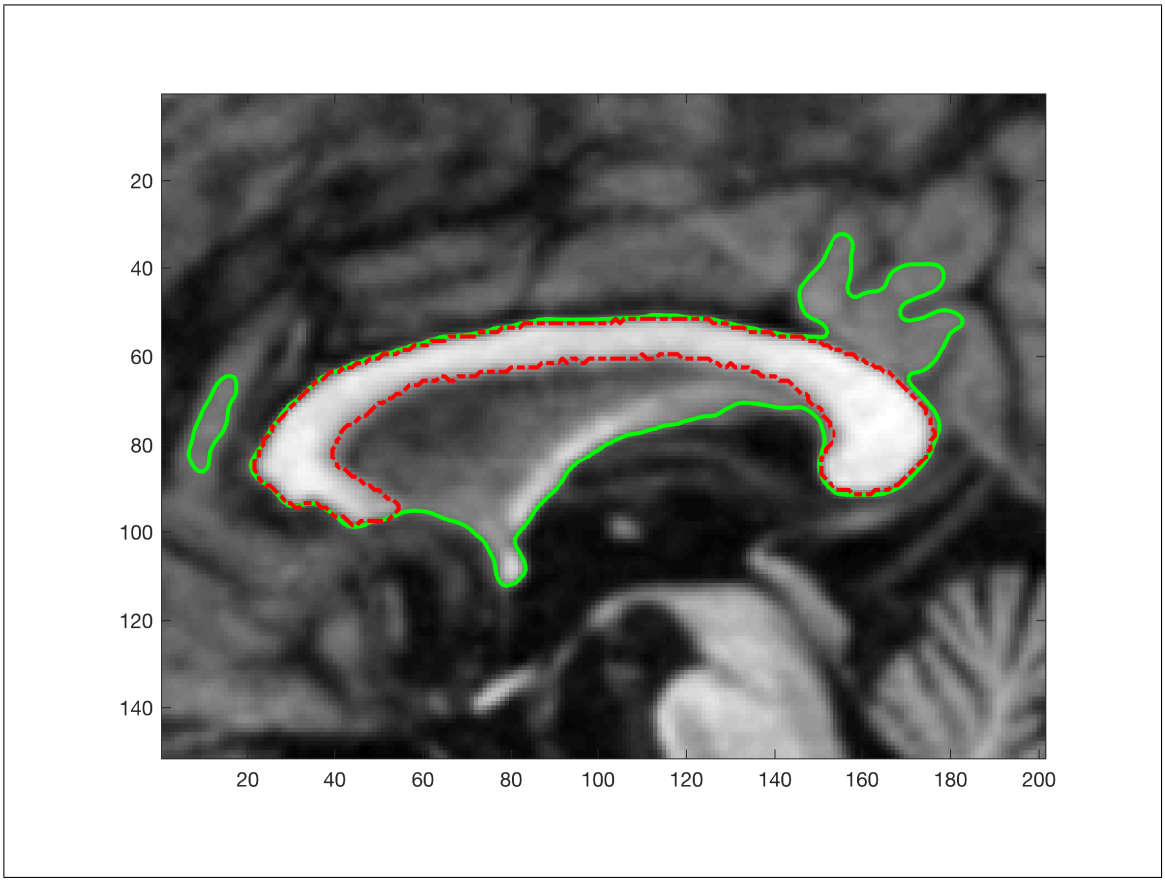

Figure 4.11. Comparing segmentation of Corpus Callosum (CC) using Chan-Vese algorithm without priors. Solid green curve shows the segmentation and dashed red curve shows the ground truth.

4.5. Experiments and results in 3D

We performed two 3D experiments. The first one was analogous to the first 2D experiment, but in this case we used a $15 \times 25 \times 5$-voxel parallelogram, with a superposed ellipsoid with semi-axes $(10, 18, 7)$ voxels, to generate the synthetic volume of $64 \times 64 \times 10$ voxels. An arbitrary rotation of 45° around the z-axis and 60° with respect to the y-axis was applied. The training shapes correspond to two different parallelograms in different positions. The second 3D experiment was the segmentation of a left ventricle from MR

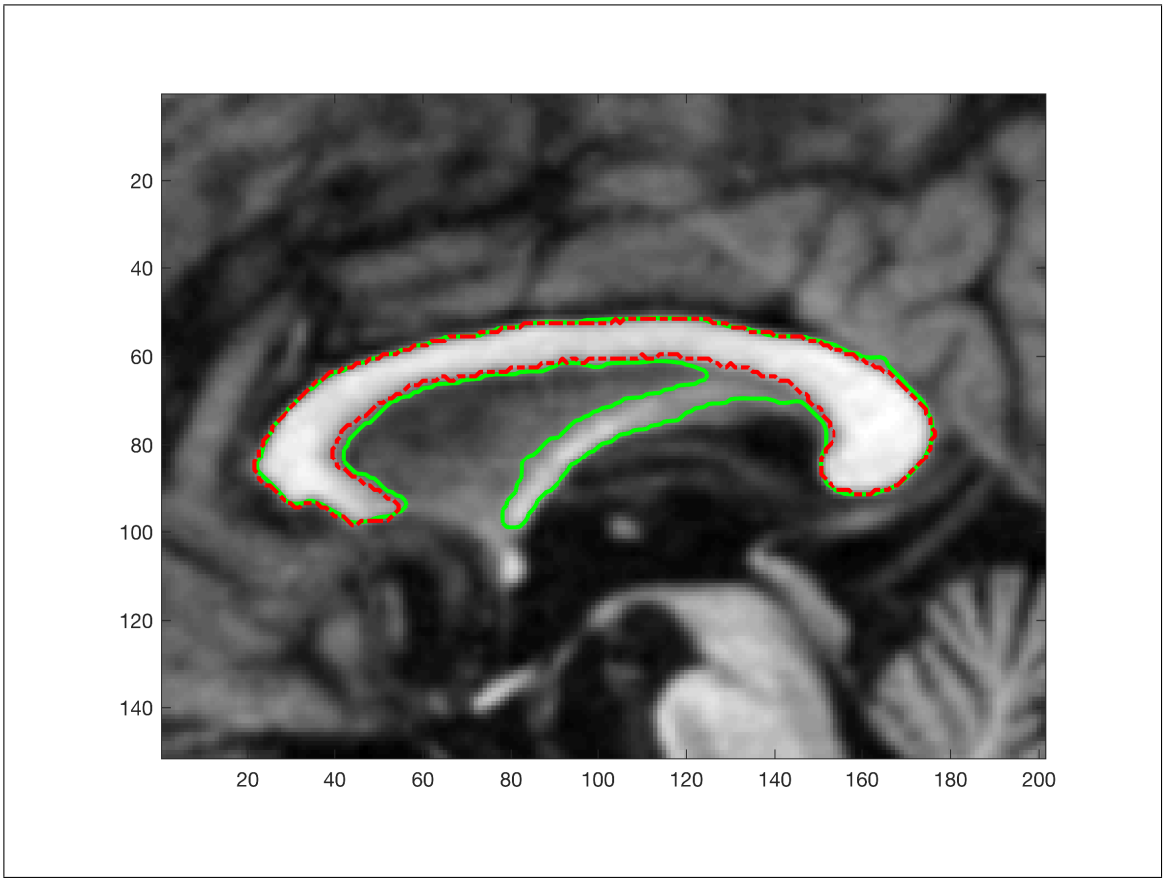

Figure 4.12. Comparing segmentation of Corpus Callosum (CC) using Cremers' regularizer (Chan-Vese + CC priors with intrinsic scaling and translation). Solid green curve shows the segmentation and dashed red curve shows the ground truth.

cardiac images, trained with three different left ventricles (Fig. 4.14) segmented by an expert.

The synthetic experiments showed that the 3D rotation was crucial in order to achieve the correct segmentation. Our approach successfully aligned data with relatively large rotations in the three axes (Figs. 4.15, 4.16, 4.17).

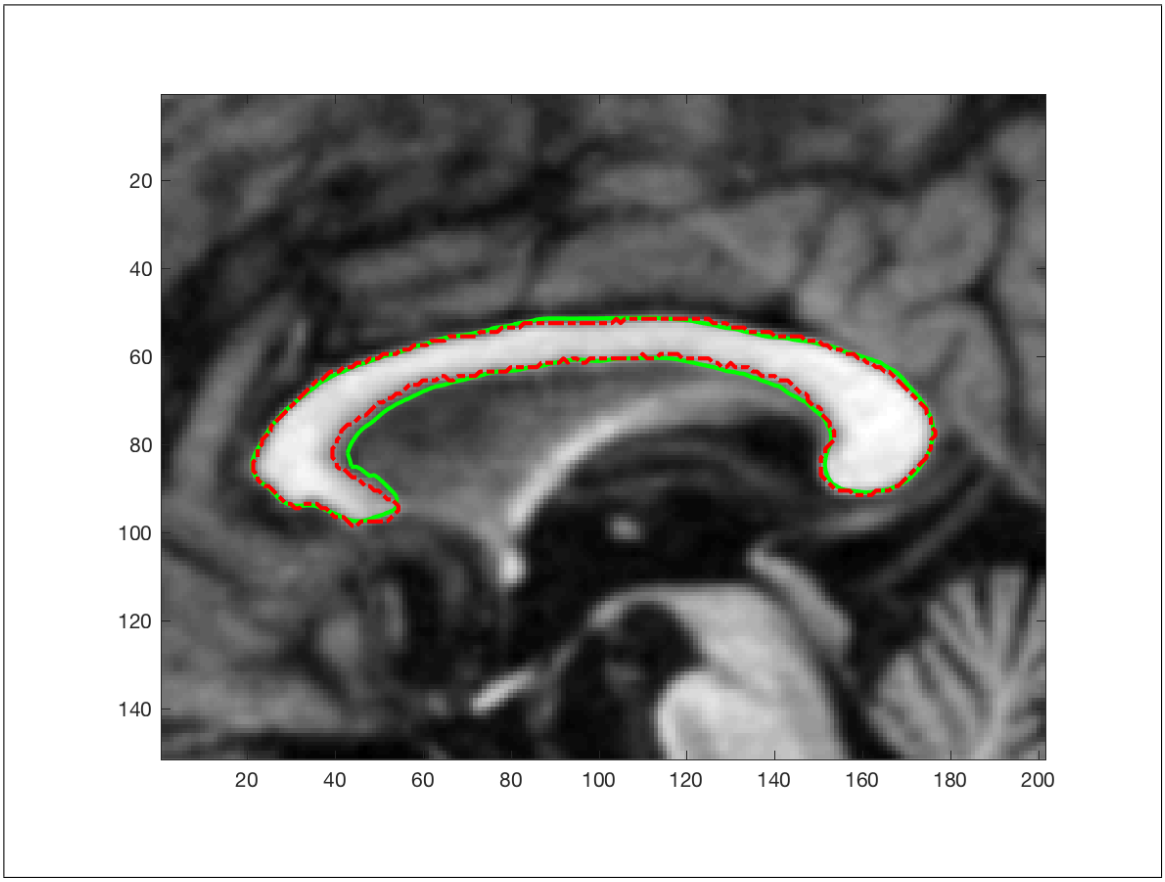

Figure 4.13. Comparing segmentation of Corpus Callosum (CC) using our approach (Chan-Vese + CC priors with intrinsic scaling, translation and rotation). Solid green curve shows the segmentation and dashed red curve shows the ground truth.

The left ventricle segmentation was a difficult task particularly at the apical and basal slices, which correspond to the bottom and the top of the ventricle. The top of the ventricle is limited by the mitral valve, which separates the left atrium from the left ventricle. Usually, this valve is not visible in MR images. In many applications, an expert performs this segmentation manually, looking several cardiac phases and exploring data using different orientations in order to define the correct position on the valve.

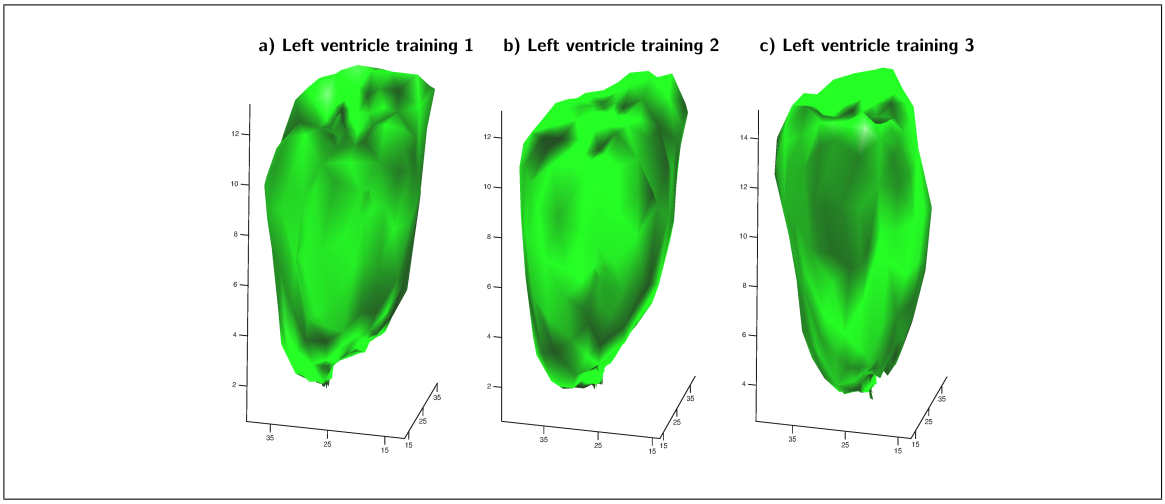

Figure 4.14. Left ventricle shape priors. The figure shows the surfaces which represents the shape of the left ventricle.

The Chan-Vese functional without any prior, is not enough to define those structures, especially at the basal slice where the surface goes into the atrium (Fig. 4.18). The combination of the Chan-Vese functional with the shape energy considering translation and scaling considerably improved the result, better delimiting the base and the apex. This approach obtained relatively successful results because all images were acquired in the same orientation (Fig. 4.19), in other words data was previously aligned. However, the Chan-Vese functional combined with the shape energy that considers translation, scaling and rotation, defined the basal region and apex even better, producing a segmentation close to the ground truth (Fig. 4.20).

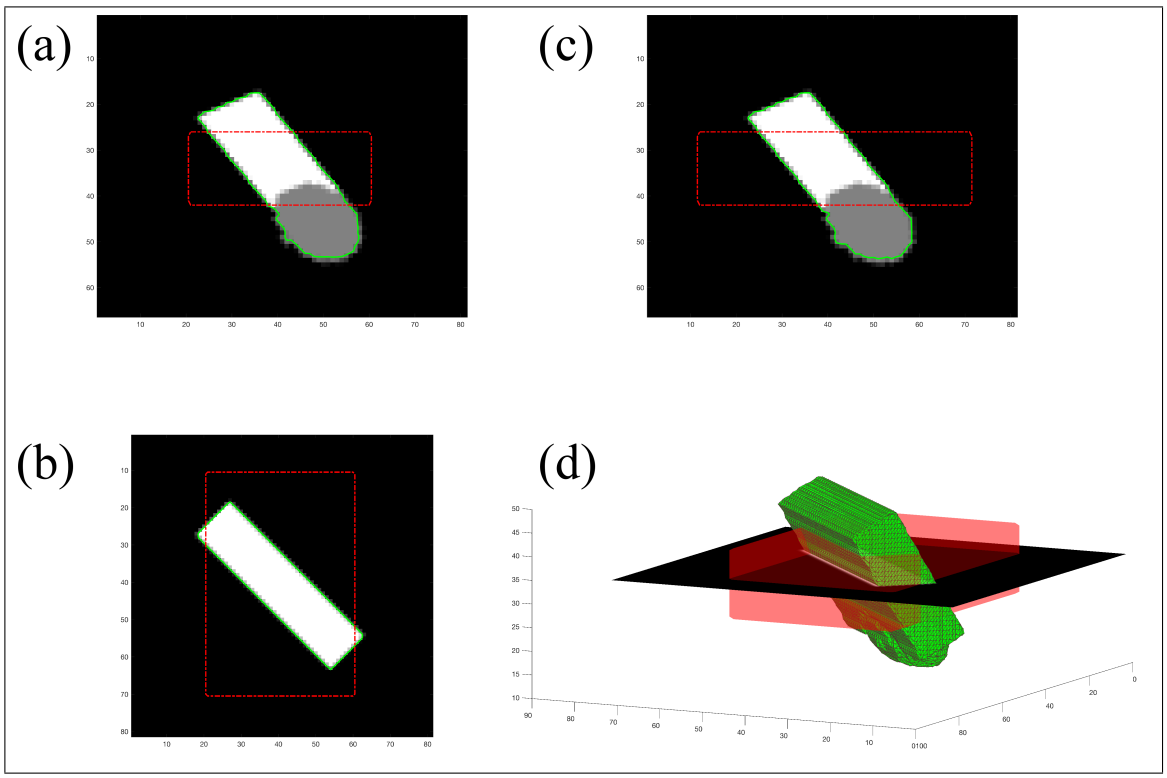

Figure 4.15. Segmentation of a synthetic parallelogram with a superposed ellipsoid using Chan-Vese (CV) without priors. (a), (b) and (c) correspond to three orthogonal slices of the synthetic shape and (d) correspond to the 3D rendering. Dashed line and surface in red show the initialization. Solid line and surface in green show the segmentation result.

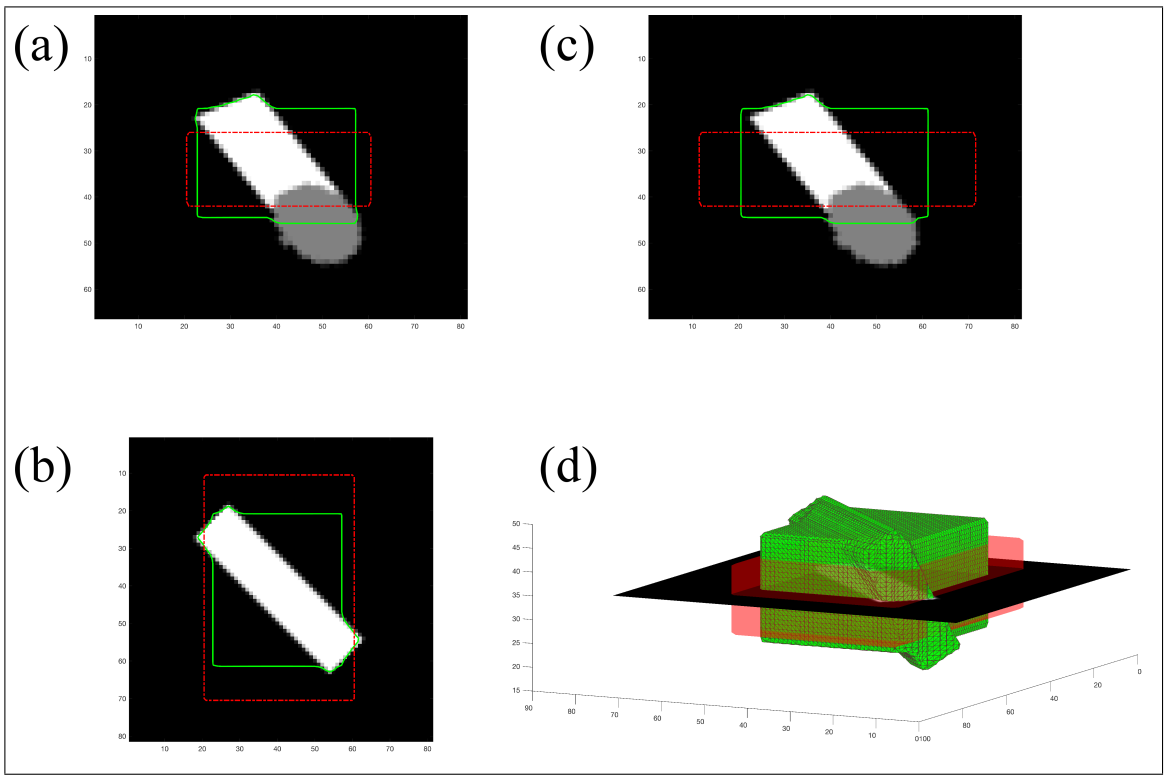

Figure 4.16. Segmentation of a synthetic parallelogram with a superposed ellipsoid using Chan-Vese and Cremers' regularizer. (a), (b) and (c) correspond to three orthogonal slices of the synthetic shape and (d) correspond to the 3D rendering. Dashed line and surface in red show the initialization. Solid line and surface in green show the segmentation result.

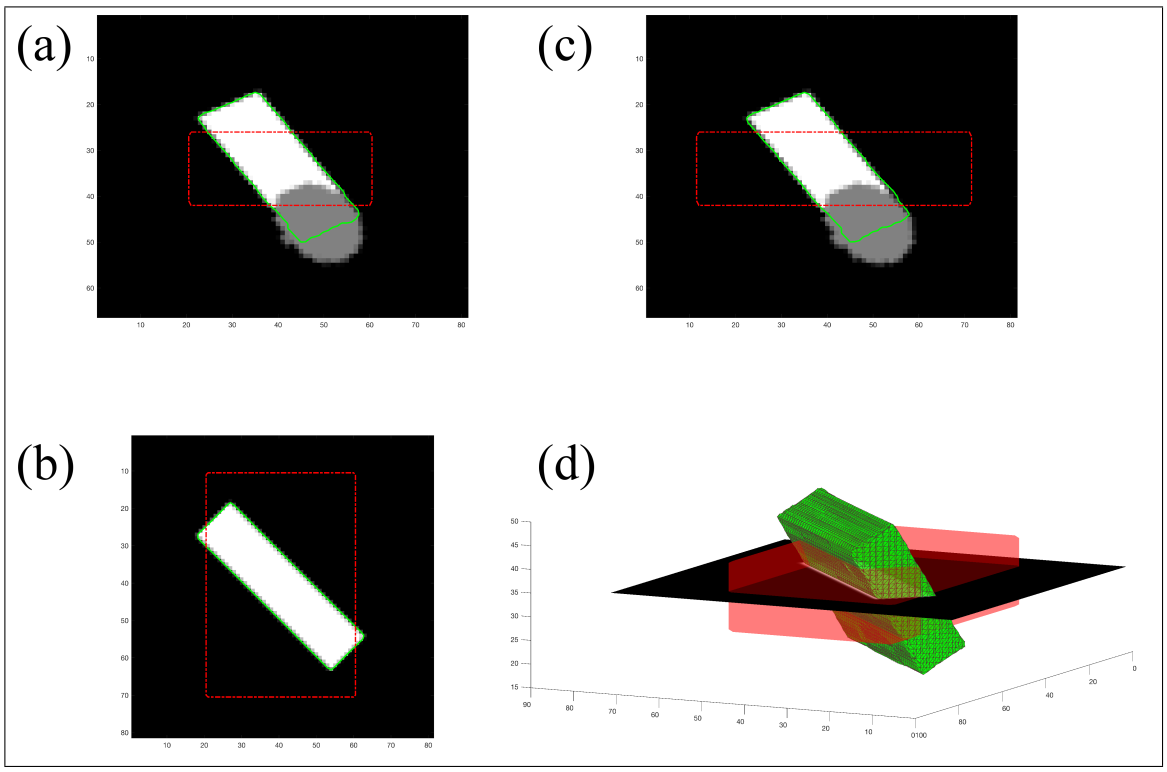

Figure 4.17. Segmentation of a synthetic parallelogram with a superposed ellipsoid using Chan-Vese and our regularizer. (a), (b) and (c) correspond to three orthogonal slices of the synthetic shape and (d) correspond to the 3D rendering. Dashed line and surface in red show the initialization. Solid line and surface in green show the segmentation result.

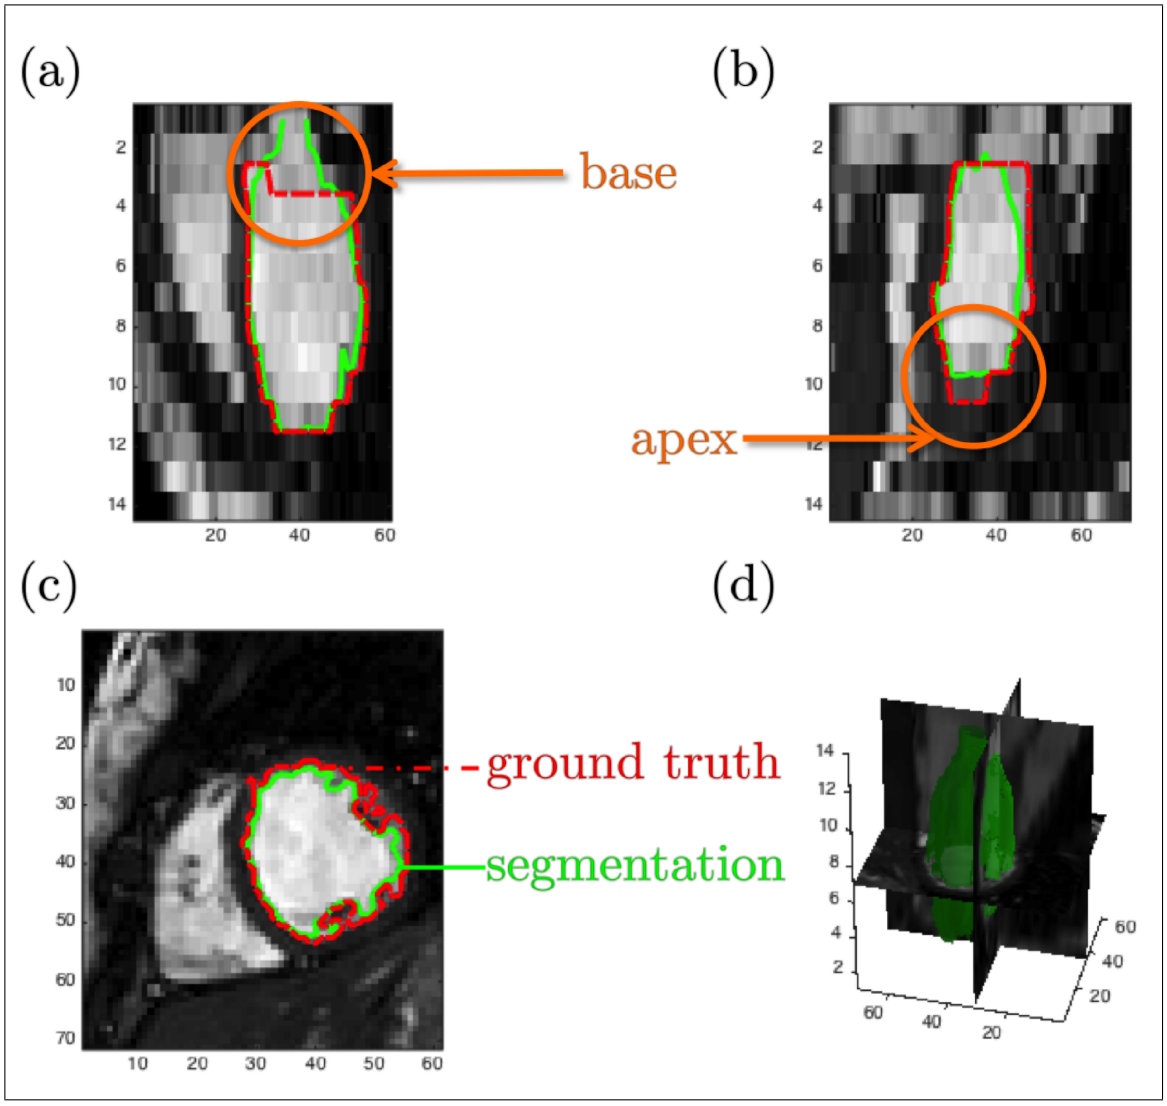

Figure 4.18. Segmentation of a 3D left ventricle using Chan-Vese (CV) without priors. (a), (b) and (c) correspond to three orthogonal slices of the heart. The red dashed line shows the Ground Truth (GT) and the green solid line shows the segmentation result. (a) shows how CV segmentation did not define well the mitral valve. (b) shows that CV segmentation cannot defined correctly the apex of the ventricle. (c) presents an axial slice that shows a better in-plane agreement between CV segmentation and the GT. Finally, (d) shows a 3D rendering of the segmented surface.

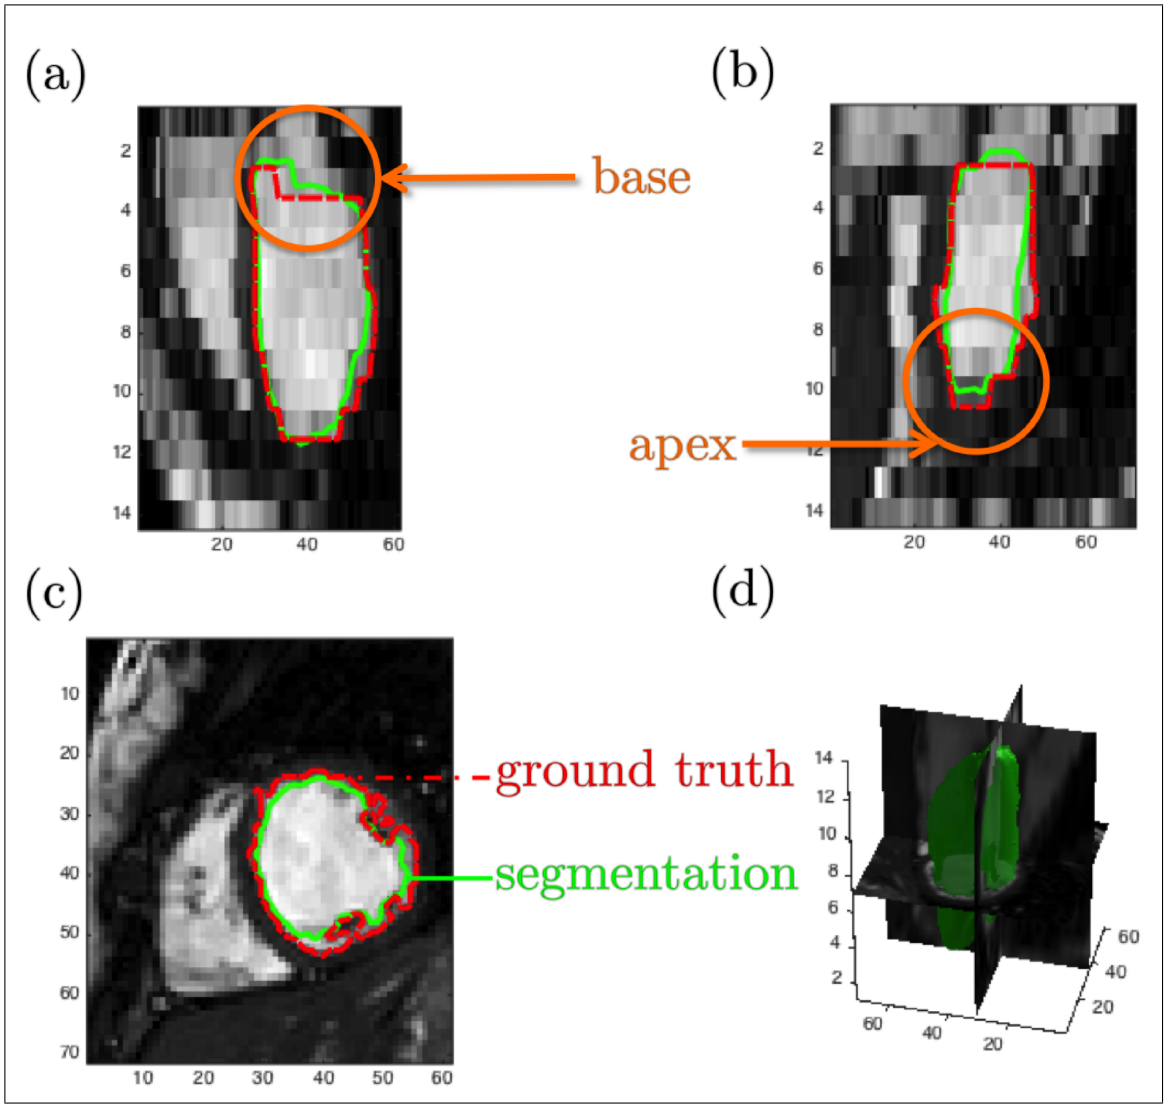

Figure 4.19. Segmentation of a 3D left ventricle using Cremers' approach. (a), (b) and (c) correspond to three orthogonal slices of the heart. The red dashed line shows the Ground Truth (GT) and the green solid line shows the segmentation result. (a) shows how Cremers' approach gets the mitral valve. (b) shows minor errors at the apex of the ventricle using Cremers' approach. (c) shows the result in an axial plane and (d) shows a 3D rendering of the segmented surface.

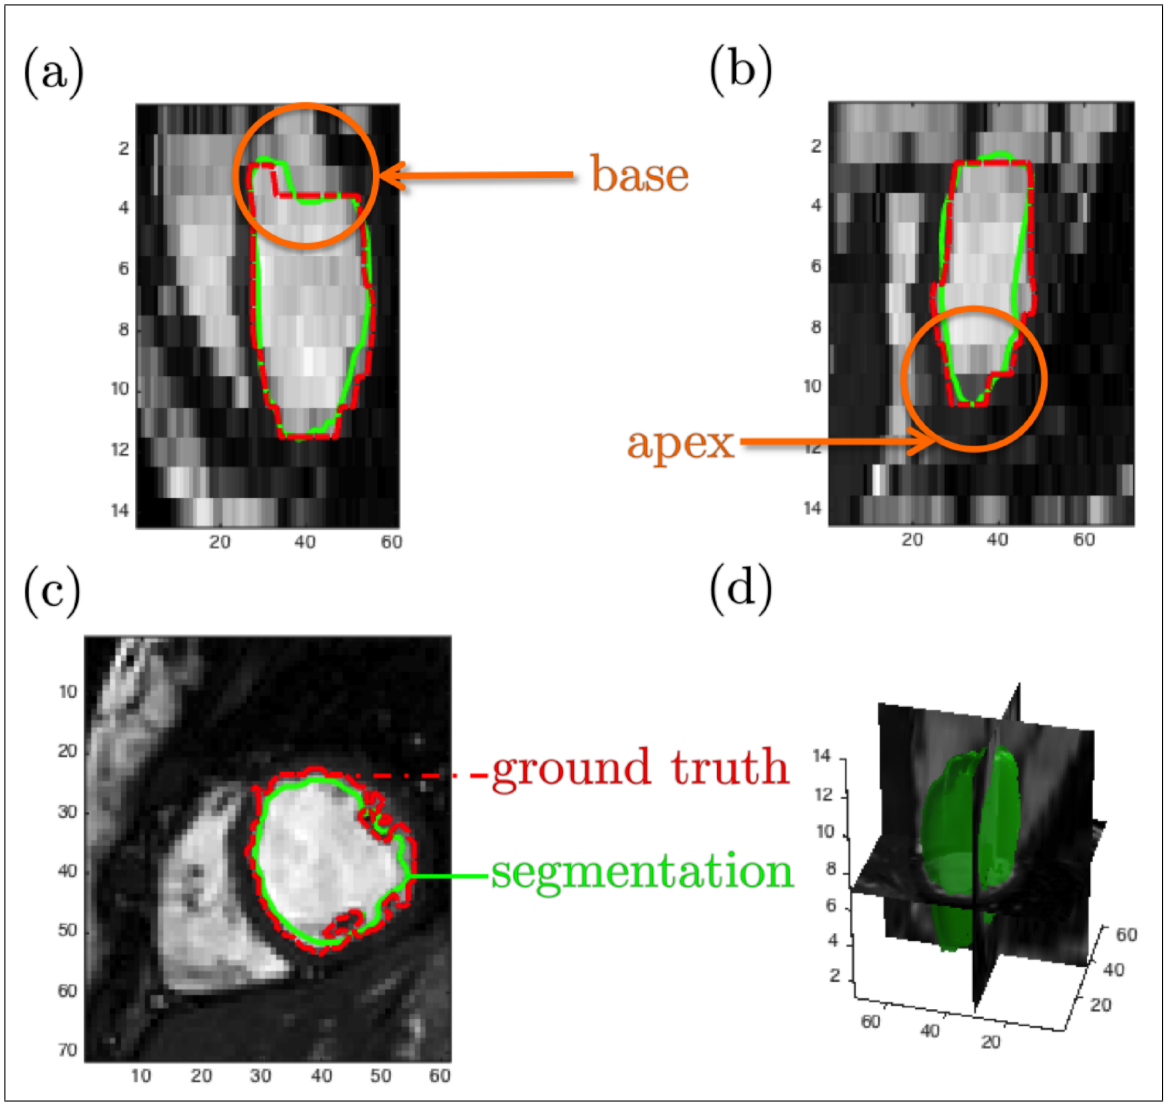

Figure 4.20. Segmentation of a 3D left ventricle using our approach. The red dashed line shows the Ground Truth (GT) and the green solid line shows the segmentation result. (a) shows how our approach matches the mitral valve. (b) shows that our approach provided a more accurate apex definition comparing with CV segmentation (Fig. 4.18) and Cremers' approach (Fig. 4.19). (c) shows the result in an axial plane and (d) shows a 3D rendering of the segmented surface.

5. CONCLUSIONS

In the context of image segmentation, the idea behind prior knowledge is to incorporate information inferred from already known characteristics or from a training database to the segmentation process. It becomes necessary when the image has low quality, high noise level, occlusion or when the visual information is not enough to find objects boundaries. In all those cases, not considering prior knowledge may lead to wrong segmentation results and tedious manual corrections are commonly needed.

To achieve the first goal of our thesis, which was to develop techniques for reliable segmenting medical image, we need to use prior knowledge. There are several kinds of prior knowledge, depending on the nature of the information considered. In this work we considered shape prior knowledge, which is useful when the shape of an object of interest is a distinguishable and relevant feature, as is, for example, for organs and tissues in medical images. In the context of pattern recognition and classification problems, there are several other features that can be used for priori knowledge, e.g. image intensity or texture of the objects of interest (Cootes, Edwards, & Taylor, 2001).

In this thesis we proposed and test two shape prior knowledge algorithms for level set segmentations. The first one corresponded to a weak shape prior knowledge method in which we integrated two already published works to create a useful tool to perform simultaneous segmentations of the right and left ventricles in cardiac magnetic resonance images. This was achieved by forcing the level set to preserve its topology during the entire deformation process. The second proposed algorithm was a new formulation for

strong shape prior knowledge based on intrinsic alignment that can deal with arbitrary translations, scaling and rotations of the training set. This algorithm is based on the eigen-decomposition of the shape covariance matrix and it was tested on 2D and 3D medical images.

The topology preserving level set algorithm enables us to segment simultaneously the left and right ventricles in patients with abnormal anatomies. Our method has a reasonable degree of automation without any need for training data sets. Simultaneous left and right ventricle segmentation is done relatively fast (about 15 seconds per cardiac phase). Two-tailored paired t-test showed no significant difference between our method and gold standards ($P < 0.05$), Pearsons correlation showed a high correlation of our measurement with gold standards (over 0.988), and Bland Altman analysis showed a high level of accuracy of our method compared with the chosen gold standards. Dice indexes also showed a high degree of accuracy (over 0.90). Minimal manual corrections were sometimes needed at regions with low contrast. This typically happens at basal slices, where valves do not appear clearly in the MR images.

As future work, the degree of automation of this first proposed method might be improved by incorporating information from other frames of the cardiac cycle. We expect this might provide a better definition of the valves at basal slices. Alternatively, a strong prior knowledge technique might produce more accurate results for these complex slices.

In our second proposed algorithm, we used the concept of shape as a strong prior. We extended the segmentation method based on level sets with shape prior knowledge proposed by Cremers et al. (2006), which uses intrinsic scaling and translation representation to account for pose invariance. As demonstrated, this is not enough in the context of medical imaging since rotation variability is commonly present. Patient positioning inside the scanner varies from one patient to another, or alternatively, intra-subject variability might incorporate rotation changes for some organs. Our contribution was therefore to extend Cremers level set framework to include an intrinsic rotation representation. This improvement resulted in a better alignment between the training shapes and the segmented object, without the need of solving additional optimization problems as in standard approaches based on registration. Our development is materialized through a shape regularization term that can be combined with any level set based segmentation algorithm. Our proposed method can work with training sets of different sizes and it also works in 2D as well as in 3D.

To deal with the intrinsic rotation and scaling factors, we used the eigenvectors and eigenvalues of the shape covariance matrix, respectively. This led to a completely new set of equations, based on the derivatives of the eigendecomposition. This formulation is independent of the dimensions of the problem (2D or higher), and allows for a closed form computation of the derivatives based on solving a linear system, which has low computational cost.

Our results show that our new shape regularization term can be used in a wide range of applications. We showed results for synthetic images with high level of noise and blurring, for occluded horse silhouettes, for real horse images and for medical images in 2D and 3D. In all cases, rotations played an important role, producing clear improvements in terms of accuracy. Medical image experiments demonstrated that, including rotations is crucial, even if the images were acquired with similar orientations, or if they were previously registered. This might be explained because of errors during the acquisition or the registration procedures, and also because of anatomical variations between subjects. Although a 3D problem has more degrees of freedom than 2D, in our approach the complexity of moving from 2D to 3D is negligible. This happens since the additional cost is to solve a slightly larger linear system.

Our shape regularization term has the ability to strongly impose prior shape knowledge in the regions where the priors get high consensus, and it lets the image information to drive the segmentation in regions where shapes have larger variability. This behavior was especially explicit in horse images and in medical images of the heart. We are currently working on more advanced application of this approach for different segmentation problems in 3D medical images such as whole heart, liver and hippocampus segmentations. Although our method can work with training data sets of any size, training sets of limited size seems to be enough to produce accurate results in several medical imaging applications.

The computation of our proposed regularizer is highly parallelizable, since each shape is processed independently. Therefore, in the future we plan to use GPU computing to substantially improve the performance of our algorithm.

REFERENCES

Ardekani, B., Bachman, A., Figarsky, K., & Sidtis, J. (2014). Corpus callosum shape changes in early alzheimers disease: an MRI study using the OASIS brain database. *Brain Structure and Function*, *219*(1), 343-352.

Aslan, M. S., Mostafa, E., Abdelmunim, H., Shalaby, A., Farag, A. A., & Arnold, B. (2011). Kernel principal component analysis and the construction of non-linear active shape models..

Bellenger, N., Burgess, M., Ray, S., Lahiri, A., Coats, A., Cleland, J., & Pennell., D. (2000). Comparison of left ventricular ejection fraction and volumes in heart failure by echocardiography, radionuclide ventriculography and cardiovascular magnetic resonance; are they interchangeable? *European Heart Journal 21 (16) (2000) 138796.*, *21(16)*, 1387-1396.

Bertrand, G., & Malgouyres, R. (1999). Some topological properties of surfaces in Z3. *Journal of Mathematical Imaging and Vision*, *11*, 207–221.

Borenstein, E., Sharon, E., & Ullman, S. (2004, June). Combining top-down and bottom-up segmentation. In *Conference on computer vision and pattern recognition workshop* (pp. 46–46).

Brown, E. S., Chan, T. F., & Bresso, X. (2012). Completely convex formulation of the chan-vese image segmentation model. *International Journal of Computer Vision*, *98*, 103-121.

Candes, E., Sing-Long, C., & Trzasko, J. (2013, Oct). Unbiased risk estimates for singular

value thresholding and spectral estimators. *IEEE Transactions on Signal Processing*, *61*(19), 4643–4657.

Carr, J., Simonetti, O., Bundy, J., & et al. (2001). Cine MR angiography of the heart with segmented true fast imaging with Steady-State Precession. *Radiology*, *219(3)*, 828–834.

Caselles, V., Catté, F., Coll, T., & Dibos, F. (1993). A geometric model for active contours in image processing. *Numerische Mathematik*, *66*(1), 1–31.

Caselles, V., Kimmel, R., & Sapiro, G. (1997). Geodesic active contours. *International Journal of Computer Vision*, *22(1)*, 61-79.

Chakraborty, A., Staib, L. H., & Duncan, J. S. (1996). Deformable boundary finding in medical images by integrating gradient and region information. *IEEE Transactions on Medical Imaging*, *15*, 859-870.

Chan, T., & Vese, L. (2001). Active contours without edges. *Image Processing, IEEE Transactions on*, *10*(2), 266–277.

Chen, S., & Radke, R. J. (2009). Level set segmentation with both shape and intensity priors..

Chen, Y., Tagare, H., Thiruvenkadam, S., Huang, F., Wilson, D., Gopinath, K., ... Geiser, E. (2002). Using prior shapes in geometric active contours in a variational framework. *International Journal of Computer Vision*, *50*(3), 315–328.

Cheng, J., & Foo, S. (2006). Dynamic directional gradient vector flow for snakes. *IEEE Transactions on Image Processing*, *15*(6), 1563–1571.

Cohen, L. D. (1991). Note on active contour models and balloons. *CVGIP: Image*

Understanding, *53*, 211-218.

Cohen, L. D., & Cohen, I. (1993). Finite-element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Transactions on Image Processing*, *15*, 1131-1147.

Cootes, T. F., Cooper, D., Taylor, C. J., & Graham, J. (1995). Active shape models their training and application. *Computer Vision and Image Understanding*, *61*(1), 38-59.

Cootes, T. F., Edwards, G. J., & Taylor, C. J. (2001). Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *23*(6), 681–685.

Cremers, D., Kohlberger, T., & Schnörr, C. (2003). Shape statistics in kernel space for variational image segmentation. *Pattern Recognition*, *36*, 1929-1943.

Cremers, D., Osher, S., & Soatto, S. (2006). Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *International Journal of Computer Vision*, *69*(3), 335-351.

Cremers, D., Tischhäuser, F., Weickert, J., & Schnör, C. (2002). Diffusion snakes: Introducing statistical shape knowledge into the mumford-shah functional. *International Journal of Computer Vision*, *50*, 295-313.

Daněk, O., Matula, P., Maška, M., & Kozubek, M. (2012). Smooth chanvese segmentation via graph cuts. *Pattern Recognition Letters*, *33*, 1405-1410.

Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology*, *26*(3), 297–302.

Dong, F., Chen, Z., & Wang, J. (2013). A new level set method for inhomogeneous image segmentation. *Image and Vision Computing*, *31*(10), 809–822.

Duan, J., Pa, Z., Yin, X., Wei, W., & Wang, G. (1994). A level set approach for computing solutions to incompressible two-phase flow. *Journal of Computational Physics*, *119*, 146–159.

Duan, J., Pa, Z., Yin, X., Wei, W., & Wang, G. (2014). Some fast projection methods based on chan-vese model for image segmentation. *EURASIP Journal on Image and Video Processing*, *7*.

Eckhardta, U., & Lateckib, L. J. (2003). Topologies for the digital spaces Z2 and Z3. *Computer Vision and Image Understanding*, *90*, 295–312.

El-Berbari, R., Bloch, I., Redheuil, A., Angelini, E., Mousseaux, E., Frouin, F., & Herment, A. (2007). An automated myocardial segmentation in cardiac MRI. In *29th annual international conference of the IEEE engineering in medicine and biology society* (pp. 4508–4511).

Geiger, D., Gupta, A., Costa, L., & Vlontzos, J. (1995). Dynamic programming for detecting, tracking, and matching deformable contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *17*(3), 294–302.

Golland, P., & Fischl, B. (2003). Permutation tests for classification: Towards statistical significance in image-based studies. In *Biennial international conference on information processing in medical imaging* (pp. 330–341).

Grosgeorge, D., Petitjean, C., Caudron, J., Fares, J., & Dacher, J. (2011). Automatic cardiac ventricle segmentation in MR images: a validation study. *International Journal of Computer Assisted Radiology and Surgery*, *6(5)*, 573–581.

Grothues, F., Smith, G., Moon, J., Bellenger, N., Collins, P., Klein, H., & Pennell, D.

(2002). Comparison of interstudy reproducibility of cardiovascular magnetic resonance with two-dimensional echocardiography in normal subjects and in patients with heart failure or left ventricular hypertrophy. *American Journal of Cardiology*, *90*, 29–34.

Gupta, A., Von Kurowski, L., Singh, A., Geiger, D., Liang, C.-C., Chiu, M.-Y., ... Wilson, D. (1993). Cardiac MR image segmentation using deformable models. In (pp. 747–750).

Han, X., Xu, C., & Prince, J. (2003). A topology preserving level set method for geometric deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *25*(6), 755–768.

He, L., Peng, Z., Everding, B., Wang, X., Han, C. Y., Weiss, K. L., & Wee, W. G. (2008). A comparative study of deformable contour methods on medical image segmentation. *Image and Vision Computing*, *26*(2), 141–163.

Heiberg, E., Sjögren, J., Ugander, M., Carlsson, M., Engblom, H., & Arheden, H. (2010). Design and validation of segment a freely available software for cardiovascular image analysis. *BMC Medical Imaging*, *10*(1).

Ino, T., Benson, L. N., Mikalian, H., Freedom, R. M., & Rowe, R. D. (1989). Determination of left ventricular volumes by simpson's rule in infants and children with congenital heart disease. *British Heart Journal*, *61(2)*, 182–185.

Ivanovska, T., Laqua, R., Wang, L., Schenk, A., Yoon, J. H., Hegenscheid, K., ... Liebscher, V. (2016). An efficient level set method for simultaneous intensity inhomogeneity correction and segmentation of MR images. *Computerized Medical Imaging*

and Graphics, *48*, 9–20.

Jacquier, A., Thuny, F., Jop, B., Giorgi, R., Cohen, F., Gaubert, J.-Y., ... Moulin, G. (2010). Measurement of trabeculated left ventricular mass using cardiac magnetic resonance imaging in the diagnosis of left ventricular non-compaction. *European Heart Journal*, *31(9)*, 1098-1104.

Jahnke, C., Fischer, J., Gerds-Li, J.-H., Gebker, R., Manka, R., Fleck, E., ... Kriatselis, C. (2011). Serial monitoring of reverse left-atrial remodeling after pulmonary vein isolation in patients with atrial fibrillation: A magnetic resonance imaging study. *International Journal of Cardiology*, *153*, 42–46.

Kass, M., Witkin, A., & Terzopoulos, D. (1988). Snakes: active contour models. *Communications on Pure and Applied Mathematics*, *1*, 321–332.

Kong, T. Y., & Rosenfeld, A. (1989). Computer vision, graphics and image processing. *Journal of Mathematical Imaging and Vision*, *48*, 357–393.

Leventon, M., Grimson, W., & Faugeras, O. (2000). Statistical shape influence in geodesic active contours. In *IEEE conference on computer vision and pattern recognition* (Vol. 1, pp. 316–323).

Liu, L., Zhang, Q., Wu, M., Li, W., & Shang, F. (2013). Adaptive segmentation of magnetic resonance images with intensity inhomogeneity using level set method. *Magnetic Resonance Imaging*, *31*(4), 567–574.

Liu, Q., Jiang, M., Bai, P., & Yang, G. (2016). A novel level set model with automated initialization and controlling parameters for medical image segmentation. *Computerized Medical Imaging and Graphics*, *48*, 21–29.

Mahapatra, D., & Buhmann, J. (2012). Cardiac LV and RV segmentation using mutual context information. In F. Wang, D. Shen, P. Yan, & K. Suzuki (Eds.), *Machine learning in medical imaging* (Vol. 7588, pp. 201–209). Berlin, Heidelberg: Springer.

Mahapatra, D., & Buhmann, J. (2013). Cardiac MRI segmentation using mutual context information from left and right ventricle. *Journal of Digital Imaging*, *5*, 898–908.

Marcus, D., Wang, T., Parker, J., Csernansky, J., Morris, J., & Buckner, R. (2007, September). Open access series of imaging studies (oasis): cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *Journal of Cognitive Neuroscience*, *19*(9), 1498–1507.

Martin Bland, J., & Altman, D. (1986). Statistical methods for assessing agreement between two methods of clinical measurement. *The lancet*, *327*(8476), 307–310.

Masutani, Y., Nemoto, M., Nomura, Y., & Hayashi, N. (2012). Clinical machine learning in action: CAD system design, development, tuning, and long-term experience. In K. Suzuki (Ed.), *Machine learning in computer-aided diagnosis: Medical imaging intelligence and analysis* (pp. 159–176). IGI Global.

Mazurowskia, M. A., Habasa, P. A., Zuradaa, J. M., Lob, J. Y., Bakerb, J. A., & Tourassib, G. D. (2008). Training neural network classifiers for medical decision making: The effects of imbalanced datasets on classification performance. *Neural Networks*, *21*(2–3), 227–436.

McInerney, T., & Terzopoulos, D. (2000). T-snakes: Topology adaptive snakes. *Medical Image Analysis*, *4*, 73-91.

Mitchell, S., Lelieveldt, B., van der Geest, R., Bosch, H., Reiver, J., & Sonka, M. (2001). Multistage hybrid active appearance model matching: segmentation of left and right ventricles in cardiac MR images. *IEEE Transactions on Medical Imaging*, *20*(5), 415–423.

Montillo, A., Metaxas, D., & Axel, L. (2003). Automated model-based segmentation of the left and right ventricles in tagged cardiac MRI. In R. Ellis & T. Peters (Eds.), *Medical image computing and computer-assisted intervention - miccai 2003* (Vol. 2878, pp. 507–515). Berlin, Heidelberg: Springer.

Mumford, D., & Shah, J. (1989). Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, *42*, 577-685.

Niessen, W., ter Haar Romeny, B., & Viergever, M. (1998). Geodesic deformable models for medical image analysis. *IEEE Transactions on Medical Imaging*, *17*(4), 634-641.

Osher, S., & Fedkiw, R. (2003). Constructing signed distance functions. In S. Osher & R. Fedkiw (Eds.), *Level set methods and dynamic implicit surfaces* (Vol. 153, pp. 63–74). Berlin, Heidelberg: Springer.

Osher, S., & Sethian, J. A. (1988). Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, *79*(1), 12–49.

Paragios, N., Mellina-Gottardo, O., & Ramesh, V. (2001). Gradient vector flow fast geodesic active contours. In *Eighth IEEE international conference on computer*

vision (Vol. 1, pp. 67–73).

Pentland, A., & Sclaroff, S. (1991). Closed-form solutions for physically based shape modeling and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *13*(7), 715-729.

Petitjean, C., & Dacher, J.-N. (2011). A review of segmentation methods in short axis cardiac MR images. *Medical image analysis*, *15*(2), 169–84.

Pieciak, T. (2012). Segmentation of the left ventricle using active contour method with gradient vector flow forces in short-axis MRIs. In E. Pitka & J. Kawa (Eds.), *Information technologies in biomedicine* (Vol. 7339, pp. 24–35). Berlin, Heidelberg: Springer.

Pluempitiwiriyawej, C., Moura, J. M. F., Wu, Y.-J. L., & Ho, C. (2005). STACS: New active contour scheme for cardiac MR image segmentation. *IEEE Transaction on Medical Imaging*, *24(5)*, 593–603.

Romdhani, S., Gong, S., & Psarrou, A. (1999). A multi-view nonlinear active shape model using kernel PCA. In (p. 483-492).

Ronfard, R. (1994). Region-based strategies for active contour models. *International Journal of Computer Vision*, *13*, 229-251.

Rousson, M., & Paragios, N. (2002). Shape priors for level set representations. In A. Heyden, G. Sparr, M. Nielsen, & P. Johansen (Eds.), *European conference in computer vision* (Vol. 2351, pp. 78–92). Springer Berlin Heidelberg.

Rousson, M., & Paragios, N. (2008). Prior knowledge, level set representations & visual grouping. *International Journal of Computer Vision*, *76*(3), 231–243.

Santarelli, M., Positano, V., Michelassi, C., Lombardi, M., & Landini, L. (2003). Automated cardiac MR image segmentation: theory and measurement evaluation. *Medical Engineering and Physics*, *25*(2), 149–159.

Sardanelli, F., Quarenghi, M., Di Leo, G., Boccaccini, L., & Schiavi, A. (2008). Segmentation of cardiac cine MR images of left and right ventricles: Interactive semiautomated methods and manual contouring by two readers with different education and experience. *Journal of Magnetic Resonance Imaging*, *27*(4), 785–792.

Souto, M., Dias, A., Trkbey, E., Marchiori, E., N. Kawel, R. A. F. M., Oliveira, R., ... Rochitte, C. E. (2013). Left atrial volume quantification using cardiac MRI in atrial fibrillation: comparison of the simpsons method with biplane area-length, ellipse, and three-dimensional methods. *Diagnostic and Interventional Radiology*, *19*, 213-220.

Souto, M., Masip, L. R., Couto, M., Surez-Cuenca, J. J., Martnez, A., Tahoces, P. G., ... Croisille, P. (2013). Quantification of right and left ventricular function in cardiac MR imaging: Comparison of semiautomatic and manual segmentation algorithms. *Diagnostics*, *3*(2), 271–282.

Suri, J., Liu, K., Singh, S., Laxminarayan, S., Zeng, X., & Reden, L. (2002). Shape recovery algorithms using level sets in 2-D/3-D medical imagery: a state-of-the-art review. *IEEE Transactions on Information Technology in Biomedicine*, *6*(1), 8-28.

Sussman, M., Smereka, P., & Osher, S. (1994). A level set approach for computing solutions to incompressible two-phase flow. *Journal of Computational Physics*, *114(1)*, 146–159.

Tajbakhsh, N., Shin, J. Y., Gurudu, S. R., Hurst, R. T., Kendall, C. B., Gotway, M. B., & Liang, J. (2016). Convolutional neural networks for medical image analysis: Full training or fine tuning? *IEEE Transactions on Medical Imaging*, *35*(5), 1299–1312.

Tavakoli, V., & Amini, A. (2013). A survey of shaped-based registration and segmentation techniques for cardiac images. *Computer Vision and Image Understanding*, *117*, 96698.

Tejos, C., Irarrazaval, P., & Cárdenas-Blanco, A. (2009). Simplex mesh diffusion snakes: Integrating 2d and 3d deformable models and statistical shape knowledge in a variational framework. *International Journal of Computer Vision*, *85*(1), 19-34.

Thiele, H., Nagel, E., Paetsch, I., & et al. (2001). Functional cardiac MR imaging with steady-state free precession (SSFP) significantly improves endocardial border delineation without contrast agents. *Journal of Magnetic Resonance Imaging*, *14(4)*, 362–367.

Tsai, A., Yezzi, J., A., Wells, W., Tempany, C., Tucker, D., Fan, A., ... Willsky, A. (2003, Feb). A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging*, *22*(2), 137–154.

Twining, C. J., & Taylor, C. J. (2001). Kernel principal component analysis and the construction of non-linear active shape models. In (pp. 23–32).

Wang, Z., Salah, M. B., Gu, B., Islam, A., Goela, A., & Li, S. (2014). Direct estimation of cardiac biventricular volumes with an adapted bayesian formulation. *IEEE Transaction on Biomedical Engineering*, *61(4)*, 12511260.

Xu, C., & Prince, J. (1998b). Snakes, shapes, and gradient vector flow. *IEEE Transactions*

on Image Processing, *7*(3), 359–369.

Xu, C., & Prince, J. L. (1998a). Generalized gradient vector flow external forces for active contours. *Signal Processing*, *71*(2), 131–139.

Yeo, S. Y., Xie, X., Sazonov, I., & Nithiarasu, P. (2011). Level set segmentation with robust image gradient energy and statistical shape prior..

Zhu, S. C., & Yuille, A. (1996). Region competition: Unifying snakes, region growing, and Bayes/MDL for multi-band image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *18*, 884-900.

Zhu, W., Ha Kang, S., & Biros, G. (2013). A geodesic-active-contour-based variational model for short-axis cardiac MR image segmentation. *International Journal of Computer Mathematics*, *90*(1), 124–139.

APPENDICES

A. TOPOLOGY PRESERVING LEVEL SET ALGORITHM

In order to clearly describe the topology preserving level set method, we need to introduce some concepts of digital topology. A more detailed discussion can be found in (Kong & Rosenfeld, 1989; Bertrand & Malgouyres, 1999; Eckhardta & Lateckib, 2003). We can start discussing the notion of simple closed curve in the digital space (or binary image) $\mathcal{Z}^2$. We focus our discussion in a two-dimensional domain since all concepts can be easily extended to a higher dimensional domain.

Firstly, each point of the digital plane or pixel, is associated to a lattice point with integer coordinates. In order to define a curve we need to introduce the concept of digital connectivity.

Two points of the digital plain are 8-adjacent if they are distinct and their corresponding coordinates differ by at most 1. All points that satisfy this distance are 8-neighbors. Two points of the digital plain are 4-adjacent if they are 8-adjacent and they differ in at most one of their coordinates. All points that satisfies this conditions are 4-neighbors (Kong & Rosenfeld, 1989). We can define mathematically those two neighborhoods $N_8(x)$ and $N_4(x)$ as follows. Let be $X \subset \mathcal{Z}^2$, and $x \in X$ with coordinates (x_1, x_2) (Eckhardta & Lateckib, 2003), then

$$N_8(x) = \{y \in X; Max(|x_1 - y_1|, |x_2 - y_2|) \leq 1\}, \tag{A.1}$$

$$N_4(x) = \{y \in X; |x_1 - y_1| + |x_2 - y_2| \leq 1\}. \tag{A.2}$$

We define n-neighborhood of x excluding x as $N_n^*(x) = N_n(x) \setminus \{x\}$. Now, an n-path is a sequence of points $x_0, ..., x_k$ in which x_i is n-adjacent to x_{i-1} for $i = 1, ..., k$. If the path is not empty, then path length is equal to k. We say that the path is closed if $x_0 = x_k$. Two points $x, y \in X$ are n-connected if there is an n-path in X between them. The entire set X is n-connected if X cannot be partitioned into two subsets that are not n-adjacent to each other. An n-component of X is a non-empty n-connected subset of X that is not n-adjacent to any other point in X.

We define $\bar{X}$ as the complement of X. X is a simple closed n-curve if it is n-connected and if each point of X is n-adjacent to exactly two points in X. This definition implies an important statement: In order to have a correspondence between the topology of X and its complement $\bar{X}$, different connectivities, n and $\bar{n}$, need to be assume, respectively. Thus, in $\mathcal{Z}^2$ we have the connetivities $(n, \bar{n}) = (8, 4), (4, 8)$ (Kong & Rosenfeld, 1989; Eckhardta & Lateckib, 2003).

An intuitive example is presented in Fig. A.1. The digital object represented in black corresponds to a one 8-component which could be a simple closed 8-curve. However, if we consider the 8-connectivity for background in white, the interior and exterior objects are

connected. However, if we consider a 4-connectivity for the background, in this case the 8-curve divides the space into two disjoint regions satisfying the Jordan property (Eckhardta & Lateckib, 2003), demonstrating that the pair of connectivities $(n, \bar{n}) = (8, 4)$ avoids the topological ambiguity.

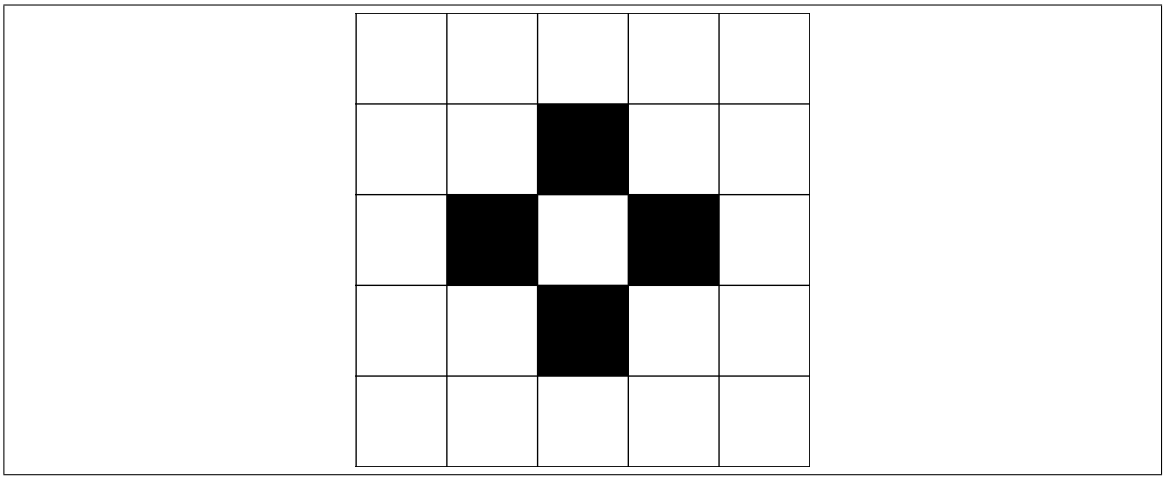

Figure A.1. Example of how the connectivity pair $(n, \bar{n}) = (8, 4)$ allows the topological correspondence between the 8-connected object (in black) and its 4-connected background (in white).

Coming back to the preserving topology algorithm, we introduce the concept of n-simple point. A point $x \in \mathcal{Z}^2$ is n-simple for $X \subset \mathcal{Z}^2$ if and only if its deletion from X (if $x \in X$) or its addition to X (if $x \in \bar{X}$) does not change the number of n-components of X and the number of $\bar{n}$-components of $\bar{X}$. Therefore, in order to preserve the topology of the object define by the zero-level curve of the signed distance function ϕ, we need to check whether the points that enter or leave the curve are simple or not. Fortunately, in the level set context, the set of points that enters or leaves the curve corresponds to sign changes of ϕ from one iteration to the other, and thus simple points of this set are the only ones that are updated.

To check if a point is n-simple, we introduce the concept of geodesic n-neighborhood of x inside X of order k, $N_n^k(x, X)$, which is the set of all points $y \in N_8^*(n) \cap X$ such that there exists an n-path from x to y of length less than or equal to k. The $N_n^k(x, X)$ can be defined recursively by A.4 (Eckhardta & Lateckib, 2003).

$$N_n^1(x, X) = N_n^*(x) \cap X, \tag{A.3}$$

$$N_n^k(x, X) = \cup\{N_n(y) \cap N_8^*(x) \cap X, y \in N_n^{k-1}(x, X)\}. \tag{A.4}$$

Considering the geodesic neighborhoods $G_4(x, X) = N_4^2(x, X)$ and $G_8(x, X) = N_8^1(x, X)$, we define the topological number $T_n(x, X)$ as the number of n-components of $G_n(x, X)$. We can evaluate if a point is simple or not, checking this local feature of the topological number, as it is established by the theorem 0.1,

Theorem 0.1. *Let $X \subset \mathcal{Z}^2$ and $x \in X$, then x is a simple point $\iff$ $T_n(x, X) = 1$ and $T_{\bar{n}}(x, \bar{X}) = 1$, being $(n, \bar{n})$ a pair of compatible connectivities*

In conclusion, the algorithm for preserving topology is detailed as follows (Han et al., 2003):

1: update ϕ^{k+1}

2: detect the sign changes of ϕ^{k+1} with respect to ϕ^k

3: **for** all pixels that have changed its sign **do**

4: **if** this pixel is an 8-simple point (i.e. $T_8(x, X) = 1$ and $T_4(x, \bar{X}) = 1$) **then**

5: accept the new value of ϕ^{k+1}

6: **else**

7: replace the value of ϕ^{k+1} by $\mathrm{sign}(\phi^k) * 0.001$

8: **end if**

9: **end for**

Some final comments about the algorithm Notice that when the pixel is not 8-simple is reject but its value in $\phi^{(k+1)}$ is replaced by a small constant in order to facilitate the probability that change its sign in the next iteration. This action produce that ϕ moves away from a signed distance function and thus the reinitialization process becomes crucial. In other hand, notice that the speed of the algorithm depends on the number candidates to analyze. In practice, the first iterations are slower since there are too many analyzed candidates but final iterations are considerably faster.

B. LINEAR SYSTEM SOLUTION FOR THE JACOBIAN OF THE EIGENDECOMPOSITION

This appendix shows the solution of the linear system presented in 4.38. For the 2D case, re-writing the right hand side of the equation give us:

$$Q^T \frac{\partial \Sigma}{\partial s_{ij}} Q = \begin{bmatrix} 0 & \gamma \\ -\gamma & 0 \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} + \begin{bmatrix} \frac{\partial \lambda_1}{\partial s_{ij}} & 0 \\ 0 & \frac{\partial \lambda_2}{\partial s_{ij}} \end{bmatrix} + \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} 0 & -\gamma \\ \gamma & 0 \end{bmatrix}.$$

Now 4.38 can be written as

$$Q^T \frac{\partial \Sigma}{\partial s_{ij}} Q = \begin{bmatrix} \frac{\partial \lambda_1}{\partial s_{ij}} & \gamma(\lambda_2 - \lambda_1) \\ \gamma(\lambda_2 - \lambda_1) & \frac{\partial \lambda_2}{\partial s_{ij}} \end{bmatrix}.$$

The left hand side of the equation changes depending on the chosen derivative of Σ (i.e., the chosen (i, j) element of $\frac{\partial \lambda_1}{\partial s_{ij}}$), as can be seen in B.1, B.2, B.5, and B.8. The right hand side of the equation will be always the same, except from the values of γ which also depend on the chosen derivative. We call γ^{ij} the solution of the system when the derivative respect to the component i, j of Σ is taken. For 2D we write the four components of the Jacobian of the eigendecomposition as:

a) $i = 1$**,** $j = 1$

$$\begin{bmatrix} q_{11} & q_{21} \\ q_{12} & q_{22} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix} = \begin{bmatrix} \frac{\partial \lambda_1}{\partial s_{11}} & \gamma^{11}(\lambda_2 - \lambda_1) \\ \gamma^{11}(\lambda_2 - \lambda_1) & \frac{\partial \lambda_2}{\partial s_{11}} \end{bmatrix}. \tag{B.1}$$

Then for the eigenvalues,

$$\frac{\partial \lambda_1}{\partial s_{11}} = q_{11}q_{11}, \frac{\partial \lambda_2}{\partial s_{11}} = q_{12}q_{12}.$$

For the eigenvectors, we have that the components of the antisymmetric matrix Γ_{11} are,

$$\gamma^{11} = \frac{q_{12}q_{11}}{\lambda_2 - \lambda_1}.$$

b) $i = 1$**,** $j = 2$

$$\begin{bmatrix} q_{11} & q_{21} \\ q_{12} & q_{22} \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix} = \begin{bmatrix} \frac{\partial \lambda_1}{\partial s_{12}} & \gamma^{12}(\lambda_2 - \lambda_1) \\ \gamma^{12}(\lambda_2 - \lambda_1) & \frac{\partial \lambda_2}{\partial s_{12}} \end{bmatrix}. \tag{B.2}$$

Then for the eigenvalues,

$$\frac{\partial \lambda_1}{\partial s_{12}} = q_{11}q_{21}, \frac{\partial \lambda_2}{\partial s_{12}} = q_{12}q_{22}.$$

For the eigenvectors, we have two solutions for the components of the antisymmetric matrix Γ_{12},

$$\gamma^{12} = \frac{q_{11}q_{22}}{\lambda_2 - \lambda_1}, \tag{B.3}$$

$$\gamma^{*12} = \frac{q_{12}q_{21}}{\lambda_2 - \lambda_1}. \tag{B.4}$$

c) $i = 2$**,** $j = 1$

$$\begin{bmatrix} q_{11} & q_{21} \\ q_{12} & q_{22} \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix} = \begin{bmatrix} \frac{\partial \lambda_1}{\partial s_{21}} & \gamma^{21}(\lambda_2 - \lambda_1) \\ \gamma^{21}(\lambda_2 - \lambda_1) & \frac{\partial \lambda_2}{\partial s_{21}} \end{bmatrix}. \tag{B.5}$$

Now for the derivatives of the eigenvalues, we have,

$$\frac{\partial \lambda_1}{\partial s_{21}} = q_{21}q_{11}, \frac{\partial \lambda_2}{\partial s_{21}} = q_{22}q_{12}.$$

And for the eigenvectors, there also exists two solutions for the components of the antisymmetric matrix Γ_{21},

$$\gamma^{21} = \frac{q_{11}q_{22}}{\lambda_2 - \lambda_1}, \tag{B.6}$$

$$\gamma^{*12} = \frac{q_{12}q_{21}}{\lambda_2 - \lambda_1}. \tag{B.7}$$

d) $i = 2$**,** $j = 2$

$$\begin{bmatrix} q_{11} & q_{21} \\ q_{12} & q_{22} \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix} = \begin{bmatrix} \frac{\partial \lambda_1}{\partial s_{22}} & \gamma^{22}(\lambda_2 - \lambda_1) \\ \gamma^{22}(\lambda_2 - \lambda_1) & \frac{\partial \lambda_2}{\partial s_{22}} \end{bmatrix}. \tag{B.8}$$

For the eigenvalues we have,

$$\frac{\partial \lambda_1}{\partial s_{22}} = q_{21}q_{21}, \frac{\partial \lambda_2}{\partial s_{22}} = q_{22}q_{22}.$$

For the eigenvectors, we have a unique solution for the components of the antisymmetric matrix Γ_{22},

$$\gamma^{22} = -\frac{q_{21}q_{22}}{\lambda_1 - \lambda_2}.$$

We define the derivatives of the eigenvalues with respect to the covariance matrix elements using 4.41,

$$\frac{\partial \lambda_k^{ij}}{\partial s_{ij}} = q_{ik}q_{jk} \quad \text{for } k = 1, 2,$$

producing the matrix defined by 4.42,

$$\frac{\partial \Lambda}{\partial s_{ij}} = L_{ij} = \begin{bmatrix} q_{i1}q_{j1} & 0 \\ 0 & q_{i2}q_{j2} \end{bmatrix}$$

For the derivatives of the eigenvectors with respect to the elements of the covariance matrix, we just arbitrarily choose one the solutions, taking care of being consistent considering the different solutions. Multiple solutions arise due to the symmetry of the problem.

We define in 4.40,

$$\gamma^{ij} := \frac{q_{i1}q_{j2}}{\lambda_2 - \lambda_1},$$

and according with that, we define the derivative of the eigenvector matrix respect to the elements of the covariance matrix using 4.39,

$$\frac{\partial Q}{\partial s_{ij}} = Q\Gamma_{ij} = \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix} \begin{bmatrix} 0 & \gamma^{ij} \\ -\gamma^{ij} & 0 \end{bmatrix}.$$

Note that when we write the Gâteaux derivative with respect to ϕ defined by 4.43, we have

$$\begin{aligned}
\left.\frac{\partial Q}{\partial \phi}\right|_{\tilde{\phi}} &= \sum_i \sum_j Q\Gamma_{ij} \left.\frac{\partial s_{ij}}{\partial \phi}\right|_{\tilde{\phi}} \\
&= Q\left(\Gamma_{11}\left.\frac{\partial \sigma_x^2(\phi)}{\partial \phi}\right|_{\tilde{\phi}} + (\Gamma_{12}+\Gamma_{21})\left.\frac{\partial \sigma_{xy}(\phi)}{\partial \phi}\right|_{\tilde{\phi}} + \Gamma_{22}\left.\frac{\partial \sigma_y^2(\phi)}{\partial \phi}\right|_{\tilde{\phi}}\right), \\
\left.\frac{\partial \Lambda}{\partial \phi}\right|_{\tilde{\phi}} &= \sum_i \sum_j L_{ij} \left.\frac{\partial s_{ij}}{\partial \phi}\right|_{\tilde{\phi}} \\
&= L_{11}\left.\frac{\partial \sigma_x^2(\phi)}{\partial \phi}\right|_{\tilde{\phi}} + (L_{12}+L_{21})\left.\frac{\partial \sigma_{xy}(\phi)}{\partial \phi}\right|_{\tilde{\phi}} + L_{22}\left.\frac{\partial \sigma_y^2(\phi)}{\partial \phi}\right|_{\tilde{\phi}},
\end{aligned}$$

the term $(\Gamma_{12} + \Gamma_{21})$, which was produced by the off-diagonal elements $i = 1, j = 2$ and $i = 2, j = 1$ are additive, which demonstrates that any of the two solutions of the linear system can be chosen, but they need to be different. We choose the solutions defined by B.3 and B.4, since they allow us for using a unique notation for all cases, which was expressed by 4.40. This framework can be easily extended to higher dimensions by considering a larger linear system.