

Incorporating shape variability in implicit template deformation.

Raphael Prevost, Remi Cuingnet, Benoit Mory, Laurent D. Cohen, and Roberto Ardon

Abstract In this chapter, we propose a method to learn and use prior knowledge on shape variability in the implicit template deformation framework. This shape prior is learnt via an original and dedicated process in which both an optimal template and principal modes of variations are estimated from a collection of shapes. This learning strategy does not require one-to-one correspondences between shape sample points and is not biased by a pre-alignment of the training shapes. We then generalize the implicit template deformation formulation to automatically select the most plausible deformation as a shape prior. This novel framework maintains the two main properties of implicit template deformation: topology preservation and computational efficiency. Our approach can be applied to any organ with a possibly complex shape but fixed topology. We validate our method on myocardium segmentation from cardiac magnetic resonance short-axis images and demonstrate segmentation improvement over standard template deformation.

Raphael Prevost, Remi Cuingnet, Benoit Mory, Roberto Ardon
Philips Research Medisys, Suresnes, France.
e-mail: {raphael.prevost,benoit.mory,remi.cuingnet,roberto.ardon}@philips.com

Laurent D. Cohen
CEREMADE UMR CNRS 7534, Paris Dauphine University, PSL, Paris, France.
e-mail: cohen@ceremade.dauphine.fr

1 Introduction

Implicit template deformation is a model-based segmentation framework that consists in deforming an initial template to segment an image. When one applies it to a particular clinical problem, the first step is to choose an adequate template. Indeed as we are working with diffeomorphisms, the template must have the same topology as the organ that we want to segment. But it also needs to be “close” to the target object, since the magnitude of the deformation is penalized. While it is always possible to build a synthetic template (*e.g.* a hyperquadric or superquadric for a ventricle [Cohen & Cohen, 1996, Bardinet et al., 1998], or an ellipsoid for the kidney in [Cuingnet et al., 2012]), one feels that this choice is probably suboptimal in other applications. The purpose of this chapter is to answer the following questions: how can we use a database to design an optimal (in some sense) template ? Can we learn the shape variability from this database so that we can take it into account within the deformation penalization ?

In this chapter, we present two variational approaches for training and segmentation, respectively. Both phases rely on the implicit template deformation framework [Saddi et al., 2007, Huang et al., 2004, Huang & Metaxas, 2008, Mory et al., 2012]. While this framework has been successfully applied in several contexts [Cuingnet et al., 2012, Prevost et al., 2012], its prior knowledge is limited to a single given shape. To introduce shape variability in the same framework, its segmentation functional is used as a measure of dissimilarity between shapes during a training step. This step estimates both a mean shape and a set of principal deformations through joint segmentation of all training shapes (Figure 1). The segmentation step then extends implicit template deformation framework by incorporating these computed statistics in the regularization term. It ensures preservation of the template topology and automatically selects the most plausible deformation as a shape prior, with very limited additional complexity.

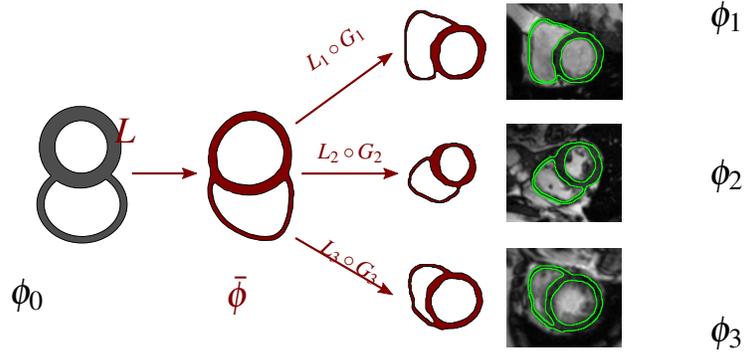


Fig. 1 Given an initial synthetic shape ϕ_0 , a set of shapes $\{\phi_n\}_n$ is simultaneously segmented via implicit template deformation while an intermediate mean shape $\bar{\phi} = \phi_0 \circ L$ is estimated. The topology of ϕ_0 is preserved during the process.

This chapter is organized as follows: Section 2 introduces the problem and lists some of the related work available in the literature. Section 3 introduces the main notations and recalls the implicit template deformation framework. In Section 4, we describe an original learning process that is tailored to the implicit template deformation framework. The learnt statistics will be used in a generalized formulation of the segmentation algorithm introduced in Section 5. Validation proving the benefits of our approach are provided in Section 6 in the application of myocardium segmentation in 2D MR images. Finally, discussion on potential improvements and conclusion conclude the chapter in Section 7 and 8.

A short version of this chapter was presented at the MICCAI 2013 conference [Prevost et al., 2013a].

2 Implicit Template Deformation

2.1 Motivation

Model-based methods are particularly effective and popular in medical image segmentation. Among them, implicit template deformation has recently been used in various applications [Saddi et al., 2007, Somphone et al., 2008, Mory et al., 2012, Prevost et al., 2012, Cuingnet et al., 2012] for its interesting properties (computational efficiency, topology preservation, compatibility with user interactions). This variational method consists in seeking a segmenting implicit function as a deformed implicit template. This template, acting as a shape prior, is therefore of paramount importance. However, in previous works the initial template was either set as a synthetic model (*e.g.* an ellipsoid for a kidney [Mory et al., 2012, Prevost et al., 2012, Cuingnet et al., 2012]) or as a segmented organ from a single arbitrary image [Saddi et al., 2007]. Despite the consensus that learning shape priors is a powerful approach to improve robustness [Cremers et al., 2007, Heimann & Meinzer, 2009], this has never been proposed in the context of segmentation by implicit template deformation. As public databases are developed and become available, it is important to think about how we can exploit them to validate but also improve our algorithms.

Here we aim at (i) learning statistics from a database of shapes (*i.e.* the most likely shape and the main variations around it) on organs that present a possibly complex shape but a consistent topology, (ii) proposing a method to exploit such learnt information within the segmentation framework of implicit template deformation. Naturally, our approaches have to maintain the interesting characteristics of implicit template deformation, namely the computational efficiency and the topology preservation. These two properties are usually incompatible but we are notwithstanding able to guarantee both by generalizing the formulation of the implicit template deformation. Combination of learnt shapes

will be used not to directly segment the images but rather within the regularization term. Thus they will act as a shape prior that is automatically updated during the segmentation.

2.2 Previous work and our shape learning Approach

The shape learning literature being considerably large, we point out here only well-known or closely related techniques.

In the early and popular active shape model [Cootes et al., 1995], objects are represented by an explicit parameterization of their boundary vertices. Statistics (mean shape and variations) are computed on these vertices coordinates, thus a suitable one-to-one vertices correspondence is needed across the database. This correspondence can be complicated to obtain: either tedious when relying on manually labeled points or lacking robustness when automatically obtained (e.g. [Besl & McKay, 1992]). Due to boundary self-intersections, shape topology may also be lost.

Implicit methods [Leventon et al., 2000, Rousson & Paragios, 2002, Tsai et al., 2003, Cremers et al., 2003] represent objects through the signed distance functions to their boundary to estimate statistics. Although vertices correspondence is no longer needed during the learning step, this representation is still inadequate for topology preservation. In [Rousson & Paragios, 2002], Rousson and Paragios built a probabilistic model in order to estimate a mean implicit function ϕ_m (and an associated variance at each pixel) from a dataset of signed distance functions. Unlike most previous works, they constrained ϕ_m to be a true distance function, i.e. $\|\nabla\phi_m\| = 1$, and not just any implicit function. Although more satisfying, this approach requires complex optimization schemes and the constraint is only enforced during the learning (and not the subsequent segmentations). Furthermore, it is still inadequate for topology preservation.

Finally, closely related with this paper, a third class of methods uses statistics on diffeomorphisms of implicit shape representations

[Arsigny et al., 2006, Vaillant et al., 2004] or on currents [Durrleman, 2010]. While they present elegant and appealing theoretical properties and do preserve shape topology, they are also much more computationally expensive. Most of their applications therefore lie in offline shape analysis and they were not used for segmentation purposes (apart from atlas-based segmentation methods [Khan et al., 2008], which are not compatible with real-time or user-interactions).

Here, we propose an approach that is closely related to this third class of methods, since implicit template deformation consists in seeking a space transformation. We thus introduce a dedicated learning approach by using the template deformation energy as a pre-metric in the shapes space. This idea was inspired by the seminal paper of Joshi *et al.*, in which they construct an unbiased mean template by minimizing a sum of distances to a database [Joshi et al., 2004]. As they were motivated by registration applications, they worked directly on images. When applied to shapes, this approach yields a co-segmentation process (sharing some ideas with [Yezzi & Soatto, 2003] to a certain extent), within which an optimal shape is estimated (see Figure 1). However, we go further than both papers by learning (and subsequently exploiting) also the variability of the shape around this mean. This means we also capture further information by building a space of main deformations around this template. Finally, we introduce a generalization of the template deformation formulation by using the computed statistics in the regularization term. The proposed framework is generic and can be applied to any organ with a possibly complex and variable shape but a fixed topology. We demonstrate its efficiency and interest by addressing the problem of myocardium segmentation in 2D cine-MR images.

3 Segmentation by implicit template deformation

Implicit template deformation [Saddi et al., 2007, Mory et al., 2012] is a variational framework for image segmentation. The segmentation is defined through the zero level-set of an implicit function $\phi : \Omega \rightarrow \mathbb{R}$, and ϕ is positive (resp. negative) inside (resp. outside) the segmentation. In this framework, the set of admissible segmentations \mathbb{S} is defined via an implicit template $\phi_0 : \Omega \rightarrow \mathbb{R}$ as the set of all implicit functions with the same topology as ϕ_0 , i.e. $\mathbb{S} = \{\phi : \Omega \rightarrow \mathbb{R} \text{ s.t. } \phi = \phi_0 \circ \psi, \psi \text{ is diffeomorphic}\}$. The unknown is thus the transformation $\psi : \Omega \rightarrow \Omega$ which is sought as a minimum of a region competition energy:

$$\min_{\psi} \left\{ \int_{\Omega} H(\phi_0 \circ \psi) r_{int} + \int_{\Omega} (1 - H(\phi_0 \circ \psi)) r_{ext} + \lambda \mathcal{R}(\psi) \right\}, \quad (1)$$

where H denotes the Heaviside function ($H(a) = 1$ if $a > 0$, 0 otherwise) while r_{int} and r_{ext} are image-based functions such as $r_{int}(\mathbf{x})$ is lower (resp. higher) than $r_{ext}(\mathbf{x})$ if voxel \mathbf{x} seems to belong to the target object (resp. background). $\mathcal{R}(\psi)$ is a constraint term on ψ that prevents the segmentation $\phi = \phi_0 \circ \psi$ to deviate too much from the initial template ϕ_0 ; it is weighted by a parameter λ . In [Mory et al., 2012], ψ is decomposed into (i) a global transformation $\mathcal{G} \in \mathbb{G}$ (e.g. a similarity) accounting for the pose of the template in the image, and (ii) a diffeomorphism $\mathcal{L} \in \mathbb{D}(\Omega)$ that yields local deformation and does change the shape of the template. This decomposition allows to define the regularization as a function of the deformation only $\mathcal{R}(\psi) = \mathcal{R}(\mathcal{L}) = \frac{1}{2} \|\mathcal{L} - Id\|_2^2$. The problem finally reads

$$\min_{\mathcal{L}, \mathcal{G}} \left\{ \int_{\Omega} H(\phi_0 \circ \mathcal{L} \circ \mathcal{G}) r_{int} + \int_{\Omega} (1 - H(\phi_0 \circ \mathcal{L} \circ \mathcal{G})) r_{ext} + \frac{\lambda}{2} \|\mathcal{L} - Id\|_2^2 \right\}. \quad (2)$$

In such a setting, ϕ_0 not only fixes the topology of the segmentation but also acts as a shape prior, which makes its choice of paramount

importance. Moreover, the term \mathcal{R} could be improved by taking into account shape variability of the considered organ. In the next sections, we develop a framework to tackle both problems by estimating statistics on a collection of shapes.

4 A learning process dedicated to template deformation

Given a training set of variables $(X_n)_{n=1\dots N} \subset \mathbb{S}^N$, one can define its *mean* (more precisely its Fréchet-mean or Karcher-mean [Karcher, 1977]) as the solution of the following problem:

$$\bar{X} = \operatorname{argmin}_{X \in \mathbb{S}} \frac{1}{N} \sum_{n=1}^N d^2(X, X_n) \quad (3)$$

This definition therefore depends both on the space \mathbb{S} that is used to represent shapes and the distance d that defines their similarity. We will use these two degrees of freedom to build a mean shape that is tailored for the implicit template deformation.

4.1 A dedicated estimation of a mean model

The first choice concerns the space of shapes \mathbb{S} . Shapes can be represented by different types of objects (i.e. vertices [Cootes et al., 1995], implicit functions [Tsai et al., 2003], deformations [Joshi et al., 2004], currents [Durrleman, 2010], etc.). Our goal here is to estimate a model for the implicit template deformation framework, so we will choose an implicit representation. However, we would like to specify and fix the topology of the considered shapes. This information will be given by an initial implicit template ϕ_0 , on which the space of admissible shapes will depend:

$$\mathbb{S}_{\phi_0} = \{\phi : \Omega \rightarrow \mathbb{R} \text{ such that } \phi = \phi_0 \circ L \text{ with } L \in \mathbb{D}(\Omega)\} \quad (4)$$

which can be thought of as the *orbit* of ϕ_0 in the set of shapes. Note that such a space is stable under any diffeomorphism. Its dependency on ϕ_0 is rather low (except the topology) as it is identical to any \mathbb{S}_ϕ such that $\phi \in \mathbb{S}_{\phi_0}$. For the sake of simplicity, in the following we will omit the index and denote this space \mathbb{S} .

In order to estimate statistics in \mathbb{S} , we then define a metric-like function in this space which should be related to our segmentation framework. To do so, we point out that any shape $\phi_1 \in \mathbb{S}$ can be warped to another shape $\phi_2 \in \mathbb{S}$ via implicit template deformation. Indeed, we can segment an image representing ϕ_2 using ϕ_1 as template. With the notations of (2), we simply have to set $\phi_0 := \phi_1$, $r_{int}^{\phi_2} := \max(-\phi_2, 0)$ and $r_{ext}^{\phi_2} := \max(\phi_2, 0)$. The definition of $r_{int}^{\phi_2}$ and $r_{ext}^{\phi_2}$ is not unique and we could have selected other functions that represents the interior and the exterior of ϕ_2 . The rationale behind this particular choice is that the difference $r_{int}^{\phi_2} - r_{ext}^{\phi_2}$ is then equal to $-\phi_2$. This leads to a tailored definition of shape dissimilarity d^2 .

Definition 1. The shape dissimilarity from shape ϕ_1 to shape ϕ_2 is defined as

$$\begin{aligned} d^2(\phi_1, \phi_2) = & \min_{\substack{L \in \mathbb{D}(\Omega) \\ G \in \mathbb{G}}} \int_{\Omega} H(\phi_1 \circ L \circ G) \max(-\phi_2, 0) \\ & + \int_{\Omega} (1 - H(\phi_1 \circ L \circ G)) \max(\phi_2, 0) \\ & + \frac{\lambda}{2} \|L - Id\|_U^2, \end{aligned} \quad (5)$$

or equivalently

$$d^2(\phi_1, \phi_2) = C(\phi_2) + \min_{\substack{L \in \mathbb{D}(\Omega) \\ \mathcal{G} \in \mathbb{G}}} - \int_{\Omega} H(\phi_1 \circ L \circ \mathcal{G}) \phi_2 + \frac{\lambda}{2} \|L - Id\|_U^2, \quad (6)$$

where $C(\phi_2)$ is a constant that only depends on ϕ_2 .

In this definition, the U -norm represents the natural norm in the Gaussian reproducing kernel Hilbert space (see [Aronszajn, 1950] for more details), which can usually be as well the L^2 norm.

Remark 1 *The shape constraint parameter λ should be chosen carefully, since a too high value will prevent ϕ_1 to be exactly matched to ϕ_2 and the learning will be biased. In practice however, it is not difficult to find a suitable value.*

This shape dissimilarity measure is not a distance but a pre-metric, as it is not symmetric and does not verify triangular inequality. The lack of symmetry is directly inherited from the segmentation process itself as the template ϕ_0 has a very particular role. Triangular inequality does not either appear as an important property for our application. Cremers *et al.* discussed these properties in [Cremers & Soatto, 2003] and point out that defining a true distance between implicit shapes is still an open problem. But anyway, this function does measure a closeness between two shapes and we can still use it to define our dedicated notion of mean shape.

Definition 2. An implicit function $\bar{\phi}$ is a mean of the set $\{\phi_n\}_{n=1..N}$ (in the sense of implicit template deformation) if it is a local minimum of the shape dissimilarity to all the elements of this set, i.e.

$$\bar{\phi} = \operatorname{argmin}_{\phi \in \mathbb{S}} \frac{1}{N} \sum_{n=1}^N d^2(\phi, \phi_n). \quad (7)$$

It is important to note that we seek the mean shape $\bar{\phi}$ as an element of \mathbb{S} . Indeed, in our application the mean shape has to preserve the topology

of the training shapes. This means that there exists $L \in \mathbb{D}(\Omega)$ such that $\bar{\phi} = \phi_0 \circ L$. The mean shape expression can thus be reformulated as

$$\bar{\phi} = \phi_0 \circ \left\{ \operatorname{argmin}_{L \in \mathbb{D}(\Omega)} \frac{1}{N} \sum_{n=1}^N d^2(\phi_0 \circ L, \phi_n) \right\}. \quad (8)$$

Expanding the segmentation costs and neglecting constant terms in Equation (8) yields the following optimization problem to solve:

$$\min_{\substack{L \in \mathbb{D}(\Omega) \\ (L_n)_n \in \mathbb{D}(\Omega)^N \\ (G_n)_n \in \mathbb{G}^N}} E_{learn} = - \sum_{n=1}^N \int_{\Omega} H(\phi_0 \circ L \circ L_n \circ G_n) \phi_n + \sum_{n=1}^N \frac{\lambda}{2} \|L_n - \mathbf{Id}\|_U^2. \quad (9)$$

This can be interpreted as segmenting simultaneously all training shapes $(\phi_n)_n$ starting from ϕ_0 while estimating an optimal common intermediate shape $\phi_0 \circ L$ (see Figure 1). In (9), the energy E_{learn} is minimized with respect to three kinds of variables

- the global transformations $(G_n)_n$, called the **poses**, that register all shapes to ϕ_0 with translation, rotation and scaling. As they are part of the optimization process, they do not bias the learning, as a fixed pre-alignment (*e.g.* [Rousson & Paragios, 2002, Tsai et al., 2003]) would do.
- the **common deformation** L , which includes the common parts of the deformations from ϕ_0 to all the training shapes;
- the local deformations $(L_n)_n$, called the **residual deformations**, are the residual components of the deformations from $\phi_0 \circ L$ to ϕ_n . Unlike L , their magnitude is penalized so that any deformation which is common to all the training set will be preferably included in L .

The optimal common deformation L^* can be used to define the optimal shape (in the sense of the segmentation algorithm) as $\bar{\phi} = \phi_0 \circ L^*$.

This shape globally minimizes the magnitude of residual deformations to each shape of the dataset. Note that the magnitude of L is not penalized so the choice of ϕ_0 defines the topology of $\bar{\phi}$ but does not affect it further (modulo the smoothness enforced to L). In our experiments, running a second time the learning process with $\phi_0 := \bar{\phi}$ did not alter the results.

Remark 2 *We assumed that the set of training shapes $\{\phi_n\}_n$ was a subset of \mathbb{S}_{ϕ_0} . It is a very natural hypothesis and ϕ_0 should be chosen accordingly. However, those training shapes will probably come from manual annotations of images and, as such, be prone to errors. As a consequence, it may occur that some training shapes do not have the correct topology. This does not question the soundness of our learning because, in such cases, we will implicitly learn the “closest shapes” with the topology of ϕ_0 .*

Details on the resolution of (9) are provided in the next subsection.

4.2 Numerical optimization

Problem (9) presents some mathematical similarities with the co-segmentation method proposed in [Prevost et al., 2013b, Prevost et al., 2013c]. It is therefore minimized similarly, with a gradient descent simultaneously on each of the unknowns. The gradient directions with respect to $\mathbf{p}_{n,i}$ (the i -th parameter of G_n), the common deformation L and the residual deformations L_n are given by the following equations¹:

¹ They are obtained with standard calculus of variation, but we omit the tedious details here.

$$\nabla_{\mathbf{p}_{n,i}} E_{learn} = \int_{\Omega_0} \delta(\phi_0 \circ L \circ L_n) \phi_n \circ G_n^{-1} \left| J_{G_n}^{-1} \right| \quad (10)$$

$$\left\langle \nabla \phi_0 \circ L \circ L_n, DL \circ L_n \cdot DL_n \circ \frac{\partial G_n}{\partial \mathbf{p}_{n,i}} \circ G_n^{-1} \right\rangle \quad (11)$$

$$\nabla_L E_{learn} = K_\sigma * \left[\sum_{n=1}^N \delta(\phi_0 \circ L) \phi_n \circ G_n^{-1} \circ L_n^{-1} \cdot \left| J_{L_n \circ G_n}^{-1} \right| \cdot DL \cdot \nabla \phi_0 \circ L \right] \quad (12)$$

$$\nabla_{L_n} E_{learn} = K_\sigma * \left[\delta(\phi_0 \circ L \circ L_n) \phi_n \circ G_n^{-1} \cdot \left| J_{G_n}^{-1} \right| \cdot DL_n \cdot \nabla \phi_0 \circ L \circ L_n \right] + \lambda (L_n - \mathbf{Id}) \quad (13)$$

The first gradient – with respect to the poses – is used in a standard gradient procedure, while the two others – with respect to the common and residual deformations – are exploited in a topology-preserving optimization scheme since the space is not stable under linear combinations. The appropriate way is to combine diffeomorphisms via composition since $(\mathbb{D}(\Omega), \circ)$ is a group. Following [Saddi et al., 2007], we therefore update any diffeomorphism \mathcal{L} in the following way:

$$\mathcal{L}^{(n+1)} \leftarrow (Id - \Delta t \nabla_{\mathcal{L}} E) \circ \mathcal{L}^{(n)}. \quad (14)$$

The regularity is enforced by a Gaussian filtering of the gradient as in [Mory et al., 2012].

All these integrands actually have a very small support (the zero level-set of an implicit function), which makes the computations fast. Moreover, the three kinds of gradients have a lot of terms in common.

Remark 3 *Some terms depends on inverses of diffeomorphisms L_n^{-1} . These deformations are built iteratively and simultaneously to the direct transformations L_n .*

Remark 4 *Although the computations needed for the training process are relatively fast, there is a high memory requirement (especially in 3D) since a high number of implicit functions and deformation fields have to be stored simultaneously. A possible solution would be to use a stochastic gradient descent [Bottou, 1998], i.e. at each iteration only consider a randomly chosen subset of the training set.*

4.3 Building a space of deformation priors

As seen in the previous subsection, minimization of (9) yields a mean shape. However, the optimal residual deformations $(L_n^*)_n$ are also available and can be used to capture further information on the variability of the training shapes.

We build a space of principal deformations \mathbb{L} to constrain future segmentation of new images. Similarly to [Rueckert et al., 2001], a principal component analysis (PCA) [Jolliffe, 1986] is applied to the residual deformations to find a suitable parametrization of such a space. The goal of this analysis is to find a reduced number of orthogonal vectors that maximize the explained variance of the residual deformations. This is accomplished by first computing the mean residual deformation

$$\bar{\ell} = \frac{1}{N} \sum_{n=1}^N L_n^* \quad (15)$$

and then performing a singular value decomposition (SVD) of the sample covariance matrix

$$S = \frac{1}{N-1} \sum_{n=1}^N (L_n^* - \bar{\ell})(L_n^* - \bar{\ell})^T . \quad (16)$$

Any deformation ℓ in agreement with the variability of the training data can then be approximated by a linear combination of the offset $\bar{\ell}$ and $(\ell_k)_{k=1..M}$ the first M singular vectors of S . It is parametrized by the vectors of its weights $w \in \mathbb{R}^M$:

$$\ell \approx \ell[w] = \bar{\ell} + \sum_{k=1}^M w_k \ell_k . \quad (17)$$

We denote \mathbb{L} the set of such transformations. M can be empirically chosen using the distribution of the modes' eigenvalues. To each singular

vector ℓ_k corresponds a singular value λ_k that represents the amount of variance of the residual deformations that is explained with this mode of variation.

Remark 5 *Note that even if the PCA is applied to the residual deformations, $\bar{\ell}$ is non-null (though with a very small magnitude) because it denotes a mean with respect to a different metric than L^* .*

The space of diffeomorphisms is not stable under linear combinations. There is therefore no guarantee that an element of \mathbb{L} is actually a diffeomorphism. Nevertheless, as shown in the next section, it is possible to use this space indirectly in a topology-preserving segmentation framework.

5 Generalized implicit template deformation

The previously estimated statistical information can be used to robustify and improve future segmentations. In order to incorporate such information in the segmentation process, we propose a generalization of the implicit template deformation framework.

5.1 An improved formulation for segmentation

A first improvement is achieved by replacing the original template ϕ_0 by the mean template $\bar{\phi} = \phi_0 \circ L^*$. Secondly, the estimation of the deformation can also be enhanced by using the space of principal deformations \mathbb{L} . In most previous works [Cootes et al., 1995, Leventon et al., 2000, Tsai et al., 2003], the learnt variable is expressed as a linear combination of modes. When dealing with deformations, this does not guarantee topological preservation. Therefore, we rather use such linear combinations indirectly. More specifically, we modify the regularization term so

that the diffeomorphism L is constrained with respect to the set \mathbb{L} instead of the identity (see Figure 2). Thus, only deformations that cannot be explained through the learnt space \mathbb{L} are penalized. The new segmentation energy therefore reads

$$E_{seg}(L, G, w) = \int_{\Omega} H(\bar{\phi} \circ L \circ G) r_{int} + (1 - H(\bar{\phi} \circ L \circ G)) r_{ext} + \frac{\lambda}{2} \|L - \ell[w]\|_U^2. \quad (18)$$

This represents a generalization of the standard template deformation formulation. The novel regularization term can be interpreted as a shape prior that depends on the image. Thus, even if the target organ has a high variability around the mean, we can learn it in order to automatically select the most plausible shape that is implicitly used to constrain the segmentation.

A related approach was proposed in [Rousson et al., 2004] with implicit functions. In this paper, the authors defined the regularization term as a distance between the segmenting implicit function and a linear combination of implicit modes previously learnt. However our method presents a major advantage over theirs. Indeed, as explained above in section 4.2, using Eqn 14 we are able to let the deformation L evolve while preserving its diffeomorphic properties (and therefore maintaining the topology of the template ϕ_0). Conversely, it is not possible to easily enforce such a constraint into the evolution of an implicit function.

5.2 Numerical optimization

Minimization of (18) can be performed with a two-step alternate scheme:

Update of the segmentation: With $\ell[w]$ fixed, the energy is minimized through a gradient descent-like scheme on L and G (see [Mory et al., 2011, Mory et al., 2012]).

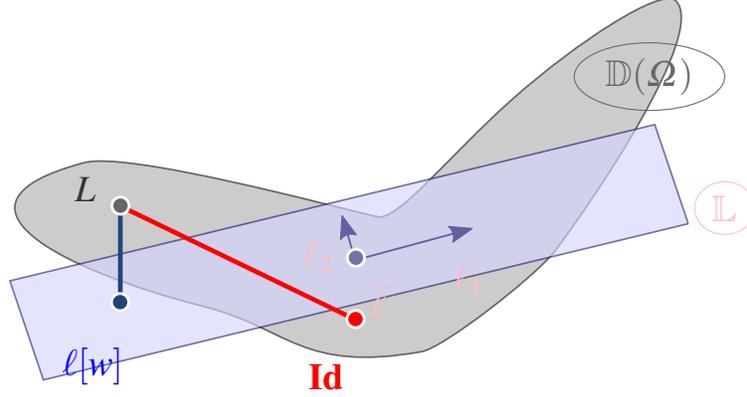


Fig. 2 Comparison of the penalization of the deformation L with the standard regularization term towards the Identity (red) and the novel term towards the space \mathbb{L} (blue) that is centered in $\bar{\ell}$ and spanned by the modes $(\ell_k)_k$. The new prior $\ell[w]$ is the projection of L onto the set \mathbb{L} . Note that all transformations of \mathbb{L} are not diffeomorphisms but L is constrained to be one.

Update of the shape prior: With L and G fixed, the update of $\ell[w]$ can be seen as a projection of L onto \mathbb{L} . Indeed the energy comes down to a simple quadratic function, whose minimizers are obtained by solving a simple linear system, as stated by the following proposition.

Proposition 1. *The minimum of $E_{seg}(L, G, \cdot)$ is reached at $w^* \in \mathbb{R}^M$ such that*

$$Aw^* = b_L \quad (19)$$

where A is a $M \times M$ matrix whose entries are $(A_{ij}) = \langle \ell_i, \ell_j \rangle_U$ while $b_L \in \mathbb{R}^M$ is defined by $(b_{L,i}) = \langle L - \bar{\ell}, \ell_i \rangle_U$.

Proof. With L and G fixed, the minimization problem comes down to

$$\operatorname{argmin}_{w \in \mathbb{R}^M} E_{seg}(L, G, w) = \operatorname{argmin}_{w \in \mathbb{R}^M} \|L - \bar{\ell} - \sum_{k=1}^M w_k \ell_k\|_U^2. \quad (20)$$

Setting to zero its derivative with respect to the weight of the mode k_0 yields

$$0 = \frac{d}{dw_{k_0}} \left\langle L - \bar{\ell} - \sum_{k=1}^M w_k \ell_k, L - \bar{\ell} - \sum_{k=1}^M w_k \ell_k \right\rangle_U (w^*), \quad (21)$$

$$0 = \left\langle L - \bar{\ell} - \sum_{k=1}^M w_k^* \ell_k, \ell_{k_0} \right\rangle_U, \quad (22)$$

$$\sum_{k=1}^M w_k^* \langle \ell_k, \ell_{k_0} \rangle_U = \langle L - \bar{\ell}, \ell_{k_0} \rangle_U. \quad (23)$$

Each k_0 yields a linear equation in w , hence the result.

Note that the matrix A has a quite small size, so this system is very easy to solve. Actually, we can even pre-compute the inverse of A since it only involves learnt variables.

However we may simplify this solution further by making some hypotheses. Recall that by construction via the PCA, the $(\ell_k)_k$ are L^2 -orthogonal. If we assume that they are also nearly U -orthogonal, then the matrix M is diagonal and the solutions are given by:

$$\forall k \in \{1, \dots, M\}, w_k^* = \frac{\langle L - \bar{\ell}, \ell_k \rangle_U}{\langle \ell_k, \ell_k \rangle_U}. \quad (24)$$

The values of w_k^* are subsequently clipped in the interval $[-3\sqrt{\lambda_k}; 3\sqrt{\lambda_k}]$, as these bounds represent the lengths of the semi-axes of the ellipsoid

estimated by the principal component analysis. Any deformation obtained with weights beyond this interval are not in agreement with the training set and thus should not be considered as possible priors. Other possibilities of computing these weights will be mentioned in the Discussion section of this chapter.

To sum up, the first step is similar to the standard implicit template deformation formulation, and the second one is straightforward. Therefore, the computational efficiency of the method is maintained while topology preservation is still guaranteed.

6 Application: Myocardium segmentation in 2D MR images

We validated our method in the context of myocardium analysis and segmentation in cardiac short-axis 2D cine-MR images. Quantitative assessment on the heart muscle is critical for diagnosis or therapy planning. This task is particularly challenging for model-based approaches because of the complex topology of the target object, i.e. a band around left and right ventricles.

6.1 Material

Our dataset is composed of 245 MR images coming from 61 different patients (for each case, several slices in the z-direction are available). The acquisitions have been synchronised so that each heart is in the same cardiac phase. The typical images size was 256×256 with resolution 1.56×1.56 mm. In every image, a myocardium segmentation has been manually performed by a radiologist. Based on the geometric information, we set for our method the scale of the deformation field σ to 10 mm. The initial synthetic template ϕ_0 used is shown in Figure

1. Global transformations are sought in the set \mathbb{G} of similarities (accounting for translation, rotation and isotropic scaling). The dataset was randomly split into a training set including 120 images from 30 patients and a testing set composed of the remaining 125 images coming from 31 patients.

6.2 Experiments on the learnt information

6.2.1 Synthetic experiments

First we conducted controlled experiments to assess quantitatively the estimation of the mean model. Random myocardium shapes were generated by applying random deformation fields to an original myocardium. We aim at recovering this original shape by estimating a mean model from subsets of these synthetically generated shapes. The efficiency of a learning process is evaluated by computing the Dice coefficient between the ground truth and the estimated mean shape. To avoid randomness bias, the experiments have been performed 100 times and the results averaged.

We reported in Figure 3 a comparison of three fully automatic methods using this metric: the implicit shape model proposed in [Tsai et al., 2003], the active shape framework [Cootes et al., 1995] with point correspondences estimated by ICP [Besl & McKay, 1992] and the proposed method. For any number of used shapes, our method provided statistically significantly better estimates of the original shape than the two others. These results can be better understood with Figure 4 showing the spatial localization of the errors. Indeed the implicit method fails to recover the entire muscle around the right ventricle: working directly on signed distance function is not adapted to thin structures. This area also causes high errors for the explicit method, which retrieves but underestimates

this part of the band. Conversely, errors for our method are lower and more evenly distributed.

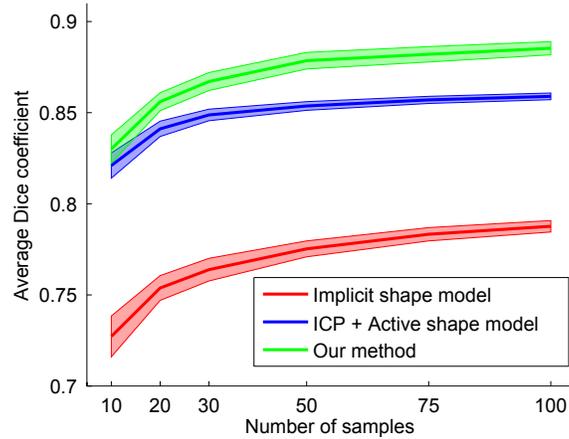


Fig. 3 Dice coefficients (averaged over the 100 experiments) between the estimated mean model and the original model as a function of number of samples using implicit shape model [Tsai et al., 2003], ICP [Besl & McKay, 1992] + active shape model [Cootes et al., 1995] and the proposed method. Bands around the curves delineate the 95% confidence interval.

6.2.2 Mean model and principal modes

We now provide a qualitative comparison between the different approaches on learnt information from clinical data. The initial dataset was randomly split into a training set including 120 images from 30 patients and a testing set composed of 125 images coming from the remaining

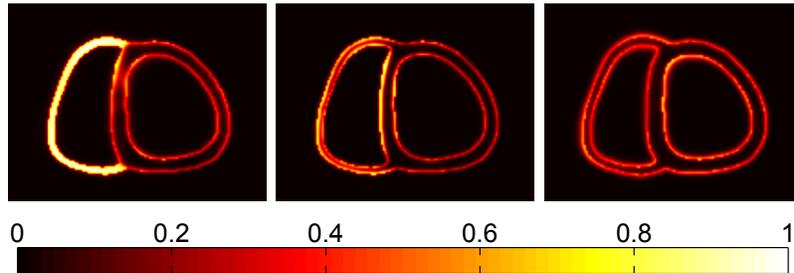


Fig. 4 Repartition of errors on the estimated model using implicit shape model [Tsai et al., 2003] (left), ICP [Besl & McKay, 1992] + active shape model [Cootes et al., 1995] (middle) and the proposed method (right). Color indicates the pixelwise empirical probability of bad classification (inside vs outside the shape).

31 patients.

The mean shape and first two modes of variation are shown in Figure 5 for each method. As expected from the results of previous subsection, the implicit method fails at recovering the true topology of the mean shape, but also with the first modes of variations. The explicit method performs better and provides a reasonable mean model. However, the modes of variation are not satisfying because very irregular and difficult to interpret. On the other hand, the results obtained with our method show a much better preservation of the topology and seem more realistic. Moving along the first principal components changes the relative size of the ventricles. This corresponds to the variability observed when moving on the axial direction of a given heart. This variation was expected because the training set include several slices of each heart. The second principal component controls the global anisotropic scaling of the hearts, which seems to rather represent an inter-subject variability.

Such variations were not taken into account by the global poses because their scaling were isotropic.

6.3 Validation of the improved segmentation

We finally evaluate how learnt information improves segmentation via implicit template deformation of unseen images. Myocardiums have been segmented in test images using (i) the synthetic model ϕ_0 as template, (ii) the estimated mean model $\bar{\phi}$ as template, (iii) the new deformation model-based regularization term in addition to the mean model $\bar{\phi}$ (with 5 modes).

The image-based classification functions r_{int} and r_{ext} were set to

$$r_{int}(\mathbf{x}) = -\log(p_{int}(\mathbf{x})) \quad \text{and} \quad r_{ext}(\mathbf{x}) = -\log(p_{ext}(\mathbf{x})). \quad (25)$$

where p_{int} and p_{ext} are of intensity probability distributions inside and outside the myocardium (estimated from the training datasets). For the intensities to be comparable, all the images were normalized beforehand.

Performance of each algorithm is quantified using Dice coefficients between the segmentation and the expert ground truth. Results on the whole testing set are summarized in Figure 6. Changing the template from ϕ_0 to the learnt $\bar{\phi}$ makes the algorithm more robust as the minimum Dice coefficient greatly increases (from 0.46 to 0.69). Modifying the regularization term by taking into account the deformation model further raises it 0.86. The proposed method globally enhances the algorithm on most images of the test database as the median goes from 0.85 for the baseline method to 0.92 with our modifications. These improvements are statistically significant with a p -value $< 10^{-4}$ for a Wilcoxon signed-rank test [Wilcoxon, 1945].

Figure 7 shows segmentation results in three different cases, for the classical regularization term with two values of the shape constraint parameter $\lambda \in \{1, 2\}$ and the new model-based regularization term. In all settings, the template was the mean model $\bar{\phi}$. Consider Case #1 (first row). Since the image term is reliable, a satisfying result is obtained with a small shape constraint. However, the myocardium deviates signif-

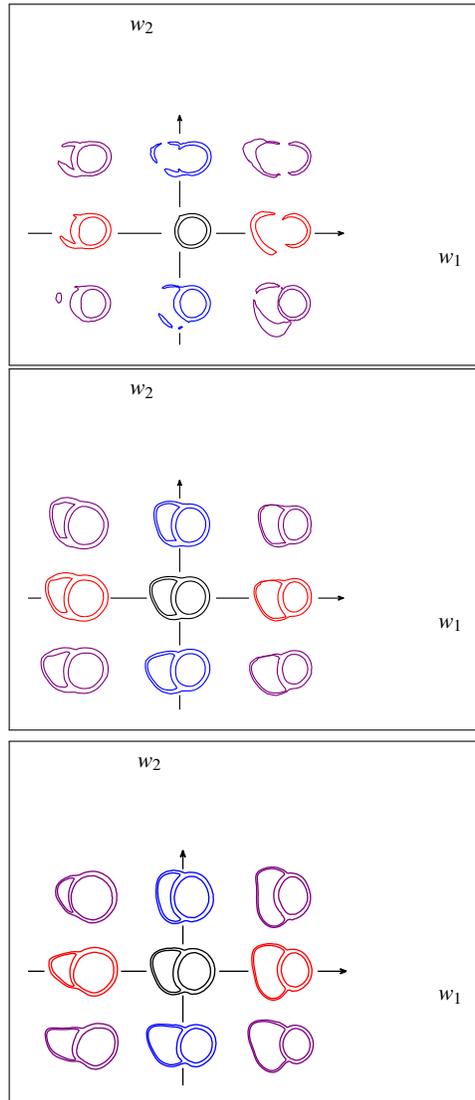


Fig. 5 Mean model and first two modes of the variation of the myocardium learnt on the training dataset using implicit shape model [Tsai et al., 2003] (left), ICP [Besl & McKay, 1992] + active shape model [Cootes et al., 1995] (right) and the proposed method (bottom). For our approach, the visualized shapes are the zero level-sets of $\bar{\phi} \circ (\bar{\ell} + w_1 \ell_1 + w_2 \ell_2)$.

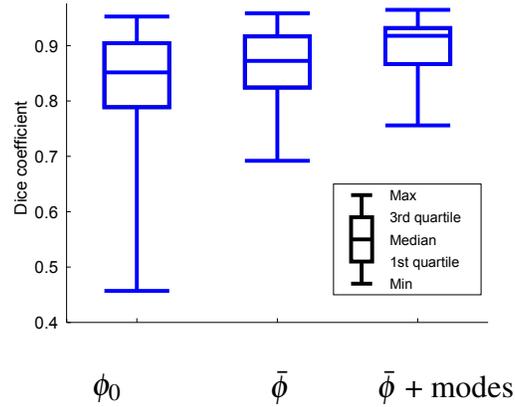


Fig. 6 Boxplot of the Dice coefficients for myocardium segmentation in MR images via implicit template deformation with synthetic model ϕ_0 (left), mean model $\bar{\phi}$ (middle), mean model $\bar{\phi}$ and deformation modes (right).

icantly from the mean shape: using a too strong constraint ($\lambda = 2$) prevents the algorithm to converge towards the right solution. Conversely in Case #2, the image information is much more ambiguous. This provokes some leaks (*e.g.* in papillary muscles of the left ventricle) with $\lambda = 1$, which shows there is no fixed value that allows a good segmentation in the first two cases. Yet by introducing the new regularization (fourth column), likely deformations are not penalized. This widens the capture range while still maintaining a high constraint on the shape and therefore avoiding unrealistic leaks. Finally, Case #3 illustrates that our method may also improve the result even if no λ was originally successful.

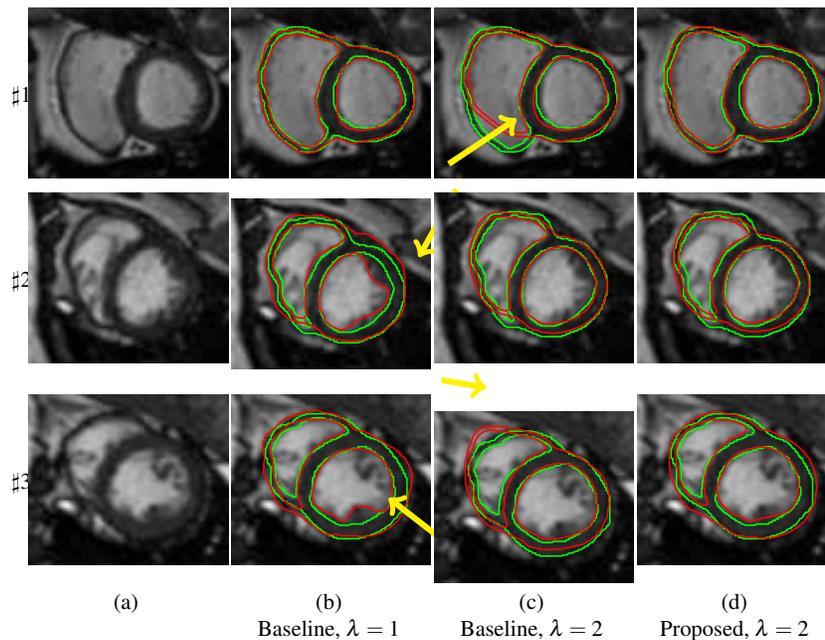


Fig. 7 Segmentation results (red) of different cases versus ground truths (green). Main failures are highlighted by yellow arrows. (a) Original images, (b,c) Standard method with small ($\lambda = 1$) and high ($\lambda = 2$) shape constraint, (d) Proposed method.

7 Discussion

We discuss herebelow some limitations of the current work and propose several ideas for future work.

From 2D to 3D: Although we first proved the potential of our approach on a 2D application, it should be noted that the whole method can be directly extended to 3D shapes, thanks to the implicit representation of shapes. Another advantage of our approach is that it does not require point-to-point correspondences between shapes, which can be

particularly challenging to obtain for three-dimensional shapes. We are thus currently investigating 3D applications, such as learning the shape variability of the liver.

Penalizing the weights of the modes: In this chapter, we proposed to replace the regularization term on the deformation L from

$$\|L - Id\|^2 \quad \text{to} \quad \min_w \|L - \ell[w]\|^2 \quad \text{with } \ell[w] \in \mathbb{L} \quad (26)$$

which basically consists in removing the penalization on all deformations in the affine space \mathbb{L} . The rationale was that \mathbb{L} is composed of deformations in agreement with the database. However the weights w should be not too large: with a Gaussian assumption, the training shapes are supposed to have their weight w_k in $[-3\sqrt{\lambda_k}; 3\sqrt{\lambda_k}]$. Our solution was to clip each weight into this interval, but there is a more elegant approach. We could decompose the original term $\|L - Id\|_U^2$ into

$$\min_{w \in \mathbb{R}^M} \|L - \ell[w]\|^2 + \|\ell[w] - \bar{\ell}\|^2 + \|\bar{\ell} - Id\|^2 \quad (27)$$

where the first term is the same as in (26). The third term is constant and can therefore be discarded in the minimization. Finally, the second term penalizes the distance between $\ell[w]$ and $\bar{\ell}$ that are both deformations in \mathbb{L} . Instead of using the standard L^2 - or U -norms, we can exploit the assumption of Gaussian distribution around $\bar{\ell}$ and define a Mahalanobis distance in \mathbb{L} as follows:

$$\|\ell[w] - \bar{\ell}\|^2 = \nu \sum_{k=1}^M \frac{w_k^2}{\lambda_k} \quad (28)$$

where ν is a normalization factor. This simply means that modes with high variance are less penalized. Minimization remains easy as (24) becomes

$$\forall k \in \{1, \dots, M\}, w_k^* = \frac{\langle L - \bar{\ell}, \ell_k \rangle_U}{\langle \ell_k, \ell_k \rangle_U + \nu \frac{1}{\lambda_k}} \quad (29)$$

The weights of the modes are slightly shifted down, according to the variance of the corresponding mode.

Choice of the dimension reductions method: In order to build the space \mathbb{L} , we used as a dimension reduction approach the principal component analysis (PCA) since it is a standard and easy-to-implement method. However we might try other methods such as independent component analysis (ICA) [Comon, 1994] which, unlike PCA, does not aim at capturing the largest variance with orthogonal vectors but with decorrelated ones.

To understand the potential benefits of ICA over PCA, consider the following simple “thought experiment”: in all the training shapes, only two disjoint zones Ω_A and Ω_B differ from the mean model. These zones are the same in each shape but they vary independently from each other. PCA will capture all the changes in a single mode, while ICA needs two modes (one for each region). Now let us imagine the segmentation of a corrupted image in which there is no information in Ω_A . With both methods, the weights of the shape prior will solely be determined by the information available in Ω_B . However, since the two regions are in the same PCA mode, *the shape prior will also change in Ω_A* if we use the PCA approach. This is clearly not desirable since the two variations were independent: there is no reason for Ω_A to influence Ω_B . The correct behaviour (i.e. the prior in Ω_A should be only the mean model *unless there exists some statistical correlation with other regions*) is obtained with the ICA modes.

Learning on logarithms: We mentioned earlier that diffeomorphisms are not stable under linear combinations. This was the reason why we did not use directly the modes of the PCA in the segmentation step. However, we may question the learning step itself: is it really sensible to perform a PCA in a space that is not linear ? The answer is probably negative from a purely theoretical point of view.

One possible solution would be to perform the PCA on the logarithms of the residual deformations $(L_i^*)_i$, since the logarithm of a diffeomor-

phism is a vector field that does lie in a linear space [Arsigny et al., 2006]. Preliminary experiments that we are currently conducting however suggest that both learning and segmentation results do not significantly change when working on logarithms, whereas the efficiency of the algorithm is lost.

8 Conclusion

In this chapter, we have presented an approach to include organ shape variability in the implicit template deformation framework. By computing a mean over a database of shapes defined with a dedicated distance, we constructed a shape template that is tailored to our algorithm. We even used further statistics by estimating (and then exploiting for the segmentation of unseen images) the main modes of variations of the deformations. The remarkable properties of this approach are its computational efficiency and the topology preservation of the initial model. A variational approach was proposed to extract statistical information (mean and principal variations) from a collection of shapes. This training method is automatic, does not require landmarks correspondance and relies upon a definition of shape dissimilarity that is directly derived from the implicit template deformation functional. We also proposed a generalization of the original segmentation algorithm in which the shape prior is automatically adapted to the current image during the deformation process while still maintaining both the computational efficiency and the topology preservation of the method (segmentation takes around one second on a standard computer). Quantitative results demonstrated the improvement over implicit template deformation for a 2D application. Our approach is very generic and can be used to segment any object with a complex shape but a fixed topology that shall be preserved. Furthermore, extension in 3D or to multiple objects (*e.g.* brain structures) is straightforward thanks to the implicit representation of shapes. Despite

its paramount importance, the image-based term was not investigated as we focused on incorporating shape information on top of any pixelwise classifier.

All in all, this approach is very promising and was proven to be both effective and efficient on the addressed clinical application. It also paves the way for numerous further investigations: for instance in [Prevost et al., 2014], we go beyond and learn not only the shape variability but also the local appearance of the organ of interest.

References

- [Aronszajn, 1950] Aronszajn, N. (1950). Theory of reproducing kernels. *Trans. Amer. Math. Soc*, 68(3), 337–404.
- [Arsigny et al., 2006] Arsigny, V., Commowick, O., Pennec, X., & Ayache, N. (2006). A log-euclidean framework for statistics on diffeomorphisms. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2006* (pp. 924–931). Springer.
- [Bardinet et al., 1998] Bardinet, E., Cohen, L. D., & Ayache, N. (1998). A parametric deformable model to fit unstructured 3D data. *Computer Vision and Image Understanding*, 71(1), 39–54.
- [Besl & McKay, 1992] Besl, P. & McKay, N. (1992). A method for registration of 3D shapes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 14(2), 239–256.
- [Bottou, 1998] Bottou, L. (1998). Online learning and stochastic approximations. *On-line learning in neural networks*, 17, 9.
- [Cohen & Cohen, 1996] Cohen, I. & Cohen, L. D. (1996). A hybrid hyperquadric model for 2-D and 3-D data fitting. *Computer Vision and Image Understanding*, 63(3), 527–541.
- [Comon, 1994] Comon, P. (1994). Independent component analysis, a new concept? *Signal processing*, 36(3), 287–314.

- [Cootes et al., 1995] Cootes, T. F., Taylor, C. J., Cooper, D. H., & Graham, J. (1995). Active Shape Models-Their Training and Application. *Computer Vision and Image Understanding*, 61(1), 38–59.
- [Cremers et al., 2003] Cremers, D., Kohlberger, T., & Schnorr, C. (2003). Shape statistics in kernel space for variational image segmentation. *Pattern Recognition*, 36(9), 1929–1943.
- [Cremers et al., 2007] Cremers, D., Rousson, M., & Deriche, R. (2007). A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *International journal of computer vision*, 72(2), 195–215.
- [Cremers & Soatto, 2003] Cremers, D. & Soatto, S. (2003). A pseudo-distance for shape priors in level set segmentation. In *2nd IEEE workshop on variational, geometric and level set methods in computer vision* (pp. 169–176).: Citeseer.
- [Cuingnet et al., 2012] Cuingnet, R., Prevost, R., Lesage, D., Cohen, L. D., Mory, B., & Ardon, R. (2012). Automatic detection and segmentation of kidneys in 3D CT images using random forests. In *MICCAI*, volume 7512 of *LNCS* (pp. 66–74). Springer.
- [Durrleman, 2010] Durrleman, S. (2010). Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution. *These de sciences (PhD thesis), Université de Nice-Sophia Antipolis (March 2010)*.
- [Heimann & Meinzer, 2009] Heimann, T. & Meinzer, H.-P. (2009). Statistical shape models for 3D medical image segmentation: a review. *Medical image analysis*, 13(4), 543.
- [Huang et al., 2004] Huang, X., Li, Z., & Metaxas, D. (2004). Learning coupled prior shape and appearance models for segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004* (pp. 60–69). Springer.
- [Huang & Metaxas, 2008] Huang, X. & Metaxas, D. (2008). Metamorphs: Deformable shape and appearance models. *IEEE Trans. PAMI*, 30(8), 1444–1459.

- [Jolliffe, 1986] Jolliffe, I. T. (1986). *Principal component analysis*, volume 487. Springer-Verlag New York.
- [Joshi et al., 2004] Joshi, S. et al. (2004). Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23(1), 151.
- [Karcher, 1977] Karcher, H. (1977). Riemannian center of mass and mollifier smoothing. *Comm. Pure Appl. Math.*, 30(5), 509–541.
- [Khan et al., 2008] Khan, A. R., Wang, L., & Beg, M. F. (2008). Freesurfer-initiated fully-automated subcortical brain segmentation in mri using large deformation diffeomorphic metric mapping. *NeuroImage*, 41(3), 735–746.
- [Leventon et al., 2000] Leventon, M., Grimson, W., & Faugeras, O. (2000). Statistical shape influence in geodesic active contours. In *Proceedings of CVPR 2000* (pp. 316–323).
- [Mory et al., 2011] Mory, B., Somphone, O., Prevost, R., & Ardon, R. (2011). Template deformation with user constraints for live 3D interactive surface extraction. In *MICCAI Workshop MeshMed*.
- [Mory et al., 2012] Mory, B., Somphone, O., Prevost, R., & Ardon, R. (2012). Real-time 3D image segmentation by user-constrained template deformation. In *MICCAI*, volume 7510 of *LNCS* (pp. 561–8). Springer.
- [Prevost et al., 2014] Prevost, R., Cuingnet, R., Mory, B., Cohen, L., & Ardon, R. (2014). Tagged template deformation. In P. Golland, N. Hata, C. Barillot, J. Hornegger, & R. Howe (Eds.), *Medical Image Computing and Computer-Assisted Intervention MICCAI 2014*, volume 8673 of *Lecture Notes in Computer Science* (pp. 674–681). Springer International Publishing.
- [Prevost et al., 2013a] Prevost, R., Cuingnet, R., Mory, B., Cohen, L. D., & Ardon, R. (2013a). Incorporating Shape Variability in Image Segmentation by Implicit Template Deformation. In *MICCAI*, volume 8151 of *LNCS* (pp. 82–89). Springer.
- [Prevost et al., 2013b] Prevost, R., Cuingnet, R., Mory, B., Correias, J.-M., Cohen, L. D., & Ardon, R. (2013b). Joint Co-segmentation and

- Registration of 3D Ultrasound Images. In *Information Processing in Medical Imaging* (pp. 268–279). Springer.
- [Prevost et al., 2012] Prevost, R., Mory, B., Correas, J.-M., Cohen, L. D., & Ardon, R. (2012). Kidney detection and real-time segmentation in 3D contrast-enhanced ultrasound images. In *Proceedings of IEEE ISBI 2012* (pp. 1559–62).
- [Prevost et al., 2013c] Prevost, R., Romain, B., Cuingnet, R., Mory, B., Rouet, L., Lucidarme, O., Cohen, L. D., & Ardon, R. (2013c). Registration of Free-Breathing 3D+t Abdominal Perfusion CT Images via Co-Segmentation. In *MICCAI*, volume 8150 of *LNCS* (pp. 99–107). Springer.
- [Rousson & Paragios, 2002] Rousson, M. & Paragios, N. (2002). Shape priors for level set representations. *Proceedings of ECCV 2002*, (pp. 416–418).
- [Rousson et al., 2004] Rousson, M., Paragios, N., & Deriche, R. (2004). Implicit active shape models for 3D segmentation in MR imaging. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2004* (pp. 209–216). Springer.
- [Rueckert et al., 2001] Rueckert, D., Frangi, A., & Schnabel, J. (2001). Automatic construction of 3D statistical deformation models using non-rigid registration. *Proceedings of MICCAI 2001*, (pp. 77–84).
- [Saddi et al., 2007] Saddi, K., Chef d’hotel, C., Rousson, M., & Chriet, F. (2007). Region-based segmentation via non-rigid template matching. *Proceedings of ICCV*, 0, 1–7.
- [Somphone et al., 2008] Somphone, O., Mory, B., Makram-Ebeid, S., & Cohen, L. D. (2008). Prior-based piecewise-smooth segmentation by template competitive deformation using partitions of unity. In *ECCV*.
- [Tsai et al., 2003] Tsai, A., Yezzi Jr, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W., & Willsky, A. (2003). A shape-based approach to the segmentation of medical imagery using level sets. *Medical Imaging, IEEE Transactions on*, 22(2), 137–154.

- [Vaillant et al., 2004] Vaillant, M., Miller, M., Younes, L., & Trouvé, A. (2004). Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23(1), 161.
- [Wilcoxon, 1945] Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics bulletin*, 1(6), 80–83.
- [Yezzi & Soatto, 2003] Yezzi, A. J. & Soatto, S. (2003). Deformation: Deforming motion, shape average and the joint registration and approximation of structures in images. *International Journal of Computer Vision*, 53(2), 153–167.