4D Ventricular Segmentation and Wall Motion Estimation Using Efficient Discrete Optimization

Ahmed Besbes[1], Nikos Komodakis[1], Ben Glocker[2], Georgios Tziritas[3], and Nikos Paragios[1]

[1] GALEN Group, Laboratoire MAS, Ecole Centrale de Paris
{ahmed.besbes,nikos.komodakis,nikos.paragios}@ecp.fr
[2] Chair for Computer Aided Medical Procedures (CAMP)
Technische Universität München
glocker@in.tum.de
[3] University of Crete, Computer Science Department
tziritas@csd.uoc.gr

Abstract. In this paper we propose a novel approach to ventricular motion estimation and segmentation. Our method is based on a MRF formulation where an optimal intensity-based separation between the endocardium and the rest of the cardiac volume is to be determined. Such a term is defined in the spatiotemporal domain, where the ventricular wall motion is introduced to account for correspondences between the consecutive segmentation maps. The estimation of the deformations is done through a continuous deformation field (FFD) where the displacements of the control points are determined using discrete labeling approach. Principles from linear programming and in particular the Primal/Dual Schema is used to recover the optimal solution in both spaces. Promising experimental results obtained on 13 MR spatiotemporal data sets demonstrate the potentials of our method.

1 Introduction

The segmentation of the left ventricle has been a problem well addressed in medical imaging. Prior art either refers to model-free approaches or model-based. Model-free methods do not make an explicit assumption on the form/geometric properties as well as the appearance of the ventricle. MRFs [1], snakes [2], level sets [3], shortest path [4] have been considered in this context. On the other hand, model-based methods often consider certain geometric priors for the ventricle which could range from simple 2D shapes [4] and 3D models which also encode local variations [1] to complex biomechanical cardiac models [5].

Ventricular wall motion estimation was often addressed through the use of MR-Tagging [6] [7] techniques that consist of introducing a rectangular pattern on the acquisition. Direct 3D motion estimation in MR is a more challenging problem since it is known that the left ventricle undergoes a rather complex deformation within the cardiac cycle. In order to account for the ill-posedness of the problem, the use of shape models towards establishing visual correspondences and tracking was often considered [8] or 4D models have been constructed

with spatial and temporal deformations being encoded [9]. Voxel-based methods often explore the visual preservation assumption [10] while being constrained to provide a smooth deformation map. More complex models use biomechanical constraints to determine such a deformation [11], an approach which might fail when processing diseased data.

In most of the cases, these methods do not relate segmentation with ventricular motion estimation. Furthermore, one can claim that they are sensitive to the initial conditions either because of the non-convexity of the designed cost function or due to the sub-optimal optimization approach. In this paper, we propose a novel approach to address both segmentation and ventricular motion estimation. We overcome the ill-posedness of the motion estimation problem through the use of interpolation techniques with higher order polynomials, while we introduce temporal segmentation consistency through the use of deformations field. In order to efficiently recover the optimal solution to the problem, we re-formulate the cost function in a fully discrete domain where the latest developments of linear programming are considered to determine the lowest potential of the cost function. Very promising results and comparisons with manual segmentation from physicians demonstrate the potentials of our approach.

2   Ventricular Segmentation and Wall Motion Estimation

2.1   Spatiotemporal Segmentation

Let us consider a spatiotemporal volume $\mathcal{V}(\mathbf{x}; t) : \Omega \times [0..\tau] \rightarrow \mathcal{R}$, with Ω being the volume domain. The task of segmenting the endocardium can be reformulated using a labeling approach, or assigning a label $\phi(\mathbf{x}; t) : \Omega \times [0..\tau] \rightarrow \{0, 1\}$. Here, label 0 corresponds to the foreground (i.e., the ventricle), whereas label 1 corresponds to the background. Without loss of generality, let us assume that certain statistical properties on the intensities of the left ventricle $p(\mathcal{V}|\phi = 0)$, as well as on the intensities of the background $p(\mathcal{V}|\phi = 1)$ are available or can be determined on the fly. Let us also assume that we have a prior left ventricle closed surface $(S^t)_{t=0}^{\tau}$ defined as:

$$\xi_S(\mathbf{x}; t) = \begin{cases} 0 & \text{if } \mathbf{x} \in S^t \\ -\mathcal{D}(\mathbf{x}, \mathcal{S}) & \text{if } \mathbf{x} \in S_{in}^t \\ \mathcal{D}(\mathbf{x}, \mathcal{S}) & \text{if } \mathbf{x} \in S_{out}^t . \end{cases} \qquad (1)$$

with $\mathcal{D}$ being the Euclidean distance between a given voxel and the surface, and$(S^t, S_{in}^t, S_{out}^t)$ being the partition of Ω defined by S^t, $\forall t \in [0..\tau]$. We define a penalization function $p_\epsilon(\xi; \phi) : \mathcal{R} \times \{0, 1\} \rightarrow \mathcal{R}$, with $\epsilon > 0$ as a decreasing (respectively increasing) function of ξ if $\phi = 0$ (respectively $\phi = 1$), and equal to identity for $\xi < \epsilon$.

In such a case, the optimal labeling should refer to the maximum conditional posterior between the decisions and the data support. If spatial and temporal

independence are assumed between voxels, that labeling can then be recovered through the minimum of:

$$E_{seg,dt}(\phi) = \sum_{t=0}^{\tau} \sum_{\mathbf{x} \in \Omega} -\log\left[p\left(\mathcal{V}(\mathbf{x};t) | \phi(\mathbf{x};t)\right) . p_\epsilon\left(\xi_S(\mathbf{x};t); \phi(\mathbf{x};t)\right)\right]$$
$$= \sum_{t=0}^{\tau} \sum_{\mathbf{x} \in \Omega} V_{dt}^p(\phi(\mathbf{x};t)) \ .$$

(2)

which is equivalent to assigning to each voxel the label which is optimally supported from the observation. Such a simplistic formulation could produce suboptimal results due the presence of noise and therefore one should introduce additional smoothness constraints on the label space, which aims to enforce regularity on the decisions, or:

$$E_{seg,sp}(\phi) = \sum_{t=0}^{\tau} \sum_{\mathbf{x} \in \Omega} \left(\sum_{\mathbf{y} \in \mathcal{N}(\mathbf{x})} \psi(\phi(\mathbf{x};t), \phi(\mathbf{y};t)) \right) = \sum_{t=0}^{\tau} \sum_{\substack{\mathbf{x} \in \Omega \\ \mathbf{y} \in \mathcal{N}(\mathbf{x})}} V_{sp}(\phi(\mathbf{x};t), \phi(\mathbf{y};t)) \ .$$

(3)

with ψ being a function measuring the dissimilarity between labels of neighboring pixels and $\mathcal{N}(\mathbf{x})$ defines the local neighborhood of $\mathbf{x}$ in the 3D spatial domain. For more robustness, one can also consider temporal constraints on the labeling if the deformations from one volume to the next are not so important, which is, however, definitely not the case for the left and right ventricular motion. On the other hand, if we assume that this deformation is known, say, $d(\mathbf{x};t)$, then one can imagine using $d(\mathbf{x};t)$ towards determining the temporal derivative on the label space and introduce a smoothness constraint of the following form:

$$E_{seg,tm}(\phi|d) = \sum_{t=0}^{\tau-1} \sum_{\mathbf{x} \in \Omega} \psi(\phi(\mathbf{x};t), \phi(\mathbf{x}+d(\mathbf{x};t); t+1)) d\mathbf{x}$$
$$= \sum_{t=0}^{\tau-1} \sum_{\mathbf{x} \in \Omega} V_{tm}(\phi(\mathbf{x};t), \phi(\mathbf{x}+d(\mathbf{x};t); t+1)) \ .$$

(4)

The interpretation of this term is straightforward, assuming known correspondences one would expect a coherent labeling between anatomical structures within the cardiac cycle. Based on this fact, we can therefore proceed as follows: we will first estimate the deformation $d(\mathbf{x};t)$, i.e. register the 3D volumes, and then we will extract the optimal segmentation (i.e. the optimal labeling $\phi(\mathbf{x};t)$)) by minimizing the total energy E_{4D} of the resulting binary 4D Markov Random Field, where the total energy is given by:

$$E_{seg}(\phi|d) = E_{seg,dt}(\phi) + \alpha E_{seg,sp}(\phi) + \beta E_{seg,tm}(\phi|d) \ .$$

(5)

Intuitively, the edges of the resulting 4D MRF will consist of regular links, connecting (in a grid-like manner) voxels belonging to the same 3D volume. On

the other hand, they refer to irregular links in the temporal domain, connecting voxels between adjacent 3D volumes, being determined via the previously estimated deformation $d(\mathbf{x}; t)$. We also note that because our MRF is binary, the exact global optimum can be easily extracted [12].

However, establishing correspondences between volumes is an ill-posed problem. Even if we assume the visual preservation assumption to be valid (not often the case for medical image modalities), one should determine three unknown variables from a single constraint. To deal with this issue, in the next section we show how we can regularize this motion estimation problem by reformulating it as another discrete MRF optimization problem.

2.2   Ventricular Motion Estimation

Let us thus assume that we wish to compute the deformation $d(\mathbf{x}; t)$ between two adjacent 3D volumes at time t. To this end, we will introduce a sparse deformation grid $\mathcal{G}$ super-imposed on the source volume (no particular assumption is made on the grid except that it is sparser than the original volume). The central idea of our approach is to deform the grid (with a 3D displacement vector $d(\mathbf{p}; t)$ for each control point $\mathbf{p}$) such that the underlying volumes are perfectly aligned. Without loss of generality, we can then assume that the displacement of a voxel $\mathbf{x}$ can be expressed using a linear or non-linear combination of the grid points, or:

$$d(\mathbf{x}; t) = \sum_{p \in \mathcal{G}} \eta(|\mathbf{x} - \mathbf{p}|) \, d(\mathbf{p}; t) \ . \tag{6}$$

where $\eta(\cdot)$ is the weighting function measuring the contribution of the control point $\mathbf{p}$ to the displacement field $d(\mathbf{x}; t)$. The use of such a model is motivated by the fact that the observations refer to anatomical structures with a rather natural temporal deformation. Furthermore, such an approach could help us to account for the ill-posedness of the problem due to the fact that the estimation of a single 3D displacement is now an over-constrained problem with many observations being available. For $\eta(\cdot)$, we use a three-dimensional Free Form Deformation (FFD) model based on cubic B-splines [13] (other interpolation models can also be considered).

Therefore, based on (6), to estimate $d(\mathbf{x}; t)$ it suffices to specify the displacements for the control points. To this end, we will consider a quantized version of the deformation space, say, $\{d^1, ..., d^i\}$ - being 3D deformation vectors - as well as a corresponding set of discrete labels, say, $\mathcal{L} = \{1, ..., i\}$. A label assignment, say, $\omega(\mathbf{p}) \in \mathcal{L}$ to a grid point $\mathbf{p}$ is associated with displacing $\mathbf{p}$ by the corresponding vector $d^{\omega(\mathbf{p})}$, i.e.:

$$d(\mathbf{p}; t) = d^{\omega(\mathbf{p})} \ . \tag{7}$$

The visual preservation imposes the constraint that the observation of the same anatomical patch should be consistent across volumes, i.e., $\mathcal{V}(\mathbf{x}; t) \approx \mathcal{V}(\mathbf{x} + d(\mathbf{x}; t); t+1)$. In our discrete framework the deformation $d(\mathbf{x}; t)$ is defined based on (6), (7), i.e. displacements are associated with labels, one can reformulate

ventricular deformation estimation as a labeling problem. Consequently, the goal is to assign a set of appropriate labels $\{\omega(\mathbf{p})\}$ (to the grid points) so that the visual preservation constraint is satisfied as much as possible, or equivalently so that the following data cost is minimized:

$$E_{mot,dt}(\omega) = \sum_{\mathbf{x} \in \Omega} |\mathcal{V}(\mathbf{x};t) - \mathcal{V}(\mathbf{x} + d(\mathbf{x};t);t+1)| \stackrel{(6),(7)}{\approx} \sum_{\mathbf{p} \in \mathcal{G}} U_{dt}^p(\omega(\mathbf{p})) \ . \qquad (8)$$

Here, the singleton potential functions $U_{dt}^p(\cdot)$ are not independent, thus the defined data term can only be approximated. Hence, we precompute the $|\mathcal{L}| \times |\mathcal{G}|$ (where $|\mathcal{G}|$ is the number of grid points) data term in a look-up table. The entry for label $\omega(\mathbf{p})$ and node $\mathbf{p}$ is determined by:

$$U_{dt}^p(\omega(\mathbf{p})) = \iint_{\Omega(\mathbf{p})} \eta^{-1}(|\mathbf{x} - \mathbf{p}|) \cdot \left| \mathcal{V}(\mathbf{x};t) - \mathcal{V}(\mathbf{x} + d^{\omega(\mathbf{p})};t+1) \right| d\mathbf{x} \ . \qquad (9)$$

with the sum of absolute differences being considered as measure of similarity (η^{-1} is the inverse projection between $\mathbf{x}$ and $\mathbf{p}$). The use of an interpolation technique to determine the deformations of the volume will inherit natural smoothness to the estimates. However, one should also expect since we aim to recover measurements for physical objects deformations that the same assumption is satisfied for the deformation of the corresponding control points. Similar to the segmentation case, one can consider a term which enforces spatial similarities across labels, or:

$$E_{mot,sm}(\omega) = \sum_{\substack{\mathbf{p} \in \Omega \\ \mathbf{q} \in \mathcal{N}(\mathbf{p})}} U_{sm}(\omega(\mathbf{p}), \omega(\mathbf{q})) \ . \qquad (10)$$

where $\mathcal{N}$ represents the neighborhood system associated with the deformation grid $\mathcal{G}$. For the distance $U_{sm}(\cdot, \cdot)$, we consider a simple piecewise smoothness term based on the Euclidean distance between the deformations corresponding to the assigned labels, i.e.:

$$U_{sm}(\omega(\mathbf{p}), \omega(\mathbf{q})) = \lambda_{pq} \left(|d^{\omega(\mathbf{p})} - d^{\omega(\mathbf{q})}| \right) \ . \qquad (11)$$

with λ_{pq} being a (spatially varying) weighting to control the influence of the smoothness/prior term. Such a smoothness term, together with the data term, allows to convert the problem of volume registration into a discrete MRF optimization problem with the following energy [14]:

$$E_{mot}(\omega) = E_{mot,dt}(\omega) + E_{mot,sm}(\omega) \ . \qquad (12)$$

2.3   4D Segmentation & Ventricular Motion Estimation

One can now consider an objective function which recovers both the 4D segmentation map as well as the corresponding deformation fields:

$$E_{seg,mot}(\phi, \omega) = E_{seg}(\phi|\omega) + \gamma E_{mot}(\omega) \ . \qquad (13)$$

(a) The primal-dual principle   (b) The primal-dual schema
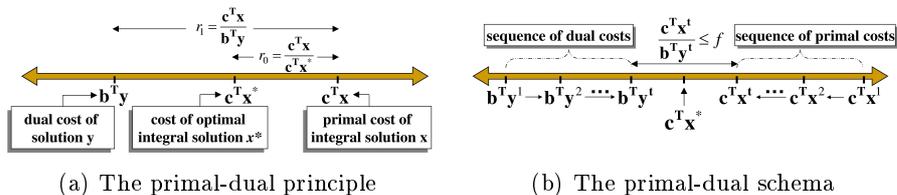
Fig. 1: (a) By weak duality, the optimal cost $\mathbf{c}^T\mathbf{x}^*$ will lie between the costs $\mathbf{b}^T\mathbf{y}$ and $\mathbf{c}^T\mathbf{x}$ of any pair $(\mathbf{x}, \mathbf{y})$ of integral-primal and dual feasible solutions. Therefore, if $\mathbf{b}^T\mathbf{y}$ and $\mathbf{c}^T\mathbf{x}$ are close enough (e.g. their ratio r_1 is $\leq f$), so are $\mathbf{c}^T\mathbf{x}^*$ and $\mathbf{c}^T\mathbf{x}$ (e.g. their ratio r_0 is $\leq f$ as well), thus proving that $\mathbf{x}$ is an f-approximation to $\mathbf{x}^*$. **(b)** According to the primal-dual schema, dual and integral-primal feasible solutions make local improvements to each other, until the final costs $\mathbf{b}^T\mathbf{y}^t$, $\mathbf{c}^T\mathbf{x}^t$ are close enough (e.g. their ratio is $\leq f$). We can then apply the primal-dual principle (as in Fig. (a)) and thus conclude that $\mathbf{x}^t$ is an f-approximation to $\mathbf{x}^*$

which is a fully discrete optimization problem. For optimizing the resulting MRF, we seek to assign a pair of labels $(\phi(\mathbf{p}), \omega(\mathbf{p}))$ to each node $p \in \mathcal{G}$, so that the MRF energy in (13) is minimized. To this end, a recently proposed method, called Fast-PD (Fast Primal Dual), will be used. This is an optimization technique, which builds upon principles drawn from the duality theory of linear programming in order to efficiently derive almost optimal solutions for a very wide class of NP-hard MRFs. For more details about the Fast-PD algorithm, the reader is referred to [12]. Here, we will just provide a brief, high level description of the basic driving force behind that algorithm.

3 Linear Programming

The driving force of the algorithm consists of the *primal-dual schema*, which is a well-known technique in the Linear Programming literature. To understand how the primal-dual schema works in general, we will need to consider the following pair of primal and dual Linear Programs (LPs):

$$\text{PRIMAL: } \min \mathbf{c}^T\mathbf{x} \qquad \text{DUAL: } \max \mathbf{b}^T\mathbf{y}$$
$$\text{s.t. } \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq 0 \qquad \text{s.t. } \mathbf{A}^T\mathbf{y} \leq \mathbf{c} \qquad (14)$$

Here $\mathbf{A}$ represents a coefficient matrix, while $\mathbf{b}, \mathbf{c}$ are coefficient vectors. Also, $\mathbf{x}, \mathbf{y}$ represent the vectors of primal and dual variables respectively. We seek an optimal solution to the primal program, but with the extra constraint of $\mathbf{x}$ being integral. Due to this integrality requirement, this problem is in general NP-hard and so we need to settle with estimating approximate solutions. A primal-dual f-approximation algorithm achieves that by use of the following principle (illustrated also in Fig. 1(a)):

Primal-Dual Principle 1 *If $\mathbf{x}$ and $\mathbf{y}$ are integral-primal and dual feasible solutions having a primal-dual gap less than f, i.e.:*

$$\mathbf{c}^T\mathbf{x} \leq f \cdot \mathbf{b}^T\mathbf{y}, \qquad (15)$$

then $\mathbf{x}$ *is an* f-*approximation to the optimal integral solution* $\mathbf{x}^*$, *i.e.* $\mathbf{c}^T\mathbf{x}^* \leq$ $\mathbf{c}^T\mathbf{x} \leq f \cdot \mathbf{c}^T\mathbf{x}^*$.

Based on the above principle, that lies at the heart of any primal-dual technique, the following iterative schema can be used for deriving an f-approximate solution (this schema is also illustrated graphically in Fig. 1(b)):

Primal-Dual Schema 1 *Keep generating pairs of integral-primal, dual solutions* $\{(\mathbf{x}^k, \mathbf{y}^k)\}_{k=1}^t$, *until the elements* $\mathbf{x}^t$, $\mathbf{y}^t$ *of the last pair are both feasible and have a primal-dual gap which is less than* f, *i.e. condition* (15) *holds true.*

In order to apply the above schema to MRF optimization, it suffices that we cast the MRF optimization problem as an equivalent integer program. The Fast-PD algorithm is then derived by applying the primal-dual schema to this pair of primal-dual LPs, while using $f = 2\frac{d_{\max}}{d_{\min}}$ ($d_{\max} \equiv \max_{a \neq b} d(a,b)$, $d_{\min} \equiv \min_{a \neq b} d(a,b)$) as the approximation factor in (15). Fast-PD is a very general MRF optimization method, which can handle a very wide class of MRFs. Essentially, it only requires that the MRF pairwise potential function is nonnegative. Furthermore, as already mentioned, it can guarantee that the generated solution is always within a worst-case bound from the optimum. In fact, besides this worst-case bound, it can also provide per-instance approximation bounds, which prove to be much tighter, i.e. very close to 1, in practice. It thus allows the global optimum to be found up to a user/application bound. Finally, it provides great computational efficiency, since it is typically 3-9 times faster than any other MRF optimization technique with guaranteed optimality properties [12].

4   Validation

In order to validate the performance of the method we have considered a set of 13 MR spatiotemporal volumes of the heart, with manual segmentation from two clinical experts being available for the diastole and the systole. These data sets had a spatial resolution of around 100x100x12 and a voxel size of around 1.77x1.77x6 millimeters. We used as prior information two learned distributions of endocardium voxels and background voxels expressed as mixture of Gaussians. These distributions were time-independent, and were used in diastole and systole as well. We also added a shape prior constraint (fixed shape S, initialized by the user) to account for the elliptic geometry of the left ventricle. In terms

Table 1: Comparison of automatic and experts' segmentations in diastole

Comparison	DSC Mean (Std)	Sensitivity	Specificity	ASD Mean (Std)
Our Method vs Expert1	0.86(±0.03)	99.06%	95.76%	1.54(±0.39)
Our Method vs Expert2	0.87(±0.02)	99.11%	96.88%	1.31(±0.37)
Expert1 vs Expert2	0.89(±0.02)	99.49%	94.46%	0.87(±0.12)
Expert2 vs Expert1	0.89(±0.02)	99.53%	94.16%	1.34(±0.47)

of segmentation performance we compare the experts' segmentation of the en-
docardium with the one obtained using the proposed method. We are interested
on several common evaluation measurements [15], and in particular the Dice
similarity coefficient (DSC), the sensitivity, the specificity, and the average sur-
face distance (ASD) from the experts segmentations. The ASD is computed in
millimeters from an anisotropic 3D Euclidean distance transform of the surfaces.
These measurements are computed in both diastole and systole and are presented
in [Tab. (1)] and [Tab. (2)]. We also compare in these tables the performances
of our method to those achieved manually by the experts.

(a) centripetal motion field   (b) centrifugal motion field
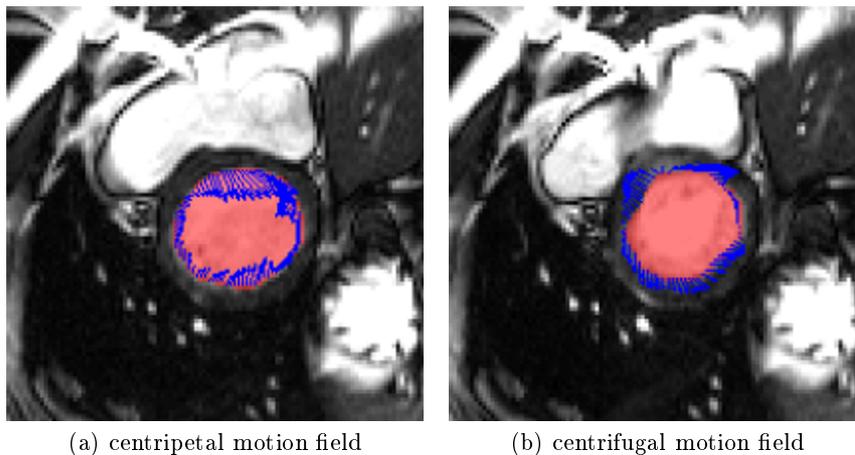
Fig. 2: Motion estimation. **(a)** beginning of systole **(b)** beginning of diastole

We achieve an ASD which is below the voxel size in both diastole and systole.
The DSC which measures the overlap between surfaces shows that our segmen-
tation is closer to the one given by Expert2 than to the one given by Expert1.
Overall, our performance is satisfactory compared to the one achieved by the
experts. We get a worse sensitivity than the experts, but a better specificity. In
terms of ventricular motion estimation, we present in [Fig. (2)] the deformation
field of the endocardium and its motion estimation. We see in particular in this
figure that the motion field is coherent with the left ventricle motion: the cen-
tripetal motion field at the beginning of systole is justified by the contraction of
the myocardium, and the centrifugal motion field at the beginning of diastole is
justified by its expansion. The 3D images in [Fig. (3)] show that we also correctly
segment the papillary muscles.

With a reasonable number of displacement labels (the complexity is linear to
the number of labels), the method takes about 10-20 seconds to converge (using a
DELL Duo with (3GHz,2GB)) assuming that a ventricle isolation has been done
and is able to produce good correspondences with a $16 \times 16 \times 16$ FFD grid. The

(a) in diastole   (b) in systole
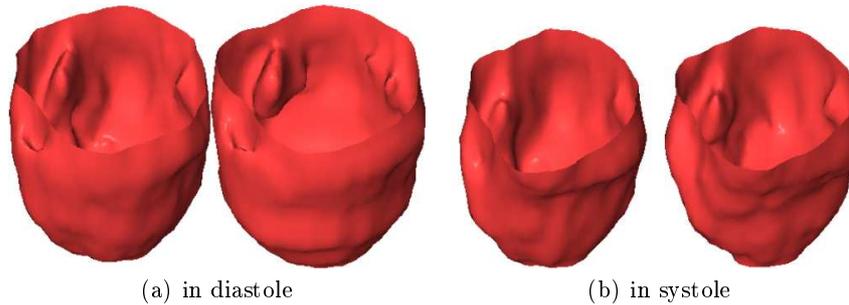
Fig. 3: Papillary muscles. In each image : automatic segmentation & experts' manual segmentation

cardiac cycle being quantized by 20-25 time points, the whole 4D segmentation and motion estimation computation takes about 70-80 seconds for a 4D volume.

5   Discussion

In this paper we have proposed a novel discrete approach to spatiotemporal segmentation and ventricular motion estimation. The strength of our approach is the coupling between the two problems and the use of a powerful combinatorial algorithm to produce their solution. In order to demonstrate the concept, we have considered a set of several heart 4D MRI exams and we have obtained quite satisfactory results. More challenging perspectives are related with the introduction of prior knowledge both in space and time related with the evolving geometry of the structures of interest. The prior information used in our approach remains quite simple, and is time-independent. That is why our results are promising and can be probably improved by the use of more complex prior information which can better capture the anatomy and the temporal dynamics of the cardiac cycle. Knowledge-based segmentation using models that encode important statistical variation of training examples within discrete optimization is a quite promising direction to be considered.

(a)   (b)
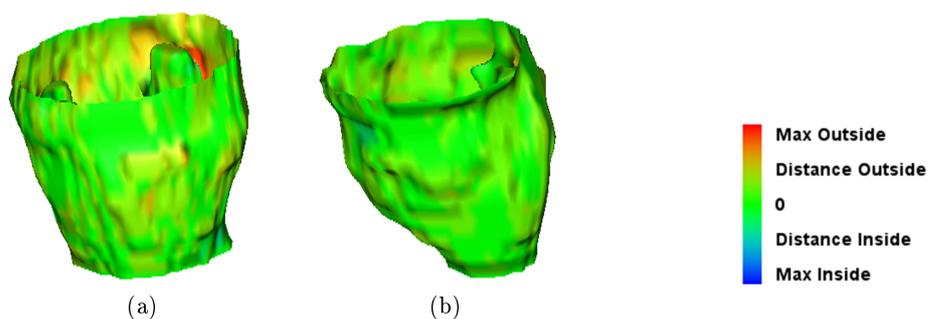
Fig. 4: Color-encoded visualization of the average surface distance for the example shown in [Fig. (2)]. **(a)** beginning of systole **(b)** beginning of diastole

Table 2: Comparison of automatic and experts' segmentations in systole

Comparison	DSC Mean (Std)	Sensitivity	Specificity	ASD Mean (Std)
Our Method vs Expert1	0.82(±0.03)	99.39%	93.34%	1.51(±0.39)
Our Method vs Expert2	0.85(±0.03)	99.46%	94.34%	1.28(±0.37)
Expert1 vs Expert2	0.86(±0.03)	99.69%	91.07%	0.86(±0.15)
Expert2 vs Expert1	0.86(±0.03)	99.66%	91.50%	1.06(±0.22)

References

1. Shi, P., Sinusas, A.J., Constable, R.T., Ritman, E., Duncan, J.S.: Point-tracked quantitative analysis of left ventricular surface motion from 3d image sequences. IEEE Trans. Med. Imaging **19**(1) (2000) 36–50
2. McInerney, T., Terzopoulos, D.: Deformable models in medical images analysis: a survey. Medical Image Analysis **1**(2) (1996) 91–108
3. Paragios, N.: A variational approach for the segmentation of the left ventricle in cardiac image analysis. Int. J. Comput. Vision **50**(3) (2002) 345–362
4. Jolly, M.P.: Automatic segmentation of the left ventricle in cardiac mr and ct images. Int. J. Comput. Vision **70**(2) (2006) 151–163
5. Sermesant, M., Forest, C., Pennec, X., Delingette, H., Ayache, N.: Deformable biomechanical bodels: Application to 4D cardiac image analysis. Medical Image Analysis **7**(4) (2003) 475–488
6. Montillo, A., Metaxas, D.N., Axel, L.: Automated model-based segmentation of the left and right ventricles in tagged cardiac mri. In: MICCAI (1). (2003) 507–515
7. Guttman, M., Prince, J., McVeigh, E.: Tag and contour detection in tagged mr images of the left ventricle. IEEE Trans. Med. Imaging **13**(1) (1994) 74–88
8. McEachen II, J.C., Duncan, J.S.: Shape-based tracking of left ventricular wall motion. IEEE Trans. Med. Imaging **16**(3) (1997) 270–283
9. Bosch, J.G., Mitchell, S.C., Lelieveldt, B.P.F., Nijland, F., Kamp, O., Sonka, M., Reiber, J.H.C.: Automatic segmentation of echocardiographic sequences by active appearance motion models. IEEE Trans. Med. Imaging **21**(11) (2002) 1374–1383
10. Horn, B.K.P., Schunck, B.G.: Determining optical flow. Artif. Intell. **17**(1-3) (1981) 185–203
11. Sermesant, M., Delingette, H., Ayache, N.: An electromechanical model of the heart for image analysis and simulation. IEEE Trans. Med. Imaging **25**(5) (2006) 612–625
12. Komodakis, N., Tziritas, G., Paragios, N.: Fast, approximately optimal solutions for single and dynamic mrfs. In: IEEE Conference on Computer Vision & Pattern Recognition (CVPR). (2007)
13. Sederberg, T.W., Parry, S.R.: Free-form deformation of solid geometric models. In: SIGGRAPH '86: Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques. (1986) 151–160
14. Glocker, B., Komodakis, N., Paragios, N., Tziritas, G., Navab, N.: Inter and intra-modal deformable registration: Continuous deformations meet efficient optimal linear programming. In: Information Processing in Medical Imaging (IPMI). (2007)
15. Gerig, G., Jomier, M., Chakos, M.: Valmet: A new validation tool for assessing and improving 3d object segmentation. In: MICCAI. (2001) 516–523