

Left Ventricle Segmentation via Graph Cut Distribution Matching

Ismail Ben Ayed¹, Kumaradevan Punithakumar¹, Shuo Li¹, Ali Islam²,
and Jaron Chong³

¹ GE Healthcare, London, ON, Canada

² St. Joseph's Health Care, London, ON, Canada

³ University of Western Ontario, London, ON, Canada

Abstract. We present a discrete kernel density matching energy for segmenting the left ventricle cavity in cardiac magnetic resonance sequences. The energy and its graph cut optimization based on an original first-order approximation of the Bhattacharyya measure have not been proposed previously, and yield competitive results in nearly real-time. The algorithm seeks a region within each frame by optimization of two priors, one geometric (distance-based) and the other photometric, each measuring a distribution similarity between the region and a model learned from the first frame. Based on global rather than pixelwise information, the proposed algorithm does not require complex training and optimization with respect to geometric transformations. Unlike related active contour methods, it does not compute iterative updates of computationally expensive kernel densities. Furthermore, the proposed first-order analysis can be used for other intractable energies and, therefore, can lead to segmentation algorithms which share the flexibility of active contours and computational advantages of graph cuts. Quantitative evaluations over 2280 images acquired from 20 subjects demonstrated that the results correlate well with independent manual segmentations by an expert.

1 Introduction

Accurate segmentation of the left ventricle (LV) cavity in magnetic resonance (MR) sequences is very important for complete diagnosis of cardiovascular diseases [1], [3]. Manual segmentation of all images is prohibitively time-consuming. Therefore, automatic or semi-automatic algorithms are highly desired. Albeit an impressive research effort has been devoted to the LV [1]–[15], current methods are still not sufficiently fast and flexible for routine clinical use, mainly because of the difficulties inherent to MR cardiac images [4]. Existing methods are based, among others, on active contours [1]–[3], [5]–[11], active appearance/shape models [12], [14], and registration [15]. Generally, the problem is stated as an energy optimization. In this connection, optimization of active contour functionals has been the most prevalent and flexible choice in the literature because it allows introducing a wide range of photometric and geometric¹ constraints on the solution [1]–[3], [5]–[11]. Generally, these constraints reference a sum over the target region or its boundary of *pixelwise* correspondences between the

¹ Geometric constraints reference object shape, position, and size.

image and geometric/photometric models learned from a training set. Pixelwise photometric information cannot distinguish connected cardiac regions having almost the same photometric profile [1]–[3], for instance the papillary muscles within the cavity and the myocardium (cf. the examples in Fig. 1). Therefore, most of existing methods bias the solution towards a finite set of shapes learned *a priori*. Although very effective in some cases, training-based algorithms may have difficulty in capturing the substantial subject variations in a clinical context [1], [3], [4]. The ensuing results are bounded to the characteristics, variability, and mathematical description of the training set. For instance, a pathological case outside the set of learned shapes may not be recovered, and photometric models have to be updated for new acquisition protocols and sequences.

To relax dependence on a statistical training, the studies in [1]–[3] investigated active contour optimization of *intensity matching* criteria, given a user-provided segmentation of one frame. In [3], the authors propose to maintain a constant photometric environment in the vicinity of the cavity boundary propagated over the sequence. Based on a global similarity measure between distributions, the method in [1], [2] maintains over a cardiac sequence a constant overlap between the intensity distributions of the cavity and myocardium, which led to promising results for mid-cavity images. Based only on the current data, these methods allow more flexibility in clinical use, although at the price of a user initialization. Furthermore, in the context of general-purpose methods, recent studies have shown that the use of *global* distribution-matching measures outperforms standard segmentation techniques based on pixelwise information, and is less sensitive to inaccuracies in estimating the models [16]–[18]. As such, it can relax the need of complex training. Unfortunately, optimization of a global measure with respect to segmentation is *NP-hard* [21], and the problem has been commonly addressed with active contours [1], [2], [16]–[18], which lead to computationally intensive algorithms. Along with an incremental contour evolution, the methods in [1], [2], [16]–[18] require a large number of updates of computationally onerous integrals, namely, the distributions of the regions defined by the contour at each iteration and the corresponding measures. Active contour methods rely on stepwise gradient descent. As a result, the algorithms are notoriously slow, converge to a local minimum, and depend on the choice of an approximating numerical scheme of contour evolution and the corresponding parameters.

This study investigates a *discrete* distribution-matching energy defined over a binary labeling, based on the Bhattacharyya kernel, and containing two priors, one geometric (distance-based) and the other photometric, each measuring a distribution similarity between the target region and a model. The ensuing problem is *NP-hard*, and the energy does not afford an analytical form amenable to graph cut optimization. To address efficiently the problem, we propose an original first-order approximation of the Bhattacharyya measure by introducing an auxiliary labeling, thereby computing a global graph cut optimum in nearly real-time. Unlike active contours, the algorithm does not compute iterative integral updates. It requires only a graph cut. The energy removes the need of a training, and prevents the papillary muscles from being included in the myocardium. Based on a distance distribution, the geometric prior handles intrinsically variations of the LV without biasing the solution towards a finite set of shapes, and relaxes optimization over geometric transformations. Furthermore, the proposed first-order analysis can be used for other intractable energies and, therefore, can lead to

segmentation algorithms which share the flexibility of active contours and computational advantages of graph cuts. Evaluations and comparisons with [1] demonstrated that the proposed method brings improvements in accuracy and computational efficiency.

2 Formulation

Consider a MR cardiac sequence containing N image functions² $\mathbf{I}_p^n = \mathbf{I}^n(p) : \mathcal{P} \subset \mathbb{N} \rightarrow \mathcal{I}$, $n \in [1..N]$, with \mathcal{P} the positional array and \mathcal{I} the space of photometric variables. Our purpose is to automatically detect the cavity of the heart in each frame $n \in [2..N]$ (cf. Fig. 1). For $n \in [2..N]$, we state the problem as the minimization of a *discrete* cost function with respect to a binary variable (labeling), $\mathcal{L}^n(p) : \mathcal{P} \rightarrow \{0, 1\}$, which defines a variable partition of \mathcal{P} : the *heart cavity* \mathbf{C}^n corresponding to region $\{p \in \mathcal{P} / \mathcal{L}^n(p) = 1\}$ and its complement, the *background* \mathbf{B}^n corresponding to region $\{p \in \mathcal{P} / \mathcal{L}^n(p) = 0\}$. The optimal labeling is sought by minimizing an original energy designed to address the problems related to cardiac MR images and containing two *kernel density matching* terms, an *intensity matching* term and a *distance matching* term. To introduce our energy, we first consider the following definitions for any labeling $\mathcal{L} : \mathcal{P} \rightarrow \{0, 1\}$, any image $\mathbf{I} : \mathcal{P} \rightarrow \mathcal{I}$, and any space of variables \mathcal{I} .

- $\mathbf{P}_{\mathcal{L}, \mathbf{I}}^{\mathcal{I}}$ is the kernel density estimate (KDE) of the distribution of image data \mathbf{I} within region $\mathbf{R}_{\mathcal{L}} = \{p \in \mathcal{P} / \mathcal{L}(p) = 1\}$

$$\forall i \in \mathcal{I}, \quad \mathbf{P}_{\mathcal{L}, \mathbf{I}}^{\mathcal{I}}(i) = \frac{\sum_{p \in \mathbf{R}_{\mathcal{L}}} K(i - \mathbf{I}_p)}{\mathbf{A}_{\mathcal{L}}}, \quad \text{with } K(y) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{y^2}{2\sigma^2}\right), \quad (1)$$

$\mathbf{A}_{\mathcal{L}}$ is the number of pixels within $\mathbf{R}_{\mathcal{L}}$: $\mathbf{A}_{\mathcal{L}} = \sum_{\mathbf{R}_{\mathcal{L}}} 1$, and σ is the width of the Gaussian kernel. Note that choosing K equal to the Dirac function yields the histogram.

- $\mathcal{B}(f, g)$ is the *Bhattacharyya* coefficient³ measuring the amount of overlap (similarity) between two distributions f and g : $\mathcal{B}(f, g) = \sum_{i \in \mathcal{I}} \sqrt{f(i)g(i)}$.

We assume that a segmentation of frame \mathbf{I}^1 , i.e., a labeling \mathcal{L}^1 defining a partition $\{\mathbf{C}^1, \mathbf{B}^1\}$, is given. Using this prior information from the first frame in the current data, the intensity/geometry model distributions of the cavity are learned, and embedded in the following distribution matching constraints to segment subsequent frames.

Intensity matching term. Given the learned model distribution of intensity, which we denote $\mathbf{M}^{\mathcal{I}} = \mathbf{P}_{\mathcal{L}^1, \mathbf{I}^1}^{\mathcal{I}}$, the purpose of this term is to find for each subsequent frame \mathbf{I}^n a region \mathbf{C}^n whose intensity distribution most closely matches $\mathbf{M}^{\mathcal{I}}$. To this end, we minimize the following intensity matching function with respect to \mathcal{L} :

$$\mathcal{B}^{\mathcal{I}}(\mathcal{L}, \mathbf{I}^n) = -\mathcal{B}(\mathbf{P}_{\mathcal{L}, \mathbf{I}^n}^{\mathcal{I}}, \mathbf{M}^{\mathcal{I}}) = -\sum_{i \in \mathcal{I}} \sqrt{\mathbf{P}_{\mathcal{L}, \mathbf{I}^n}^{\mathcal{I}}(i) \mathbf{M}^{\mathcal{I}}(i)} \quad (2)$$

Distance matching term. The purpose of this term is to constrain the segmentation with prior geometric information (shape, scale, and position of the cavity) obtained

² The number of frames N is typically equal to 20 or 25.

³ Note that the values of \mathcal{B} are always in $[0, 1]$, where 0 indicates that there is no overlap, and 1 indicates a perfect match between the distributions.

from the learning frame. Let c be the centroid of cavity \mathbf{C}^1 in the learning frame and $\mathbf{D}(p) = \frac{\|p-c\|}{N_{\mathbf{D}}} : \mathcal{P} \rightarrow \mathcal{D}$ a *distance image* measuring at each point $p \in \mathcal{P}$ the normalized distance between p and c , with \mathcal{D} the space of distance variables and $N_{\mathbf{D}}$ a normalization constant. Let $\mathbf{M}^{\mathcal{D}} = \mathbf{P}_{\mathcal{L}^1, \mathbf{D}}^{\mathcal{D}}$ the model distribution of distances within the cavity in the learning frame. We propose to find a region \mathbf{C}^n whose distance distribution most closely matches $\mathbf{M}^{\mathcal{D}}$ by minimizing:

$$\mathcal{B}^{\mathcal{D}}(\mathcal{L}, \mathbf{D}) = -\mathcal{B}(\mathbf{P}_{\mathcal{L}, \mathbf{D}}^{\mathcal{D}}, \mathbf{M}^{\mathcal{D}}) = -\sum_{d \in \mathcal{D}} \sqrt{\mathbf{P}_{\mathcal{L}, \mathbf{D}}^{\mathcal{D}}(d) \mathbf{M}^{\mathcal{D}}(d)} \quad (3)$$

Note that this geometric prior is invariant to rotation, and embeds *implicitly* uncertainties with respect to scale via the kernel width σ in (1). The higher σ , the more scale variations allowed. In our experiments, $\sigma = 2$ was sufficient to handle effectively variations in the scale of the cavity (cf. the examples in Fig. 1). The proposed geometric prior relaxes (1) complex learning/modeling of geometric characteristics and the need of a training set and (2) *explicit* optimization with respect to geometric transformations.

The proposed energy function. We propose to minimize an energy containing the intensity/distance matching terms and a regularization term for smooth segmentation boundaries. For each $n \in [2..N]$, the algorithm computes the optimal labeling \mathcal{L}_{opt}^n minimizing the following discrete cost function over all $\mathcal{L} : \mathcal{P} \rightarrow \{0, 1\}$:

$$\mathcal{F}(\mathcal{L}, \mathbf{I}^n) = \underbrace{\mathcal{B}^{\mathcal{I}}(\mathcal{L}, \mathbf{I}^n)}_{\text{Intensity Matching}} + \underbrace{\mathcal{B}^{\mathcal{D}}(\mathcal{L}, \mathbf{D})}_{\text{Geometry Matching}} + \underbrace{\lambda \mathbf{S}(\mathcal{L})}_{\text{Smoothness}} \quad (4)$$

where $\mathbf{S}(\mathcal{L})$ is related to the length of the partition boundary given by [22]:

$$\mathbf{S}(\mathcal{L}) = \sum_{\{p,q\} \in \mathcal{N}} \frac{1}{\|p-q\|} \delta_{\mathcal{L}(p) \neq \mathcal{L}(q)}, \quad \text{with } \delta_{x \neq y} = \begin{cases} 1 & \text{if } x \neq y \\ 0 & \text{if } x = y \end{cases} \quad (5)$$

and \mathcal{N} is a neighborhood system containing all unordered pairs $\{p, q\}$ of neighboring elements of \mathcal{P} . λ is a positive constant that balances the relative contribution of \mathbf{S} .

Global and efficient graph cut optimization. Optimization of the distribution matching functions in $\mathcal{F}(\mathcal{L}, \mathbf{I}^n)$ is not directly amenable to graph cut computation. It is an *NP-hard* problem. Furthermore, gradient-based optimization procedures are computationally very expensive and difficult to apply. To solve this problem efficiently, we propose a *first-order approximation* of the Bhattacharyya measures in $\mathcal{F}(\mathcal{L}, \mathbf{I}^n)$ by introducing an *auxiliary*⁴ labeling which corresponds to an arbitrary, fixed partition. For any labeling \mathcal{L} , we rewrite the intensity matching term minus a constant as follows:

$$\mathcal{B}^{\mathcal{I}}(\mathcal{L}, \mathbf{I}^n) - \underbrace{\mathcal{B}^{\mathcal{I}}(\mathcal{L}_a, \mathbf{I}^n)}_{\text{Constant}} \approx \underbrace{\sum_{p \in \mathcal{P}} \delta \mathcal{B}_{p, \mathcal{L}_a, \mathcal{L}}^{\mathcal{I}}}_{\text{Variations of } \mathcal{B}^{\mathcal{I}}} \approx -\frac{1}{2} \sum_{p \in \mathcal{P}} \sum_{i \in \mathcal{I}} \sqrt{\frac{\mathbf{M}^{\mathcal{I}}(i)}{\mathbf{P}_{\mathcal{L}_a, \mathbf{I}^n}^{\mathcal{I}}(i)}} \delta \mathbf{P}_{p, \mathcal{L}_a, \mathcal{L}}^{\mathcal{I}}(i), \quad (6)$$

⁴ Note that \mathcal{L}_a is an arbitrary fixed labeling which can be obtained from a given segmentation of the first frame.

where $\delta\mathcal{B}_{p,\mathcal{L}_a,\mathcal{L}}^{\mathcal{I}}$ (respectively $\delta\mathbf{P}_{p,\mathcal{L}_a,\mathcal{L}}^{\mathcal{I}}(i)$) is the *elementary variation* of $\mathcal{B}^{\mathcal{I}}(\mathcal{L}_a, \mathbf{I}^n)$ (respectively $\mathbf{P}_{\mathcal{L}_a, \mathbf{I}^n}^{\mathcal{I}}(i)$) that corresponds to changing the label of pixel p from $\mathcal{L}_a(p)$ to $\mathcal{L}(p)$. Elementary variation $\delta\mathcal{B}_{p,\mathcal{L}_a,\mathcal{L}}^{\mathcal{I}}$ is computed in the rightmost approximation of (6) with the *first-order expansion* of the Bhattacharyya measure $\mathcal{B}^{\mathcal{I}}(\mathcal{L}, \mathbf{I}^n)$. Now we compute elementary variations $\delta\mathbf{P}_{p,\mathcal{L}_a,\mathcal{L}}^{\mathcal{I}}(i)$, $i \in \mathcal{I}$, using the expression of the kernel density estimate in (1), which yields after some algebraic manipulations:

$$\delta\mathbf{P}_{p,\mathcal{L}_a,\mathcal{L}}^{\mathcal{I}}(i) = \begin{cases} \delta_{\mathcal{L}_a(p) \neq 1} \frac{K(i - \mathbf{I}_p^n) - \mathbf{P}_{\mathcal{L}_a, \mathbf{I}^n}^{\mathcal{I}}(i)}{\mathbf{A}_{\mathcal{L}_a} + 1} & \text{if } \mathcal{L}(p) = \mathbf{1} \\ \delta_{\mathcal{L}_a(p) \neq 0} \frac{\mathbf{P}_{\mathcal{L}_a, \mathbf{I}^n}^{\mathcal{I}}(i) - K(i - \mathbf{I}_p^n)}{\mathbf{A}_{\mathcal{L}_a} - 1} & \text{if } \mathcal{L}(p) = \mathbf{0} \end{cases} \quad (7)$$

where $\delta_{x \neq y}$ given by (5). Finally, using (7) in (6) and after some manipulations, the intensity matching term reads as the sum of *unary penalties* plus a constant:

$$\mathcal{B}^{\mathcal{I}}(\mathcal{L}, \mathbf{I}^n) \approx \text{constant} + \sum_{p \in \mathcal{P}} \mathbf{b}_{p, \mathbf{I}^n}^{\mathcal{I}}(\mathcal{L}(p)), \quad (8)$$

with $\mathbf{b}_{p, \mathbf{I}}^{\mathcal{I}}$ given, for any image $\mathbf{I} : \mathcal{P} \rightarrow \mathcal{I}$ and any space of variables \mathcal{I} , by

$$\begin{aligned} \mathbf{b}_{p, \mathbf{I}}^{\mathcal{I}}(\mathbf{1}) &= \frac{\delta_{\mathcal{L}_a(p) \neq 1}}{2(\mathbf{A}_{\mathcal{L}_a} + 1)} \left(\mathcal{B}^{\mathcal{I}}(\mathcal{L}_a, \mathbf{I}) - \sum_{i \in \mathcal{I}} K(i - \mathbf{I}_p) \sqrt{\frac{\mathbf{M}^{\mathcal{I}}(i)}{\mathbf{P}_{\mathcal{L}_a, \mathbf{I}}^{\mathcal{I}}(i)}} \right) \\ \mathbf{b}_{p, \mathbf{I}}^{\mathcal{I}}(\mathbf{0}) &= \frac{\delta_{\mathcal{L}_a(p) \neq 0}}{2(\mathbf{A}_{\mathcal{L}_a} - 1)} \left(\sum_{i \in \mathcal{I}} K(i - \mathbf{I}_p) \sqrt{\frac{\mathbf{M}^{\mathcal{I}}(i)}{\mathbf{P}_{\mathcal{L}_a, \mathbf{I}}^{\mathcal{I}}(i)}} - \mathcal{B}^{\mathcal{I}}(\mathcal{L}_a, \mathbf{I}) \right) \end{aligned} \quad (9)$$

Using a similar computation for the distance matching term, adopting the same notation in (9) for distance image \mathbf{D} , and ignoring the constants, our problem reduces to optimizing the following sum of unary and pairwise (submodular) penalties:

$$\mathcal{L}_{opt} = \arg \min_{\mathcal{L} : \mathcal{P} \rightarrow \{0,1\}} \sum_{p \in \mathcal{P}} \{ \mathbf{b}_{p, \mathbf{I}^n}^{\mathcal{I}}(\mathcal{L}(p)) + \mathbf{b}_{p, \mathbf{D}}^{\mathcal{D}}(\mathcal{L}(p)) \} + \lambda \mathbf{S}(\mathcal{L}) \quad (10)$$

In combinatorial optimization, a *global optimum* of the sum of unary and pairwise (submodular) penalties can be computed efficiently in *low-order polynomial time* by solving an equivalent max-flow problem [20]. In our case, it suffices to build a *weighted graph* $\mathcal{G} = \langle \mathbf{N}, \mathbf{E} \rangle$, where \mathbf{N} is the set of nodes and \mathbf{E} the set of edges connecting these nodes. \mathbf{N} contains a node for each pixel $p \in \mathcal{P}$ and two additional terminal nodes, one representing the foreground region (i.e., the cavity), denoted \mathbf{T}_F , and the other representing the background, denoted \mathbf{T}_B . Let $w_{p,q}$ be the weight of the edge connecting neighboring pixels $\{p, q\}$ in \mathcal{N} , and $\{w_{p, \mathbf{T}_F}, w_{p, \mathbf{T}_B}\}$ the weights of the edges connecting each pixel p to each of the terminals. By setting the edge weights as follows:


$$w_{p, \mathbf{T}_F} = \mathbf{b}_{p, \mathbf{I}^n}^{\mathcal{I}}(\mathbf{0}) + \mathbf{b}_{p, \mathbf{D}}^{\mathcal{D}}(\mathbf{0}); \quad w_{p, \mathbf{T}_B} = \mathbf{b}_{p, \mathbf{I}^n}^{\mathcal{I}}(\mathbf{1}) + \mathbf{b}_{p, \mathbf{D}}^{\mathcal{D}}(\mathbf{1}); \quad w_{p,q} = \frac{\lambda}{\|p - q\|},$$

we compute, using the max-flow algorithm of Boykov and Kolmogorov [20], a minimum cut \mathcal{C}_{opt}^n of \mathcal{G} , i.e., a subset of edges in \mathbf{E} whose removal divides the graph into


two disconnected subgraphs, each containing a terminal node, and whose sum of edge weights is minimal. This minimum cut, which assigns each node (pixel) p in \mathcal{P} to one of the two terminals, induces an optimal labeling \mathcal{L}_{opt}^n ($\mathcal{L}_{opt}^n(p) = 1$ if p is connected to \mathbf{T}_F and $\mathcal{L}_{opt}^n(p) = 0$ if p is connected to \mathbf{T}_B), which minimizes globally the approximation in (10) and, therefore, the proposed energy function.

3 Experimental Evaluations and Comparisons

We applied the method to 120 short axis cardiac cine MR sequences acquired from 20 subjects: a total of 2280 images including apical, mid-cavity and basal slices were automatically segmented, and the results were compared to independent manual segmentations by an expert. Using the same datasets, we compared the accuracy and computational load/time of the proposed method with the recent LV segmentation in [1]. Similar to [1], the proposed method relaxes the need of a training, and model distributions were learned from a user-provided segmentation of the first frame in each sequence. The regularization and kernel width parameters were unchanged for all the datasets: λ is fixed equal to 0.15, and kernel width σ is set equal to 2 for distance distributions, and equal to 10 for intensity distributions. In Fig. 1, we give a representative sample of the results for 2 subjects. Although it uses information from only one frame in the current data, the method handles implicitly variations in the scale/shape of the



Subject 1, a sample of the results with mid-cavity (1st row) and apical (2nd row) frames



Subject 2, a sample of the results with mid-cavity (1st row) and basal (2nd row) frames

Fig. 1. A representative sample of the results for 2 subjects

Table 1. Dice metrics (DM expressed as mean \pm standard deviation) and Reliability function ($\mathcal{R}(d) = Pr(DM > d)$). The higher the DM and \mathcal{R} , the better the performance.

Performance measure	DM	$\mathcal{R}(0.80)$	$\mathcal{R}(0.85)$	$\mathcal{R}(0.90)$
Proposed method	0.91 ± 0.04	0.97	0.93	0.80
Method in [1]	0.88 ± 0.09	0.89	0.85	0.72

Table 2. Computation time (CPU) and number of kernel density estimations (KDEs) for the proposed method and the curve evolution method in [1]

Computation time/load	Average CPU/frame	Average CPU/subject	Nb of KDEs/frame
Proposed method	0.08 sec	9.62 secs	1
Method in [1]	4.33 secs	494.45 secs	300

cavity and prevents the papillary muscles from being included erroneously in the myocardium. Therefore, it relaxes the need of complex training and optimization over geometric transformations.

Dice metric. We evaluated the *Dice Metric* (DM) commonly used to measure the similarity (overlap) between manual and automatic segmentations [1], [6], [7]. Let \mathbf{V}_a , \mathbf{V}_m and \mathbf{V}_{am} be the volumes⁵ of, respectively, the automatically segmented cavity, the corresponding hand-labeled cavity, and the intersection between them. DM is given by⁶: $DM = \frac{2\mathbf{V}_{am}}{\mathbf{V}_a + \mathbf{V}_m}$. The proposed method yielded a DM equal to 0.91 ± 0.04 for all the data analyzed (DM is expressed as mean \pm standard deviation). Table 1 reports DM statistics for the proposed method and [1]. Using the same data, the method in [1] yielded a DM equal to 0.88 ± 0.09 . Fig. 2 (a) depicts the DM for a representative sample of the analyzed volumes. The proposed method led to a significant improvement in average accuracy. Note that an average DM higher than 0.80 indicates an excellent agreement with manual segmentations [7], and an average DM higher than 0.90 is, generally, difficult to obtain because the small structure of the cavity at the apex decreases significantly the DM [6]. For instance, the study in [6] reports a DM equal to 0.81 ± 0.16 whereas the authors in [7] report a DM equal to 0.88 ± 0.06 .

Reliability. we examined quantitatively and comparatively the *reliability* of the algorithm by evaluating the *reliability function*—i.e., the complementary cumulative distribution function (ccdf)—of the obtained Dice metrics, defined for each $d \in [0, 1]$ as the probability of obtaining DM higher than d over all volumes: $\mathcal{R}(d) = Pr(DM > d) = (\text{number of volumes segmented with } DM \text{ higher than } d) / (\text{total number of volumes})$. $\mathcal{R}(d)$ measures how reliable the algorithm in yielding accuracy d . For the proposed method and [1], we report in Table 1 the reliability of different accuracy levels ($d = 0.80, d = 0.85, d = 0.90$), and plot \mathcal{R} as a function of d in Fig. 2 (c). Our algorithm led to a higher reliability curve and 8% improvement in reliabilities. For instance, we obtained $\mathcal{R}(0.80) = 0.97$, i.e, an excellent agreement ($DM > 0.80$) in 97% of the cases, whereas the method in [1] achieved 89% of the cases with a similar accuracy.

⁵ The volume is measured by the sum of areas of the cavity region in 6 segmented images.

⁶ DM is always in $[0, 1]$. DM equal to 1 indicates a perfect match between segmentations.




Fig. 2. Comparisons of manual and automatic segmentations of 2280 images (380 volumes) acquired from 20 subjects. (a) Dice metric (DM) in a representative sample of the tested volumes for the proposed method and [1]. (b) Automatic volumes versus manual volumes for the proposed method: the obtained *correlation coefficient* is $r = 0.9779$. (c) Reliability ($\mathcal{R}(d) = Pr(DM > d)$) for the proposed method and [1]. The proposed method led to a *higher* reliability curve.

Correlation coefficient. The proposed method yielded a high correlation between manual and automatic segmentations: $r = 0.9779$. The linear regression plot, displayed in Fig. 2 (b) with the identity line, illustrates this correlation.

Computation time/load. Although based on distribution measures, the algorithm led to nearly *real-time* segmentation. Running on a 2 GHz machine, it needs 0.08 *sec* to process a frame (12.5 *images/sec*). Table 2 reports the computation time/load for the proposed method and [1]. The proposed approximation leads to a significant decrease in computation load because it requires only a joint kernel density estimation (KDE) and graph cut, whereas the curve evolution in [1] requires approximately 300 KDEs.

References

1. Ben Ayed, I., Li, S., Ross, I.: Embedding Overlap Priors in Variational Left Ventricle Tracking. *IEEE Trans. on Medical Imaging* (available online) (in press, 2009)
2. Ben Ayed, I., Lu, Y., Li, S., Ross, I.: Left ventricle tracking using overlap priors. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I. LNCS*, vol. 5241, pp. 1025–1033. Springer, Heidelberg (2008)
3. Hautvast, G., Lobregt, S., Breeuwer, M., Gerritsen, F.: Automatic Contour Propagation in Cine Cardiac Magnetic Resonance Images. *IEEE Trans. on Medical Imaging* 25(11), 1472–1482 (2006)
4. Jolly, M.-P.: Automatic recovery of the left ventricular blood pool in cardiac cine MR images. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I. LNCS*, vol. 5241, pp. 110–118. Springer, Heidelberg (2008)
5. Fradkin, M., Ciofolo, C., Mory, B., Hautvast, G., Breeuwer, M.: Comprehensive segmentation of cine cardiac MR images. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I. LNCS*, vol. 5241, pp. 178–185. Springer, Heidelberg (2008)
6. Lynch, M., Ghita, O., Whelan, P.F.: Segmentation of the Left Ventricle of the Heart in 3-D+t MRI Data Using an Optimized Nonrigid Temporal Model. *IEEE Trans. on Medical Imaging* 27(2), 195–203 (2008)
7. Pluempitwiriyawej, C., Moura, J.M.F., Lin Wu, Y.-J., Ho, C.: STACS: new active contour scheme for cardiac MR image segmentation. *IEEE Trans. on Medical Imaging* 24(5), 593–603 (2005)

8. Sun, W., Çetin, M., Chan, R., Reddy, V., Holmvang, G., Chandar, V., Willsky, A.S.: Segmenting and tracking the left ventricle by learning the dynamics in cardiac images. In: Christensen, G.E., Sonka, M. (eds.) IPMI 2005. LNCS, vol. 3565, pp. 553–565. Springer, Heidelberg (2005)
9. Kausa, M.R., von Berga, J., Weesea, J., Niessenb, W., Pekar, V.: Automated segmentation of the left ventricle in cardiac MRI. *Medical Image Analysis* 8(3), 245–254 (2004)
10. Paragios, N.: A level set approach for shape-driven segmentation and tracking of the left ventricle. *IEEE Trans. on Medical Imaging* 22(6), 773–776 (2003)
11. Fritscher, K.D., Pilgram, R., Schubert, R.: Automatic cardiac 4D segmentation using level sets. In: Frangi, A.F., Radeva, P.I., Santos, A., Hernandez, M. (eds.) FIMH 2005. LNCS, vol. 3504, pp. 113–122. Springer, Heidelberg (2005)
12. Andreopoulos, A., Tsotsos, J.K.: Efficient and Generalizable Statistical Models of Shape and Appearance for Analysis of Cardiac MRI. *Medical Image Analysis* 12(3), 335–357 (2008)
13. Jolly, M.-P.: Automatic Segmentation of the Left Ventricle in Cardiac MR and CT Images. *International Journal of Computer Vision* 70(2), 151–163 (2006)
14. Zambal, S., Hladůvka, J., Bühler, K.: Improving segmentation of the left ventricle using a two-component statistical model. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 151–158. Springer, Heidelberg (2006)
15. Zhuang, X., Rhode, K.S., Arridge, S.R., Razavi, R., Hill, D.L.G., Hawkes, D.J., Ourselin, S.: An atlas-based segmentation propagation framework using locally affine registration – application to automatic whole heart segmentation. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 425–433. Springer, Heidelberg (2008)
16. Ben Ayed, I., Li, S., Ross, I.: A Statistical Overlap Prior for Variational Image Segmentation. *International Journal of Computer Vision* (available online) (in press, 2009)
17. Zhang, T., Freedman, D.: Improving performance of distribution tracking through background mismatch. *IEEE Trans. on Pattern Anal. and Machine Intell.* 27(2), 282–287 (2005)
18. Freedman, D., Zhang, T.: Active contours for tracking distributions. *IEEE Transactions on Image Processing* 13(4), 518–526 (2004)
19. Boykov, Y., Funka-Lea, G.: Graph Cuts and Efficient N-D Image Segmentation. *Int. J. of Computer Vision* 70(2), 109–131 (2006)
20. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. on Pattern Anal. and Machine Intell.* 26(9), 1124–1137 (2004)
21. Rother, C., Kolmogorov, V., Minka, T., Blake, A.: Cosegmentation of Image Pairs by Histogram Matching–Incorporating a Global Constraint into MRFs. In: CVPR(1), pp. 993–1000 (2006)
22. Boykov, Y., Kolmogorov, V.: Computing geodesics and minimal surfaces via graph cuts. In: ICCV, pp. 26–33 (2003)