LEFT VENTRICLE SEGMENTATION FROM CARDIAC MRI COMBINING LEVEL SET METHODS WITH DEEP BELIEF NETWORKS

Tuan Anh Ngo *Gustavo Carneiro*

Australian Centre for Visual Technologies
University of Adelaide, Australia

ABSTRACT

This paper introduces a new semi-automated methodology combining a level set method with a top-down segmentation produced by a deep belief network for the problem of left ventricle segmentation from cardiac magnetic resonance images (MRI). Our approach combines the level set advantages that uses several a priori facts about the object to be segmented (e.g., smooth contour, strong edges, etc.) with the knowledge automatically learned from a manually annotated database (e.g., shape and appearance of the object to be segmented). The use of deep belief networks is justified because of its ability to learn robust models with few annotated images and its flexibility that allowed us to adapt it to a top-down segmentation problem. We demonstrate that our method produces competitive results using the database of the MICCAI grand challenge on left ventricle segmentation from cardiac MRI images, where our methodology produces results on par with the best in the field in each one of the measures used in that challenge (perpendicular distance, Dice metric, and percentage of good detections). Therefore, we conclude that our proposed methodology is one of the most competitive approaches in the field.

1. INTRODUCTION

The leading cause of death in the world is cardiovascular disease [17], and one of the best methods to improve the survival rate is based on the early diagnosis using imaging technologies. Over the last few years, there has been significant developments of imaging technologies that have enabled physicians to analyze better some parameters to assess the health of the heart (e.g., ejection fraction, wall motion, etc.). One of the current dominant imaging technologies is the cardiac magnetic resonance imaging (MRI), using the short axis view, but the segmentation of the left ventricle (LV) is a crucial first step in this analysis. Manual LV segmentation is still the standard clinical practice, but it suffers from operator bias, poor reproducibility and relatively large inter- and intra-observer variability. One possible solution to these issues is the development of a (semi-)automated LV segmentation. However, there are a few issues that must be solved before it can be accepted in a clinical setting, such as [8]: 1) the precise segmentation of the LV when the outflow tract is present, reducing the strength of edge information; and 2) the variability of the LV shape across slices, phases and patients.

Automated and semi-automated LV segmentation from cardiac MRI images has been an intensive area of research, and it is possible to classify current techniques into three categories: 1) active contour models, 2) machine learning models, and 3) combined active contour and machine learning models. Active contours with explicit contour representation [13] segments an object by minimizing an energy function with internal constraints denoting contour smoothness, and external constraints usually represented by strong edges. The use of implicit contour representation with active contours, known as the level set method [20], allowed the implementation of a similar optimization function directly on the fixed Cartesian grid without having to parameterize the curve representing the segmentation, which also allowed the delineation of objects that change topology. The main issue with active contour models is the fact that the energy function must contain all terms that are needed to segment an object, requiring substantial hand-engineering of functions (and their parameters). It is important to mention that the task of writing such energy function addressing all possible shape and appearance variations of the LV from cardiac MRI is a complicated, if not impossible task. The introduction of machine learning models has addressed exactly this issue with the use of an annotated training set to automatically learn the parameters of statistical appearance and shape models [4]. However, the automatic learning of the model parameters either requires a large training set or models with relatively low capacity, so it is clear that it would be advantageous to include some of the segmentation priors used by active contour models in order to alleviate the model learning issues. This is the idea behind the combination of active contour and machine learning models, such as the combination of Markov random field with active contour models [9], or conditional random field with active contour models [2,23]. Nevertheless, the proper training of the random fields is usually considered intractable, but there has been some progress regarding the implementation of an efficient training of random field models [22]. Our proposal focus on solving the LV segmentation problem from cardiac MRI images using this combination of active contour and machine learning models, but the machine learning model used is based on deep belief network [7], which offers: 1) efficient training and inference approaches; 2) advantages in terms of model flexibility (which means that it can be easily adapted to different types of classification and segmentation problems); and 3) relatively small annotated training sets for a robust estimation of model parameters.

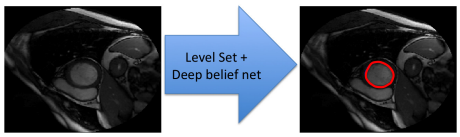

Fig. 1. Proposed methodology for cardiac MRI segmentation.

In this paper we propose a novel semi-automated LV segmentation from MRI images. Our method uses a level set method, which has a constraint provided by an LV segmentation estimated by a deep belief network (DBN), as depicted in Fig. 1. The main novelties of our paper are:

- The combination of level sets and DBN for image segmentation,
- The DBN segmentation model that produces a segmenta-

tion directly from raw pixels.

We test our approach on the MICCAI grand challenge on left ventricle segmentation from cardiac MRI images using the 15 training, 15 validation and 15 testing datasets [21]. The results show that our approach produces results on par with the best in the field in each one of the measures used in that challenge (perpendicular distance, Dice metric, and percentage of good detections). Since other approaches never achieve the top results in all three measures, we conclude that our proposed method is one of the most competitive approaches in the field.

2. PROBLEM FORMULATION

In this section, we first explain the level set method used in our methodology and then we explain how the DBN segmentation model is formulated. Assume that a database of annotated cardiac MRI images is denoted by $\mathcal{D} = \{(I, \mathbf{c})_i\}_{i=1}^{|\mathcal{D}|}$, where $I : \Omega \rightarrow \mathbb{R}$ represents an image (with $\Omega \subseteq \mathbb{R}^2$ denoting the pixel address space) and $\mathbf{c} : [0, 1] \rightarrow \Omega$ denotes the explicit contour representation of the segmentation.

2.1. Distance Regularized Level Set (DRLS)

In the original level set formulation [20], the evolution of the level set function tends to develop irregularities in the signed distance function, which are fixed with periodic re-initializations of this distance function, presenting practical and theoretical issues [14], such as numerical problems and the scheduling of re-initializations. By including a term in the level set formulation that guarantees that the signed distance function remains regularized, Li et al. [14] eliminates the need for re-initializations and consequently the issues involved with them. Because of this advantage, we use this level set implementation, and the implicit contour representation is the zero level set of a signed distance function $\phi : \Omega \rightarrow \mathbb{R}$, as in $\mathcal{C} = \{\mathbf{x} \in \Omega | \phi(\mathbf{x}) = 0\}$, with $\phi(\mathbf{x})$ denoting the signed Euclidean distance from $\mathbf{x}$ to $\mathcal{C}$ taking negative values for points inside the contour and positive values outside this contour.

In order to find the contour $\mathcal{C}$, we define an energy functional

$$\mathcal{E}(\phi, \phi_{\text{DBN}}, L) = \mu \mathcal{R}_p(\phi) + \mathcal{E}_{\text{ext}}(\phi, \phi_{\text{DBN}}, L), \tag{1}$$

where $\mathcal{R}_p(\phi) = \int_\Omega p(|\nabla\phi|) d\mathbf{x}$ is the level set regularization term, $\mu > 0$ is a constant, and $\mathcal{E}_{\text{ext}}$ is the external energy term with ϕ_{DBN} representing the shape produced by the DBN model explained later in Sec. 2.2. By solving the curve evolution equation $\frac{\partial \phi}{\partial t} = -\frac{\partial \mathcal{E}}{\partial \phi}$ using a temporal variable $t \in [0, \infty)$ (note that we assume that ϕ is parameterized not only by $\mathbf{x}$, but also by t), we can find the zero level set using an initial guess ϕ_0. The evolution of the time-dependent function ϕ follows the steepest descent direction of the energy functional $\mathcal{E}$. The derivative of the regularization term $\frac{\partial \mathcal{R}_p}{\partial \phi} = -\text{div}(d_p(|\nabla\phi|)\nabla\phi)$, where $d_p(|\nabla\phi|) = p'(|\nabla\phi|)/|\nabla\phi|$ and div is the divergence operator. The idea of using a distance regularizer is based on the fact that $p(|\nabla\phi|)$ should have two local minima at $|\nabla\phi| = \{0, 1\}$, which maintains the signed distance property $|\nabla\phi| = 1$ in a vicinity of the zero level set, and $|\nabla\phi| = 0$ at locations far away from the zero level set [14].

The second term in (1) is defined as:

$$\mathcal{E}_{\text{ext}}(\phi, \phi_{\text{DBN}}, L) = \lambda \mathcal{E}_{\text{lng}}(\phi) + \alpha \mathcal{E}_{\text{area}}(\phi) + \gamma \mathcal{E}_{\text{shp}}(\phi, \phi_{\text{prior}}, L) + \beta \mathcal{E}_{\text{shp}}(\phi, \phi_{\text{DBN}}, L), \tag{2}$$

where $\lambda, \gamma, \beta > 0$, $\alpha \in \mathbb{R}$, $\mathcal{E}_{\text{lng}}(\phi) = \int_\Omega g\delta(\phi)|\nabla\phi| dx$ returns small value for g and $|\nabla\phi|$ at edges (with $\delta(\phi)$ denoting the Dirac delta function), $\mathcal{E}_{\text{area}}(\phi) = \int_\Omega gH(-\phi) dx$ speeds up the level set evolution process by quickly increasing or decreasing the contour area (depending on the value of α, with $H(-\phi) = 1$ when $\phi < 0$, and $H(-\phi) = 0$ otherwise), $g = \frac{1}{1+|\nabla G_\sigma * I|^2}$ is a function that is small at edges and close to one elsewhere (with ∇G_σ being the gradient of a Gaussian kernel, and $*$ being the convolution operator). Also in (2), we add two energy terms to take into consideration the prior geometrical shape [5,16] learned from the annotated training set and the result of the segmentation produced by the DBN, with $\mathcal{E}_{\text{shp}}$ defined as:

$$\mathcal{E}_{\text{shp}}(\phi, \phi_{\text{k}}, L) = \int_\Omega (\phi - \phi_{\text{k}})^2 (L + 1)^2 dx, \tag{3}$$

for $k \in \{\text{prior}, \text{DBN}\}$, where $L : \Omega \rightarrow \{-1, +1\}$ is known as the dynamic labeling function that assumes the values $+1$ or -1, indicating that the prior must be enforced or not [5], respectively. Note that the the size of the window where $L = +1$ is a rectangle of $M \times N$ pixels. In practice, this dynamic labeling defines a window around the region of interest where the object of interest is believed to be localized, and this means that initially, $L = +1$ will be around the initialization $\phi_0 < 0$, and after each iteration, $L = +1$ will be around the updated $\phi < 0$. Finally, the ϕ_{prior} is computed from the training set $\mathcal{D}$ defined above by: 1) centering the training contours $\mathbf{c}$ at the origin $(0, 0)$, and 2) defining a bounding box of size $M \times N$ pixels around the centered contours. This means that all contours will have the same center, which are represented by $\widetilde{\mathbf{c}}_i$ (for $i \in \{1, 2, ..., |\mathcal{D}|\}$). The ϕ_{prior} is then the distance function computed from $\bar{\mathbf{c}}$, which is the mean aligned contour in the $M \times N$ window, calculated as $\bar{\mathbf{c}} = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \widetilde{\mathbf{c}}_i$.

2.2. Deep Belief Network (DBN)

One of the main recent advances that has happened in machine learning is the development of deep learning techniques, consisting of a hierarchical representation that can learn complicated functions, representing several levels of abstractions [1]. One of the breakthroughs that enabled the exploration of deep learning architectures was the development of the contrastive-divergence learning algorithm [6] that could estimate reliably the parameters of these deep hierarchies, with several levels of non-linear operators. Deep learning architecture has been applied not only to classic learning problems, producing better results than competing methodologies [7], but also to new learning problems, previously too difficult to be handled by traditional machine learning methodologies. For instance, image segmentation from raw pixel data is a problem that can be effectively handled by deep learning methodologies, and we propose a solution based on DBN in this section. Moreover, this solution is used to build the distance function ϕ_{DBN} for the level set energy function (2).

Specifically, we exploit the model depicted in Fig. 2 with the following joint probability:

$$\begin{aligned} &P(I_L, \mathbf{h}_1, ..., \mathbf{h}_K, \mathbf{y}) \\ &\quad = P(\mathbf{h}_K, \mathbf{h}_{K-1}, \mathbf{y}) \left(\prod_{k=1}^{K-2} P(\mathbf{h}_{k+1} | \mathbf{h}_k) \right) P(\mathbf{h}_1 | I_L) \\ &\quad = P(\mathbf{h}_K, \mathbf{h}_{K-1}, \mathbf{y}) \left(\prod_{k=1}^{K-2} P(\mathbf{h}_k | \mathbf{h}_{k+1}) \right) P(I_L | \mathbf{h}_1) \end{aligned} \tag{4}$$

where I_L represents the raw pixel extracted from the window defined by $L = +1$ (3), $\mathbf{y} \in \{0, 1\}^{M \times N}$ represents the segmentation map of I_L, $\mathbf{h}$ denotes the hidden variables, and

$$\begin{aligned} -\log P(\mathbf{h}_K, \mathbf{h}_{K-1}, \mathbf{y}) \propto &\mathcal{E}_{\text{RBM}}(\mathbf{h}_K, \mathbf{h}_{K-1}, \mathbf{y}) \\ = &-\mathbf{b}_K^\top \mathbf{h}_K - \mathbf{a}_{K-1}^\top \mathbf{h}_{K-1} - \mathbf{a}_y^\top \mathbf{y} - \\ &(\mathbf{h}_K)^\top \mathbf{W} \mathbf{h}_{K-1} - (\mathbf{h}_K)^\top \mathbf{W}_y \mathbf{y} \end{aligned} \tag{5}$$

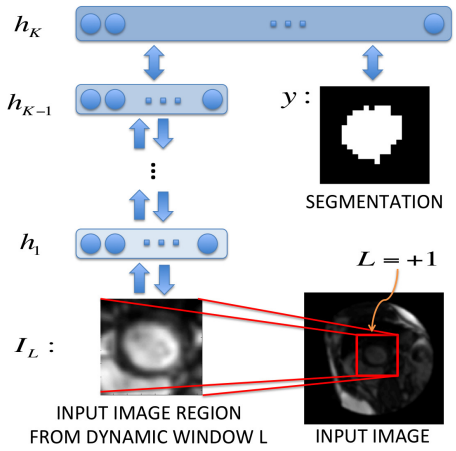

Fig. 2. Deep belief network model.

is known as a restricted Boltzmann machine (RBM) [7], where the energy function in (5) is defined by the bias vectors $\mathbf{b}_K, \mathbf{a}_{K-1}, \mathbf{a}_y$ and weight matrices $\mathbf{W}, \mathbf{W}_y$. Note from (4) and Fig. 2 that the conditional probabilities outside the top pair of layers (representing the RBM) can either be top-down or bottom-up. Also in (4), the remaining terms are related to the probability of hidden given visible variables and vice versa, which in general has the form

$$P(\mathbf{h}_{k+1}|\mathbf{h}_k) = \prod_j P(\mathbf{h}_{k+1}(j) = 1|\mathbf{h}_k),$$
$$\text{where } P(\mathbf{h}_{k+1}(j) = 1|\mathbf{h}_k) = \sigma(\mathbf{b}(j) + \mathbf{h}_k^\top \mathbf{W}(:, j))$$
$$P(\mathbf{h}_k|\mathbf{h}_{k+1}) = \prod_i P(\mathbf{h}^k(i) = 1|\mathbf{h}_{k+1}),$$
$$\text{where } P(\mathbf{h}_k(i) = 1|\mathbf{h}_{k+1}) = \sigma(\mathbf{a}(i) + \mathbf{W}(i, :)\mathbf{h}_{k+1}), \quad (6)$$

where $\sigma(x) = \frac{1}{1+e^{-x}}$, the operator (j) returns the j^{th} vector value, $(i, :)$ returns the i^{th} matrix row, and $(:, j)$ returns the j^{th} matrix column. The definition for $P(\mathbf{h}_1|I_L)$ is the same as $P(\mathbf{h}_{k+1}|\mathbf{h}_k)$ and $P(I_L|\mathbf{h}_1)$ is the same as $P(\mathbf{h}_k|\mathbf{h}_{k+1})$ in (6).

This DBN is trained layer by layer in an unsupervised way by stacking RBMs up to layer $K-1$ [7]. The error being minimized during this unsupervised training is the reconstruction error of the visible input. Note that as each layer k is added to the network, the result obtained from the first layer I_L up to layer $k-1$ is used as the "visible" input for training the RBM formed by layers $k-1$ and k. The supervised training takes place only at the highest layer K, when the segmentation $\mathbf{y}$ is provided as visible inputs to the top RBM, as depicted in Fig. 2. Note that the segmentation map is computed from the annotation $\mathbf{c}$, where pixels inside the contour are labeled "1", and outside are labeled "0". Each RBM is trained with contrastive divergence (CD) [6], which provides a maximum likelihood estimation of the network parameters (i.e., weights and biases) using a stochastic gradient descent algorithm (thus very efficient for large scale problems).

The inference process that produces the segmentation and classification is achieved by first taking an input test image at the input visible layer I_L, and then computing the probability of activation up until the layer $K-1$ using the bottom-up conditional probabilities in (6). Then the algorithm performs Gibbs sampling in order to achieve a stable value for the segmentation $\mathbf{y}$, and hidden layers $\mathbf{h}_{K-1}$ and $\mathbf{h}_K$. The initialization of this sampling process is based on the probability distribution for layer $K-1$ with $P(\mathbf{h}_{K-1}|\mathbf{h}_{K-2})$ and $\mathbf{y} = \mathbf{0}$ (for all input nodes in the segmentation layer).

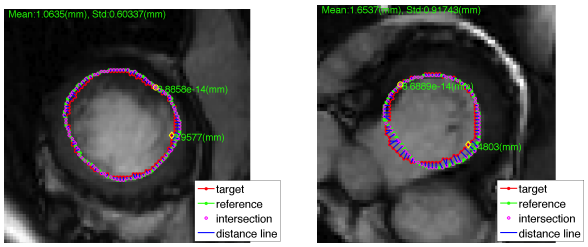

Fig. 3. Segmentation results. The left image shows a relatively simple case, but the image on the right shows a challenging case with the outflow tract present. In the legend, 'target' (red) denotes the detection and 'reference' (green) shows the manual annotation.

2.3. Segmentation Algorithm Combining DRLS and DBN

The segmentation process (detailed in Alg. 1), consists of a level set evolution explained in Sec. 2.1, where we assume that ϕ_{prior} has been computed from the training images. The first step takes the user-defined input center and scale, which forms a circle ϕ_0 that is used to initialize the level set evolution. With ϕ_0, we compute the dynamic window L in (2) by taking the region where $\phi_0 < 0$ and extend it with a fixed margin. With the window L, we also form the input window I_L for the DBN and run an inference process in order to find the map $\mathbf{y}$, as explained in Sec. 2.2. We then use $\mathbf{y}$ to compute the distance function ϕ_{DBN}. At this point, we run the level set iteration, minimizing the energy in (1), which updates the distance function ϕ.

Algorithm 1 Combined Level Set and DBN Segmentation

- Given test image I, ϕ_0 from I, and ϕ_{prior} from $\mathcal{D}$

for t = 1:T **do**

- Compute the dynamic window L from ϕ_{t-1}
- From L extract image region I_L for the DBN, and infer $\mathbf{y}$
- Compute distance function ϕ_{DBN} from map $\mathbf{y}$
- Run DRLS using $\phi_{t-1}, \phi_{\text{prior}}, \phi_{\text{DBN}}, L$ to produce updated distance function ϕ_t

end for

- Segmentation is the zero level set $\mathcal{C} = \{\mathbf{x} \in \Omega | \phi_T(\mathbf{x}) = 0\}$

3. EXPERIMENTAL RESULTS

In order to assess the performance of our algorithm, we use the MICCAI 2009 challenge database [21], consisting of three data sets (online, validation and training) obtained from the Sunnybrook Health Sciences Centre, Toronto, Canada. Each data set contains 15 cases (4 ischemic heart failures, 4 non-ischemic heart failures, 4 LV hypertrophies and 3 normal cases), thus forming 45 cardiac short axis cine-MR (SAX-MR) datasets with expert contours for the endocardial and epicardial contours in all slices at end diastole (ED) and end systole (ES) phases[1], but note that in this paper, we focus on the endocardial segmentation problem. All the images were obtained during 10-15 second breath-holds with a temporal resolution of 20 cardiac phases over the heart cycle, and scanned from the ED phase. Six to 12 SAX images were obtained from the atrioventricular ring to the apex. Three measures

[1] Only endocardial contours are available for ES.

Table 1. Quantitative experiments comparing the performance of several competing approaches on the MICCAI 2009 challenge database [21]. Each cell is formatted as "mean (standard deviation) [min value - max value]". The best result for each measure on each dataset is highlighted, and '?' means that the result is not available.

Method	Endocardial AVP	Endocardial ADM	"Good" Percentage
Training set (15 sequences)			
DBN+LS	$\mathbf{1.96(0.35)[1.43-2.55]}$	$\mathbf{0.90(0.03)[0.84-0.94]}$	$\mathbf{98.45(3.11)[91.66-100]}$
LS ONLY	$2.58(0.27)[2.28-3.08]$	$0.86(0.03)[0.79-0.91]$	$\mathbf{98.61(3.57)[88.88-100]}$
Jolly [12]	$2.09(0.53)[1.35-3.23]$	$0.88(0.06)[0.75-0.95]$	$96.93(7.59)[72-100]$
Validation set (15 sequences)			
DBN+LS	$2.22(0.46)[1.69-3.30]$	$\mathbf{0.89(0.03)[0.83-0.93]}$	$\mathbf{96.58(9.58)[63.15-100]}$
LS ONLY	$2.91(0.35)[2.41-3.73]$	$0.84(0.04)[0.77-0.90]$	$\mathbf{97.01(6.97)[73.68-100]}$
Jolly [12]	$2.26(0.59)[1.35-3.68]$	$0.88(0.04)[0.75-0.95]$	$95.62(8.83)[62-100]$
Wijnhout [24]	$2.29(0.57)[1.67-3.93]$	$\mathbf{0.89(0.03)[0.82-0.94]}$	$86.47(11)[68.4-100]$
Lu [15]	$\mathbf{2.07(0.61)[1.32-3.77]}$	$\mathbf{0.89(0.03)[0.84-0.94]}$	$72.45(18.86)[?-?]$
Huang [10]	$2.10(0.44)[?-?]$	$\mathbf{0.89(0.04)}[?-?]$	?
Marak [18]	?	$0.86(0.04)[?-?]$	?
O'Brien [19]	?	$0.81(?)[?-?]$	?
Online set (15 sequences)			
DBN+LS	$\mathbf{2.04(0.35)[1.53-2.67]}$	$\mathbf{0.90(0.04)[0.83-0.95]}$	$\mathbf{98.71(3.66)[86.66-100]}$
LS ONLY	$2.66(0.38)[2.24-3.49]$	$0.85(0.04)[0.80-0.92]$	$\mathbf{99.33(2.58)[90-100]}$
Full set (45 sequences)			
DBN+LS	$2.08(0.40)[1.43, 3.30]$	$\mathbf{0.90(0.03)[0.83, 0.95]}$	$\mathbf{97.91(6.18)[63.15, 100]}$
LS ONLY	$2.72(0.36)[2.24, 3.73]$	$0.85(0.04)[0.77, 0.92]$	$\mathbf{98.31(4.78)[73.68, 100]}$
Constantinides (full) [3]	$2.44(0.56)[1.31-4.20]$	$0.86(0.05)[0.72-0.94]$	$80(16.00)[29-100]$
Constantinides (semi) [3]	$\mathbf{1.94(0.42)[1.47-3.03]}$	$0.89(0.04)[0.80-0.96]$	$91.00(8.00)[61-100]$
Hu [8]	$2.24(0.40)[?-?]$	$0.89(0.03)[?-?]$	$91.06(9.42)[?-?]$
Huang [11]	$2.03(0.34)[?-?]$	$\mathbf{0.90(0.04)}[?-?]$	?

are computed for each data set in order to assess the performance of the proposed methodology: percentage of "good" contours, average perpendicular distance (AVP) and the average Dice metric (ADM). A contour is considered good if its AVP is less than 5mm, where each measure was computed for the annotated slices and a mean value for all the slices is given. However, AVP and ADM are computed only for good contours.

For the combined model proposed in this paper, the DBN parameters/structure and level set weights are learned using the training set, and validated with the online set. The validation set is used exclusively for testing. Note that this setup is implemented to enable a comprehensive comparison with other approaches that used the validation set for testing. With this setup, the DBN configuration achieved is: 2 hidden layers with 100 nodes in the first layer and 1000 in the second, the segmentation layer has size 20×20 in (2). The level set weights learned are: $\mu = 0.12, \lambda = 4, \alpha = -2, \gamma = 0.0005$, and $\beta = 0.001$, and the window size of $L = +1$ in (3) is $M \times N = 73 \times 73$.

In Table 1, we show quantitative results (mean, standard deviation and range) for the proposed approach combining level sets and DBN (labeled as "DBN+LS") and for the approach that uses the original DRLS formulation [14] (i.e., without the term $\mathcal{E}_{\text{shp}}$ in Eq. 2), which is labeled as "LS ONLY". The goal of comparing "DBN+LS" and "LS ONLY" is to show the influence of the DBN in the level set formulation. Moreover, we also show the results of several methodologies proposed in the literature for comparison purposes. In general, most of the approaches can be considered to be some variation of the active contour model [3,10–12,15,18], and few can be classified as machine learning methods [19,24], and one can be classified as a combination of both methods [8]. In Fig. 3, we show a couple of segmentation results produced by our approach.

From these results, we can conclude that the influence of DBN in the level set evolution is important, producing significantly more accurate results in terms of AVP and ADM, and decreasing slightly (but not significantly so) the "Good" percentage results. In general, our method is comparable or superior to all other competing methods in almost all measures, except for the AVP in the valiadation and full sets. These results place our approach among the most competitive in the field. In terms of running time per patient, the approaches vary from one minute [3,12,24] to anything in between two and three minutes [8,15]. Our approach currently takes 167.25 ± 32.71 seconds per patient (i.e., between two and three minutes), which is similar to the state of the art.

4. CONCLUSIONS AND FUTURE WORK

In this paper we present a technique for the LV segmentation in cardiac MRI images that combines level sets with deep belief networks. This is the first time such combination is proposed. Moreover, the DBN segmentation model proposed is also new, with promising results. We apply our methodology on the MICCAI 2009 challenge database [21], and the results show that the proposed methodology is one of the most accurate among the approaches that have used such database. We plan to extend our approach in several ways. First, we are currently working to make it fully automatic, with a method to detect the LV blood pool. Second, we are also extending the methodology for the detection of the epicardial contour. Finally, we also plan to work on the reduction of the running time of our approach to be closer to one minute per patient.

5. REFERENCES

[1] Y. Bengio, "Learning deep architectures for AI," *Foundations and Trends in Machine Learning*, Vol. 2. no.1 pp. 1-127, 2009. 2

[2] D. Cobzas and M. Schmidt, "Increased discrimination in level set methods with embedded conditional random fields," *CVPR*, 2009. 1

[3] C. Constantinides et al., "Fully automated segmentation of the left ventricle applied to cine MR images: Description and results on a database of 45 Subjects," *Engineering in Medicine and Biology Society (EMBC)*, 2012. 4

[4] T. Cootes, C. Taylor, D. Cooper, J. Graham et al. "Active shape models-their training and application," *Computer vision and image understanding*, 61(1):3859, 1995 1

[5] D. Cremers, "Dynamical statistical shape priors for level set-based tracking," *IEEE TPAMI* 28(8):12621273, 2006. 2

[6] M. Carreira-Perpina and G. Hinton, "On contrastive divergence learning," *Artificial Intelligence and Statistics*, vol. 2005, p. 17, 2005. 2, 3

[7] G. Hinton and R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, Vol. 313. no. 5786, pp. 504 - 507, 2006. 1, 2, 3

[8] H. Hu, H. Liu, Z. Gao and L. Huang, "Hybrid segmentation of left ventricle in cardiac MRI using gaussian-mixture model and region restricted dynamic programming," *Magnetic Resonance Imaging*, 2012. 1, 4

[9] R. Huang, V. Pavlovic, and D. Metaxas, "A graphical model framework for coupling mrfs and deformable models." *CVPR*, 2004. 1

[10] S. Huang et al., "Segmentation of the Left Ventricle from Cine MR Images Using a Comprehensive Approach," *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge*, 2009. 4

[11] S. Huang et al., "An image-based comprehensive approach for automatic segmentation of left ventricle from cardiac short axis cine MR images," *Journal of Digital Imaging* Vol. 24, No.4, pp. 598-608, (2011). 4

[12] M. Jolly, "Fully automatic left ventricle segmentation in cardiac cine MR images using registration and minimum surfaces," *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge*, 2009. 4

[13] M. Kass, A. Witkin, D. Terzopoulos, "Snakes: Active contour models," *International Journal of Computer Vision*, 1(4):321331, 1988. 1

[14] C. Li, C. Xu, C. Gui and M. Fox, "Distance regularized level set evolution and its application to image segmentation," *IEEE TIP*, Vol. 19, no. 12, pp. 3243-3254, 2010. 2, 4

[15] Y. Lu et al. "Automatic Image-Driven Segmentation of Left Ventricle in Cardiac Cine MRI, *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge*, 2009. 4

[16] M. Lynch et al. "Left-ventricle myocardium segmentation using a coupled level-set with a priori knowledge, *Computerized Medical Imaging and Graphics*, 30(4):255-262, 2006. 2

[17] J.P. Mackenbach et al., "Socioeconomic inequalities in health in 22 European countries," *N Engl J Med*, 358, pp. 24682481, 2008. 1

[18] L. Marak et al., "4D Morphological segmentation and the MICCAI LV-segmentation grand challenge, *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge*, 2009. 4

[19] S. O'Brien et al., "Segmenting the Left Ventricle in 3D Using a Coupled ASM and a Learned Non-Rigid Spatial Model, *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge*, 2009. 4

[20] S. Osher and J. Sethian, "Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations," *Journal of Computational Physics*, 79(1):1249, 1988. 1, 2

[21] P. Radau et al., "Evaluation framework for algorithms segmenting short axis cardiac MRI, *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge*, 2009. 2, 3, 4

[22] M. Szummer, P. Kohli and D. Hoiem, "Learning CRFs using Graph Cuts," *ECCV*, 2008. 1

[23] G. Tsechpenakis and D. Metaxas, "CRF-driven implicit deformable model," *CVPR*, 2007. 1

[24] J. Wijnhout et al., "LV Challenge LKEB contribution: fully automated myocardial contour detection, *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge*, 2009. 4