

Optimization for Multi-Region Segmentation of Cardiac MRI

Johannes Ulén, Petter Strandmark and Fredrik Kahl

Centre for Mathematical Sciences
Lund University, Sweden
{ulen, petter, fredrik}@maths.lth.se

Abstract. We introduce a new multi-region model for simultaneous segmentation of the left and right ventricles, myocardium and the left ventricular papillary muscles in MRI. The model enforces geometric constraints such as inclusion and exclusion between the regions, which makes it possible to correctly segment different regions even though the intensity distributions are identical. We efficiently optimize the model using Lagrangian duality which is faster and more memory efficient than current state of the art. As the optimization is based on global techniques, the resulting segmentations are independent of initialization. We evaluate our approach on two benchmarks with competitive results.

1 Introduction

Automatic segmentation of cardiac MR images is an acknowledged difficult task. Many successful approaches concentrate on segmenting the left ventricle (LV) as this part is the most interesting for diagnostic purposes. Still, quantifiable information about the cardiac function is gained from segmenting the right ventricle (RV) as well. In this paper, we use a joint model of the whole heart where the final result is improved compared to segmenting the parts independently.

This paper introduces a new mathematical model for cardiac MR segmentation. It is based on the following list of desiderata. Firstly, the human heart is composed of several interacting geometric parts — this fact should be reflected in the model. Secondly, the model should be complete in the sense that every voxel of the image should be modeled, both in terms of geometry/shape and appearance using statistical principles. This avoids many ad hoc procedures. Finally, it should be possible to estimate a global solution to the resulting optimization problem which is not dependent on a good initialization. There is always a compromise between the complexity of the model and the tractability of the optimization problem. For example, image-driven methods do not have a strong model and are typically designed to be efficient; however, they rely on good heuristics. The main contribution of this paper is that we advance the state of the art by showing that despite a rather sophisticated model, we can in an efficient manner compute solutions very close to the global optimum for the segmentation problem.

Most segmentation approaches in medical imaging rely on local descent techniques, e.g. [11,12,14], and may get stuck in local optima. It has been

shown that it is possible to apply global optimization techniques to make the segmentation more robust to poor initialization, e.g. [2,10]. We follow this trend. Our framework builds on the multi-region scheme presented in [5] where it is shown that geometric relationships, e.g. when one object is included in another, can be modeled and globally optimized via graph cuts. The key property that makes this possible is that the resulting energy minimization problem is submodular [9]. We also identify submodular relationships; however, we go beyond submodularity to enable other geometric relationships and priors to be incorporated into the model. The standard technique for solving non-submodular energies of this type is roof duality (RD) [15]. However, the method is quite memory intensive and may fail in giving a complete segmentation without time-consuming probing. Instead we develop a Lagrangian dual approach that uses half of the memory compared to RD and it always produces a segmentation. The method can easily be parallelized as described in [16]. Our algorithm is tested on two different data sets, one of which was used in the MICCAI 2009 Cardiac MR Left Ventricle Segmentation Challenge [1], on which we achieve results on par with the competing methods.

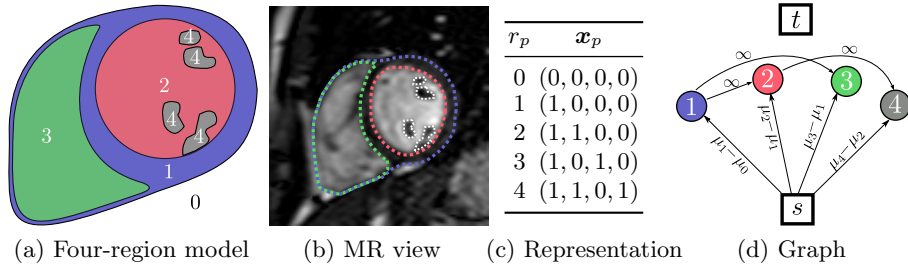


Fig. 1. (a) A constructed short-axis view showing how the heart is modeled. Region 0 is the background, region 1 contains myocardium *and* the left and right ventricular cavities. Region 2 is the left ventricular cavity and region 3 the right ventricular cavity. Region 4 is the papillary muscles of the left ventricle. (b) An example of a slice from a short-axis image acquired with MRI where all four regions have been manually delineated. (c) The Boolean representation of the four regions reflect their geometric relationships as given in (a). (d) Graph construction for one voxel. The circled number corresponds to a vertex associated with the region number. The directed arrows are the directed edges in the graph.

2 Model

In our model, the heart below the atrioventricular plane consists of four different regions as shown in Fig. 1(a) and (b). The joint model describes both the geometry of the different regions and their appearances in the MR images. An energy minimization approach is proposed in which a minimizer of the energy function gives the desired segmentation.

Let \mathcal{L} be the set of region indices and let \mathcal{P} be the set of voxel indices. Each voxel should be assigned a region index $r_p \in \mathcal{L}$. We introduce $\mathbf{x} \in \mathbb{B}^{|\mathcal{L}| \times |\mathcal{P}|}$, where $\mathbb{B} = \{0, 1\}$ and \mathbf{x} is indexed as \mathbf{x}_p^i with $i \in \mathcal{L}$ and $p \in \mathcal{P}$. Further, \mathbf{x}^i represents all Boolean variables associated with region i and \mathbf{x}_p represents all Boolean variables associated with voxel p . Each voxel in the image is represented by $|\mathcal{L}|$ Boolean variables, making it possible to directly encode geometric relationships between regions, like inclusion and exclusion. Fig. 1(c) shows the correspondence between r_p and \mathbf{x}_p .

The energy function to be minimized is $E(\mathbf{x}) = D(\mathbf{x}) + V(\mathbf{x}) + W(\mathbf{x})$, whose three components are, in order, the unary terms, the pairwise terms (regularization) and the geometric interaction terms. For every voxel p , the unary terms introduce a cost for each labeling of \mathbf{x}_p :

$$D(\mathbf{x}) = \sum_{p \in \mathcal{P}} \sum_{i \in \mathcal{L}} D_p^i(\mathbf{x}_p^i). \quad (1)$$

The pairwise terms use a connectivity \mathcal{N} to favor smooth and correctly located boundaries:

$$V(\mathbf{x}) = \sum_{i \in \mathcal{L}} \sum_{p, q \in \mathcal{N}} V_{p, q}^i(\mathbf{x}_p^i, \mathbf{x}_q^i). \quad (2)$$

The geometric interaction terms associate a cost with labeling voxel p with different combinations of the two regions i and j and are used either to attract or repel different regions to each other:

$$W(\mathbf{x}) = \sum_{p \in \mathcal{P}} \sum_{\substack{i, j \in \mathcal{L} \\ i \neq j}} W_p^{i, j}(\mathbf{x}_p^i, \mathbf{x}_p^j). \quad (3)$$

Unary terms. The unary terms are constructed from the probability, P , of each voxel belonging to any of the four defined regions, cf. Fig. 1. We define $\mu_i(p) = -\log(P(r_p = i))$, for voxel p and region i . The probability P is estimated from training data under the assumption that the spatial location and the intensity of a voxel are independent. We split the possible locations into four categories: left ventricle, right ventricle, myocardium and background. Similarly the intensity is split into three categories: blood, muscle and background.

The spatial distribution is estimated by first resizing each heart in the training data to the same size by linear interpolation. Then a binary mask is constructed for each category and each heart. The masks are enlarged and smoothed and then they are all added together constructing the final probability mask.

The intensity distribution for each region is estimated by collecting all intensities from the examples in the training data. The histogram of intensities is then smoothed and a distribution is constructed. For both the location and intensity probability a lowest probability is set, in order to capture occurrences unseen in the training data. An example of the final μ_i 's can be found in Fig. 2.




Fig. 2. Example of μ_i for the slice shown in (a). Recall that $\mu_i(p) = -\log(P(r_p = i))$. A lower intensity corresponds to higher probability.

Having estimated μ_i , the unary terms $D_p^i(\mathbf{x}_p^i)$ need to be constructed to reflect Fig. 1(c):

$$\begin{aligned} D_p^1(1) &= \mu_1(p) - \mu_0(p), & D_p^2(1) &= \mu_2(p) - \mu_1(p), \\ D_p^3(1) &= \mu_3(p) - \mu_1(p), & D_p^4(1) &= \mu_4(p) - \mu_2(p), \end{aligned} \quad (4)$$

and $D_p^i(0) = 0$ for all i and p . For example, region 4 is according to Fig. 1(c) represented as $\mathbf{x}_p = (1, 1, 0, 1)$. The cost of this \mathbf{x}_p is $(\mu_1(p) - \mu_0(p)) + (\mu_2(p) - \mu_1(p)) + (\mu_4(p) - \mu_2(p)) = \mu_4(p) - \mu_0(p)$.

Pairwise terms. The regularization weights are chosen differently for each region in a method related to the discussion in [6]. For each region i we choose the pairwise terms as:

$$V_{p,q}^i(p, q) = w_{p,q} \frac{1}{1 + \beta (P(r_p = i) - P(r_q = i))^2}, \quad (5)$$

where β can be used to tune the regularization. The neighborhood \mathcal{N} for the regularization is chosen as 18-connectivity. The multipliers $w_{p,q}$ give different weights to different types of edges. One common choice is $w_{p,q} = 1/\text{dist}(p, q)$; however, we instead use the arguably more correct way described in [3] based on solid angles. Since MRI have anisotropic resolution it is very important to take that into consideration both when calculating the distance between voxels and when using the method from [3].

Geometric interaction terms. In our model region 1 contains both region 2 and region 3. This is modeled by the use of geometric interaction terms as $W_p^{1,2}(0, 1) = \infty$ and $W_p^{1,3}(0, 1) = \infty, \forall p \in \mathcal{P}$. Furthermore, the left ventricular papillary muscle must be inside the left ventricle. This is modeled as $W_p^{2,4}(0, 1) = \infty, \forall p \in \mathcal{P}$, see Fig. 1(d). By this construction any \mathbf{x}_p -labeling not listed in Fig. 1(c) will have ∞ cost except for $\mathbf{x}_p = \{1, 1, 1, 1\}$ and $\mathbf{x}_p = \{1, 1, 1, 0\}$, these two cases are handled in Section 3.

User interaction. The method needs the user to select which slices to be segmented and it also needs one click in the center of the right and left ventricle in one slice. The two center points are used to roughly align the hearts in order to get good spatial statistics. The algorithm can handle slices lacking any of the regions.

Pre-processing. Badly captured MRI are identified by looking at the distribution of the intensities. If there are multiple peaks in the histogram close to each other for the lower intensities, the image is assumed to be too bright and the intensity distribution is shifted to fit an average histogram.

Post-processing. In all ground truth data we come across, only the left ventricular epicardium is delineated. In our model we do not have this restriction — we segment the full myocardium. In order to compare our results with the ground truth we must remove all myocardium which is not part of the left ventricular epicardium. To do this, the thickness of the septum is approximated as the shortest distance between the left and right ventricles in the resulting segmentation. Then outlying myocardium is removed based on this thickness approximation, cf. Fig. 4(a).

We also assume that the left ventricle and the myocardium are convex. The resulting segmentation is taken as the convex hull in each slice.

The regularization can sometimes make the segmentation miss the most apical slice. In this case either the segmentation from the same slice at another time step or, if it is not available, the segmentation from a more basal slice is shrunk and fitted at the bottom.

3 Solving the Optimization Problem

The energy function is minimized using graph cuts by associating each binary variable with a vertex in an s-t graph. The global minimum is then found as the minimum cut of the graph. We use the maximum flow implementation [4] to compute minimum cuts.

It is well-known that an energy $E(\mathbf{x})$ can be minimized exactly by finding the minimum cut of a graph as long as all energy terms are submodular [9], but this is not the case for our energy. The unary and pairwise terms are of standard type and well-known to be submodular. All geometric interactions except the separation of region 2 and 3 can be represented with a submodular function. The corresponding graph construction for one voxel is shown in Fig. 1(d). As illustrated in the figure, we want region 1 to contain regions 2 and 3 and at the same time we want region 2 and 3 to be separated. Unfortunately, this last constraint leads to a frustrated cycle and cannot be modeled by a submodular energy function, see [5].

3.1 Using the Lagrangian Dual

Minimizing $E(\mathbf{x})$ is a difficult problem, since it contains the non-submodular separation of region 2 and 3. If we let $E'(\mathbf{x})$ be our energy without the separation constraint, E' will be easy to minimize. Our model has four different kinds of vertices, $(\mathbf{x}^1, \mathbf{x}^2, \mathbf{x}^3, \mathbf{x}^4) = \mathbf{x}$, where the superindex denotes the region label. The separation constraints dictate that \mathbf{x}^2 and \mathbf{x}^3 cannot be equal to 1 at the same

time. Adding this constraint gives us the new problem

$$\begin{aligned} \min_{\mathbf{x}} \quad & E'(\mathbf{x}) \\ \text{subject to} \quad & \mathbf{x}^2 + \mathbf{x}^3 \leq 1. \end{aligned} \tag{6}$$

We note that this problem can be solved as an integer programming problem. However, this is not tractable due to the large number of variables. Instead, we look at the Lagrange dual problem:

$$\begin{aligned} \max_{\boldsymbol{\lambda}} \quad & d(\boldsymbol{\lambda}) \\ \text{subject to} \quad & \boldsymbol{\lambda} \geq 0, \end{aligned} \tag{7}$$

where $d(\boldsymbol{\lambda}) = \min_{\mathbf{x}} (E'(\mathbf{x}) + \boldsymbol{\lambda}^T (\mathbf{x}^2 + \mathbf{x}^3 - 1))$ is the Lagrange dual function. Let d^* denote the optimal value for (7) and p^* the optimal value for (6). By weak duality we then have that $d^* \leq p^*$.

The Lagrange dual function d is always concave. However, it is not differentiable in general, which precludes gradient ascent methods for its maximization. We can, however, use the projected supergradient method [13].

This looks very similar to a gradient ascent method but has some key differences. Specifically, the method is easy to implement, but in general has worse convergence properties than first-order gradient-based methods. We refer the reader to [13,16] for details.

Definition 1. A supergradient to a function f at a point \mathbf{x}_0 is a vector \mathbf{v} fulfilling $f(\mathbf{x}) - f(\mathbf{x}_0) \leq (\mathbf{x} - \mathbf{x}_0)^T \mathbf{v}$, for every point \mathbf{x} .

Lemma 1 (from [16]). Let $\boldsymbol{\lambda}$ be given and let \mathbf{x}^* be the optimal solution to $d(\boldsymbol{\lambda}) = \min_{\mathbf{x}} (f_1(\mathbf{x}) + \boldsymbol{\lambda}^T f_2(\mathbf{x}))$. Then $f_2(\mathbf{x}^*)$ is a supergradient to f at $\boldsymbol{\lambda}$.

Proof. For any $\boldsymbol{\lambda}$ it holds that

$$\begin{aligned} d(\boldsymbol{\lambda}) &\leq f_1(\mathbf{x}^*) + \boldsymbol{\lambda}^T f_2(\mathbf{x}^*) = f_1(\mathbf{x}^*) + \boldsymbol{\lambda}_0^T f_2(\mathbf{x}^*) + (\boldsymbol{\lambda} - \boldsymbol{\lambda}_0)^T f_2(\mathbf{x}^*) \\ &= \min_{\mathbf{x}} (f_1(\mathbf{x}) + \boldsymbol{\lambda}_0^T f_2(\mathbf{x})) + (\boldsymbol{\lambda} - \boldsymbol{\lambda}_0)^T f_2(\mathbf{x}^*) = d(\boldsymbol{\lambda}_0) + (\boldsymbol{\lambda} - \boldsymbol{\lambda}_0)^T f_2(\mathbf{x}^*). \end{aligned}$$

The projected supergradient method is very simple. Let $\boldsymbol{\lambda}_0$ be an initial guess of the optimal value of the concave function d . Then, in each step i a new possible solution $\boldsymbol{\lambda}_{i+1}$ is calculated as $\boldsymbol{\lambda}_{i+1} = [\boldsymbol{\lambda}_i + \tau_i \mathbf{v}_i]^+$, where \mathbf{v}_i is any supergradient to d at $\boldsymbol{\lambda}_i$, τ_i is a step-length and $[\cdot]^+$ is a projection onto the feasible set $\{\boldsymbol{\lambda} \geq 0\}$.

From Lemma 1 we directly choose a supergradient for a given $\boldsymbol{\lambda}$ as $\mathbf{v} = (\mathbf{x}^2)^* + (\mathbf{x}^3)^* - \mathbf{1}$, where $(\mathbf{x}^i)^*$ are the vertices belonging to label i for the optimal solution of $d(\boldsymbol{\lambda})$.

In each step, the optimal solution \mathbf{x}^* for a chosen $\boldsymbol{\lambda}_i$ can be calculated via a minimum graph cut. Furthermore, as the edges will be very similar in each step, the graph structure can be reused reducing the running time [8].

In the experiments the step size from [16] is used. Since supergradient methods do not guarantee improved value in each step the best solution thus far is always saved and once the *relative duality gap* $(p - d)/|p|$ is small enough the algorithm terminates. Here p and d are the currently best primal and dual energies.




Fig. 3. Example segmentation from LUND.

LV endo.	LV epi.	RV endo.	LV endo.	LV epi.	RV endo.
ED 0.96 ± 0.02	0.93 ± 0.03	0.91 ± 0.07	ED 0.62 ± 0.12	0.90 ± 0.03	0.57 ± 0.14
ES 0.87 ± 0.05	0.88 ± 0.05	0.80 ± 0.11	ES 0.47 ± 0.25	0.86 ± 0.04	0.42 ± 0.14
(a) Full multi-region model.			(b) Each region segmented separately.		

Table 1. Results in the Dice metric for LUND reported as mean \pm one standard deviation. “ED” is end diastole and “ES” is end systole. Note that the multi-region model has a huge influence on the segmentation results.

4 Experiments

The segmentation is only performed on the slices of the heart which are fully below the atrioventricular plane. The quality of the segmentation is measured by the *Dice metric*, which is given by $2|A \cap B| / (|A| + |B|)$, where A and B are the ground truth and the computed segmentation, respectively. The algorithm is evaluated on two data sets: LUND and SUNNYBROOK. Each data set is trained and evaluated separately.

LUND consists of cine short-axis steady state free precession MR images of 62 healthy normals captured on a Philips Interera CV 1.5T with five channel cardiac synergy coil. Each heart has the left and right ventricular endocardium and the left ventricular epicardium manually delineated by an expert. The data set is split into two equally sized parts, one used for training and one used for evaluation. Results are given in Table 1(a) and an example segmentation in Fig. 3. We also evaluate three clinical parameters. The left ventricular mass has an error of 15.6 ± 11.5 g. The left and right ventricular ejection fraction errors are $5.6 \pm 2.9\%$ and $7.1 \pm 5.2\%$, respectively.

We also compare our method to a simplified version where we run the segmentation for each region separately, see Table 1(b). Without the complete multi-region model, the localization of the ventricles becomes very difficult and the blood pools are often overestimated. Two examples where the multi-region model improves the segmentation are given in Fig. 4.




Fig. 4. Examples of how modeling multiple regions improve the segmentation of the ventricular epi- and endocardium. The color scheme is the same as in Fig. 1.

Our model was also optimized with roof duality (RD) [15]. Our method was constantly faster than RD and at the same time giving a very small relative duality gap. The small gap gives us certificates that the solutions are very close to (and in many cases exactly) the global minimum, see Table 2.

SUNNYBROOK consists of 30 patients with different heart diseases and is split up into two equally sized parts, one for training and one for evaluation. The data set was used in the 2009 MICCAI segmentation challenge [1]. SUNNYBROOK lacks ground truth for the right ventricles, so this was manually constructed by a non-expert. Therefore, this ground truth was only used for training and not for evaluation. The results given by the evaluation code used in the challenge is given in Table 3 along with results from competing methods. In the challenge, the Dice metric is calculated per slice and averaged over all slices.

Note that the small training data of SUNNYBROOK gives our method a disadvantage as there just 15 hearts spanning over three different diseases and one group of normals. Image-driven methods do not suffer from the small training set as they do not need to be trained.

Method	Mean	Method	Mean	After 25 iterations	
Our	46 ± 27	Our	30 ± 27	SUNNYBROOK	0.0014 ± 0.0033
RD	6109 ± 12451	RD	1934 ± 7984	LUND	0.000055 ± 0.0021
(a) SUNNYBROOK [1]		(b) LUND		(c) Relative duality gap	

Table 2. Running time in seconds per heart (ED and ES) and relative duality gap of our algorithm after 25 iterations run on an Intel i5 2500K CPU. For RD we canceled any computation taking longer than 12 hours.

Method	Dice		LV Mass (g)	LV ejection fraction (%)
	LV endo.	LV epi.		
Our	0.86 \pm 0.05	0.92 \pm 0.02	27.1 \pm 28.3	12.5 \pm 8.7
Marák et al.	0.86 \pm 0.04	0.93 \pm 0.01	23 \pm ?	14 \pm ?
Lu et al.	0.89 \pm 0.03	0.94 \pm 0.02	21.6 \pm 14.6	8.08 \pm 5.06
Wijnhout et al.	0.89 \pm 0.03	0.93 \pm 0.01	28.7 \pm 18.7	7.02 \pm 4.78
Casta et al.	?	0.93 \pm ?	†	?
O’Brien et al.	0.81 \pm ?	0.91 \pm ?	?	?
Constantinides et al.	0.89 \pm 0.04	0.92 \pm 0.02	†	†
Huang et al.	0.89 \pm 0.04	0.94 \pm 0.01	?	?
Jolly	0.88 \pm 0.04	0.93 \pm 0.02	31.8 \pm 17.7	8.35 \pm 5.78

Table 3. Results for SUNNYBROOK. “?” means not reported in the corresponding paper. “†” means that the result is not directly comparable. Mass and ejection fraction is reported as difference between manual and automatic value.

5 Future work and limitations

Extending with one more region. It is possible to extend the model to also include papillary muscles in the right ventricle; we need only to introduce one more variable per voxel. If we let $\mu_5 = P(r_p = 5)$ for this new region and follow the notation in Fig. 1(d) we need only to add one vertex corresponding to the new variable and two edges: one s-t edge with value $\mu_5 - \mu_3$ and one edge going from region 5 to region 3 with ∞ weight. Initial experiments gave worse results for both the right ventricle and myocardium segmentation with the added region. The new region had a tendency to overflow into the septum since this would give region 3 a rounder shape giving a lower regularization cost. We have not yet found a good way of tackling this without relying too heavily on heuristics.

Short- and long-axis images. The LUND data set was manually delineated using both short- and long-axis images. For a number of hearts the most basal slice for the short-axis images containing the left ventricular cavity also cut through to the atrium. For these slices it was hard or even impossible to even manually delineate the left ventricle solely based on information from the short-axis images. When the ground truth was produced, long-axis images were used to properly segment them. It would be desirable for our algorithm to incorporate information from long-axis images as well so we could handle these few slices as well.

Reducing memory usage. In the current implementation we use a standard max-flow implementation [4], but the structure of the graph is highly repetitive. For instance, all geometric interaction terms are equal and they need not be explicitly stored in the graph. Also, if we were to use a pairwise term based on voxel intensity we would just need to save the pairwise terms in one “layer” reducing the memory used by the pairwise terms to 1/4.

6 Concluding Discussion

We have demonstrated that it is possible to apply global optimization techniques for segmentation of cardiac MRI using a sophisticated model of the heart. The model is optimized with a new method which is both fast and memory effective. The added complexity of the model is motivated by improvements in the segmentation results.

Acknowledgments We thank the Cardiac MR group at the University Hospital of Lund for providing us with the LUND data set and expert delineations. We used Segment [7] to read the DICOM images.

References

1. MICCAI Cardiac MR Left Ventricle Segmentation Challenge (2009), http://smial.sri.utoronto.ca/LV_Challenge/Home.html
2. Boykov, Y., Jolly, M.: Interactive organ segmentation using graph cuts. In: MICCAI (2000)
3. Boykov, Y., Kolmogorov, V.: Computing geodesics and minimal surfaces via graph cuts. In: ICCV (2003)
4. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. PAMI (2004)
5. Delong, A., Boykov, Y.: Globally optimal segmentation of multi-region objects. In: ICCV (2009)
6. Grady, L., Jolly, M.: Weights and topology: A study of the effects of graph construction on 3d image segmentation. In: MICCAI (2008)
7. Heiberg, E., et al.: Design and validation of Segment- freely available software for cardiovascular image analysis. BMC Medical Imaging 10(1), 1 (2010)
8. Kohli, P., Torr, P.: Dynamic graph cuts for efficient inference in markov random fields. PAMI (2007)
9. Kolmogorov, V., Zabih, R.: What energy functions can be minimized via graph cuts? PAMI (2004)
10. Lin, X., Cowan, B., Young, A.: Model-based graph cut method for segmentation of the left ventricle. In: IEEE-EMBS (2006)
11. Lorenzo-Valdés, M., et al.: Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm. Medical Image Analysis 8(3), 255–265 (2004)
12. Mitchell, S., et al.: Multistage hybrid active appearance model matching: segmentation of left and right ventricles in cardiac MR images. IEEE Transactions on Medical Imaging 20(5), 415–423 (2002)
13. Nesterov, Y.: Introductory Lectures on Convex Optimization. Kluwer Academic Publishers (2004)
14. Paragios, N.: A variational approach for the segmentation of the left ventricle in cardiac image analysis. Int. Journal Computer Vision 50(3), 345–362 (2002)
15. Rother, C., Kolmogorov, V., Lempitsky, V., Szummer, M.: Optimizing binary MRFs via extended roof duality. In: CVPR (2007)
16. Strandmark, P., Kahl, F.: Parallel and Distributed Graph Cuts by Dual Decomposition. In: CVPR (2010)