

Right Ventricle Segmentation with Probability Product Kernel Constraints

Cyrus M.S. Nambakhsh^{1,3}, Terry M. Peters^{1,3}, Ali Islam^{1,3},
and Ismail Ben Ayed^{1,2}

¹ The University of Western Ontario, London, ON, Canada

² GE Healthcare, London, ON, Canada

³ Robarts Research Institute, London, ON, Canada

⁴ St. Joseph's Health Care, London, ON, Canada

Abstract. We propose a fast algorithm for 3D segmentation of the right ventricle (RV) in MRI using shape and appearance constraints based on probability product kernels (PPK). The proposed constraints remove the need for large, manually-segmented training sets and costly pose estimation (or registration) procedures, as is the case of the existing algorithms. We report comprehensive experiments, which demonstrate that the proposed algorithm (i) requires only a single subject for training; and (ii) yields a performance that is not significantly affected by the choice of the training data. Our PPK constraints are non-linear (high-order) functionals, which are not directly amenable to standard optimizers. We split the problem into several surrogate-functional optimizations, each solved via an efficient convex relaxation that is amenable to parallel implementations. We further introduce a scale variable that we optimize with fast fixed-point computations, thereby achieving pose invariance in real-time. Our parallelized implementation on a graphics processing unit (GPU) demonstrates that the proposed algorithm can yield a real-time solution for typical cardiac MRI volumes, with a speed-up of more than 20 times compared to the CPU version. We report a comprehensive experimental validations over 400 volumes acquired from 20 subjects, and demonstrate that the obtained 3D surfaces correlate with independent manual delineations.

1 Introduction

Quantification of right ventricle (RV) function, including ejection fraction, ventricular enlargement, aneurysms, wall motion and contraction/motion analysis can be useful in the diagnosis of various cardiovascular abnormalities. RV segmentation in 3D magnetic resonance images (MRI) is an essential step towards such quantifications, yielding RV dynamics which translate into extensive clinical information [1]. However, manual segmentation of 4D (3D+time) cardiac volumes is tedious and time-consuming. As pointed out in the recent cardiac-segmentation review in [2], RV segmentation in MRI is still acknowledged to be a difficult and completely unsolved problem, unlike left ventricle (LV) segmentation which has been intensively researched during the last decade [2]. The main difficulties arise from the complex deformations of the RV chamber in 3D, its highly variable, crescent-shaped structure and the presence of papillary muscles.

Most of the existing RV/LV segmentation algorithms, for instance, those based on statistical shape models [3], registration [4], and probabilistic-atlas classification [5], require an intensive learning from a large, manually-segmented training set. Although they can lead to outstanding performance in cases that befit the training set, these algorithms may have difficulty in capturing the substantial variations in a clinical context, with the results often being dependent on the choice of a specific training set. Furthermore, most of the existing shape-based algorithms require additional costly pose optimization procedures (rotation, translation, and scaling), which are often based on slow gradient-descent techniques. To remove the dependence on a training set, several recent cardiac image segmentation studies have attempted to build subject-specific models from a user-provided segmentation of a single 2D frame in a cardiac sequence [6,7,8,9]. Unfortunately, these solutions are designed for LV segmentation, and require an intensive user input, e.g., the algorithm in [7] uses manual segmentations of several 2D slices in one subject data. Moreover, these approaches are difficult to extend beyond the 2D case.

This study investigates rapid 3D segmentation of the right ventricle (RV) in cardiac MRI with shape and appearance constraints based on probability product kernels (PPK), which relax the need for large, manually-segmented training sets and costly pose estimation (or registration) procedures. We report comprehensive experiments, which demonstrate that the proposed algorithm (i) requires only a single subject for training; and (ii) yields a performance that is not significantly affected by the choice of the training data. Our PPK constraints are non-linear (high-order) functionals, which are not directly amenable to standard optimizers. We split the problem into several surrogate-functional optimizations, each solved via an efficient convex relaxation that is amenable to parallel implementation. We further introduce a scale variable which we optimize with fast fixed-point computations, thereby achieving scale-invariance in real-time. Our parallelized implementation on a graphics processing unit (GPU) demonstrates that the proposed algorithm can yield a real-time solution for a typical cardiac MRI volume, with a speed-up of more than 20 times in comparison to the CPU version. We report a performance evaluation over 400 volumes acquired from 20 subjects, and demonstrate that the obtained 3D surfaces correlate with independent manual delineations.

2 Formulation

The Functional: Let $I : \Omega \subset \mathbb{R}^3 \rightarrow \mathcal{Z}_I \subset \mathbb{R}$ be an image function which maps 3D domain Ω to a finite set of intensity values \mathcal{Z}_I . Let $D : \Omega \rightarrow \mathcal{Z}_D \subset \mathbb{R}$ be a function that measures the distance between each point $\mathbf{x} = (x, y, z) \in \Omega$ and a given anatomical landmark (i.e., a point) $\mathcal{O} \in \Omega$, which will be used to build a translation-invariant shape prior and to learn an intensity prior. $D(\mathbf{x}) = \|\mathbf{x} - \mathcal{O}\|$, with $\|\cdot\|$ the standard L2 norm. \mathcal{Z}_D is a finite set of distance values. \mathcal{O} is obtained from a very simple user input that amounts to the manual identification of the centroid of the RV cavity within a middle slice with a single mouse click. Let $A : \Omega \rightarrow \mathcal{Z}_A \subset \mathbb{R}$ be a function measuring the angle between the vector

pointing from each point $\mathbf{x} \in \Omega$ to \mathcal{O} and the fixed x -axis unit vector v : $A(\mathbf{x}) = \frac{\langle \mathbf{x}, \mathcal{O}, v \rangle}{\|\mathbf{x}\| \|v\|}$. Our objective functional is:

$$\begin{aligned} \hat{u} &= \arg \min_{u \in \{0,1\}} \mathcal{E}(u) \quad \text{with} \\ \mathcal{E}(u) &:= - \underbrace{\langle \mathbf{P}^I(u, \cdot), \mathbf{M}^I \rangle_\rho}_{\text{Intensity Prior}} - \underbrace{\alpha_D \langle \mathbf{P}^D(u, \cdot), \mathbf{M}^D \rangle_\rho}_{\text{Distance prior}} - \underbrace{\alpha_A \langle \mathbf{P}^A(u, \cdot), \mathbf{M}^A \rangle_\rho}_{\text{Angle prior}} \\ &\quad \underbrace{\hspace{10em}}_{\text{Shape prior}} \\ &\quad + \underbrace{\gamma \int_{\Omega} C |\nabla u| \, d\mathbf{x}}_{\text{Smoothness/Edges}} \end{aligned} \quad (1)$$

The following is a detailed description of the notations and variables that appear in the optimization problem we define in equation (1):

- $u : \Omega \rightarrow \{0, 1\}$ is a binary function, which defines a variable partition of Ω : $\{\mathbf{x} \in \Omega / u(\mathbf{x}) = 1\}$, corresponding to the target RV segment, and $\{\mathbf{x} \in \Omega / u(\mathbf{x}) = 0\}$, corresponding to the complement of the target segment in Ω .
- For image data $J \in \{I, D, A\} : \Omega \subset \mathbb{R}^3 \rightarrow \mathcal{Z}_J$, and for any binary function $u : \Omega \rightarrow \{0, 1\}$, $\mathbf{P}^J(u, \cdot)$ is a vector encoding the probability density function (pdf) of data J within the segment defined by $\{\mathbf{x} \in \Omega / u(\mathbf{x}) = 1\}$:

$$\mathbf{P}^J(u, z) = \frac{\int_{\Omega} \mathcal{K}_z(J) u \, d\mathbf{x}}{\int_{\Omega} u \, d\mathbf{x}} \quad \forall z \in \mathcal{Z}_J \quad (2)$$

with \mathcal{K}_z a Gaussian window: $\mathcal{K}_z(y) = \frac{1}{(2\pi\sigma^2)^{(1/2)}} \exp\left(-\frac{\|z-y\|^2}{2\sigma^2}\right)$, with σ the width of the window.

- $\langle f, g \rangle_\rho$ is the *probability product kernel* [10], which evaluates the affinity between two pdfs f and g :

$$\langle f, g \rangle_\rho = \sum_{z \in \mathcal{Z}} [f(z)g(z)]^\rho \quad \rho \in]0, 1], \mathcal{Z} \in \{\mathcal{Z}_I, \mathcal{Z}_D, \mathcal{Z}_A\} \quad (3)$$

The higher $\langle f, g \rangle_\rho$, the better the affinity between f and g . Notice that the PPK in (3) can be viewed as a generalization of the Bhattacharyya coefficient [7]. Minimization of the PPKs in (1) aims at finding a target region whose shape and intensity pdfs most closely match *a priori* learned models:

- \mathbf{M}^I is a model of intensity. We learn \mathbf{M}^I from intensity data within a cylinder centered at \mathcal{O} . The radius of the cylinder, d , is a free parameter which has to be fixed experimentally.
- \mathbf{M}^D and \mathbf{M}^A are models of distances and angles respectively, describing a RV shape invariant with respect to translation. We learn these models from a single training subject different from the testing subject.
- $C : \Omega \rightarrow \mathbb{R}$ is an edge-indicator function given by $C(\mathbf{x}) = \frac{1}{1 + \|\nabla I(\mathbf{x})\|}$. γ , α_D and α_A are positive constants that balance the contribution of each constraint in (1).

Introducing a Scale Variable: The shape prior in (1) is not invariant with respect to scale (or size) of the RV regions. To illustrate this, we plotted in Fig. 2 (c) the distance pdfs corresponding to the ground-truth segmentations of 20 different subjects (two different volumes for each subject). The figure demonstrates that the distance pdfs have similar Gaussian shapes, but shifted supports. This shift is due to inter-subject variations in scale (or size). To account for such shifts, we further introduce a scale variable in the model of distances: $\mathbf{M}^D(\cdot, s) : \mathcal{Z}_D \times \mathbb{R} \rightarrow [0, 1] / \mathbf{M}^D(z, s) = \mathbf{M}^D(z + s), s \in \mathbb{R}$. Thus, to account for this new variable (s), we replace the distance-based prior in (1) by a scale-dependent prior: $\langle \mathbf{P}^D(u, \cdot), \mathbf{M}^D \rangle_\rho \rightarrow \langle \mathbf{P}^D(u, \cdot), \mathbf{M}^D(\cdot, s) \rangle_\rho$. Therefore, s becomes a variable which has to be optimized along with the segmentation region. With this new variable, our problem becomes: $\{\hat{u}, \hat{s}\} = \min_{u, s} \mathcal{E}(u, s)$.

Two-step Optimization: Our model has two different types of variables, the target region described by indicator function u and the scale variable s . We therefore adopt an iterative two-step procedure, by first fixing the scale variable and optimizing the proposed functional with respect to u via convex relaxed surrogate functionals, and then optimizing over the scale variable via fixed-point computations, with u fixed.

Step 1—Optimization with Respect to the Segment Via Surrogate Functionals and Convex Relaxation: To simplify further development, let us assume that our functional contains only one probability product kernel and has the following general form:

$$\min_{u(x) \in \{0,1\}} \left\{ \mathcal{E}(u) := - \langle \mathbf{P}(u, \cdot), \mathbf{M} \rangle_\rho + \gamma \int_{\Omega} C |\nabla u| \, d\mathbf{x} \right\} \quad (4)$$

Observe that we omitted the superscripts that we defined previously for the pdfs to simplify further presentation and notations. Once the problem in (4) is solved, extension to a weighted sum of probability product kernels, as is the case in (1), becomes straightforward. Unfortunately, the probability product kernel in (4) is a non-linear (high-order) functional, which results in a difficult (non-convex) optimization problem that is not directly amenable to standard solvers. We split the problem into several surrogate-functional optimizations, each solved via an efficient convex relaxation.

Surrogate Functionals: We proceed by constructing and optimizing iteratively surrogate functionals of \mathcal{E} (whose optimization is easier than the original functional):

Definition 1. *Given a fixed labeling u^i (i is the iteration number), $\mathcal{S}(u, u^i)$ is a surrogate functional of \mathcal{E} if it satisfies the following conditions [11]:*

$$\mathcal{E}(u) \leq \mathcal{S}(u, u^i) \quad (5a)$$

$$\mathcal{E}(u) = \mathcal{S}(u, u) \quad (5b)$$

Rather than optimizing directly \mathcal{E} , we optimize the surrogate functional over the first variable at each iteration:

$$u^{i+1} = \min_u \mathcal{S}(u, u^i), \quad i = 1, 2, \dots \quad (6)$$

Using the constraints in (5a) and (5b), and by the definition of minimum in (6), we can show that the solutions in (6) yield a decreasing sequence of \mathcal{E} : $\mathcal{E}(u^i) = \mathcal{S}(u^i, u^i) \geq \mathcal{S}(u^{i+1}, u^i) \geq \mathcal{E}(u^{i+1})$. Therefore, if \mathcal{E} is lower bounded, sequence $\mathcal{E}(u^i)$ converges to a minimum of \mathcal{E} . Now, consider the following proposition ¹ :

Proposition 1. *Given a fixed $u^i : \Omega \rightarrow \{0, 1\}$, the following functional is a surrogate of functional \mathcal{E} defined in (4):*

$$\mathcal{S}(u, u^i) = -\langle \mathbf{P}(u^i, \cdot), \mathbf{M} \rangle_\rho + \int_\Omega f^i u^- \, d\mathbf{x} + \int_\Omega g^i u^+ \, d\mathbf{x} + \int_\Omega C |\nabla u| \, d\mathbf{x} \quad (7)$$

where

$$u^-(\mathbf{x}) := \begin{cases} 1 - u(\mathbf{x}), & \text{for } u^i(\mathbf{x}) = 1 \\ 0, & \text{otherwise} \end{cases}, u^+(\mathbf{x}) := \begin{cases} u(\mathbf{x}), & \text{for } u^i(\mathbf{x}) = 0 \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

and

$$f^i = \sum_{z \in \mathcal{Z}} \frac{D_{z,i,\rho} \mathcal{T}_{\mathbf{M},z}}{\int_\Omega \mathcal{T}_{\mathbf{M},z} u^i \, d\mathbf{x}}; \mathcal{T}_{\mathbf{M},z} = \mathcal{K}_z(J)\mathbf{M}(z); \quad (9)$$

$$D_{z,i,\rho} = \left[\frac{\int_\Omega \mathcal{T}_{\mathbf{M},z} u^i \, d\mathbf{x}}{\int_\Omega u^i \, d\mathbf{x}} \right]^\rho; g^i = \sum_{z \in \mathcal{Z}} \frac{D_{z,i,\rho}}{\rho \int_\Omega u^i \, d\mathbf{x}} \quad (10)$$

Convex Relaxation: Now, note that $\mathcal{S}(u, u^i)$ has a linear form, which is amenable to powerful global solvers. At each iteration, we optimize $\mathcal{S}(u, u^i)$ with the convex-relaxation technique recently developed in [12]. The optimizer in [12] is amenable to parallel implementations on graphics processing units (GPU). Therefore, it can yield real-time solutions for 3D grids.

Step 2–Fixed-point Optimization with Respect to the Scale Variable: We fix labeling variable u and optimize \mathcal{E} with respect s . Considering a variable change $z \leftarrow z - s$ and the fact that only the distance-distribution prior depends on s , we have:

$$\begin{aligned} \frac{\partial \mathcal{E}}{\partial s} &= -\alpha_D \frac{\partial \langle \mathbf{P}^D(u, \cdot), \mathbf{M}^D(\cdot, s) \rangle_\rho}{\partial s} \\ &= -\rho \alpha_D \sum_{z \in \mathcal{Z}_D} \frac{\partial \mathbf{P}^D(u, z - s)}{\partial s} [\mathbf{M}^D(z)]^\rho [\mathbf{P}^D(u, z - s)]^{\rho-1} \end{aligned} \quad (11)$$

Using the pdf expression in (2), we also have:

$$\frac{\partial \mathbf{P}^D(u, z - s)}{\partial s} = \frac{\int_\Omega \frac{\partial \mathcal{K}_z(s+D)}{\partial s} u \, d\mathbf{x}}{\int_\Omega u \, d\mathbf{x}} = \frac{\int_\Omega (z - s - D) \mathcal{K}_z(s + D) \, d\mathbf{x}}{\sigma^2 \int_\Omega u \, d\mathbf{x}} \quad (12)$$

¹ The proof of proposition 1 is given in the supplemental material available at: <http://externe.emt.inrs.ca/users/benayedi/BenAyed-Miccai13-Supp.pdf>

Table 1. Quantitative evaluations over 400 volumes acquired from 20 subjects. The statistics are expressed as mean \pm std.

<i>RMSE (mm)</i>	<i>DM</i>	GPU time/volume	CPU time/volume
2.30 \pm 0.12	0.84 \pm 0.07	0.129 sec	2.72 sec

Embedding (12) in (11), setting the obtained expression to zero, and after some manipulations, we obtain the following necessary condition for a minimum of \mathcal{E} with respect to s :

$$s - g(s) = 0 \text{ where}$$

$$g(s) = \frac{\sum_{z \in \mathcal{Z}_D} \int_{\Omega} (z - D) \mathcal{K}_z(D + s) [\mathbf{M}^D(z)]^{\rho} [\mathbf{P}^D(u, z - s)]^{\rho-1} u d\mathbf{x}}{\sum_{z \in \mathcal{Z}_D} \int_{\Omega} \mathcal{K}_z(D + s) [\mathbf{M}^D(z)]^{\rho} [\mathbf{P}^D(u, z - s)]^{\rho-1} u d\mathbf{x}} \quad (13)$$

Note that since the necessary condition in (13) has the form of a fixed-point equation, the solution can be obtained by fixed-point iterations:

$$s^{n+1} = g(s^n), \quad n = 1, 2, \dots \quad (14)$$

Let s^{opt} be the limit of sequence s^n at convergence. We have: $s^{opt} = \lim_{n \rightarrow +\infty} s^{n+1} = \lim_{n \rightarrow +\infty} g(s^n) = g(\lim_{n \rightarrow +\infty} s^n) = g(s^{opt})$. Consequently, s_{opt} is a solution of the necessary condition obtained in (13).

3 Experiments

We evaluated the algorithm over a data set containing short axis cardiac cine MRI volumes of 20 subjects (20 volumes per subject, each corresponding to a cardiac phase, i.e., we used 400 volumes in total). We performed three of types of experiments (The parameters were invariant for all the subjects $\alpha_D = 0.2$; $\alpha_A = 0.1$; $\gamma = 65$):

- Standard quantitative evaluations, which compare the results with independent manual segmentations approved by an expert;
- Comprehensive evaluations which demonstrate: (i) the performance of the proposed algorithm is not significantly affected by the choice of the training subject and (ii) the shape description we propose does not change significantly from one subject to another; and
- Computational evaluations, which demonstrate that the parallelized computations can bring a significant speed-up of more than 20 times.

Example: Fig. 1 depicts a typical example of the results, and demonstrates a high conformity between the manual and automatic segmentation.

Computational Evaluations: The parallelized implementation was run on an NVIDIA Tesla C1600 GPU, and the non-parallelized version on a 2.13 GHz Xeon (E5506), with 6 GB of RAM. Table 1 reports the average GPU/CPU times per volume. The parallelized implementation requires about 0.129 seconds for a typical volume, a speed-up of more than 20 times compared to the CPU version.




Fig. 1. An typical example using a $125 \times 125 \times 6$ volume. (a-b): Manual and automatic surfaces; (c-d): The corresponding 2D contours/slices.

Quantitative Performance Evaluations: We proceeded to a leave-one-out validation, where one subject was used for training and the rest of the subjects were used for testing. We assessed the similarities between the ground truth and the obtained segmentations using a surface-based measure, the Root Mean Squared Error (*RMSE*), and a region-based measure, the Dice Metric (*DM*). Here following an description of these measures.

- *Dice Metric (DM)*: Let V_m and V_a be the automated and manually segmented volumes, respectively. *DM* is given by $DM = 2 \frac{V_a \cap V_m}{V_a + V_m}$, and is always in $[0 \ 1]$, 1 indicating a perfect match and 0 a total mismatch.
- *RMSE*: *RMSE* evaluates a distance between automated surfaces and the corresponding manual ones. The *RMSE* over N points is given by: $RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{x}_i - \tilde{x}_i)^2 + (\hat{y}_i - \tilde{y}_i)^2}$ where (\hat{x}_i, \hat{y}_i) is a point on the automatically detected surface and $(\tilde{x}_i, \tilde{y}_i)$ is the corresponding point on the manually traced surface. The lower *RMSE*, the better the conformity of the results to the ground truth.

Table 1 reports the results, and demonstrates that the obtained 3D surfaces correlate with manual delineations. Note that a *DM* higher than 0.80 indicates an excellent agreement between manual and automatic segmentations [7].




Fig. 2. (a) and (b): Robustness of the proposed algorithm with respect to the choice of the training subject, (c): Invariance of the distance-based shape model.

Robustness with Respect to the Choice of Training Subject: We proceeded to a comprehensive leave-one-in evaluation method consisting of 20 tests, each corresponding to the choice of a different training subject. Then, we segmented the entire dataset and measured the corresponding average DM and $RMSE$. Figs. 2 (a) and (b) plot the obtained average DM and $RMSE$ as functions of the index of the training subject, demonstrating a very low variation.

Invariance of the Shape-prior Models: Using ground-truth segmentations, we plotted in Fig.2 (c) the distance distributions corresponding to the 20 subjects in the dataset (We used two volumes for each subject). The figure demonstrates that the distributions have very similar shapes, but slightly different supports. These slight shifts, which are due to inter-subject variations in scale (or size), are handled efficiently with the proposed fixed-point computations.

4 Conclusion

We proposed a real-time 3D MRI segmentation of the right ventricle based on probability product kernel constraints. The proposed algorithm removes the need for large, manually-segmented training sets and costly pose estimation procedures. We reported comprehensive experiments, which support the fact that a single subject is sufficient for training our algorithm and demonstrate that the obtained performance is independent of the choice of training data.

References

1. Grothues, F., Moon, J.C., Bellenger, N.G., Smith, G.S., Klein, H.U., Pennell, D.J.: Inter-study reproducibility of right ventricular volumes, function, and mass with cardiovascular magnetic resonance. *Am. Heart. J.* 147(2), 218–223 (2004)
2. Petitjean, C., Dacher, J.N.: A review of segmentation methods in short axis cardiac mr images. *Medical Image Analysis* 15, 169–184 (2011)
3. Zhang, H., Wahle, A., Johnson, R.K., Scholz, T.D., Sonka, M.: 4-D cardiac MR image analysis: left and right ventricular morphology and function. *IEEE Transactions on Medical Imaging* 29(2), 350–364 (2010)
4. Zhuang, X., Rhode, K.S., Arridge, S.R., Razavi, R.S., Hill, D., 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)
5. Lorenzo-Valdés, M., Sanchez-Ortiz, G.I., Elkington, A.G., Mohiaddin, R.H., Rueckert, D.: Segmentation of 4d cardiac mr images using a probabilistic atlas and the em algorithm. *Medical Image Analysis* 8(3), 255–265 (2004)
6. 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)
7. Ben Ayed, I., Chen, H.M., Punithakumar, K., Ross, I., Li, S.: Max-flow segmentation of the left ventricle by recovering subject-specific distributions via a bound of the bhattacharyya measure. *Medical Image Analysis* 16, 87–100 (2012)

8. Zhu, Y., Papademetris, X., Sinusas, A.J., Duncan, J.S.: Segmentation of the left ventricle from cardiac mr images using a subject-specific dynamical model. *IEEE Transactions on Medical Imaging* 29(4), 669–687 (2010)
9. Hautvast, G., Lobregt, S., Breeuwer, M., Gerritsen, F.: Automatic contour propagation in cine cardiac magnetic resonance images. *IEEE Transactions on Medical Imaging* 25(11), 1472–1482 (2006)
10. Jebara, T., Kondor, R.I., Howard, A.: Probability product kernels. *Journal of Machine Learning Research* 5, 819–844 (2004)
11. Zhang, Z., Kwok, J.T., Yeung, D.Y.: Surrogate maximization/minimization algorithms and extensions. *Machine Learning* 69, 1–33 (2007)
12. Yuan, J., Bae, E., Tai, X.C.: A study on continuous max-flow and min-cut approaches. In: *CVPR*, pp. 2217–2224 (2010)