

TRIC: Trust Region for Invariant Compactness and its application to abdominal aorta segmentation

Ismail Ben Ayed¹, Michael Wang², Brandon Miles³, and Gregory J. Garvin⁴

¹ GE Healthcare, London, ON, Canada

² McGill University, Montreal, QC, Canada

³ Simon Fraser University, Burnaby, BC, Canada

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

Abstract. This study investigates segmentation with a novel invariant compactness constraint. The proposed prior is a high-order *fractional* term, which is not directly amenable to powerful optimizers. We derive first-order *Gateaux* derivative approximations of our compactness term and adopt an iterative trust region paradigm by splitting our problem into constrained sub-problems, each solving the approximation globally via a Lagrangian formulation and a graph cut. We apply our algorithm to the challenging task of abdominal aorta segmentation in 3D MRI volumes, and report quantitative evaluations over 30 subjects, which demonstrate that the results correlate well with independent manual segmentations. We further show the use of our method in several other medical applications and demonstrate that, in comparison to a standard level-set optimization, our algorithm is one order of magnitude faster.

1 Introduction

Embedding shape priors in medical image segmentation is necessary in numerous applications [15,13,12,16], more so when the target segment has an intensity profile very similar to other parts in the image. Based on standard techniques such as statistical shape models [16] and probabilistic atlases [12], most of the existing medical image segmentation algorithms require (i) an intensive learning from a large, manually-segmented training set; and (ii) registration or pose optimization procedures (i.e., w.r.t rotation, translation, and scaling). Although they can yield excellent results in some applications, training-based algorithms may have difficulty in capturing the substantial variations that occur in a clinical context, with the results often being dependent on the choice of a specific training set. This is due to the fact that they enforce a strict *pixelwise* consistency between the solution and the template shapes in a training set.

Recently, there has been an ongoing research effort towards embedding *global* shape constraints in segmentation [15,14,13,2,6,18]. These include constraints on segment convexity [18], axial symmetry [15], area [14] and compactness [6], as well as geometric inter-segment relationships [2]. Several recent studies have

shown that such global shape constraints can lead to excellent performances in various medical applications [15,14,13,2], while removing the need for intensive training and pose estimation. Unfortunately, such constraints are high-order functionals, which result in difficult optimization problems that are not directly amenable to standard powerful optimizers, e.g., graph cuts or convex-relaxation techniques. The following summarizes the contributions of this study.

Contributions in the general-purpose context: We propose a novel global shape constraint, which measures segment compactness w.r.t a point set. Our compactness constraint can be invariant w.r.t scale, rotation and translation, unlike the shape compactness in [6] (which is pose dependent). Unfortunately, our prior is a high-order *fractional* functional, which is not directly amenable to powerful optimizers. We derive first-order *Gateaux* derivative approximations of our compactness term and adopt an iterative trust region paradigm [9] by splitting our problem into constrained sub-problems, each solving the approximation globally via a Lagrangian formulation and a graph cut. We show the use of our method in several medical applications and demonstrate that, in comparison to a standard level-set optimization, our algorithm is one order of magnitude faster.

Contributions in the application context: Abdominal aorta segmentation is an essential step towards accurate assessments of abdominal aortic aneurysms (AAA) [7,8,5]. Most of the existing works addressed the problem in CTA, except the recent interactive-segmentation study in [7] (which also considered MRI). In MRI, the task is seriously challenged by the intensity similarities and very weak edges between the aorta and its neighboring structures. Furthermore, this tubular structure may have sudden/unpredictable changes in the scale (size) of the 2D aorta cross-sections (e.g., due to aneurysms). With the scale-invariance property, our method can handle such unpredictable changes in scale. We report comprehensive evaluations over a set of 30 MRI volumes acquired from 30 subjects, which show a Dice metric of 0.91 ± 0.03 , an excellent agreement with independent manual segmentations.

2 Formulation

Proposed Functional: Let $I(\mathbf{p}) : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$ be an image function defined over a domain Ω . Our objective is to find an optimal region in Ω , so that the region is compact with respect to a point set and follows some appearance and boundary priors. We optimize a functional containing three terms:

$$\min_{u \in \{0,1\}} E(u) = \alpha \mathcal{A}(u) + \beta \mathcal{S}(u) + \gamma \mathcal{C}(u) \quad (1)$$

The following details each of the variables and terms that appear in (1):

- $u : \Omega \rightarrow \{0,1\}$ is binary function, which defines a variable segmentation of Ω : $\{\mathbf{p} \in \Omega / u(\mathbf{p}) = 1\}$, corresponding to the target segment, and $\{\mathbf{p} \in \Omega / u(\mathbf{p}) = 0\}$, corresponding to the complement of the target segment in Ω .

- We introduce the following invariant compactness prior:

$$\mathcal{C}(u) = \frac{\int_{\Omega} u(\mathbf{p}) \mathcal{D}^2(\mathbf{p}, \mathcal{Q}) d\mathbf{p}}{\left(\int_{\Omega} u(\mathbf{p}) d\mathbf{p}\right)^2} \quad (2)$$

where $\mathcal{D}(\mathbf{p}, \mathcal{Q})$ is the shortest-path distance between each point $\mathbf{p} \in \Omega$ and a set of reference points \mathcal{Q} : $\mathcal{D}(\mathbf{p}, \mathcal{Q}) = \min_{\mathbf{q} \in \mathcal{Q}} \|\mathbf{p} - \mathbf{q}\|$, with $\|\cdot\|$ denoting the L_2 norm. Depending on the application, point set \mathcal{Q} can be either a variable, which depends on the segmentation (i.e., on function u), or fixed (e.g., obtained from user inputs). To understand the meaning of the proposed compactness prior, let us first consider the particular case of a single reference point, which depends on u and corresponds to the centroid of the target segment, i.e., \mathcal{Q} is a singleton containing $\frac{\int_{\Omega} \mathbf{p} u(\mathbf{p}) d\mathbf{p}}{\int_{\Omega} u(\mathbf{p}) d\mathbf{p}}$. In this case, $\mathcal{C}(u)$ becomes one of the well known *Hu moment invariants* (invariant to translation, rotation and scaling). The recent study in [19] has shown that this invariant can be effectively used to measure shape circularity (i.e., the deviation of a given shape from the most compact a shape—a circle). The lower the value of this invariant, the closer the shape to a circle. Our extension of this compactness measure to multiple reference points can accommodate a more general class of shapes that deviate significantly from a circle (See the liver example in Fig. 2, where reference points are obtained from a user scribble). It is worth noting that, in applications where point set \mathcal{Q} is fixed (e.g., obtained from user inputs), $\mathcal{C}(u)$ loses its invariance w.r.t translation but remains invariant w.r.t scaling and rotation.

- \mathcal{A} is a log-likelihood appearance term [3]: $\mathcal{A}(u) = \int_{\Omega} u(\mathbf{p}) \log \frac{M_t(I(\mathbf{p}))}{M_c(I(\mathbf{p}))} d\mathbf{p}$, where M_t and M_c are fixed (learned *a priori*) model distributions of intensity within the target segment and its complement in Ω , respectively.
- $\mathcal{S}(u)$ is a pairwise term, which regularizes the segmentation boundary and biases it towards strong edges [3]: $\mathcal{S}(u) = \sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathcal{N}} \frac{\exp(-\sigma \|I(\mathbf{p}) - I(\mathbf{q})\|^2)}{\|\mathbf{p} - \mathbf{q}\|} \delta_{u(\mathbf{p}) \neq u(\mathbf{q})}$, with $\delta_{x \neq y}$ equal to 1 if $x \neq y$ and 0 otherwise. \mathcal{N} is a 16-neighborhood system containing all unordered pairs $\{\mathbf{p}, \mathbf{q}\}$ of neighboring elements of Ω .
- α, β, γ and σ are positive constants, which have to be fixed empirically.

Trust region optimization: Our prior is a *fractional* term, which is not directly amenable to powerful optimizers, e.g., graph cuts [3]. We derive first-order *Gateaux* derivative approximations of the compactness term and adopt an iterative trust region paradigm [9]: We split our problem into a sequence of easier sub-steps, each approximating the functional within a trust region around the current solution (i.e., a region where the approximation can be trusted).

Constrained sub-problems: At each iteration k , we solve the following constrained sub-problem via a Lagrangian formulation and a graph cut:

$$\begin{aligned} \min_{u: \Omega \rightarrow [0;1]} \tilde{E}_k(u) \quad \text{s.t.} \quad \|u - u_k\| < d_k, \quad \text{For } k = 0, 1, 2, \dots \\ \text{with } \tilde{E}_k(u) = \tilde{\mathcal{C}}_k(u) + \alpha \mathcal{A}(u) + \beta \mathcal{S}(u) \end{aligned} \quad (3)$$

where u is relaxed in $[0; 1]$ and $\tilde{\mathcal{C}}_k(u)$ is the first-order *Gateaux* derivative approximation of the compactness term near current solution u_k . d_k defines the size of the trust region, which is adjusted automatically at previous iteration k (see Algorithm 1, line 7). We derive the approximation in the case \mathcal{Q} is fixed (scale and rotation invariant case):

$$\tilde{\mathcal{C}}_k(u) = K_1 + \frac{\int_{\Omega} u(\mathbf{p}) d\mathbf{p}}{(\int_{\Omega} u_k(\mathbf{p}) d\mathbf{p})^2} - \frac{2 \int_{\Omega} u_k(\mathbf{p}) \mathcal{D}^2(\mathbf{p}, \mathcal{Q}) d\mathbf{p}}{(\int_{\Omega} u_k(\mathbf{p}) d\mathbf{p})^3} \int_{\Omega} u(\mathbf{p}) \mathcal{D}^2(\mathbf{p}, \mathcal{Q}) d\mathbf{p} \quad (4)$$

where K_1 a constant independent of u .

Lagrangian formulation of the sub-problems: We state each sub-problem in (3) as an unconstrained optimization: $\min_{u: \Omega \rightarrow [0; 1]} \tilde{E}_k(u) + \lambda \|u - u_k\|$. For binary functions u and u_k in $\{0, 1\}$, $\|u - u_k\|$ can be approximated with a unary potential [4]: $\|u - u_k\| \approx \int_{\Omega} \phi_k(p)(u(\mathbf{p}) - u_k(\mathbf{p})) d\mathbf{p}$, where ϕ_k is the signed distance function corresponding to the boundary defined by u_k , i.e., $\{\mathbf{p} \in \Omega | \nabla u_k \neq 0\}$.

Algorithm: A summary of the procedure is given in *Algorithm 1*. Each sub-problem in line 3 can be solved globally with a graph cut [3]. Note that the use of graph cuts is an option, among others, to solve the trust-region sub-problems. One can use other global optimization techniques, e.g., those based on convex relaxation [13]. Once candidate solution u^* is computed, the merit of the approximation is evaluated by the ratio between the actual and approximate reduction in the functional. Based on this ratio, the solution is updated in line 6 and the trust region is adjusted in line 7. We set parameter τ_2 in line 7 to 0.25.

Algorithm 1: TRUST REGION FOR INVARIANT COMPACTNESS (TRIC)

```

1 Repeat
2   //Solve trust region sub-problem
3    $u^* \leftarrow \min_{u: \Omega \rightarrow \{0, 1\}} \tilde{E}_k(u) + \lambda \int_{\Omega} \phi_k(p)(u(\mathbf{p}) - u_k(\mathbf{p})) d\mathbf{p}$ 
4    $\tilde{\mathcal{R}} = \tilde{E}_k(u_k) - \tilde{E}_k(u^*)$  //Approximate functional reduction
5    $\mathcal{R} = E(u_k) - E(u^*)$  //Actual functional reduction
6   //Update current solution:  $u_{k+1} \leftarrow \begin{cases} u^* & \text{if } \frac{\mathcal{R}}{\tilde{\mathcal{R}}} > 0 \\ u_k & \text{otherwise} \end{cases}$ 
7   //Adjust the trust region  $d_{k+1} \leftarrow \begin{cases} d_k \cdot \gamma & \text{if } \frac{\mathcal{R}}{\tilde{\mathcal{R}}} > \tau_2 \\ d_k / \gamma & \text{otherwise} \end{cases}$ 
8 Until Convergence

```

3 Experiments

We report several sets of experiments to demonstrate the benefit of the proposed prior, including: (i) Quantitative evaluations of supervised abdominal aorta segmentations over a set of 30 MRI volumes acquired from 30 subjects; (ii) Additional examples, which show the use of our method in other applications, e.g.,

Dice metric (with compactness)	Dice metric (without compactness) [3]
0.91 \pm 0.03	0.33 \pm 0.05

Table 1. Quantitative evaluations of abdominal aorta segmentations over 30 subjects.

liver segmentation in CT and delineations of cardiac structures (the cavity and myocardium) in MRI; and (iii) Comparisons in regard to speed/optimality with a standard level-set optimization applied to the same compactness functional, which show that our method is one order of magnitude faster and is less likely to get trapped into weak local minima (See the plot in Fig. 3).

Abdominal aorta segmentation:

Typical examples: Fig. 1 depicts aorta segmentations in 3D T2-weighted MRI volumes. The problem is challenging due to the intensity similarities and very weak edges between the aorta and its neighboring structures, as well as the sudden/unpredictable changes in the sizes of the 2D aorta cross-sections. The first row of Fig. 1 depicts the cross-sectional results for one subject at different slices, along with the ground truth and the segmentations obtained *without* compactness (this corresponds to the well known method in [3]). The solution of [3] leaked into several parts of the background, whereas our method yielded segmentations that are very close to the ground truth. The second row of Fig. 1 depicts 3D results obtained for another subject. The scale invariance of our method handled well the sudden/unpredictable changes in the size of the 2D aorta cross-sections. For this application, we assume that the centroid of the aorta cross-section within each 2D slice is given. Such a reference point is used to define the shortest-path distance and build the appearance models. The method belongs to the class of tubular/vascular structure segmentation techniques that use the centerline of the structure as input, e.g., [17]; See [10] for a complete categorization of prior works. Standard semi-automated centerline extraction techniques [10], which use minimum user inputs (e.g., one or two seed points for the whole volume), can further automate the process. The parameters were fixed as follows: $\alpha = 5 \times 10^{-3}$, $\beta = 10^{-2}$ and $\gamma = 10^8$. The appearance models are learned from inside/outside a disc centered at the reference point, with a radius equal to 10 pixels.

Quantitative evaluations: The evaluation was carried out over a data set of 30 T2-weighted MRI volumes acquired from 30 subjects. We segmented a total of 1968 cross-sectional 2D slices. The results were compared to independent manual segmentations performed by an expert using the well known Dice metric (*DM*) measure. *DM* evaluates the similarity (overlap) between the automatically detected and ground-truth segments: $DM = \frac{2\mathbf{S}_{am}}{\mathbf{S}_a + \mathbf{S}_m}$, with \mathbf{S}_a , \mathbf{S}_m , and \mathbf{S}_{am} corresponding respectively to the sizes of the segmented aorta volume, the corresponding hand-labeled volume, and the intersection between them. The parameters were fixed for all subjects, and were similar to those used in the examples above. Table 1 reports the *DM* mean and standard deviation *with* and *without* compactness (i.e., using the method in [3]). The plot in Fig. 1 (bottom-

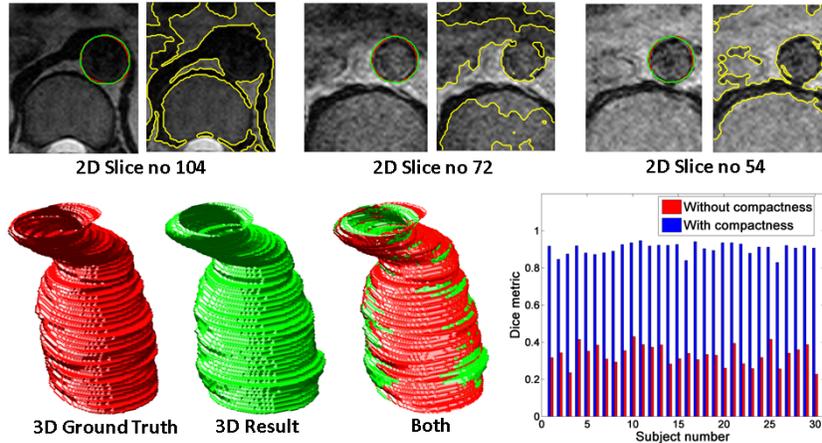


Fig. 1. Abdominal aorta segmentation in 3D MRI volumes. The first row depicts cross-sectional results for one subject at different slice levels, with the red curve showing the ground truth, the green curve showing the result with our compactness prior and the yellow curve showing the result without compactness (i.e., using [3]). Second row, left: 3D result obtained for another subject. Second row, right: the Dice metrics obtained for 30 subjects *with* and *without* compactness.

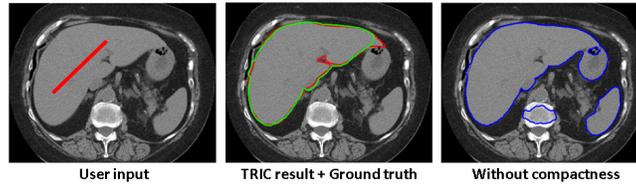


Fig. 2. Interactive liver segmentation in CT. $\alpha = 5 \times 10^{-3}$, $\beta = 10^{-1}$ and $\gamma = 5 \times 10^8$.

right) depicts the Dice metrics obtained for the 30 subjects. The results confirm that the compactness term brings a substantial improvement in accuracy.

Liver segmentation in CT: Fig. 2 illustrates the application of our compactness term to interactive liver segmentation in CT. This example shows how our compactness term can accommodate a more general class of shapes that deviate significantly from a circle. Here, the shortest-path distance is defined w.r.t multiple reference points (a segment), which were obtained from a simple two-point user input (Fig. 2, first column). The second column depicts the obtained result (green curve) along with the ground truth (red curve). The last column depicts the result obtained without compactness (i.e., with [3]).

Cardiac segmentation in MRI: Fig. 3 shows delineations of the cavity/myocardium boundaries in MRI, a well-known problem in cardiac image analysis [1]. The first column depicts cavity segmentation starting from a single-

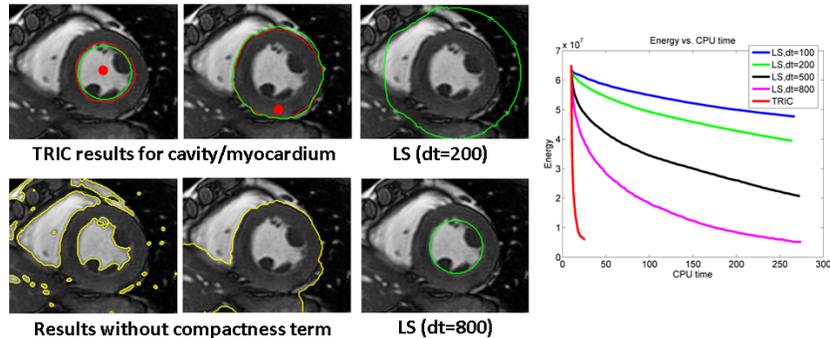


Fig. 3. Delineation of the cavity/myocardium boundaries in cardiac MRI. First column: delineation of the cavity boundary; Second column: delineation of the external boundary of the myocardium; Third column: results obtained with a level set optimization [11] for two different time steps (dt). Last column: Plots of the evolution of the cavity-segmentation functional w.r.t CPU time for TRIC and level-set optimization (LS) with different time steps.

point user input, which serves as a reference to define the shortest-path distance and build the appearance models. The upper figure depicts the user input (red point), obtained result (green curve) and ground truth (red curve), whereas the lower one depicts the result without compactness. The parameters were fixed as follows: $\alpha = 5 \times 10^{-3}$; $\beta = 0$; $\gamma = 5 \times 10^8$; the appearance models are learned from inside/outside a disc centered at the reference point, with a radius of 10 pixels. The second column depicts epicardium delineation using an additional user-specified point to learn myocardium appearance. The upper figure shows the obtained result (green curve) along with the ground truth (red curve) and the additional user input (red point). The lower figure depicts segmentation without compactness. In this example, we added a hard constraint [3] to ensure that the solution includes the cavity segment (which was computed at the previous step), and fixed the parameters as follows: $\alpha = 5 \times 10^{-3}$; $\beta = 10^{-1}$; $\gamma = 10^8$.

Comparisons with level sets in regard to speed/optimality: For the cavity segmentation example, we plotted the evolution of our functional w.r.t CPU time for TRIC and level-set gradient-descent optimization (LS), implemented via the well known scheme in [11] (See the right-hand side of Fig. 3). For the LS method, we varied the artificial time step (dt) in interval [100; 800] so as to ensure that the curve evolution is stable¹. The plot shows that, in the case of our compactness prior, the LS method is highly sensitive to the gradient-descent time step and is one order of magnitude slower than our method.

¹ In comparison to level set implementations based on re-initialization procedures, the scheme in [11] allows much larger time steps (and therefore faster curve evolution). However, the value of dt still have to be fixed carefully; a very large value would cause instability in curve evolution.

References

1. Ben Ayed, I., Lu, Y., Li, S., Ross, I.G.: Left ventricle tracking using overlap priors. In: MICCAI (1). pp. 1025–1033 (2008)
2. Ben Ayed, I., Punithakumar, K., Garvin, G.J., Romano, W., Li, S.: Graph cuts with invariant object-interaction priors: Application to intervertebral disc segmentation. In: IPMI. pp. 221–232 (2011)
3. Boykov, Y., Funka-Lea, G.: Graph cuts and efficient n-d image segmentation. *International Journal of Computer Vision* 70(2), 109–131 (2006)
4. Boykov, Y., Kolmogorov, V., Cremers, D., Delong, A.: An integral solution to surface evolution pdes via geo-cuts. In: ECCV (3). pp. 409–422 (2006)
5. de Bruijne, M., van Ginneken, B., Viergever, M.A., Niessen, W.J.: Interactive segmentation of abdominal aortic aneurysms in cta images. *Medical Image Analysis* 8(2), 127–138 (2004)
6. Das, P., Veksler, O., Zavadsky, V., Boykov, Y.: Semiautomatic segmentation with compact shape prior. *Image and Vision Computing* 27(1-2), 206–219 (2009)
7. Duquette, A.A., Jodoin, P.M., Bouchot, O., Lalonde, A.: 3d segmentation of abdominal aorta from ct-scan and mr images. *Computerized Medical Imaging and Graphics* 36(4), 294–303 (2012)
8. Freiman, M., Esses, S.J., Joskowicz, L., Sosna, J.: An iterative model-constrained graph-cut algorithm for abdominal aortic aneurysm thrombus segmentation. In: ISBI. pp. 672–675 (2010)
9. Gorelick, L., Schmidt, F., Boykov, Y.: Fast trust region for segmentation. In: CVPR. pp. 1714–1721 (2013)
10. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3d vessel lumen segmentation techniques: Models, features and extraction schemes. *Medical Image Analysis* 13(6), 819–845 (2009)
11. Li, C., Xu, C., Gui, C., Fox, M.D.: Level set evolution without re-initialization: A new variational formulation. In: CVPR. pp. 430–436 (2005)
12. Linguraru, M.G., Pura, J.A., Pamulapati, V., Summers, R.M.: Statistical 4d graphs for multi-organ abdominal segmentation from multiphase ct. *Medical Image Analysis* 16(4), 904–914 (2012)
13. Nambakhsh, C.M.S., Peters, T.M., Islam, A., Ben Ayed, I.: Right ventricle segmentation with probability product kernel constraints. In: MICCAI (1). pp. 509–517 (2013)
14. Niethammer, M., Zach, C.: Segmentation with area constraints. *Medical Image Analysis* 17(1), 101–112 (2013)
15. Qiu, W., Yuan, J., Ukwatta, E., Sun, Y., Rajchl, M., Fenster, A.: Fast globally optimal segmentation of 3d prostate mri with axial symmetry prior. In: MICCAI (2). pp. 198–205 (2013)
16. Rousson, M., Cremers, D.: Efficient kernel density estimation of shape and intensity priors for level set segmentation. In: MICCAI (2). pp. 757–764 (2005)
17. Schaap, M., Neefjes, L., Metz, C., van der Giessen, A.G., Weustink, A.C., Mollet, N., Wentzel, J., van Walsum, T., Niessen, W.J.: Coronary lumen segmentation using graph cuts and robust kernel regression. In: IPMI. pp. 528–539 (2009)
18. Veksler, O.: Star shape prior for graph-cut image segmentation. In: ECCV (3). pp. 454–467 (2008)
19. Zunic, J.D., Hirota, K., Rosin, P.L.: A hu moment invariant as a shape circularity measure. *Pattern Recognition* 43(1), 47–57 (2010)