

Active Contours with Group Similarity

Xiaowei Zhou[†], Xiaojie Huang[‡], James S. Duncan[‡], Weichuan Yu[†]

[†] Department of ECE, The Hong Kong University of Science and Technology

[‡] Image Analysis and Processing Group, Yale University

Abstract

Active contours are widely used in image segmentation. To cope with missing or misleading features in images, researchers have introduced various ways to model the prior of shapes and use the prior to constrain active contours. However, the shape prior is usually learnt from a large set of annotated data, which is not always accessible in practice. Moreover, it is often doubted that the existing shapes in the training set will be sufficient to model the new instance in the testing image. In this paper, we propose to use the group similarity of object shapes in multiple images as a prior to aid segmentation, which can be interpreted as an unsupervised approach of shape prior modeling. We show that the rank of the matrix consisting of multiple shapes is a good measure of the group similarity of the shapes, and the nuclear norm minimization is a simple and effective way to impose the proposed constraint on existing contour models. Moreover, we develop a fast algorithm to solve the proposed model by using the accelerated proximal method. Experiments using echocardiographic image sequences acquired from acute canine experiments demonstrate that the proposed method can consistently improve the performance of active contour models and increase the robustness against image defects such as missing boundaries.

1. Introduction

Image segmentation is a fundamental task in many applications. Among various techniques, the active contour model is widely used. A contour is evolved by minimizing certain energies to match the object boundary while preserving the smoothness of the contour [2]. The active contour is usually represented by landmarks [18] or level sets [20, 8]. A variety of image features have been used to guide the active contour, typically including image gradient [7, 31], region statistics [34, 8], color and texture [14].

In real applications, the performance of the active contour model is prone to be corrupted by missing or misleading features. For example, segmentation of the left ventricle in ultrasound images is still an unresolved problem due




Figure 1. A toy example illustrating our idea. The images in the top row show the segmentation results of the region-based active contour applied on each image separately. The images in the bottom row show the corresponding results with a group similarity constraint on the shapes, which makes the shapes consistent with each other and less affected by local misleading features. This property is desired in many applications, such as segmentation of the left ventricle from a cardiac image sequence.

to the characteristic artefacts in ultrasound such as attenuation, speckle and signal dropout [23]. To improve the robustness of active contours, the shape prior is often used. The prior knowledge of the shape to be segmented is modeled based on a set of manually-annotated shapes to guide the segmentation. Previous deformable template models [32, 27, 17, 21] can be regarded as the early efforts towards knowledge-based segmentation. In more recent works, the shape prior was applied by regularizing the distance from the active contour to the template in a level-set framework [10, 24, 9]. Another category of methods popularly used for shape prior modeling is the active shape model or point distribution model [11]. Briefly speaking, each shape is denoted by a vector and regarded as a point in the shape space. Then, the principal component analysis is carried out to obtain the mean and several most significant modes of shape variations, which establish a low-dimensional space to describe the favorable shapes. During the segmentation of a new image, the candidate shape is constrained in the shape space [19, 29]. Also, dynamic models can be integrated to model the temporal continuity when tracking an object in a sequence [12, 35]. Other extensions of the active shape

model include manifold learning [15] and sparse representation [33], to name a few.

While the shape prior has proven to be a powerful tool in segmentation, it has two limitations:

1. Previous methods for shape prior modeling require a large set of annotated data, which is not always accessible in practice.
2. It is often doubted that the existing shapes in the training set will be sufficient to model the object shape in a new image.

In this paper, we propose to use the similarity among object shapes as a prior for segmentation. A practical application is the segmentation of cardiac images, where the shape of the left ventricle shall keep consistent throughout a cardiac cycle although it deforms globally due to the heart beating. Figure 1 gives a simple illustration. The three hearts in the top row share the similar shape, while there are some local artefacts such as occlusion and missing boundaries. If the active contour is applied on each of them separately, the segmentation result will be affected by the misleading features. Our results are presented in the bottom row, where the similarity among the object shapes is used as a constraint to make the active contours more robust. In summary, the contributions of this paper are:

1. We showed that the vectors representing a group of similar shapes would form a low-rank matrix, even if they are different with each other due to certain global coordinate transformation.
2. Based on the low-rank property of similar shapes, we proposed to use the nuclear norm (convex surrogate of rank) to regularize the group similarity of shapes in segmentation. The regularizer could be conveniently integrated into existing active contour models.
3. We developed a fast algorithm to solve the proposed model by using the accelerated proximal gradient method. The experiments showed that the proposed constraint made the active contour model better regularized and require fewer iterations to converge.
4. We applied the proposed method to ultrasound image segmentation and demonstrated that the group-similarity regularization could significantly improve the robustness of the active contour model.

The rest of this paper is organized as follows: Section 2 introduces the basic theory and the formulation of our method. Section 3 describes the algorithm to solve our model. Section 4 demonstrates the merits of our method by experiments. Finally, Section 5 concludes the paper with some discussions.

2. Formulation

2.1. Group similarity measure

To apply a group similarity constraint to active contours, a proper similarity measure is desired. Typically, the similarity between two contours is measured by calculating the distances between the corresponding points on the contours, and the group similarity can be measured by the sum of pairwise distances between contours. The main drawback of this method is that the contour distance is not invariant under similarity transformation. For example, two contours with identical shapes but different sizes may be regarded as dissimilar based on the distance between them, which is not desired in our task of segmenting the left ventricle.

Here, we propose to use the matrix rank to measure the group similarity of shapes. Suppose each shape is represented by a vector. Multiple shapes form a matrix. Intuitively, the rank of the matrix measures the correlation among the shapes. For example, the rank equals to 1 if the shapes are identical, and the rank may increase if some shapes change. Moreover, we can show that the shape matrix is still low-rank if the shape change is due to the similarity transformation such as translation, scaling and rotation.

For example, let vector $\mathcal{C} = [x_1, \dots, x_p, y_1, \dots, y_p]^T \in \mathbb{R}^{2p}$ represent a digitized parametric curve in the 2-D plane, where (x_i, y_i) is a landmark on the curve. Suppose there are n curves $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_n$, and for each $i \neq 1$, \mathcal{C}_i is generated from \mathcal{C}_1 through affine transformation. Then, the matrix $[\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_n] \in \mathbb{R}^{2p \times n}$ has the following property

$$\text{rank}([\mathcal{C}_1, \dots, \mathcal{C}_n]) \leq 6, \quad (1)$$

for any large n and p . The proof is straightforward. For any \mathcal{C} generated from \mathcal{C}_1 by certain affine transformation $\mathbf{x}' = \mathbf{M}\mathbf{x} + \mathbf{T}$, where \mathbf{x} (or \mathbf{x}') denotes the coordinates of a landmark before (or after) transformation, we can represent \mathcal{C} as $\mathcal{C} = \Phi\mathbf{w}$, where

$$\Phi = \begin{bmatrix} \mathcal{C}_1^x & \mathbf{0} & \mathcal{C}_1^y & \mathbf{0} & \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathcal{C}_1^x & \mathbf{0} & \mathcal{C}_1^y & \mathbf{0} & \mathbf{1} \end{bmatrix},$$

$$\mathbf{w} = [m_{11}, m_{21}, m_{12}, m_{22}, t_1, t_2]^T.$$

Here, \mathcal{C}_1^x and $\mathcal{C}_1^y \in \mathbb{R}^p$ denote the first half and the second half of vector \mathcal{C}_1 , and m_{ij} and t_i are the elements of \mathbf{M} and \mathbf{T} , respectively. Since Φ only depends on \mathcal{C}_1 for any \mathcal{C} , $[\mathcal{C}_1, \dots, \mathcal{C}_n] = \Phi[\mathbf{w}_1, \dots, \mathbf{w}_n]$. The dimension of the column space of Φ is at most 6. Therefore, $\text{rank}([\mathcal{C}_1, \dots, \mathcal{C}_n]) \leq 6$.

Intrinsically, the rank of the shape matrix describes the degree of freedom of the shape change. The low-rank constraint will allow the global change of contours such as translation, scaling, rotation and principal deformation to fit the image data while truncating the local variation caused by image defects.

2.2. Energy function

Given a sequence of images I_1, \dots, I_n , we try to find a set of contours $\mathcal{C}_1, \dots, \mathcal{C}_n$ to segment the object in these images. To keep the contours similar to each other, we propose to segment the images by

$$\min_{\mathbf{X}} \sum_{i=1}^n f_i(\mathcal{C}_i), \quad \text{subject to } \text{rank}(\mathbf{X}) \leq K, \quad (2)$$

where $\mathbf{X} = [\mathcal{C}_1, \dots, \mathcal{C}_n]$ and K is a predefined constant. $f_i(\mathcal{C}_i)$ is the energy of an active contour model to evolve the contour in each frame, such as snake [18], geodesic active contour [7], and region-based models [34, 8]. For example, the region-based energy in [8] reads

$$f_i(\mathcal{C}_i) = \int_{\Omega_1} (I_i(\mathbf{x}) - u_1)^2 d\mathbf{x} + \int_{\Omega_2} (I_i(\mathbf{x}) - u_2)^2 d\mathbf{x} + \beta \text{length}(\mathcal{C}_i), \quad (3)$$

where Ω_1 and Ω_2 represent the regions inside and outside the contour, and u_1 and u_2 denote the mean intensity of Ω_1 and Ω_2 , respectively.

Since rank is a discrete operator which is both difficult to optimize and too rigid as a regularization method, we propose to use the following relaxed form as the objective function:

$$\min_{\mathbf{X}} \sum_{i=1}^n f_i(\mathcal{C}_i) + \lambda \|\mathbf{X}\|_*. \quad (4)$$

Here, $\text{rank}(\mathbf{X})$ in (2) is replaced by the nuclear norm $\|\mathbf{X}\|_*$, *i.e.* the sum of singular values of \mathbf{X} . Recently, the nuclear norm minimization has been widely used in low-rank modeling such as matrix completion [6] and robust principal component analysis [5]. As a tight convex surrogate to the rank operator [16], the nuclear norm has several good properties: Firstly, the convexity of the nuclear norm makes it possible to develop fast and convergent algorithms in optimization. Secondly, the nuclear norm is a continuous function, which is important for a good regularizer in many applications. For instance, in our problem, the small perturbation in the shapes may result in a large increase of $\text{rank}(\mathbf{X})$, while $\|\mathbf{X}\|_*$ may rarely change.

3. Algorithm

In this section, we will discuss how to solve the optimization problem in (4). If there is no regularizer $\|\mathbf{X}\|_*$, (4) can be locally minimized by gradient descent, which gives the curve evolution steps in typical active contour models. In our model, it is difficult to apply gradient descent directly due to the nuclear norm, which is nonsmooth and its subgradient is hard to compute.

Recently, the Proximal Gradient (PG) method [1, 22] is used to solve the following category of problems

$$\min_{\mathbf{X}} F(\mathbf{X}) + \lambda \mathcal{R}(\mathbf{X}), \quad (5)$$

where $F(\mathbf{X})$ is a differentiable function and $\mathcal{R}(\mathbf{X})$ corresponds to a convex penalty which can be nonsmooth. Our problem is in this category with $F(\mathbf{X}) = \sum_{i=1}^n f_i(\mathcal{C}_i)$ and $\mathcal{R}(\mathbf{X}) = \|\mathbf{X}\|_*$. The basic step in PG is to make the following quadratic approximation to $F(\mathbf{X})$ based on the previous estimate \mathbf{X}' at each iteration

$$\begin{aligned} Q_\mu(\mathbf{X}, \mathbf{X}') &= F(\mathbf{X}') + \langle \nabla F(\mathbf{X}'), \mathbf{X} - \mathbf{X}' \rangle + \frac{\mu}{2} \|\mathbf{X} - \mathbf{X}'\|_F^2 + \lambda \mathcal{R}(\mathbf{X}), \\ &= \frac{\mu}{2} \|\mathbf{X} - [\mathbf{X}' - \frac{1}{\mu} \nabla F(\mathbf{X}')] \|_F^2 + \lambda \mathcal{R}(\mathbf{X}) + \text{const}, \end{aligned} \quad (6)$$

where $\langle \cdot, \cdot \rangle$ means the inner product, $\|\cdot\|_F$ denotes the Frobenius norm, and μ is a constant. It is shown in [22] that, if $F(\mathbf{X})$ is differentiable with Lipschitz continuous gradient, the sequence generated by the following iteration will converge to a stationary point of the function in (5) with a convergence rate of $\mathcal{O}(\frac{1}{k})$

$$\begin{aligned} \mathbf{X}^{k+1} &= \arg \min_{\mathbf{X}} Q_\mu(\mathbf{X}, \mathbf{X}^k), \\ &= \arg \min_{\mathbf{X}} \frac{1}{2} \|\mathbf{X} - [\mathbf{X}^k - \frac{1}{\mu} \nabla F(\mathbf{X}^k)] \|_F^2 + \frac{\lambda}{\mu} \mathcal{R}(\mathbf{X}). \end{aligned} \quad (7)$$

The next question is how to solve the update step in (7). For our problem, the following lemma is required, which has been proven in [4].

Lemma 1 Given $\mathbf{X} \in \mathbb{R}^{m \times n}$, the solution to the problem

$$\min_{\mathbf{X}} \frac{1}{2} \|\mathbf{X} - \mathbf{Z}\|_F^2 + \alpha \|\mathbf{X}\|_* \quad (8)$$

is given by $\mathbf{X}^* = \mathcal{D}_\alpha(\mathbf{Z})$, where

$$\mathcal{D}_\alpha(\mathbf{Z}) = \sum_{i=1}^{\min(m,n)} (\sigma_i - \alpha)_+ \mathbf{u}_i \mathbf{v}_i^T, \quad (9)$$

\mathbf{u}_i and \mathbf{v}_i are the left and right singular vectors of \mathbf{Z} , σ_i is the singular value, and $(\cdot)_+ = \max(\cdot, 0)$.

The operator $\mathcal{D}_\alpha(\mathbf{Z})$ refers to the *singular value thresholding operator* [4]. Hence, the update step of our algorithm reads

$$\mathbf{X}^{k+1} = \mathcal{D}_{\frac{\lambda}{\mu}}(\mathbf{X}^k - \frac{1}{\mu} \nabla F(\mathbf{X}^k)), \quad (10)$$

The remaining issue is to find $\nabla F(\mathbf{X}^k)$. Here, $F(\mathbf{X})$ is decomposable as a sum of quantities over each frame.




Figure 2. Toy example. Top: the results of region-based segmentation without the similarity constraint proposed in this paper. Bottom: the results of the same model with the similarity constraint.

Hence, $\nabla F(\mathbf{X}^k) = [\nabla f_1(\mathcal{C}_1^k), \dots, \nabla f_n(\mathcal{C}_n^k)]$. For example, the gradient of the region-based energy defined in (3) is [34]

$$\nabla f_i(p) = [(I(p) - u_2)^2 - (I(p) - u_1)^2] \mathbf{n}_p + \beta \kappa_p \mathbf{n}_p, \quad (11)$$

where p is a landmark on \mathcal{C} , and \mathbf{n}_p and κ_p are the normal vector and the curvature at p , respectively.

To accelerate the convergence of the proximal method, the Nesterov method can be used [22]. Instead of updating \mathbf{X} from the previous estimate \mathbf{X}^k , the approximation in (6) is made at an intermediate point \mathbf{Y}^k , which is the linear extrapolation from \mathbf{X}^{k-1} and \mathbf{X}^k

$$\mathbf{Y}^k = \mathbf{X}^k + \frac{t^{k-1} - 1}{t^k} (\mathbf{X}^k - \mathbf{X}^{k-1}), \quad (12)$$

where t^k is determined online. This modification will give a convergence rate of $\mathcal{O}(\frac{1}{k^2})$. Please refer to [1, 22] for more details.

The intuition of our algorithm is that, at each iteration, we first evolve the active contours according to the image-based forces and then impose the group similarity regularization via singular value thresholding. The overall algorithm is summarized in Algorithm 1.

Algorithm 1 Accelerated proximal gradient algorithm to solve our model in (4)

1. **Initialize:** $\mathbf{X}^0 = \mathbf{X}^{-1}$, $t_0 = t_{-1} = 1$
 2. **for** $k = 0 \rightarrow$ maximum number of iterations **do**
 3. $\mathbf{Y}^k = \mathbf{X}^k + \frac{t^{k-1} - 1}{t^k} (\mathbf{X}^k - \mathbf{X}^{k-1})$
 4. **for** $i = 1 \rightarrow n$ **do**
 5. $\mathbf{y}_i^k \leftarrow \mathbf{y}_i^k - \frac{1}{\mu} \nabla f_i(\mathbf{y}_i^k)$
 6. **end for**
 7. $\mathbf{X}^{k+1} = \mathcal{D}_{\frac{\Delta}{\mu}}(\mathbf{Y}^k)$
 8. $t^{k+1} = \frac{1 + \sqrt{1 + 4(t^k)^2}}{2}$
 9. **if** $\|\mathbf{X}^{k+1} - \mathbf{X}^k\|_F <$ tolerance **then**
 10. **return**
 11. **end if**
 12. **end for**
-

4. Results

In this section, we evaluate the proposed method on both synthesized data and ultrasound images. To demonstrate the advantages of the group similarity constraint, we compare the results of the same active contour model before and after applying the proposed constraint. We select the region-based active contour in (3) as the basic model, which is less sensitive to initialization and has fewer parameters to tune compared with edge-based methods.

In our implementation, we initialize the active contours as $\mathbf{X}^0 = [\mathcal{C}_0, \dots, \mathcal{C}_0]$, where \mathcal{C}_0 is a rough outline of the object placed manually in an image. Three parameters need to be selected in our algorithm. β in (3) controls the smoothness of each contour, λ in (4) controls the group similarity of contours, and μ in (7) controls the step-length of curve evolution in each iteration. We choose the parameters empirically and use the same set of values for all experiments. The effect of λ will be further elaborated in Section 4.3.

4.1. Toy example

To better illustrate our idea, we synthesized eight images as shown in Figure 2. The dark region in a heart shape is the object to be segmented, which has different sizes and positions in different images. Occlusion or deletion is simulated and added to the different part of each shape. Gaussian noise is also added. Our goal is to find the heart shapes in all images based on the prior that the shapes are similar to each other in these images. As shown in Figure 2, our results with the shape similarity constraint are more robust against the local defects in images compared to the results without such constraint. The rank of output \mathbf{X} is 4.

4.2. Ultrasound image segmentation

We apply our method to five 2-D short-axis sequences of canine echocardiography. The echocardiography segmentation is a very challenging problem due to various misleading features in ultrasound images. Please refer to Table 1 and the appendix for more information on the data sets.




Figure 3. Two examples of left ventricle segmentation. For each sequence, 7 frames uniformly spaced through the sequence are selected to demonstrate the results. For each panel, the top row and the bottom row present the results of region-based active contours without and with the proposed constraint, respectively. The green solid curve is the output contour. The yellow dashed curve is the manual reference.

seq.	(#pixel) \times #frame	rank	#iter	time (sec)
1	(328 \times 302) \times 20	4	28 (50)	9.6 (16.4)
2	(267 \times 257) \times 26	7	35 (50)	12.5 (17.3)
3	(306 \times 279) \times 26	5	14 (50)	5.2 (17.4)
4	(276 \times 271) \times 24	4	30 (34)	10.7 (12.0)
5	(275 \times 278) \times 22	3	26 (27)	8.0 (8.1)

Table 1. Sequence information including the data size, the rank of output \mathbf{X} , the number of iterations for the algorithm to stop, and the computational time. The value in parentheses corresponds to the result without the shape constraint.

4.2.1 Qualitative comparison

Uniformly-selected frames of two sequences are displayed in Figure 3 to qualitatively evaluate the segmentation.

The results of the region-based active contour without the proposed constraint are given in the top rows. The results are corrupted in several images. For example, in frames 7, 9 and 13 of sequence 1, the active contour fails to stop near the true endocardial border on the right due

to the low contrast over there. Similar errors can be found in frames 11 and 13 of sequence 2. Moreover, the active contour is prone to be trapped by the misleading features, *e.g.* the upper part of the contour in frame 1 of sequence 2 is attracted to the bright pattern in the blood pool.

Our results are given in the bottom rows. There are two observations worth mentioning. Firstly, the contour shapes are globally consistent with each other throughout the sequence, which is attributed to the group similarity constraint. Hence, the contours are more resistant to local misleading features. Secondly, the constrained shape model is still flexible enough to adapt the deformation of the object shape. For example, the segmented shapes in frame 1 and frame 13 of sequence 2 are largely different and the shape change is far away from the affine transformation. The drawback of our method is that it cannot address the global bias of the model. For example, the image contrast of sequence 2 is poor near the upper endocardial border. Therefore, the region-based active contours cannot attach closely to the true boundary. In practice, more appealing results can

Sequence 1			
Method	MAD (mm)	HD (mm)	Dice (%)
w/	0.80 ± 0.20	3.15 ± 0.98	93.1 ± 1.3
w/o	1.29 ± 0.60	7.47 ± 3.22	88.7 ± 5.7
Sequence 2			
Method	MAD (mm)	HD (mm)	Dice (%)
w/	0.75 ± 0.15	3.45 ± 0.79	92.2 ± 1.6
w/o	1.02 ± 0.17	5.81 ± 1.08	88.8 ± 3.4
Sequence 3			
Method	MAD (mm)	HD (mm)	Dice (%)
w/	0.50 ± 0.10	2.40 ± 0.70	95.1 ± 0.5
w/o	0.58 ± 0.20	2.81 ± 1.23	94.4 ± 1.2
Sequence 4			
Method	MAD (mm)	HD (mm)	Dice (%)
w/	0.75 ± 0.19	2.52 ± 0.68	92.9 ± 1.3
w/o	1.10 ± 0.78	4.15 ± 2.18	89.6 ± 6.2
Sequence 5			
Method	MAD (mm)	HD (mm)	Dice (%)
w/	1.16 ± 0.34	4.27 ± 1.17	88.4 ± 2.1
w/o	1.26 ± 0.66	5.01 ± 2.22	87.5 ± 5.1

Table 2. Quantitative evaluation of the segmentation results w/ and w/o the proposed constraint on the region-based active contours. The values in each cell indicate mean ± standard deviation calculated over all frames in each sequence.

be obtained by including more energy terms such as edge-based energies, which is out of the scope of this paper.

4.2.2 Quantitative evaluation

We use the following metrics to quantitatively evaluate the segmentation results compared to the manual segmentation: Mean Absolute Distance (MAD), Hausdorff Distance (HD), and the Dice coefficient. Let C_1 and C_2 be two contours to be compared, then

$$\begin{aligned}
\text{MAD}(C_1, C_2) &= \frac{\int_0^1 d(C_1(s), C_2) |C_1'(s)| ds}{2|C_1|} \\
&\quad + \frac{\int_0^1 d(C_2(s), C_1) |C_2'(s)| ds}{2|C_2|}, \\
\text{HD}(C_1, C_2) &= \max\{\sup_s d(C_1(s), C_2), \sup_s d(C_2(s), C_1)\}, \\
\text{Dice}(C_1, C_2) &= \frac{2|\Omega_{C_1} \cap \Omega_{C_2}|}{|\Omega_{C_1}| + |\Omega_{C_2}|}. \tag{13}
\end{aligned}$$

Here, $d(p, C)$ is the minimum distance from point p to contour C , $|C|$ represents the contour length, Ω_C denotes the region inside C , $|\Omega_C|$ means the area of the region, and \sup indicates the supremum.

The results are summarized in Table 2. Regarding the mean of the metrics, a smaller MAD/HD or a larger Dice




Figure 4. The effect of the regularization parameter λ . The markers and bars denote the mean values and standard deviations of MAD, respectively. The curve shows the cubic fitting of the means. The results are averaged over all sequences.

coefficient indicates a more accurate segmentation. Generally, the performance with the proposed constraint is better than that without the constraint. The improvement in the Hausdorff distance is the most notable, which measures the largest error for each contour. This is due to the fact that part of the segmentation result is corrupted by the missing boundary while this error can be corrected by adding the shape constraint. Regarding the standard deviation of the metrics, a smaller standard deviation indicates the more stable performance. As shown in Table 2, the standard deviation with the proposed constraint is distinctly lower than that without the constraint, which shows the significance of the proposed constraint to improve the robustness of the active contour model.

4.3. Effect of λ

The most important parameter in our method is the weight λ of the nuclear norm regularization in (4). The larger λ is, the lower $\text{rank}(\mathbf{X})$ will be, which makes the output contours more similar to each other. Figure 4 shows the influence of λ to the segmentation accuracy averaged over the tested sequences in our experiments. As the λ increases in a proper range, both of the accuracy and the variance are improved, which demonstrates the advantage of the proposed constraint. However, after the turning point the accuracy decreases, since the over-regularization gives a large bias in shape estimation.

In our experiments, we selected λ empirically and applied the same λ to all sequences. The curve in Figure 4 shows that the accuracy changes smoothly over λ and the performance is stable in a wide range. Another alternative way is to choose a constant K specifying the degree of freedom allowed for shape variation and then solve the model with a decreasing sequence of λ until $\text{rank}(\mathbf{X})$ reaches K .

4.4. Convergence and computational time

Our algorithm is implemented in MATLAB and tested on a desktop with a Intel i7 3.4GHz CPU and 8G RAM. The computational time and the number of iterations required are listed in Table 1. The experiments showed that the algorithm with the shape constraint converged faster than that without shape constraint. This can be explained by the fact that the added constraint will make the active contour model better regularized, which results in faster convergency and fewer iterations. For instance, in the case of missing boundary, the active contour without the shape constraint may evolve further away from the ground truth and converge slowly, where we have to stop it manually after a number of iterations, *e.g.* 50 iterations in our experiments. Instead, the algorithm with the constraint converged within 40 iterations for all sequences. Moreover, the only additional computation brought by adding the constraint is the singular value decomposition on \mathbf{X} . Since the size of \mathbf{X} is merely $2p \times n$, where p is the number of landmarks and n is the number of frames, the singular value decomposition is not very time-consuming. Therefore, the algorithm with the proposed constraint is even faster in computation compared to that without the constraint.

5. Discussion

In this paper, we proposed a simple and effective way to regularize the group similarity of shapes in the active contour model based on low-rank modeling and rank minimization. Following are some discussions.

Relation to other works. In the active shape model [11], a candidate shape is constrained in a shape space described by $\mathcal{C}(\mathbf{w}) = \bar{\mathbf{C}} + \Phi\mathbf{w}$, where $\bar{\mathbf{C}}$ is the mean shape, Φ is a matrix with columns representing different modes of shape variation, and \mathbf{w} is a vector of coefficients. Since the number of columns of Φ is often small, the candidate shape is limited in a low-dimensional space. Hence, the active shape model also admits a low-dimensional assumption, which is similar to our low-rank assumption. The difference is that the shape space in the active shape model is constructed from offline training, while the low-rank model in our method is constructed unsupervisedly along with the segmentation. The low-rank constraint has also been adopted in other vision problems such as structure from motion [3], tracking [28], and motion segmentation [30]. Matrix factorization was used in these works to extract structure or motion information from detected feature points. Another topic closely related to our work is shape analysis such as shape clustering [26, 25]. The difference is that sample shapes are provided as inputs in shape analysis while they are the outputs of segmentation.

Why not using level sets. We use the landmarks to represent the contour instead of level sets. The reason is that

the low-rank property in (1) will not hold if the level-set representation is used. For instance, if there are n contours represented by the zero-level sets of n signed distance functions (SDFs), and the contours are identical in shape but different in location, the matrix consisting of the vectorized SDFs has a rank of n , which is full-rank. Other variational methods for image segmentation [13] also have this issue.

Limitations. A limitation of using the shape similarity constraint is the possibility of removing frame-specific details of the shapes. The trade-off between noise removal and signal preserving is a fundamental challenge in many problems. A possible solution in our problem is to refine the segmentation by running an active contour model that is more sensitive to local features with our results being both initialization and templates to constrain the curve evolution.

Extension to 3-D segmentation. In this paper, we demonstrate the advantages of our method based on 2-D active contours. In practice, it can be naturally applied to 3-D deformable models. The implementation for 3-D segmentation will be our future work.

Acknowledgments

We gratefully acknowledge Dr. Albert J. Sinusas from the Yale University Section of Cardiology for providing access to the echocardiographic data that was used here for testing and which was acquired under NIH grant R01 HL082640. We also gratefully acknowledge the members of Dr. Sinusas' team who set up the experiments and acquired these data, especially Dr. Ben Lin and Mr. Donald Dione.

Appendix

All ultrasound images were acquired using an open chest canine preparation at the Yale Translational Research Imaging Center and the studies were performed with appropriate institutional approvals. 2D short-axis images were acquired at a frame rate of 122-149 fps using a Philips iE33 ultrasound system with X7-2 phased array transducer at a nominal frequency of 4.4 MHz (Philips Healthcare, Andover, MA). The exported B-mode data sets were temporally downsampled to generate the sequences used here with a frame rate of 30-40 fps. The two sets of data in Figure 3 represent continuous runs of data selected to begin at two different times during the cardiac cycle. Manual tracings were performed by computer vision researchers working with these data.

References

- [1] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009. 3, 4
- [2] A. Blake and M. Isard. *Active contours*. Springer, 2000. 1

- [3] C. Bregler, A. Hertzmann, and H. Biermann. Recovering non-rigid 3d shape from image streams. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2000. 7
- [4] J. Cai, E. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20:1956, 2010. 3
- [5] E. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11, 2011. 3
- [6] E. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009. 3
- [7] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997. 1, 3
- [8] T. Chan and L. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001. 1, 3
- [9] T. Chan and W. Zhu. Level set based shape prior segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2005. 1
- [10] Y. Chen, H. Tagare, S. Thiruvankadam, F. Huang, D. Wilson, K. Gopinath, R. Briggs, and E. Geiser. Using prior shapes in geometric active contours in a variational framework. *International Journal of Computer Vision*, 50(3):315–328, 2002. 1
- [11] T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models – their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995. 1, 7
- [12] D. Cremers. Dynamical statistical shape priors for level set-based tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(8):1262–1273, 2006. 1
- [13] D. Cremers, T. Pock, K. Kolev, and A. Chambolle. Convex relaxation techniques for segmentation, stereo and multiview reconstruction. In *Markov Random Fields for Vision and Image Processing*. MIT Press, 2011. 7
- [14] D. Cremers, M. Rousson, and R. Deriche. A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *International Journal of Computer Vision*, 72(2):195–215, 2007. 1
- [15] P. Etyngier, F. Ségonne, and R. Keriven. Shape priors using manifold learning techniques. In *Proceedings of IEEE International Conference on Computer Vision*, 2007. 2
- [16] M. Fazel. *Matrix rank minimization with applications*. PhD thesis, Stanford University, 2002. 3
- [17] A. Jain, Y. Zhong, and S. Lakshmanan. Object matching using deformable templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(3):267–278, 1996. 1
- [18] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988. 1, 3
- [19] M. Leventon, W. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2000. 1
- [20] R. Malladi, J. Sethian, and B. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(2):158–175, 1995. 1
- [21] D. Metaxas. *Physics-based deformable models: applications to computer vision, graphics, and medical imaging*. Kluwer Academic Publishers, 1996. 1
- [22] Y. Nesterov. Gradient methods for minimizing composite objective function. *CORE Discussion Papers*, 2007. 3, 4
- [23] J. Noble and D. Boukerroui. Ultrasound image segmentation: A survey. *IEEE Transactions on Medical Imaging*, 25(8):987–1010, 2006. 1
- [24] N. Paragios. A level set approach for shape-driven segmentation and tracking of the left ventricle. *IEEE Transactions on Medical Imaging*, 22(6):773–776, 2003. 1
- [25] O. Sidi, O. van Kaick, Y. Kleiman, H. Zhang, and D. Cohen-Or. Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. *ACM Transactions on Graphics*, 30(6):126, 2011. 7
- [26] A. Srivastava, S. H. Joshi, W. Mio, and X. Liu. Statistical shape analysis: Clustering, learning, and testing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):590–602, 2005. 7
- [27] L. Staib and J. Duncan. Boundary finding with parametrically deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(11):1061–1075, 1992. 1
- [28] L. Torresani and C. Bregler. Space-time tracking. In *Proceedings of the European Conference on Computer Vision*, 2002. 7
- [29] A. Tsai, A. Yezzi Jr, W. Wells III, C. Tempany, D. Tucker, A. Fan, W. Grimson, and A. Willsky. Model-based curve evolution technique for image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2001. 1
- [30] R. Vidal and R. Hartley. Motion segmentation with missing data using powerfactorization and gpca. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2004. 7
- [31] C. Xu and J. Prince. Snakes, shapes, and gradient vector flow. *IEEE Transactions on Image Processing*, 7(3):359–369, 1998. 1
- [32] A. Yuille, P. Hallinan, and D. Cohen. Feature extraction from faces using deformable templates. *International Journal of Computer Vision*, 8(2):99–111, 1992. 1
- [33] S. Zhang, Y. Zhan, M. Dewan, J. Huang, D. N. Metaxas, and X. S. Zhou. Towards robust and effective shape modeling: Sparse shape composition. *Medical image analysis*, 16(1):265–277, 2012. 2
- [34] S. Zhu and A. Yuille. Region competition: Unifying snakes, region growing, and bayes/mdl for multiband image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(9):884–900, 1996. 1, 3, 4
- [35] Y. Zhu, X. Papademetris, A. Sinusas, and J. Duncan. Segmentation of the left ventricle from cardiac mr images using a subject-specific dynamical model. *IEEE Transactions on Medical Imaging*, 29(3):669–687, 2010. 1