Shape Regularized Active Contour using Iterative Global Search and Local Optimization *

Tianli Yu[1], Jiebo Luo[2], and Narendra Ahuja[1]

[1]Beckman Institute & ECE Department
Univ. of Illinois at Urbana-Champaign
Urbana, IL 61801

[2]Kodak Research Laboratories
Eastman Kodak Company
Rochester, NY 14650-1816

Abstract

Recently, nonlinear shape models have been shown to improve the robustness and flexibility of segmentation. In this paper, we propose Shape Regularized Active Contour (ShRAC) that incorporates existing nonlinear shape models into the classical active contour approach. ShRAC uses a discrete representation of the contour to allow efficient combinatorial search. The search for optimal contour is performed by a coarse-to-fine algorithm that iterates between combinatorial search and gradient-based local optimization. First, Multi-Solution Dynamic Programming (MSDP) is used to generate initial candidates by minimizing only the image energy. In the second step, a combination of image energy and shape energy determined by a given prior shape model is minimized for the initial candidates using a local optimization method and the best one is selected. To have diverse initial candidates, we employ a Clustered Solution Pruning procedure in the MSDP search space. Finally, Local Shape Regularization is used to feed shape constraints back into the new MSDP search space of the next iteration. Our search strategy combines the advantages of global combinatorial search and local optimization, and has shown excellent robustness to local minima caused by distracting suboptimal segmentations. Experimental results on segmentation of different anatomical structures using ShRAC are provided.

1. Introduction

Segmentation of anatomical structures is often a critical component of medical imaging systems, such as a Computer-Aided Diagnosis (CAD) system and a Patient Information System (PIS). In chest radiography, researchers have developed numerous methods for segmenting the lung fields, rib cage, heart, clavicles, blood vessels, as well as abnormal structures such as lung nodules. However, given the projective nature of chest radiography, superimposed anatomical structures make images complicated and challenging for diagnosis for both radiologists and computerized systems. The clinical importance of chest radiography and its complicated nature continue to drive researches on developing robust segmentation algorithms to assist radiologists and improve automation.

While the nature of chest radiography images poses many challenges for robust algorithms, the limited domain of the problem also provides an opportunity to incorporate prior knowledge of the shape, and in some cases the appearance, of the anatomical structures of interest. After the early years of research, which focused on rule-based reasoning and hand-crafted shape models [4], more recent approaches are focusing on statistically learning the shape (and appearance when applicable) of an object from a set of training examples, and use the learned model to constrain the search. Active Shape Model (ASM) [2] is such a successful segmentation method that uses linear shape and appearance models extracted by principle components analysis (PCA). The problem with linear models is that they admit invalid shapes when the shape distribution is in fact nonlinear, and therefore multiple models are necessary to span a nonlinear shape distribution while excluding invalid shapes.

It is desirable to develop a method to incorporate prior nonlinear shape constraints into segmenting different anatomical structures (e.g., lung fields, clavicles, ribs) in medical images. This method could also serve as a flexible segmentation tool for those objects whose shapes may vary nonlinearly, giving rise to what appear to be multiple subclasses, without the need to build separate models. For the specific problem of lung field segmentation, it is desired that the segmentation algorithm starts with a fixed initial shape

*The support of the National Science Foundation under grant IBN 04-22073 is gratefully acknowledged. Part of this work was done while Tianli Yu was an intern in Kodak Research Laboratories.

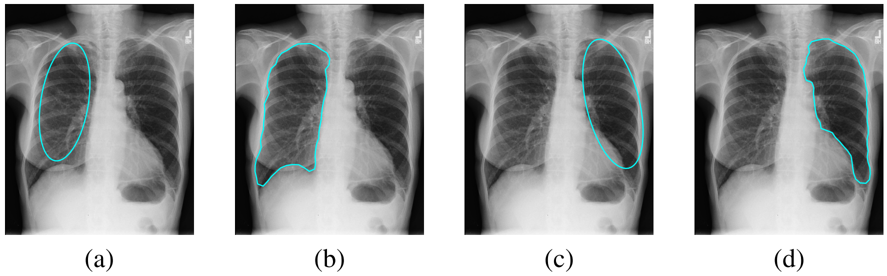

Figure 1. Starting with the same elliptic shape, a segmentation algorithm should converge to the correct shape without model selection (e.g. to adapt to whether it is the left or the right lung field). Here we show the result of the proposed ShRAC algorithm: (a) Initial contour on the right lung. (b) Final contour on the right lung. (c) Initial contour on the left lung (note that it is the same shape as in (a)). (d) Final contour on the left lung.

(Fig. 1), and automatically adapt and converge to the correct shape, without any model selection in advance, e.g., to decide whether it is the left or right lung field. In addition, the method should also be robust to noise and clutter that are commonly seen in medical images.

Nonlinear shape priors have been incorporated into nonlinear ASM [9] and variational image segmentation with nonlinear shape statistics [3]. These algorithms have been shown to be more flexible and powerful in model building and therefore more successful in segmenting objects with nonlinear shape variations. However, they also pose many difficulties to the design of a search strategy that balances the use of shape prior constraints and actual image structures to find a globally optimal solution.

In the work by Cremers et al. [3], a Mumford-Shah functional type image energy and a shape energy derived from Kernel Space Density Estimation (KSDE) are combined to construct an energy function for segmentation. The KSDE-based shape energy enforces object-specific nonlinear shape constraints, in contrast to general smoothness constraints in conventional snakes [5]. However, this shape energy term makes it difficult to minimize the energy function globally. Instead, gradient-descent is employed in their work, which often converges to a local minimum.

Algorithms developed for ASM [2] address this problem by not minimizing a single energy function. Instead, the search iterates between two steps. In one step, ASM searches for the best landmarks, which is equivalent to minimizing an image energy term. In the other step, the contour is regularized to the space of training shapes, which can be viewed as minimizing a shape energy term. Although this alternating process provides certain capabilities for the algorithm to jump out of a local minimum and find a better solution, the lack of a global energy function may introduce oscillations when trying to minimize two separate targets. Therefore, a more robust method is needed to avoid potential oscillations.

In this paper, we propose the Shape Regularized Active Contour (ShRAC) and the associated search strategy under the framework of energy minimization. ShRAC incorporates nonlinear shape priors through Kernel Space Density Estimation, similar to [3]. Our search strategy combines the advantages of combinatorial search, which provides global optimality, and gradient-based local optimization. In particular, we propose to first use Multi-Solution Dynamic Programming (MSDP) to minimize the image energy as the initial candidate selection step, and then use a local optimization method that eventually minimizes a combination of the image energy, and shape energy determined by the prior shape models. In addition, we propose a Clustered Solution Pruning process in the MSDP search space to ensure the diversities of the selected candidates. Local Shape Regularization is also used to help feed the shape information from local optimization back into the MSDP search space for the next iteration. Our combined search strategy results in improved robustness to image noise and various distracting structures that present in medical images.

Although the main problems solved in this paper are in the context of medical image segmentation, the framework and the associated search strategies should be applicable to other object segmentation tasks involving joint constraints from image features and prior shape models.

This paper is organized as follows: Section 2 gives our formulation of ShRAC. Section 3 presents the iterative search strategy. Section 4 uses two experiments to demonstrate the robustness of ShRAC, and Section 5 gives some conclusions and outlines our related future work.

2. Problem Formulation

We formulate ShRAC based on a nonlinear shape distance measure instead of the smoothness term used in a conventional active contour. It is inspired by the work of Cremers et al. [3] that uses nonlinear shape priors for variational image segmentation. In particular, we are interested in a discrete formulation so that efficient combinatorial search methods can be applied.

2.1. A Discrete Formulation for ShRAC

Similar to the conventional Active Contour, ShRAC tries to find a contour c that minimizes an energy function $E(c)$, which is the weighted sum of two terms:

$$E(c) = E_{shape}(c) + wE_{image}(c) \tag{1}$$

where $E_{image}(c)$ is the image energy term, and $E_{shape}(c)$ is the shape energy term that measures the difference between the contour c and the training shapes. We adopt the shape energy measure used in [3], which uses a distance measure based on Kernel Space Density Estimation (KSDE):

$$\begin{aligned} E_{shape}(c) &= \sum_{j=1}^{r}\left(\sum_{i=1}^{m}\alpha_i^j \tilde{k}(c_i, c)\right)^2 \cdot (\lambda_j^{-1} - \lambda_{\perp}^{-1}) \\ &\quad + \lambda_{\perp}^{-1} \cdot \tilde{k}(c, c) \end{aligned} \tag{2}$$

where α_i^j is the jth eigenvector of the centered kernel matrix, $\tilde{k}(\cdot,\cdot)$ is the centered kernel function, and c_i is the ith training example. Furthermore, $\lambda_j = (1/m)\tilde{\lambda}_j$, where $\tilde{\lambda}_j$ is the jth eigen value of the centered kernel matrix, m is the number of training examples, and $\lambda_{\perp}$ is a constant value chosen to replace all the smaller eigenvalues.

We define the image energy in terms of the intensity differences on two side of a contour. This is because many anatomical structures do not have homogeneous intensities, therefore a piecewise constant intensity model is not always appropriate. $E_{image}(c)$ can be computed as the integration of local edge strength $h_e(\cdot)$ along the contour.

$$E_{image}(c) = \frac{\int h_e[v(s)]ds}{L(c)} \tag{3}$$

where $v(s)$ can be viewed as a parametric representation of points on the contour. $L(c)$ is the length of the contour which is used as a normalization term. The term $h_e(\cdot)$ can be computed as the absolute or signed intensity difference between pixels on the two sides of the contour. In many medical applications, such as lung field segmentation or rib segmentation, a signed intensity difference will give better performance since the intensity difference generally has the same sign along the contour.

To use efficient combinatorial optimization methods, we represent the contour as a set of control points $c = \{v_i | i = 1...N\}$ sequentially connected by line segments. Then we rewrite Equation (3) in a discrete form as:

$$E_{image}(c) = \frac{\sum_{i=0}^{n-1} h_e[v_i \vec{v_{i+1}}] * l(v_i \vec{v_{i+1}})}{L(c)} \tag{4}$$

where $h_e[v_i \vec{v_{i+1}}]$ represents the average edge strength along $v_i \vec{v_{i+1}}$ and $l(v_i \vec{v_{i+1}})$ is the length of $v_i \vec{v_{i+1}}$. To further simplify the image energy term, we ignore the length

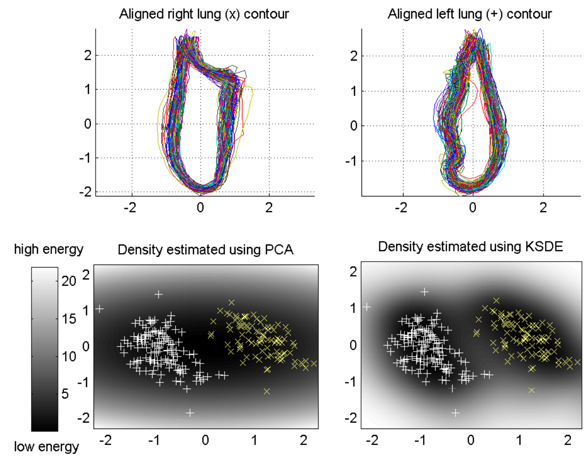

Figure 2. Comparison of distance functions learned by PCA (lower left) and KSDE (lower right) from the training shapes of both left (+) and right (x) lungs.

difference between different contour segments and treat $l(v_i v_{i+1})/L(c)$ as a constant. This simplification is appropriate because the control points in all the training examples are equally spaced, and therefore the shape energy will push these control points towards the equally spaced positions. This enables us to write the discretized energy function as:

$$E(c) = E_{shape}(c) + w' \sum_{i=0}^{n-1} h_e[v_i \vec{v_{i+1}}] \tag{5}$$

2.2. Nonlinear Shape Learning in ShRAC

The nonlinear shape model adopted in ShRAC allows us to capture a set of shapes containing nonlinear variations. For example, we can put together training shapes from both left and right lungs and build a single shape model to represent the distribution of the acceptable shapes.

Fig. 2 contrasts the performance of PCA and KSDE in modeling the shape energy for a set of lung shapes. Both left and right lung shapes (Fig. 2, first row) are aligned into a common coordinate space and represented as a sequence of control points. The learned Mahalanobis distances from PCA and from KSDE using (2) are plotted in the second row, along the axes spanned by the eigenvectors that correspond to the largest two eigenvalues of PCA. Clearly, KDSE shows two clusters, separated by a high-energy ridge (in light gray) in the shape space, for left and right lung fields. It gives a better characterization of the training shape distribution than the single cluster energy from PCA, which actually assigns the lowest energy (highest likelihood) to the invalid shapes in between the two clusters.

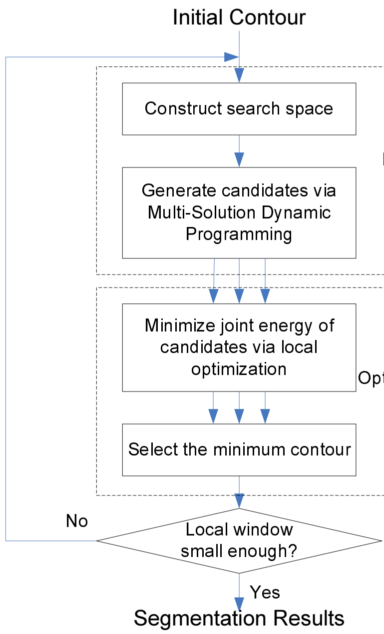

Figure 3. The flow diagram of the proposed search strategy.

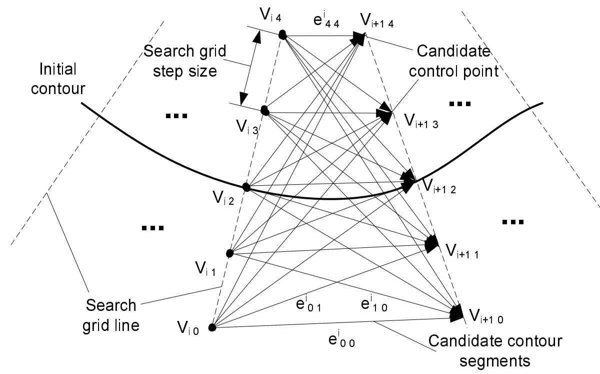

Figure 4. A MSDP search space constructed from an initial contour. $\{V_{ij}\}$ are candidate points for control point i and $\{e^i_{jk}\}$ are candidate contour segments between control point i and $i+1$.

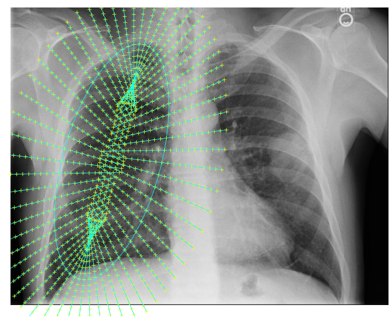

Figure 5. A search grid constructed from the initial elliptical shape for lung field segmentation. Each line is a search direction for a control point and each + sign is a candidate control point.

3. Search Strategies for ShRAC

3.1. Combinatorial Search v.s. Local Optimization

Note that the image energy term in the energy function (5) has a good property: it can be expressed as a sum of functions of two consecutive control points. Functions of this form can be optimized efficiently using combinatorial search methods such as Dynamic Programming. Dynamic programming guarantees finding a global minimum within polynomial time [1]. This ensures the robustness of the algorithm to initial conditions and local minima.

The shape energy term has a more complex form. The addition of this term makes the whole function no longer decomposable into a sum of functions of two consecutive control points. Except for stochastic minimization (often impractical), we can only rely on local gradient-based optimization to solve this problem. Unless the energy function is convex, local gradient-based optimization usually finds a local minimum that is close to the initial contour. With the presence of imaging noise and distracting structures in typical medical images, a local optimization method can be easily trapped in local minima.

Previous research has shown that combinatorial search strategies can be used to generate good initializations for various snake algorithms [7]. We adopt a hybrid search strategy for ShRAC that combines the advantages of combinatorial search and local optimization. The algorithm iterates between the following two stages. In the first stage, we use Multi-Solution Dynamic Programming (MSDP) to generate initial candidates by minimizing E_{image} alone. In the second stage, we add the shape regularization term and use a local optimization method to minimize the entire energy in (5), starting with these candidates. After local optimization, a best contour is selected and used as the input of the next iteration. We also use a coarse-to-fine scheme involving multiple cycles: initally we use a large window to compute the edge energy and a large step size in MSDP, which can be viewed as a coarse search stage, and then reduce both of them to perform finer searches. The whole algorithm is illustrated in Fig. 3.

3.2. Multi-Solution Dynamic Programming

First, we generate initial candidates for local optimization by minimizing the image energy E_{image} alone. This is done by first constructing a search space arround the initial contour and then using Dynamic Programming to select a closed contour that minimizes the image energy.

Fig. 4 illustrates the construction of a search space from an initial contour. First, we choose a search direction for each control point. These directions can be the normal to

1. For each candidate point V_{1j} on the search direction of the first control point V_1, find the top p contours with minimum weights that pass through this candidate point.
    1.1 Set the local weight of V_{1j} as 0, local weight of other candidates V_{1m} (m <> j) to +infinity

    1.2 For control point V_2 to V_n,
        For each candidate point, compute all the possible local weights by adding the local weights of candidates of previous control point and the weights of the connecting edges. Select the top p minimum weight among these sums as the local weights of the current candidate and record the edges selected.

    1.3 compute the final weights by adding the local weights from candidates of V_n and the last set of edge weights from V_n to V_{1j}, select the first p mininum one.

2. Find top p minimum weight contours from the n*p contours from step 1

3. Trace back the local edge record and output the p minimum contours

Figure 6. The Multi-Solution Dynamic Programming Algorithm.

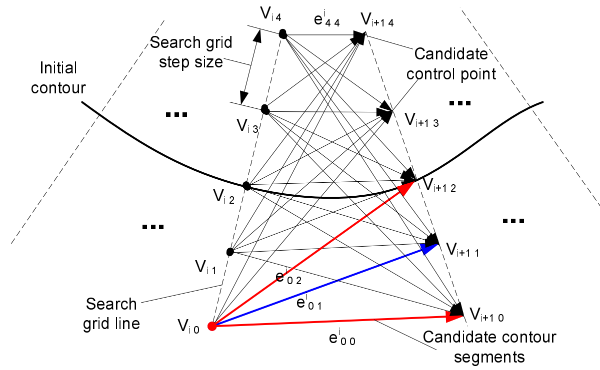

Figure 7. Clustered Solution Pruning: blue edge $e^i{}_{01}$ is only preserved if its image energy is less than both of the red edges $e^i{}_{00}$ and $e^i{}_{02}$.

the initial contour, radial directions from the center of mass of the initial contour, or other object-specific directions. Next, we construct a search grid by placing equally spaced control point candidates along the search direction on both sides of the initial positions of the control points. Directed edges are created to connect neighboring control point candidates, and they become the candidates for the contour segments. In this directed graph, each edge is assigned a weight $h_e(V_{ia}V_{(i+1)b})$ which is the local edge strength. This way, we construct a directed graph with vertices representing control point candidates and edges representing contour segment candidates. All possible closed contours in this directed graph constitute the search space for Dynamic Programming. Fig. 5 shows a search grid (constructed for lung field segmentation), where the search direction at a control point is chosen to be the average of the normal direction and the radial direction.

Dynamic Programming [1] is used to efficiently find a minimum-weight contour in such a search space. In order to generate several different initial candidates for the subsequent local optimization, we modify the dynamic programming algorithm to find the first p closed contours that minimize the image energy E_{image}. The modified algorithm, named Multi-Solution Dynamic Programming (MSDP), is summarized in Fig. 6. Note that standard Dynamic Programming is a degenerate case of MSDP when $p = 1$.

The computational complexity of the original dynamic programming algorithm is $O(nm^2)$, where n is the number of control points and m is the number of candidates for each control points. The MSDP's complexity is $O(pnm^2)$, which means it grows linearly with the number of candidate contours needed.

3.3. Clustered Solution Pruning for MSDP

The solutions from MSDP alone consist of top ranking contours that minimize the image energy E_{image}. Since these results are used as initial candidates for the subsequent local optimization, we argue that it is often critical that these candidates be as diverse as possible for the following reasons.

First, the image energy is a continuous function, which means minimum contours from MSDP will often cluster together and might not include the optimal contour for the entire energy. Second, medical images are intrinsically noisy and ambiguous, and often many competing interpretations exist. It is critical to preserve multiple, and more importantly, distinct solutions when only image energy is minimized. If we apply local optimization starting from similar candidates, the final solution often converges to the same minimum, which may or may not be global, thus defeating the purpose of using MSDP.

We propose to prune the search space of MSDP so that with a moderate p we can obtain a diverse set of initial candidates. We refer to this process as Clustered Solution Pruning, which is similar in spirit to the procedure employed in [10] where non-distinct proposals are pruned by minimizing the approximate Kullback-Leibler (KL)-divergence. In practice, we propose a simple yet effective scheme. For each directed edge e^i_{jk} in the search graph (Fig. 7), we compare it with its two neighboring edges e^i_{jk+1} and e^i_{jk-1}; the edge is pruned from the search graph if it is not the minimum among the three edges.

This pruning reduces the search space to that only consists of candidate segments with locally minimum image energy. In essence, this makes our search focus on local minima of the image energy. By starting from multiple locally minimum solutions to the image energy, the subsequent local optimization of the whole energy will have a better chance to reach a global minimum solution. Fig. 8 shows an example where the proposed pruning method helps MSDP to find more diversified initial candidates that eventually lead to the desired solution.

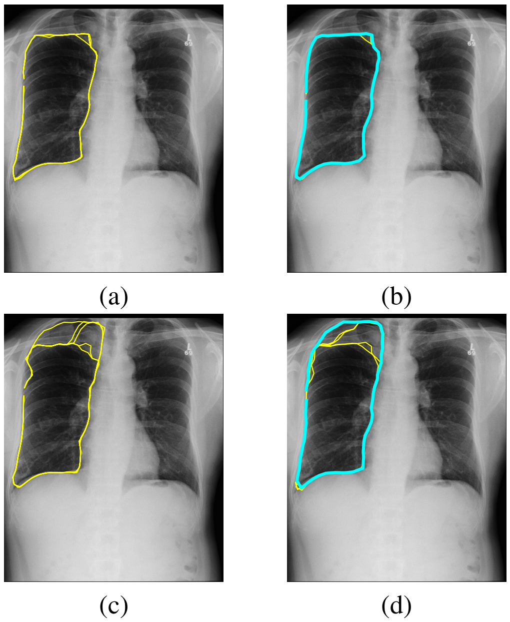

Figure 8. Effect of Clustered Solution Pruning. (a) Initial candidates selected by MSDP without pruning. (b) Local optimization result of the candidates from (a). (c) Initial candidates selected by MSDP with pruning. (d) Local optimization result of candidates from (c). The bold blue contours are the finally selected solutions.

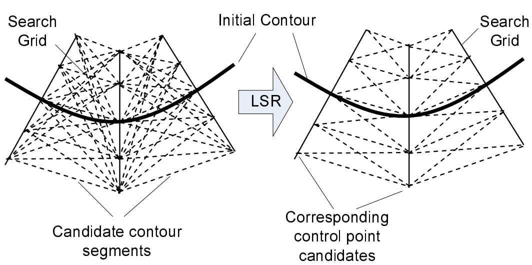

Figure 9. MSDP Search graph before (left) and after (right) Local Shape Regularization.

3.4. Local Shape Regularization for MSDP

The MSDP search space constructed in Section 3.2 covers a wide range of complex shapes. This allows the search results from MSDP depend less on the initialization and are more robust to local minima. However, usually after the first iteration of the search algorithm, the contour selected as input for the next iteration already has fairly good information about the expected shape in the image. Therefore, it is not necessary to search again the fully connected search graph as in Fig. 4. For example, it is not necessary to include the edge $V_{i0}V_{(i+1)4}$ of Fig. 4 in the search graph since it is very different in orientation from the initial contour.

We use the following Local Shape Regularization process in constructing the MSDP search space. All the edges in the search graph are compared to the corresponding edge in the initial contour. Those edges with orientation differences larger than a threshold are removed from the search graph (see Fig.9). This procedure further prunes the search space and allows computation to be focused on those candidates that have similar shapes to the initial contour. It helps feed the shape information from local optimization back into the MSDP search space.

4. Experimental Results

4.1. Lung Field Segmentation

First, we applied ShRAC to the lung field segmentation to demonstrate the power of using the nonlinear shape model to capture widely different set of shapes as a single class. We trained a single shape model, using manually segmented image masks of both left and right lungs (200 shapes). A set of 60 control points is used to represent the contour of a lung field (left or right). Each shape is aligned with a coordinate system that has the origin at the shape center and y-axis coincides with the shape's major axis. It is also scaled to have unit area. Such normalization ensures that the shape training uses common reference coordinates. The two intersection points of a contour with the x-axis are used as anchor points and the other 58 control points are evenly distributed between these anchor points (29 on each side).

Our segmentation algorithm follows the flow diagram in Fig. 3. We chose to perform two iterations of the coarse-to-fine search. The search area covered by MSDP is large enough to account for possible poor initializations (Fig. 5). In MSDP, we empirically chose the top 5 contours as the initial candidates for local optimization. The local optimization method we used is the BFGS Quasi-Newton method for unconstrained nonlinear minimization in Matlab [8].

To evaluate the performance of the algorithm, we implemented another algorithm as a proxy of the method in [3]. It is a coarse-to-fine iterative algorithm, which at each iteration uses the same local optimizer as used in ShRAC. The algorithm starts with a large windows size to compute edge energy and image gradient, and gradually reduces them at each iteration. We refer to this algorithm as IQN (Iterative Quasi-Newton) method.

Both of these algorithms are initialized by a program that computes a rough center line of each lung field (by searching for valleys in the image). The center line is used to scale and orient an elliptical shape to roughly cover the lung field. Note the initialization algorithm is not supposed to be robust and it sometimes fails to cover the entire lung field correctly.

Fig. 10 shows two x-ray images segmented by ShRAC and IQN. The results show that IQN is more susceptible to local minima and therefore to initialization. On the other

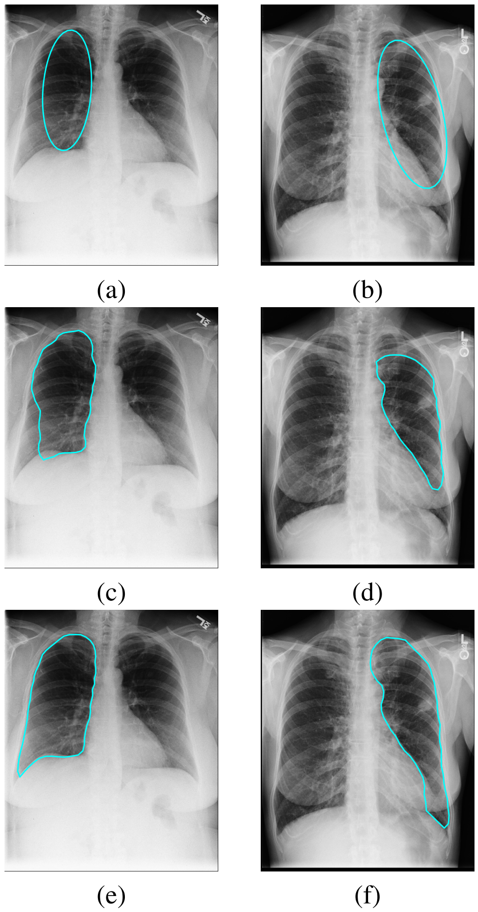

Figure 10. Segmentation results of IQN and ShRAC for lung fields. (a, b) Initial contours, (c, d) Segmentation result by IQN, (e, f) Segmentation result of ShRAC.

hand, ShRAC shows good robustness with respect to initialization. To compare the performances quantitatively, we ran these two algorithms on two data sets (CR - 136 images and JSRT - 247 images) for which manually segmented ground truth is available. We use the following overlap ratio Ω to measure performance:

$$\Omega = TP/(TP + FP + FN) \tag{6}$$

where TP stands for true positive (the area correctly classified as object), FP for false positive (area incorrectly classified as object), and FN for false negative (area incorrectly classified as background).

The overlap ratios obtained by these two algorithms on the two data sets are listed in Table 1 (excluding training images). For these two data sets, we also list the performance data of an ASM implementation from another study. Recall that the ASM algorithm uses linear models, so two separate models are built for left lung and right lung, and a model selection stage is used to select the correct model before the algorithm begins. The ASM algorithm also requires more elaborate training because many landmarks need to be identified and meticulously located to properly train the object appearance model.

Table 1. Accuracy comparison (mean and standard deviation of the overlap ratio Ω) of the three segmentation algorithms on lung field segmentation for two data sets (CR - 136 images, JSRT - 247 images).

	IQN	ShRAC	ASM
CR	0.841±0.081	0.891±0.037	0.893±0.060
JSRT	0.881±0.070	0.907±0.033	0.920±0.022

In Table 1, ShRAC clearly outperforms IQN. The use of MSDP in ShRAC leads to 2.6% - 5.0% increase in overall segmentation performance. Visually, JSRT is considerably easier than CR, partially explaining why the increase is smaller. It is noteworthy that the improvement is due to capturing of shape details near the lung boundary, and therefore translates into substantial visual improvement in the segmentation. For a blob-like object such as the lung field, getting the bulk of the lung field right is almost trivial compared to locating the tips and corners accurately; for example, the dramatic improvement shown in Fig. 8 amounts to only a 10% change in the overlap ratio. Our ShRAC algorithm achieves almost the same performance as the two-model ASM algorithm that requires accurate landmarks in training. Along with the other advantages mentioned earlier, this shows the promise of ShRAC as a more flexible and powerful segmentation tool.

4.2. Clavicle Segmentation

To verify the generality of ShRAC, we also applied it to segmenting clavicles. Segmenting clavicles is quite challenging because there are many similar structures in the vicinity, including the first three ribs and shoulder joints. Since the two clavicles are always symmetric, we merge them together and represent them as a single contour (Fig. 11).

We used 30 manually segmented contours as the training examples. The search directions for clavicles in MSDP are restricted to the vertical direction since the bones are very thin and lack of horizontal features. The same lung field centerlines are used to estimate a rough position and scale of the initial mean shape. Fig. 11 shows three results obtained by ShRAC and IQN. Locating bone structures is even more difficult for a local optimization algorithm such as IQN because there are many local minima created by other bones in the neighborhood. However, the MSDP-based ShRAC is capable of jumping out of local minima and finding superior solutions. For quantitative evaluation, we computed the vertical displacement of the extracted control points to

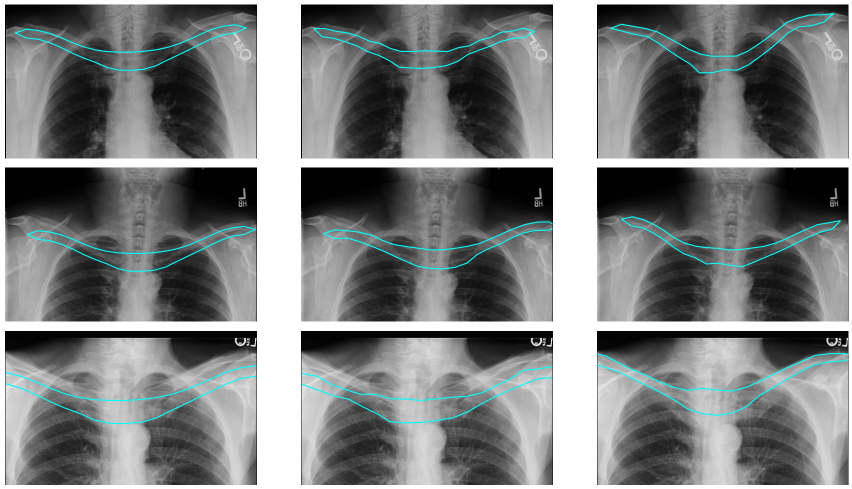

Figure 11. Segmentation results of IQN and ShRAC for clavicles. (a) Initial contours, (b) Segmentation results of IQN, (c) Segmentation results of ShRAC. The average vertical displacement over 100 images (size 512x625, only upper half are shown in the figure) using IQN and ShRAC is 13.6 pixels and 2.3 pixels, respectively.

their projections in the vertical direction on the manually segmented ground truth contours. The average vertical displacement over 100 images (size 512x625) using IQN and ShRAC is 13.6 pixels and 2.3 pixels, respectively.

5. Conclusions and Future Work

We have presented the Shape Regularized Active Contour approach along with an associated robust search strategy. ShRAC uses a nonlinear prior shape model obtained using Kernel Space Density Estimation, which makes it suitable for segmenting a class of objects that has nonlinear within-class shape variations. Our search strategy for ShRAC is implemented as a coarse-to-fine iterative algorithm that combines the advantages of combinatorial search and local optimization. It uses Multi-Solution Dynamic Programming to generate initial candidates that have minimal image energies, and then uses local gradient-based minimization of the entire energy to select the final optimal contour. A Clustered Solution Pruning process is applied to MSDP search space to reduce the number of redundant candidates. Local Shape Regularization is also used to feed shape information back into the MSDP search space for the next iteration. Our experiments demostrate robustness of our search strategy to initialization and distracting structures in medical images.

There are a few directions to extend ShRAC. Incorporating powerful appearance models, preferrably nonlinear, will be of major interest. Another direction is to extend ShRAC to 3D object segmentation, where Dynamic Programming can be replaced by other combinatorial search methods such as Graph Cuts [6].

References

[1] A. A. Amini, T. E. Weymouth, and R. C. Jain. Using dynamic programming for solving variational problems in vision. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12(9):855–867, 1990.

[2] T. F. Cootes, C. J. Taylor, D. Cooper, and J. Graham. Active shape models-their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

[3] D. Cremers, T. Kohlberge, and C. Schnorr. Shape statistics in kernel space for variational image segmentation. *Pattern Recognition*, 36(9):1929–1943, 2003.

[4] B. V. Ginneken, B. M. H. Romeny, and M. A. Viergever. Computer-aided diagnosis in chest radiography: A survey. *IEEE Trans. on Medical Imaging*, 20(12):1228–1241, 2001.

[5] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. Comput. Vision*, 1(4):321–331, 1988.

[6] V. Kolmogorov and R. Zabih. What energy functions can be minimizedvia graph cuts? *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(2):147–159, 2004.

[7] J. Liang, T. McInerne, and D. Terzopoulos. United snakes. In *ICCV '99: Proceedings of the International Conference on Computer Vision-Volume 2*, page 933, Washington, DC, USA, 1999. IEEE Computer Society.

[8] The MathWorks, Inc., http://www.mathworks.com/access/helpdesk/help/toolbox/optim/. *Matlab Optimization Toolbox*, 2004.

[9] S. Romdhani, S. Gong, and A. Psarrou. A multi-view non-linear active shape model using kernel pca. In *Proc. 10th British Machine Vison Conference*, pages 483–492, Sep 1999.

[10] Z. Tu and S.-C. Zhu. Image segmentation by data-driven markov chain monte carlo. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(5):657–673, 2002.