

GRAPH CUT SEGMENTATION USING AN ELLIPTICAL SHAPE PRIOR

Greg Slabaugh, Gozde Unal

Intelligent Vision and Reasoning Department
Siemens Corporate Research
Princeton, NJ 08540

ABSTRACT

We present a graph cuts-based image segmentation technique that incorporates an elliptical shape prior. Inclusion of this shape constraint restricts the solution space of the segmentation result, increasing robustness to misleading information that results from noise, weak boundaries, and clutter. We argue that combining a shape prior with a graph cuts method suggests an iterative approach that updates an intermediate result to the desired solution. We first present the details of our method and then demonstrate its effectiveness in segmenting vessels and lymph nodes from pelvic magnetic resonance images, as well as human faces.

1. INTRODUCTION

Segmentation is a fundamental task in image processing, and numerous methods have been developed. An interesting class of segmentation methods relies on energy minimization to partition an image into different regions. Included in this class are active contour methods [6, 3] as well as the recently popular graph theoretic techniques [1, 5].

In this first class, active contours, the energy is typically comprised of image terms (regional and/or boundary based) as well as intrinsic regularization terms. Based on the energy minimization, an initial contour iteratively deforms to move to the structures of interest. Although quite successful, these methods can be sensitive to the initialization of contour, since the energy minimization is subject to local minima, as well as “leaking,” which occurs when, due to noise, clutter, poor contrast, etc., the image data does not provide enough information to stop the contour at the desired location.

To increase robustness, researchers working on active contour methods have recently looked to shape priors [4, 9, 10, 11] to incorporate a priori shape information into the active contour evolution in order to further constrain the segmentation. Shape priors can be modeled by a known class of shapes or through statistical training.

In the second class of energy minimization segmentation techniques, those based on graph cuts [1, 5], the problem is formulated on a discrete graph. The graph is com-

posed using vertices representing the image pixels, as well as edges connecting the vertices, typically using 4 or 8 neighborhood connectivity. For example, in [1], the energy is comprised of a region term that assigns penalties based on labeling a pixel as foreground or background, as well as a boundary term that assigns a penalty based on the dissimilarity of adjacent pixels. Usually, the energy function to be minimized is the summation of the weights of the edges that are cut. Graph cut methods differ from active contour methods in that they are not iterative, and achieve global minimization for certain classes of energy functions.

Recently, iterative graph cut methods have been introduced in an attempt to combine active contours with graph theoretic techniques [12]. In this paper, we extend this method by including a shape constraint to the contour update.

1.1. Our Contribution

Despite their advantages, graph cut segmentation approaches can also lead to erroneous segmentations as shown in Figure 3(d). In this paper we segment an image using a graph cuts approach, but additionally include an elliptical shape prior to constrain the segmentation. Incorporating this shape constraint results in more robust segmentations.

We choose to work with an elliptical shape prior for several reasons. First, an ellipse is a powerful, descriptive shape that can model a multitude of objects, including a wide variety of anatomical structures like blood vessels and lymph nodes, the segmentation of which is our primary application. Second, by making use of a geometric primitive, we do not have to perform statistical shape training on a database of shapes. Such training is often difficult and time-consuming. In addition, an ellipse has a simple parametric equation that can be efficiently estimated from samples [7]. We take advantage of this in our approach.

To our knowledge this is the first paper to consider the use of shape priors in a graph cuts-based algorithm. We demonstrate that the use of such information results in more accurate segmentations than those resulting from standard graph cut methods.

2. SEGMENTATION APPROACH

We begin by briefly reviewing the graph theory that is used to minimize our energy function. We then describe how we build this function using the image data and the elliptical shape prior. We then list each step of our algorithm.

2.1. Graph Cut Theory

Consider an undirected graph $G = \langle V, E \rangle$ that is composed of vertices V and undirected edges E that connect the vertices. Each edge $e \in E$ is assigned a non-negative cost. There are two special vertices (also called terminals) in the graph that are identified as the source s and the sink t . With the exception of the terminals, the vertices are comprised of pixels P in the image. An example is shown in Figure 1(a). A cut C on the graph is a partition of V into two disjoint sets S and $T = V - S$ such that $s \in S$ and $t \in T$, as shown in Figure 1(b). The cost of the cut is the sum of the costs of all edges that are severed by the cut. The minimum cut problem is to find the cut with the smallest cost. There are numerous algorithms that solve this problem in polynomial time, see [2] for more details.

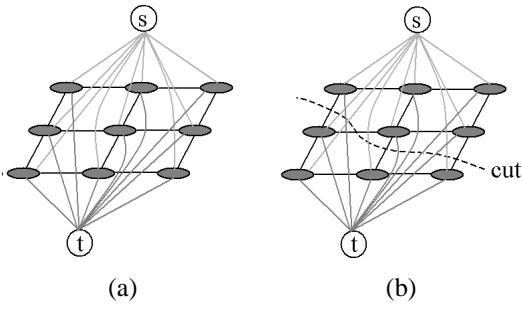


Fig. 1. A simple 2D graph for a 3x3 image (a) and its minimal cut (b). Figure based on [1].

2.2. Energy Function

Our goal then is to take a set of pixels P and labels L and compute a labeling f that minimizes an energy function. This function takes the standard form [8]

$$E(L) = \sum_{p \in P} D_p(f_p) + \sum_{p, q \in N} V_{p,q}(f_p, f_q), \quad (1)$$

where E is the energy, p and q are pixels, N is the neighborhood formed from the vertex connectivity. $D_p(f_p)$ measures the cost of assigning the label f_p to pixel p , while $V_{p,q}$ measures the cost of assigning the labels f_p, f_q to the adjacent pixels p, q . In our implementation, both $D_p(f_p)$ and $V_{p,q}$ are comprised of two terms, one from the image data and the other from the shape constraint.

Given an ellipse, we compute the mean intensity of the pixels inside and outside the ellipse, u_i and u_o , respectively. Our image-based energy term is inspired by [3] and is

$$D_p(\text{foreground}) = |I(p) - u_i| \quad (2)$$

$$D_p(\text{background}) = |I(p) - u_o| \quad (3)$$

and $V_{p,q}$ is the same as described in [1].

For the shape prior term, we first form a binary image M we call a *shape mask*, which is 0 inside the ellipse and 1 outside the ellipse. Our shape-based energy term is then

$$D_p(\text{foreground}) = |M(p) - 1| \quad (4)$$

$$D_p(\text{background}) = |M(p) - 0| \quad (5)$$

and $V_{p,q}$ again is the same as [1]. We weight the contribution of the shape terms by a factor λ , which was set to 25 for all experiments in this paper. A larger λ results in less deviation of the graph cut solution from the current elliptical shape.

2.3. Implementation

We compute the graph over a set of pixels in a narrowband [12] around the ellipse, as shown in Figure 2. The minimum cut is shown as a dark contour in the figure. After the cut is complete, we update the model by finding the best fitting ellipse [7] to the points of the segmentation result. We then form a new band around the updated ellipse, and then set up a new graph in the new narrowband. We iterate like this until the solution converges. An iterative solution is necessary, since the estimated shape using shape prior influences the graph cut solution and the graph cut solution is used to update the estimated shape using the shape prior.

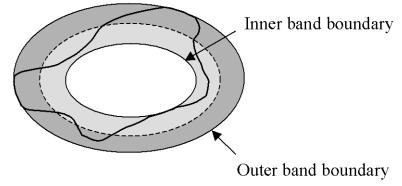


Fig. 2. Narrowband used in the method.

The complete algorithm then consists of the following steps:

1. Identify an initial ellipse, either automatically or through user interaction.
2. Form the shape mask M from the current ellipse.
3. Compute the mean intensities (u_i, u_o) of the pixels inside and outside the mask.

4. Form the narrowband around the ellipse via dilation.
5. Set up the graph using the vertices in the narrowband. Compute the energy using the image intensity, the mean intensities (u_i, u_o) , and the shape mask.
6. Compute the minimum cut of the graph.
7. Fit an ellipse to the points on the minimum cut. Display ellipse.
8. Go back to step 2 until convergence.

3. EXPERIMENTAL RESULTS

We used our algorithm to segment a blood vessel from a pelvic magnetic resonance (MR) image, shown in Figure 3. In part (a), a user initializes the segmentation by clicking on the image to generate a seed point, around which we fit a small circle. We evolve the ellipse using our approach; an intermediate result after 7 iterations is shown in (b), and the final converged result after 20 iterations is shown in (c). The segmentation completes in less than one second on a machine with a Pentium 4 2.66 GHz processor. The shape prior constrains the solution to an elliptical region, and the method is able to accurately segment the blood vessel, even though another blood vessel with a similar intensity is nearby. For comparison, we executed a graph cut algorithm without a shape prior, shown in (d). Without the shape prior the segmentation leaks through nearby dark structures and produces an undesirable result. Figure 4 shows the segmentation of lymph nodes. The rightmost image of the figure is particularly challenging since the lymph node is adjacent to a blood vessel.

As a non-medical example, in Figure 5 we segment a human face. A person was photographed in a controlled room with a dark background. We initialize the segmentation automatically by placing a small circle around the brightest part of the image. For these images, this occurs where there are specular highlights on the subject's forehead. We then run our segmentation algorithm on the luminance of the image. The results are shown for three face poses in Figure 5. While more sophisticated face segmentation and tracking methods exist, for example those that consider skin tone, facial features, temporal coherence, physical motion models, etc., this example demonstrates the ability of our technique to segment challenging data using a simple intensity model and a shape prior. Improved results could be obtained by considering more domain-specific knowledge.

A more in-depth analysis and validation of the method is ongoing, however, current results show that the use of a shape prior in a graph-cuts segmentation algorithm has much promise, particularly in its application to detection of vessels and lymph nodes from pelvic MR images.

4. CONCLUSION

This paper presented a method to incorporate a shape prior into a graph cuts segmentation. We described our approach and experimentally demonstrated its usefulness.

While we consider elliptical shape priors, it should be possible to extend the method to other parametric shapes or shape spaces formed from statistical shape training. This is left for future work. In addition, we are interested in including other region and boundary-based image terms to further assist in the segmentation. Finally, we plan on implementing the approach in 3D using ellipsoids to segment volume data.

5. ACKNOWLEDGEMENTS

We thank Irwin Sobel of HP Laboratories for helping us acquire the data used in Figure 5. We also thank Vladimir Kolmogorov for graph cut software.

6. REFERENCES

- [1] Boykov, Y. and Jolly, M.P., "Interactive Graph Cuts for Optimal Boundary and Region Segmentation of Objects in N-D Images," Proc. Intl. Conf. on Comp. Vision, Vol. I, pp. 105-112, 2001.
- [2] Boykov, Y. and Kolmogorov, V., "An Experimental Comparison of Min-Cut/Max-Flow Algorithms for Energy Minimization in Vision," in IEEE Trans. on Patt. Anal. and Machine Intel. (PAMI), vol. 26, no. 9, pp. 1124-1137, Sept. 2004.
- [3] Chan, T., Vese, L., "Active Contours Without Edges," IEEE Trans. on Image Processing, 10(2), pp. 266-277, Feb. 2001.
- [4] Cootes, T., Beeston, C., Edwards, G., and Taylor, C., "Unified Framework for Atlas Matching Using Active Appearance Models," in the Intl. Conf. on Information Processing in Medical Imaging, pp. 322-333, 1999.
- [5] Grady, L, Funka-Lea, G, "Multi-label Image Segmentation for Medical Applications Based on Graph-Theoretic Electrical Potentials," ECCV Workshops CVAMIA and MMBIA, pp. 230-245, 2004.
- [6] Kass, M. Witkin, A. and Terzopoulos, D., "Snakes: Active Contour Models," Int. J. Computer Vision, Vol. 1, No. 4, pp. 321-331, 1987.
- [7] Fitzgibbon, A., Pilu, M. and Fisher, R., "Direct Least-Squares Fitting of Ellipses," IEEE Trans. on Patt. Anal. and Machine Intel. (PAMI), 21(5), pp. 476-480, May 1999.

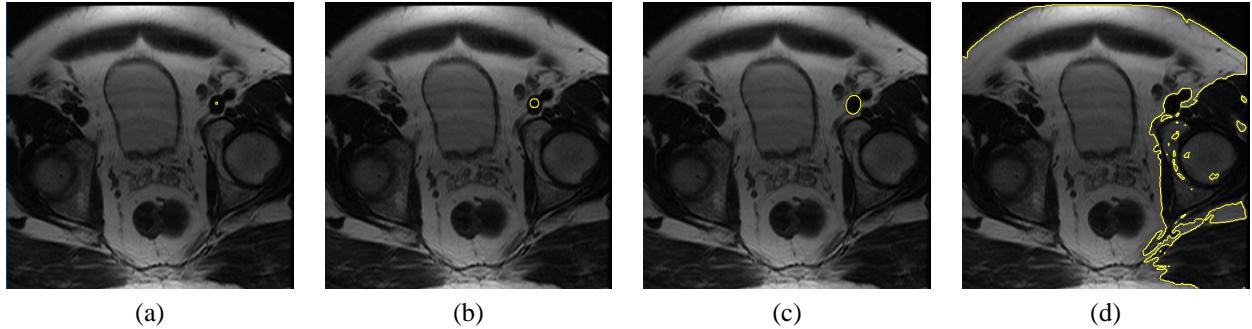


Fig. 3. Segmentation of a blood vessel in a pelvic MR image. In (a), the user selects a click point where we initialize a small circle. The ellipse then evolves (b) until it converges (c) to achieve the segmentation result. For comparison, in (d) we show the result using a graph cut segmentation without using a shape prior.

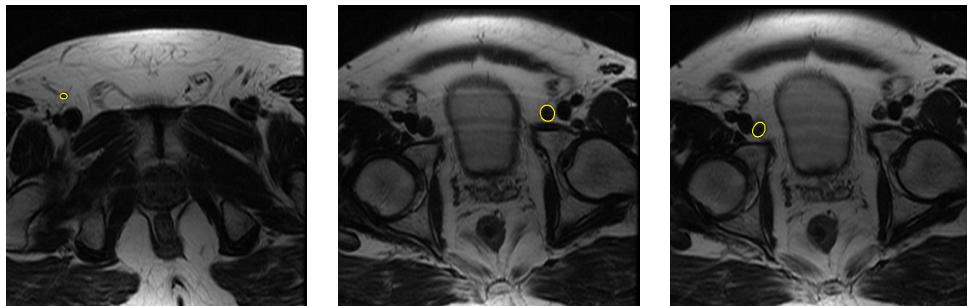


Fig. 4. Segmentation of lymph nodes in pelvic MR images.

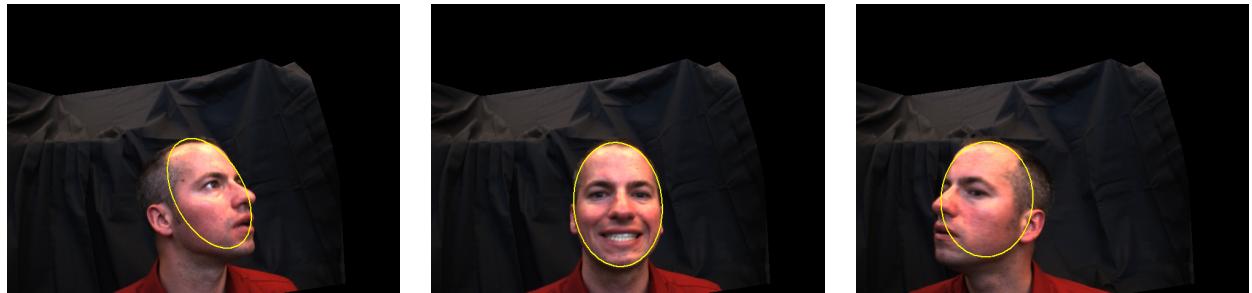


Fig. 5. Automatic segmentation of a human face.

- [8] Kolmogorov, V., Zabih, R., "What Energy Functions Can Be Minimized via Graph Cuts?", *IEEE Trans. on Patt. Anal. and Machine Intel. (PAMI)*, Vol. 26, No. 2, pp. 147-159, 2004.
- [9] Levinton, M., Grimson, E., Faugeras, O., "Statistical Shape Influence in Geodesic Active Contours, in Proc. CVPR, Vol. 1, pp. 316-323, 2000.
- [10] Rousson, M. and Paragios, N., "Shape Priors for Level Set Representations," in Proc. ECCV, Vol. 2, pp. 89-93, 2002.
- [11] Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, E., Willsky, A., "A Shape-Based Approach to Curve Evolution for Segmentation of Medical Imagery," *IEEE Trans. on Med. Img.* Vol. 22, No. 2, 137-154, February 2003.
- [12] Xu, N., Bansal, R., and Ahuja, N., "Object Segmentation Using Graph Cuts Based Active Contours," *Proc. CVPR*, Vol. 2, pp. 46-53, June 2003.