A Unifying View of Contour Length Bias Correction

Christina Pavlopoulou   Stella X. Yu

Computer Science Department
Boston College, Chestnut Hill, MA 02469, USA
{pavlo,syu}@cs.bc.edu

Abstract. Original active contour formulations may become ill-posed especially for boundaries characterized by prominent features. Attempts to yield well-posed formulations lead to bias towards short contours. We provide a framework to unify existing bias correcting energy methods and propose a novel local bias correcting scheme similar to non-maximum suppression. Our method can be seen as an approximation of a well-known algorithm that transforms a graph with positive and negative weights to a graph with only positive weights while preserving the shortest paths among the nodes.

1 Introduction

One of the most well-known energy criteria for modeling and extracting object boundaries is that of Snakes, initially proposed in [1]:

$$E[C(s)] = \int_{C(s)} \frac{1}{2}(\alpha|C'(s)|^2 + \beta|C''(s)|^2) ds - \lambda \int_{C(s)} \|\nabla I\| ds \qquad (1)$$

$C(s)$ denotes the contour parametrized by s. The first two terms favor smooth contours, whereas the third favors contours adhering to prominent image features like strong discontinuities. The above energy has no intrinsic preference towards short boundaries, however it may become ill-posed. Good boundary segments receive negative cost and the minimum of the objective may become $-\infty$. Past approaches that attempted to correct the formulation ([2,3,4]) led to criteria strongly biased towards short segments. An example is shown in Figure 1. Given two points on the object boundary, the criterion in [2] will extract the shortest possible curve instead of the actual boundary.

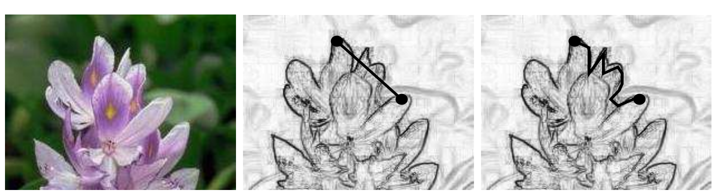

(a) original image (b) biased criterion (c) optimal solution

Fig. 1. Traditional energy criteria suffer from bias towards segments of short length. Given two points on the object boundary, the criterion of [2] will produce a straight line (shown in (b)). The desired boundary is shown in (c).

We offer a novel interpretation of the bias problem and introduce a framework for correcting it. Our framework unifies existing approaches like min ratio cycles [5], piecewise extension of the contour [6], non maximum-suppression [7], and our probabilistic formulation in [8].

The length bias is a result of converting Criterion 1 to positive by adding a large constant. Such a transformation leads to a well-posed functional, however the minima are not preserved. The new optimum solution is strongly biased towards short boundaries.

To remove the bias we turn to a discrete representation of Criterion 1. We represent the image with a graph where each node corresponds to a pixel and neighboring pixels are connected. The edge weights are derived from the biased criterion that is, Equation 1 plus constant, and are positive. The goal is to find the quantity α to remove from the weights so that they remain positive and the bias is eliminated. We show that earlier bias elimination approaches follow this framework and provide different choices regarding α. We additionally propose a local bias correction scheme, which is an approximation of a well-known algorithm of converting a graph with positive and negative weights to a graph with only positive weights while preserving the shortest paths among nodes.

The problematic nature of Functional 1 has been recognized early on and some of the problems consistently appearing in the literature include: The contour fails to latch to prominent image discontinuities and shrinks to a point. The contour produced is often too smooth and can not model geometrically complex boundaries. Self-intersecting contours are allowed and cannot be easily avoided. These problems have been mostly attributed to either the suboptimal nature of the optimization method, or the parametric form of the energy functional. Rarely have they been linked to the bias introduced when the energy criterion is converted to positive.

Earlier approaches required initialization of the contour very close to the actual boundary [1,9]. The intelligent scissors method described in [2,3,4] provided a novel way for the user to guide the delineation process. Usually, these approaches require a lot of user interaction to delineate the boundary. Level set methods [10] employ an intrinsic representation of the curve and thus are not prone to problems related to parametrization. However, it is difficult to impose topological constraints, for example extraction of a single region. Methods like the ones in [11,6,7] have incorporated heuristics in the optimization process; they essentially extract the boundary in a piecewise manner. Additional image features ([12,13,14]) and stronger contour priors ([15,16,17]) have also been explored. Such methods impose additional constraints but do not correct the built-in bias of the original criterion. The most direct attempt to address the bias problem has been to normalize the quality score of the contour by the length of the contour [5,18].

2 The Boundary Length Bias Problem

To better understand the nature of the bias we will employ a discrete version of Functional 1 and we will omit the second-order derivative. We assume that a curve C is discretized into n points. Let c_i be the i-th point. Then, the energy 1 is given by:

$$E[C] = \sum_{i=1}^{n} \{ d(c_{i+1}, c_i) - \lambda \left\| \nabla I \right\|_{c_i} \} \tag{2}$$

where $d(c_{i+1}, c_i)$ is an approximation of the first derivative of the curve and $\|\nabla I\|_{c_i}$ is the gradient intensity at point c_i. $d(c_{i+1}, c_i)$ can be defined as the Euclidean length of the linear segment connecting neighboring points c_{i+1} and c_i.

Criterion 2 can be globally optimized with dynamic programming. To this end, the image is represented with a graph. Each arc (u, v) is weighted according to 2:

$$w(u, v) = d(u, v) - \lambda f(u, v) \tag{3}$$

where $f(u, v)$ refers to the image-derived features term.

The weights of Eq. 3 become negative at image locations with prominent image features. In the case where negatively weighted cycles are formed, the minimum of Eq. 2 is $-\infty$ and the problem becomes ill-posed. A negative cost cycle acts as a black hole in the energy landscape and forces all candidate boundary segments to include that cycle. Such an example is illustrated in Fig. 2. When the weights are positive, the shortest paths from S to all the other nodes include the bold edges. However, when negatively weighted cycles are introduced (Fig. 2(c)), the shortest paths are altered entirely so that they include negative cycles.

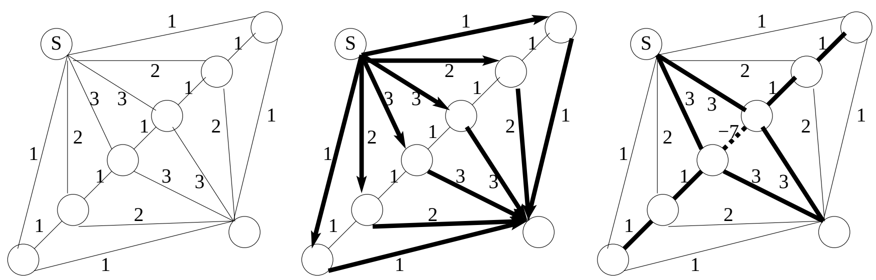

(a) graph with positive weights (b) shortest paths from S (c) shortest paths for neg. cycle

Fig. 2. Negatively weighted cycles act as black holes in the energy landscape. (a) Original graph with positive weights. (b) Shortest paths from S to all the other nodes for the weights of (a) (shown with bold arrowed lines). (c) The edge indicated with dashed line has obtained negative weight -7 and negative cycles have been created. The shortest path from S (shown with bold lines) are forced to include such cycles.

Removing negatively weighted cycles is computationally very difficult. An image will typically consist of many prominent features which will cause the creation of an exponential number of such cycles. Further, because the optima of the solution change drastically when such a cycle is created, it is difficult to impose simple constraints that will ensure the extraction of contours adhering to the object boundaries.

In practice, algorithms like [2,3,4] assume that the weights are positive. This is equivalent to adding a large positive constant M to the original weights such that:

$$M = \{\max_{w(u,v)<0} |w(u, v)|\} + c \tag{4}$$

where $c > 0$. The weights obtained will be

$$w_M(u, v) = w(u, v) + M \tag{5}$$

Such a transformation does not preserve the optima of the objective criterion since the length of the contour is implicitly part of the optimization process. The objective criterion optimized instead, is:

$$E_M(\mathbf{C}) = \sum_{(u,v)\in C} \{d(u,v) - \lambda f(u,v)\} + nM \tag{6}$$

This difference between Equations 1 and 6 is the term nM which is an additional smoothing term proportional to the length of the contour. Its introduction is arbitrary and its effect can be significant when long and geometrically complex contours are to be extracted. When such criteria are used for interactive contour extraction a large amount of human input is required, as has been observed in [19].

3 Removing the Bias

To remove the bias introduced by adding a constant (Eq. 5), we seek $\hat{w}$ of the form:

$$\hat{w}(u,v) = w_M(u,v) - \alpha(u,v) \tag{7}$$

Our goal is to estimate $\alpha(u,v)$ so that $\hat{w}(\cdot) > 0$ and we will do so in a local fashion. Previously proposed bias correction methods provide different choices for $\alpha(u,v)$.

3.1 Local Bias Correction

The role of negative weights is to encourage the inclusion of boundary segments in the final solution. Thus, we need to assign very low positive weights to good boundary segments. The quality of a segment can be assessed based on the quality of its neighbors: a segment should receive low value if it is significantly better than nearby segments. The simplest segment is the edge between two nodes and we define:

$$w^+(u,v) = w_M(u,v) - \max_w w_M(u,w) \tag{8}$$

where w and v are adjacent to u.

Non-maximum suppression and piecewise extension of the boundary are very similar to this transformation. Non-maximum suppression assigns high values to locally best pixels. Piecewise boundary extension, extracts a boundary in an incremental fashion so that it is composed from high-score segments.

Converting Negative Weights to Positive *Provided there are no negatively weighted cycles*, a graph with negative and positive weights can be converted to a graph with positive weights so that the shortest paths among the nodes are preserved. Such a transformation is part of Johnson's all pairs shortest paths algorithm ([20]) and defines a new weighting function $\hat{w}(u,v)$ as:

$$w^+(u,v) = w(u,v) + h(u) - h(v) \tag{9}$$

The function $h(\cdot)$ is computed as follows. We create a new graph G' consisting of all the nodes of the original graph G and an additional dummy node s. Node s is connected to all the other nodes with weights equal to 0. Then, $h(u)$ is defined as the cost of the shortest path from s to u. The weights thus defined are positive. Figure 3 shows an example of such a transformation.

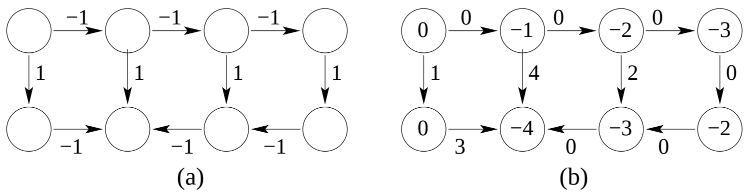

Fig. 3. (a) Original graph with positive and negative weights. (b) Graph with transformed weights. The numbers inside the nodes indicate the shortest path costs from a dummy node s.

Since a criterion with both positive and negative weights (but without negatively weighted cycles) does not suffer from the length bias, it follows that there exist criteria with positive terms which do not have an implicit bias.

In practice however, this algorithm cannot be applied since the weights induced by the image features will lead to negative cycles. Our local correction method can be seen as an approximation to the optimal algorithm. Instead of a single dummy node s, we use as many dummy nodes as the nodes of the graph, and find the shortest paths in a small neighborhood of each node. If there is a single best contour, our method yields the same contour as the optimal method.

3.2 Ratio Weight Cycles

The ratio weight criterion in [5] minimizes a normalized version of the original energy functional given by:

$$w(C) = \frac{\sum_e w(e)}{\sum_e n(e)} \tag{10}$$

where $w(C)$ is the weight of a contour C.

Finding the minimum of Eq. 10 is equivalent to converting the original graph weights $w(e)$ to $w(e) - \lambda n(e)$ and finding zero cost cycles, i.e.:

$$\hat{w}(C) = w(e) - \lambda n(e) = 0 \tag{11}$$

This is equivalent to *finding the largest λ such that no negatively weighted cycles are created*. The approach as presented in [5] does not model open curves and does not admit user interaction. It provides a way of estimating λ given a fixed $n(e)$.

The authors explore two types of $n(e)$. When $n(e) = 1$, the shortest mean cycle is found and the bias of Eq. 6 is reduced to M. The data term is also altered so that the contour extracted has on average good features. When $n(e) = 1/|\nabla I|(e)$, the criterion minimized is very similar to the original snakes criterion. The stronger the gradient intensity, the shorter the contour extracted.

3.3 Probabilistic Formulation

In [8], we proposed a probabilistic formulation which is capable of extracting geometrically complex boundaries. The weights defined by this method are of the form:

$$w(u, v) = x_j - \log \sum_i e^{-x_i} \tag{12}$$

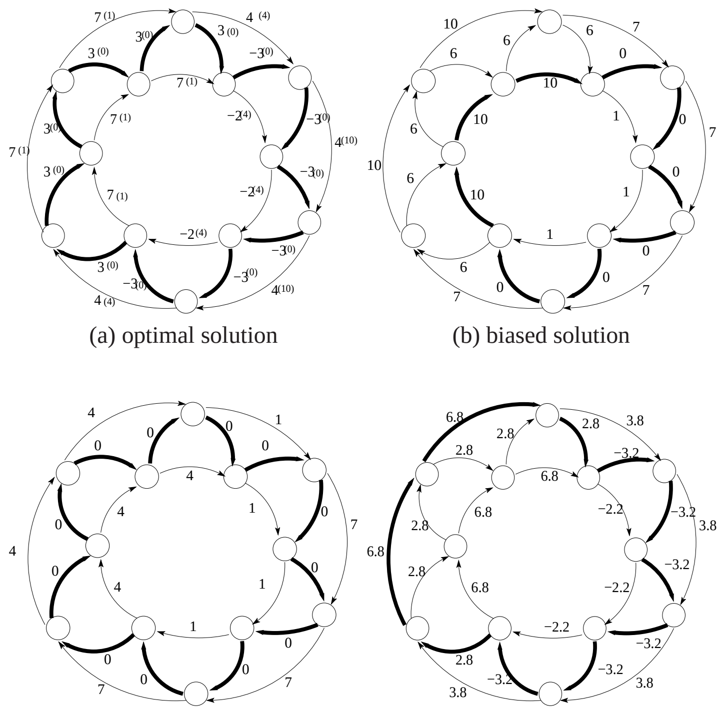

(a) optimal solution (b) biased solution

(c) local bias correction (d) min ratio

Fig. 4. Bold lines indicate the optimal contour. In (a)the weights in parentheses are obtained by optimally converting the negative weights to positive. Graph (b) is obtained by adding 3 to the original weights of graph (a). In (c) the local bias correction is applied to the weights of (b). In (d) the weights of (b) are transformed according to the min ratio cycle method.

where x_i, x_j are the scores of neighboring segments, in our case edges emanating from the same node. The approximation scheme used to compute the $\log$ of the summation of exponentials results in weights similar to the ones of the local correcting method. That is, locally best segments receive very low costs.

4 Evaluation

4.1 Toy Example

Figure 4 demonstrates how the methods presented transform the weights and alter the optima of the criteria. Bold lines indicate the desired boundary. In the optimal case, (Fig. 4(a)) the graph contains both positive and negative weights and the desired boundary is the petal-shaped one. Inside the parentheses, the weights obtained from the optimal transformation are indicated. If the weights are translated by a constant, as is shown in Figure 4(b), then the optima of the criterion are not preserved. As a result, the optimum contour has been smoothed out (part of the inner cycle is included). The local bias transformation is shown in Figure 4(c); for this example the weights calculated are not similar with the optimal ones but the optimal contour is the desired boundary. Finally, in Figure 4(d) the weights obtained from the mean ration are shown. In this case, the optimal contour includes part of the outer cycle.

4.2 Contour Completion

Figure 5 shows results obtained for some real images for the task of contour completion. The gradient of the image was only used to guide the process. Seed points were selected by identifying consecutive strong gradient intensity points. The shortest paths among all seed points were found using biased weights and the locally corrected weights. In the case of biased weights one can see the tendency towards simpler contours. Further, boundaries which are not characterized with high intensity gradient are not always followed, as for example in the black and white flower. On the other hand, the locally corrected weights produce more detailed edge maps and oftentimes complete the contours in a more conceptually compatible fashion. On the downside, they may lead to irregular boundaries as in the case of the woman example.

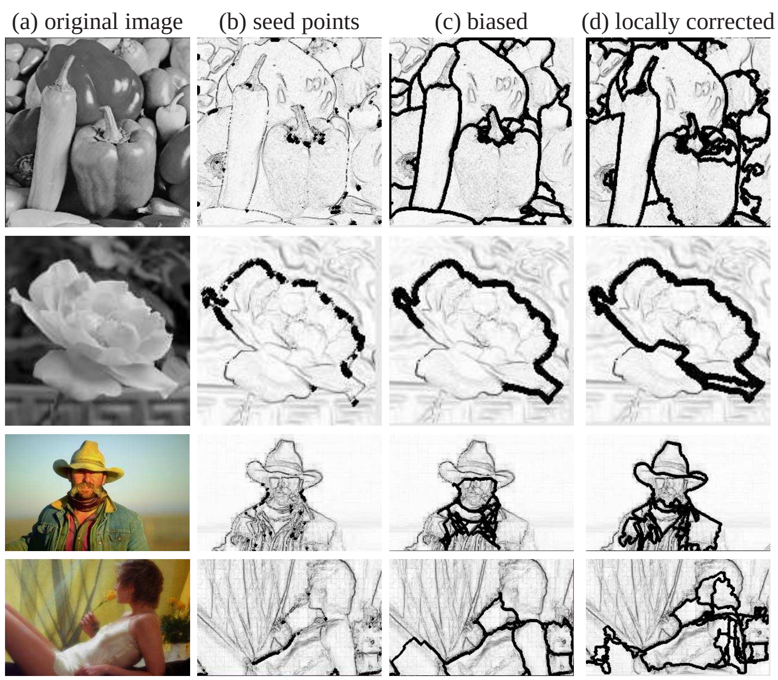

Fig. 5. Seed points were selected as strong edge points on the object boundary (b). Shortest paths between all seed points were found and displayed for the biased method (c) and the locally corrected method (d).

Acknowledgments: This research is funded by NSF CAREER IIS-0644204 and a Clare Boothe Luce Professorship to Yu.

References

1. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active Contour Models. In: Int'l Conference on Computer Vision, IEEE (1987) 1, 2
2. Mortensen, E., Barrett, W.: Intelligent Scissors for Image Composition. In: SIGGRAPH. (1995) 1, 2, 3
3. Mortensen, E., Barrett, W.: Interactive Segmentation with Intelligent Scissors. Graphical Models and Image Processing **60** (1998) 1, 2, 3
4. Falcao, A., Udupa, J., Samarasekera, S., Sharma, S.: User-Steered Image Segmentation Paradigms: Live Wire and Live Lane. Graphical Models and Image Processing **60** (1998) 233–260 1, 2, 3
5. Jermyn, I., Ishikawa, H.: Globally optimal regions and boundaries as minimum ratio weight cycles. IEEE Transactions on Pattern Analysis and Machine Intelligence **23** (2001) 1075–1088 2, 5
6. Mortensen, E., Barrett, W.: A Confidence Measure for Boundary Detection and Object Selection. In: Proc. Computer Vision and Pattern Recognition, IEEE (2001) 2
7. Mortensen, E., Jia, J.: A Bayesian Network Framework for RealTime Object Selection. In: Proc. Workshop on Perceptual Organization on Computer Vision, IEEE (2004) 2
8. Pavlopoulou, C., Yu, S.X.: Boundaries as Contours with Optimal Appearance and Area of Support. In: 7th Int'l Conf. on Energy Minimization Methods in Computer Vision and Pattern Recognition. (2009) 2, 5
9. Geiger, D., Gupta, A., Costa, L., Vlotzos, J.: Dynamic Programming for Detecting, Tracking, and Matching Deformable Contours. IEEE Transactions on Pattern Analysis and Machine Intelligence **17** (1995) 294–302 2
10. Malladi, R., Sethian, J.A., Vemuri, B.C.: Shape Modeling with Front Propagation: A Level Set Approach. IEEE Transactions on Pattern Analysis and Machine Intelligence **17** (1995) 158–175 2
11. Neuenschwander, W., Fua, P., Iverson, L., Szekely, G., Kubler, O.: Ziplock Snakes. Int'l Journal of Computer Vision **25** (1997) 191–201 2
12. Cohen, L.: On Active Contour Models and Balloons. Computer Vision Graphics and Image Processing: Image Understanding **52** (1991) 211–218 2
13. Paragios, N., Deriche, R.: Geodesic Active Contours and Level Set Methods for Supervised Texture Segmentation. International Journal On Computer Vision **46** (2002) 223–247 2
14. O. Gérard and T. Deschamps and M. Greff and L. D. Cohen: Real-time Interactive Path Extraction with on-the-fly Adaptation of the External Forces. In: European Conference in Computer Vision. (2002) 2
15. Sullivan, J., Blake, A., Isard, M., MacCormick, J.: Bayesian Object Localisation in Images. Int. J. Comput. Vision **44** (2001) 111–135 2
16. Allili, M., Ziou, D.: Active contours for video object tracking using region, boundary and shape information. Signal, Image and Video Processing **1** (2007) 101–17 2
17. Joshi, S.H., Srivastava, A.: Intrinsic bayesian active contours for extraction of object boundaries in images. Int. J. Comput. Vision **81** (2009) 331–355 2
18. Schoenemann, T., Cremers, D.: Globally optimal image segmentation with an elastic shape prior. In: IEEE International Conference on Computer Vision (ICCV). (2007) 2
19. Rother, C., Kolmogorov, V., Blake, A.: Grab-cut Interactive Foreground Extraction Using Iterated Graph Cuts. ACM Trans. Graph. (SIGGRAPH) **23** (2004) 309–314 4
20. Cormen, T., Leiserson, C., Rivest, R.: Introduction to Algorithms. McGraw Hill (1990) 4