

Combining Local and Global Cues for Closed Contour Extraction

Vida Movahedi

<http://www.cse.yorku.ca/~vida>

James H. Elder

<http://elderlab.yorku.ca/~elder/>

Centre for Vision Research

York University

Toronto, Canada

Abstract

Algorithms for computing closed contours are generally based upon local Gestalt cues relating pairs of oriented elements, and a Markov assumption to then group these elements into chains. Without additional global constraints, these algorithms generally do not perform well on general natural scenes. Such global cues could include symmetry, shape priors or global colour appearance. A key challenge is to combine these local and global cues in a statistically optimal way. Here we propose a novel, effective method for rigorously combining local and global cues, both at the stage of forming new closed contour hypotheses, and at the stage of evaluating and ranking these hypotheses. We also demonstrate the importance of promoting the diversity of hypotheses. We evaluate our results on a standard public dataset, and demonstrate a substantial performance improvement over prior methods.

1 Introduction

We address the problem of extracting the simple closed contours that bound the salient objects in a natural image [1, 2]. Existing methods can be classified into three main categories: i) *region segmentation methods*, which optimize the labeling of pixels (or superpixels [3]) as figure or ground using methods such as graph cuts [4], ii) *active contour models*, which search for the optimal parametric deformable contour bounding the object, e.g. [5], and 3) *contour grouping methods*, which search for the optimal cycle of local oriented primitives (e.g. edgels, line segments) forming the boundary. The method proposed in this paper falls in this third category.

Contour grouping algorithms are generally based upon classical Gestalt cues such as proximity, good continuation and similarity [2]. Many algorithms base grouping heuristics on these principles [6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16]. There have also been a number of efforts to pose the problem of contour grouping as one of probabilistic inference, and to learn the required statistical parameters from labelled data [17, 18, 19, 20], although not all of these methods compute closed contours.

Gestalt cues such as proximity and good continuation are often called ‘local’ cues, as they are defined on pairs of local oriented features, and a Markov assumption (either explicit or implicit) is made in order to infer the probability or plausibility of longer chains of these features. While this Markov approximation has some statistical justification [9], it is an

incomplete model, and more global cues must also be brought to bear in order to achieve good results [6, 11, 23]. Accurate integration of these local and global cues is a challenge, partly because of deviations from the Markov assumption. *The biggest contribution of our paper* is a novel, effective method for combining local and global cues, both at the stage of forming new closed contour hypotheses, and at the stage of evaluating and ranking these hypotheses.

A second challenge faced by contour grouping algorithms stems from the exponential nature of the search space, which tends to yield many very similar high probability hypotheses, while neglecting other more diverse hypotheses with slightly lower probability. This tendency lowers performance as partial hypotheses that look slightly less promising are weeded out too early. *A second major contribution* of the present paper is a novel method for promoting diversity in the formation and ranking of contour hypotheses, leading to substantial improvements in performance.

The local cues used by our method have been used by others in prior work, but typically have been assumed conditionally independent and have been modelled parametrically as factored univariate distributions. *A third contribution* of our paper is to show that even though these local cues are to a large degree uncorrelated, use of a multivariate mixture model for the joint distribution is able to capture enough statistical dependencies to render substantial improvements in performance.

2 Dataset

For training and evaluation we used a subset of the Salient Object Dataset (SOD) [18], which in turn is based upon the Berkeley Segmentation Dataset (BSD) [17]. While the BSD provides exhaustive segmentations of each image, SOD identifies the boundaries of a small number of salient objects in each of these images. These boundaries were identified by multiple human subjects, and each subject could potentially identify multiple salient objects, thus for each image there are multiple simple, closed ground truth contours corresponding to one or more objects. The SOD dataset consists of 200 training images and 100 test images.

For this paper we randomly selected 30 training images for our training dataset. For our test dataset we selected the 63 images from the SOD test dataset that included at least one clear object entirely contained within the image frame. Each of these contours is represented as a cycle of pixels. We will refer to these as ground truth contours. For reasons of efficiency, we perform all training and evaluation for our algorithm on images that have been down-sampled by a factor of 4.

3 Forming the Association Graph

The first stage of processing in our method extracts from the image a graph representation of the oriented structure in the image. Local oriented edges are first detected using a standard multi-scale method [9]. These edges are then grouped locally into subpixel-localized line segments of variable length [9]. As in prior work, these N line segments are duplicated to separately represent the two possible trajectories through each line segment [16]. Each of the resulting segments forms a vertex in our association graph, and each edge in this graph represents a grouping hypothesis between specified endpoints of two segments.

For any particular segment, typically only a small number of segments in the image form plausible continuations, and so to limit computational requirements we construct a sparse directed graph with fixed maximal out-degree k , limiting the size of the graph to $|E| \in \mathcal{O}(kN)$. In this paper we use $k = 20$.

In order to learn the statistics required to group these segments, we conducted a hand-labelling exercise on the segments extracted for our training images. For each training image, one ground truth boundary was randomly selected from SOD and 3 subjects selected the cycle of segments that they felt best approximated the object boundary. We will refer to these as ground truth cycles.

We used these ground truth cycles to learn models for local association between segments. In particular, at inference we wish to assign to each edge in our graph a weight equal to the log likelihood ratio L_{ij} of the geometric and photometric relationship between two segments i and j , conditioned on whether they are neighbours on a ground truth cycle [3, 6, 7] (we will refer to these as the ON and OFF conditions in the sequel). Statistics for the ON condition are learned from neighbouring segments in our ground truth cycles, while the OFF statistics are learned from randomly selected segment pairs. The relational cues include: i) Proximity, i.e., the tip to tail distance between the two segments, ii) Parallelism, measured as the sum of the two angles formed by the linear interpolant connecting the segments, iii) Cocircularity, measured as the difference of these two angles, and iv) Brightness difference, measured as the difference between the mean luminance at the two segments.

While prior methods [3, 6, 7] have approximated these local cues as independent and modelled the marginals parametrically, we use EM to learn a Gaussian mixture model for the 4-dimensional likelihood for both ON and OFF conditions, which allows us to capture small dependencies between the cues. The sparse graph is then formed by deleting all but the k outgoing edges with highest likelihood L_{ij} for each vertex i . We optimized the number of mixture components by minimizing the average number of edges in the ground truth cycles not represented in the sparse graph, using a cross-validation method. At the optimal values of 8 components for both ON and OFF mixture models, less than 3% of edges in the ground truth cycles are unrepresented, compared to about 5% loss using the naive Bayes assumption and parametric univariate models [3, 6, 7].

4 Extracting Closed Contours

Object boundaries in the image correspond to cycles in our sparse graph, *i.e.* paths in which the first and last nodes are identical, and no other node appears twice. Furthermore, in the image, bounding contours must be simple (non-intersecting). If these contours could be modelled as first-order Markov, then this is a shortest-path problem that can be solved with Dijkstra’s algorithm or dynamic programming [7]. However the Markov assumption cannot be strictly correct, since the constraints of closure and simplicity induce global dependencies. These deviations from the Markov assumption render exact polynomial methods such as Dijkstra’s algorithm ineffective, since they cannot incorporate constraints like self-intersections. For example, the simplicity constraint breaks the required optimal substructure property of the shortest path problem: subpaths of simple shortest paths are not necessarily the shortest simple subpaths.

In pursuing a probabilistic approach, it is important to note [6, 7] that deviations from the strict Markov assumption do not mean that the naive Bayes model (factoring of the likelihood over edges in the cycle) cannot be used as a useful basis for approximate search algorithms

that can i) detect and avoid self-intersections, ii) dynamically incorporate global cues and iii) yield multiple candidate closed contour solutions that can then be evaluated using more accurate probabilistic models. In prior work, this approach has proven successful when domain-specific priors are available [6], and also in the context of a multi-scale coarse-to-fine approach, where priors from coarse scales can be fed down into finer scales to narrow the approximate search [7].

Here we explore whether this approximate search approach can yield competitive results at a single scale and without domain-specific priors. We focus in particular on the problems of 1) how to combine local grouping cues with global cues in forming partial contour hypotheses and 2) how to rank candidate closed contours once they are computed.

Our algorithm is based upon a greedy constructive search [6, 7]:

Algorithm: Closed Contour Extraction

Input: Sparse association graph

Output: A set C of cycles in the graph corresponding to candidate simple closed curves.

Initialize $C = \{\}$; $m = 1$; $S =$ set of all paths of length 1 (all nodes).

While $m < M$

1. Extend all paths in S by one node. Set $m = m + 1$.
2. Discard all paths corresponding to non-simple curves.
3. Add all closed paths to set C , and remove them from S .
4. Calculate the cost of each path in S (see section 4.1).
5. Discard all but the lowest-cost N_m paths.

The maximum curve length M is determined from the ground truth cycles in the training data, and the number of paths N_m to retain at each stage is determined by a specified memory budget. In the following we address the key question of how the cost of each path is determined.

4.1 Integrating Local and Global Cues in Path Costs

Assuming conditional independence between local cues for edges in a path c_k , the log likelihood ratio for the path based on the set of these local cues is proportional to the average log likelihood ratio over its constituent edges:

$$f_l(c_k) = \frac{1}{m} \sum_{(i,j) \in c_k} \log L_{ij}, \text{ where } m \text{ is the length of } c_k \quad (1)$$

It has been noted, however [8], that due to their local nature, relying solely on these local cues can result in "missing the forest for the trees". In prior work this problem has been addressed by incorporating top-down priors [6, 7]. However, in the context of other approaches, more generic global cues such as symmetry of parts [14, 23], convexity [26], compactness and color homogeneity [9] have been used to good effect. However, typically these global cues are translated into local components in a heuristic fashion, or used only to rank closed contours. Here we explore how local and global cues can be integrated in a more rigorous way, and during the critical phase when closed contour hypotheses are being formed. We develop and test this method using a global colour contrast cue.

We define the global colour contrast $f_g(c_k)$ of path c_k as the χ^2 -distance between normalized colour histograms of the pixels within two adjacent bands of width w_c on either side

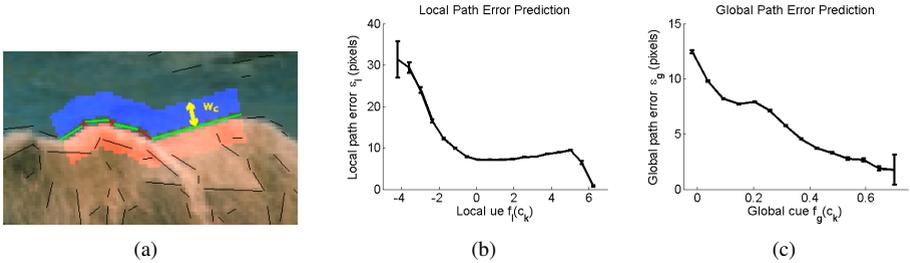


Figure 1: (a) Figure/ground bands used to compute global color contrast cue. (b) and (c) Learned nonparametric predictors for path error based upon local and global cues.

of the path (Figure 1(a)). We explored a number of colour spaces and histogram resolutions, evaluating each by computing a signal-to-noise ratio for $f_g(c_k)$, as a discriminant between paths on ground truth cycles and paths computed using only local cues. We found that using only the L^* and a^* channels of Lab space with a relatively coarse sampling of 8 bins per channel maximized SNR.

This global colour contrast cue is quite distinct from a local colour contrast cue computed at each segment and then converted to a log likelihood ratio that is then summed over the path. The global colour contrast cue selects for paths where there is a coherent and consistent difference in the colours found in the figure from the colours found in the background. In the local cue this coherence is lost: a path with identical global figure/ground colour distributions could be selected as long as locally colour contrast is high. This distinction is not academic: empirically we find that the global cue carries much more information. The key question is how to combine this global cue with the local cues of proximity, good continuation (parallelism and cocircularity) and brightness contrast.

Modeling local grouping decisions using maximum likelihood methods is natural, as there is a clear division between the ON and OFF classes. However, combining these cumulative local cues probabilistically with global cues is a non-trivial problem, partly because the effects of deviations from the naive Bayes assumption accumulate, making the absolute value of the average log likelihood ratio (Eqn. 1) unreliable. Here we observe that at the stage of evaluating paths, the division between ON and OFF is no longer clear: paths can be partially correct. For this reason, it is more appropriate at the path stage to adopt a regression approach: given a measure of distance between two paths, learn a regressor that will predict the distance of a candidate path from ground truth based on both local and global cues.

The tricky part of this regression problem is to acquire appropriate training data: we need plausible candidate paths in order to learn appropriate weights for combining cues, however we need to use these cues in order to determine plausible paths. We solve this chicken and egg problem by computing one set of paths based only on local cues and a second set based only on the global colour contrast cue. We then use both sets of paths to generate plausible candidate hypotheses, and learn regression weights for both local and global cues from these.

In order to make learning efficient we used a relatively inexpensive measure of the distance of candidate paths from ground truth: the average distance of the pixels on the path from the closest ground truth contour. We take a nonparametric approach to the regression problem, binning the local cue $f_l(c_k)$ and global cue $f_g(c_k)$, and then estimating the mean $\hat{\epsilon}_l$ and $\hat{\epsilon}_g$ and variance σ_l^2 and σ_g^2 of predicted error for each. Figure 1 shows the learned nonparametric error predictors for both local and global cues.

Assuming independence between the local and global cues, the least-squares error prediction ε_{ML} (ML for normal distributions) is then given by

$$\hat{\varepsilon}_{ML} = \frac{\hat{\varepsilon}_l/\sigma_l^2 + \hat{\varepsilon}_g/\sigma_g^2}{\sigma_l^{-2} + \sigma_g^{-2}} \quad (2)$$

This predicted error is used as the path cost in Step 4 of our closed contour extraction algorithm (Section 4).

4.2 Promoting Path Diversity

Due to the exponential size of the search space, as paths grow there tends to be many highly similar paths with low predicted error. This can have a negative effect on performance, for two reasons: i) Due to the limited memory budget, distinctive paths that look slightly less promising at an early stage may be pruned out early, before their strengths are recognized. ii) The resulting closed contour candidates may all be highly similar, limiting the ability of the ranking stage (Section 5) to gainfully apply more sophisticated measures of salience on the closed contours.

To address this problem, at each iteration of the closed contour extraction algorithm we soft-cluster the candidate paths and then select representatives from each cluster. To make clustering efficient, we represent each path by an indicator vector \mathbf{v}_k that identifies the segments on the path (but not their order), and then do PCA on these indicator vectors, weighted by the inverse of their predicted errors. By sequentially selecting paths with maximum projection on the distinct leading eigenvectors, we promote paths that are the leading representatives of low-error clusters, while ensuring that all leading clusters are represented.

Algorithm: Path Diversity

Inputs: A large set of candidate paths $S = \{c_k\}$ and their predicted errors $E = \{\hat{\varepsilon}(c_k)\}$

Output: A small set of N_m selected paths $S' = \{c'_k\}$

Initialize $S' = \{\}$

1. Represent each path $c_k \in S$ as an N -vector \mathbf{v}_k , where N is the number of segments in the image. If segment $i \in c_k$, set $\mathbf{v}_k(i) = \frac{1}{\hat{\varepsilon}_k}$, otherwise $\mathbf{v}_k(i) = 0$.
2. Compute the principal components $U = \{\mathbf{u}_k\}$ of $V = \{\mathbf{v}_k\}$.
3. Until $\|S'\| = N_m$:
 - For $k = 1$ to N
 - i. Find fragment $c_k \in S$ with the largest projection on \mathbf{u}_k .
 - ii. Add c_k to S' and remove from S .

5 Ranking Closed Contours

Once closed contours have been computed, there is opportunity for refining our error predictions, since 2D features such as the area of the figure can now be computed. We repeat the approach we took to integrating local and global cues in path costs (Section 4.1): nonparametric regression of the predicted error against complementary cues. Here, however, we use the CM measure [13], normalized by the length of the ground truth contour, as a measure of error, as it corresponds closely with human perception. Since far fewer closed contour

candidates are computed than open path candidates, the computational cost at learning time is manageable.

As predictors, we use the following cues: i) log of the figure area, ii) log of the contour perimeter, ii) root mean square of the turning angles on the contour after smoothing with a Gaussian (we used variance=3), and iv) the Kullback-Leibler divergence of color probability distribution of outside region from color distribution of inside region of the contour. The predicted error is then given by

$$\hat{\epsilon}_{ML} = \frac{1}{\sum_i \sigma_i^{-2}} \sum_i \hat{\epsilon}_i / \sigma_i^2. \quad (3)$$

where $\hat{\epsilon}_i$ is the estimated error for cue i and σ_i^2 is the variance of this estimate. Please see the supplementary material for more details.

6 Experiments

We evaluated our method on the test component of 63 images from the SOD dataset. We compared our local/global (LG) method with the Ratio Contour (RC) algorithm of Stahl & Wang [22], the Adaptive Grouping (AG) method of Estrada & Jepson [9], the Multiscale (MS) method of Estrada & Elder [9], and the Superpixel Closure (SC) method of Levinstein & Dickinson [13], in each case using the parameters recommended by the authors. All methods except for the RC method use colour cues in some way. Test images were down-sampled by a factor of 4 prior to running our LG method to match training; resulting closed polygons were then up-sampled by a factor of 4 to evaluate error. All other methods were run and evaluated at full resolution, as in the original papers.

6.1 Quantitative Results

We used two evaluation measures: i) A region error ϵ_r based on the popular intersection to union measure, and a contour error ϵ_c based on a normalized version of the contour mapping (CM) measure [18], defined as the average distance between corresponding pixels on contours A and B , normalized by the perimeter of the ground truth boundary B . The CM measure determines the monotonic mapping that minimizes this distance, and has been shown to accurately capture human judgements of shape segmentation error [18].

For each image in the SOD dataset there are a number of salient ground truth object boundaries generated by multiple human subjects. The error of an algorithm-generated contour is defined as the *minimum* error over all ground truth contours for that image. When algorithms are allowed to report multiple contours per image, we report the error as the *minimum* error over all algorithm contours for the image, as is the norm (e.g., [18]).

Figure 2 shows the quantitative results. In panels (a) and (b) we report the average region error ϵ_r and contour error ϵ_c , respectively, as a function of the number of contours n the algorithm is allowed to report. We see that our LG method outperforms other methods by the region measure when the number of output contours is 12 or greater, and by the contour measure when the number of output contours is 2 or greater. However, the region error is not a good measure of performance for contour-based methods, as shown in [18], since it is not sensitive to spikes, wiggles, and major shape differences. Please see examples in the supplementary material.

Panel (c) shows how the best contours computed by each method compare, using the contour error measure ϵ_c . While the superiority of the LG method derives in part from the

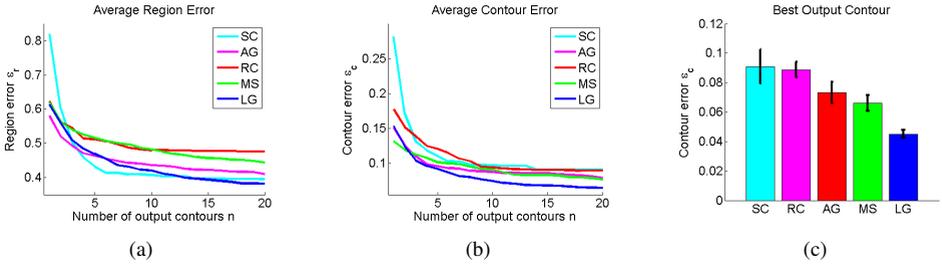


Figure 2: Quantitative evaluation on SOD test dataset. (a) Minimum region error using intersection over union measure as a function of the number of output contours allowed. (b) Minimum contour error using contour mapping measure as a function of the number of output contours allowed. (c) Minimum contour error over all output contours. See Section 6 for a listing of the algorithms tested.

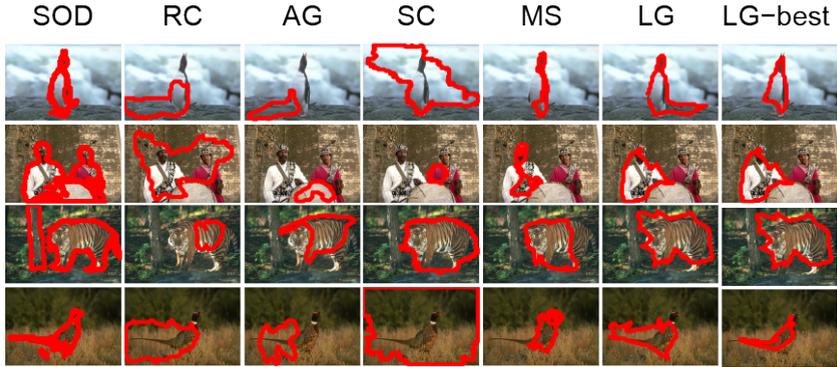


Figure 3: Sample qualitative results. See supplementary material for complete results on SOD test dataset. The first column shows ground truth. Columns 2-6 show the best of 20 contours output by each algorithm. The last column shows the best overall output from our algorithm LG. See Section 6 for a listing of the algorithms tested.

larger number of contours computed (1444 vs. 9 - 150 for the other methods), this result does suggest that refinements to the method for ranking computed closed contours could lead to further gains for small- n evaluation (future work).

6.2 Qualitative Results

Figure 3 shows the best of $n = 20$ contours for each method on a selection of test images from the SOD dataset. Relative to the other methods, our LG algorithm does a good job of consistently selecting the main salient object, although the boundary often incorporates parts of the background. This suggests that further improvements to the LG algorithm can potentially be obtained through better methods for ranking the closed contours extracted by the algorithm.

6.3 Run Time

Our LG method has an average run time of 2.58 minutes per image on a 3.4 GHz Quad Core Intel CPU (without edge/line detection) implemented in unoptimized MATLAB code. The MS method has an average run time of about 20 minutes on the same machine, while the run time of the other 3 methods are less than a minute. While not yet competitive in speed with the RC, AG and SC methods, an advantage of the LG method is that it is highly parallelizable, as at each stage of grouping, evaluation of continuation hypotheses can be done independently for each contour fragment.

7 Conclusions and Future Work

We have presented a novel method for combining local and global cues for improved extraction and ranking of the closed bounding contours of salient objects in natural images. We have also introduced a novel algorithm for promoting the diversity of both partial and complete contour hypotheses, demonstrating substantial performance gains. Our method outperforms other competing state-of-the-art methods by a contour mapping error measure which has been shown to accurately capture human judgements of shape segmentation error. Using a region-based measure, the Superpixel Closure (SC) method achieves better or similar performance to our LG model when the number of output contours is limited. This suggests future work on using richer regional information could lead to better performance of our method. Our LG method performs better when the number of hypotheses is not limited.

Currently, our LG method is implemented in unoptimized MATLAB code, so we believe that significant gains can be made through more careful vectorization and implementation in a compiled language where necessary. We believe an additional factor of 2 improvement is achievable by adaptively detecting when the process of closed contour extraction can be terminated. But perhaps the greatest gain could be achieved by generalizing the contour extraction stage to allow hierarchical grouping of partial contour paths, rather than simply extending each path by a single segment at each iteration. Such a hierarchical computation could in principle lead to a logarithmic reduction in computation time, and would more closely match the hierarchical architecture of the object pathway in primate visual cortex [2, 3].

References

- [1] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. on Pattern Anal. and Machine Intell.*, 26(9):1124–1137, 2004.
- [2] C.E. Connor, S.L. Brincat, and A. Pasupathy. Transformation of shape information in the ventral pathway. *Current Opinion in Neurobiology*, 17(140-147), 2007.
- [3] J. H. Elder and R. M. Goldberg. Ecological statistics of gestalt laws for the perceptual organization of contours. *Journal of Vision*, 2(4):324–353, 2002.
- [4] J. H. Elder and S. W. Zucker. Computing contour closure. *Proc. 4th European Conference on Computer Vision*, 1:399–412, 1996.

- [5] J. H. Elder and S. W. Zucker. Local scale control for edge detection and blur estimation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 20(7): 699–716, 1998. ID: 1.
- [6] J. H. Elder, A. Krupnik, and L. A. Johnston. Contour grouping with prior models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(6):661–674, JUN 2003.
- [7] F. J. Estrada and J. H. Elder. Multi-scale contour extraction based on natural image statistics. In *Computer Vision and Pattern Recognition Workshop, 2006. CVPRW '06*, page 183, 2006.
- [8] F. J. Estrada and A. D. Jepson. Perceptual grouping for contour extraction. In *Proceedings of International Conference on Pattern Recognition*, volume 2, pages 32–35, 2004.
- [9] F. J. Estrada and A. D. Jepson. Robust boundary detection with adaptive grouping. In *Computer Vision and Pattern Recognition Workshop, CVPRW '06*, volume 2006, pages 184–184, 2006.
- [10] P. Felzenszwalb. A min-cover approach for finding salient curves. *Workshop on Perceptual Organization in Computer Vision*, 2006:185, 2006.
- [11] F. Ge, S. Wang, and T. Liu. Image-segmentation evaluation from the perspective of salient object extraction. *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, 1:1146–1153, 2006.
- [12] D. W. Jacobs. Robust and efficient detection of salient convex groups. *IEEE transactions on pattern analysis and machine intelligence*, 18(1):23–37, 1996.
- [13] A. Levinshtein, C. Sminchisescu, and S. Dickinson. Optimal contour closure by super-pixel grouping. *Proceedings of European Conference on Computer Vision*, 2010.
- [14] A. Levinshtein, C. Sminchisescu, and S. Dickinson. Multiscale symmetric part detection and grouping. *International Journal of Computer Vision*, pages 1–18, 2013.
- [15] T. Liu, J. Sun, N. Zheng, X. Tang, and H. Shum. Learning to detect a salient object. In *Proceedings of IEEE Computer Society Conference on Computer and Vision Pattern Recognition (CVPR)*, pages 1–8, 2007.
- [16] S. Mahamud, L.R. Williams, K.K. Thronber, and K. Xu. Segmentation of multiple salient closed contours from real images. *IEEE transactions on pattern analysis and machine intelligence*, 25(4):433, 2003.
- [17] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings of the 8th IEEE International Conference on Computer Vision*, volume 2, pages 416–423, 2001.
- [18] V. Movahedi and J.H. Elder. Design and perceptual validation of performance measures for salient object segmentation. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*, pages 49–56, 2010.

- [19] X. Ren, C. C. Fowlkes, and J. Malik. Scale-invariant contour completion using conditional random fields. In *Proceedings of the 10th IEEE International Conference on Computer Vision*, volume 2, pages 1214–1221, 2005.
- [20] X. Ren, C. C. Fowlkes, and J. Malik. Learning probabilistic models for contour completion in natural images. *International journal of computer vision*, 77(1-3):47–63, 2008.
- [21] M. Riesenhuber and T. Poggio. Hierarchical models of object recognition in cortex. *nature neuroscience*, 2(11):1019–1025, nov 1999.
- [22] J. S. Stahl and S. Wang. Edge grouping combining boundary and region information. *IEEE Transactions on Image Processing*, 16(10):2590–2606, 2007.
- [23] J. S. Stahl and S. Wang. Globally optimal grouping for symmetric closed boundaries by combining boundary and region information. *IEEE transactions on pattern analysis and machine intelligence*, 30(3):395–411, 2008.
- [24] J. S. Stahl, K. Oliver, and S. Wang. Open boundary capable edge grouping with feature maps. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPR Workshops 2008. IEEE Computer Society Conference on*, pages 1–8, 2008.
- [25] S. Wang, T. Kubota, J. M. Siskind, and J. Wang. Salient closed boundary extraction with ratio contour. *IEEE transactions on pattern analysis and machine intelligence*, 27(4):546–561, 2005.
- [26] S. Wang, J. S. Stahl, A. Bailey, and M. Dropps. Global detection of salient convex boundaries. *International Journal of Computer Vision*, 71(3):337–359, 2007.
- [27] M. Wertheimer. Untersuchungen zur lehre von der gestalt. ii. *Psychologische Forschung*, (1):301–350.
- [28] F. Yanbin, L. Caixia, Z. Bingsen, and P. Zhenkuan. An improved algorithm of contour initialization in active contour model. *ICIG '07 Proceedings of the Fourth International Conference on Image and Graphics*, pages 337–341, 2007.