Convexity Grouping of Salient Contours

Padraig Corcoran[1], Peter Mooney[1], and James Tilton[2]

[1]Department of Computer Science,
National University of Ireland Maynooth
padraigc@cs.nuim.ie
http://www.cs.nuim.ie/~padraigc/
[2]NASA Goddard Space Flight Center

Abstract. Convexity represents an important principle of grouping in visual perceptual organization. This paper presents a new technique for contour grouping based on convexity and has the following two properties. Firstly it finds groupings that form contours of high convexity which are not strictly convex. Secondly it finds groupings that form both open and closed contours of high convexity. The authors are unaware of any existing technique which exhibits either of these properties. Contour grouping is posed as the problem of finding minimum cost paths in a graph. The proposed method is evaluated against two highly cited benchmark methods which find strictly convex contours. Both qualitative and quantitative results on natural images demonstrate the proposed method significantly outperforms both benchmark methods.

Keywords: Salient Contour, Graph Searching, Convexity.

1 Introduction

Contours defining object shape are some of the most important features in visual object recognition. However finding such object contours from real scenes is an extremely difficult task. Most contours are regularly fragmented by occlusion, shadows, and low reflectance contrast. In order to infer shape from contours the human visual system must selectively group contours projecting from a common object while keeping contours from different objects separate [3]. This grouping process in the human visual system is known as perceptual organization and the resulting contours are referred to as salient. Computational modelling of perceptual organization presents the following two challenges [13]. Firstly a function which can determine the saliency of a particular contour grouping must be defined. Secondly an effective algorithm for finding such groupings must be used. Many existing contour grouping techniques are based on local Gestalt properties such as closure, good continuation, and proximity [4]. In this work we propose a technique for finding contour groupings of high convexity. Convexity represents an important factor in grouping for many reasons. It is generally accepted that the parts of an objects contour with high convexity correspond to object parts [8]. Borra and Sarkar [2] compared several grouping methods in the context of object recognition and found that grouping subject to a convex constraint

X. Jiang, M. Ferrer, and A. Torsello (Eds.): GbRPR 2011, LNCS 6658, pp. 235–244, 2011.

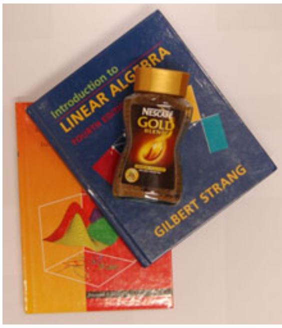

Fig. 1. The book in the background is occluded by two foreground objects. The foreground jar is fully visible but is not strictly convex.

gave best performance. They showed in some cases that convexity dominates the effects of Gestalt properties such as good continuation. Convexity is also a nonaccidential property which can distinguish structure from noise in real images [5]. For these reasons convexity, as a grouping property, has been considered by many researchers [16,5].

Despite this progress existing techniques for finding groupings which form convex contours are constrained in two ways. Firstly all techniques can only find closed contours. Although closure is an important grouping property not all salient contours form closed contours. This is case when one object is occluded by another. For example consider the image containing three objects in Figure 1 where a book in the background is occluded by two foreground objects. Finding a complete contour of this occluded book is not possible without prior knowledge of its complete shape. If such knowledge is not available, employing the principle of maximum entropy, we should aim to find the parts of the object's contour which are visible. In this example this would correspond to the open contour of the occluded book which exists between the background and the book in question. Attempting to find a complete contour of the occluded object using exiting techniques would most likely fail [1]. To overcome the problem of occlusion many authors use local features which may still be visible if the object is not completely occluded [9]. The second constraint exhibited by existing techniques is that they can only finding groupings that form strictly convex contours. For example consider the object of the jar in the foreground of Figure 1. This object is not occluded so finding a complete contour is possible. Application of existing techniques will fail to find this contour because it exhibits high convexity but is not strictly convex.

In this paper we propose a new technique for finding open and closed contours of high convexity. This technique involves three steps. We begin by extracting a set of suitable elements for grouping. A graph representation of these elements is then constructed. Finally an efficient graph mining algorithm is applied to find groupings which form salient contours where saliency is defined as a function of convexity. The layout of this paper is as follows. Section 2 describes the method used to extract elements for grouping and the conversion of these elements to a graph representation. Section 3 details the mining algorithm used to group elements into salient contours. Finally sections 4 and 5 present results and draw conclusions respectively.

2 Grouping Element Extraction and Graph Construction

The elements used for grouping are extracted from an image by the following steps. First edge detection is performed. The resulting edge pixels are then grouped into 8-connected contours. Then these contours are segmented into elements of a suitable scale for grouping. We now describe each of these steps.

Colour edge detection is performed using the photometric invariant technique of van de Weijer et al. [18] and the resulting edge image is thinned. The linking of edge pixels into 8-connected contours is performed using a region growing approach. When an edge junction is encountered the contour in question is terminated and separate contours are created for each branch. In many existing grouping algorithms the grouping elements correspond to 8-connected contours [15,16]. Often the scale of such elements is too great. Consider Figure 2a which shows a single 8-connected contour. This contour contains parts of two individual salient contours corresponding to the bird and tree branch. In order to extract these individual salient contours the scale of the original contour must be reduced before a grouping process is applied. To overcome this issue of scale Jacobs [5] proposed the following strategy, which was later used by Stahl et al. [13]. When grouping two contours Jacobs removed the constraint that the entire second contour must be contained in the grouping. This strategy is not suitable in the context of extracting open contours. In cases where the scale of the original contour is too great and no actual grouping is performed it will fail. Consider again the single contour in Figure 2a. A closed contour of the bird cannot be formed due to occlusion by the tree branch. In order to extract the part of this contour corresponding to the bird the scale of the original contour must be reduced. Since no grouping is required the method of Jacobs cannot reduce the scale. To overcome this problem we propose a novel solution. We represent each 8-connected contour as a series of smaller scale line segments. These line segments are generated through a process of line fitting [6]. Grouping is then performed at the line segment scale. Figure 2b shows the result of this line fitting process applied to the contour in Figure 2a. Grouping at this scale allows the extraction of the open contour corresponding to the bird object

To facilitate the search for salient contours we convert the above line segment representation to a graph representation known as an image graph [17,12]. To construct this graph each line segment endpoint is represented as a node in the graph. All nodes have degree 1 or greater. If two endpoints are connected by a line segment a *real edge* is constructed between them [12]. To allow grouping of segments originating from different 8-connected contours *virtual edges* must be constructed between them [12]. A *virtual edge* is constructed between two segment endpoints if and only if both endpoints in question have degree 1 and the spatial distance between them in the image domain is less than a specified threshold. A path in the image graph corresponds to a contour in the image

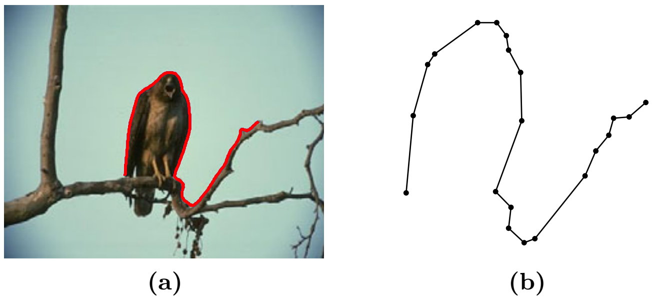

Fig. 2. A single 8-connected edge component is plotted in red in (a). Each connected pair of points in (b) represent a line segment.

domain. This transforms the task of finding salient contour groupings to one of finding suitable paths in the image graph.

3 Salient Contour Grouping

In order to find salient paths in the image graph two elements are required. These are a measure of path saliency and an effective graph searching algorithm. A novel measure of saliency, which is a function of contour convexity, is proposed. The graph searching algorithm used is an implementation of Beam Search [14]. A single application of Beam search starting from a specified initialization or root edge, if successful, will find a single contour. Therefore to find multiple salient contours the search must be repeated multiple times from different root edges. In the following three sections we describe the proposed measure of saliency, the Beam search implementation and the algorithm used to determine a suitable partially ordered set of root edges for searching.

3.1 Contour Saliency Measure

The saliency of a contour i is determined using the function CRL in Equation 1.

$$CRL(i) = C(i) \times R(i) \times L(i) \tag{1}$$

This measure is a product of three terms which we now describe. The function C measures the convexity of the contour in question which lies in the range $[0, 1]$ and approaches 1 for convex contours. This is determined using the convexity measure of Corcoran et al. [11] which has computational complexity of $O(n)$ [7]. The function R in Equation 1 equals the percentage of the contour length which is constructed from real edges and lies in the range $[0, 1]$. This term is important because it models the Gestalt property of good continuation where contours with less breaks appear more salient. It is common for an image to contain many short contours having high convexity which are not perceived by the human visual system as being salient. In order to prevent the extraction of such contours the function L in Equation 1, which is defined by Equation 2, returns a value in the range $[0, 1]$ which approaches 1 for longer contours.

$$L(i) = 1 - \frac{1}{1 + \left(\frac{\Gamma_i}{K}\right)^2} \tag{2}$$

The variable Γ_i equals the total length of the contour in question. The variable K is a specified parameter which determines the point at which this function approaches unity. The shape of this function is such that its first derivative decreases as contour length increases. This models the fact that short contours are generally perceived as being non-salient but when contours are sufficiency long, length becomes a less important factor. All three terms in the product of Equation 1 are in the range $[0, 1]$. The resulting CRL saliency value will also be in this range with higher values signifying greater saliency. Representing a closed contour as an open contour with equal first and last vertices allows CRL to be applied to both types of contour. The time complexity required to compute CRL is $O(n)$. This makes the task of finding image contours of high convexity tractable.

3.2 Beam Search

Beam search uses a breadth-first search to build its search tree. At each level of this tree it generates all successors but only stores a predetermined number of these which are determined to be the best. This number is referred to as the beam size. Due to the fact that a path may be extended in both directions the search tree is bi-directional. A search begins with a path containing two nodes joined by a single edge where this edge is known as the root edge. Consider the image graph in Figure 3a where nodes are represented by circles, real edges by solid lines and virtual edges by dashed lines. A path in this graph, represented by the colour red, contains the four nodes (A, B, C, D, E) of which A and E are leaf nodes in the search tree. The edge connecting nodes B and C is the root edge and is represented by a thicker line. Due to the bidirectional nature of the search the path may be extended along the edge connecting to A to G, and the edge connecting to E to F.

The Beam search terminates when no further extension can be made to any path. The most salient contour encountered during the entire search procedure is then returned. This contour may correspond to only part of a longer path which was discovered. For example the most salient contour encountered during the creation of the path in Figure 3a may correspond to the path containing the nodes (A, B, C, D) which is only a part of an overall longer path discovered. When a path forms a cycle all edges not contained in the cycle are removed. For example consider the path in Figure 3b. If this path was extended by adding the edge connecting H to D this would form a cycle and in turn result in the edge connecting D to E being removed from the path. When a cycle is formed the path in question cannot be extended any further. The proposed Beam search

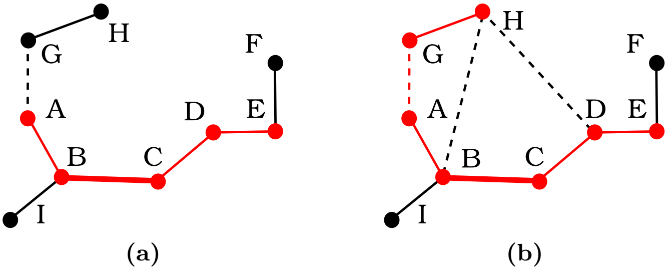

Fig. 3. The paths in (a) and (b) contain the nodes (A, B, C, D, E) and (H, G, A, B, C, D, E) respectively

is subject to two constraints which we now discuss. Any extension to a current path resulting in a contour which self intersects is prohibited. The second constraint imposed is that if an extension results in a cycle, this cycle must contain the corresponding root edge. In the context of Figure 3b extending the path in question by adding the edge connecting H to B would be prohibited. This constraint prevents searches starting at different root edges from consistently finding the same path and was originally proposed by Saund [12].

3.3 Multiple Contour Extraction

The search for salient contours must be performed efficiently. To reduce redundancy Beam searches should commence from root edges, belonging to these contours, as early as possible [12]. Edges corresponding to a single line segment in the image makes it difficult to predict, without considering additional information, if a given edge belongs to a salient contour. To overcome this difficulty we propose the following strategy presented in Algorithm 1. For each line segment we create a corresponding Beam search tree with that line segment as its root edge (line 1). We refer to this set of search trees as ST. Next we increase the depth of each tree to a relatively small size specified by parameter RD (line 2). We then determine the most salient sub-path in each tree and remove all trees from ST for which this value is below a specified threshold RST (line 3). All remaining trees are then ordered by saliency (line 4). The Beam search is then applied to each element of ST in this order until completion (line 6). Upon completion of a given Beam search the following steps are executed: First the most salient contour, MSC, encountered during the search is determined (line 7). If its saliency is greater than a specified threshold CST it is accepted and added to a set entitled SC (line 9). If accepted, the contour MSC is used to reduce the number of elements in ST by deleting any element if its corresponding root edge is contained in MSC (line 10). The most recently processed tree is then removed from ST (line 12). Once all elements in ST have been processed the complete set of contours found is returned (line 14).

Algorithm 1. Multiple Contour Extraction

```
$ST$ = CreateTrees()
increaseDepth($ST$, $RD$)
removeUnsalient($ST$, $RST$)
sort($ST$)
while $ST$.size() $> 0$ do
    completeSearch($ST$[1])
    $MSC = SetTrees[1].MSC$
    if $MSC.Saliency > CST$ then
        $SC$.add($MSC$)
        $ST$.delete($MSC$)
    end if
    $ST$.erase(1)
end while
return $SC$
```

4 Results

To evaluate the proposed technique for finding salient contour groupings the Berkeley segmentation dataset [10] was used. To allow a comparative study the contour grouping techniques of Wang et al. [16] and Jacobs [5] were used. Both these techniques are constrained by the fact that they can only extract strictly convex closed contours. We qualify this statement with the fact that the method of Jacobs [5] is robust to noise and therefore can extract closed contours which are not strictly but are almost convex. Also the method of Wang et al. [16] is robust to occlusion due to the image boundary. These two techniques represent state of the art convexity based grouping methods. When applied to images containing unoccluded object contours which are strictly convex both techniques perform well. Examples of this can be found in the original publications. It was discovered that when these techniques were applied to most natural images they performed poorly. For example consider Figure 4 which displays the results of applying both techniques to a single image. Both techniques fail to extract the most salient contours in this image which correspond to the contours of the bird and the tree branches. This failure can be attributed to the fact that these contours are occluded by the a tree branch and the image boundary respectively, and are not strictly convex. For example the contour of the bird contains two concave parts just below its head.

Figure 5 displays the contours extracted using the proposed technique for a set images. These results were achieved using a Beam width of 1, RD value of 15, RST value of 0.6, CST value of 0.7 and K value of 20. When the Beam width was greater than 1 the algorithm discovered a greater number of contours which were salient with respect to the proposed metric but these in general did not correspond to genuine salient contours or object parts. From this analysis we drew the conclusion that a best first search performs the task of filtering noise. In all cases the proposed method accurately finds the most salient contours. These include the contours of the bird and deer in Figure 5a and Figure 5b respectively.

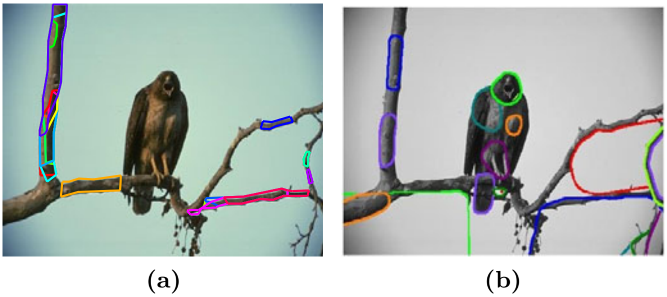

(a) (b)

Fig. 4. Convex contours extracted using the methods of Wang et al. and Jacobs are shown in (a) and (b) respectively

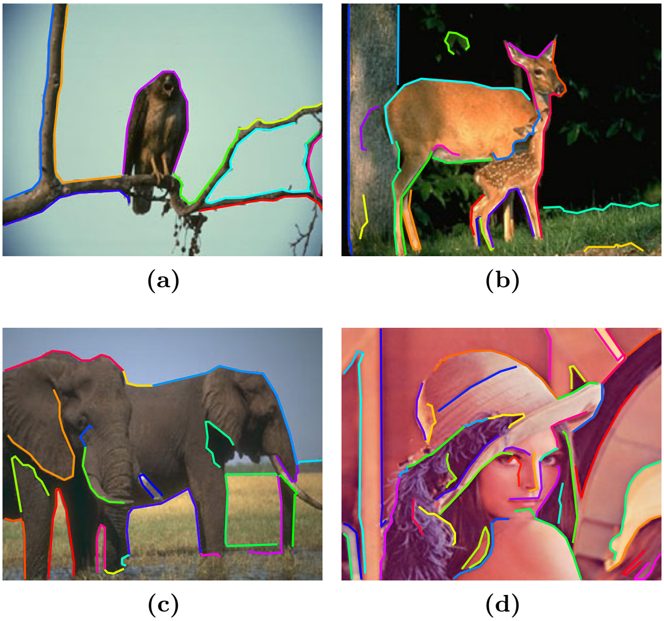

(a) (b)

(c) (d)

Fig. 5. Contours extracted using the proposed technique are shown

It is evident from this qualitative evaluation the proposed method outperforms both benchmark methods. To quantify this performance, we compared precision-recall curves on the Berkeley benchmark set [10] which contains 100 test images. The method of Wang et al. [16] is supervised in the sense that the number of contours extracted must be specified by the user. Therefore it cannot be subject to this analysis. The curves in question are shown in Figure 6 and were generated by varying the threshold applied to the gradient image processed by both algorithms. The dip on the left hand side of the curves is due to the fact that with a high threshold the salient contours contain many breaks in the edge image. The curve of the proposed technique is superior to the curve corresponding to

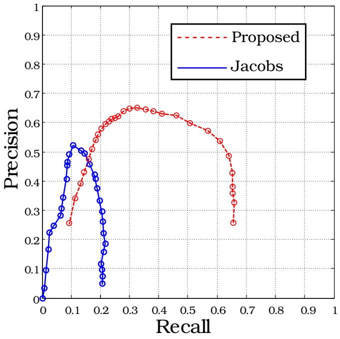

Fig. 6. The PR curves for the proposed and Jacobs method are shown

the method of Jacobs signifying its superior performance. In all test cases the running time of the proposed algorithm was under 1 second.

5 Conclusions

This paper presents a new technique for finding salient image contours of high convexity. The novelty of this technique lies in its ability to find both open and closed contours which exhibit high convexity but may not be strictly convex. Qualitative and quantitative results achieved on natural images show the proposed technique to outperform two current methods for finding convex image contours.

Acknowledgments

Research presented in this paper was part-funded by a Strategic Research Cluster grant (07/SRC/I1168) and a Research Professor Award (07/RPR/I1177) from Science Foundation Ireland under the National Development Plan.

References

1. Ansari, N., Delp, E.: Partial shape recognition: a landmark-based approach. IEEE International Conference on Systems, Man and Cybernetics 2, 831–836 (1989)
2. Borra, S., Sarkar, S.: A framework for performance characterization of intermediate-level grouping modules. IEEE Transactions on Pattern Analysis and Machine Intelligence 19(11), 1306–1312 (1997)
3. Elder, J., Elder, J., Zucker, S., Zucker, S.: A measure of closure. Vision Research 34, 3361–3369 (1994)
4. Geisler, W.S., Perry, J.S., Super, B.J., Gallogly, D.P.: Edge co-occurrence in natural images predicts contour grouping performance. Vision Research 41(6), 711–724 (2001)
5. Jacobs, D.: Robust and efficient detection of salient convex groups. IEEE Transactions on Pattern Analysis and Machine Intelligence 18(1), 23–37 (1996)

6. Jain, R., Kasturi, R., Schunck, B.G.: Machine vision. McGraw-Hill, Inc., New York (1995)
7. Latecki, L.: Shape similarity measure based on correspondence of visual parts. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(10), 1185–1190 (2000)
8. Latecki, L.J., Lakmper, R.: Convexity rule for shape decomposition based on discrete contour evolution. Computer Vision and Image Understanding 73(3), 441–454 (1999)
9. Lowe, D.G.: Object recognition from local scale-invariant features. IEEE International Conference on Computer Vision 2, 1150–1157 (1999)
10. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: Proc. 8th Int'l Conf. Computer Vision, vol. 2, pp. 416–423 (July 2001)
11. Corcoran, P., Mooney, P., Tilton, J.: Convexity Measure for Partial Contours by Shape Similarity. In: British Machine Vision Conference. Dundee, Under Review (2011)
12. Saund, E.: Finding perceptually closed paths in sketches and drawings. IEEE Transactions on Pattern Analysis and Machine Intelligence 25(4), 475–491 (2003)
13. Stahl, J., Wang, S.: Convex grouping combining boundary and region information. In: IEEE International Conference on Computer Vision, vol. 2, pp. 946–953 (2005)
14. Russell, S., Norvig, P.: Artificial Intelligence: a Modern Approach, 3rd edn. Prentice-Hall, Upper Saddle River (2009)
15. Wang, S., Kubota, T., Siskind, J., Wang, J.: Salient closed boundary extraction with ratio contour. IEEE Transactions on Pattern Analysis and Machine Intelligence 27(4), 546–561 (2005)
16. Wang, S., Stahl, J., Bailey, A., Dropps, M.: Global detection of salient convex boundaries. International Journal of Computer Vision 71(3), 337–359 (2007)
17. Wang, S., Wang, J., Kubota, T.: From fragments to salient closed boundaries: an in-depth study. In: IEEE Conference on Computer Vision and Pattern Recognition, vol. 1, pp. I–291– I–298 (2004)
18. van de Weijer, J., Gevers, T., Smeulders, A.: Robust photometric invariant features from the color tensor. IEEE Transactions on Image Processing 15(1), 118–127 (2006)