EXTRACTING REGIONS OF SYMMETRY

Abhinav Gupta, V. Shiv Naga Prasad and Larry S. Davis

Perceptual Interfaces and Reality Laboratory,
Univ. of Maryland - Inst. of Advanced Computer Studies (UMIACS),
Univ. of Maryland - College Park, College Park 20742, MD, USA.

ABSTRACT

This paper presents an approach for extending the normalized-cut (n-cut) segmentation algorithm to find symmetric regions present in natural images. We use an existing algorithm to quickly detect possible symmetries present in an image. The detected symmetries are then individually verified using the modified n-cut algorithm to eliminate spurious detections. The weights of the n-cut algorithm are modified so as to include both symmetric and spatial affinities. A global parameter is defined to model the tradeoff between spatial coherence and symmetry. Experimental results indicate that symmetric quality measure for a region segmented by our algorithm is a good indicator for the significance of the principal axis of symmetry.

1. INTRODUCTION

Symmetry is one of the basic features characterizing natural shapes and objects. It is believed to enhance recognition and reconstruction, and is likely to be employed in pre-attentive vision. According to Gestalt psychologists, it forms an important feature in perceptual grouping and saliency problems [1]. Bilateral symmetries are invariant under reflections about a line, called the axis of symmetry, passing through the centroid of the object [2, 3]. We present an approach for combining cues from bilateral symmetries and spatial coherence into a single segmentation process based on the n-cuts algorithm [4]. This can be used for enforcing spatial coherence constraints on the symmetries detected by methods like [5].

The problem of detecting symmetries is difficult when prior information about the image structure is unavailable; this is due to the high computational complexity involved in comparing different parts of the image. An approach was proposed in [5] for detecting multiple bilateral symmetries present in natural images. In order to make the computation feasible, it considers local features only, and uses linear pairwise voting to obtain confidence values for the various axes of symmetry present in an image. Due to the reliance on local features and pairwise comparison, the method detects spurious symmetries along with the significant ones. These spurious detections are usually caused by disconnected pairs of points exhibiting high symmetry but having very low spatial coherence. It is possible to reduce the number of spurious detections by requiring the symmetric points to form a sizeable and spatially coherent region in the image. Fig. 1(a) shows an image with two candidate axes of symmetry. The blue axis is the principal axis of symmetry. The symmetry axis shown in green corresponds to a spurious detection produced by the clutter of spatially incoherent edges present in the crowd. The net symmetry of the segment corresponding to the spurious axis is low as most of its member pixels do not exhibit symmetry.

(a) Original image | (b) Proposed Algorithm

Fig. 1. a) Original Image with two candidate axis of symmetry. b) Symmetric Object with respect to principal axis of symmetry

In some cases, pixels from the background may also vote for certain axes of symmetry. This is undesirable as the detected symmetries should be produced by coherent objects and not the background. Fig. 2(a) shows an image with three axes of symmetry. The voters for the red axis of symmetry are the pixels inside the face and the background. By enforcing spatial coherence, we are able to obtain a segment containing only the face (Fig. 2(b)). Fig. 2(c) shows a segmentation of the image with respect to the axis shown in green (detected due to background pixels). As can be seen, the obtained segment is asymmetric, thus enabling us to determine that the detected symmetry was spurious.

N-cuts [4] and their close relative - multiscale image

The authors wish to thank David Jacobs for discussions on the work.

(a) Original Image	(b) Region w.r.t. red axis	(c) Region w.r.t. green axis

Fig. 2. a) The original image and the axes of symmetry. b) Spatially coherent region segmented out by the algorithm corresponding to the red axis of symmetry. c) Region corresponding to the green axis of symmetry.

segmentation [6], have become popular techniques for spatial segmentation. We describe how to modify the edge weight definitions in the n-cuts algorithm to include symmetry considerations and apply it to extract symmetric regions in images. This will be followed by a discussion on measuring the inter-pixel affinities and the symmetric quality of the segments produced by the modified algorithm (Section 2.1 and 2.2 respectively). We present the experimental results in Section 3 before concluding in Section 4.

2. SYMMETRY CONSTRAINTS IN N-CUTS

We refer the reader to [4] for a description of the n-cuts algorithm. Given an image and an axis of symmetry, our aim is to extend the definition of the edge weights so that minimization of $Ncut(A, B)$ will cause one of the segments to form a coherent symmetric region in the image. A simple approach would be to include additional graph edges based on the symmetry exhibited by pairs of pixels about the given axis of symmetry. Consider the case when some asymmetric parts of the image have strong spatial coherence with the parts producing the detected symmetry. When segmentation is carried out, these asymmetric portions will also be included in the region of symmetry. However, this might not always be desirable; in some cases we might want to include only symmetric portions of the object in the region of symmetry, overcoming the spatial coherence. To do this, we must lower the weights between the symmetric and asymmetric portions. We define a parameter θ to capture the various possibilities - high values of θ will enable symmetric points to "drag" their neighbors into their segments, whereas low values of θ will cause a "leaking" of symmetric points into asymmetric segments. Thus, low values of θ tend to favor spatial coherence and high values of θ tend to favor symmetry.

Consider two pixels i and j which are spatial neighbors of one another, and i_s and j_s, their respective symmetry counterparts. Depending upon the spatial coherence affinities of i and j, and whether they exhibit symmetry with their counterparts, they should be assigned differently to the segments produced by n-cuts. Let s_i and s_j be their symmetry affinities with their counterparts, and c_{ij} be the spatial coherence affinity between them. Let $w(i, j)$ be the net affinity between i and j - a high value will force them into the same segment, whereas a low value is likely to place them in different segments.

- If s_i and s_j are both low (no symmetry) then coherence affinity, c_{ij}, dominates $w(i, j)$.
- If both s_i and s_j are high then θ influences $w(i, j)$'s dependence on c_{ij}. When θ is high then irrespective of c_{ij}, i and j are bound together by making $w(i, j)$ high. On the other hand, for low θ's, $w(i, j)$ is proportional to c_{ij}.
- If exactly one of s_i and s_j is high then i and j are "cut loose" for medium values of θ, i.e., $w(i, j)$ is low; for extreme values of θ, $w(i, j)$ is proportional to c_{ij}.

These motivations can be captured by defining w_{ij} as

$$\begin{aligned} w(i,j) &= \sqrt{(1-s_i)(1-s_j)}c_{ij} + \sqrt{s_i s_j}[f_2(\theta) - c_{ij} f_1(\theta)] \\ &\quad + \frac{s_i + s_j - 2\sqrt{s_i s_j}}{2} c_{ij} f_3(\theta) \end{aligned} \tag{1}$$

where

$$f_1(\theta) = \begin{cases} 0 & : \quad \theta \geq 0.25 \\ -1 & : \quad \text{otherwise} \end{cases}$$

$$f_2(\theta) = \begin{cases} 1 & : \quad \theta \geq 0.25 \\ 0 & : \quad \text{otherwise} \end{cases} \text{, and}$$

$$f_3(\theta) = \begin{cases} 0 & : \quad 0.25 < \theta \leq 0.75 \\ 1 & : \quad \text{otherwise} \end{cases}$$

Now consider $w(i, i_s)$, the weight between i and its symmetry counterpart i_s. This will be low if s_i itself is low to begin with. For high values of s_i and low θ's, asymmetric spatial neighbors of i will lower $w(i, i_s)$, independent of s_i. This will enable i's spatial neighbors to "drag" it into asymmetric segments for low values of θ. We can encode this by defining $w(i, i_s)$ as

$$w(i, i_s) = s_i \frac{1}{\|N(i)\|} \sum_{j \in N(i)} \left[1 + \sqrt{(1-s_j)c_{ij}} f_1(\theta)\right] \tag{2}$$

We guarantee that the weights are symmetric by defining the graph's weights as $\tilde{w}(p, q) = w(p, q) + w(q, p)$.

2.1. Measuring the Affinities

We use intervening gradients to quantify the spatial coherence affinities (c_{ij}) [7]. The Gradient Vector Flow (GVF) field [5] and responses from Haar filters (complimentary to GVF) are used for quantifying the symmetry affinities (s_i). The final symmetric affinity between two pixels is the maximum of the affinities derived from GVF and Haar filters.

2.2. Measuring the Quality of the Segmentation

We would like to measure the quality of the segments in terms of spatial cohesiveness and symmetry so as to eliminate spurious symmetry detections. Let $\mathbf{y}$ define the segmentation obtained upon applying n-cuts on the modified graph $\mathbf{G}$ produced by proposed approach. For quantifying spatial cohesiveness we generate a new graph $\mathbf{G}'$ which only has spatial affinities as is typically used in n-cut approaches [7]. If $\mathbf{y}$ is used to define a cut on $\mathbf{G}'$ then it will have high cut values where the spatial cohesiveness constraints are violated. If $\mathbf{D}'$ and $\mathbf{W}'$ denote the degree and weight matrices for $\mathbf{G}'$, then $\frac{\mathbf{y}^T(\mathbf{D}'-\mathbf{W}')\mathbf{y}}{\mathbf{y}^T\mathbf{D}'\mathbf{y}}$ would provide an inverse measure of the spatial cohesiveness of the segments [4]. Measuring the degree of symmetry of the segment is relatively easy as we already know the axis of symmetry and also know that all points in the segment should be symmetric. We measure the symmetry of a segment by the average symmetry of its member pixels measured using the GVF and texture features.

3. EXPERIMENTAL RESULTS

We first show the effect of parameter θ on the segmentation results. Fig. 3(a) shows the face of a man and the principal axis of symmetry. We see from Fig. 3(b) that for small values of θ, only the forehead is classified as a symmetric region because spatial coherence is given preference over symmetry (spatial coherence breaks at the eyes). However, as the value of θ is increased, symmetry is given preference over spatial coherence - the symmetric structure of the eyes causes them to fall into the symmetric segment (Fig. 3(c)). Thus, θ enables us to shift emphasis between spatial coherence and symmetry.

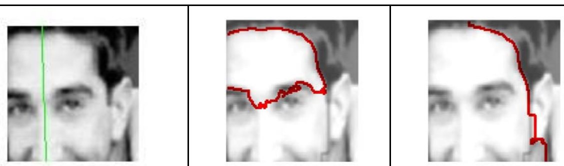

(a) Original image	(b) $\theta < 0.25$	(c) $\theta > 0.75$

Fig. 3. Segmentation variation with θ. Low values of θ tend to favor spatial coherency while higher values of θ favor symmetry.

Next, we show how the symmetric quality of a segment is used to remove spurious symmetry detections by [5]. Fig. 4(a) shows an image of a house and the three candidate axes of symmetry. The principal axis of symmetry of the house is colored blue; the other axes of symmetry - shown in red and green - are caused by parts of the house. In case of n-cuts [4], the spatial incoherence between the door and the rest of the house causes the door to lie in a different segment from the rest of the house. Now consider the case when the proposed algorithm is used to segment the image w.r.t. to the blue axis. High values of θ reduce the weights between the tree's pixels and those of the house as the tree's pixels do not exhibit symmetry. The segmentation results are shown in Fig. 4(c). The symmetric qualities of segmentations w.r.t. the three candidate axes of symmetry (blue, red and green) were 0.35, 0.22 and 0.27 respectively - the segmentation w.r.t. to the principal axis is assigned a higher value.

We investigated the effectiveness of the symmetric quality measure using 27 natural images. The 54 segments (each image having a foreground and background) were manually classified as belonging to one of the following four classes.

- Symmetric Foreground: The foreground segments belonging to this class are symmetric objects. The symmetric quality of such segments should be high.
- False Line of Symmetry: The foreground segments in this class were formed because of the spurious axis of symmetry which is detected by [5].
- Symmetric Backgrounds: This class consisted of the background segments which were themselves symmetric.
- Asymmetric Backgrounds: This class represented background segments which were asymmetric.

In order to enforce spatial coherence, θ was kept low at 0.2 while segmenting. Fig. 5 shows some representative thumbnails of the segmented images and a scatter plot of the symmetry qualities assigned to the segments. (The y-axis represents the index number of the segment and the x-axis represents the symmetric quality of the segment.) On average, the segments corresponding to symmetric objects are assigned significantly higher values - showing the effectiveness of the measure.

Fig. 6 illustrates the application of symmetry constraints on an image with occlusions. Segmentation by the proposed approach puts the face into one segment as edges of symmetry join the face's pixels across the bars. The graph can further be coarsened based on quality of segmentation, merging the blue and magenta(Symmetric Qualities = .31 and .25) into one segment.

4. SUMMARY

We proposed an approach for extracting symmetric regions in natural images. It was based on a modified version of n-cuts algorithm to include spatial and symmetry affinities. Experimental results show that when the initial hypothesis of symmetry is correct then the approach is able to extract the region corresponding to the symmetric object. On the other hand, when the symmetry hypotheses lack spatial coherence, the obtained segmentations are asymmetric. By

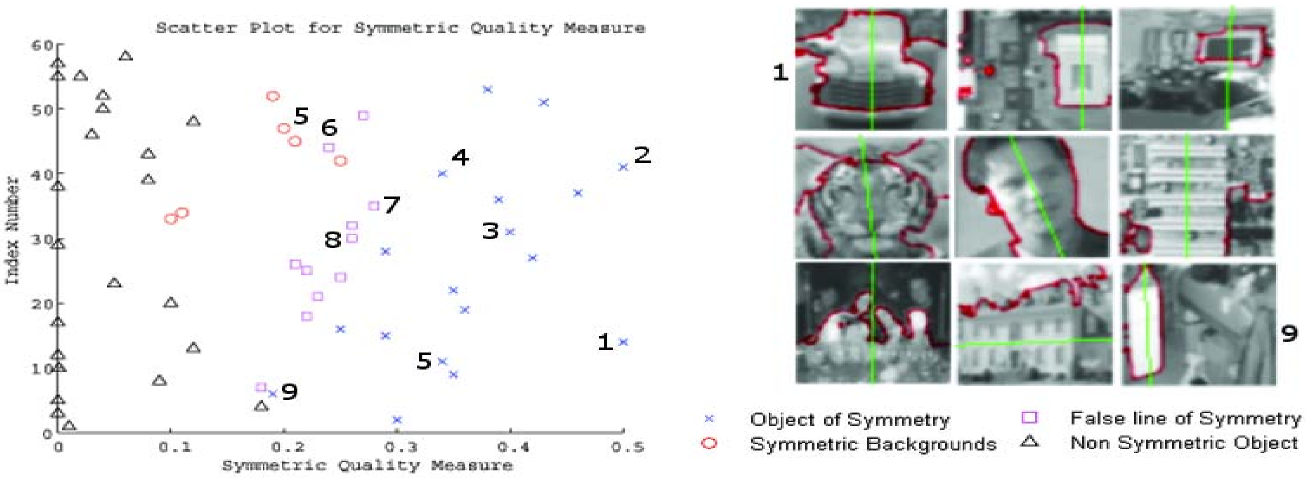

Fig. 5. Scatter plot of the symmetric quality for the 54 segments derived from 27 natural images using the proposed algorithm. The thumbnails show the images containing the segments which have been numbered in the plot.

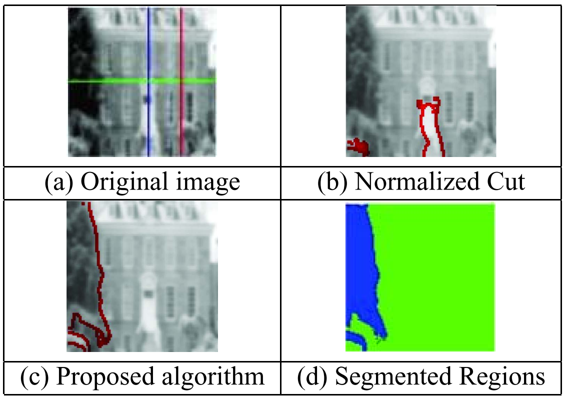

Fig. 4. For high values of θ, the proposed algorithm tends to put the door in the same segment as the rest of the house - the tree is put in a different segment.

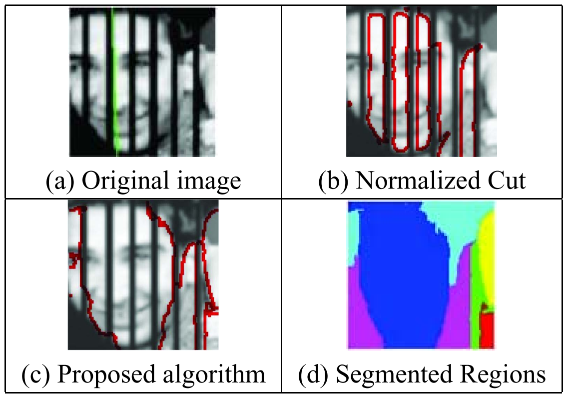

Fig. 6. The Normalized cut algorithm tends to put different parts of face in different segments. However, the additional symmetry weights tend to put whole face in one segment. The graph created can be coarsened using the symmetric values.

measuring the symmetry of the segments, we are able to distinguish between significant and spurious detections.

5. REFERENCES

[1] M. A. Arbib, Ed., *The Handbook of Brain Theory and Neural Networks*, The MIT Press, Cambridge, Massachusetts, London, 1995.

[2] H. Zabrodsky, S. Peleg, and D. Avnir, "Symmetry as a continuous feature," *IEEE Trans. Pattern Anal. and Machine Intell.*, vol. 17, no. 12, pp. 1154–1166, Dec. 1995.

[3] H. J. A. M. Heijmans and A. V. Tuzikov, "Similarity and symmetry measures for convex shapes using Minkowski addition," *IEEE Trans. Pattern Anal. and Machine Intell.*, vol. 20, no. 9, pp. 980–993, Sep. 1998.

[4] J. Shi and J. Malik, "Normalized cut and image segmentation," *IEEE Trans. Pattern Anal. and Machine Intell.*, vol. 22, no. 8, pp. 888–905, 2000.

[5] V. S. N Prasad and B. Yegnanarayana, "Finding axis of symmetry from potential fields," *IEEE Trans. Image Proc.*, vol. 13, no. 12, pp. 1559–1566, 2004.

[6] M. Galun, E. Sharon, R. Basri, and A. Brandt, "Texture segmentation by multiscale aggregation of filter responses and shape elements," in *Proc. IEEE Int'l Conf. Computer Vision (ICCV 2003)*, 2003, pp. 716–723.

[7] J. Malik, S. Belongie, T Leung, and J. Shi, "Contour and texture analysis for image segmentation," *Int'l J. Computer Vision*, vol. 43(1), pp. 7–27, 2001.