Curvature Covariation as a Factor in Perceptual Salience

Thomas M. Murphy[1], Mark Matlin[1,2], Leif H. Finkel[1]

[1]Department of Bioengineering, University of Pennsylvania, 301 Hayden Hall, 3320 Smith Walk, Philadelphia, PA 19104, USA

[2]Department of Physics, Sarah Lawrence College, 1 Mead Way, Bronxville, NY 10708, USA

Abstract-The salience of a contour depends upon several factors, including continuity, closure and curvature consistency. We analyze the statistics of curvature variation using a single image from Shimon Ullman's [1] original work on contour salience. We develop a measure based on the arc length of a contour segment over which curvature variation remains within a constrained range. Locally, all contours in the image are similar with respect to curvature consistency. However, when the entire contour is considered, the most salient contours are found to have the most consistent curvatures. This finding reinforces Ullman's point that salience is a global property of the object.

We interpret these results in view of Rosenholtz's [2] recent model of salience as a statistical measure of outliers from a population. In addition, we speculate on the visual cortical mechanisms in striate and extrastriate cortex required to carry out salience measurements on this class of images.

Keywords - Salience, curvature, cocircularity, visual cortex

I. INTRODUCTION

The perceptual salience of an object is the degree to which it pops-out from the background and captures attention. The salience of a target determines the difficulty of locating it in search tasks, and the speed of recognizing it in rapid presentations. The Gestalt psychologists identified several properties that confer salience upon objects, such as continuity, colinearity or cocircularity, and closure. However, the relative degree to which each of these properties contributes to overall salience remains unclear, as does the manner in which these factors are integrated.

As Ullman [1,3] pointed out, salience is a global property that integrates the Gestalt factors across an entire object. In Fig. 1, taken from Ullman's original paper, the three circular contours pop out and are more salient than the background squiggles. We are interested, in this paper, in understanding quantitatively what factors render the circles salient.

One Gestalt factor that distinguishes the circular contours from the background is closure. Closure is itself a global property (Yen & Finkel [4]), and Kovács & Julesz [5] have shown, using roughly circular contours, that closure leads to a marked increase in salience.

In Ullman's original algorithm, salience was determined by the length, continuity, and curvature of the contour. Long, smooth contours with little change of curvature and no gaps are calculated to be most salient. In a detailed study of this algorithm, Alter and Basri [6] found that, for many images, it robustly predicted salience values in accord with human perception. To our knowledge, a detailed analysis of the original Ullman image (Fig. 1) has not been performed.
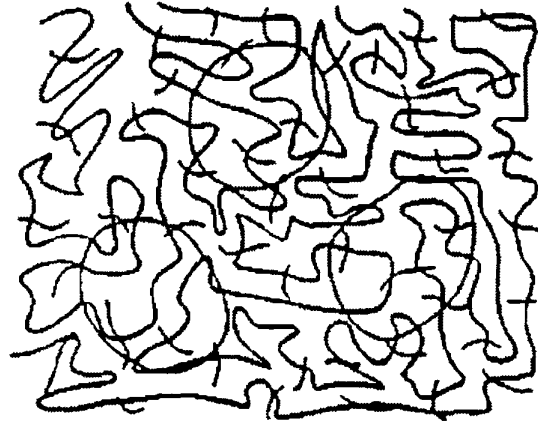

Fig. 1. Shashua & Ullman image.

We sought to determine, in this image, the degree to which the circular contours are more or less cocircular than the background squiggles, over a range of spatial scales. In other words, to what degree does curvature covary across the circular contours as compared to across the squiggles. This study therefore represents an attempt to gather image statistics on a single image.

II. METHODOLOGY

By deliberate construction, the three closed contours in Fig. 1 are physically similar to the background contours in terms of contrast, thickness, and range of orientations. The background squiggles are not all continuous, but terminate at the borders of the figure, forming 3 open contours. The closed contours are shorter (239, 244, 243 pixels) than the background contours (1238, 2011, 880 pixels), but it is not clear that the terminations of the background contours are perceptually significant.

An erosion algorithm was first employed to reduce image contours down to single-pixel widths (using the MATLAB [7] bwmorph function). We were then able to decompose the image into closed and background contours and investigate each separately.

Orientations were computed using Freeman & Adelson's [8] G2/H2 steerable filters. At locations where contours cross, we ascertained that the correct orientation was assigned to each contour.
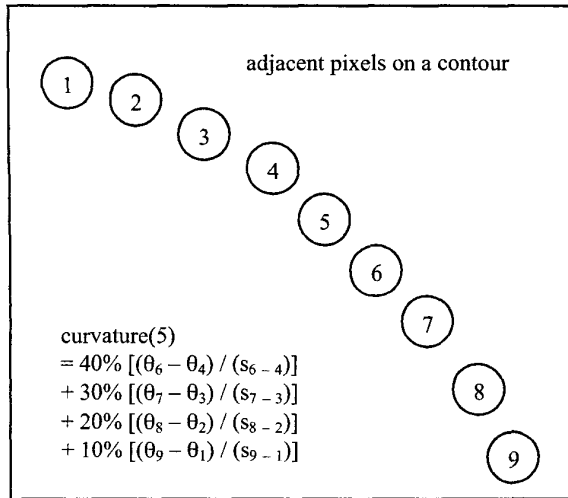
adjacent pixels on a contour

$\text{curvature}(5)$
$= 40\% \left[(\theta_6 - \theta_4) / (s_{6-4}) \right]$
$+ 30\% \left[(\theta_7 - \theta_3) / (s_{7-3}) \right]$
$+ 20\% \left[(\theta_8 - \theta_2) / (s_{8-2}) \right]$
$+ 10\% \left[(\theta_9 - \theta_1) / (s_{9-1}) \right]$

Fig. 2. Method of calculation of curvature at each pixel.

Fig. 3. Threshold at 45% of curvature range.

Curvature is a scale-dependent quantity. Therefore, as illustrated in Fig. 2, we used a weighted average of the curvatures computed over several arc lengths, and assigned the averaged value as the curvature at each pixel. This weighted average was chosen to provide some minimal smoothing while retaining the true nature of the contours. The angle subtraction formulas are corrected modulo 360 so that, for example, the difference between a 350 orientation and a 10 orientation is 20 , not 340 .

Once the curvature had been calculated at each pixel, cocircular segments were defined. A range (max curvature – min curvature) of curvature values for all contours was computed. Adjacent pixels on the same contour were deemed cocircular if the difference in their curvatures was below a set threshold, expressed as a percentage of this total curvature range.

All simulations were carried out using the MATLAB application development environment (version 6 R12) and the associated Image Processing Toolbox (version 3.1).

III. RESULTS

We first investigated a number of technical image processing issues. We carried out the measurements with and without an initial contour-thickening step in an attempt to smooth out the thinned contours. Results were similar in both cases.

Our algorithm is parameterized to allow subsampling of the contours (for example, start at pixel x of the contour and consider only every y pixels). However, the results that we present represent a consideration of all pixels. We also considered calculating curvatures using the slope of the line connecting the pixels for θ (instead of the steerable filter result) and the Euclidean distance between pixels for s (instead of the arc length). Both of these approaches were
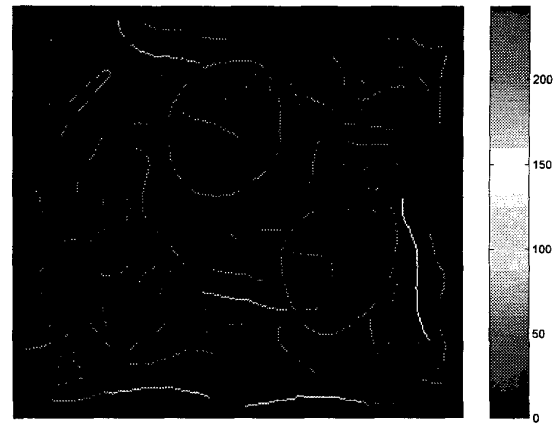
abandoned because of poor results. Other qualitatively poorer results were obtained using Gabor filters instead of steerable filters, performing calculations with overlapping segments of the contours, and mid-pixel averaging if greater than threshold. All of these techniques were subsequently dropped.

Decomposition of the image yielded three closed contours, three long background contours, and a number of short, open cross-hatches. The mean and standard deviation of the curvatures for the closed and background contours were computed. As expected, the mean curvature for each closed contour (0.0202, 0.0204, 0.0185) was close to the inverse of its approximate radius (0.0263, 0.0258, 0.0259, respectively), computed by dividing the contour length by 2π. The sign of the curvature is generally determined by the direction of contour traversal (clockwise or counter-clockwise). The standard deviation of the curvature over the closed contours (0.0284, 0.0310, 0.0343) was significantly less than that of the background contours (0.0840, 0.0869, 0.0877). Thus, overall, the circular contours are more cocircular than the background contours. However, it remained to be determined whether portions of the circular contours were more cocircular than portions of the background. This question can be assessed by determining, for each pixel, how far one can move along the contour until the curvature deviates past a fixed threshold. For example, in Fig. 3, the color of the pixel (and accompanying color bar) indicates the number of consecutive adjacent pixels on each contour that remain within a threshold of 45% of the total curvature range of all of the circular contours. It is apparent that at this threshold the closed contours contain much longer stretches of cocircularity, compared to the background contours and cross-hatches.

Fig. 4 provides a similar computation when the threshold has been set to 35% of the maximal curvature range. This stricter threshold cuts down on the extent of contours conforming to the curvature constraints. Although the closed
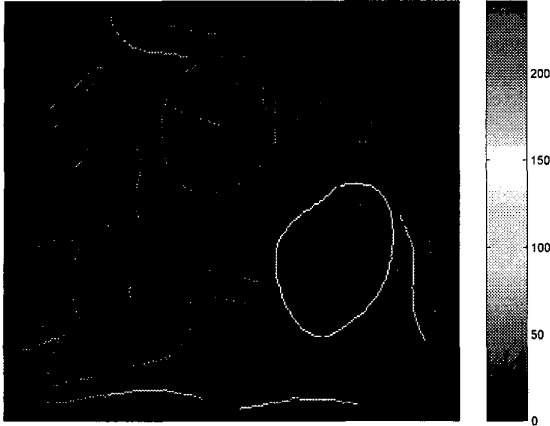
Fig. 4. Threshold at 35% of curvature range.   Fig. 5. Threshold at 15% of curvature range.

contours still contain longer segments meeting this criterion, the difference between the closed contours and background squiggles is less dramatic.

Fig. 5 shows the results with a threshold set at 15% of the maximal curvature range. As the definition of cocircularity becomes stricter, i.e., curvature is constrained to a narrower range of values, the differences between the circular and background contours diminishes. These results therefore show that at a local scale, all contours in the image are similar. What distinguishes the circular contours (apart from closure) is that they maintain a similar curvature over a much longer extent, compared to the squiggles, where "similar" can be quantitatively defined. In other results, not shown here, we have found that further loosening of the curvature constraints (ie., allowing variation over 50% or more of the maximum range), yields similar results to the 45% case, up to a limit. At this point (within a threshold of approximately 75% of the maximal curvature range), the definition of cocircularity is so loose that its discriminatory power begins to diminish. Large portions of all contours (circular and background) begin to appear cocircular.

IV. DISCUSSION

If curvature covariation is a factor in determining perceptual salience then Fig. 3 argues that the circular contours should be more salient than the background. However, as Figs. 4 and 5 demonstrate, the degree to which curvature covariation contributes to salience will depend upon the mechanisms and the scale over which curvature information is computed in visual cortex.

A number of studies have proposed neural and computational mechanisms for computing curvature [9]. These lower level curvature calculations are consistent with mechanisms, such as end-stopping, available in primary visual cortex. At higher cortical levels, Van Essen, Connor and colleagues have shown that cells in areas V2 and V4
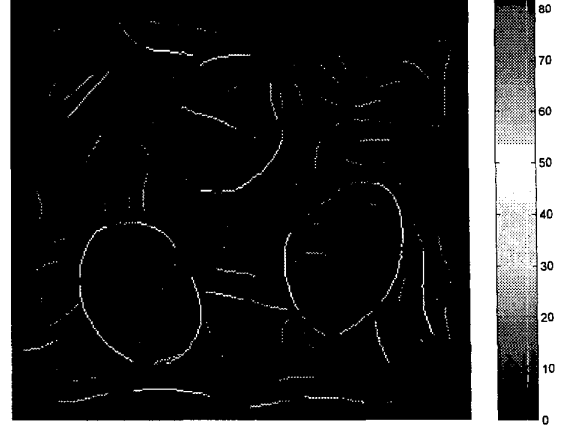
are selective both for the magnitude and direction of curvature [10, 11].

Our results show that curvature values on the circles are much more similar over longer extents. A cortical mechanism for distinguishing the circles could thus be based on the similarity of firing rates of curvature sensitive cells in V4. For example, Hopfield & Brody [12] have proposed a mechanism in which groups of cells with similar firing rates synchronize. Synchronization occurs naturally in the types of cortical architectures studied by Connors and colleagues [13]. Hopfield and Brody make the point that in a large ensemble of cells, a large fraction of cells firing at the same rate is statistically unlikely. Thus, a set of connected V4 cells, each sensitive to magnitude and direction of curvature, that are coupled by horizontal connections, would rapidly synchronize in response to the circles, but to a much lesser degree to the background squiggles.

Salience depends upon the degree to which a target differs from the background. Rosenholtz [2] has proposed that a possible metric for salience is to consider the Mahalanobis distance between a target and the background. Thus, salient targets are those whose features are statistical outliers from the background population.

For the Ullman figure, the circular contours are statistical outliers by virtue of their consistent curvature relative to the fluctuating curvatures of the background. The contours are also outliers with respect to closure, as discussed above. One might ask, however, why the background squiggles are not salient since they statistically differ compared to the circles. One possible answer relates to Rosenholtz's [2] explanation of salience asymmetries. As illustrated in Fig. 6, the standard deviation of curvatures on a circle is much less than on a squiggle. Thus, curvatures on the circle are more consistent, and lie several (circle) standard deviations away from the mean curvatures of the squiggles. But the mean curvature of the squiggle lies close (in terms of the standard deviation of the squiggle distribution) to the mean of the circle curvatures.
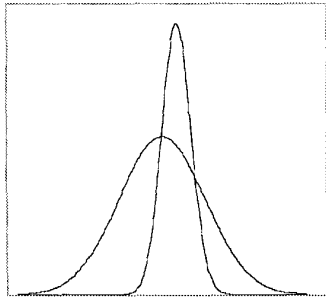
18

Fig. 6. Gaussian model of circle vs. background curvatures.

The relative consistency of the circle curvatures distinguish it from the background.

Area V4 is well situated to carry out computations of the type envisioned here. Connor and colleagues have found cells in macaque V4 that respond selectively to the magnitude and direction of curvature. V4 contains an extensive network of horizontal connections, which span significant portions of the visual field [14]. Through these connections, possibly together with top-down information from higher temporal areas, V4 is thought to mediate contextual population-based interactions within the scene, such as those required in color constancy. Finally, V4 is known to play a critical role in determining salience. Lesions of V4 render animals incapable of detecting a less salient target in the presence of more salient distractors [15].

V. CONCLUSION

For the single image considered here, the consistency of curvature on a contour is correlated with increased perceptual salience. Salience appears to depend upon the statistical distribution of feature properties, such as curvature, over the object, in comparison to the background. These findings support the earlier work of both Ullman and Rosenholtz. The need to compare curvature measurements over a significant image region, to segment contours into target vs. background, and possibly to discriminate the direction of figure, suggest the need for both horizontal as well as top-down information, possibly at the level of V4, in determining contour salience.

ACKNOWLEDGEMENT

REFERENCES

[1] A. Shashua and S. Ullman, "Structural saliency," Proceedings of the International Conference on Computer Vision, pp. 482-488, 1988.

[2] R. Rosenholtz, "A simple saliency model predicts a number of motion popout phenomena," Vision Research, vol. 39, pp. 3157-3163, 1999.

[3] S. Ullman, High-level Vision – Object recognition and Visual Cognition, The MIT Press, Cambridge, MA, 1996.

[4] S.-C. Yen and L. H. Finkel, "Extraction of perceptually salient contours by striate cortical networks," Vision Research, vol. 38, no. 5, pp. 719-741, 1998.

[5] I. Kovács and B. Julesz, "A closed curve is much more than an incomplete one: effect of closure in figure-ground segmentation," Proceedings of the National Academy of Sciences, USA, vol. 90, pp. 7495-7497, 1993.

[6] T. D. Alter and R. Basri, "Extracting salient contours from images: an analysis of the saliency network," Proceedings of The IEEE Conference on Computer Vision and Pattern Recognition, June 1996.

[7] The MathWorks, Inc., 3 Apple Hill Drive, Natick, MA 01760, USA.

[8] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 13, no. 9, pp. 891-906, September 1991.

[9] A. Dobbins, S. W. Zucker, and M. S. Cynader, "Endstopping and curvature," Vision Research, vol. 29, no. 10, pp. 1371-1387, 1989.

[10] J. Hegdé and D. C. Van Essen, "Selectivity for complex shapes in primate visual area V2," The Journal of Neuroscience, vol. 20 RC61, pp. 1-6, 2000.

[11] A. Pasupathy and C. E. Connor, "Shape representation in area V4: position-specific tuning for boundary conformation," The Journal of Neurophysiology, vol. 86, pp. 2505-2519, November 2001.

[12] J. J. Hopfield and C. D. Brody, "What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration," Proceedings of the National Academy of Sciences, USA, vol. 98, no. 3, pp. 1282-1287, 30 January 2001.

[13] M. Beierlein, J. R. Gibson, and B. W. Connors, "A network of electrically coupled interneurons drives synchronized inhibition in neocortex," Nature Neuroscience, vol. 3, no. 9, pp. 904-910, September 2000.

[14] Y. Amir, M. Harel, and R.J. MalachJ, "Cortical hierarchy reflected in the organization of intrinsic connections in macaque monkey visual cortex," J. Comp Neurol., vol. 334, pp.19-46, 1993.

[15] P.H. Schiller, "The effects of V4 and middle temporal (MT) area lesions on visual performance in the rhesus monkey," Vis Neurosci; vol. 10, pp. 717-46, 1993