

# Segmentation of Multiple Salient Closed Contours from Real Images

Shyjan Mahamud, Lance R. Williams, Karvel K. Thornber, and Kanglin Xu

**Abstract**—Using a saliency measure based on the global property of contour closure, we have developed a segmentation method which identifies smooth closed contours bounding objects of unknown shape in real images. The saliency measure incorporates the Gestalt principles of proximity and good continuity that previous methods have also exploited. Unlike previous methods, we incorporate contour closure by finding the eigenvector with the largest positive real eigenvalue of a transition matrix for a Markov process where edges from the image serve as states. Element  $(i, j)$  of the transition matrix is the conditional probability that a contour which contains edge  $j$  will also contain edge  $i$ . In this paper, we show how the saliency measure, defined for individual edges, can be used to derive a saliency relation, defined for pairs of edges, and further show that strongly-connected components of the graph representing the saliency relation correspond to smooth closed contours in the image. Finally, we report for the first time, results on large real images for which segmentation takes an average of about 10 seconds per object on a general-purpose workstation.

**Index Terms**—Perceptual organization, contours, Markov chains, eigenvectors.

## 1 INTRODUCTION

VISUAL perception evolved in a world of objects, many of which are bounded by smooth closed contours. We hypothesize that these contours obey a stochastic distribution that is utilized by perceptual processes in finding contours bounding objects. In prior work, [23], [25], [26], this distribution has been modeled and used to derive a saliency measure which exploits the closure of contours bounding objects. It was found that this measure provides a significant improvement over previous approaches in highlighting edges lying on contours bounding objects in small synthetic scenes created from contours of real objects and natural background texture [26]. However, no method was presented for actually segmenting out the salient closed contours. Despite the effectiveness of the saliency measure, we will show that a simple threshold on the saliencies is not sufficient for segmentation, especially in cases where two or more contours contain edges of similar saliency. In this paper, we generalize the saliency measure described in [26] and use it as the basis for a system which segments out multiple smooth closed contours in real images. In previous work, a routine from a standard numerical library was used to solve the eigenproblem required to compute the saliency measure. However, due to the number of edges involved, this is infeasible for large real images. To address this problem, we have developed an efficient technique that exploits the sparseness and symmetry of representations intrinsic to the problem,

using, which we report for the first time, results on large real images.<sup>1</sup> Given an edge image as in Fig. 1, we would like to extract out separately the individual contours bounding the two pears. We would like to achieve such a segmentation with no a priori knowledge of the specific objects that generate these contours. Such a task is one of the goals of *perceptual grouping*. In lieu of any specific knowledge about the objects generating the contours, we impose a subset of the Gestalt principles for perceptual organization. Most previous approaches to perceptual grouping of edges have incorporated the Gestalt principles of proximity and good continuation in some form (e.g., [2], [9], [6], [12]). These approaches assume that adjacent edges of an object boundary are close together and can be smoothly interpolated. In addition to these two local properties, we exploit the global property that contours bounding objects must be closed. Unlike proximity and good continuation, closure cannot be reduced to a local property defined for pairs of edges in isolation.

Previous approaches [1], [6], [11], [17] have used graph-based search techniques to find closed contours. A graph of affinities between edges is constructed where the affinities model proximity and good continuation. The affinity between two edges is a purely local relation that is positively correlated with the likelihood that a smooth contour joins the pair of edges. Closure is imposed by searching the graph for cycles that minimize a cost function (possibly subject to other constraints like convexity [11] or winding number [6]). Our approach differs from these others because we first use the local affinity relation (which, in our case, is a transition matrix) to compute a global saliency relation (which gives the probability that a closed contour joins the pair of edges). This relation, defined for pairs of edges, is based on a saliency measure, defined for individual edges, which was first proposed and compared extensively with other saliency measures (including [9], [18], [21]), which do not incorporate

- S. Mahamud is with the Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213.
- L.R. Williams and K. Xu are with the Department of Computer Science, University of New Mexico, Albuquerque, NM 87131.  
E-mail: williams@cs.unm.edu.
- K.K. Thornber is with NEC Research Institute, Inc., 4 Independence Way, Princeton, NJ 08540.

Manuscript received 10 Jan. 2002; revised 14 Aug. 2002; accepted 23 Oct. 2002.  
Recommended for acceptance by D. Jacobs and M. Lindenbaum.  
For information on obtaining reprints of this article, please send e-mail to: tpani@computer.org, and reference IEEECS Log Number 117716.

1. An initial account of the work described in this paper was presented in [13].




Fig. 1. An example of an edge image. The image was synthetically created by superimposing two copies of edges from the boundary of a real pear on a background texture.

closure in [26]. Only after the computation of this global saliency relation do we employ graph search to isolate individual closed contours. Furthermore, we show that using a saliency relation based on contour closure leads naturally to a specific type of graph search, namely, the determination of *strongly connected components*. The close relationship between the strongly connected component computation and the closure property of the global saliency relation distinguishes our work from previous approaches, where generic graph search techniques have been applied to graphs representing a local affinity relation defined for pairs of edges. To illustrate the crucial role played by the global property of contour closure, we show that a method based on a purely local affinity relation produces poor segmentations.

Computing the saliency measure (and relation) requires identifying the eigenvector with the largest positive real eigenvalue of a sparse, positive matrix possessing a particular symmetry. Ordinary techniques for the computation of eigenvectors and eigenvalues are infeasible for large real images. We have developed efficient techniques that exploit the sparseness and symmetry of the matrix to significantly reduce the time required to compute this eigenvector. In this paper, we report the first results on real images with a large number of edges. Our technique reduces the time taken to compute the segmentation for each object contour from an average of around 2 1/2 hours to around 10 seconds.

## 2 PROBLEM FORMULATION

Since the Gestalt principles of proximity and good continuation can be reduced to local properties of the positions and orientations of two edges, we can model them using only local information. Following [15], [23], [25], proximity and good continuation can be modeled by a distribution of smooth curves traced by particles moving with constant speed in directions undergoing Brownian motion. In our work, the transition probability between edge  $i$  and edge  $j$  is denoted by  $P_{ji}$ . It is the sum of the probabilities of all paths that a particle can take between the two edges (see [23] for details). Two parameters control the motion of the particle and embody the Gestalt principles of proximity and good continuation. Each particle has a half-life,  $\tau$ , which determines the distance over which pairs of edges are likely to be linked by random walks. Hence,  $\tau$  models proximity. The variance,  $T$ , of the Gaussian random variable representing change in direction models the principle of good continuation. A third parameter,  $\gamma$ , represents the speed of the particle and, hence, determines the effective scale at which the scene is analyzed, since the transition probabilities vary with speed. At larger speeds, the

distance between a pair of edges is effectively smaller, while at slower speeds, the same distance is effectively larger. In our initial experiments, we chose a fixed speed that was judged to give good results for most images. In a later section, we present results where the optimal  $\gamma$  for each object in the scene is identified using an optimization method. This leads to a scale-invariant segmentation.

The smooth continuation of a curve between two edges requires that the tangent at any point along the curve be continuous. If we wish to extend the curves to include additional edges, then tangent continuity must be enforced at the edges themselves. A particle visiting an edge, and traveling in a given direction, must continue along in that same direction to preserve tangent continuity. This requirement can be ensured by replacing each *oriented* edge, where the orientation is an angle in the range,  $[0 - \pi)$ , with two oppositely directed edges, where the directions are angles in the range,  $[0 - 2\pi)$ . A particle must enter and exit a directed edge in the same direction. If we do not impose tangent continuity at the edges, it is possible to get contours with cusps (i.e., reversals in direction) at the edges, which are not judged to be salient in practice. For more details, see [26]. Since every directed edge  $i$  has a sibling edge at the same position, but pointing in the opposite direction, it will be convenient to denote the sibling edge by  $\bar{i}$ .

Imposing tangent continuity through the use of directed edges has an important implication for the structure of the transition matrix  $\mathbf{P}$ . From symmetry, the probability that any particle travels along a curve starting from edge  $i$  and ending in edge  $j$  is the same as the probability of a particle traveling from edge  $\bar{j}$  to edge  $\bar{i}$  in the reverse direction. Hence,  $P_{ji} = P_{\bar{j}\bar{i}}$ . We call this special symmetry of the transition matrix *reversal symmetry*, which is distinct from the usual symmetry  $P_{ji} = P_{ij}$ , which need not hold in general. Reversal symmetry has important implications for both the form of the expressions that define the saliencies and for the problem of efficiently computing them.

In the rest of the paper, we will have occasion to associate a vector  $s$  with the set of *directed* edges (e.g., the vector of saliencies for each directed edge), one component for each directed edge. Analogous with the case for edges, a component of such a vector  $s_i$  associated with edge  $i$  will have a sibling component denoted by  $\bar{s}_i = s_{\bar{i}}$  associated with edge  $\bar{i}$ .

## 3 EDGE AND LINK SALIENCIES

In this section, we first motivate the expression for the saliency measure introduced in [26]. We then show that the saliency measure can be computed by solving an eigenproblem associated with the transition matrix  $\mathbf{P}$ . Given an edge image, we define a closed contour as a finite closed sequence of edges. By a closed sequence, we mean that, if we start from any edge in the sequence and trace out the contour, we will return to the same edge. Each closed contour  $\alpha$  has a likelihood (or probability) associated with it, which we denote by  $p(\alpha)$ . This probability is the product of the transition probabilities along the path defined by the edges of the contour.

### 3.1 Edge Saliency

We would like to define the saliency of an edge so that it is directly related to the likelihood that a closed contour contains that edge. We begin by considering the set of infinite closed

contours containing the edge. In order to calculate the relative saliencies of infinite closed contours, we start by considering the relative saliencies of closed contours of finite length and take the limit as the length goes to infinity. Restricting ourselves to finite contours for now, the saliency of an edge should be proportional to the expected number of closed contours which contain that edge. The expected number of closed contours of length  $n$  which contain edge  $i$  is simply the sum of the probabilities of all such closed contours:

$$E_i^n = \sum_{\alpha} p(\alpha \mid i \in \alpha, |\alpha| = n). \quad (1)$$

Since we are interested in the relative saliencies of the various infinite contours which contain different edges, we take the limit  $n \rightarrow \infty$  for the expected number of closed contours which contain a given edge  $i$  relative to the expected number which contain any edge and obtain the formal definition for the saliency of edge  $i$ :

$$c_i = \lim_{n \rightarrow \infty} \frac{E_i^n}{\sum_j E_j^n}. \quad (2)$$

This definition suggests that there is a simple relationship between edge saliencies and the eigenvector corresponding to the largest positive real eigenvalue of the transition matrix  $\mathbf{P}$ .

**Theorem 1 (First Saliency Theorem).** *The saliency for edge  $i$  is given by:*

$$c_i = s_i \bar{s}_i, \quad (3)$$

where the  $s_i$ s are the components of the eigenvector (normalized so that  $\sum_i s_i \bar{s}_i = 1$ ) corresponding to the largest positive real eigenvalue,  $\lambda$ , of the transition matrix  $\mathbf{P}$ , i.e.,  $\mathbf{P}s = \lambda s$ .

**Proof.** See Appendix A and, also, [26] for an earlier proof.  $\square$

It is important to note that since  $\mathbf{P}$  is *positive* (all entries are positive), Perron's theorem [10] guarantees that the largest eigenvalue of  $\mathbf{P}$  will be real and positive. The components of the corresponding eigenvector  $s_i$  will all be positive (i.e.,  $s_i > 0$ ). Note that, due to reversal-symmetry, we would expect  $c_i = c_i$  as can be verified from the expression above.

### 3.2 Link Saliency

In this section, we use the saliency measure, defined in the last section for individual edges, to derive a *saliency relation*, and defined for pairs of edges. Because it is associated with a "link" between two edges, we term an element of the saliency relation a *link saliency*. The link saliency,  $C_{ij}$ , equals the probability that a closed contour passes through edge  $j$  and then (without visiting another edge) passes through edge  $i$ . In a way which is analogous to the definition for edge saliency, we have:

$$C_{ij} = \lim_{n \rightarrow \infty} \frac{E_{ij}^n}{\sum_l E_l^n}, \quad (4)$$

where  $E_{ij}^n$  is the expected number of closed contours of length  $n$ , which pass through edges  $j$  and  $i$  in succession and  $E_l^n$  is as defined before in (1). Like the edge saliencies, the link saliencies also have a simple relationship with the eigenvector corresponding to the largest positive real eigenvalue of  $\mathbf{P}$ .

**Theorem 2 (Second Saliency Theorem).** *The link-saliencies between any two edges  $j$  and  $i$  are given by:*

$$C_{ij} = \frac{\bar{s}_i P_{ij} s_j}{\lambda}, \quad (5)$$

where the  $s_i$ s are the components of the eigenvector (normalized such that  $\sum_i s_i \bar{s}_i = 1$ ) corresponding to the largest positive real eigenvalue,  $\lambda$ , of the transition matrix,  $\mathbf{P}$ .

**Proof.** See Appendix A.  $\square$

As in the case of the edge saliencies, due to reversal symmetry, we would expect  $C_{ij} = C_{ji}$  as can be verified from the expression above (recall that  $P_{ij} = P_{ji}$  and  $\bar{s}_i = s_i$ ).

Since we are concerned with closed contours, an important conservation property holds for all edges. Any *closed* contour that goes from some edge  $k$  to a second edge  $i$  must continue on to some third edge  $j$ . This is not necessarily true in the case of open contours. We confirm this conservation property and at the same time use it as a consistency check on the expressions for the  $C_{ij}$ s and  $c_i$ s:

$$\sum_k C_{ik} = \sum_k \frac{\bar{s}_i (P_{ik} s_k)}{\lambda} \quad (6)$$

$$= \frac{\bar{s}_i (\lambda s_i)}{\lambda} \quad (7)$$

$$= \bar{s}_i s_i \quad (8)$$

$$= c_i. \quad (9)$$

Doing a similar calculation for  $\sum_j C_{ji}$ , we find

$$\sum_k C_{ik} = c_i = \sum_j C_{ji}. \quad (10)$$

### 3.3 Contour Saliency

Ideally, our segmentation algorithm should extract closed contours in order of increasing saliency. A possible definition for the saliency of a closed contour would be to define it as the probability of a particle tracing a path through the same edges, i.e., the product of the conditional probabilities along the contour's path. However, this definition is dependent on the length of the contour. A closed contour,  $\alpha$ , and another closed contour formed by traversing the edges in  $\alpha$  twice, i.e.,  $\alpha \cdot \alpha$ , should be judged to have the same saliency. However, it is clear that the probability of the second contour will be much less than that of the first. In fact, it will be the square of the first. A more natural definition for the saliency of a closed contour  $\alpha$ , a definition which is invariant to repetition, is the geometric mean of the conditional probabilities along the contour's path:

$$\lambda(\alpha) = p(\alpha)^{1/|\alpha|}, \quad (11)$$

where  $|\alpha|$  is the length of the closed contour and  $p(\alpha)$  is the product of the conditional probabilities which comprise it. In other words, if the length normalized probability of one contour is greater than that of a second contour, then, we consider the first contour to be more salient than the second. This definition of contour saliency has an interesting relationship with the transition matrix  $\mathbf{P}$  constructed from the given contour (see [26]). If we imagine a scene containing just the closed contour  $\alpha$ , and where the probabilities between nonadjacent edges are zero, then the saliency of the contour is just the largest positive real eigenvalue of  $\mathbf{P}$ .

### 3.4 Importance of Directionality

We conclude this section by demonstrating how well the saliency measure performs for a simple example consisting of edges from the silhouettes of two pears artificially

superimposed on a background texture. See Fig. 1. The saliency measure for each edge was computed using the expression for  $c_i$  given in (3) after solving for the largest positive real eigenvalue of  $\mathbf{P}$  and its corresponding eigenvector. The saliency plot is shown in Fig. 2a. The length of an edge in the plot is proportional to its saliency. It can be plainly seen that the edges bounding both pears have high (and comparable) saliencies. The saliencies of all other edges have been suppressed. Numerically, their saliencies are 20 orders of magnitude smaller than those of the pears.

Using the same example, we demonstrate the importance of using pairs of directed edges to form a transition matrix  $\mathbf{P}$ , of size  $2N \times 2N$  as opposed to simply using the  $N$  edges to form a symmetric transition matrix  $\mathbf{A}$ , of size  $N \times N$ . Recall that this mechanism is required so that closed contours do not include reversals in direction at the locations of the edges. For the purpose of this demonstration, we construct a symmetric transition matrix  $\mathbf{A}$  from  $\mathbf{P}$  by setting  $A_{ij} = P_{ij} + P_{i\bar{j}} + P_{\bar{i}j} + P_{\bar{i}\bar{j}}$ . It can be verified that  $\mathbf{A}$  is symmetric because  $P_{ij} = P_{\bar{j}\bar{i}}$ . Fig. 2b shows the squared magnitude of the components of the eigenvector with largest positive real eigenvalue of  $\mathbf{A}$ . Two edges in the background texture that which, simply by chance, are proximal and very nearly collinear, are extremely salient while edges forming the closed boundary of the pears are ignored. It follows that using a nonsymmetric transition matrix  $\mathbf{P}$  and pairs of oppositely directed edges,  $i$  and  $\bar{i}$ , is essential to satisfactory performance of the saliency measure.

In order to distinguish the contours bounding the two pears, one might try to simply threshold the edge saliencies, i.e., the  $c_i$ s. However, as is illustrated by this example, edges from different objects can have saliencies of comparable magnitude. It is therefore likely that such a simple strategy will group together edges bounding distinct objects. In the next section, we develop a more robust approach that uses the link saliencies, i.e., the  $C_{ij}$ s, to group together sets of edges belonging to individual objects.

## 4 SEGMENTATION

The goal of segmentation is to group together into distinct sets, edges bounding distinct objects in the scene. To motivate our segmentation algorithm, consider the hypothetical case where some oracle provides us with a set  $S$  of closed contours in the scene whose saliencies are above some threshold. We can construct a graph whose vertices correspond to the edges in our scene. We create a directed link in this graph from edge  $i$  to edge  $j$  if  $i$  and  $j$  are successive edges of some salient contour in  $S$ . The Third Saliency Theorem (see Appendix A) tells us that such a construction induces a partition of the graph into a set of isolated *strongly-connected* components. A strongly-connected component [5] is a set of edges in which any pair of edges  $i$  and  $j$  have a path from one to the other, i.e.,  $i \rightarrow j$  as well as  $j \rightarrow i$ . In general, each strongly-connected component will contain multiple salient contours that share common edges. It is shown in Appendix A that the partition into a set of strongly-connected components is a direct consequence of the property of closure of the contours in  $S$ . As noted in the introduction, the strong dependence between the nature of the partition and the property of closure is a distinguishing feature of our approach, as compared with other approaches [6], [11] which employ generic graph search. More precisely, in our approach, the determination




Fig. 2. Saliency plots for the 2-pear example. (a) Our measure with directed edges. (b) Our measure with undirected edges. The length of each edge is proportional to its saliency value.

of strongly-connected components makes sense only in the context of a graph derived using a saliency relation based on contour closure.

In practice, of course, we do not know the salient contours beforehand. Nevertheless, since the links in the salient contours become the links in the graph, all we need to know is which of the links are salient, i.e., the likelihood that some salient contour passes through a given link. The link-saliencies, i.e., the  $C_{ij}$ s, encode precisely this information.

Ideally, the set of edges will be partitioned into isolated components. However, in practice, not all of the components provide reliable segmentations. The dominant contours tend to suppress the saliencies of all other contours to the degree that the saliencies of these nondominant contours are insufficient to induce components that can be isolated reliably. Hence, in practice, we begin by extracting the most salient contours, and since such contours will normally contain the most salient edge, we first identify the contours corresponding to the strongly-connected component containing the most salient edge. Having identified the most salient contours, we suppress their link saliencies in order to reveal the next set of dominant contours. We suppress the current set of dominant contours by deflating *all* transition probabilities between edges of the strongly-connected component. Specifically, if  $i$  and  $j$  are edges in the component, then the link  $i \rightarrow j$  is deflated by setting  $P_{ji} = 0$  (as well as setting the reversal-symmetric “sibling”  $P_{\bar{j}\bar{i}} = 0$ ). We then iterate this process to reveal multiple salient contours.

Ideally, the strongly-connected component containing the most salient edge will be isolated from the other components. In practice, due to noise, some of the  $C_{ij}$ s might wrongly indicate that the strongly-connected component containing the most salient edge is connected to one or more other strongly-connected components. Nevertheless, we can extract the component of interest by utilizing an important property of strongly-connected components: *The set of edges in a strongly-connected component containing a given edge is the intersection of the set of edges reachable from the given edge and the set of edges reachable if all links are reversed*. Because of reversal symmetry, the above property reduces to a particularly simple form. Let  $reachable(j)$  be the set of edges reachable from a given edge  $j$ . Due to reversal symmetry, it can be verified that the set of edges reachable from  $j$  when all links are reversed is the same as the *reversal* of the set of edges reachable from edge  $\bar{j}$ . The reversal of the set  $reachable(\bar{j})$  is defined to be

$$\overline{reachable(\bar{j})} = \{\bar{k} \mid k \in reachable(\bar{j})\}.$$

Hence, in order to identify the strongly-connected component containing the most salient edge  $j$ , we find

$$reachable(j) \cap \overline{reachable(\bar{j})}.$$

See Algorithm 1. Interestingly, the above is analogous to the expression for edge saliency,  $c_i = s_i \bar{s}_i$ . One needs simply to replace the the eigenvector  $s_i$  with the set  $reachable(i)$ , the reversal operator for vectors with the reversal operator for sets, and component-wise multiplication of vectors with intersection of sets. In our case, due to reversal symmetry, the above property reduces to a particularly simple form.

---

**Algorithm 1:** Extract an object.

---

```

extract(A)

begin

  P ← affinities(A)

  s ← eigenvector(P)

  j = argmax_i (s_i s_i)

  return reachable(j) ∩ overline{reachable(j)}
end

```

---

In order to decide if a link is salient or not, we need to threshold the  $C_{ij}$ s. We could use a single threshold for the entire graph. However, we can do better by choosing an adaptive threshold for the set of links which originate from edge  $j$ , i.e., the  $j$ th column of the link saliency matrix,  $C$ . To threshold these  $C_{ij}$ s in a natural manner, we sort the  $j$ th column in decreasing order. In this sorted list,  $z$ , we find the  $k$ th largest value (in all of the experiments in this paper  $k$  equals two). Edges joined by links from  $j$  with magnitude larger than the  $z_k$  are assumed to lie on salient closed contours, and are therefore selected. Such a thresholding scheme might misclassify certain links as salient. However, we have observed that the extraction of strongly-connected components is usually robust to such misclassifications. See Algorithm 3.

---

**Algorithm 2:** Segment an image.

---

```

segment(A)

begin

  S ← ∅

  for i ← 1 to N do

    B ← extract(A)

    S ← S ∪ {B}

  A ← A - B
  end

  return S
end

```

---



---

**Algorithm 3:** Compute edges reachable from  $j$ .

---

```

reachable(j)

begin

  B ← {j}

  if not visited(j) then

    for i ← 1 to |A| do

      C_ij ← s_i P_ij s_j

    end

    z ← sort(C[:, j])

    for i ← 1 to |A| do

      if C_ij > z_k then

        B ← B ∪ reachable(i)

      end
    end

  end

  return B
end

```

---

The termination criterion we use for the current implementation is to simply stop after reporting some fixed number of components. See Algorithm 2. One way to justify the use of such a simple criterion is to imagine a higher-level module that, for example, performs object-recognition, and which employs the segmentation algorithm to highlight regions where the presence or absence of some object can be determined using domain knowledge available to the recognizer. It is up to the recognizer to determine the number of salient contours that it wants to process (based on, for example, real-time constraints). If the module determines that the segmentation algorithm is reporting garbage after a certain number of iterations, then it can decide to terminate the search for additional contours. Alternatively, we could stop when the largest positive real eigenvalue  $\lambda$  becomes negligibly small.

As a demonstration, we apply the segmentation algorithm to the two pear example from Fig. 1. The segmentation in the first and second iterations is shown in Figs. 3a and 3b, respectively. As previously noted, a segmentation based on simply thresholding the edge saliencies would not be able to separate the two pears.

## 5 RESULTS

In this section, we show results of our segmentation algorithm on a few real images. All the images were taken using a Kodak DC50 480 × 480 pixel digital color camera. The Canny edge detector [4] was run on the images after converting them to gray scale, with the parameters  $\sigma = 3.0$ , low hysteresis threshold = 0.2 and high hysteresis threshold = 0.8. The set of edges returned by the Canny edge detector were found to be quite redundant. The edges are sampled to improve running times with almost no sacrifice




Fig. 3. Segmentation results for the two pear example. (a) First iteration. (b) Second iteration.

in performance. In our experiments, we sample the edges such that no two edges are closer than 5 pixels apart.

The entries of the transition matrix  $\mathbf{P}$  were calculated with parameter settings (see Section 2 for their descriptions and also [23])  $\gamma = 0.15$ ,  $T = 0.004$ , and  $\tau = 5.0$ . All edge images are remapped to a  $64 \times 64$  image size. Since the transition matrix  $\mathbf{P}$  has a special symmetry (the reversal symmetry), we had previously developed an algorithm that finds the eigenvector corresponding to the largest positive real eigenvalue of  $\mathbf{P}$  (required for the computation of the saliency relation) by exploiting the reversal symmetry. See [24] for details.

### 5.1 Example Segmentation

In our first example, we chose a simple scene where nonoccluding objects (fruits) were placed on a textured background (concrete). Fig. 4 shows four fruits on a concrete background in gray scale, Fig. 4a and the corresponding edge image, Fig. 4b (with 2,800 directed edges after the sampling process described above). Notice that the contrast between the texture of the fruit on the top-left (a cantaloupe) and that of the background is quite low. As a result, few edges are detected along some parts of the boundary of the cantaloupe. Figs. 4a, 4c, 4e, 4g, and 4i show the edge saliencies, i.e., the  $c_i$ s, computed during the first five iterations of the segmentation algorithm. Figs. 4b, 4d, 4f, 4h, and 4j show the corresponding contours that are extracted during those same iterations. It is interesting to note that the contour bounding the cantaloupe has been extracted despite the fact that there are large gaps in some parts of the contour.

Fig. 5a shows the variation of the saliency of the dominant contour across iterations. The dominant contour is extracted at each iteration by tracing out the most salient links starting from the most salient edge until we return to the most salient edge again. Its saliency is measured by the expression in (11). Since the saliency of the dominant contour decreases as we extract out successive contours, we see that the contours are indeed extracted in the order of their saliencies.

Finally, we give the time requirements for our algorithm for this example. It takes  $\approx 13$  seconds to generate a total of  $\approx 111,000$  entries in the sparse transition matrix  $\mathbf{P}$  on an SGI R10000. Since generating the entries in the matrix is easily parallelizable, it is useful to know the time per entry which is  $\approx 0.11$  msec. In Fig. 5b, we show the time taken to isolate successive objects. The eigensolver described above (see [24] for details) is adaptive, the time roughly varying according to the complexity of the contours extracted and

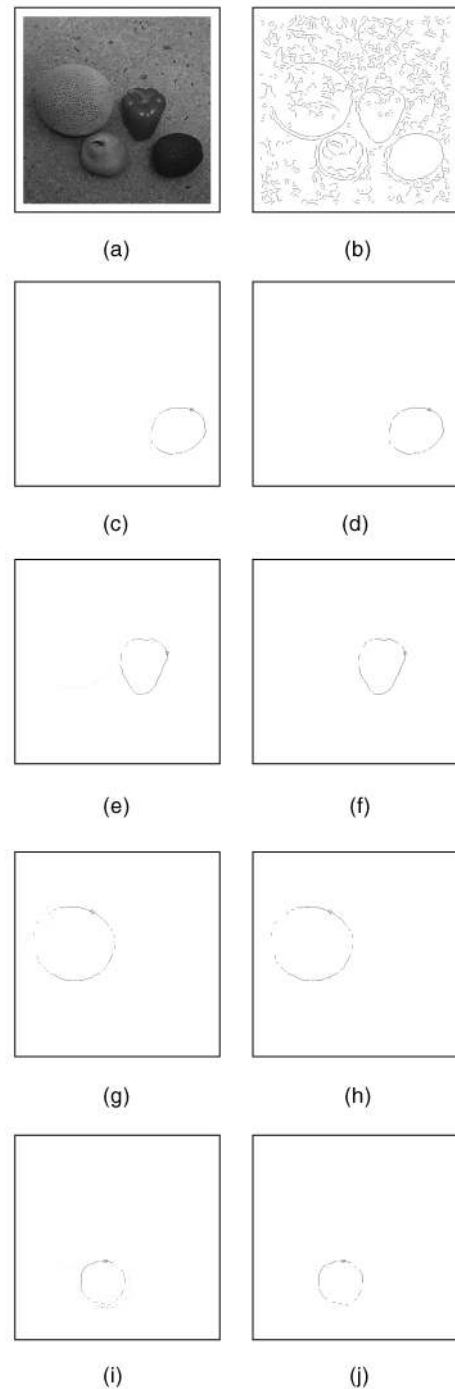


Fig. 4. Fruits on concrete. (a) Gray-scale image. (b) Canny edge output. (c), (d), (e), (f), (g), (h), (i), and (j) First four iterations of the segmentation algorithm. For each iteration, the saliencies are shown on the left and the segmentation is shown on the right. In each saliency plot, the length of an edge is proportional to the saliency of that edge. The most salient edge in both the saliency and segmentation plots are shown inside the small circle.

the number of edges each contains. As expected, the first iteration took the least time of  $\approx 3$  seconds since the contour extracted is relatively salient (as seen from the dominant contour saliency plot above) compared to the other contours in the scene. The third iteration took the longest time of  $\approx 20$  seconds—possibly because of the large gaps in the contour being extracted (bounding the cantaloupe). The average time for all iterations is  $\approx 9.9$  seconds.

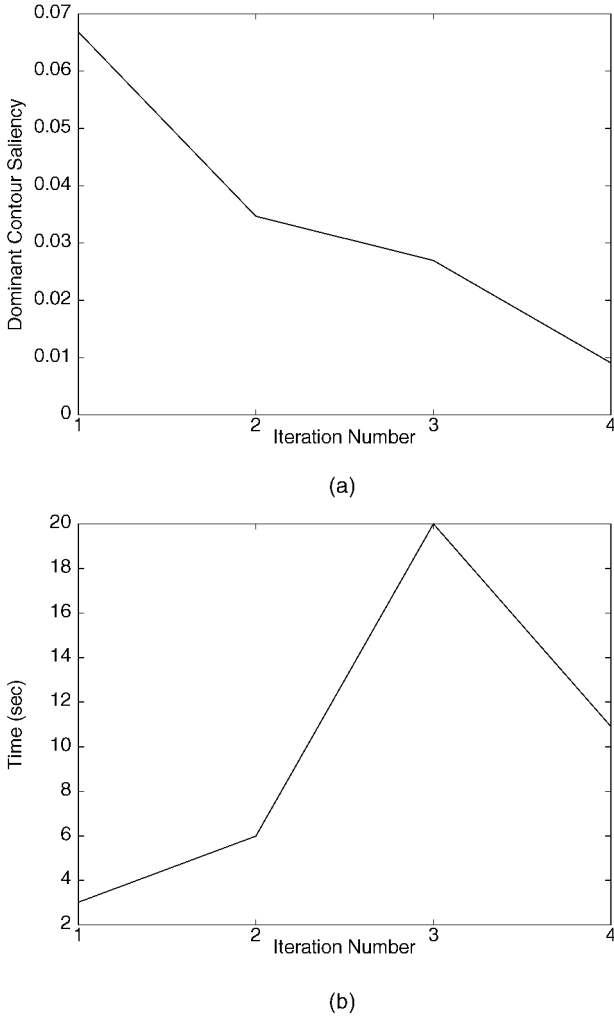


Fig. 5. (a) Dominant contour saliency versus iteration number. (b) Time per iteration.

## 5.2 Importance of Global Information

In this section, we will show the importance of the global information encoded in the link saliency matrix  $\mathbf{C}$  by replacing it with the transition matrix  $\mathbf{P}$  which encodes only local information. The edge saliency vector  $\mathbf{c}$  is left unchanged. With this replacement, the segmentation algorithm extracts out the same contour in the first iteration as the algorithm using the  $C_{ij}$ s. Note that this contour is easy to trace out since there are no large gaps present between successive edges of the contour. However, the hard part is to get a starting edge (i.e., the most salient edge in the current iteration) which (for this demonstration) is still being provided by the  $c_i$ s. Fig. 6 shows the segmentation after the second iteration. As can be seen, the segmentation completely breaks down. The  $P_{ij}$ s are sufficient as long as we start off from the most salient edge in each iteration and there are no large gaps in the contours being traced. The breakdown in the second iteration shows the need for the more global information encoded in the  $C_{ij}$ s in cases where there are large gaps in the contours being traced.

In a previous paper [26], a purely local saliency measure (termed WJ) was judged to be more effective in isolating smooth closed contours in the presence of background clutter than three other well-known saliency measures (those termed




Fig. 6. Fruits on concrete. (a) When the global link saliency matrix  $\mathbf{C}$  is replaced by the local transition matrix  $\mathbf{P}$ , from which it is derived, the segmentation algorithm fails in the second iteration. (b) When the global link saliency matrix  $\mathbf{C}$  is replaced by a local link saliency matrix  $\mathbf{W}$ , based on the WJ saliency measure, the segmentation algorithm fails in the first iteration.

GM, SB, and SU and based on [8], [18], and [21]). Analogous to the expression,  $c_i = s_i \bar{s}_i$  (where  $s$  and  $\bar{s}$  are right and left eigenvectors with largest positive real eigenvalue of  $\mathbf{P}$ ) the saliency of an edge according to the WJ measure is  $w_i = x_i \bar{x}_i$ , where  $x_i = \sum_j P_{ij}$ . We observe that the WJ measure can be seen as a single step of the power-method iteration necessary for computing the eigenvector with largest positive real eigenvalue of  $\mathbf{P}$ . It follows that, by comparing the performance of a segmentation algorithm based on the saliency measure derived in this paper (the measure termed WT in [26]) to one based on the WJ measure, we can ascertain the value of power-method iterations beyond the initial step. This speaks directly to the important issue of iterative versus noniterative (i.e., voting) methods in perceptual organization, an issue that is explored extensively in [14].

Using reversal symmetry,  $P_{ij} = P_{ji}$ , we can easily show that  $w_i = \sum_{j,k} P_{ki} P_{ij}$ , which implies that the saliency for edge  $i$  is the sum of the probabilities of contours of length two centered on  $i$ . Because the saliency is determined solely by the probabilities of length two contours, it follows that the global property of contour closure plays no role in determining edge saliency. Consequently, edges forming a closed contour can be of low saliency despite the fact that they contain many closed contours of relatively high probability.

In our second demonstration of the importance of using global information, the global link saliency matrix  $\mathbf{C}$ , based on the WT measure, is replaced by a local link saliency matrix  $\mathbf{W}$ , based on the WJ measure. The expression for local link saliency is analogous to the expression for global link saliency, (5). Specifically,

$$W_{ij} = \bar{x}_i P_{ij} x_j. \quad (12)$$

Using reversal-symmetry, it can be shown that

$$W_{ij} = \sum_{k,l} P_{ki} P_{ij} P_{jl}. \quad (13)$$

Thus, the  $W_{ij}$ s are proportional to the probability that a contour of length three is centered on the link  $j \rightarrow i$ . However, when we use this local relation, the segmentation algorithm breaks down in the first iteration as shown in Fig. 6b. The most salient edge according to the  $w_i$ s (indicated by the circle) lies on the cantaloupe, which was the third fruit extracted using the global saliency relation. While tracing the contour bounding the cantaloupe, the algorithm loses its way when

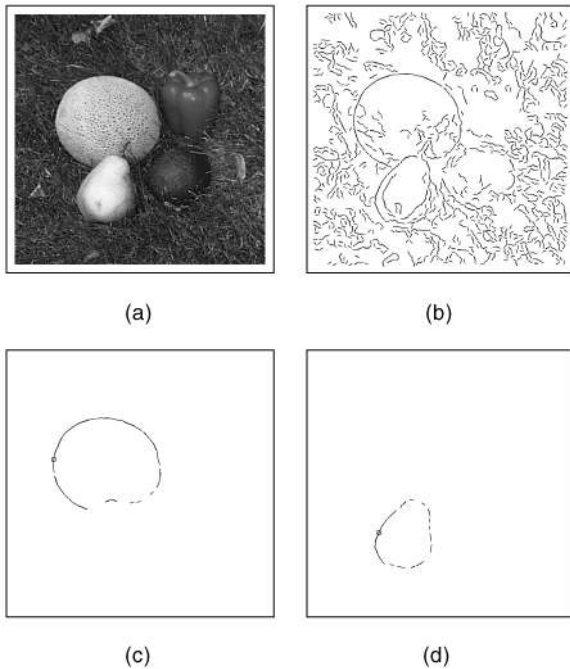


Fig. 7. Fruits on grass. (a) Gray-scale image. (b) Canny edge output. (c) and (d) First and second segmentations.

it encounters the large gaps. In summary, the  $c_i$ s are essential for reliably determining the starting edge for the segmentation algorithm, and the  $C_{ij}$ s are essential for bridging large gaps. Both are functions of the eigenvector the with largest positive real eigenvalue of the  $\mathbf{P}$  matrix.

### 5.3 Additional Segmentation Examples

Fig. 7 shows the same four fruits with grass as the background and with one of the fruits occluding another. Due to poor contrast between the two dark fruits and the background, the Canny edge detector does not reliably detect the edges bounding the two fruits. The fruits are hardly salient in the edge image (not shown), even for human observers. Our algorithm can be expected to extract out contours only when provided with reliable edge information. In this case, the algorithm picks out only the other two fruits in the image. Of the two fruits that it does pick out, one partly occludes the other. Due to the poor contrast between the two fruits, the edge information (especially the orientation) is quite poor in the region around the occlusion. However, despite this fact, and the fact that the contour bounding the occluded fruit contains a large gap at the occlusion, the algorithm segments out both fruits individually.

Fig. 8a shows an example where there are significant shadows that produce strong smooth contours adjacent to the stones. However, since they are not closed, the shadow contours are not as salient as the contours that actually bound the stones. Consequently, they do not confuse the algorithm.

Finally, Fig. 9a is an image of a large number of coins on a tabletop. Although this is an image which would be relatively easy to segment using image brightness, the segmentation which is shown in Fig. 9c, has been computed solely using the Canny edges shown in Fig. 9b.

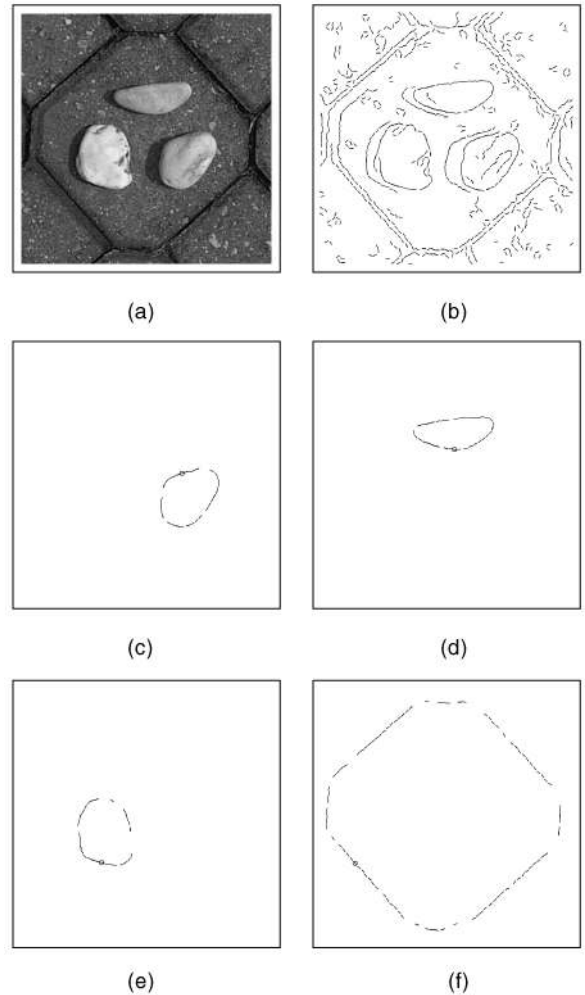


Fig. 8. Stone on pavement. (a) Gray-scale image. (b) Canny edge output. (c), (d), (e), and (f) First four segments.

## 6 FINDING THE OPTIMUM SPEED

The shape distribution that is used to build the  $\mathbf{P}$  matrix is defined by three parameters  $T$ ,  $\tau$ , and  $\gamma$ . Although there has been some interesting recent work on learning parameter settings for grouping algorithms (see [19]), we have simply selected values for  $T$ ,  $\tau$ , and  $\gamma$  which we have found yield good results in practice. In this section, we describe preliminary work on choosing the value of one of these parameters,  $\gamma$ , the particle's speed, automatically.

The segmentation algorithm which we have described assumes a fixed value for  $\gamma$ . However, there are two properties of the shape distribution that are directly affected by the particle's speed. First, the distance a particle travels before it decays *increases* with increasing speed. Second, the variance in a particle's direction of motion relative to the distance it travels *decreases* with increasing speed. Consequently, the choice of  $\gamma$  effectively determines both the curvature of the closed contours which will be classified as most salient and the optimum distance between adjacent edges. A more principled approach would be to isolate closed contours irrespective of their average curvature and irrespective of the average distance between adjacent edges. One way to do this would be to systematically vary speed within Algorithm 1 so that the contour which is most salient, i.e., the






Fig. 9. *Coins*. (a) Original image. (b) Edge input obtained by Canny detector. (c) Segmented objects, numbered in the order in which they are extracted.

contour with maximum eigenvalue, among contours of all possible  $\gamma$  is extracted. In principle, the saliency of contours of different average curvature and different edge sampling rates would be maximized at different speeds, resulting in a more robust segmentation algorithm. See Algorithm 4.

---

**Algorithm 4:** Extract an object (scale-invariant).

---

*extract*( $A$ )

**begin**

$\gamma_{\max} \leftarrow \underset{\gamma}{\operatorname{argmax}} \operatorname{eigenvalue}(\operatorname{affinities}(A, \gamma))$

$s \leftarrow \operatorname{eigenvector}(\operatorname{affinities}(A, \gamma_{\max}))$

$j = \underset{i}{\operatorname{argmax}} (s_i; s_i)$

**return**  $\operatorname{reachable}(j) \cap \overline{\operatorname{reachable}(\bar{j})}$

**end**

---

Fig. 10a is an image of three fruits on a wooden table with prominent wood grain background. The segmentation shown in Fig. 10c was computed using Algorithm 4 instead of Algorithm 1. The optimal speed within the range  $[0.01, 1.5]$  for each object was computed using Brent's method [3], which does not require analytic derivatives, and is able to locate a local optimum in  $\lambda(\gamma)$ .

In our initial attempt to run the modified segmentation algorithm on this image, we used the same values for  $T$  and  $\tau$  used to compute the other results in this paper. Unfortunately, the  $\gamma$  which maximized the eigenvalue was 1.5. Because this value is on the boundary of the search interval, it is not actually a local optimum of  $\lambda(\gamma)$  and the resulting segmentation was of very low quality. After increasing the diffusion constant,  $T$ , from 0.004 to 0.007 and the decay constant,  $\tau$ , from 5.0 to 6.5, the optimization procedure returned values of  $\gamma$  which were within the range  $[0.01, 1.5]$

which implies that they are true local optima. The optimal speed,  $\gamma_{\max}$ , for the first object was 1.23 and 1.27 for the second. The eigenvalue,  $\lambda(\gamma_{\max})$ , at the optimal speed was 0.0099 for the first object and 0.0071 for the second. The closed contours are of good quality. See Fig. 10c. The algorithm failed to find the third object because the Canny edge detector returned very few edges that lie on its boundary.

## 7 CONCLUSION

We have demonstrated how a saliency relation based on the global property of contour closure can form the basis of a segmentation algorithm able to identify multiple salient closed contours in real images. More specifically, we have demonstrated that computing the connected-components of a graph representing a saliency relation based on the relative number of closed contours containing pairs of edges, is more effective than searching a graph based on a purely local relation based on geometric properties of the pair of edges in isolation.

Our approach to grouping edges into salient closed contours involves the solution of an eigenvector/eigenvalue problem. Recently, other researchers [16], [18], [19], [22] have also proposed grouping image features by solving eigenvector/eigenvalue problems.

The normalized min-cut approach described in [22] can group more general image features than our approach can. However, since we restrict ourselves to grouping edges, we are able to impose the important constraint of tangent continuity, which has no counterpart for nonedge features such as texture or brightness. Furthermore, any approach enforcing tangent continuity using the mechanism of edge-directionality requires a nonsymmetric transition matrix  $\mathbf{P}$  for which the min-cut approach proposed in [22] does not apply. As previously noted, the use of a symmetric transition matrix makes contours containing cusps salient




Fig. 10. *Fruits*. (a) Original image. (b) Edge input obtained by Canny detector. (c) Segmented contours, numbered in the order in which they are extracted.

(see the discussion in Section 2). Hence, we would expect poor performance on edge grouping problems with a min-cut approach since tangent continuity cannot be enforced.

Like [22], the dominant eigenvector-based method described by [18] is applicable to features other than edges. Also like [22], this method assumes that the transition matrix is symmetric and, therefore, cannot use edge-directionality to enforce tangent continuity.

Of course, the generic grouping algorithms of [1], [7], [16], [22] make much weaker assumptions about the input image than does the algorithm which we describe here. The algorithm we describe is specifically designed to group edges into smooth closed contours. When an image does not contain closed contours, when the contours it contains are not smooth, or when local edge detection processes fail because of lack of contrast, i.e., when our assumptions are violated, our method will fail. It is possible that, in such cases, generic-grouping methods which are able to organize a wider variety of image features, and which make weaker assumptions about them, may succeed in such cases.

However, we believe that grouping-methods that are designed to solve a specific grouping problem, such as grouping edges into smooth closed contours, will outperform general-purpose methods on images for which their assumptions hold. This is because generic methods cannot fully exploit domain specific constraints such as contour closure and tangent continuity, which have no counterparts for non-edge features such as texture or brightness.

## APPENDIX A

First, we prove some preliminary lemmas.

**Lemma 1.** *If  $\mathbf{s}$  is a (right) eigenvector of  $\mathbf{P}$  with eigenvalue  $\lambda$ , then  $\bar{\mathbf{s}}$  is a left eigenvector of  $\mathbf{P}$  or equivalently a (right) eigenvector of  $\mathbf{P}^T$  with the same eigenvalue.*

**Proof.** Taking the  $i$ th component of  $\mathbf{P}^T \bar{\mathbf{s}}$ ,

$$\sum_j P_{ji} \bar{s}_j = \sum_j P_{ij} s_j \quad (14)$$

$$= \sum_j P_{ij} s_j \quad (15)$$

$$= \lambda \bar{s}_i. \quad (16)$$

Hence,  $\bar{\mathbf{s}}$  is an eigenvector for  $\mathbf{P}^T$  or equivalently a left eigenvector for  $\mathbf{P}$  with the same eigenvalue  $\lambda$ .  $\square$

**Lemma 2.** *For an irreducible positive matrix  $\mathbf{P}$  that is reversal-symmetric,*

$$\lim_{n \rightarrow \infty} \left( \frac{\mathbf{P}}{\lambda} \right)^n = \mathbf{s} \cdot \bar{\mathbf{s}}^T, \quad (17)$$

where  $\lambda$  is the largest eigenvalue of  $\mathbf{P}$  and  $\mathbf{s}$  is the corresponding eigenvector.

**Proof.** For a general irreducible positive matrix  $\mathbf{A}$ , it is shown in [10] that

$$\lim_{n \rightarrow \infty} \left( \frac{\mathbf{A}}{\lambda} \right)^n = \mathbf{x} \cdot \mathbf{y}^T, \quad (18)$$

where  $\mathbf{x}$  and  $\mathbf{y}$  are, respectively, the right and left eigenvectors of  $\mathbf{A}$  corresponding to the largest eigenvalue

for  $\lambda$  normalized such that  $\mathbf{x}^T \mathbf{y} = 1$ . From the previous lemma, we know that if  $\mathbf{s}$  is a left eigenvector of a reversal-symmetric matrix  $\mathbf{P}$ , then  $\bar{\mathbf{s}}$  is the corresponding right eigenvector with the same eigenvalue. Hence, the proof.  $\square$

**Proof of the First Saliency Theorem (Theorem 1).** First, we note the simple relationship between the diagonal elements of the powers of the transition matrix  $\mathbf{P}$  and probabilities of closed contours.  $(\mathbf{P}^k)_{ii}$  is the sum of the probabilities of all closed contours of length  $k$  that pass through edge  $i$ . Using this relationship and letting  $\lambda$  be the largest eigenvalue of  $\mathbf{P}$ , the definition for  $c_i$  in (2) can be rewritten in terms of the powers of  $\mathbf{P}$  as:

$$c_i = \lim_{n \rightarrow \infty} \frac{(\mathbf{P}^n)_{ii}}{\sum_j (\mathbf{P}^n)_{jj}} \quad (19)$$

$$= \lim_{n \rightarrow \infty} \frac{\left( \frac{\mathbf{P}}{\lambda} \right)_{ii}^n}{\sum_j \left( \frac{\mathbf{P}}{\lambda} \right)_{jj}^n}. \quad (20)$$

The above limit exists if the limit for both the numerator and the denominator exists and the limit for the denominator is nonzero. Using Lemma 2,

$$\lim_{n \rightarrow \infty} \left( \frac{\mathbf{P}}{\lambda} \right)_{ii}^n = s_i \bar{s}_i, \quad (21)$$

where the  $s_i$ s are the components of the eigenvector of  $\mathbf{P}$  corresponding to its largest eigenvalue  $\lambda$  and assuming that the eigenvector is normalized such that  $\sum_j s_j \bar{s}_j = 1$ . Hence, both the limits for the numerator and denominator in (20) exist and is equal to respectively,  $s_i \bar{s}_i$  and  $\sum_j s_j \bar{s}_j$ . Finally, we note that the limit for the denominator in the ratio is nonzero. For our problem there is a nonzero probability that a contour passes through any two edges in the image in succession.<sup>2</sup> Hence,  $\mathbf{P}$  is positive [10] and according to Perron's theorem [10] for positive matrices, all the components of the eigenvector corresponding to the largest eigenvalue are positive. Hence,  $\sum_j s_j \bar{s}_j > 0$  and, hence, the limit of the denominator in (20) is nonzero. Since we assume that the eigenvector is normalized so that  $\sum_j s_j \bar{s}_j = 1$ , the expression for  $c_i$  becomes:

$$c_i = \lim_{n \rightarrow \infty} \frac{\left( \frac{\mathbf{P}}{\lambda} \right)_{ii}^n}{\sum_j \left( \frac{\mathbf{P}}{\lambda} \right)_{jj}^n} \quad (22)$$

$$= \frac{\lim_{n \rightarrow \infty} \left( \frac{\mathbf{P}}{\lambda} \right)_{ii}^n}{\lim_{n \rightarrow \infty} \sum_j \left( \frac{\mathbf{P}}{\lambda} \right)_{jj}^n} \quad (23)$$

$$= \frac{s_i \bar{s}_i}{\sum_j s_j \bar{s}_j} \quad (24)$$

$$= s_i \bar{s}_i. \quad (25)$$

$\square$

**Proof the Second Saliency Theorem (Theorem 2).** The probability that closed contours of length  $n$  pass through edges  $j$  and  $i$  successively is given by  $P_{ji}^{n-1} P_{ij}$  since all such contours pass through the link from edge  $j$  to edge  $i$

2. This is not necessarily true when we consider sparse representations of the matrix  $\mathbf{P}$ , but for such a case, all that is required is that there is a nonzero probability that a contour start from edge  $i$  and end in edge  $j$  with the possibility of threading through intermediate edges.

at least once. Hence, we can rewrite the definition for the link saliencies (4) as:

$$C_{ij} = \lim_{n \rightarrow \infty} \frac{\left(\frac{P}{\lambda}\right)_{ji}^{n-1} \cdot \left(\frac{P_{ij}}{\lambda}\right)}{\sum_l \left(\frac{P}{\lambda}\right)_{ll}^n}. \quad (26)$$

Again, using the limit theorem in Lemma 2, and arguments similar to that made in the proof of Theorem 1 on the existence of limits, we have:

$$C_{ij} = \frac{\lim_{n \rightarrow \infty} \left(\frac{P}{\lambda}\right)_{ji}^{n-1} \cdot \left(\frac{P_{ij}}{\lambda}\right)}{\sum_l \lim_{n \rightarrow \infty} \left(\frac{P}{\lambda}\right)_{ll}^n} \quad (27)$$

$$= \frac{s_j \bar{s}_i \left(\frac{P_{ij}}{\lambda}\right)}{\sum_l s_l \bar{s}_l} \quad (28)$$

$$= \frac{\bar{s}_i P_{ij} s_j}{\lambda}. \quad (29)$$

□

**Theorem 3 (Third Saliency Theorem).** *Given a set  $\mathcal{C}$  of closed contours in an image, consider the induced graph  $G$  whose vertices are edges from the image. The only links between vertices correspond to the directed links between successive edges of the contours in  $\mathcal{C}$ . Then,  $G$  is partitioned into isolated strongly-connected components, no two of which have any link between them.*

**Proof.** It is easily seen that  $G$  has isolated strongly-connected components iff for two edges  $i$  and  $j$ , there is a path  $i \rightsquigarrow j$  iff there is a path  $j \rightsquigarrow i$ . Hence, in our case, we need to prove that the above condition between two edges  $i$  and  $j$  always holds. It is enough to show this for simple paths where there are no loops. Any path  $i \rightsquigarrow j$  can be decomposed into a sequence of subpaths, each of which is a subsequence fully contained in some closed contour of  $\mathcal{C}$ . The subsequences are constructed in the following manner. Consider the set  $A_i$  of all the contours in  $\mathcal{C}$  that contain edge  $i$ . Starting from edge  $i$  we trace out the the path  $i \rightsquigarrow j$ . As we move along this path, we remove from the set  $A_i$  any closed contours that does not contain the whole subsequence seen so far. Then, either we reach edge  $j$  before exhausting the contours in  $A_i$ , or  $A_i$  becomes empty at some intermediate edge. In the former case, any of the remaining closed contours in  $A_i$  provides a return path to edge  $i$  from  $j$  by tracing out the rest of such a closed contour. In the latter case, let  $k$  be the last intermediate edge after which the set  $A_i$  becomes empty. The path from  $i \rightsquigarrow k$  is the first subsequence that we construct. At  $k$ , there still exists some closed contour in  $A_i$ . Thus, there is a return path  $k \rightsquigarrow i$  by completing any such closed contour remaining in  $A_i$ . We recursively construct the remaining subpaths by considering the path  $k \rightsquigarrow j$  and starting with the set  $A_k$  which is the set of all contours in  $\mathcal{C}$  that pass through edge  $k$ . As argued above, each such subpath has a returning path from the end of the subpath to its beginning. Hence, all the returning paths can be concatenated together in reverse order to get a returning path from  $j \rightsquigarrow i$ . Hence,  $i \rightsquigarrow j$  iff  $j \rightsquigarrow i$ . A strongly connected component of a graph is defined [5] as a subset of nodes where for any two nodes  $i$  and  $j$  in the subset, there exists a path from  $i \rightsquigarrow j$  and from  $j \rightsquigarrow i$ . Hence, in

our case, since  $i \rightsquigarrow j$  iff  $j \rightsquigarrow i$ , if there exists any path  $i \rightsquigarrow j$ , then edges  $i$  and  $j$  belong to some strongly connected component. Hence, the whole graph can be partitioned into a set of isolated strongly connected components with no links between any two of the components. □

## APPENDIX B

In this appendix, we give an analytic expression characterizing the probability distribution of boundary-completion shapes derived in [23]. For a derivation of a related function, see [20]. We define the transition probability,  $P_{ji}$ , between two directed edges,  $i$  and  $j$ , to be

$$P(j | i) = \int_0^\infty P(j | i; t) dt \approx FP(j | i; t_{opt}), \quad (30)$$

where  $P(j | i; t)$  is the probability that a particle which begins its stochastic motion at  $(x_i, x_i, \theta_i)$  at time 0 will be at  $(x_j, y_j, \theta_j)$  at time  $t$ . This probability is defined to be the sum over the probabilities of all paths that a particle can take between the two edges. This integral is approximated analytically. The approximation is the product of  $P$  evaluated at the time at which the integral is maximized, i.e.,  $t_{opt}$ , and a weighting factor,  $F$ . The expression for  $P$  at time  $t$  is

$$P(j | i; t) = \frac{3 \exp\left[-\frac{6}{T^3} (at^2 - bt + c)\right] \exp\left(-\frac{t}{\tau}\right)}{\sqrt{\frac{\pi^3 T^3 t^2}{2}}}, \quad (31)$$

where

$$a = \frac{2 + \cos(\theta_j - \theta_i)}{3} \quad (32)$$

$$b = \frac{x_{ji}(\cos \theta_j + \cos \theta_i) + y_{ji}(\sin \theta_j + \sin \theta_i)}{\gamma} \quad (33)$$

$$c = \frac{(x_{ji}^2 + y_{ji}^2)}{\gamma^2} \quad (34)$$

for  $x_{ji} = x_j - x_i$  and  $y_{ji} = y_j - y_i$ . The distribution of shapes is determined by the half-life,  $\tau$ , the variance,  $T$ , and the speed of the particle,  $\gamma$ . The expression for  $P$  should be evaluated at  $t = t_{opt}$ , where  $t_{opt}$  is real, positive, and satisfies the cubic equation

$$-\frac{7t^3}{4} + \frac{3(at^2 - 2bt + 3c)}{T} = 0. \quad (35)$$

If more than one real, positive root exists, then the root maximizing  $P(j | i; t)$  is chosen. Finally, the extra factor  $F$  is

$$F = \sqrt{\frac{2\pi t_{opt}^5}{\frac{12(3c - bt_{opt})}{T} + \frac{7t_{opt}^3}{2}}}. \quad (36)$$

## ACKNOWLEDGMENTS

The authors would like to thank Satish Rao for first suggesting a connection between our segmentation algorithm and strongly-connected components. Shyjan Mahamud and Lance R. Williams were partially supported by NEC Research Institute. Lance R. Williams and Kanglin Xu were partially supported by Los Alamos National Laboratory.

## REFERENCES

- [1] A. Amir and M. Lindenbaum, "A Generic Grouping Algorithm and Its Quantitative Analysis," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 20, no. 2, pp. 168-185, Feb. 1998.
- [2] M. Boldt, R. Weiss, and E. Riseman, "Token-Based Extraction of Straight Lines," *IEEE Trans. Systems, Man, and Cybernetics*, vol. 19, no. 6, pp. 1581-1594, 1989.
- [3] R.P. Brent, *Algorithms for Minimization without Derivatives*. Prentice-Hall, 1973.
- [4] J. Canny, "A Computational Approach to Edge Detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 9, no. 6, pp. 679-698, 1986.
- [5] T.H. Cormen, C.E. Leiserson, and R.L. Rivest, *Introduction to Algorithms*, chapter 23, MIT Press, 1989.
- [6] J.H. Elder and S.W. Zucker, "Computing Contour Closure," *Proc. European Conf. Computer Vision*, vol. I, pp. 14-18, 1996.
- [7] Y. Gdalyahu, D. Weinshall, and M. Werman, "A Randomized Algorithm for Pairwise Clustering," *Advances in Neural Information Processing Systems*, vol. 11, 1999.
- [8] G. Guy and G. Medioni, "Inferring Global Perceptual Contours from Local Features," *Int'l J. Computer Vision*, vol. 20, pp. 113-133, 1996.
- [9] L. Herault and R. Horaud, "Figure-Ground Discrimination: A Combinatorial Optimization Approach," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, pp. 899-914, 1993.
- [10] R.A. Horn and C.R. Johnson, *Matrix Analysis*, chapter 8, Cambridge, U.K.: Cambridge Univ. Press, 1990.
- [11] D.W. Jacobs, "Robust and Efficient Detection of Convex Groups," *Proc. Computer Vision and Pattern Recognition*, 1993.
- [12] D.G. Lowe, *Perceptual Organization and Visual Recognition*. Boston: Kluwer, 1985.
- [13] S. Mahamud, K.K. Thornber, and L.R. Williams, "Segmentation of Salient Closed Contours from Real Images," *Proc. Int'l Conf. Computer Vision*, 1999.
- [14] G. Medioni, M. Lee, and C. Tang, *A Computational Framework for Segmentation and Grouping*. Elsevier Science B.V., 2000.
- [15] D. Mumford, "Elastica and Computer Vision," *Algebraic Geometry and Its Applications*, C. Bajaj, ed., New York: Springer-Verlag, 1994.
- [16] P. Perona and W. Freeman, "A Factorization Approach to Grouping," *Proc. European Conf. Computer Vision*, vol. I, pp. 655-670, 1998.
- [17] S.V. Raman, S. Sarkar, and K.L. Boyer, "Hypothesizing Structures in Edge-Focused Cerebral Magnetic Resonance Images Using Graph-Theoretic Cycle Enumeration," *Computer Vision, Graphics, and Image Processing*, vol. 57, no. 1, pp. 81-98, 1993.
- [18] S. Sarkar and K. Boyer, "Quantitative Measures for Change Based on Feature Organization: Eigenvalues and Eigenvectors," *Computer Vision and Image Understanding*, vol. 71, pp. 110-136, 1998.
- [19] S. Sarkar and P. Soundararajan, "Supervised Learning of Large Perceptual Organization: Graph Spectral Partitioning and Learning Automata," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 5, May 2000.
- [20] E. Sharon, A. Brandt, R. Basri, "Completion Energies and Scale," *Proc. Computer Vision and Pattern Recognition*, 1997.
- [21] A. Shashua and S. Ullman, "Structural Saliency: The Detection of Globally Salient Structures Using a Locally Connected Network," *Proc. Int'l Conf. Computer Vision*, 1988.
- [22] J. Shi and J. Malik, "Normalized Cuts and Image Segmentation," *Proc. Computer Vision and Pattern Recognition*, 1997.
- [23] K.K. Thornber and L.R. Williams, "Analytic Solution of Stochastic Completion Fields," *Biological Cybernetics*, vol. 75, pp. 141-151, 1996.
- [24] K.K. Thornber, S. Mahamud, and L.R. Williams, "The Eigenvalue Problem for Reversal Matrices," NEC Technical Report 97-162, 1998.
- [25] L.R. Williams and D.W. Jacobs, "Stochastic Completion Fields: A Neural Model of Illusory Contour Shape and Saliency," *Neural Computation*, vol. 9, no. 4, pp. 837-858, 1997.
- [26] L.R. Williams and K.K. Thornber, "A Comparison of Measures for Detecting Natural Shapes in Cluttered Backgrounds," *Int'l J. Computer Vision*, vol. 34, no. 2/3, pp. 81-96, 2000.
- [27] L.R. Williams and K.K. Thornber, "Orientation, Scale and Discontinuity as Emergent Properties of Illusory Contour Shape," *Neural Computation*, vol. 13, no. 8, pp. 1683-1711, 2001.



**Shyjan Mahamud** received the BTech degree in computer science and engineering from the Indian Institute of Technology, Mumbai, India in 1992. He received the MS degree in computer science from the University of Massachusetts at Amherst in 1995 and the PhD degree in computer science from Carnegie Mellon University in 2002. His research has focused on computer vision, including perceptual grouping, structure from motion, and object recognition.



**Lance R. Williams** received the BS degree in computer science from the Pennsylvania State University in 1985 and the MS and PhD degrees in computer science from the University of Massachusetts at Amherst in 1988 and 1994, respectively. He spent four years as a post-doctoral scientist at NEC Research Institute. He is now an assistant professor in the Computer Science Department at the University of New Mexico in Albuquerque. His principal research interests are human visual information processing, perceptual organization for artificial vision systems, and neural computation.



**Karvel K. Thornber** received the PhD degree from Caltech in 1966 in electrical engineering, physics, and applied mathematics. He worked at Bell Laboratories for 20 years, where he worked on carrier transport in electromagnetic fields in dissipative media with applications like the design of submicron, field effect transistors. He left Bell Laboratories and joined NEC Research Institute (NECI) at its inception in 1989. At NECI, he has published papers on continuous valued

logic and human and computer vision. His greatest joy has been and remains working on unsolved problems at the interface of physics and computation.



**Kanglin Xu** received the MS degree in mechanical engineering from Shanghai University of Science and Technology in February 1990. He was a faculty member at Shanghai University from 1990 to 1996, working in the areas of robotics, flexible manufacturing systems, and computer integrated manufacturing systems. Since 1998, he has been studying at the University of New Mexico and expects to finish the MS degree in computer science in Decem-

ber 2002. His research interests include robotics, image processing, and applications of computers in manufacturing engineering.

► For more information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.