PERGAMON

Pattern Recognition 36 (2003) 2349–2362

PATTERN RECOGNITION
THE JOURNAL OF THE PATTERN RECOGNITION SOCIETY

www.elsevier.com/locate/patcog

Inferring region salience from binary and gray-level images

Yossi Cohen, Ronen Basri*,[1]

Department of Computer Science and Applied Math, The Weizmann Institute of Science, Rehovot 76100, Israel

Received 18 March 2002; received in revised form 19 February 2003; accepted 19 February 2003

Abstract

We introduce a method that uses contour fragments to highlight regions of interest. Our method obtains as input either a binary image or the gradient map of a gray-level image. It produces a saliency map that reflects for every point in the image our belief that it belongs to a salient region. Saliency is determined by criteria such as closure, convexity, and size. In addition, gaps in the boundaries of regions diminish their saliency. Explicit scale parameter determines the size of interest. The method is implemented by a convolution of the input edge image with a linear filter that specifies the region of influence of a contour point over the image. Experiments demonstrate the utility of the method for saliency and segmentation.

Keywords: Perceptual grouping; Region enhancement; Region salience

1. Introduction

Computational approaches to perceptual grouping group together contour fragments in an attempt to detect regularities in images. Grouping processes attempt to rank image structures according to their perceptual saliency, they identify curves in images, and they complete curves over gaps. The outcome of such processes is used to detect objects of interest ("attention"), separate them from their background, and perform image segmentation.

In this paper we introduce a method that uses contour fragments to highlight regions of interest. Our method produces, given a contour image, a saliency map that reflects for every point in the image our belief that it belongs to a salient region. The saliency values are determined by criteria such as closure, convexity, and size. In addition, gaps in the boundaries of regions diminish their saliency. Our experiments demonstrate the utility of our method for saliency and segmentation.

* Corresponding author. Tel.: +972-8-934-2809; fax: +972-8-934-2945.

E-mail address: ronen.basri@weizmann.ac.il (R. Basri).

[1] Research was supported in part by the Israeli Ministry of Science, Grant No. 2104. The vision group at the Weizmann Institute is supported in part by the Moross Foundation.

Existing techniques for grouping curve fragments largely consider criteria such as smoothness, curvilinearity, and the length of curves and disregard region considerations. Most of these techniques optimize measures based on Elastica and its variations ([1–7], see a comparison of applying several such methods to closed contours in Ref. [8]). These measures prefer an open, straight line over a closed curve. Only a few methods combine curve with region considerations. Shashua and Ullman's Saliency Network ([9], see an analysis in Ref. [10]) ranks curves according to their length and smoothness. Closed curves are encouraged by treating them as infinitely long. Huttenlocher and Wayner [11] and Jacobs [12] developed methods for detecting curve fragments that form convex configurations. Elder and Zucker [13] search for optimal cycles in a graph whose nodes represent curve fragments and links are weighted according to a likelihood measure of contiguity. Leung and Malik [14] combine intensity-based region segmentation with curve completion. Their method determines whether two pixels belong to the same region by considering both the difference in their intensities as well as the likelihood that they are separated by a subjective contour. Jermyn and Ishikawa [15] optimize an energy function that combines region with boundary information through an application of Green's theorem. Geiger et al. [16–18] propose a diffusion process in which saliency

0031-3203/03/$30.00 © 2003 Pattern Recognition Society. Published by Elsevier Ltd. All rights reserved.
doi:10.1016/S0031-3203(03)00120-1

values diffuse from junctions to the regions in the image. Finally, Williams and Thornber [8] propose a method for computing the salience of closed curves through a mechanism of random walks.

The method we present assigns saliency values to points according to region considerations. The method obtains as input either a contour image or the gradient map of a gray-level image. It produces a dense saliency map in which every point in the image is given a saliency value reflecting its likelihood to belong to a salient region. The method enables focusing attention to regions of different scales by setting an explicit scale parameter. The saliency map is computed by convolving the gradient image with a linear filter that specifies the region of influence of a contour point over the image. This convolution can also be viewed as a process that diffuses the saliency from the edge points to the regions. We show below that the most salient curve according to the method is a circle of the selected scale and that the largest saliency value is obtained at its center. Our method differs from most previous approaches by considering *explicitly* region criteria. (For instance, similar to our method, Guy and Medioni [6] also encourage circular extensions, but their method still prefers an open straight line over a closed circular curve.) The method is most related to the diffusion process proposed in Refs. [16–18]. Both methods use convolution to detect regions in contour images. However, Refs. [16–18] diffuse the saliency values from junction points, whereas in our method the saliency values diffuse from the edge elements. This leads to quite significant differences in the methods. Also, by diffusing the saliency values from the edge elements we avoid the need to accurately detect and classify the junctions in the image.

The paper contains the following sections. In Section 2 we present our method for computing the saliency values. In this section we introduce a filter whose objective is to highlight, given an edge element, the points that are likely to be included in a region whose boundaries include the element. The saliency map is obtained by convolving the edge image with this filter. In Section 3 we introduce the boundary field, which highlights the boundaries of salient regions. We show in addition that the most salient region obtained with our method is a circle of the selected scale. In Section 4 we exploit contrast and convexity to reduce the blur in the saliency map. Experimental results are shown in Section 5.

2. The model

Our scheme obtains as input a gradient image; every pixel u has two values, the magnitude of the gradient at u and the direction of the gradient. Given a gray level image the gradient image can be produced from the image by applying a standard gradient operator (e.g., Sobel). Alternatively, if a contour image is given, we produce a gradient image by setting the gradient magnitude to 1 at every contour pixel and 0 elsewhere. The direction of the gradient in this case is set along the normal to the contour at the pixel. Note that in this case there is a two-way ambiguity in the direction of the gradient. For simplicity we assume below that only a single edge may pass through any image pixel. Thus, we interchangeably refer to a pixel u also as an edge element u. It is straightforward to modify this assumption to allow multiple edges to pass through a pixel.

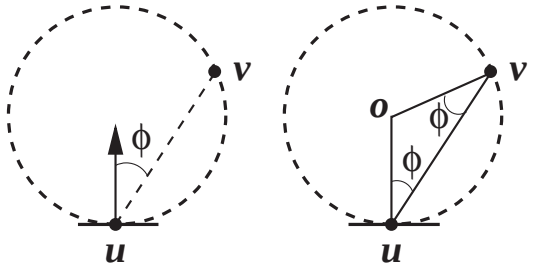

Fig. 1. Constructing an induction field around an edge element at u. The left figure shows a point v and the circle through v that is tangent to u. ϕ denotes the angle between v–u and the normal to u. The right figure shows that the triangle $\triangle ouv$ is isosceles.

Given an edge element u, we determine the saliency induced by u over the image to reflect the likelihood of every pixel in the image, in the absence of other information, to belong to a region whose boundary includes u. This saliency is set according to two principles. First, neighboring points to u are considered more salient than farther points. Second, all the points that lie along a circle that is tangent to u (excluding u itself) are equally salient. These two principles are achieved by defining the saliency induced by u as a Gaussian function that decays with the diameter of a circle whose tangent is u.

Formally, let u denote the position of an edge element, and let $\theta(u)$ denote its normal orientation. Given an edge element u we define a coordinate system around u as follows. The center of the coordinate system is set at u, the X-axis is aligned with the tangent direction to u, and the Y-axis is aligned with the normal direction, $\theta(u)$. In this coordinate frame the family of circles tangent to u are described by the formula:

$$x^2 + (y \pm r)^2 = r^2 \tag{1}$$

for $0 < r < \infty$. Given a point $v = (x, y)$, let $h = 2r$ denote the diameter of the circle through v that is tangent to u. By rearranging Eq. (1), we obtain that

$$h = \frac{x^2 + y^2}{|y|}. \tag{2}$$

Alternatively, we can express the diameter h such that it is independent of the choice of coordinate system as follows. Let ϕ be the angle between the normal to u and the vector $v - u$ (see Fig. 1 (left)). Then,

$$\cos\phi = \frac{y}{\sqrt{x^2 + y^2}}. \tag{3}$$

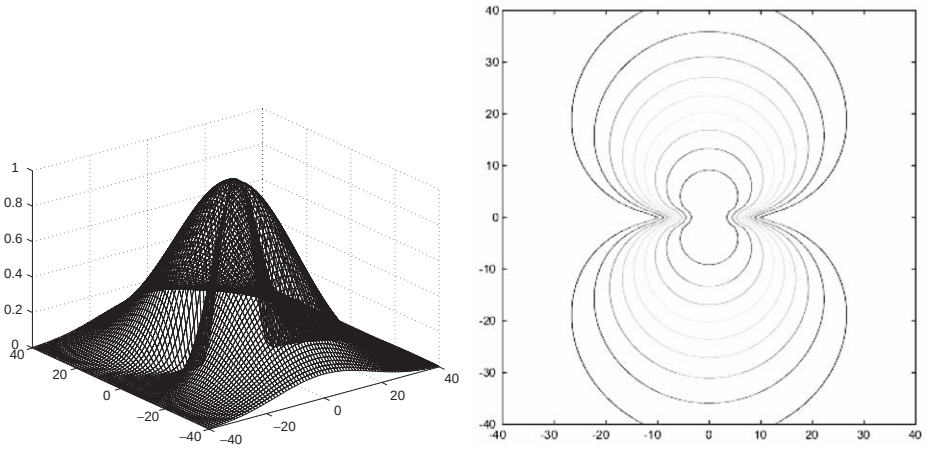

Fig. 2. The induction field associated with an edge element (left), and its level sets (right), $\sigma = 30$.

Substituting Eq. (3) into Eq. (2) and replacing $\sqrt{x^2 + y^2}$ by $\|v - u\|$ we obtain

$$h = \frac{\|v - u\|}{|\cos\phi|}. \tag{4}$$

Note that $h = \infty$ along the X-axis ($y = 0$), except at the origin where $h = 0$.

The salience induced by an edge element u, which we refer to as the *induction field* of u, is a Gaussian function that decays with the diameter of the circle whose tangent is u:

$$k(u, v) = \frac{1}{\pi\sigma^2} \exp\left(-\frac{h^2}{2\sigma^2}\right), \tag{5}$$

where h can be replaced by either of the expressions in Eqs. (2) or (4). The constant factor $1/\pi\sigma^2$ normalizes the induction field so that its integral over the entire 2-D plane is one (see Appendix A). Note that although h diverges along the X-axis k is finite everywhere in the plane. In particular,

$$\lim_{y\to 0} k(u, v) = \begin{cases} \dfrac{1}{\pi\sigma^2} & x = 0, \\ 0 & x \neq 0. \end{cases} \tag{6}$$

The scale parameter σ determines the rate of decay of the induction field and consequently, as is shown in Section 3 below, the size of the highlighted regions. In application of the scheme the scale parameter should vary according to the size of the viewed objects, the distance of the objects from the viewer, and the sampling rate of the image. Applying the method at several different scales is often desired in order to highlight all the objects of interest.

Fig. 2 shows a 3-D plot of the induction field of an edge element placed at the origin and the level sets of this induction field for a particular choice of σ.

The salience of a point v is obtained by integrating the induction fields at v with respect to all edge pixels u as follows:

$$R(v) = \int k(u, v) f(u)\, \mathrm{d}u, \tag{7}$$

where $f(u)$ is a weight assigned to the edge element u. In a gradient image, for example, the weight $f(u)$ may be proportional to the magnitude of the gradient at u, whereas in a binary image it may be constant for all edges.

The process of producing the saliency map can be described also as a process in which saliency values diffuse from the edge elements to the entire image. In our scheme the saliency values propagate along straight lines, and the amount diffused varies according to the angular deviation of the propagation direction from the normal to the source element.

3. The shape of salient regions

The saliency map produced with our method highlights the regions of interest, but the boundaries of these regions are not given explicitly in the saliency map. To extract the shape of the highlighted regions we may detect edges in the saliency map, e.g., by considering the gradient of the saliency map. Based on the gradient we define below the boundary field, which indicates the locations of the boundaries of highlighted regions. We show further that this boundary field may be obtained by convolving the image with the boundary field induced by every edge element separately, and derive an explicit expression for this boundary field. In addition, by analyzing the boundary field induced by the edge elements we claim that the most salient region is a circle whose diameter equals the selected scale σ.

The *boundary field* is defined to be a 2-D vector field whose magnitude at every point reflects the likelihood that the point lies on the boundary of a salient region and whose

orientation indicates the normal direction to the boundary. The boundary field is defined as

$$\vec{B}(v) = \sigma^3 \vec{\nabla} R(v), \tag{8}$$

where $\vec{\nabla} = (\partial/\partial x, \partial/\partial y)$ denotes the gradient operator. The constant factor σ^3 multiplies the gradient in order to make the magnitude of the boundary field comparable at different scales, as will become apparent shortly.

The boundary field can be computed in two ways. The first is to apply the definition directly, namely, the saliency map $R(v)$ is computed, and then its gradient is derived. Alternatively, the boundary field can be inferred directly from the boundary fields induced by every edge element separately. Because of the linearity of the gradient operator, the boundary field can be expressed simply as a sum of the boundary fields induced by the individual elements.

Below we provide analytic expressions to the boundary field induced by an edge element. We first determine the direction of the boundary field. Denote the boundary field induced by an element u by $\vec{b}(u,v) = \sigma^3 \vec{\nabla} k(u,v)$. The level sets of $k(u,v)$ are the circles that are tangent to u, and $k(u,v)$ decays with the diameter of the circle. This implies that the gradient vectors are orthogonal to those circles, and that they point inward toward the centers of the circles.

Next, we determine the magnitude of the boundary field. Let v lie along a ray l through u that forms an angle ϕ with the normal to u (see Fig. 1 (left)). Denote the distance between u and v by $q = \|v - u\|$, according to Eq. (4)

$$\frac{\mathrm{d}q}{\mathrm{d}h} = |\cos\phi|. \tag{9}$$

Denote by o the center of the circle. Note that the triangle $\triangle ouv$ is isosceles (see Fig. 1 (right)). Its base is v–u, the side o–u coincides with the normal to u, and the side o–v coincides with the gradient vector at v. Consequently, the angle between the direction of the gradient of the induction field $k(u,v)$ at v and the line l connecting u and v is ϕ. Thus,

$$|\vec{\nabla} k(u,v)| = \frac{1}{|\cos\phi|}\,\frac{\mathrm{d}k}{\mathrm{d}q}, \tag{10}$$

and from Eq. (9),

$$|\vec{\nabla} k(u,v)| = \frac{1}{\cos^2\phi}\,\frac{\mathrm{d}k}{\mathrm{d}h}. \tag{11}$$

Finally, using Eq. (5) we obtain

$$|\vec{b}(u,v)| = -\frac{h}{\pi\sigma\cos^2\phi}\exp\left(-\frac{h^2}{2\sigma^2}\right). \tag{12}$$

(Note that as $\phi \to \pm\pi/2$, $h \to \infty$ and consequently $\vec{b}(u,v) \to 0$.) Notice that the magnitude of the boundary field is invariant to the sampling rate of the image since increasing the diameter h and the scale parameter σ by the same proportion maintains the magnitude constant.

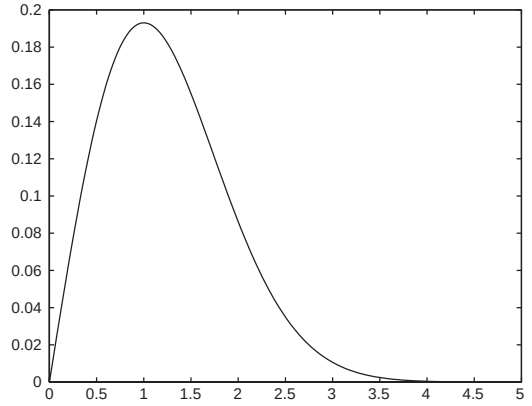

Fig. 3. $|\vec{b}(u,v)|$ as a function of h/σ.

It is not difficult to show that $|\vec{b}(u,v)|$ obtains its maximum value when $h = \sigma$. Thus, given a single edge element u, the most likely regions according to our scheme are the two circles of diameter σ that are tangent to u. Said another way, the most likely continuations for this element u are the boundaries of these two circles. Fig. 3 shows $|\vec{b}(u,v)|$ as a function of h/σ.

The analysis above gives an insight to the type of regions highlighted by our scheme. Every edge element u contributes most of its saliency to the circles of diameter σ that are tangent to u. Consequently, every cross section of a region that is consistent with a circle whose diameter is roughly σ contributes to the saliency of the region. The boundaries of the highlighted region is determined by a vector summation of the boundary fields induced by every contour element. Note that a vector summation of influences was used previously in Ref. [6] to highlight smooth curves.

Finally, it can be shown that when the method is applied to an image of an ellipse its center point obtains the maximal saliency. To see this notice that at the center of the ellipse the boundary field induced by any element along the boundary of the ellipse is canceled by the boundary field induced by the opposite element. Thus, the gradient of the saliency map at the center of the ellipse vanishes. Symmetry considerations imply that this maximum is global. For a circle of radius r the saliency at the center point c_r is given by

$$R(c_r) = \frac{2r}{\sigma^2}\exp\left(-\frac{r^2}{2\sigma^2}\right). \tag{13}$$

4. Enhancing the saliency map

When we consider the bounding contour of an object we need to take into account the two regions that lie on both sides of the contour. Assuming the object is in general position, one side contains the bounded object, while the other corresponds to the occluded background. Given an edge element it is unknown a-priori in which side of the edge the

bounded object lies. Therefore, our scheme associates with an edge element an induction field that extends symmetrically to both sides of the edge. In general, this gives a preference to convex regions because the saliency emitted from the boundaries of a convex region tends to concentrate over a relatively small area, whereas in non-convex regions the saliency tends to be spread over larger areas. However, as we shall see in our examples, this tends to somewhat blur the saliency map.

One possible way to reduce this blur is to measure the local curvature of the edge and limit the associated induction field to the convex side of the edge. This solution, however, involves an estimation of second derivatives of the contour. This might be a problem when the boundaries of objects contain nearly straight line segments. In this section we discuss two other methods to enhance the saliency map. The first method uses the (signed) direction of the gradient of edge elements and is appropriate for both binary and gray-level images. The second method uses pairwise interactions of edge elements. These methods are described in the rest of this section.

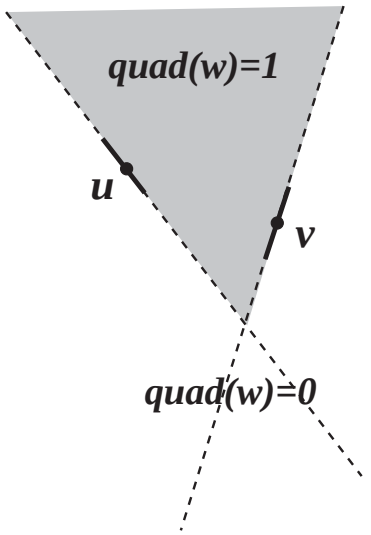

Fig. 4. quad(w) defined by the two elements u and v.

4.1. Signed induction fields

The first method, which we call *signed induction fields*, is effective when regions of interest contain intensities that are either all brighter or all darker than their immediate background. The method is based on preserving the sign of the contrast between a region and its background while the saliency is accumulated. This method is often appropriate for gray-level and binary images. The method is inappropriate for contour images or when a region includes intensities that are both brighter than parts of its immediate background and darker than other parts of its immediate background. An example for this condition is when a gray object appears in front of a checkerboard.

In this method we modify the induction field (5) as follows:

$$k_s(u,v) = \mathrm{sign}(v)k(u,v), \tag{14}$$

where sign(v) is defined as 1 if v lies in the brighter side of u (in other words, if the angle between the gradient direction at u and the vector $v-u$ is less or equal to π), and -1 if v is in the darker side of u. The saliency map is obtained in a similar manner to Eq. (7), namely, by integrating k_s over all edge elements u. The result of this summation would attribute large positive values to points inside bright salient regions and low negative values to points inside dark salient regions. The saliency of a point, therefore, is now determined by the *absolute value* of the obtained saliency.

4.2. Pairwise induction fields

In the second method, which we refer to as *pairwise induction fields*, the induction fields are determined jointly by pairs of elements. By considering pairs of elements we can confine the saliency induced by the elements to the part of the plane where the bounded region is most likely to appear. This enhances the saliency map and creates preference to convex regions. The method is based on the following principles. Recall that in the original model the induction field of an element is produced by diffusing saliency values from the element to the entire image. Roughly speaking, in the new method when this saliency reaches a contour point it is reflected and continues diffusing along a new direction. The new directions are limited to the portion of the plane where the construction of a convex shape through the source and the reflecting edge elements is possible. Below we define this model formally and show that it can be computed by a second application of our original model, after a slight modification.

Given two edge elements in general position, u and v, consider the lines obtained by extending the elements to infinity. These lines divide the plane into four sections (three sections if the elements are parallel). One of these sections is bounded by the two half lines that include u and v. In our model only points in this section obtain saliency from u and v. Define the function quad(w) to be 1 inside this section and 0 elsewhere (see Fig. 4). In order to define the induction field for this model we need to determine the diameter h that corresponds to the distance traveled by the diffusing saliency. We define this diameter by

$$h_{uv}^2(w) = h_u^2(v) + h_v^2(w), \tag{15}$$

where $h_u(v)$ denotes the diameter of the circle through v that is tangent to u and $h_v(w)$ denotes the diameter of the circle through w that is tangent to v. The pairwise induction field is defined by

$$k_p(u,v,w) = \mathrm{quad}(w)\frac{1}{\pi^2\sigma^4}\exp\left(-\frac{h_{uv}^2(w)}{2\sigma^2}\right), \tag{16}$$

and the pairwise saliency map by

$$R_p(w) = \int\int k_p(u,v,w) f(u) f(v)\,\mathrm{d}u\,\mathrm{d}v, \tag{17}$$

where the integration is performed over all pairs of edge elements.

We now show that the pairwise saliency map can be computed as a two-step application of our original model with some modification. First, notice that $k_p(u,v,w) = \mathrm{quad}(w)k(u,v)k(v,w)$. The term $\mathrm{quad}(w)$ limits us to consider points w that lie in one of the four sections of the plane determined by u and v. To incorporate this constraint we think of an element v as defining two elements, v and $\bar{v}$, with opposite normals (the positive orientation is set arbitrarily). Consider the saliency map $R'(v) = \int k(u,v)\,\mathrm{d}u$, where the integration is performed over all edge elements u on the positive side of v, that is, elements for which the angle between u–v and the normal to v is less or equal to $\pi/2$. Similarly, $R'(\bar{v})$ is obtained by integrating over all elements u on the negative side of v. We can now express the pairwise saliency map $R_p(w)$ as

$$R_p(w) = \begin{cases} \int k(v,w)g(v)\,\mathrm{d}v & w \text{ on the positive side of } v, \\ \int k(v,w)g(\bar{v})\,\mathrm{d}v & w \text{ on the negative side of } v, \end{cases} \tag{18}$$

where $g(v) = f(v)R'(v)$. Notice that $g(v)$ is the saliency map obtained with the original model (for an arbitrary choice of orientation for v) scaled by $f(v)$. This shows that the pairwise saliency map can be obtained by a two-step application of the original model with the assignment of sides.

Finally, a third model can be constructed by combining the two models, the signed and the pairwise induction fields. In this model we consider only pairs of elements u and v whose gradients point simultaneously either to the direction of their common section or away from this section. This model corresponds to a two-step application of a variation of the signed induction fields.

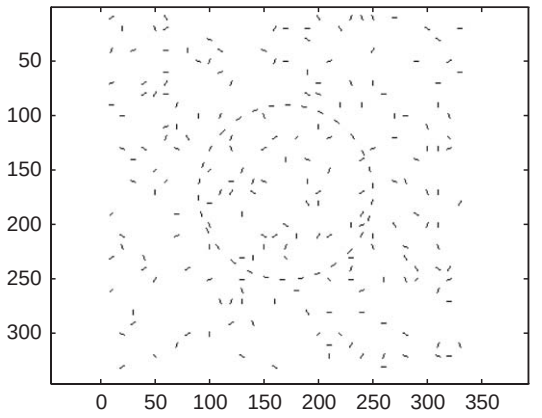

Fig. 5. A dashed circle in a noisy background.

5. Experiments

We have implemented our proposed scheme and applied it to several contour images and gray-level images. In the first experiment we applied the method to a dashed circle in a background of noisy edge elements (Fig. 5, for a comparison of the performance of saliency methods on this type of input see Ref. [8]). Fig. 6 shows the saliency map and the magnitude of the boundary field obtained with our method. It can be seen that the pixels near the center of the circle obtain the highest saliency values, and the saliency values decrease as we move away from the center of the circle. Furthermore, it can be seen that the circle stands out in the boundary field. We also applied the pairwise induction fields to this image. The results are shown in Fig. 7. As can be seen, this method results in a higher contrast between the

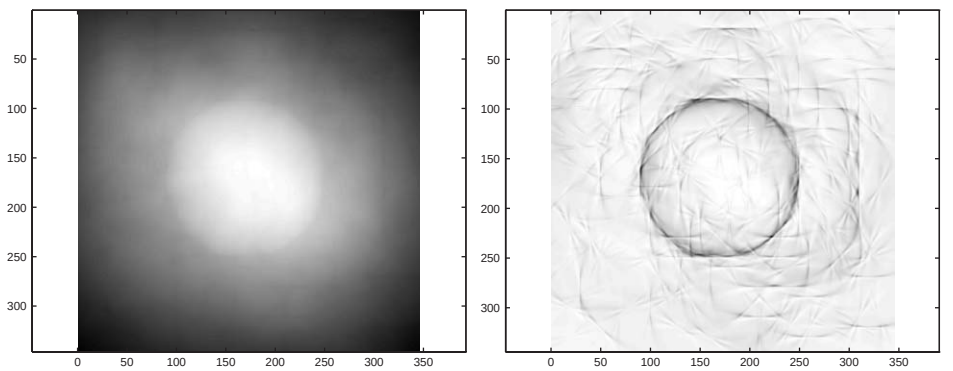

Fig. 6. Applying our method to the circle ($\sigma = 110$). Left: the saliency map. Right: The magnitude of the boundary field (shown in inverse video).

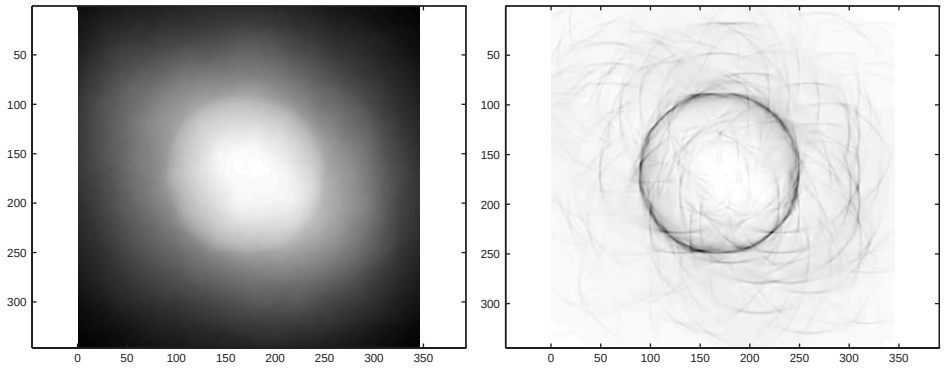

Fig. 7. Applying the pairwise receptive fields to the circle ($\sigma = 110$). Left: the saliency map. Right: the magnitude of the boundary field.

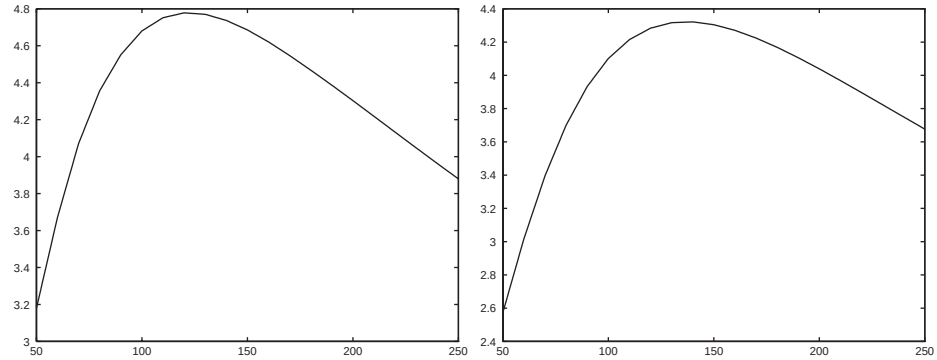

Fig. 8. The maximal (left) and average (right) saliency values as a function of scale obtained inside the circle using the pairwise induction fields.

circle and its surroundings. The circle remains salient in a fairly large range of scales. Fig. 8 shows the change in the maximal and average saliency values of the circle as a function of scale.

In a second experiment we applied the signed induction fields to a binary picture of a club sign (Fig. 9). Figs. 10 and 11 show the saliency maps obtained with two choices of scale parameters. It can be seen that at a smaller scale the convex parts are apparent, whereas at a larger scale the division to parts disappears. The magnitude of the boundary field is shown in Fig. 12.

In the next two experiments we applied the method to gray-level images. We first produced the corresponding gradient maps by applying a 3×3 Sobel operator to the images. Then, we applied the signed induction fields to the gradient maps. As the weight of an edge ($f(u)$ in Eq. (7)) we used the gradient magnitude. The first input image and the magnitude of its gradient are shown in Fig. 13. Fig. 14 shows the saliency map obtained with two scale parameters. The goats stand out in these figures. Fig. 15 shows the respective magnitude of the boundary field.

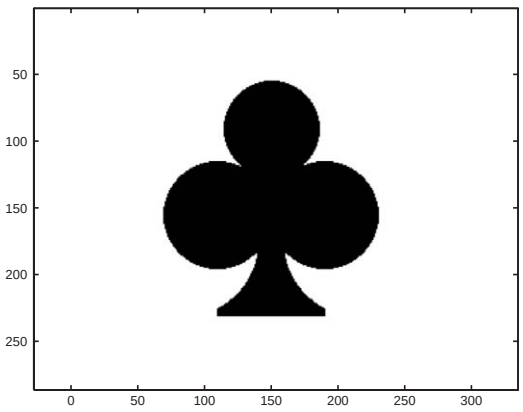

Fig. 9. A club sign.

A second gray-level image and its gradient magnitude are shown in Fig. 16. The saliency map for several choices of scale parameters and the corresponding magnitude of the boundary field are shown in Figs. 17 and 18 respectively.

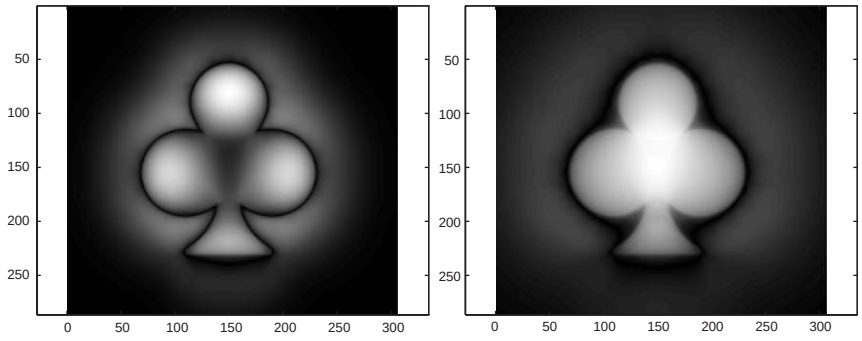

Fig. 10. The saliency maps for $\sigma = 30$ (left) and 60 (right).

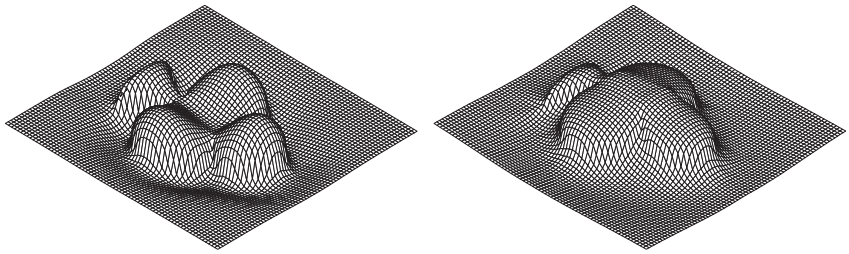

Fig. 11. The saliency maps presented as a mesh for $\sigma = 30$ (left) and 60 (right).

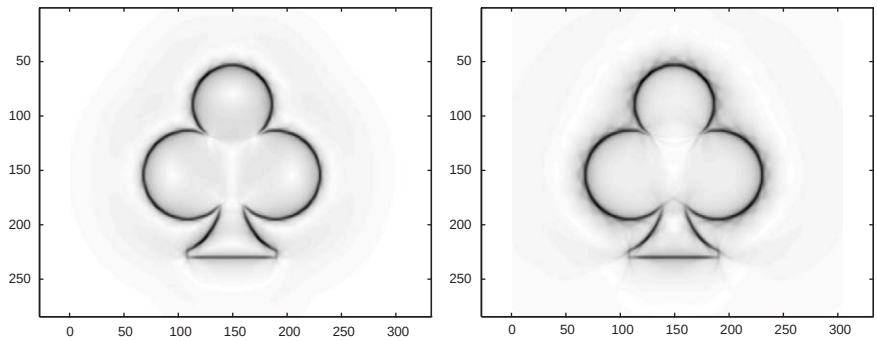

Fig. 12. The magnitude of the boundary field for $\sigma = 30$ (left) and 60 (right).

Next, we used the method to segment the images by simply thresholding the magnitude of the boundary fields. For every region obtained in this process we assigned its saliency to be the average saliency of its points. The results of this segmentation process are shown in Fig. 19. It can be seen that regions corresponding to the circle, the goats, the telephone set, and the monitor were accurately segmented and considered most salient.

A unique property of our method is that it encourages closed, convex regions. To demonstrate this we compared our method to two popular algorithms. In the first example we compare our algorithm with this of Guy and Medioni [6]. Fig. 20 (left) shows a line drawing containing both a closed circle and an open curve. As is apparent from the saliency map, the boundary field, and the segmentation results (Figs. 20 and 21), our method prefers the circle over the open curve. In contrast, Guy and Medioni's algorithm prefers the open curve (Fig. 22).

Finally, we tested our method on the image in Fig. 23 (left). Our results are shown in Figs. 23 and 24. As can be seen, the telephone is the most salient region. In contrast, the saliency map obtained with Shashua and Ullman's algorithm

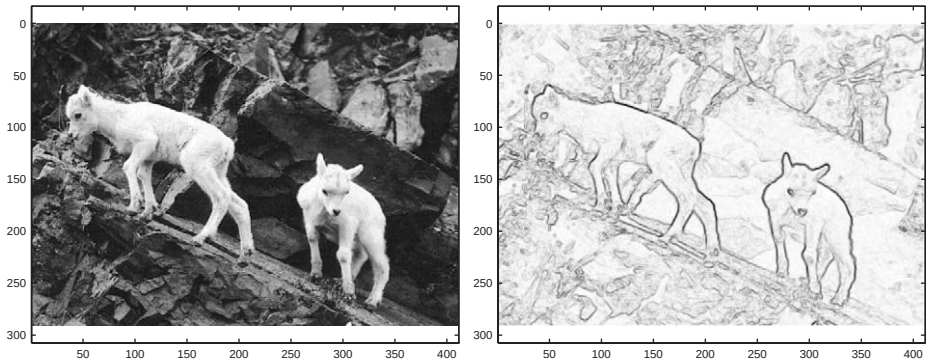

Fig. 13. A gray-level image (left), and the magnitude of its gradient (right).

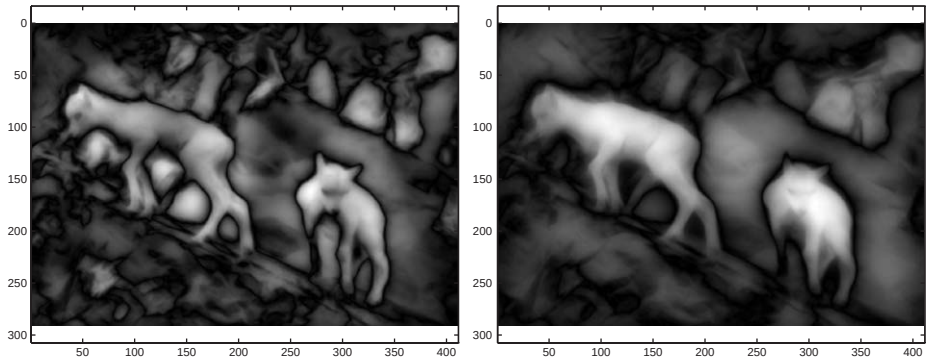

Fig. 14. The saliency map for $\sigma = 20$ (left) and 50 (right).

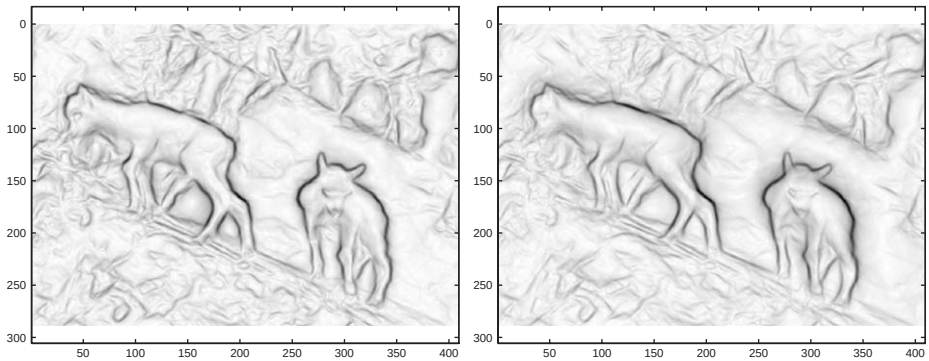

Fig. 15. The magnitude of the boundary field for $\sigma = 20$ (left) and 50 (right).

[9] highlights a curve that contains parts of the stapler in addition to the telephone (Fig. 25).

6. Summary

We have presented a method that uses contour fragments to highlight salient regions in images. The method produces for every point in the image a saliency value reflecting our belief that the point belongs to a salient region. The salient regions can be segmented by considering the boundary field, the gradient field of the saliency map adjusted with respect to scale. The method can be applied both to contour and to gray-level images. We have presented experiments with both contour and gray-level images which demonstrate the perfor-

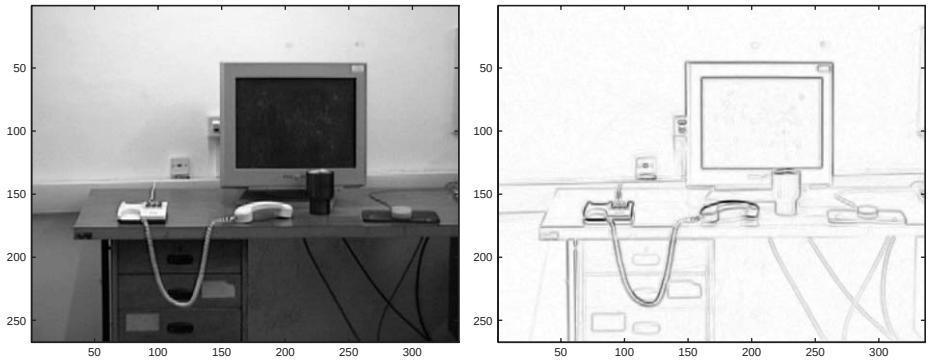

Fig. 16. A gray-level image (left), and the magnitude of its gradient (right).

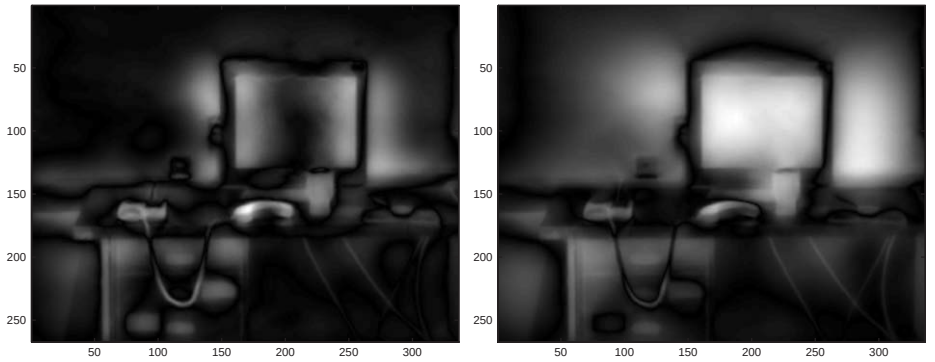

Fig. 17. The saliency map for $\sigma = 20$ (left) and 40 (right).

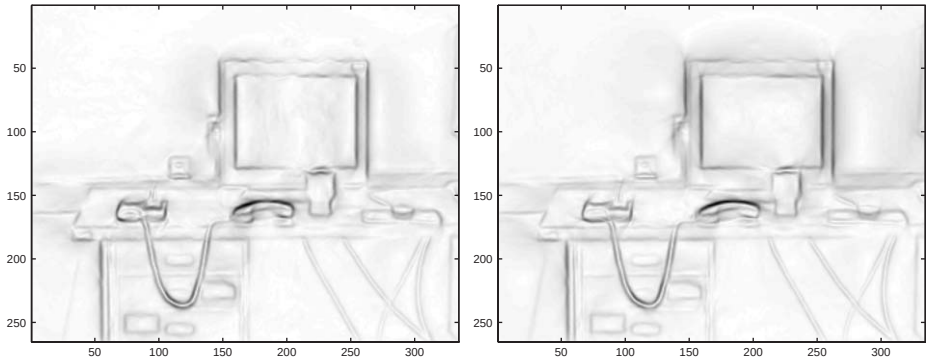

Fig. 18. The magnitude of the boundary field for $\sigma = 20$ (left) and 40 (right).

mance of the method and its utility for saliency and segmentation.

Our method is unique since, unlike most grouping methods that consider curve fragments, it uses the fragments to highlight *regions* of interest. This is achieved using criteria that include closure, convexity, and size. In contrast, most other methods prefer an open, straight line over a closed curve. The method is implemented by convolving the edge image with a linear filter that specifies the region of influence of a contour point over the image. Note that when the scale parameter is large a straightforward application of the scheme will lead to a runtime complexity that is quadratic in the size of the image. Nevertheless, by exploiting the smoothness of the filter we expect to be

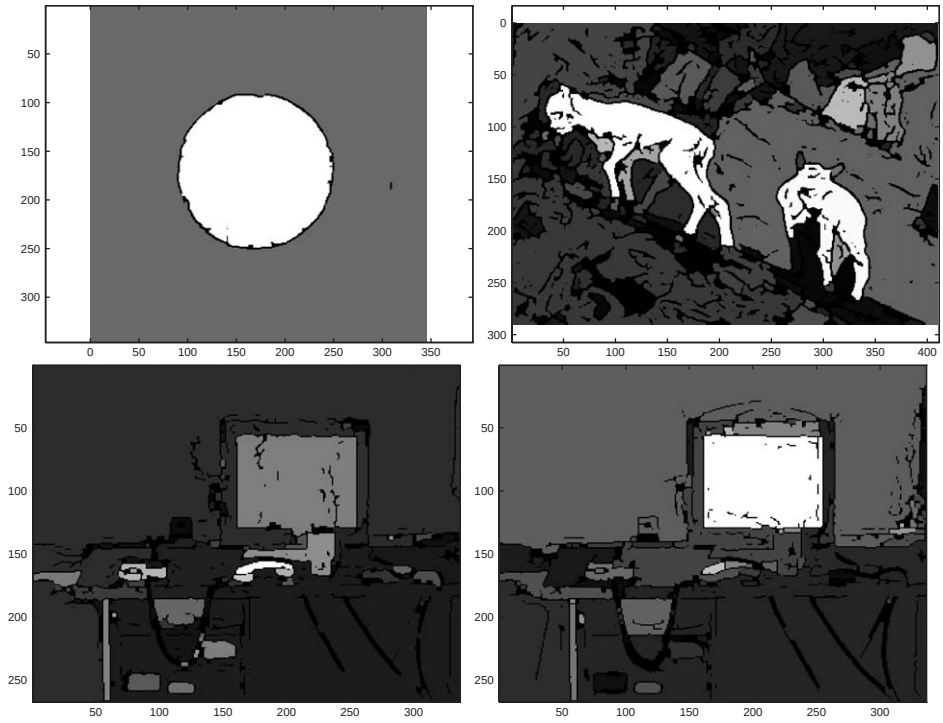

Fig. 19. Image segmentation using the saliency map obtained in Fig. 7 (top left), Fig. 14 (top right), and Fig. 17 (bottom left, $\sigma = 20$, bottom right $\sigma = 40$). The brightness of a region is proportional to its average saliency.

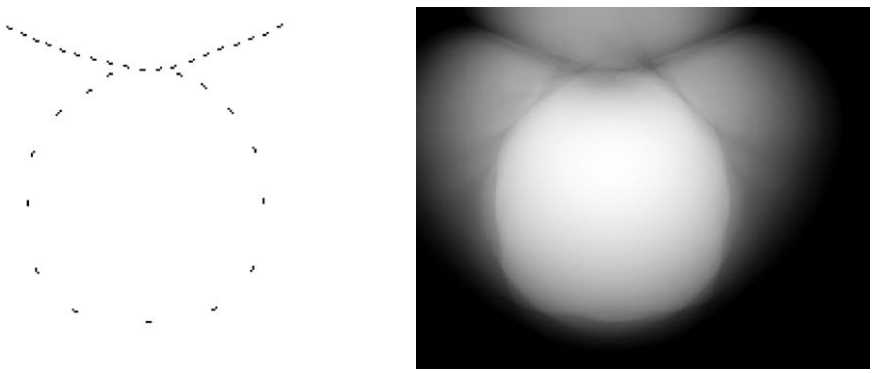

Fig. 20. An image containing a dashed circle and a dashed open curve (left) and its saliency map ($\sigma = 80$, right).

able to reduce this runtime complexity considerably (see Ref. [7]).

Appendix A. Normalization

In this appendix we show that

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} k(u,v)\,\mathrm{d}x\,\mathrm{d}y = 1, \tag{A.1}$$

where

$$k(u,v) = \frac{1}{\pi\sigma^2}\exp\left(-\frac{h^2}{2\sigma^2}\right), \tag{A.2}$$

and

$$h = \frac{\sqrt{x^2+y^2}}{|\cos\phi|}. \tag{A.3}$$

We first show that

$$\int_0^{\infty} \exp\left(-\frac{\rho^2}{2\sigma^2\cos^2\phi}\right)\rho\,\mathrm{d}\rho = \sigma^2\cos^2\phi. \tag{A.4}$$

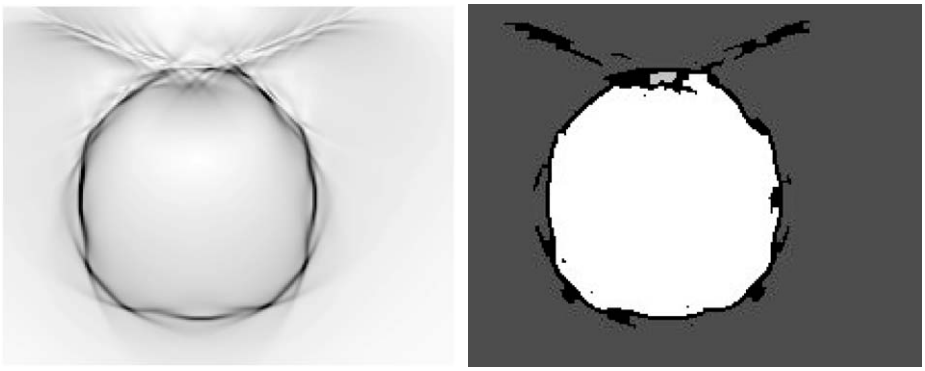

Fig. 21. The magnitude of the boundary field (left) and the segmentation obtained from this boundary field (right).

Fig. 22. Saliency map obtained using Guy and Medioni's method (with $\sigma = 80$, software downloaded from http://iris.usc.edu/tensorvt). Darker pixels represent more salient locations. Note that the open curve is more salient than the circle.

From normal distributions we know that

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \exp\left(-\frac{x^2+y^2}{2\bar{\sigma}^2}\right) \mathrm{d}x\,\mathrm{d}y = 2\pi\bar{\sigma}^2. \tag{A.5}$$

In polar coordinates this becomes

$$\int_{0}^{2\pi}\int_{0}^{\infty} \exp\left(-\frac{\rho}{2\bar{\sigma}^2}\right) \rho\,\mathrm{d}\rho\,\mathrm{d}\theta = 2\pi\bar{\sigma}^2. \tag{A.6}$$

Since the integrated function is independent of θ we get that

$$\int_{0}^{\infty} \exp\left(-\frac{\rho^2}{2\bar{\sigma}^2}\right) \rho\,\mathrm{d}\rho = \bar{\sigma}^2. \tag{A.7}$$

Eq. (A.4) is obtained by substituting $\sigma\cos\phi$ for $\bar{\sigma}$. Using Eq. (A.4) we obtain

$$\int_{0}^{2\pi}\int_{0}^{\infty} \exp\left(-\frac{\rho^2}{2\sigma^2\cos^2\phi}\right) \rho\,\mathrm{d}\rho\,\mathrm{d}\phi$$

$$= \sigma^2 \int_{0}^{2\pi} \cos^2\phi\,\mathrm{d}\phi = \pi\sigma^2, \tag{A.8}$$

and this implies Eq. (A.1).

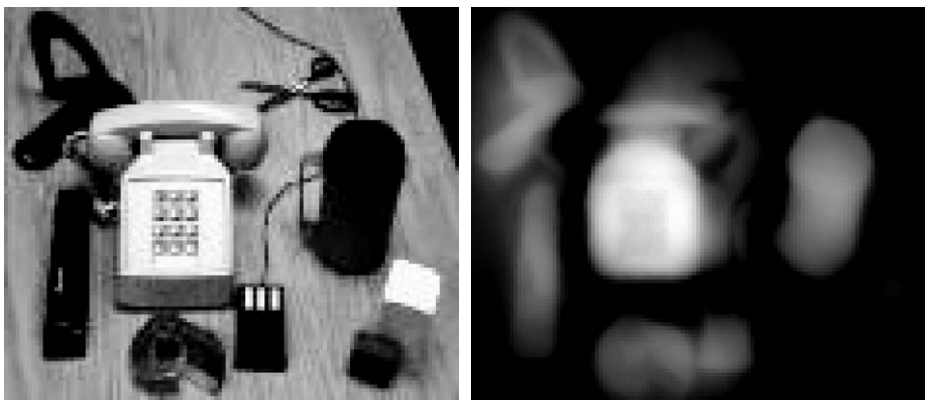

Fig. 23. An image (left, from Ref. [10], used with permission) and its saliency map ($\sigma = 30$, right).

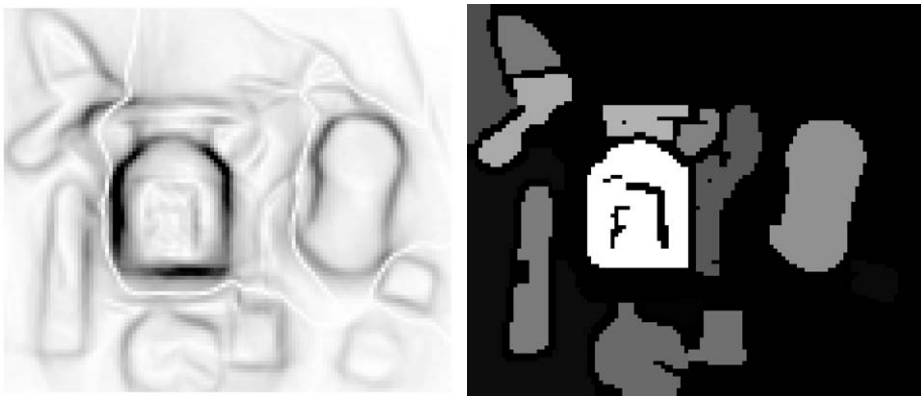

Fig. 24. The magnitude of the boundary field (left) and the segmentation obtained from this boundary field (right).

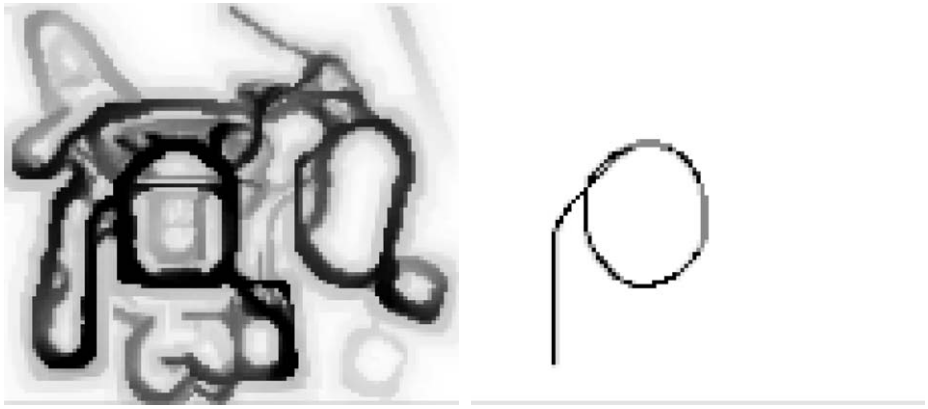

Fig. 25. Saliency map obtained using Shashua and Ullman's method (left) and the most salient curve (right). Darker pixels represent more salient locations. Note that the most salient curve includes parts of the telephone and the stapler (from Ref. [10], used with permission).

References

[1] S. Ullman, Filling-in the gaps: the shape of subjective contours and a model for their generation, Biol. Cybernet. 25 (1976) 1–6.

[2] B.K.P. Horn, The curve of least energy, ACM Trans. Math. Software 9 (4) (1983) 441–460.

[3] A.M. Bruckstein, A.N. Netravali, On minimal energy trajectories, Comput. Vision Graphics Image Process. 49 (1990) 283–296.

[4] D. Mumford, Elastica and computer vision, in: C. Bajaj (Ed.), Algebraic Geometry and Its Applications, Springer, Berlin, 1994, pp. 491–506.

[5] L.R. Williams, D.W. Jacobs, Stochastic completion fields: a neural model of illusory contour shape and salience, Neural Comput. 9 (4) (1997) 837–858.

[6] G. Guy, G. Medioni, Inferring global perceptual contours from local features, Int. J. Comput. Vision 20 (1/2) (1996) 113–133.

[7] E. Sharon, A. Brandt, R. Basri, Completion energies and scale, IEEE Trans. Pattern Anal. Mach. Intell. 22 (10) (2000) 1117–1131.

[8] L.R. Williams, K.K. Thornber, A comparison in measures for detecting natural shapes in cluttered backgrounds, International Journal of Computer Vision 34 (2/3), August 1999, pp. 1–16.

[9] A. Sha'ashua, S. Ullman, Structural saliency: the detection of globally salient structures using a locally connected network, in: Second International Conference on Computer Vision (ICCV), Tarpon Springs, FL, December 5–8, 1988, pp. 321–327.

[10] T.D. Alter, R. Basri, Extracting salient curves from images: an analysis of the saliency network, Int. J. Comput. Vision 27 (1) (1998) 51–69.

[11] D. Huttenlocher, P. Wayner, Finding convex edge grouping in an image, Int. J. Comput. Vision 8 (1) (1992) 7–29.

[12] D. Jacobs, Robust and efficient detection of convex groups, IEEE Trans. Pattern Anal. Mach. Intell. 18 (1) (1996) 23–37.

[13] J.H. Elder, S.W. Zucker, Computing contour closure, in: B. Buxton and R. Cipolla (Eds.) Fourth European Conference on Computer Vision (ECCV), Springer 1996, Vol. 1, pp. 399–412.

[14] T. Leung, J. Malik, Contour continuity in region based image segmentation, in: Fifth European Conference on Computer Vision (ECCV), 1998, pp. 544–559.
[15] I.H. Jermyn, H. Ishikawa, Globally optimal regions and boundaries, as Minimum Ratio Weight Cycles, IEEE Trans. on Pattern Analysis and Machine Intelligence, 23 (10) (2001) 1075–1088.
[16] D. Geiger, K. Kumaran, L. Parida, A computational view of visual organization for figure/ground separation, in: IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), San Francisco, Ca, June 18–20, 1996, pp. 155–160.
[17] K. Kumaran, D. Geiger, L. Gurvits, Illusory surfaces and visual organization, Network: Computation in Neural Systems, 7 (1) (1998) 36–60.
[18] D. Geiger, H. Pao, N. Rubin, Salient and multiple illusory surfaces, in: IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), San Francisco, Ca, June 23–25, 1998, pp. 118–124.

About the Author—YOSSI COHEN received his B.Sc. degree in Computer Science and Mathematics from Ben Gurion University, Beer Sheba, Israel in 1994 and M.Sc. degree in Computer Science and Mathematics from the Weizmann Institute of Science, Rehovot, Israel in 1997. He is presently working in the Hi-Tech industry as a software team leader.

About the Author—RONEN BASRI received the B.Sc. degree in Mathematics and Computer Science from Tel Aviv University in 1985, where he graduated Summa Cum Laude. He received the Ph.D. degree in Computer Science from the Weizmann Institute of Science in 1990. From 1990 to 1992 he has been a post-doctoral fellow at the Massachusetts Institute of Technology in the Department of Brain and Cognitive Science and the Artificial Intelligence Laboratory under the McDonnell-Pew and the Rothchild programs. Since then he has been affiliated with the Weizmann Institute of Science in the Department of Computer Science and Applied Mathematics, where he is currently holding the position of Associate Professor. Between 1999 and 2000 he spent a sabbatical at NEC Research Institute in Princeton, New Jersey. Ronen Basri's research has focused on computer vision, especially in the areas of object recognition, perceptual organization, and visually guided navigation.