

Globally Optimal Regions and Boundaries

Ian H. Jermyn and Hiroshi Ishikawa
Courant Institute of Mathematical Science
New York University, New York, NY, U.S.A.
jermyn@cs.nyu.edu, ishikawa@cs.nyu.edu

Abstract

We propose a new form of energy functional for the segmentation of regions in images, and an efficient method for finding its global optima. The energy can have contributions from both the region and its boundary, thus combining the best features of region- and boundary-based approaches to segmentation. By transforming the region energy into a boundary energy, we can treat both contributions on an equal footing, and solve the global optimization problem as a minimum mean weight cycle problem on a directed graph. The simple, polynomial-time algorithm requires no initialization and is highly parallelizable.

1. Introduction

Image segmentation methods generally fall into two classes, being either region-based or boundary-based. The former class uses properties of areas of the image to choose among possible segmentations, while the latter looks at the properties of the image only on the boundary of the proposed segmented regions.

Both methods have their advantages and drawbacks. Region-based methods tend to be global, optimizing a functional of the image segmentation. On the other hand, they often ignore important boundary properties such as smoothness. Boundary-based approaches can treat such properties very naturally, but suffer from their own difficulties. First, most algorithms find only local minima, and thus have no measure of the significance of the extracted boundary for the image as a whole. Second, although there do exist algorithms guaranteed to find global minima, using graph techniques such as dynamic programming and Dijkstra's algorithm, these do not adapt easily to closed contours. Unfortunately, open contours do not segment regions in the image, so that further processing is needed to group the contours into proto-surfaces. Third, boundary-based methods cannot incorporate region information such as texture easily. In addition, many of the extant algorithms

require initialization by the user in some way, by specifying the end-points of the contour, or by defining an initial contour that then evolves to a solution.

In this paper we propose a new form of energy functional for the segmentation of regions in images, and an efficient method for finding its global optima in energy order. The energy is of a very general form (equation (1)), being the modulus of the integral over the region of any integrable function, divided by a measure of the length of the boundary of the region. The solution to the optimization problem is the global maximum of this energy over all regions. The region integral can be transformed to a boundary integral, and then combined with boundary-dependent terms. In this way the energy can have contributions from both the region and its boundary, allowing region information such as texture and homogeneity to be combined with boundary information such as intensity gradients. Once expressed as a boundary integral, we can cast the global optimization problem into the form of a minimum mean weight cycle problem in a digraph. This problem has a simple, polynomial-time algorithm that requires no initialization, and is highly parallelizable, with each pixel able to perform its computations independently, reading from, but never writing to, its neighbours.

The paper is laid out as follows. In the next section, we discuss related work. In section 3, we describe the general form of the energy functional and its properties, and give some examples of possible uses. We discuss the algorithm that globally optimizes such energies and its relation to our problem in section 4. In section 5 we describe some specific models of regions, and show the results of experiments with these models.

2. Related work

Image segmentation has a huge literature, and here we only touch on some of the work more closely related to ours.

Contour-based grouping methods include Parent and Zucker's [20] work using relaxation methods, Sha'ashua and Ullman's [21] work on saliency networks, and Guy

and Medioni’s [10] work using voting schemes. There is also the work of Cox et al. on a Bayesian sequential tracking scheme [4]. Elder and Zucker [8] have developed a method for finding closed contours using chains of tangent vectors. Williams and Thornber [25, 23] and Williams and Jacobs [24] discuss contour closure using stochastic completion fields. Closest to our work however, because it starts from an energy optimization criterion, is the work on active contours. The seminal works in this area are Kass et al. [14] and Blake and Zisserman [2], and much subsequent work follows this both in the form of the energy functionals used, and in algorithmic techniques. Another body of work applies dynamic programming techniques to minimize the contour energy. Amini et al. [1] use dynamic programming as part of a gradient descent procedure. Montanari [19] uses dynamic programming to find the minimum energy path between given end-points. Geiger et al. [9] use initialization with a series of points, and a choice of window around those points, to delineate the space of contours considered. Much of this work uses initialization and restricted regions of the image to limit the space of contours over which the optimization proceeds, and most algorithms find local minima, or approximations to global minima over a limited set of contours. Globally minimum closed contours are not found.

The paper by Cox, Rao, and Zhong [3] is particularly related to our work. They use a graph algorithm known as the pinned ratio algorithm to find closed contours in an image. The method can be made initialization-free, and finds a global minimum under some weak constraints. Their method is not as general as ours however, as they cannot combine region and boundary information, and the region information they use must be positive everywhere in the image.

Related work in the area of region-based segmentation is that of Shi and Malik [22]. They use a generalized eigenvalue method to find normalized cuts. The denominator in our equation (1) plays a similar role to the cut normalization. Leung and Malik [17] extend their work by incorporating weak contour continuity information into the region-based model.

Psychological work has emphasized the importance of closure in perception since the Gestalt movement. Work in illusory contours has also shown the importance of the Gestalt concept of closure to the perceptual organization involved in these phenomena [11, 12]. More recent work by Kovács and Julesz, and Elder and Zucker has demonstrated that closure is a very important determinant of contour saliency [15, 6, 7].

3. Theoretical framework

An image is a real-valued function I on a domain in \mathbb{R}^2 . A simple region is denoted \mathcal{R} and its boundary $\partial\mathcal{R}$. A

Gaussian is denoted G , and convolution $*$.

3.1. Combining regions and boundaries

The form of energy functional with which we deal is

$$E(I, \mathcal{R}) = \frac{\left| \int_{\mathcal{R}} f \, dx \, dy \right|}{\int_{\partial\mathcal{R}} g \, ds} \quad (1)$$

where s is the arc length parameter, f is any real-valued function on \mathbb{R}^2 , and g is any positive real-valued function on $\partial\mathcal{R}$. We define the solution to the optimization problem as the global maximum of E over all regions \mathcal{R} . For reasons that will become clear, we note that by assigning two energies to each region, $\pm E(I, \mathcal{R})$, and then minimizing over all such energies, we achieve the same solution. This can be viewed as an assignment of two orientations to each region, and then a minimization over all oriented regions. The denominator is a (possibly weighted) measure of the length of the boundary, and has the effect of damping the scaling behaviour of the energy, which would otherwise have a strong preference for large regions. It also functions as a boundary smoothing term, as follows. If f and g were unity we would be maximizing the area over the length, and fixing the length of the boundary would produce a disc as the solution to the optimization problem. This is also the solution to an active contour model with a fixed length and a smoothing term that is the square of the curvature. In general, the effect of the dependence on area divided by that on length will be to produce smoother boundaries.

The numerator of equation (1) can always be rewritten as an integral over the boundary $\partial\mathcal{R}$ of the region \mathcal{R} :

$$E(I, \mathcal{R}) = \frac{\left| \int_{\partial\mathcal{R}} \hat{n} \cdot \vec{A} \, ds \right|}{\int_{\partial\mathcal{R}} g \, ds} \quad (2)$$

where \hat{n} is the normal vector to the boundary, and \vec{A} is defined by the equation $\vec{\nabla} \cdot \vec{A} = f$. Such an \vec{A} always exists. It can for instance be given by the following integrals:

$$\begin{aligned} \vec{A}_x(x, y) &= \frac{1}{2} \int_0^x f(x', y) dx' \\ \vec{A}_y(x, y) &= \frac{1}{2} \int_0^y f(x, y') dy' \end{aligned}$$

There is a choice of constant functions that can be added to this vector field, but the choice does not affect the value of the boundary integral in equation (2). Indeed, we can add any divergence-free vector field and still have the same boundary integral.

Similarly to equation (1), we can view the boundary as having two possible orientations corresponding to the bounding curve running in the two possible directions. Removing the modulus signs and minimizing over all oriented boundaries is then equivalent to the original maximization problem over un-oriented regions.

Remarks

- As advertized, the form of equation (2) allows us elegantly to include boundary as well as region information in our model. Indeed the present work shows that they are essentially the same, although one description may be more appropriate than the other. We can add to \vec{A} any other vector field \vec{B} and still compute the global optimum.
- Averaging the weight of the boundary over a measure of its length has at least two important advantages over unaveraged contour models. First, it removes the uncontrollable dependence on contour length that such models inevitably exhibit. This is most noticeable if the energy has a gradient term and a length term for example. The length term is normally positive, while the gradient term is negative. Depending on the parameters, the global solution could be trivial in one of two ways: infinitely long or infinitely short. These are extreme examples, but the implicit dependence on length is always present. Second, a simple polynomial-time algorithm exists for the minimum mean weight cycle problem, whereas the minimum weight cycle problem is NP-hard.

3.2. Forms of region function

The function f in equation (1) can be any integrable function. In particular, it can be the convolution of the image with any filter F : $f(p) = I * F(p)$. In this case, equation (2) also takes the form of a convolution:

$$E(I, \mathcal{R}) = \frac{\int_{\mathcal{R}} (I * F) dx dy}{\int_{\partial\mathcal{R}} g ds}$$

$$E(I, \partial\mathcal{R}) = \frac{\int_{\partial\mathcal{R}} \hat{n} \cdot (I * \vec{v}) ds}{\int_{\partial\mathcal{R}} g ds}$$

where $\vec{\nabla} \cdot \vec{v} = F$.

The function f need not be a linear filter however. Some choices of f and their meaning in our model are given below. Throughout we include the possibility of a Gaussian smoothing of the image, or in other words the possibility of examining the image at various scales. The examples are intended to include the case of zero variance, when G is a delta function.

$I * G$. In this case the model is looking for globally maximum intensity regions. It will find bright spots such as specular reflections, as well as large regions of high intensity.

$I * \vec{\nabla}^2 G$. Viewed as a region integral, this function finds the region with the largest absolute value of the integrated Laplacian of the smoothed intensity. Such

regions correspond to ‘‘lumps’’ or ‘‘dips’’ in the intensity function, since regions with undulations in the intensity will make both positive and negative contributions to the region integral, reducing its absolute value. Converted to the form of equation (2), $\vec{A} = I * \vec{\nabla} G$. \vec{A} is thus the vector field of wavelet coefficients. Viewed in this way, the model finds regions whose boundaries pass through points with high smoothed intensity gradient, in a direction perpendicular to the gradient. The model averages over length, thus removing the dependence on scale, and our algorithm finds the globally optimal region and boundary. We investigate such a model in section 5.

$|I * \vec{\nabla}^2 G|$. The previous function does not deal well with the case of contrast-reversing boundaries, which introduce both positive and negative contributions to the region integral. We can deal with the case of general boundaries (including contrast-reversing) using the absolute value of the Laplacian, at the expense of losing a simple boundary interpretation. This region function is a better way to deal with contrast-reversing boundaries than the normal method of taking the magnitude of the gradient, since it preserves the notion that intensity change should be normal to the boundary.

$|I * \vec{\nabla} G|^{-1}$. f is a positive, monotonically decreasing function of the magnitude of the wavelet coefficients $|I * \vec{\nabla} G|$, such as $1/|I * \vec{\nabla} G|$ or $M - |I * \vec{\nabla} G|$ (M is an upper bound on the magnitude of the wavelet coefficients in the image). The integral of such a term over a region will be large if the gradient has small magnitude everywhere. It will thus seek out the region with globally most homogeneous intensity.

$I * T$. A filter T (or linear combination of filters) that responds strongly to a particular class of textures can be used to segment globally optimal regions of that texture.

Most interestingly, the function can be a combination of these examples, so that we could search for the region with the best response to a given texture and that had a high intensity gradient boundary for example, or that had a homogeneous intensity surrounded by a high intensity gradient boundary.

Before passing to a description of the algorithm that we use to solve the global optimization problem, we make two observations about the form of energy functional.

As mentioned, equation (2) has an interesting invariance. If we add to \vec{A} any vector field with zero divergence, we can see by transforming to the form of equation (1) that the energy will not change. When $f = I * \vec{\nabla}^2 G$, this corresponds to adding a harmonic function to the intensity.

Equation (1) is the most general form of energy that we can optimize globally at present (although see section 6). In

the experiments we use a slightly restricted form, in which the function g is unity and the length is approximated by an edge count. The algorithms to find the global maximum of the more general case [16, 5, 18] are more complex than the one we describe in section 4. For the sake of clarity, we restrict ourselves to this case.

4. Algorithmic solution

To find the global maximum of the energy in equation (2) (or equivalently equation (1)) we use a graph algorithm due to Richard Karp [13]. This algorithm finds the minimum mean weight cycle in a directed graph.

The algorithm requires no initialization by the user. It is also highly parallelizable, with each pixel able to perform its computations independently, reading from, but never writing to, its neighbours. In the experiment, we iteratively apply the algorithm. After each iteration, we remove from the graph those vertices through which the previous solution passed. We thus find a series of regions of increasing energy, which can be viewed as a series of hypotheses about regions in the image of gradually decreasing probability.

We first describe the algorithm and then clarify its relation to our problem.

4.1. Algorithm

We begin with a weighted directed graph G , with weight function w . We wish to find the minimum mean weight simple cycle, where the mean weight of an edge progression composed of edges $\{e_i : i \in \mathcal{I}\}$ is defined as $\frac{\sum_{i \in \mathcal{I}} w(e_i)}{|\mathcal{I}|}$.

First, define the function $F_k(v)$ taking each vertex $v \in V$ (V is the vertex set) to the weight of the minimum weight path of length $k \geq 0$ to v from an arbitrary start vertex s , and define it to be ∞ if no path exists of length k . Then it can be shown (proof is given in [13]) that the weight λ^* of the minimum mean weight cycle is given by

$$\lambda^* = \min_{v \in V} \max_{k \in [0..(n-1)]} \left\{ \frac{F_n(v) - F_k(v)}{n - k} \right\} \quad (3)$$

where $n = |V|$.

$F_k(v)$ can be computed using the recurrence

$$\begin{aligned} F_k(v) &= \min_{(u,v) \in E} F_{k-1}(u) + w((u,v)) \\ F_0(s) &= 0 \\ F_0(v) &= \infty, \forall v \neq s \end{aligned}$$

where E is the edge set of G . With in

The computation of F for all $k \in [0..(n-1)]$ can be performed using dynamic programming in time $O(n|E|)$.

The minimum weight paths can be computed simultaneously. Using a further $O(n^2)$ time we can compute λ^* from $F_k(v)$, leading to an overall computation time of $O(n|E|)$. The cycle itself can be extracted by selecting the minimizing v and k in equation (3), and finding a cycle of length $n - k$ in the minimum weight path from s to v .

4.2. Application

We recall that, as discussed in section 3, if we remove the modulus signs from equations (1) and (2) and view the region and boundary as having two possible orientations, minimizing over all oriented regions or boundaries is equivalent to maximizing equations (1) and (2) over all un-oriented regions or boundaries.

To cast our problem in the form of a minimum mean weight cycle problem, we embed a directed graph in the image, with the property that for every two vertices, u and v , if (u, v) is an edge then (v, u) is also. Thus each cycle can have two possible orientations. The embedding η takes each vertex v to a point $\eta(v)$, and each edge $e = (s, t)$ to a tangent vector $\eta(e)$ located at the median point of $\eta(s)$ and $\eta(t)$, and directed from $\eta(s)$ to $\eta(t)$. The unit normal vectors $\hat{n}(e)$ required by equation (2) can then be defined from the tangent vectors by a fixed rotation. A region boundary is then by definition a simple cycle in this graph.

The weight of an edge $e = (u, v)$ is defined as $\Delta s \{ \hat{n}(e) \cdot \vec{A} \}$, where the vector field \vec{A} is evaluated at the midpoint of the edge, and so lies in the same tangent space as $\hat{n}(e)$. Δs is the Euclidean distance between $\eta(u)$ and $\eta(v)$, and plays the role of the measure ds in equation (2). Note that because $\hat{n}(e)$ is defined using a fixed orientation from the tangent vectors, the weights of edges between the same two points but in opposite directions will have the same absolute value but opposite sign. This ensures that the weights of cycles that differ only in orientation will have the same absolute value but opposite sign, as required when we remove the modulus signs and minimize over oriented boundaries in equation (2). Summing the edge weights so defined over a cycle in the graph then gives a discrete version of the numerator in equation (2).

We can now apply the minimum mean weight cycle algorithm to find the solution to our problem on this discrete domain. If the graph is dense enough in the image plane, we will have a good solution to the continuous problem.

5. Experiments

For the bulk of the experiments, we chose $f = I * \vec{\nabla}^2 G$ in equation (1). Thus we are finding regions over which the absolute value of the integral of the Laplacian is as large as possible. These correspond to ‘‘lumps’’ or ‘‘dips’’ in the intensity function, since regions with undulations in the intensity will make both positive and negative contributions

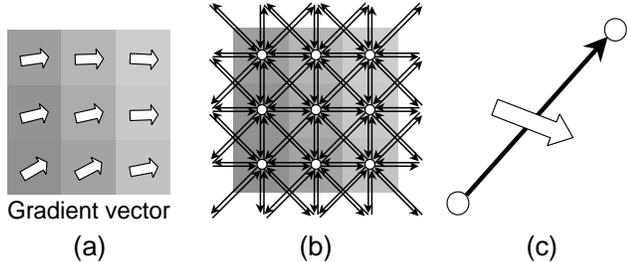


Figure 1. (a) For each pixel we compute the gradient vector. (b) The graph has a node for each pixel and eight outgoing edges for each node (except at the boundary.) (c) The edge weight is calculated by taking a cross product of the gradient vector and the edge vector.

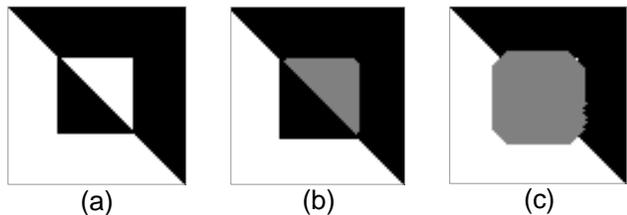


Figure 2. (a) A synthetic contrast-reversing boundary. (b) The result of applying the region energy $I * \vec{\nabla}^2 G$. (c) The result of applying the region energy $|I * \vec{\nabla}^2 G|$. The region found is shown in grey.

to the region integral, reducing its absolute value. With this choice of f , the vector field \vec{A} in equation (2) becomes $I * \vec{\nabla} G$, the wavelet coefficients of the image at a scale dictated by the width of G . (We took the width to be small, of the order of a few pixels, so that we are dealing with a very slightly blurred estimate of the gradient at pixel scale.) As a boundary integral, the energy becomes

$$E(I, \partial\mathcal{R}) = \frac{\int_{\partial\mathcal{R}} \hat{n} \cdot (I * \vec{\nabla} G) ds}{\int_{\partial\mathcal{R}} ds} \quad (4)$$

The integrand in equation (4) is minimal when the boundary tangent vector is perpendicular to the intensity gradient.

To apply the algorithm of section 4, we used a directed graph with an eight-valent node for each pixel (Figure 1). For each node, we computed the gradient vector at the pixel

by taking a wavelet coefficient:

$$\begin{aligned} \vec{A} &= (I * \vec{\psi}_s)(\vec{x}) \\ \vec{\psi}_s(\vec{x}) &= s^{-1} \vec{\psi}(s^{-1} \vec{x}) \\ \vec{\psi}(\vec{x}) &= \pi^{-1} e^{-|\vec{x}|^2} \vec{x}, \end{aligned}$$

where “*” denotes a convolution and $\vec{\psi}$ is the derivative of a Gaussian. For an edge going from node $n_{\vec{u}}$ to node $n_{\vec{v}}$ corresponding to pixels \vec{u} and \vec{v} , the edge weight is computed as

$$(\vec{v} - \vec{u}) \times \frac{\vec{A}(\vec{u}) + \vec{A}(\vec{v})}{2}.$$

This is the cross product of the tangent vector with length equal to the Euclidean distance between the nodes with the wavelet coefficients. This is the same as taking the dot product of the coefficients with the appropriately oriented normal vector, and weighting by the Euclidean distance between the nodes.

We iteratively applied the algorithm explained in section 4. After each iteration, we removed from the graph those vertices through which the previous solution passed. In this way a series of regions of increasing energy was extracted. This can be viewed as a series of hypotheses about regions in the image of gradually decreasing probability. Results are shown in Figure 3. The numbers indicate the order in which the regions were found. The finding of specularities (and their reverse - dark spots) by the method is to be expected as these are isolated peaks or troughs in the image. Although potentially interesting, these tiny regions can be eliminated by the addition of a term that favors larger areas (for example homogeneity).

In order to illustrate that the method can also deal with the case of contrast-reversing boundaries, we took a synthetic image and used the region function $|I * \vec{\nabla}^2 G|$. The results are shown in Figure 2. They illustrate that replacing the Laplacian by its absolute value finds contrast-reversing boundaries. The expansion of the region found beyond the contrast-reversing boundary is a consequence of using an edge count instead of the geometrical length. It results in multiple degenerate solutions, one of which is illustrated. Use of the more sophisticated algorithm mentioned at the end of section 3.2 would break the degeneracy and pick out the (correct) solution of minimum length.

We note that the region functions used in the experiments have no parameters. The scale at which we compute the gradient is variable, but we chose a small pixel-size scale beforehand, and stayed with it throughout the experiments.

6. Conclusion

In this paper we proposed a new form of energy functional for the segmentation of regions in images, and an efficient method for finding its global optima in energy order.

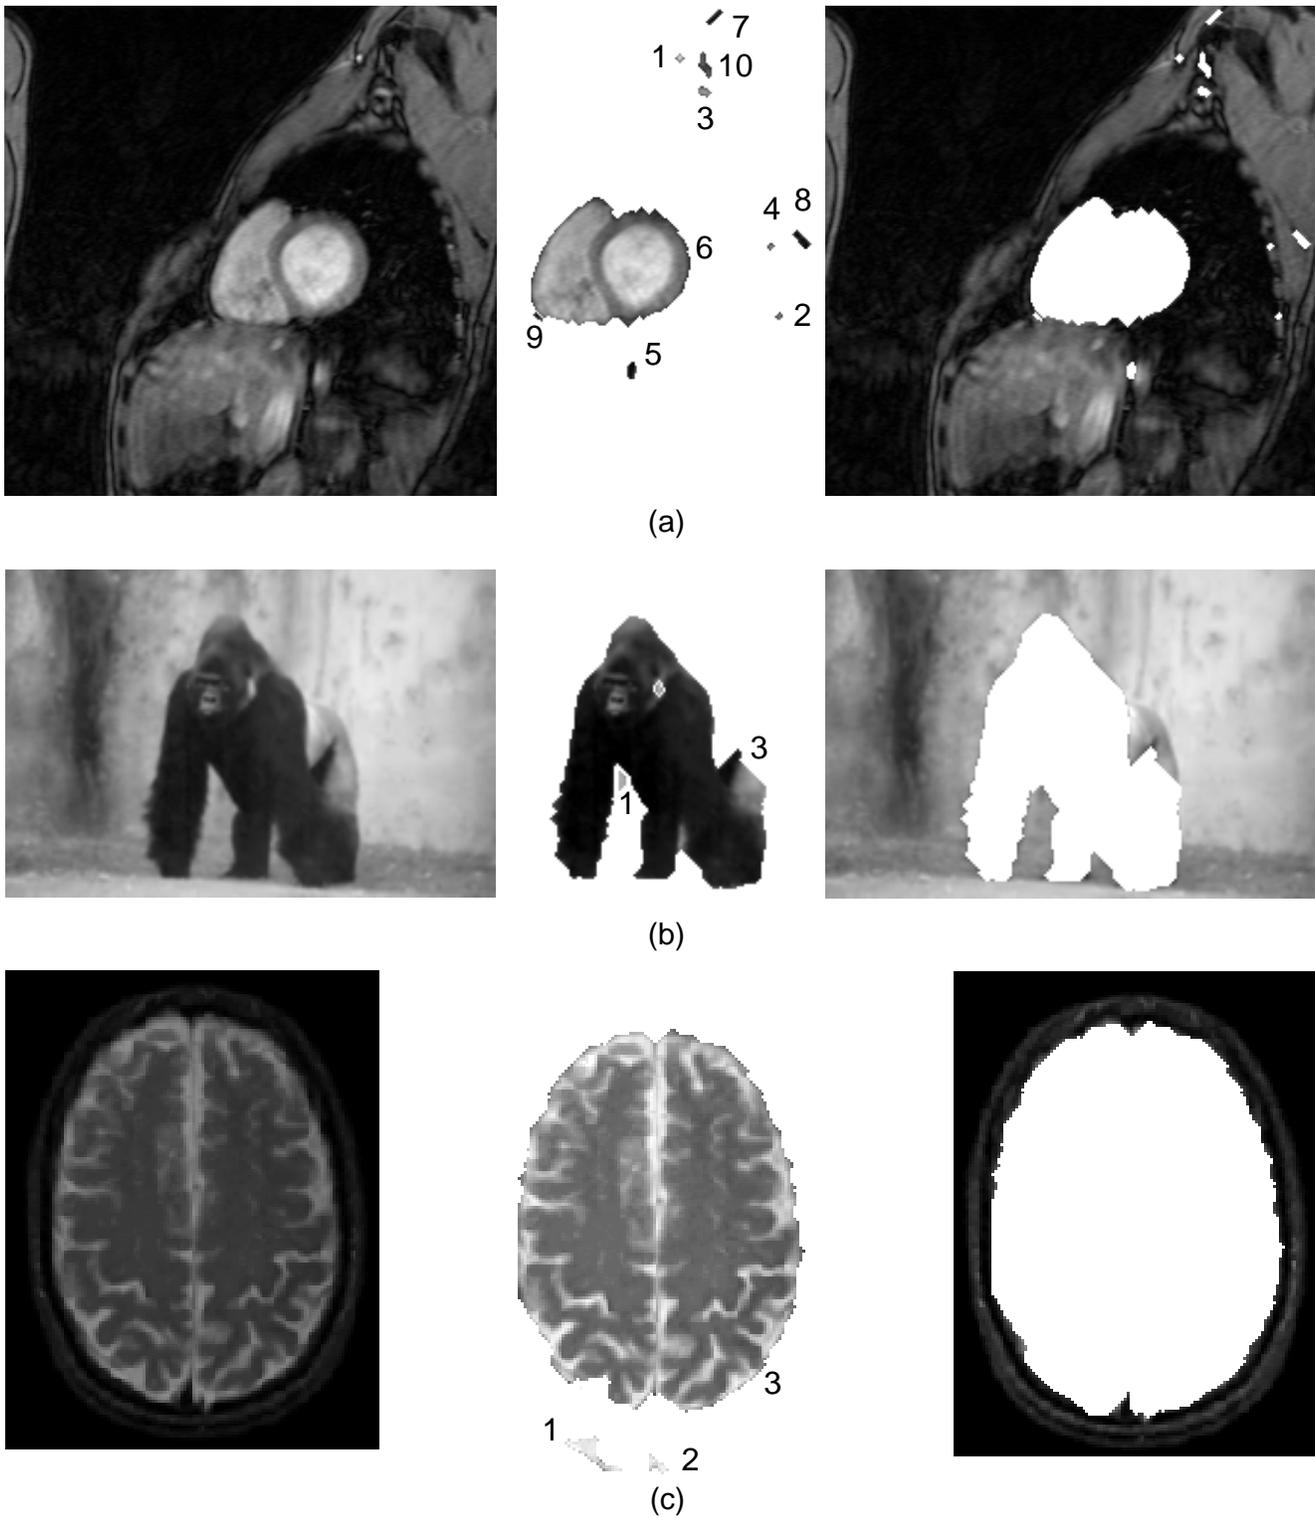


Figure 3. (a) A 256×256 pixel image. Ten regions are shown. (b) A 200×134 pixel image. Shown are three least energy regions (the left ear, under the right arm, and the gorilla.) (c) A 124×166 pixel image. Three regions are shown. The numbers indicate the order in which the regions were found.

The energy can have contributions from both the region and its boundary, thus allowing typical region information such as texture and homogeneity to be combined with typical boundary information such as intensity gradients. The two types of energy are transformable into each other however, and by transforming the region energy into a boundary energy we can cast the global optimization problem into the form of a minimum mean weight cycle problem in a digraph. This problem has a simple, polynomial-time algorithm that requires no initialization, and is highly parallelizable. We described experiments using combinations of region and boundary information that illustrate the strength of the method. The energy is of a very general form, although always globally optimizable by the same algorithm, and offers many other possibilities for further modeling.

Acknowledgments. The authors wish to thank the Instituto de Matemática Pura e Aplicada in Rio de Janeiro for their generous hospitality, and Professor Davi Geiger for his advice and encouragement. This work was partially supported by grant number MURI-AFOSR 25-74100-F1837.

References

- [1] A. Amini, S. Tehrani, T. Weymouth, "Using dynamical programming for minimizing the energy of active contours in the presence of hard constraints," *Proc. Int'l Conf. Comp. Vis.*, pp. 95–99, 1988.
- [2] A. Blake, A. Zisserman, *Visual Reconstruction*, MIT Press, 1987.
- [3] I. J. Cox, S. B. Rao, Y. Zhong, "Ratio regions: a technique for image segmentation," *Proc. Int'l Conf. Patt. Rec.*, 2, pp. 557–564, 1996.
- [4] I. J. Cox, J. M. Rehg, S. Hingorani, "A bayesian multiple-hypothesis approach to edge grouping and contour segmentation," *Int'l J. Comp. Vis.*, Vol. 11 (1), 1996.
- [5] G. B. Dantzig, W. Blatner, M. R. Rao, "Finding a cycle in a graph with minimum cost to time ratio with application to the ship routing problem," in *Theory of Graphs: International Symposium*, Dunod, Paris, and Gordon and Breach, New York, pp. 209–213, 1966.
- [6] J. Elder, S. W. Zucker, "The effect of contour closure on the rapid discrimination of two-dimensional shapes," *Vision Res.*, Vol. 33, pp. 981–991, 1993.
- [7] J. Elder, S. W. Zucker, "A Measure of Closure," *Vision Res.*, Vol. 34 (24), pp. 3361–3369, 1994.
- [8] J. Elder, S. W. Zucker, "Computing Contour Closure," *Proc. Euro. Conf. Comp. Vis.*, pp. 399–412, June, 1996.
- [9] D. Geiger, A. Gupta, L. Costa, J. Vlontzos, "Dynamic programming for detecting, tracking, and matching deformable contours," *IEEE Trans. Patt. Anal. Mach. Intell.*, Vol. 17 (3), pp. 294–302, 1993.
- [10] G. Guy, G. Medioni, "Inferring global perceptual contours from local features," *Int'l J. Comp. Vis.*, Vol. 20 (1-2), 1996.
- [11] G. Kaniza, "Contours without gradients or cognitive contours," *Italian Jour. Psych.*, Vol. 1, pp. 93–112, 1971.
- [12] G. Kaniza, *Organization in Vision: essays on gestalt perception*, Praeger, New York, p. 150, 1979.
- [13] R. Karp, "A characterization of the minimum cycle mean in a digraph," *Dis. Math.*, Vol. 23, pp. 309–311, 1978.
- [14] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active Contour Models," *Int. J. Comp. Vis.*, Vol. 1 (4), pp. 321–331, 1988.
- [15] I. Kovács, B. Julesz, "A closed curve is much more than an incomplete one: effect of closure in figure-ground segmentation," *Proc. Natl. Acad. Sci. USA*, Vol. 90, pp. 7495–7497, 1993.
- [16] E. L. Lawler, "Optimal cycles in doubly weighted linear graphs," in *Theory of Graphs: International Symposium*, Dunod, Paris, and Gordon and Breach, New York, pp. 209–213, 1966.
- [17] T. Leung, J. Malik, "Contour continuity in region based image segmentation," *Proc. Euro. Conf. Comp. Vis.*, Germany, 1998.
- [18] N. Meggido, "Combinatorial optimization with rational objective functions," *Mathematics of Operations Research*, Vol. 4, pp. 414–424, 1979.
- [19] U. Montanari, "On the optimal detection of curves in noisy pictures," *Comm. ACM*, Vol. 15 (5), pp. 335–345, 1971.
- [20] P. Parent, S. W. Zucker, "Trace inference, curvature consistency, and curve detection," *IEEE Trans. Patt. Anal. Mach. Intell.*, Vol. 11 (8), 1989.
- [21] A. Sha'ashua, S. Ullman, "Structural saliency: the detection of globally salient structures using a locally connected network," *Proc. Second. Int'l Conf. Comp. Vis.*, Florida, USA, 1988.
- [22] J. Shi, J. Malik, "Normalized Cuts and Image Segmentation," *Proc. IEEE Conf. Comp. Vis. Patt. Rec.*, Puerto Rico, pp. 731–737, 1997.
- [23] K. K. Thornber, L. R. Williams, "Analytic Solution of Stochastic Completion Fields," *Biological Cybernetics*, Vol. 75, pp. 141–151, 1996.
- [24] L. R. Williams, D. W. Jacobs, "Stochastic Completion Fields: A Neural Model of Contour Shape and Saliency," *Neural Computation*, Vol. 9, pp. 849–870, 1997.
- [25] L. R. Williams, K. K. Thornber, "A Comparison of Measures for Detecting Natural Shapes in Cluttered Backgrounds," *Proc. Euro. Conf. Comp. Vis.*, pp. 432–448, June, 1998.