Region and Edge-Adaptive Sampling and Boundary Completion for Segmentation

Scott E. Dillard[1], Lakshman Prasad[2], and Jacopo Grazzini[2]

[1] Pacific Northwest National Laboratory
[2] Los Alamos National Laboratory

Abstract. Edge detection produces a set of points that are likely to lie on discontinuities between objects within an image. We consider faces of the Gabriel graph of these points, a sub-graph of the Delaunay triangulation. Features are extracted by merging these faces using size, shape and color cues. We measure regional properties of faces using a novel shape-adaptive sampling method that overcomes undesirable sampling bias of the Delaunay triangles. Instead, sampling is biased so as to smooth regional statistics within the detected object boundaries, and this smoothing adapts to local geometric features of the shape such as curvature, thickness and straightness. We further identify within the Gabriel graph regions having uniform thickness and orientation which are grouped into directional features for subsequent hierarchical region merging.

1 Introduction

Perceptual organization (*aka* grouping and segmentation) is a process that computes regions of the image that come from different objects with little detailed knowledge of the particular objects present in the image [9]. One possible solution to the image segmentation problem, which requires to construct a partition of the image into perceptually meaningful parts, is to perform it subsequently to edge detection [27,16,26,1]. Early works in computer vision have emphasized the role of edge detection and discontinuities in segmentation and recognition [19,18,29]. This line of research stresses that edge detection should be done at an early stage on a brightness, colour, and/or texture representation of the image and segmentation (likewise other early vision modules) should operate later on [17].

An edge detector [30,6] yields a set of pixels which are likely to lie on object boundaries, but the union of these pixels may not form complete boundaries that partition the image [5]. The goal of the segmentation is to then complete the object boundaries by linking pixels together into closed loops [16,25,2]. A well-known solution to this problem is to consider the output of the edge detector as a real-valued function (e.g., magnitude of the image gradient) and then perform the watershed transform [28,12,11]. This produces a so-called *over-segmentation*, a partition of the image that contains too many regions, but from which the desired segmentation can hopefully be obtained by region merging.

A number of methods have been presented which work with a binary edge detector, one which indicates that a pixel is or is not to be treated as part of

G. Bebis et al. (Eds.): ISVC 2010, Part II, LNCS 6454, pp. 64–74, 2010.

an object boundary. One strategy is to employ the Delaunay triangulation of the edge pixels and then select some subset of the triangle edges to complete the object boundaries. Criteria for selecting this subset include region properties such as triangle size and average of color [14,15], or contour properties such as continuity of direction [23,22].

In this paper we improve upon existing methods in two ways. First, we propose a method for shape-adaptive sampling of regional properties (e.g. color) that simultaneously smoothes properties within objects and sharpens them across object boundaries. Second, we incorporate directionality information into the regional properties, not just contour properties, to extract long, straight features of nearly-constant thickness. These two components are incorporated into a pre-segmentation method that parses the image into visually significant regions which are then combined into a hierarchy by merging.

2 Background

We begin with an edge detector, a method for identifying those edge pixels which are likely to coincide with the boundaries of objects depicted in an image. Edge detection is a well-studied problem and an active area of research [20]. We use the Canny detector [3] although we do not rely on any particular properties of this detector and the general strategy presented here is applicable to any detector. For clarity we call the set of points produced by the edge detector *edgels*, whereas a line segment between two such points is an *edge*.

Proximity is paramount among the cues responsible for grouping edgels into salient object contours [7], and so proximity graphs such as the Delaunay triangulation and Gabriel graph provide good candidate edges for completing object boundaries. The *Delaunay triangulation* of a point set $P \subset \mathbb{R}^2$ is a triangulation of the convex hull of P such that the interior of every triangle's circumcircle is disjoint from P. The *Gabriel graph* is a sub-graph of the Delaunay triangulation containing every edge that is a diameter of a circle whose interior is disjoint from P. For uniformity we also refer to this circle as the edge's circumcircle, with circumradius equal to half the edge's length. More complete definitions of these constructions can be found in the text by Goodman & O'Rourke [10].

Let $D_P(x) = \min_{p \in P} \|x - p\|$, the distance from a given point to the closest member point in the set P. Taken as an elevation map, the faces and edges of the Gabriel graph are uniquely identified with peaks and saddle points of $D_P(x)$, respectively [8]. A Gabriel face f contains a point m which is a local maximum of D_P, maximally far from all the nearby points P and hence considered to lie in the middle of the shape defined by P. The value of D_P gives the local width of the shape, which is the radius of the largest circumcircle of any triangle in the Gabriel face. We define the *thickness* of a Gabriel face to be twice this radius, equal to $2D_P(m)$. Figure 1 illustrates a Delaunay triangulation (dashed lines) and the Gabriel graph embedded within it (solid lines.) The circles labeled f and g are the largest that can fit inside their respective Gabriel faces and so define the thickness of those faces. If we consider both faces as parts of the whole shape,

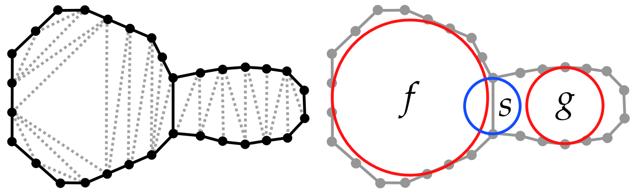

Fig. 1. Left: A Delaunay triangulation (dashed lines) and Gabriel graph (solid lines.) Right: Circles f and g are the circles of greatest radius centered within their respective faces. The distance between the radii of f and s gives the total variation in thickness of the shape.

then the Gabriel edge between them, labeled s, identifies the thinnest part of the shape. We measure the variation in a shape's thickness by the ratio between it's maximum and minimum local thickness, in this case f and s respectively.

3 Related Work

Prasad & Skourikhine [22] proposed a method to link edgels into boundaries that begins by grouping them into chains that are adjacent according to an 8-pixel neighborhood connectivity. A system of filters for the Delaunay triangulation of edge chains was formulated, for example, "Do all three triangle vertices belong to different chains?" or "Does a triangle edge connect to chain endpoints?"

The methods proposed by Köthe *et al.* [14] and Letscher & Fritts [15] are similar in that they both were inspired by α-shapes [4], so-called because all Delaunay triangles and edges of circumradius greater than α are removed. Additionally, both groups of authors proposed to use a pair of thresholds, with segmentation regions being "seeded" by triangles having circumradius greater than the larger threshold and region boundaries delineated by edges and triangles with circumradius below the smaller threshold. Their methods differ in how they deal with connected components of triangles with radii between α and β. Köthe *et al.* proved that, under certain assumptions, a topological thinning of the remaining triangles produces an accurate boundary reconstruction. Letscher & Fritts proposed to merge the remaining regions according to color information, with mergers being subject to topological constraints. Both sets of authors advocate using color information to merge triangles. In this paper we explore the effect of the sampling method used to assign a color to triangles.

4 Region and Edge-Adaptive Sampling

Multiple authors [15,14] have suggested the use of circumradius to discern large triangles in the middle of objects from small triangles on their boundary, with these small triangles being subsequently merged with others based on regional properties such as color. It then becomes important how one measures the pixel color distribution under a triangle. Monte Carlo sampling has been suggested [22]

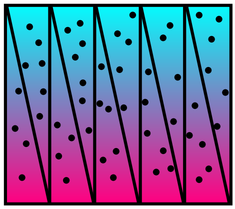

(a) Monte Carlo sampling

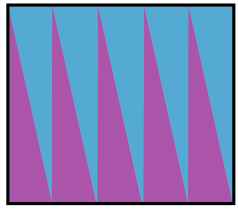

(b) Sample mean

Fig. 2. (a) Triangle colors are estimated by uniform sampling over their area, either by a Monte Carlo method or rasterization. (b) The shape of triangles biases the sampling.

to obtain triangle color information, and enumerating the pixels under a triangle by rasterization is an obvious choice. However, uniform sampling of pixel intensity over the triangle interior inherits a bias from the shape of the triangle. If the triangle has one edge much shorter than the other two then the barycenter shifts toward that edge and biases the sampling in that direction, as shown in Figure 2. The desired behavior of an edge detector is to yield a high density of points along the boundary of an object and few in its interior, so we expect that many such skinny triangles will occur.

To counter these triangle sampling artifacts we bias the sampling in a different direction. Rather than sample uniformly within the triangle, we propose to sample within its circumcircle. We observe that a Gabriel face contains the circumcenters of all the Delaunay triangles it comprises. The amount which a circumcircle extends outside of the Gabriel face depends on the local properties of the shape at the Gabriel edge. In Figure 3(a) the Gabriel edge is long compared to the radii of the maximal circumcircles in the adjacent faces and does not obstruct the circumcircles from penetrating across the edge. In Figure 3(b) the Gabriel edge is shorter, forming a narrowing of the shape which restricts the amount which the neighboring circumcircles overlap. In Figure 3(c) the Gabriel edge obviously bounds the shape and so the circumcircle is contained well within the interior of the face. By sampling regional properties, such as average pixel color, within the circumcircles of Delaunay triangles we not only alleviate the undesirable sampling bias of skinny triangles but introduce desirable bias by sampling away from Gabriel edges which are likely to bound shapes and smoothing the sampling across Gabriel edges which are likely to lie in the middle of shapes.

We can exaggerate the influence of boundary shape on the sampling bias by shrinking the circumcircles. Samples are then shifted away from the edgels, which are those pixels indicated by the edge detector to be regions of significant image variability, e.g. high gradient or fluctuating image statistics. The effect is similar to anisotropic diffusion [21] but is decoupled from the image gradient by way of the edge detector. Figure 4 demonstrates the circumcircle-based sampling method and contrasts it with triangle-based sampling.

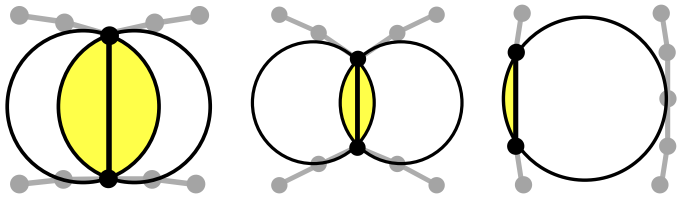

Fig. 3. (a) Circumcircles overlap across Gabriel edges in the interior of shapes, creating a smoothing effect when color is sampling within these circles. (b) Circumcircles do not overlap as much in regions where the shape is varying, and hardly at all (c) across edges on the shape boundary.

To evaluate the effect of triangle color sampling on region merging we use the binary scale-climbing region merging heuristic proposed by Guigues *et al.* using the Mumford-Shah two-term energy function [11]. Figure 5(a) shows the Delaunay triangles that remain after removing edges that are longer than 5 pixels. We arbitrarily chose to merge regions until 40% of edgels were removed. Figure 5(b) shows the merged regions resulting from uniform sampling of triangle colors, and in Figure 5(c) the triangles are given the colors sampled under their circumcircles. Notice in Figure 5(b) how some long, thin triangles poke out of the regions they have been merged with. The region boundaries in Figure 5(c) correspond more closely with the Gabriel graph, as we expect from the deliberate sampling bias. Also note that the dynamic range of colors is preserved better with the proposed sampling method, as color is not sampled along region boundaries.

It should be noted that the sum of the areas of the circumcircles is not linearly related to the sum of the areas of their triangles. Consider the Delaunay triangulation of points placed along two parallel lines at unit intervals. Move the lines apart and the area of each triangle grows in proportion to the distance moved, as do the radii of the circumcircles and thus their areas grow quadratically. To mitigate what is potentially a source of quadratic computational complexity, we use a constant-time sampling procedure for each circumcircle. We generate a pyramid of downsampled images, where each pixel at level i is the average of the four pixels below it at level $i+1$. A circumcircle centered at c with radius r is given the color obtained from the image at level $\lfloor \log_2 r \rfloor$ using bilinear interpolation to weight the four downsampled pixels in the neighborhood of c.

5 Directional Features

Large, significant features can readily be identified by eliminating long Gabriel edges, which are unlikely to bound objects because they have very little corroboration from the edge detector. It does not suffice, however, to consider edge length alone when identifying smaller features, which may still be visually and semantically significant. Rather than fall back to color-based region merging in

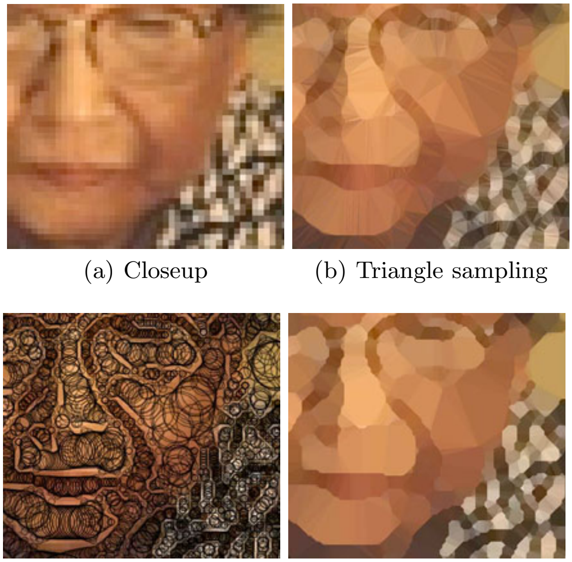

(a) Closeup (b) Triangle sampling

(c) Shrunken circumcircles (d) Circumcircle sampling

Fig. 4. Closeup on the face of the woman in Figure 8(a), left. (a) Original detail. In (b), Delaunay triangles are given the average color of pixels underneath them using Monte Carlo sampling and bilinear interpolation. Circumscribing circles of Delaunay triangles are shown in (c), with 75% of their original radius. In (d) triangles are given the average color of pixels under their shrunken circumcircles.

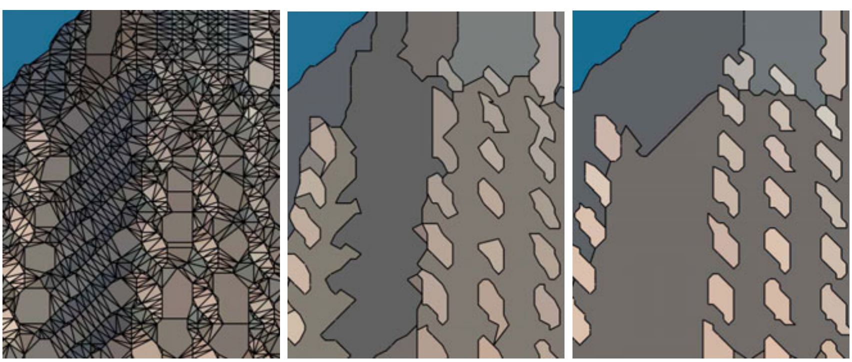

(a) Delaunay triangles (b) With triangle sampling (c) With circumcircle sampling

Fig. 5. Closeup on the building in Figure 8(d), right. (a) Delaunay triangles remaining after removing edges longer than 5 pixels. Small triangles are merged using the heuristic of Guigues *et al.* [11] until 40% of edgels have been removed. In (b) triangles colors were obtained by uniform sampling over their interiors, and in (c) by the proposed method of sampling in their circumcircles.

the small-radius regime, we propose to extract long, thin features using directional properties. A *directional feature* is a connected set of Gabriel faces with minimal variation in thickness and orientation.

To identify these features we first classify all Gabriel edges as *contour* or *chord*. A contour edge is definitely part of the shape boundary, whereas a chord edge may or may not be contained in the shape interior. An edge e is classified as contour if it's length is not greater than $\sqrt{2}$ pixels, or a chord otherwise. This corresponds to the usual notion of 8-neighbor pixel connectivity, although different thresholds may be used with sub-pixel edge detection [14]. Additionally, Gabriel edges may be classified as contour if shape thickness varies too much across the edge. This variation is measured by the ratio of edge length to the maximum thickness of either adjacent Gabriel face. This additional criterion for discriminating between contour and cord edges will be justified shortly. The *degree* of a Gabriel face is then number of Gabriel chords on its boundary. Candidate groups of degree-2 faces are then identified. A *sequence* of Gabriel faces is one such that faces f_i and f_{i+1} share a common Gabriel edge. We find all sequences of degree-2 faces and then recursively split them until there is no significant variation of thickness or directionality within each sequence.

The variation of thickness is measured by the ratio $T_{\min}/T_{\max}$, where $T_{\min}$ is the length of the shortest Gabriel edge between subsequent faces of the sequence, and $T_{\max}$ is the greatest thickness of any face in the sequence. If this ratio smaller than some fixed threshold (we use 0.5 in the results presented below) then the sequence is split at the shortest Gabriel edge. Figure 6(a) shows a pair of Gabriel faces that fail this condition. Because we only consider sequences of Gabriel faces with variation in thickness above this threshold, it is justified to classify as contour those Gabriel edges across which thickness varies too much, as they cannot possibly be part of a sequence. This results in more faces having degree of 2 and the formation of longer candidate sequences.

The variation of directionality of the sequence is measured by considering how well the sequence is approximated by a straight line segment between the circumcenters of the first and last Gabriel faces (centers of their maximal circumcircles.) Call these centers p_0 and p_1. Let e be a Gabriel edge between subsequent Gabriel faces, with endpoints q_0 and q_1. Let t be the value for which $q_0+t(q_1-q_0)$ intersects line p_0p_1, and let $C_e = |0.5 - t|$. When $C_e = 0$ the line p_0p_1 cuts e directly in half. If $C_e < 0.5$ then p_0p_1 intersects e somewhere in its interior. $C_e > 0.5$ means that p_0p_q strays outside the "channel" defined by e. We bound the directional variation of Gabriel face sequences by placing a threshold on C_e for each edge e encountered in the sequence. If some C_e value is greater than this threshold, we split the sequence at the edge with the largest value of C_e. Figure 6(b) shows a sequence of Gabriel faces with small values of C_e for the intervening edges, whereas Figure 6(c) exhibits larger values of C_e.

Figure 7 shows the effect of varying the threshold on C_e from 0.125 to 2. To illustrate the directional features we show in white their *centerlines*, which are the line segments connecting circumcenters of Delaunay triangles encountered

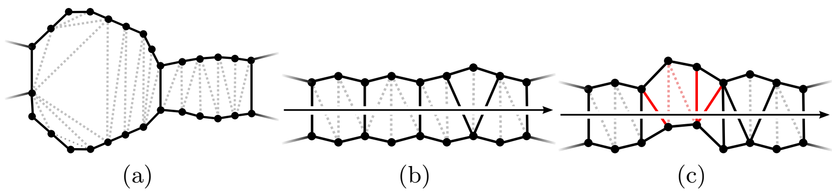

Fig. 6. (a) Thickness varies too much between the larger Gabriel face and the shared Gabriel edge, and so these two faces are not a candidate directional feature. (b) The line from the first Gabriel edge midpoint to the last cuts all intervening edges close to their middles, so these faces constitute a directional feature. (c) The line strays outside of the middle of some of the intervening edges (colored) and prevents the faces from forming a directional feature.

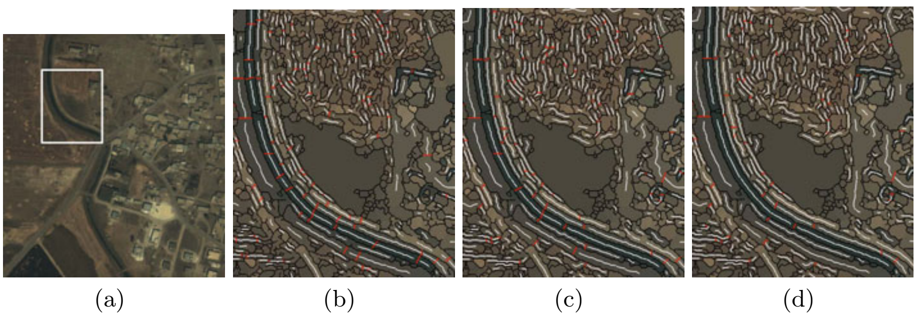

Fig. 7. (a) Satellite image. Directional features are found using varying thresholds: (b) 1/8, (c) 1/2 and (d) 2. Contour edges are drawn in black and the centerline approximations of directional features are drawn in white. Gabriel edges which cut directional features are drawn in red. Note how larger thresholds cause fewer cuts.

on a walk from the first Gabriel edge to the last. For low thresholds of C_e these centerlines will be nearly straight while larger thresholds allow them to curve.

6 Results

Figure 8 shows the proposed method applied to four images from the Berkeley segmentation dataset [20]. A 2x-upsampled Canny edge detector was used with Gaussian smoothing $\sigma = 1.5$, high and low hysteresis thresholds set to 80th and 70th percentile of image gradient magnitude, respectively. Delaunay edges longer than 5 pixels were removed. Directional features were identified with $C_e \leq 1$ and $T_{\min}/T_{\min} \geq 1/2$. Images in the middle row of Figure 8 show the result of the proposed pre-segmentation methods. Initial regions with Gabriel faces having thickness not less than 2.5 pixels are shown with their average color. Smaller regions are left white, except for directional features whose centerlines are drawn in black. Region colors were obtained with the proposed circumcircle-sampling

Fig. 8. Top row: original images. Middle row: Large segmentation regions, those containing Gabriel faces of thickness greater than 2.5 pixels, are drawn with their color. Small regions are left white, except for the centerlines of directional features, which are drawn in black. Bottom row: Regions remaining after merging until 60% of edgels are removed.

method. The bottom row of Figure 8 shows regions that remain after removing 60% of edgels using the binary scale-climbing heuristic [11].

7 Conclusion

Completion of edge detection boundaries by Delaunay triangulation allows shape information to be explicitly incorporated into the pre-segmentation process.

Shape-adaptive color sampling tangibly improves region merging results. Early identification of salient image features based on color, size, directionality and uniformity of thickness allows for these features to be captured as discrete perceptual units within the region merging hierarchy, rather than relying on the region merging heuristic to hopefully coalesce them at some point in the merging process. The effect is to increase the density of information attributed to regions contained in the hierarchy, which in turn increases the value of hierarchical region merging as a tool for higher-level machine vision tasks such as parts-based object recognition and scene description [24,13].

References

1. Arbelaez, P., Cohen, L.: Constrained image segmentation from hierarchical boundaries. In: Proc. IEEE Computer Vision and Pattern Recognition (2008)
2. Arbelaez, P., Maire, M., Fowlkes, C.C., Malik, J.: From contours to regions: An empirical evaluation. In: Proc. IEEE Computer Vision and Pattern Recognition, pp. 2294–2301 (2009)
3. Canny, J.: A computational approach to edge detection. Readings in computer vision: issues, problems, principles, and paradigms, 184 (1987)
4. Edelsbrunner, H., Mücke, E.: Three-dimensional alpha shapes. In: Proceedings of the 1992 workshop on Volume visualization, p. 82. ACM, New York (1992)
5. Elder, J., Zucker, S.: Computing contour closure. In: Buxton, B.F., Cipolla, R. (eds.) ECCV 1996. LNCS, vol. 1065, pp. 399–412. Springer, Heidelberg (1996)
6. Elder, J.: Are edges incomplete? International Journal of Computer Vision 34(2/3), 97–122 (1999)
7. Elder, J., Goldberg, R.: Ecological statistics of Gestalt laws for the perceptual organization of contours. Journal of Vision 2(4), 5 (2002)
8. Giesen, J., John, M.: The flow complex: A data structure for geometric modeling. Computational Geometry 39(3), 178–190 (2008)
9. Gonzalez, R., Woods, R.: Digital Image Processing, 2nd edn. Prentice-Hall, Upper Saddle River (2002)
10. Goodman, J., O'Rourke, J.: Handbook of Discrete and Computational Geometry. Chapman & Hall, Boca Raton (2004)
11. Guigues, L., Cocquerez, J., Le Men, H.: Scale-sets image analysis. International Journal of Computer Vision 68(3), 289–317 (2006)
12. Haris, K., Efstratiadis, S., Maglaveras, N., Katsaggelos, A.: Hybrid image segmentation using watersheds and fast region merging. IEEE Transactions on Image Processing 7(12), 1684–1699 (1998)
13. Hoiem, D., Efros, A., Hebert, M.: Closing the loop in scene interpretation. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (2008)
14. Köthe, P., Stelldinger, H.M.: Provably correct edgel linking and subpixel boundary reconstruction. In: Franke, K., Müller, K.R., Nikolay, B., Schäfer, R. (eds.) DAGM 2006. LNCS, vol. 4174, pp. 81–90. Springer, Heidelberg (2006)
15. Letscher, D., Fritts, J.: Image segmentation using topological persistence. In: Kropatsch, W.G., Kampel, M., Hanbury, A. (eds.) CAIP 2007. LNCS, vol. 4673, pp. 587–595. Springer, Heidelberg (2007)
16. Leung, T., Malik, J.: Contour continuity in region based image segmentation. In: Burkhardt, H., Neumann, B. (eds.) ECCV 1998. LNCS, vol. 1406, pp. 544–559. Springer, Heidelberg (1998)

17. Manjunath, B., Chellappa, R.: A unified approach to boundary perception: edges, textures and illusory contours. IEEE Transactions on Neural Networks 4, 96–108 (1993)
18. Marr, D.: Vision. W.H. Freeman & Co, New York (1982)
19. Marr, D., Hildreth, E.: Theory of edge detection. Proceedings of the Royal Society London B 207, 187–217 (1980)
20. Martin, D., Fowlkes, C., Malik, J.: Learning to detect natural image boundaries using local brightness, color, and texture cues. IEEE Transactions on Pattern Analysis and Machine Intelligence 26(5), 530–549 (2004)
21. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. IEEE Transactions on pattern analysis and machine intelligence 12(7), 629–639 (1990)
22. Prasad, L., Skourikhine, A.: Vectorized image segmentation via trixel agglomeration. Pattern Recognition 39(4), 501–514 (2006)
23. Ren, X., Fowlkes, C., Malik, J.: Learning probabilistic models for contour completion in natural images. International Journal of Computer Vision 77, 47–63 (2008)
24. Shotton, J., Blake, A., Cipolla, R.: Multiscale categorical object recognition using contour fragments. IEEE Transactions on Pattern Analysis and Machine Intelligence 30(7), 1270–1281 (2008)
25. Stelldinger, P., Köthe, U., Meine, H.: Topologically correct image segmentation using alpha shapes. In: Kuba, A., Nyúl, L.G., Palágyi, K. (eds.) DGCI 2006. LNCS, vol. 4245, pp. 542–554. Springer, Heidelberg (2006)
26. Sumengen, B., Manjunath, B.: Multi-scale edge detection and image segmentation. In: Proc. of European Signal Processing Conference, pp. 4–7 (2005)
27. Tabb, M., Ahuja, N.: Multiscale image segmentation by integrated edge and region detection. IEEE Transactions on Image Processing 6(5), 642–655 (1997)
28. Vincent, L., Soille, P.: Watersheds in digital spaces: an efficient algorithm based onimmersion simulations. IEEE Transactions on Pattern Analysis and Machine Intelligence 13(6), 583–598 (1991)
29. Walters, D.: Selection of image primitives for general-purpose visual processing. Computer Vision, Graphics, and Image Processing 37, 261–298 (1987)
30. Ziou, D., Tabbone, S.: Edge detection techniques: an overview. International Journal on Pattern Recognition and Image Analysis 8(4), 537–559 (1998)