LA-UR-

10-03018

Approved for public release;
distribution is unlimited.

Title:

Region and Edge-adaptive Sampling and Boundary
Completion for Segmentation

Author(s):

Scott E. Dillard
Lakshman Prasad
Jacopo Grazzini

Intended for:

British Machine Vision Conference, Aberystwyth, Wales, UK. Aug
30 - Sep 2 2010.

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Region and Edge-adaptive Sampling and Boundary Completion for Segmentation

BMVC 2010 Submission # 405

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-dependant 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.

1 Introduction

One way to approach the image segmentation problem, which asks to construct a partition of the image consisting of perceptually meaningful parts, is to perform it subsequently to edge detection. An edge detector 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. The goal of the segmentation is to then complete the object boundaries by linking pixels together into closed loops. 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 [5, 6, 12]. 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.

Recently, 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 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 triangle size and edge length, regional properties such as average of image color over the triangles [7, 8], or contour properties such as directionality [11].

In this paper we improve upon existing methods in two ways. First, we propose a method for shape-dependant 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.

Figure 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. To move either circle out of its face it must be shrunk to the radius of s . This difference in radii measures how well separated a Gabriel face is from the rest of the graph.

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 [9]. We use the Canny detector [1] 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 image points produced by the edge detector *edgels*. A line segment between two such points is an *edge*.

Proximity is paramount among the cues responsible for grouping edgels into salient object contours, 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. More complete definitions of these constructions can be found in the text by Goodman & O'Rourke [4].

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 edges and faces of the Gabriel graph are uniquely identified with peaks and saddle points of $D_P(x)$, respectively [3]. 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 can be considered to lie in the middle of the shape bounded by P . The value of D_P gives the local girth of the shape, which is the radius of the largest circle that can be centered in f without touching any point in P . This maximal circle is the circumcircle of either an acute Delaunay triangle or a Delaunay edge which is the hypotenuse of two right triangles, of which there is exactly one occurrence in every Gabriel face. We refer to this circle as the Gabriel face's *maximal circumcircle*. Of all the Gabriel edges bounding f , let e be the longest. If one wants to move the circle centered at m outside of f without touching P , one must shrink the diameter to be no more than the length of e . The *separation threshold* of a Gabriel face is the difference between the radii of its maximal circumcircle and half the length of its longest bounding Gabriel edge. Referring to Figure 1, note that Gabriel face f has a high separation threshold while g does not.

3 Related work

Prasad & Skourikhine [11] proposed a method to link edgels into boundaries that begins by grouping edgels into chains that are adjacent according to an 8-pixel neighborhood connectivity. A system of filters was then given in terms of binary triangle edge predicates, for example, "Do all three triangle vertices belong to different chains?" or "Is a vertex the endpoint of a chain?"

The methods proposed by Köthe *et al.* [7] and Letscher & Fritts [8] are similar in that they both were inspired by α -shapes [2], 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 topologically 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.

From a broader standpoint, the strategy of using the Delaunay triangulation to complete edge detector boundaries can be considered an *pre-segmentation* strategy. The results are only as good as the edge detector, which usually results in too many regions, and so in practice some subsequent region merging method is required. Haris *et al.* [6] and Guigues *et al.* [5] proposed region merging methods based on an initial pre-segmentation, derived from the watershed transform of the image gradient.

4 Regional sampling

Multiple authors [7, 8] have suggested to use 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 [11] 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 neighbor-

Figure 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.

Figure 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.

ing 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 shape-dependant 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 [10] 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.

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 [5]. Figure 5(b) 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(c) shows the merged regions resulting from uniform sampling of triangle colors, and in Figure 5(d) the triangles are given the colors sampled under their circumcircles. Notice in Figure 5(c) how some long, thin triangles poke out of the regions they have been merged with. The region boundaries in Figure 5(d) 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.

Figure 4: (a) BSDS image and (b) closeup on the face. In (c) Delaunay triangles are given the average color of pixels underneath them using Monte Carlo sampling and bilinear interpolation. Gabriel edges of pixels found by the Canny detector are shown in (d). Circumscribing circles of Delaunay triangles are shown in (e), with 75% of their original radius. In (f) triangles are given the average color of pixels under their shrunk circumcircles.

Figure 5: (a) BSDS image. (b) Delaunay triangles remaining after removing edges longer than 5 pixels. Small triangles are merged using the heuristic of Guigues *et al.* [5] until 40% of edges have been removed. In (c) triangles colors were obtained by uniform sampling, and in (d) by sampling in circumcircles.

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 be identified within the Delaunay triangulation by considering triangle and edge circumradius. A large triangle, whose circumradius is above some threshold, seeds an object and the boundary of that object is delineated by short edges. This approach works well when the edge detector reveals large open regions of the image. However, visually significant features may be thinner than the threshold, containing no triangle above the threshold and otherwise indistinguishable from its neighbors by considering circumradius alone. Rather than fall back to color-based region merging we propose to extract long, thin features using directional properties.

We classify all Gabriel edges as *contour* or *chord*. An edge e is a contour edge if $\|e\|^2 \leq 2$, a chord otherwise. This corresponds to the usual notion of 8-neighbor pixel connectivity. The *degree* of a Gabriel face is the number of chords bounding it. We first identify candidate sequences of degree-2 faces, which are then recursively split until there is no significant variation of circumradius and directionality within the component. The variation of circumradius is measured by the ratio $(\hat{m} - \check{s})/\hat{m}$ where $\hat{m}$ is the largest circumradius of Gabriel faces in sequence and $\check{s}$ is the circumradius of the smallest Gabriel edge in the sequence which is not the first or last Gabriel edge. If this ratio larger than some fixed threshold (we use 0.75 in the results presented below) then the sequence is split at the shortest Gabriel edge.

The variation of directionality of the sequence is measured by considering how well the sequence is approximated by a straight line segment between the midpoints of the first and last Gabriel chords. We use the internal Gabriel chords to define an adaptive threshold for measuring this approximation quality. Let s_0 and s_n be the first and last Gabriel chords in the sequence, respectively. Let p_0 and p_n be their midpoints. Take each Gabriel chord s_i , $1 \leq i < n$, and scale it about its midpoint by some constant factor (we use 0.5 in the results presented below.) If the line segment from the midpoint of s_0 to the midpoint of s_n does not intersect each of the scaled interior Gabriel chords then the sequence is split and each subsequence considered recursively. The split performed at the interior Gabriel chord whose intersection point with p_0p_n is farthest from its midpoint. Figure 6 illustrates this grouping criterion.

We combine these directional criteria for Gabriel face merging with circumradius-based criteria similar to that proposed by others [7, 8, 11], specifically we remove any Gabriel edge whose radius is above some fixed threshold. This threshold depends on the class of images considered and the edge-detector used. For images from the Berkeley Segmentation Dataset [9] we use a threshold of 5 pixels.

Figure 6: In the top figure, the two Gabriel faces cannot be merged because the line segment between the midpoints of the leftmost and rightmost Gabriel edges does not intersect all interior Gabriel edges, whereas in the bottom figure it does.

Figure 7: (a) Gabriel graph, contours in black and chords in red. (b) Chords longer than 5 pixels are removed. (c) Merging of Gabriel faces based on directionality criterion. (d) Merging of remaining Gabriel faces with most similar neighbor.

6 Results

Figure 7 shows the successive results of applying the criteria to group Gabriel faces. Starting with the Gabriel graph in Figure 7(a), we first remove any edge longer than 5 pixels in Figure 7(b), then grouping the remaining Gabriel faces by the directionality criterion in Figure 7(c). Finally, faces which were not merged with any other are merged with their neighbor having the most similar color.

Figure 8 shows the proposed method applied to satellite imagery.

7 Conclusion

Completion of edge detection boundaries by Delaunay triangulation allows the shape information to be explicitly incorporated into the pre-segmentation process. While results of this process are ultimately limited by the discriminatory power of the edge detector, early identification of salient image features can improve region merging results, and feature definitions based on directionality and thickness, in addition to color and size, provide additional additional object descriptors that can be incorporated into the region merging hierarchy to be fed into subsequent object detection and scene analysis methods.

Figure 8: (a) Satellite image (b) Segmentation regions. Contours (from Canny edge detector) in black, Remaining Gabriel chords after merging in red.

References

- [1] J. Canny. A computational approach to edge detection. *Readings in computer vision: issues, problems, principles, and paradigms*, page 184, 1987.
- [2] H. Edelsbrunner and E.P. Mücke. Three-dimensional alpha shapes. In *Proceedings of the 1992 workshop on Volume visualization*, page 82. ACM, 1992.
- [3] J. Giesen and M. John. The flow complex: A data structure for geometric modeling. *Computational Geometry*, 39(3):178–190, 2008.
- [4] J.E. Goodman and J. O’Rourke. *Handbook of discrete and computational geometry*. Chapman & Hall, 2004.
- [5] L. Guigues, J.P. Cocquerez, and H. Le Men. Scale-sets image analysis. *International Journal of Computer Vision*, 68(3):289–317, 2006.
- [6] K. Haris, S.N. Efstratiadis, N. Maglaveras, and A.K. Katsaggelos. Hybrid image segmentation using watersheds and fast region merging. *IEEE Transactions on Image Processing*, 7(12):1684–1699, 1998.
- [7] U. Köthe, P. Stelldinger, and H. Meine. Provably Correct Edgel Linking and Subpixel Boundary Reconstruction. *Pattern Recognition*, pages 81–90.
- [8] D. Letscher and J. Fritts. Image Segmentation Using Topological Persistence. In *Computer Analysis of Images and Patterns*, pages 587–595. Springer.
- [9] DR Martin, CC Fowlkes, and J. Malik. 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.
- [10] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on pattern analysis and machine intelligence*, 12(7):629–639, 1990.
- [11] L. Prasad and A.N. Skourikhine. Vectorized image segmentation via trixel agglomeration. *Pattern Recognition*, 39(4):501–514, 2006.

- [12] L. Vincent and P. Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE transactions on pattern analysis and machine intelligence*, 13(6):583–598, 1991.