

Extraction of Lakes from Satellite Imagery

Amnon Krupnik
Department of Civil Engineering
Technion – Israel Institute of Technology
Haifa 32000, Israel
E-mail: krupnik@tx.technion.ac.il

James H. Elder
Centre for Vision Research
York University
Toronto, ON M3J 1P3, Canada
E-mail: jelder@yorku.ca

ABSTRACT

We address the problem of finding the bounding contour of a lake in high-resolution satellite imagery. The proposed solution consists of a novel approach for probabilistic contour grouping when some prior (possibly rough) knowledge about the lake is known. Such knowledge is available, for example, from an existing GIS.

The grouping process is based on a Bayesian approach and a constructive algorithm. We use four object cues that are calculated for each edge segment in the image, and four grouping cues between consecutive edges on the lake boundary. We derived statistical models for these cues by using data from a training set of lakes, for which an IKONOS image and NTDB (Canadian National Topographic Database) polygonal data are available, and for which the contours have been traced manually. We then test the algorithm on new lakes. For the latter, we obtained independent measurements of the actual contour by eight geomatics experts. These measurements are used for the evaluation of the results.

A quantitative analysis of the results shows that our algorithm improved on the accuracy of the prior GIS models by an average of 41%. The accuracy of our algorithm is comparable to human expert accuracy.

KEYWORDS: contour grouping; perceptual organization; high-resolution satellite imagery; GIS; lake boundary extraction

1. Introduction

Intelligent extraction of cartographic details from aerial and high-resolution satellite imagery is an essential step in the automation of the mapping process. While data for modern digital mapping environments are usually originated in digital imagery, most object extraction tasks are still performed interactively. In order to perform such tasks automatically, with reliability that will satisfy the high standards of the mapping community, extensive knowledge about the objects of interest must be employed.

Many research studies have been based on employing common (“expert”) knowledge about the objects of interest, in order to transform image features into a meaningful arrangement. Thus, for example, a method for extracting roads would assume that the width of a road is within certain bounds, or that its curvature must be smaller than a predefined value. Unfortunately, formulating the knowledge representation as such might lead to solutions that are good only for a particular set of data. A more general approach would be to employ objective knowledge about the objects and image features. Such knowledge can be formulated in probabilistic terms. Examples for employing probabilistic inference to mapping tasks can be found in Barzohar and Cooper (1996) and Geman and Jedyak (1996).

We propose a perceptual-grouping approach for extracting object contours* from high-resolution satellite images. There are numerous existing algorithms for partially grouping contour elements into disjoint contour fragments (Lowe, 1989; Sound, 1992; Huttenlocher and Wayner, 1992; Jacobs, 1993; Cox *et al.*, 1993; Jacobs, 1996; Crevier, 1999), and such partially-grouped representations have been shown to be useful for object recognition (Jacobs, 1996). However the goal of computing *complete* bounding contours has proven to be more elusive. While recent approaches exploiting the global property of contour closure have yielded limited success (Elder and Zucker, 1996; Mahamud *et al.*, 1999), the general problem of computing the complete bounding contour of an object of arbitrary shape in a complex natural image remains essentially unsolved.

In this paper we explore the question of whether this difficult problem may be made tractable by application of approximate prior knowledge about the object of interest. We pose the problem of contour grouping as a problem of probabilistic inference: the goal of the computation is to compute the most probable closed contour that completely bounds the object of interest. We formulate a Bayesian inference problem, based on a combination of probabilistic object cues, which provide

* The term *contour* refers to an object boundary, and not to its more common meaning of an equal-elevation line.

information about whether individual edge segments (referred to in the following as *tangents*) lie on the object boundary, with probabilistic grouping cues providing evidence about whether sequences of tangents should be grouped together. In order to use this probabilistic information to estimate the most probable object boundary, we introduce an approximate, constructive search algorithm that allows critical global constraints to be applied throughout the search.

The shortest-path formulation of the grouping problem, suggested by Elder and Zucker (1996), Elder and Goldberg (1998) and Crevier (1999 & 2000), is attractive because it guarantees an exact solution in polynomial time. The problem is that the assumptions required to formulate the problem in this way may be so severe that the optimal solution to the approximated problem may not closely resemble the optimal solution to the exact problem. In particular, the shortest-path approach demands that local inferences be conditionally independent, making it impossible to incorporate important global constraints. For example, the constraint that object boundaries contain no geometric self-intersections is an obvious and important one, but since it is a global constraint (distant components of the contour may intersect), it cannot be incorporated into the shortest path formulation. In this paper we propose as an alternative an approximate, constructive search technique, which finds a good (not necessarily optimal) solution, and which can accommodate such important global cues and constraints.

The constraint that the complete object boundary be *closed* is also critical to the success of our approach. Contour closure has been shown to be a powerful cue to perceptual grouping in human vision (Koffka, 1935; Elder and Zucker, 1993; Kovacs and Julesz, 1993). Jacobs (1993 and 1996) has studied the problem of inferring highly-closed convex cycles of line segments from an image, to be used as input for a part-based object recognition strategy. More recently, closure has been demonstrated as a very potent cue in computer vision algorithms for grouping more general contours (Elder and Zucker, 1996; Mahamud *et al.*, 1999).

An interesting aspect of the proposed approach is the utilization of existing GIS data. Many GIS projects are facing nowadays the challenging stage of database updating. In this process, databases already contain a large amount of information. While this information might be outdated and even erroneous, it can provide important clues for the extraction procedure. The connection between digital images and a geographical database can be therefore viewed as a two-way relation, with respect to object recognition and extraction. Digital images are used as a source of up-to-date information, while existing databases are utilized as a source of prior knowledge and for guiding the extraction procedure. This concept has been suggested recently in several studies (e.g., Vosselman, 1996; Agouris *et al.*, 1998).

Perceptual organization algorithms are often evaluated informally on somewhat arbitrary natural images. One goal of the present work is to objectively and quantitatively evaluate algorithm performance on real data. We therefore apply the proposed approach to the problem of extracting precise lake boundaries from panchromatic, high-resolution satellite images. Rough polygons that describe the boundaries were available from a GIS, and constituted part of the prior information. For the evaluation process, we compiled a ground truth database derived from hand segmentations of the tested lakes by eight human mapping experts, and use this database to compare the performance of our algorithm to expert human performance.

Section 2 describes the methodology of the perceptual organization approach and its application to the problem of extracting lake boundaries. In section 3 we present and analyze results obtained by experimenting the derived algorithm with satellite images and GIS.

2. Methodology

2.1 Definitions

We assume an input representation in which contours are locally represented by variable length tangents. These tangents are obtained by an adaptive multi-scale edge detection algorithm (Elder and Zucker, 1998), and a greedy algorithm for grouping local edges into relatively short linear segments of variable length. In addition to carrying position and orientation information, these tangents are augmented with estimates of the luminance intensity on both sides of the contour. Details of the representation can be found in Elder and Zucker, 1996.

Let $T = \{t_1, \dots, t_N\}$ be the set of all tangents in an image. The set S of possible contours may then be represented as tangent sequences:

$$s = \{t_{\alpha_1}, \dots, t_{\alpha_m}\} \in S, \text{ iff :} \quad (1)$$

$$\alpha_i \in \{1, \dots, N\}, \forall i \in \{1, \dots, m\}; \alpha_i \neq \alpha_j, \forall i \in \{1, \dots, m\}, \forall j \in \{2, \dots, m-1\}, i \neq j$$

This definition restricts the mapping to be injective (tangents cannot be repeated), with the exception that the first and last tangent may be the same. In this case the contour is *closed*.

We assume that there exists a correct organization of the image $C \subset S$. C includes all tangent sequences that correspond to actual contours in the image, and all subsequences of these. In other words, contours in C need not be complete. C includes the contour bounding the object of interest, but also many other 'distracting' contours that might exist in the image. We do not restrict these contours to be disjoint, since there are many instances in real images where contours bounding distinct structures may share one or more contour elements. For example, in our application, the visible boundary of a lake is in fact composed largely of trees and bushes that surround the lake. Thus, there are many contour elements that may be considered either part of the boundary of a single tree, or part of the boundary of an entire lake.

We define the set of *object fragments* S^0 such that

$$s = \{t_{\alpha_1}, \dots, t_{\alpha_m}\} \in S^0 \subset S \text{ iff } t_{\alpha_i} \in T^0, \forall i \in \{1, \dots, m\}, \quad (2)$$

where T^0 is the set of tangents on the boundary of the object of interest. We define the set of *object contours* C^0 as the set of correctly organized object fragments, i.e., $C^0 = S^0 \cap C$. We assume that the object boundary c^{0*} and therefore all object contours are simple (not self-intersecting).

2.2 Constructive algorithm

We propose a constructive method for estimating the complete object boundary c^{0*} from contours that are likely to be object contours. We begin by constructing the set $S_2 = \{s_{2_1}, \dots, s_{2_{n_2}}\}$ of contours s_{2_i} of length $m = 2$ tangents, and estimating the posterior probability $p(s_{2_i} \in C^0 | D)$ that each is an object contour, given observed data D (see Section 2.4).

To limit memory and time complexity, we discard all but the N_2 most probable contours \hat{C}_2^0 , and from these we construct a new set $S_3 = \{s_{3_1}, \dots, s_{3_{n_3}}\}$ of length $m = 3$ contours. This process is iterated to some maximum contour length $m = M$ tangents. Approximate prior knowledge about the object of interest allows the maximum contour length M to be chosen in a principled manner (Section 2.7). Self-intersecting and closed contours may form at all stages $m > 2$; these are subtracted from \hat{C}_m^0 and the closed contours \hat{C}_m^{0*} are stored. Note that self-intersection is a global property that could not be handled easily using a shortest-path algorithm.

2.3 Approximations

In order to estimate the posterior probability of an object contour $p(s \in C^0 | D)$, where $s = \{t_{\alpha_1}, \dots, t_{\alpha_m}\}$, we make a series of simplifying approximations. First we assume that object membership and contour correctness are independent events, so that $p(s \in C^0 | D) = p(s \in S^0 | D)p(s \in C | D)$. At first blush this assumption may seem questionable, since one might think that if one tangent lies on the object and another lies off the object, the probability that they group should be 0. Recall, however, that we do not assume that the correct image organization C is disjoint (Section 2.1), hence this probability does not need to be zero.

We next assume that the object membership of any tangent is independent of all other tangents, so that $p(s \in S^0 | D) = \prod_{i=1}^m p(t_{\alpha_i} \in T^0 | D)$. Similarly, we assume that pairwise groupings are mutually independent, so that $p(s \in C | D) = \prod_{i=1}^{m-1} p(\{t_{\alpha_i}, t_{\alpha_{i+1}}\} \in C | D)$.

We further assume that object membership depends only upon a subset D^0 of the observable data D , and that grouping depends upon a disjoint subset D^c . D^0 consists of a set of observations for each tangent: $D^0 = \{d_{\alpha_1}^0, \dots, d_{\alpha_m}^0\}$, and we assume that object membership for a particular tangent depends only upon the corresponding observation. Similarly, D^c consists of a set of observations for each tangent pair in the sequence, $D^c = \{d_{\alpha_1 \alpha_2}^c, \dots, d_{\alpha_{m-1} \alpha_m}^c\}$, and we assume that each pairwise grouping depends only upon the corresponding observation. Thus we have that:

$$p(s \in C^0 | D) = p(s \in S^0 | D)p(s \in C | D), \quad (3)$$

where

$$\begin{aligned}
p(s \in S^0 | D) &= \prod_{i=1}^m p(t_{\alpha_i} \in T^0 | d_{\alpha_i}^0) \\
p(s \in C | D) &= \prod_{i=1}^{m-1} p(\{t_{\alpha_i}, t_{\alpha_{i+1}}\} \in C | d_{\alpha_i, \alpha_{i+1}}^c).
\end{aligned} \tag{4}$$

2.4 Bayesian formulation

The individual posterior probabilities are estimated using Bayes' theorem:

$$\begin{aligned}
p(t_{\alpha_i} \in T^0 | d_{\alpha_i}^0) &= (1 + L_{\alpha_i}^0 P_{\alpha_i}^0)^{-1} \\
p(\{t_{\alpha_i}, t_{\alpha_{i+1}}\} \in C | d_{\alpha_i, \alpha_{i+1}}^c) &= (1 + L_{\alpha_i, \alpha_{i+1}}^c P_{\alpha_i, \alpha_{i+1}}^c)^{-1},
\end{aligned} \tag{5}$$

where

$$\begin{aligned}
L_{\alpha_i}^0 &= \frac{p(d_{\alpha_i}^0 | t_{\alpha_i} \notin T^0)}{p(d_{\alpha_i}^0 | t_{\alpha_i} \in T^0)} & L_{\alpha_i, \alpha_{i+1}}^c &= \frac{p(d_{\alpha_i, \alpha_{i+1}}^c | \{t_{\alpha_i}, t_{\alpha_{i+1}}\} \notin C)}{p(d_{\alpha_i, \alpha_{i+1}}^c | \{t_{\alpha_i}, t_{\alpha_{i+1}}\} \in C)} \\
P_{\alpha_i}^0 &= \frac{p(t_{\alpha_i} \notin T^0)}{p(t_{\alpha_i} \in T^0)} & P_{\alpha_i, \alpha_{i+1}}^c &= \frac{p(\{t_{\alpha_i}, t_{\alpha_{i+1}}\} \notin C)}{p(\{t_{\alpha_i}, t_{\alpha_{i+1}}\} \in C)}
\end{aligned} \tag{6}$$

The observable data for a tangent ($d_{\alpha_i}^0$) or pair of tangents ($d_{\alpha_i, \alpha_{i+1}}^c$) consist of multiple distinct cues, which typically can be related to classical Gestalt grouping cues (Wertheimer, 1938; Koffka, 1935). For example, $d_{\alpha_i, \alpha_{i+1}}^c$ may consist of one cue for proximity, another for good continuation, etc. Each of these cues is assumed to be independent so that the joint likelihood for the data vector for a given tangent or pair of tangents is simply the product of the individual likelihoods.

The distributions of Equation (6) can be estimated from a set of training images, in which instances of a specific object or class of objects appear in various contexts. The distributions can then be used to estimate novel instances in novel contexts.

In the application described here, we used four object cues d_i^0 for tangent t_i :

- The intensity on the dark side of t_i : lakes tend to have a predictable (dark) intensity in panchromatic IKONOS imagery.
- The distance between t_i and the nearest tangent in the direction opposite to the local intensity gradient (i.e., toward the interior of the lake if t_i is on the lake boundary): lakes tend to have few interior contours.
- The distance between t_i and the nearest point on the polygonal model.
- The angle of t_i with respect to the nearest polygonal model segment.

As grouping cues $d_{i,j}^c$ between tangents t_i and t_j we used:

- The distance (tip to tail) between the tangents (proximity cue).
- A measure of *good continuation* between the two tangents, approximated as the absolute value of the sum of the angles formed by a linear interpolant.
- The change in the mean intensity (brightness) between the two tangents.
- The change in edge contrast between the two tangents.

Figure 1 (b) shows a map of the posterior probability of object membership $p(t_i \in T^0 | d_i^0)$ for the four object cues combined. The gray values of the tangents is proportional to the probability that they lie on the lake boundary.




Figure 1: Example of a test lake. (a) Initial GIS polygon shown in red, registered polygon shown in blue; (b) Map of posterior probabilities of object membership ($t_{\alpha} \in T^0 | d_{\alpha}^0$) for the four object cues combined. Edge gray values are proportional to the posterior probability; (c) Results: yellow – computer-generated contour; red – manually-traced contour; white – non-lake edges; (d) Disagreement between computer-generated and manually-traced contours (black and white areas).

2.5 Selecting the most probable complete boundary

Our independence assumptions allow the posterior probability of an object contour to be approximated as a product of the probabilities that individual tangents lie on the object, and the probabilities of pairwise tangent groupings. This is somewhat troubling; it implies that the posterior probability of an object contour is a monotonically decreasing function of its length, suggesting a strong bias toward shorter contours. We find empirically in our own application that the vast majority of candidate closed contours are very short, and so this bias could be critical.

Our incorporation of object knowledge in the grouping process overcomes this problem. The constructive algorithm yields a set $\hat{C}^{0*} = \{\hat{C}_3^{0*} \cup \hat{C}_4^{0*} \cup \dots \cup \hat{C}_M^{0*}\}$ of closed, non-intersecting contours that are relatively likely to lie on the boundary of the object of interest. From this set, we wish to select the contour most likely to form the *complete* boundary. If this contour indeed represents the entire object boundary, it must be the case that all tangents that are not part of the contour *do not* lie on the object boundary. This last completeness condition is the crucial factor in preventing a bias to short contours.

Specifically, the maximum *a posteriori* (MAP) estimate \hat{c}^{0*} of the lake boundary is given by

$$\hat{c}^{0*} = \arg \max_{c \in \hat{C}^{0*}} p(c \in c^{0*} | D). \quad (7)$$

Given $c = \{t_{\alpha_1}, \dots, t_{\alpha_m}\}$,

$$p(c = c^{0*} | D) = p(c \in C^0, t_i \notin T^0, \forall i \in I^D | D) = p(c \in C^0 | D) \prod_{i \in I^D} (1 - p(t_i \in T^0 | D)), \quad (8)$$

where

$$I^D = \{1, \dots, N\} \setminus \{\alpha_1, \dots, \alpha_m\}.$$

In other words, the probability that a particular contour c forms the entire object boundary is given by the product of the probability that c is correctly organized and lies on the boundary, and the probability that all other tangents do not lie on the boundary.

2.6 Estimating the priors

Statistical models for the prior terms can be derived from lake instances for which the bounding contours were hand-traced by a mapping expert on the edge map. The prior ratios P_i^0 and $P_{i,j}^c$ for object membership and grouping, respectively, are determined solely from the number of tangents in the image and the length of the prior model.

The prior ratio for object membership P_i^0 is simply the ratio of the number of tangents off the lake boundary to the number on the lake boundary. Of course, we do not know the exact number of tangents on the boundary, but we can use the prior model to estimate this. We use the training data to estimate the mean $\bar{\alpha}$ of the ratio of the number of tangents on the lake boundary to the length of the polygonal model. Given a novel image we estimate the number of tangents N_l on the lake boundary to be $N_l = \bar{\alpha}L$, where L is the length of the polygonal model (in pixels). We then have that

$$P_i^0 = \frac{(N - N_l) / N}{N_l / N} = N / N_l - 1, \quad (9)$$

where N is the total number of tangents in the image.

We estimate the prior ratio $P_{i,j}^c$ for grouping by assuming that on average each tangent groups with one other tangent in each direction. Note that this is simply an estimate of the expected case, and does not imply that every tangent in the image will group with exactly one tangent in both directions. This assumption implies that

$$P_{i,j}^c = N - 2. \quad (10)$$

2.7 Selecting the maximum contour length

The approximate polygonal model can be used to select the maximum contour length M considered in the search for the object boundary. In Section 2.6 we noted the linear relationship between the number N_l of tangents on the lake boundary and the length L of the prior model. We can also estimate the standard deviation σ_α . Given a new image and new prior model of length L pixels, we estimate the 95th percentile upper bound M on the number of tangents on the boundary as

$$M = L(\bar{\alpha} + 2\sigma_\alpha). \quad (11)$$

3. Experiments and results

In order to demonstrate the proposed approach we test it on the problem of computing exact lake boundaries from high-resolution IKONOS imagery. The images cover an area over the Gatineau Hills in Québec, which contains lakes in different sizes and shapes. We use panchromatic, 1 m resolution georeferenced imagery. In addition to the images, we have rough polygonal models of the boundaries of these lakes from the Canadian Topographic Database (NTDB).

The experiments consist of four stages:

- Registration: As the IKONOS data is only georeferenced (and not true orthophotographs), this stage is essential in order to bring the two sources into a common coordinate system.
- Training: A first set of lakes was used as training data for estimating the priors and to model the likelihoods for each of the object and grouping cues.
- Grouping: The proposed approach was tested on a second set of lakes. None of these lakes participated in the training stage.
- Evaluation: The results of the grouping were compared to contours that were traced manually by human experts.

3.1 Registration

As lakes are horizontal objects, the fact that the IKONOS images are not true orthophotographs does not affect the relative accuracy among different parts of the boundary. However, the absolute position of the lake might be shifted by several tens of meters. The absolute position of the NTDB polygonal model, on the other hand, is considered to be fairly accurate, despite the inaccurate shape of the polygon, compared to the true shape of the lake.

In order to apply an algorithm that utilizes the two sources, we performed a rough registration. This registration was based on the observation that lakes encompass relatively few contours (save for those generated by islands, marshlands, etc.). Registration thus involves a search over the set of 2D translations and (possibly) rotations to minimize the number of edge pixels within the polygon. Figure 1 (a) shows an example of the polygonal model, superimposed on the image, before (red) and after (blue) the registration.

3.2 Training

We carried out a training session with a set of seven lakes. The contour of each of these lakes was hand-traced on the edge map, resulting a sequence of edges on the lake boundary. Based on this tracing, the following were calculated:

- Average ratio of the number of tangents on the lake boundary to the length of the polygonal model: We find that the relationship between the length of the polygonal model and the number of tangents on the human-traced object boundary for our training data is nearly linear, with a Pearson correlation of 0.97. The average ratio is $\bar{\alpha} = 0.28$ tangents/pixel. This value is used for calculating the prior P_i^0 according to Equation (9).
- Maximum contour length: We further used the ratio between the length of the polygonal model and the number of tangents on the human-traced boundary and found that the standard deviation of α is $\sigma_\alpha = 0.06$ tangents/pixel. With new images, this value is used for calculating the maximum contour length according to Equation (11).
- Likelihoods and posteriors for each of the objects and grouping cues: The numerators and denominators of the likelihood ratios defined in Equation (6) can be obtained by fitting a certain probability distribution function to the data collected about the lake boundaries. Based on the shape of the data histograms, we employed a generalized Laplacian distribution that has been used in the past to model kurtotic wavelet response histograms (Mallat, 1989; Simoncelli and Adelson, 1996; Simoncelli, 1999). Figure 2 (top and middle on the left) displays the data and the fitted probabilistic model of $p(INTENSITY | t_{\alpha_i} \notin T^0)$ and $p(INTENSITY | t_{\alpha_i} \in T^0)$ for the intensity cue, respectively. The bottom graph is the posterior $p(t_{\alpha_i} \in T^0 | INTENSITY)$ for one of the images. Similar graphs for the brightness change grouping cue are shown on the right side of Figure 2.

3.3 Grouping

We tested the grouping approach proposed in Section 2 on five novel instances of lakes. These images are relatively complex: the number of tangents in an image averaged 19,000, ranging from roughly 3,000 to 42,000. The number of tangents on the lake boundary averaged 490, ranging from roughly 100 to 1000.

Results of the grouping process for four of the five lakes are shown in Figure 1 and Figure 3a (overlaid on the edge map), and Figure 4 (overlaid on the image). The yellow polygon was constructed from the actual sequence of grouped edges. Gaps were completed by straight segments.

3.4 Ground truth

In order to objectively evaluate the quality of the results, ground truth is required. An ideal ground truth database would be obtained by a field survey of the actual lake perimeters. However, To obtain such data at the resolution of our satellite imagery (1 m) is impractical. Typically, such high-resolution contours are mapped by hand from remote-sensed imagery. Thus, for this study we constructed a ground truth database from hand segmentations of lake boundaries on the image data by eight human mapping experts*.

* Note that these test lakes were traced in the image domain. In contrast, the lake contours used to estimate the object and grouping statistics were traced by a single expert (one of the authors) in the tangent domain, using an in-house software-tracing tool. This gave us direct statistics on the estimated tangents of the contours.

We were concerned that these human expert segmentations might be less accurate for low contrast images. To control for this, segmentations were performed on two sets of images: the original set, and a second, histogram-equalized set. Each expert was instructed to trace the perimeter of each of the test lakes using the software tool of their choice*.

In order to compare two estimates of a closed object boundary, we represent both contours as binary images, with interior pixels assigned the value of 1 and exterior pixels assigned the value 0. We then define the difference between the contours as the integral of the absolute value of the difference map, normalized by the mean area of the contours (see Figure 1d for example). This difference metric is non-negative, and we find that when pooled over experts or lakes, the distribution of segmentation differences is roughly log-normal. Thus all statistical tests were performed in the log domain.

Using our difference metric, we found segmentations by the same expert on the original images to differ from those on the histogram-equalized images by an average of 9.4%. Across experts, segmentations differed by an average of 12.6%, for both the original and histogram-equalized image sets. Figure 3b shows an example of the diversity between human experts for segmentations of the same lake.

3.5 Performance of the grouping algorithm

Lake boundaries computed by our algorithm were found to differ from the human expert segmentations by an average of 15.2% for the original image set and 15.9% for the histogram-equalized set. This difference is statistically insignificant ($t(72)=0.26, p>.05$)**. Thus, the contrast of the images appears not to play a role in the consistency of our algorithm's results with human expert judgments. We therefore restrict further analysis to the human expert segmentations on the original images. Figure 1c and Figure 3a show examples of the differences between our algorithm and the human expert segmentations. Differences are not independently distributed around the lake perimeters, but rather are concentrated at ambiguous areas where marshlands and lake inlets may or may not be considered part of the lake.

The mean difference between our algorithm's segmentations and the human expert segmentations (15.2%) is not significantly greater than differences between experts (12.6%), $t(61)=1.32, p>.05$. Nevertheless, there are a number of possible contributions to the differences between the algorithm and the human experts. For example, the human experts segmented the images themselves, whereas the algorithm 'viewed' the data in the edge domain. In addition, the algorithm exploited prior information in the form of the GIS model, which the human experts did not have access to. To evaluate the importance of these factors, we also computed differences between the algorithm's segmentations and those generated by a single human expert (one of the authors) employing the edge-based tracing tool used to obtain the training segmentations. This expert also had access to the prior GIS models. We found that the mean difference between our algorithm and this edge-based segmentation was 14.1%, not significantly different from the differences between our algorithm and the image-based human expert segmentations (15.2%): $t(4)=0.84, p>.05$.

We have shown that the mean difference between our algorithm's segmentations and human expert segmentations is not significantly greater than differences between human experts. It would be useful to know by how much our algorithm improves on prior knowledge embodied in the original GIS models. We find that the mean difference between the prior models and the human segmentations is 25.6%, whereas the mean difference between the algorithm's segmentations and human expert segmentations is only 15.2%. Thus our algorithm significantly improves on the accuracy of the prior models, $t(59)=3.62, p<.05$. The magnitude of this improvement ranges from 0-67%, averaging 41% over the 5 lakes tested. The worst case (0% improvement) is shown in Figure 3. The lack of improvement is due to ambiguity over whether a large section of marshland (in the lower part of the lake) should be considered part of the lake. This ambiguity is also reflected in the diversity of human expert segmentations.

* One of the 8 observers did not segment the histogram-equalized images. For one of the 5 original images, a second observer's segmentation was corrupted. Thus, we obtained 8 expert segmentations for 4 of the 5 original images, and 7 for the fifth. We obtained 7 expert segmentations for all the histogram-equalized images.

** All statistical tests reported are independent measures *t*-tests, using the Welch-Satterthwaite approximation for the number of degrees of freedom.

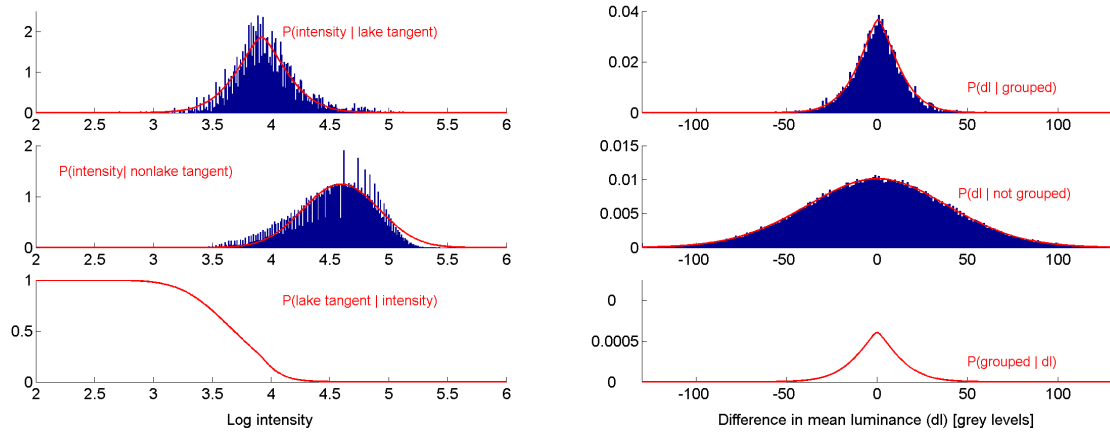


Figure 2: Left: Likelihood distributions for the intensity cue of tangents on and off the lake boundaries, and the resulting posterior distribution; Right: Likelihood distributions for the intensity difference cue of grouped tangents and tangents that are not grouped, and the resulting posterior distribution.

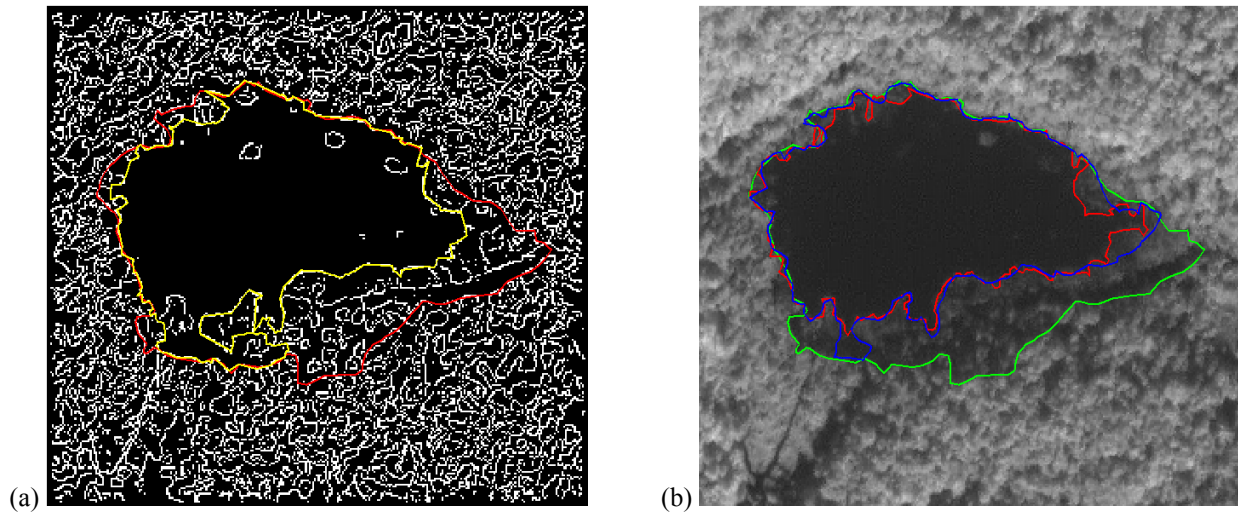


Figure 3: Results of matching over marshland: (a) Yellow – computer-generated contour; Red – manually-traced contour; (b) Differences among contours that were traced by different observers.




Figure 4: Results of two test areas: Yellow – computer-generated contour; Blue – registered GIS polygon.

4. Summary

We have introduced the problem of grouping complete bounding contours using prior object models, and have shown how it can be formulated as a problem of Bayesian inference. We proposed an approximate constructive search technique on the set of possible bounding contours. This approach permits the incorporation of important global constraints on the simplicity and completeness of the bounding contour.

We demonstrated the effectiveness of the method on the problem of computing exact lake boundaries from high-resolution satellite imagery given approximate polygonal models from a GIS database. We compiled a ground truth database from hand segmentations by human mapping experts, allowing us to quantitatively and objectively evaluate our results. Our analysis shows that the algorithm improved on the accuracy of prior GIS models by an average of 41%. Furthermore, The accuracy of our algorithm is comparable to human expert accuracy.

A number of independence assumptions were made to simplify the formulation (Section 2.3). The accuracy of our results suggests that these assumptions were reasonable for the problem at hand. However, relaxing these assumptions will allow the method to be generalized to problems involving far weaker priors. We believe the constructive search technique to be general enough to allow many of these assumptions to be relaxed in future work.

Acknowledgment

This work was supported by research grants from GEOIDE, CRESTech and NSERC.

References

- Agouris P, Gyftakis S and Stefanidis A (1998) Using a fuzzy supervisor for object extraction within an integrated geospatial environment. *International Archives of Photogrammetry and Remote Sensing*, 32(III/1):191-195.
- Barzohar M and Cooper DB Automatic finding of main roads in aerial images by using geometric-stochastic models and estimation. *IEEE Trans. Pattern Analysis and Machine Intelligence* 18(7):707-721.
- Cox JJ, Rehg JM and Hingorani S (1993) A Bayesian multiple-hypothesis approach to edge grouping and computer segmentation. *Int. J. Comp. Vision* 11(1):5-24.
- Crevier D (1999) A probabilistic method for extracting chains of collinear segments. *Computer Vision and Image Understanding* 76(1):36-53.
- Elder JH and Zucker SW (1993) The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research* 33(7):981-991.
- Elder JH and Goldberg R (1998) The statistics of natural image contours. *Proceedings, IEEE Computer Society Workshop on Perceptual Organization in Computer Vision*.
- Elder JH and Zucker SW (1996) Computing contour closure. *Proceedings, 4th European Conference on Computer Vision* 399-412.
- Elder JH and Zucker SW (1998) Local scale control for edge detection and blur estimation. *IEEE Trans. Pattern Analysis and Machine Intelligence* 20(7):699-716.
- Geman D and Jedynak B (1996) An active testing model for tracking roads in satellite images. *IEEE Trans. Pattern Analysis and Machine Intelligence* 18(1):1-13.
- Huttenlocher DP and Wayner PC (1992) Finding convex edge groupings in an image. *Int. J. Comp. Vision* 8(1):7-27.
- Jacobs DW (1993) Robust and efficient detection of convex groups. *Proceedings, IEEE Conference on Computer Vision and Pattern Recognition* 770-771.
- Jacobs DW (1996) Robust and efficient detection of salient convex groups. *IEEE Trans. Pattern Analysis and Machine Intelligence* 18(1):23-37.
- Koffka K (1935) *Principles of Gestalt Psychology*. Harcourt, Brace & World, New York.
- Kovacs I and Julesz B (1993) A closed curve is much more than an incomplete one: Effect of closure in figure-ground discrimination. *Proc. Natl. Acad. Sci. USA* 90:7495-7497.
- Lowe DG (1989) Organization of smooth image curves at multiple scales. *Int. J. Comp. Vision* 3:119-130.
- Mahamud S, Thornber KK and Williams LR (1999) Segmentation of salient closed contours from real images. *Proceedings, International Conference on Computer Vision* 891-897.
- Mallat SG (1989) A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Analysis and Machine Intelligence* 11(7):674-693.
- Simoncelli EP and Adelson EH (1996) Noise removal via Bayesian wavelet coring. *Proceedings, IEEE International conference on Image Processing* pp 379-389.
- Simoncelli EP (1999) Modeling the joint statistics of images in the wavelet domain. *Proceedings, SPIE 44th Annual Meeting*, 3813:188-195.
- Sound E (1992) Labeling of curvilinear structure across scales by token grouping. *Proceedings, IEEE Conference on Computer Vision and Pattern Recognition* 817-830.
- Vosselman G (1996) Uncertainty in GIS Supported Road Extraction. *International Archives of Photogrammetry and Remote Sensing* 31(B3):909-917.
- Wertheimer M (1938) Laws of organization in perceptual forms. In Ellis WD (ed) *A sourcebook of Gestalt Psychology* Routledge and Kegan Paul, London pp 71-88.