1997

Mine boundary detection using partially ordered Markov models

Xia Hua
Iowa State University

Jennifer Davidson
Iowa State University

Noel A. Cressie
Iowa State University, ncressie@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/infopapers>

Part of the [Physical Sciences and Mathematics Commons](#)

Recommended Citation

Hua, Xia; Davidson, Jennifer; and Cressie, Noel A.: Mine boundary detection using partially ordered Markov models 1997, 152-163.
<https://ro.uow.edu.au/infopapers/2380>

Mine boundary detection using partially ordered Markov models

Abstract

Detection of objects in images in an automated fashion is necessary for many applications, including automated target recognition. In this paper, we present results of an automated boundary detection procedure using a new subclass of Markov random fields (MRFs), called partially ordered Markov models (POMMs). POMMs offer computational advantages over general MRFs. We show how a POMM can model the boundaries in an image. Our algorithm for boundary detection uses a Bayesian approach to build a posterior boundary model that locates edges of objects having a closed loop boundary. We apply our method to images of mines with very good results. 2004 Copyright SPIE - The International Society for Optical Engineering.

Keywords

detection, boundary, partially, mine, ordered, markov, models

Disciplines

Physical Sciences and Mathematics

Publication Details

Hua, X., Davidson, J. & Cressie, N. A. (1997). Mine boundary detection using partially ordered Markov models. *Proceedings of SPIE - The International Society for Optical Engineering*, 3167, Statistical and Stochastic Methods in Image Processing II, (pp. 152-163).

Mine Boundary Detection Using Partially Ordered Markov Models

Xia Hua and Jennifer Davidson

Department of Electrical & Computer Engineering
Iowa State University

Noel Cressie

Department of Statistics
Iowa State University

Abstract

Detection of objects in images in an automated fashion is necessary for many applications, including automated target recognition. In this paper, we present results of an automated boundary detection procedure using a new subclass of Markov random fields (MRFs), called partially ordered Markov models (POMMs). POMMs offer computational advantages over general MRFs. We show how a POMM can model the boundaries in an image. Our algorithm for boundary detection uses a Bayesian approach to build a posterior boundary model that locates edges of objects having a closed-loop boundary. We apply our method to images of mines with very good results.

Keywords: automated target recognition (ATR), boundary detection, Markov random fields, partially ordered Markov models.

1. Introduction

This paper presents results for solving the mine-detection problem using boundary identification of objects. Our data typically have low signal-to-noise ratio and there is background clutter and noise. Consequently, approaches to edge detection that use gradient methods, such as the Sobel edge transform, proved to be unsuccessful. These approaches use local pixel-intensity information to identify whether a pixel location is part of a boundary. Filters are used to collect local gradient information and, if the magnitude of the local gradient is large enough, the pixel is declared an edge pixel. Unfortunately, such techniques are sensitive to noise even after thinning and linking, which hinders the determination of closed-object boundaries. We define a closed-object boundary as one that forms a loop, and has no gaps as the boundary is traced.

In this research, we develop a statistical method to obtain statistically optimal one-pixel wide closed-object boundaries in gray-level images, in a computationally fast way. We use a new subclass of Markov random fields, called partially ordered Markov models, to specify our *a priori* boundary model. A fast algorithm based on statistical modeling and Bayes' theorem is developed to identify a closed-object boundary. We first perform a high-pass filter on the original image to increase the ratio of signal to noise, and consequently it provides a better input to the segmentation process. A segmentation algorithm (Helterbrand *et al.*, 1994; Hong and Rosenfeld, 1984) is then applied, allowing us to obtain an initial estimate of one-pixel-wide closed boundaries in image data. Next, we applied the maximum-a-posteriori (MAP) approach to search for the final estimate of the boundary image. This is implemented using iterated conditional modes (e.g., Cressie, 1993, p. 514), which we shall subsequently refer to as ICM. As applications of this algorithm to real data show, it is a fairly robust automated detection procedure that locates the boundaries of objects in mine data.

Recently, a new subclass of Markov random fields called partially ordered Markov models have been developed (Cressie and Davidson, 1994), (Davidson and Cressie, 1996). Partially ordered Markov models are a generalization of Markov mesh models (Abend *et al.*, 1965) and allow the explicit description of the joint probability in terms of spatially local conditional probabilities. Due to the partially ordered Markov model's nice product form of the joint probability, inference and simulation is quite straightforward. POMMs have been successfully applied in fast texture analysis including texture synthesis (Cressie and Davidson, 1994), (Davidson *et al.*, 1996a), texture modeling (Davidson *et al.*, 1996a), (Davidson *et al.*, 1995b), (Davidson *et al.*, 1996b), and image segmentation (Davidson *et al.*, 1995a), (Talukder, 1994).

2. Boundary Models

In this section, we derive the posterior boundary model by specifying an intensity model and a class of prior boundary models. The posterior boundary model will have support only on a subset of the total boundary configuration space (Helterbrand *et al.*, 1994).

Let D denote an $M \times N$ rectangular array of pixel locations. Let $Y(s)$ represent a random variable that denotes the observed intensity at pixel location $s \in D$, and let $X(s)$ denote the unobserved, true underlying intensity at pixel location s . It is assumed that $\{X(s) : s \in D\}$ is a random field with realizations that are constant within four-connected regions (see Gonzalez and Woods, 1992, for connectivity definitions) of D .

We assume an object in the image is a region whose intensity values are (approximately) constant. Thus, based on the premise that objects are defined as regions of statistical homogeneity, it is reasonable to model the intensity probability mass function (pmf) as,

$$P(y | \omega) \equiv P(Y(s) = y(s) : s \in D | \omega). \quad (1)$$

Here, $y \equiv \{y(s) : s \in D\}$ and $w \equiv \{w(s) : s \in D\}$, where $w(s) = 1$ if a boundary is present at location s , and $w(s) = 0$ otherwise.

In Equation (1), we restrict the boundary ω to a subset of all the possible boundary images. Specifically, restrictions relating to the width of the boundaries (one-pixel wide), connectivity properties (eight-connected closed boundaries), and other restrictions are placed on the final properties we want for the boundaries, which we collectively term "permissible" boundaries. Recall that a pixel A is an eight-neighbor of pixel B at location (i,j) if A lies in one of the eight closest locations to B, that is, if A's location is one of $(i-1,j-1)$, $(i-1,j)$, $(i-1,j+1)$, $(i,j-1)$, $(i,j+1)$, $(i+1,j-1)$, $(i+1,j)$, or $(i+1,j+1)$. An eight-connected path consists of a sequence of pixels each consecutive pair of which are eight neighbors; boundary paths are required to be eight-connected. We denote the set of permissible boundaries by Ω_p . The set Ω_p allows us to search a much smaller set for good boundary estimates and, at the same time, guarantee that each such estimate have the desired properties of closed-loop and single-pixel-width. See Helterbrand *et al.*, 1994, for more details.

Further, a Gaussian model is assumed for the observed intensities, $Y(s) = X(s) + \epsilon(s)$; $s \in D$, where ϵ is a normal white noise process with zero mean and variance σ^2 , $\epsilon(s) \sim N(0, \sigma^2)$, and the ϵ -process is independent of the X -process. We note that each $w \in \Omega_p$ partitions D into disjoint four-connected regions. Let $d_l(\omega)$ denote the set of generic four-connected regions for a particular ω . The number of disjoint connected regions depends on ω and will be denoted by $K(\omega)$. Then, each $\omega \in \Omega_p$ implies a partition of D into disjoint connected regions $\{d_l(\omega) : l=1, \dots, K(\omega)\}$, where it is assumed that $X(s)$ is constant on connected regions. That is, $X(s) = X(t)$ if $s, t \in d_l(\omega)$; $l=1, \dots, K(\omega)$.

Define $\mu_l(\omega) \equiv X(s)$; $s \in d_l(\omega)$. For $\omega \in \Omega_p$, it follows that

$$P(Y(s) : s \in D | \omega) = \prod_{l=1}^{K(\omega)} P(Y(s) : s \in d_l(\omega) | \omega),$$

where $Y(s) | \omega \sim N(\mu_l(\omega), \sigma^2)$; $s \in d_l(\omega)$. Thus, upon letting $n_l(w)$ denote the number of distinct sites in $d_l(w)$, we have $s_{lm} \in d_l(\omega) \subseteq D$, $m=1, \dots, n_l(w)$, and the observed intensity model is specified as

$$P(Y(s) : s \in D | \omega) = \prod_{l=1}^{K(\omega)} \prod_{m=1}^{n_l(w)} \left(2\pi\sigma^2 \right)^{-\frac{1}{2}} \exp \left\{ -\frac{(Y(s_{lm}) - \mu_l(\omega))^2}{2\sigma^2} \right\},$$

where $\{\mu_l(w) : l=1, \dots, K(w)\}$ and $\sigma^2 > 0$ are parameters.

Next, we describe a class of partially ordered Markov prior boundary models with support on the set of permissible boundaries. First, we discuss the set of configurations of pixels which ensures permissible boundary identification. Throughout this paper, we restrict the boundary to $w \in \Omega_p$, one of the set of permissible boundary images. The set Ω_p allows us to search a much smaller set for good boundary estimates and, at the same time, guarantee that each such estimate has the desired properties.

We expect and design the partially ordered Markov prior boundary models to ensure the necessary characteristics of a boundary estimate such as the boundary connectivity and also to endorse desired properties such as boundary smoothness and robustness to noise. A good closed-boundary model should be capable of penalizing the excessive placement of boundaries due to noise in the observed image and also of allowing the placement of boundaries where data strongly suggest that boundaries are present.

Now we describe the class of prior closed-boundary models that is supported on the set of permissible boundary configurations and that penalizes configurations with many small connected regions that are not expected to be objects. The various configurations can be defined using lower adjacent neighborhoods as follows. We use a hash mark to express a target pixel location $s = (i, j) \in D$. We use a 5x5 neighborhood as shown in Fig. 1 in which we define three types of undesirable configurations that produce unpermissible boundaries. The figures corresponding to the discussions in (I), (II), and (III) below show the type of undesirable configuration. Where the configuration shown is smaller than the 5x5 neighborhood as given in Figure 1, it is assumed that those pixel values outside the smaller window but inside the 5x5 window are unimportant and are ignored during the actual computation process.

Figure 1. The 5x5 lower adjacent neighborhood $adj_{\prec}y(i, j)$ used to describe undesirable configurations.

(I) w is not a closed boundary if there exists at least one “lit” location (i.e., presence of a boundary segment) that does not have two lit neighbors, as shown in Figure 2. We notate this set of undesired configurations as S_1 .

(II) Object boundaries with the “square” corner configurations shown in Figure 3 are also undesirable. The x’s in Figure 3 express either 1 or 0. These configurations include boundaries that are broader than one-pixel-wide as well as those that do not satisfy the eight-connected permissibility requirement. Requiring eight-connectivity typically results in a smoother boundary. See Figure 4 for an example of both an undesired and a desired configuration. Notate this set of undesired configurations as S_2 .

(III) Due to noise in the observed image, a permissible boundary configuration, $\omega \in \Omega_p$, may include many small connected regions which are not expected to be objects. In this paper, we penalized three small connected regions with an area of one pixel, two pixels and three pixels. We show these undesired configurations in Figures 5, 6, and 7, respectively notated by S_3 , S_4 , and S_5 .

Figure 2. Boundary configurations that are not closed (S_1).

Figure. 3. The "square" corner configurations (S_2).

Figure. 4. (a) Undesirable four-connected boundary. (b) Desirable, eight-connected boundary that is smoothed.

0	1	0
1	0	1
0	1	0
		$\langle i, j \rangle$

Figure 5. The configurations corresponding to a boundary around a single pixel (S_3).

0	0	0	0
0	1	1	0
1	0	0	1
0	1	1	0
			$\langle i, j \rangle$

0	0	1	0
0	1	0	1
0	1	0	1
0	0	1	0
			$\langle i, j \rangle$

Figure 6. The configurations corresponding to a boundary around two pixels (S_4).

0	0	0	0	0
0	0	1	0	0
0	1	0	1	0
0	1	0	0	1
0	0	1	1	0
				$\langle i, j \rangle$

0	0	0	0	0
0	0	0	1	0
0	0	1	0	1
0	1	0	0	1
0	0	1	1	0
				$\langle i, j \rangle$

0	0	0	0	0
0	0	1	1	0
0	1	0	0	1
0	1	0	1	0
0	0	1	0	0
				$\langle i, j \rangle$

0	0	0	0	0
0	0	1	1	0
0	1	0	0	1
0	0	1	0	1
0	0	0	1	0
				$\langle i, j \rangle$

0	0	0	0	0
0	0	0	0	0
0	1	1	1	0
1	0	0	0	1
0	1	1	1	0
				$\langle i, j \rangle$

0	0	0	1	0
0	0	1	0	1
0	0	1	0	1
0	0	1	0	1
0	0	0	1	0
				$\langle i, j \rangle$

Figure 7. The configurations corresponding to a boundary around three pixels (S_5).

The five sets S_k , $k = 1, \dots, 5$ account for 17,433,606 of the 2^{25} possible binary configurations on the 5×5 neighborhood. Thus, the remaining 16,120,826 binary configurations are the “good” ones. We denote the set of remaining “good” configurations by S_6 . Note also that the sets S_k , $k = 1, \dots, 6$, are mutually exclusive, that is, they do not overlap.

We now define a class of partially ordered Markov prior boundary models on the permissible configuration space Ω_p . The principle is to assign low or zero probabilities to the configurations in sets S_k , $k = 1, \dots, 5$, and high probabilities to the remaining configurations (S_6). If desired, the probabilities in the sets S_k , $k = 1, \dots, 5$, could be different. Our prior boundary model is

$$P(w) = \prod_{k=1}^6 p_k^{c_k},$$

where $\{p_k\}$ and $\{c_k\}$ are defined below. This prior can be motivated by formally assigning local conditional probabilities $\{P(y(i, j) | \text{adj}_{\prec} y(i, j))\}$ as follows. Given a particular binary image, whose image values are known at each location (i, j) in the array D , the value for $P(y(i, j) | \text{adj}_{\prec} y(i, j))$ is deemed to be determined by the actual pixel values that occur in $\text{adj}_{\prec} y(i, j)$. Since the six sets S_k , $k = 1, \dots, 6$ completely exhaust all possible configurations in $\text{adj}_{\prec} y(i, j)$, and these six sets do not overlap, there is exactly one set S_k that the 5x5 window consisting of the pixels in $\text{adj}_{\prec} y(i, j)$ is contained in. Define

$$P(y(i, j) = 0 | \text{adj}_{\prec} y(i, j)) \equiv p_{kij},$$

where (i, j) denotes the location in the image of the target pixel, and k is the index of the set S_k in which $\text{adj}_{\prec} y(i, j)$ is contained. The formulation is unusual in that $\{p_{kij} : i = 1, \dots, M, j = 1, \dots, N\}$ will be determined by the pixel values $\{y(i, j) : i = 1, \dots, M, j = 1, \dots, N\}$ that are actually observed.

We now list explicitly the values for $P(y(i, j) | \text{adj}_{\prec} y(i, j))$.

If $y(i, j) = 0$, then define

$$p_{kij} = p_k, \text{ if } \text{adj}_{\prec} y(i, j) \in S_k.$$

Consequently,

$$P(y(i, j) = 1 | \text{adj}_{\prec} y(i, j)) = 1 - p_k, \text{ if } \text{adj}_{\prec} y(i, j) \in S_k. \quad (2)$$

Since $y(i, j)$ is not equal to 1 in the image, then the conditional probability (2) will never actually occur in the likelihood.

Similarly, if $y(i, j) = 1$, then define

$$p_{kij} = 1 - p_k, \text{ if } \text{adj}_{\prec} y(i, j) \in S_k.$$

Consequently,

$$P(y(i, j) = 1 | \text{adj}_{\prec} y(i, j)) = p_k, \text{ if } \text{adj}_{\prec} y(i, j) \in S_k. \quad (3)$$

Since $y(i, j)$ is equal to 1 in the image, the conditional probability (3) is always used in the likelihood. This is true for $k=1, \dots, 6$.

Finally, we formally use the partially ordered Markov model to construct a prior boundary probability, given by the following equation:

$$P(w) = \prod_{l=1}^M \prod_{m=1}^N P(y(l, m) | \text{adj}_{\prec} y(l, m)) = \prod_{k=1}^6 p_k^{c_k},$$

where c_k is the number of times that configurations from set S_k occur in the given image.

The construction of the prior model using a POMM is straightforward and allows the user to specify those configurations that are not wanted. This is a general approach that can be used for other similar types of stochastic modeling. The main advantage of using the partially ordered Markov model as the prior probability of mass function for $w \in \Omega_p$ is that the joint prior probability distribution of w has a product form. In contrast, the more general Markov random field models have a major disadvantage in that an explicit form of the joint probability is not obtainable. Because of the product form of the joint probability, the partially ordered Markov prior boundary model yields a fast and implementable approach to obtain the maximum-a-posteriori boundary configuration w^* .

By Bayes' Theorem, for $w \in \Omega_p$ we can write a posterior closed boundary model:

$$\begin{aligned} \mathcal{P}_r(w | Y(s) : s \in D) &\propto \mathcal{P}_r(Y(s) : s \in D | w) \times \mathcal{P}_r(w) \\ &= \prod_{l=1}^{K(w)} \prod_{m=1}^{n_l(w)} \left(2\pi\sigma^2 \right)^{-\frac{1}{2}} \exp \left\{ -\frac{(Y(l, m) - \mu_l(w))^2}{2\sigma^2} \right\} \\ &\times \{p_1^{c_1} \times p_2^{c_2} \times p_3^{c_3} \times p_4^{c_4} \times p_5^{c_5} \times p_6^{c_6}\} \end{aligned}$$

where $\{\mu_l(w) : l = 1, \dots, K(w)\}$ and $\sigma^2 > 0$ are parameters, and p_k and $c_k, k = 1, \dots, 6$ are defined above.

3. Boundary Models for Mine Detection

In this section, we present the algorithm we used to identify the boundaries of mines and mine-like objects using the prespecified posterior boundary pmf on Ω_p . The algorithm we performed on mine-detection begins by using a multi-resolution image segmentation algorithm to obtain a mostly data-driven initial labeling of boundary and nonboundary pixels. Statistical label models usually require *a priori* knowledge of the number of labels to use and their corresponding parameters; however, this information is not known in our case. Thus, the segmentation algorithm output is used to estimate the number of labels and the corresponding model parameters. Given the output from the segmentation process, the statistical labeling algorithm uses a *maximum-a-posteriori* approach to search for a statistically optimal labeling. Because partially ordered Markov models have the product form of the joint probability, to compute a joint probability using partially ordered Markov models can be much faster than using the more general Markov random field models. We discuss this detail later.

The data we used was one of six multi-spectral images of one scene. The scene contained mines and mine-like objects. The bandwidth of the original data we used was in the range 375 – 425 nanometers, and was 480 x 720 pixels in area, shown in Figure 9. From this we extracted a 256 x 256 subimage, and from that produced a 64 x 64 image by local averaging in 4 x 4 windows of the 256 x 256 image. The 64 x 64 is shown in Figure 10(a). Next we give a description of the algorithm to detect mine and mine-like objects.

We applied a segmentation algorithm to the 64 x 64 image that is detailed in Helterbrand *et al.*, 1995. The segmentation algorithm needs five parameter values that are critical to obtain a good initial estimate to input to the ICM algorithm. The algorithm uses a multi-resolution or pyramid image representation. The root node of the pyramid is the original image, that is, the image data at its finest resolution. The succeeding levels of the pyramid are constructed by grouping a 4 x 4 neighborhood of pixels from the preceding level. The main idea in using pyramids is that at each level of reduced resolution, homogeneous regions (regions with approximately constant gray value) become more compact and hence easier to segment. The argument for partitioning the image this way is that region classification is generally more powerful than pixel classification and, at each level in the pyramid, information is collected over neighborhoods and passed up to the next level. "Parent" nodes at higher levels are connected by link strengths to their 16 "children" at the next lower level. The basic idea of the algorithm is to define link strengths between parent/child pairs on adjacent levels of the pyramid, combining information on gray value similarity and spatial proximity of the pairs. Link strengths are computed iteratively, from the root node to the leaf nodes, for the number of iterations specified. Typically, only 8-10 iterations are necessary. While the segmentation algorithm does not output a permissible boundary, the output is relatively close to permissible, and with a small amount of effort a permissible boundary image $w \in \Omega_p$ is obtained. See Helterbrand *et al.*, 1995, for more details. In principal, any algorithm that produces a permissible boundary can be used as a starting value for the ICM algorithm.

The five parameters that the user must specify to run the segmentation algorithm are:

1. Variance: The variance between levels in the pyramid.
2. Alone: The "Alone" parameter governs the level at which a region becomes defined, which is directly proportional to the number of pixels in that region.
3. Merge: Merge two regions if the gray value difference between the two regions is less than the "Merge" value.
4. Iteration: The number of passes through the entire pyramid.
5. Delta: A large value means that the region is more compact and less spread out.

With this initial estimate, we next applied the iterated conditional modes (ICM) algorithm to the image and obtained a (local) MAP boundary estimate. We remark that in the context of the MAP closed boundary problem, the ICM algorithm automatically selects the mode of the current candidate boundaries. Such a choice guarantees that $P(w^{(t)} | Y) \geq P(w^{(t-1)} | Y)$, where Y represents the image data and t represents the number of iterations. However, if the ICM algorithm converges, it is most likely to a local maximum. The next step was to remove line segments between regions with the same gray values. A line in the image that has (approximately) the same gray values on both sides is removed. The effect of removing line segments results in a better estimate, as we are looking for nontrivial closed-loop boundaries.

Next we give a specific discussion of the algorithms for using partially ordered Markov models. The algorithm we used here is shown in Figure 8.

We wanted to compare the results with a similar experiment (Davidson *et al.*, 1995b) to detect mine boundaries that used a Markov random field as the prior boundary model. Thus, in our experiment here, we used the same 64 x 64 subimage (Figure 10(a)) as given in Davidson *et al.*, 1995b, as our starting image. Before applying the segmentation algorithm, we performed a high-pass filter on this 64 x 64 image. Because edges and other abrupt changes in gray levels are associated with high-frequency components, image sharpening can be achieved in the frequency domain by a highpass filtering process, which attenuates the low-frequency components without disturbing high-frequency information in the Fourier transform. For this particular image, the cutoff frequency was set at the frequency $f_{c_k} = 0.125 Hz$, that is, we keep information containing all frequencies greater than the frequency f_{c_k} . This cutoff frequency was chosen after a visual investigation of the frequency image. This step could be automated. The resulting image from the highpass filter is shown in Figure 10(b). Then we applied the segmentation algorithm to the highpass filtered image, followed by the adjustment to produce a permissible closed boundary image. The five parameters needed in the initial segmentation algorithm are listed as follows: Variance = 110, Alone = 0.016, Merge = 15, iteration = 8, and Delta = 2.7. This segmented image is a permissible closed initial boundary estimate and is shown in Figure 11(a). Finally, we applied the ICM algorithm with the partially ordered Markov boundary model as described above to the initial boundary image and searched for a MAP boundary estimate. Here we used $p_k = 5 \times 10^{-2}$, $k = 1, \dots, 5$ and $p_6 = 0.75$. The MAP estimate of the boundary image, obtained after a large number iterations, is shown in Figure 11(b).

Figure 8. Flowchart describing algorithm for mine boundary detection.

Figure 9. 480 x 720 original image

Figure 10. (a) 64 x 64 local-averaged image. (b) High-passed filtered image

Figure 11. (a) Initial estimate after segmentation (b) final MAP boundary estimate

In the paper, Hua *et al.*, 1995, a general Markov random field model (MRF) was used to model the boundary image. That algorithm for boundary detection combines a Bayesian approach with a histogram specification technique to locate edges of objects that have a closed-loop boundary. However, the algorithm took on the order of hours to compute a final answer, as several passes through the image using the ICM algorithm was needed. Using the POMMs instead of Markov random fields to model the boundary image was believed to be a computationally faster way to identify object boundaries. Our final answer as presented in this paper verified this, as the algorithm using POMMs took less than half the time than the algorithm using MRFs. While both the MRF and POMM boundary models can be used to advocate desired characteristics of a boundary estimate, using the POMM to model the boundary so improved the estimate of the boundary that only one pass was needed to obtain a good final result.

4. Conclusions

In this research, we have developed a simple process that can successfully and quickly detect noisy targets in a fairly automated manner, even with a local optimization technique such as the ICM algorithm. The new class of statistical models, partially ordered Markov models, was applied to detect the boundary of mines and mine-like objects in real data. We believe that this method is a general and potentially powerful one that warrants further investigation and applications.

5. Acknowledgments

This work was supported from grants by ONR (N00014-93-1-0001) and NSF (DMS-9204521). The authors gratefully acknowledge the use of J. Helderbrand's code to perform segmentation on the mine data in Section 3.

Bibliography

(Abend *et al.*, 1965) K. Abend, T.J. Harley, and L.N. Kanal. Classification of binary random patterns. *IEEE Transactions on Information Theory*, IT-11:538-544, Oct. 1965.

- (Cressie and Davidson, 1994) N.A.C. Cressie and J.L. Davidson. Image analysis with Partially Ordered Markov Models. Preprint No. 94-15, Statistical Laboratory, Iowa State Univ., Ames, IA, 1994.
- (Cressie, 1993) N.A.C. Cressie. *Statistics for Spatial Data*. Wiley and Sons, NY, revised edition, 1993.
- (Davidson and Cressie, 1996) J.L. Davidson and N.A.C. Cressie. Partially ordered Markov models and their use in texture synthesis. Manuscript, Department of Electrical and Computer Engineering, Iowa State Univ., Ames, IA, 1996.
- (Davidson *et al.*, 1995a) J. L. Davidson, X. Hua, and D. Ashlock. Genetic algorithms and model selection for texture analysis. In *Proceedings of the 4th Annual Midwest Electro-Technology Conference*, pages 56–59, Ames, IA, 1995.
- (Davidson *et al.*, 1995b) J. L. Davidson, X. Hua, and D. Ashlock. Texture analysis using genetic algorithms and partially ordered Markov models. In E. Dougherty, F. Prêteux, and S. Shen, editors, *Neural, Morphological, and Stochastic Methods in Image and Signal Processing*, volume 2568, pages 197–208, San Diego, CA, 1995.
- (Davidson *et al.*, 1996a) J. L. Davidson, Xia Hua, and D. Ashlock. A comparison of genetic algorithm, regression, and Newton's method for parameter estimation of texture models. In *Proceedings, IEEE Southwest Symposium on Image Analysis and Interpretation*, pages 201–206, San Antonio, TX, 1996.
- (Davidson *et al.*, 1996b) J.L. Davidson, X. Hua, and D. Ashlock. Parameter estimation of partially ordered Markov models for textures recognition. *Proceedings of the 5th Annual Midwest Electro-Technology Conference, ISU*, 1996.
- (Gonzalez and Woods, 1992) R.C. Gonzalez and R.E. Woods. *Digital Image Processing*. Addison-Wesley, New York, 1992.
- (Helterbrand *et al.*, 1994) J. D. Helterbrand, N. Cressie, and J. L. Davidson. A statistical approach to identifying closed object boundaries in images. *Advances in Applied Probability*, 26:831–854, Dec. 1994.
- (Helterbrand *et al.*, 1995) J. D. Helterbrand, J. L. Davidson, and N. Cressie. Optimal closed boundary identification in gray-scale imagery. *Journal of Mathematical Imaging and Vision*, 5:179–205, 1995.
- (Hong and Rosenfeld, 1984) T. H. Hong and A. Rosenfeld. Compact region extraction using weighted pixel linking in a pyramid. *IEEE Trans. on Patt. Analys. and Mach. Intell.*, PAMI-6:222–229, 1984.
- (Hua *et al.*, 1995) X. Hua, J. L. Davidson, and N. Cressie. Mine boundary detection using Markov random field models. In A. C. Dubey, I. Cindrich, and J. M. Ralston, editors, *SPIE's International Symposium on Aerospace/Defense Sensing and Control and Dual-Use Photonics*, volume 2496, Orlando, FL, 1995.
- (Talukder, 1994) A. Talukder. Partially ordered Markov models for texture synthesis and classification. Master's thesis, Dept. of Electrical and Computer Engineering, Iowa State Univ., Ames, 1994.