

AD-A213 854

ICE FLOE IDENTIFICATION IN SATELLITE IMAGES USING
MATHEMATICAL MORPHOLOGY AND CLUSTERING ABOUT
PRINCIPAL CURVES

by

Jeffrey D. Banfield
Adrian E. Raftery


TECHNICAL REPORT No. 172

August 1989

Department of Statistics, GN-22

University of Washington

Seattle, Washington 98195 USA



DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

89 10 31 21E

Ice Floe Identification in Satellite Images Using Mathematical Morphology and Clustering about Principal Curves

Jeffrey D. Banfield

Montana State University

Adrian E. Raftery

University of Washington

ABSTRACT

Identification of ice floes and their outlines in satellite images is important for understanding physical processes in the polar regions, for transportation in ice-covered seas and for the design of offshore structures intended to survive in the presence of ice. At present this is done manually, a long and tedious process which precludes full use of the great volume of relevant images now available.

We describe an automatic and accurate method for identifying ice floes and their outlines. Floe outlines are modeled as closed principal curves (Hastie and Stuetzle, 1989), a flexible class of smooth non-parametric curves. We propose a robust method of estimating closed principal curves which reduces both bias and variance. Initial estimates of floe outlines come from the erosion-propagation (EP) algorithm, which combines erosion from mathematical morphology with local propagation of information about floe edges.

The edge pixels from the EP algorithm are grouped into floe outlines using a new clustering algorithm. This extends existing clustering methods by allowing groups to be centered about principal curves rather than points or lines. This may open the way to efficient feature extraction using cluster analysis in images more generally. The method is implemented in an object-oriented programming environment for which it is well suited, and is quite computationally efficient.

KEYWORDS: Erosion; Feature extraction; LANDSAT; Non-parametric curves; Object-oriented programming; Robustness.

Jeffrey D. Banfield is Assistant Professor, Department of Mathematical Sciences, Montana State University, Bozeman, MT 59717. Adrian E. Raftery is Associate Professor, Department of Statistics, GN-22, University of Washington, Seattle, WA 98195. Banfield's research was supported in part by the Office of Naval Research under Contract no. N-00014-89-1-1114. Raftery's research was supported in part by the Office of Naval Research under Contract no. N-00014-88-k-0265. The authors are grateful to David A. Rothrock for providing the data on which this work was based and for many helpful discussions, and to John A. McDonald and Werner X. Stuetzle for helpful comments and discussions.

1. INTRODUCTION

Knowledge of the shapes, sizes and spatial distribution of ice floes is important for understanding the physical processes operating on the ice pack in the polar regions. It is also important for practical problems associated with transportation in ice-covered seas and for the design of offshore structures intended to survive in the presence of ice.

Such information can be found in satellite images of the polar regions such as Figure 1, which exist in large and rapidly increasing numbers. Practical use of such images requires identification of the outlines of ice floes above a certain size. To date this has been done manually (Rothrock and Thorndike, 1984), a slow and tedious process that often takes a day or more to record the data from a single image and effectively precludes full use of the data. Automating the process is inherently difficult. Problems include the presence of many smaller floes and of melt ponds on the surface of floes which ensure that floes often do not appear as homogeneous blocks of ice in the image.

In this article we describe an automatic method for identifying the outlines of ice floes. The outcome of this is shown in Figure 2, and is almost the same as the result of very careful manual digitization. We model ice floe outlines as closed principal curves (Hastie and Stuetzle, 1989 - hereafter HS), a flexible family of one-dimensional non-parametric curves in a higher-dimensional space. Our method consists of identifying a set of edge pixels and grouping them into clusters which are centered about a principal curve. Each cluster corresponds to a floe and the corresponding principal curve is the estimated floe outline.

The method involves several new statistical techniques:

- (1) A way of estimating closed principal curves that reduces both bias and variance and is robust to outliers. Here, outliers take the form of melt ponds on the surface of ice floes (Section 2).
- (2) The erosion-propagation (EP) algorithm provides initial estimates of floe outlines. This combines the existing idea of erosion from mathematical morphology (Matheron 1975; Serra 1982) with that of local propagation of information about floe boundaries (Section 3).




Figure 1. A polar LANDSAT image showing ice floes. This is a 200×200 pixel image, where each pixel is 80m square; it thus represents a 15×15 km area.

For	
WALS GRA&I	☒
DIIC TAP	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	




Figure 2. The ice floe outlines, larger than a fixed minimum size, found by our procedure for the data in Figure 1.

- (3) A method for clustering about principal curves. Existing clustering algorithms separate data into groups, each of which is clustered about some central point (Gordon 1981, 1987; Murtagh 1985; Committee on Applied and Theoretical Statistics 1989). Here we generalize this to allow each group to be clustered about a different principal curve. This opens the possibility that cluster analysis may be useful more generally for fast feature extraction in images (Section 4).

The method is implemented in an object-oriented programming environment for which it is well suited, and seems computationally efficient.

2. ESTIMATING CLOSED PRINCIPAL CURVES

In this section, we first review the definition and basic properties of principal curves (Section 2.1). We then describe a new algorithm for estimating closed principal curves that reduces both bias and variance and is robust (Section 2.2).

2.1 Principal curves

A principal curve is a smooth one-dimensional curve that passes through the middle of an m -dimensional data set, such as that shown in Figure 3, for which $m=2$. It is non-parametric and its shape is suggested by the data; it thus provides a non-linear summary of the data. The idea was introduced and developed by Hastie (1984), Hastie and Stuetzle (1985) and HS.

A one-dimensional curve in m -space is an m -vector consisting of m functions of a single variable λ , called coordinate functions. The variable λ parameterizes the curves and provides an ordering along it; λ will often be arc-length along the curve. Let $X \in \mathbf{R}^m$ be a continuous random vector. Then $f(\lambda)$ is a *principal curve* of X if

$$E[X | f^{-1}(X) = \lambda] = \lambda,$$

where

$$f^{-1}(x) = \max_{\lambda} \{ \lambda : \|x - f(\lambda)\| = \inf_{\mu} \|x - f(\mu)\| \}.$$




Figure 3. Data for illustrating the fitting of closed principal curves. The data were obtained by generating points uniformly on the circumference of a circle, and perturbing them randomly along the normal to the circle according to a Gaussian distribution.

Given the distribution of X , HS proposed the following algorithm for finding f :

$$f_{i+1}(\lambda) = E[x | f_i^{-1}(X) = \lambda], \quad (2.1)$$

where f_i is the i th iterate.

When the distribution of X is unknown, this extends to estimation of f from a data set $\{x_1, \dots, x_n\}$ by estimating $E[x | f_i^{-1}(X) = \lambda]$. HS do this by means of scatterplot smoothing, using neighborhoods of each point defined by their projections onto the current estimate of the principal curve, rather than by their position in $\mathbb{R}^m$.

Let $\hat{f}_i$ be the i th iterate and let $\tilde{\lambda}_j^i = \hat{f}_i^{-1}(x_j)$ ($j = 1, \dots, n$). Let λ_j^i be the j th order statistic of the set $\{\tilde{\lambda}_1^i, \dots, \tilde{\lambda}_n^i\}$, and let $x_{(j)}^i$ be the data point that projects onto λ_j^i . Then let N_j^i be the set of data points $x_{(l)}^i$ such that λ_l^i is in a neighborhood of λ_j^i ; the size of the neighborhood is controlled by the *span*, equal to the fraction of all the data points which are in the neighborhood. The next iterate is then

$$f_{i+1}(\lambda) = \hat{E}[X | N_j^i]. \quad (2.2)$$

This is calculated using a coordinatewise scatterplot smoother, and various possibilities are discussed by HS.

There is no formal proof that the algorithm converges, but HS report that they have had no convergence problems with more than 40 real and simulated examples. Figure 4 shows the result of applying the HS estimation procedure to the data in Figure 3.

2.2 An algorithm for estimating closed principal curves that reduces both bias and variance and is robust

Scatterplot smoothers generally produce curves that are biased towards the center of curvature. It is clear from Figure 4 that the estimated principal curve suffers from this problem, and that the bias increases with the span. In most statistical problems there is a tradeoff between bias on one hand, and smoothness and variance on the other. For closed curves, the center of curvature is interior to the curve (except for small local regions in non-convex curves). We can use this fact to modify the HS principal curve estimation algorithm so as to reduce bias and variance at once.




Figure 4. Estimated principal curves for the data in Figure 3. The principal curves were estimated using the algorithm (2.2) of HS with spans of 0.2 (thin line), 0.3 (medium line) and 0.5 (thick line).

If the distribution of X is known, let $\delta^i(X, \lambda) = f_{i+1}(\lambda) - f_i(\lambda)$ be the change in the calculated curve from iteration i to iteration $i+1$ in the algorithm (2.1). Then

$$\delta^i(X, \lambda) = E[X - f_i(\lambda) | f_i^{-1}(X) = \lambda].$$

This suggests that when the distribution is unknown and is estimated from the data using an iterative algorithm such as (2.2), the *projections* of the data in N_j^i , rather than the data themselves, should be used to calculate $\hat{f}_{i+1}(\lambda)$. Let $\delta_j^i = \hat{f}_{i+1}(\lambda_j^i) - \hat{f}_i(\lambda_j^i)$ be the change in the estimate of f from the i th to the $(i+1)$ th iteration, let $p_j^i = x_{(j)}^i - \hat{f}_i(\lambda_j^i)$ be the projection of $x_{(j)}^i$ onto $\hat{f}_i$, and let $\bar{p}_j^i$ be the coordinatewise average of the projections of the data in N_j^i onto $\hat{f}_i$. Then we calculate $\hat{f}_{i+1}$ by setting $\delta_j^i = \bar{p}_j^i$.

Figure 5 compares this algorithm with that of HS for various spans. Our algorithm does not suffer from the bias problem inherent in that of HS. Figure 6 shows that our algorithm produces a much smoother curve when the span is small enough for that of HS to yield a relatively unbiased curve.

Outliers arise in the form of shallow but sometimes large melt ponds on the surface of the ice floe. An example of this is Figure 7, which shows the edge pixels of one floe in Figure 1 identified by the EP algorithm. The points near the left side interior to the floe are from melt ponds. Since they do not belong to the edge of the floe, we need to ensure that they not affect the estimate of the principal curve.

To eliminate the effect of outliers we use a slight modification of an approach suggested by HS. In the scatterplot smoothers we use a weighted average, where the weight of a point depends on its distance from the current estimate of the principal curve. We calculate the standard deviation of the lengths of the projections and set the weight for a point to zero if it is more than three standard deviations from the current estimate of the principal curve, and to unity otherwise. Figure 8 shows the result of this robust procedure for the data in Figure 7, as well as that of the non-robust procedure which uses the mean of all the data in each neighborhood. The robust procedure has clearly achieved its goal.




Figure 5. Comparison of the HS algorithm for estimating closed principal curves with that proposed here. This shows the principal curves resulting from the HS algorithm (thin line) and the algorithm proposed here with spans of (A) 0.2, (B) 0.3 and (C) 0.5.




Figure 6. The principal curve from the HS algorithm (thin line) at a span small enough to eliminate most of the bias, compared with the estimate proposed here (thick line) with a span of 0.2. The smaller the span, the less the bias and the rougher the estimated curve. Notice how rough the HS curve is, compared with the present estimate.




Figure 7. Ice floe edges with melt ponds. The small circles are the edge pixels for one of the floes in Figure 1, as identified by the EP algorithm. The points interior to the floe are from melt ponds.




Figure 8. Principal curve estimated using the robust procedure described in the text (thick line), compared with the estimate from the non-robust procedure (thin line). The robust estimate is unaffected by the melt ponds, while the non-robust estimate is pulled towards them.

3. THE EROSION-PROPAGATION (EP) ALGORITHM

To select the potential edge pixels and provide an initial grouping of them into floe outlines, we use the EP algorithm. This operates on binary images. However, images of ice floes such as Figure 1 are usually greyscale. The marginal distribution of pixel intensities is highly bimodal, and so we work with the simpler binary image obtained by thresholding the original image; see Figure 9. The final result is relatively insensitive to the precise choice of threshold.

The erosion part of the EP algorithm, which identifies the potential edge elements, is a standard application of ideas in mathematical morphology (Serra, 1982). The propagation part of the EP algorithm keeps track of the floe to which an edge pixel belongs by locally propagating the information about edge elements into the interior of the floe as it is eroded. This is facilitated by the object-oriented programming environment.

The algorithm is iterative and operates on a binary image consisting of figures (ice floes) on a contrasting background (water). At the first iteration, if a pixel is ice and a specified subset of its neighbors is water, the pixel is "melted" and becomes water, so that the figure to which it belongs is eroded. In our implementation, a pixel is "melted" if any of the eight neighboring pixels is water. At the second iteration, the same operation is performed on the image resulting from the first iteration, and so on. This can be formally described in terms of structuring elements using the terminology of mathematical morphology (Banfield 1988). Any edge locating operator can be used to provide the initial set of potential edge pixels. We use the pixels "melted" at the first iteration of the EP algorithm.

Some results are shown in Figure 10. We can control the minimum size of the floes by waiting until a specified number of iterations, $i_{\min}$, have passed before recording a floe. The smallest floe which can be recorded is then a square of side $(2i_{\min} + 1)$ pixels. Smaller floes "melt" and are not recorded.

The idea of the propagation part of the EP algorithm is that the locations of the edge pixels are propagated towards the interior of the figure as it is eroded. At the end of the process, a single interior point of the figure will "know" the locations of all the edge pixels to which it corresponds. The location information is passed to only a few pixels which are taken as far from the eroded pixel as possible subject to them not belonging to a different floe. This




Figure 9. Binary version of Figure 1 for two threshold levels. The results are similar. However, (A) has a lower threshold level than (B) and therefore has more clutter in the water but less noise interior to the floes.




Figure 10. The results of applying the EP algorithm to Figure 9(A) after (A) 3 iterations, (B) 7 iterations and (C) 12 iterations.

ensures that the amount of location information to be processed does not become unmanageable. It also prevents loss of information due to irregularity of the floe, melt ponds, or pixel misclassification at the thresholding stage.

The key to the propagation part of the EP algorithm is that it is never necessary for any pixel to know the direction of the interior of the floe. If a pixel is eroded at the i th iteration, then it was the center of a $(2i + 1)$ by $(2i + 1)$ square of ice pixels before the start of the erosion process. Thus its location may be passed anywhere within the square of side $(2i + 1)$ pixels surrounding it without any risk of the information being passed to another floe. We have found that it is enough to pass the location information in two opposite directions within the square. All of the eroded pixels are processed in exactly the same manner and it is this uniform processing that allows the algorithm to be implemented on parallel processing machines.

In Figure 11 we show the results of the EP algorithm applied to the data in Figure 1, as thresholded in Figure 9(B) with a minimum floe size of 15×15 pixels (i.e. 1.2 km. square). The results are reasonably good: of the 35 floes identified by the EP algorithm, 23 are "right" in the sense of being close to floes identified by careful manual digitization.

However, the EP algorithm tends to subdivide floes. This can occur when the floes are non-convex or when they have noise in the interior. Figure 12 shows an example of the non-convex case. As the floe shown in Figure 12(A) was eroded, the narrow middle section was pinched in and the floe was divided into two partial floes, shown in Figures 12(B, C). Figure 13(A) is an example of a floe with melt ponds that cause the EP algorithm to produce three partial floes, as seen in Figures 13(B, C, D).

4. CLUSTERING ABOUT CLOSED PRINCIPAL CURVES

The EP algorithm tends to subdivide floes. We have therefore developed a method for determining which of the floes identified by the EP algorithm should be merged, based on an algorithm for clustering about closed principal curves. Since we want to find out whether to merge tentatively identified floes, this is hierarchical and agglomerative.







Figure 11. Result of the EP algorithm applied to Figure 9(A), where floes are not recorded unless they have survived at least 7 iterations. This corresponds to a minimum floe size of 15×15 pixels, or 1.2 km square. The points are the edge elements identified by the EP algorithm and the numbers interior to each floe are the centers found by the EP algorithm. Note that centers 1 and 4 are on the same floe, which was subdivided because of the melt ponds. Other floes were also subdivided.



A




B




C


Figure 12. When a non-convex floe (A) is eroded the narrow region can be "pinched off" resulting in two partial floes (B) and (C).




A



B



C



D

Figure 13. Noise in the interior of a floe can erode outwards and cause the floe to be subdivided. In this case, the melt ponds in (A) caused the floe to be subdivided into three partial floes (B), (C) and (D)

The objective of cluster analysis is to group a set of observations into "interesting" subsets. In practice, this has usually meant grouping observations which are close to one another. Ward (1963) proposed a hierarchical agglomerative algorithm for dividing data into g groups such that the sum of the within-group sums of squares is minimized. The algorithm starts by assigning each observation to a separate group. At each agglomeration two groups are merged, chosen so as to minimize the increase in the sum of within-group sums of squares. This clustering criterion is optimal if the data are generated by a finite mixture of spherical Normal distributions. This corresponds to clusters which tend to be of the same size and spherical

Murtagh and Raftery (1984) proposed decomposing the within-group sum of squares into parts and using a weighted sum of the parts as the clustering criterion, with weights chosen so as to emphasize aspects of interest in the application. For two-dimensional data, they suggested decomposing the within-group sum of squares into parts parallel and perpendicular to the first principal component of the group and downweighting the parallel part. This criterion was generalized by Banfield and Raftery (1989) who also showed that it is optimal when the data are generated by a mixture of Normal distributions with covariance matrices whose eigenvalues are constant across clusters. This corresponds to clusters which tend to be elliptical with the same size and shape but different orientations.

We now apply the idea of decomposing and reweighting the within-group sum of squares to the present problem. The edge pixels for an ice floe that has not been subdivided should be

- (a) tightly clustered about the floe outline, as estimated by the principal curve, and
- (b) regularly spaced along the outline, so that the variance of the distances between neighboring edge pixels should be small; see Figure 14 and Figure 15.

We decompose the variance, V , for a group of edge pixels into parts corresponding to (a) and (b) and a residual part, as follows. Let the locations of the edge pixels be x_j ($j = 1, \dots, n$), ordered so that $\hat{f}(\lambda_j) = x_j$ and $\lambda_1 \leq \dots \leq \lambda_n$, where $\hat{f}$ is the principal curve of the group, estimated by the method of Section 2.2. Let $\epsilon_j = \hat{f}(\lambda_j) - \hat{f}(\lambda_{j+1})$ be the vector defined by two adjacent projections onto the estimated principal curve. Since we are working with closed curves, all subscripting is modulo n .




Figure 14. A complete floe, showing the edge pixels identified by the EP algorithm (open circles) and their projections onto the estimated principal curve (dark circles). The projections onto the principal curve are regularly spaced, and so the variance of the distance between adjacent projections is small.




Figure 15. A partial floe. There are large gaps between the projections of the edge pixels onto the estimated principal curve, and so the variance of the distance between adjacent projections is large.

We may now write the within group variance as

$$V = \sum_{j=1}^n \|x_j - \bar{x}\|^2$$

$$= V_{about} + V_{along} + V_{residual},$$

where $V_{about} = \sum_{j=1}^n \|x_j - \hat{f}(\lambda_j)\|^2$, and $V_{along} = 1/2 \sum_{j=1}^n \|\epsilon_j - \bar{\epsilon}\|^2$. V_{about} is a measure of lack of tightness of the distribution of the edge pixels *about* the floe outline, and thus of the extent to which requirement (a) is not satisfied. V_{along} is a measure of the lack of regularity of the distribution of the pixels *along* the floe outline, and thus of the extent to which requirement (b) is not satisfied. It is shown by Banfield (1988) that

$$V_{residual} = 1/2 \sum_{j=1}^n \|\epsilon_j\|^2 + \sum_{j=1}^n r_j \bullet r_{j+1},$$

where $r_j = \hat{f}(\lambda_j) - \frac{1}{n} \sum_{k=1}^n \hat{f}(\lambda_k)$ and $\bullet$ denotes the dot product. It follows that $V_{residual}$ is a measure of the overall size of the floe. The general reweighted within-group variance is thus

$$\alpha V_{about} + \beta V_{along} + \gamma V_{residual}.$$

$V_{residual}$ contains no information of interest to us here, so we set $\gamma=0$. Also, without loss of generality we set $\beta=1$. Our clustering criterion is therefore

$$V^* = \alpha V_{about} + V_{along}.$$

To determine whether a set of groups should be merged, we first calculate V^* for each of the individual groups and then we calculate V^* for the union of the groups. If the union has a smaller value of V^* than any of the individual groups we merge them, and otherwise not. To make the number of candidate mergers manageable, we note that if a floe has been subdivided the partial floes will have common edge pixels, but not conversely. Therefore, only mergers of groups with common edge pixels are considered.

To determine α , we note that by arguments similar to those of Banfield and Raftery (1989) V^* will be an optimal criterion, conditional on the estimated principal curves, if the edge pixels are normally distributed about the floe outlines and if $E[V_{along}] = \alpha E[V_{about}]$. We therefore estimate α as the average of V_{along} / V_{about} for the floes that we know were not subdivided, namely those which have no shared edge elements. Using a span of 0.3 to estimate the principal curves, this yielded $\hat{\alpha} = 0.39$.

Table 1 shows the results of merging the partial floes in Figure 11, and Table 2 shows the results of trying to merge floes that should not be merged. In each case our method gives a result which is both correct and clearcut. Figure 16 shows the final results of the procedure, together with the identified edge pixels.

5. DISCUSSION

We have developed an automatic method for finding the outlines of ice floes in satellite images. It is accurate and computationally efficient. It involves three new statistical techniques: a way of estimating closed principal curves that reduces both bias and variance and is robust, the EP algorithm, and a method for clustering about principal curves.

The approach would seem to be applicable more generally to the detection of non-linear features in images. It extends cluster analysis to the case where similar pixels tend to be grouped about arbitrary curved features, open or closed, using the idea of decomposing and reweighting the within-group sum of squares proposed by Murtagh and Raftery (1984). This suggests that cluster analysis may be useful for feature extraction in images more generally.

The procedure is implemented in an object-oriented programming environment. One of the advantages of this environment is that each floe resulting from the procedure can be represented as an instance of a "floe object" and can carry with it, in the form of instance variables, specific information about the floe to be used in further analysis. It is relatively fast: a 512x512 8-bit image can be analyzed in approximately one hour on a Symbolics 3600 and should be considerably faster on some of the newer workstations (e.g. about 20 minutes on a Sun-3/80 and under 10 minutes on a Sun SPARCstation 1). The processing time is linear in the number of pixels, but does depend upon the complexity of the image.




Figure 16. Principal curves of the floes found after using the clustering method to merge the partial floes found by the EP algorithm (solid curves). The circles are the edge elements found by the EP algorithm.

In our implementation of the EP algorithm we erode an ice pixel if any of its eight neighbors are water. Other rules for eroding a floe may be used and they can change the rate of erosion, the effect of pixel misclassification and the shape of the floes that can be found. The rules may also be changed as the erosion proceeds. For example, dilating the image, or "refreezing" the ice, at an early iteration could eliminate some pixel misclassifications and other noise in the interior of the floes. The EP algorithm has the potential of being implemented on parallel processing machines.

To date, the development of automated techniques for the analysis of polar satellite images has been limited to ice floe tracking (Ninnis, Emery and Collins 1986; Fily and Rothrock 1986, 1987; Vesecky, Samadani, Smith, Daida and Bracewell 1988). The primary tool in these automated tracking methods is cross-correlation, which provides the ability to match regions in two different images, but does not give any information about the morphology of the individual ice floes or the spatial structure of the ice pack. Vesecky *et al.* (1988) use segments of ice floe boundaries to track ice floe movements, but this does not provide the type of information needed to study ice floe morphology and spatial structure. The need for more information on both morphology and spatial structure was clearly shown by the 1984 Marginal Ice Zone Experiment (Burns *et al.* 1987; Campbell *et al.* 1987).

The problem considered here does not fall neatly into one of the problem areas in image understanding that have been intensely studied in recent years, namely *image restoration*, *classification*, *segmentation* and *feature extraction*. It does, however, combine elements from all of them. Image restoration attempts to reconstruct a degraded image. Image classification tries to assign each pixel to one of several predetermined categories; it may be regarded as a special case of restoration. Image segmentation (Rosenfeld and Kak, 1982; Chellappa and Sawchuk, 1985) seeks to identify areas of contiguous pixels that are, for example, devoted to the same crop. The aim of feature extraction is to find linear or curvilinear features in images.

Our problem shares goals with feature extraction, segmentation and classification. While restoration is not an explicit goal, we would expect the methods developed here to work well in the presence of degradation. Restoration and classification methods do not, by themselves, address the present problem. Current feature detectors would seem to have difficulty locating features as arbitrary as ice floes. For example, the Hough transform (Hough 1962; Ballard

Table 1. *Floe merger results. This table shows the values of the merger criterion V^* for the subdivided floes from Figure 11. The value of the criterion for the individual partial floes is shown under the "Individual" column. The value under the "Merged" column is the criterion value when the floes in the left column are merged. The floe numbers refer to the center numbers in Figure 11. The fact that the value under the "Merged" column is smaller than any of the values under the "Individual" column indicates that the floes should be merged.*

Floes	Criterion Values	
	Individual	Merged
Floe 1	11.36	
Floe 4	5.71	0.64
Floe 5	4.07	
Floe 8	10.86	2.36
Floe 12	8.29	
Floe 2	1.32	
Floe 3	1.36	.86
Floe 10	6.95	
Floe 6	4.37	
Floe 11	3.71	.60
Floe 25	1.33	
Floe 34	1.90	.33

Table 2. *Floe non-merger results. This table shows the values of the merger criterion V^* for non-subdivided floes from Figure 11. The value of the criterion for the individual partial floes is shown under the "Individual" column. The value under the "Merged" column is the criterion value when the floes in the left column are merged. The floe numbers refer to the center numbers in Figure 11. The fact that the value under the "Merged" column is larger than any of the values under the "Individual" column indicates that the floes should not be merged.*

Floes	Criterion Values	
	Individual	Merged
Floe 27	.24	
Floe 13	.39	10.70
Floe 34	1.90	
Floe 25	1.33	7.41
Floe 7	.20	
Floe 30	.59	
Floe 31	.93	2.27

1981; Davis 1982), an obvious candidate for locating closed curves, requires an initial pattern description which it then tries to find in the image. It would be difficult to provide an initial pattern description that is general enough to accommodate the wide range of commonly found ice floe shapes. Our approach may also be applicable to segmentation problems, especially those concerned with identifying not only regions but also the shapes of their outlines.

The Bayesian and stochastic relaxation approach of Geman and Geman (1984) may well be applicable to the present problem, although it has to date been used mainly for restoration. It would require extensive modeling assumptions for the ice floe problem, and experience

suggests that it would be computationally expensive. The computational burden might be reduced by using as an approximation the ICM algorithm of Besag (1986), although this does not yet appear to have been applied to problems such as the present one. Our procedure on the other hand requires only the assumption that the ice floe boundaries be closed curves, and it is relatively fast.

REFERENCES

- Ballard, D.H., (1981) "Generalizing the Hough Transform to Detect Arbitrary Shapes", *Pattern Recognition*, 13, 111-112.
- Banfield, J. (1988), "Constrained Cluster Analysis and Image Understanding," Ph.D. Dissertation, Department of Statistics, University of Washington.
- Banfield, J.D., and Raftery, A.E. (1989), "Model-based Gaussian and non-Gaussian clustering," Unpublished manuscript.
- Besag, J. (1986), "On the statistical analysis of dirty pictures (with Discussion)," *Journal of the Royal Statistical Society, series B*, 48, 259-302.
- Burns, B.A., Cavaleri, D.J., Keller, M.r., Campbell, W.J., Grenfell, T.C., Maykut, G.A. and Gloersen, P. (1987), "Multisensor Comparison of Ice Concentration Estimates in the Marginal Ice Zone", *Journal of Geophysical Research*, 92, 6843-6856
- Burns, J., Hanson, A., Riseman, E. (1986) "Extracting Straight Lines," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8, 425-455
- Campbell, W.J., Gloersen, P., Josberger, E.G., Johannessen, P.S., Guest, P.S., Mognard, N., Shuchman, R., Burns, B.A., Lannelongue, N. and Davidson, K.L. (1987), "Variations of Mesoscale and Large Scale Sea Ice Morphology in the 1984 Marginal Ice Zone Experiment as Observed by Microwave Remote Sensing," *Journal of Geophysical Research*, 92, 6805-6824.
- Chellappa, R. and Sawchuk, A. (1982) *Digital Image Processing and Analysis* (2 vols.), IEEE Computer Society Press.
- Committee on Applied and Theoretical Statistics (1989), "Discriminant Analysis and Clustering," *Statistical Science*, 4, 34-69.

- Davis, L.S. (1982), "Hierarchical generalized Hough transforms and line-segment based generalized Hough transforms," *Pattern Recognition*, **15**, 277-285.
- Fily, M. and Rothrock, D. A. (1986), "Extracting Sea Ice Data from Satellite SAR Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, **24**, 849-854.
- Fily, M. and Rothrock, D. A. (1987), "Sea Ice Tracking by Nested Correlations," *IEEE Transactions on Geoscience and Remote Sensing*, **25**, 570-580.
- Geman, S., and Geman, D. (1984), "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **6**, 721-741.
- Gordon, A.D. (1981), *Classification*, London: Chapman and Hall.
- Gordon, A.D. (1987), "A review of hierarchical classification," *Journal of the Royal Statistical Society, series A*, **150**, 119-137.
- Hastie, T. (1984), "Principal Curves and Surfaces" Ph.D. Dissertation, Department of Statistics, Stanford University.
- Hastie, T., and Stuetzle, W.X. (1985), "Principal Curves and Surfaces," Technical Report No. 56, Department of Statistics, University of Washington.
- Hastie, T., and Stuetzle, W.X. (1988), "Principal curves," *Journal of the American Statistical Association*, **84**, 502-516.
- Hough, P.V.C. (1962), "Method and mean for recognizing complex patterns," U.S. Patent 3069654.
- Matheron, G. (1975), *Random Sets and Integral Geometry*, New York: Wiley.
- Murtagh, F. (1985), *Multidimensional Clustering Algorithms*, COMPSTAT Lectures 4, Vienna: Physica-Verlag.
- Murtagh, F., and Raftery, A.E. (1984), "Fitting straight lines to point patterns," *Pattern Recognition*, **17**, 479-483.
- Ninnis, R. M., Emery, W. J., and Collins, M. J. (1986), "Automated Extraction of Pack Ice Motion From Advanced Very High Resolution Radiometer Imagery," *Journal of Geophysical Research*, **91**, 10725-10734.

- Rosenfeld, A., and Kak, A.C. (1982), *Digital Picture Processing* (2 vols.), New York: Academic Press.
- Rothrock, D., and Thorndike, A. (1984), "Measuring the Sea Ice Floe Size Distribution," *Journal of Geophysical Research*, **89**, 6477-6486.
- Serra, J. (1982) *Image Analysis and Mathematical Morphology*, New York: Academic Press.
- Vesecky, J. F., Samadani, R., Smith, M. P., Daida, J. M. and Bracewell, R. N. (1988), "Observation of Sea-Ice Dynamics Using Synthetic Aperture Radar Images: Automated Analysis," *IEEE Transactions on Geoscience and Remote Sensing*, **26**, 38-48.
- Ward, J.H. (1963), "Hierarchical grouping to optimize an objective function," *Journal of the American Statistical Association*, **58**, 236-244.

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 172	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Ice Floe Identification in Satellite Images Using Mathematical Morphology and Clustering about Principal Curves		5. TYPE OF REPORT & PERIOD COVERED Technical Report 9/1/88-8/31/89
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Jeffrey D. Banfield Adrian E. Raftery		8. CONTRACT OR GRANT NUMBER(s) N-00014-88-k-0265
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Statistics, GN-22 University of Washington Seattle, WA 98195		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR-661-003
11. CONTROLLING OFFICE NAME AND ADDRESS ONR Code N63374 1107 NE 45th Street Seattle, WA 98105		12. REPORT DATE August 1989
		13. NUMBER OF PAGES 32
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report)
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Erosion; Feature extraction; LANDSAT: Non-parametric curves; Object-oriented programming; Robustness;		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Identification of ice floes and their outlines in satellite images is important for understanding physical processes in the polar regions, for transportation in ice-covered seas and for the design of offshore structures intended to survive in the presence of ice. At present this is done manually, a long and tedious process which precludes full use of the great volume of relevant images now available. (cont. on next page)		

We describe an automatic and accurate method for identifying ice floes and their outlines. Floe outlines are modeled as closed principal curves (Hastie and Stuetzle, 1989), a flexible class of smooth non-parametric curves. We propose a robust method of estimating closed principal curves which reduces both bias and variance. Initial estimates of floe outlines come from mathematical morphology with local propagation of information about floe edges.

The edge pixels from the EP algorithm are grouped into floe outlines using a new clustering algorithm. This extends existing clustering methods by allowing groups to be centered about principal curves rather than points or lines. This may open the way to efficient feature extraction using cluster analysis in images more generally. The method is implemented in an object-oriented programming environment for which it is well suited, and is quite computationally efficient.