The University of Kansas

Information and Telecommunication Technology Center

A Technical Report of the
Networking and Distributed Systems (NDS) Laboratory

Automated Multiyear Ice Extraction Using Multiresolution Tracking and Binary Clustering

LeenKiat Soh
Costas Tsatsoulis

ITTC-FY98-TR-11810-03

April 1998

Project Sponsor:

Naval Research Laboratory

Automated Multiyear Ice Extraction Using Multiresolution Tracking and Binary Clustering

ABSTRACT

This technical report presents an algorithm that extracts multiyear ice class from Synthetic Aperture Radar (SAR) sea ice imagery using image processing and machine learning techniques. Multiyear ice is ice which has survived more than one year of freeze-thaw cycle, and its extraction is important for calculating ice coverage percentage, ice growth and trend analysis in the polar regions for global climatology and marine navigation. Our technique combines multiresolution tracking and binary clustering to identify the multiyear ice class in an image automatically. The clustering algorithm is called the Aggressive Consolidation of Aggregated Populations, which identifies the most dominant class or population in a spatial environment and consolidates that class by merging other populations with it. This clustering technique assumes there are only two classes of importance: the target class (multiyear ice in our case) and the non-target class. We have applied our technique to more than 50 SAR sea ice images, including those of ERS-1, ERS-2, and RADARSAT. In an evaluation session by sea ice geophysicists, our technique was judged to be able to extract multiyear ice class from a variety of sea ice images of different geographical locations and seasons exceptionally well.

1 INTRODUCTION

The extraction of multiyear ice is important for both global climatology and polar navigation issues. Any significant change in multiyear ice coverage percentage of the polar regions indicate possible changes in global climate and oceanic activities, especially in long-term trend analysis. This is because multiyear ice is ice that has survived several freeze-thaw cycles and has occupied the polar regions for a long period of time. Also, multiyear ice has consistent backscatter or intensity values which lead to more reliable identification and classification than one can do with, for example, new or young ice. This feature encourages geophysicists to study multiyear ice. For operational purposes, polar navigators prefer to plan routes to avoid multiyear ice due to the thickness of the ice that might cause problems for ice breakers. Therefore, our technique, called AMIE for Automated Multiyear Ice Extraction, has a real-world application and is a helpful tool in sea ice analysis.

The basic segmentation structure of AMIE is the same as ASIS (Soh and Tsatsoul 1997b). AMIE uses dynamic local thresholding and Multiresolution Peak Detection to generate a preliminary segmentation of the original image. Then, instead of applying the Aggregated Population Equalization (APE) spatial clustering (Soh and Tsatsoul 1997b) to determine the number of classes, we assume two classes: multiyear ice class and non-multiyear ice class. Subsequently, we have designed a binary clustering algorithm called the *A*ggressive *C*onsolidation of *A*ggregated *P*opulation (ACAP). This algorithm locates the dominant class in the image and consolidates the class by aggressively recruiting members of other classes; members not consolidated form another class. This algorithm is a variant of the APE spatial clustering; however, we do not strive for an equilibrium of strengths among populations in our spatial environment in this technique; instead, we promote the strengthening of a single population to dominate the environment.

In what follows, we first discuss the underlying methodology of our approach in multiyear ice extraction. Then, we describe the implementation choices and issues of the approach. Finally, we present the results and conclude the report.

2 METHODOLOGY

The underlying methodology is based on an unsupervised clustering concept outlined in detail in (Soh and Tsatsoul 1997b). The idea is to first create N number of classes or populations in the image, creating a preliminary segmentation. Given these classes or populations, we identify the most dominant class and apply our binary clustering technique, called Aggressive Consolidation of Aggregated Populations (ACAP), to recruit as many as possible other populations to join this dominant class. This integration process results in two classes: the multiyear ice (MY) class and the non-MY class. To obtain the initial segmentation, we use dynamic local thresholding and Multiresolution Peak Detection (MRPD) to derive a set of significant thresholds from the image. Given the set of values, we apply interpolations to segment the image into different classes, and we build a spatial environment that relates the populations based on spatial characteristics derived from the classes. The binary clustering algorithm operates on the spatial environment to finally obtain the multiyear ice class. Figure 1 shows the program modules of AMIE. As we can see, AMIE also has a feature that computes the percentage of multiyear ice coverage.

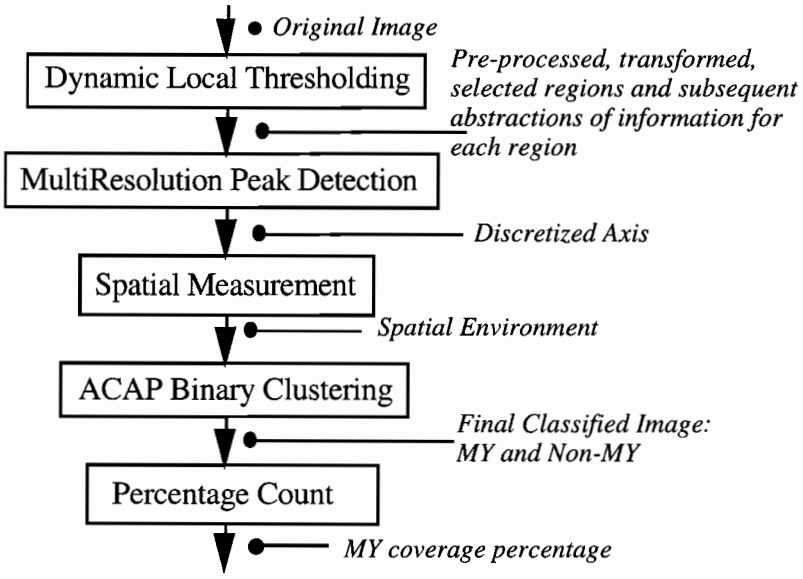

Figure 1. The block diagram of AMIE.

3 Implementation of AMIE

3.1. Dynamic Local Thresholding

The backbone of our implementation is dynamic local thresholding. It has been well documented in (Haverkamp *et al.* 1995, Soh and Tsatsoulis 1997a) for applications in SAR sea ice classification and segmentation in remotely-sensed imagery, respectively. Thus, in this report, we do not go into details of dynamic local thresholding. The underlying concept behind this technique is that by dividing the image into smaller regions for individual analysis, local features can be identified, while the interpolation scheme used to coordinate the local results together preserves the global information. Briefly, the steps of dynamic local thresholding are (1) Divide the image into smaller, overlapping regions, (2) Compute the histogram of each region, (3) Compute the variance of each histogram, and use a filter to qualify high-variance regions for further processing, (4) Perform a bimodal Gaussian curve approximation for qualified histograms, (5) Given the parameters of the approximation, compute the optimal bisector of the curve using maximum likelihood, (6) Compute the peak-to-valley ratio of each curve and qualifies those of high ratio values, and (7) From qualified curves, group its optimal bisector into the set of threshold values. A set of thresholds are detected as a result of dynamic local thresholding; however, instead of grouping these values into a pre-specified number of threshold classes, we perform multiresolution significant threshold selection. Note that after significant threshold selection, the threshold values are spread across the whole image through regional and pointwise interpolations to ensure a complete segmentation.

3.2. Multiresolution Significant Threshold Selection

Our new significant threshold selection strategy is based on data reduction adapted from (Sezan 1990) and it detects peaks in the histogram. Multiresolution has been used in image processing to tackle the issues of granularity (Eklundh and Rosenfeld 1979, Rosenfeld 1984, Bergholm 1987). When an image is blurred to a lower resolution, features that survive must have been significant. Thus, the occurrences of such features can be used to inform us of their existence. On the other hand, when an image is of a high resolution, the localization of features is accurate. Thus, by following the occurrences of a feature from low to high resolutions, we can determine its existence and localization accurately (Kakarala and Hero 1992, Jackway and Deriche 1996). This is the underlying concept of multiresolution. Our histogram of threshold values is a 1-D signal that has peaks and valleys. The objective of this module is to extract, from the signal, a set of significant peaks as thresholds for segmenting the image. To handle noise effects, we have employed the concept of multiresolution for its advantages as discussed.

A. Background

The basis of our technique is the utilization of the cumulative distribution function (cdf) of the histogram to detect peaks, which has been proposed to eliminate noise effects in histogram manipulation. For example, Boukharouba *et al.* (1985) located peaks using the zero-crossings and local extrema of a peak detection signal determined from the curvature of the cdf. To obtain the peak detection signal, the authors fitted the cdf to a Chebyshev polynomial. Subsequently, Sezan (1990) refined the method to lessen computational intensity and reduce the number of user-specified parameters. As reasoned in (Sezan 1990), the concept of using the cumulative distribution function to detect peaks is as follows. First, generate a peak detection signal from the image histogram. Then, locate the histogram peaks using the zero-crossings of the peak detection signal and the local extrema between the zero-crossings. To obtain the peak detection signal, the histogram is convolved with the peak detection kernel with a window size indicator called Ω , the *peak-detection parameter*. This parameter is particularly important in determining the sensitivity of the algorithm to noisy peaks in the histogram. Sezan (1990) further postulated several principles to estimate the *start*, *end*, and *maximum* points of the peaks as the following.

- A zero-crossing (similar to that of Laplacian second order derivative), of the peak detection signal, η_Ω, to negative values—a negative crossover—signifies the *start* of a peak. The bin value, ω, at which the negative crossover occurs is defined to be the estimate of a *start* point. Thus, the start point of the ith peak is ω_i^s; the next negative crossover at the bin ω_{i+1}^s estimates the start of the next peak, and so on.

- A zero-crossing of the peak detection signal, η_Ω, to positive values—a positive crossover—following a negative crossover estimates the bin values at which the peak reaches its maximum. The maximum of the ith peak is ω_i^m.

- The bin value between two successive negative crossovers, ω_i^s and ω_{i+1}^s, at which the detection signal, η_Ω, achieves its local maximum is defined to be the estimate of the end point of the peak. The end point of the ith peak is ω_i^e.

Practically, since the histogram is discrete and so is the detection signal, the zero-crossing of a signal can occur in two situations. If the signal attains the zero value, then the exact point of the zero-crossing can be obtained. For a crossover, the bin value of the leader of the transition becomes the crossover point location. Given these considerations, the ith peak can be represented by a triplet of $\langle \omega_i^s, \omega_i^m, \omega_i^e \rangle$.

B. Multiresolution Peak Detection (MRPD)

One inherent disadvantage of Sezan's algorithm is the determination of the peak detection parameter, Ω. If this number is small, then the algorithm responds to local variations and noise. On the other hand, if this number is large, then the algorithm may overlook legitimate peaks because of the averaging effect. We propose a new method called the Multiresolution Peak Detection (MRPD) to improve Sezan's algorithm. Image analysis at a fine resolution yields noise and unnecessary detail and at a coarse level distorts local deviations. Therefore, multiple scales of resolution are useful to counter these two problems (Rosenfeld 1984). Such advantages can also be obtained when applied to 1-D signal-type data such as a histogram as in our case.

The concept underlying our approach is as follows: First, we create a multiresolution map (or a scale space (Witkin 1983, 1984)) of the histogram peaks. Second, we track each peak through the scale space to assess each peak's significance. Third, we filter out spurious peaks and localize significant peaks. The essential part of this concept lies with the assessment of peak significance—we use a weighted neighborhood to fuse together different scales of information. The MRPD algorithm has been designed based on several assumptions:

- Peaks found at a low-level resolution are more significant than the peaks found at a high-level resolution. This concurs with the normal usage of multiresolution in the literature.

- Peaks found in high-level resolution are more accurate in terms of localization than the peaks found in low-level resolution. This also concurs with the concept of multiresolution.

- Neighboring peaks suggest a dominant peak. A peak is more significant if it is surrounded (in a 2- or higher-dimensional space) or flanked (in 1D space) by other peaks than a peak without such neighbors. Viewing this assumption from the noise modeling angle, a single peak is more likely to be a noisy assertion than a cluster of peaks.

- The significance of a peak is proportional to its height.

With the above assumptions, we have designed a multiresolution peak detection algorithm, with its block diagram illustrated in Figure 2. The first module generates the cumulative distributed function from the histogram. Then, the second module computes the maximum peak detection parameter, Ω_{max}, such that a multiresolution map can be created using $\Omega = 3,5,7,\ldots,\Omega_{max}$. The fourth module tracks the peaks to assess their significance. The tracking is performed by first identifying all peaks at each resolution level. Then, each peak's weight is modified based on its own height, its neighboring peaks, and its number of occurrences throughout the scale space. Peaks detected at lower resolutions are given higher scores than those at higher resolutions—this strategy complies with the objectives of multiresolution. At the end of the tracking, each peak accumulates a weight value that represents its significance. The fifth module merges adjacent peaks. This provision is installed to prevent a continuum of adjacent peaks from all becoming significant, which would result in erroneous segmentation because of these spurious peaks. In our definition, a significant peak is well defined and flanked by smaller peaks. If there exists a continuum in which all peaks are significant, that means a large peak has been detected and due to the tracking over scale space, its neighbors within the continuum have been exaggerated in significance. Hence, these neighbors must be purged. The merging is carried out by combining the weights of the adjacent peaks, removing all peaks except the tallest, and replacing the weight of that surviving peak with the sum of the weights of the continuum. Finally, a filter is used to extract a set of peaks to become the set of significant thresholds. Details can be found in (Soh 1997).

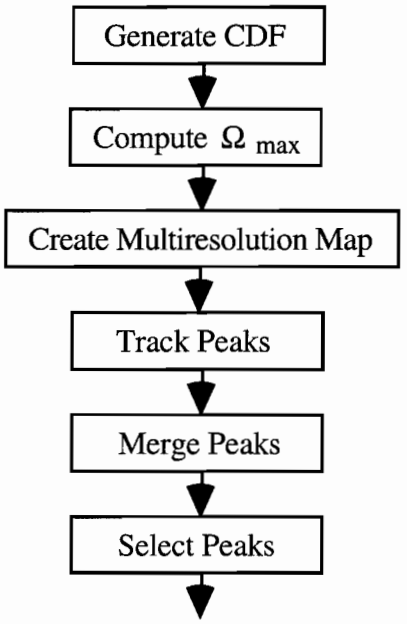

Figure 2. Block diagram of our multiresolution peak detection algorithm.

3.3. Aggressive Consolidation of Aggregated Populations (ACAP)

Similar to the APE spatial clustering in which we locate the dominant population, here we seek a multiyear core to which neighboring populations will be consolidated. Given a spatial environment, Π_S, we compute the potential of each population label, π_i, as a multiyear core, as

$$\varpi(\pi_i) = \#_{j=0}^{N_S}\left(SD(j,i) = \max_{k=0}^{N_S}\left(SD(j,k)\right)\right), \tag{1}$$

where SD, the spatial descriptor, is:

$$SD(i,j) = V_{T_{S(k-1)}(l-1)} = \frac{\sum_{i=1}^{\#pixels}\sum_{q\in\theta(p_i,1)}\xi\left(c(p_i),T_{S(k-1)}\right)*\xi\left(c(q),l-1\right)}{\sum_{i=1}^{\#pixels}\xi\left(c(p_i),T_{S(k-1)}\right)}, \tag{2}$$

where $\xi(a,b) = \{1 \quad if\ a=b; 0 \quad if\ a\neq b\}$, $c(p)$ is the label of pixel p, $i = T_{S(k-1)}$, a significant threshold, and $j = l-1$. So the descriptor describes how one population behaves with another population. The function of a population and itself is the probability of that population having itself as a neighbor in a defined neighborhood. The stronger the core of a population, the higher this probability is. Conversely, for a population that has a weak core, the probability of the population having itself as a neighbor is low. For a pair of interspersing populations with labels i

and j, $SD(i,j)$ and $SD(j,i)$ might be high. For a parasitic population with a label i at the fringes of another population j, $SD(i,i)$ should be low and $SD(i,j)$ should be high. Note that the spatial descriptor is non-commutative.

Equation (1) says that the potential of a population label, π_i, as a multiyear core is the number of occurrences where another population label has π_i as the most frequent neighbor. Our minimal requirement of a candidate for a multiyear core is $\varpi(\pi_i) \geq 2$. That is, a candidate must have at least one attached *parasitic* population class, assuming it already has itself as a parasitic population. To select the multiyear core from the list of candidates, we compare their core strengths. Suppose the list of candidates is

$$\mathrm{N}(\Pi_S) = \left\{ \pi_k \middle| \varpi(\pi_k) \geq 2 \right\}. \tag{3}$$

Define the multiyear core as $MY(\Pi_S)$. We have

$$MY(\Pi_S) = \pi_i \middle| SD(i,i) = \max_{\pi_k \in \mathrm{N}(\Pi_S)} SD(k,k). \tag{4}$$

Once we have obtained $MY(\Pi_S)$, we build the list of neighbors that will be consolidated to the multiyear core. This list of multiyear classes is defined as

$$my(\Pi_S) = \left\{ \pi_i \middle| SD\left(i, MY(\Pi_S)\right) = \max_{k=0}^{N_S} SD(i,k) \right\}. \tag{5}$$

Thus, every population label that has the multiyear core as its most frequent neighbor will become a member of $my(\Pi_S)$. This strategy assumes that every parasitic population attached to the multiyear core is a multiyear class. Each pixel of a parasitic population will become a multiyear ice pixel with a computed chance based on the degree of spatial reliance of the class on the multiyear core, as shown in a later section.

To more aggressively consolidate multiyear classes, we have introduced another parameter called the aggressivity, α_{SD}. For the two neighboring populations of $my(\Pi_S)$, we examine the following. Suppose we denote the neighboring population as $\pi_{neighbor}$. If

$$\left[\max_{j=0}^{N_S} SD\left(l\left(\pi_{neighbor}\right), j\right)\right] - SD\left(l\left(\pi_{neighbor}\right), MY\left(\Pi_S\right)\right) \leq \alpha_{SD}, \qquad (6)$$

then $\pi_{neighbor}$ will be added to $my\left(\Pi_S\right)$. This condition states that if the spatial reliance of a neighbor *A* on its most frequent neighbor is similar to the spatial reliance of that neighbor *A* on the multiyear core (less than α_{SD}), then the neighbor *A* becomes a multiyear class. This aggressive consolidation of multiyear class allows the multiyear core to overtake control of a neighboring population. The larger the value of α_{SD} is , the more aggressive the consolidation process is.

Our observation of multiyear ice in SAR sea ice imagery has inspired the design of ACAP. Normally, multiyear ice regions in an image form strong cores with parasitic classes occupying the fringes of these regions or as speckles within these regions. These parasitic classes may be a mixture of other ice classes and multiyear ice, due to overlapping of backscatter that occurs when various ice classes co-exist. At these *fringe* regions, one notice the change in intensity away from the homogeneity of multiyear ice regions. The change results in significant classes as extracted by the MRPD algorithm. Here, we use the ACAP to more aggressively regain the *multiyear ice classes lost to the fringe interference*. The aggressivity factor is now experimentally set at 0.025.

B. ACAP and RADARSAT Data

For the RADARSAT data, the observation of the multiyear ice class needs to be refined. The RADARSAT data is noisier in presentation than either the ERS-1 or the ERS-2 data. For the ERS-1 and ERS-2 data, the strongest class is almost always the $MY\left(\Pi_S\right)$ since the multiyear ice regions are highly homogeneous. For the RADARSAT data, because of the noise factor, that observation might not be true as often. To ensure that the multiyear core, $MY\left(\Pi_S\right)$, has a strong core and inclusive of other classes, we have added the following constraint. Note that for the RADARSAT data, the event of having a candidate (for multiyear core) with only one parasitic population and with a higher core strength than the true multiyear core (corrupted with noise) is more likely, compared to the ERS data. Thus, if the potential of the multiyear core found by the Equation (1) above is 2, and the multiyear core does not have the strongest core in the spatial environment, then redefine the multiyear core as the population that has the strongest core, $\pi_{\max}$:

$$\pi_{\max} = \pi_i | SD(i,i) \geq SD(j,j) \quad j = 0 \ldots N_S. \qquad (7)$$

This condition also assumes that a true multiyear core has more than one parasitic populations. In addition, if the potential of the multiyear core found by the Equation (1) above is larger than 2, it

might still be a weak class due to high noise in the RADARSAT data. The event of having a weak class with a respectably high core strength that has more than one parasitic populations is more likely to occur in a RADARSAT sea ice image than in an ERS image. Under noisy conditions, the strongest class, $\pi_{\max}$, might have $\varpi(\pi_{\max}) < 2$ because of the behavior of scattered populations. This prevents $\pi_{\max}$ from becoming a member of $\mathrm{N}(\Pi_S)$. In most cases, Equation (6) covers this constraint; however, when Equation (6) fails, we needed to be able to consolidate the current multiyear core class by merging it with the strongest population. As a remedy to counter such noise effects, we have imposed another constraint. If $\pi_{neighbor}$ is also $\pi_{\max}$, and $SD(MY(\Pi_S), MY(\Pi_S))$ is less than *compactness*, then $\pi_{neighbor}$ becomes a member of $my(\Pi_S)$. This condition is a variant of the aggressivity clause of Equation (6). The higher the value of *compactness* is, the more likely that the neighbor will be merged to the multiyear ice classes. Currently, the *compactness* factor is experimentally set at 0.500.

C. Non-Deterministic Conversion

Each pixel of a population class in the list of $my(\Pi_S)$ is converted to a multiyear ice pixel based on its membership class' degree of spatial reliance on the $MY(\Pi_S)$. Suppose the label of the pixel is $l\langle p\rangle = l(\pi_i)$. The probability of a conversion is defined as

$$prob_{conversion}\langle p\rangle \propto \frac{1}{SD(i,i)} \cdot \frac{1}{SD(MY(\Pi_S), MY(\Pi_S)) - SD(i, MY(\Pi_S))}. \qquad (8)$$

The first denominator acts as a normalization factor. This normalization factor also implies that the stronger the core of a class, the lower $prob_{conversion}\langle p\rangle$ is. In addition, the second denominator serves as the reliance factor. The higher the value of $SD(i, MY(\Pi_S))$ is, the greater the spatial reliance of the class π_i on the multiyear core is; as a result, the $prob_{conversion}\langle p\rangle$ is also higher. To implement the non-deterministic conversion, we compute the number of non-multiyear core pixels surrounding the current pixel. Then we generate a random number limited by the value of $prob_{conversion}\langle p\rangle$. If the random number is larger than the number of non-multiyear core pixels, then the pixel is converted. Logically, if the $prob_{conversion}\langle p\rangle$ is higher, then the probability of getting a larger random number is higher, leading to a higher conversion rate, as desired. In addition, if the pixel is surrounded by a higher number of non-multiyear core neighbors, then $prob_{conversion}\langle p\rangle$ is smaller.

4 RESULTS

To illustrate the working of AMIE, we present two examples. Figure 3 shows an original ERS-1 SAR image. Figure 4 shows the multiyear ice class found by AMIE. The multiyear ice coverage was 76.69%. Figure 5 shows a portion of a RADARSAT ScanSAR image. As we compare Figures 4 and 6, we can see that the later is more grainy and the regions are less homogeneous. The constraints imposed by AMIE to counter noise took effect. Figure 6 shows the result of AMIE. The multiyear ice coverage was 83.74%.

We have applied AMIE to more than 50 SAR sea ice images successfully. During an evaluation session at the National Snow and Ice Data Center (NSIDC) at Boulder, Colorado, based on the multiyear ice extraction results on 16 SAR sea ice images, AMIE was judged to be exceptionally well in handling various image characteristics due to differences in geographical locations and seasons. In addition, the fully automated implementation of AMIE allows mass processing of SAR sea ice images without human intervention.

5 CONCLUSIONS

We have presented a technique that extracts multiyear ice class from SAR sea ice imagery. This technique uses image processing and binary clustering methodologies. Specifically, AMIE employs multiresolution peak detection to obtain significant peaks in the image to create a preliminary segmentation. Further, AMIE computes a spatial environment of the segmentation classes or populations that describes the relationships among them. Finally, AMIE uses a binary clustering technique called the Aggressive Consolidation of Aggregated Populations to identify a dominant population in the environment and strengthen that population with parasitic neighbors. We have encountered implementation issues when handling ERS and RADARSAT data. As shown, constraints could be imposed conveniently to improve the consolidation strategy.

We have applied AMIE to more than 50 ERS-1, ERS-2, and RADARSAT images successfully. AMIE is fully automated once the two parameters, aggressivity and compactness, have been fine-tuned. Thus, AMIE on one hand allows customization and specification to accommodate other types of imagery data. AMIE has been evaluated by sea ice geophysicists to be an exceptional technique in multiyear ice extraction for SAR sea ice imagery.

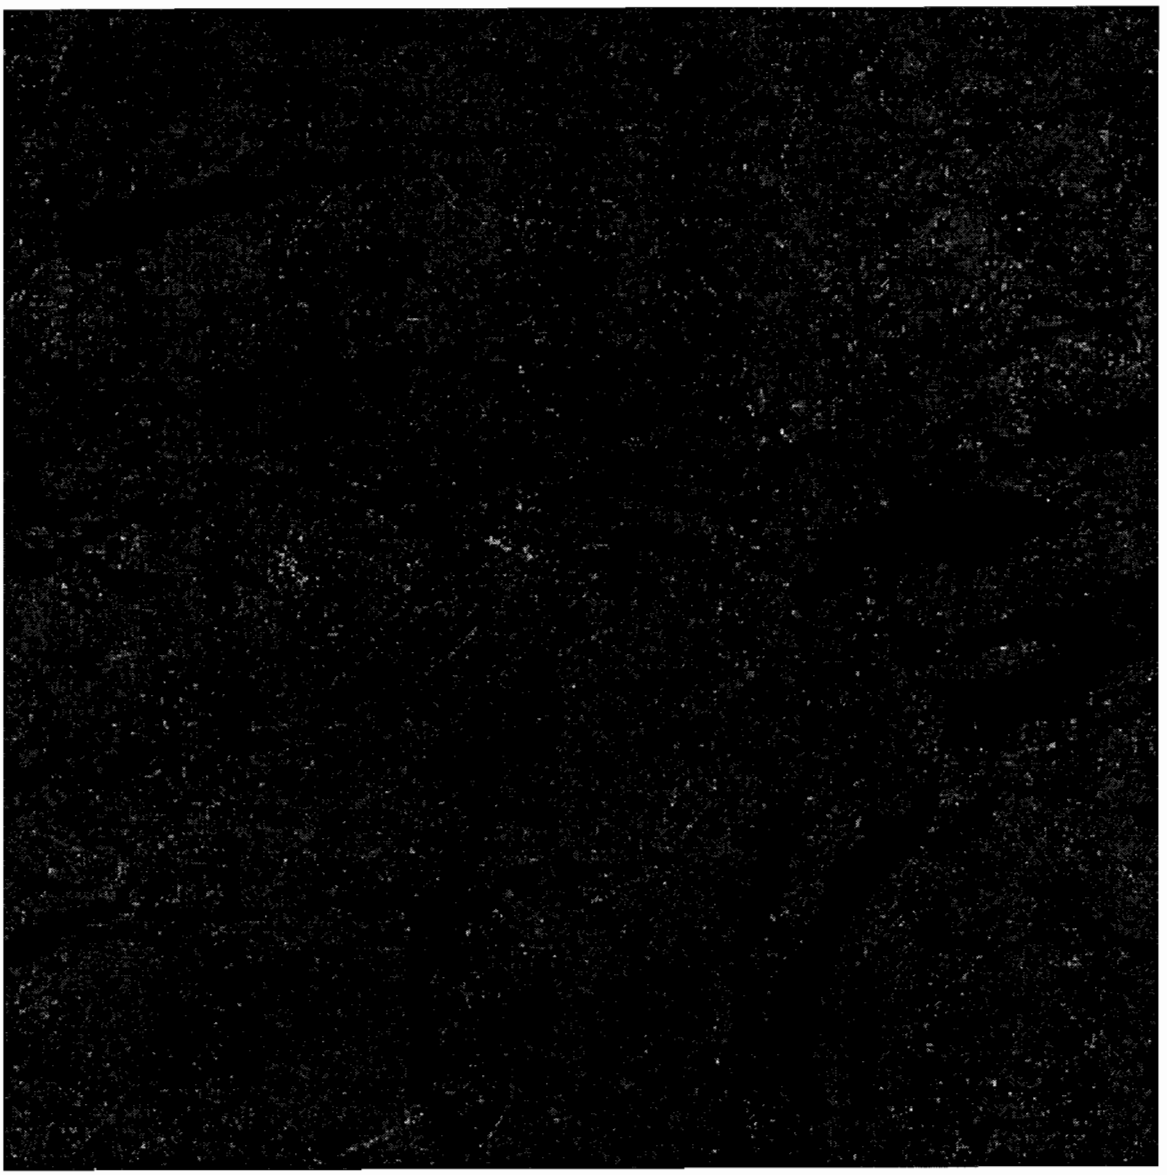

Figure 3. The original image of a quick-look processed ERS-1 SAR sea ice image. Copyright European Space Agency.

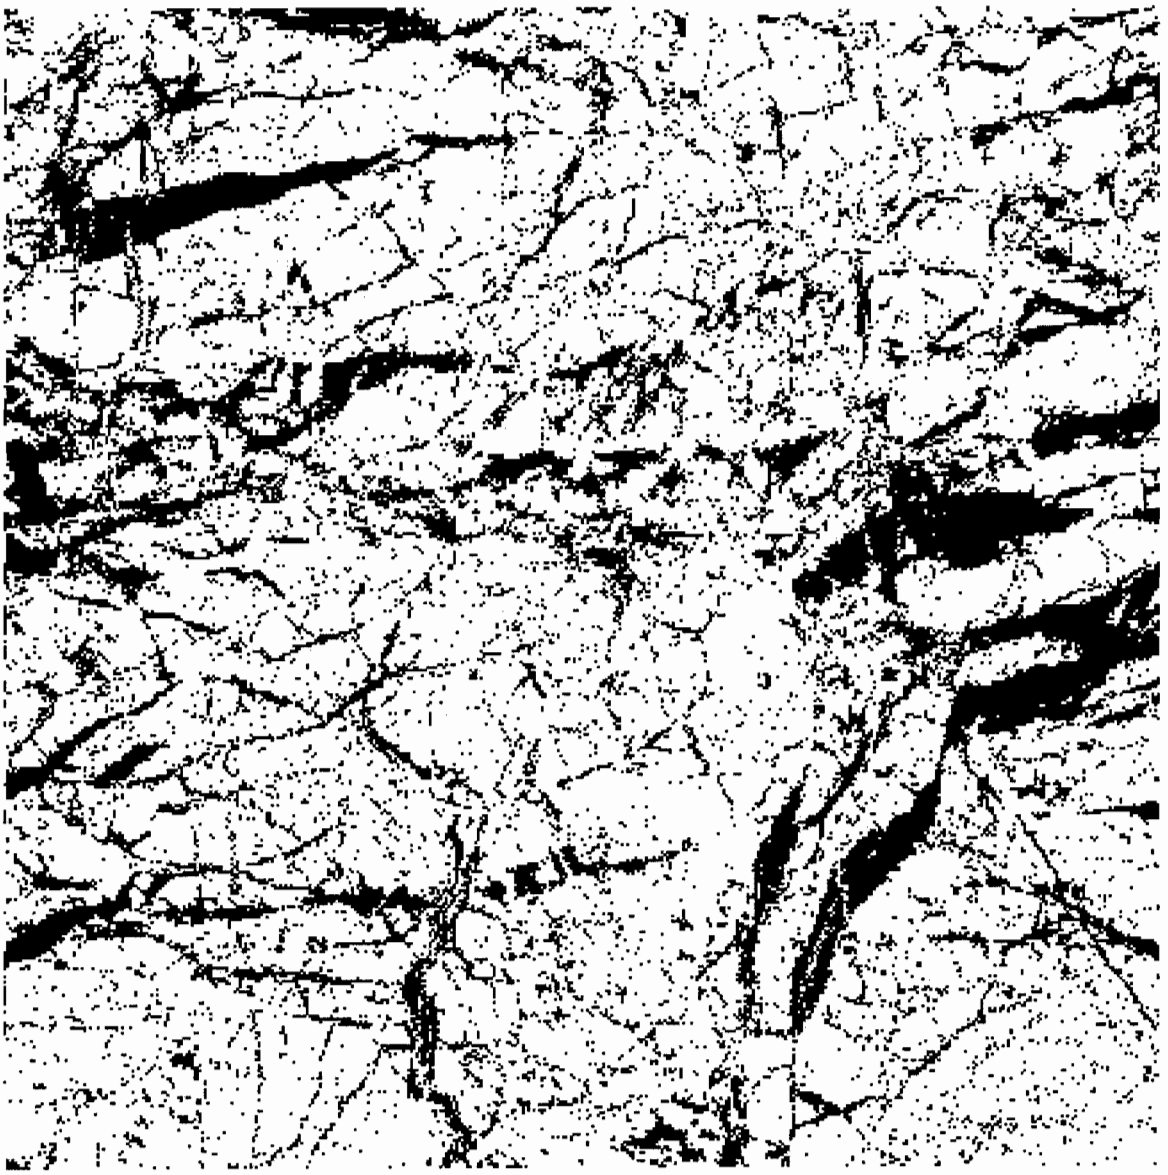

Figure 4. The AMIE-processed image. White pixels denote multiyear ice class; black denote non-multiyear ice class.

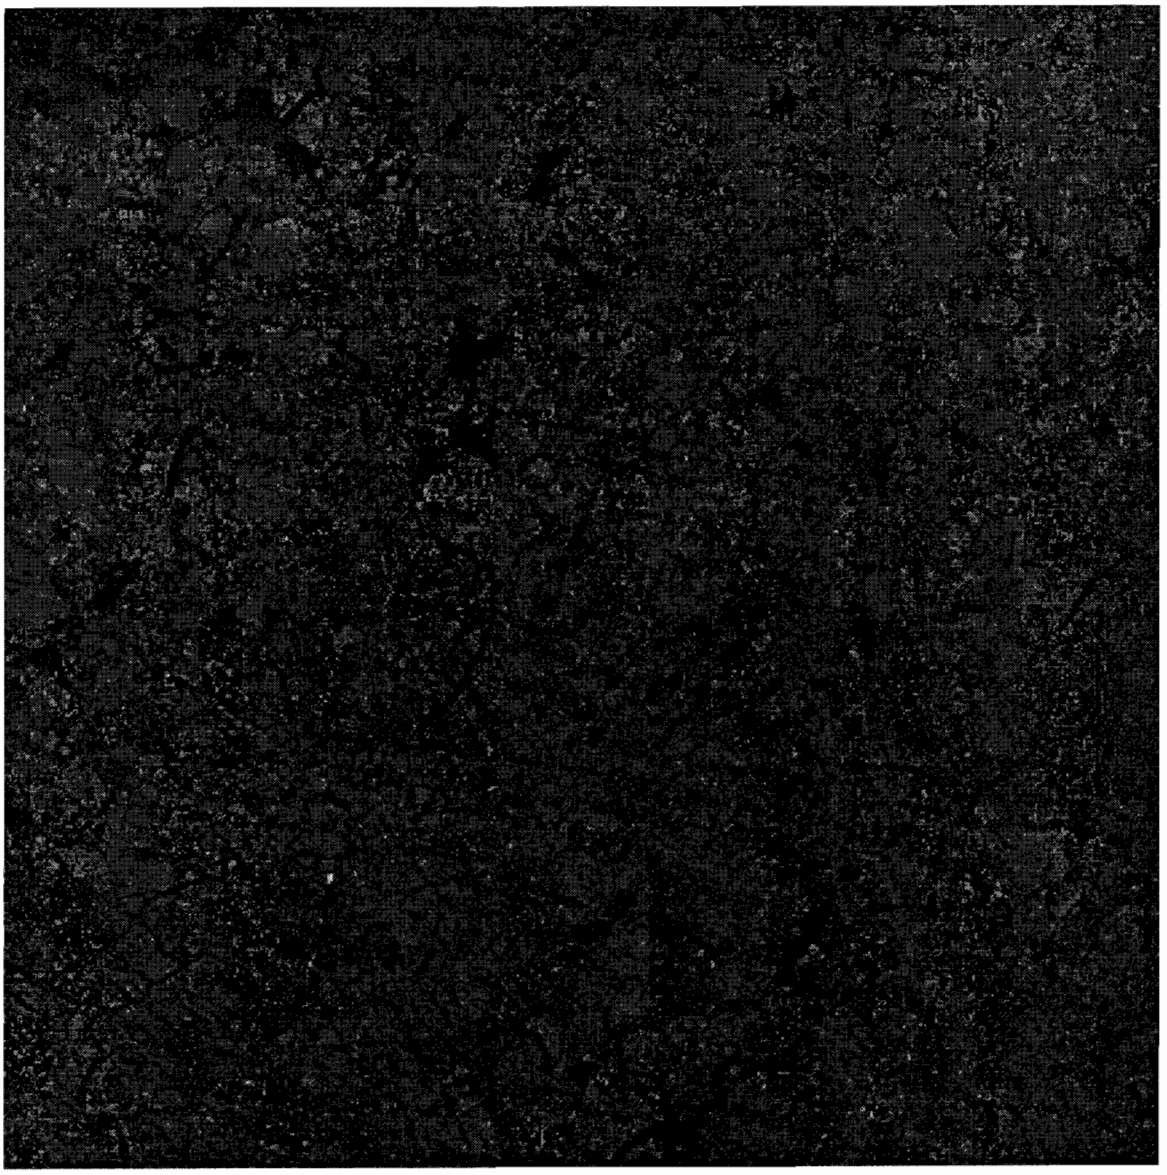

Figure 5. The original image of a RADARSAT ScanSAR sea ice image.
Copyright Canadian Ice Service.

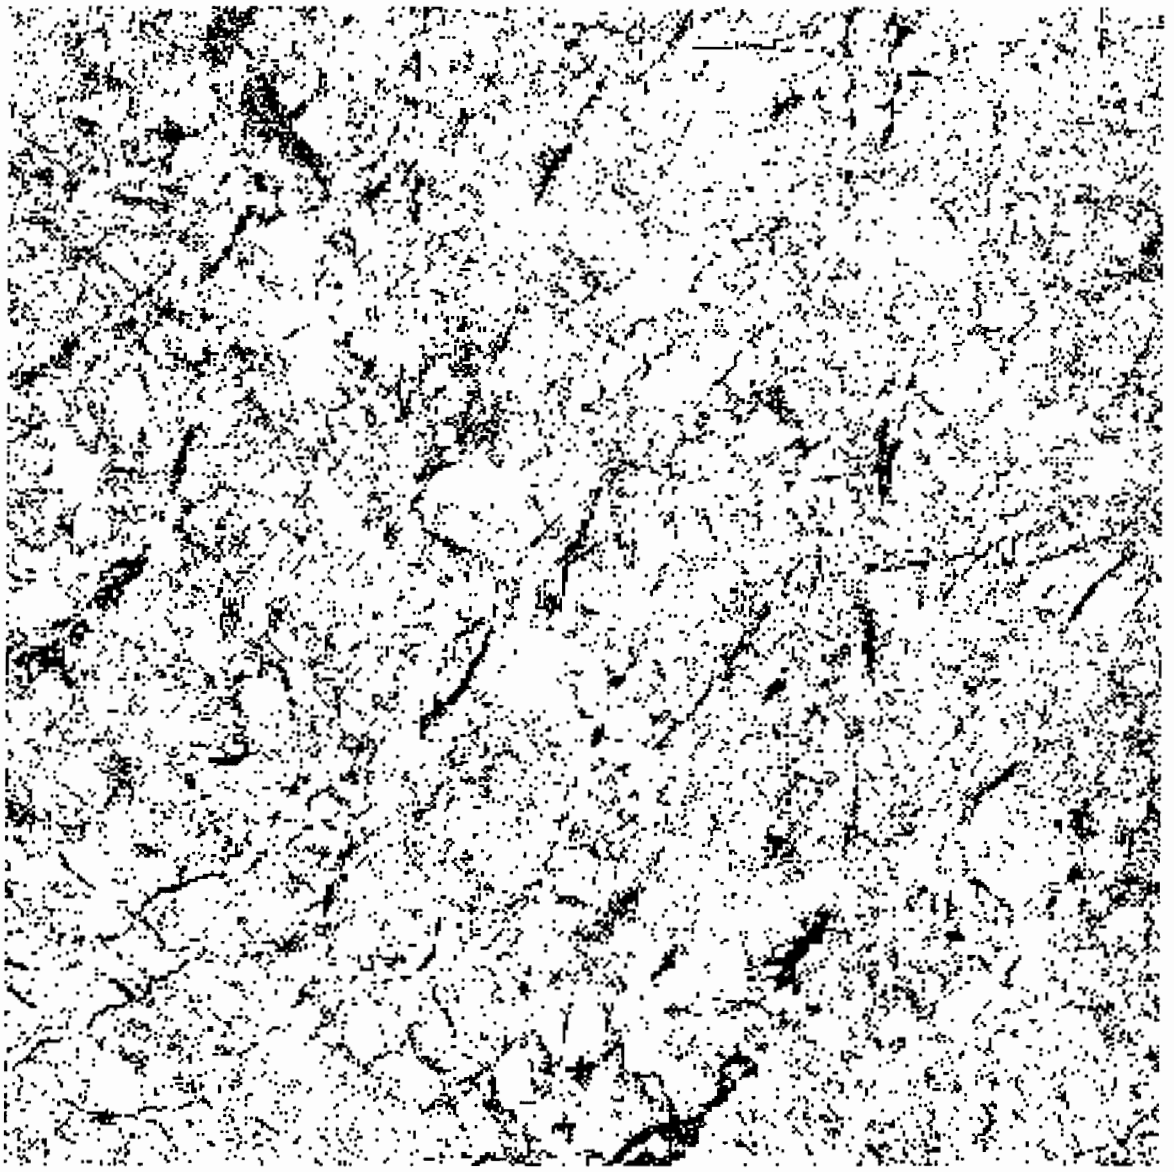

Figure 6. The AMIE-processed image. White pixels denote multiyear ice class; black denote non-multiyear ice class.

6 BIBLIOGRAPHY

Bergholm, F. (1987). Edge Focusing, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **9**(6), 726–741.

Boukharouba, S., J. M. Rebordao, and P. L. Wendel (1985). An Amplitude Segmentation Method Based on the Distribution Function of an Image, *Computer Vision, Graphics, and Image Processing*, **29**, 47–59.

Eklundh, J.-O. and A. Rosenfeld (1979). Peak Detection Using Difference Operators, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **1**(3), 317–325.

Haverkamp, D., L.-K. Soh, and C. Tsatsoulis (1995). A Comprehensive, Automated Approach to Determining Sea Ice Thickness from SAR Data, *IEEE Transactions on Geoscience and Remote Sensing*, **33**(1), 46-57.

P. T. Jackway and M. Deriche (1996). Scale-Space Projection of the Multiscale Morphological Dilation-Erosion, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18**(1), 38–51.

Kakarala, R. and A. O. Hero (1992). On Achievable Accuracy in Edge Localization, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **14**(7), 777–781.

Rosenfeld, A. (editor), (1984). *Multiresolution Image Processing and Analysis*, Berlin: Springer-Verlag.

Sezan, M. I. (1990). A Peak Detection Algorithm and Its Application to Histogram-Based Image Data Reduction, *Computer Vision, Graphics, and Image Processing*, **49**(1), 36–51.

Soh, L.-K. (1997). *Automated Image Segmentation: A Data Investigation Model and SAR Sea Ice Applications*, Ph.D. Dissertation, Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS.

Soh, L.-K. and C. Tsatsoulis (1997a). Segmentation for Satellite Imagery of Natural Scenes Using Data Mining, accepted for revision for publication by *IEEE Transactions on Geoscience and Remote Sensing*.

Soh, L.-K. and C. Tsatsoulis (1997b). Automated Sea Ice Segmentation Using Spatial Clustering, ITTC Technical Report *#ITTC-FY98-TR-11810-01*, University of Kansas, Lawrence.

Witkin, A. P. (1983). Scale-Space Filtering, in *Proceedings of IJCAI*, Karlsruhe, Germany, August, 1019–1022.

Witkin, A. P. (1984). Scale Space Filtering: A New Approach to Multi-Scale Description, in *Image Understanding*, edited by S. Ullman and W. Richards, New Jersey: Ablex Publishing, 79–95.