Naval Research Laboratory
Stennis Space Center, MS 39529-5004

NRL/MR/7240-93-7019

The Joint Ice Center SAR Workstation: Algorithm Evaluation

DENISE J. GINERIS
Sverdrup Technology
Stennis Space Center, MS

FLORENCE M. FETTERER
Remote Sensing Applications Branch
Remote Sensing Division

March 1993

DISCLAIMER NOTICE

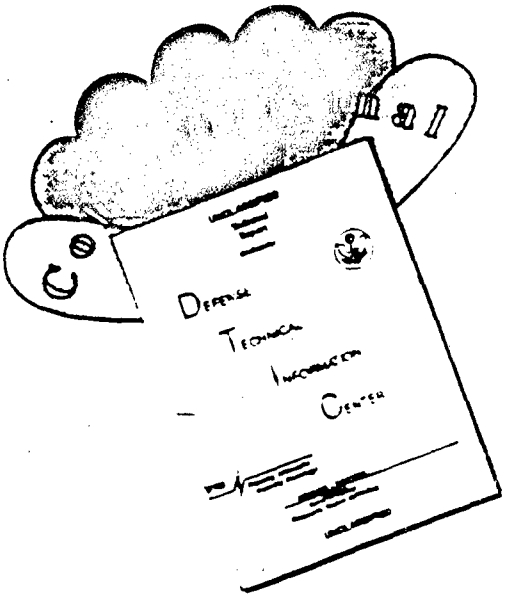

THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF COLOR PAGES WHICH DO NOT REPRODUCE LEGIBLY ON BLACK AND WHITE MICROFICHE.

REPORT DOCUMENTATION PAGE

Form Approved
OBM No. 0704-0188

1. Agency Use Only *(Leave blank).*	2. Report Date.	3. Report Type and Dates Covered.
	March 1993	Final

4. Title and Subtitle.

The Joint Ice Center SAR Workstation:
Algorithm Evaluation

5. Funding Numbers.

Program Element No. 0603704N
Project No.
Task No.
Accession No. DN258116
Work Unit No. 72519803

6. Author(s).

Denise J. Gineris* and Florence M. Fetterer

7. Performing Organization Name(s) and Address(es).

Naval Research Laboratory Detachment
Remote Sensing Applications Branch
Stennis Space Center, MS 39529-5004

8. Performing Organization Report Number.

Memorandum Report 7019

9. Sponsoring/Monitoring Agency Name(s) and Address(es).

Space and Naval Warfare Systems Command
2451 Crystal Drive
Arlington, VA 22202

10. Sponsoring/Monitoring Agency Report Number.

Memorandum Report 7019

11. Supplementary Notes.

12a. Distribution/Availability Statement.

Approved for public release; distribution is unlimited. Naval Research Laboratory, Washington, DC 20375-5320.

12b. Distribution Code.

13. Abstract *(Maximum 200 words).*

Great improvements in remote sensing can be achieved through the use of active sensors. The European Space Agency's synthetic aperture radar (SAR) on board the ERS-1 satellite is one example of such a sensor. Operating at 5.5 GHz, VV polarization, it provides excellent detection of and discrimination between types of arctic sea ice. Data from this satellite are downlinked to the Alaska SAR Facility (ASF) located at the University of Alaska in Fairbanks. The ASF in turn disseminates raw and processed SAR products to the scientific and operational communities. Among those users interested in SAR data for operational ice analysis is the Navy/NOAA Joint Ice Center (JIC). The JIC analyzes SAR data on a microcomputer-based workstation specifically designed by the Jet Propulsion Laboratory for sea ice analysis. The workstation includes software tools and algorithms for detailed ice classification, ice motion estimation, and lead and ridge analysis. Before the workstation can be used with confidence, however, it must first be tested for accuracy of analysis and ease of use. General workstation performance was evaluated over about a 6-month period (March to October 1992), while ice algorithm performance was evaluated by comparing SAR results to results from other sensors over a much shorter time period. The system operated excellently, but there is still room for improvement. Details of the evaluation and recommendations for improvements are contained in this document.

14. Subject Terms.

Remote Sensing, SAR, ERS-1, Sea Ice

15. Number of Pages.
58

16. Price Code.

17. Security Classification	18. Security Classification of Report.	19. Security Classification of This Page.	20. Limitation of Abstract. of Abstract.
Unclassified	Unclassified	Unclassified	SAR

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Contents

The Joint Ice Center SAR Workstation: Algorithm Evaluation

1.0 Introduction

The analysis of arctic sea ice is of great interest to many. The scientific, industrial, and defense communities are all working to understand the large and small scale geophysics of the Arctic. On a large scale, scientists wish to understand the relationship between the movement of ice masses, the weather systems which act upon the ice, and the oceans currents which circulate below the ice. An understanding of the relationships between these environmental factors leads to a better understanding of heat exchange and ice dynamics in the region. On a smaller scale, scientists wish to calculate ice concentrations, locate ice edges, follow ice motion, and analyze lead and ridge formation. Information about these characteristics is important to phenomenological studies of the ice, as well as to ship routing and oil rig safety.

The primary sources of operational information on arctic sea ice to date have been passive satellite sensors operating in the visible, infrared, and microwave wavelength bands. The sensors most often used by analysts are the Advanced Very High Resolution Radiometer (AVHRR) on the NOAA polar orbiting series, and the Operational Linescan System (OLS) and the Special Sensor Microwave Imager (SSM/I) systems of the Defense Meteorological Satellite Program. The advantages of these systems are well known. They provide synoptic coverage at reasonably fine resolutions. In addition, they provide daily, even sometimes overlapping, coverage of the Arctic. These sensors, though, are not without disadvantages. The visible, infrared, and higher frequency radiometer sensors are plagued by the inability to sense the surface through cloud cover. As the Arctic is often cloud covered, this represents quite a problem to those who wish to study this area. In addition, six months of darkness in the Arctic make the visible sensors unusable half the time. Passive microwave systems, such as the SSM/I, although unaffected by cloud cover or darkness, tend to be of too coarse a resolution (on the order of 25 km) to be used for detailed analysis of an area.

Great improvements in remote sensing can be achieved through the use of active sensors. With such sensors, the user controls all aspects of data collection. Frequency and polarization can be selected to optimize the sensor's response to various surface types. Synthetic aperture radars operate at frequencies which penetrate clouds. They achieve high azimuth and range resolution through the use of digital processing techniques. These techniques serve to simulate the resolving capability of a radar antenna much larger than that which can be carried in space. The European Space Agency's synthetic aperture radar (SAR) on board the ERS-1 satellite is one example of such a sensor. This sensor is one of a group of instruments launched on July 17, 1991 into a sun-synchronous orbit. This C-band radar, which operates at 5.5 GHz (or 5.3 cm wavelength), transmits and receives vertically polarized data. This choice of frequency and polarization provides excellent detection of and discrimination between types of arctic sea ice. The raw signal data collected by this satellite are transmitted to ground stations capable of processing SAR data. At the ground stations, the data are processed into a variety of products depending upon the desires of the user community. Data from the Beaufort and East Siberian Sea area, as illustrated by the station mask in Figure 1, are downlinked to the Alaska SAR Facility (ASF) located at the University of Alaska in Fairbanks. This facility serves the science community by providing information varying from raw signal data to processed images and derived ice motion and concentration estimates.

Among those users interested in SAR data for operational ice analysis is the Navy/NOAA Joint Ice Center (JIC) in Suitland, MD. The Naval Polar Oceanography Center is colocated with the JIC. The JIC receives SAR data from the ASF over a satellite link called SARCOM. JIC analyzes SAR data on a micro-computer based workstation specifically designed for sea ice analysis. The system receives both high (30 m) and low (240 m) resolution imagery, catalogs the data, and provides detailed ice classification and ice motion analysis products upon request. The

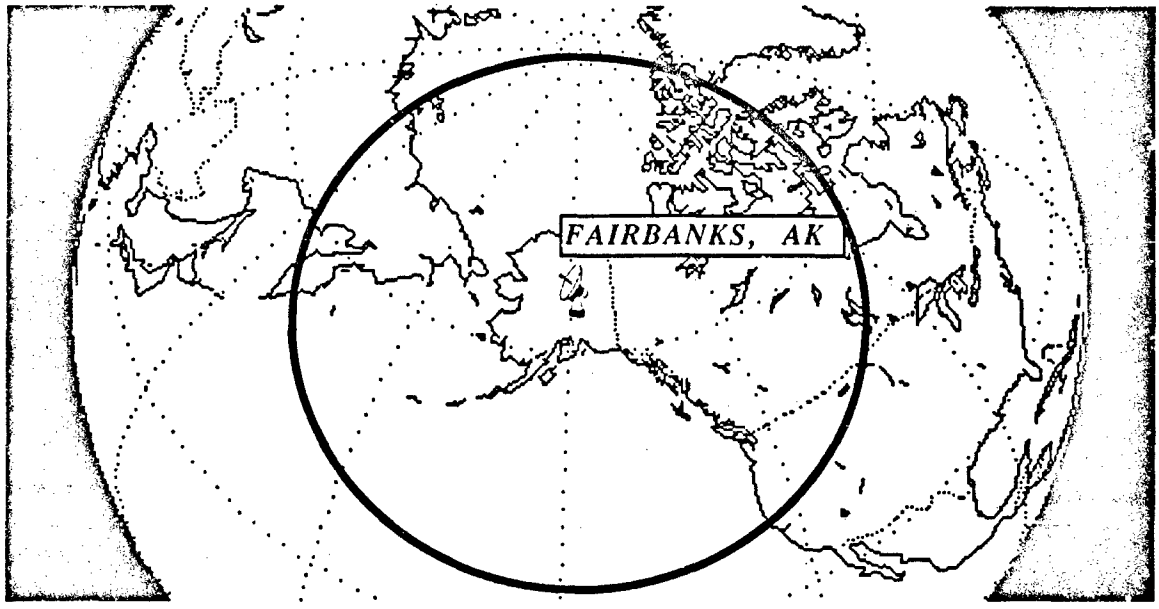

Figure 1. The station mask of the Alaska SAR Facility.

workstation was designed by the Jet Propulsion Laboratory (JPL) and delivered to JIC and to the NRL Remote Sensing Applications Branch with software which included algorithms for ice motion and ice classification. These algorithms are similar to those JPL provided to the ASF for use by the scientific community. NRL has added a graphics editor and a ship routing tool to the workstation software. NRL has also developed a set of software tools for image processing called the Navy Satellite Image Processing System (NSIPS). SAR images can be exported from the workstation environment into NSIPS. NSIPS has additional tools for ice analysis, including a ridge and lead algorithm developed at Stanford University and the University of Michigan. The presence at NRL of a SAR workstation identical to the workstation at JIC makes it possible to test new software on the NRL system without unduly disrupting SAR workstation operations at JIC. In this note, "SAR workstation" refers to both the NRL and JIC systems.

The use of these products heralds the first attempt to employ satellite SAR data in an operational setting. We anticipate that the workstation will help the JIC to provide more accurate and timely ice analyses. Before the workstation can be used with confidence, however, it must be tested for accuracy of analysis and ease of use. The format for the formal testing and transition to operational use of the workstation products was arrived at through the discussions of the SAR Working Group. This group includes representatives of JIC, NRL, and the program sponsors. Testing procedures are laid forth in the "JIC SAR Transition and Implementation Plan" (*Fetterer and Bertoia,* 1992). In brief, we have evaluated the overall performance of the workstation (excluding SARCOM) as well as the performance of the ice classification algorithm, the ice motion algorithm, and the ridge and lead algorithm. Workstation performance was evaluated over about a six month period (March to October 1992), while algorithm performance was evaluated by comparing SAR ice products to products from other sensors over a much shorter time period. For the algorithm performance tests, a data set was devised which consisted of SAR imagery collected during the LEADEX experiment (15 March 1992 to 15 April 1992) and covering an area of the Arctic from about 100° W to 155° E (Figure 2). This imagery was selected because of its wide range of coverage, and because of the availability of ground truth and imagery from other sensors during this time. The results of the algorithm evaluations are presented in this note. For a more

complete description of SARCOM, the SAR workstation hardware, ice and ocean applications, and managerial aspects of the SAR program, see *Fetterer et al.(1992)*.

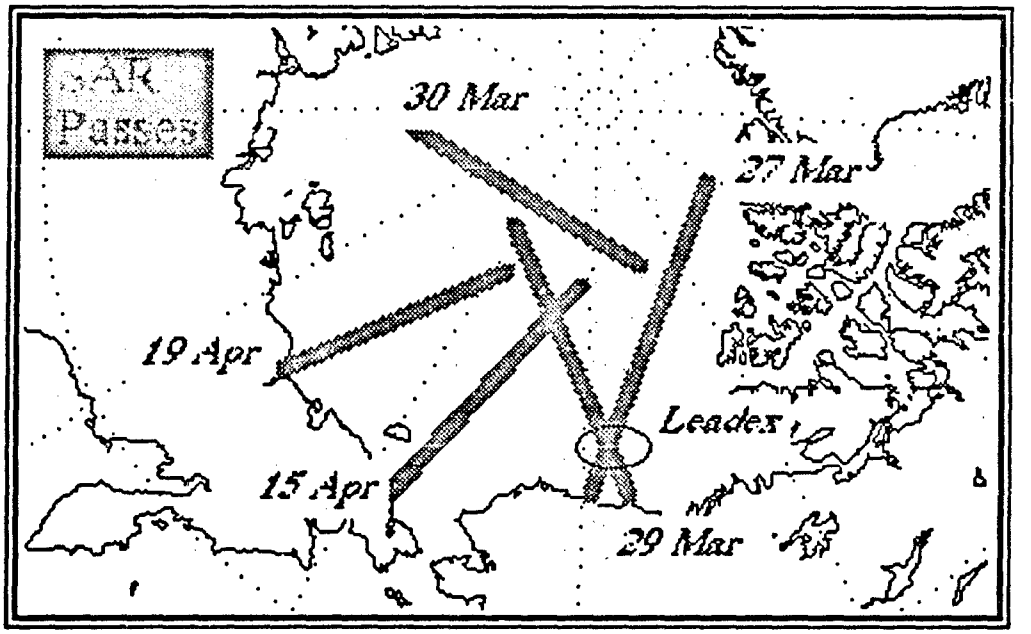

Figure 2. SAR passes used in the evaluation.

2.0 Daily Operational Use

2.1 Using the Workstation

The JIC/SAR workstation consists of a SUN computer system, SKYBOLT array processor, and various hard and optical disk drives. The user interacts with the workstation software through a series of menus. Selections from the menus bring up a series of INGRES database forms through which the user can choose the process and imagery they desire. An example of an INGRES form is presented in Figure 3. The system is quite easy to use once the user becomes familiar with the INGRES database system.

In order to analyze imagery, the user must first order it from the ASF and set up the workstation to receive transmitted images. Receiving images is achieved by selecting data ingest from the main menu and then completing the INGRES database form which appears. A user must be logged onto the workstation account to receive data, but the actual workstation software need not be running. The user can find out if the data have been transmitted by inspection of the LINK.LOG file on the /JIC/SARCOM/RECEIVE directory and by inspection of the low and high resolution image caches, /JIC/IMAGESTOR and /JIC/FULLSTOR. Inspection of the LINK.LOG file tells the user that data were transmitted to the workstation. Inspection of the low resolution and high resolution image caches tells the user that imagery has been properly geocoded (that is, earth located) and put into the INGRES database. Imagery which has been improperly transmitted by the ASF or improperly received by the workstation will reside in the error cache /JIC/SARCOM/RECEIVE/ERRCACHE) or in the low resolution holding cache /JIC/SARCOM/RECEIVE/LOHOLD). Images in the error cache should be disposed of. Retransmission of the images should be requested if desired. Images in the "lohold" cache have been transmitted correctly but need geocoding. Simply leaving the workstation software running

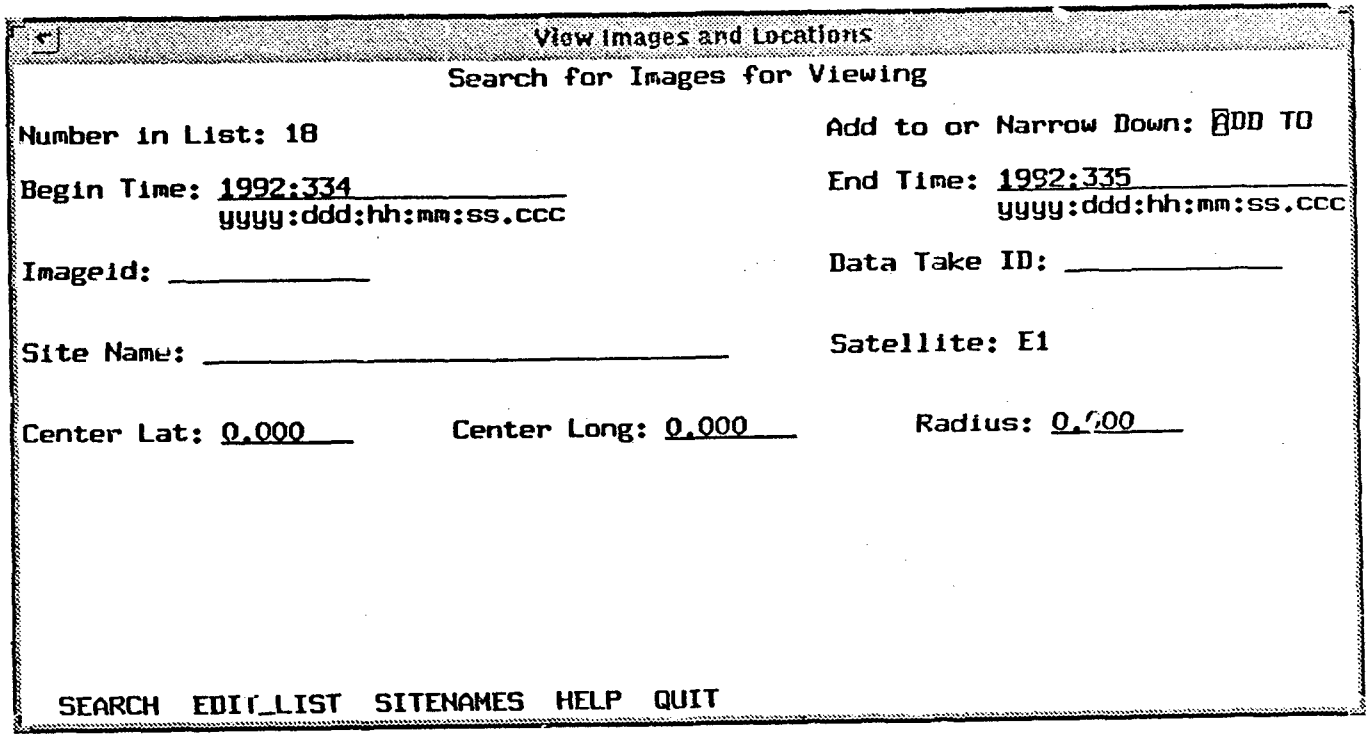

Figure 3. Example of an INGRES database form.

for a while will cause these images to be properly geocoded in time. Details on the image transmission and geocoding procedures can be found in the "Joint Ice Center SAR Workstation User's Guide" (*Baltzer et al., 1992*).

Received data can be viewed by the user by selecting the "view data" option in the main menu. An INGRES form will appear through which the user can select the images they wish to view. The system will then present maps of the image coverage, as well as the selected images at 3:1, 2:1, or 1:1 image to screen pixel resolution. (The imagery itself is 240 m resolution and 100 m pixel spacing) The user can request up to sixty images at a time for viewing. Many other functions such as drawing scales, locating a latitude and longitude, and annotating are also available. Figure 4 shows the image map, an image, and the list of additional functions which can be performed on the image.

Approximately 25 images are received by the workstation daily via satellite link when the systems at ASF and the ERS-1 satellite are fully functional. The number of images received and the area covered by the images will vary in accordance with changes in requests for coverage from the JIC to the ASF. These images are received on the SAR workstations at NRL and at JIC (occasionally, one workstation will be down at time of transmission, and imagery will have to be transmitted to the other workstation via telnet link). Using the "view data" option, the user can survey these images and produce a map of image coverage in about an hour. The user can then proceed to further analysis of the imagery if so desired or can choose to quit. Generally, it is wise to perform analysis on the data as soon as it arrives at the workstation. Only sixty images may reside in the image cache at a time, and waiting too long may cause the images to be flushed from the image cache (/JIC/IMAGESTOR) to an optical disk. Analysis will then take much longer since the images must be retrieved from the optical disk.

To date approximately 1400 images have been received by the SAR workstation here at NRL. They reside in the image cache and on fourteen optical disks along with their corresponding product data. Standard product data, which are routinely backed up onto optical disks by the system, consist of the results of unsupervised ice analysis: ice classification maps, ice concentration list files, and ice motion vectors files. Non-standard products consist of the results of supervised ice analysis. These products must be saved by the user.

2.2 Problems Encountered and Recommendations

Generally, problems encountered were not serious flaws in the software but rather inconveniences which should be corrected in future updates. The most severe problem concerns the entry and deletion of images from the image cache. Every analysis performed or product generated by the system requires that the image under consideration be loaded into the image cache. As the image cache reaches its limit of 60 images, the oldest images (by date of acquisition) in the cache are deleted from the cache when the user wants to work with imagery not in the cache. This causes significant problems when the user is analyzing older data. The oldest data is always removed to make room in the cache, even though it may be the most recent data to be used. This is a frustrating occurrence, as jobs must be continually restarted. For instance, if a user submits a motion request using older data, and then attempts to view newer data, the older data will be removed from the cache before any motion analysis can be performed. This in turn causes any jobs in the queue using older data to bomb. This problem affects all analysis performed on the workstation. Another annoyance is the speed with which the optical disk is accessed. The optical disk is far too slow to handle the amount of data which must pass between it and the hard disk. A request to view ten images which reside on the disk can take anywhere from 20 to 30 minutes to load. The daily backup of new images and products from the hard disk to the optical disk takes about the same amount of time. Finally, an option should be added to the workstation which would allow the user to produce a map of image coverage or a printed listing of selected images without having to load imagery into the image cache first. In many instances the user only wants to survey available data either through a listing or a map but in the workstation's present configuration, no matter what is desired, the user must wait for the system to load imagery into the

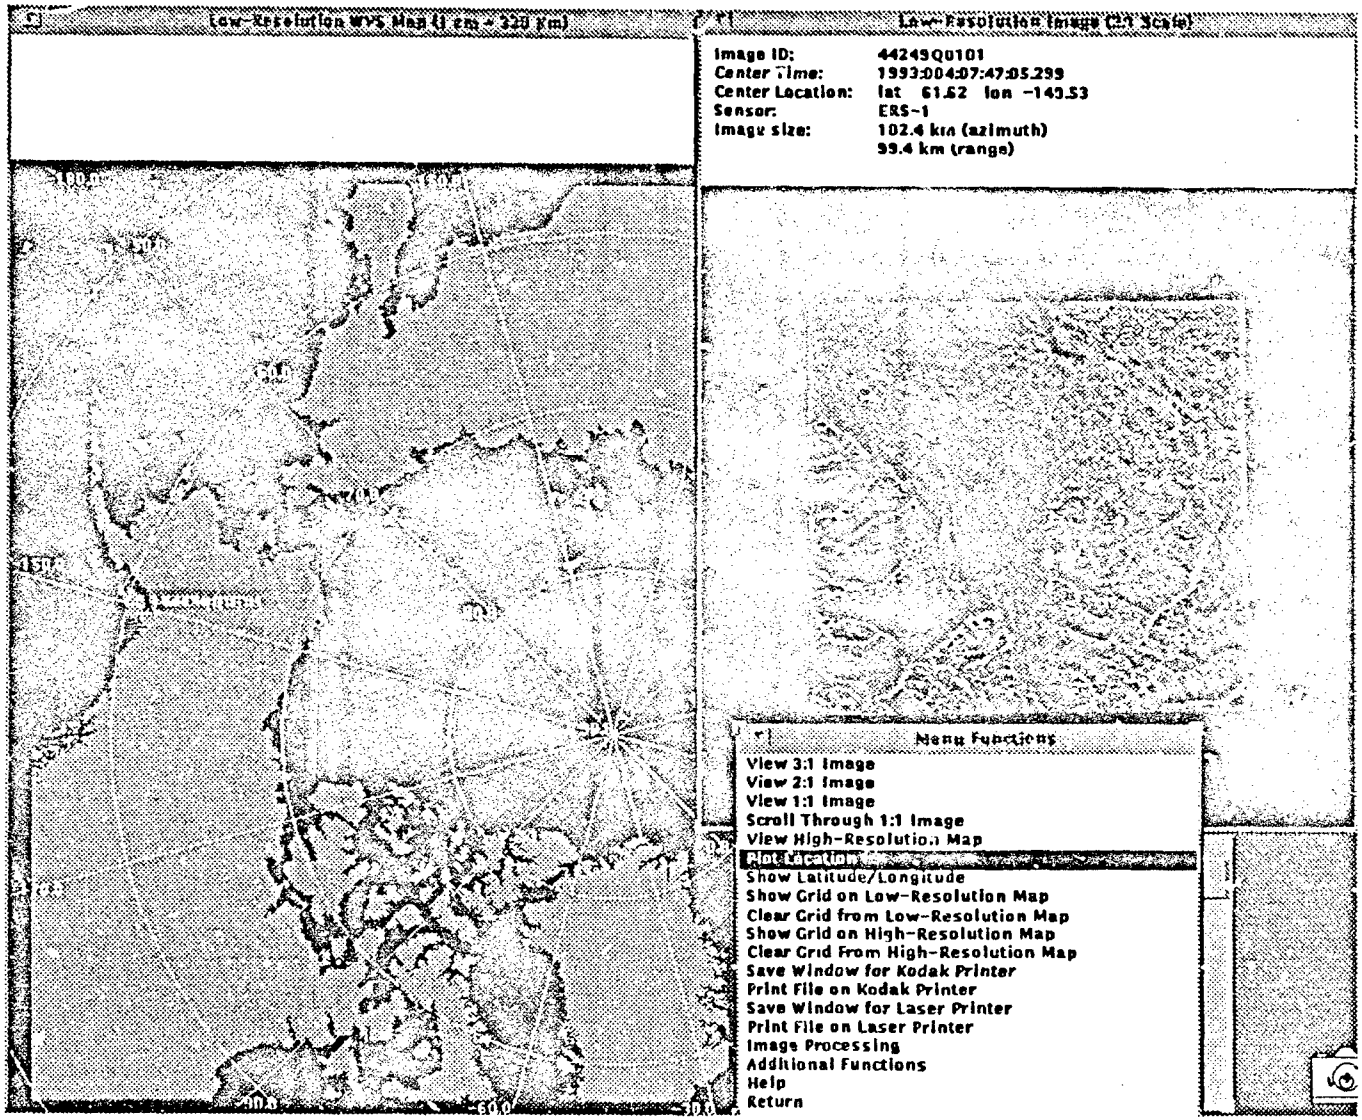

Figure 4. Example of a workstation screen.

image cache. This is quite time consuming. Recommendations for improvements in the general function of the workstation are, therefore, to:

1. base the order of which images are removed from the image cache on date of an image's last use instead of date of image acquisition. Date of last use can easily be added as an entry to the INGRES databases. System clean-up facilities which now access the date of acquisition in the database as a clean-up parameter can then be modified to key in on date of last use.

2. change the order of the functions in the image loading procedure on the workstation. Currently, when a user wishes to access images, the sequence of processes is to first load the imagery into the image cache and then display a map and list of image locations. Simply changing the order of the processes so that a map and list are generated before the imagery is loaded would enable a user to easily survey the imagery, without disrupting the view image functions. This is a desirable feature, as often a user wishes only to see what imagery is available, rather than viewing the imagery itself.

3. install faster optical disk equipment as soon as technology permits. The loading of data from the optical disk currently in use is far too slow to allow for the type of synoptic analysis operational users routinely perform.

3.0 Ice Classification Algorithm Evaluation

3.1 Algorithm Description

The unsupervised classification algorithm used by the workstation is described in detail in *Kwok et al. (1992)*, but a synopsis of the method is presented here for the reader. First, the radar data are calibrated. This calibration step maps the pixel values of a radar image into radar backscattering coefficients by removing the effects of noise, antenna pattern, and system gains from the data. In the course of compensating for residual antenna effects, natural incidence angle variation in the backscatter of the ice is also removed. Theoretically then, image pixels throughout the image containing the same type of ice are represented by the same backscatter cross section values.

The calibrated image then undergoes cluster analysis using the ISODATA clustering technique. This cluster analysis is performed in the image intensity domain on 5% of the image pixels. The cluster with the most points, or dominant cluster, is then labeled by comparing the cluster centroid (equivalent to the mean backscatter value of the cluster) to the mean backscatter values of various ice types in a pre-determined look-up table of ice backscatter means and standard deviations. Cluster centroids for other ice types are not obtained from the image, but are instead obtained using the expected differences in mean backscatter, as determined from the look-up table, between each ice type and the dominant ice type. The relative differences calculated from the look-up table are then added or subtracted to the dominant cluster centroid in order to find the values of the cluster centroids for other ice types. In this way, the algorithm applies a sliding scale for ice type to backscatter, instead of simply classifying the image by a straight comparison to backscatter values in the look-up table alone. This technique makes the algorithm more robust by compensating for changing overall backscatter values from image to image, but it depends upon constant differences in the backscatter distributions of different ice classes. Simply put, it does not account for between class variability.

Once the cluster centroids are located, the image pixels are then labeled using either minimum distance or maximum likelihood classification techniques. The minimum distance classifier works in a fashion similar to a thresholding of the data. The backscatter values from individual pixels are compared to the backscatter values of the cluster centroids. Pixels are then labeled as being a certain ice type when the Euclidean distance from the pixel value to the cluster centroid value for that ice type is a minimum.

Minimum distance classification does not take into account the natural variability of backscatter within the ice types. It assumes that all ice types identified in the data have identically shaped backscatter distributions. Maximum likelihood classification, on the other hand, makes use of all of the statistically derived backscatter distribution information in the look-up table. Since the distributions, and not just the means, are used, the maximum likelihood method takes the natural variability of the backscatter signature within ice types into account. The maximum likelihood classifier also uses probability to label pixels. A pixel is labelled as belonging to a particular ice type if the product of the a priori probability that any pixel belongs to that ice type, and the probability density function (i.e. the normalized backscatter distribution) of that ice type is maximum at that pixels value. In simpler terms, the statistical distributions of backscatter for the various ice types are weighted by the probabilities that the ice types exist in the image, and the pixel is labeled according to the maximum of this weighted value. The maximum likelihood classifier provides better results than the minimum distance version, but misclassification can result it ice type distributions in the look-up table are not realistic, or if a priori probabilities are not correct. It also takes about a minute longer to run than the minimum distance classifier.

The largest source of error in classification is the natural variability in the backscatter of ice not taken into account by the seasonal look-up tables which are used by the algorithm. The original look-up tables were based extensively on scatterometer data collected over a ten year period in the Beaufort Sea area. Early estimates from ERS-1 showed that these look-up tables overestimated backscatter for all ice types by about 2 dB. Revised winter look-up tables were formulated by JPL using backscatter measurements from ERS-1 imagery. The winter time look-up table has entries for multi-year ice (MY), first year ice (FY), open water or new ice (OW/NI), and "ice type 4", which is first year ice with a darker signature (I#4). While the geophysical interpretation of "ice type 4" remains to be validated, it has been called "thin first year ice", "undeformed first year ice", or simply "dark first year ice". The winter look-up table used by the algorithm is presented in Table 1. Backscatter values are for ERS-1 parameters of 5.25 GHz, VV polarization, and an incidence angle of 25°. The look-up tables do not take into account factors such as geographic location or ambient air temperature, both of which may be related to radical variations in backscatter values. Although look-up tables have been assembled for all ice seasons, good algorithm performance is only assured from Julian day 280 (late October) to day 120 (late April), when surface temperatures are usually below -5° centigrade in the Beaufort Sea.

Table 1. Look-up table for winter to early spring.

Ice Type	Mean (dB)	Standard Dev (dB)	Slope (dB/degrees)
Multi-year(MY)	-10.0	1.01	-0.08
Thick First Year(FY)	-14.0	1.06	-0.24
Thin First Year(I#4)	-17.0	1.26	-0.24
New Ice/Open Water (NI/OW)	-21.8	1.66	0.00

Improper calibration of the SAR imagery can also cause errors to occur in the cluster labeling of the classification algorithm, but the long term accuracy of the calibration of the ERS-1 SAR has been on the order of 3 dB, which is within the error limits of the algorithm. Some error may be caused by the location uncertainty of the dominant cluster centroid during cluster analysis. The degree of uncertainty is expected to produce a maximum misclassification of pixels on the order of 5% to 10%.

Supervised classification is an additional ice classification option. In supervised classification, the analyst uses the mouse to draw polygons around "training areas" of different ice types. The image is then classified using the backscatter statistics of the training sets. The analyst dictates the number of classes. Supervised classification is desirable when the user observes that the ice types present in an image do not easily correlate with those types listed in the look-up tables. Examples of such ice types would be multi-year ice with a high concentration of melt

ponds, which would appear much darker than the backscatter for multi-year ice in the look-up table, and wind roughened open water, which would have a much higher backscatter cross section than that of water in the look-up table (which is assumed calm). Supervised classification is also desirable when the number of ice types to be classified is greater than that in the look-up table. Of course, when the backscatter distribution of ice type classes overlaps, some classification error is unavoidable, regardless of whether supervised or unsupervised classification is used.

3.2 Running the Algorithm

The methods used in both unsupervised and supervised classification are similar. In both cases the user employs INGRES forms to select the imagery they wish to analyze. Any number of images can undergo classification at the same time. When unsupervised analysis is requested, the system simply asks whether the minimum distance or maximum likelihood version of the algorithm will be used. The classification requests are then submitted to the system in batch form. When a classification is completed an alert will appear in the operator alert manager window. As the user "handles" this alert the classification product is entered into the INGRES database and the classification appears on the screen for the user's inspection. When supervised classification is requested the system first calibrates the image and then presents it to the user for the determination of the classes and training areas. At this point the user can either select classes and outline training areas on the image to create a look-up table for classification, or can elect to use a look-up table saved from a prior session. As with unsupervised classification, the user must choose between minimum distance and maximum likelihood classifiers. The classification is performed by the system in real time and the results are immediately presented to the user.

Both unsupervised and supervised classification run quickly once an image has been selected and loaded. Unsupervised classification takes approximately three minutes per image and supervised classification takes about five. It takes approximately five minutes to select the desired data. Loading of the desired data into the system may take anywhere from five to twenty minutes depending on the number of images requested and the storage location of the images. Images already in the image cache are quick to load, while images stored on an optical disk take much longer.

3.3 Comparison of Unsupervised and Supervised Classification Results

A classifier is of no use at all to the operational ice community if it does not produce accurate, consistent results from image to image. The wintertime performance of the JIC workstation ice classifier was tested for accuracy by comparing the result of unsupervised classification for any given image with the result of supervised classification for the same image. For these tests, an image set was chosen which spanned approximately three weeks (from March 27, 1992 to April 19, 1992) and an area of the Arctic from approximately 90°W to 130°E and from 70°N to 85°N. Table 2 summarizes pertinent satellite pass information and Figure 2 provides a map of the swaths which were used in this analysis. The test area was selected to provide maximum coverage within the ASF station mask. Unsupervised classification using the minimum distance classifier, unsupervised classification using the maximum likelihood classifier, and supervised classification using the maximum likelihood classifier were performed on 64 images.

Table 2. SAR data used in ice classification algorithm evaluation .

Date	Time	Ascending or Descending	Pass Start and End Points (Latitude, Longitude)
27Mar92	21:25Z	D	83.33,-103.45 to 70.12,-149.21
29Mar92	06:49Z	A	71.89,-142.66 to 81.98,-171.88
	22:02Z	D	73.74,-153.15 to 71.50,-156.05
30Mar92	04:38Z	A	81.62,-135.79 to 82.68,139.49
15Apr92	21:21Z	D	73.58,-142.88 to 70.82,-146.33
	23:00Z	D	78.97,-156.12 to 68.79,-173.53
19Apr92	00:47Z	D	80.26,-177.66 to 69.08,160.14
	20:55Z	D	73.86,-135.99 to 72.38,-138.06

Using supervised classification results as "truth", the accuracy of the unsupervised classification algorithm was examined through two methods. The first method of analysis employed plotting the ice type concentration estimates for each image determined by unsupervised classification against those of supervised classification. If all points on these plots fell on a line with a slope of one and an intercept of zero, the unsupervised classification results could be said to be completely accurate. Departures from this perfect linear fit indicate error. The results of this analysis, with best-fit regression lines and 45° lines of no error, are presented in Figures 5 through 9.

Points in the data set which indicated gross errors, or outliers, were removed from the data set prior to analysis if they were found to be cases outside the stated spatial domain of the algorithm, which is the Beaufort Sea north of the fast ice zone. This step was performed to insure a characterization of the general accuracy of the algorithm. The outliers removed from the data set are presented in Table 3. They fall into two categories: 1) outliers caused by the presence of fast ice in the image which is mistaken for multi-year ice by the algorithm, and 2) outliers caused by the presence of first year ice with abnormally high backscatter values in the Chukchi sea. These outliers represented only six of the sixty-four images analyzed for this evaluation.

Table 3. Classification outliers.

Image ID	Date	Location (Latitude, Longitude)	MY conc (Supervised classification)	MY conc (Algorithm)	Reason
13028	27Mar92	70.82,-146.33	9%	44%	heavily ridged fast ice
15206	15Apr92	70.87,-148.41	10%	58%	heavily ridged fast ice
18557	15Apr92	71.05,-171.23	6%	25%	anomalously bright FY
18558	15Apr92	70.30,-172.82	2%	57%	anomalously bright FY
18559	15Apr92	69.54,-172.82	1%	57%	anomalously bright FY
18560	15Apr92	68.79,-173.53	0%	60%	anomalously bright FY

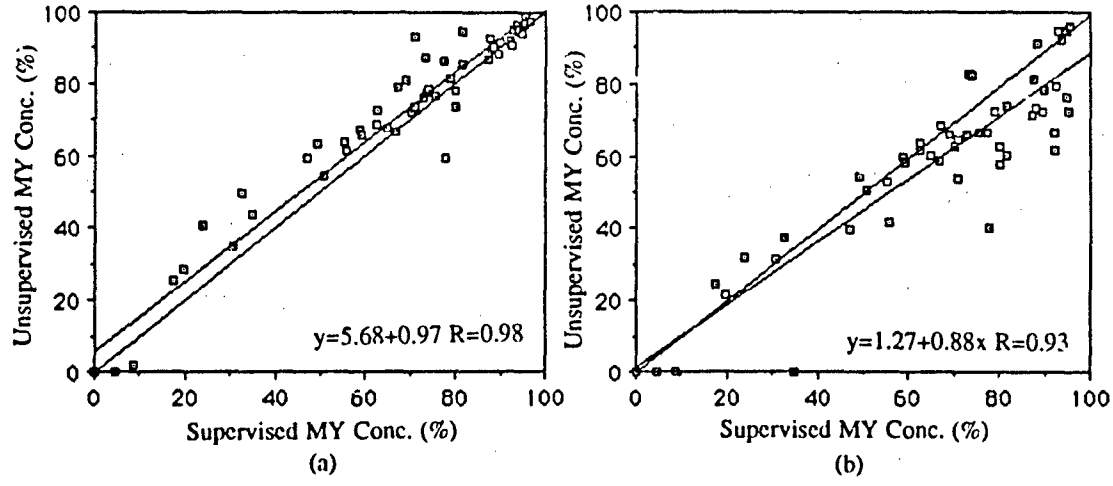

Figure 5. Results for multi-year ice. a) Supervised classification versus algorithm maximum likelihood classification. b) Supervised classification versus algorithm minimum distance classification.

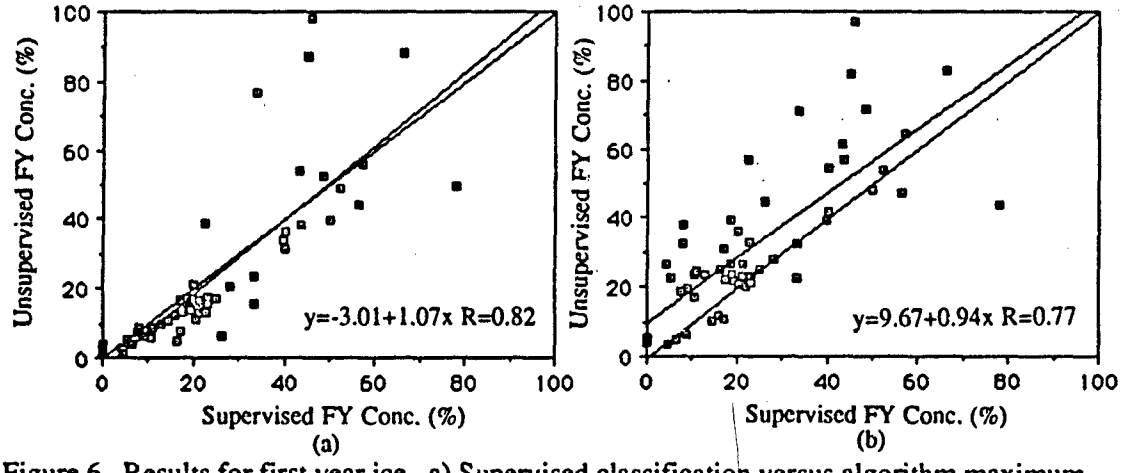

Figure 6. Results for first year ice. a) Supervised classification versus algorithm maximum likelihood classification. b) Supervised classification versus algorithm minimum distance classification.

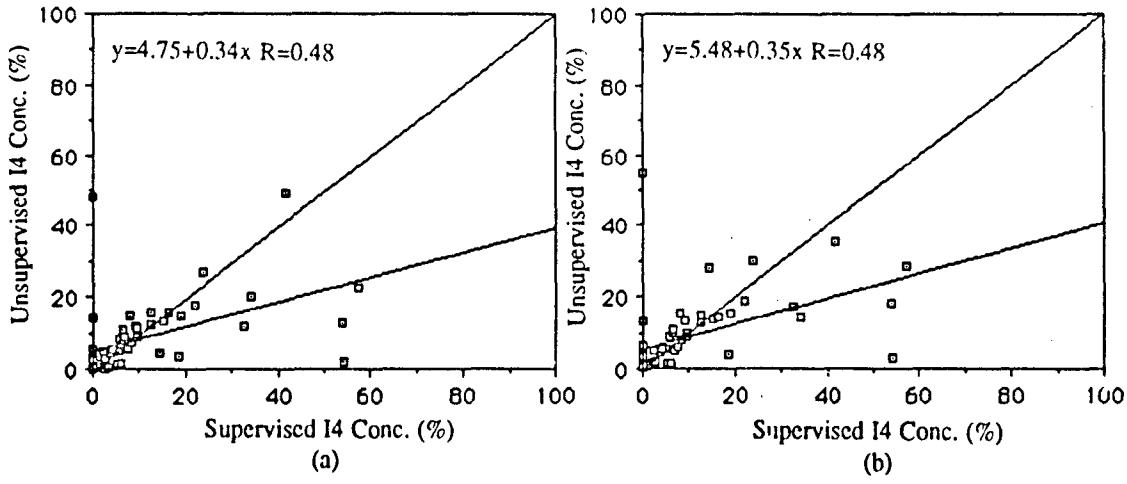

Figure 7. Results for ice type #4. a) Supervised classification versus algorithm maximum likelihood classification. b) Supervised classification versus algorithm minimum distance classification.

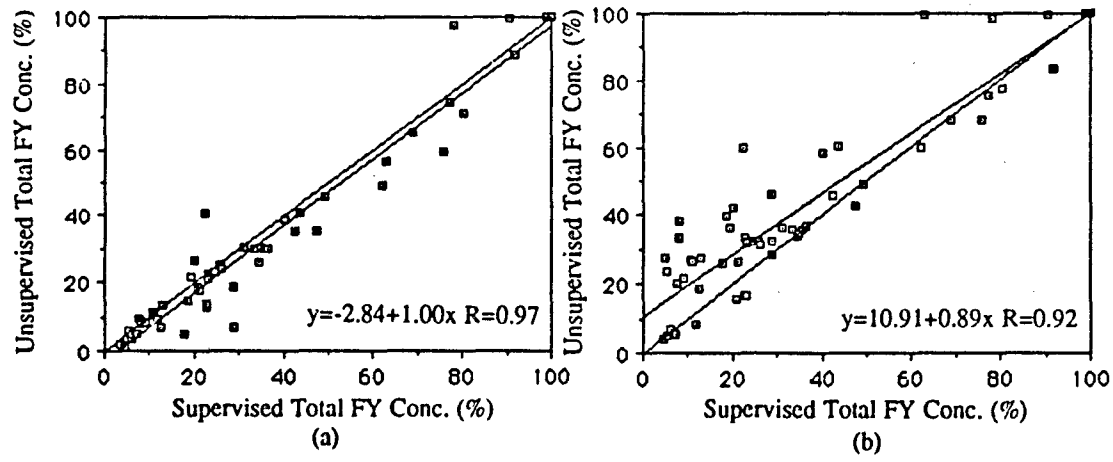

Figure 8. Results for all FY (first year and ice type #4). a) Supervised classification versus algorithm maximum likelihood classification. b) Supervised classification versus algorithm minimum distance classification.

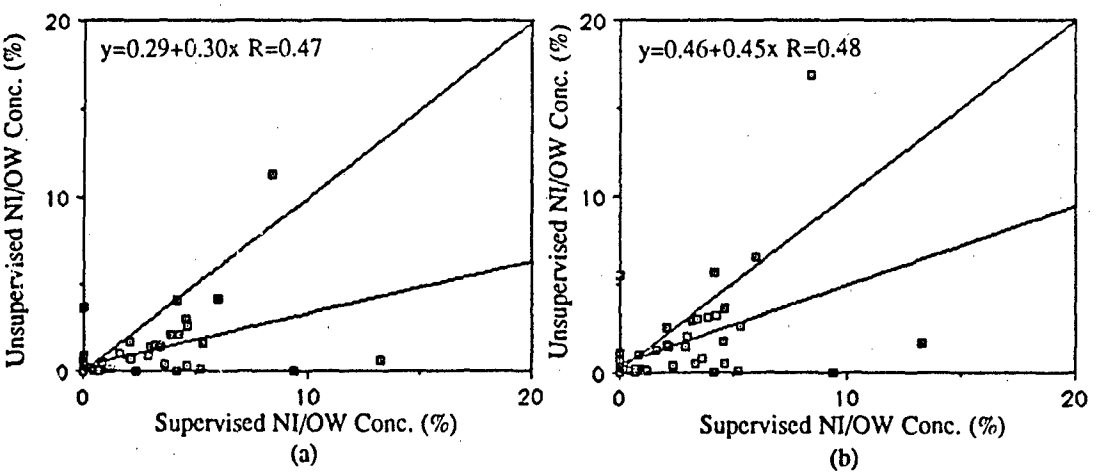

Figure 9. Results for new ice and open water. a) Supervised classification versus algorithm maximum likelihood classification. b) Supervised classification versus algorithm minimum distance classification.

In general, the data trends were linear with varying amounts of scatter around a linear fit, depending upon the ice type in question. In the first year ice and multi-year ice cases, the maximum likelihood classifier performed better than the minimum distance classifier, as evidenced by the regressed slopes having values close to one and the high correlation values of the fits. This is not the case with ice type #4, although the data appear to have a linear trend for concentrations below 40%. Some of the scatter in ice type #4 values is due to confusion between the first year and ice type #4 ice types in the supervised classification. While an effort was made to choose training areas with "dark" and "bright" first year ice, in some images backscatter of first year ice was intermediate, and it was very difficult for the eye to distinguish two first year ice types. When a more accurate supervised measurement is created by taking the sum of the ice type #4 and first year concentrations and considering a total first year concentration rather than separate concentrations, as in Figure 8, the data trend is linear with a slope close to one throughout the range of concentration values. The new ice data (Figure 9) exhibits clustering of the values, but this is due to the lack of data samples with a wide enough range of concentrations to show a trend.

For first year and multi-year ice, the minimum distance best fit line is offset by less than 10% from the line of no error. For the maximum likelihood classifier, the offset is approximately 5% or less. It must be remembered that some error is introduced into this analysis by the subjective nature of supervised classification, but this error was minimized by having a single user select training areas and by performing all the classifications in the course of about two days, thus keeping the imagery fresh in the user's mind.

In the second method of analysis, the difference between ice concentration estimates given by supervised and unsupervised classification was found for each ice type in every image. Statistics for these differences, by ice type, were then calculated in order to produce values which would provide measures of misclassification or error. Here the outliers in Table 3 were included. Table 4 presents the results of these calculations. Column 1 of the table describes the measurement being derived. For example, "Sup-Min MY" refers to the difference between multi-year ice concentrations as estimated by supervised classification and by the unsupervised minimum distance classifier. Columns 2, 3 and 4 give the mean, standard deviation, and (absolute) maximum value of this difference, respectively. In every case, the mean difference between unsupervised and supervised ice estimates is less than 10%, which is within the expected error of the algorithm. However, mean difference is not a good measure of error, since it may be deceptively small when data points are on both sides of the line of no error (see Figures 5 through 8) and both positive and negative differences are present. While the mean difference may be a good measure of error for

maximum likelihood classification of multi-year ice (since the algorithm gives results which are fairly consistently biased compared with supervised classification), the standard deviation of the differences is a better measure of error for the other cases.

The last four rows of Table 4 present the mean differences in concentration for images with multi-year concentrations greater than 60% and first year concentrations less than 40% (as determined by supervised classification). From the plots in Figures 5 through 9 it is apparent that the majority of the data samples lie within these regions. Therefore it is for these regions that the most accurate representation of algorithm accuracy can be obtained. (These regions exclude the outliers in Table 3). Here mean difference for maximum likelihood classification of multi-year ice is -3%, and the standard deviation for the maximum likelihood classification of first year ice is plus or minus 6%. These measures are the best measures of the algorithm error in classifying multi-year and first year ice. Note that the maximum likelihood classifier outperformed the minimum distance classifier, although both algorithms did produce some very blatant misclassifications, as evidenced by the large values of the maximum differences.

The errors (or percent difference in supervised and unsupervised results for all sixty-four images) were plotted against date of acquisition and location. As evidenced by Figures 10 and 11, the errors appear to be random with respect to both of these factors.

Table 4. Error, or percent difference in supervised and unsupervised results.

Error Measurement	Mean (%)	Stand. Dev.(%)	Max Diff. (%)
Sup-Min MY	4	13	36
Sup-Max MY	-8	15	60
Sup-Min FY	-9	15	63
Sup-Max FY	2	14	44
Sup-Min I4	4	18	63
Sup-Max I4	5	18	71
Sup-Min Total FY	-5	14	37
Sup-Max Total FY	7	15	60
Sup-Min NI	1	3	25
Sup-Max NI	1	2	20
Sup-Min MY>60%	9	11	38
Sup-Max MY>60%	-3	6	18
Sup-Min Total FY<40%	-8	10	37
Sup-Max Total FY<40%	3	6	22

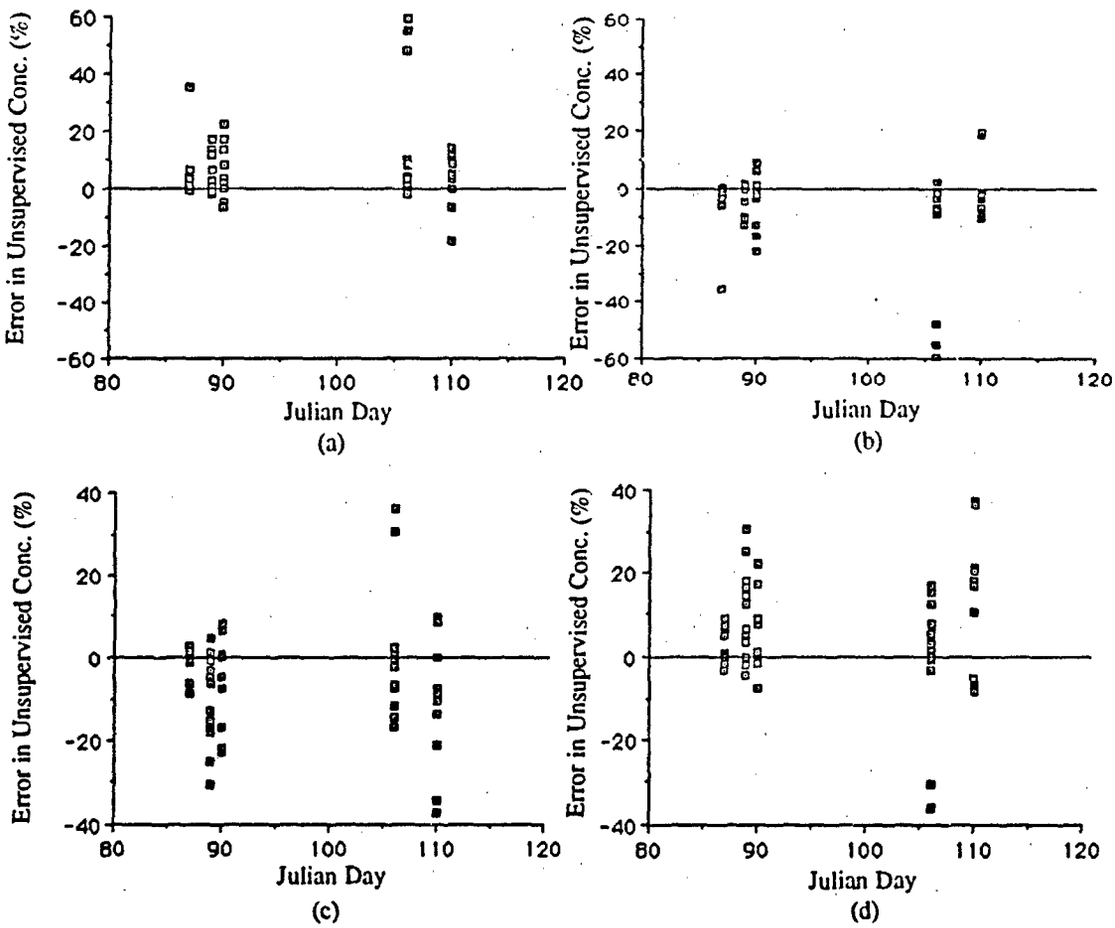

Figure 10. Classifier error versus date. For a) multi-year ice and maximum likelihood algorithm classification, b) first year ice and maximum likelihood algorithm classification, c) multi-year ice and minimum distance algorithm classification, d) first year ice and minimum distance algorithm classification.

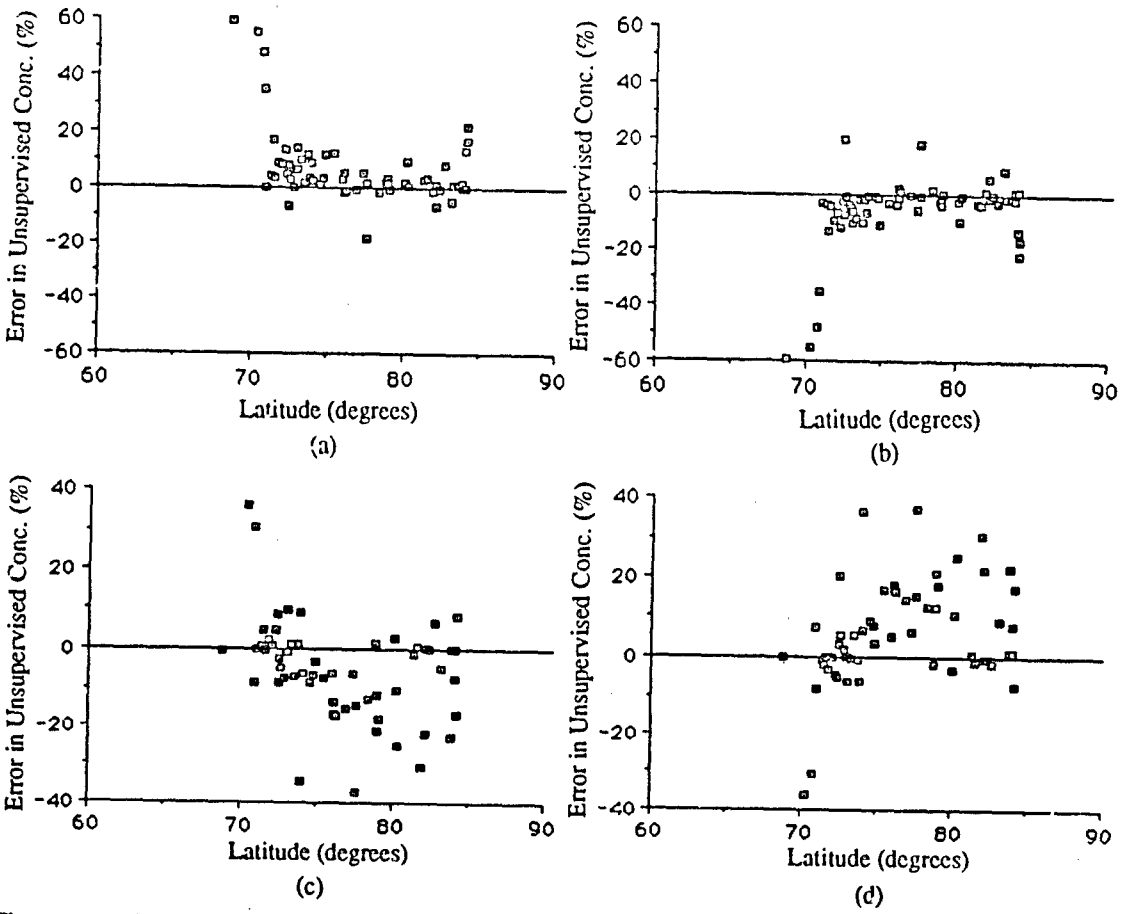

Figure 11. Classifier error versus latitude. For a) multi-year ice and maximum likelihood algorithm classification, b) first year ice and maximum likelihood algorithm classification, c) multi-year ice and minimum distance algorithm classification, d) first year ice and minimum distance algorithm classification.

In general, the unsupervised classification algorithm performed accurately within the 10% inherent uncertainty of the algorithm. The minimum distance version tended to over-estimate the amount of first year ice in the image, as can be observed by the slight bias in the plots in Figures 6 and 8. This error was most obvious in images which had large (above 90%) concentrations of multi-year ice. An example of this over-estimation, and the corresponding maximum likelihood classification, is given in Figure 12. In this particular case, the minimum distance version of the classifier over-estimated the amount of first year ice by about 40%. The misclassification may be due to regional variations in backscatter of the multi-year ice. Most of the high multi-year concentration cases occurred at higher latitudes and the multi-year ice at these latitudes was slightly lower in mean backscatter coefficient and slightly higher in standard deviation than multi-year ice at lower latitudes. Visually, the ice appeared to have a grainy look, as evidenced in Figure 13. In these cases, the backscatter value of some of the multi-year ice pixels was probably slightly closer to the mean of the first year ice in the look-up table than to the mean of multi-year ice. In some cases of extremely high (above 97%) multi-year concentrations, the minimum distance classifier failed completely with an abnormal termination program error.

Conversely, the maximum likelihood classifier tended to over-estimate the amount of multi-year ice as evidenced by the bias in the plot in Figure 5. This error was most obvious when the

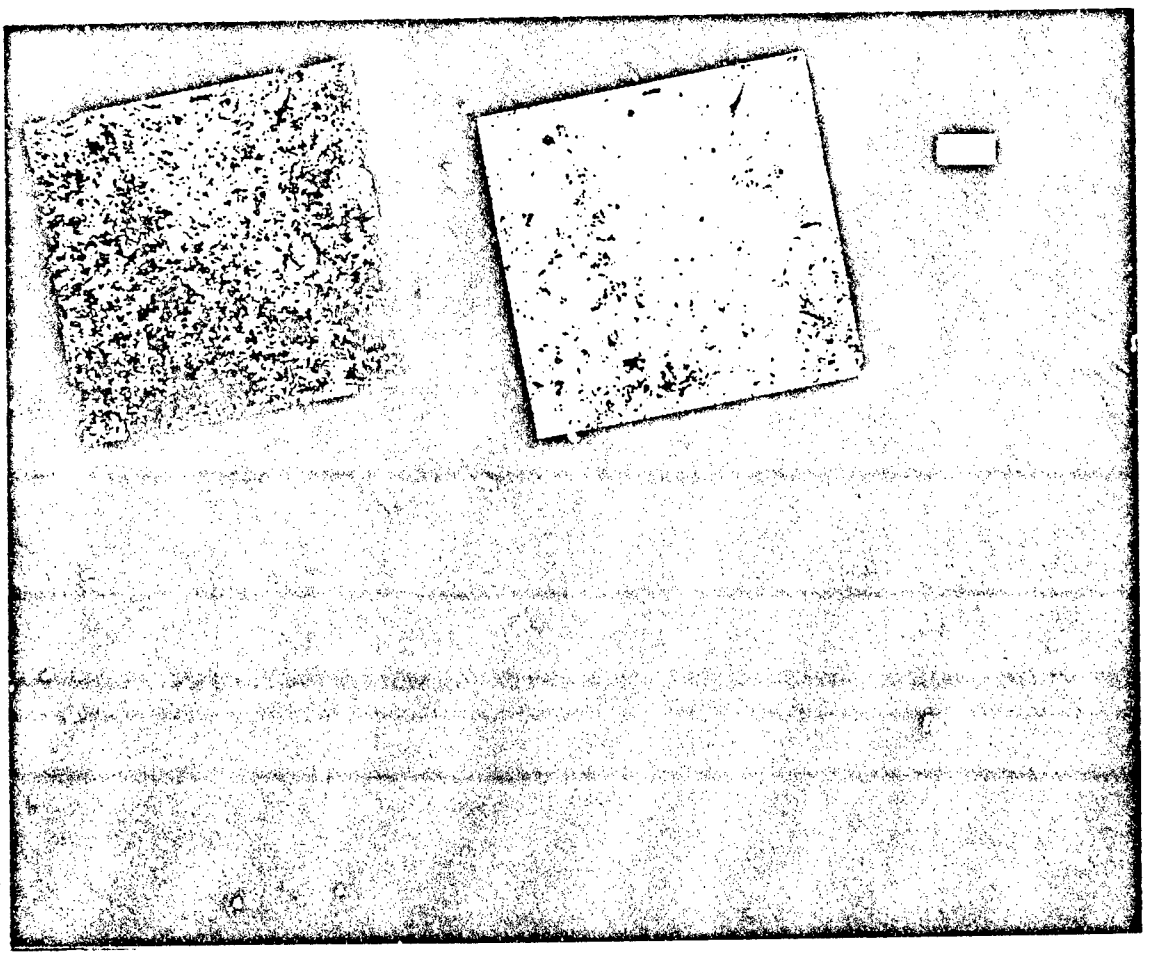

Figure 12. Example of a minimum distance classifier misclassification, overestimation of FY ice.

Figure 13. Comparison of "grainy" MY ice to normal MY ice.

amount of first year ice was large (50 to 90 %), as illustrated in Figure 14. In this particular case the maximum likelihood version of the classifier over-estimated the amount of multi-year ice by about 60%. The error in this case is probably caused by inaccurate estimation of the a priori probability for multi-year ice. If the probability of the existence (expressed as a probable concentration) of multi-year ice is higher than the actual amount of multi-year ice, the maximum likelihood classifier would tend to classify more pixels as multi-year ice than it should.

Both minimum distance and maximum likelihood versions of the algorithm tended to misclassify dark (-19 to -20 dB) first year ice as new ice or open water. Examples of this type of error are presented in Figure 15. This error is the result of some first year ice in the images having backscatter coefficients closer in value to that typical of new ice rather than that of typical first year ice. The error is unavoidable with the algorithm in its present configuration, since widening the distribution of the "dark" first year ice class in the look-up table to correct the error would overlap the distribution of the new ice class, resulting in misclassification of new ice. The algorithm also misclassifies other surface types which have backscatter distributions which overlap the ice types in the look-up table. Examples of this are wind-roughened water (-8 dB to -5 dB), which is confused with multi-year ice, and freezing new ice in leads (-13 dB to -10 dB), which can be confused with either first year ice or multi-year ice. We did not have enough open water or new ice in this data set to properly evaluate the performance of the algorithm for the open water/new ice type.

3.4 Comparison of SAR Unsupervised Classification With Ice Type Classification From Other Sources

Additional evaluation of the unsupervised ice classification algorithm was performed through comparing SAR classification results with classification results from other sources. The primary sensor used in these comparisons was the SSM/I passive microwave radiometer. SSM/I data were collected and processed by NRL during the LEADEX experiment. Table 5 presents a summary of the data used in this comparison. Figure 16 presents a map of the coverage of each sensor. The maximum likelihood version of the unsupervised ice classification algorithm was used for the SAR data. For the SSM/I data, the Navy CAL/VAL sea ice algorithm was employed. The JIC currently receives ice concentration information obtained by running this algorithm at Fleet Numerical Oceanography Center (FNOC) and plotting the results at NRL.

Table 5. Data used for SAR - SSM/I comparison.

SAR Data Date	Time	SSM/I Data Date	Time
27Mar92	21:25Z	27Mar92	18:18Z
29Mar92	06:49Z	29Mar92	17:52Z
	22:02Z		
30Mar92	04:38Z	30Mar92	17:40Z
15Apr92	21:21Z	15Apr92	17:38Z
	23:00Z		
19Apr92	00:47Z	19Apr92	16:47Z
	20:55Z		

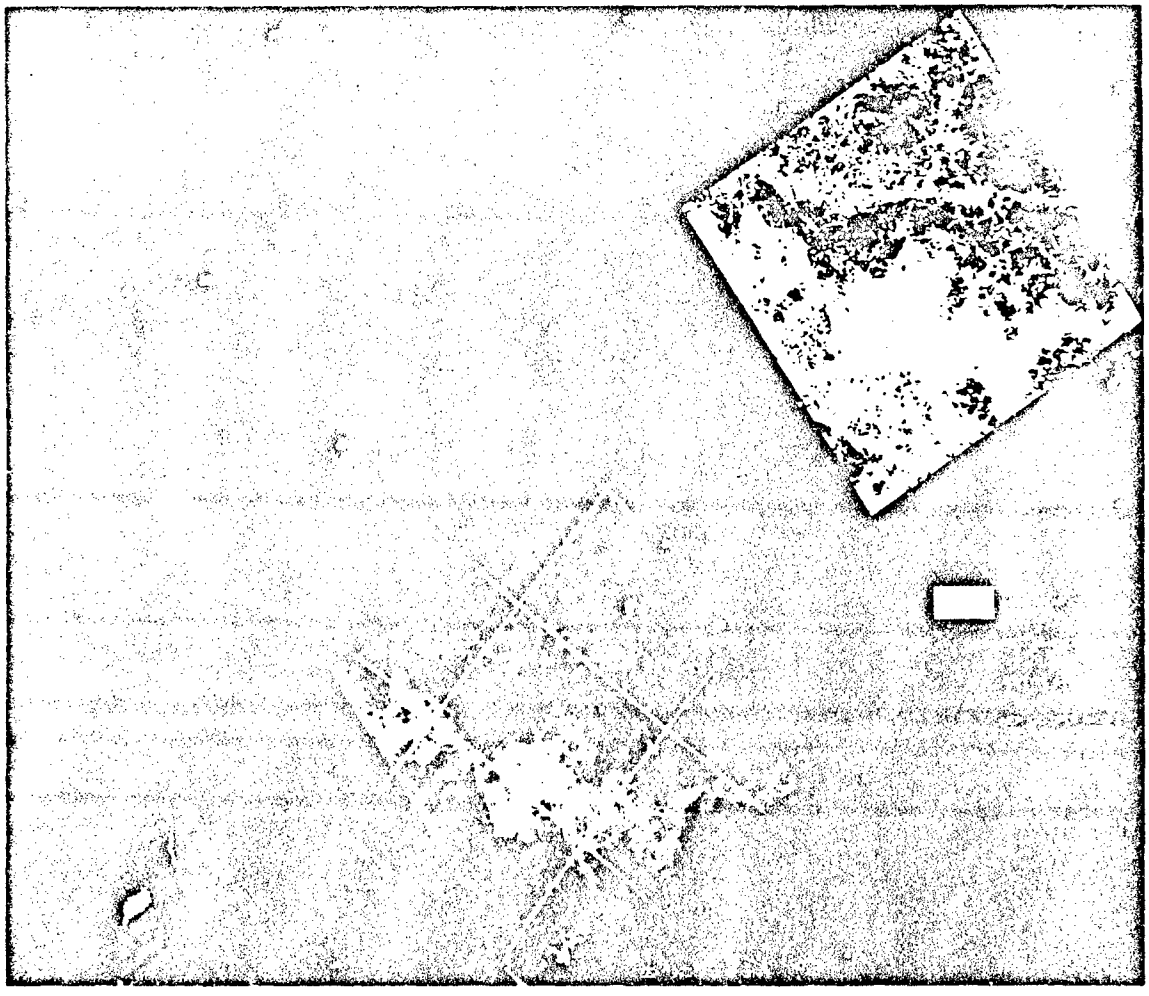

Figure 14. Example of a maximum likelihood classifier misclassification, overestimation of MY Ice.

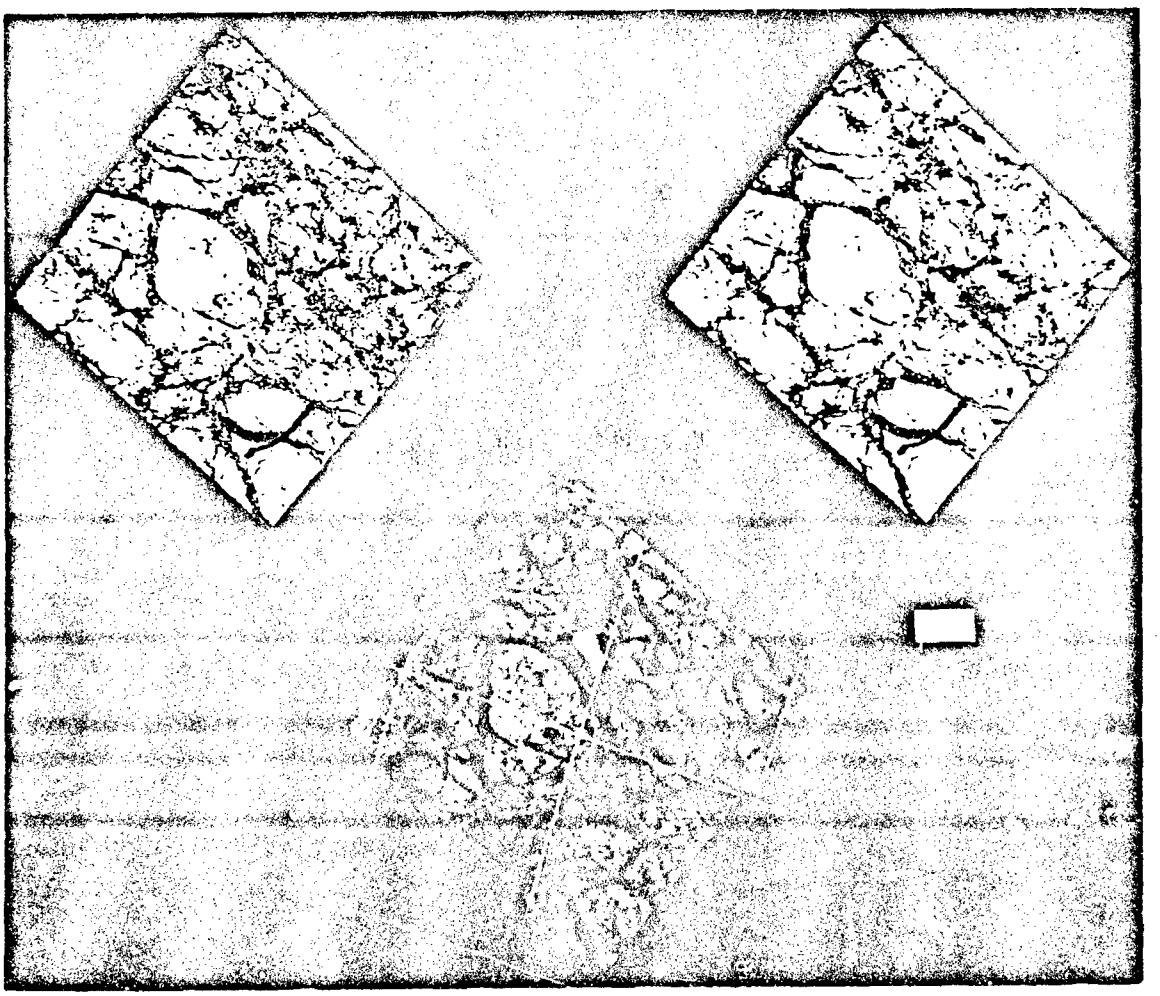

Figure 15. Example of new ice / dark FY ice misclassification.

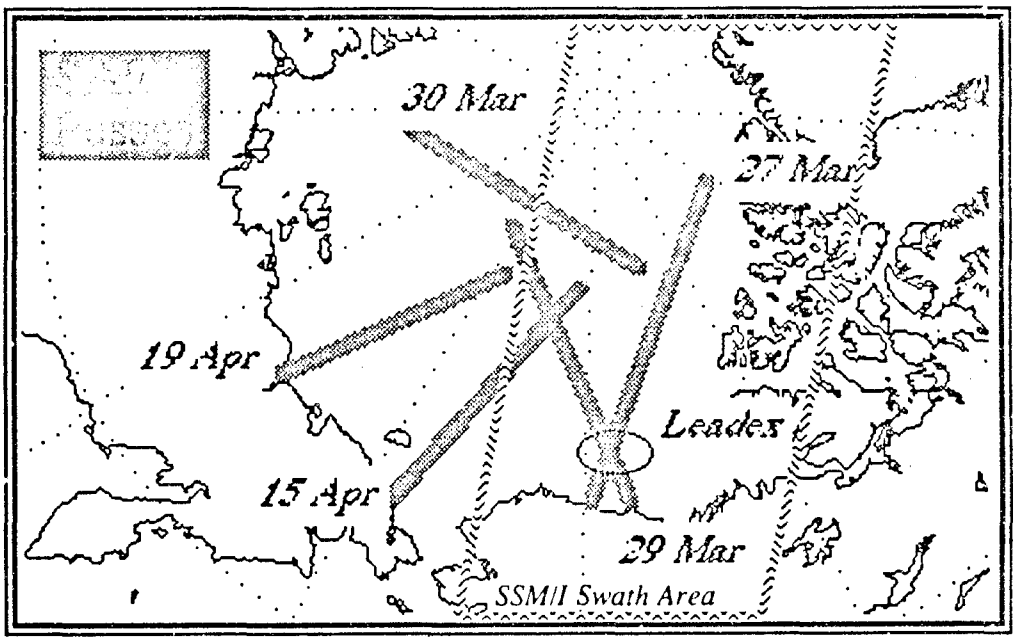

Figure 16. Imagery overlap for ice classification comparisons.

The SSM/I sea ice algorithm estimates ice concentration from brightness temperatures the sensor measures at 19 and 37 GHz vertical polarization. Under certain conditions, the 37 GHz horizontal polarization temperatures are also used. If the total ice concentration (TOTICE) estimated by the algorithm is greater or equal to 25%, the algorithm also makes an ice type identification. The identification is based on the values of the total ice concentration and the brightness temperature of the 19 GHz channel (T19) and is estimated from:

$$TBI = [T19-TC-(1.0-TOTICE)*180.]/TOTICE$$

where TBI = ice type identifier
TC = an atmospheric offset threshold that changes with season

If TBI is less than or equal to 238, the ice is considered predominantly multi-year. If TBI is greater than 238, the ice is considered predominantly first year. Essentially, when the brightness temperature is low (as would be expected for thick, less saline ice) and the total ice concentration is high, multi-year ice would be expected. When the brightness temperature is high (as would be expected for thin, more saline ice) and the total ice concentration is low, first year ice is expected. While a version of the algorithm which calculates actual percent concentration by type exists, it was not implemented operationally because of its longer running time. See *Hollinger, 1991* for a full description of the algorithm and its validation.

SSM/I ice concentrations and type identifications were found for the area covered by each SAR image, and the ice concentration results from each source were then plotted against each other. The results of this analysis are presented in Figure 17. They highlight the disadvantage of the SSM/I algorithm regarding ice type. Both SAR and SSM/I algorithms correctly estimated 100% ice concentration for all areas of analysis (this can be visually verified in the SAR imagery), but the SSM/I algorithm failed to consistently identify the dominant ice type properly. The SSM/I CAL/VAL algorithm estimated all regions of analysis except for one as being predominantly multi-year ice, even those areas which are visually 100% first year ice. Because ice concentration was

close to 100% for every case, an evaluation of how concentration from SAR compared with concentration from SSM/I was not performed.

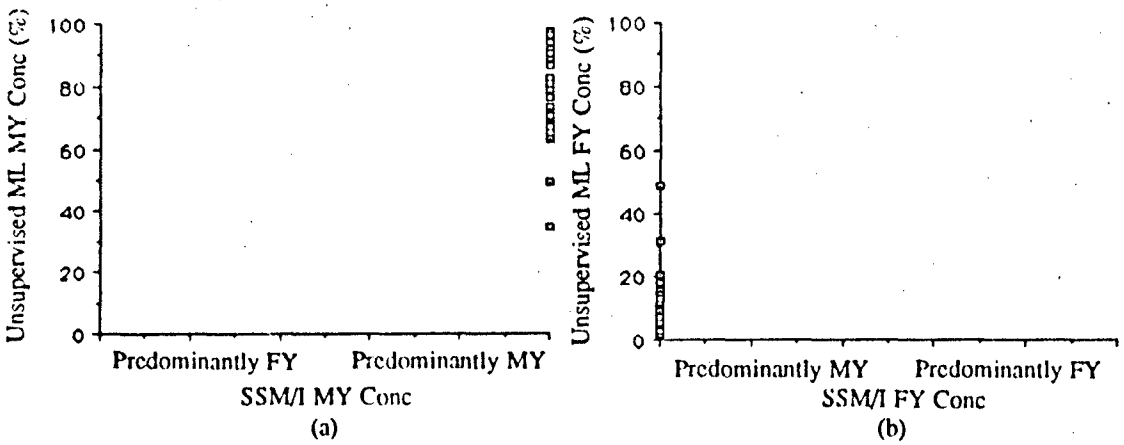

Figure 17. Ice concentrations from SAR compared with concentrations from SSM/I using the Navy CAL/VAL algorithm. a) Multi-year ice b) First year ice.

Ice type maps from the SAR ice classification algorithm were also compared with JIC Western Arctic Analysis charts and aerial data collected during the LEADEX experiment. The JIC had recently begun to use SAR imagery as a data source at the time of this analysis. The ice types and concentration estimates generated by unsupervised maximum likelihood classification were compared to those of the JIC charts, as well as to LEADEX ground truth (*Shuchman and Onstott, 1992*), when a good correlation in location could be achieved. The results of the comparison to the JIC charts are presented in Table 6 and in Figures 18a through c. Since the JIC ice analyst estimates concentrations to the nearest 10%, the plots appear discrete. The SAR concentration estimates are fairly well correlated to the JIC data for most cases, although differences occasionally reached 30%. In a few cases, as on 30 March 1992 (see Table 6), the SAR and JIC estimates were radically different.

Table 6. Comparison of SAR and JIC Ice Estimates.

		SAR Concentrations			JIC Concentrations		
Lat.	Lon.	MY	FY	NI	MY	FY	NI
27Mar92							
71.62	-147.56	54.2	45.5	0.3	70.0	20.0	10.0
72.36	-146.63	68.3	30.5	1.2	70.0	20.0	10.0
73.85	-144.62	70.6	26.6	2.8	70.0	20.0	10.0
75.30	-142.89	82.6	17.1	0.3	90.0	10.0	0.0
76.74	-139.27	88.9	10.9	0.2	90.0	10.0	0.0
80.12	-128.75	90.3	9.3	0.4	90.0	10.0	0.0
83.33	-103.45	97.7	2.2	0.1	90.0	10.0	0.0
29Mar92							
71.89	-142.66	65.1	33.4	1.4	40.0	50.0	10.0
72.25	-155.14	63.1	35.5	1.4	40.0	50.0	10.0
72.64	-143.56	67.6	30.4	2.1	40.0	50.0	10.0
72.99	-154.16	65.5	30.3	4.2	40.0	50.0	10.0
73.38	-144.58	67.0	30.3	2.8	40.0	50.0	10.0
73.74	-153.15	79.1	18.9	2.1	90.0	10.0	0.0
74.84	-146.88	80.7	18.5	0.8	90.0	10.0	0.0
78.38	-154.82	90.8	9.1	0.1	90.0	10.0	0.0
81.98	-171.88	91.0	9.0	0.0	90.0	10.0	0.0
30Mar92							
82.18	134.78	73.8	26.2	0.0	90.0	10.0	0.0
82.19	-140.05	96.9	3.1	0.0	90.0	10.0	0.0
83.52	150.84	0.0	81.2	0.1	90.0	10.0	0.0
83.82	157.51	0.0	99.9	0.1	90.0	10.0	0.0
84.09	-172.15	94.8	5.1	0.1	90.0	10.0	0.0
84.16	172.44	40.8	59.2	0.0	90.0	10.0	0.0
15Apr92							
68.03	-174.21	0.0	100.0	0.0	0.0	100.0	0.0
71.05	-171.23	25.8	74.2	0.0	40.0	50.0	10.0
71.34	-145.74	34.7	65.0	0.0	40.0	50.0	10.0
71.80	-170.41	28.4	71.3	0.3	40.0	50.0	10.0
72.55	-169.47	63.8	35.5	0.7	40.0	50.0	10.0
72.84	-143.92	67.3	30.1	2.6	90.0	10.0	0.0
73.58	-142.88	71.9	24.1	4.0	90.0	10.0	0.0
76.19	-163.44	78.1	21.4	0.5	90.0	10.0	0.0
78.97	-156.12	91.7	8.2	0.1	90.0	10.0	0.0
19Apr92							
72.38	-138.06	79.4	18.6	3.0	90.0	10.0	0.0
72.53	163.92	1.8	97.6	0.6	0.0	100.0	0.0
74.00	166.02	43.4	56.6	0.0	0.0	100.0	0.0
75.46	168.51	59.4	40.6	0.0	90.0	10.0	0.0
76.89	171.43	59.3	40.7	0.0	90.0	10.0	0.0
78.96	177.23	85.2	14.8	0.0	90.0	10.0	0.0
80.26	-177.66	86.5	13.5	0.0	90.0	10.0	0.0

The LEADEX ground truth used as a comparison covered only a small fraction of the area analyzed but still provided some useful information. In cases where comparisons could be performed, the SAR classification products correlated well with ice reconnaissance observations (Table 7).

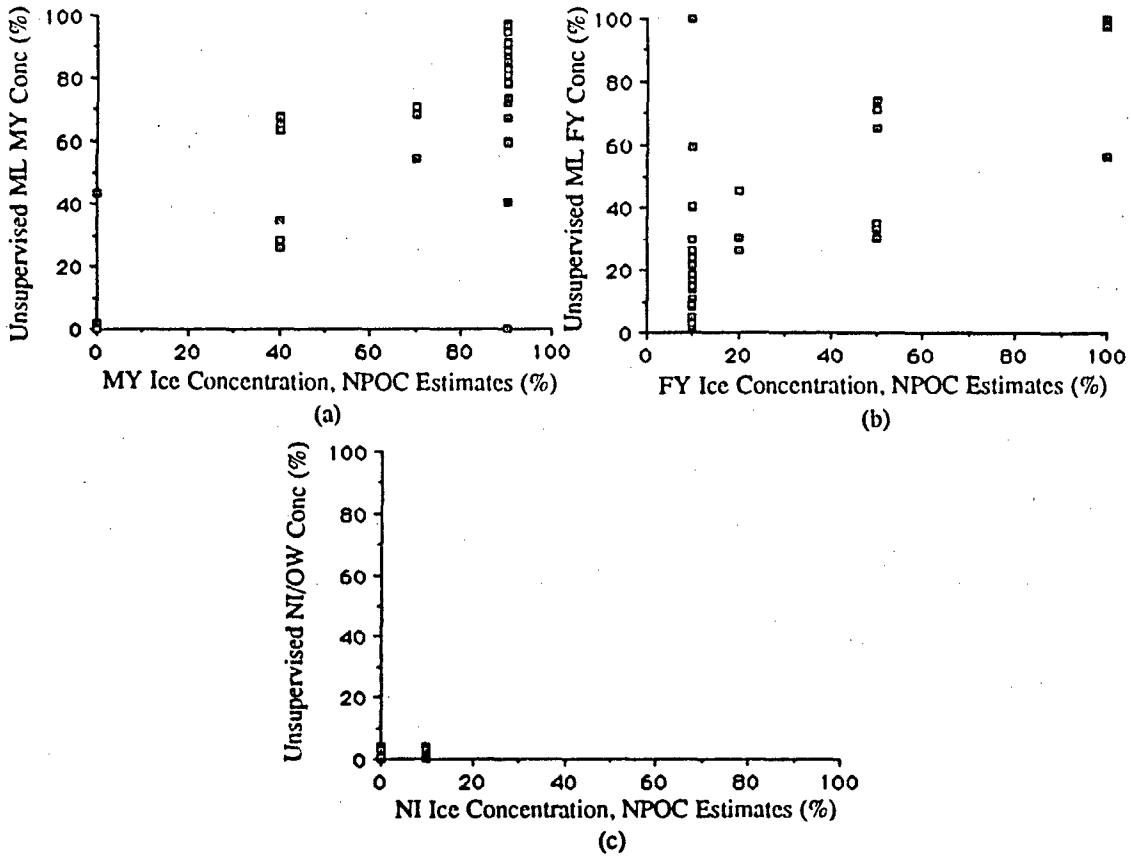

Figure 18. Ice concentration from JIC ice charts compared to concentration from SAR for a) multi-year ice, b) first year ice, c) new ice and open water.

Table 7. Comparison of Unsupervised SAR Classification to LEADEX Ground Truth

Date	Latitude, Longitude	Ice Reconnaissance	SAR Ice Classification
29Mar92	76.58,-144.87	250m refrozen lead	linear feature I#4
29Mar92	72.80,-143.88	300m refrozen lead	linear feature I#4
29Mar92	72.95,-143.88	thin FY with snow	ice type 4
29Mar92	73.03,-144.05	thin FY with snow	ice type 4
29Mar92	73.08,-144.53	thin FY with snow	ice type 4
29Mar92	72.73,-144.90	300m older lead with snow	linear feature I#4
29Mar92	72.88,-144.53	multi-year pan	MY
29Mar92	72.88,-144.55	multi-year floe	MY
29Mar92	72.88,-144.54	refrozen lead	ice type #4
15Apr92	72.65,-145.83	FY NE/SW shear zone	pattern of linear FY NE/SW features throughout image

3.5 Problems Encountered and Recommendations

Few problems were encountered while using the ice classification algorithm. The algorithm ran quickly and with a minimum of effort on the part of the user. This algorithm is easily incorporated into an operational scenario, with the user queueing up images to be classified

in the morning and having results available by mid-morning. However, users could benefit from more accessible ice classification results. At this time the user must load all the images into the image cache before any results can be presented. This creates a problem if the number of images to be examined is larger than what the image cache can hold. The availability of a listing of ice classification results, with image location, date, and ice concentrations, would make synoptic studies and comparisons to other data sources much easier.

With a few exceptions, performance of the algorithm was such that an analyst would be able to have confidence in the percent by type estimates for each 100 km square classification map without having to examine the image from which the map came. There were, however, two consistent errors: the algorithm under-estimated the amount of first year ice in the images by confusing those areas of first year ice having low backscatter with new ice or open water, and low backscatter pixels within multi-year floes were misclassified as first year ice. For the images we evaluated these misclassifications did not have a significant impact on the error expressed as the difference between supervised and unsupervised (algorithm) classification. However, they are potentially serious. For instance, an image with 100% multi-year ice would be classified as having a significant amount of first year ice, and broad areas of first year ice can appear in a classification map as new ice or open water.

These problems can be traced to the backscatter distributions of the ice types in the look-up table used by the algorithm. Although many data points went into the development of the look-up table, the distributions in the look up table do not entirely reflect the natural variability in backscatter signature exhibited by ice types. For instance, multi-year ice with a high concentration of melt ponds will have a backscatter value lower than that of "normal" multi-year ice, while first year ice, if it is undeformed, can have a signature much darker than that of "normal" first year ice. In short, all ice types (but particularly first year and young types) can exhibit a wide range in backscatter depending on a variety of factors including meteorological conditions and ice history. Additionally, the backscatter value of open water can overlap that of all ice types depending on wind strength. Improving the estimation of the mean for these distributions by collecting more samples over more of the Arctic is recommended. Widening the standard deviation of the classes to accommodate all natural variability is not, because it would result in much larger overlaps between classes, which would defeat the purpose of improving classification.

Turning aside from this problem to focus on improving algorithm performance, it is recommended that the maximum likelihood version of the algorithm be used in place of minimum distance. In addition, the accuracy of the algorithm could be improved by judicious "tuning" of the classifier. Tuning involves adjusting the decision parameters of the classifier, which are the backscatter distributions and the a priori probabilities for the ice types. Recommendations are to:

1. adjust backscatter distributions on the basis of an expanded backscatter signature data base. Constructing the data base should be part of a study which has additional objectives of better defining the regional and seasonal domain of the algorithm by documenting backscatter variability over place and time.

2. adjust a priori probabilities based on a study of Joint Ice Center Western Arctic analyses, which can serve as a climatology of concentration by type for different regions and seasons. Investigate additional methods of adjusting a priori probabilities, for instance by examining the size of ISODATA clusters produced by the algorithm.

3. include spatial information. A single pixel classified as first year ice, for example, is probably multi-year ice if it is surrounded by many pixels classified as multi-year ice. It should be fairly straightforward to modify the algorithm to include spatial information, and to document any improvement which results.

While none of the above recommendations can be implemented immediately, they only require study which can be carried out with existing data over a year or less under the 6.3 SAR Applications in Polar Regions project. In contrast, addressing the problem of errors caused by the

natural variability of surface types requires more creative approaches of the kind usually developed under 6.2 research programs. This is because a simple backscatter-based algorithm will always fail when backscatter for the surface being classified mimics that of another type (e. g. wind-roughened open water or frost flowers on new ice mimicking multi-year ice), and additional information must be considered. Three lines of inquiry are recommended. While the first is being incorporated into the algorithm in a limited way now, the second and third require substantial development time. These lines of inquiry are to:

1. add weather fields. It is important to know what the surface temperature is in order to know which seasonal look-up table to use. The backscatter of ice changes dramatically as the surface temperature rises from -5°C to near zero. In addition, wind influences the backscatter of open water. If it is known that the wind is blowing more than a few meters per second, the backscatter value of water in the look-up table can be set to exceed that of multi-year ice, and water can therefore be correctly identified.

2. use an expert system. Expert or rule-based systems range from simple to extremely complex. For instance, rules can tell the algorithm what backscatter distributions to use in certain regions in order to accommodate regional variability in backscatter signatures. A much more comprehensive system can approach classification the way an analyst would, by considering factors such as floe shape, climatology, and the recent history of the area. The current backscatter algorithm would then be only a part of much larger and necessarily more complex system. It is not clear that the increased complexity of a such comprehensive expert system over the current backscatter-based algorithm would buy significantly improved results. A simple rule-based system, however, could improve classification accuracy for a relatively low development cost.

3. improve knowledge of the link between backscatter signature and ice type. For instance, the backscatter of first year ice forms two distinct clusters. This was unanticipated before the first ERS-1 data were received. As a consequence, JPL increased the number of ice types classified by the algorithm from 3 to 4 by adding "ice type 4". While "ice type 4" is first year ice with a relatively dark signature, ground truth from Leadex indicated that it is not always thinner than "normal" first year ice. The best ice type label for "ice type 4" is yet to be determined. In addition, the look-up tables do not contain, nor could be expected to contain, the wide variety of ice types which exist in nature. At present, ice types which have no match in the look-up tables are automatically classified as ice types which do exist in the table. Further study using existing satellite data must be performed in order to identify any additional backscatter clusters and tie them, when possible, to ice types.

4.0 Ice Motion Product Evaluation

4.1 Algorithm Description

The ice motion algorithm (*Kwok et al., 1990*), while straightforward like the ice classification algorithm, is more complex. First, two images are selected by the user. Of the two images, the one acquired earliest in time is labeled the source image, while the later image is the target image Ice motion analysis, like ice classification, is performed on geometrically and radiometrically corrected images. The fidelity of these corrections is essential to the algorithm, as errors in either can cause significant errors in the ice motion results. Errors in the geometric rectification of the data can cause errors in ice motion estimates by causing the ice floes in the image to be misshapen or incorrectly located in the image. Errors in image calibration can cause ice floes to have inconsistent backscattering coefficients which, in turn, can cause the algorithm to misidentify features.

Different locations in an ice field are characterized by different types of ice motion. In the central ice pack, the motion is mostly translational. In the coastal zone and ice margin areas, the ice floes may have a large rotational component to their motion, as well as a large translational

component. Different procedures are followed by the algorithm to provide for the expected differences in ice motion due to location within the ice field. The position of the ice edge is taken from climatology.

Within the central ice pack, two passes (high and low resolution) of an area-based matching technique are used for the estimation of ice motion. In the first pass, a two-dimensional cross correlation between source and target image is performed at a reduced resolution. The correlation peak obtained from this correlation directly gives the spatial lag of image patches, which are simply subimages taken from the main image. The approximate translation vectors produced by this pass then become a guide to a final pass which also employs a two-dimensional cross-correlation, but at full resolution. The similarity measures of this correlation are two-dimensional normalized cross correlation coefficients. Fourier transform processing makes the procedure computationally efficient.

For the marginal ice zone, where rotation of ice floes is expected, the procedure varies slightly. First, ice in the source and target images is classified and floes are identified. The boundaries of the ice floes are then parameterized as curves. Using a technique called psi-s curve matching, the algorithm uses the parameterized curves to match floes in the source image with floes in the target image. Rotation and translation vectors calculated from this technique serve as hints for an area matching technique identical to that used within the ice pack. For both pack ice and marginal ice, resultant ice motion vectors are checked for consistency and filtered to remove bad matches. The finished product is a map of ice motion vectors sampled on a 5 km grid, which can be overlain on the source image.

For supervised ice motion analysis, the user provides a first estimate of the ice motion. The rest of the procedure is identical to that of the unsupervised ice motion analysis. The supervised ice motion algorithm on the JIC workstation was not working at the time of this analysis and so was unavailable for testing. Users have voiced no requirement for this function.

4.2 Running the Algorithm

Approximately 50 pairs of images have undergone unsupervised ice motion analysis. Like the ice classification algorithm, ice motion employs the INGRES database. The user first selects candidate source images via an INGRES database form. Once the source images are selected, the user proceeds to select a target image. In general, the user does a rough selection of candidate target images by viewing them in the "view low resolution images" process before entering the ice motion application. The objective of the user is to find images from different days which have the same ice floes or other features in them. The user can also let the system select candidate target images by specifying the location, the period of time which is allowed to pass between source and target image, or an ice motion estimate. With either method, the user must look at the images to make sure that common ice features exist. Once a source and target image are selected the pair is submitted to the system for processing. Several motion pairs may be submitted at one time, but submitting more than about seven or eight can cause the system to run very slowly. When the algorithm has finished with a motion pair, an alert will appear in the operator alert manager window. As the user "handles" this alert the ice motion results are ingested into the INGRES database and the product is displayed on the screen for the user's inspection. The product appears as pair of source and target images with either relative or absolute motion vectors overlain on the source image. The mean motion and direction of motion are also displayed in x-y coordinates.

The unsupervised ice motion algorithm runs slowly. It takes about 20 minutes to complete one ice motion pair. Data selection using the INGRES databases is straightforward and easy to do. It takes about ten minutes to select both source and target images. Loading the desired data into the system may take anywhere from five to fifteen minutes depending on the location of the data on the optical disks.

4.3 Comparison of Unsupervised Ice Motion Results to Ice Motion Results From Other Sources

As with the ice classification algorithm, results from the SAR ice motion algorithm were compared to results from other data sources. Evaluation of the unsupervised ice motion algorithm was performed through a comparison of SAR ice motion results to those obtained from AVHRR and OLS imagery. AVHRR Channel 4 and OLS IR imagery of the Beaufort Sea collected during the LEADEX experiment were processed to pixel sizes of 1 km and 600 m respectively. (The resolution of AVHRR varies from 1 km at nadir to 4 km at the scan limbs, while that of OLS is 0.6 km across the scan.) Images were then selected from this set of data based on date of collection and degree of cloud cover. Only those images which were mostly cloud free and were close in time to the dates of collection of SAR data were selected. Unfortunately, most of the AVHRR and OLS imagery had significant cloud cover, so the data set used in this comparison is small. Table 8 provides a list of the images which were used in this comparison and Figure 19 illustrates the areas of common coverage. SAR motion results were also compared to predicted ice motion calculated using the Polar Ice Prediction System (PIPS) model (*Preller and Posey, 1989*). The PIPS model provides arctic wide predictions. For each time entry in Table 8, the first number in the column is the time interval between source and target images, and the second number is the time of acquisition of the source image. PIPS estimates are from models which predict fields at 0000Z.

Table 8. Image sources and acquisition times for motion comparison.

Case	SAR delta T	AVHRR delta T	OLS delta T	PIPS delta T
Mar26-29	72hrs, 06:49Z	71hrs, 14:11Z	-	72hrs, 00:00Z
Mar26-29	72hrs, 06:49Z	-	73hrs, 15:10Z	72hrs, 00:00Z
Mar26-29	72hrs, 22:03Z	71hrs, 14:11Z	-	72hrs, 00:00Z
Mar24-26	32hrs, 21:28Z	48hrs, 14:35Z	-	48hrs, 00:00Z
Mar24-27	72hrs, 21:28Z	71hrs, 14:35Z	-	72hrs, 00:00Z
Mar24-29	104hrs, 21:28Z	119hrs, 14:35Z	-	120hrs, 00:00Z

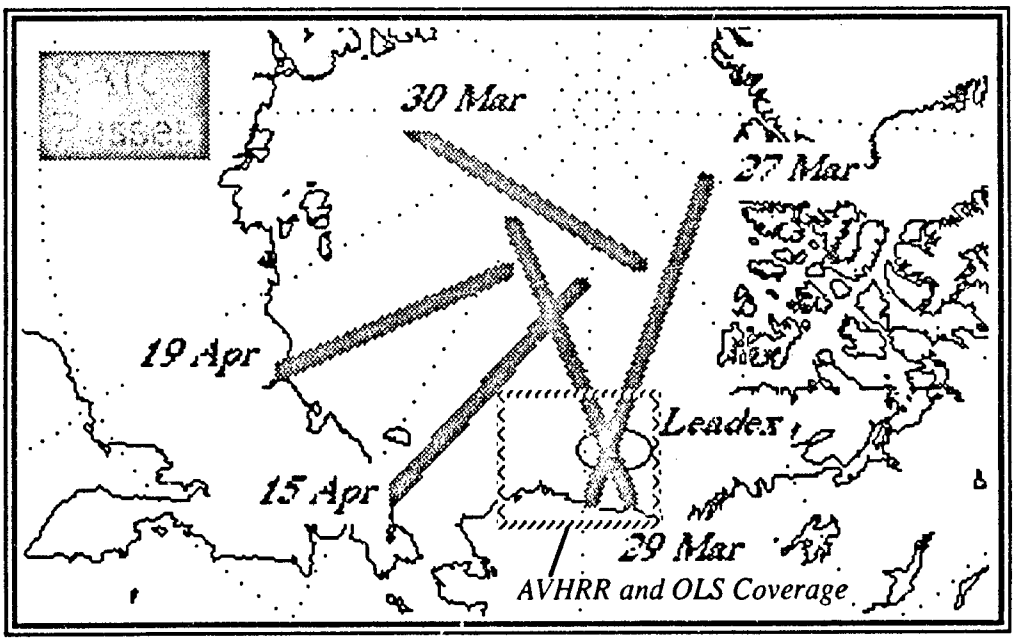

Figure 19. Map of data coverage for ice motion comparison

To make the comparisons, motion vectors for the OLS and AVHRR image pairs were generated using NRL's cross-correlation motion algorithm (*Emery et al., 1991*). For each source and target image pair, areas of SAR coverage were located on the AVHRR or OLS imagery, and the motion vectors within those areas were averaged to produce a mean velocity measurement. In cases where, due to clouds, no motion vectors were obtained in the area of SAR coverage, an average of vectors in close proximity to the area was taken. The value and heading of the mean velocity were then compared to values calculated in the same way for the SAR image pairs. Results are presented in Table 9. In general, the headings calculated from the various imagery sources are consistent. With the exception of the March 24/March 29 data pair, the mean difference in estimated heading is only eight degrees. The estimations of ice velocity, however, are not consistent, although the SAR imagery produced larger estimates of ice velocity in every case. The discrepancies between the ice velocity estimates are largely due to geolocation error in the AVHRR and OLS imagery. Often the Alaskan coast was cloudy, which made it impossible to accurately shift imagery for the best geolocation. Experience with a larger AVHRR data set from the region showed that, on average, images were shifted 4 km in both x and y (image pixel row and column) directions when manual correction to geolocation was possible. A 4 km average error, then, for unshifted images would produce errors in the velocity estimates of up to 5.7 km/day for a two day time interval, 3.8 km/day for a three day time interval, and 2.3 km/day for a five day time interval. When these uncertainties are taken into account, the discrepancy between the SAR velocity estimates and those produced from AVHRR and OLS is explained. An additional cause of the discrepancy between SAR and AVHRR or OLS motion estimates is different acquisition times. In general, the acquisition time for the SAR images lead that of the AVHRR and OLS images by about eight hours. SAR resolves deformation and shear in the ice pack at a scale much smaller than that obtainable with AVHRR or OLS. Since the resolution of the SAR imagery is much finer than that of the other data sources, the detail in the ice motion calculated from the SAR imagery is much greater than that of the other sources as well. A visual comparison of the difference in detail between the ice velocity estimates from different sources is presented in Figures 20 through 25. In the AVHRR and OLS cases, the images are oriented with north at the top, while in the SAR cases north is to the right.

Results of the comparison to PIPS estimates are also presented in Table 9. The SAR and PIPS ice motion estimates agree well. The average difference between SAR ice velocity estimates and PIPS ice velocity estimates is only 0.70 km per day, which is small, given the radical difference in resolution of the two motion analyses (5 km grid spacing for SAR and 64 km for PIPS). The heading estimates are also comparable in most of the cases.

Table 9. Comparison of ice motion estimates from various sources.

		SAR		AVHRR		OLS		PIPS	
Time Interval	Location	km/day	hdg	km/day	hdg	km/day	hdg	km/day	hdg
Mar. 26-29	72.86,-143.84	1.52	149°	1.19	169°	-	-	2.27	114°
Mar. 26-29	73.60,-144.89	1.92	146°	-	-	0.53	143°	1.69	97°
Mar. 26-29	72.30,-155.92	1.65	40°	0.19	50°	-	-	2.23	82°
Mar. 24-26	73.26,-145.47	9.90	317°	4.07	312°	-	-	10.69	326°
Mar. 24-27	74.01,-144.43	3.53	323°	1.66	315°	-	-	4.88	329°
Mar. 24-29	73.26,-145.47	1.88	318°	0.43	270°	-	-	2.35	333°

The accuracy of motion estimates from SAR is being evaluated by JPL using comparisons with drift from on-ice GPS locations. There is no doubt that by this measure, motion estimates from SAR exceed the accuracy of estimates from OLS, AVHRR, or PIPS. SAR's finer resolution allows users not only to observe average motion of the ice pack, but also provides information about deformation within the ice pack. In addition, with SAR, the ice motion in an area can be monitored regardless of cloud cover or lack of daylight, which is a feature no other sensor can provide. SAR does not, however provide the synoptic view which is obtainable with AVHRR, OLS, or a model, and this is a major drawback for ice analysts. Operationally, it is difficult to obtain sequential imagery for ice motion in a timely fashion, since the user must order imagery ahead of time and if the ice is moving quickly, it may be difficult to "catch" it in the ordered pass.

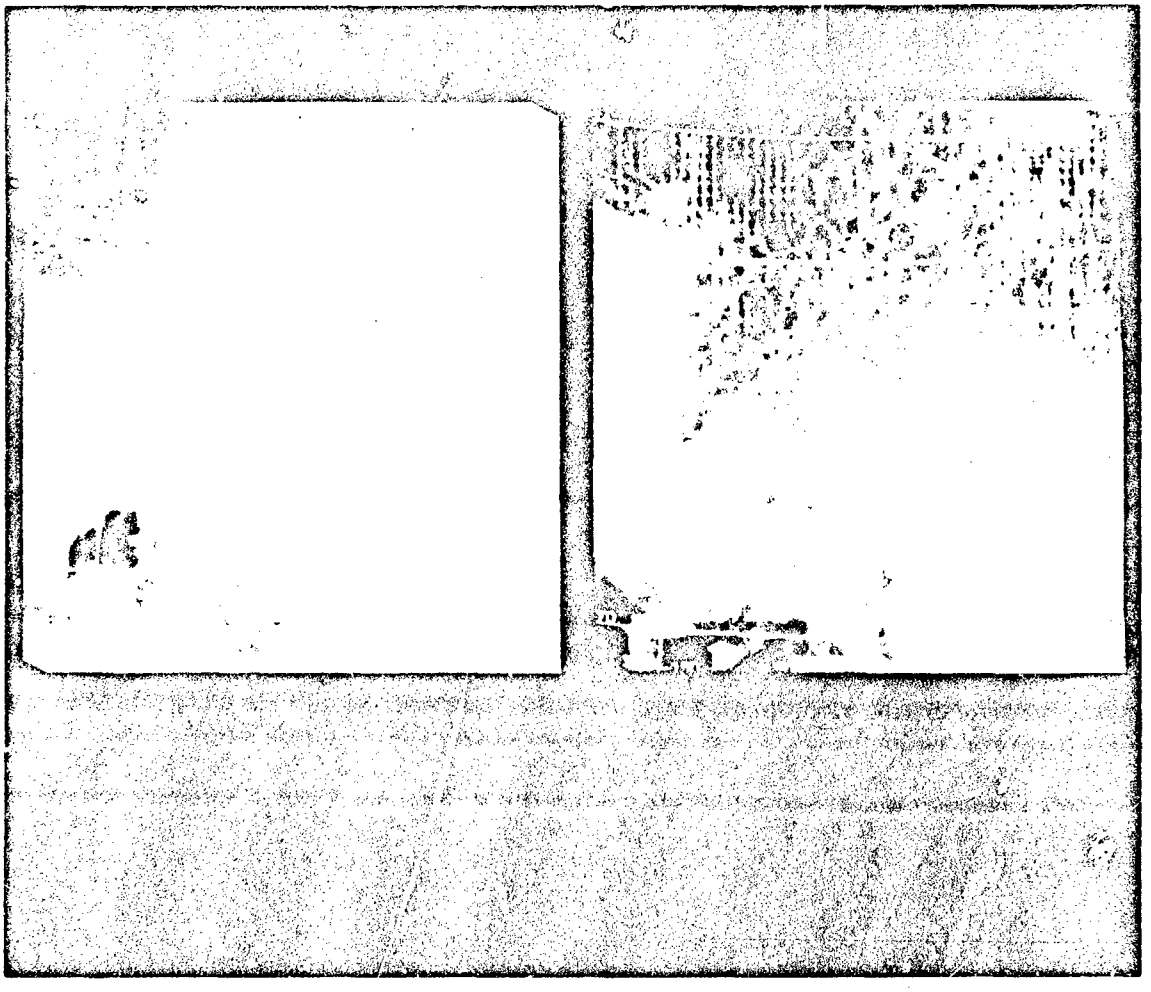

Figure 20a. AVHRR ice motion vectors, March 26-29, 1992.

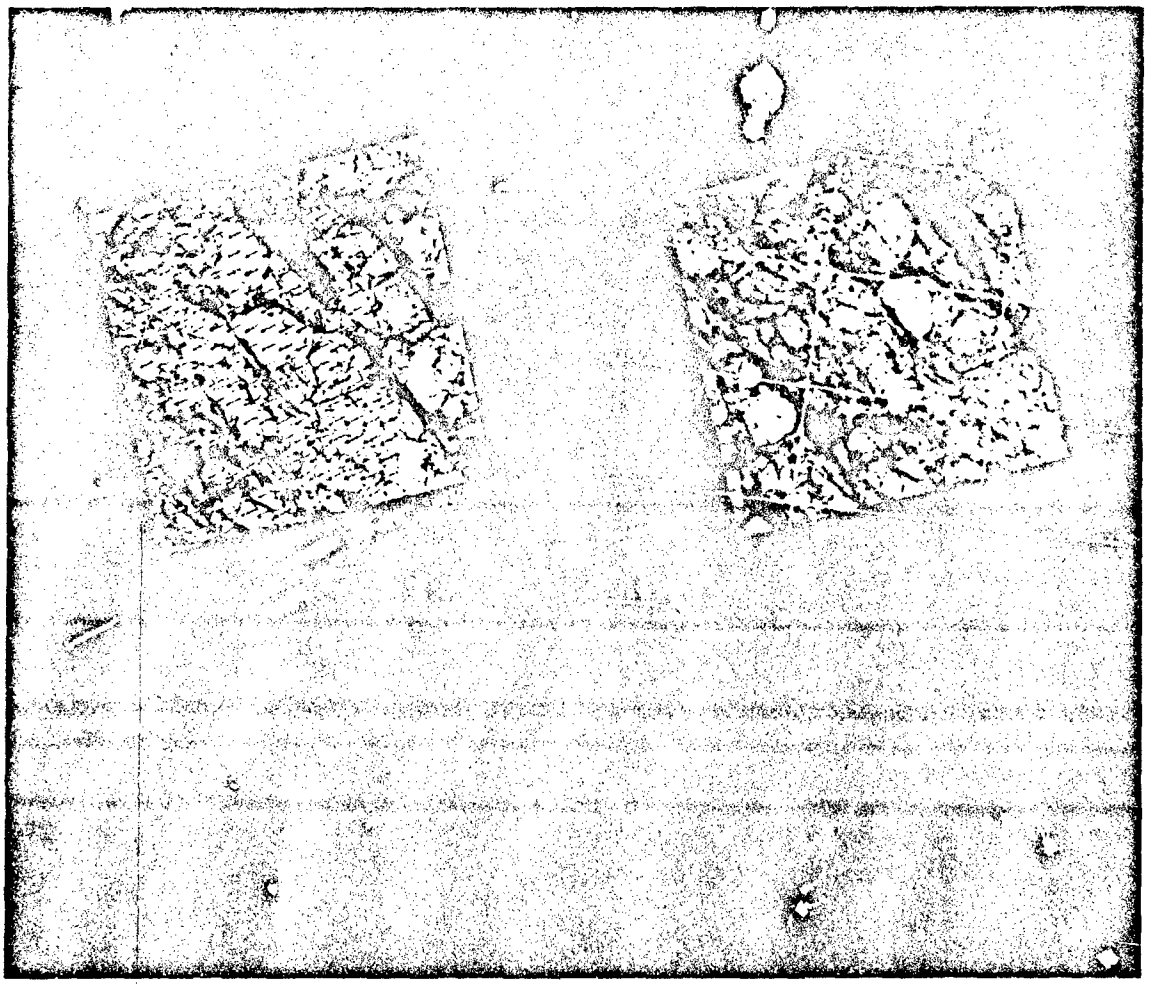

Figure 20b. SAR ice motion vectors, March 26-29, 1992.

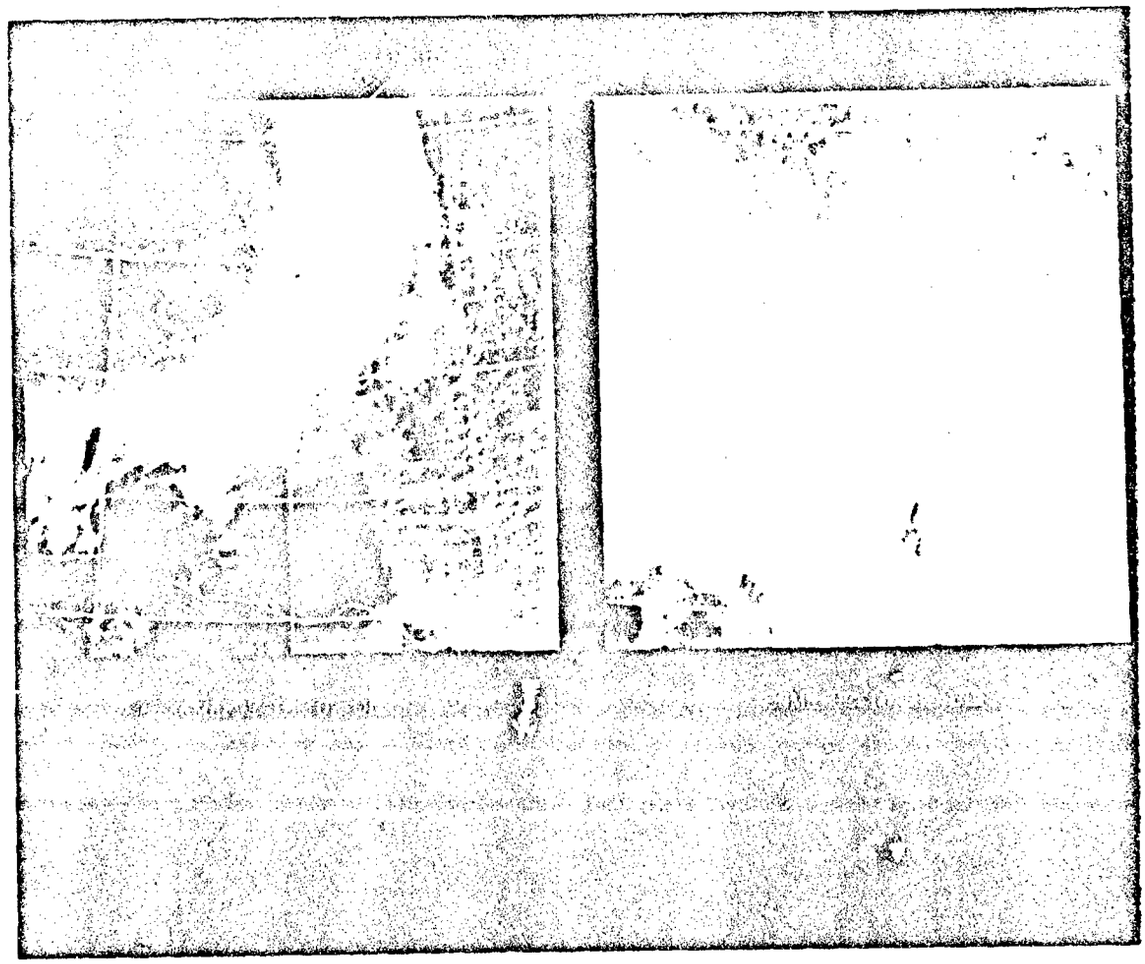

Figure 21a. OLS ice motion vectors, March 26-29, 1992.

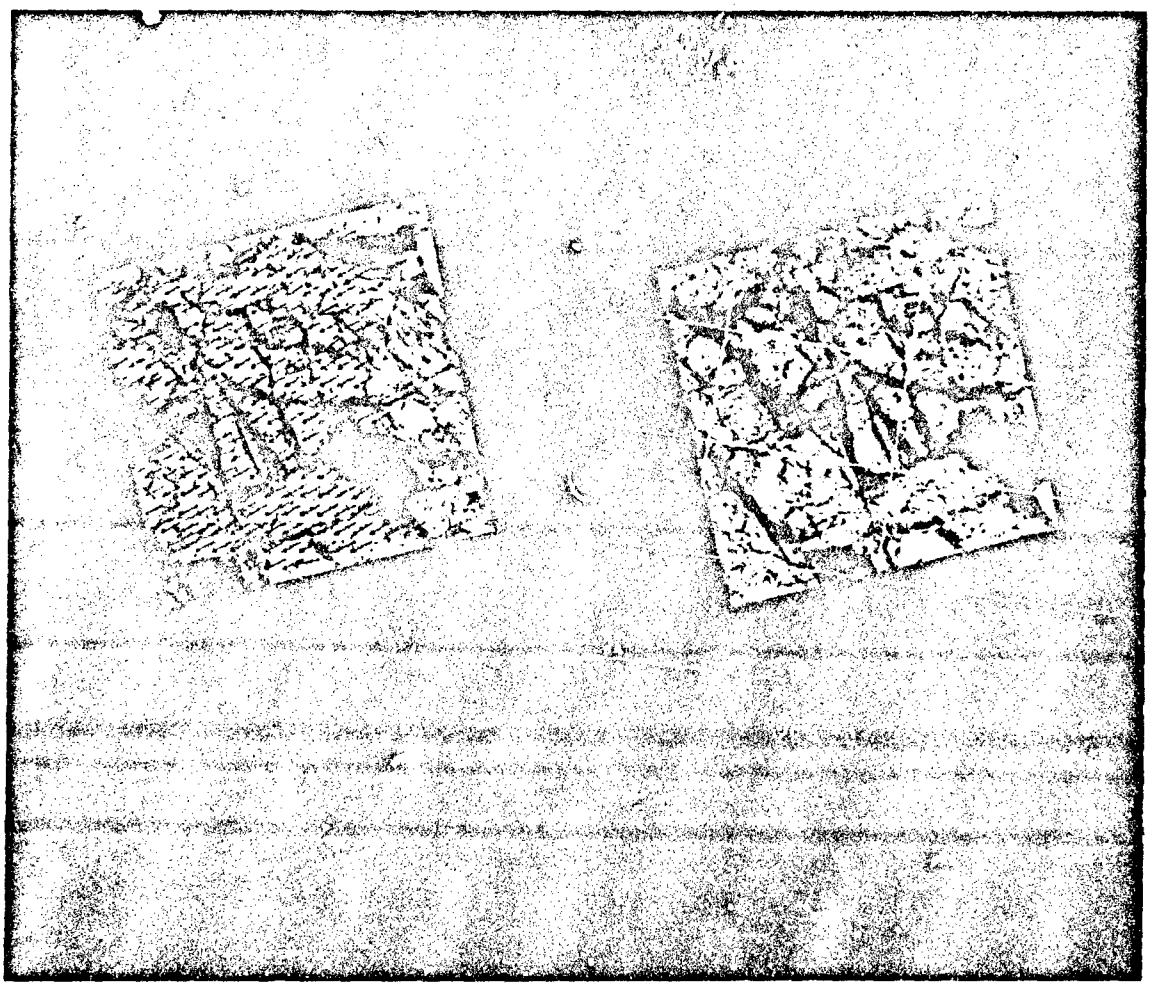

Figure 21b. SAR ice motion vectors, March 26-29, 1992.

Figure 22a. AVHRR ice motion vectors, March 26-29, 1992.

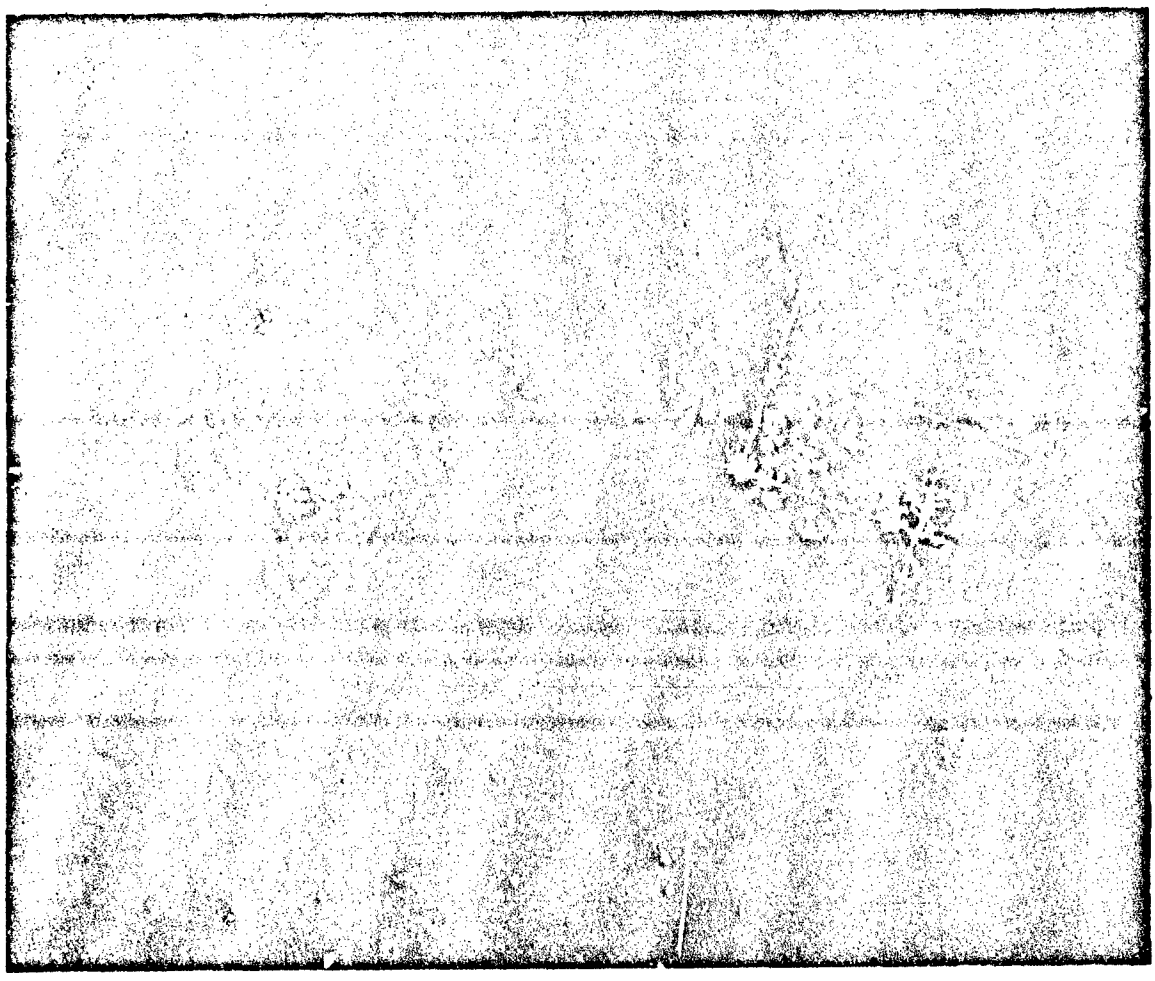

Figure 22b. SAR ice motion vectors, March 26-29, 1992.

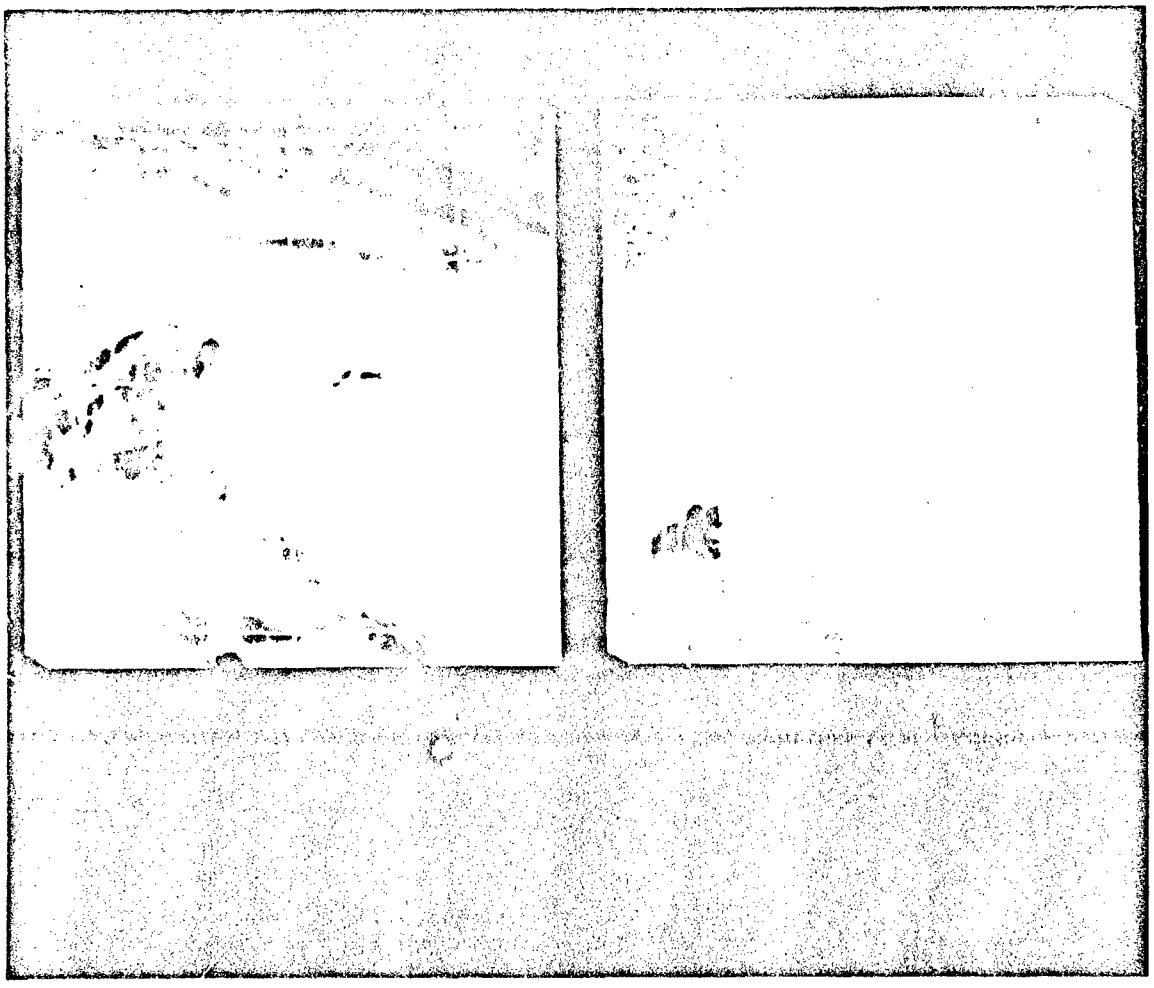

Figure 23a. AVHRR ice motion vectors, March 24-26, 1992.

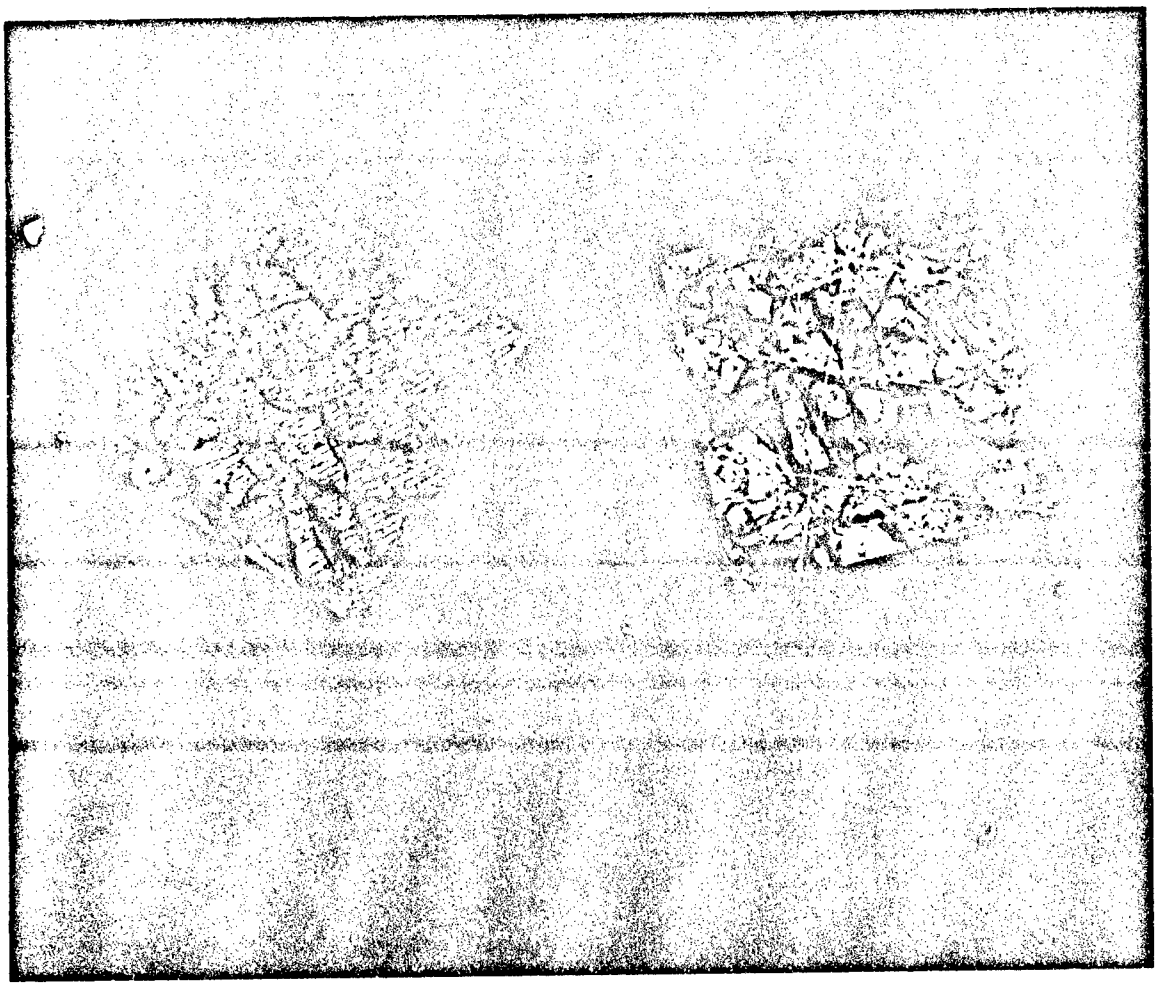

Figure 23b. SAR ice motion vectors, March 24-26, 1992.

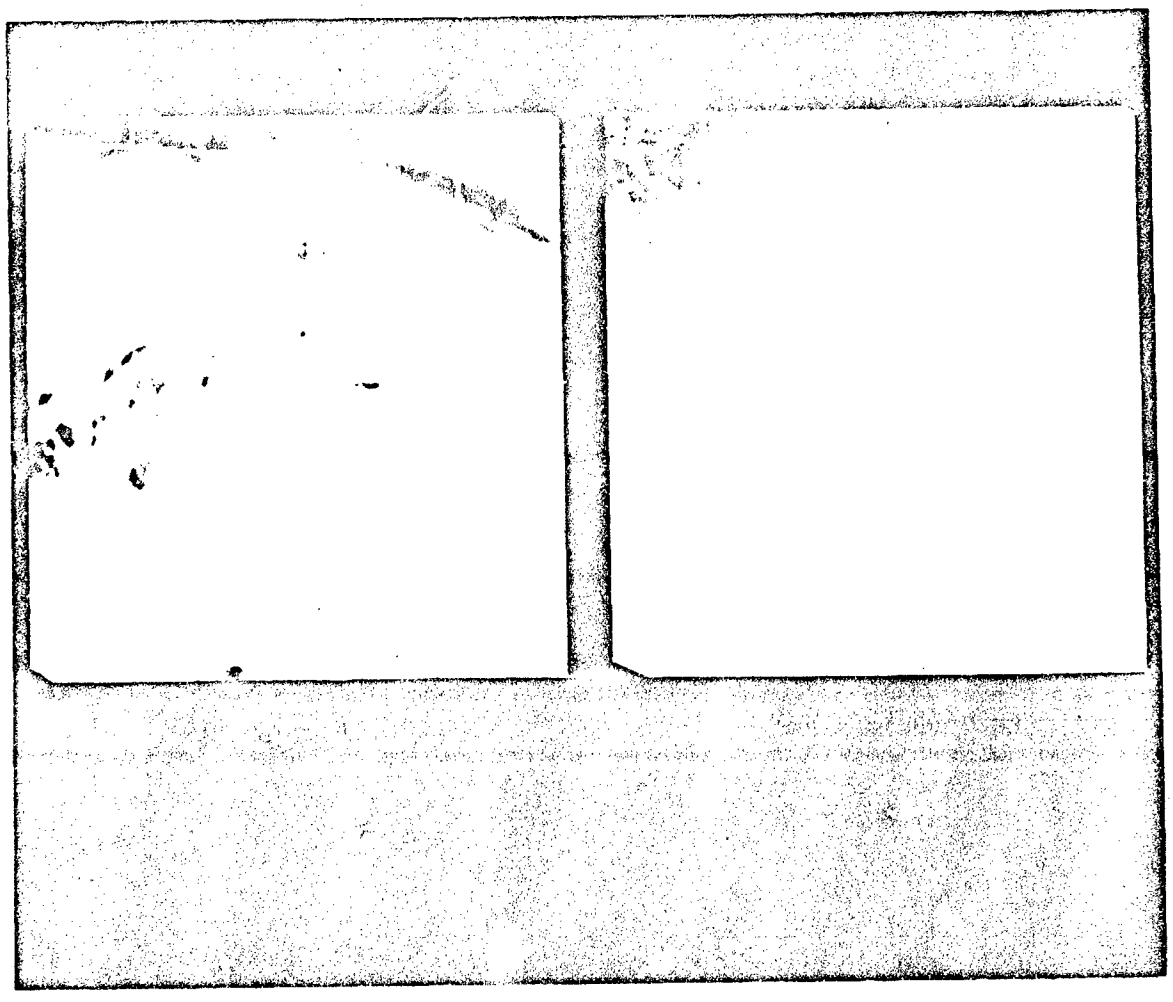

Figure 24a. AVHRR ice motion vectors, March 24-27, 1992.

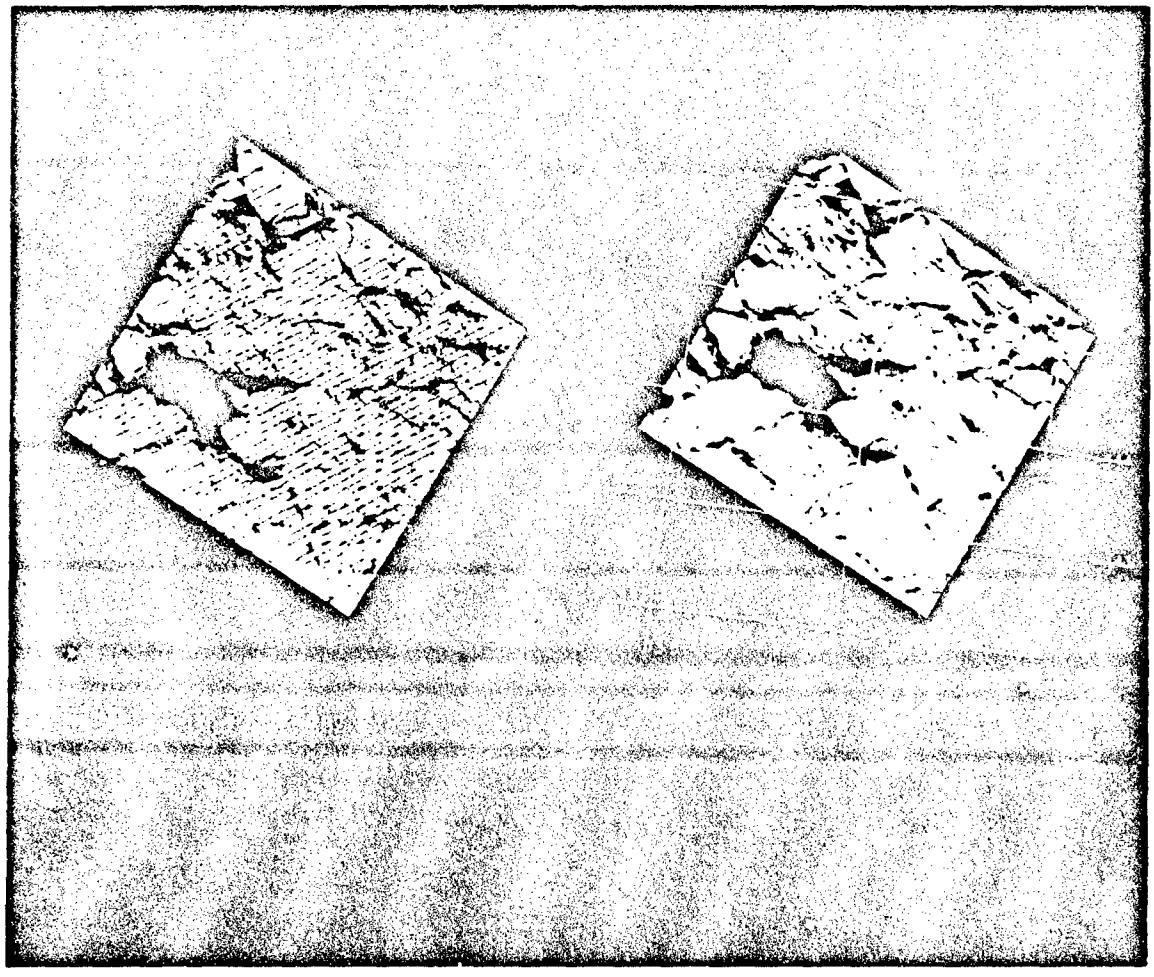

Figure 24b. SAR ice motion vectors, March 24-27, 1992.

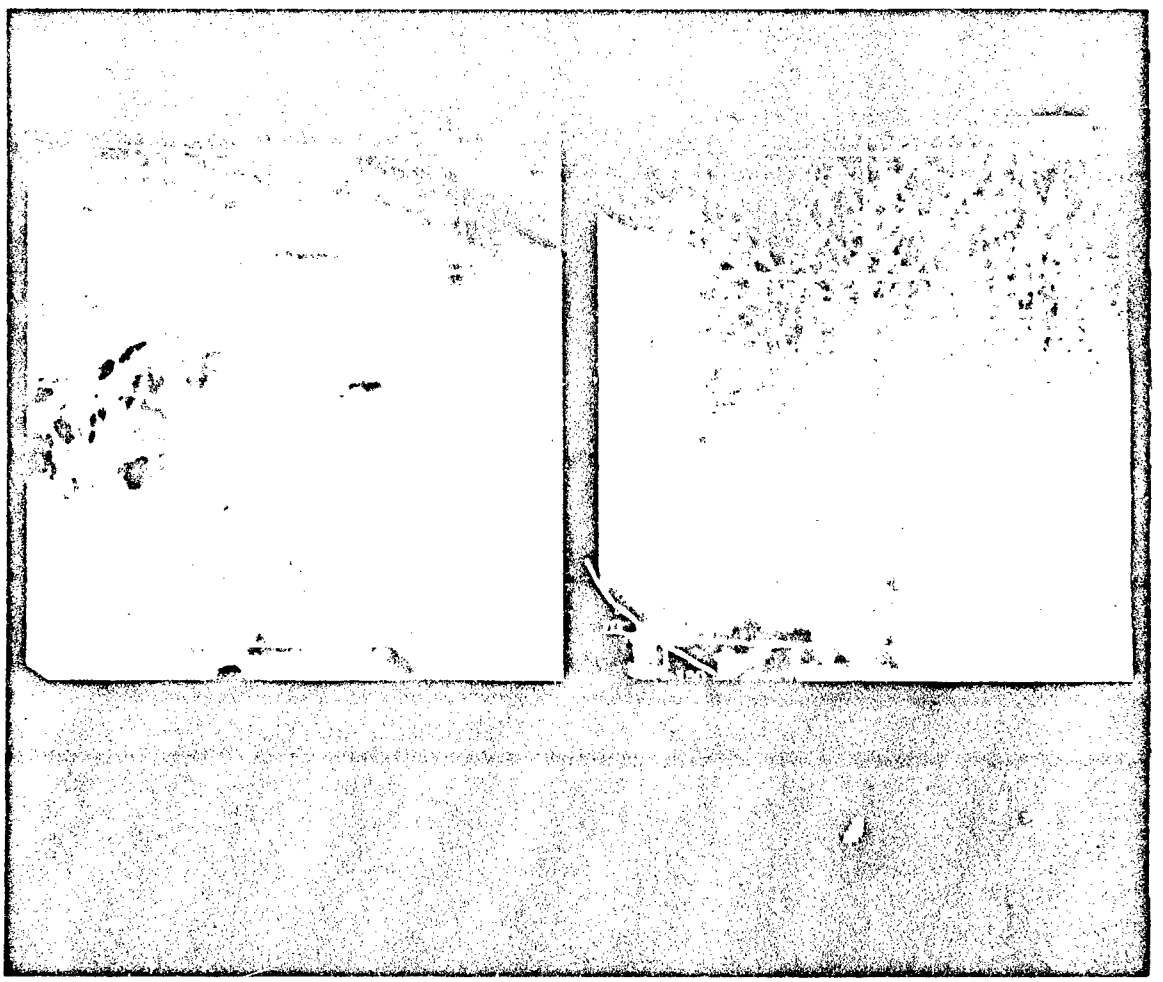

Figure 25a. AVHRR ice motion vectors, March 24-29, 1992.

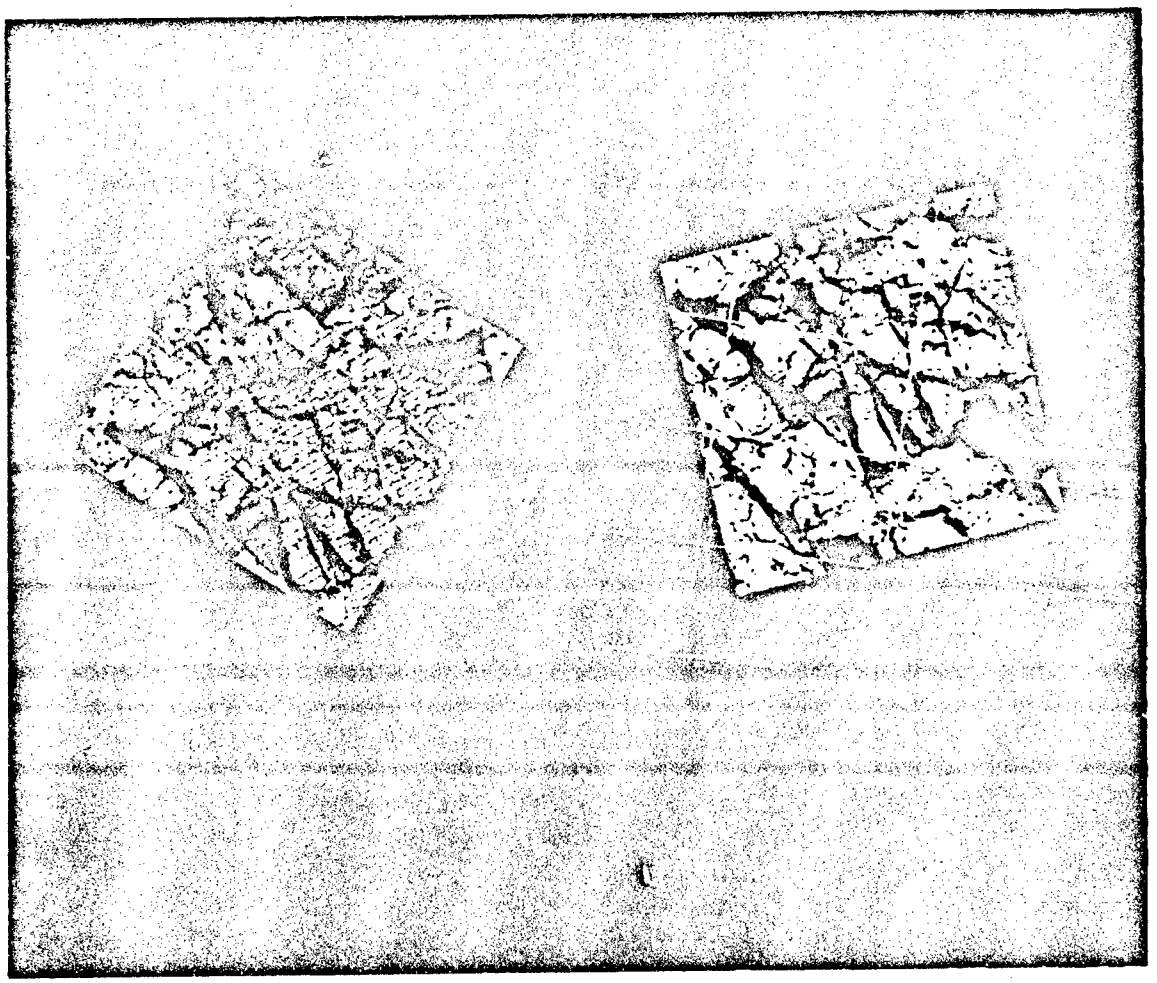

Figure 25b. SAR ice motion vectors, March 24-29, 1992.

4.4 Problems Encountered and Recommendations

Although generally straightforward to run, the ice motion algorithm on the SAR/JIC workstation is not as user friendly as the ice classification algorithm. Operationally, the supervised ice motion process should be made to work correctly or should be completely removed from the workstation to avoid user confusion. Another problem is the amount of time it takes to process imagery through the algorithm. Twenty minutes per image pair is an operationally unacceptable amount of time. The JIC has indicated a willingness to have motion vectors on a coarser grid than 5 km. Going to a coarser grid will speed up the processing time considerably. Long running programs make it more difficult to perform synoptic studies since less geographic area can be analyzed in a given time frame. Another limiting factor is the inability to batch large numbers of ice motion processes. If too many jobs are set up the processor slows considerably, and images needed by the motion program end up being removed from the image cache before they can be used. This causes fatal errors in the program.

The SAR motion algorithm has filters to remove vectors in which little confidence is placed. However in some cases, such as that shown in Figure 26, completely erroneous vectors remain in the final product. (In the evaluation above, images with erroneous vectors were not included). All motion vectors calculated by the algorithm go into the calculation of average motion for an image, so these erroneous vectors can cause the average motion calculation to be inaccurate as well. The addition of some method for the user to edit bad vectors out of the ice motion results and recalculate the average ice motion using only correct vectors is necessary. This would provide much more accurate average estimation of ice motion. It would also be more informative if the average motion could be provided as a speed and geographic heading rather than as the x and y grid displacement it is now. Speed and geographic heading can be much more easily compared to motion from other data sources. Recommendations are therefore to:

1. remove the supervised ice motion function from the workstation. This function has not been used at either NRL or JIC. In its present configuration, it is too difficult to use operationally. If it can be made to be more user friendly, it can be reinstalled in a future update.

2. decrease the amount of CPU time that it takes to run the algorithm. Calculating motion vectors on a grid coarser than 5 km is an easy alteration that would decrease operating time considerably. Changing to a 10 km grid, for instance, would decrease the CPU time to one fourth of what it is at present.

3. provide a point and click function by which users can remove erroneous vectors that the motion filters do not catch. Speed and geographic heading should be calculated from the final edited vector file after the user has removed all undesirable vectors.

3. provide the ice motion estimates as an average speed and geographic heading instead of an average x and y displacement.

5.0 High Resolution Data Analysis

5.1 Processing Specifics

The SAR/JIC workstation has the capability to receive and view full resolution ERS-1 data. The 30 m resolution data are processed at ASF to form an image with 15 m pixel spacing. The data are received by the system in much the same way that low resolution images are, and are placed in the /JIC/FULLSTOR directory. Full resolution data are usually sent over SARCOM in compressed form. However, uncompressed full resolution data can be sent as

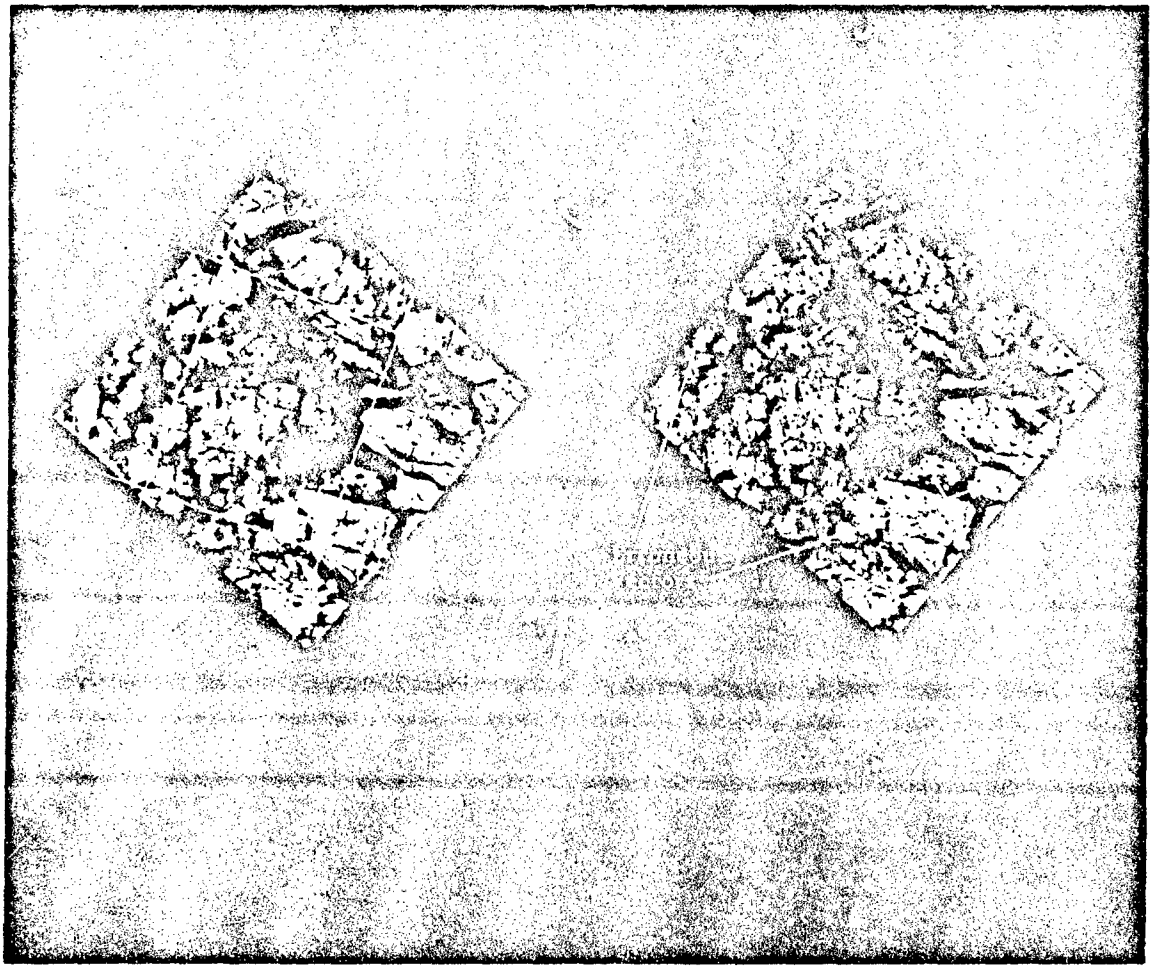

Figure 26. Example of erroneous motion vectors.

well. Uncompressed full resolution data, which occupies about 64 Mb of memory per image, is available for viewing immediately. Compressed 32:1 data must first be reconstructed using the JIC workstation software before viewing. The user selects the image to be reconstructed using standard INGRES forms and submits the reconstruction to the array processor. Only one reconstruction can run at a time, and it is advisable not to run any other processes, as this will cause the system to run unacceptably slowly. Unlike other processes, which will halt when the user exits the workstation but resume when the user enters it again, reconstruction will fail when the software is exited and must be restarted.

Once the user obtains either an uncompressed full resolution image or a reconstructed full resolution image , they can view the data using the "select framelet for geocoding" option in the "process standard products" menu. When this option is selected the entire full resolution image is displayed at an 8:1 screen pixel to image pixel ratio in a non-geocoded format. As the image is not geocoded, it may be somewhat distorted and flipped in time relative to its true earth location. The user selects a smaller portion of the image, or framelet, to be viewed at full resolution. The framelet is 1536 by 1536 pixels, or about 19.2 km by 19.2 km. If the user wants to keep the framelet for further analysis, they can select the "geocode framelet" option from the main full resolution menu. Selection of this function causes the framelet to be geocoded and stored in the framestor cache of the /JIC directory. The geocoded framelets can then be viewed using the "view framelets" option of the "view standard products" menu. Options such as image intensity scaling and graphics are also available in both the "select framelets for geocoding" and "view framelet" modes.

At this time, the only ice analysis function available for high resolution imagery is supervised ice classification. The methods of supervised ice classification are identical to those for low resolution data, and the results are much the same.

5.2 Problems Encountered and Recommendations

The most serious problems which arise when dealing with full resolution data are image artifacts due to the 32:1 image compression. Unlike their full resolution counterparts which have not undergone compression, full resolution reconstructed images have a distinctly "wormy" character, and image tile effects are evident as linear features in the imagery (Figure 27). In addition, the reconstructed images can contain defocused areas due to compression filtering. Although these artifacts do not appear to change the statistics of the image, they do appear to change the texture. Imagery compressed at an 8:1 ratio displays fewer artifacts than a 32:1 ratio and should be examined as an option (see Figure 28).

The artifacts of image compression can also cause errors in the ice analysis algorithms. When supervised classifications are performed on a full resolution image, an 8:1 compressed and reconstructed version of the image, and a 32:1 compressed and reconstructed version of the image, the classification results differ by amounts up to 5% (Table 10). A decrease in the dynamic range of the data can also be observed in the backscatter values in Table 11.

Table 10. Comparison of ice classification results using different compression ratios.

Case	MY Conc (%)	FY Conc (%)	I#4 Conc (%)	NI Conc (%)
Full Res	48.31	27.73	17.10	6.86
8:1 Comp	46.90	27.11	18.36	7.62
32:1 Comp	47.03	25.97	15.46	11.55

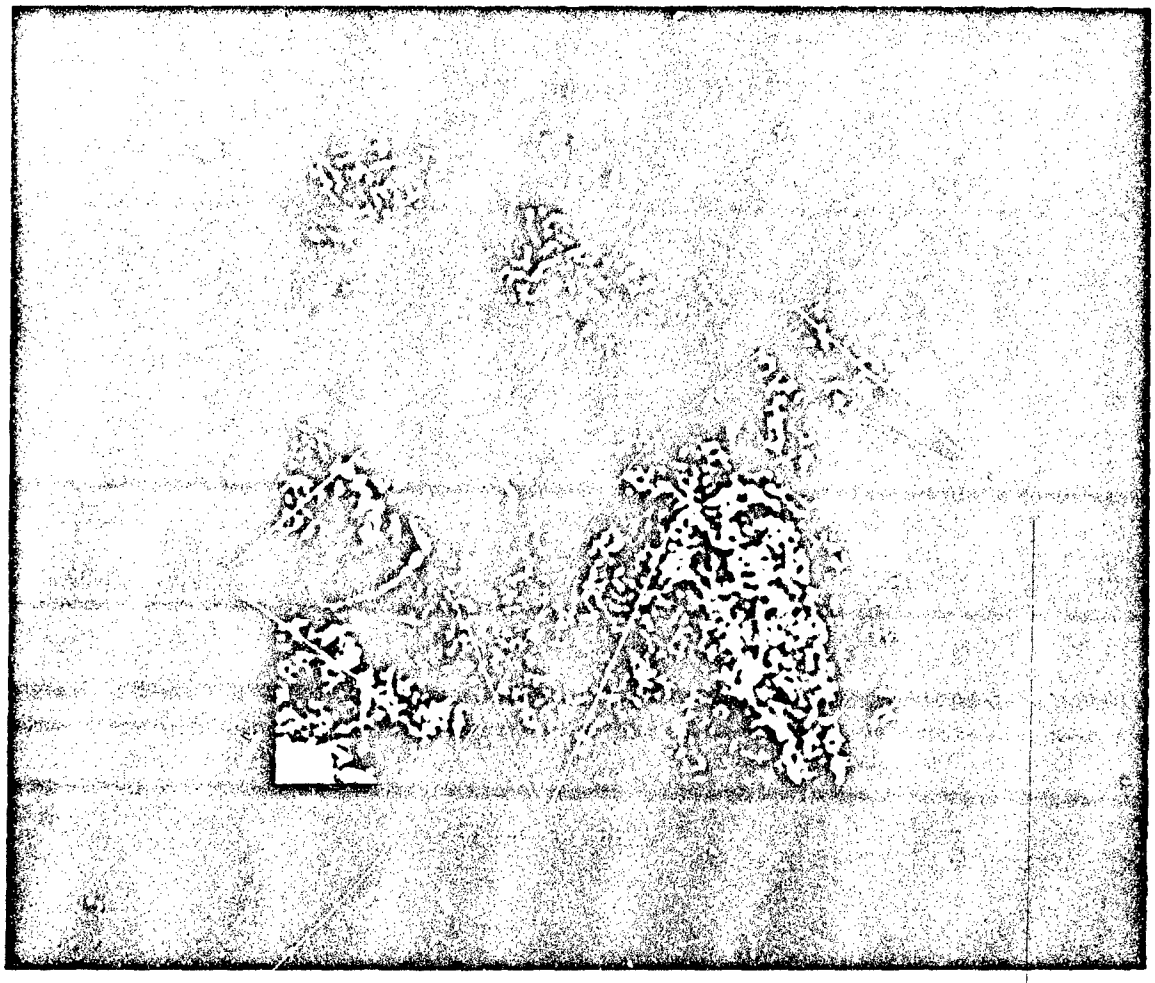

Figure 27. Image compression artifacts.

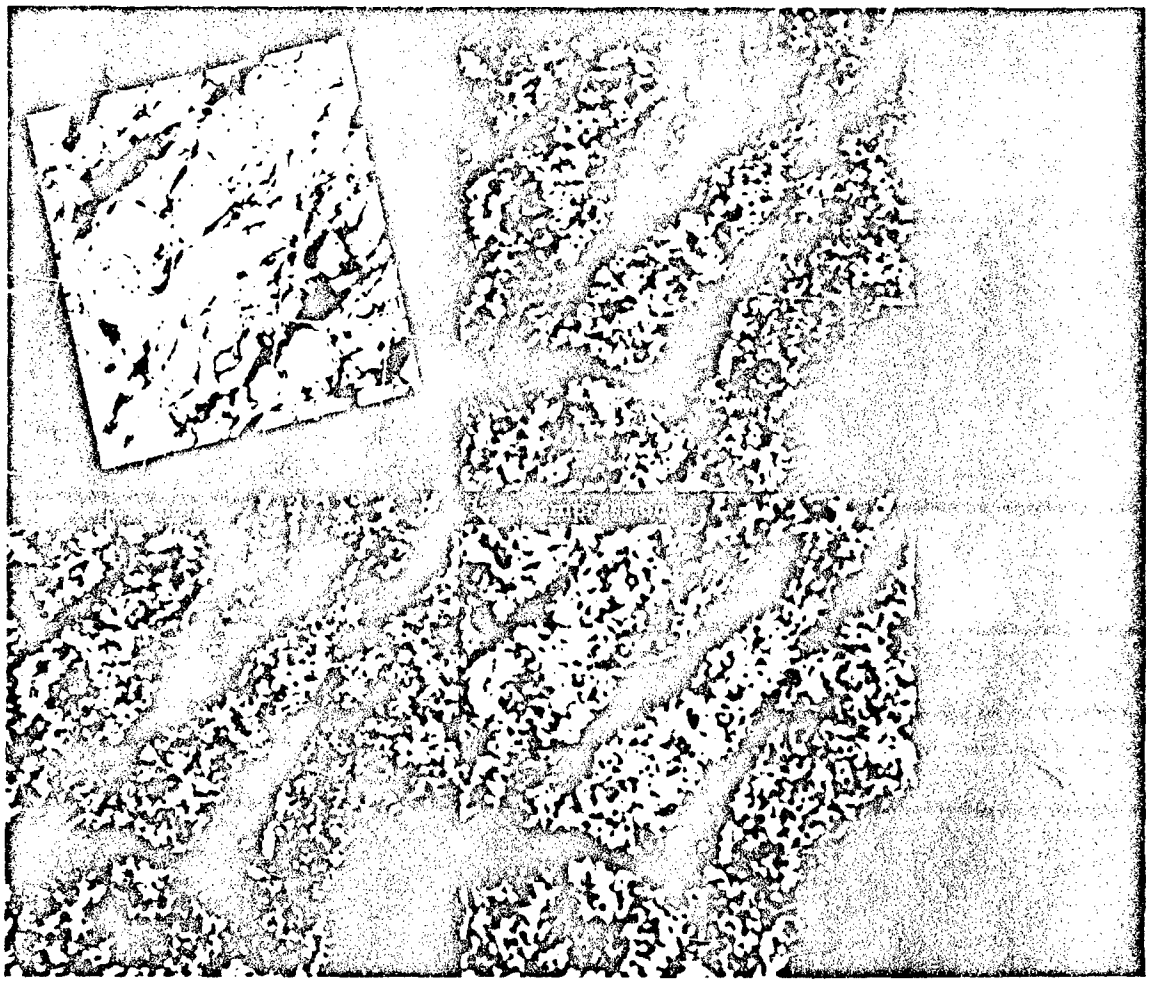

Figure 28. Comparison of 8:1 and 32:1 compression ratio reconstructed images to a full resolution image. The image in the upper left corner is a low resolution image showing the area (in red) covered by the reconstructed and full resolution images.

Table 11. Comparison of mean backscatter values for different compression ratios

Case	MY(dB)		FY(dB)		I#4(dB)		NI(dB)	
	mean	stddev	mean	stddev	mean	stddev	mean	stddev
Full	-9.89	2.22	-14.14	2.41	-18.10	2.66	-21.18	2.49
8:1	-9.94	2.19	-14.08	2.46	-17.92	2.87	-20.60	2.98
32:1	-9.99	2.02	-14.12	2.37	-17.38	3.37	-19.94	3.37

Visually, the more "wormy" the image, the more inconsistent the ice classification (Figure 29). Floes which are visually 100% of one ice type are often classified as being a salt and pepper mixture of two ice types. In addition, compression artifacts can also cause errors in ridge analysis (Figure 30). Ridge analysis will be described fully in section 6.0. The "wormy" nature of the compressed imagery can cause the algorithm to over estimate the number of ridges in an image by misinterpreting the boundaries of the "wormy" features to be ridges as well. Lead analysis is similarly affected. Recommendations for improvements are, therefore, to:

1. develop an image compression algorithm which does not create as many compression artifacts as the algorithm currently in use does. This task may be as simple as adopting the 8:1 compression which was tested during this evaluation, or it could require the development of a new compression scheme. The 8:1 compression is in place on the system, but the compressed images take up four times the space of the 32:1 compressed images. Another compression scheme could save space, but would take considerable time to develop. The results which the current compression scheme provides may be of a high enough quality for operational use and no change may be necessary.

2. make high resolution data fully accessible by all algorithms and workstation functions. Both supervised and unsupervised algorithms should be able to be performed on either low or high resolution data.

6.0 Lead and Ridge Product Evaluation

6.1 Algorithm Descriptions

At present no algorithms to detect leads and ridges in SAR imagery reside within the framework of the SAR/JIC workstation software. However, lead and ridge extraction algorithms are available in the NSIPS image processing software which is resident on both the JIC and NRL SAR workstation systems. The algorithms are based on the work of *Vesesky et al.(1989)*. They are designed to work with SAR imagery of any resolution and up to 512 x 512 pixels in size. For lead analysis, a 5 x 5 pixel smoothing filter is applied to the data and then a threshold is applied to the image. Choosing the proper threshold for the image is accomplished by eye. The resulting image is a binary image, with black representing the leads and white representing non-lead areas. Lead statistics of lead density, lead area, and lead orientation are then calculated from this smoothed binary image. Lead density is calculated to be the percentage of lead pixels per grid cell, where the grid size is defined by the user. A map of lead density for the image is created by overlaying the binary lead image with an image of shaded grid cells, each shade representing the lead density within that grid cell. Lead area is equal to the sum of pixels within a lead feature. This measurement is represented as a histogram which displays the number of occurrences of leads in size bins. Lead orientation is determined from the orientation of the line that is a least squares fit to the lead pixels. This measurement, also represented by a histogram, shows the number of leads which have a particular orientation. Before lead area, orientation, and density are calculated leads of fewer than 6 pixels are removed from the data set.

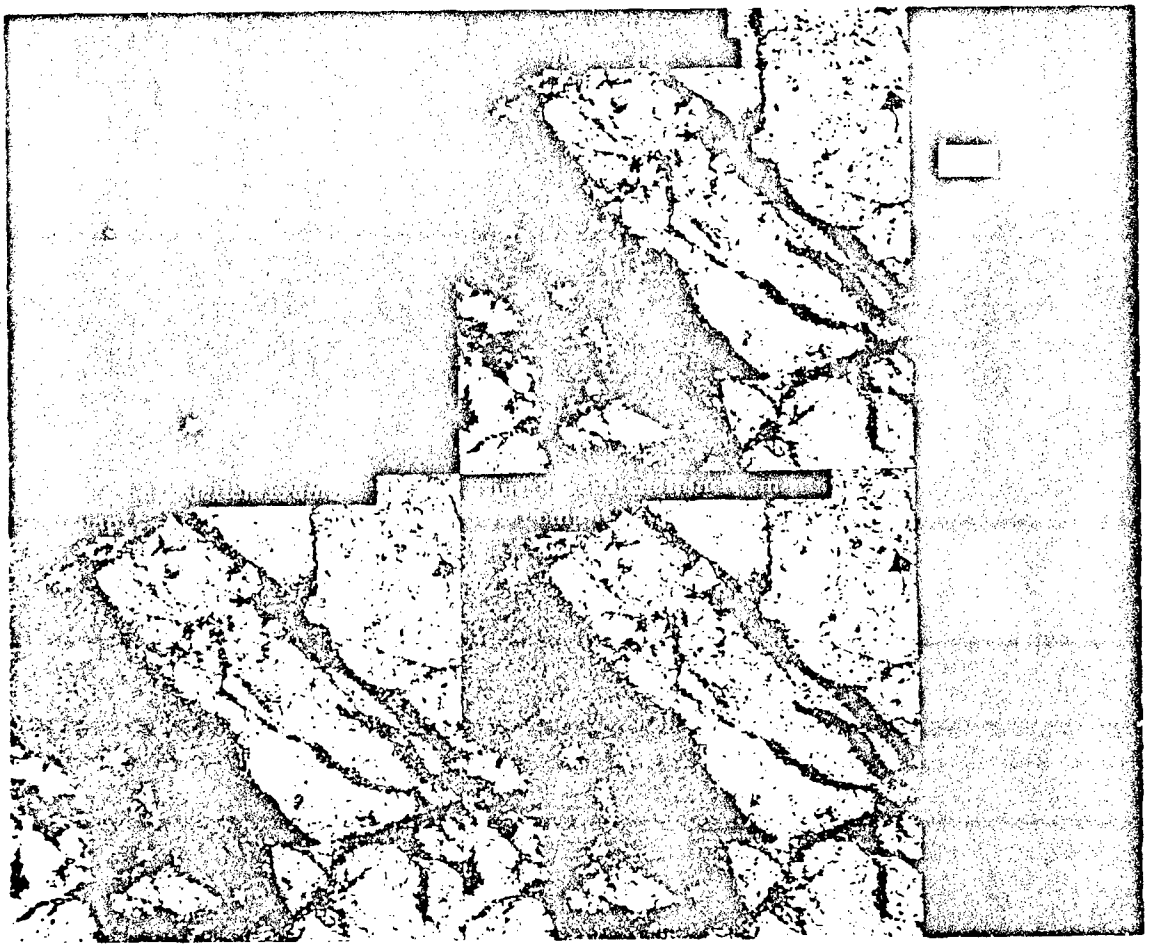

Figure 29. Comparison of classification results for a reconstructed and uncompressed image.

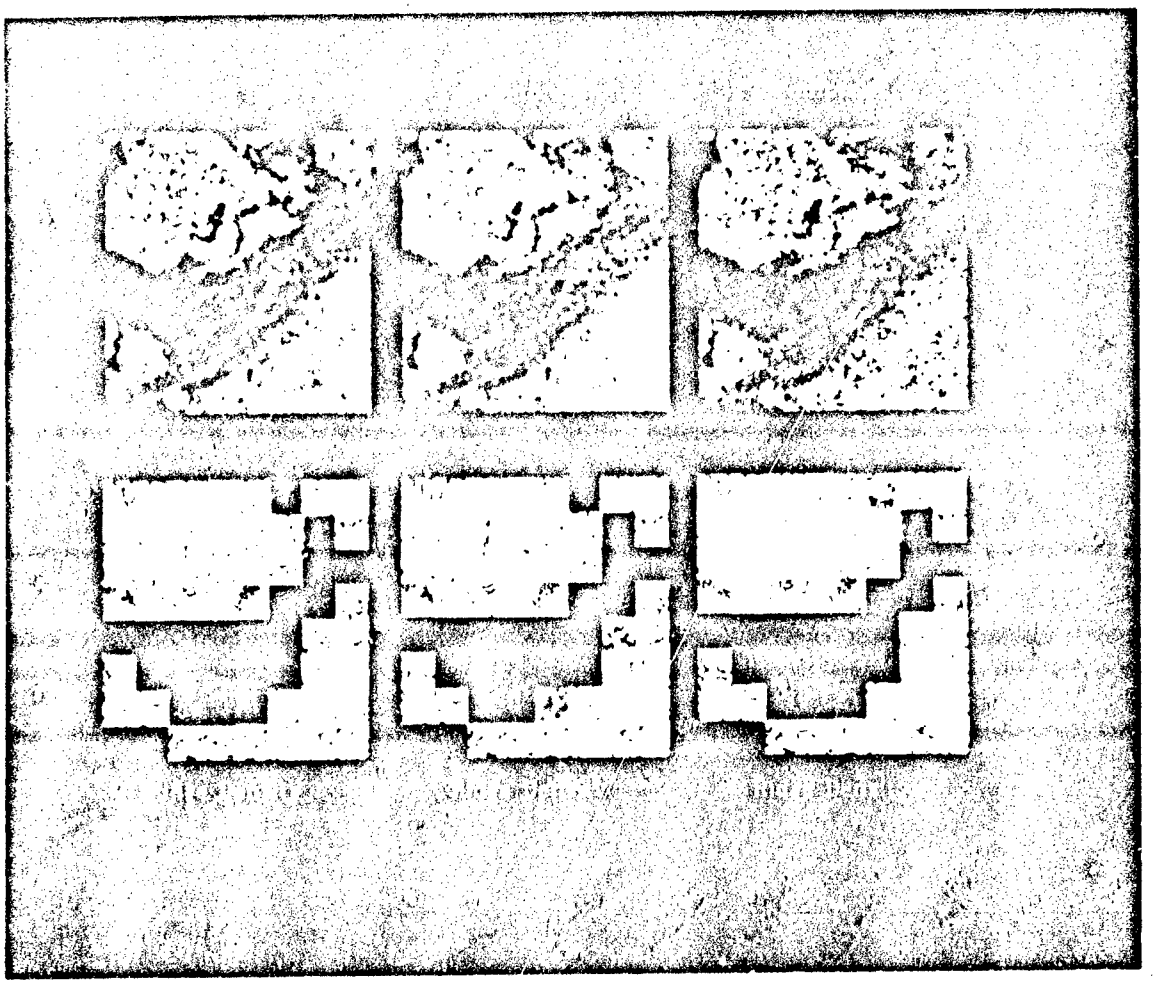

Figure 30. Comparison of ridge analysis results for a reconstructed and uncompressed image.

Ridge analysis is slightly more complicated. Ridges are defined as bright linear features, so the algorithm seeks to filter out all features which do not meet this criteria. It achieves this goal by applying a threshold and a line detector to the image. The threshold isolates the brightest 20% of the pixels. A binary image is created, with values exceeding the threshold as white and values less than the threshold as black. The algorithm then attempts to isolate linear features by applying series of Kasvand masks to the original image. Linear features of less than 6 pixels are removed (but this threshold is adaptable). The bright feature binary image is then combined with the image of linear elements with a logical 'and' operation. The algorithm then computes ridge density measurements from the image of bright linear features. Like lead density, ridge density is defined to be the percentage of bright pixels within a given grid area. This measurement is represented as a binary image overlain with an image of shaded grid areas. For ridge length and orientation, the bright linear features in the image are first thinned to a width of one pixel. Ridge nodes (intersections between ridge features) and ridge ends are then identified, and individual ridges are defined as the line segments connecting ridge nodes. The length of the line segment between two nodes, or a node and an end, represents the total length of the ridge segment, and the orientation of the ridge segment is defined as the orientation of a line that is a least squares fit to the line segment. These ridge lengths and ridge orientations are then tabulated like those in the lead statistics and presented in histogram form.

6.2 Running the Algorithms

The user performs lead and ridge analysis through NSIPS. Using a menu, an image is selected for analysis. The image can be no larger than 512 x 512 pixels in size. The user selects either the lead or ridge algorithm. Unfortunately, ridge analysis is time consuming, and takes up to an hour to run on a 512 x 512 image. Lead analysis is quicker. In both cases the user must supply the algorithm with the size of the grid to be used and the number of contiguous pixels below which a feature is too small to be considered a lead or a ridge. For lead analysis the resulting products are a smoothed binary map of leads, a map of lead density, and histograms of lead area and orientation. For ridge analysis, the resulting products are a binary map of ridges, map of thinned ridges, a map of ridge density, and histograms of ridge length and ridge orientation. The algorithm is simple to use and requires no preparation of the data for analysis. In addition, the algorithm can be performed on either low resolution or high resolution imagery although the results are quite different.

6.3 Lead and Ridge Analysis Results

Examples of algorithm results are presented in Figures 31 and 32. In general, both lead and ridge algorithms identify some features accurately when compared with a visual inspection of the imagery. However, both algorithms also tend to identify more of the image as a lead or ridge than really exists. In the case of the lead algorithm patches of first year ice in the image are mistakenly identified as leads (Figure 31). Although they may have been leads at one time, they are not proper leads, which are only those fractures in the ice which contain new ice or open water. These first year ice areas are identified as leads because the algorithm finds leads using pixel value alone. As some first year ice has low backscatter, it is hard to select a threshold which separates leads from ice. This is particularly true with freezing ice in a lead, which can have a range of backscatter values, or with wind-roughened open water in a lead, which can have backscatter higher than multi-year ice.

The ridge algorithm, as evidenced by Figure 32, misidentifies image speckle as ridges. Although increasing the minimum feature length may remove some of this effect, it does so at the expense of removing true ridges as well. The ridge algorithm also has difficulty locating ridges in first year ice (Figures 30 and 32). Ridges in first year ice are ignored by the algorithm because they are not among the brightest 20% of the image pixels, even though they are clearly apparent in the imagery itself.

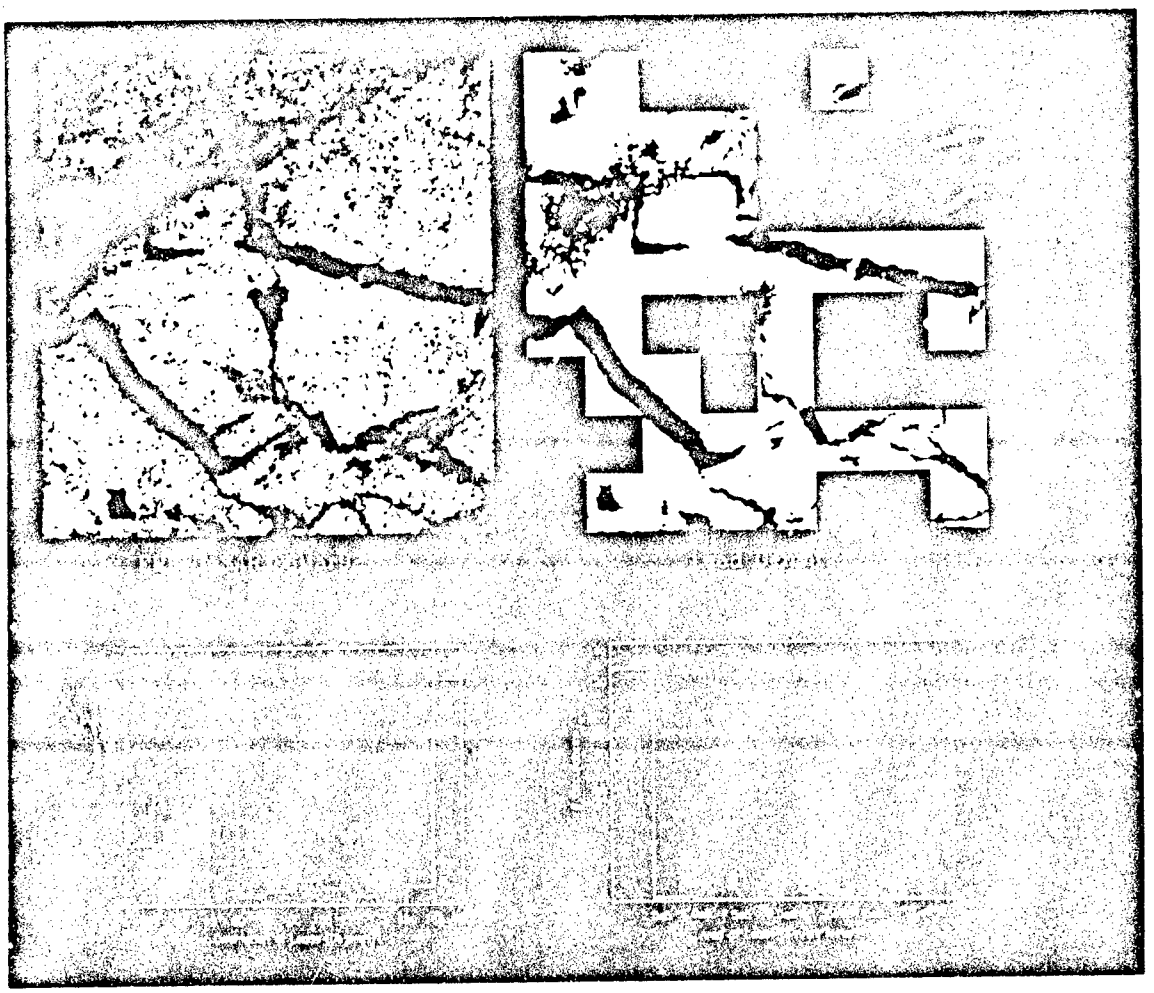

Figure 31. Example of lead analysis results.

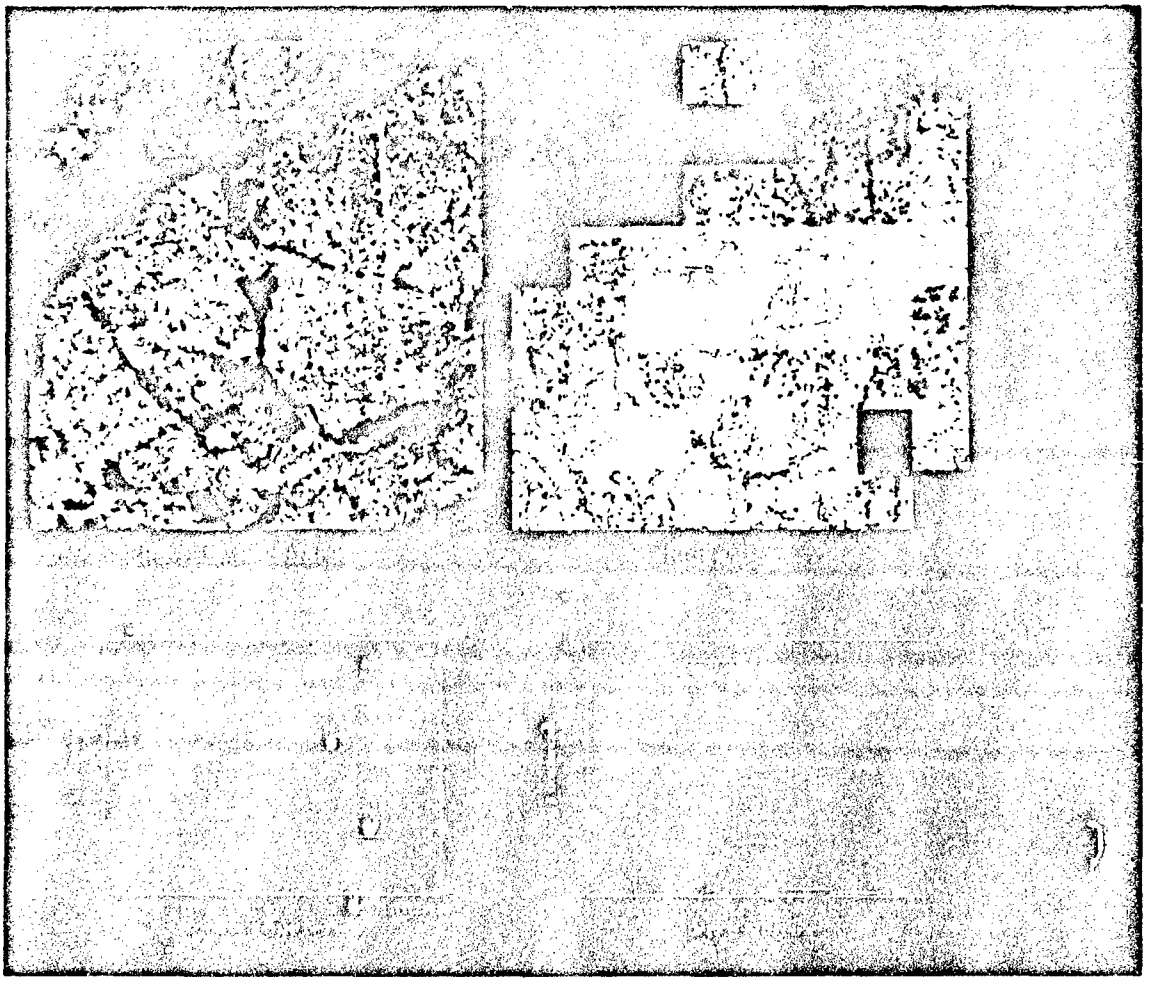

Figure 32. Example of ridge analysis results.

At this time, there are no good sources of lead and ridge data with which these results can be compared. However, extensive ground truth was acquired by ERIM during LEADEX, and will be used in a validation of the algorithm.

6.4 Problems Encountered and Recommendations

Operationally, running the algorithms is very time consuming. The algorithms should be modified to run more efficiently and on images larger then 512 X 512 pixels. Most SAR imagery is larger than this size. The algorithms often fail at both lead and ridge identification. First year ice areas are misidentified as leads. Ridges which are not among the brightest 20% of the pixels, such as those which occur in first year ice areas, are not detected, while image speckle is misidentified as ridges. The ridge algorithm results should be viewed with caution, since many ridges will not be resolved by even the full resolution SAR imagery. Recommendations for improvements to these algorithms are to:

1. modify the algorithms so that they can operate on images much larger than 512 by 512 pixels.

2. improve the techniques used by the algorithm for lead identification. At present, the choice of a user determined threshold does not allow leads to be separated from first year ice. Higher order image statistics, such as standard deviation and texture parameters, may aid in the separation of leads with calm water from other low backscatter regions. Taking advantage of feature shape (i.e. leads are linear) may help as well. Many of the considerations for ice type expressed in section 3.5 apply here as well.

3. improve ridge identification through the use of a sliding window threshold instead of a constant value threshold. This change would enable the algorithm to isolate the brightest 20% of the pixels within an image sub-region, and identify ridges in first year ice which may not necessarily be among the brightest 20% of all image pixels.

4. develop methods to remove erroneous ridges caused by speckle. A spatial filter which compares a ridge feature to its surrounding area may achieve this end, but care must be taken not to remove small ridges as well.

7.0 Conclusions

SAR workstations were installed at the Navy/NOAA Joint Ice Center and at NRL's Remote Sensing Applications Branch at Stennis Space Center near the beginning of 1992. After initial difficulties setting up smooth communications, imagery has been requested of ASF by JIC and received on both systems in a fairly steady stream. In this report we have evaluated the performance of the algorithms designed by JPL and by Stanford University to operate on these SAR data, and have concluded that:

• when compared with supervised classification, the ice classification algorithm gives results which are generally accurate to within 10%, although there are notable exceptions and some consistent errors.

• ice motion vectors from SAR far exceed the accuracy of vectors from other satellite image sources due to better resolution and geolocation accuracy.

• ridge density appears to be vastly overestimated (although ridges found by manual image inspection are correctly identified) and first year ice is identified as lead ice by the lead and ridge algorithm.

We did not attempt to validate the algorithms. Algorithm validation is being carried out by ERIM using measurements acquired during the LEADEX field experiment.

Of the three algorithms, it is the lead and ridge algorithm which is the most problematic. Recommendations for improving all algorithms fall into 2 categories: those for which development costs are low and prospective benefits high, and those for which additional data gathering and developmental research is necessary. These recommendations are described in detail in sections 3.5, 4.4, and 6.4. We will implement the former while looking for opportunities to begin work on the latter through existing and planned programs.

8.0 Acknowledgements

This evaluation would not have been possible without effort on the part of many people to insure that the workstation, its peripherals, and the SARCOM communications link to Alaska ran smoothly as a system. John Schmidt and Conrad Johnson, NRL Remote Sensing Applications Branch, have been instrumental in solving hardware, software, and communication link problems since the initial installation of the workstation in late 1991. Tim Gremillion (Sverdrup Technology) provided creative software tools to meet user requests for new workstation capabilities. CDR Brad Smith, the commanding officer of the Naval Polar Oceanography Center, supported the use of the workstation at the Joint Ice Center and clarified the role of satellite SAR in an operational environment. We depended upon Gary Wohl, Joint Ice Center Senior Ice Scientist, and Cheryl Bertoia, Joint Ice Center Operations, to highlight concerns of the operational user, to collaborate on data acquisition, and for their knowledge of sea ice. As project manager at JPL for the SAR workstation contract and for the Alaska SAR Facility's Geophysical Processing System, Ron Kwok enabled an instrument for polar research to serve the operational community as well. This success was the result of careful oversight of the workstation implementation on the part of JPL and subcontractor Vexcel Corporation. Tom Baltzer (Vexcel), Joanne Shimada (JPL) and Dean Myerson (Vexcel) provided excellent documentation and answered numerous requests concerning workstation software with solutions or recommendations. Finally, Don Montgomery, Office of the Oceanographer of the Navy, and CDR Pete Ranelli, SPAWAR, provided continued interest, guidance, and support. This work was sponsored by the Chief of Naval Operations under program element 603704N, CDR P. Ranelli, program manager.

9.0 References

Baltzer T., T. Gremillion, R. Kwok, D. Myerson, J. Shimada (1992). *Joint Ice Center SAR Workstation User's Guide*, JPL Report D-8861

Emery, W., C. Fowler, J. Hawkins, and R. Preller (1991). Fram Strait Satellite Image-Derived Ice Motions, *Journal of Geophysical Research*, Vol 96, No C3, 4751-4768.

Fetterer F. and C. Bertoia (1992). *JIC SAR Transition and Implementation Plan*, NRL memo

Fetterer F., W. Pichel, D. Montgomery, G. Wohl, C. Bertoia, R. Shuchman, and R. Kwok (1992). *Navy/NOAA SAR Oceans and Ice Applications Plan*, NRL Technical Note 226

Hollinger J., ed., (1991). *DMSP Special Sensor Microwave/Imager Calibration/Validation*, NRL Final Report , Vol II

Kwok, R., E. Rignot, B. Holt, and R. Onstott (1992). Identification of Sea Ice Types in Spaceborne Synthetic Aperture Radar Imagery. *Journal of Geophysical Research* , Vol 7, No C2, 2391-2402.

Kwok, R., J. Urlander, R. McConnell, and S. Pang (1990). An Ice Motion Tracking System at the Alaska SAR Facility, *IEEE Journal of Oceanic Engineering*, Vol 15, No 1, 44-54.

Preller, R. and P. Posey (1989). *The Polar Ice Prediction System - A Sea Ice Forecasting System*, Naval Ocean Research and Development Activity, Stennis Space Center, MS, NORDA Rpt. 212, 42pp.

Shuchman, R. and R. Onstott (1992). *Characterization of ERS-1 SAR Snow and Sea Ice Signatures During LEADEX*, ERIM Interim Report

Vesesky J., M. Smith, and R. Samadani (1990). Extraction of Lead and Ridge Characteristics from SAR Images of Sea Ice, *IEEE Transactions of Geoscience and Remote Sensing*, GE-28, No 4, 740-745.