

AUTOMATED TRACKING OF ARCTIC ICE FLOES IN MULTITEMPORAL SAR IMAGERY

R. McConnell, W. Kober, F. Leberl
Vexcel Corporation
2477 55th Street
Boulder, CO 80301, USA
Tel: 303/444-0094
FAX: 303/444-0470

R. Kwok, J. Curlander
Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109

ABSTRACT

We discuss the use of two algorithms to perform shape matching on the boundaries of ice floes in SAR images in order to produce an ice motion map. The algorithms match a shape descriptor known as the psi-s curve. The first algorithm uses normalized correlation to match the psi-s curves, while the second uses dynamic programming to compute an elastic match that better accommodates deformation of the ice floe boundary.

As a result, there has been increasing attention to algorithms for tracking ice floes by matching feature shapes (Vesecky, 1988).

We describe here a set of algorithms developed at jointly by Vexcel Corporation and the Jet Propulsion Laboratory (JPL) which have successfully overcome these obstacles. Rather than correlating raw pixel values, these algorithms first extract features from the images. The shapes of the extracted features are then compared using shape descriptors known as psi-s curves. To match the psi-s curves, the algorithms use both normalized correlation and dynamic programming. An ice tracking system that uses these algorithms is currently being developed jointly by Vexcel and JPL under a grant from NASA. The system, known as the Geophysical Processing System (GPS), is scheduled for installation at the Alaska SAR facility in Fairbanks in April of 1990.

I. Introduction

There has been considerable interest in recent years in the use of SAR imagery to study movement of arctic ice floes. The movement of ice floes is of interest to shipping, oil drilling, and military operations. In addition, the motion of ice floes plays a large role in the world's weather, because this motion exposes large expanses of unfrozen ocean water to the much more frigid arctic air, and is therefore responsible for a large amount of heat transfer between the ocean and the atmosphere. SAR imagery is suited for this task because it allows continuous coverage through clouds that prevail in the arctic, as well as during the dark winter months.

II. Psi-s curves

Suppose f(s) = (x(s),y(s)) is a continuously differentiable parametric curve such that s is arc length along the curve. At each point, the vector f'(s) = (x'(s),y'(s)) gives the vector that is tangent to the curve.

A common approach to the automated tracking of arctic ice has been to select a patch from an early image (the source image) and to cross correlate it with a later image (the target image) at each position that could plausibly correspond to the same patch of ice (Fily, 1986). The position that maximizes the computed correlation coefficient is deemed likely to contain the corresponding patch of ice. Sometimes consistency checks between several matches are used to identify false matches.

Let theta(s) be equal to the angular difference between f'(s) and the unit vector (1,0). theta is a real-valued, function of s which is continuous everywhere, except where phase wrapping introduces discontinuities, i.e., where it wraps around from 2pi to 0 or from 0 to 2pi. The psi-s curve is derived by adding or subtracting multiples of 2pi to portions of the theta-s curve as needed to remove these discontinuities. In the case where the curve has points where it is not differentiable, multiples of 2pi are added or subtracted in order to minimize the magnitudes of the discontinuities.

This method is known as area correlation. The problem with it is its computational expense, especially when ice floes rotate. To accommodate rotation, the patch must be rotated and correlated several times at each potential match position in the target image. The search space thus becomes very large, increasing the likelihood of false matches and increasing enormously the computational burden.

The boundary of a binary region in a digital image has a limited number of orientations, depending on the tessellation. It is therefore necessary to interpolate a smooth boundary through the jagged artifacts of digitization before computing the psi-s curve.

III. Matching of psi-s curves using correlation

One of the methods GPS ice tracking system uses for matching shapes is normalized correlation of the psi-s curves of the shapes. The GPS ice tracking system employs correlation by extracting arbitrary segments of fixed size from the psi-s curves from the source image and finding the subsegment of the same size in the psi-s curves from the target image that maximizes the correlation coefficient.

Correlating psi-s curves is an effective method of matching features that rotate strongly from one image to another. The reason is that rotation of a feature causes a constant to be added to its psi-s curve. The correlation coefficient between two functions is invariant to addition of a constant to one of the functions. Therefore, the ability of correlation to identify a match is unhampered by rotation of the features. The rotation can be estimated at the best match by linear regression of the psi values from the two sets.

Correlation is also invariant to scalar multiplication of one of the sets of samples being correlated. This occurs when one of the psi-s functions can be derived from the other by scalar multiplication with a factor other than 1.0. This would mean that the variation in the orientations of the tangents to one feature is greater than it is for the other, and hence, one of the features is a coiled up version of the other. This does not correspond to what humans would consider resemblance, nor are two such curves likely to represent the same ice floe in two images. The scalar multiple can be estimated the linear regression slope coefficient, and matches rejected where this estimate is significantly different from 1.0.

The final task is to determine which of the matches obtained by correlation are false and which are correct. The solution to this problem used in the GPS starts with the observation that most ice floes have at least two correct matches on them. The features involved in these matches move as part of a rigid body. The length of the axis joining them will be unchanged from one image to the next, and the estimates of the rotations of each obtained by their linear regression intercepts will coincide with each other and with the rotation of the axis joining the two features. It is unlikely that a pair of matches that are not both correct will satisfy these criteria. The GPS examines all pairs of matches obtained from psi-s correlation, tests for these criteria, and keeps all pairs of matches that satisfy them. This is a much more effective way of separating false matches from correct matches than is thresholding the correlation coefficients of the matches.

Figures 1, 2, 3, and 4 depict an ice image pair, the extracted features from each pair, and the motion vectors derived using psi-s correlation followed by the bad match filter.

IV. Problems with matching psi-s curves using correlation

Figure 5 illustrates a pairing between corresponding elements of two psi-s curves derived from segmented ice images. It is clear that the mapping of the elements is not one-to-one, and the mapping of the elements cannot be known until the curves are matched up.

In fact, implicit in correlation is the assumption that the psi values are disturbed from one image to the next, but that the mapping between the s domains remains linear. However, any disturbance of the psi values usually results from disturbance of the shape being matched, and therefore also results in distortion of arc length over some intervals of the boundary. The correlation matching process is therefore misspecified, and this can lead to failure of correlation matching to find correct matches when distortion of arc length is severe.

V. Dynamic Programming

The field of sequence comparison deals with the comparison of similar sequences, where the correspondence between the elements is not one-to-one and not known in advance. The techniques developed in this field provide a way of measuring the similarity between such sequences, as well as computing the optimal and most natural mapping between the elements of one sequence and those of the other. Most of the techniques in this field are based on a class of algorithms known as "dynamic programming". Dynamic programming is a process whereby a recursive problem with an exponential search tree can be solved in polynomial time by using a table to retain intermediate results that are shared by different branches of the tree. A survey of this area can be found in (Sankoff, 1983).

Dynamic programming has been used in the past for shape comparison in computer vision, but this has been restricted largely to handwriting analysis (Burr, 1983).

A. A Dynamic Programming Solution to the Matching of Psi-s Curves

A procedure known as "dynamic time warping" is a variant of dynamic programming for matching real-valued sequences. Dynamic time warping examines all sets of mappings between elements of one sequence and those of the other, subject to the constraints that the mappings do not cross each other, and that every element from one sequence is paired with at least one element from the other sequence. It assigns to each of these mappings the sum of absolute differences of paired elements. It then produces the mapping with a minimum sum of absolute differences, or "cost."

This can be computed recursively as follows:

Let x_n and y_n be two sequences of length m and n . Then x_i and y_j are the prefixes of these sequences containing i and j elements, respectively, and x_i and y_j are the i 'th and j 'th elements of the respective sequences.

The distance between the two sequences can be expressed as follows:

$$d(x_m, y_1) = \epsilon |x_1 - y_1|$$

$$d(x_1, y_n) = \epsilon |x_1 - y_n|$$

$$d(x_m, y_n) = \min \left(\begin{array}{l} d(x_{m-1}, y_n) - |x_m - y_n| \\ d(x_{m-1}, y_{n-1}) - |x_m - y_n| \\ d(x_m, y_{n-1}) - |x_m - y_n| \end{array} \right)$$

The relation maps x_m to y_n at cost $|x_m - y_n|$, and resorts to recursion to compute the minimum-cost mappings of the preceding elements.

The mapping can be computed cheaply by keeping a table, where element (i, j) of the table contains $d(x_i, y_j)$. By starting with element $(1, 1)$ the table can be built up inductively without resorting to recursion. When the table is completed, element (x_m, y_n) contains the cost of matching the two sequences. Backtracking in the array from element (x_m, y_n) through the elements giving rise to the minimum at each step of the recurrence relation gives the mapping between the elements. The time and space required to complete the table is proportional to its size, or $m \times n$.

The GPS uses a variant of this approach. The following recurrence relation is used:

$$d(x_m, y_1) = \epsilon |x_1 - y_1|$$

$$d(x_1, y_n) = \epsilon |x_1 - y_n|$$

$$d(x_m, y_n) = \min \left(\begin{array}{l} d(x_{m-1}, y_n) - r * |x_m - y_n| \\ d(x_{m-1}, y_{n-1}) - |x_m - y_n| \\ d(x_m, y_{n-1}) - r * |x_m - y_n| \end{array} \right)$$

where ϵ is a parameter that penalizes excessive arc-length warping.

After the table is completed, the bottom row is searched for a minimum value. Backtracking from this location gives the interval over x_m that matches y_n the best. When y_n is a fragment of a psi-s curve from one image, and x_m is the set of psi-s curves from the other, this procedure finds the best match.

Figure 6 illustrates the use of the table on a small example, while Figures 7 and 8 show dynamic programming match results on ice image data.

B. Problems with the Dynamic Time Warping Approach

As explained above, the primary advantage of the dynamic time warping approach over the correlation approach is that it is not blinded to a match by distortion of the s dimension of the psi-s curve. The chief disadvantage of the approach is that the underlying similarity measure is based on the sum of differences. Unlike correlation, the sum of differences measure of similarity is sensitive to the addition of a constant to

one of the sets of samples. In psi-s matching, this is exactly what happens when there is rotation of the features from one image to the next. Therefore, although the approach is not blinded by distortion in the s dimension of the psi-s curves, it can be blinded by strong rotation of the features to be matched.

To remedy this problem, a rotation-invariant descriptor of a curve can be derived from the psi-s curves by subtracting a running mean from them before the dynamic programming match is performed.

VI. Conclusion

Two methods of tracking arctic ice floes SAR images have been presented. Both of the methods do shape matching on psi-s curves. The first method, based on normalized correlation, is indifferent to rotation of the ice floes, and therefore excels in matching features whose rotation is unknown.

The second method, based on dynamic time warping, excels in matching features that resemble each other less closely, but whose rotation can be estimated.

REFERENCES

1. Burr, D.J., "Designing a Handwriting Reader", IEEE Trans. Patt. Anal. Mach. Intelligence, Vol. PAMI-5, No. 5, pp554-559, 1983.
2. Fily, M., and D. A. Rothrock, "Extracting Sea Ice Data from Satellite SAR Imagery", IEEE Trans. on Geoscience and Remote Sensing, Vol. GE-24, No. 6, pp849-854, 1986.
3. Sankoff, D., and J. B. Kruskal, eds., Time Warps, and Macromolecules: the Theory and Practice of Sequence Comparison. Addison Wesley: 1983.
4. Vesecky, J.F., et. al., "Observation of Sea-Ice Dynamics Using Synthetic Aperture Radar Images: Automated Analysis", IEEE Transactions on Geoscience and Remote Sensing, Vol. 26, No. 1, pp38-47, 1987.



Figure 1. October 1978 Seasat SAR image of an area of central pack ice in the Beaufort Sea.




Figure 3. Features extracted automatically from the image in Figure 1. To extract the features, the gray values are clustered to get a binary classification. The boundaries of the regions in the resulting binary image are vectorized, and all boundaries with less than a minimum perimeter are discarded.




Figure 2. Seasat SAR image of the same ice three days later.




Figure 4. Correct matches obtained by correlation of the psi-s curves of features extracted from the images in Figures 1 and 2.




Figure 5. Two matching segments of sea ice feature boundaries and their psi-s curves. The matches of the psi values were produced using the GPS's dynamic programming algorithm. Deformation of the axis along which the curves match frequently inhibits a correct match when correlation of psi-s curves is used.




Figure 7. Matches produced by dynamic programming from the data used in Figure 4.

	1	2	1	2	3	3	2	5	0	1	1
2	1	0	1	0	1	1	0	3	2	1	1
3	3	1	1	1	0	0	1	2	5	3	3
2	4	1	2	1	1	1	0	3	4	4	4
4	7	3	4	3	2	2	2	1	5	7	7
4	10	5	6	6	3	3	4	2	5	8	10
0	11	7	6	8	6	6	5	7	2	3	4

Implied mapping:

1	2	1	2	3	3	2	5	0	1	1
	2	3	2	4	4	0				

Figure 6. Illustration of the table produced by a dynamic programming match of the sequence 232440 against the sequence 12123325011, which produces the best-match subsequence 233250. The table is filled in row-by-row using the recurrence relation used by the GPS and a value of 1.0 for r . When the table is completed, the bottom row is scanned for a minimum value. Scanning vertically from the minimum value of 2, one finds the location of the end of the best-match subsequence. By backtracking from that location through the elements giving rise to the minimum in the recurrence relation, one finds the beginning of the best-match subsequence, as well as the mapping between the elements.