Glacier Surface Motion Computation from Digital Image Sequences

Adrian N. Evans

***Abstract*— A technique for computing the field of short-term glacier surface motion on a local scale is presented. Time-lapsed image negatives, digitized to a high resolution, provide the raw data for the three-stage technique. First, cross-correlation is used to establish a number of candidate displacement vectors for a series of regularly spaced templates. A relaxation-labeling routine is then applied to select the most appropriate candidate vectors, according to the local flow. Novel aspects of the relaxation algorithm include a new, efficient form of the support function and the absence of a null-match category. A new development is the application of a post filter to the resultant flow field, providing suitable displacement vectors for templates that were originally unmatched and correcting vectors that are still inconsistent with the local flow. Results from an image sequence from New Zealand's Mount Cook National Park show the superiority of the technique over the maximum cross-correlation method and demonstrate the effectiveness of the post filter in improving correlation-relaxation labeling.**

***Index Terms*—Correlation-relaxation labeling, glacier surface motion, motion estimation, remote sensing.**

I. Introduction

DETERMINING glacier surface movement is of scientific value both intrinsically and as part of integrated investigations into the relationship between the hydraulic conditions at the glacier bed and glacier motion. To estimate glacier motion, direct distance measurements of markers sunk into the ice can be taken from the bedrock close to the ice margin [4]. When the ice margin is inaccessible by terrestrial survey using electromagnetic distance measurement (EDM) or theodolite, sites overlooking the glacier can be used in conjunction with targets placed or air-dropped onto the ice [5], [6].

Alternatively, remote sensing has been used to track surface features in the ice itself, using Landsat visible spectrum [2], [13] or SPOT images [8], [14]. Rignot *et al.* [10] have applied synthetic aperture radar (SAR) interferometry techniques to analyze the dynamics of the Greenland ice sheet using ERS-1 SAR data and to differentiate the ice velocity of a floating glacier from the underlying tidal motion [9]. A major advantage of remotely sensed data is that it is in a suitable form for automated analysis. However, while the motion estimates produced have a high accuracy, the pixel size of 10–30 m limits the size of features that can be matched.

Manuscript received June 15, 1998; revised February 25, 1999. This work was supported in part by a grant from Massey University, Palmerston North, New Zealand, Research Fund.

The author is with the Department of Electronic and Electrical Engineering, University of Bath, Claverton Down, Bath, BA2 7AY U.K. (e-mail: a.n.evans@bath.ac.uk).

Publisher Item Identifier S 0196-2892(00)02490-6.

Consequently, although the spatial resolution is sufficient for large scale and regional analysis, it is too coarse for local scale measurements, where the size of the glacier under study is in the order of meters rather than kilometers, and the spacing between velocity estimates may only be a few meters. Furthermore, the use of remote sensing methods is restricted to those geographic regions that are well covered by orbiting satellites.

Therefore, it would be advantageous to combine the spatial resolution achieved by local surveying methods with the ease of automated interpretation and ability to calculate motion at numerous points of remote sensing methods. An approach of this type has the potential of achieving a dense, local motion field for glaciers of any spatial scale and is widely applicable, regardless of the availability of remotely sensed data. This is the subject of this paper.

A new technique is presented that uses digital image processing algorithms to rapidly and automatically compute complete motion estimates from a time-lapsed sequence of digital images of a glacier. The technique can be described as follows. First, a series of images taken from a fixed camera site overlooking the glacier are digitized to a high resolution, providing data in a suitable form for automated motion analysis. Cross-correlation is then used to establish multiple match positions, termed "candidate matches," for each of a series of templates regularly spaced on the vertices of a rectangular grid. A computationally efficient probabilistic relaxation routine selects those candidate matches that provide the best fit with the local motion field. This acts as a smoothness constraint and produces a motion field that is regularized to some degree. Finally, a post filter is used to provide a velocity vector that is indicative of the local flow for any templates for which no candidate matches were initially found or whose resultant velocity vectors are inconsistent with the local flow. This paper does not attempt to relate the motion vectors found to absolute physical distances, a problem previously considered [5], [6], but instead describes the image processing stages required to generate a dense motion field in the pixel domain.

The arrangement of the paper is as follows. Section II describes the process used to acquire, digitize, and register the glacier images. The technique for generating candidate matches for each template is discussed in Section III. Section IV presents the probabilistic relaxation routine, and in Section V, the post-filtering technique is described. Results and discussion are presented in Section VI, and conclusions are given in Section VII.

II. Image Acquisition and Alignment

The first step in the image acquisition process is to obtain a sequence of photographs from a fixed camera sight overlooking

0196–2892/00$10.00 © 2000 IEEE

the glacier. A Leica camera was used with Kodak Technical Pan 100×125 mm film. The camera was fastened to a kinematic mount, allowing its positioning on the tripod to be accurately repeated. To minimize the effect of variation caused by shadow, the photographs were taken at the same time each day.

Digital images were formed by digitizing the photographic negatives. A Kodak photo-CD Pro system was used, giving a resolution of 40 pixels per mm. Variations in the position of the negative within the camera and during the digitization process resulted in digital images that were not perfectly aligned. To overcome this problem, fixed image points can be used to find the transform parameters between images. The fixed points corresponded to features such as rock faces on valley sides and were aligned using the affine transformation

$$[x, y, 1] = [u, v, 1] \begin{bmatrix} a_{11} & a_{12} & 0 \\ a_{21} & a_{22} & 0 \\ a_{31} & a_{32} & 1 \end{bmatrix}. \tag{1}$$

The affine transform relates points (x, y) from the first of an image pair to corresponding points (u, v) in the second image via the six parameters a_{11} to a_{32}, which quantify the amount of interframe rotation, magnification, scale, and shear. These parameters can be inferred by using the correspondences between pairs of points and simple matrix inversion [16]. A minimum of three points and their matches are required, but usually an overdetermined set of equations is used to produce more accurate results. In addition, a further nonlinear stage for reducing the image registration error was used, as described below.

Each fixed region in the subject image was divided into a number of nonoverlapping templates, typically sized 8×8 pixels. Matches for these templates were then found using the cross-correlation coefficient (CCC) ρ, given by (2), shown at the bottom of the page, where $f(x, y)$ is an $N \times N$ template in the original image, and $g(x+u, y+v)$ is the template shifted by (u, v) in the time-lapsed image. The values $\overline{f}$ and $\overline{g}$ are the corresponding regional means, and the value of ρ ranges from -1 (no match) to 1 (perfect match). The fixed regions consisted of exposed rock faces with intensity profiles that produced good matches for the vast majority of the templates.

The image capture and digitization stages produces images whose interframe transformation is predominantly translational, with small amounts of rotation and scale. As a result, each fixed region has approximately the same horizontal and vertical displacement vectors. Using the matches found, the median horizontal and vertical displacements were then calculated for each fixed region, and any points whose displacements were excessively far from the median were removed from the affine parameter calculation.

III. Generation of Candidate Matches

Using a pair of images from the sequence, template matching can be applied to establish correspondences between positions in each image. In this technique, templates from the first image are matched against all positions within a search area in the second image, giving rise to a two-dimensional (2-D) correlation surface. The goodness of match can be determined by using either a similarity or a dissimilarity measure.

The first step is to determine points in the first image upon which the templates to be matched are centered. To do this, two approaches have been proposed. When the images are low contrast, local operators such as the Moravec operator can be applied to detect prominent feature points that are likely to produce good matches [1]. These methods produce few feature points in homogeneous regions, which is advantageous, for example, when tracking the evolution of clouds [17], as featureless regions typically correspond to the absence of cloud cover. Alternatively, the templates can be centered on the vertices of a regular grid [13].

The aim of this study is to present a detailed picture of the glacier surface velocity. Therefore, achieving a full and regular coverage of the glacier surface is the most important factor. Furthermore, bare and debris-covered glacial ice produces images with features at both extremes of the intensity range, making the application of global thresholds to select feature points inappropriate. For these reasons, the template centers were chosen to lie on an evenly spaced grid. Although this did result in some templates that were too featureless to be successfully matched, the majority of the feature points were matchable, and a post filter is applied at a later stage to provide suitable motion vectors for those templates without any matches.

Candidate match positions for each template can be found by gauging the goodness of match for every position within a search area of radius R_{MAX} using the CCC (2). The size of R_{MAX} directly determines the number of potential match positions that must be evaluated. Other more computationally efficient match measures have been proposed, such as the sum of absolute value of differences (SAVD) or the sum of square differences (SSD). Wu *et al.* [18] investigated the relative advantages of these three measures in the context of a correlation-relaxation labeling technique, concluding that in the presence of high level distortion, the high quality of the initial candidates produced by the CCC required fewer iterations and produced an overall better quality result than that produced using candidates generated by the SAVD or the SSD. In these cases, the extra computational effort required by the CCC is more than compensated for in the latter stages of the technique and hence,

$$\rho = \frac{\sum_{x,y} \left((f(x, y) - \overline{f}) \times ((g(x+u, y+v) - \overline{g}(x+u, y+v)) \right)}{\left(\sum_{x,y} \left((f(x, y) - \overline{f} \right)^2 \right)^{1/2} \times \left(\sum_{x,y} \left((g(x+u, y+v) - \overline{g}(x+u, y+v) \right)^2 \right)^{1/2}}. \tag{2}$$

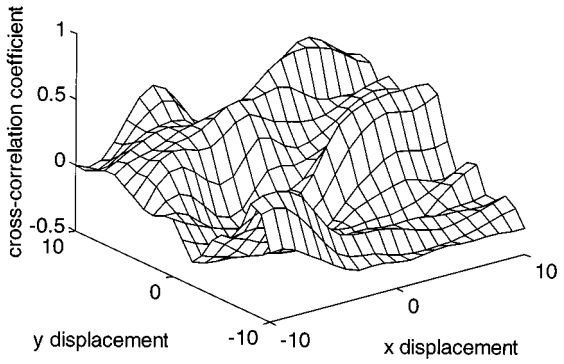

Fig. 1. Typical correlation surface for the image pair used in Section VI, which is multimodal and suffers from indistinct peaks.

the CCC is used here. However, it should be noted that the SSD is adequate when less distortion is present.

Selecting the position with the highest CCC value and calculating the corresponding displacement vector gives rise to the well-known maximum cross-correlation (MCC) method. The assumption behind this approach is that the correlation surface has a single sharp peak. This is typically not the case for the glacier image sequences of this application where nonrigid motion, ablation, poor contrast, atmospheric disturbances, and errors in the matching process result in a correlation surface that is often multimodal and has indistinct peaks. Fig. 1, a typical correlation surface for the glacier images presented in Section VI, illustrates this point. The result is that the best match position does not necessarily occur at the peak of the correlation surface.

To overcome this problem, rather than considering only the highest value in the correlation surface, a number of candidates can be established by ranking the cross-correlation coefficients and selecting the N highest, where N is the number of candidates required [17]. Only those positions that have a CCC value above a predetermined threshold qualify as potential candidates. For example, if the threshold for Fig. 1 is 0.5, only those positions with a CCC exceeding this value qualify as potential candidates. Consequently, the actual number of candidate matches for each template lies in the range of 0 to N, depending on the number of significant matches found. This process eliminates the vast majority of the possible match positions within the search window, leaving a small number of potential matches.

Displacement vectors are generated for each template from its set of candidate match positions. To select the most appropriate of the candidate matches for each template, a smoothness constraint can be applied to the displacement vectors. This is achieved by using probabilistic relaxation.

IV. Refinement of Matches Using Probabilistic Relaxation

Probabilistic relaxation is an iterative technique that uses contextual information to refine label probabilities. In this respect, the templates in the reference image are nodes whose possible labels are the velocity vectors derived from the candidate matches. The number of possible vectors for each template lies between 0 and N, the maximum number of candidate matches allowed. The relaxation process uses the initial vector probabilities and velocity vectors in the local neighborhood to select the most appropriate vector for each template. By this process, the velocity field is smoothed, and many of the problems posed by correlation surfaces that are multimodal or have indistinct peaks are resolved.

The first step of the relaxation process is to assign initial probabilities to the velocity vectors for each template. The initial probability that template J has a velocity vector j denoted $P^{(0)}(J \rightarrow j)$, can be approximated by normalizing the correlation coefficients of the candidate vectors such that

$$P^{(0)}(J \rightarrow j) = \frac{\rho(J \rightarrow j)}{\sum_{\lambda \in \Omega_{2J}} \rho(J \rightarrow \lambda)} \tag{3}$$

where $\rho(J \rightarrow j)$ is the CCC value for template J and candidate match j, and Ω_{2J} is the set of velocity vectors for the template.

Previous approaches [3], [17] have accommodated the possibility that a template does not have an appropriate vector by augmenting each set of velocity vectors with a no-match label that is updated in the relaxation routine. Here, the no-match possibility is not considered, and those templates without any candidate matches play no part in the relaxation process. Instead, a post filter is used to provide appropriate displacement vectors for unmatched and inconsistent templates (see Section V). The absence of the no-match class allows the candidate vectors to evolve without the danger of a mass migration toward the no-match category, which can occur in extant techniques if the free parameters are incorrectly set.

The probability values of the set of velocity vectors for each template are iteratively updated using the nonlinear relaxation formula of Rosenfeld *et al.* [11]

$$P^{(n+1)}(J \rightarrow j) = \frac{P^{(n)}(J \rightarrow j) Q(J \rightarrow j)}{\sum_{\lambda \in \Omega_{2J}} P^{(n)}(J \rightarrow \lambda) Q(J \rightarrow \lambda)}. \tag{4}$$

In the updating formula, the support function $Q(J \rightarrow j)$ assesses how compatible the velocity vector $J \rightarrow j$ is with those in the local neighborhood G_J. Following [7], support is calculated by computing the product of the sum of the velocity vectors probabilities for all templates in the local neighborhood

$$Q(J \rightarrow j) = \prod_{I \in G_j} \sum_{i \in \Omega_{2J}} P^{(n)}(I \rightarrow i) R(I, J, i, j). \tag{5}$$

The term $R(I, J, i, j)$ has the same function as the compatibility coefficient of Rosenfeld and Kak [12] and gauges consistency between the velocity vectors $J \rightarrow j$ and $I \rightarrow i$. In [19], the compatibility between two velocity vectors was assessed as being a function of their magnitudes, orientations, and the distance between them. Here, a computationally simpler approach is used, based on the horizontal and vertical components of the candidate vectors. The information required for this is given directly by the candidate match generation procedure. Furthermore, as the templates are centered on a regular grid, the distance between them can be given by their relative grid positions

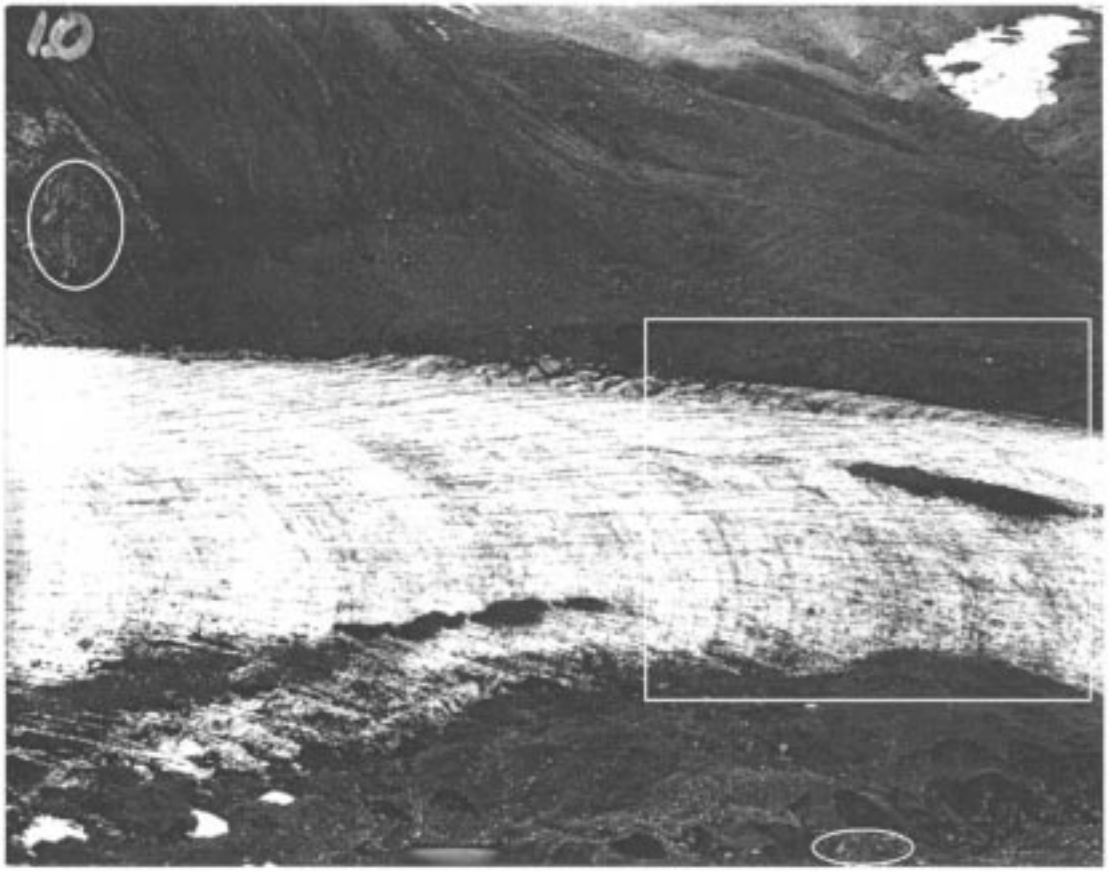

Fig. 2. First of image pair: the main branch of the Godley glacier, flowing left to right across the image, is joined by a tributary flowing from the top right corner of the image. Also shown are the fixed areas (ellipses) and the rectangular region used in Figs. 3–6.

rather than their absolute separation, which is computationally simpler and easily scaleable . This gives the function

$$R(I, J, i, j) = \exp\left[\frac{|\Delta x_{I,i} - \Delta x_{J,j}|}{\sigma}\right] \cdot \exp\left[\frac{|\Delta y_{I,i} - \Delta y_{J,j}|}{\sigma}\right] D(I, J) \quad (6)$$

where $\Delta x_{I,i}$ and $\Delta x_{J,j}$ are the vertical displacements of vectors being compared, and $\Delta y_{I,i}$ and $\Delta y_{J,j}$ are the horizontal displacements. The parameter σ is the global displacement variance that controls the rate of convergence. $D(I, J)$ is a function that depends on the relative position of the templates I and J, $D_{I,J}$, and is given by

$$D(I, J) = \max\{0, D_0 - D_{I,J}\} \cdot g \quad (7)$$

where D_0 is the distance within which templates are considered neighbors, and g is a constant in the approximate range of 1–5. $D_{I,J}$ is simply the sum of the horizontal and vertical separation between the two templates, measured on a scale where 1 is equal to the template spacing in pixels. This effectively makes the parameters σ and g independent from the actual pixel spacing between templates, enabling different spacings to be used without the need to alter the parameter values. In practice, the sensitivity of the results to the parameter values has been found to be relatively low, with variations in the parameter values not producing wide fluctuations in the results.

The number of iterations performed can be controlled using a termination criteria, given by the average absolute probability difference between successive iterations [12]. More simply, a predetermined number of iterations can be performed, found by experimentation or set using *a priori* knowledge.

V. Post Filtering

Although the motion field produced by the relaxation routine is smoothed, it still has some undesirable features. Templates with no candidates play no part in the relaxation process and thus remain unmatched. More critically, the best result that the relaxation routine can achieve for templates without any candidate displacement vectors consistent with the local flow is to select the most acceptable of the available candidates. Where this occurs, an undersmoothed flow with inconsistent vectors results. Previous relaxational matching techniques [3], [19] have simply assigned these templates a null label. The danger with such an approach is that large null clusters often result, especially if a strong smoothness constraint is imposed.

In contrast, the philosophy behind the approach presented here is that as the glacier is known to move in a slowly time-varying manner, every image point should in theory have a displacement vector associated with it and further, these vectors should be consistent with the local flow. With these aims, the post filter is used to improve the motion field produced by the relaxation routine in two ways. First, if there is sufficient support

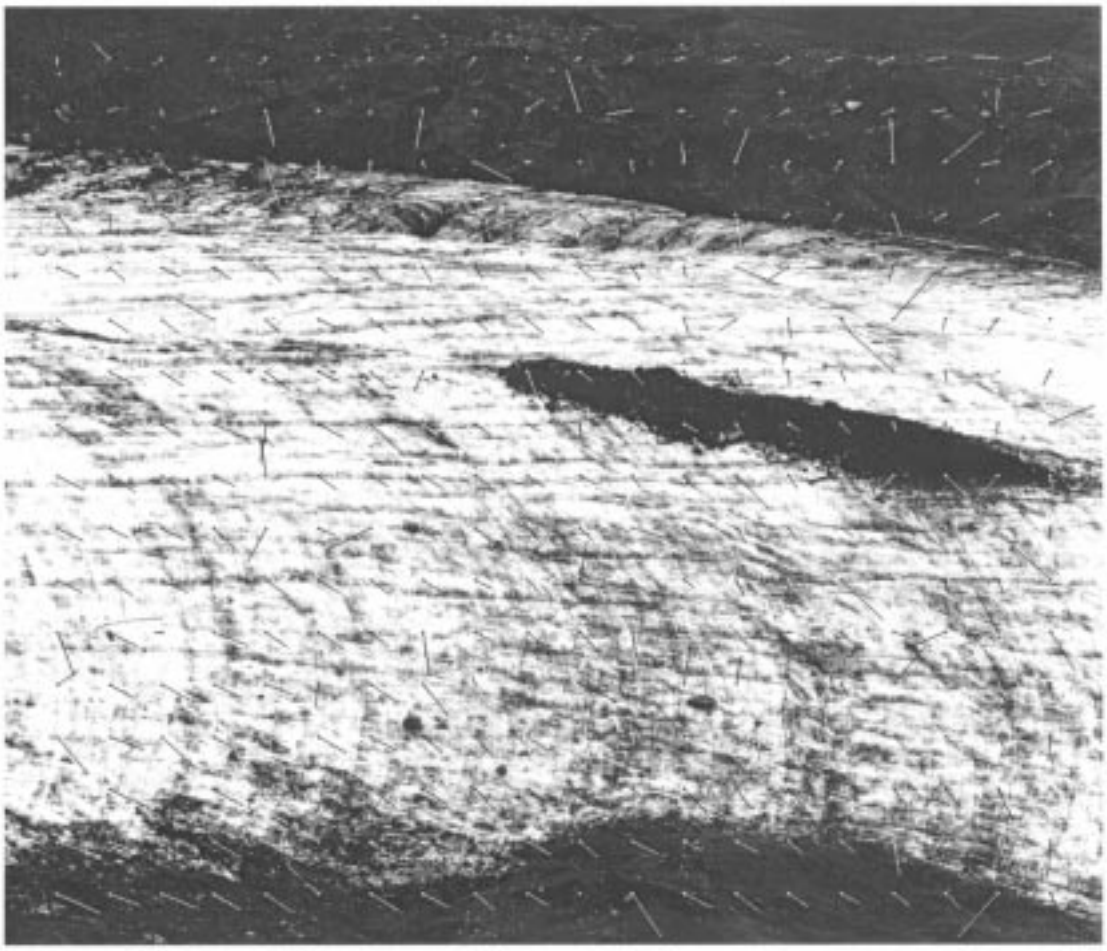

Fig. 3. Displacement vectors produced by the maximum cross-correlation method for the region from Fig. 2. The vectors are shown at ten times their actual size, and their origins are marked with a small cross. The choice of color is purely for reasons of clarity.

in the local neighborhood, templates with no candidate matches are assigned a velocity vector that is consistent with the surrounding region. Second, when the velocity vector selected by the relaxation scheme is the best of the available candidates but is inconsistent with the surrounding motion field, the assigned label is replaced by one that is characteristic of the local flow if there is sufficient support, or assigned unmatched otherwise. In this way, post filtering provides a velocity vector for points with no matches and corrects inconsistent matches, producing a smoother, more regular motion field. In the context of relaxation labeling, constraint filtering [15] has previously been used to assign nodes to a null class, using a model graph. The post filter proposed here is more powerful, as it uses the resultant displacement vectors within a local neighborhood to provide labelings other than the null category.

The operation of the post filter can be described as follows. As the feature points are centered on a regularly spaced grid, a local neighborhood for each point can be defined using a mask operator. For example, a 3×3 mask contains the central feature point and its eight nearest neighbors. The median horizontal and vertical components of the displacement vectors within the mask, denoted Δy_{MED} and Δx_{MED}, respectively, are found. Those neighborhood points with no significant matches are not included in the median calculations, and if more than half the feature points in the neighborhood are unmatched, the median displacements are classified as undefined, and the filtering process terminates.

The vector formed from the median horizontal and vertical displacements Δy_{MED}, Δx_{MED} provides a motion vector for those templates that is unmatched or inconsistent with the local velocity field. The former case results from no significant matches being found at the matching stage. For the latter, consistency with the local flow is determined by measuring the difference between the median and central displacement Δy_J, Δx_J vector, and if

$$\begin{aligned} &abs(\Delta y_J - \Delta y_{MED}) + abs(\Delta x_J - \Delta x_{MED})) \\ &\quad > k(abs(\Delta y_{MED}) + abs(\Delta x_{MED})) \end{aligned} \tag{8}$$

the proposed vector is deemed incompatible and is replaced by the median vector. The constant k is a smoothness constraint that relates the difference between the match and median displacements to the average horizontal and vertical values of the median displacements. A typical value is 0.5.

VI. Experimental Results and Discussion

The above method is applied to a pair of images from a time-lapse photograph sequence of the confluence of two branches of

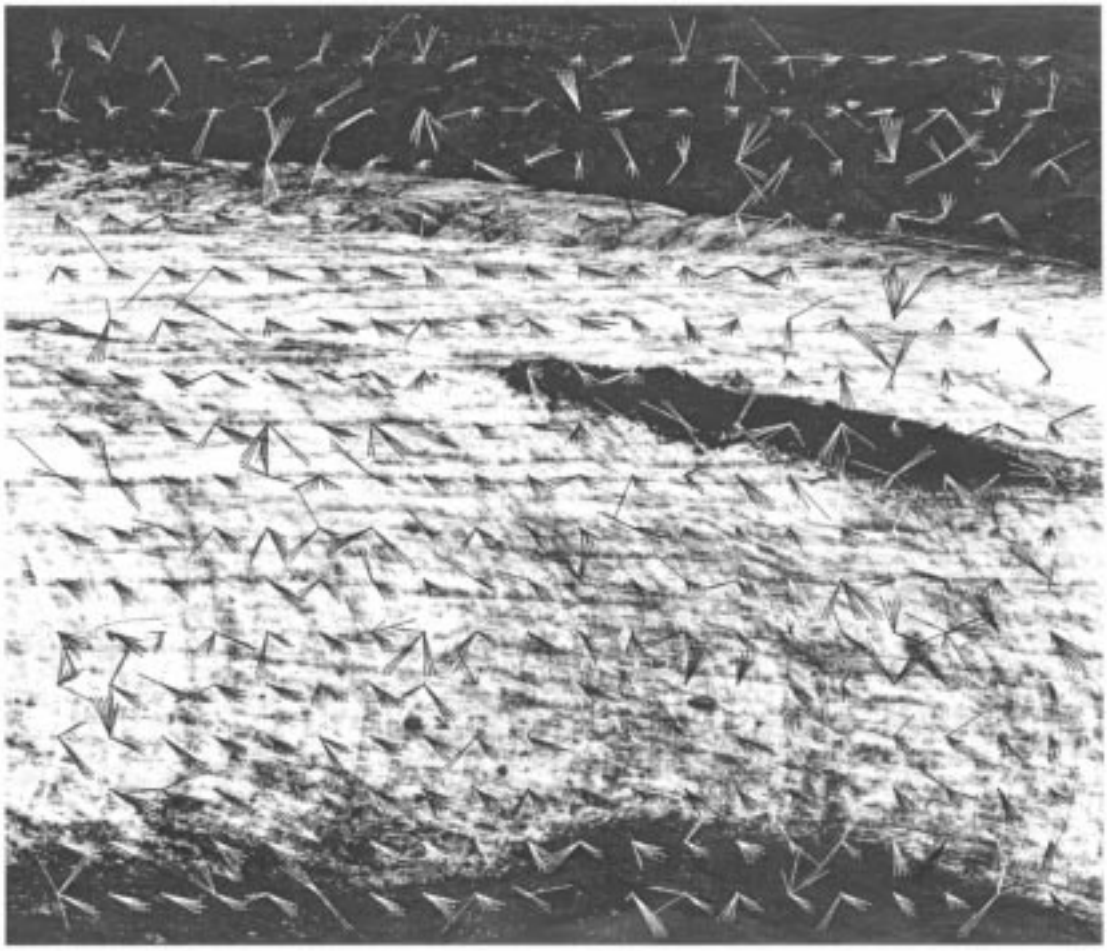

Fig. 4. Candidate displacement vectors for the region shown in Fig. 3, displayed using the same conventions. The maximum number of candidates for each template is nine. Those templates with no vectors are marked with a small cross at their center.

the Godley glacier, Mount Cook National Park, New Zealand. Fig. 2 presents the first image of the pair, showing the bare-ice surface of the glacier's main branch flowing left to right across the image. A predominantly debris-covered tributary extends from the top right corner of Fig. 2, where some bare ice is visible, to the junction with the main glacier in the center right of the image. The second image of the pair was taken 48 h later from the same site.

The 100×125 mm negatives were digitized to a resolution of 4096×5120 pixels onto a photo-CD, giving a resolution of 40 pixels per mm. Image registration was performed using multiple points from the fixed areas marked with white ellipses in Fig. 2. The fixed areas, identified during the field trip, consisted of a rock face and a large stationary rock located in the side of a ravine away from the glacier. The registration vector was predominantly translational, with only a small amount of distortion and rotation.

Using the CCC, matches were found for 11×11 templates, regularly spaced on a 100 pixel grid. The radius of the search area R_{MAX} was 9 pixels. The displacement vectors produced by the maximum cross-correlation (MCC) method for a region of at the confluence of the two glaciers, denoted by the white rectangle in Fig. 2, are given in Fig. 3. Although there is some consistency between vectors, there are many spurious results, and it is difficult to discern an overall picture of the optical flow. There are also a small number of feature points for which no significant candidate matches were found, marked by a small cross at the template center. This result demonstrates the inadequacy of the MCC method for this application, as it fails to cope with the inter-image differences caused by changes in the surface, nonrigid motion, and other factors.

Displacement vectors for the nine highest cross-correlation values for each template are shown in Fig. 4 and provide the input for the relaxation stage. The task for the relaxation routine is to select the most appropriate of these vectors according to the local flow. Fig. 5 presents the result produced after six iterations of the relaxation procedure, with the global displacement variance parameter σ of (6) equal to 5 and the parameters in (7) set as $D_0 = 2$ and $g = 1.5$. The result is clearly an improvement on the MCC method with a smoother, more consistent motion flow, from which it is possible to draw some conclusions about the vector field at the confluence of the glaciers. However, there are a significant number of vectors that are still inconsistent with the local flow. For example, the fourth vector from the right, sixth row down, is clearly incompatible with those around it. Examination of the corresponding point in Fig. 4 reveals that, although inconsistent, this vector is the most appropriate candidate displacement. A similar situation is found with the other

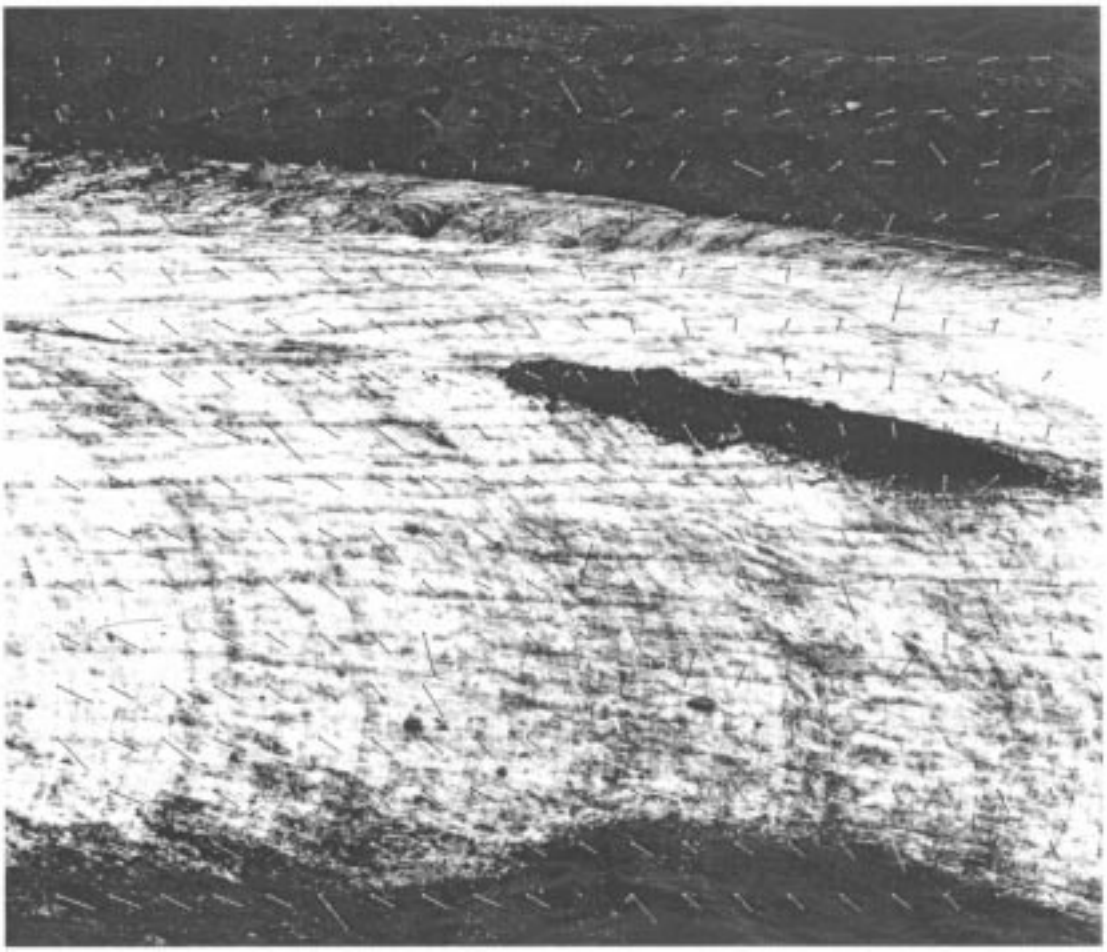

Fig. 5. Displacement vectors produced by relaxation labeling routine, using the candidate vectors and display conventions of Fig. 3.

inconsistent points, making the interpretation of the motion field more difficult. Those feature points with no candidate matches persist, as they play no part in the relaxation procedure.

After application of the post filter with a 3×3 mask size and $k = 0.67$, the result is improved to that of Fig. 6, which has a smooth, locally consistent motion field. For the 340 points, the post filter has provided displacement vectors for five unmatched points and changed 11.3% of the resultant vectors to that of the median displacements. This percentage can be adjusted by altering the value of k in (8). The overall result is a motion field that maintains the motion estimates produced by the relaxation routine for the majority ($\approx$90%) of templates, but which has replaced those inconsistent vectors and unmatched templates that reduce the accuracy of the vector field and hence presents a more easily interpreted picture of the surface motion of the glacier. The general direction of the vectors is consistent with the anticipated downstream movement.

This result was typical of that found over the whole of the glacier surface from Fig. 2. Over a 38 row $\times$ 52 column grid of templates with 100 pixel spacing, using the same parameters values as above, application of the post filter produced a median vector that was undefined at 31 points and provided appropriate vectors for 71 unmatched and 361 inconsistent points (3.6% and 18.3%, respectively, of the full grid). Overall, the percentage of matched points changed was higher than that for the subregion, indicating that the post filter had more effect in areas where the candidate matches showed less consistency. The only areas in which the technique failed to produce a motion field was in localized regions where the image was saturated, giving rise to clusters of unmatched points where the median displacement was undefined. This only occurred for 1.6% of the points.

A final comment concerns the problem of quantifying the precision of the velocity vectors. This is far more problematic than for the MCC method, where the correlation surface can be used to provide an estimate of confidence in the vectors, as the final velocities result from interactions between the candidate vectors at different locations. Obviously, the precision of those velocity vectors provided by the post filter is lower than for those resulting from the relaxation stage, but as the overall motion field is a better representation of the underlying motion, the average measurement error must therefore be lower than that of the MCC method.

VII. Conclusions and Future Work

A new technique for computing glacier surface motion has been presented. The technique uses digital images from scanned photographs to extend the application of automated image analysis from a regional scale, calculated from remotely sensed data, to a local one. To cope with the problems posed by the low

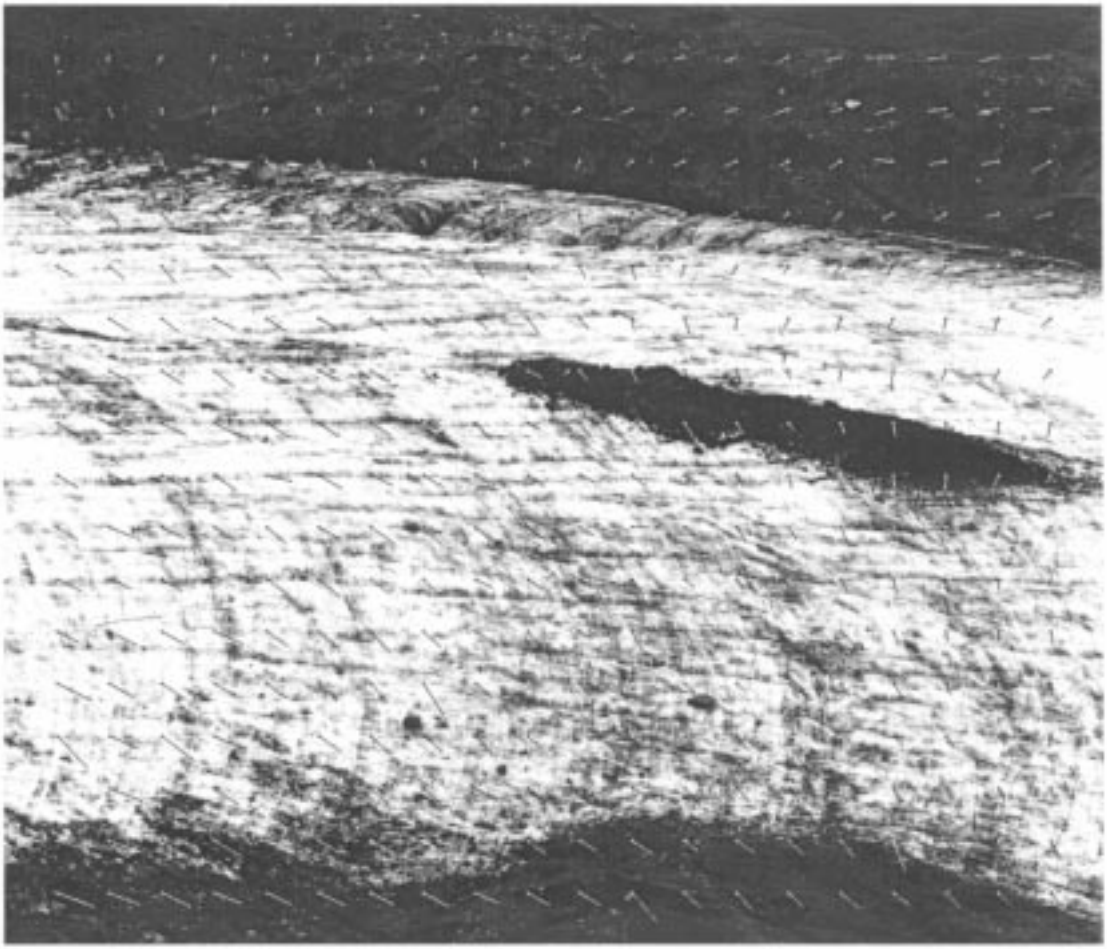

Fig. 6. Result of post filter applied to the displacement vectors of Fig. 5. The display conventions are the same as those in Figs. 3–5.

contrast images characterized by nonrigid motion, featureless regions, and surface ablation, a three-stage technique has been developed that consists of cross-correlation, relaxation labeling, and a post filter. Features of the operation of the post filter include the retention of the majority of the displacement vectors selected by the relaxation routine, while replacing incongruent vectors with ones consistent with local neighborhood.

The results produced are much improved over those produced by MCC technique and are also in agreement with the known direction of the glacier surface flow. The benefits conferred by the post filter can be seen in the smooth motion field produced. An area of further work is the development of a measurement error metric for the vectors, allowing accuracy estimates to be made for additional fields produced from the velocity field, such as mass flux.

In addition to the glacier study presented here, other application areas where the assessment of nonrigid motion is important include soil mounds, dam walls, and railway or road cuttings, where the movement over time is of interest for scientific or safety reasons. The process of data capture is one in which the rapidly developing field of digital camera technology is anticipated to make a major contribution and will enable high resolution digital images to be obtained directly in the field.

ACKNOWLEDGMENT

The author would like to thank Dr. E. W. Smith, A. J. Poad, and D. Scott for their contribution to the field trip.

REFERENCES

[1] S. T. Barnard and S. W. Thompson, "Disparity analysis of images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 2, no. 4, pp. 333–340, 1980.
[2] R. A. Bindschadler and T. A. Scambos, "Satellite-image-derived velocity field of an Antarctic ice stream," *Science*, vol. 252, no. 5003, pp. 242–246, 1991.
[3] A. N. Evans, "Full field motion estimation for large nonrigid bodies using correlation/relaxation labeling," in *Proc. IEE 6th Int. Conf. Image Processing and Its Applications*, vol. 2, London, U.K., July 1997, IEE Conf. Pub. 443, pp. 473–477.
[4] A. Goudie, *Geomorphological Techniques*, 2nd ed. London, U.K.: Unwin Hyman, 1990.
[5] W. D. Harrison, K. A. Echelmeyer, D. M. Cosgrove, and C. F. Raymond, "The determination of glacier speed by time-lapse photography under unfavorable conditions," *J. Glaciol.*, vol. 38, no. 129, pp. 257–265, 1992.
[6] W. D. Harrison, C. F. Raymond, and P. MacKeith, "Short period motion events on variegated glacier as observed by automatic photography and seismic methods," *Ann. Glaciol.*, vol. 8, pp. 82–89, 1986.
[7] J. Kittler and E. R. Hancock, "Combining evidence in probabilistic relaxation," *Int. J. Pattern Recognit. Artif. Intell.*, vol. 3, no. 1, pp. 29–51, 1989.

[8] B. Lefauconnier, J. O. Hagen, and J. P. Rudant, "Flow speed and calving rate of Kongsbreen glacier, using SPOT images," in *Proc. Remote Sensing of Arctic Environments 2nd Circumpolar Symp.*, Tromso, Norway, 1992, pp. 59–66.
[9] E. Rignot, "Tidal motion, ice velocity and melt rate of Petermann Gletscher, Greenland, measured from radar interferometry," *J. Glaciol.*, vol. 22, no. 142, pp. 476–485, 1996.
[10] E. Rignot, K. C. Jezek, and H. G. Sohn, "Ice flow dynamics of the Greenland ice sheet from SAR interferometry," *Geophys. Res. Lett.*, vol. 22, no. 5, pp. 575–578, 1995.
[11] A. Rosenfeld, R. A. Hummel, and S. W. Zucker, "Scene labeling by relaxational operations," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC–6, pp. 420–433, Dec. 1976.
[12] A. Rosenfeld and A. C. Kak, *Digital Picture Processing*, 2nd ed. New York: Academic, 1982, vol. 2.
[13] T. A. Scambos, M. J. Dutkiewicz, J. C. Wilson, and R. A. Bindschadler, "Application of image cross-correlation to the measurement of glacier velocity using satellite image data," *Remote Sens. Environ.*, vol. 42, no. 3, pp. 177–186.
[14] I. M. Whillans and Y.-H. Tseng, "Automatic tracking of crevasses on satellite images," *Cold Regions Sci. Technol.*, vol. 23, pp. 201–214, 1995.
[15] R. C. Wilson, A. N. Evans, and E. R. Hancock, "Relational matching by discrete relaxation," *Image Vis. Comput.*, vol. 13, no. 5, pp. 411–421, 1995.
[16] G. Wolberg, *Digital Image Warping*. Los Alamitos, CA: IEEE Comput. Soc., 1990.
[17] Q. X. Wu, "A correlation-relaxation-labeling framework for computing optical flow—Template matching from a new perspective," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 17, pp. 843–853, Aug. 1995.
[18] Q. X. Wu, S. J. McNeill, and D. Pairman, "Correlation and relaxation labeling: An experimental investigation on fast algorithms," *Int. J. Remote Sensing*, vol. 18, no. 3, pp. 651–662, 1997.
[19] Q. X. Wu and D. Pairman, "A relaxation labeling technique for computing sea surface velocities from sea surface temperature," *IEEE Trans. Geosci. Remote Sensing*, vol. 33, pp. 216–220, Jan. 1995.

Adrian N. Evans was born in Winchester, U.K., on December 7, 1967. He received the B.Eng. degree from Loughborough University of Technology, Loughborough, U.K., in 1990, and the Ph.D. degree in medical image processing from the University of Southampton, Southampton, U.K., in 1994.

From 1993 to 1994, he was with the University of York, U.K., as a Postdoctoral Research Associate. From 1994 to 1997, he was with Massey University, Palmerston South, New Zealand, as an Assistant Lecturer and Lecturer in information engineering. Since September 1997, he has been a Lecturer in the Department of Electronic and Electrical Engineering, University of Bath, U.K. His current research interests include noise modeling, nonlinear filtering, feature extraction and motion estimation in application to medical and remotely sensed images, and still-image and video coding.