PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Applications of Remote Image Sensing with the NASA Massively Parallel Processor [and Discussion]

M. Halem, D. Lane and J.-P. A. L. Muller

Phil. Trans. R. Soc. Lond. A 1988 **324**, 365-372
doi: 10.1098/rsta.1988.0025

Phil. Trans. R. Soc. Lond. A **324**, 365–372 (1988)
Printed in Great Britain

Applications of remote image sensing with the NASA massively parallel processor

By M. Halem

Space Data and Computing Division, NASA/Goddard Space Flight Center, Greenbelt, Maryland 20771, *U.S.A.*

[Plates 1–4]

Satellite observing systems are producing image observations of the Earth's surface and atmosphere with spectral and spatial resolutions that result in data rates that current general-purpose computing systems are incapable of processing and analysing. As a result, current processing systems have been able to analyse only limited amounts of image data with less than optimal algorithms for generating high-quality geophysical parameters. A massively parallel processor (MPP) is operationally available at NASA/GSFC for routine image-analysis applications. Research studies with the MPP are being pursued in the area of interactive spatial contextual classifications for the land thematic mapper data, automatic SIR-B stereo terrain mapping, ice-motion detection, faint-object image restoration and other general purpose ocean and land image-processing systems. Several applications are presented comparing the MPP products with enhancements of imaging data with standard image-processing methods. Finally, a work-station parallel processor for space station on-board image processing will be described.

1. Introduction

Satellite observing systems are capable today of spectrally imaging the Earth's surface and atmosphere at data rates far exceeding our ability to process or extract their full information content in any practical manner. This has resulted in limited duty-cycle use of our instruments with high data rates and, very often, archiving vast amounts of unprocessed data for possible future access. With advances in charge-coupled device (CCD) technology, more of the passive sensors with lower data rates will evolve into imaging sensors with even higher spectral and spatial resolutions provided by their densely packed arrays of detectors. As we move into the permanently manned space station era, an Earth-observing system is planned to carry these facility scale instruments so as to provide continuous images that will monitor the long-term changes in the Earth's bio-physical and chemical cycles. Along with developments in passive sensing, massive image data sets will also be produced from space-borne active sensors that use very large real or synthetic apertures. The impact of these observing systems could lead to a revolution in our understanding of Earth system science, but to achieve that goal a radically new technology based on parallel computing architectures and algorithms will have to be developed.

Recognizing the enormous image processing requirements for the Earth-science sensors and similar imaging requirements emerging in the planetary, solar-terrestrial and astrophysical sciences, NASA embarked, in 1979, on a pilot project to design and build a massively parallel processor (Potter 1985) that will begin addressing these problems. The system was delivered

in 1983 by the Goodyear Aerospace Corporation (Batcher 1980) and shortly thereafter was made available on a selective basis to a wider community. A research programme was initiated in 1985, through a national solicitation, that established an MPP working group that is exploring the algorithmic potential of this uniquely powerful architecture. In this report, I present selected applications from the MPP working group investigations that dealt with remote Earth-sensing image algorithms. Section 2 contains a short description of the MPP system, and §3 describes automated Shuttle Imaging Radar-B (SIR-B) stereo terrain mapping and the application of similar techniques for sea-ice motion detection. Section 4 presents multispectral classification algorithms applied to *Landsat* thematic mapper data. Section 5 outlines potential future developments and applications of massively parallel architectures that could be applied to on-board space-station processing.

2. Description of MPP

The MPP is an integrated system of 16384 processors, each containing its own local memory of 1024 bits, and a global data staging memory of 32 megabytes. All processors perform the same single-bit arithmetic and logical operations at instruction speeds of 100 ns (10 MHz). The processors are arranged in a grid connected array 128×128 with single-bit transfer to nearest neighbours in north, south, east or west directions simultaneously executed at the same cycle speeds. Automatic east–west and north–south wrap-around symmetry occurs also at the same clock cycle.

Transfer from the global data staging memory to the array memory is at 80 megabytes per second and consists of 128 bits moving into or out of a column or row simultaneously in two cycles. Three controllers govern the execution and synchronization, and global broadcast of instructions to all processors. The system is accessed through a front-end Digital Equipment Corporation VAX 11/780 processor and a DR-780 interface to the MPP with a transfer rate of 5 megabytes per second. A functional sketch of the hardware layout is shown in Fischer (1987). A high-level parallel extension of the PASCAL language is available for general scientific programming on the MPP. Assembly coded macroinstructions for full floating point operations are transparently provided under MPP PASCAL. Programs are compiled and loaded through the VAX front-end.

Arithmetic performance speeds for addition and multiplication can range from 40×10^9 operations per second for 1-bit addition to 1.68×10^9 operations per second for 32-bit addition and approximately a factor of ten slower for fixed point multiplication. Floating point 32-bit addition and multiplication operations are performed at 450 and 215 MFLOPS (million floating-point operations per second), respectively. Overall, MPP machine computing performance relative to a CYBER 205 and Cray-1 for vector modelling computations is three to five times faster, respectively, and upwards of 100–1000 times faster for appropriate fixed-point algorithms.

3. Automated stereo elevation mapping from space

In this section, we present results of a fully automated stereo matching algorithm for the determination of terrain elevation data from space-borne synthetic aperture radar (SAR) imagery data. This is work that was developed at Goddard Space Flight Center and

implemented on the MPP (Strong & Ramapriyan 1987). Their MPP algorithm is based on a low-level vision technique (Marr & Poggio 1979 as extended by Quam 1984).

The algorithm was applied to a pair of overlapping SIR-B images acquired during the October 1984 Space Shuttle Challenger flight. The images were taken over a plateau region of Bangladesh at incidence angles of 25° and 42°, respectively. As a result of the viewing geometry, a linear scale change in the slant range (stereopsis) direction is first applied to the latter (i.e. test) image to produce the same pixel resolution as the former (i.e. reference) image. For each pixel in the reference image, a rectangular neighbourhood surrounding it is correlated with a set of neighbourhoods corresponding to each point in a search area in the test image. The disparity function is then determined by associating in the stereopsis direction the location of the corresponding centre pixel of the neighbourhood within the search area with the maximum correlation value. Badly matching regions of the disparity function are defined by adjacent pixels with jumps greater than one and filled in by an interpolation procedure based on an analogue to a finite-difference heat-flow equation applied to the region. The disparity function is then used as a geometric correction function to warp the test image so that it more closely matches the reference image. The algorithm is iterated for the entire image with successively smaller neighbourhoods. Figure 1, plate 1, shows the reference and test images overlayed after the initial linear warping and edge alignment. The two images are displayed with red and green, respectively, as maximum brightness values. One readily observes a broad plateau region with a surrounding river valley on the right. Figure 2, plate 1, shows the results of the matching algorithm after two iterations starting with rectangular neighbourhoods of 25×25 pixels and 13×13 pixels, respectively. The coincidence of the red and green in contrast with figure 1 indicates the agreement of the match. Figure 3*a*, plate 3, shows the disparity function with dark regions representing pixels in the reference image to the right of corresponding pixels in the test image and white regions to the left. Disparity is linearly proportional to elevation with dark areas corresponding to low elevations. By a three-dimensional graphics-rendering capability developed for the MPP (Dorband 1986), figure 3*b*, plate 3, displays a three-dimensional view of this region illuminated from a source approximately at the location of the shuttle at 42° from vertical. Overlayed on the disparity functions are the observed brightness values. An examination of a cartographic contour map of the same region available at NASA shows very good feature agreements between plateau, river, valleys and height gradients for the entire region. This calculation required more than 10^{11} operations in a little less than a minute which could not have been computed in less than a week of dedicated, around-the-clock VAX 11/780 CPU time. More importantly, these results establish the viability of generating reasonably accurate estimates of terrain height from space.

4. Sea-ice motion detection from time-varying SAR imagery

Unlike elevation mapping by stereo matching, where surface features are stationary in time and all disparities in multi-incident viewing image pairs are attributed to variations in elevation, detection of ice-floe motion assumes 'identical viewing geometries' in pairs of images and attributes local differences to motion within the scene. In this study, we report on algorithms for detection of ice motion developed for the MPP (Strong & Ramapriyan, 1987) and some preliminary results achieved thus far for two-dimensional ice motion detection.

The study is conducted with data acquired in 1978 from nearly identical viewing angles of

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

the SAR sensor flown on *Seasat* as kindly provided by C. Elachi, of the Jet Propulsion Laboratory. We assume that a three-day interval as used in this study is sufficiently small that the dynamics of motion will not cause the distortion of the ice surface to the extent that an area in one image will not be recognizable in the other.

A major difference to be considered in extending the ideas of the automated stereo matching algorithm is the need to account for rotation and the fact that large changes are possible between image pairs. Because of the latter condition, we first apply a global cross correlation between the reference image and test image. This is performed efficiently by assuming periodicity in both horizontal and vertical directions and performing a two-dimensional fast Fourier transformation. Because translations of large coherent regions produce local maxima or peaks in the global cross correlation, the locations of these peaks can be used to define an approximate translation vector between the corresponding pixel of a given peak and the centre pixel of the global correlation. After translating the test images by an amount corresponding to a given local maximum, the local correlation technique of the type used for stereo matching is then used to identify the corresponding regions of points in the test image which match those in the reference image. A threshold value is used to identify those points in this region which will be stored in a two-dimensional array. The process is applied to all local peaks creating a sparsely filled, two-dimensional array representing the global displacement field. In accounting for rotation of large regions, we apply a sequence of small finite rotations to all points of the test image. We can fill the boundary elements of the test image if the original test image contains pixels outside the 512×512 region or else fill in by zeros. The process is then repeated for the rotated test image thereby building up the displacement functions. At the end of a complete rotation sequence, gaps in the displacement array are filled in, as in the case of stereo matching, by a finite-difference diffusion approximation. This algorithm is much more computer-intensive than stereo matching. At the time of this report, only the global correlation and a single translation have been completed. The investigators are in the process of assessing the size of regions which can be detected with the global correlation functions. Figure 4*a*, *b*, plate 3, shows two nearly identical views of Banks Island, located northeast of Alaska, taken by *Seasat* with a three-day interval. The local horizontal and vertical correlation values are calculated after performing a global cross correlation and translating the test image by an amount given by the correlation peak. In figure 5*a*, *b*, plate 4, the large solid dark regions represent small values of horizontal and vertical displacements as one expects after the test image has been translated. Figure 5*c*, plate 4, shows the maximum local correlation value at each pixel indicating a good match in the areas with a low amount of translation. The pixels in this region contributed to the local maximum in the global correlation function. The scheme does seem to be working as expected as far as we have thus far gone and holds extreme promise of being reasonably successful for automatic ice-motion detection.

5. Contextual classification algorithms

A major goal of multispectral remote sensing has long been the desire to segment and classify scenes of the Earth's surface according to their physical and biological spectral characteristics. The inherent difficulty of this problem lies principally in two sensor-related areas, namely, the non-uniqueness or even lack of wide band spectral signatures and the inadequate temporal and spatial resolution and sampling for any given scene. A third and often overlooked limitation has

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

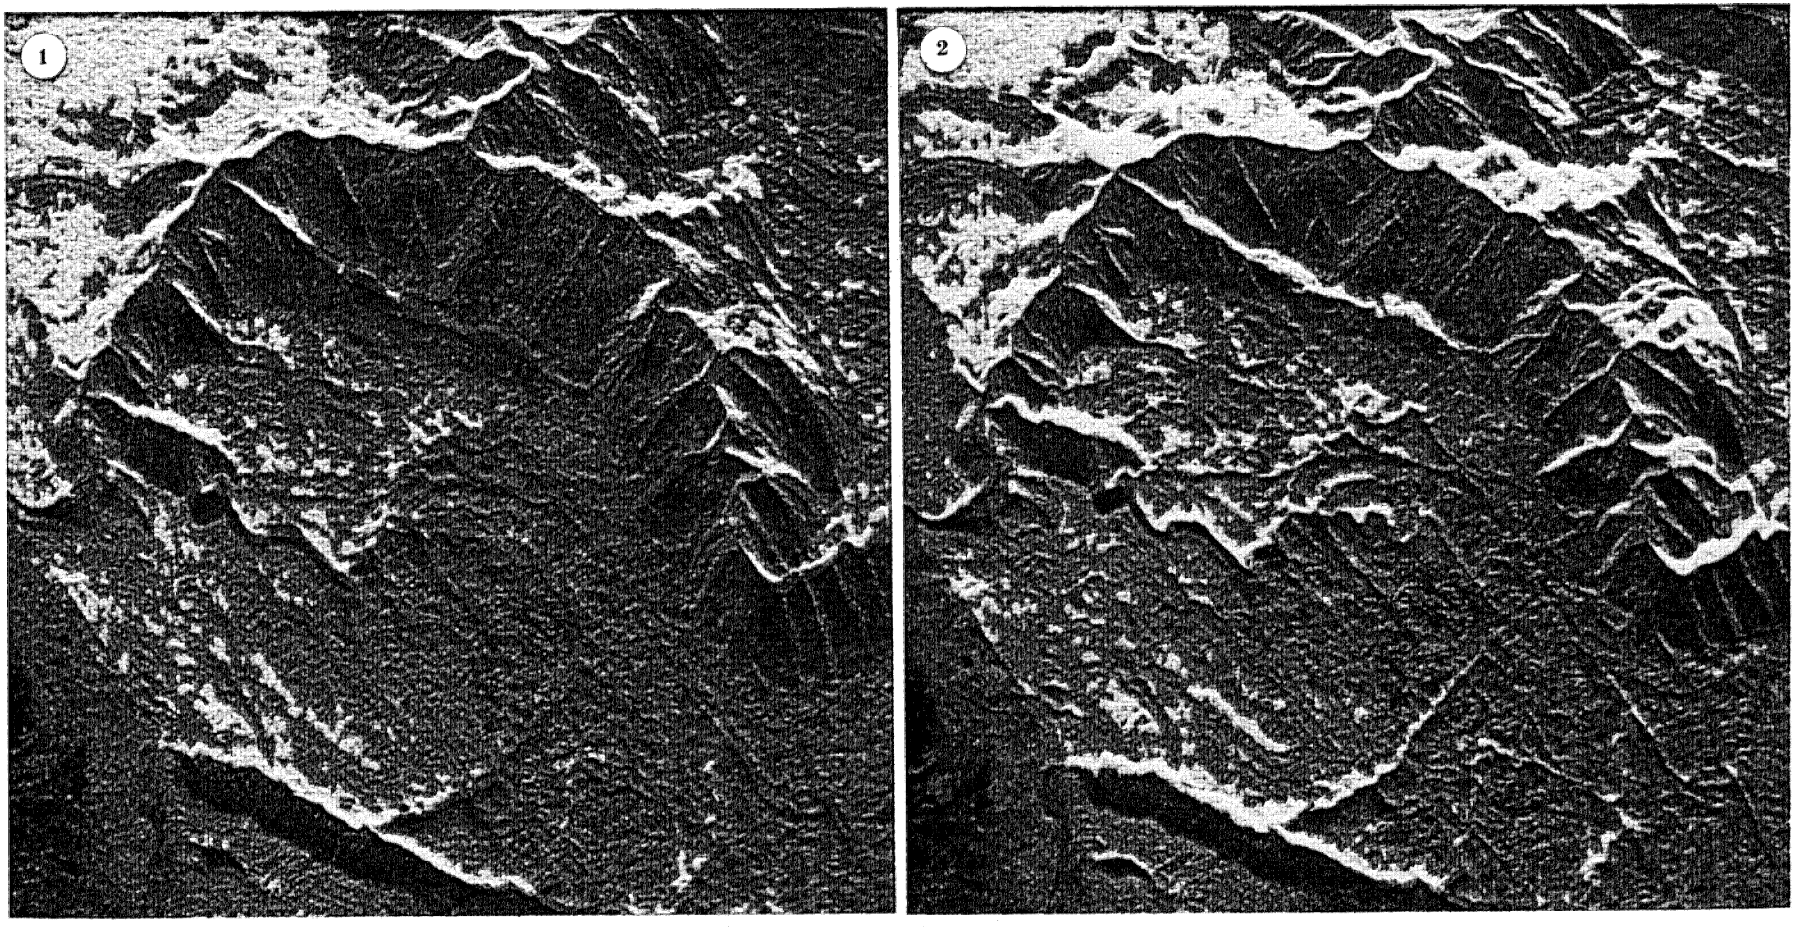

FIGURE 1. Pseudo-stereo overlay of SIR-B shuttle images. Red and green indicate maximum radar reflection values, respectively.

FIGURE 2. Overlay after application of two iterations matching algorithm. The blending of red and green into yellow indicates good agreement resulting from warping.

Phil. Trans. R. Soc. Lond. A

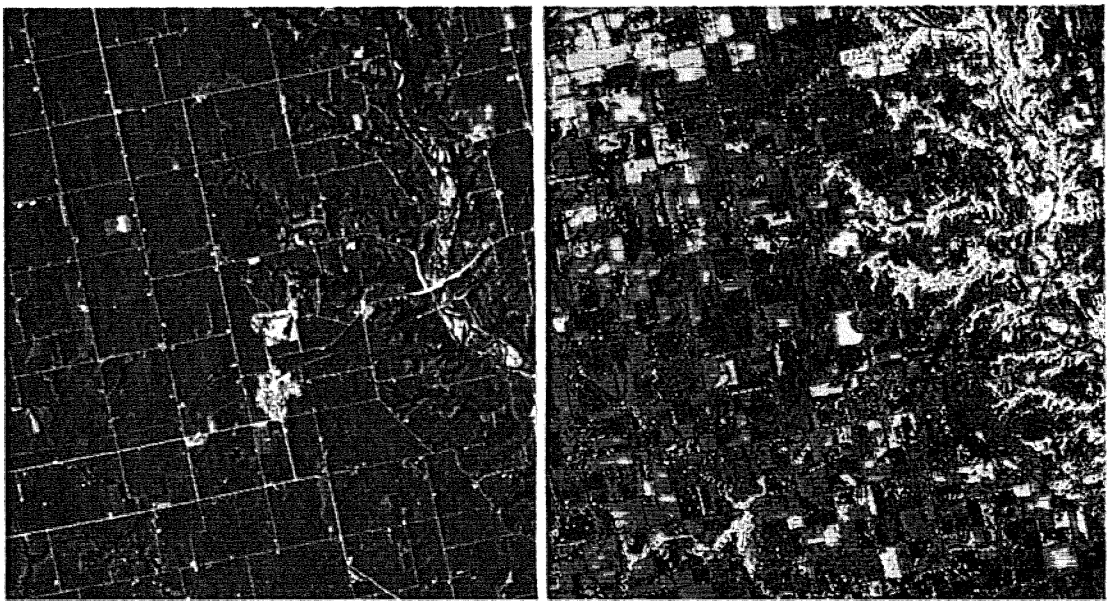

FIGURE 6. Portions of a *Landsat* scene before (left) and after (right) the application of 100 iterations of unsupervised ISODATA classification algorithm.

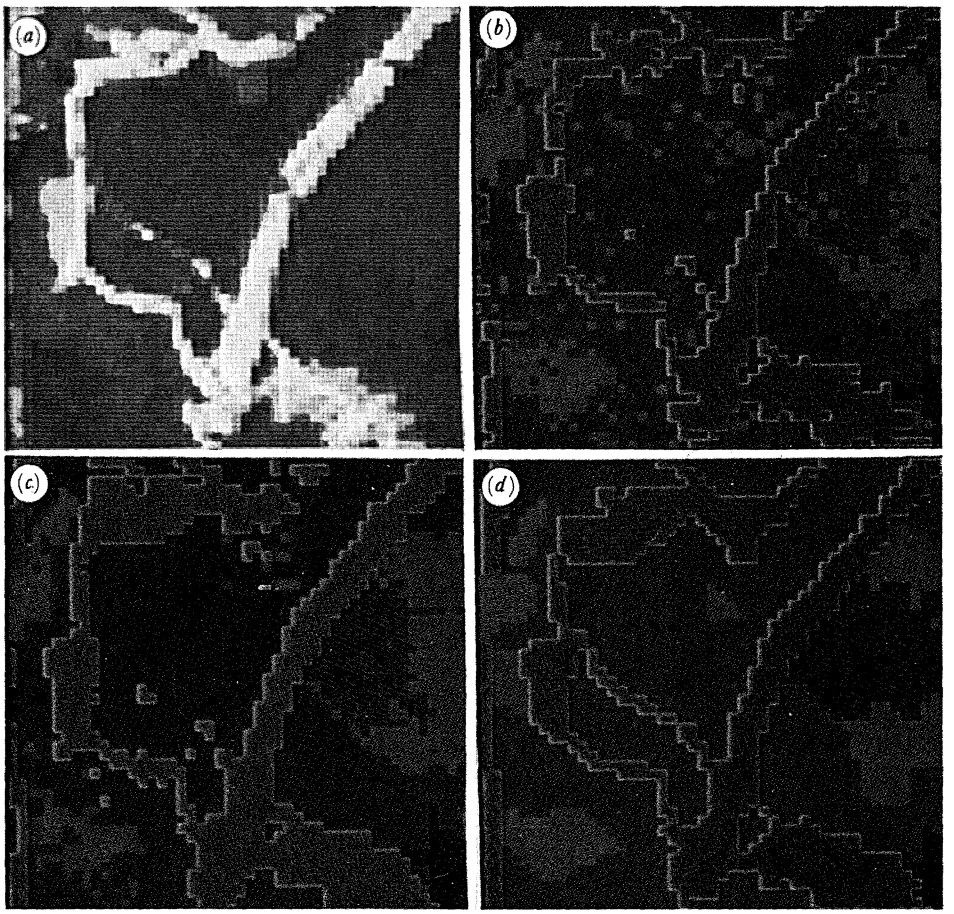

FIGURE 7. Three image representations of a *Landsat* scene: (*a*) unclassified; (*b*) after application of maximum-likelihood classification algorithm; (*c*) after application of five-point contextual extension of maximum-likelihood classifier; (*d*) ground-truth classification.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

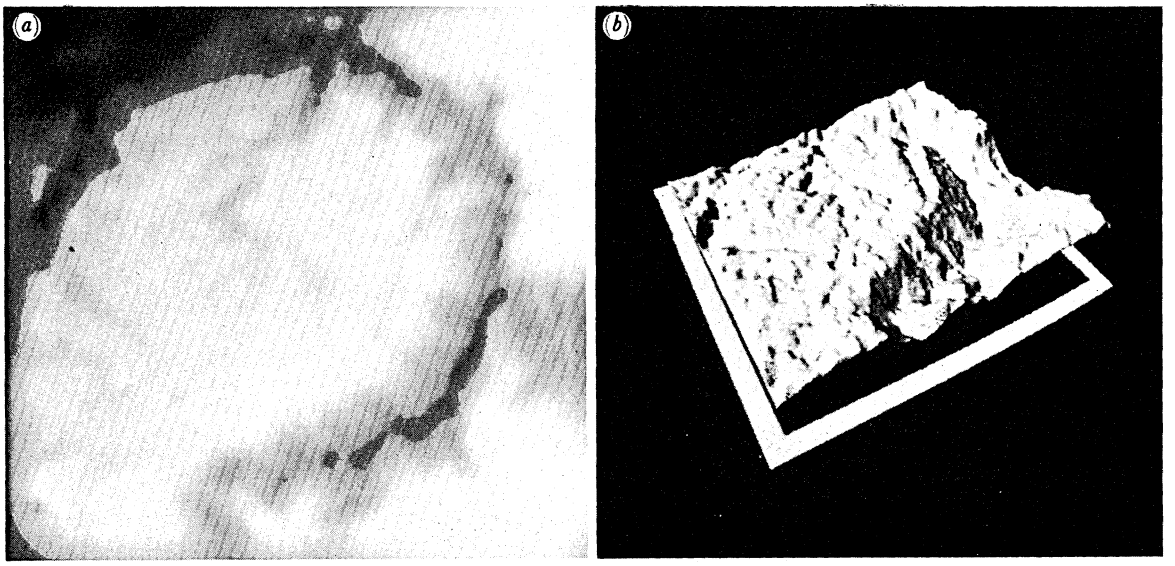

FIGURE 3. (*a*) *SIR-B* stereo disparity function after second iteration. Dark regions indicate low elevation and white regions indicate high elevation (*b*). Three-dimensional perspective rendering with stereo disparity function as height surface illuminated from a light source at an angle 42° from vertical.

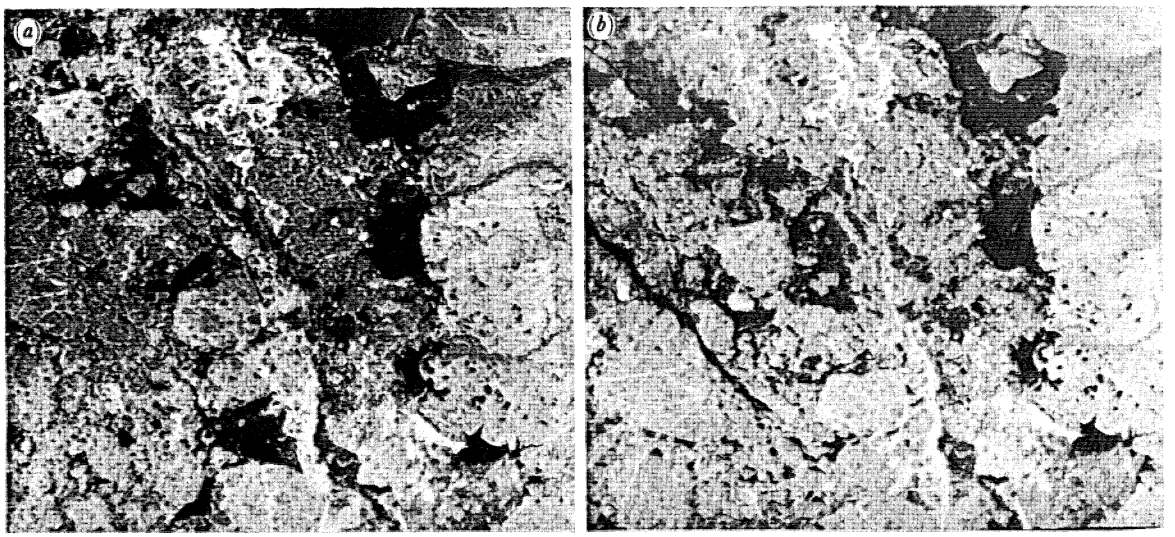

FIGURE 4. Near identical views of Bank Island, Alaska, taken by *SEASAT* synthetic aperture radar. (*a*) Image taken three days before image (*b*).

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

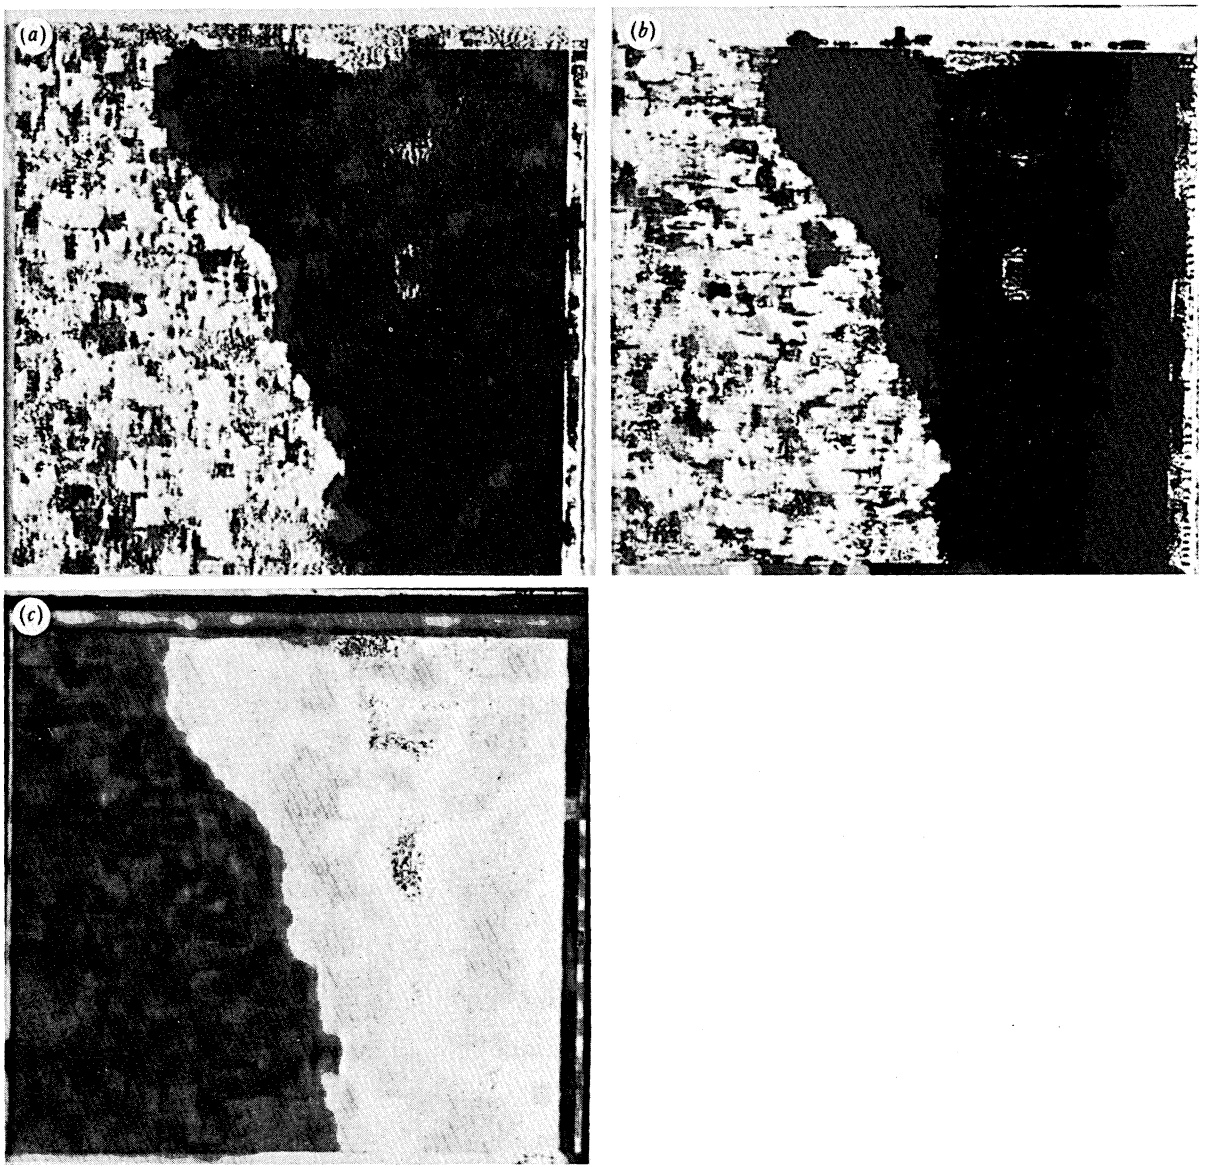

FIGURE 5. (*a*) Horizontal and (*b*) vertical translation functions, respectively. Large dark areas indicate little or no translation. (*c*) Match score values of displacement array. Brightness shows good matches within overlapping ice floes.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

also been the inadequacy of current computational resources required for complex processing, analysis and classification of the massive volumes of pixel information available per scene.

Whereas both of the former factors are being addressed through the design of more advanced multispectral sensors with narrower spectral intervals, slower progress is being made in the development of improved image-analysis techniques. Although conceptual methods holding great promise have been available for some time, detailed testing progress has been limited until now because of the great demands on computational power that are needed to implement these techniques.

In this section, I report mainly on the use of the MPP to implement a contextual classification algorithm (Tilton & Strong 1984). Results are first presented for an unsupervised classification algorithm on the MPP which is needed as the initialization step for the generalized maximum likelihood classification algorithm. The unsupervised algorithm is known as ISODATA and is a straightforward spectral clustering method using the speed of the MPP to carry out the enormous number of calculations for a very large class of classification categories. Assume we are interested in M classification categories of a multispectral scene which has $N \times N$ pixels and where each pixel has L associated observed spectral bands. The algorithm consists of initially selecting M points in the spectral phase space of L-tuples labelled $m_i^{(0)}$, $i = 1, \ldots, M$ considered to be the spectral centroids of all pixels in a certain category. The $m_i^{(0)}$ can be chosen by performing a principal eigenvector decomposition over all pixels in the scene and selecting the $m_i^{(0)}$ centroids along this principal direction. Each pixel in the scene p_k, $k = 1, \ldots, N^2$ is associated with the centroid category $m_i^{(0)}$ if the condition that $\|p_k - m_i^{(0)}\| < \|p_k - m_j^{(0)}\|$ holds for all $j \neq i$. The norm $\| \ \|$ is taken as the usual euclidean L distance measure. The algorithm then consists of iterating for a new set $m_i^{(1)}$ obtained by calculating the mean for all pixels associated with $m_i^{(0)}$ and defining the new associated classes. Elaborations to incorporate cluster splitting, cluster combining and pixel weighting are ISODATA modifications planned, but not yet implemented, on the current MPP version. Figure 6, plate 2, illustrates the classifications starting from the image shown in figure 6*a* for a 512×512 *Landsat* TM scene over Iowa, after 100 iterations. The time to do this algorithm on the MPP was 18 s as compared with about 7 h on a VAX 11/780. This already demonstrates that the MPP can be used to convert ISODATA into an interactive tool for use by a scientist at a workstation remotely connected to the MPP.

Next, the maximum likelihood classifier (MLC) is applied to resolve the scene further according to *a priori* spectral class statistics obtained from training site data. Two sets of statistics were used in the study. One set characterizes the spectral information for each ground-cover class and is obtained from a per-pixel analysis of data from training sites within the TM scene being analysed. The training-site data had previously been correlated with field data giving the true ground-cover label for each pixel in each training site. The other set characterizes the contextual information for each pixel location and is compiled for individual pixels and their four nearest neighbours. The intuitive procedure whether doing MLC or generalized MLC algorithms consists in developing a decision rule for classifying each pixel. In the generalized MLC we applied the rule that associates the generalized pixel P_k with category i if $g_i(P_k) \geqslant g_j(P_k)$ for all $j \neq i$ and where $g_i(P_k)$ is given as a discriminant function for category i by $g_i(P_k) = \ln(D_i) - (P_k - M_i)^T C_i^{-1} (P_k - M_i)$, where P_k and M_i are single or spatially arranged vectors, C_i is the statistical covariance matrix associated with category i, D_i is the associated determinant of covariance matrix C_i.

The underlying rationale for embedding contextual spatial statistical information into the

generalized maximum likelihood classifier rests on the assumption that *a priori* probabilities can be established for likely and unlikely geometric distributions of certain category types. For example, consider three classes of ground cover (forest, corn and water) represented in three different spatial configurations. Clearly, on a per-pixel basis, one would not expect to find a field of corn surrounded by water or a forest at a 30 m resolution. Although this example is simplistic, many situations do arise where we can incorporate probabilistic spatial distributions. Thus, we extend the associated phase-space vector tuple by considering as components of the vector the spectral band information of all four nearest neighbours in a prescribed order. This not only increases the pixel vector size linearly by the product of the number of spatial pixels and their bands, but also increases the complexity of the covariant matrix by the square of the increased vector size. For example, a five-point pixel stencil with seven bands increases the size of the covariant matrix from 7×7 to 35×35. For a *Landsat* TM scene this involves 6967×5969 pixel elements, well beyond the capability of any of today's mini processors.

Preliminary results have been obtained (Tilton 1986) on the application of this algorithm to an airborne multispectral scanner (ABMSS) scene acquired on 29 July, 1978, which covers a 2.8 km $\times$ 2.8 km area in British Columbia, Canada near the Anderson River. This data set is a part of a SAR–MSS data set that was acquired, preprocessed, and loaned to us by the Canada Centre for Remote Sensing (CCRS), Department of Energy, Mines, and Resources, of the Government of Canada. The data set covers a forested area consisting of areas of pure Douglas fir, Douglas fir mixed with lodgepole pine, Douglas fir mixed with cedar, Douglas fir mixed with hemlock, hemlock mixed with Douglas fir, hemlock mixed with cedar, forest clearings (including clearcuts and logging roads) and rock outcroppings. Figure 7*a*, *b*, *c*, *d*, plate 2, shows the unclassified image, a maximum-likelihood classification assuming each class has equal *a priori* probabilities, a five-point contextual maximum-likelihood classification, and the ground-truth classification, respectively. Figure 7*c* clearly shows much improved spatial continuity within the classification categories. Performing a statistical comparison with ground truth for the entire 2401 pixels which were classified, the spatial MLC had a 79.4% correct classification compared with 75.5% for the standard MLC for the eight classes listed above. Although only a small 3.9% improvement was obtained for this statistical validation, the contextual classification result looks visually much closer to the ground truth classification, having a much 'cleaner' appearance. We are pursuing more extensive and varied test regions where extensive ground truth data are available. On a larger data set (where only limited ground truth data were available), the time required to perform a 512×390 contextual classification on the MPP was approximately 1 h. This compares favourably with doing the per-pixel MLC on a VAX. The MPP therefore affords us an opportunity to explore spatial classification techniques that until now have been considered too time consuming.

6. Future projects

As pointed out in the introduction, space-borne imaging sensors for remotely observing the Earth will continue to evolve in capabilities and continue to add more massive volumes of valuable scientific data. Already, the planned capability of instruments for acquiring data on the earth-observing system platform of the space station greatly outstrips our ability even to transmit all this data back to our ground processing systems. In addition, it is expected that scientists in the manned space-laboratory modules will want to carry out research campaigns

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

on events of opportunity requiring on-board image processors of extensive performance capabilities. As a result, studies are underway at NASA to explore new technologies required for greater on-board processing capabilities, to capture greater amounts of potentially useful data.

The CMOS (complementary metal–oxide semiconductor) technology that was used in the development of the MPP, for which I have here attempted to demonstrate a unique potential for remote sensing applications, was based upon the latest custom chip capabilities of the late 1970s and early 1980s. The MPP array uses 88 boards each containing 24 chips. On each chip are embedded 8000 transistors allowing the design of 8 complete processors that use circuits of line width 4.5 μm. In addition, adjacent to each chip on the board are two memory chips of 4 K bits each with appropriate pin connections to the processor chips thus providing 1 K bits of memory per processor. Parallel computers with arrays of 16000 to 64000 processors are now commercially available at prices only slightly greater than mini-processors of the VAX-11/780 class. Parallel processors with somewhat smaller arrays and deeper memories are also becoming available at costs comparable with SUN or micro-VAX workstations. Clearly, parallel processing is the likely way to go in the pursuit of new image processing capabilities.

In this regard, NASA supported research studies have already shown that one can produce chips with approximately one million transistors allowing for the design of 128 processors on a single chip with each processor containing over 1 K bits of on-chip memory. Memory chips are also available with 4 megabits allowing for an additional 16 K to 64 K bits of off-chip memory to each processor. Thus, well before 1995 when the space station is expected to be operational, MPP-like workstations in a single board should be available with more power than today's MPP. It is similarly expected that a fully space qualified MPP-like workstation will also be available for research aboard the manned laboratory modules of the space station. In all probability, one could also expect to find at the large ground-processing facilities, parallel processors with arrays of more than a million processors and internal memories of more than a gigabyte.

Thus, the immediate challenge facing the remote-sensing community is not in the hardware arena, but rather the development of greatly improved parallel image algorithms that will exploit these new computing architectures to maximize the extraction and use of the full scientific information content of image data.

I express my deepest appreciation particularly to Dr Strong, Dr Ramapriyan and Dr Tilton not only for allowing me to present elements of their current research at this discussion, but for taking considerable time from their busy work schedules to explain and often tutor me in the methodologies they used.

References

Batcher, K. E. 1980 Design of a massively parallel processor. *IEEE Trans. Computers* **C-29**, 836–840.

Dorband, J. E. 1987 Ray tracing on the MPP. In *Proc. 1st Symposium on the Frontiers of Massively Parallel Scientific Computation, Greenbelt, 24–25 September 1986*. Conference Publication 2478, NASA/Goddard Space Flight Center.

Fischer, J. 1987 In *Proc. 1st Symposium on the Frontiers of Massively Parallel Scientific Computation, Greenbelt, 24–25 September 1986*. NASA Conference Publication 2478, NASA/Goddard Space Flight Center.

Marr, D. & Poggio, T. 1979 A computational theory of human stereo vision. *Proc. R. Soc. Lond.* B **204**, 301–328.

Potter, J. L. 1985 *The massively parallel processor*. Cambridge, Massachusetts: MIT Press.

Quam, L. 1984 Hierarchical warp stereo. In *Proc. Image Understanding Workshop, New Orleans, October 1984*, pp. 149–155. DARPA.

Strong, J. P. & Ramapriyan, H. P. 1987 Massively parallel correlation techniques to determine local differences in pairs of images. In *Proc. 2nd Int. Supercomputer Conference. Santa Clara, 5–8 May 1987*. St Petersburg, Florida: International Supercomputing Institute, Inc.

Tilton, J. C. 1986 Contextual classification on the massively parallel processor. In *Proc. 1st Symposium on the Frontiers of Massively Parallel Scientific Computation, Greenbelt, 24–25 September 1986*. NASA Conference Publication, 2478, NASA/Goddard Space Flight Center.

Discussion

D. LANE (*Intelligent Automation Laboratory, Department of Electrical and Electronic Engineering, Heriot-Watt University, Edinburgh, U.K.*). The hardware architecture of the MPP provides a processing tool that is potentially very powerful. However, its usefulness will depend upon the ease with which users can express the solution to their problem in a machine format. How, therefore, is software written for the MPP, and to what extent does it allow maximum exploitation of the MPP's capabilities?

M. HALEM. There is, of course, an assembly language, but in addition there are several fully developed and supported high-level languages with parallel extensions. The most popular and frequently used language is MPP PASCAL in addition to FORTH. An implementation of pure LISP has been demonstrated. The MPP PASCAL is compiled on the host DEC 8800 or VAX 11/780 and FORTH is interpreted while running on the MPP. The languages are also supported with highly efficient built in standard mathematical library routines such as sin, cos, max and min as well as data reordering routines based on sorting. A transportable simulator is also available for use by investigators on their own local VAX/VMS computers for simulating the MPP codes at the hardware level. Development of a parallel extension of 'C' is in progress. MPP PASCAL has been found to be comparably efficient on the MPP to serial compilers on standard architectures.

J.-P. A. L. MULLER (*Department of Photogrammetry and Surveying, University College London, U.K.*). Your examples all have one thing in common: the fact that the problem is content- and context-independent so that a SIMD architecture, such as the MPP, is excellent for processing the same algorithm on every pixel in the image scene. It would be useful to know whether mapping less well-behaved algorithms on the MPP has been looked at, and how relevant Dr Halem's examples are to future automated image understanding systems.

M. HALEM. A number of less well-behaved algorithms have been mapped and efficiently implemented on the SIMD architecture of the MPP. Fast Fourier transforms are implemented by assigning block sub-arrays of processors to perform FFTs on lines of an image as described in our SAR processing. In the graphic three-dimensional rendering of the scene in elevation mapping, sort computation which are global in terms of routing data to arbitrary processor were required and effected with a two-dimensional bitonic sort. In general, numerous scientists are mapping non-mesh oriented algorithms onto the MPP and other SIMD machines.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

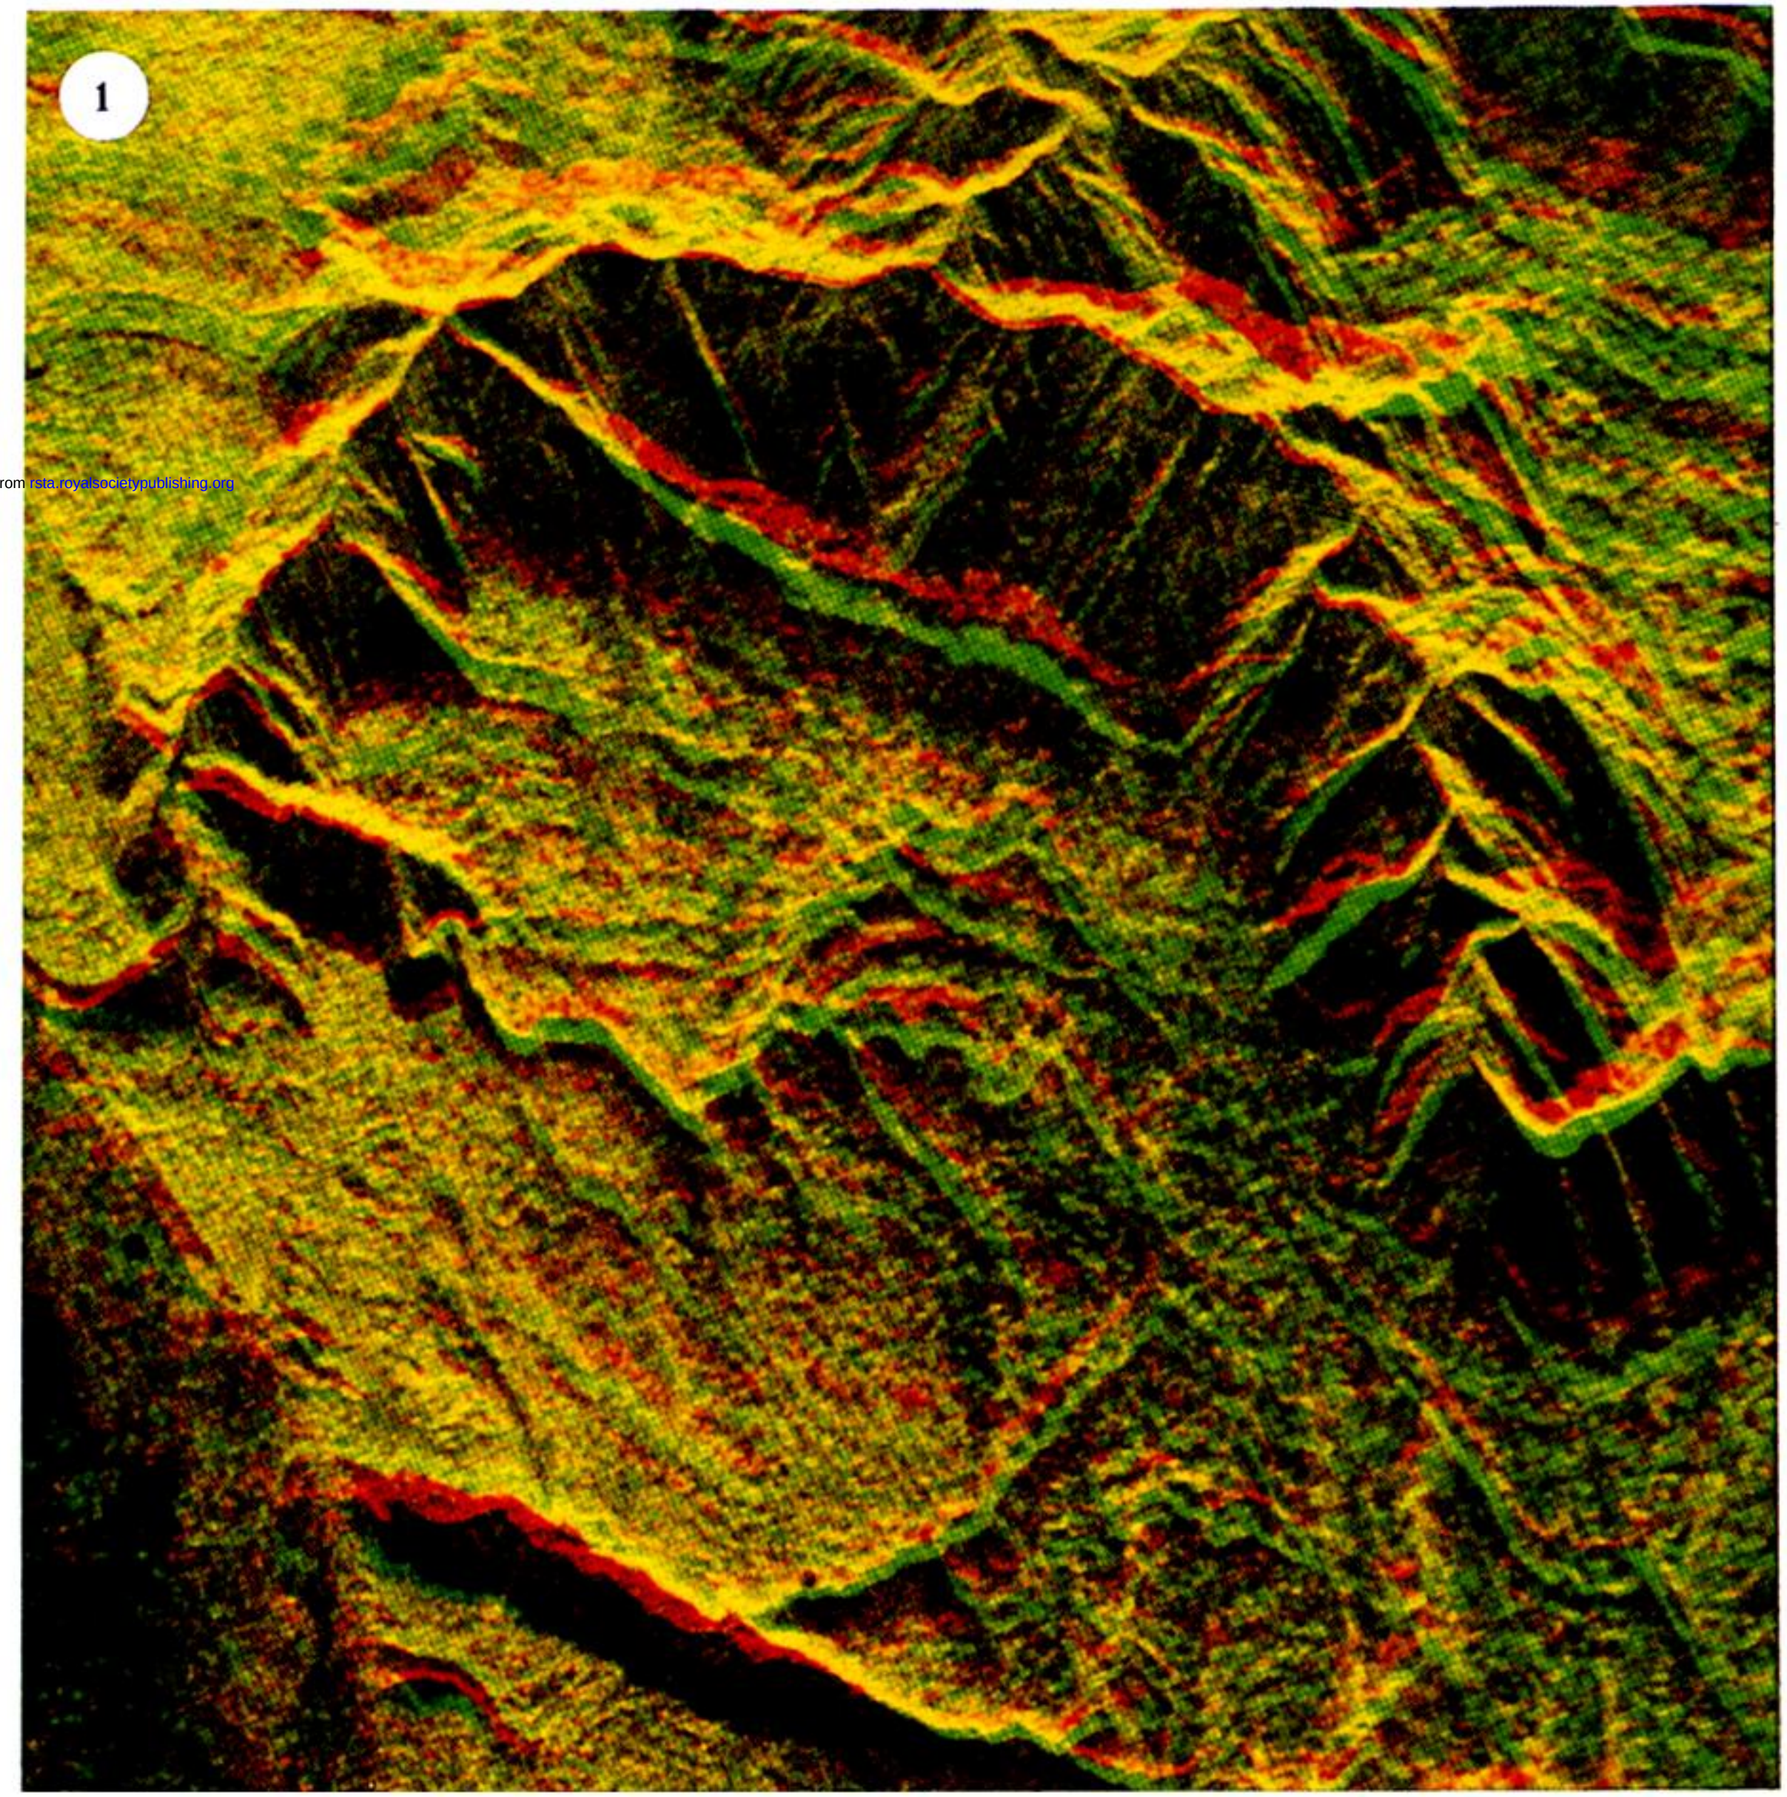

FIGURE 1. Pseudo-stereo overlay of SIR-B shuttle images. Red and green indicate maximum radar reflection values, respectively.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

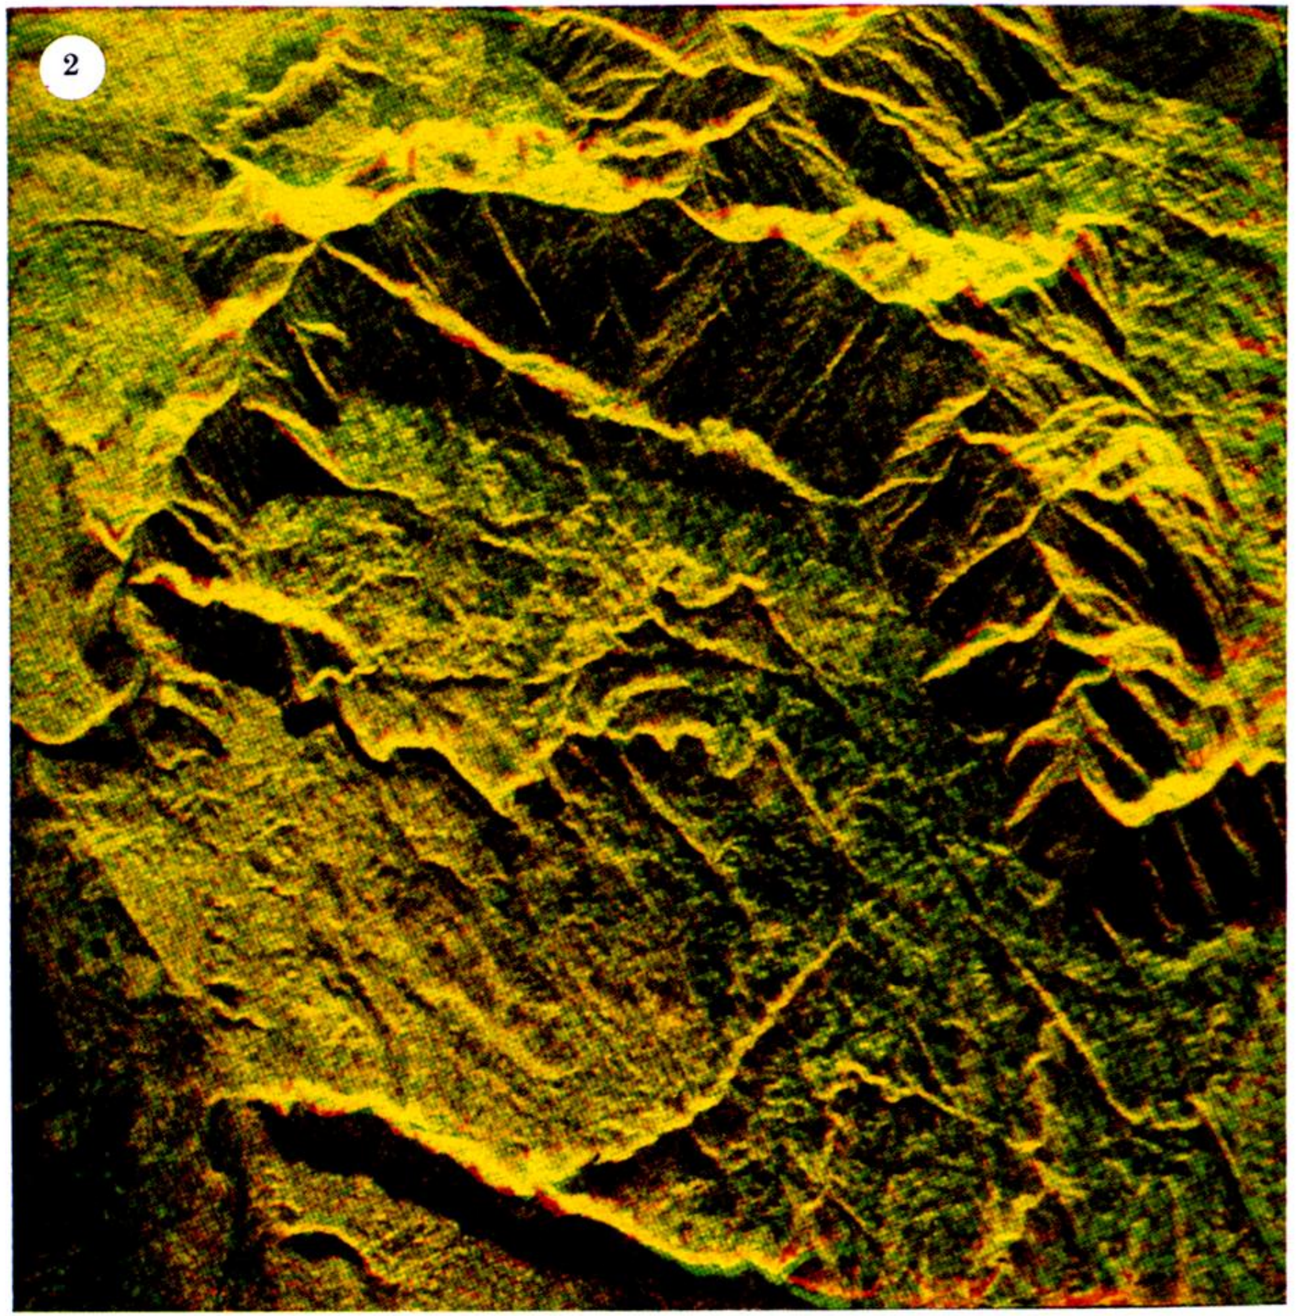

FIGURE 2. Overlay after application of two iterations matching algorithm. The blending of red and green into yellow indicates good agreement resulting from warping.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

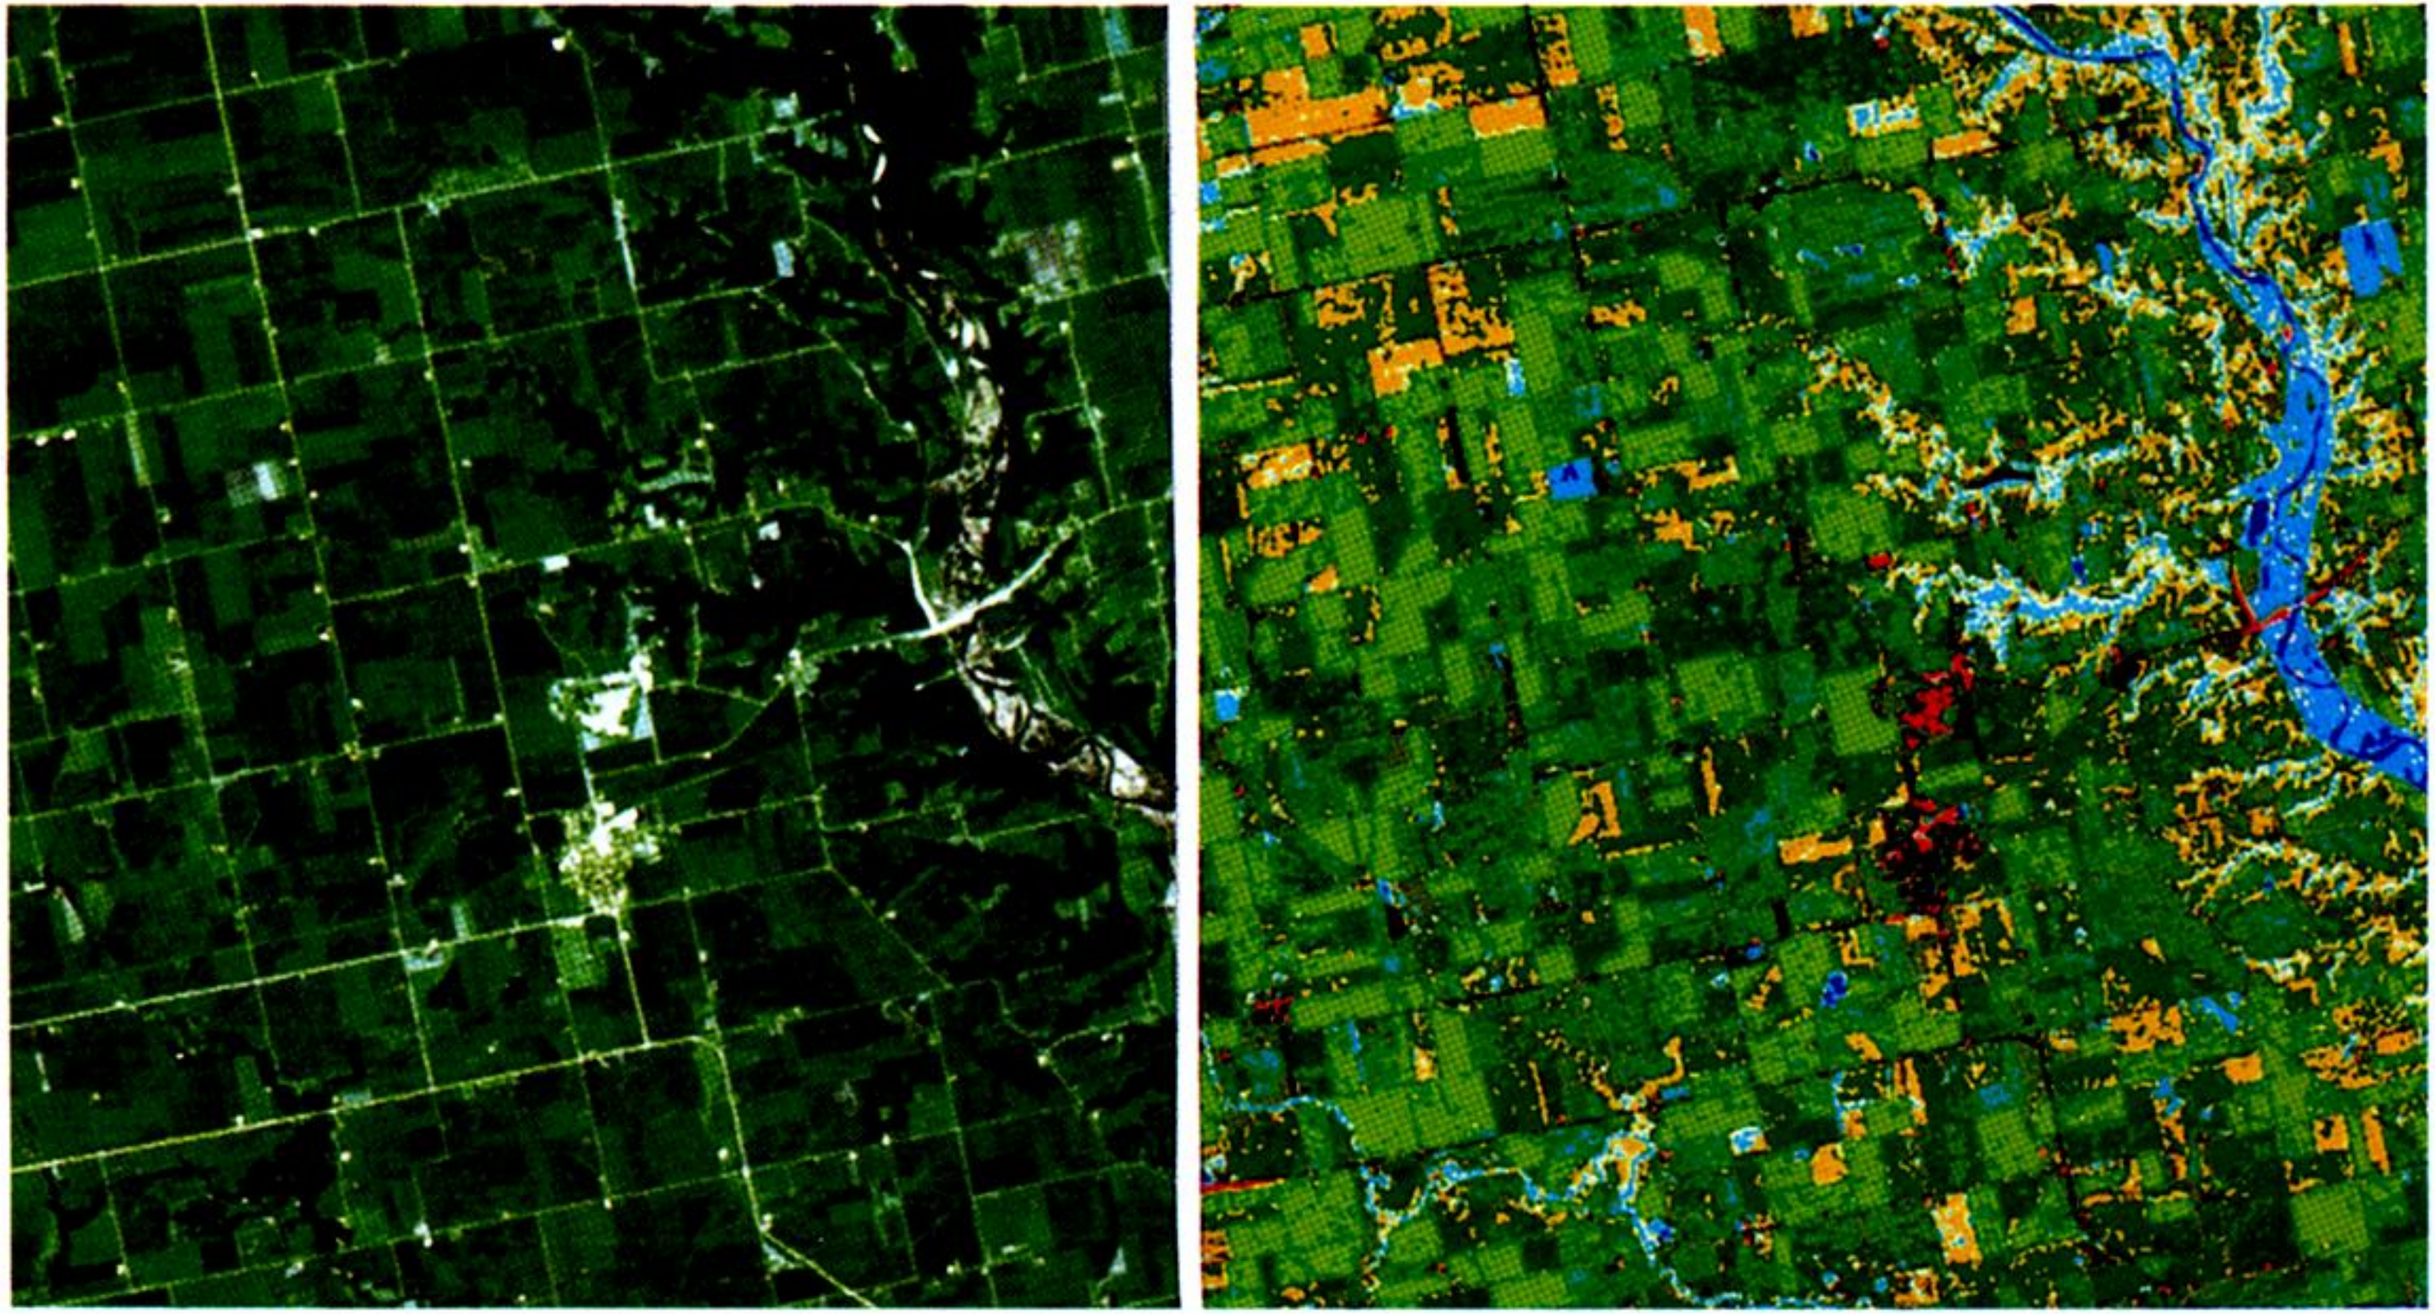

IGURE 6. Portions of a *Landsat* scene before (left) and after (right) the application of 100 iterations of unsupervised ISODATA classification algorithm.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

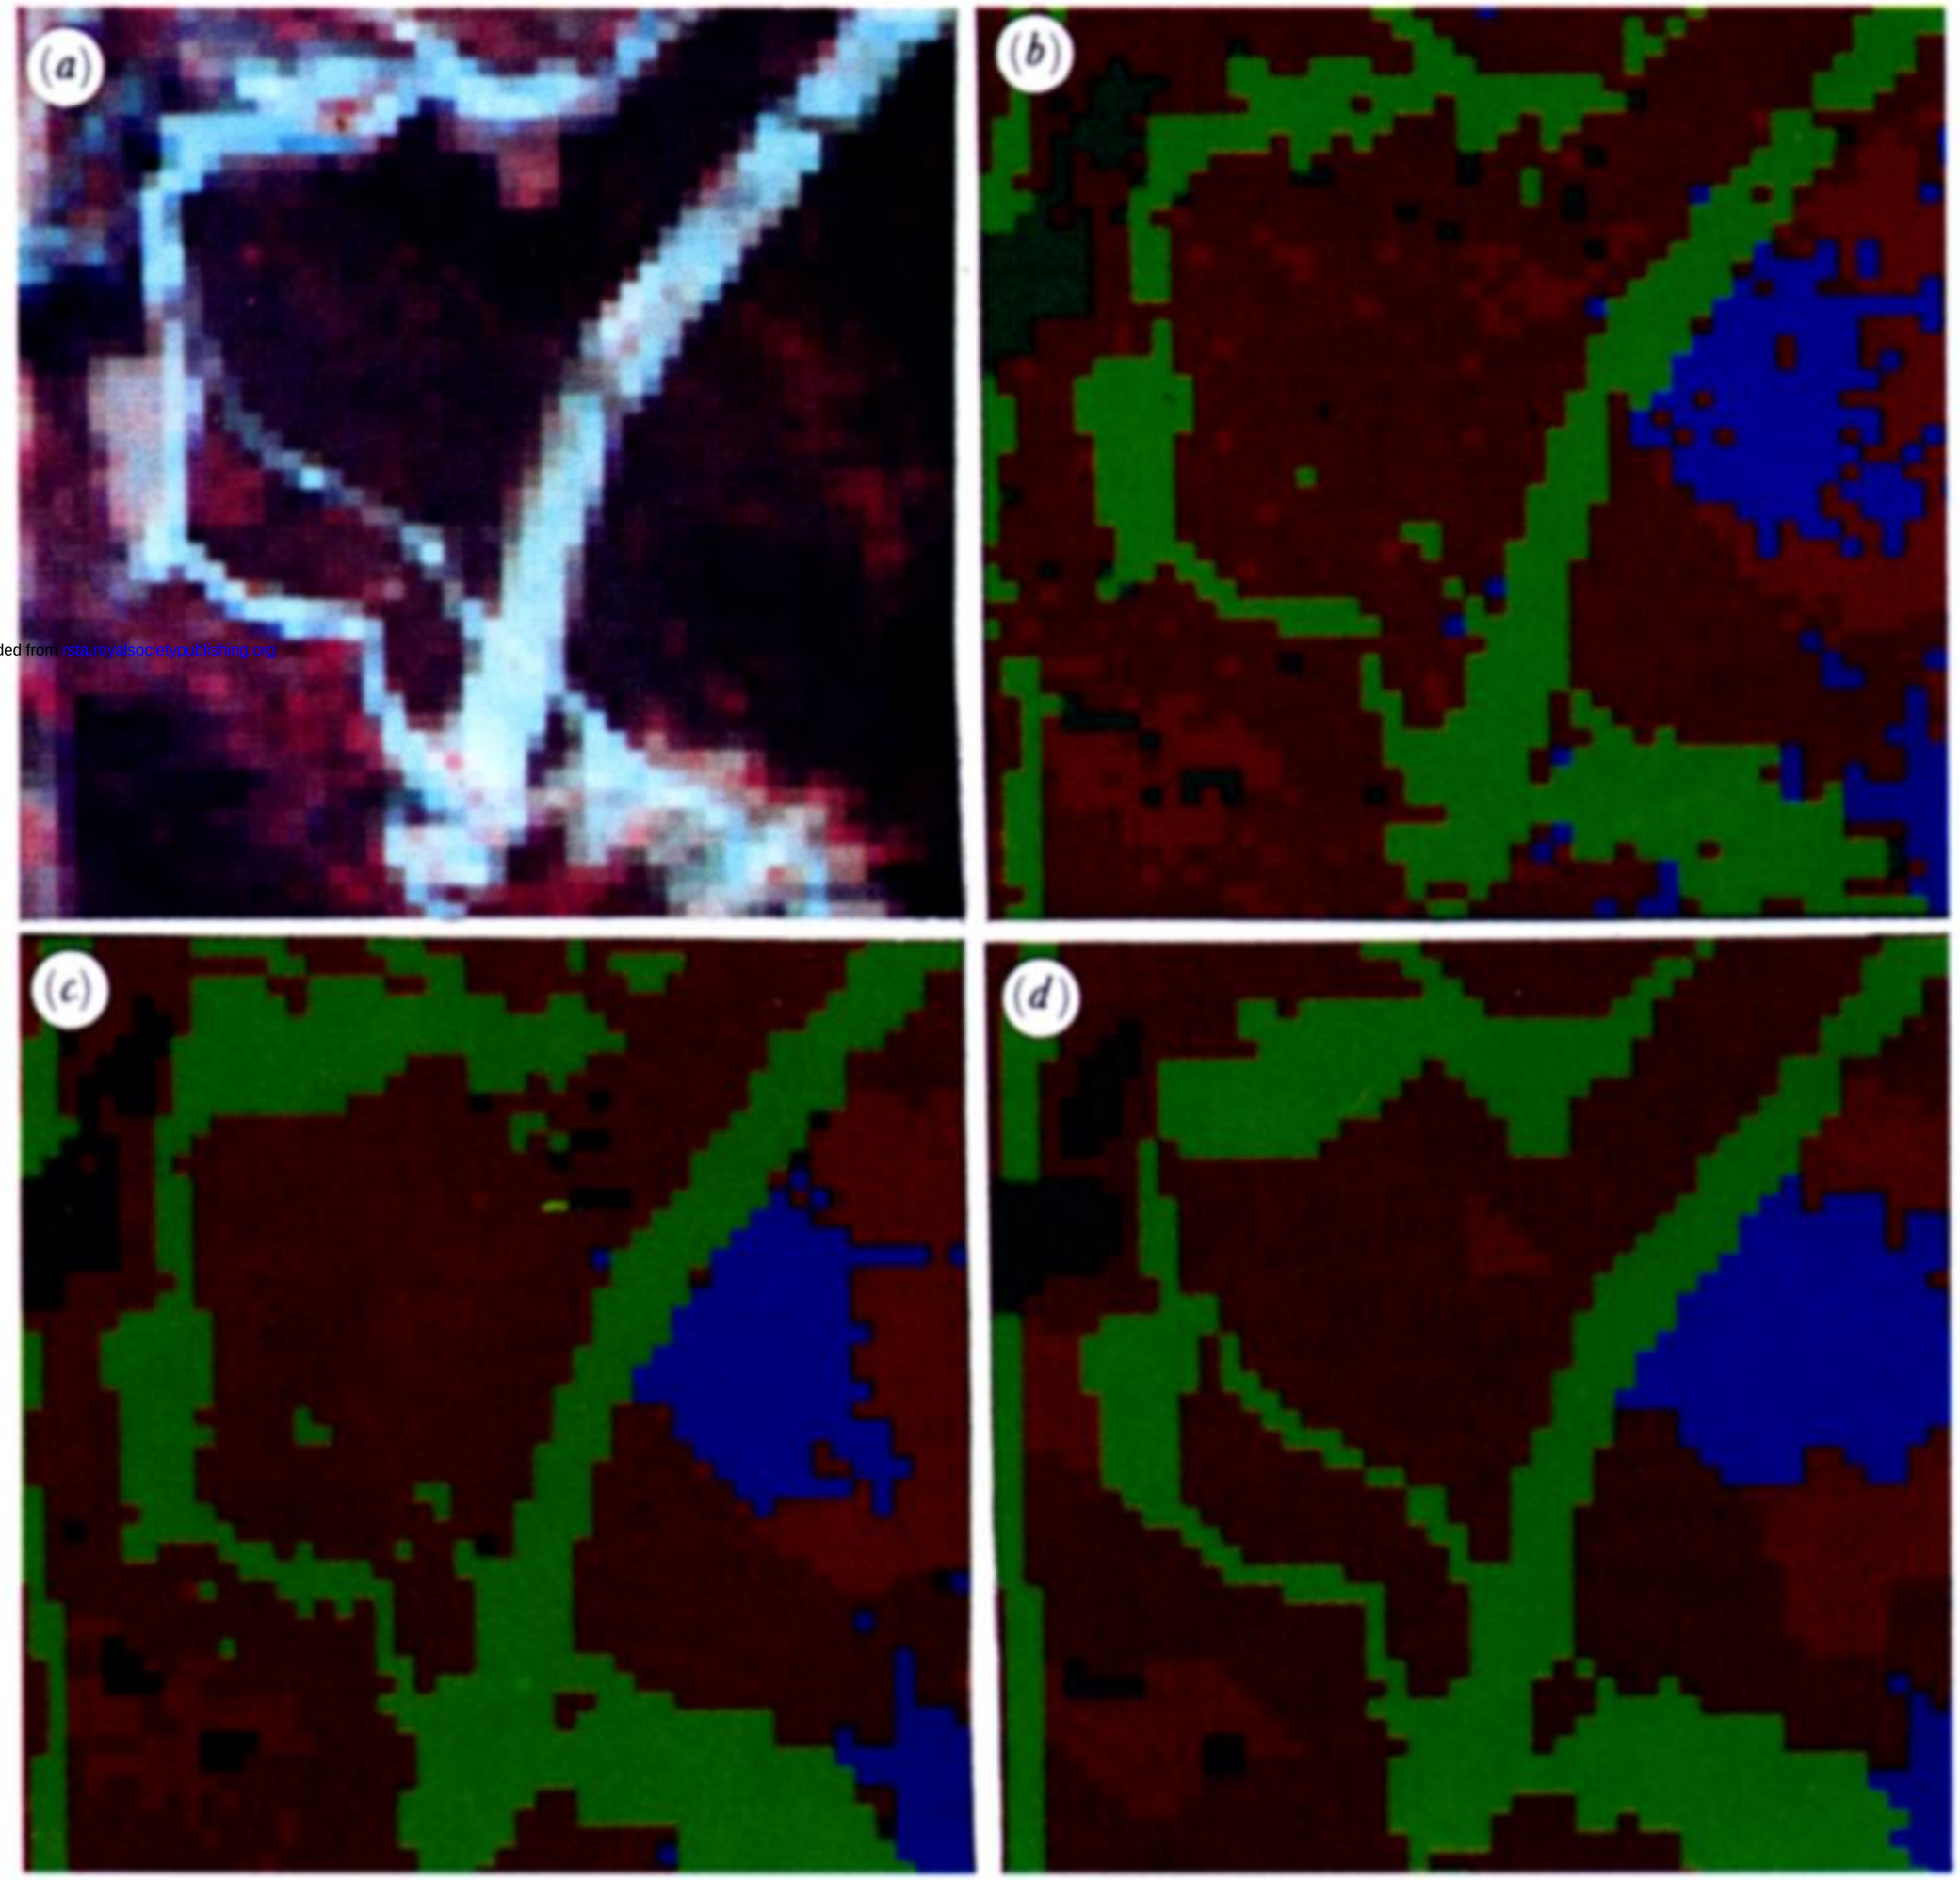

FIGURE 7. Three image representations of a *Landsat* scene: (*a*) unclassified; (*b*) after application of maximum-likelihood classification algorithm; (*c*) after application of five-point contextual extension of maximum-likelihood classifier; (*d*) ground-truth classification.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

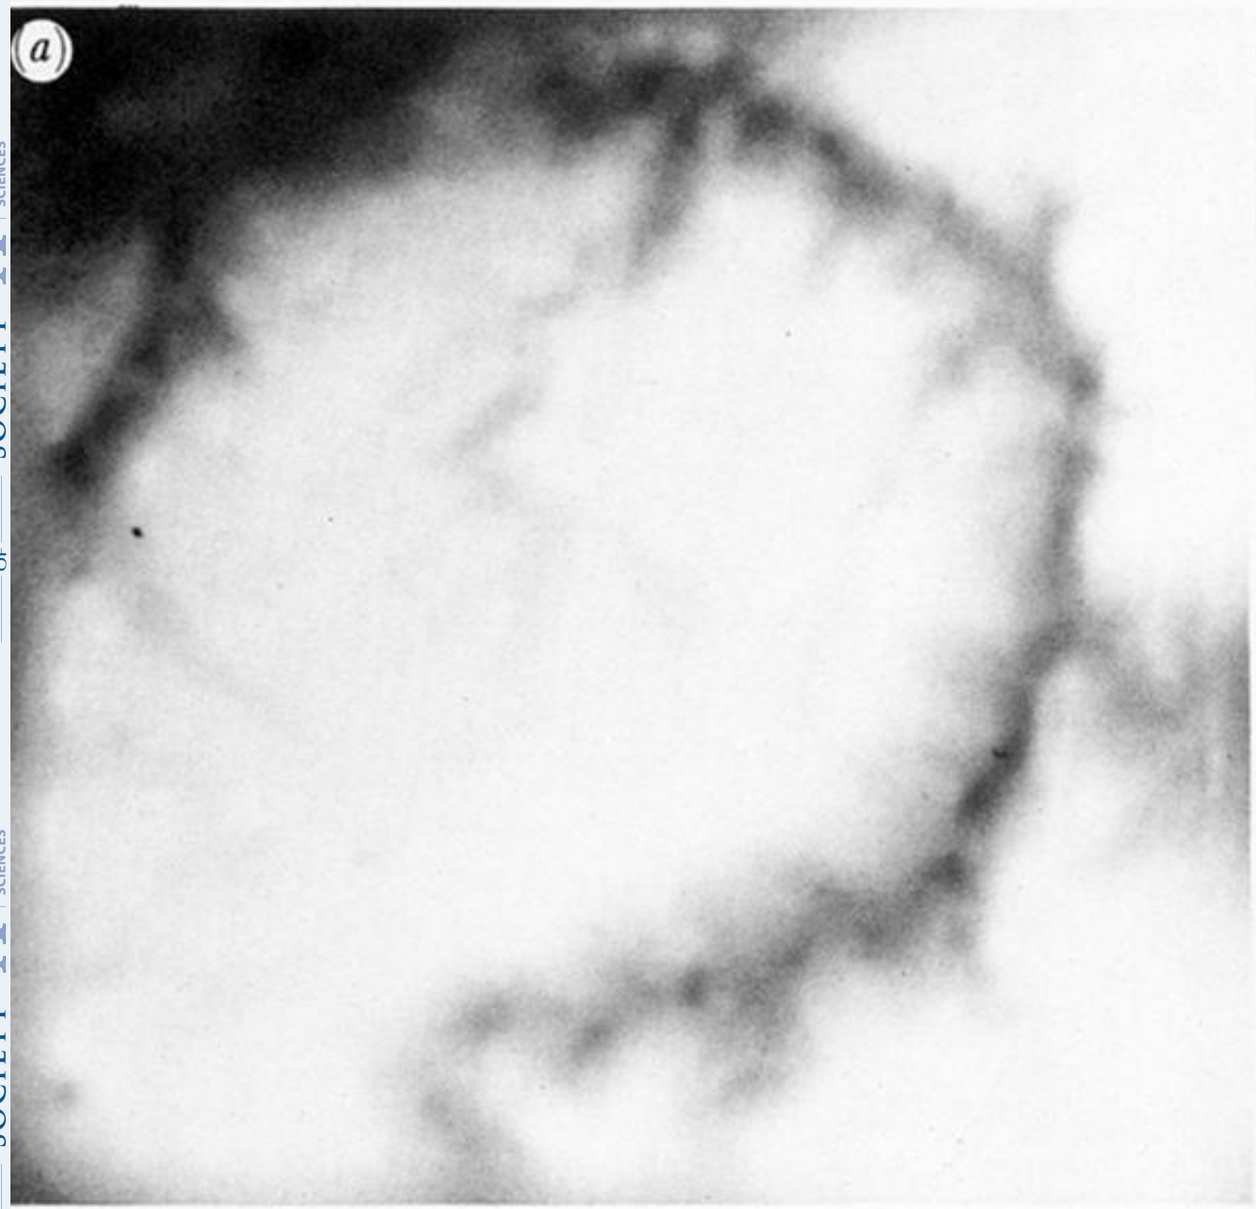

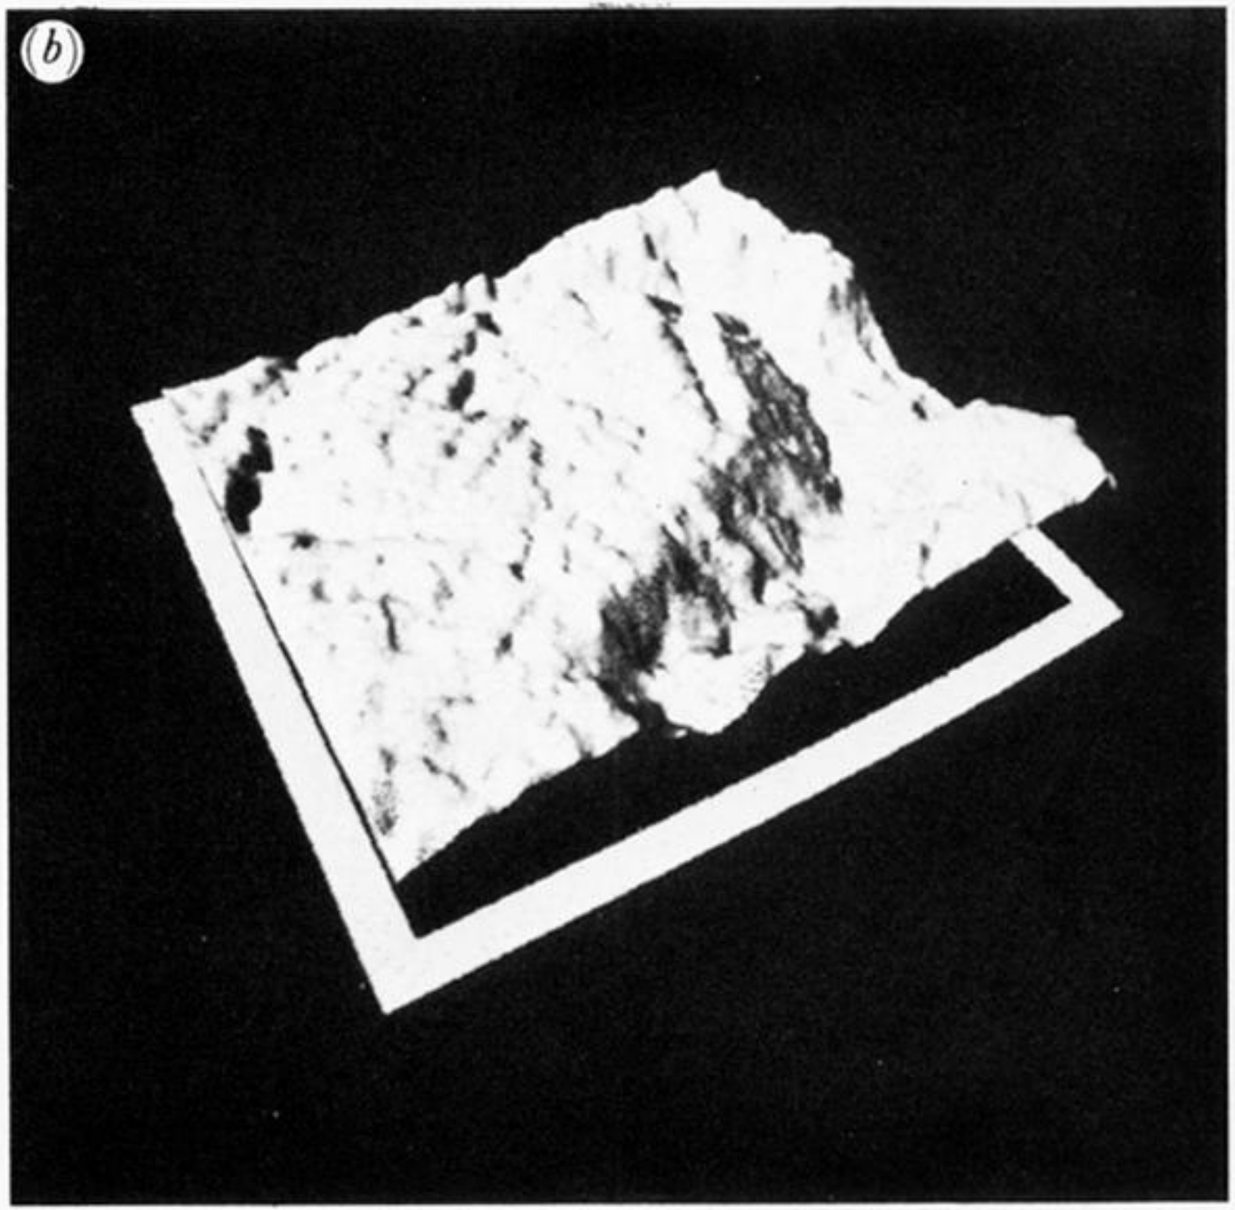

FIGURE 3. (*a*) *SIR-B* stereo disparity function after second iteration. Dark regions indicate low elevation and white regions indicate high elevation (*b*). Three-dimensional perspective rendering with stereo disparity function as height surface illuminated from a light source at an angle 42° from vertical.

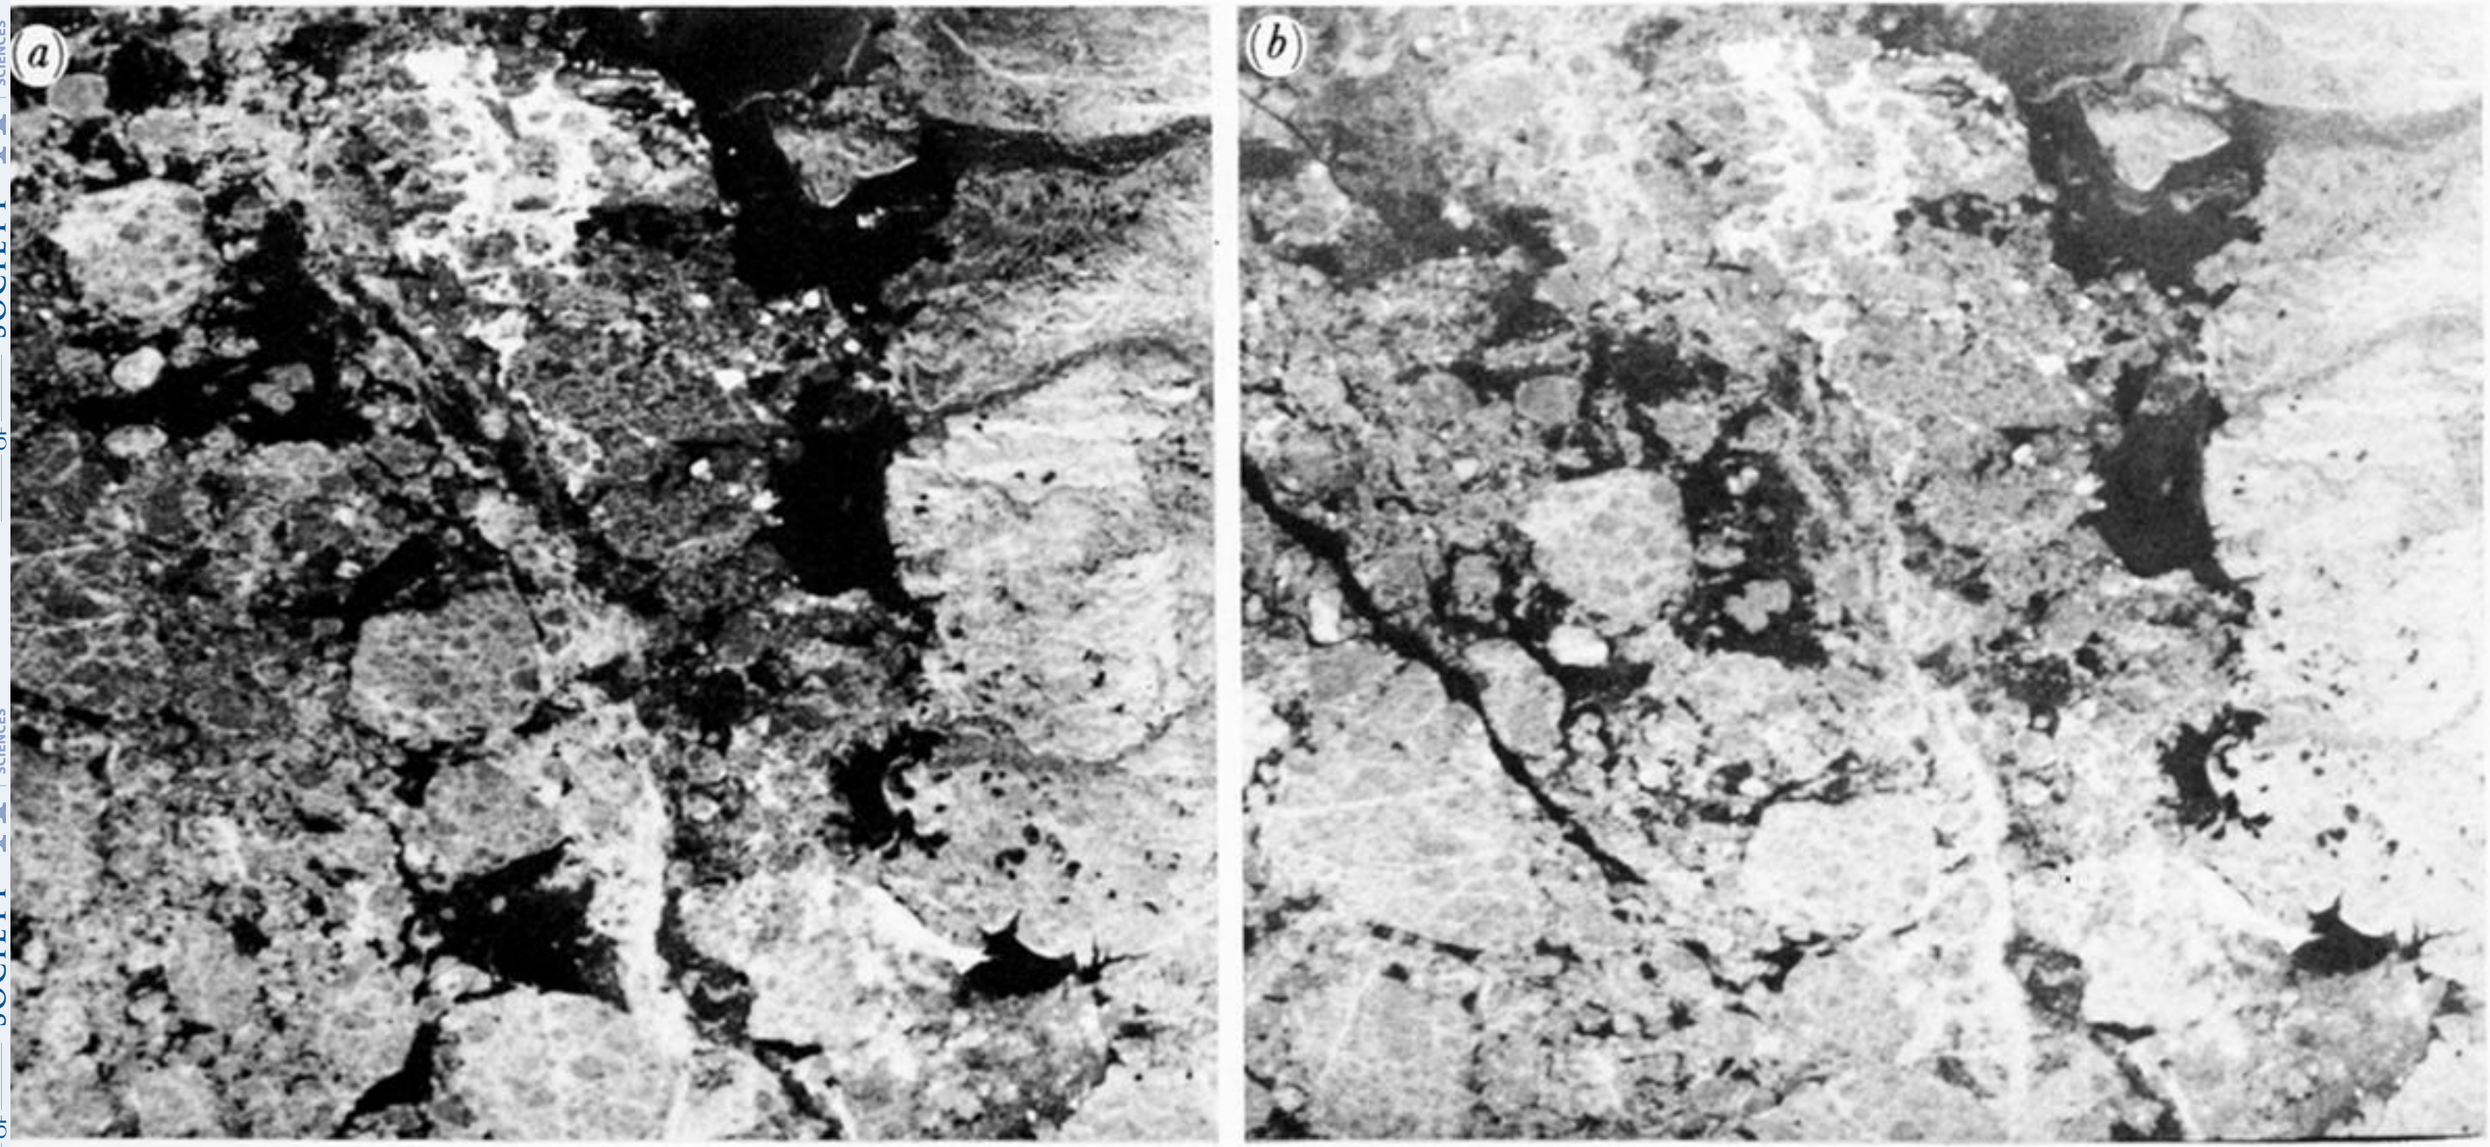

FIGURE 4. Near identical views of Bank Island, Alaska, taken by *SEASAT* synthetic aperture radar. (*a*) Image taken three days before image (*b*).

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

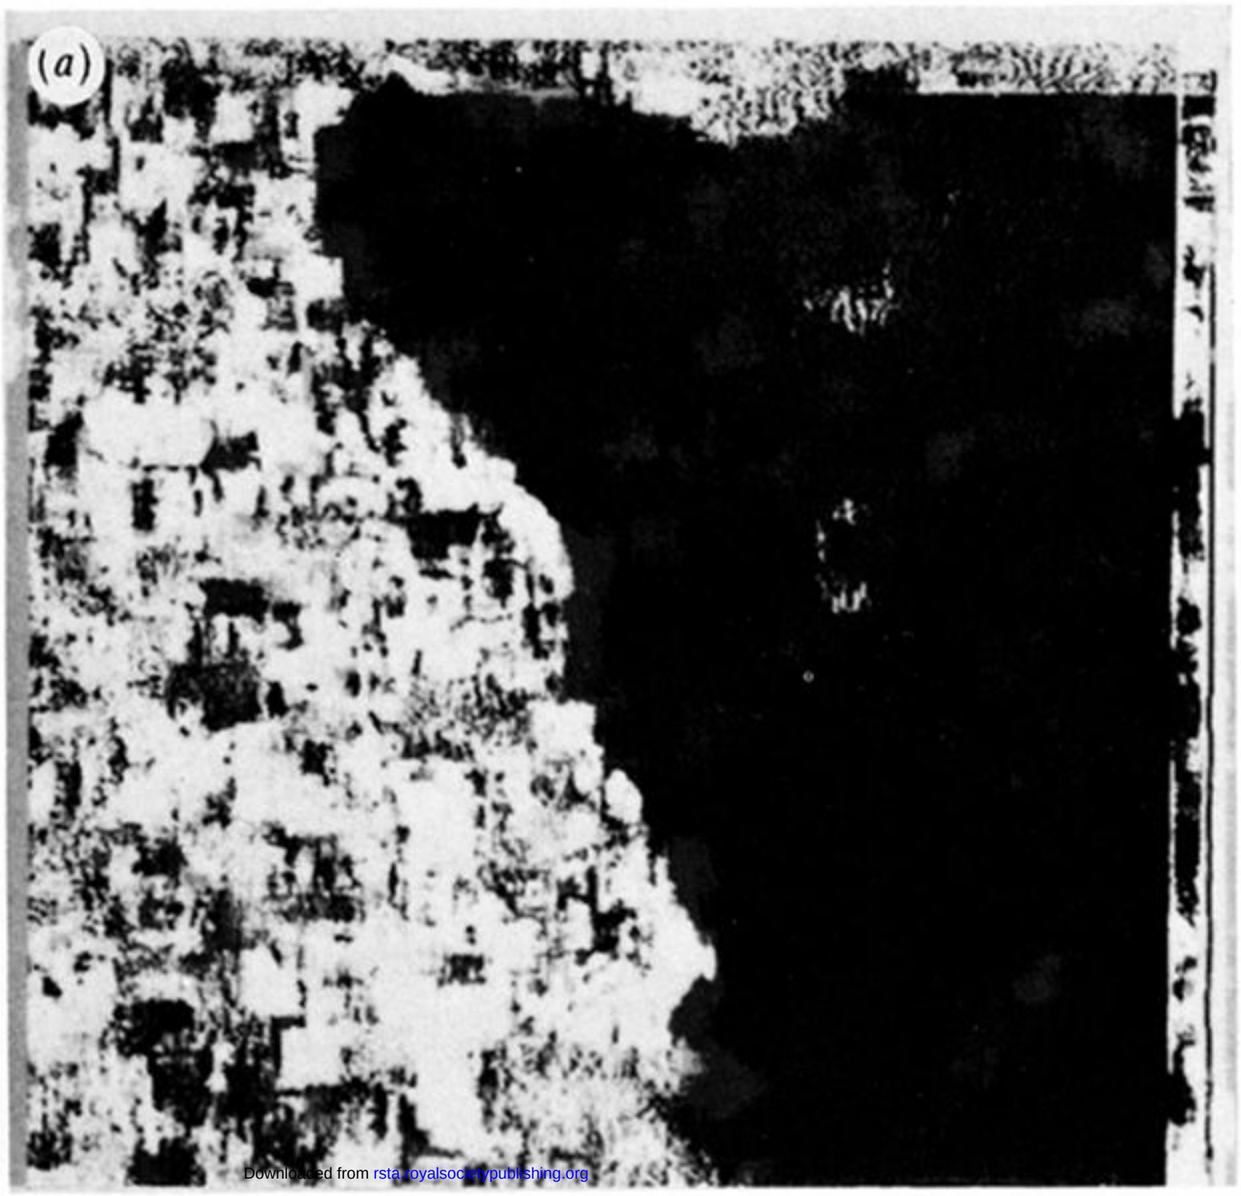

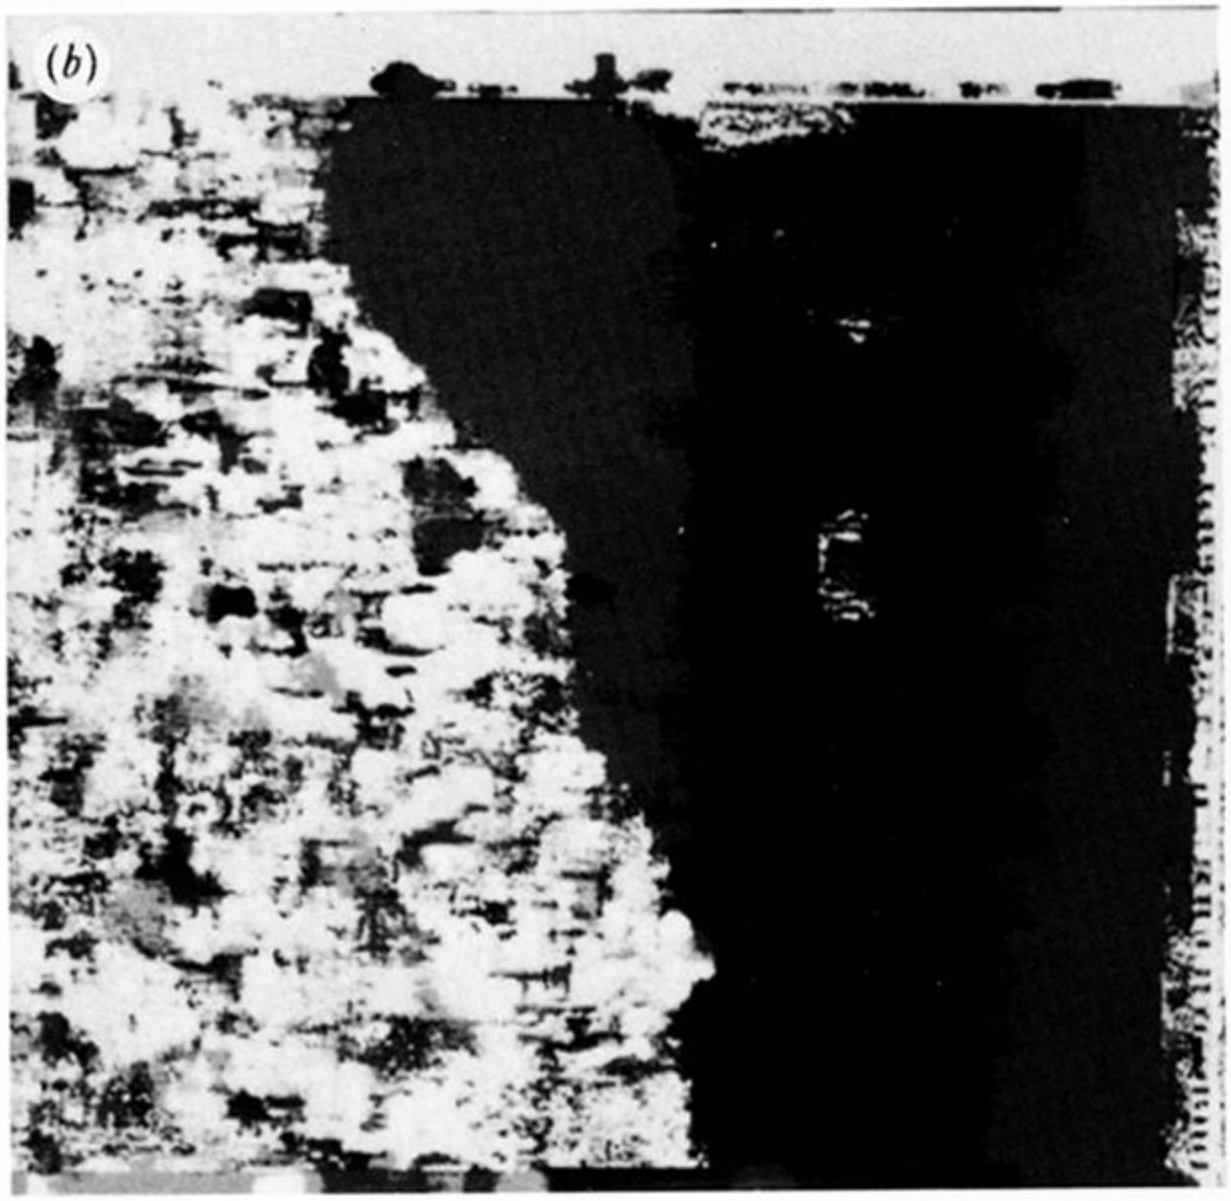

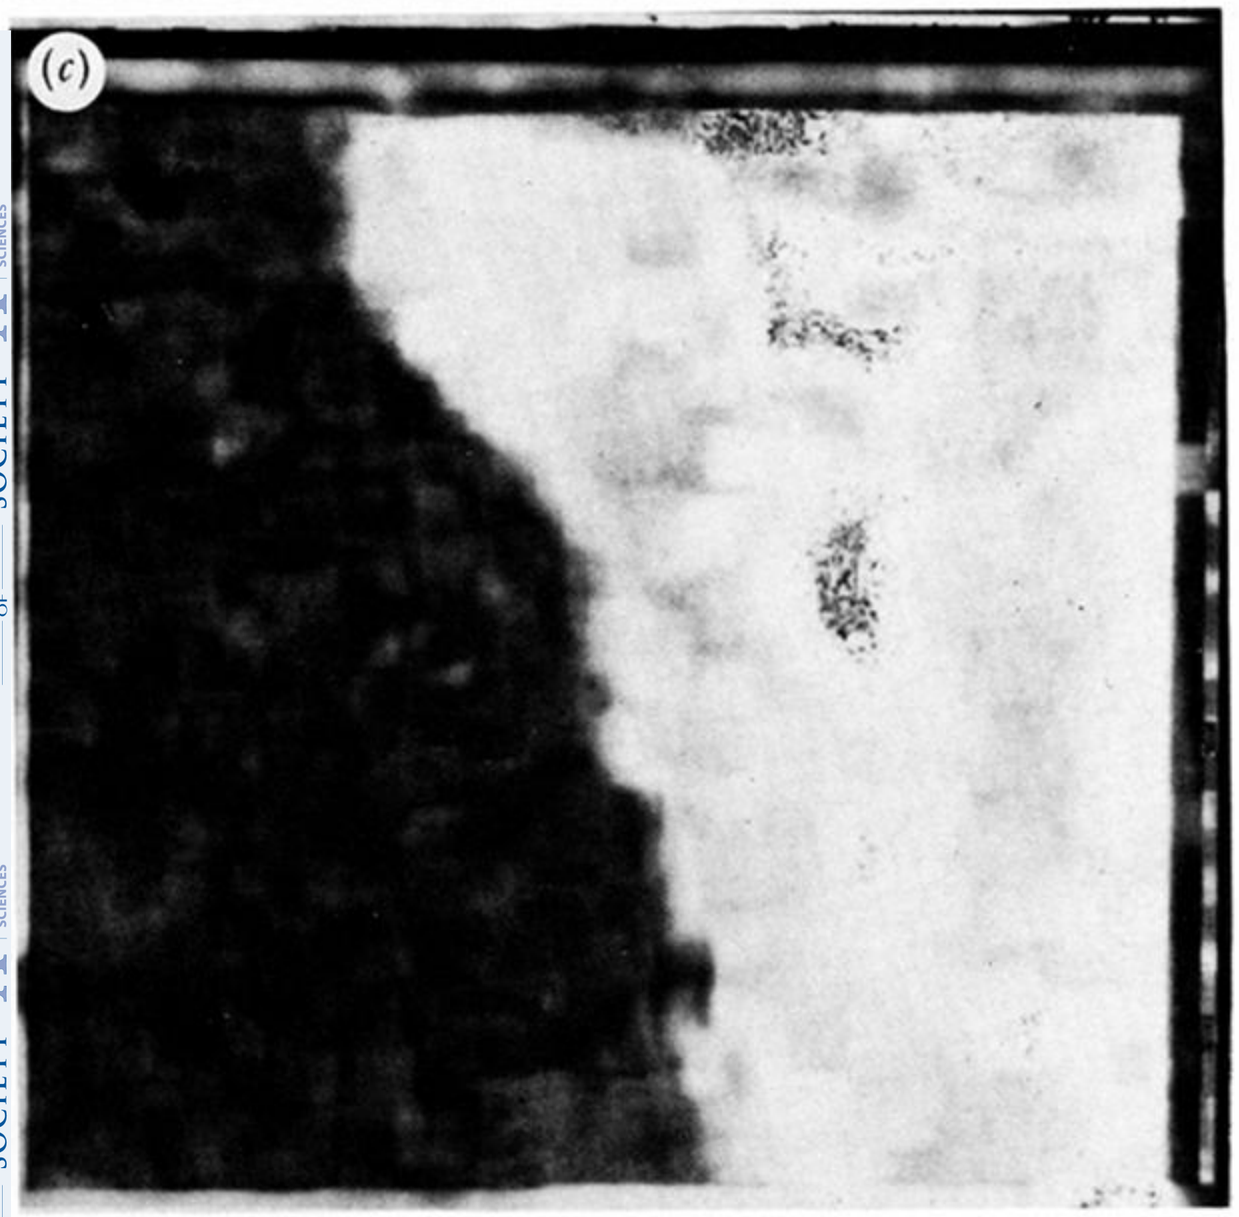

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

FIGURE 5. (*a*) Horizontal and (*b*) vertical translation functions, respectively. Large dark areas indicate little or no translation. (*c*) Match score values of displacement array. Brightness shows good matches within overlapping ice floes.