Reprinted from

Eighth International Symposium

Machine Processing of

Remotely Sensed Data

with special emphasis on

Crop Inventory and Monitoring

July 7-9, 1982

Proceedings

Purdue University
The Laboratory for Applications of Remote Sensing
West Lafayette, Indiana 47907 USA

LANDSAT IMAGE REGISTRATION FOR AGRICULTURAL APPLICATIONS

R.H. WOLFE, JR.

IBM Federal Systems Division
Houston, Texas

A.G. WACKER

University of Saskatchewan
Saskatoon, Saskatchewan

R.D. JUDAY

National Aeronautics and Space Administration/Johnson Space Center
Houston, Texas

T. KANEKO

IBM Communication Systems Div.
San Jose, California

I. ABSTRACT

An image registration system has been developed at the NASA Johnson Space Center (JSC) to spatially align multi-temporal Landsat acquisitions for use in agriculture and forestry research. Working in conjunction with the Master Data Processor (MDP) at the Goddard Space Flight Center, it functionally replaces the long-standing LACIE Registration Processor as JSC's data supplier. The system represents an expansion of the techniques developed for the MDP and LACIE Registration Processor, and it utilizes the experience gained in an IBM/JSC effort evaluating the performance of the latter. These techniques are discussed in detail. Several tests were developed to evaluate the registration performance of the system. The results indicate that 1/15-pixel accuracy (about 4m for Landsat MSS) is achievable in ideal circumstances, sub-pixel accuracy (often to 0.2 pixel or better) was attained on a representative set of U. S. acquisitions, and a success rate commensurate with the LACIE Registration Processor was realized. The system has been employed in a production mode on U. S. and foreign data, and a performance similar to the earlier tests has been noted.

II. INTRODUCTION

Image registration forms a major link in the chain of processes through which Landsat MSS data flow toward agriculture and forestry interpretation. That link is necessitated by the merging of images acquired at different times, different orbital paths, and different Landsat sensors. Manual registration is impractical for the volume of image data investigated in such large-scale experiments as LACIE and AgRISTARS, and it is necessary to resort to automated techniques. Such an approach was developed during LACIE, and the LACIE Registration Processor was established at NASA's Goddard Space Flight Center (GSFC) to provide the Johnson Space Center (JSC) with multitemporal acquisitions. That system edge-correlated each Landsat acquisition of a given LACIE ground segment with a base or reference acquisition and performed nearest-neighbor resampling to fulfill as pure translation the required realignment. The resulting alignments were shown by Kaneko[2,3] to fall within the specified 1-pixel accuracy, but the presence of residual affine distortions (scale, rotation, and shear) was noted. With heavily-clouded scenes removed, the system realized a success rate of about 2/3, having failed to register the remainder because of false correlations or other anomalies. The LACIE Registration Processor remained the workhorse for registration of JSC-bound data until as recently as the 1980-81 growth year.

Meanwhile, the Master Data Processor (MDP)[4,5] was developed at GSFC to geometrically correct whole 185x170-km Landsat frames. The registration process employed by MDP geometric correction performs cross-correlation of scattered image patches or control points, in contrast to the full-image correlation used in the LACIE Registration Processor. The control points are essentially hand-picked according to their distinguishability and expected stability (see Niblack's description[6]). Under optimal conditions, the MDP has demonstrated subpixel (10 to 20-meter) accuracy, and in such cases the imagery would require no further registration at JSC or, at most, application of a prescribed resampling to render the imagery in the desired coordinate system. Often, however, the conditions are not so favorable because of the sparseness of usable control points caused by cloud cover or correlation failures attributable to temporal instability. At the time the JSC Registration Processor was installed, MDP control points were not yet available for

a substantial portion of the world. Sparseness or complete absence of control points has resulted in residual MDP registration errors often ranging from a few pixels to 45 pixels (2.5 km) or more. It became obvious that a residual, local registration could nevertheless be accomplished on the MDP products at JSC by selecting a concentration of control points in the sample segment of interest. In that manner, the JSC residual registration would act as a vernier correction to the MDP process.

The JSC system was conceived to facilitate working toward an eventual goal of consistent 0.2-pixel registration accuracy, while also complementing the MDP by performing residual registration. The system is based on a merging of the techniques employed in the MDP and LACIE Registration Processor as augmented by the experience gained in Kaneko's evaluation of the latter. A number of enhancements are included, as are apparent in the following two sections. These sections describe respectively the basic structure of the JSC Registration Processor and the detailed step-by-step registration process. Following them, tests which have been developed and used for evaluating the registration performance are described, and the performance results are reported.

III. STRUCTURE OF THE JSC REGISTRATION PROCESSOR

Our design was driven by the desire to have a system which could be used immediately in a production mode, while serving as a research tool for implementing and testing new techniques. The design was set up to facilitate expansion to other sensors besides the Landsat MSS and to allow easy replacement of algorithms. To satisfy these requirements, the system was built so that the processing functions are modularized, that is self-contained with straightforward data interfaces. Thus, each process, to be described below, is contained as one FORTRAN subroutine whose control data interface is solely through the calling arguments. Image data are handled by each subroutine by calling a special set of data base access routines which were designed to facilitate eventual transfer to a standard data base system.

A number of alternative choices for processing were built into the JSC Registration System. For example, two types of normalized cross-correlation are available at user discretion. Multiple choices were accommodated because it was uncertain at the design stage which ones would prove most effective. Similarly, many of the constants had not yet been settled upon, and it was necessary to keep them under user control. Thus, all control parameters and most constants can be changed by the user, so that a fairly flexible reconfiguration capability is maintained. Much of the experimentation done on the system has involved trying different values for control constants and different existing algorithm alternatives (as described later).

At present, the image frames delivered by the MDP are too big to handle by the JSC Registration Processor and a smaller area of interest (632x632 pixels, or less) is first extracted by the Landsat Image Verification and Extraction System (LIVES)[7] and delivered in the same basic (MDP) format. LIVES also adds some data of local interest, such as the geodetic coordinates of the area of interest, its intensity range, etc. The image data occur in one of two forms called P or A-data (or P or A-tapes). P-data consist of the resampled image resulting from the MDP geometric correction process. That is, the pixels in a P-format image represent nominally a uniform spacing in a standard map projection with geometric distortion caused by sensor effects and satellite pointing errors removed (to within the precision possible, depending on the number and quality of control points used). A-data consist of the original pixels, radiometrically corrected, but not resampled, but data are included to allow the recipient to effect the same resampling as the MDP would have performed. That is, the additional data, in the form of geometric transformation tables and sensor distortion constants, facilitate conversion of the A-data to the corresponding P-data. LIVES can extract areas of interest from frames of A or P-data, and those areas are passed along unchanged. The registration processor accepts either type, initially executing the required resampling in the case of A-data.

The processes to be executed in performing registration are shown in Figure 1. The progression of these functions is basically of the form Geometric Correction-Matching-Geometric Correction. Two images--the reference and registrant-- are handled, the latter being conformed to match the former. The first geometric correction is performed possibly on both images to render them nominally in the same coordinate system to simplify the matching process. We have elected to match the images by performing cross-correlation in translation only; we accommodate higher-order distortions or mismatches by correlating on sub-image patches distributed throughout the area of interest. The patches must be small

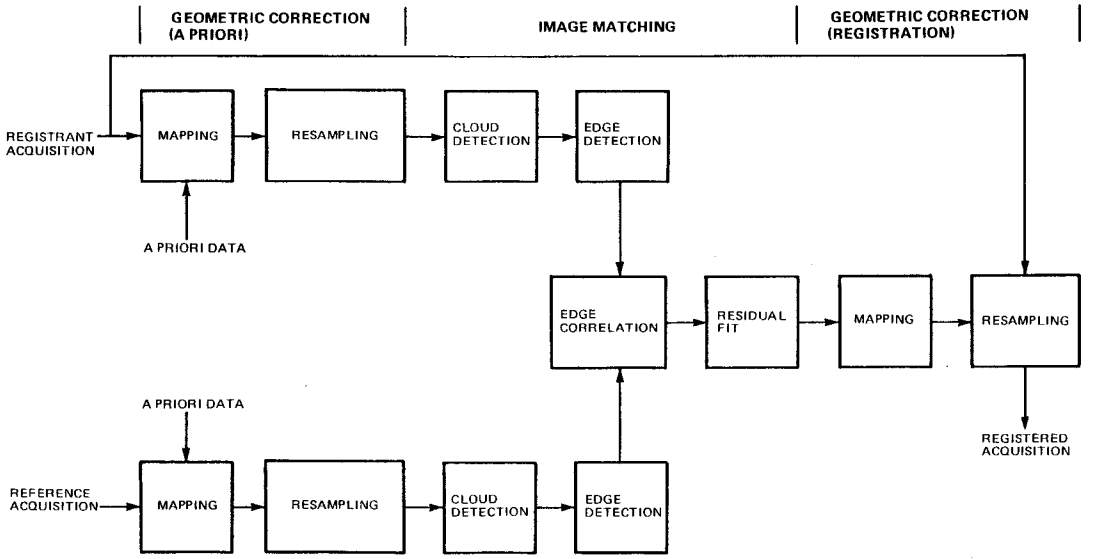

Figure 1. Registration processing flow.

enough that only translational error need be accounted for over their dimensions, yet large enough to yield cross-correlation measures of sufficient singal-to-noise ratio (SNR). The initial geometric correction facilitates this compromise by removing any previously known distortions by use of available a priori data. Examples of this process include conversion to compatible map projections (e.g., sidelap rotation or Hotine Oblique Mercator to Universal Transverse Mercator projections) and the A-to-P conversion required for A-data.

Next, the images are matched to estimate the geometric distortion of the registrant with respect to the reference. Finally, the estimated distortion function is incorporated in the overall geometric transformation on the registrant, and the second geometric correction is performed on the original registrant data (A or P) to produce the final registered image. Geometric correction is divided into Mapping and Resampling functions in Figure 1. The former generates for the latter the appropriate geometric transformation function based the desired map projections, MDP-supplied A-to-P conversion data, and the residual distortion function estimated in image matching. Matching is divided into a number of functions as shown in Figure 1. These functions form the basis of the registration system modularization.

IV. REGISTRATION STEPS

A. MAPPING

The geometric coordinate transformation used for pixel resampling is constructed in the same manner as the MDP, viz. the bivarate mapping function is constructed as a pair of tables from which the input space position of a desired output space sample can be found adequately by bilinear interpolation. The transformation is separated, as illustrated in Figure 2, into horizontal (sample-to-sample) and vertical (line-to-line) shear distortion transformations to allow decoupling of the resampling convolution in those directions. The "hybrid" space occuring betwen the vertical and horizontal shear transformations shares its horizontal coordinate with output space and its vertical coordinate with input space. To build these tables, a uniform grid is established in output space; it is transformed through a residual distortion polynominal if applicable, thence transformed from the output map projection to the input map projection (Hotine Oblique Mercator and Universal Transverse Mercator are presently accommodated). The characteristics of the map projection depend on the particular World Reference System (WRS) format center coordinates associated with the image. Different WRS points can be associated with the input and output spaces by virtue of different satellite paths for the reference and registrant acquisitions; this difference manifests itself in sidelap regions in the form of a sidelap rotation.

If the input image is A-data, the mapped grid is further transformed through the MDP-provided grid to yield finally the output space grid in sensor space coordinates. This last transformation uses third order interpolation, rather than bilinear, to preserve the precision inherent in the MDP tables. The vertical coordinate values of the transformed grid constitute directly the tabular vertical shear transformation to hybrid space; the horizontal counterpart (hybrid-to-input space) is attained by interpolating among the horizontal coordinate values of the grid in input space to produce values aligned with input space horizontal lines. Again, third order interpolation is used to preserve grid precision. Our reward for all this complexity comes in the form of an effective consolidation of all transformations, which eliminates the need to successively resample the image for each component of the transformation. For example, the second mapping includes the residual distortion function resulting from image matching, and our procedure folds that portion into the overall mapping so that the final resampling utilizes the original image rather than resorting to double resampling.

B. RESAMPLING

The tabular mapping function is utilized in the same manner as the MDP to effect separate horizontal (along scan line) and vertical (along sample column)

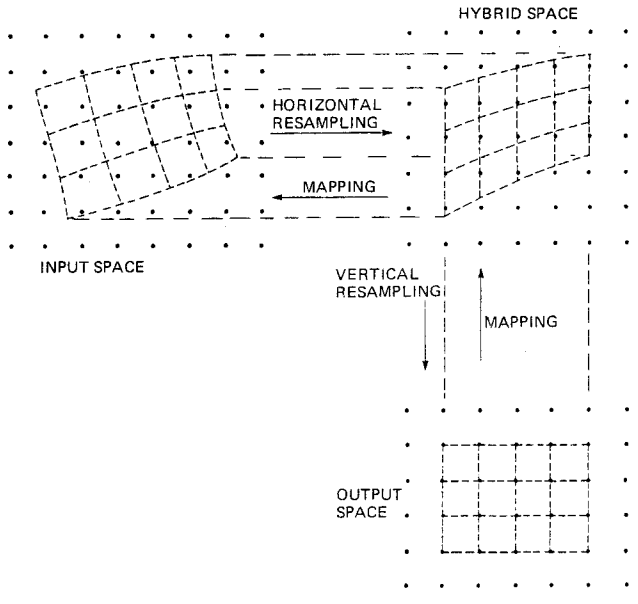

Figure 2. Separation of the geometric transformation into horizontal and vertical shear components.

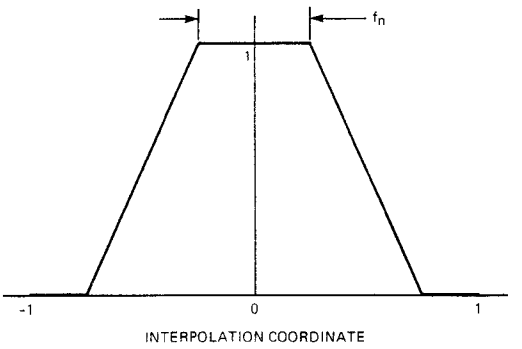

Figure 3. Convolving function for a mixture of linear interpolation and nearest-neighbor resampling. The parameter f_N governs the relative size of the nearest neighbor region.

resamplings. This separation saves considerable computation time but is predicated on the assumption of a separable convolution filter. Two convolution filters are provided, one being the TRW cubic convolution[8] used by the MDP. Alternatively, a combination nearest neighbor/linear interpolation is provided with a parameter to allow variable mixing of the two. The corresponding convolving function is presented in Figure 3. Nearest neighbor resampling occurs when the interpolated sample is within a relative proximity $f_N/2$ to an original sample; otherwise, linear interpolation applies. The parameter f_N can be varied from 0 to 1 to effect pure linear, a combination, or pure nearest neighbor resampling, repectively. The horizontal and vertical resampling are decoupled, so that one may specify, for example, cubic convolution in the scan direction and linear interpolation from line-to-line.

Several MSS effects, including relative detector misplacement, differential sampling delay, earth rotation differential skew (staircasing), and varying sample spacing, create geometric distortions that vary from line to line. Since these distortions cannot readily be accommodated by the tabular transformation functions, their compensation is built directly into horizontal resampling as the so-called high-frequency correction. The correction need only be applied to A-data; the MDP has already performed it when creating P-data. The high-frequency correction parameters are provided in the A-data.

C. CLOUD DETECTION

The ERIM SCREEN algorithm is used to generate a map of cloud, cloud shadow, and data radiometrically inconsistent with the expected range of earth data (i.e., garbled data). ERIM's method requires the data to be in an absolute radiometric scale; therefore, differences in radiometric adjustments or compensations between current MDP data and ERIM's LACIE-era data (supplied by the LACIE Registration Processor) must be accounted for. The required transformation thus far has eluded accurate specification, and an approximate recalibration has been estimated[10] by comparing recent data processed by both systems. The recalibration is applied before executing the SCREEN algorithm. The resulting cloud/shadow/garble map is used to eliminate such areas from image matching.

D. EDGE DETECTION

It is desirable to match edge features in agricutural scenes because of the great radiometric variability from season to season. Our method resembles that of the LACIE Registration Processor but is supplemented with the cloud map. An absolute gradient is evaluated at each pixel, or it is marked as ineligible if associated with cloud, shadow, or garble. The absolute gradient, a non-linear filter, is computed as shown in Figure 4. The weights are included to allow for different scaling in the line and column direction and for treating the diagonal values differently. The ad hoc value of $2^{-1/2}$ has been used for the diagonal values, so far. Edge detection is performed separately on MSS bands 5 and 7, and the two binary edge maps are united (OR-ed) together to form a composite in the same manner as in the LACIE Registration Processor.

E. EDGE CORRELATION

The reference and registrant binary edge images are matched by normalized cross-correlation. A given reference image sub-area or patch is overlaid on a corresponding registrant search area (patch plus border), and the number n_c of coincidences of edge pixels is tabulated for each translational offset. Two options are provided for normalization; viz., the normalized cross-correlation N is found from

$$N = \begin{cases} n_c/n_s & \text{(Template Matching)} \\ n_c/\sqrt{n_p n_s} & \text{(Classical)} \end{cases}$$

where n_s is the number of registrant search area edge pixels within the reference patch overlay and n_p is the number of edge pixels in the reference patch. The first option was used in the LACIE Registration Processor and effectively treats the reference edge map as a known template; the second option, which more closely resembles classical cross-correlation, has been used generally in the JSC Registration Processor.

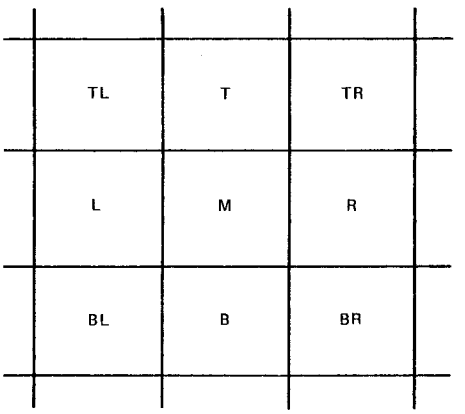

Figure 4. Pixel template for computing the edge gradient. The gradient at Pixel M is $w_{10}|L-R| + w_{01}|T-B| + w_{11}|TL-BR| + w_{11}|TR-BL|$, the w's being weights.

Often the input registrant is so greately misaligned that an efficient correlation technique is necessary to reduce the basic misregistration to a few pixels. The LACIE Registration Processor full-image correlation method is used; viz., the mean and standard deviation (S.D.) of correlations at every 32nd (fourth line by eighth sample) offset are computed to establish a rejection threshold at mean plus thrice S.D.; an approximate correlation, determined from one-ninth of the data (every third line by every third sample), is compared with that threshold; and the full correlation is evaluated only at offsets for which the threshold was exceeded. The highest value then gives a reference offset from which the finer, patch-level correlations can be executed over a considerably smaller range of offsets. The image-level correlation is executed over an offset range up to 120 pixels; whereas the patch-level correlation range has generally been limited to 10 or 20 pixels. A regular grid (nominally three rows by five columns) of patches is established over the image to yield a geometrically configured set of correlation functions from which image geometric distortions of higher-order than pure linear translation can be inferred.

F. OFFSET DETERMINATION

The sub-pixel position of the patch-level correlation peak is located by the same method as the MDP to estimate the peak offset from the correlation function origin. Up to a fifth-order bivariate polynomial is fit by least squares to the array of correlation samples and the polynominal's peak is determined by iteration.

Assuming that a fit uncertainty variance can be estimated from the least squares residuals, we can estimate the covariance of uncertainty in the polynominal coefficients by multiplying the fit uncertainty variance by the reverse of the least squares normal matrix. By use of perturbation techniques, small displacements in the polynominal peak location can be related to perturbations in the polynomial coefficients, and it can be shown that the uncertainty covariance in peak location is

$$\underline{\underline{E}}_p = \underline{\underline{T}}\ \underline{\underline{E}}\ \underline{\underline{T}}^T,$$

where

$$\underline{\underline{T}} = (\underline{\underline{P}}_{xx}{}^T \underline{a})^{-1} \underline{P}_x{}^T$$

is a 2x15 matrix (2 peak coordinates by 15 fit coefficients for fifth-order), $\underline{a}$ is the vector of fit coefficients, $\underline{P}_x$ is a 15x2 array of partial derivatives of the polynomial terms $\underline{P}$ (such that $\underline{P}^T\underline{a}$ is the polynomial) with respect to the two spatial coordinates, and $\underline{\underline{P}}_{xx}$ is the corresponding array of second partial derivatives. These partial derivatives are evaluated at the estimated peak location. The surface is fit to a 5x5 array of correlation samples centered on the largest one, and the resulting 25 residuals are hardly adequate for estimating the uncertainty covariance of the 15 fit coefficients. Therefore, a fit uncertainty variance has been selected a priori based on some initial tries. The peak uncertainty covariance so estimated can be used (i.e., its inverse) to weight the patch in the least squares solution for the residual distortion function.

Besides weighting each patch, a peak-to-background ratio (an effective SNR) is estimated and compared with a threshold for possibly rejecting the patch. That ratio is computed as the value of the correlation surface peak divided by the standard deviation of correlation samples away from the peak, consistent with the analysis of Mostafavi and Smith[11,12]. Preliminary analyses by Wacker et al[13] indicate adequate rejection with a peak-to-background threshold of 4.2.

G. RESIDUAL DISTORTION FIT

Having produced an array of residual translational offsets over the registrant image based on the corresponding patch correlations, we can fit a bivariate polynomial to this array to characterize the residual distortion. An affine (first-order) model is assumed which will account for translation, scale change, rotation, and shear (rectangle-to-parallelogram) distortion. The affine model was adjudjed sufficient from early considerations of the sample segment size, accuracy requirements, and observed MSS residual geometric distortions. The peak uncertainty covariances are used as weights in the fit, viz. the reciprocals of the two diagonal elements are used as weights for the two coordinates, respectively. The uncertainty covariance of the distortion polynomial coefficients is also estimated to give a mesaure of registration accuracy.

Similarly, the MDP provides an uncertainty covariance (actually, its inverse in the form of the weighted least squares normal matrix) to measure its registration accuracy. The two measures can be compared to grade the residual registration performance of the JSC Registration Processor against the MDP's initial registration performance. If the latter is significantly better than the former (as would occur if the MDP had many good control points, while JSC's area of interest showed poor feature definition), then the residual distortion estimate should be ignored, or at least downweighted. On the other hand, if the MDP's estimate is poor compared to the residual solution, the latter should be relied upon in full. This concept can be implemented heuristically as follows. The MDP implies there is zero residual distortion in its output imagery to a level of confidence characterized by its covariance. The JSC system implies there is a residual distortion to a level of confidence characterized by its covariance. A better estimate of the distortion is attained by computing a weighted average of the MDP's zero estimate and the JSC system's estimate, where the weights are the inverses of the respective covariances. Hence, the final vector $\underline{X}$ of affine coefficient estimates is

$$\underline{X} = (a\underline{\underline{C}}_{MDP}^{-1} + \underline{\underline{C}}_{JSC}^{-1})^{-1}\underline{\underline{C}}_{JSC}^{-1}\underline{X}_{JSC}$$

from the MDP covariance $\underline{\underline{C}}_{MDP}$ (reinterpreted as a covariance of affine uncertainty) and JSC-estimated residual affine coefficients $\underline{X}_{JSC}$ and covariance $\underline{\underline{C}}_{JSC}$. A scalar parameter a is included above to rescale the MDP covariance if desired. So far, problems with control points have

caused the MDP covariance to underestimate its error significantly, requiring a to be set very small, and the above algorithm has not really been usable. Recent improvements in MDP control point handling should make it useful in the future.

The residual affine coefficients are delivered to the Mapping function so that the residual distortion can be included in the final mapping and resampling.

V. OBSERVED PERFORMANCE

The registration processor is currently in a period of confidence testing during which results are carefully scrutinized, operational procedures are refined, operating parameters are adjusted, and programming and algorithmic problems are resolved. The system is in production status, and a file of experience is being accumulated. Although we are not yet ready to issue definitive statements regarding its performance, certain trends and quantitative results are available. Work is underway to characterize fully the behavior of the processor.

A control test was devised in which registration error was introduced into a selected Landsat MSS image according to a specified transformation. The spatial resolution of the transformed image was degraded slightly by using bilinear interpolation in the resampling. A 196x177-pixel area of the original image was defined as the reference, the pair of images was registered by the processor, and the resulting mapping estimate was compared with the known one. The comparison was converted into an RMS value over the image, and a value of 1/15 pixel was obtained. The image was a very clean one, having good field definition and no atmospheric difficulties (though clouds would not have deleterious effect in this unitemporal, unisensor registration). The induced affine distortion was one exactly modellable by the processor. Thus, the test would tend to show a maximum accuracy of the processor.

Comparing that test's 1/15-pixel result with the 1/5-pixel production goal indicates to us that accommodating the vagaries and vicissitudes of real data should be the major focus of our immediate successive efforts. We note here that the sharpness in determining the peak of the cross-correlation surface is limited by variations in the data, rather than by the technique. One of us (AGW) found that the sharpness inherent in the method is on the order of ten times better than the result of our known-distortion test, as evidenced by the results obtained from autocorrelation. Thus, the finite size of the results is thought to originate in the discrete nature of the cross-correlation raster, the forced binary nature of the edge images being cross-correlated, and the fact that the cross-correlations were done between the uncorrected images. We have considered modifications to the system that would reduce this value by doing cross-correlation iteratively (i.e., repeat cross-correlation of the registrant after correction based on the first correlation).

However, limitations of time and resouces will delay attention in favor of more pressing problems. The operational improvements (i.e., sensor and preprocessing idiosyncracy accommodation) that are in work as of this writing comprise the following: selection of thresholds for peak-to-background ratio that optimally balance success rate and false registrations; determination of the optimum patch size; improved spectral recognition of clouds (necessary because of unknown changes in the absolute ratiometry between LACIE and MDP processing); and recognition and deletion of outliers in the family of cross-correlation peaks amongst the patches. The methodological improvements to be developed and tested as time permits include: allocation of "edge" pixels to be done at the patch level, so as to distribute control of the remapping more evenly over the image; iteration, in which a roughly-determined registration drives a patch resampling (followed by finer cross-correlation that takes advantage of the reduced differential rotation, scale, etc.); and methods of sharper edge definition and cross-correlation.

Visual examination of a substantial number of registrations has been done by blink comparison on a CRT. The accuracy is typically high enough to make nugatory the use of the color in the CRT, the registration accuracy usually being better than the color alignment; so the comparison is done by blinking in only one color. To permit semi-quantitative analysis, there are also displayed images that have been resampled onto warped grids of known displacement. The images with the known warpings are put up on the CRT to serve as "calibration" images; their motions of analytically-known distances are used for visual comparison to estimate the size of motions in the imagery being examined. The threshold of detectability is somewhat smaller than 1/10 pixel by this technique. In an alternate embodiment, subimage magnification by interpolation and replication gives similar sensitivity. The strong feature of this evaluation method is that the eye is used to do whole areas of image intercomparison; the context enables sharper discrimi-

nation than does a concentration on edge-labelled pixels alone. Additionally, image morphology of order different than the limits programmed into the processor is evident to the eye. And finally, the visual technique avoids the circularity of using the same metric methods to register and verify. Consistent results are the more believable if the methods are independent. The data reviewed were taken by Landsats 2 and 3 over agricultural areas in the central and eastern U.S. They represent a dispersion of field types and sizes, seasonal epoch, and weather conditions. Typical results are a ½ to 1-pixel maximum relative displacement with RMS values around ¼, but with an occasional gross miss. In the production data currently being run (calendar year 1981 MSS acquisitions) we see frequent obviously incorrect values of "line length majority", a value delivered by Landsat and passed through a sequence of GSFC processors ending with the MDP. It causes fractional-pixel motion of selected six-line strips (i.e., sweeps) in the images. Some of the gross misses are attributable to poor initial location and the periodic structure of agricultural scenes. During the blinking of those examples, the eye notices first the edge structures that the cross-correlator caught, and the registration at first appears satisfactory, when suddenly, as a dramatic revelation, it becomes clear that the miss is large. In only one (of many tens reviewed) was there a registration that was disappointing by reason of a poor geometric shape resulting from the residual mapping. The solution was incorrect in shear, by up to 3 or 4 pixels at one diagonal pair of corners, with the other pair of corners well in conjunction between the images.

Production registration of 1981 data sets from the U.S. and Argentina have indicated results similar to the test results. The initial misregistration of Argentina scenes often reached 45 pixels (2.5 km), due to the lack of control points in the MDP correction; yet, the JSC registrations were usually successful. Detailed statistics have not yet been compiled, but it is apparent that at least 2/3 of the cloud-free scenes are registering properly to an accuracy roughly in line with the foregoing discussion. That success rate is comparable to the one realized by the LACIE Registration Processor.

VI. CONCLUSIONS

The available evidence indicates the JSC Registration Processor has surpassed our initial performance expectations; however, room for improvement is evident. It is apparent that reaching a consistent 0.2-pixel registration accuracy is feasible. Emphasis should be placed on increasing the success rate, and a number of improvements in that direction are underway. We envision an evolutionary growth consisting of successions of tests of new parameters and algorithms on the existing system, followed by system upgrades reflecting those test results.

VII. ACKNOWLEDGEMENTS

The effective programming leadership of W. Lee of IBM contributed significantly to the successful completion of the JSC Registration Processor, and his participation is gratefully acknowledged. Conscientious program testing and evaluation support by R. J. Smart and W. H. McCumber of IBM is gratefully acknowledged. This work was performed in part under NASA Contract NAS9-14350.

VIII. REFERENCES

1. G. J. Grebowsky, "LACIE Registration Processing," in Proceedings of the Technical Sessions, The LACIE Symposium, pp. 87-97 (1979).

2. T. Kaneko, "Evaluation of Landsat Image Registration Accuracy," Photogr Eng Remote Sensing, 42, 1285-1299 (1976).

3. T. Kaneko, "Image Registration Accuracy Evaluation by an Edge Method," IBM memorandum to NASA JSC Mission Planning and Analysis Division, Contract NAS9-14350 (June 4, 1976).

4. R. Berstein, "Digital Image Processing of Earth Observation Sensor Data," IBM J. Res. Develop., 20, 40-56 (1976).

5. "MSS Data Processing Description," IBM Federal Systems Division report to NASA GSFC, Contract NAS5-22999 (November 1978).

6. W. Niblack, "The Control Point Library Building System," Photogr Eng Remote Sensing, 47, 1709-1715 (1981).

7. P. M. Hinson and C. H. Jeffress, "Earth Observations Division Landsat Imagery Preprocessing System," in Proceedings of the Sixth Symposium on Machine Processing of Remotely Sensed Data, p. 48 (1980).

8. K. W. Simon, "Digital Image Reconstruction and Resampling for Geometric Manipulation," in Proceedings of

Symposium on Machine Processing of Remotely Sensed Data, pp. 3A-1 - 3A-11 (1975); reprinted in Digital Image Processing for Remote Sensing, edited by R. Bernstein, IEEE Press, New York, pp. 84-94 (1978).

9. R. J. Kauth, P. J. Lambeck, W. Richardson, G. S. Thomas, and A. P. Pentland, "Feature Extraction Applied to Agricultural Crops as Seen by Landsat," in Proceedings of the Technical Sessions, the LACIE Symposium, pp. 705-721 (1979).

10. A. G. Wacker, "Empirically Determined Calibration Differences Between MDP-LIVES and LACIE Processed Data," AgRISTARS Report SR-J1-04133, JSC-17412 (June 1981).

11. H. Mostafavi and F. W. Smith, "Image Correlation with Geometric Distortion Part I: Acquisition Performance," IEEE Trans. Aerospace and Electronic Systems, AES-14, 487-493 (1978).

12. H. Mostafavi and F. W. Smith, "Image Correlation with Geometric Distortion Part II: Effect on Local Accuracy," IEEE Trans. Aerospace and Electronic Systems, AES-14, 494-500 (1978).

13. A. G. Wacker, R. H. Wolfe, Jr., and R. D. Juday, to be published.

Robert H. Wolfe, Jr. Dr. Wolfe received his B.A. and M.A. degrees in astronomy from the University of Kansas in 1962 and 1963, respectively, and his Ph.D. degree in physics from the University of Houston in 1973. Shortly thereafter, he joined the IBM Federal Systems Division where he investigated performance modeling of various imaging scanners used in NASA's aircraft project. He has since investigated performance and data reduction for a number of orbiting earth and celestial-viewing instruments, in addition to his work on the JSC Registration Processor. He is a member of the American Society of Photogrammetry, the Optical Society of America, and the American Astronomical Society.

Richard D. Juday Mr. Juday received a B.A. degree in physics from Rice University in 1965 and an M.S. degree in remote sensing from the University of Houston in 1979. He has been active in earth observations since 1970 with the Johnson Space Center. His chief interests are color displays and registration of digital imagery.

Arthur G. Wacker Dr. Wacher received his B.Sc. from Queens University in 1955, his M.Sc. from the University of Saskatchewan in 1962 and his Ph.D. from Purdue University in 1972. Following some industrial experience with Northern Electric, he joined the Department of Electrical Engineering at the University of Saskatchewan in 1957, and currently holds the rank of Professor. From 1967-1971, while on educational leave from the University of Saskatchewan, he was associated with Purdue University and its Laboratory for Applications of Remote Sensing. His research interests are in the area of pattern recognition and picture processing, particularly in relation to LANDSAT imagery and cartography. During 1980 to 1981 he spent his sabbatical performing registration research in the Earth Observations Division at Johnson Space Center. He is a member of the Association of Professional Engineers of Saskatchewan and of the Institute of Electrical and Electronics Engineers.

Toyohisa Kaneko Dr. Kaneko received his B.E. and M.E. degrees in electrical engineering from the University of Tokyo, Japan, in 1962 and 1964, and his Ph.D. in the same area from Princeton University in 1970. He joined IBM's Thomas J. Watson Research Center in 1970 where he worked on image processing, digital processing, and storage hierarchies. In 1974 he moved to the Federal Systems Division to work in computer processing of remotely sensed data, automatic image quality assessment, color film generation, and multi-crop classification. He is currently with the Communication Systems Division in the area of new office automation equipment. He is a member of the Institute of Electrical and Electronics Engineers and the Institute of Electronics and Communications Engineers in Japan.