

Research issues in image registration for remote sensing

Roger D. Eastman
Loyola College in Maryland
Baltimore, Maryland
reastman@loyola.edu

Jacqueline Le Moigne
NASA Goddard Space Flight Center
Greenbelt, Maryland
Jacqueline.LeMoigne@nasa.gov

Nathan S. Netanyahu
Bar-Ilan University
Ramat-Gan, Israel &
University of Maryland
College Park, Maryland
nathan@cfar.umd.edu

Abstract

Image registration is an important element in data processing for remote sensing with many applications and a wide range of solutions. Despite considerable investigation the field has not settled on a definitive solution for most applications and a number of questions remain open. This article looks at selected research issues by surveying the experience of operational satellite teams, application-specific requirements for Earth science, and our experiments in the evaluation of image registration algorithms with emphasis on the comparison of algorithms for subpixel accuracy. We conclude that remote sensing applications put particular demands on image registration algorithms to take into account domain-specific knowledge of geometric transformations and image content.

1. Introduction

Image registration, the alignment of images and other data for comparison and fusion, is an important element in data processing for remote sensing. It finds many applications, including but not limited to the operational production of satellite image products, validation and precision correction of satellite navigation data, radiometric validation of new sensors, data fusion for classification and thematic map production, temporal image comparison for change detection, atmospheric velocity computations, and land use analysis and planning.

This varied set of application needs, as well as similar needs in allied fields like medical imaging, have driven the development of a large number of approaches for image registration (see Brown [5], Fonseca and Manjunath [12] and Zitova and Flusser [35]). In an editorial on medical image registration Pluim and Fitzpatrick [22] note that since 1988 the number of related articles per year has grown about 10-fold, from approximately 15 to over 140.



Figure 1. Multisensor registration of ETM+ at 30m (surround) to IKONOS at 1m (inset)

Image registration for remote sensing, particularly in support of operational production of imagery products, can be quite automatic, efficient and effective. However, many end users still depend on labor-intensive manual selection of control points (CPs) or, to quote one researcher, "registration by graduate student," with times cited from a few hours to 4-5 days for registering a single image [8, 13]. This article explores this disparity in the use of advanced techniques, discusses selected issues in evaluating automatic methods and fostering their greater adoption, and outlines the research program in image registration conducted by the image registration working group at NASA Goddard Space Flight Center (GSFC).

In this article we describe image registration as a process that aligns one image to another of the same area taken at the same time or later, and with the same sensor or a different one. An image registration method can be defined

by three main steps: (1) extraction of features to be used to match the imagery; (2) feature matching strategy and metrics, commonly the optimization of a similarity measure between images; and (3) resampling or indexing of the data to a common coordinate system [5]. The transformations allowed for alignment can vary, from translation only in x and y , to rotation, scale and translation (RST), to fully-affine, to homographic, or to a more complex model [32, 26]. Many current systems use two levels of transformation as they match CPs as small image regions (chips) using only translation or a low order transformation, and then use the CPs to compute a more complex transformation. Some algorithms use Digital Elevation Models (DEMs) to correct for parallax caused by terrain relief.

2. Application requirements for image registration

In part, the wide variation in image registration techniques in remote sensing is a result of the wide variation in applications. Fonseca and Manjunath [12] noted this, proposed a rule-based system that would automatically select an appropriate method for an application and reviewed the literature to characterize the available methods to that end. Our own work [16] supports this goal. In particular, we have implemented a number of methods to evaluate them in a consistent framework and to develop an online toolbox of techniques for end users.

This program also requires understanding application needs to get an end-to-end view of how image registration fits into overall data processing. The literature on image registration is dominated by articles advancing or proposing a new technique, with far fewer articles that document user requirements or survey current practice. In this section we look at issues in user requirements, dividing users into two groups: operational groups that use image registration to produce image products, and Earth scientists (and others) who use the products.

2.1. Operational system requirements and practices

Operational satellite science teams that support remote sensing sensors rely on image registration for a number of purposes, including updating a geometric model that relates a sensor to the navigation data (position and orientation) for a satellite; updating the navigation data; image geolocation; sensor-to-sensor and intra-sensor calibration; and validation of model and data quality. Geolocating a satellite image relative to a ground coordinate system can be done in two primary ways: systematic correction using satellite navigation and calibration data, and precision correction through registration of image features to previously located ground features. Precision correction can be done indirectly, by updating the camera models used for

systematic correction, or directly by relating a new image to a ground coordinate system.

The teams have an end-to-end view of data processing requirements for image production. From an early point in sensor design they develop specifications that include image registration reliability and accuracy requirements, and they typically implement in-house software to those requirements. Their knowledge of the satellite and sensor enables them to model the radiometric and geometric transformations between images so the image registration algorithms can be tuned and integrate satellite navigation data. Below we survey six representative operational systems to look for characteristics and commonalities.

ASTER is a 14-band multispectral imager launched in 1999 on the Terra (EOS-AM1) satellite. It has three separate image subsystems with individual telescopes that sample at 15m, 30m and 90m. Iwasaki and Fujisada 2005 [14] describe the image registration system. Image registration was intended to improve geolocation of imagery to within 100m (nadir), adjust intratelescope band-to-band position to within 0.2 pixels and adjust intertelescope position to within 0.3 pixels. Registration is done to a database of about 300-600 CPs which were mapped on topographic maps to about 2-4 meters in x , y and z . The similarity measure is normalized correlation with transformation limited to translation. Subpixel estimation is calculated by fitting a second-order polynomial to the correlation values. Matches are rejected for correlation less than 0.7 or when clouds are detected. Validation of satellite performance was done in a retrospective study and the authors cite accuracies of 50m, 0.2 pixels and 0.1 pixels, respectively. Problems were noted with parallax errors in high-elevation areas due to time delays between bands, with cross-track performance due to satellite jitter, failures by the cloud detection algorithm, and cross-band matching in scenes with considerable vegetation.

GOES I-M is a series of weather satellites launched from 1994 to 2001 with seven bands that sample at 1km, 4km and 8km. Madani, Carr and Shoser 2004 [20] describe the AutoLandmark system for image registration. GOES geolocation requirements vary from night to day, from 4km to 10km, to accommodate solar heating. Registration is done to a database of coastlines represented in vector format which is used to avoid radiometric issues in correlation. A landmark of 24x24 or 96x96 pixels is extracted from a sensed image and edges detected using the Sobel operator. This is matched to a binary image rasterized from the coastline vector database, with a correlation measure of the sum of the edge strengths and transformations limited to translation. Cloud areas are masked out. Bisection search is used to find subpixel maximums. Matches are rejected when a fuzzy logic quality metric (QM) falls below a threshold of 0.9. The QM takes into account the quality of the edge cor-

Instrument	Satellite	Resolution	Similarity	Subpixel approach
ASTER	Terra	15-90m	Correlation to DEM corrected CPs	Fit to correlation surface
GOES	GOES I-M	1-8km	Correlation to vector coastline database	Bi-section search
MISR	Terra	275m	Correlation to DEM corrected CPs	Least squares optimization
MODIS	Terra	250m – 1km	Correlation to DEM corrected CPs	Fixed regular sampling
HRS	SPOT	2.5 m	Correlation to DEM corrected data	Not described
ETM+	Landsat-7	15-60m	Correlation to arid region CPs*	Fit to correlation surface*
VEGETATION	SPOT	1km	Correlation to DEM corrected CPs	Not described

Table 1. Operational Image Registration Systems. * ETM+ information in table refers to band-to-band registration

relation, the fraction of cloud contamination, and the contrast and possibly total illumination in the image. In the infrared bands the daily warming/cooling cycle means land and water can reverse in radiation, causing coastal edges to invert, migrate and even disappear.

MISR is a second instrument on the Terra satellite. Jovanovic *et al.* 2002 [15] describe the geometric correction system. It has 9 pushbroom cameras with four bands each, all sampling at 275m, intended to image atmospheric elements from multiple angles for climate analysis. This represents a particular challenge for geolocation of the multiple cameras with the objective of achieving better than 140m uncertainty in absolute geolocation at nadir and between cameras. Registration has been done to a database of 120 ground CPs, each a collection of 9 co-located 64x64 image chips from Landsat Thematic Mapper (TM) images of an identifiable ground feature located to 30m. Each CP was selected for seasonally invariant features and then mapped into terrain-corrected imagery. A ray-casting algorithm is used to warp each chip to the appropriate geometry for the appropriate MISR camera. Chip matching is done to sub-pixel accuracy in translation, potentially to 1/8 pixel, by least squares optimization [1]. Cloud masking is used and outliers are rejected.

MODIS is a third instrument on the Terra satellite. Wolfe *et al.* 2002 [18] describe the geolocation system (see also [31]). MODIS has 36 spectral bands that sample at 1km, 500m and 250m. The MODIS Land Science Team set a requirement of geolocation of 150m with an operational goal of 50m. Registration is done to a database of 605 land CPs known to 15m in three dimensions. For each CP, 24km² chips of 30m sampling were constructed from Landsat TM bands 3 and 4 taken from cloud-free images predominately along coastlines and waterways. The higher-resolution TM chips are resampled to MODIS resolution using the MODIS point spread function and nominal MODIS position information. Area-based correlation is used for matching, with a 25m sampling interval to get subpixel accuracy and a threshold of 0.6. Matches are used to update a camera geometric model that accounts for internal and external calibration and gives per pixel geometric corrections.

HRS, or High Resolution Stereoscopic with sampling at 2.5m, is an instrument on the SPOT 5 satellite launched in 2002. Baillarin *et al.* 2005 [3] describe the ANDORRE system with the TARIFA core for image fusion and rectification. HRS requirements are 90% of points within 16m circular accuracy and 10m elevation accuracy. ANDORRE uses an extensive database of orthoimagery tiles integrated with DEM data to generate, using ray tracing, a simulated image for matching. Matching is done using multiresolution search with the number of levels set to keep a 5x5 pixel size search window. CPs are automatically found, matched by correlation and used to calibrate a parametric model. Geometric outliers and CPs with correlation coefficient below 0.80 are rejected. In the case of an inadequate number of tie points the correlation threshold is lowered to 0.70 and the number of points increased ("densification"); a subscene of the image still without a sufficient number is rejected. The authors cite a final circular error of 2m for 88% of the points in 100 test images. Problems occur due to seasonal changes, low intensity images and homogenous subscenes.

ETM+ is an instrument on the Landsat-7 satellite launched in 1999. Lee *et al.* [19] describe the extensive ongoing set of assessment and calibration activities for ETM+ including image registration for geolocation. ETM+ requirements are for 250m absolute geodetic accuracy, 0.28 pixels band-to-band resolution and 0.4 pixels temporal registration. The three requirements lead to multiple registration approaches for calibration and assessment. Geodetic accuracy is updated by correlating systematically corrected panchromatic band regions against a database of CP image chips extracted from USGS digital orthophoto quadrangles with vertical information from USGS DEMs. About 75-100 CP chips are used for assessing accuracy. Band-to-band registration is used to update a parameter file with corrected off-sets, so selected images in arid regions with little vegetation and significant high spatial frequency content can improve cross-band correlation. Subpixel registration is done by second-order fit to the correlation surface. Temporal resolution is assessed by correlating a reference image with a precision and terrain corrected panchromatic band of the test image. Success is reported for absolute geodetic accu-

racy as 99.9% of test CPs within 250m and a 54m RMS error, while for band-to-band it is approximately 1/4 to 1/2 of the operational requirement. Problems have been seen in decay of satellite attitude data due to aging gyros, and band-to-band calibration due to post-launch and seasonal operating temperatures. Band-to-band correlation of band 6 (infrared) has been problematic due to radiometric differences including seasonal temperature variation in arid regions.

VEGETATION is an instrument on the SPOT series of satellites, operational since SPOT4 in 1999. Sylander *et al.* 2000 [24] describe the geolocation system. The requirements are 1km in absolute location, 500m in multi-temporal registration with an operational objective of 300 m, and multispectral registration of 300m. Initial work indicated that existing coastline databases were not accurate enough, and reliance on systematic correction was not adequate so a database of approximately 3500 CPs was built from VEGETATION images. Each CP has four image chips from different seasons and orientations. Matching using correlation is done under the control of human operators who insure there are 10 matched points per orbit to compute satellite position corrections. Absolute location was measured to 300m RMS, multitemporal to 325m in RMS, and multispectral was reported as 0.11 pixels, which at that time was slightly outside the requirements. Problems occur due to seasonal changes, images with inadequate distribution of CPs, rejection of images such as islands where CPs cannot be uniformly distributed, and CPs that may have terrain features poorly suited for correlation.

From this survey we can see several common features and needs. The systems all share the use of pre-constructed databases, either of CPs or vector coastlines; the use of normalized intensity (or vector-edge) correlation as the similarity measure between features; DEM-correction of features before correlation; matching in local regions rather than globally; the use of translation since it dominates in small regions; cloud masking or thresholds to eliminate cloudy regions. Most systems use subpixel estimation but vary in the method used to compute it. The systems also vary in how local matches are used in precision correction – whether they are combined to estimate a global image transformation or an update to the geometric camera model for systematic correction. High-resolution sensors, on the order of 1 m, face specific challenges in accounting for terrain elevation effects and temporal changes caused by humans and other Earth processes. The problems reported by the groups show some differences but add up to the following picture of current registration issues: effectiveness of normalized correlation in cross-band registration; adaption to thermal changes in satellite geometry and minor problems in orbit data; uniform distribution of CPs adequate for further numerical computations; CPs that are poorly suited for correlation for a number of reasons, including seasonal

changes in temperature or vegetation or homogenous terrain features, and properly modeling of satellite position variations.

2.2. End user application requirements

In contrast to the operational teams, end users do not always have the same resources such as intimate knowledge of the sensor and data processing, the full set of sensor imagery, the ability or need to maintain a large database of CPs, access to navigation modeling and other correction data, or statistical models of geolocation error. In addition, they may need to register heterogeneous data sources such as maps and older aerial images for which the supporting information is not available or complete. But, they do bring to image registration tasks knowledge of their application needs, most specifically accuracy needs and application-specific models.

Accuracy needs of image registration for change detection has been one of the most studied application needs [27, 28, 6]. Townshend *et al.* 1992 [27] found that for four regions studied in Landsat TM imagery, to keep NDVI (Normalized Difference Vegetation Index) pixel-based change estimations down to an error of 10% required subpixel registration of 0.2 pixels. However, for three more homogeneous, arid regions they found that registrations of 0.5 to 1.0 pixels were adequate for 10% error. Dai and Khorram 1998 [10] report similar results with about 0.2 pixels registration required for 10% error on average in four test sites, with a slight but similar dependency on the spatial frequency content of the imagery. Other confounding issues considered for misregistration impact on change detection has been the number of classes and region heterogeneity (Verbyla and Boles 2000 [28]), as well as change detection resolution and more through measures of region heterogeneity in high resolution imagery (Wang and Ellis 2005 [30]). Authors have also considered how to adapt to misregistration errors by estimating them or lowering the resolution of change detection [28, 30, 6].

Knowledge and use of appropriate transformation models can assist in improving registration accuracy. Toutin 2004 [26] extensively reviews transformation models used for satellite imagery ortho-rectification. Physical models that take into account satellite geometry and the sensor imaging model, such as those used by the operational groups, can be the most accurate but are not always available. Empirical models, such low order affine or RST, higher order polynomial, or piecewise transformations, do not require detailed sensor and satellite models but may require more CPs. Three-dimensional rational functions (RFs), ratios of polynomials, can approach the accuracy of physical models and are distributed by some imagery providers in meta-data but are not commonly used in automatic registration algorithms.

There are good studies that give extensive analyses of end user image registration experience and accuracy. Wang and Ellis 2005 [29] give such a detailed description and analysis for the registration by manual CPs of IKONOS and historical aerial photographs in China, detailing the acquisition by GPS of precise ground CPs for the imagery, studying the impact of various factors including terrain elevation variation and number of CPs, and using independent check points (ICPs) to measure accuracy. Armston *et al.* 2002 [2] gives a similar analysis of geometric correction of Landsat imagery over Australia, giving comparisons of physical and RF transformation models (the latter produced higher error) and looking at the spatial distribution of errors for control and check points. They give the rules used for selection of CPs, which included syntactic properties (choose high contrast in two orthogonal directions, within 5x5 window) and semantic properties (avoid cleared land and bodies of water and choose bare ground among vegetation.)

The assessment of transformations produced by manual CP selection is often judged by precision rather than formally by accuracy, as measures like RMS error are used that evaluate tightness of variance of the residuals for CPs Brovelli *et al.* 2006 [4]. This is often done with the residual under the computed transformation to a separate set of independent check points using holdout validation, or with leave-one-out cross-validation (LOOCV) which recomputes the transformation with $N - 1$ subsets and computes the residual on the remaining point. Spatial and orientation distributions of residuals have also been considered (Armston *et al.* 2002 [2], Buiten and van Putten 1997 [7]). In general, accuracy increases with the number of CPs up to a diminishing return [26, 29].

From this brief survey of end user requirements for registration there can be seen connections with the experience of the operational groups. Homogenous regions that cause problems for image registration correlation may have compensating effects in reducing the accuracy requirements in applications like change detection; it may be possible to develop models for end users that, given the application requirements and the images to be processed, will specify updated and realistic accuracy requirements.

3. Studies for evaluation of image registration

Work of the NASA image registration group has emphasized the study and evaluation of automatic image registration algorithms for several purposes: precision correction of satellite imagery, band-to-band calibration, dimension and data reduction for ease of transmission, temporal multi-sensor fusion and on-board processing. Despite the wide variety of algorithms available for image registration, no commercial software seems to respond to the needs of Earth and space data registration. Our objective is to carry out systematic studies to support image registration users

in selecting appropriate techniques for a remote sensing application based on accuracy and suitability for that application. We carry this out by surveying, designing and developing different components of the registration process so we can evaluate their performance on well-chosen multiple source data, provide quantitative inter-comparison and to build an operational image registration toolbox. Following these goals, we performed several sets of experiments, and according to the observations described in the previous sections, the challenges that were addressed in our experiments were concerned with different aspects of the same features in various temporal, various spatial resolutions and various spectral data, and obtaining subpixel accuracy.

Assessing an image registration algorithm for subpixel accuracy on remote sensing data presents some difficulty since ground truth is not available. Interleaving two images for visual assessment can detect gross mismatches and global alignment but is difficult to extend to quantitative subpixel evaluation. Manual registration can be used to calculate the unknown transformation but it is uncertain if it accurate enough to test subpixel accuracy on small regions. A second approach is to generate synthetic image pairs by matching one image against a transformed and resampled version of itself with or without added noise. To avoid some resampling issues this can be done by using high resolution imagery and downsampling to a lower resolution using an appropriate point spread function to generate both images in a pair. However, while useful this approach is limited in realistically modeling noise, temporal scene changes or cross-sensor issues. A third approach is to use circular registration results on natural imagery when three or more overlapping images are available. In this case the transformations should compose to yield the identity— for three images registered pairwise by T_1 , T_2 and T_3 , the composition $T_1 \circ T_2 \circ T_3$ should be close to the identity transformation [34].

In all our experiments, we assume that the data have already been corrected according to a navigation model and we can assume this systematic correction is accurate within a few or a few tens of pixels. Our precision-correction algorithms utilize selected image features or CPs to refine this geolocation accuracy within one pixel or a subpixel at the CPs. We first performed preliminary experiments summarized in [16] looking at different components individually [23, 33] and then developed a modular image registration framework [17] that enables systematic studies comparing these components and to test new components in a rigorous fashion.

The first series of experiments investigated the use of different features with exhaustive correlation search: gray levels, edges, orthonormal/orthogonal wavelets (such as Haar and Daubechies) and over-complete decompositions (such as the Simoncelli steerable pyramid). The results showed that, as expected, edge or edge-like features (such

as wavelets) were more robust to noise and local intensity variations than pure intensity values. In comparing edges and wavelets, a full exhaustive search using edge features proved to be potentially more accurate but was also time-prohibitive compared to an approach based on features extracted from an orthogonal wavelet.

The next series of experiments studied the potential limits of using wavelets, and particularly the effects of image translation on wavelet-based image registration [23]. It is well-known that wavelet representations are not translation-invariant since when a separable orthogonal wavelet transform is used, information about the signal changes within or across sub-bands. Within these restrictions we quantitatively assessed the usefulness of orthogonal wavelets. The main result was that the normalized correlation coefficients of low-pass orthogonal wavelet subbands are essentially insensitive to translations for features larger than twice the wavelet filter size. The third-level low-pass subbands produce a correlation peak that varies with translation from 0.7 to 1.0 with an average of 0.9. On the other hand, high-pass subbands are more sensitive to translation, although this sensitivity is limited.

A follow-on series of experiments focused on evaluating several wavelet-type pyramids that may be utilized both for invariant feature extraction and for representing images at multiple spatial resolutions to accelerate registration. Using an optimization-based scheme, Daubechies orthogonal wavelets, spline wavelets and the Simoncelli steerable Pyramid were considered. We found that the band-pass wavelets obtained from the Steerable Pyramid due to Simoncelli perform best in terms of accuracy and consistency, while the low-pass wavelets obtained from the same pyramid give the best results in terms of the radius of convergence. Although not as reliable, spline wavelets show the best accuracy when convergence occurs.

A more general set of experiments investigated combinations of the major image registration components: features, similarity metrics and search strategy. We built a framework to selectively combine the components in order to reach optimum registration on a given type of data and under given circumstances. Another use of the framework is to test each element of registration independently (so we know the edge detection implementation, for example, is the same in all experiments.) The purpose of this framework is triple-fold: (1) assess various combinations of components as a function of the application, (2) assess a new registration component compared to other known ones, and (3) serve as a registration tool where a user can select components as a function of the application at hand, the available computational resources and the required accuracy.

In our current framework:

- features can be either gray levels, Low-Pass features from Simoncelli steerable filters decomposition

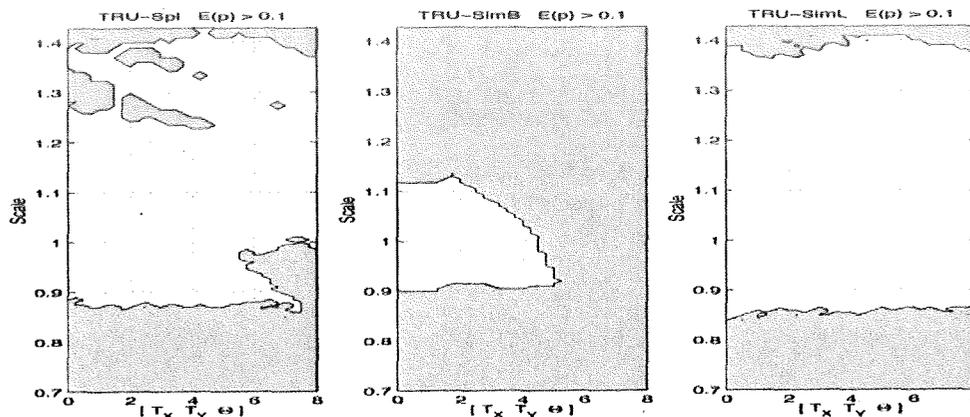
or from a Spline decomposition, or Simoncelli Band-Pass features, similarity metrics can be either cross-correlation, the L_2 Norm, Mutual Information or an Hausdorff distance,

- matching strategies are either based on a Fast Fourier Correlation, three different types of optimization - pure gradient Descent [11], a Marquard-Levenberg approach [25] or a stochastic gradient approach [9], and a Robust Feature Matching approach [21].

Studies using this framework are described in [17], and some of the first results are the following:

- Fast Fourier correlation is very fast but limited to transformations with scales close to 1 (in the range [0.95, 1.2]) and images containing very small amounts of noise.
- When using a Marquard-Levenberg approach, features such as Simoncelli Band-Pass are more accurate and more robust to noise, with errors in the range of [0.032, 0.25] pixel. At the same time, Simoncelli-Low Pass features are less sensitive to the initial guess than Band-Pass features.
- An approach using a stochastic gradient approach and a Mutual Information metrics is overall the most robust to initial conditions and the most accurate with errors in the range of [0.12, 0.18] pixel.

Figure 2 illustrates the types of results we obtained in one of these studies, investigating the sensitivity of components to initial conditions. In this experiment components of the framework were tested using synthetic data as well as Landsat multi-temporal data and IKONOS, Landsat, MODIS and SeaWiFS data over four well-chosen EOS Land Validation Core Sites. Synthetic images were created using a controlled process [33] designed to emulate real data. Starting from a given source image, three types of modifications are applied in various combinations to produce synthetic test data: (1) geometric warping, (2) radiometric variations and (3) addition of noise. Figure 2 shows the results of the Marquard-Levenberg optimization strategy combined with a Mutual Information metrics and three different feature choices, low-pass Spline-wavelet features, Simoncelli band-pass features and Simoncelli low-pass features, when applied to synthetic data. Shifts are varied from 0 to 8 pixels, rotations are varied from 0 to 8 degrees and scales are varied from -0.7 to 1.4. The white areas depict the regions of convergence of the algorithms with an error less than 1 pixel. This particular experiment demonstrates how Simoncelli Band-Pass features are more sensitive to the initial guess than Splines or Simoncelli Low-Pass features, while these last features exhibit the largest area of convergence.



Sensitivity of TRU Algorithms to Initial Guess

Figure 2. Sensitivity to initial conditions

4. Conclusion

In order to better understand requirements for the development and evaluation of image registration algorithms, this article has examined selected research issues by surveying the experience of operational satellite teams, application-specific requirements for Earth science, and our framework and experiments in the evaluation of image registration algorithms. From the experience and current practice of operational groups, as well as documented requirements for end user applications, we can see the importance of subpixel registration for image subregions with increasing need to perform it robustly between sensors, particularly with the advent of very high resolution and hyperspectral imagery and more demanding requirements for applications. We also see the need for image registration techniques that are tailored to the needs of remote sensing applications with specific attention to sensor-related geometric transformations and knowledge of accuracy requirements and image content.

Acknowledgments

We would like to acknowledge the invaluable contributions of collaborators with the NASA image registration group over the years: Arlene Cole-Rhodes and Kisha Johnson of Morgan State University, Tarek El-Ghazawi of George Washington University, Peyush Jain, Amiee Joshua, Jeffrey Masek, Jeffery Morsissette, Nargess Memarsadeghi and Ezinne Uko-Ozoro of NASA/GSFC, and David Mount of the University of Maryland at College Park.

References

- [1] F. Ackerman. Digital image correlation: Performance and potential application in photogrammetry. *Photogrammetric Record*, 11(64):429–439, 1984.
- [2] J. D. Armston, T. J. Danaher, B. M. Goulevitch, and M. I. Byrne. Geometric correction of landsat mss, tm, and etm+ imagery for mapping of woody vegetation cover and change detection in queensland. In *Proceedings of the 11th Australasian Remote Sensing and Photogrammetry Conference*, pages 1–23, Brisbane, Australia, 2002.
- [3] S. Baillarin, A. Bouillon, M. Bernard, and M. Chikhi. Using a three dimensional spatial database to orthorectify automatically remote sensing images. In *Proceedings of the ISPRS Hangzhou 2005 Workshop*, pages 89–94, Hangzhou, China, 2005.
- [4] M. A. Brovelli, M. Crespi, F. Fratarcangeli, F. Giannone, and E. Realin. Accuracy assessment of high resolution satellite imagery by leave-one-out method. In *7th International Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Science*, pages 533–542, 2006.
- [5] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(24):325–376, 1992.
- [6] L. Bruzzone and R. Cossu. An adaptive approach to reducing registration noise effects in unsupervised change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 41(11):2455–2465, 2003.
- [7] H. J. Buiten and B. v. Putten. Quality assessment of remote sensing image registration – analysis and testing of control point residuals. *ISPRS Journal of Photogrammetry and Remote Sensing*, 52(2):57–73, 1997.
- [8] D. Carrion, M. Gianinetto, and M. Scaioni. Georef: a software for improving the use of remote sensing images in environmental applications. In *Proceedings of the First Biennial Meeting of the International Environmental Modelling and Software Society*, pages 360–365, 2002.

- [9] A. Cole-Rhodes, K. Johnson, and J. Le Moigne. Multiresolution registration of remote sensing imagery by optimization of mutual information using a stochastic gradient. *IEEE Transactions on Image Processing*, 12(12):1495–1511, 2003.
- [10] X. Dai and S. Khorram. The effects of image misregistration on the accuracy of remotely sensed change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 36(5):1566–1577, 1998.
- [11] R. D. Eastman and J. Le Moigne. Gradient-descent techniques for multi-temporal and multi-sensor image registration of remotely sensed imagery. In *Fourth International Conference on Information Fusion*, Montreal, Canada, 2001.
- [12] L. M. G. Fonseca and B. S. Manjunath. Registration techniques for multisensor remotely sensed imagery. *Photogrammetric Engineering and Remote Sensing*, 62(9):1049–1056, 1996.
- [13] A. Ganas and E. Athanassiou. A comparative study on the production of satellite orthoimagery for geological remote sensing. *Geocarto International*, 14:51–59, 2000.
- [14] A. Iwasaki and H. Fujisada. Aster geometric performance. *IEEE Transactions on Geoscience and Remote Sensing*, 43(12):2700–2706, 2005.
- [15] V. M. Jovanovic, M. A. Bull, M. M. Smyth, and J. Zong. MISR in-flight camera geometric model calibration and georectification performance. *IEEE Transactions on Geoscience and Remote Sensing*, 40(7):1512–1519, 2002.
- [16] J. Le Moigne, A. Cole-Rhodes, R. Eastman, P. Jain, A. Joshua, N. Memarsadeghi, D. Mount, N. S. Netanyahu, J. Morsissette, and E. Uko-Ozoro. Image registration and fusion studies for the integration of multiple remote sensing data. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 100–105, Toulouse, France, 2006.
- [17] J. Le Moigne, H. S. Stone, A. Cole-Rhodes, R. D. Eastman, P. Jain, K. Johnson, J. T. Morisette, and N. S. Netanyahu. Registration of multiple sensor earth science data. In *NASA Earth Science Technology Conference (ESTC)*, pages 100–105, Palo Alto, CA, 2004.
- [18] D. S. Lee, J. C. Storey, M. J. Choate, and R. W. Hayes. Achieving sub-pixel geolocation accuracy in support of modis land science. *Remote Sensing of Environment*, 83(1–2):31–49, 2002.
- [19] D. S. Lee, J. C. Storey, M. J. Choate, and R. W. Hayes. Four years of landsat-7 on-orbit geometric calibration and performance. *IEEE Transactions on Geoscience and Remote Sensing*, 42(12):2786–2795, 2004.
- [20] H. Madani, J. L. Carr, and C. Schoeser. Image registration using autolandmark. In *Proceedings of Fusion of Earth Data*, pages 234–778, Sophia-Antipolis, France, 2004.
- [21] N. S. Netanyahu, J. Le Moigne, and J. G. Masek. Georegistration of Landsat data via robust matching of multiresolution features. *IEEE Transactions on Geoscience and Remote Sensing*, 42(7):1586–1600, 2004.
- [22] J. P. W. Pluim and J. M. Fitzpatrick. Image registration (editorial). *IEEE Transactions on Medical Imaging*, 22(11):1341–1343, 2005.
- [23] H. S. Stone, J. Le Moigne, and M. McGuire. Image registration using wavelet techniques. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(10):1074–1081, 1999.
- [24] S. Sylvander, P. Henry, C. Bastien-Thiry, F. Meunier, and D. Fuster. Vegetation geometrical image quality. In *Proceedings of the VEGETATION 2000 Conference*, Balgirate, Italy, 2000.
- [25] P. Thévenaz, U. Ruttiman, and M. Unser. A pyramid approach to sub-pixel registration based on intensity. *IEEE Transactions on Image Processing*, 7(1):27–41, 1998.
- [26] T. Toutin. Review article: Geometric processing of remote sensing images: models, algorithms and methods. *International Journal of Remote Sensing*, 25(10):1893–1924, 2004.
- [27] J. R. G. Townshend, C. O. Justice, C. Gurney, and J. McManus. The impact of misregistration on change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 30(5):1054–1060, 1992.
- [28] D. L. Verbyla and S. H. Boles. Bias in land cover change estimates due to misregistration. *International Journal of Remote Sensing*, 21(18):3553–3560, 2000.
- [29] H. Wang and E. C. Ellis. Image misregistration error in change measurements. *Photogrammetric Engineering and Remote Sensing*, 71(9):1037–1044, 2005.
- [30] H. Wang and E. C. Ellis. Spatial accuracy of orthorectified ikonos imagery and historical aerial photographs across five sites in china. *International Journal of Remote Sensing*, 26(9):1893–1911, 2005.
- [31] X. Xiong, N. Che, and W. Barnes. Terra modis on-orbit spatial characterization and performance. *IEEE Transactions on Geoscience and Remote Sensing*, 43(2):355–365, 2005.
- [32] L. Zagorchev and A. Goshtasby. A comparative study of transformation functions for nonrigid image registration. *IEEE Transactions on Image Processing*, 15(3):529–538, 2006.
- [33] I. Zavorin and J. Le Moigne. Use of multiresolution wavelet feature pyramids for automatic registration of multisensor imagery. *IEEE Transactions on Image Processing*, 14(6):770–782, 2005.
- [34] I. Zavorin, H. S. Stone, and J. Le Moigne. Evaluating performance of automatic techniques for subpixel registration of remotely sensed imagery. In *Proceedings of SPIE – Volume 5014 Image Processing: Algorithms and Systems II*, pages 306–316, 2003.
- [35] B. Zitova and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, 2003.