Multi-image registration for an enhanced vision system

Glenn Hines[a], Zia-ur Rahman[b], Daniel Jobson[a],Glenn Woodell[a]

[a]NASA Langley Research Center, Hampton, VA 23681;
[b]College of William & Mary, Department of Applied Science, Williamsburg, VA 23187

ABSTRACT

An Enhanced Vision System (EVS) utilizing multi-sensor image fusion is currently under development at the NASA Langley Research Center. The EVS will provide enhanced images of the flight environment to assist pilots in poor visibility conditions. Multi-spectral images obtained from a short wave infrared (SWIR), a long wave infrared (LWIR), and a color visible band CCD camera, are enhanced and fused using the Retinex algorithm. The images from the different sensors do not have a uniform data structure: the three sensors not only operate at different wavelengths, but they also have different spatial resolutions, optical fields of view (FOV), and bore-sighting inaccuracies. Thus, in order to perform image fusion, the images must first be co-registered. Image registration is the task of aligning images taken at different times, from different sensors, or from different viewpoints, so that all corresponding points in the images match. In this paper, we present two methods for registering multiple multi-spectral images. The first method performs registration using sensor specifications to match the FOVs and resolutions directly through image resampling. In the second method, registration is obtained through geometric correction based on a spatial transformation defined by user selected control points and regression analysis.

Keywords: image registration, control points, resampling

1. INTRODUCTION

Improving aviation safety is a primary mission of the National Aeronautics and Space Administration (NASA). Synthetic Vision Systems (SVS) is an initiative at the NASA Langley Research Center (LaRC) to improve the level of aviation safety. One component of the SVS is the Enhanced Vision Systems (EVS) project where several sensor technologies are being utilized to provide vision-enabling information to commercial and general aviation pilots to assist them in poor visibility conditions.

Infrared sensors are routinely used in remote sensing applications. Coupling an infrared sensor with a visible band sensor — for frame of reference or for additional spectral information — and properly processing the two information streams has the potential to provide valuable information in night and/or poor visibility conditions. A multi-spectral infrared and visible EVS[1] was developed at LaRC to test this concept. This system contains a short wave infrared (SWIR) camera, a long wave infrared (LWIR) camera, a color, visible band, charge coupled device (CCD) camera, and the hardware and software to process the data streams from the cameras. The system has been deployed and flown on the NASA LaRC B757 ARIES research aircraft.

The digitized images obtained from the data streams of the visible and infrared cameras are processed using the Multi-scale Retinex (MSR) algorithm. The MSR is a general-purpose image enhancement algorithm for producing good visual representations of scenes. It performs a non-linear spatial/spectral transform that synthesizes strong local contrast enhancement and color constancy.[2,3] Though it was developed for color CCD images, the MSR has been applied to infrared data from the Airborne Visible Infrared Imaging Spectrometer (AVIRIS) and the LANDSAT thematic mapper to pre-process the images for better classification.[4,5] The dynamic range compression, color constancy and sharpening provided by the MSR makes it an ideal algorithm to use for image enhancement in poor weather conditions. Since the MSR is used to enhance infrared as well as visible spectral bands, different features in different bands can be accentuated. From this a composite or fused image can be created that contains more information than any individual spectral band. For the EVS, the MSR serves as both an enhancement processor and fusion engine.[6,7] A set of images consisting of an image from each of the cameras taken during one time-aligned frame is fused into a single image. This process is then repeated for all the image frames making up a video sequence.

To properly perform fusion it is critical to ensure that the information from each sensor refers to the same features in the environment.[8,9] The different sensors of the EVS have different acquisition lattices and optics, therefore they capture information in data structures that are substantially different from each other. Thus, the images must first be registered before any fusion is performed. Several authors[10–12] have addressed image registration problems with innovative, but often complex, general solutions. In this paper we describe two straightforward solutions for registering the EVS images.

The first solution is a registration algorithm based on simply correcting camera sensor specification differences through cropping, scaling and image resampling, assuming that the cameras are bore-sighted and hence share a common center. We will show and discuss an example of registration with this algorithm. Tests based upon this method of registration showed that the assumption that the cameras were perfectly bore-sighted was invalid and adjustments had to be made to correct for the differences in sensor orientation and alignment. A second algorithm was implemented that factored in these obvious distortions and gave better results than the original algorithm. However, whereas the first algorithm was completely automatic — the only inputs were the characteristics of the different sensors and which sensor provided the baseline — the second algorithm does require a manual selection of control points. It should be pointed out, though, that if this correction was to be performed on a video sequence, then only one frame is needed for manual selection of control points. The coefficients that are obtained from this correction can be used to correct all the other frames, assuming that the nature of the distortion does not change. We will also show and discuss examples of this algorithm.

2. BACKGROUND

2.1. Image registration

Image registration is the task of aligning images taken at different times, from different sensors, or from different viewpoints so that all corresponding points in the images match. A transform must be defined that relates the points in one image to their corresponding points in another. This transform depends upon the characteristics of the differences between the images being registered, and is computed with respect to a reference or baseline image. The images that are to be matched to the reference are called the sensed, or, distorted image.

More particularly, image registration is defined as a mapping between two or more images both spatially (geometrically) and with respect to intensity. Expressed mathematically we have:

$$I_2 = g\big(I_1(f(x_1, x_2))\big),$$

where I_1 and I_2 are two-dimensional images (indexed by x_1, x_2), $f : (x_1, x_2) \to (x'_1, x'_2)$ maps the indices of the distorted frame to match those of the reference frame, and g is a one-dimensional intensity or radiometric transform.[13] In this paper, we assume that we do not need to make any radiometric adjustments, so $g = \mathcal{I}$, the identity transform. Hence we are concerned only with the spatial transformation, f. In generating a spatial transform for the EVS, our primary difficulty is the lack of fiducial markers within the images generated by the EVS sensors. The cameras are, however, assumed to be bore-sighted so they are expected to have a common center of alignment. The spatial transform should, then, properly align the images, but should not affect any characteristic differences that should be exposed by registration.

Spatial transforms may take on different forms depending upon the application. Simple, common transforms specified by analytic expressions include rigid-body, affine, projective or perspective, and polynomial.[14,15] The distortions between the images of the EVS in general seem constrained to those correctable by affine transforms. They also appear to be characterizable by a global (versus local[16]) transform where a single transform correctly maps all the points on the distorted image to match the corresponding points on the reference image. An affine transform fulfills the requirements for the needed transform.

An affine transform can perform rotation, translation, scaling and shearing operations. It offers six degrees of freedom when selecting six unknown coefficients and solving a system of six linear equations. In general, it can

perform triangle-to-triangle mappings. A general representation of an affine transform is $[y_1, y_2, 1] = [x_1, x_2, 1]T$ where

$$T = \begin{bmatrix} a_{11} & a_{12} & 0 \\ a_{21} & a_{22} & 0 \\ a_{31} & a_{32} & 1 \end{bmatrix},$$

x_1 and x_2 reference the input coordinate system, y_1 and y_2 reference the output coordinate system, and a_{ij} are transform coefficients.[17]

The forward mapping functions are

$$y_1 = a_{11}x_1 + a_{21}x_2 + a_{31}$$

and

$$y_2 = a_{12}x_1 + a_{22}x_2 + a_{32}.$$

Geometric, image-to-image registration can be summarized in three general steps:

1. Feature identification and matching is performed to establish a correspondence between features in the distorted image to those in the reference image;

2. A spatial transformation is selected and the transformation coefficients are computed based upon the feature matching criteria;

3. The distorted image is inverse-mapped using the computed transformation and resampled to register it with the reference image.

Feature identification and matching are often performed by selecting pixel locations called control points. Identification of control points can be accomplished in several different ways.[16, 18] Manual identification of control points is commonly performed. The images are displayed, normally side-by-side, and corresponding points usually based on features such as lines, edges, or contours are selected from both images.

The spatial transform coefficients that represent the unknown image distortions are determined from the control points. A minimum of three non-collinear control points are required to determine the six unknown coefficients of an affine transformation. Wolberg and Jensen[17, 19] describe several techniques to solve for unknown coefficients including pseudo-inverse solutions, least squares with ordinary and orthogonal polynomials, and weighted least squares with orthogonal polynomials.

Image resampling is the process of transforming a sampled image from one (input pixel grid) coordinate system to another (output pixel grid), where a sampled image is the digitization of the spatial coordinates of an image function $f(y_1, y_2)$ — a two-dimensional intensity function.[17, 20, 21] The two coordinate systems are related to each other by the mapping function of a spatial transformation.

To perform image resampling, initially, the output pixels are inverse mapped using the transformation function to a new grid which (usually) doesn't correspond to the input grid. Thus an interpolation (image reconstruction) procedure is used to generate a continuous surface through the samples of the new grid. Then the input image is sampled (digitized) at these points to provide the discrete output pixel values of the process. Three common methods of interpolation are nearest neighbor, bilinear, and parametric cubic convolution.[17, 22]

2.2. Sensor specifications

The cameras chosen for the EVS are commercial-off-the-shelf (COTS) devices. The SWIR camera sensor captures wavelengths between 1-3 μm. The LWIR (thermal) camera sensor captures wavelengths between 8-12 μm. And, the color visible band CCD camera operates between 0.4-0.7 μm. Table 1 shows the relevant manufacturer characteristics of the sensors.[1] The images from the three sensors obviously need to be be registered because of the differences in these characteristics. The solutions developed to resolve these differences are discussed in Section 3.

	SWIR	LWIR	CCD
Image Dimensions (pixels)	$320H \times 240V$	$320H \times 240V$	$542H \times 497V$
Optics FOV	$34°H \times 25°V$	$39°H \times 29°V$	$34°H \times 25°V$
Detector Readout Frame Rate	60Hz (typical)	60Hz	30Hz (interlaced)

Table 1: Sensor Specifications

The characteristics of the actual images obtained for registration differ from the initial manufacturer specifications because of data acquisition and storage to tape. First, all images have a nominal image size of 640 x 480 pixels corresponding to the NTSC format of the recorded images. However, the actual size of the images is quite different after the images are cropped so that the FOVs match the "visible" part of the images (see Section 3). Second, ground test measurements of the cameras' FOVs differed from the manufacturer provided values. These updated characteristics are shown in Table 2, and need to be included in the computations for proper registration of the data streams.

	SWIR	LWIR	CCD
Image Dimensions (pixels)	$640H \times 480V$	$640H \times 480V$	$640H \times 480V$
Optics FOV	$31.5°H \times 23.5°V$	$41°H \times 30.75°V$	$33.5°H \times 25°V$
Detector Readout Frame Rate	60Hz (typical)	60Hz	30Hz (interlaced)

Table 2: Updated Sensor Specifications

2.3. Algorithm overview

The algorithms operate on a set of three, time-aligned images where each image is acquired by an individual camera of the EVS. Each of the video streams is recorded, or post-processed, with video timecode information in each frame. The frames are time-aligned simply by finding the frames with matching time codes. This set of time-aligned frames is then used to obtain the registration parameters with respect to the baseline frame. All other frames of the video sequence can be processed with the same parameters. Each frame, including the ones from the color CCD sensor, is converted to grayscale before registration and further processing.

3. REGISTRATION ALGORITHMS

Our first solution for image registration algorithm is based solely on camera sensor specifications. The cameras were assumed to be properly bore-sighted at installation thus the only distortion parameters to account for in registration are the differences in FOVs and resolutions. This algorithm, called the SS (sensor specifications) algorithm, performs registration by first equalizing the FOVs and then resampling the distorted image to match reference resolutions. Based upon the lessons learned from the SS algorithm, a geometric image-to-image registration algorithm was implemented. Both of these algorithms are discussed below. For each of the algorithms, we use the SWIR image as the baseline since it has the "worst" image parameters (the smallest FOV and poorest spatial resolution). The size of an image can be modified through interpolation but we cannot increase the FOV.

3.1. SS algorithm

The first step of the SS algorithm is to equalize the instantaneous FOVs (IFOV)s of the sensors. The FOV is the angular extent of the full image on the sensor and the IFOV is the angular extent on an individual detector element, i.e., the solid angle through which a detector element is sensitive to radiation.

From Figures 1, 2, and 3, we observe that the visible portion of the images is actually smaller than the full image capture window. The FOVs listed in Table 2 are assumed to correspond to the visible portion and

not the capture window. Thus, the first stage of processing is to crop the images to the visible portions. The second stage of processing is to ensure that the two images are representing the same portion of the scene. Since the FOVs of the SWIR and the LWIR sensors differ — LWIR has the greater FOV and hence captures a wider swath of the scene — the LWIR image needs to be cropped so that it encompasses the same FOV as that encompassed in the SWIR image. The dimensions of the cropped LWIR images — the number of columns and rows — are determined by a simple scaling operation. The horizontal and vertical IFOVs of the LWIR image are obtained using

$$\texttt{IFOV-LWIR-HORIZONTAL} = \frac{\texttt{FOV-LWIR-HORIZONTAL}}{\texttt{LWIR-COLS}}$$

and

$$\texttt{IFOV-LWIR-VERTICAL} = \frac{\texttt{FOV-LWIR-VERTICAL}}{\texttt{LWIR-ROWS}}$$

respectively. The number of cropped columns and rows for the LWIR image is then determined by

$$\textbf{columns} = \frac{\texttt{FOV-SWIR-HORIZONTAL}}{\texttt{IFOV-LWIR-HORIZONTAL}}$$

$$\textbf{rows} = \frac{\texttt{FOV-SWIR-VERTICAL}}{\texttt{IFOV-LWIR-VERTICAL}}$$

.

After cropping, the SWIR and LWIR FOVs are equal, but since the dimensions of the cropped LWIR are different from the dimensions of the SWIR, the IFOVs of the LWIR and the SWIR images are still different. To make the IFOVs the same we must resample the cropped LWIR image so that it is the same size as the SWIR image. This entails: (1) computing an expansion factor that will make the dimensions of the cropped LWIR image greater than the dimensions of the SWIR (2) pixel replicating the cropped LWIR based on the expansion factor and (3) downsampling the expanded LWIR image to the SWIR dimensions. We use the bi-linear interpolation method.[23] Nearest neighbor interpolation can also be selected if desired but bilinear interpolation is more spatially accurate and results in images that are slightly smoother.

A similar sequence of operations is performed between the SWIR image and the visible image. If the FOVs are the same, as in Table 1 then the visible image is simply downsampled to match the SWIR resolution. The initial results from the SS algorithm clearly indicated that the distortions present in the images were not excessive, but they also were not limited to FOV and resolution differences.

3.2. MLR algorithm

Based on the results obtained from the SS algorithm a more general, geometric image to image registration algorithm is implemented. The distortions between the images seem to be due to sensor translation, (slight) rotation, scale change, and, possibly, shear. An affine transform is, thus, used to model the spatial transformation. Control points are manually selected for identifying and matching corresponding features between the reference and distorted images. Since we assume that the sensors do not change alignment over time, we only need to register one baseline set of images that can subsequently be used for the rest of the image frames.

We use point mapping without feedback[13] to approximate the global affine transformation. The first stage of the MLR algorithm is to select a minimum of three non-collinear control points from two input images. More points can be chosen to make the coefficients more representative of the distortions throughout the overall image if the points are well distributed. Global distortion representation is also improved by choosing pixels on the perimeter if possible. The control points are then analyzed using multiple linear regression[24, 25] to approximate the coefficients of the affine transform. Residuals to determine the accuracy of the regression model obtained are calculated. The defined affine transform provides a mapping between the baseline and distorted images. The distorted image is then resampled using the transform parameters to create the registered image. Bilinear interpolation is used for resampling.

Figure 1: Original SWIR **Figure 2**: Original LWIR **Figure 3**: Original Visible

4. RESULTS

To demonstrate the performance of the algorithms we processed a set of videos taken by the EVS cameras during a flight test at Patrick Henry airport in Newport News, Virginia. The video sequence was of the B757 aircraft approach to a runway, and was digitized using a Canopus Video Board. Three images (one from each camera) time-aligned at 00:26:14:18 were used for registration. As stated earlier, the SWIR image is used as the baseline for registration since it has the poorest spatial resolution and FOV. The SWIR, LWIR and visible images are shown in Figures 1, 2 and 3 respectively.

To provide a similarity metric to validate the performance of the registration algorithms we take the absolute difference of the reference and corrected images. This provides a visible validation of the registration process since features such as runway edges should align if registration is performed correctly.

4.1. SS algorithm

Applying the SS algorithm with the SWIR image as the baseline, and the LWIR and visible images as distorted images yields the "registered" SWIR, LWIR and visible images shown in Figures 4, 5 and 6 respectively. The FOV of the LWIR image has been made smaller to match the FOV of the SWIR image. This change in FOV can clearly be seen in the horizontal direction of Figure 5, by observing that the blurred artifact (which is an antenna in the FOV of the camera) in the upper left corner of the original LWIR is now almost completely removed in the registered image. In the vertical direction, the decrease in FOV is noted by the missing timecode at the top and the missing ground features at the bottom of Figure 5 that are in the original image. The IFOVs have also been matched though resampling. The general effects of resampling can be seen by observing the expansion of image features from Figure 2 to Figure 5. The FOV of the original visible image in Figure 3 has been made slightly smaller, again to match the FOV of the SWIR, in Figure 6. Since the FOVs nearly match and the image dimensions are the same, there is only a small expansion to match IFOVs, hence the registered image features are only slightly increased from the original.

Figure 7 is the differenced SWIR and SS registered LWIR, and Figure 8 is the differenced SS registered LWIR and SS registered visible image. The misalignment between the images after registration can clearly be seen in Figure 7 by observing the difference in the outline of the runway from the LWIR component of the image, and the runway lights from the SWIR image. There is at least a large translation and a small rotation difference between the SWIR and registered LWIR. Similarly, the misalignment between the registered LWIR and visible images differenced in Figure 8 can also be seen by noting the difference in the outline of the runway from the LWIR image, and the runway lights from the visible image. Again, there is an obvious translation between the images. Figures 7 and 8 clearly display the misalignment between the images thus indicating that differences in sensor design characteristics are not the only cause of distortion between the images.

4.2. MLR algorithm

We first applied the MLR algorithm to the original (uncropped) SWIR and visible images, again using the SWIR as the baseline. Due to the lack of features around the perimeter of the SWIR image we used the runway lights as control points. Note that we are only using three control points for demonstration purposes. Figure 9

Figure 4: Cropped SWIR   Figure 5: SS Reg. LWIR   Figure 6: SS Reg. visible

Figure 7: SWIR and SS Registered LWIR   Figure 8: SS Registered LWIR and visible

repeats the original SWIR image for reference. Figure 10 shows the registered visible image and Figure 11 is the differenced SWIR and registered visible image. The coefficients obtained are given in Table 3.

	b_0	b_1	b_2
x'	-0.546156	1.021212	-0.004578
y'	-20.440557	-0.007477	0.972837

Table 3: Visible to SWIR MLR Coefficients

A close look at the runway and the runway lights in the two images shows that they are now registered. In particular, in the lower right corner of the SWIR image there are four runway lights lined up horizontally. In the visible image there are three runway lights in the same position, except the second light from the left is not visible. Figure 11 shows the four lights differenced in a horizontal line with the missing visible light filled in from the SWIR image. It is clear to see the warp performed during registration by observing the timecode size and location differences in the differenced images. The timecodes are the same size and at the same location in the original images. Figure 9 and Figure 10 could now be equally cropped to remove disjoint pixels around the perimeter to obtain the final images to be fused.

Next we applied the MLR algorithm to the registered visible and LWIR images using the visible image as the baseline. Since the runway lights are not visible in the LWIR image, we use the intersecting lines at the bottom and top of the runway, and a stripe at the beginning of the runway towards the right in the LWIR image as control points. Figure 12 repeats the MLR registered visible image for reference. Figure 13 is the registered LWIR image and Figure 14 is the differenced registered visible and the just registered LWIR images. The coefficients obtained are given in Table 4.

Figure 9: Repeated SWIR

Figure 10: MLR Reg. visible

Figure 11. SWIR and MLR Reg. visible

	b_0	b_1	b_2
x'	8.347350	0.850628	0.037684
y'	9.637629	-0.012015	0.779082

Table 4: LWIR to visible SWIR MLR Coefficients

As is evident in Figure 14 the runway portion of the LWIR image is aligned with the runway portion of the registered visible image. Also, the runway lights from the visible image border the perimeter of the LWIR runway. The one runway stripe selected as a control point is aligned. The taxiways on the right side of the image and the horizon across the image are also aligned. Again, any disjoint pixels around the perimeter could be removed by cropping. At this point all three original images are registered.

As a final test of the MLR algorithm we applied the same control point coefficients to a later frame in the video sequence. Figures 15, 16 and 17 are the SWIR, LWIR and visible images at time 00:26:14:28, 10 seconds later in the sequence. Figure 18 shows the MLR registered visible image. Figure 19 is the differenced SWIR and MLR registered visible images. Figure 20 is the MLR registered LWIR image. Figure 21 is the differenced MLR registered visible and MLR registered LWIR image. As in the previous set of images, the registration can be observed by noting the alignment of the runway and runway lights in Figures 19 and 21.

4.3. Discussion

The registration images shown visually demonstrate the performance of the two algorithms on typical image data from the EVS. The registration inaccuracies of the SS algorithm are obvious. The differing FOV and resolution specifications given do not take into account the other distortions within the images. With this much discrepancy there seems to be either a fundamental problem in the bore-sighting or alignment of the cameras, or the alignment is changing during flight. If the sensors were actually bore-sighted and aligned, the

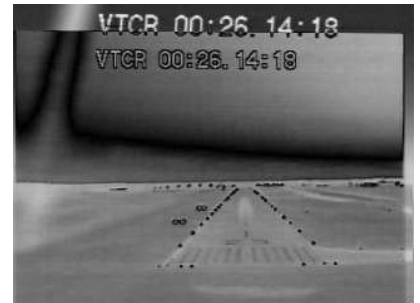**Figure 12.** Repeated MLR Registered visible

Figure 13: MLR Reg. LWIR

Figure 14. MLR Reg. visible and LWIR

Figure 15. Orig. SWIR at Time 26:14:28

Figure 16. Orig. LWIR at Time 26:14:28

Figure 17. Orig. visible at Time 26:14:28

Figure 18. MLR Registered visible at Time 26:14:18

Figure 19: SWIR and MLR Registered visible

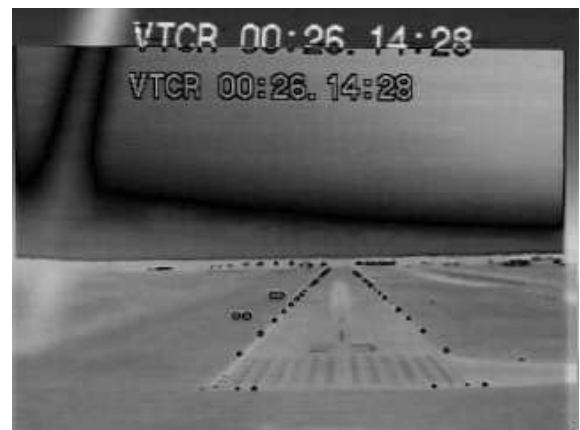**Figure 20.** MLR Registered LWIR at Time 26:14:18

Figure 21: MLR Registered visible and LWIR

SS algorithm should be able to match the performance of the MLR algorithm and in addition, not require any manual intervention like selection of control points. True FOV values could be obtained from a thorough ground calibration, and non-interpolated pixels of the actual image dimensions could be obtained from raw digital data streams from the cameras.

The MLR algorithm provides better registration of the images than the SS algorithm configured with the current set of specifications and with the current EVS alignment. In our examples the runway and runway lights are clearly aligned. The coefficients obtained with only three points indicates that there are rotation, translation, scale and possibly shear distortion components found between the images. These distortions can be seen by viewing the timecode warps at the top of the differenced images. The application of the same MLR control points to a set of time-aligned images later in the same video sequence produced the same level of registration. This indicates that we could successfully use the registration coefficients obtained from one set of time-aligned images to apply to, at least, a group of frames from the video sequence. If the alignment is not changing substantially during flight then all frames could be processed with the same transform.

5. CONCLUSIONS

Image registration is an essential prerequisite to subsequent image fusion. This research effort produced two algorithms to perform multi-image registration for the EVS. The SS algorithm uses EVS camera specifications and performs registration based solely on these parameters. The performance of this algorithm indicates that there is a severe inaccuracy in the boresighting or alignment of the cameras. Correction of these issues should improve the performance of the algorithm and allow it to be used to automatically register all images across the cameras, or as validation of the MLR algorithm.

The MLR algorithm uses control point selection and linear regression to compute the coefficients of an affine spatial transformation. This transformation is then used to register the LWIR and visible images to the SWIR image. In addition, the MLR registration algorithm provides a means to generate a base set of coefficients for post processing of the full video stream across all cameras. We have subsequently used a set of baseline coefficients to process an entire 20 second video clip from each of the three cameras.

In addition, the coefficients obtained could also be used to back out the actual distortion values (translation amount, rotation angle, etc.) for feedback to the EVS designers. Improvements could also be made in the computation of the coefficients by using point selection with feedback or other more robust feature selection mechanisms. Manual control point selection can be improved by MSR enhancement of the images to emphasize and sharpen features prior to registration. This was done for another EVS data set and greatly improved the ability to select corresponding points. Most importantly, the actual boresighting and alignment can be checked against the values obtained from MLR and SS registration, and adjusted appropriately. This procedure could be performed both before, and after, EVS flight opportunities and used to verify and validate system alignment.

REFERENCES

1. C. Tiana, J. Kerr, and S. Harrah, "Multispectral uncooled infrared enhanced vision system for flight test," in *Proceedings of SPIE*, **4363**, April 2000.
2. D. J. Jobson, Z. Rahman, and G. A. Woodell, "Properties and performance of a center/surround retinex," *IEEE Trans. on Image Processing* **6**, pp. 451–462, March 1997.
3. D. J. Jobson, Z. Rahman, and G. A. Woodell, "A multi-scale Retinex for bridging the gap between color images and the human observation of scenes," *IEEE Transactions on Image Processing: Special Issue on Color Processing* **6**, pp. 965–976, July 1997.
4. B. D. Thompson, Z. Rahman, and S. K. Park, "Retinex preprocessing for improved multi-spectral image classification," in *Visual Information Processing VIII*, Proc. SPIE 3716, 1999.
5. B. D. Thompson, Z. Rahman, and S. K. Park, "A multi-scale retinex for improved performance in multi-spectral image classification," in *Visual Information Processing IX*, Proc. SPIE 4041, 2000.
6. Z. Rahman, D. Jobson, G. Woodell, and G. Hines, "Multi-sensor fusion and enhancement using the retinex image enhancement algorithm," in *Visual Information Processing XI, Proceedings of SPIE 4736*, April 2002.

7. Rahman. see http://dragon.larc.nasa.gov/fog_haze for examples.

8. R. Brooks and S. Iyengar, *Multi-Sensor Fusion: Fundamentals and Applications*, Prentice Hall, 1998.

9. R. Luo and M. Kay, *Multisensor Integration and Fusion for Intelligent Machines and Systems*, Ablex, Norwood, NJ, 1995.

10. H. Li and Y. Zhou, "Automatic visual/ir image registration," in *Optical Engineering*, **35(2)**, pp. 391–400, February 1996.

11. R. Sharma and M.Pavel, "Multisensor image registration," in *SID Digest, Society for Information Display*, **XXVIII**, pp. 951–954, May 1997.

12. H. Li, B. Manjunath, and S. Mitra, "A contour-based approach to multisensor image registration," in *IEEE Transactions on Image Processing*, **4, No. 3**, March 1995.

13. L. Brown, "A survey of image registration techniques," in *ACM Computing Surveys*, **24, No.4**, December 1992.

14. R. Schowengerdt, *Remote Sensing: Models and Methods for Image Processing*, Academic Press, 1997.

15. K. Novak, "Rectification of digital imagery," in *Photogrammetric Engineering and Remote Sensing*, **58, No. 9**, pp. 399–344, March 1992.

16. A. Goshtasby, "Image registration by local approximation methods," in *Image Vision Computing*, **6**, pp. 255–261, November 1988.

17. G. Wolberg, *Digital Image Warping*, IEEE Computer Society Press, 1990.

18. L. Fonseca and B. Manjunath, "Registration techniques for multisensor remotely sensed imagery," in *Photogrammetric Engineering and Remote Sensing*, **62, No. 9**, pp. 1046–1056, September 1996.

19. J. Jensen, *Introductory Digital Image Processing*, Prentice Hall, 1996.

20. N. Dodgson, "Image resampling," Tech. Rep. 261, University of Cambridge, United Kingdom, August 1992.

21. R. Gonzalez and R. Woods, *Digital Image Processing*, Addison-Wesley, 1993.

22. S. Park and R. Schowengerdt, "Image reconstruction by parametric cubic convolution," in *Computer Vision, Graphics, and Image Processing*, **23**, pp. 258–272, 1983.

23. H. Burdick, *Digital Imaging*, McGraw Hill, 1997.

24. G. Snedecor, *Statistical Methods*, Iowa State University, 8 ed., 1989.

25. D. Wittink, *The Application of Regression Analysis*, Allyn and Bacon, 1988.