In proc. SPIE Vol 7835. (c) 2010 SPIE Digital Library. http://dx.doi.org/10.1117/12.865034

Co-aligning aerial hyperspectral push-broom strips for change detection

Erik Ringaby[a], Jörgen Ahlberg[b], Niclas Wadströmer[b], Per-Erik Forssén[a]

[a]Computer Vision Lab, Linköping University, Sweden;
[b]FOI, Swedish Defence Research Agency, Linköping, Sweden

ABSTRACT

We have performed a field trial with an airborne push-broom hyperspectral sensor, making several flights over the same area and with known changes (e.g., moved vehicles) between the flights. Each flight results in a sequence of scan lines forming an image strip, and in order to detect changes between two flights, the two resulting image strips must be geometrically aligned and radiometrically corrected. The focus of this paper is the geometrical alignment, and we propose an image- and gyro-based method for geometric co-alignment (registration) of two image strips. The method is particularly useful when the sensor is not stabilized, thus reducing the need for expensive mechanical stabilization. The method works in several steps, including gyro-based rectification, global alignment using SIFT matching, and a local alignment using KLT tracking. Experimental results are shown but not quantified, as ground truth is, by the nature of the trial, lacking.

Keywords: Hyperspectral, push-broom, co-alignment, registration, change detection

1. INTRODUCTION

Hyperspectral sensors are electro-optical sensors (cameras) that sample the incoming radiation (light) in a large number (tens, hundreds) of narrow spectral bands. A common design principle is to let the sensor observe one line of the world at each time instant. This line is split into the various wavelengths by the optics, and sampled by a rectangular sensor chip. The resulting image then has one spatial and one spectral axis. In order to create a data cube with two spatial axes, the sensor is equipped with a line scanner, i.e., a moving mirror, allowing the sensor to scan the scene.

When installed nadir-looking on an aircraft, the sensor scans the land surface by exploiting the ego motion of the aircraft instead of using a line scanner, i.e., as a push-broom sensor. The sensor collects a large number of *scan lines*, forming an image *strip*. Ideally, these scan lines are parallell and equidistant. In practice, they are not, due to the varying speed and angle of the aircraft. Mechanical sensor stabilization, making the sensor to be constantly nadir-looking even when the aircraft is slightly tilted, will align the scan lines to some degree.[1,2]

If the 3D position and rotation of the aircraft is exactly known, and an elevation map of the ground is available, the position of each pixel's footprint can be calculated (georeferenced). Navigation aids such as GPS and gyro are often available to determine the aircraft's position, and mechanical sensor stabilization might, as mentioned, simplify the problem.

If the targeted application is object/target detection, in the form of anomaly detection or signature-based detection,[3] exact georeferencing of each pixel is not always necessary. For example, assume a pixel footprint size of half a meter. To detect a vehicle with a position error of, say, 10 pixels (5 meters), is usually well within the required precision. However, in order to detect changes between two strips, the co-alignment (registration) error

Further author information: (Send correspondence to E.R.)
E.R: E-mail: ringaby@isy.liu.se, Telephone: +46 (0)13 28 13 02
J.A.: E-mail: jorgen.ahlberg@foi.se, Telephone: +46 (0)13 37 80 68
N.W.: E-mail: niclas.wadstromer@foi.se, Telephone: +46 (0)13 37 85 79
P.-E.F: E-mail: per-erik.forssen@liu.se, Telephone: +46 (0)13 28 56 54

Figure 1. The sensor installation in the aircraft. Left: The sensors are installed in the metal box beside the pilot. Right: Computer and storage equipment in the back seat.

must be much smaller, preferrably subpixel size. This is hard to achieve even with navigation aids mentioned above, and, of course, even harder if these are not available or fail.

The work in this paper is motivated by such a situation; the sensor was installed without stabilization in a small (and wobbly) aircraft. GPS and gyro were available, but in some of the flights the GPS refused to deliver data, without any apparent reason except Murphy's law. We have thus studied how well we can co-align two strips using image- and gyro-based methods. Note that this is a quite different problem than the image registation problem extensively studied in the literature, where images from a staring sensor are to be registered (popularly called "stitching"). The problem at hand is more similar to stabilization of rolling shutter video.[4]

1.1 Outline

The outline of the paper is as follows. The scenario, the sensor, and the data are described in section 2, the co-alignment method is described in section 3, results are shown in section 4, and conclusions are drawn in section 5.

2. SENSORS AND DATA

2.1 Sensor

The used imager, an ImSpec sensor, is a visual and near-infrared (391–961 nm) hyperspectral imager from SpecIm,[5] with 1024 pixels in each scan line and a maximum of 240 spectral bands. Due to limitations in the read-out electronics' data rate, the number of spectral bands might need to be reduced to meet requirements on the number of lines to be acquired per second. In this experiment 60 spectral bands were recorded, which is more than enough for our applications. The dynamic range is up to twelve bits, but in this case the used dynamic range was only around eight bits. The sensor's integration time was set to 7 ms and the scan line frequency was set to 100 Hz. The IFOV is approximately 0.51 mrad and the FOV 29.9 degrees (1024 pixels).

2.2 Data collection

The imager was mounted nadir-looking in a small fixed-wing aircraft as shown in Fig. 1. The aircraft was equipped with gyro and GPS, but the GPS data was not logged to the hard disk during all flights. At a flight speed of 166 km/h and an altitude of 1000 meters, the distance between the pixels is about 0.51 and 0.46 meters in the flight direction and in the cross-flight direction respectively. At 400 meters, the distances are 0.2 and 0.46 meters.

Data was collected at the Swedish Army Ground Combat School's premises at Kvarn outside Linköping, Sweden. Data was collected over two areas, about 1800 meters long and 580 meters wide at a flight altitude of 1000 meters (and 230 meters wide at 400 meters altitude). In one area, four military and two civilian vehicles were placed and moved between two positions between the data acquisitions. In the other area, around 20 boards

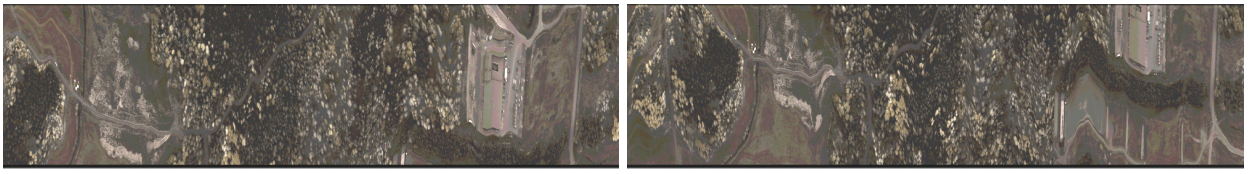

Figure 2. Two strips from approximately the same area.

of sizes from 7×7 m to 0.2×0.2 m large were placed and moved between two positons between the acquisitions. The data set is available for download at www.foi.se/hsi/. The data can be used freely for non-commercial use provided that the source is acknowledged. Two data strips are shown in Fig. 2.

3. CO-ALIGNMENT METHOD

The co-alignment algorithm can be divided into a number of steps:

1. Rectification using gyro data
2. Radiometric correction
3. Global strip alignment
4. Local alignment
   (a) Local point correspondences search
   (b) Estimation of a sparse grid of affine transformations
   (c) Transformation parameter interpolation and extrapolation
   (d) Co-alignment using interpolated transformation map

In this paper, the radiometric correction is omitted.

3.1 Rectification using a differential gyro

In contrast to systems such as Hymap[1] our sensor is mounted without mechanical stabilisation. Instead it is rigidly attached to a gyro, and the gyro signal is used to afterwards compensate for the rotation component of the aircraft wobble. This requires that both the the gyro signal and the hyperspectral sensor have accurate timestamps, and that the gyro signal is expressed in the same coordinate frame as the images (in practice we need to know the relative orientation of the two frames in 3D).

Our gyro outputs a differential signal, which needs to be integrated and subsequently bias corrected before it can be used. For passages where the gyro signal is absent, we simply set the gyro derivative to zero. We feed the aligned and integrated gyro signal into a rotation correction method that we have previously developed for cellphone imagery.[4] The result of this gyro-based correction is shown in figure 3.

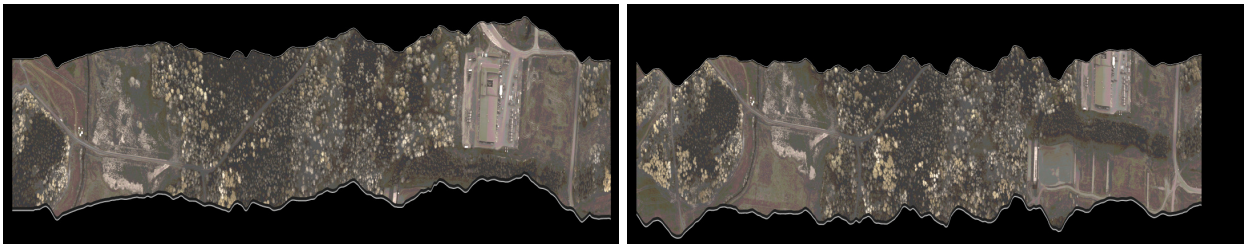

Figure 3. The two strips from figure 2, after rectification using gyro data.

Note that the gyro-based compensation removes errors due to aircraft rotation, but not errors due to an uneven translation component.

3.2 Global strip alignment

After the gyro rectification the strips are more similar, but due to different flight paths, they still need to be aligned to each other to enable comparison.

This is done globally by estimating a homography between the two strips. First, scale-invariant feature transform (SIFT) features[6] are detected in both images. SIFT is an algorithm that converts image patches to 128-dimensional descriptor vectors. It finds extrema in scale-space to detect interest points and calculates gradients and assigns orientations to them. The descriptor is created by accumulating orientations into histogram bins based on both orientation and position. SIFT operates on grey scale images and a mean of three wavelengths approximately corresponding to red, green and blue have been used.

When the SIFT descriptors have been calculated in both images, those that lie near the border are removed. This is done because they cause many false matches.

The SIFT features with the lowest ratio score (ratio between best and second best match) are used as tentative correspondences as suggested by Lowe.[6] These are then fed into RANSAC[7] which estimates a homography. The descriptors considered as inliers are shown in figure 4. As can be seen in the figure, the descriptors only lie in the left part of the image and only points from this part are used to calculate the homography. This means that the alignment will be worse in the right part of the image. The homography is applied to the first image and the result can be seen in figure 5, top. In figure 5, bottom, the first image is shown as red and the second image shown as blue. The alignment works pretty well to the left, but as can be seen it fails to the right.

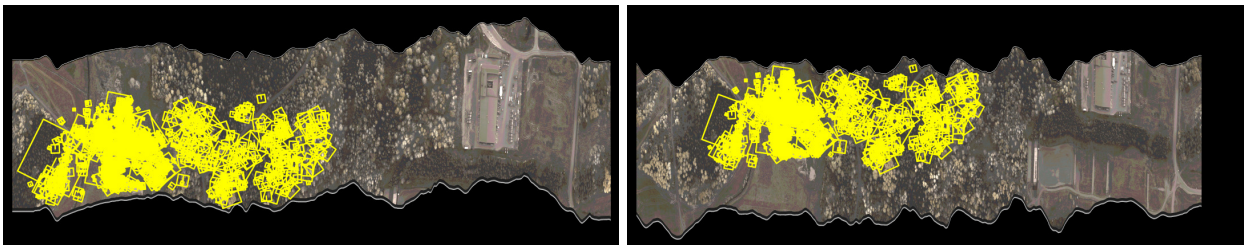

Figure 4. SIFT features consistent with homography

3.3 Local alignment

The reason why the global alignment does not work may be due to inaccuracy and drift in the gyro. The correction described in section 3.1 only corrects for rotations, and may cause some errors if the aircraft had different velocities while acquiring the different strips or if the velocity was non-constant. To compensate for this, the alignment needs to vary over the strip. A regular grid is placed on the unrectified strip and transformed to the new coordinate system from the gyro-based correction section 3.1, see figure 6.

Each grid point position in the first image can then be coarsely transformed to the second image using the global homography described in section 3.2. To improve the accuracy, the same procedure as for the global homography is used except that only a part of each image around the point is used.

When each grid point in the first image has a well estimated position in the second image, an image patch is extracted at each point. The KLT-tracker[8,9] is then used on the image patches to calculate a dense point-to-point correspondence within the patch. The KLT tracker has a much higher spatial accuracy than SIFT, but is sensitive to rotations and large translations so it can not be used initially. We use a track and re-track procedure which means that when a point is tracked from the first image patch to the other one, it is tracked back to the first image patch again, and if it returns to the original position the point is regarded as a match, otherwise it is discarded.[10]

After this step we have accurate point correspondences that are used to calculate a local affinity around the current point in the strip. We use an affine transformation here, instead of the homography used in section 3.2, as the homography easily becomes unstable when the support is small. Due to non-overlapping strip areas or insufficient number of point correspondences the local affinity can be difficult to calculate accurately. If the

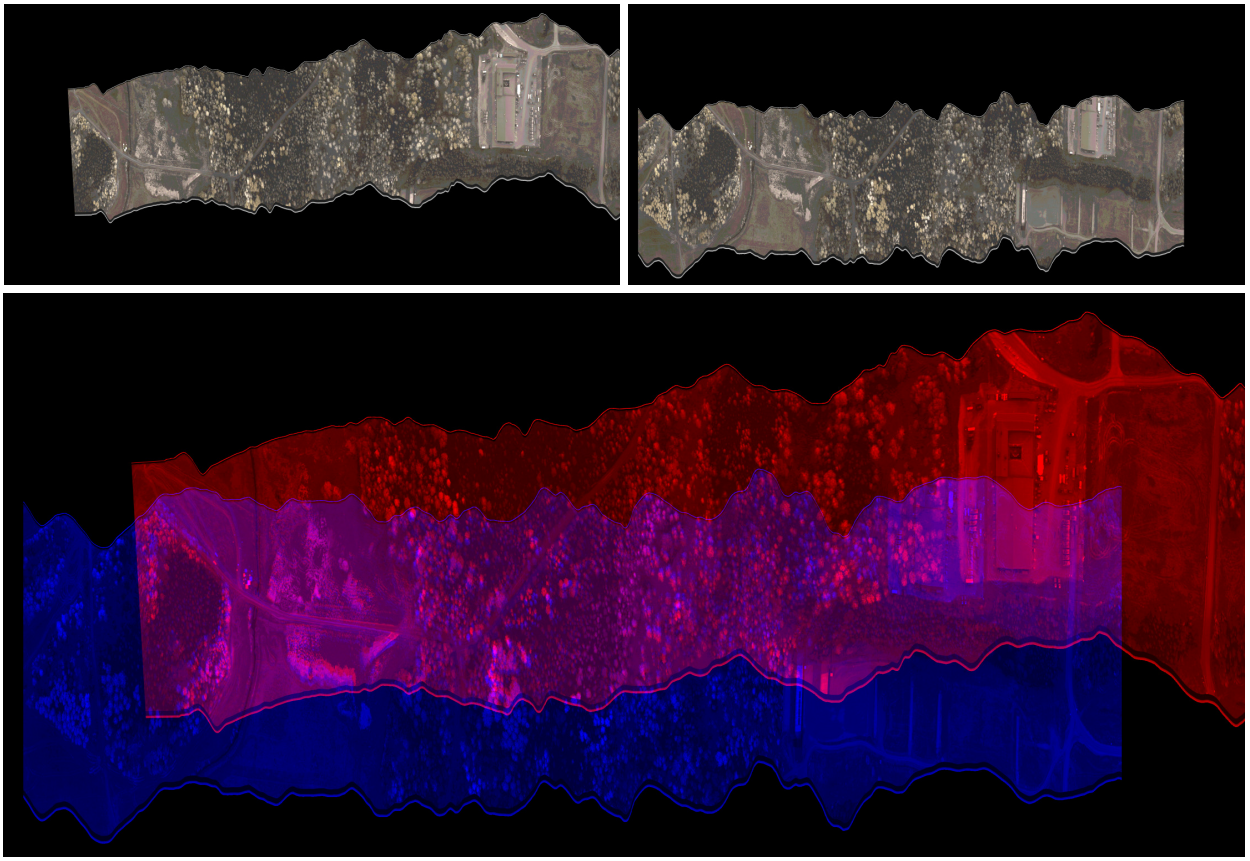

Figure 5. Global alignment

number of points that satisfy the affinity is below a certain threshold, the affinity is rejected. This is indicated with red plus signs in figure 6, and the yellow circles are the accepted affinities.

The grid used in the previous step is sparse and in order to get values for all pixels, the values describing the local affinities are interpolated in a cubic manner. Only pixels within the convex hull of the points can be interpolated. Values outside are extrapolated by constant propagation of the affinity parameters away from the edge of the convex hull.

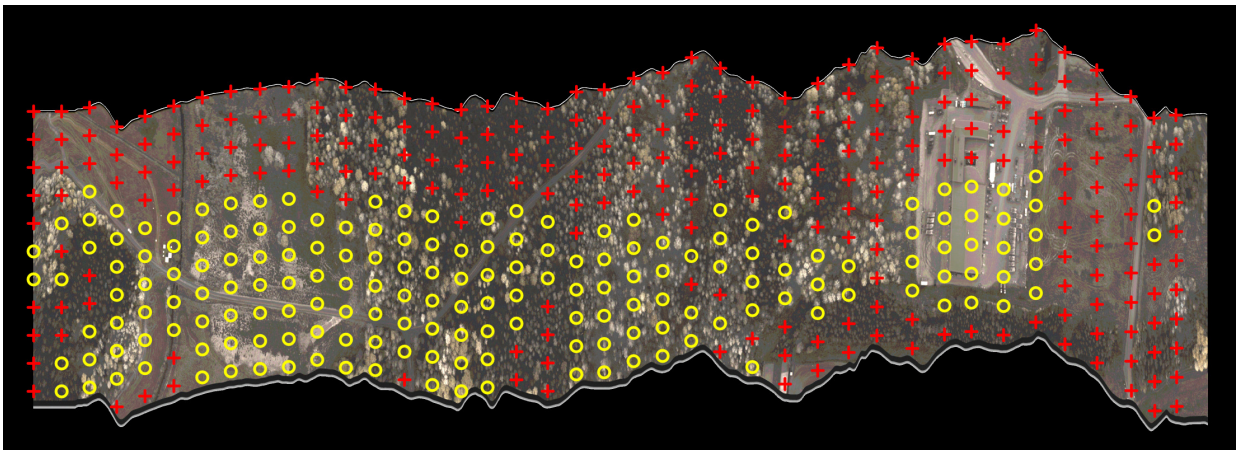

Figure 6. Positions where local affinities are calculated

4. EXPERIMENTAL RESULTS

When the local affinities have been calculated and interpolated the two strips can be aligned. The result of this procedure is shown in figures 7 and 8. If compared with 5 we can see that the alignment is now fairly accurate. Note however that in areas where the gyro based rectification had problems, the alignment is still not perfect. Near the edges of a strip, the alignment accuracy is also degraded, the reson for this is that here extrapolation rather than interpolation has been used to find the affinity parameters.

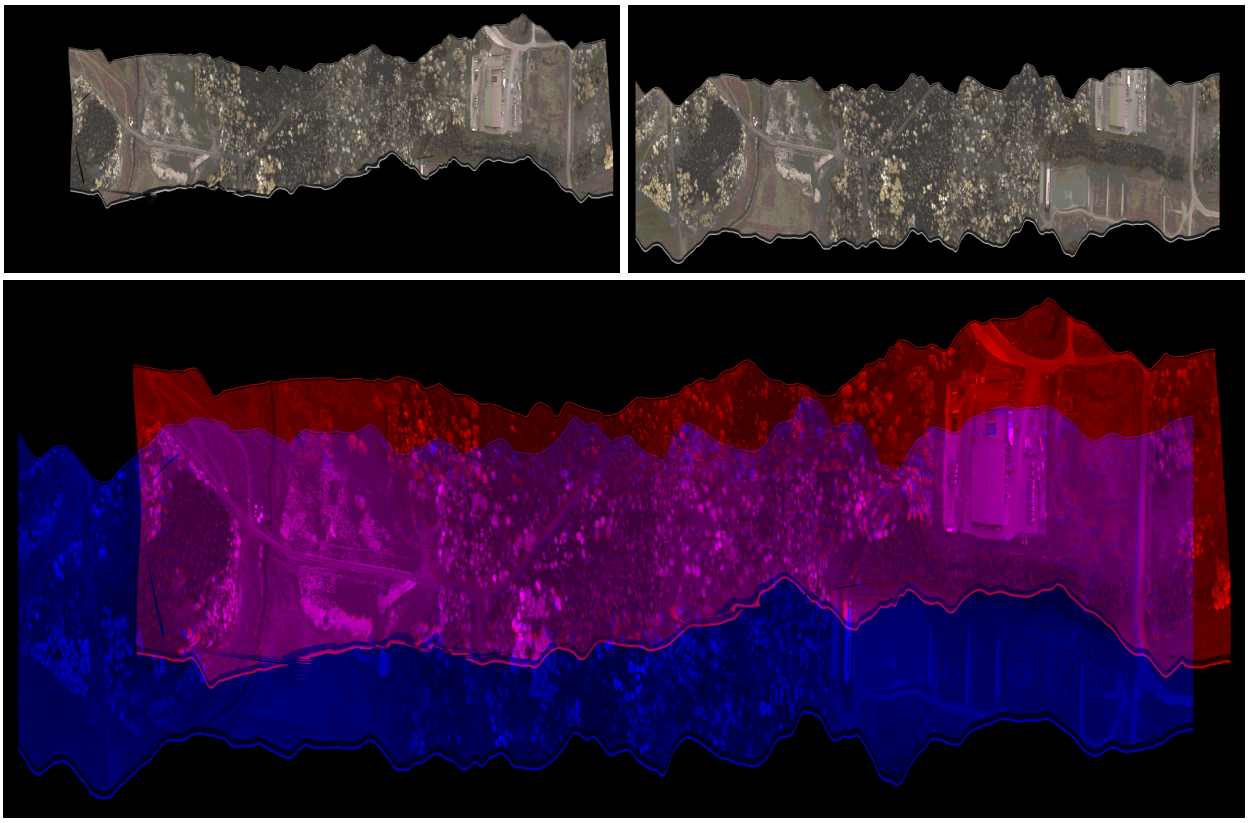

Figure 7. Result of the proposed co-alignment procedure.

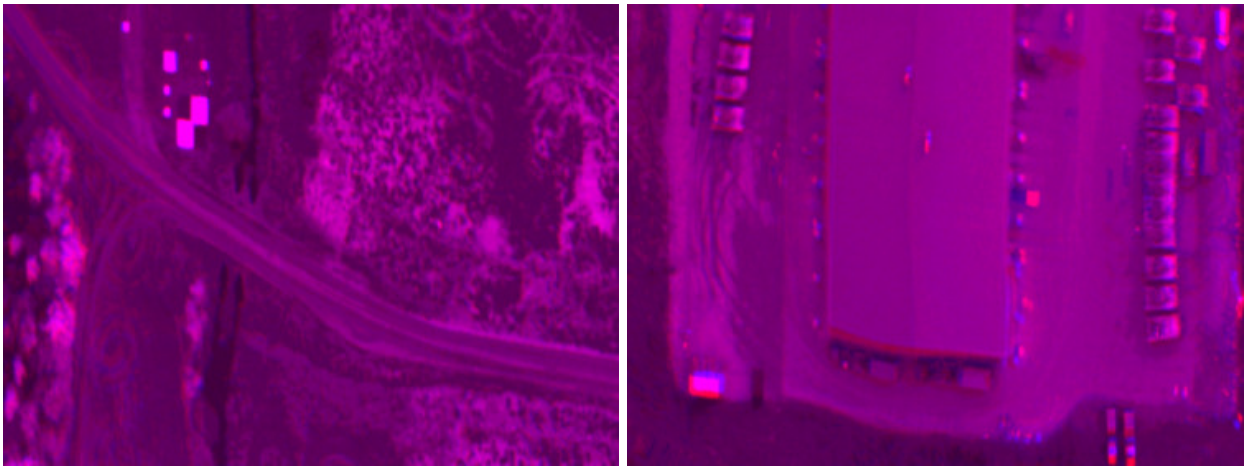

Figure 8. Result of the proposed co-alignment procedure, zoomed in.

5. CONCLUSION

We have shown that we are able to co-align the scan lines of two partly overlapping data strips from flights with a hyperspectral sensor. The method relies on gyro based rectification post capture, and does not require mechanical stabilisation.

ACKNOWLEDGEMENT

The authors would like to thank Dr. Thomas Svensson and his colleagues at the IR Systems Group at FOI for the data. The flights were performed by SkyMovies AB with Dr. Svensson in the back seat.

REFERENCES

[1] Kruse, F. A., Boardman, J. W., Lefkoff, A. B., Young, J. M., Kierein-Young, K., Cocks, T. D., Jenssen, R., and Cocks, P. A., "HyMap: An Australian hyperspectral sensor solving global problems – Results from USA HyMap data acquisitions," *Proc. Australasian Remote Sensing and Photogrammetry Conference*, (2000).

[2] Snyder, D., Kerekes, J., Fairweather, I., Crabtree, R., Shive, J., and Hager, S., "Development of a web-based application to evaluate target finding algorithms," *Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, **2**, 915–918 (2008).

[3] Ahlberg, J. and Renhorn, I., "Multi- and hyperspectral target and anomaly detection," Scientific report FOI-R--1526--SE, Swedish Defence Research Agency (FOI) (2004).

[4] Forssén, P.-E. and Ringaby, E., "Rectifying rolling shutter video from hand-held devices," *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (2010).

[5] Chevalier, T., "Samverkande sensorer. Illustrerat exempel med hyperspektral kamera och 3D-laserradar," Tech. Rep. FOI-R--2325--SE, Swedish Defence Research Agency (FOI) (2007).

[6] Lowe, D. G., "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision* **60**(2), 91–110 (2004).

[7] Hartley, R. and Zisserman, A., [*Multiple View Geometry in Computer Vision*], Cambridge University Press (2000).

[8] Lucas, B. and Kanade, T., "An iterative image registration technique with an application to stereo vision," *Proc. International Joint Conferences on Artificial Intelligence*, 674–679 (1981).

[9] Shi, J. and Tomasi, C., "Good features to track," *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (1994).

[10] Baker, S., Scharstein, D., Lewis, J. P., Roth, S., Black, M. J., and Szeliski, R., "A database and evaluation methodology for optical flow," *Proc. IEEE International Conference on Computer Vision (ICCV)*, (2007).