Data Quality Report - 2014

Hyperspectral

ARSF - Data Analysis Node

Updated on: June 3, 2014

Contents

1	Overview					
2	Geo-referencing accuracy					
3	Timing Errors					
4	Sensor calibration					
	4.1	Winte	r 2014 calibration	2		
		4.1.1	Wavelength calibration accuracy	2		
		4.1.2	Radiometric calibration	3		
5	Overflowed Pixels 5					
6	Smear Correction					
7	Bac	I CCD	Pixels	8		
	7.1	Eagle	Sensor SN110001	8		
	7.2	Fenix	Sensor	10		
		7.2.1	Detection method A - Constant input variable output .	11		
		7.2.2	Detection method B - Constant input constant invalid			
			output (CICO)	11		
		7.2.3	Detection method C - Linear input non-linear output			
			(LINO)	11		
		7.2.4	Detection method D - Rapid saturation	12		
		7.2.5	Detection method E - Visual inspection	12		

1 Overview

This report describes issues that should be considered when further processing any of the 2014 Airborne Research Survey Facility (ARSF) datasets. The document may be updated over the course of the year, with the latest version available at:

http://arsf-dan.nerc.ac.uk/trac/wiki/Reports

2 Geo-referencing accuracy

ARSF currently deliver data at level 1 (calibrated sensor data) and are now also delivering mapped level 1 data (level 3). This allows users to get quick access to georeferenced data but also maintains the capability of being able to apply user-developed algorithms and generating level 2 products (e.g. atmospherically corrected radiances) prior to mapping to a projection or datum that suits.

The quality of the geocorrection for each project is described in the documentation supplied with the project and is normally of the order of a couple of metres (approximately 1 pixel). Where a vector overlay or other ground truth information is available, ARSF provide an indication of the average error, included in the screenshot images. If you need higher accuracy, please contact us at: arsf-processing@pml.ac.uk. It may be possible to tune specific flight lines for higher accuracy or we can provide instructions on how to make your own alignments.

3 Timing Errors

Due to an error in the handling of synchronisation between the navigation system and the Specim sensors, small timing errors (order of 0.05s for Eagle and Hawk, 1s for Fenix) may occur. The consequence of timing errors is to cause scan lines to be positioned incorrectly and manifest visually as "wobbles" in the imagery. The wobbles are correlated to, but out of sync with, movements of the aircraft. An example is shown in Figure 1 and Figure 2 below.

This issue has been extensively investigated and demonstrated to be a fault in the Specim systems. Specim are working with ARSF to provide upgrades and improvements to correct this issue, but have not yet succeeded.

Therefore we endeavour to correct all timing errors prior to delivery. As this is a manual process and relies on finding suitable visible features in the imagery, some errors may still remain. If any are found, please contact us at arsf-processing@pml.ac.uk.



Figure 1: timing error in an Eagle line



Figure 2: corrected version of above (0.13 seconds difference)

4 Sensor calibration

Calibration of the Fenix sensor is undertaken annually at ARSF's calibration lab at Gloucester Airport in collaboration with the NERC Field Spectroscopy Facility. The old Eagle and Hawk instruments have been retired from service and replaced with the Fenix, which covers the spectral range previously covered separately by the Eagle and Hawk.

4.1 Winter 2014 calibration

4.1.1 Wavelength calibration accuracy

Wavelength calibration was undertaken by viewing a number of spectral lamps using the Fenix. The lamps viewed were: Hg-Ar, He, H, Kr, O, Ne and CO2. These provide a number of spectral emission features at known wavelengths that can be seen with the sensor, this allows specific pixel numbers to be confirmed as viewing particular wavelengths.

Wavelength calibration was undertaken in February 2014. Differences from known spectral features can be found in table 1.

Spectral	Measured Wave-	FWHM (nm)	Error (nm)	Pass/Fail (<2nm Err)
404.7	404.81	3.16	-0.15	Pass
435.8	436.13	2.86	-0.29	Pass
546.1	546.19	2.65	-0.12	Pass
578.0	578.17	4.13	-0.17	Pass
667.8	667.81	2.99	-0.01	Pass
706.5	706.69	3.21	-0.19	Pass
1083.0	1083.09	9.62	-0.09	Pass
1181.9	1181.94	10.06	-0.04	Pass
1363.4	1364.58	13.39	-1.18	Pass
1442.7	1441.07	15.18	1.63	Pass
1816.7	1816.82	11.13	-0.12	Pass
1875.1	1875.37	11.98	-0.27	Pass
2058.7	2057.73	10.11	0.97	Pass
Means:		7.73	0.00	Pass

Table 1: Wavelength calibration offsets for Fenix, February 2014 calibration. Horizontal line denotes break between VNIR and SWIR sections.

The FWHM values obtained in these tables have been obtained by fit-

ting gaussian curves to the spectral features and then measuring the best-fit FWHM. While it is believed that this gives a reasonable estimate of the FWHM values of the sensors, this has not been conclusively demonstrated - use with caution. Note also that the FWHM as labelled in the data header (hdr) files is actually bandwidth of each band.

4.1.2 Radiometric calibration

Following the wavelength calibration, radiometric calibration is undertaken by viewing an integrating sphere provided by the NERC Field Spectroscopy Facility using the Fenix. The sphere is calibrated to NPL standards. This provides a light source of known radiance at each wavelength in use, allowing a calibration curve for the instruments to be calculated.

Radiometric calibration of the Fenix was undertaken both in December 2013 and in February 2014. The sensor was found to be generally radiometrically stable over that period (average 0.15% change), though in water absorption bands in the SWIR the stability was slightly poorer (2% change). Data from these bands should not be used in any case, since the water in the atmosphere will block the signal. A graph showing the change between the two calibrations is shown in figure 3.



Figure 3: Fenix calibration multiplier percentage differences between Dec. 2013 and Feb. 2014 calibrations

5 Overflowed Pixels

The instruments have a limited dynamic range and must be set to capture data over the appropriate range of signal strength. For example, if the area of interest is dark, then the instrument will be configured to capture as much low light detail as possible. This configuration is set based on operator experience, the principal investigator's indication of the areas of importance and the prevailing conditions. Inevitably, some pixels are unexpectedly bright - e.g. sunglint over water or part of a cloud. These pixels may exceed the maximum capture level and overflow. Typically they are not in areas of interest, but should be accounted for. The accompanying mask file will contain an overflow flag value in the level 1 equivalent pixel.

In Fenix data, overflows are marked for just the pixel/band in question. However, Eagle uses a frame transfer CCD, where data are read out in rows. Incoming light continues to accumulate in unread rows during the transfer and is removed by "smear correction" software, which relies on data from one row to correct the next. If a pixel overflows, information is lost and all subsequent pixels in that column cannot be fully corrected. In Eagle, the net effect is that an overflow at 600nm will cause all bluer bands (600nm -> 400nm) to be under-corrected for that spatial pixel. In this case, the mask file will contain a "smear affected" flag value for the equivalent pixel position. When Eagle data with overflows are delivered, we mask all bands (in the mask file) following an overflow as they will incorporate some unknown additional light. If you would prefer your actual level 1 files to be masked out rather than use the separate mask file please contact arsf-processing@pml.ac.uk.

Figures 4 and 5 show an Eagle level 1 band and equivalent mask band in the blue part of the spectrum (450nm). Note that although these data may appear good in the level 1 image, the bright mask values mark pixels that have been adversely affected due to overflows occurring in a higher band.

6 Smear Correction

The Eagle uses a CCD that shifts data out line by line at the end of a frame. While this readout process is quick, additional light still falls onto the detector during the readout period. Currently this is corrected for by subtracting a small amount of light measured in the previous line(s) as they are read-out. This procedure assumes the light input is unchanged during the integration and read-out, but this is a good approximation. However, the sensor is often run with a bandset that doesn't record all of the lines.

The Eagle CCD is 1024×1024 , with nominal sensitivity from ~200nm to 1200nm, with readout progressing from red (1200nm) to blue (200nm). In operation, only the middle ~500 bands are recorded (~450-950nm), partly due to low sensitivity in the other regions, but also because there are significant internal reflections/second order effects (which is normal). Figure 6 shows a view of the amount of light falling on different parts of the detector. The red box shows the approximate area that is recorded in normal conditions. Internal reflections can be clearly seen, although they have been highly enhanced to make them visible. The amount of light in the central region greatly overwhelms that of the reflection, although their contribution to error can still be significant in weakly illuminated bands.

Consequently, any light falling in the 950-1200nm area cannot be corrected for and smear resulting from that light will remain in the final image. This erroneous light will have the greatest effect where the "true" signal is lowest (e.g. absorption bands). We have run some simulations and the error is naturally worse in bands where there is little light (the red and blue



Figure 4: level 1 image of bright ice at 450nm



Figure 5: mask image of same band as Figure 4. White pixels are adversely affected.

ends of the spectral range, where signal versus noise is lowest, which possibly contributes to why these are poorer quality) and worse the shorter the integration time (below 10ms integration times, errors rapidly increase).

Figure 7 shows estimates of the error introduced by smear into a realworld dataset. The values for unrecorded smear were taken from a calibration experiment and applied to real data from Ethiopia (day 299a/2008) at a variety of integration times. This will likely cause overestimates of the error, as the calibration lamp is brighter than real-world collection conditions, but



Figure 6: highly enhanced view of the Eagle CCD, showing recorded bands (inside the red box) and light outside this region.

it is indicative of the relative magnitudes of the error to be expected. As can be seen, the error is dependent on the signal at a particular band, with higher error (2-3%) at the edges of the spectral range, but also with spikes of 1% error in the absorption bands.

The estimated (probably overestimated) error ranges from <0.1% (long integration times, high signal bands) to $\sim70\%$ error (worst possible case of short integration time, low signal bands). The following table 2 shows the best and worst-case error estimates for real-world data at a variety of integration times. Figure 8 shows the distribution of integration times over several years of data collection, to give an indication of the likely impact.

7 Bad CCD Pixels

7.1 Eagle Sensor SN110001

The first pixel of each line of each band is 'noisy' and should not be used. The generated mask files will have this pixel marked as failing quality control. As a result, if you mask the data using the supplied mask files, the noisy pixels



Figure 7: estimated real-world error due to uncorrected smear, using calibration lamp data to provide estimates of the smear and real Ethiopia data (day 299a/2008). Note that the x-axis is band-number, not wavelength (band 0 =~1000nm, band 250 = ~400nm).



Figure 8: distribution of integration times for flights in 2008-2010.

will be zeroed in the masked level-1 file.

Integration	Best case error	Worst-case error
$Time \ (ms)$	(% @ peak signal	(% @ weakest
	$\mathbf{strength})$	signal)
1	5%	70%
4	$\sim 2\%$	20%
6	2%	10%
8	0.6%	9%
12	0.5%	6%
16	0.3%	4%
18	0.2%	4%
23	0.2%	3%
26	0.1%	2.5%
28	0.1%	2.5%
31	0.1%	2%

Table 2: Best and worst-case error estimates for real-world data

7.2 Fenix Sensor

The Fenix instrument has a number of bad pixels that give inaccurate values. There are many different types of error (e.g. constant pixel values, uncorrected offset, duplicating neighbouring pixels, etc), and $\sim 1\%$ (about 600) of pixels are to be expected to be bad on the type of CCD used in the Fenix instrument. A list of known bad pixels has been included within the mask files. The bad pixels will appear in level 1 datasets as straight lines along the direction of flight and as undulating lines in level 3 following the motion of the aircraft (e.g. Figure 9). Typically, they will only affect a single band and are difficult to detect. A complete solution for detecting and removing these was finalised in 2012 for the Hawk instrument and updated in 2014 for Fenix.

The final list of bad pixels uses 5 methods of detection, run on a set of test data. If a pixel fails any method then it is marked as bad. The methods are tested sequentially, and if a pixel fails an earlier method then it is still tested in later methods, so that it is marked bad for every test that it fails. The method used to flag the bad pixels can be identified using the mask files, which in turn allows masking of only certain bad pixel types.

The methods used are summarised below:

- Constant input variable output
- Constant input constant invalid output
- Linear input non-linear output
- Rapid saturation
- Visual inspection

7.2.1 Detection method A - Constant input variable output

This method marks pixels bad when they vary significantly given a constant light as input. This is tested by considering the median raw value (DN) for the pixel over time, and testing individual epochs against a percentage threshold. If the value exceeds this threshold it is flagged as bad.

7.2.2 Detection method B - Constant input constant invalid output (CICO)

This method marks pixels bad when their response to a constant light greatly varies from the response of their spatial and/or spectral neighbours. Again, this is done using the raw values. To determine when a pixel's response varies it has to be compared to its close neighbours. A moving window is used to detect responses that differ significantly.

The mean and standard deviation of pixel values are calculated for the moving window and then used within the formula below.

$$\frac{\mu_{ofCCDpixel} - \mu_{windowcentredat(sample,band)}}{\sigma_{windowcentredat(sample,band)}} \tag{1}$$

where the μ are means and σ the standard deviation. When this exceeds the theshold the bad-counter for this pixel is incremented. When the counter reaches a maximum allowed value the pixel being tested is marked bad.

7.2.3 Detection method C - Linear input non-linear output (LINO)

This method takes the average DN for each pixel over time, for multiple data captured at increasing integration times. The increasing integration times should correspond to increasing sensor response in a linear fashion. Using regression over the average values versus integration time it is possible to get a measurement of linearity from the Pearson Product-Moment Correlation Coefficient - the closer the value to 1.0 the better the fit. The formula is given as:

$$\sum_{i=1}^{N} \frac{(X_i - \mu_X) * (Y_i - \mu_Y)}{(N-1) * \sigma_X * \sigma_Y}$$
(2)

with X as the integration times and Y as the mean value of the CCD pixel, μ and σ are the sample mean and standard deviation, N the number of integration times tested (sample size). A pixel is marked bad if the pearson product-moment correlation coefficient is less than a threshold value.

7.2.4 Detection method D - Rapid saturation

The rapid saturation detection works similarly to CICO in order to detect linear but invalid responses to different integration times. This method uses the gradients of the regression fit of mean pixel values versus integration times. Once again, a moving window of certain spatial and spectral width will iterate over each pixel of each band. If the slope at the centre of the window, which is being tested, varies greatly in relation to its neighbours, then the pixel in that position will be called bad. The function used to scan over each pixel is:

$$\frac{slope_{ofCCDpixel} - \mu_{windowcentredat(sample,band)}}{\sigma_{windowcentredat(sample,band)}}$$
(3)

when the above value exceeds the threshold then the pixel will be called bad. In other words, if a pixel is saturating more rapidly than its neighbours then it is detected as bad.

7.2.5 Detection method E - Visual inspection

This method is purely a visual examination of the test data and marking any pixels as bad that are clearly giving erroneous results.



Figure 9: A bad pixel on Hawk data band 187, in a scene over water (images inverted to improve contrast on paper)