EPS Product Validation Report: IASI L1 PCC PPF

Doc.No.	:	EUM/OPS-EPS/REP/10/0148
Issue	:	v2
Date	:	5 August 2010

EUMETSAT
Eumetsat-Allee 1, D-64295 Darmstadt, Germany
Tel: +49 6151 807-7
Fax: +49 6151 807 555
http://www.eumetsat.int

Document Change Record

Issue / Revision	*Date*	*DCN No.*	*Changed Pages / Paragraphs*
v1	27 Apr 2010		First Issue
v2	5 Aug 2010		Minor editorial updates.

Table of Contents

List of Figures

1 INTRODUCTION

1.1 Purpose and Scope

The purpose of this document is to summarise the results of the validation of the Principal Component (PC) compressed IASI Level 1C data. The products under validation have been processed with the IASI L1 PCC Product Processing Facility (PPF) version 1.1.1 configured with the following eigenvector files:

```
IASI_EV1_xx_M02_20100428000000Z_xxxxxxxxxxxxxxZ_20100212000102Z_xxxx_xxxxxxxxxx
IASI_EV2_xx_M02_20100428000000Z_xxxxxxxxxxxxxxZ_20100212000102Z_xxxx_xxxxxxxxxx
IASI_EV3_xx_M02_20100428000000Z_xxxxxxxxxxxxxxZ_20100212000102Z_xxxx_xxxxxxxxxx
```

1.2 Document Structure

In Section 2 the PC Compression methodology and terminology is introduced. The evolution of the IASI L1C product and the band separation are discussed in Section 3. Section 4 describes the noise normalisation matrix, the training set and the number of scores, which went into the generation of the eigenvectors under validation. In Section 5 the configuration of the processor is discussed. Section 6 contains the validation results and finally in Section 7 the conclusion is presented.

1.3 Abbreviations and Acronyms

EARS	EUMETSAT Advanced Retransmission Service
EPS	EUMETSAT Polar System
GEADR	Global External Auxiliary Data Record
GIADR	Global Internal Auxiliary Data Record
GPFS	Generic Product Format Specification
GS	Ground Segment
HDF	Hierarchical Data Format
IASI	Infrared Atmospheric Sounding Interferometer
MDR	Measurement Data Record
Noise	Random measurement error
PCC	Principal Component Compression
PCR	Principal Component Residuals
PCS	Principal Component Scores
PFS	Product Format Specification
PPF	Product Processing Facility
RMS	Root Mean Square

1.4 Reference Documents

Ref.	Document Title and Reference Number
[RD1]	IASI Level 1 PCC Product Generation Specification, EUM/OPS-EPS/SPE/08/0199
[RD2]	Future Dissemination Approach for IASI level 1 Products, EUM/STG-OPSWG/23/08/DOC/11
[RD3]	Operational Dissemination of IASI Data using PC Compression, EUM/OPS-EPS/TEN/08/0202
[RD4]	Generation of eigenvector files for the IPCC PPF, EUM/OPS-EPS/SPE/08/0200
[RD5]	EPS Programme Auxiliary Data Inventory, EUM/EPS/SYS/LIS/00/002
[RD6]	EARS Operational Service Specification, EUM/OPS/SPE/01/0839
[RD7]	Product Validation Review Board for IASI L1 Day-2 Product, EUM/OPS-EPS/MIN/10/0103
[RD8]	Introduction of IASI L1 Data reduction in EPS, EUM/OPS-EPS/TEN/08/0185
[RD9]	Engineering Change Proposal EPS_AB_ECP_319, in EUMETSAT Dimensions data base.
[RD10]	OPS(L1) Change Request, CNES RF_001 (Hummingbird #300724)
[RD11]	N. C. Atkinson, F. I. Hilton, S. M. Illingworth, J. R. Eyre, and T. Hultberg: “Potential for the use of reconstructed IASI radiances in the detection of atmospheric trace gases”, Atmos. Meas. Tech. Discuss., 3, 501-529, 2010
[RD12]	Technical Note. Noise Covariance Matrix. IA-TN-0000-3271-CNE
[RD13]	Compression des Données IASI par Composantes Principales - Analyse d'Impact sur les Gaz Trace, NOV-3788-NT-9112

2 PRINCIPAL COMPONENT COMPRESSION METHODOLOGY

Principal component compression works by representing multidimensional data, like IASI spectra, in a lower dimensional space, which accounts for most of the variance seen in the data. This space is spanned by a truncated set of the eigenvectors of the data covariance matrix. By noise-normalising the spectra prior to the application of the compression technique, the ability to fit the data is enhanced by avoiding giving too much weight to variance caused by noise.

The coefficients used for the compression are determined by: 1) the training set of IASI L1C spectra, $X \in \mathbb{R}^{mxn}$, where m is the number of channels and n is the number of spectra in the training set, 2) the noise normalisation matrix, $N \in \mathbb{R}^{mxm}$ and 3) the number, s, of eigenvectors to retain.

Let $N^{-1}\bar{x} \in \mathbb{R}^{m}$ and $C \in \mathbb{R}^{mxm}$ be the mean and covariance of the normalised training dataset $N^{-1}X$ and let $E \in \mathbb{R}^{mxs}$ be the s most significant eigenvectors of C. The compressed representation (the PC scores), $p \in \mathbb{R}^{s}$, of a IASI spectrum, $x \in \mathbb{R}^{m}$, can now be computed as

$$p = E^T N^{-1}(x - \bar{x})$$

from which a noise-reduced approximation, $\tilde{x} = NEp + \bar{x} \in \mathbb{R}^m$ of x can be reconstructed. For this noise-reduced approximation of the radiances we use the terms "reconstructed" and "noise filtered" radiances synonymously.

The radiances measured by IASI are spectrally highly correlated and the atmospheric information they contain can therefore be explained by a small subset of the leading eigenvectors. On the contrary, this is not the case for the noise. In fact, if the random noise is Gaussian, uniform and uncorrelated, all possible directions (including all eigenvectors) explain exactly the same amount of noise.

We use the term "residual" for the difference between the original and reconstructed spectra, i.e. $x - \tilde{x}$. Sometimes, if it is clear from the context, we also use the term "residual" as short hand notation for the noise-normalised residual, i.e. $N^{-1}(x - \tilde{x})$.

If we assume that each measured spectrum, x, can be written as a sum, $x = x_0 + \varepsilon_x$, of true radiances (the atmospheric signal), x_0, and random measurement error (the instrument noise), ε_x, we can write the residual as

$$x - \tilde{x} = x_0 + \varepsilon_x - (NEE^T N^{-1}(x_0 + \varepsilon_x - \bar{x}) + \bar{x}) = x_0 - \tilde{x}_0 + (I - NEE^T N^{-1})\varepsilon_x,$$

where $\tilde{x}_0$ denotes the reconstructed true radiances, i.e. $\tilde{x}_0 = NEE^T N^{-1}(x_0 - \bar{x}) + \bar{x}$.

The first term in the expression for the residual, $x_0 - \tilde{x}_0$, corresponds to the loss of atmospheric signal and is known as the reconstruction error. The second term, $\varepsilon_x(I - NEE^T N^{-1})$, is the part of the random measurement error which is removed by the reconstruction process.

The reconstruction score of a spectrum is defined as the root mean square of the noise normalised residual, i.e. $\sqrt{\frac{1}{m}\sum_{i=1}^{m} r_i^2}$, where $r = N^{-1}(x - \tilde{x})$. If the noise normalisation corresponds well to the actual noise seen in the data, the average reconstruction score is close to one with a relatively small standard deviation. If the reconstruction score of a particular spectrum is significantly higher than the average, this is a sign that the residual does not consist of residual noise only, but is affected by non-negligible reconstruction error. This can happen if the spectrum contains features which are not well represented in the training set and is easily detected by the user since the reconstruction score is disseminated along with the PC scores.

3 ASSUMPTIONS AND OPEN ISSUES

3.1 Evolution of IASI L1C Product

The training set of IASI L1C spectra, which was used for the computation of the eigen-vectors, were all produced in the EPS GS1 with the IASI Level 1 PPF v4.0.3. With the Day-2 version v5.0.x, a pixel-dependent threshold for the calculation of the spectral shift was added to the IASI L1 processor to improve the spectral calibration. The IASI L1 Day-2 processor became operational 18 May 2010. Independently of this, the configuration auxiliary data of the IASI L1C PPF was updated, which also had an impact on the spectra. The most recent versions of the auxiliary data are:

IASI_BRD_xx_M02_20100303160000Z_xxxxxxxxxxxxxxZ_20100303151426Z_IAST_0000000012
IASI_GRD_xx_M02_20100303160000Z_xxxxxxxxxxxxxxZ_20100303153132Z_IAST_0000000020
IASI_ODB_xx_M02_20100303160000Z_20100903160000Z_20100303151607Z_IAST_0000000008

This configuration started to be used on GS1 on 18 May 2010 at the same time as the switch to the IASI L1 Day-2 processor (we will refer to this configuration as ‘new’ and to any previous configuration as ‘old’). On GS2 the Day-2 processor was running for some time with the previous configuration. To assess the impact on the compression of the changing IASI L1C characteristics we gathered three different versions of the IASI L1C data corresponding to the same sensing time interval (based on the same L0 data).

Configuration ID	IASI L1 PPF version	Aux Data
V1	4.0.3	Old
V2	5.0.2	Old
V3	5.0.2	New

For a single PDU (sensing time 20100311062057Z to 20100311062352Z) we computed the mean and the standard deviation of the radiance differences between V2/V1 and V3/V2 respectively. In the following two figures (Figure 1 and Figure 2) these means and standard deviations divided by the standard deviation of the random instrument noise in each channel are shown for both raw and reconstructed radiances.

It is seen that the mean V2/V1 difference is largely preserved by the compression, whereas the standard deviation is reduced.

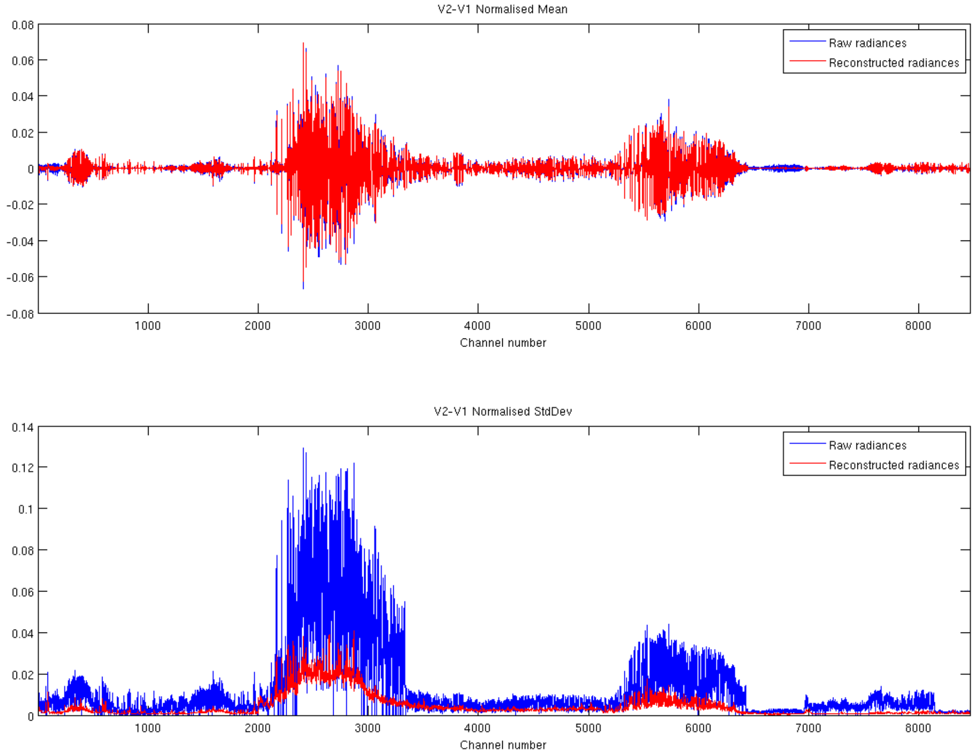

Figure 1: Radiance difference statistics for V2 type radiances minus V1 type radiances

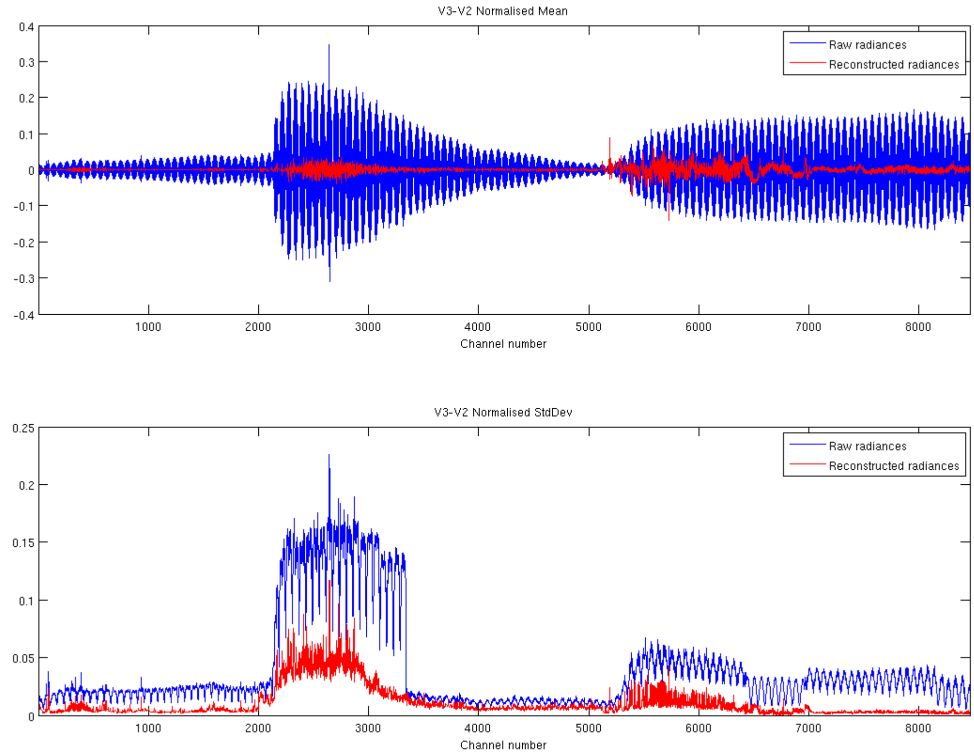

Figure 2: Radiance difference statistics for V3 type radiances minus V2 type radiances

From the V3/V2 radiance difference statistics it can be observed that the new configuration introduces a bias which is not captured by the compression, which is based on a training set consisting of spectra computed with the old configuration. By selecting a training dataset consisting of V3 spectra it would be possible to get rid of the bias. However, the ability to represent spectral features associated with rare events depends on the inclusion of such rare events in the training set [RD11]. An experiment was conducted by adding about 10,000 V3 type spectra to the existing training set, which includes rare events. But since the observed decrease in the bias was very small, we decided not to update the eigenvectors by adding V3 type spectra to the training set. We note that in any case the bias is well below noise level. If and when a full dataset of reprocessed V3 type data becomes available, a regeneration of the eigenvectors should be considered.

3.2 Band Separation

It was decided to perform the compression independently in each of the three spectral bands. The following band limits are used:

Band 1 covers channel	0 to 1996	(wavenos. 645 – 1144 cm^{-1})	(1997 channels)
Band 2 covers channel	1997 to 5115	(wavenos. 1144.25 – 1923.75 cm^{-1})	(3119 channels)
Band 3 covers channel	5116 to 8460	(wavenos. 1924 – 2760 cm^{-1})	(3345 channels)

The most important reason to do the compression separately in each band is to avoid affecting the entire spectrum in the event of single band failures. In addition it makes both compression and reconstruction faster.

4 GENERATION OF EIGENVECTORS

As described in Section 2, the eigenvectors used for compression are fully determined by:

1. The noise normalisation matrix.
2. The training set of spectra.
3. The number of eigenvectors to retain.

In this section the three ingredients (the noise normalisation matrix, the training set and the number of scores), which went into the computation of the eigenvectors, are described.

4.1 Choice of Noise Normalisation Matrix

Ideally the noise normalisation matrix, N, should be equal to the matrix square root of the noise covariance matrix. This would result in a de-correlation of the noise and ensure that the noise, provided the size of the training set is big enough, would not affect the direction of the eigenvectors. Nevertheless, for simplicity, we have used a diagonal noise normalisation matrix. This has proven to work well in practice and does not introduce any reconstruction error (provided the eigenvectors are not distorted). On the other hand it is no longer guaranteed that an equal amount of noise will be projected to each of the eigenvectors, so the noise filtering might become slightly less efficient.

For a previous set of eigenvectors, the standard deviation of the noise obtained from the IASI_NCM matrix provided by CNES [RD12] was used for the diagonal of N. However it turned out that the spatial correlation observed in the residuals could be reduced in some channels by using the standard deviation of the noise as estimated by the PC methodology itself.

To estimate the noise covariance matrix, we note that it can be written as a sum of the covariance of the noise within the residuals and the covariance of the noise within the reconstructed radiances. Under the assumption that the reconstruction error is negligible, the residuals consist of transformed noise only, so we simply compute the covariance of the residuals to get an estimate of the noise covariance within the residuals. To estimate the noise covariance, $NEE^TN^{-1}RN^{-1}EE^TN$, within the reconstructed radiances, the knowledge of the overall noise covariance matrix, R, which we are trying to estimate, is required. Therefore, in the first iteration, we used the IASI noise covariance matrix from CNES as a first guess of R in this computation. In a second iteration the previously obtained overall noise estimate was used for R.

The noise estimation was done independently in each of the three bands. In the following six figures the square root of the diagonal of the noise covariance matrix estimates (i.e. the standard deviations) are shown.

- Black: The noise guess at the start of the iteration. (In the first iteration this is the noise from CNES.)
- Green: The noise in the residual.
- Blue: The noise in the reconstructed radiances.
- Red: The new noise estimate (sum of the noise in the residual and in the reconstructed radiances). Note that the variances are added, not the standard deviations which are the ones shown in the figures.

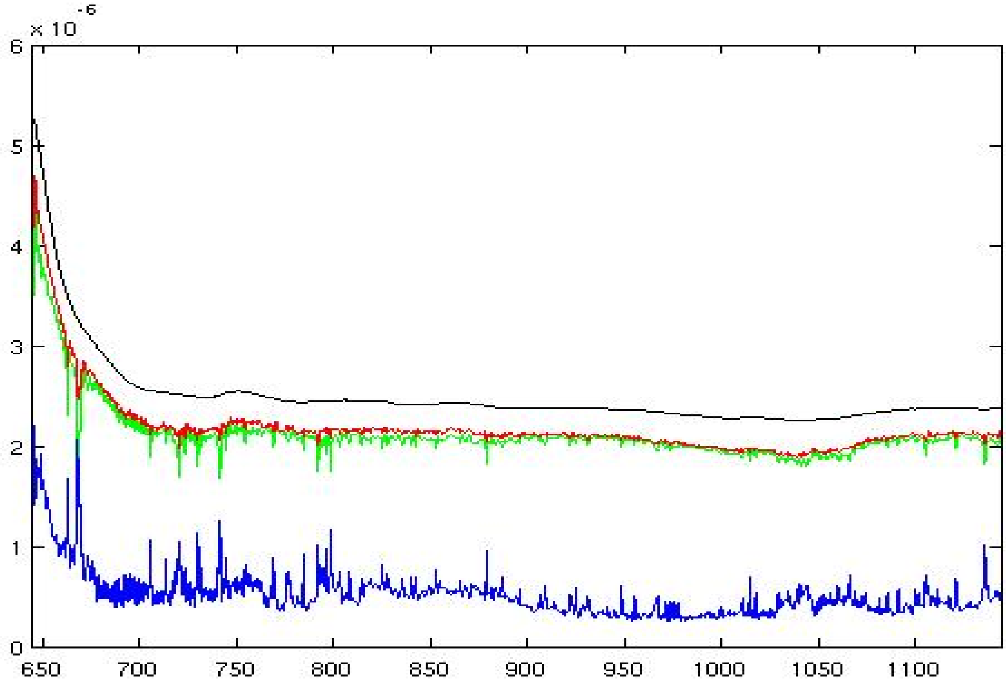

Figure 3: Band 1, noise estimation, first iteration (x-axis unit: cm^{-1}, y-axis unit: $W/m^2/sr/m^{-1}$)

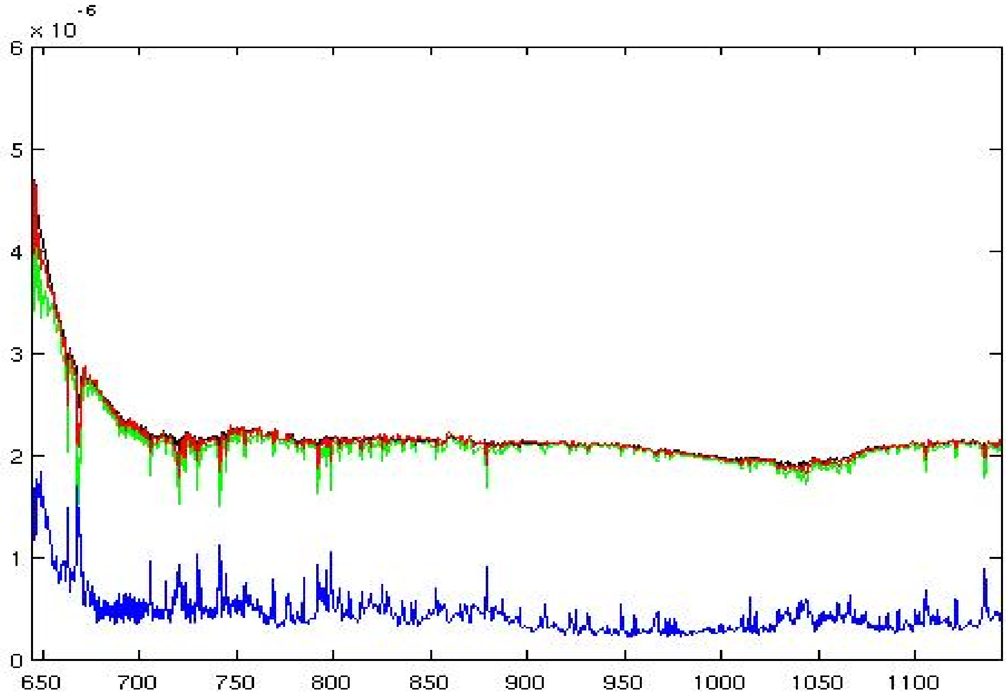

Figure 4: Band 1, noise estimation, second iteration (x-axis unit: cm^{-1}, y-axis unit: $W/m^2/sr/m^{-1}$)

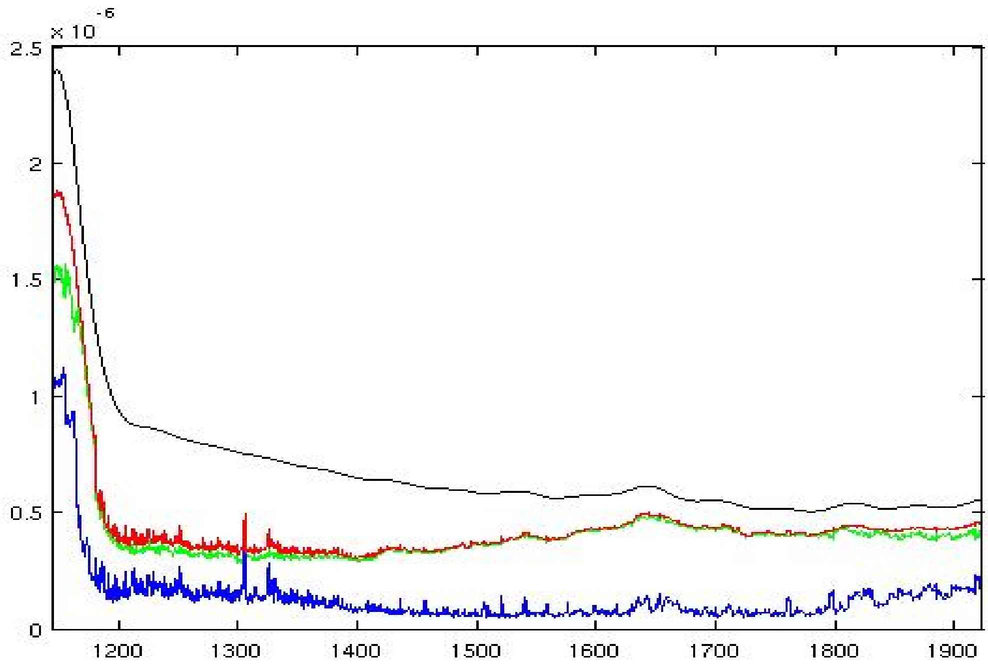

Figure 5: Band 2, noise estimation, first iteration (x-axis unit: cm^{-1}, y-axis unit: $W/m^2/sr/m^{-1}$)

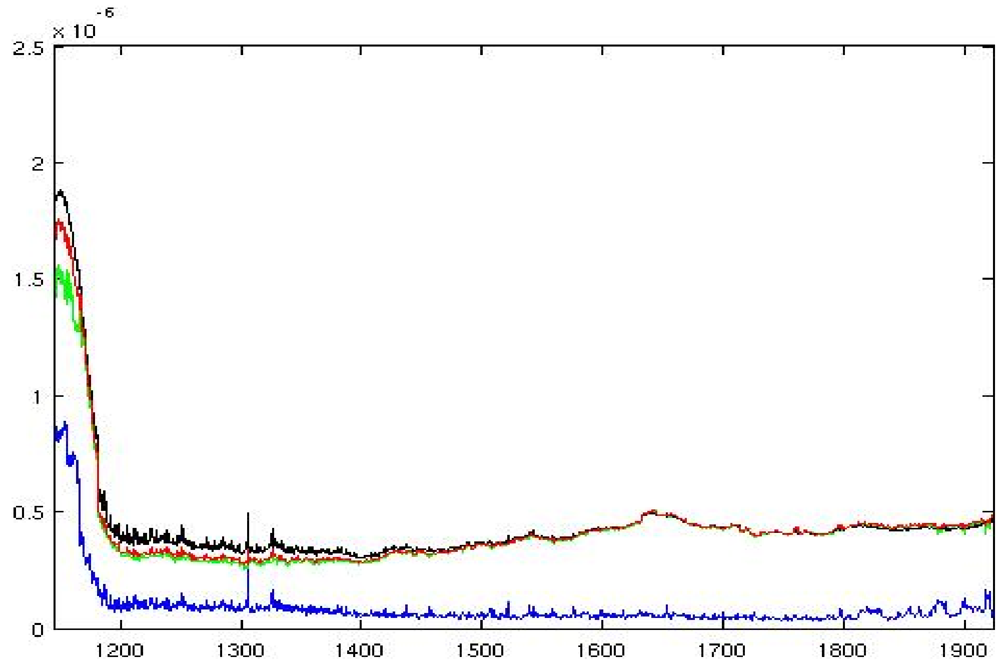

Figure 6: Band 2, noise estimation, second iteration (x-axis unit: cm^{-1}, y-axis unit: $W/m^2/sr/m^{-1}$)

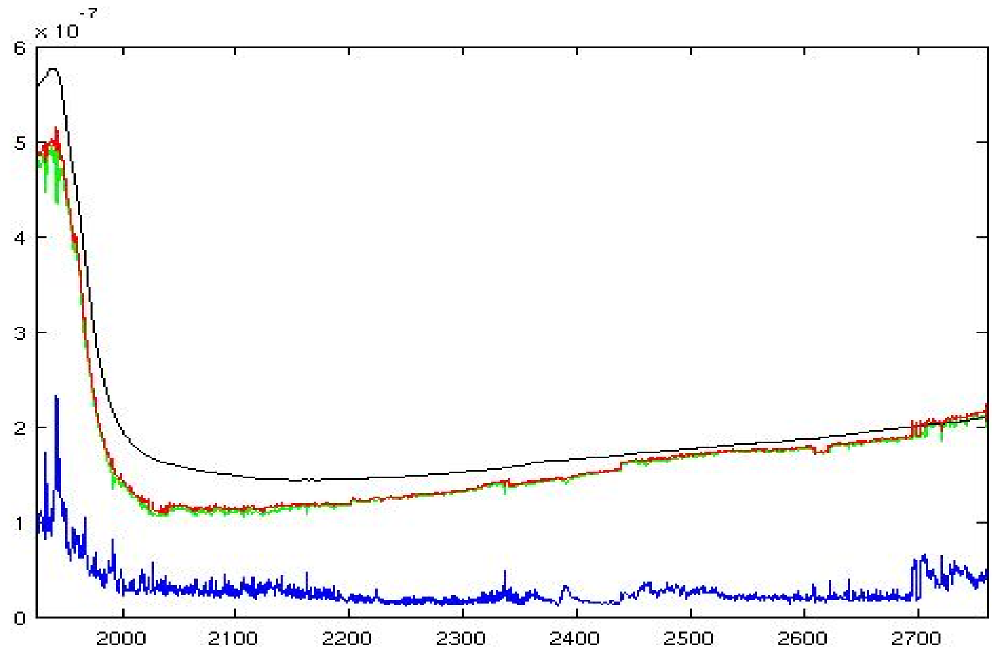

Figure 7: Band 3, noise estimation, first iteration (x-axis unit: cm^{-1}, y-axis unit: W/m^2/sr/m^{-1})

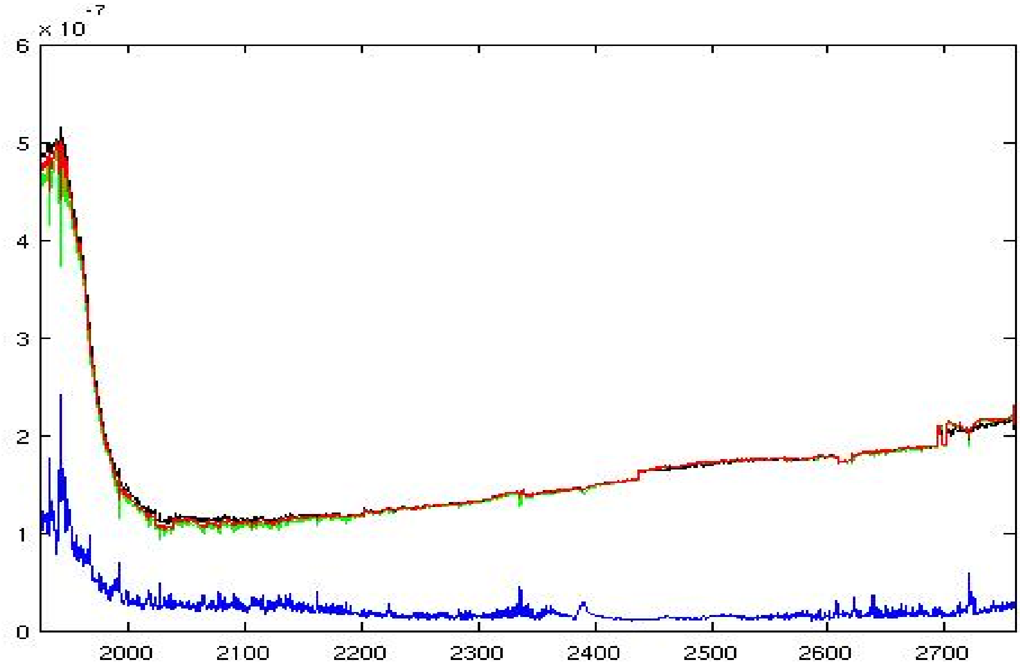

Figure 8: Band 3, noise estimation, second iteration (x-axis unit: cm^{-1}, y-axis unit: W/m^2/sr/m^{-1})

It can be seen that the estimate of the noise standard deviation is generally a bit lower than the estimate from CNES, except between 1200-1400 cm^{-1} where it is a lot lower, and above 2700 cm^{-1} where it is about the same. This estimate of the noise standard deviation has been used for the noise normalisation matrix and is included in the eigenvector files. It is important to note that this noise estimate does not pretend to represent the 'true' noise or be 'better' than the estimate from CNES. It should also be noted that a single noise normalisation matrix was estimated to be used for all four detectors, although the observed noise figures are different for the four detectors – and furthermore depend on the strength of the signal.

4.2 Choice of Training Set

The training set used for the generation of the eigenvectors under validation consists of 101,829 real L1C spectra from IASI observations. These spectra can be divided into two parts:

1. **The base spectra**. The base training set consists of 74,719 spectra which were picked randomly (one per scan line) from seven different days (7 December 2007, 12 March 2008, 9 May 2008, 26 July 2008, 26 August 2008, 25 September 2008 and 7 January 2009).
2. **The outlier spectra**. These 27,110 spectra were selected in an ad-hoc process with the purpose of including spectra with features corresponding to rare events which were not well captured in the base spectra. A total of 15 months of IASI data was available for the selection of the outlier spectra, which were mainly identified by the reconstruction score.

4.3 Choice of Number of Scores

It is important that the number of PC scores is high enough to make the reconstruction error negligible. We have used two methods to choose the number of PC scores to retain:

1) By looking at the second derivative of the average reconstruction score for an ensemble of spectra as a function of the number of PC scores, it is possible to determine at what number of PC scores the rate of decrease of the average reconstruction score stabilises.
2) If an eigenvector carries some non-negligible amount of atmospheric information, this is expected to result in some spatial correlation of the corresponding observed PC score. By plotting the PC score spatial correlation as a function of the eigenvector rank it is possible to select the number of PC scores to be high enough such that any significant spatial correlation gets included.

In Figure 9, the second derivative of the residual RMS as a function of the number of scores is shown for each of the three bands (the curves have been smoothed by applying a running average of length 5). It can be seen that, in all bands, the second derivative stabilises around zero from a certain number of scores, which means that the rate of decrease of the residual RMS starts to get constant, which again would be consistent with adding eigenvectors representing only noise (i.e. orthogonal to space of the atmospheric signal).

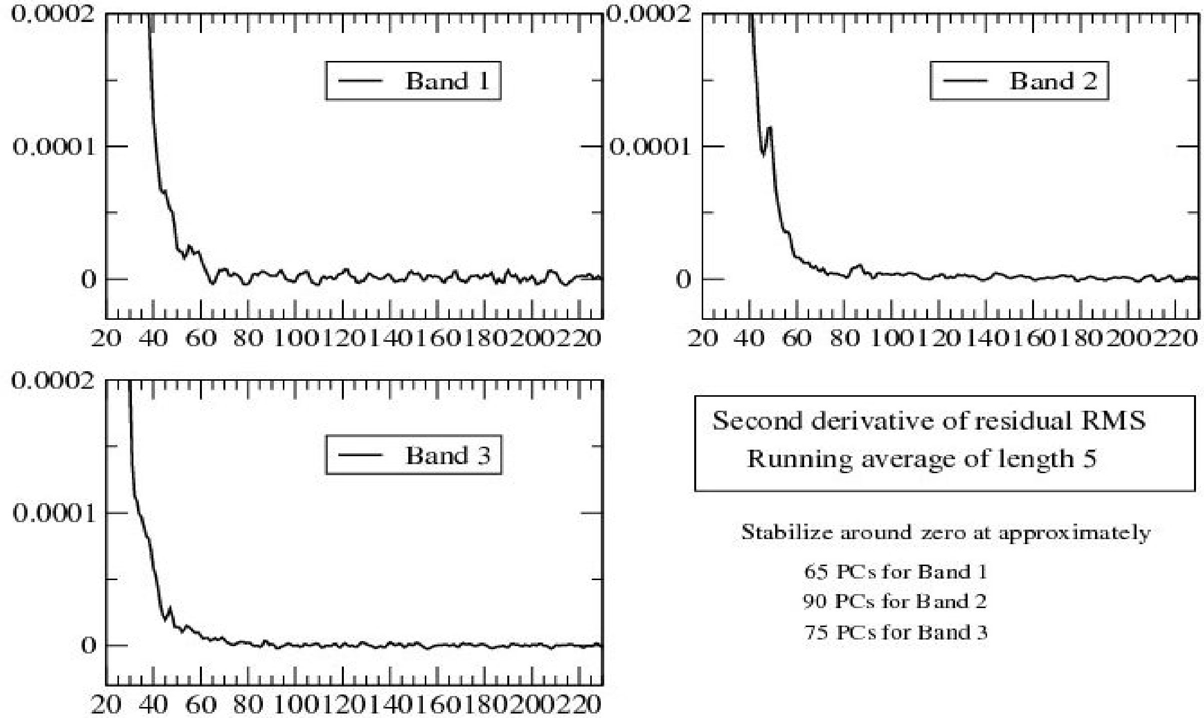

Figure 9: Second derivative of residual RMS as a function of the number of scores

Although the number of required scores estimated by the first method would probably be enough for most purposes, we decided to look at the spatial correlation of the PC scores, which is shown in the following three figures as a function of PC rank. The different colours correspond to eight different orbits (from different seasons and different areas).
By choosing **90**, **120** and **80** scores in Band 1, 2 and 3 respectively, all PC score ranks with significant spatial correlation are included.

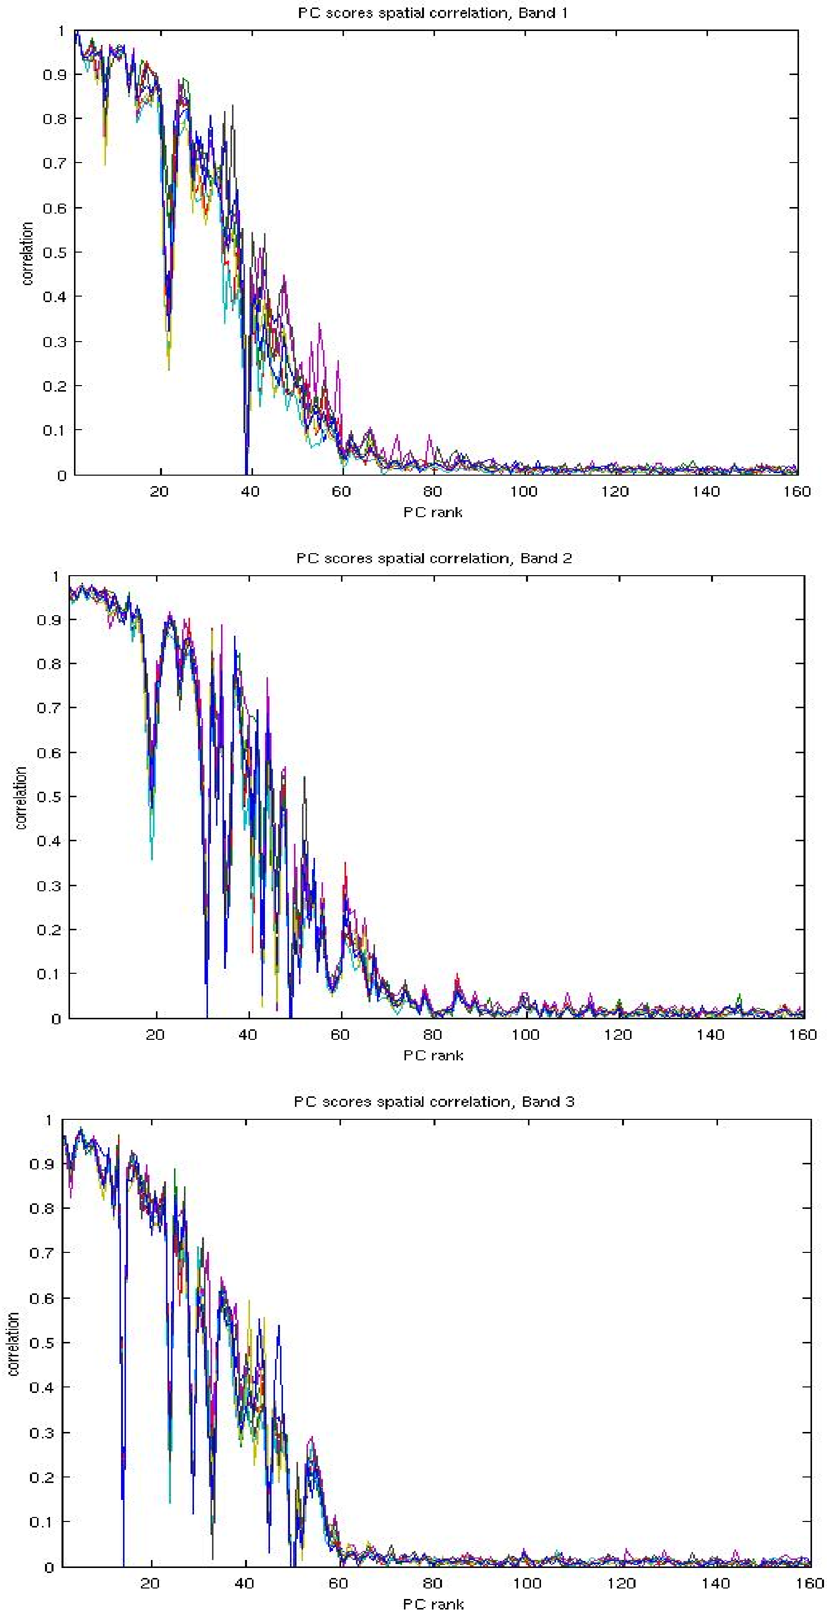

Figure 10: PC score spatial correlation

5 CONFIGURATION

5.1 Thresholds for Outlier Detection

The identification of outliers is based on thresholds for the residual RMS, where these thresholds depend on the sum of the radiances within the band as well as the pixel number. The computation of the thresholds was based on a full day (30 April 2010) of IASI L1C data, which consists of a total of 1,284,459 spectra of good quality (according to the GQisFlagQual quality flag). In Figure 11, Figure 12 and Figure 13 scatter plots of the residual RMS versus the radiance sum are presented.

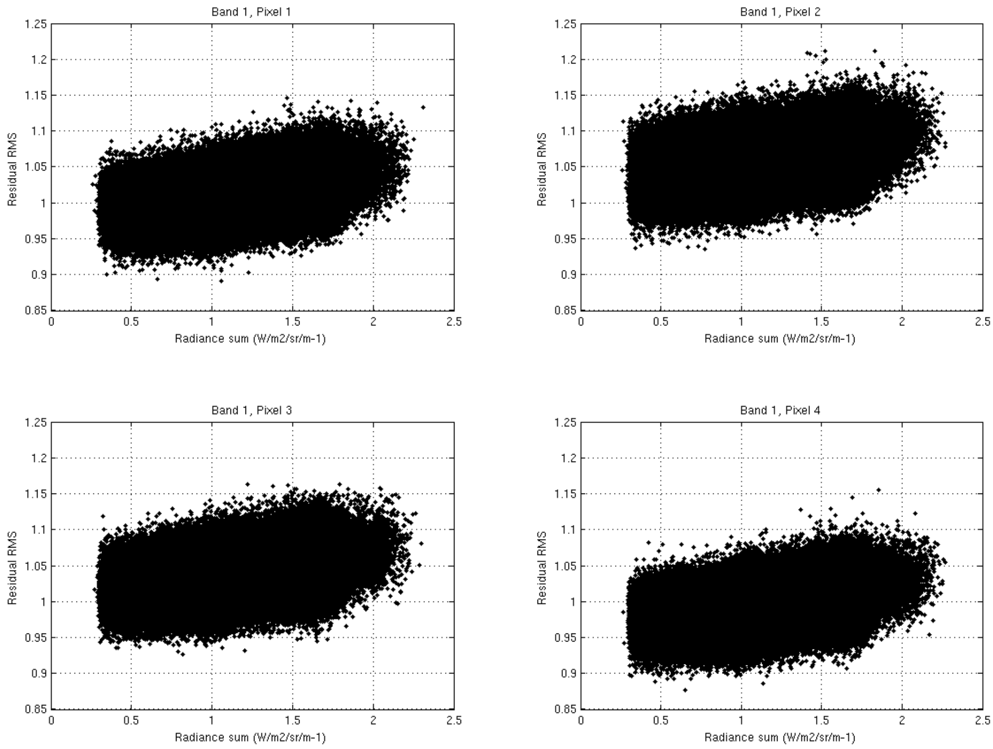

Figure 11: Band 1 residual RMS / radiance sum scatter plot

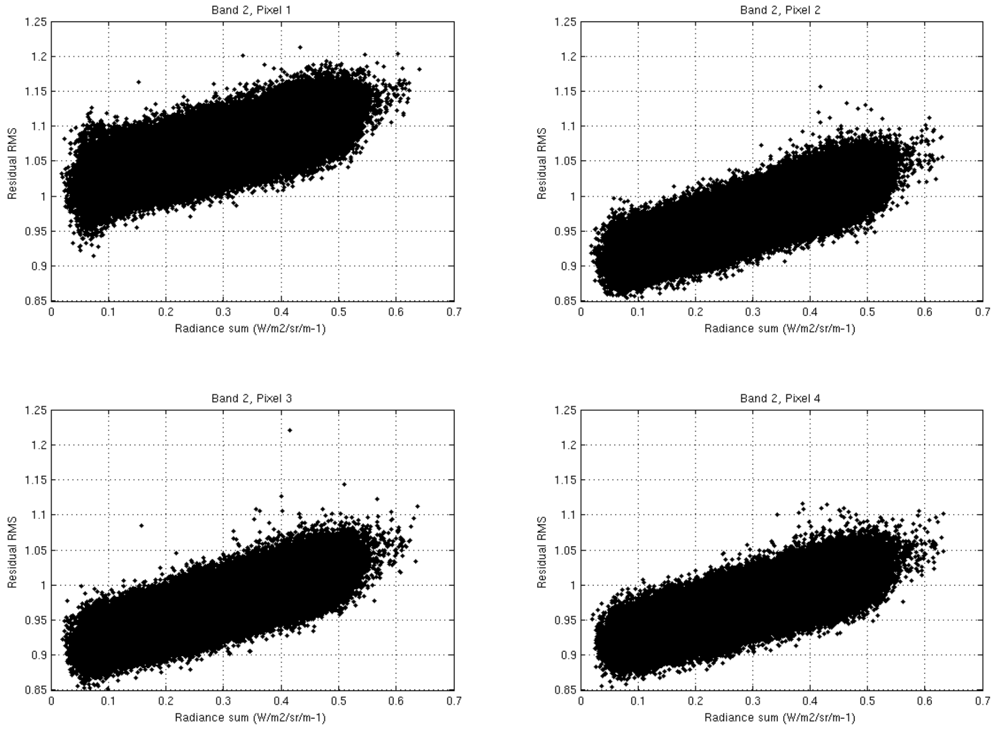

Figure 12: Band 2 residual RMS / radiance sum scatter plot

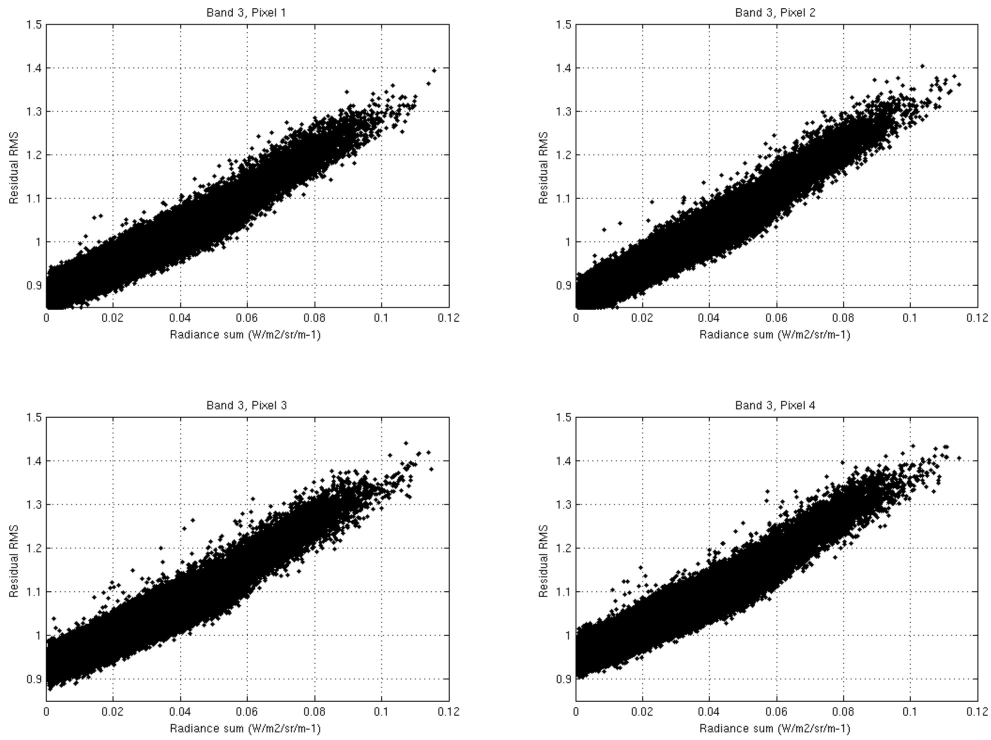

Figure 13: Band 3 residual RMS / radiance sum scatter plot

The increase of the residual RMS with the signal, as well as the difference between the four pixels, is clearly seen in the scatter plots.

In the first step the so-called outlier-slopes are computed by linear regression between the residual RMS and the radiance sum, where all four pixels are treated together (i.e. the outlier-slopes do not depend on the detector). The following slopes are obtained:

```
<outlierSlopeB1> 0.031742 </outlierSlopeB1>
<outlierSlopeB2> 0.23556 </outlierSlopeB2>
<outlierSlopeB3> 3.7215 </outlierSlopeB3>
```

To compute what we call the residual RMS anomaly for each field of view, we first subtract the outlier-slope multiplied by the radiance sum. We then compute the average of the resulting values for each of the four detectors individually, and finally these means are subtracted. Figure 14 shows histograms of the resulting residual RMS anomaly.

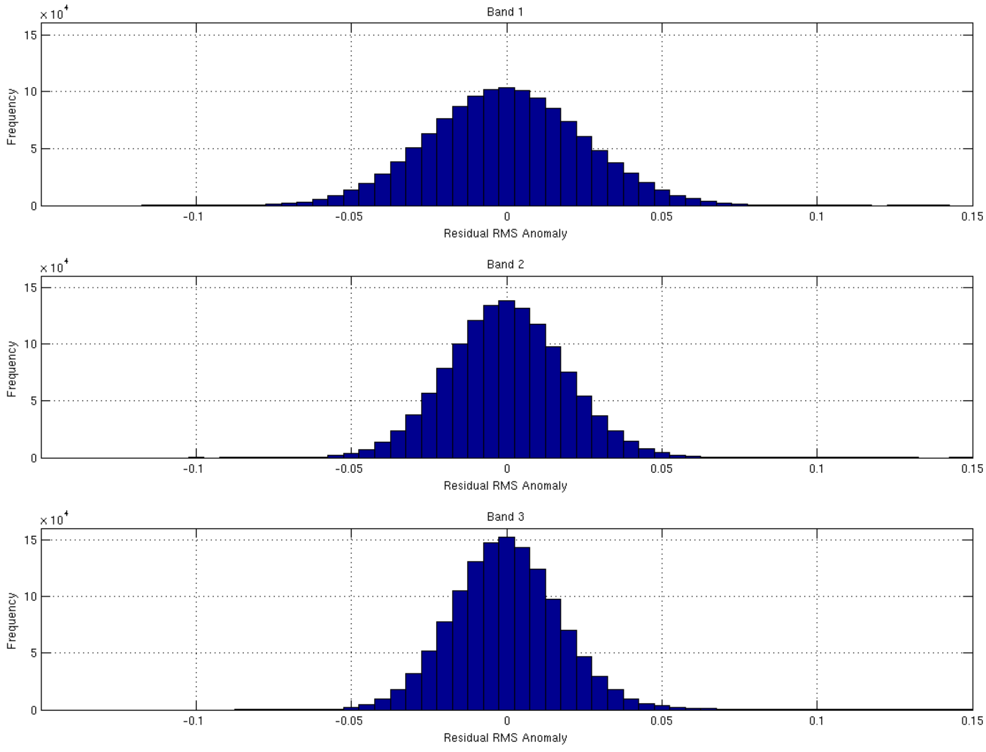

Figure 14: Residual RMS anomaly

We see that the distribution in Band 1 is flatter, which is consistent with the fact that there are fewer channels in Band 1. After inspection of this figure we chose the thresholds 0.12, 0.1 and 0.12 (in Band 1, 2 and 3 respectively) on the residual RMS anomaly in order to consider a spectrum as outlier and store it in the IASI_IPO auxiliary product. These anomaly thresholds were then converted to thresholds on the signal-corrected residual RMS for each individual detector by adding the relevant mean. The resulting thresholds are shown below:

```
<outlierThresholdB1D1> 1.0947 </outlierThresholdB1D1>
<outlierThresholdB1D2> 1.1427 </outlierThresholdB1D2>
```

```
<outlierThresholdB1D3> 1.1215 </outlierThresholdB1D3>
<outlierThresholdB1D4> 1.0768 </outlierThresholdB1D4>

<outlierThresholdB2D1> 1.0993 </outlierThresholdB2D1>
<outlierThresholdB2D2> 1.0006 </outlierThresholdB2D2>
<outlierThresholdB2D3> 1.0024 </outlierThresholdB2D3>
<outlierThresholdB2D4> 1.0019 </outlierThresholdB2D4>

<outlierThresholdB3D1> 0.9902 </outlierThresholdB3D1>
<outlierThresholdB3D2> 0.9899 </outlierThresholdB3D2>
<outlierThresholdB3D3> 1.0418 </outlierThresholdB3D3>
<outlierThresholdB3D4> 1.0628 </outlierThresholdB3D4>
```

5.2 Allocation of Number of Bytes per PC Score

After careful examination of the dynamic range of the PC scores the following configuration was chosen:

```
<nbrScoresB1P1>   1  </nbrScoresB1P1>
<nbrScoresB1P2>  41  </nbrScoresB1P2>
<nbrScoresB1P3>  48  </nbrScoresB1P3>

<nbrScoresB2P1>   2  </nbrScoresB2P1>
<nbrScoresB2P2>  61  </nbrScoresB2P2>
<nbrScoresB2P3>  57  </nbrScoresB2P3>

<nbrScoresB3P1>   1  </nbrScoresB3P1>
<nbrScoresB3P2>  44  </nbrScoresB3P2>
<nbrScoresB3P3>  35  </nbrScoresB3P3>
```

So far there have been no occurrences of overflow (i.e. PC scores too high or too low to fit within the relevant number of bytes).

6 VALIDATION RESULTS

6.1 Visual Inspection of Reconstructed Radiances

A first check of the soundness of the compression can be achieved by comparing plots of raw and reconstructed radiances (or brightness temperatures computed from raw and reconstructed radiances). Figure 15 shows a plot of brightness temperature at 657.5 cm^{-1} computed from raw and reconstructed radiances and Figure 16 at 1775.75 cm^{-1}. What might look like spatial smoothing is in fact purely a result of noise filtering within each field of view.

The same type of plot is shown for the channel at 2380.5 cm^{-1} in Figure 17 and Figure 18, for daytime (descending) and night-time (ascending) respectively. The effect of the noise filtering is striking and reveals a peculiar hidden ‘snakeskin’ pattern in the radiances, which might be caused by micro-vibrations of IASI’s beam splitter.

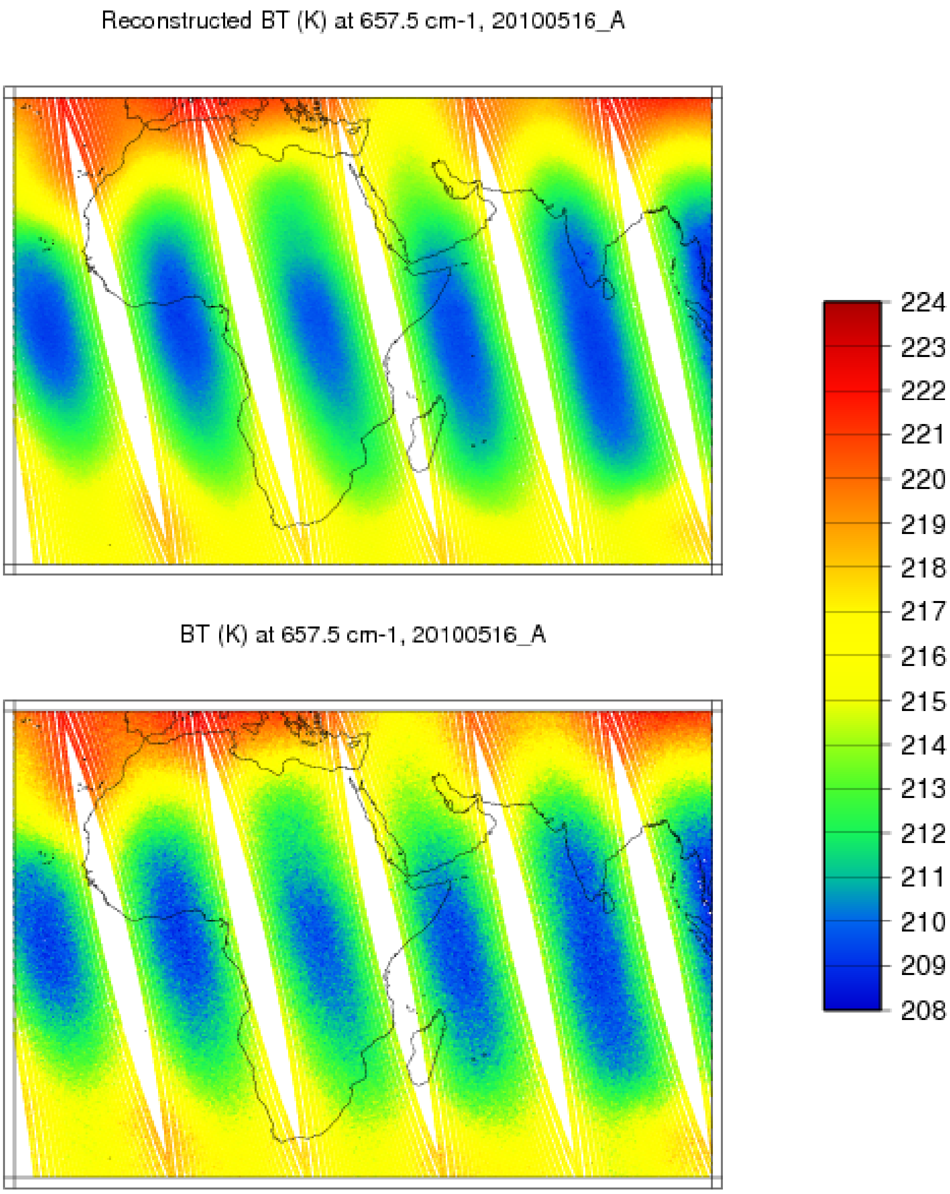

Figure 15: Brightness temperatures from reconstructed and raw radiances at wavenumber 657.5 cm^{-1}

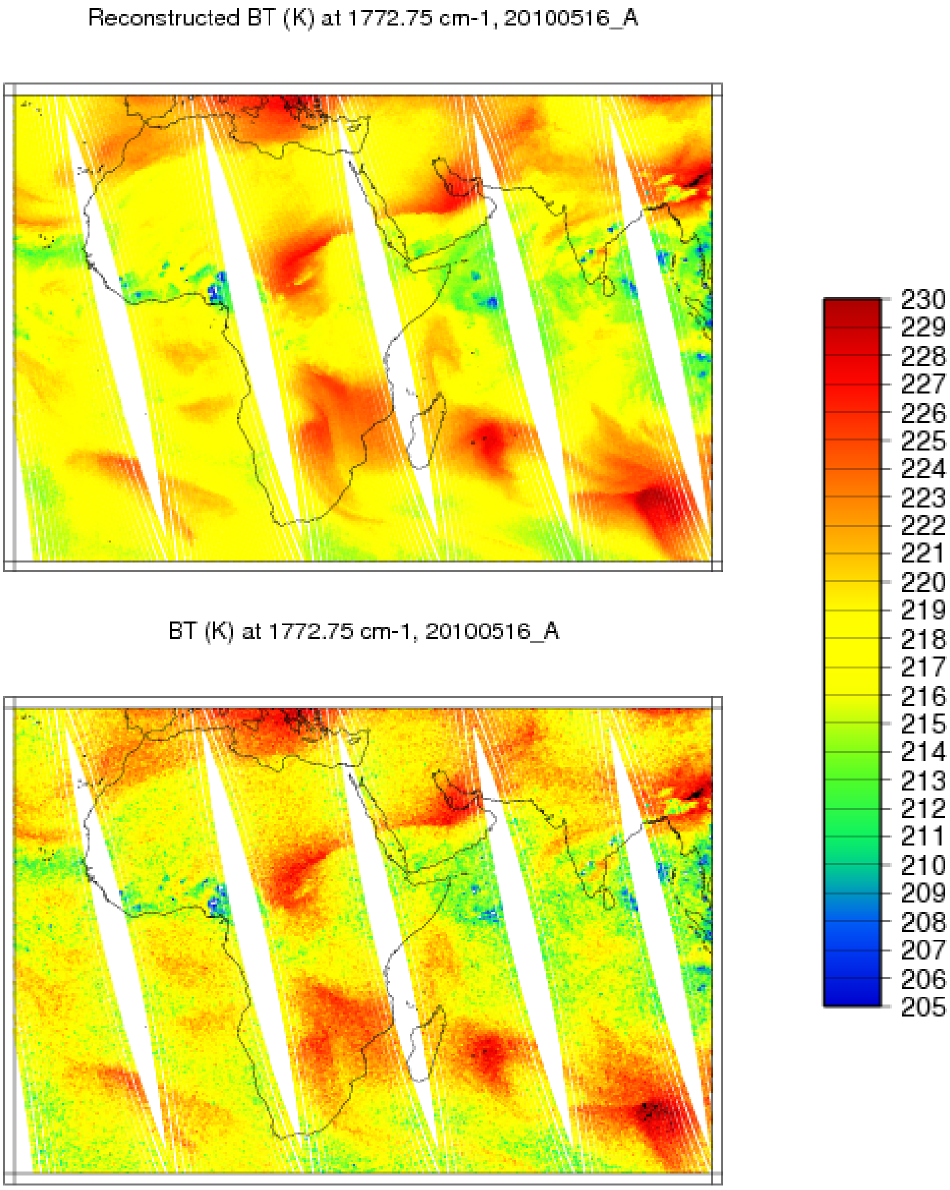

Figure 16: Brightness temperatures from reconstructed and raw radiances at wavenumber 1772.75 cm^{-1}

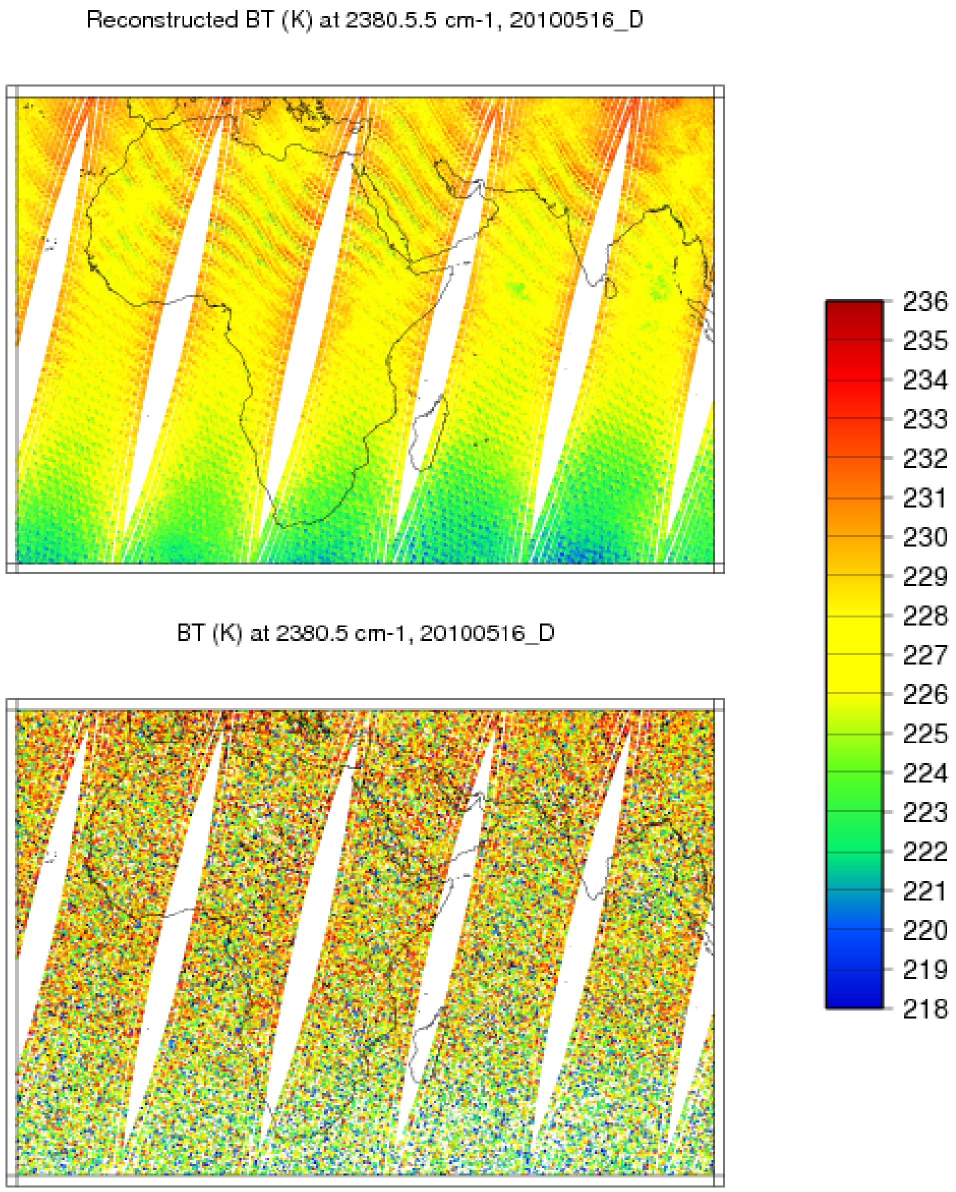

Figure 17: Brightness temperatures from reconstructed and raw radiances at wavenumber 2380.5 cm^{-1} (descending orbit)

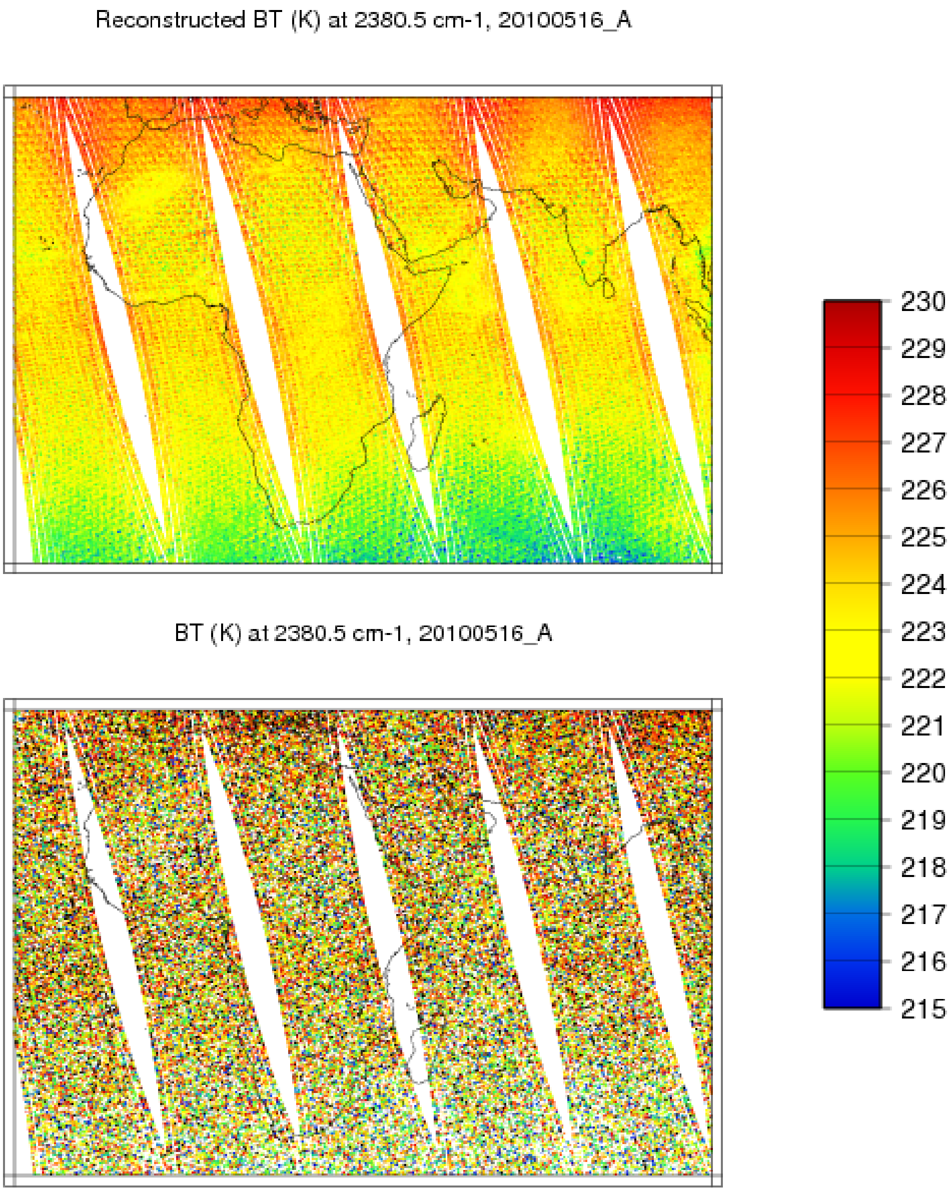

Figure 18: Brightness temperatures from reconstructed and raw radiances at wavenumber 2380.5 cm^{-1} (ascending orbit)

6.2 Examination of the Residuals

In Figure 19, the mean, standard deviation, minimum and maximum values of the noise normalised residuals, computed over a full day (30 April 2010), are shown. The figure looks as expected: the mean is close to zero, the standard deviation is close to one and there are no peculiar features in the minimum and maximum values.

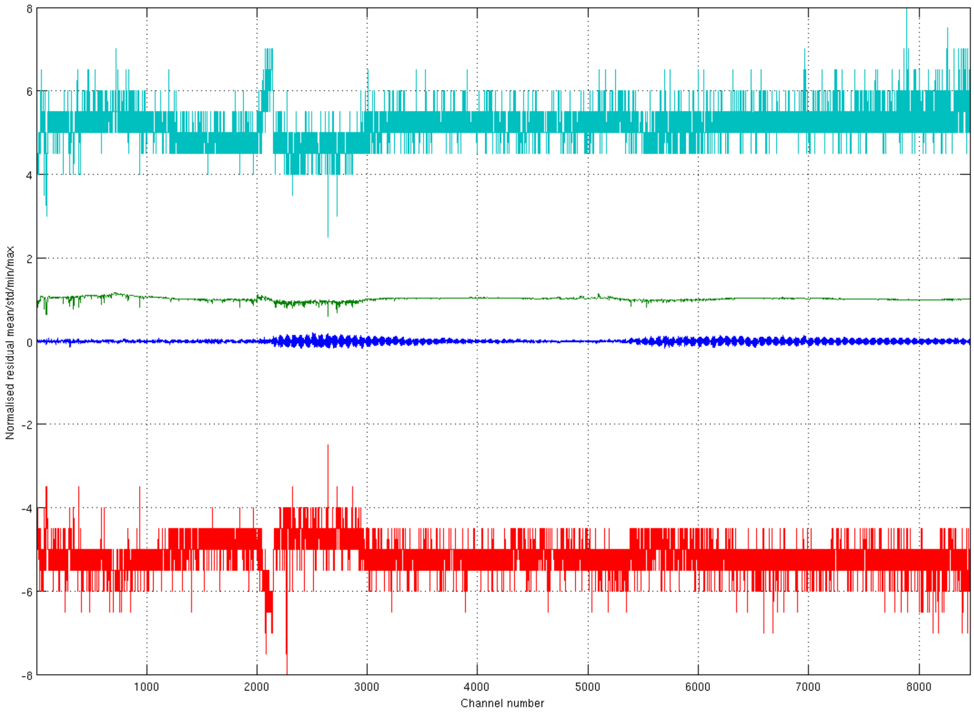

Figure 19: Residual mean, standard deviation, minimum and maximum value

Figure 20 show the same means and standard deviation as above, but with a different scale which make the features clearer. The non-zero mean is caused by the recent changes to the IASI L1C spectra, as discussed in Section 3.1. A few downward spikes can be noticed in the standard deviation; those channels where the standard deviation of the residual is clearly lower than the noise are the ones where the noise filtering is less efficient and a relatively high proportion of the original noise is still contained in the reconstructed radiances.

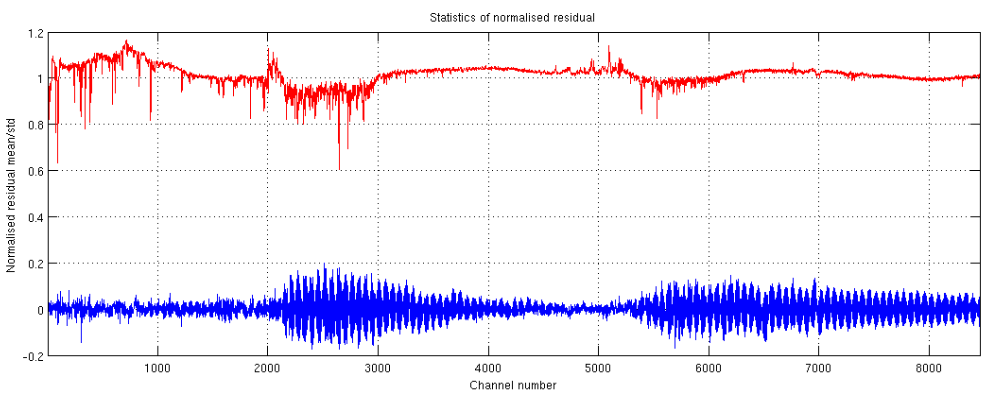

Figure 20: Residual mean and standard deviation

While the residual mean is very similar among the four pixels, there are clear differences between the residual standard deviations for each of the four detectors, as can be seen in Figure 21.

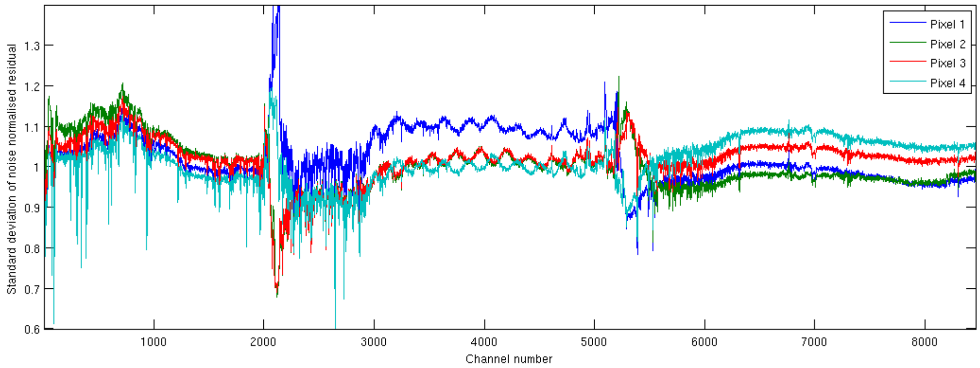

Figure 21: Residual standard deviation for each detector

It can be seen that the noise properties vary considerably from pixel to pixel, especially in the band overlap regions.

The correlation structure of the residuals can be used to check whether the crucial assumption, that the reconstruction error is negligible, is fulfilled. Recall that the residuals consist of two terms: the reconstruction error, $x_0 - \tilde{x}_0$, and transformed noise, $(I - NEE^T N^{-1})\varepsilon_x$. An estimate, R, of the raw radiance noise covariance matrix can be used to compute an estimate of the covariance of the transformed noise part of the residuals, $(I - NEE^T N^{-1})R(I - NEE^T N^{-1})^T$. In the absence of reconstruction error, the observed covariance of the residuals is expected to be very close to this theoretical expression which we computed by assuming the CNES noise covariance matrix as R for this comparison. In order to compensate for the imperfect knowledge of R, we computed the correlation matrices corresponding to both covariance matrices prior to the comparison. The comparison of the theoretical and observed correlation matrices showed a very good agreement between the two for all three bands. Here (in Figure 22 and Figure 23) we only show the comparison for some selected channels in Band 1.

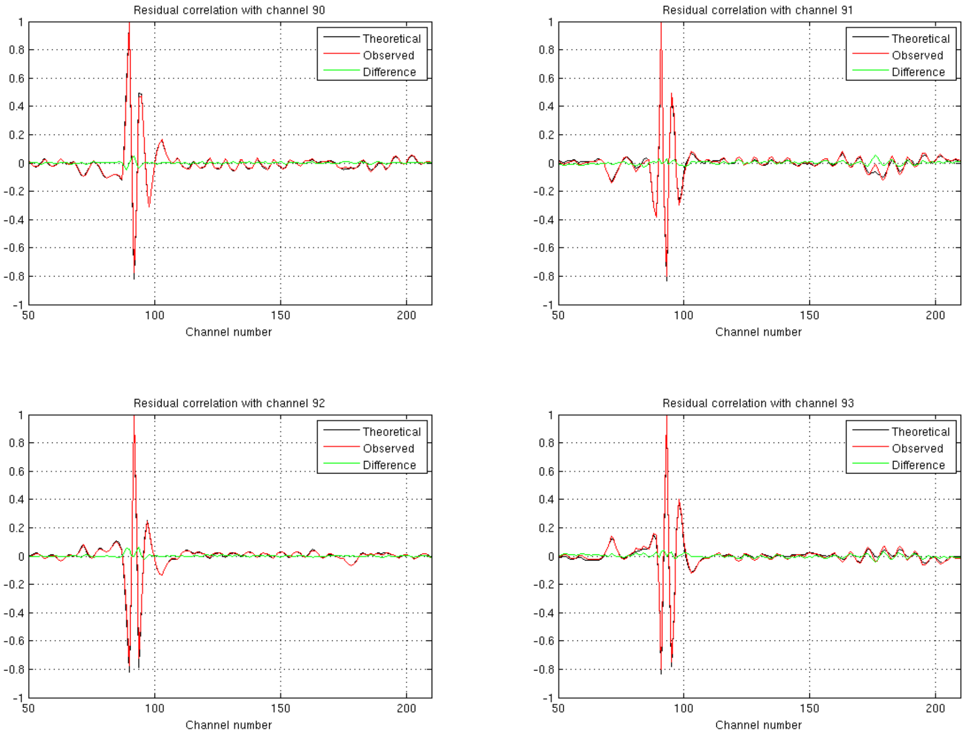

Figure 22: Residual correlation in selected channels

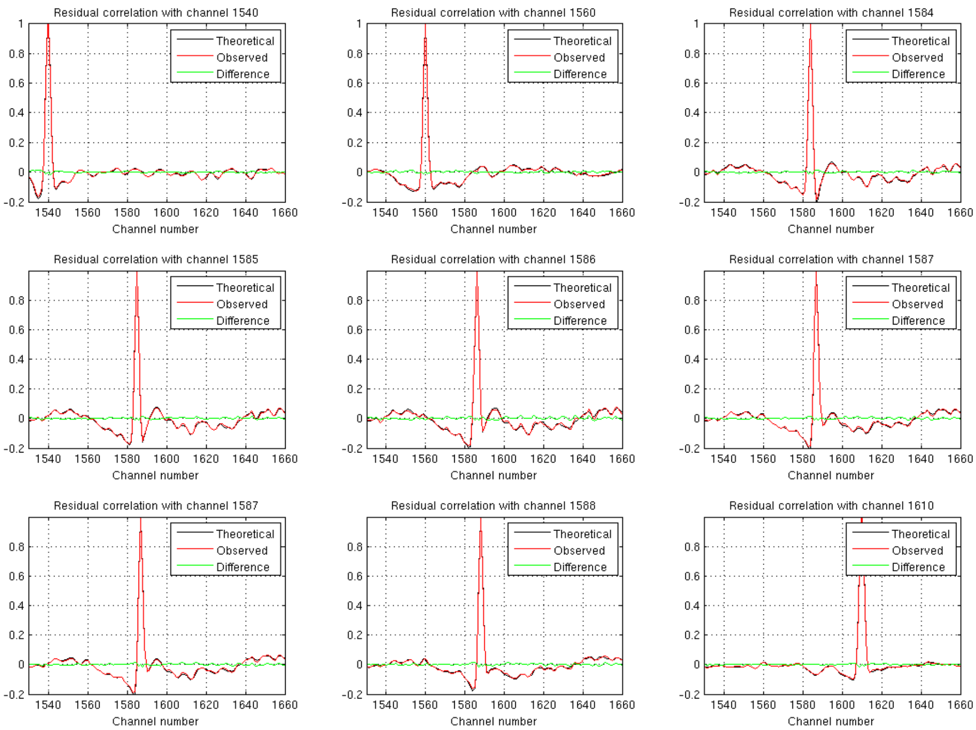

Figure 23: Residual correlation in selected channels

The minimum and maximum values of the difference between the two correlation matrices are:

	Min	Max
Band 1	-0.1415	0.0653
Band 2	-0.1188	0.1294
Band 3	-0.0942	0.0857

7 CONCLUSION AND RECOMMENDATION

Visual inspection of maps of reconstructed radiances obtained from the compressed products under evaluation showed that the spatial features observed in the raw radiances are retained or even enhanced as a result of the noise filtering.

Examination of the statistical properties of the residuals showed that they are consistent with the expected behaviour and that no systematic reconstruction error is contained in the residuals, except for the small bias introduced by the configuration change of the IASI L1C processing. Of course, this does not exclude the possibility that occasional individual spectra might be affected by reconstruction error, but we hope that our method for detecting outliers will be capable of reliably detecting such spectra. Very narrow unrepresented features (affecting only very few channels) might be hard to detect if the impact on the overall reconstruction score is not big enough. Looking at the min/max values of the residual, there is nothing that indicates the existence of such narrow unrepresented features, but in parallel to the monitoring of the reconstructions scores, the residual min/max values will be monitored.

In the event that new features which are not well represented in the truncated eigenvector space are encountered, the eigenvectors can be regenerated based on an updated training set.

It is recommended to install the IASI PCC PPF V1.1.1 on GS1 and release the IASI Level 1 PCS products for trial dissemination to users. This will allow users of IASI data to evaluate the impact of the compression for their particular application over a long time period.