EUMETSAT ECMWF

EUMETSAT/ECMWF Fellowship Programme Research Report

51

Assimilation of radiance observations from geostationary satellites: second year report

Chris Burrows

October 2019

Series: EUMETSAT/ECMWF Fellowship Programme Research Reports

A full list of ECMWF Publications can be found on our website under:
http://www.ecmwf.int/en/research/publications

Contact: library@ecmwf.int

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, UK

Executive summary

This report summarises the work carried out during the second year of this EUMETSAT Research Fellowship.

All-sky radiance (ASR) data from the SEVIRI instruments on Meteosat-8 and Meteosat-11 remain key inputs to the ECMWF operational data assimilation system and as such great care is taken to ensure observation quality and continuity. In this reporting period we have had to manage a number of potentially disruptive real-time events. Firstly a large, sudden bias change in the $7.3\mu m$ water vapour channel on Meteosat-8 occurred in January 2019, which remains unexplained, but was sufficiently short-lived that the impact on operations was minimal. The second event was a scheduled decontamination of the SEVIRI instrument onboard Meteosat-11. Such exercises (from previous experience) are known to change the bias characteristics of the observed radiances, particularly in the $13.4\mu m$ channel, which is used for cloud initialisation, and this was again the case. The rapid adjustment of the ECMWF adaptive bias correction system (VarBC) ensured a smooth transition and there was no adverse impact on the assimilation system. The final issue was the dissemination of unphysical ASR radiance data during a sun-satellite-ground-station co-linearity event. This highlighted some weaknesses in our real-time data screening processes and has now resulted in the implementation of additional upstream checks to protect against similar occurrences in the future.

Efforts have continued to increase the exploitation and impact of SEVIRI ASR data, as well as clear sky radiance (CSR) data from other instruments in the ECMWF assimilation and forecasting system. One important area is to extend the assimilation of geostationary radiances to include window channels. These can convey important information about boundary layer humidity (in the tropics) and assist the vertical localisation of humidity from the sounding channels above. Results so far suggest that there is additional skill to be gained from window channels, but that the screening of cloud/aerosol as well as handling time varying biases is challenging. Another important development area is that of further improvements to the specification of observation errors for ASR data. Following on from the success of the implementation of inter-channel correlations for radiances from geostationary satellites (Burrows [2018]), focus has turned to the possibility of varying observation errors throughout the 4D-Var assimilation time window. Currently, all ASR/CSR observations are assigned the same observation error regardless of when they are observed. While the observation quality does not change, observation errors should include errors of representation and account for the growth of model error over the assimilation window, suggesting that data at the end of the window should be down-weighted (Howes et al. [2017]). However, a previous EUMETSAT sponsored study by McNally [2019] has demonstrated that observations at the end of the window are the most critical and influential, and suggests that down-weighting these would likely result in degraded forecast skill. This issue is clearly most acute for geostationary radiances as they report with high temporal frequency over the 4D window. A number of different time-varying observation error models have now been tested in dedicated assimilation experiments. While the results of these do demonstrate that the assimilation system is very sensitive to such modifications, none have emerged (so far) as a clear improvement and investigations in this area continue.

In addition to maintaining and optimising the use of current geostationary radiance data from EUMETSAT, a significant amount of effort has been directed towards preparing for the assimilation of radiance data from the the future MTG-IRS mission. Some of the enhancements that will be delivered with MTG-IRS can actually be investigated using existing data from a number of different sources. Firstly, The

Chinese geostationary satellite FY-4A carries an experimental hyperspectral infrared instrument (called GIIRS) and data from this instrument became available in 2019. An initial assessment of the radiances has been performed and uncovered a number of anomalies including horizontal banding, spectral shift, inter-dwell variability, intra-dwell variability and spurious geometric features. However, despite these issues, early indications suggest that, with highly selective screening, enough good quality GIIRS data can be be recovered to allow sensible prototype assimilation trials to begin. Secondly, MTG-IRS will deliver radiance data with very high time and spatial resolution and we can begin to gain some experience of how best to exploit these particular characteristics of the future IRS by working with existing radiance observations from the GOES-16 Advanced Baseline Imager (albeit with low spectral resolution non-hyperspectral data). Unfortunately, to be in a position to perform data assimilation experiments, a number of significant data issues had to be investigated including spurious image striping and cloud/aerosol contamination in the window channels. However, these have now been largely resolved by NOAA/NESDIS/STAR, and GOES-16 half-hourly data have been activated as part of the ECMWF operational assimilation system. Experiments are now underway to explore if radiance data with a repeat frequency in the range 10-20 minutes (IRS will have 15 minute repeat) can be successfully digested by the 4D-Var assimilation system.

Contents

1 GOES-16

In January 2018, the satellite GOES-13 was retired as the operational GOES-EAST full-disk service. Its replacement was GOES-16, launched in 2016 as "GOES-R", and it carries the Advanced Baseline Imager (ABI), a very similar instrument to the Advanced Himawari Imager (AHI) which is onboard Himawari-8 and Himawari-9. GOES-16 became operational without the provision of operational clear sky radiance (CSR) data, which ECMWF would assimilate as it does for the pre-existing geostationary satellites. An unofficial GOES-16 CSR product was, however, provided in near-real-time by NOAA/NESDIS/STAR early in 2018, which allowed these data to undergo a pre-assimilation assessment. In fact, STAR have been providing two versions of the CSR products, using different cloud masks; "baseline" and "enterprise". Here, only the baseline product has been assessed as the indications are that this is more likely to become the algorithm used when the product becomes operational, although the enterprise algorithm may have better performance (Haixia Liu, private communication). Also, the enterprise product is produced with a significant temporal delay compared to the baseline product, which is not desirable for operational data assimilation as timeliness is beneficial.

Compared to GOES-13, GOES-16 has many more channels than its predecessor, including three water vapour channels compared to just one on GOES-13. Therefore, the instrument is able to probe the atmospheric water vapour in several broad, overlapping layers in the troposphere. Furthermore, the CSR products are provided much more frequently. From GOES-13, CSRs from one full disk were received every three hours, with northern-hemisphere scans provided hourly in the intervening periods. GOES-16 started by providing full-disk scans every 15 minutes, and at the time of writing this has been improved upon, and full-disks are provided every 10 minutes, which is unprecedented for CSRs.

Currently, only the radiances from water vapour channels from geostationary satellites are assimilated at ECMWF, and in the sections that follow, priority will be given to consideration of these channels. The following sub-sections describe characteristics of the GOES-16 CSR data, its impact when assimilated and some preparations that will be required in the future.

1.1 Striping in $7.4\mu m$ channel

During the period 19-22 May 2018, the $7.4\mu m$ channel exhibited horizontal striping which was noticeable in the brightness temperatures as received in the BUFR files. These stripes extended the full width of the disk, and appeared to affect every 17th row of CSRs (in fact, the adjacent rows to these also appear to be affected, but to a lesser extent). The reason for this is likely due to an issue with a small subset of elements within the detector array. Within the detector array there are 332 rows of elements, so given the CSR averaging box size of 15x15 pixels, this would suggest that every 22nd row of CSR "super-obs" should be affected, but there is some vertical overlap in the swaths (Tim Schmit, private communication). These stripes are apparent in Figure 1.

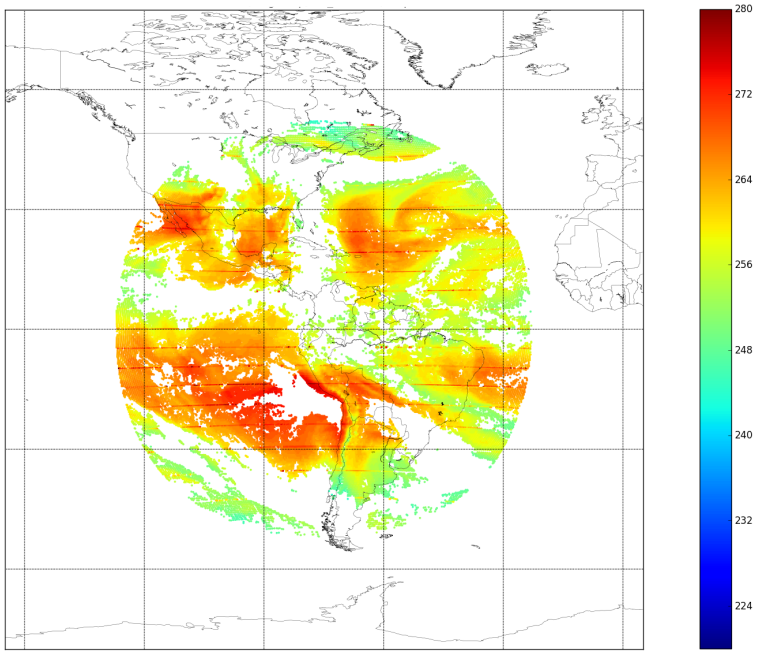

(a)

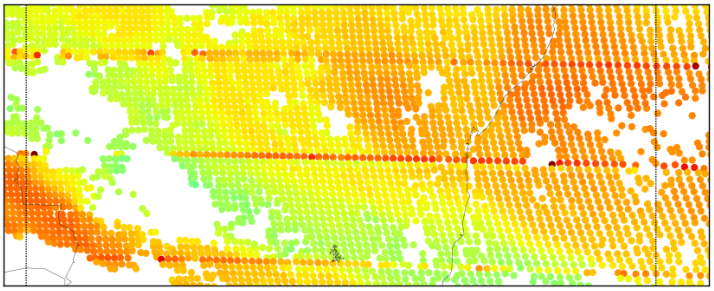

(b)

Figure 1: GOES-16 CSR brightness temperatures for a full-disk scan from May 2018, displaying horizontal striping (7.4μm channel). Shown are the full disk scan (a) and a zoomed region of the eastern coastline of Brazil (b).

It can be seen that although the stripes represent a clear systematic problem, the CSR values are not unphysical, and hence would not necessarily be rejected by usual quality control methods. A suggestion provided by Jim Jung (private communication) was to develop a quality control check using the standard deviation of the brightness temperatures of the clear pixels that contribute to each CSR. This is supplied in the BUFR files and represents the inhomogeneity of the scene within the averaging box. Clearly, if a small number of clear pixels are systematically different to the remaining pixels in the averaging box, then the standard deviation should be sensitive to this. Some experimentation resulted in a threshold of 4K being selected. This was a trade-off between removing as much of the striping as possible, whilst minimising the rejections of unaffected data. Figure 2 shows that the 4K threshold removes the affected observations quite cleanly, with only a small number of erroneous rejections.

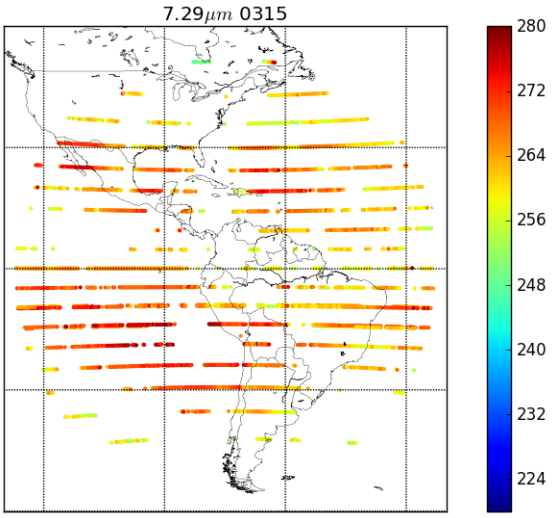

(a)

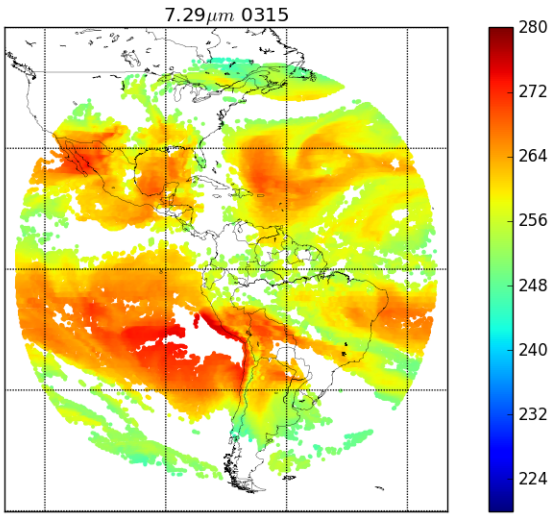

(b)

Figure 2: The GOES-16 CSR brightness temperatures from a single scan identified as affected by striping using the standard deviation check (a) and the remaining brightness temperatures deemed to be unaffected (b). These observations are from the 7.29μm channel.

The striping issue appears to have been an isolated incident, but because the other quality control checks could easily fail to identify these degraded observations, the check on the standard deviations was added to the operational ECMWF code in anticipation of GOES-16 radiances being assimilated operationally.

1.2 Diurnal cloud issue

At the end of May 2018, shortly after the striping ceased, there was another brief issue, whereby large diurnal biases appeared in the data. Specifically, at night, it appeared that the cloud mask was not identifying cloud-affected pixels correctly, and hence very low brightness temperatures were seen in the cloudy regions, see Figure 3.

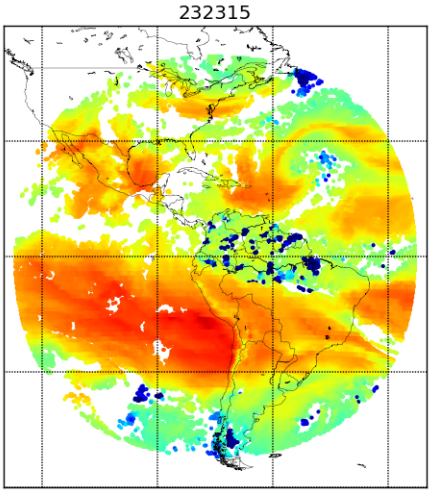

(a)

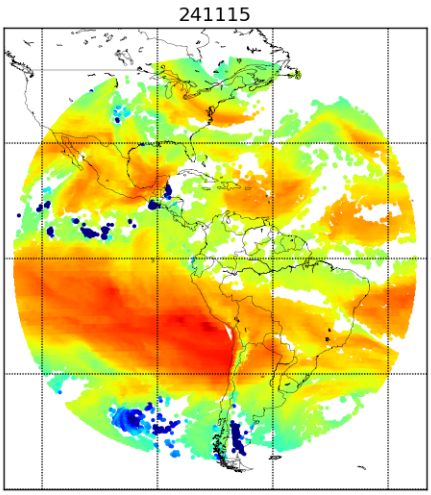

(b)

Figure 3: GOES-16 CSR brightness temperatures ($7.3\mu m$ channel) at 23:15UTC, 23/05/18 (a) and 12 hours later (b). In both cases, the blue pixels correspond to cloud contamination and only appear on the night side of the terminator, which crosses the two disks diagonally. The colour scale is the same as for Figure 2.

This issue was resolved quickly, but it was encouraging to note that the quality control in the ECMWF system would have rejected most of these contaminated observations thanks to a check on the window channel first guess departures. This can be clearly seen in Figure 4, where the standard deviation of the first guess departures of "all" data looks to be perturbed only slightly due to the striping, but greatly due to the cloud contamination. Conversely however, after quality control, the "used" data shows that the cloud contamination is not apparent, but the striping persists, indicating that it is important to provide additional quality control to protect against potential reoccurrences of the striping problem.

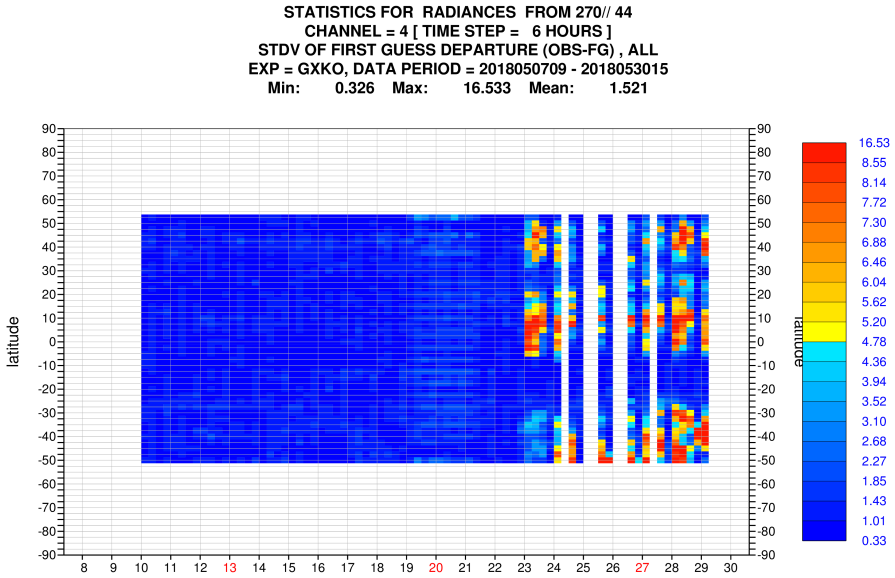

(a)

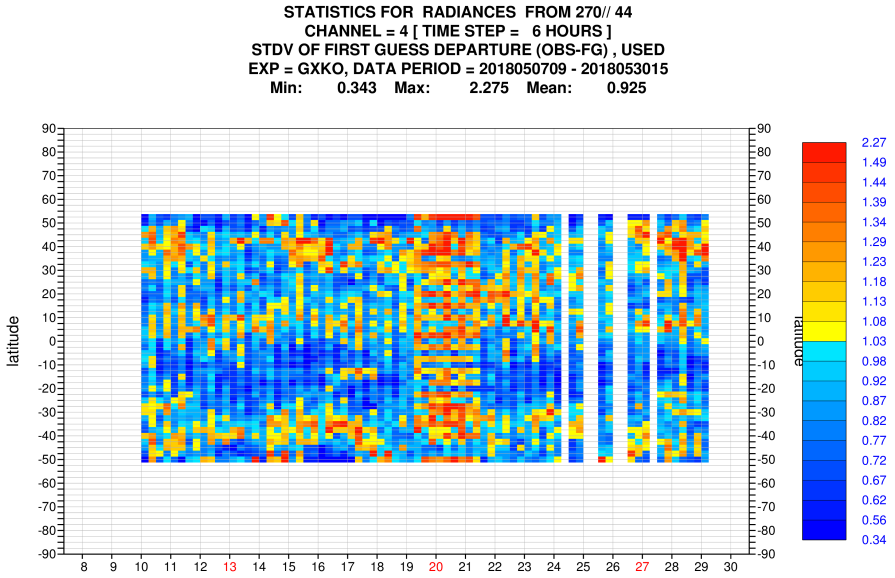

(b)

Figure 4: Zonal time series of the standard deviations of GOES-16 CSR first guess departures ($7.3\mu m$ channel) . The striping issue occurred 19-22 May, and the cloud contamination began on the 23rd. Before quality control is applied (a), the signal from the cloud contamination dominates, and the that from the striping is relatively small. After quality control is applied (b), the cloud contamination is removed, but the signal from the striping persists (note the different colour scales).

1.3 Assimilation of GOES-16 radiances

Extensive testing was performed to assess the impact GOES-16 observations would have on forecasts when they were assimilated in the ECMWF 4D-Var system. Because there was no overlap period with GOES-13, the baseline for this testing was a full observing system, but with no radiances from GOES-EAST. With the additional clear sky radiances, large improvements are not anticipated to be observed in the medium range forecast skill, but it is possible to infer improvements at shorter ranges by examining the first-guess fits to independent observations. For these experiments, GOES-16 was assimilated in the same way as Himawari-8, with observation errors of 2K, and with inter-channel error correlations neglected (although at cycle 46r1 of the ECMWF IFS, the other geostationary satellites started to be assimilated with correlated observation errors (Burrows [2018]). The check for striping was also implemented (Section 1.1). Although the data are supplied every 15 minutes, temporal thinning results in an effective temporal sampling of 30 minutes.

Some changes in first guess fits to independent observations are shown in Figure 5. Here, we can see that for ATMS and IASI, the addition of GOES-16 results in an improved short-range forecast fit for the channels sensitive to water vapour. First guess fits to other satellite radiances (not shown) corroborate this. We also see an improved fit to conventional wind observations at the heights at which the GOES-16 water vapour channels have sensitivity to atmospheric humidity. Finally, first guess fits to other geostationary radiances improve; notably the satellites whose disks overlap with that of GOES-16, suggesting consistency between these instruments.

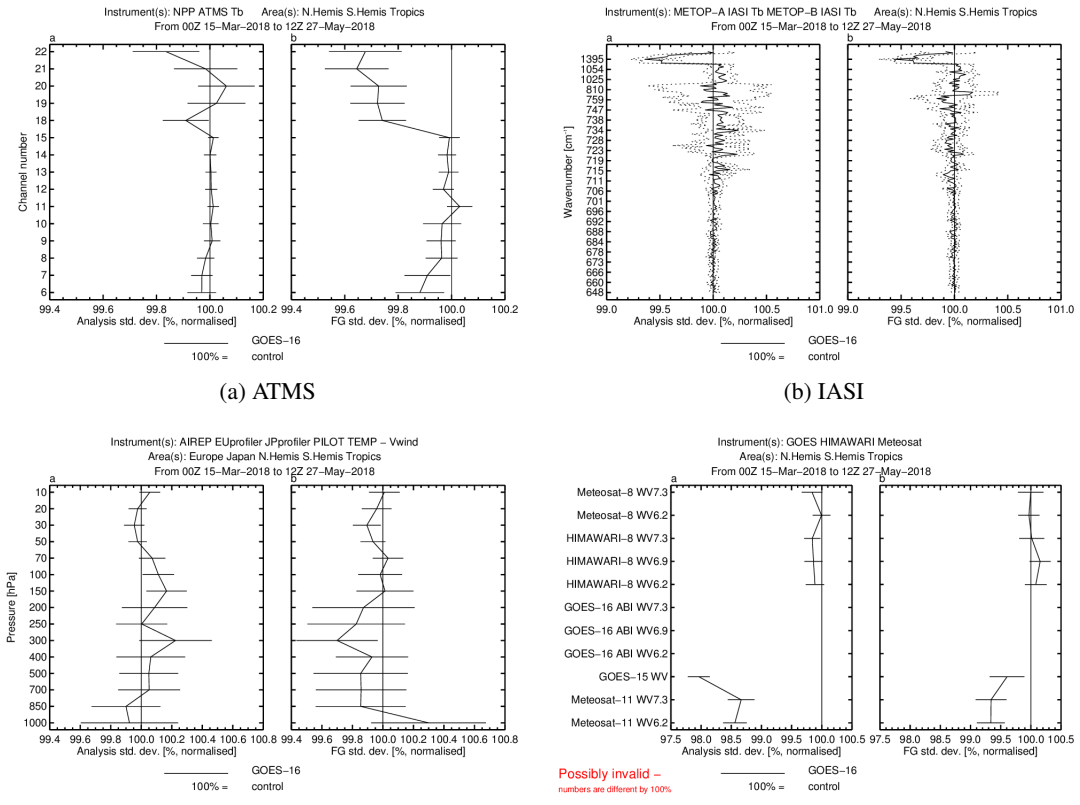

(a) ATMS

(b) IASI

(c) Conventional wind obs

(d) GEOs

Figure 5: Changes in analysis and first guess fits to independent observations when GOES-16 CSRs are introduced into the 4D-Var assimilation system.

Given the evidence from this experiment, the decision was made to begin assimilating GOES-16 CSRs on 26/7/2018, following a short period during which bias correction coefficients were allowed to converge to stable values.

Another tool to assess the impact of observations in the operational system is the forecast sensitivity to observation impact (FSOI), which uses adjoint techniques to determine the degree to which each observation type contributes to the total error reduction in a 24-hour forecast (Langland and Baker [2004], Cardinali [2009]). The ECMWF implementation (A. Geer, C. Lupu, B. Ingleby, A. Bonet) splits the impact from infrared sensors into those channels of each instrument that are sensitive to temperature and water vapour, and this is shown in Figure 6 for the month of May 2019.

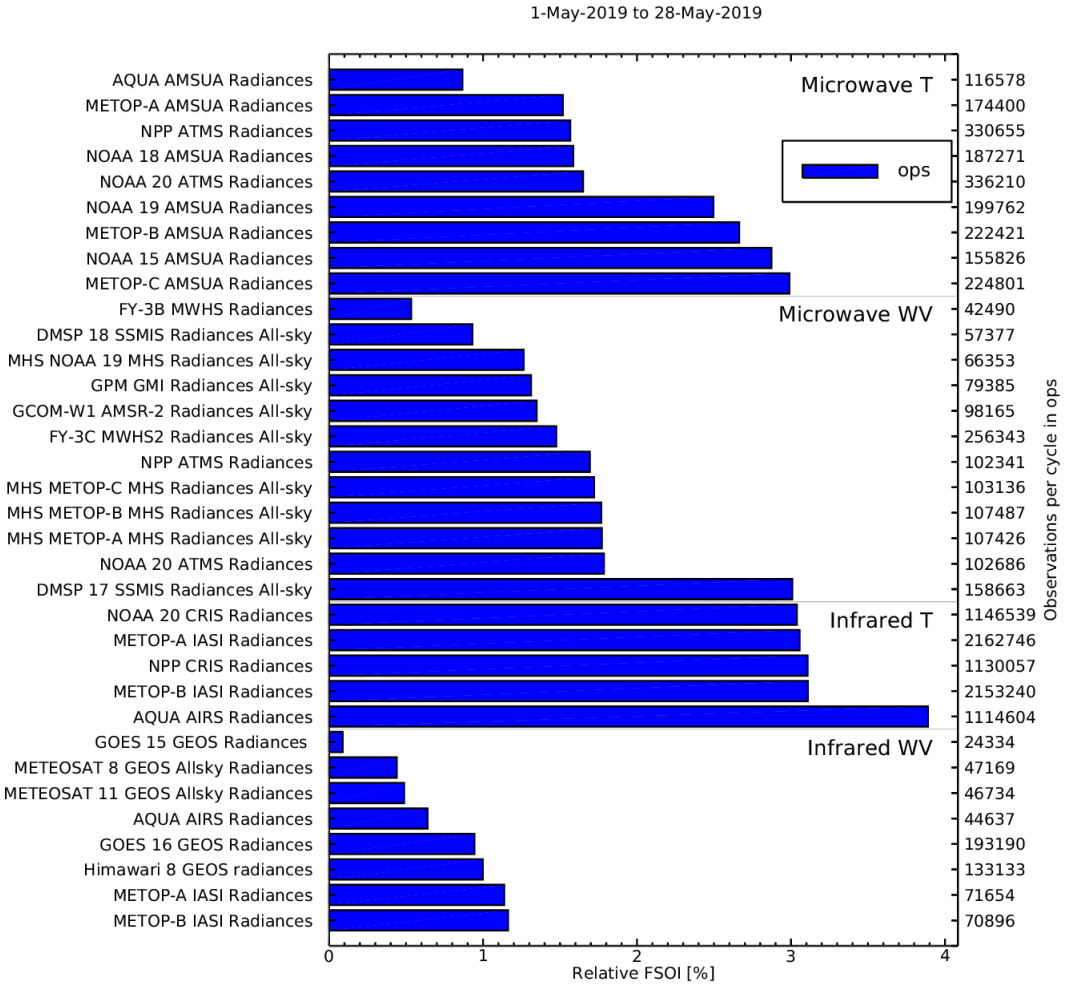

Figure 6: Relative FSOI from May 2019, for satellite radiance observations.

In the section for the water vapour channels, IASI provides the largest impact, closely followed by Himawari-8 and GOES-16. The relative order of Himawari-8 and GOES-16 varies from month to month. It is worth noting that Himawari-8 has less overlap with neighbouring satellites (see Figure 7), and hence more of the Himawari-8 observations are likely to provide unique information than those from GOES-16. On the other hand, ECMWF can effectively use half-hourly observations from GOES-16, whereas only hourly data are supplied from Himawari-8. The FSOI for the Meteosats and GOES-15 are smaller, likely due to them having fewer water-vapour sensitive channels.

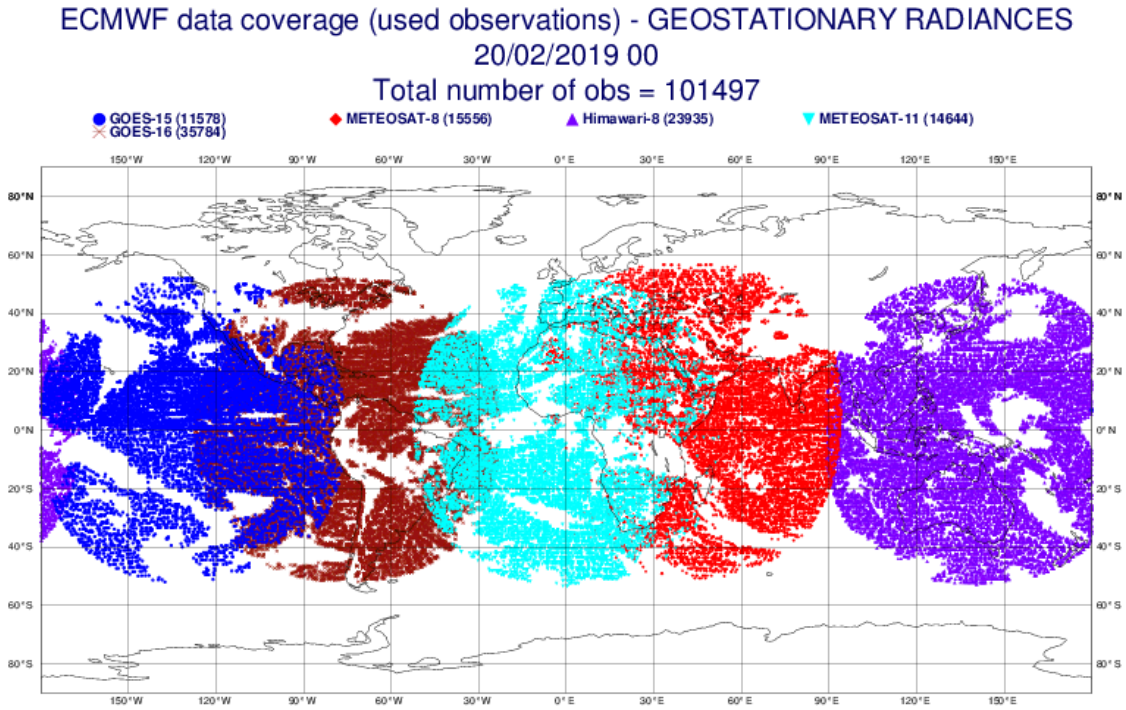

Figure 7: Coverage of geostationary radiances that were assimilated over a 6-hour period centred on 00Z, 20/2/2019.

1.4 Aerosol contamination in window channels

Some preliminary investigations were underway to test the assimilation of window-channel observations over ocean from the geostationary satellites, when it was noticed that a degraded forecast impact was produced in the region of the GOES-16 disk. A careful analysis found that this degradation was related to a plume of Saharan dust over the Atlantic ocean. This can be seen as the cold "tongue" on the eastern side of the disk in Figure 8a, which shows first guess departures for the $10.3\mu m$ channel. Meteosat-11 ASRs, however, did not show the same feature (Figure 8b), suggesting that the aerosol-contaminated observations had been successfully screened.

It may be worth noting that this region of aerosol contamination was in the daytime hemisphere, and co-incident with the sun-glint region, so routine visible-channel checks for aerosol may have been adversely affected by this, although this has not been confirmed.

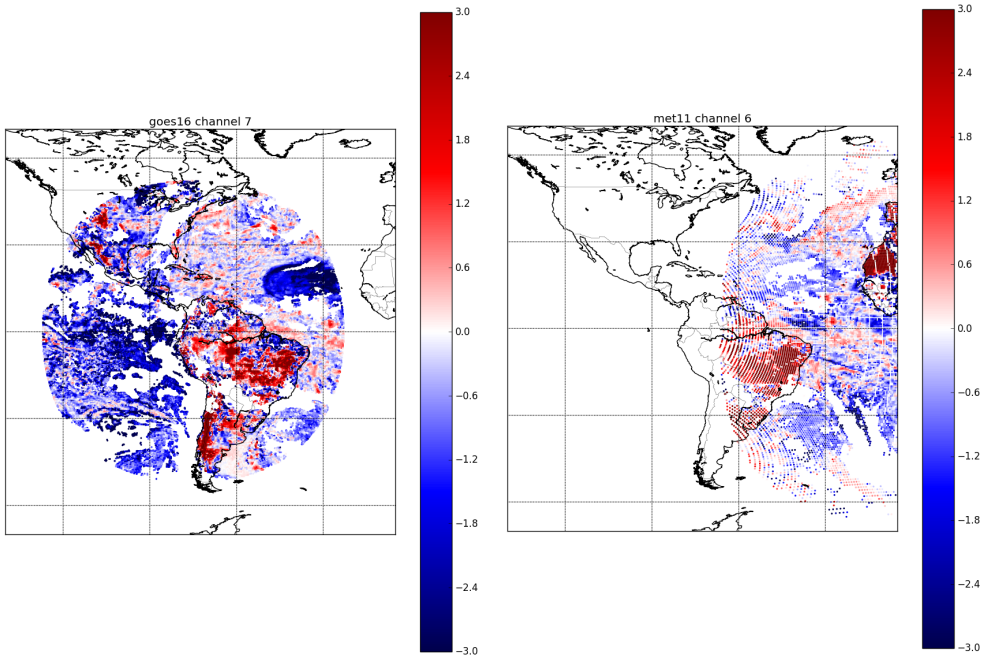

(a) GOES-16 (b) Meteosat-11

Figure 8: Window channel first guess departures on 2/8/2018, from GOES-16 (10.3μm, a) and Meteosat-11 (10.8μm, b).

In order to identify aerosol for hyperspectral infrared observations, a scheme has been proposed whereby particular channel differences are computed, and if a threshold is exceeded, then the scene is determined to be affected by aerosol (Letertre-Danczak [2016]). With geostationary radiances, this is not possible because the channels have broad spectral response functions (SRF), and the computed differences between any pair of channels will provide more information about the atmospheric state than about aerosol alone. Therefore it is proposed that instead of taking the difference of the observations from two channels, we instead examine the difference of first guess departures of two channels. The two channels are selected for having different responses to aerosol, but responses to the remaining atmospheric/surface variables that are as similar as can be attained. The double-differencing aims to cancel out the differences in sensitivity to other components of the non-aerosol state. If we separate the signal into that part which is due to aerosol, and that which comes from the rest of the atmosphere, for the two channels numbered 1 and 2, we have the following, where O denotes observation, and B denotes the simulated equivalent using the co-located background profile (which is aerosol-free) as input:

$$\begin{aligned}(O_1 - B_1) - (O_2 - B_2) &= (O_1^{atm} + O_1^{aer} - B_1^{atm}) - (O_2^{atm} + O_2^{aer} - B_2^{atm}) \\ &= \left(O_1^{atm} - B_1^{atm}\right) - \left(O_2^{atm} - B_2^{atm}\right) + (O_1^{aer} - O_2^{aer}) \\ &\simeq O_1^{aer} - O_2^{aer}\end{aligned}$$

That is, if the model background brightness temperature is sufficiently close to the atmospheric component of the observation (for each channel), the terms specific to the atmosphere become small, and it is seen that we are effectively measuring the difference in the aerosol-component of the observations. The two channels selected were the 10.3μm and 11.2μm channels. A threshold was determined empiri-

cally to reject as much contaminated data as possible, but minimising the number of false rejections; we assume that the scene is free of aerosol if the following inequality is satisfied:

$$(O_{10.3} - B_{10.3}) - (O_{11.2} - B_{11.2}) \quad > \quad -0.5K$$

The effect of applying this aerosol detection can be seen in Figure 9, where the contaminated observations in the window channel are successfully removed. The threshold is subject to refinement, and a stricter value of $-0.3K$ or $-0.2K$ may be better at ensuring that the less-badly affected observations are still screened, without too many false rejections.

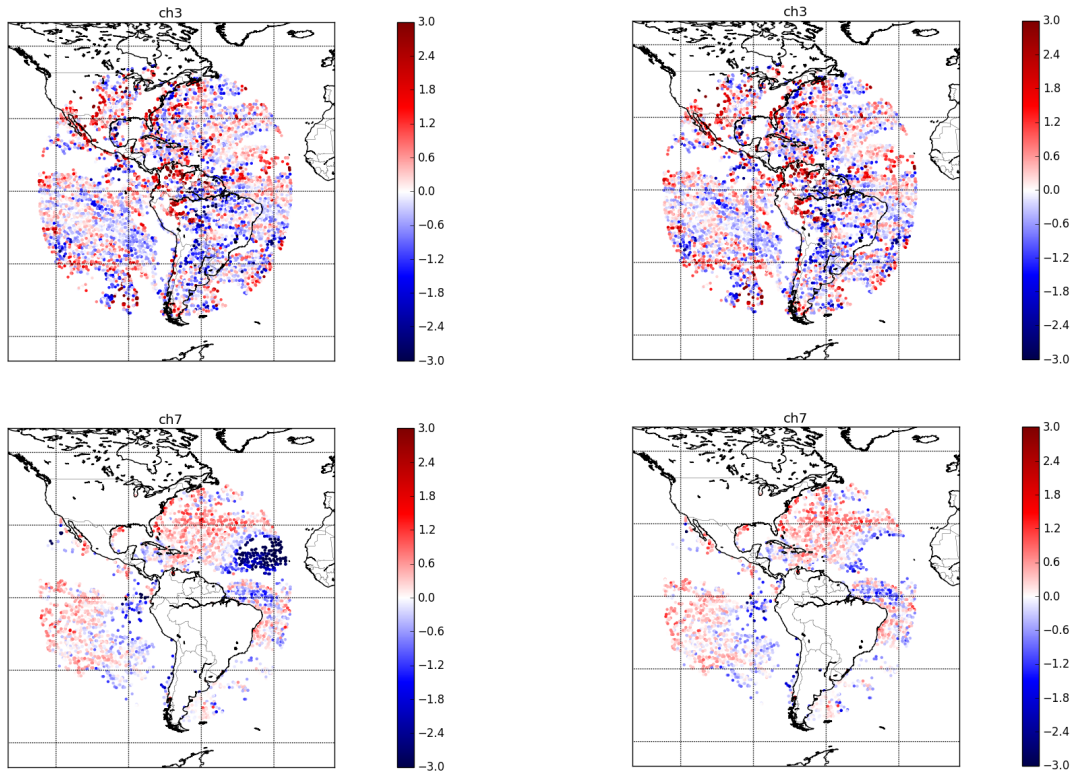

(a) (b)

Figure 9: GOES-16 first guess departures from 00Z on 2/8/18 that were assimilated in trial experiments without aerosol detection (a) and with aerosol detection (b). The two channels shown at the $7.0\mu m$ channel (top) and the $10.3\mu m$ channel (bottom).

If window channel radiances from GOES-16 are to be assimilated successfully, it will be necessary to implement such a check.

1.5 Diurnal bias of window channels

The GOES-16 window channel radiances show a significant diurnal bias which manifests as observations being too cold on the night side of the terminator. This signal has similarities to what was described in Section 1.2, and although the signal here is much smaller, it is always present, but only in the window

channels. When surface-sensitive satellite radiances are assimilated at ECMWF, a skin-temperature sink variable is added to the control vector for each observation, and this is adjusted during the minimisation. In the future, this information on sea surface temperature (SST) may be retained, and hence geostationary observations will play an important role in constraining the diurnal behaviour of the SST, so it is crucial that the observations themselves are not subject to diurnal biases.

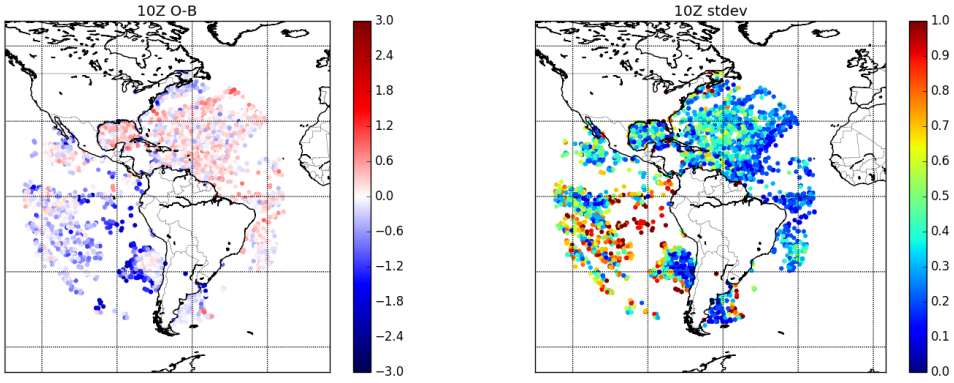

(a)

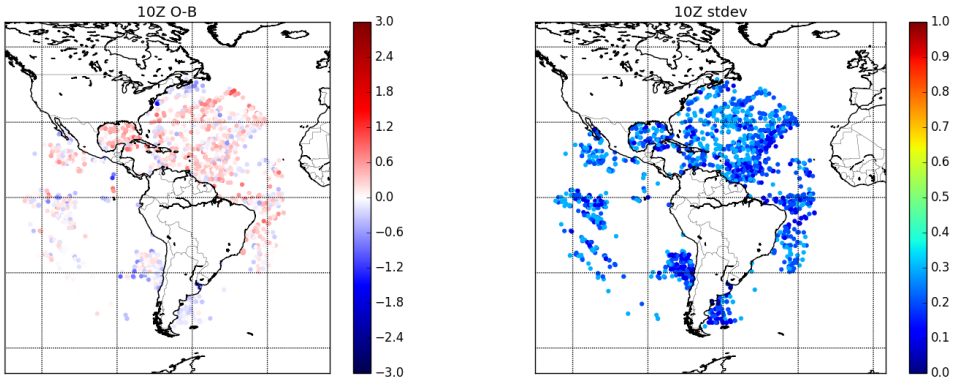

(b)

Figure 10: GOES-16 first guess departures from 10.3μm channel observations at 10Z on 24/8/18 that were assimilated in trial experiments without additional cloud detection (a) and with additional cloud detection (b). The left panel of each case show the first guess departures, and the right panel shows the standard deviation of the clear pixels that contribute to the CSR. The terminator runs approximately north-west to south-east, and it is night on the south-west side of this line.

It has been suggested (Sharon Nebuda and Haixia Liu, private communication) that these biases are less pronounced in the "enterprise" version of the GOES-16 CSRs, and that for the baseline product, the standard deviation of the clear pixels may help to screen out cloud-affected CSRs, particularly in the vicinity of marine stratocumulus regions. The effect of applying a 0.4K threshold on the standard deviations for the window channels can be seen in Figure 10.

It can be seen that the largest diurnal biases have been successfully removed, but that a change in sign

persists either side of the terminator. This may be related to a residual signal from undetected cloud or possibly a model bias, although such a stark day/night difference does not appear as strongly for Himawari-8 (not shown).

1.6 GOES-16 CSR version 2

Some preparations are in progress to make use of the version 2 of the GOES-16 CSR product. This contains the following changes (Tom King, Sharon Nebuda, Haixia Liu, private communication):

- Data are provided up to zenith angles of 74° rather than 60°.
- Observation locations are at the geometric centre of the averaging boxes, not at the "centre of gravity".
- "Sigma clipping" has been introduced to eliminate outliers in the averaging boxes.
- The BUFR format is changed to remove two unused near-infrared channels.
- Negative radiance values are set to MISSING.
- Satellite zenith angle calculation is improved.

Most of these have only slight impacts on the usage of the data, but the change to the BUFR format will require the data to be removed from operations for a period during which the system can be modified to account for the change. The increase in zenith angle has led to greater coverage for GOES-16, although an assimilation experiment has shown that this additional data has a neutral (i.e. not statistically significant) impact on the accuracy of short-range forecasts as verified against independent observations (results not shown here). The transition to version 2 will be made later in 2019.

1.7 Scanning mode change

On April 2nd 2019, the scanning mode of GOES-16 changed from Mode-3 to Mode-6. This resulted in full-disk scans being provided every 10 minutes instead of every 15 minutes. Currently, there is a threshold on the number of GOES-16 radiance observations that can be ingested into the ECMWF system, and this would be exceeded if Mode-4 scanning is tested, as it occasionally is (i.e. continuous 5-minute full-disk scanning). In order to protect against this, the filenames of the BUFR files are filtered such that they are only extracted if the observation times correspond to minutes of '00', '15', '30' or '45'. As mentioned in Section 1.3, the thinning applied at ECMWF effectively results in capping the temporal sampling of the observations to a minimum of 30 minutes. Therefore, approximately only half of the 15-minute data was being used. In the transition to Mode-6, the data will be timed with minute values of '00', '10', '20', '30', '40' and '50'. Consideration was given as to whether to modify the filtering to account for this change, but it was decided to take no action as the result would be that observations with minutes of '00' and '30' will be retained, and the loss of the observations at '15' and '45' should make little difference as the thinning would halve the number of used observations anyway. When the mode change took place, the number of *used* observations remained approximately constant despite the number of ingested observations reducing by a factor of two.

2 SEVIRI issues

Three short-lived data events occurred with SEVIRI all sky radiance (ASR) data in the past year, and these will be documented in this section.

2.1 Meteosat-8 $7.3\mu m$ channel issue — 8/1/2019

On 8/1/2019, the $7.3\mu m$ channel underwent a sudden bias shift. At 16Z, the first guess departures showed a large jump to positive values (approximately 4K) across the disk. The following hour, the departures became negative, with almost the same magnitude. Over the subsequent 5 to 6 hours, the bias gradually relaxed to typical values. This can be seen in a time sequence of first-guess departures, Figure 11.

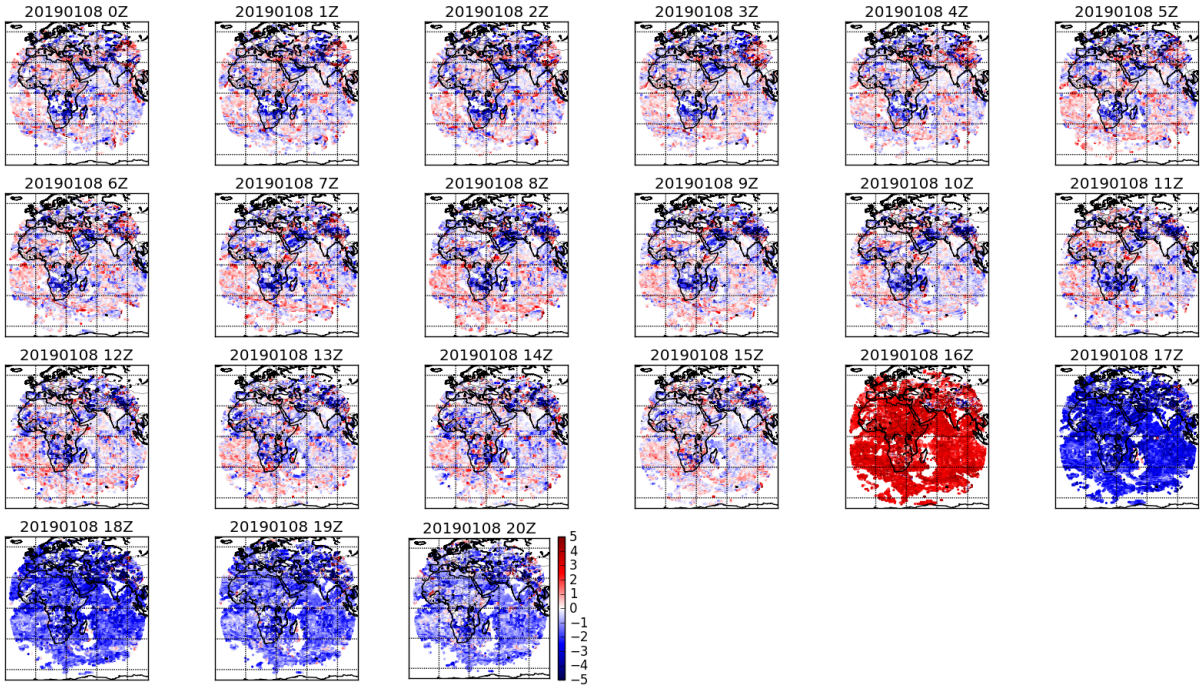

Figure 11: Time sequence of ASR first guess departures for the $7.3\mu m$ channel on Meteosat-8 on 8/1/19.

It is observed that the bias jump affected the disk uniformly. A histogram also shows the uniformity of the shift, and the behaviour on relaxation, Figure 12.

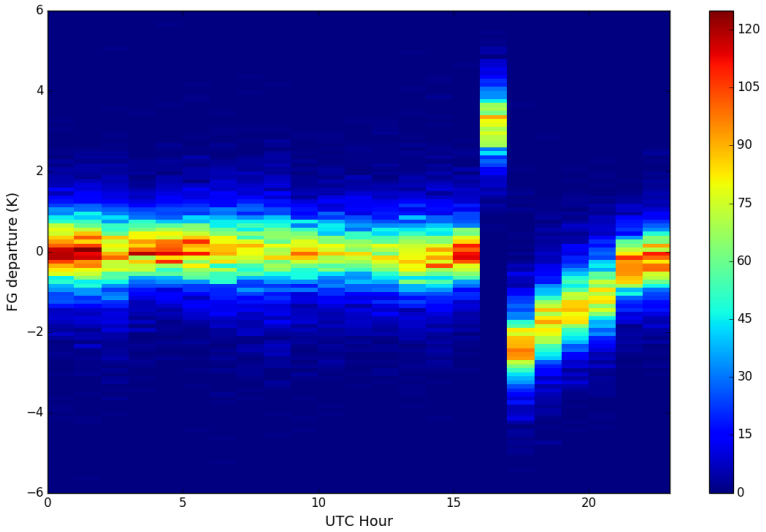

Figure 12: Temporal histogram of CSR first guess departures for the 7.3μm channel on Meteosat-8 on 8/1/19.

The cause of this has been investigated by EUMETSAT, but no conclusion had been drawn as of 10/1/19, although we are assured that this is unlikely to reoccur (EUMETSAT helpdesk, private communication).

2.2 Meteosat-11 decontamination

From 21/1/19 to 28/1/19, Meteosat-11 underwent routine decontamination to remove ice from the SEVIRI instrument. As a precaution, Meteosat-11 was rejected from being assimilated for this period. Before it was rejected, the first guess departures for the 13.4μm channel were becoming increasingly negative; this is a known feature of ice contamination (Hewison and Müller [2013]). The other channels are affected to a much lesser effect. Therefore, the variational bias correction acted to increase the brightness temperatures in order that the observations better matched the model background; compare the dashed and solid lines in the top panel of Figure 13 until 20/1/19. After the decontamination was performed, it can be seen that the raw first guess departures were close to zero, but the bias correction was still acting to increase the brightness temperatures. In the few days following the data stream being resumed, the behaviour of the bias correction can be examined, in particular the speed with which it adjusts to a step change in bias.

Although the 13.4μm channel is not assimilated directly, it is used as part of the cloud analysis scheme in the assimilation of the water-vapour channel ASRs. Prior uncertainty as to the data quality on recommencement of the stream led to the decision to resume the assimilation of Meteosat-11 data on 7/2/19, by which time the bias correction was fully spun-up.

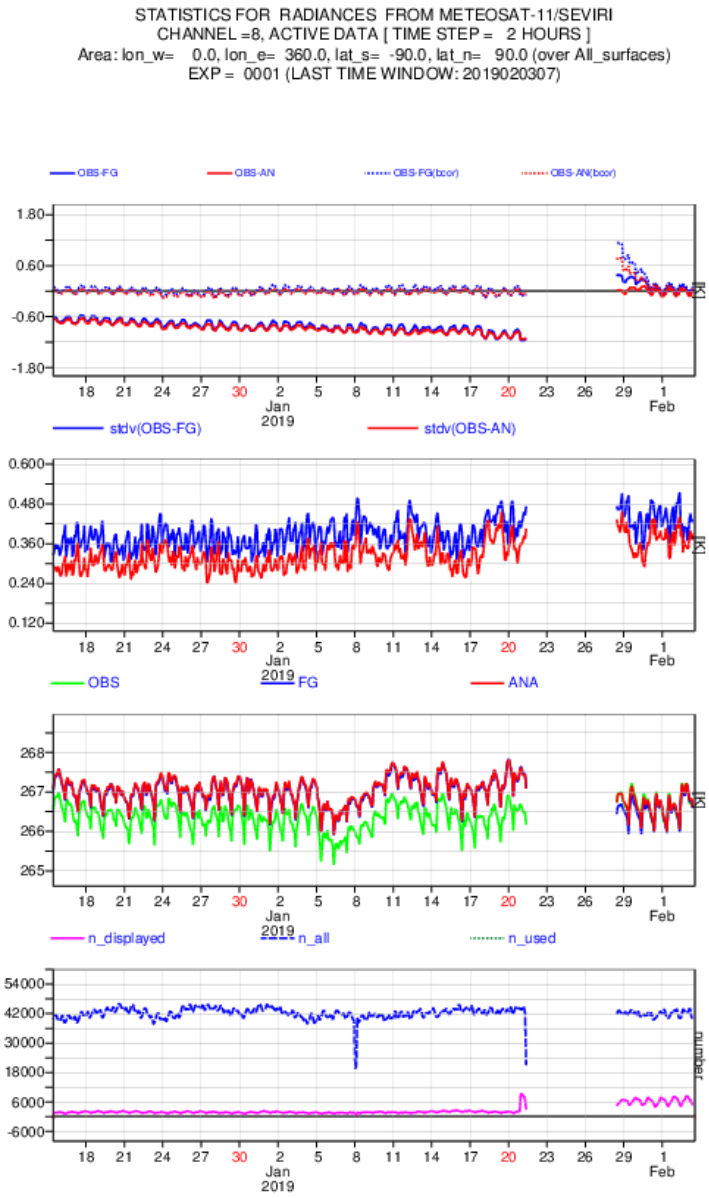

Figure 13: Time series of ASR first guess departure statistics for the 13.4μm channel on Meteosat-8 in January 2019, showing, in the top panel, the change in bias characteristics following the decontamination procedure.

2.3 Meteosat-11 eclipse issue

On 26/9/18, during an eclipse period, Meteosat-11 produced some unphysical ASR data for the 6.25μm channel. This was at the time of a sun-satellite-ground-station colinearity event. For the previous full-disk scan, there was no evidence of anything wrong with the data (see Figure 14). Here, all the observations are within normal physical bounds.

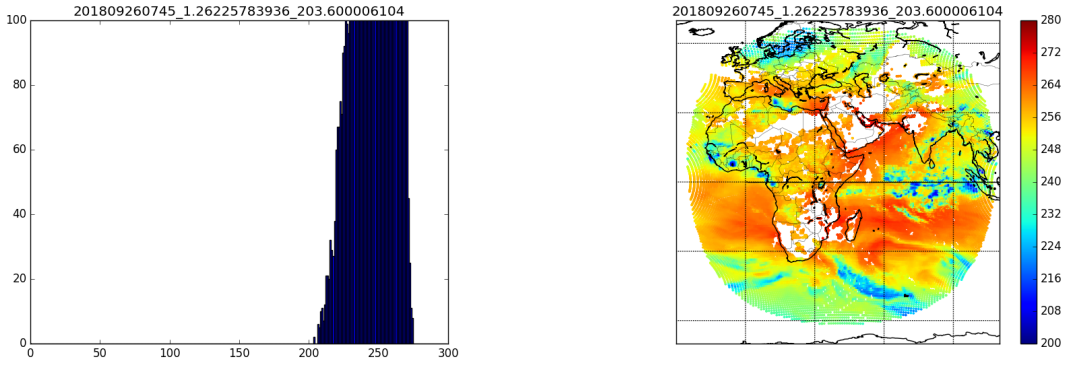

(a) (b)

Figure 14: Zoomed histogram of observed brightness temperatures at 0745Z on 26/9/18 from the 6.25μm channel (a) and the corresponding map of observed brightness temperatures (b).

However, an hour later (Figure 15), a large part of the disk was missing, and at the north edge of the southern portion of remaining observations, some extremely small brightness temperatures were provided, which can be seen in the histogram.

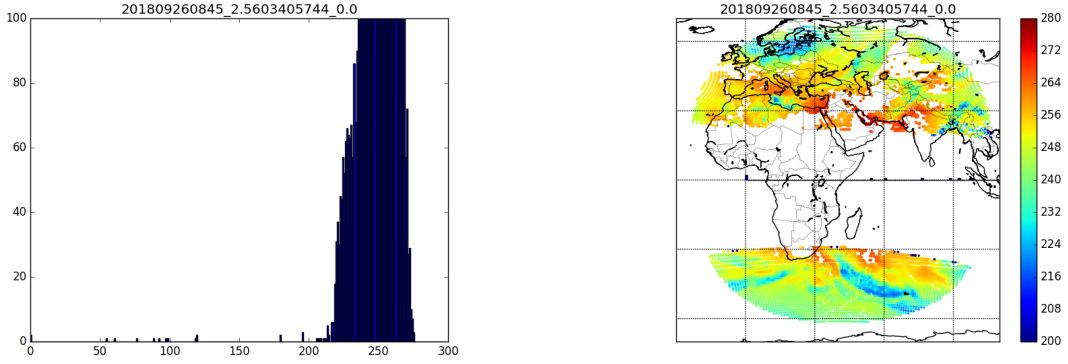

(a) (b)

Figure 15: Zoomed histogram of observed brightness temperatures at 0845Z on 26/9/18 from the 6.25μm channel (a) and the corresponding map of observed brightness temperatures (b).

At ECMWF, physical bounds checking is performed, but only for a single water vapour channel, which was assumed to be representative of all water vapour channels of a given sounding. In this instance, the channel that was checked was not affected as badly as the 6.25μm channel which led to unphysical observations getting into operations. These subsequently failed further quality control checks and did not cause any problems, but it exposed the potential for operational failures to occur in particular circumstances. Part of the ASR processing involves initialising the cloud analysis scheme, and this uses bias-corrected brightness temperatures as input. From Figure 15a it can be seen that brightness temperatures of exactly zero were supplied. Now, the cloud analysis scheme assumes that the brightness temperatures are positive, and a numerical exception will result if they are not. Therefore, for these zero-valued observations, a failure would occur if the bias correction takes the brightness temperature

below zero. For the channel affected here, the bias correction fluctuates around zero with time. On the day of this issue, the operational bias corrections were positive, so a failure did not occur, but a check has since been implemented to screen all the channels that will be used, and EUMETSAT has expressed the intention to implement a fix to prevent such affected observations being supplied as part of the ASR product.

3 Varying observation errors throughout the assimilation window

When observations are assimilated in strong-constraint 4D-Var, as is the case at ECMWF (at least in the troposphere), they must be assigned observation errors. Consider first the 3D-Var cost function, where J is the scalar cost function, to be minimised, $\mathbf{x}$ is the model state vector, $\mathbf{x}_b$ is the background state vector, $\mathbf{y}$ is the observation vector (which are all valid at the same time), $\mathbf{B}$ is the background error covariance matrix, $\mathbf{R}$ is the observation error covariance matrix and $H(\cdot)$ is the nonlinear observation operator:

$$J(\mathbf{x}) = (\mathbf{x}-\mathbf{x}_b)^T\mathbf{B}^{-1}(\mathbf{x}-\mathbf{x}_b) + (\mathbf{y}-H(\mathbf{x}))^T\mathbf{R}^{-1}(\mathbf{y}-H(\mathbf{x})) \tag{1}$$

The value of $\mathbf{x}$ which minimises this function will be that which provides an optimal fit to the background and observations given their respective uncertainties. In the first term, the background error covariance matrix represents the uncertainty of the background state alone, but in the second term, the observation error covariance matrix must represent the uncertainty in the observations *and* the observation operator. To understand this, consider a scalar case where the observation error is perfect (denoted by a prime):

$$J(x) = \frac{(x-x_b)^2}{\sigma_b^2} + \frac{(y-H'(x))^2}{\sigma_o^2} \tag{2}$$

In this case, the observation error must match the uncertainty in the observations themselves. But if the forward model is allowed to impart uncertainty, when the second term is evaluated, it is impossible to determine how much of the uncertainty comes from the observations and how much comes from the forward model and hence the "observation" error, σ_o must also include the component of uncertainty originating from suboptimalities in the observation operator.

In 4D-Var, the forecast model is used to account for time-dependence, in order that more representative fields may be compared to observations which are made at various times throughout the assimilation window. The non-linear forecast model $M(\cdot)$ can be thought of as part of the observation operator, and hence the strong-constraint 4D-Var cost function becomes:

$$J(\mathbf{x}) = (\mathbf{x}-\mathbf{x}_b)^T\mathbf{B}^{-1}(\mathbf{x}-\mathbf{x}_b) + (\mathbf{y}-H(M(\mathbf{x})))^T\mathbf{R}^{-1}(\mathbf{y}-H(M(\mathbf{x}))) \tag{3}$$

Importantly, the first term is unchanged; the background fields and the optimal state to be computed are both valid at time $T = 0$, and the time dependence is only present in the second term. "Strong constraint" 4D-Var assumes that the model uncertainty does not vary with lead-time, and in "weak constraint" 4D-Var, an additional term is added to the cost function which explicitly penalises the model at longer lead times.

Radiances from geostationary satellites enable statistics to be computed with relatively high temporal sampling and thus it is possible to obtain meaningful first guess departure statistics for different time bins within the assimilation window. Figure 16 shows the standard deviation of the first guess departures of the $6.25\mu m$ channel for Himawari-8 as a function of the hourly bins throughout the 12-hour assimilation window.

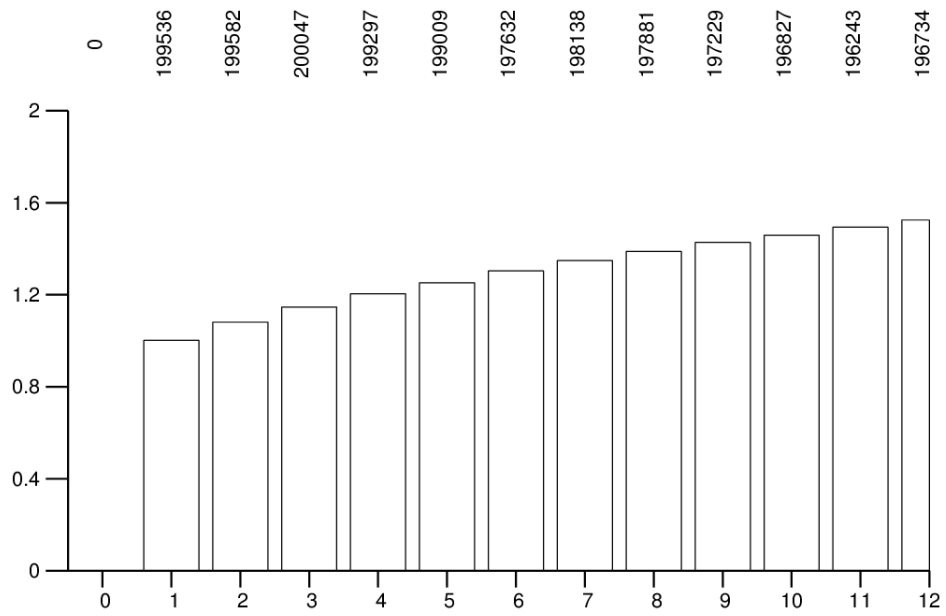

Figure 16: Standard deviation of first guess departures of Himawari-8 CSRs from the $6.25\mu m$ channel as a function of the time in hours from the start of the assimilation window. The values above the bars indicate the numbers of observations that contributed to the statistics of each bin.

These standard deviations increase monotonically throughout the assimilation window, and this is due to the fact that model error is increasing with time, and the forecast model forms a part of the generalised observation operator. There is no reason to expect that the observations themselves have errors that vary in this way.

As an aside, we note that it is possible to diagnose observation errors using the Desroziers method (Desroziers et al. [2005]) as a function of time throughout the window. A condition imposed in the Desroziers method is that the departure variances will be equal to the sum of the diagnosed observation error variances and background error variances (mapped into observation space of the form $\mathbf{HBH}^T$). For Himawari-8, the resulting errors are plotted in Figure 17.

Interestingly, the diagnosed observation errors remain approximately constant throughout the assimilation window, but the background errors grow monotonically. This is to be expected as the "backgrounds" later in the window are not strictly backgrounds but are forecasts from the background field, which is valid at the start of the window, therefore, the "background errors" at the end of the window are larger due to the contribution from the model error. Typically, when the Desroziers method is applied, observations from throughout the window are included in the statistics, and given that for strong-constraint 4D-Var, the R-matrix should contain forward-model error, we can surmise from Figure 17 that our diagnosed observation errors may be under-estimated. This would be consistent with the practical necessity of inflating the diagnosed observation errors to give the optimum forecast impact (Bormann et al. [2015]) (Pete Weston, Tony McNally, Sean Healy and Neill Bowler, private communications).

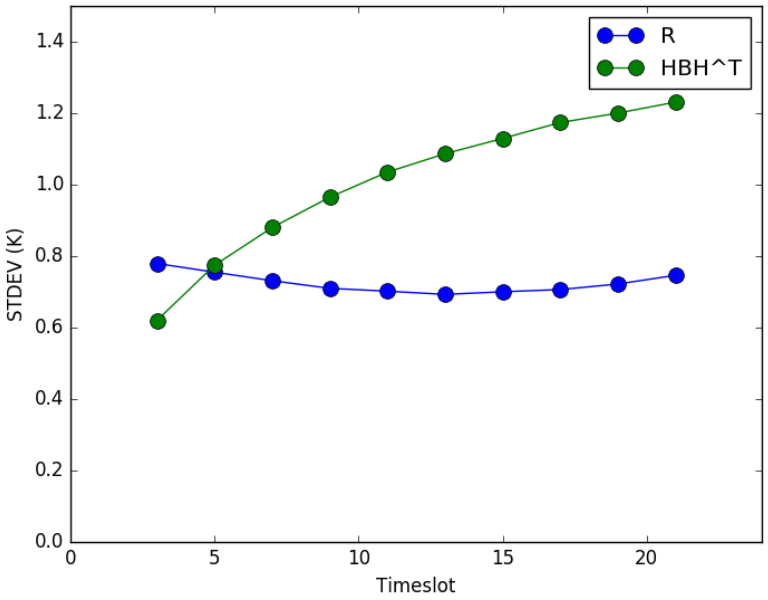

Figure 17: Diagnosed observation and background errors for Himawari-8 CSRs from the $6.25\mu m$ channel as a function of the time in hours from the start of the assimilation window.

As the model uncertainty grows with increasing lead time, one possible approach within the strong-constraint formulation would be to account for model error as part of the observation error covariance matrix (Howes et al. [2017]). This would require lead-time-dependent R-matrices, with larger values towards the end of the assimilation window, and in theory should result in improved analyses.

The time-dependence of the observation errors would be such that the later observations are down-weighted towards the end of the assimilation window. Is this desirable? Indications are that the later observations are among the most important for forecast skill McNally [2019], and of course these are the observations that are closer to the times at which the forecasts are being verified. So, it seems that there may be a discrepancy between what is optimal for the analysis and what is optimal to improve forecast skill.

Fundamentally, there should be no reason that the time throughout the window should affect the quality of the *actual* observations, but by imposing observation errors that increase with forecast lead time in the window, the later observations become less able to influence the analysed state. In terms of the optimality of the analysis, this is the correct thing to do; we are less certain of the state at the end of the window, so if the poorly-forward-modelled observation equivalents are close to the observations themselves, small observation errors will result in a state at the end of the window that matches the observations well, but is worse at the start of the window. In NWP, what we are actually interested in optimising is the trajectory given the forecast operator. Therefore, if we break down the 4D-Var problem into a series of 3D-Var problems, solved at points distributed throughout the window, we have

$$J_t = (\mathbf{x}_t - \mathbf{x}_{b,t})^T \left(\mathbf{M}_t \mathbf{B} \mathbf{M}_t^T\right)^{-1} (\mathbf{x}_t - \mathbf{x}_{b,t}) + (\mathbf{y}_t - H(\mathbf{x}_{b,t}))^T \mathbf{R}^{-1} (\mathbf{y}_t - H(\mathbf{x}_{b,t})) \tag{4}$$

where

$$\mathbf{x}_{b,t} = M_t\left(\mathbf{x}_b\right) \tag{5}$$

This is slightly contrived, but illustrates that if we are interested in the solution at the end of the window $\mathbf{x}_T$, the background term will be penalised by a larger error and hence the observations will automatically be given more weight comparatively.

In order to explore the effect of scaling observation errors throughout the assimilation window, a series of experiments were run. These were full-system 4D-Var experiments, but the observation errors for the geostationary radiances alone were varied throughout the assimilation window. Note that here, inter-channel observation error correlations are neglected, and the default observation error standard deviation is 2K. The time-variation of the observation error is parametrised according to the following equation:

$$\sigma_o = \sigma_{default} + \frac{2ct}{T} - c \tag{6}$$

where t is the time within the window, T is the length of the window and c is a parameter that indicates the value by which the default observation error, $\sigma_{default}$, is increased by at the end of the window (of course, the magnitude of c must be less than the value of σ_o to avoid generating negative observation errors). The parametrisation is chosen such that the observation errors vary linearly throughout the window. The mean of the standard deviation of the error averaged across the window is unchanged, although the net weight given to the observations will not be the same for each experiment as it is the variance that determines this weight. Despite this, for a given value of c, if the sign is changed, the net weight of the observations will be unchanged, so these pairs may be compared fairly (assuming that the observations are distributed uniformly throughout the window). This is demonstrated in Figure 18 for a range of values of c.

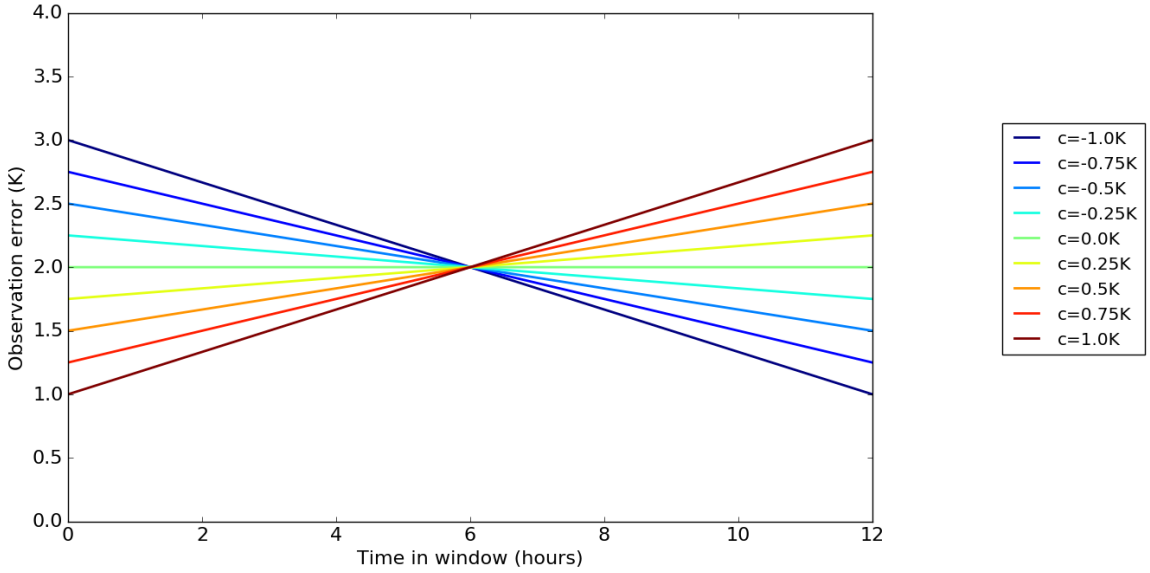

Figure 18: Observation errors as scaled throughout the 12-hour assimilation window, by choosing a range of values of the parameter c in Equation 6 and a value of $\sigma_{default} = 2K$.

The experiments using these observation error specifications for the geostationary radiances have been

run and the forecast impact is fairly neutral, but significant signals are present in the change of the first-guess fits to independent observations, as can be seen in Figure 19 for ATMS. The first-guess fit to other humidity-sensitive observations shows the same trend, but no impact is discernible for temperature-sensitive observations.

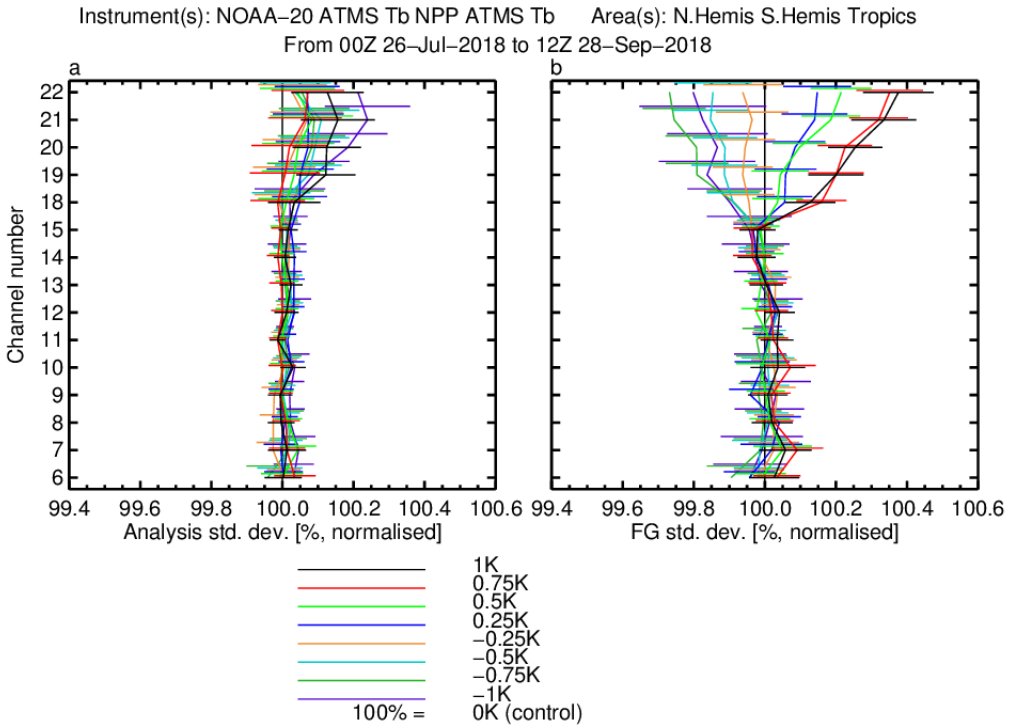

Figure 19: Change in standard deviations of ATMS first guess departures for a series of experiments when the observation errors are varied throughout the assimilation window using a number of values of the parameter c.

By comparing Figures 18 and 19 it can be seen that for the ATMS water-vapour channels (18-22), the first guess fits are improved when the observation errors are larger at the start of the window and smaller at the end. Similarly, the fit is degraded when the observation errors are larger at the start of the window. This suggests that short-range forecasts can be improved when the observations are given more weight closer to the end of the assimilation window (i.e. closer to the time for which the verifying first-guess departures are computed). The mechanism for this improvement is yet to be determined, although it could be a compensatory effect for another deficiency in the assimilation system.

This sensitivity needs to be investigated further in order to develop any firm conclusions, but the initial indications are that this approach may be encouraging, and consistent with the reasons given earlier in this section, although these ideas will need to be consolidated in due course, and thus these results should only be taken as a starting point for future investigations. (Thanks to Tony McNally and Elias Holm for interesting discussions.)

4 GIIRS

4.1 Introduction

Launched in December 2016, the Chinese geostationary satellite FY-4A is the first geostationary satellite to carry a hyperspectral infrared sounder; GIIRS. It is located at $105°E$ and also carries an advanced imaging instrument, AGRI.

GIIRS comprises two spectral bands; the long-wave band covers 700-1130cm^{-1} and the mid-wave band covers 1650-2250cm^{-1}. The maximum optical path difference (OPD) is 0.8cm, and this determines the spectral resolution of the instrument which is 0.625cm^{-1}, and is the same as that of CrIS. Given these parameters, we can determine that there are 1650 channels in total.

The GIIRS detector array, and hence the field of regard (FOR) is approximately rectangular and comprises 128 detector elements with 32 in the north-south orientation and 4 in the east-west orientation. Only the northern hemisphere is currently being sampled, with the complete domain being covered every two hours (with the exception of the 16Z scan, which is omitted). During this two hour period, the domain is sampled with 60 FORs in the east-west, and 7 in the north-south.

Radiance data from GIIRS have been available in near-real-time since January 2019, via the CMA/NSMC website and, from mid-March, via EUMETCast-Terrestrial, and the analysis presented here uses observations obtained from the two sources. The coverage of GIIRS is shown in Figure 20.

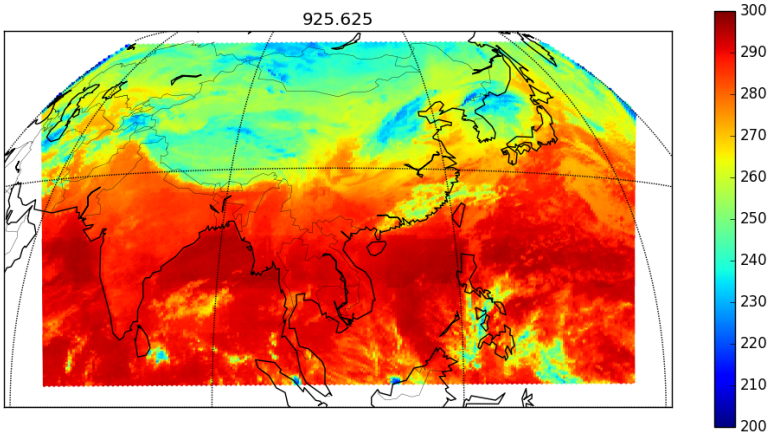

Figure 20: Coverage scanned by GIIRS. Data shown are the brightness temperatures from the 925.625cm^{-1} channel. Data are from the scan starting at 00Z on 3/2/2019.

Certain characteristics of the data need to be ascertained in order to be able to make a meaningful comparison with radiative transfer simulations. Firstly, the supplied radiances need to be converted into brightness temperature equivalents using the Planck function:

$$T_B(\nu) = \frac{hc\nu}{k_B \ln\left(1 + \frac{2hc^2\nu^3}{B(\nu)}\right)} \tag{7}$$

where T_B is the brightness temperature, ν is the wavenumber, h is the Planck constant, c is the speed of light, k_B is the Boltzmann constant and B is the radiance. Official RTTOV coefficients have not been produced for GIIRS yet, although coefficients have been made available by CMA/SSEC (Di et al. [2018]). We have been informed that these coefficients assume that the observations have been apodised with a Hamming function (Qifeng Lu, private communication), so in order to use these data, it is important to know if the observations we receive have been apodised or not.

4.2 Apodisation

Initially, an observed brightness temperature spectrum was compared against a simulation produced using RTTOV along with the CMA/SSEC RTTOV coefficients, see Figure 21. The model input to RTTOV was a co-located profile from an operational ECMWF analysis.

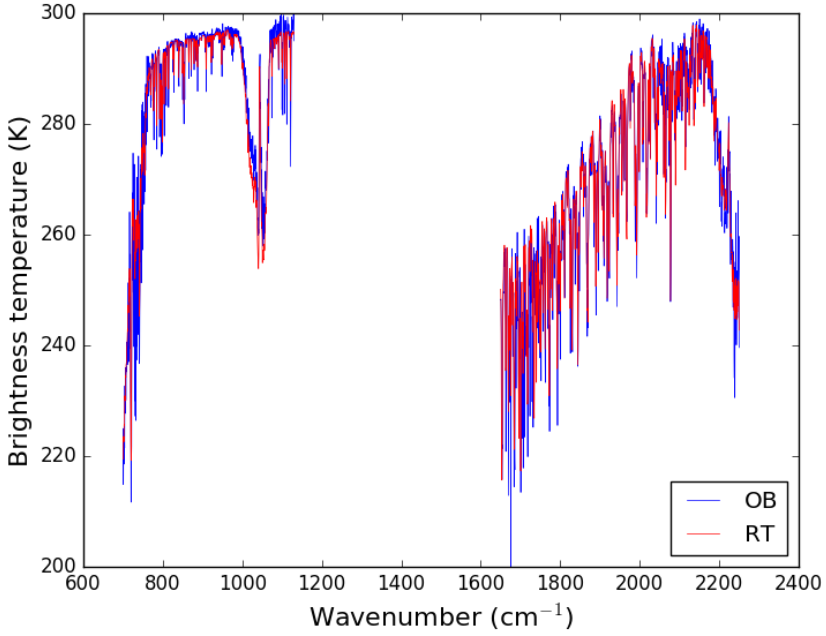

Figure 21: Observed and simulated GIIRS spectra, with no apodisation applied to the observations.

The observations exhibit significantly larger variability than the simulations, and hence the spectral resolution of the observations appears to be higher. This is consistent with the hypothesis that the observations received are unapodised. Unapodised spectra are subject to a "sinc" function instrument lineshape (ILS), which is the Fourier transform of the top hat function which defines the optical path difference (OPD) limits due to the motion of the moveable mirror, in the absence of attenuation or self-apodisation effects (Griffith and de Haseth [1986]). This is shown in Figure 22 (zero-filling has been applied); note the large negative sidelobes in the ILS. The top hat function is essentially a filter which we multiply the *true*, infinite interferogram with. When the Fourier transform is performed, the result is the true spectrum, *convolved* with the ILS. The sinc function represents the maximum possible spectral resolution for the given maximum OPD. In this example (using the parameters for GIIRS), the maximum OPD is 0.8cm, and the first zero-crossing of the ILS occurs at $0.625 cm^{-1}$. The larger the maximum OPD, the narrower

the central lobe of the ILS.

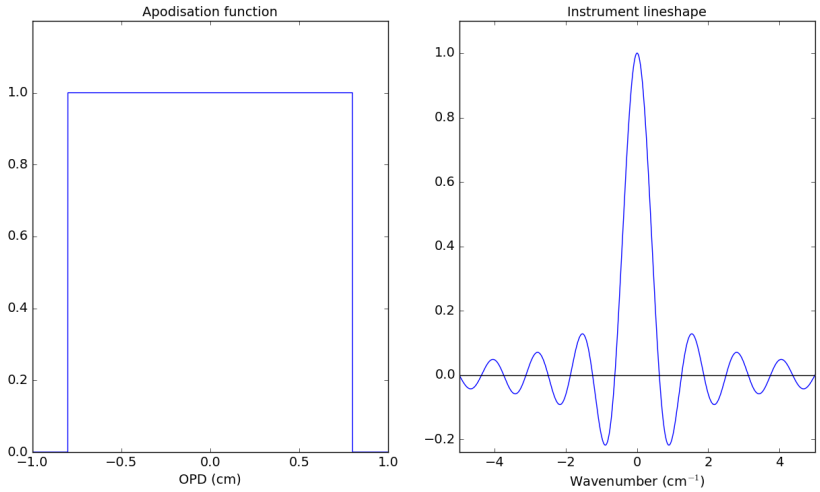

Figure 22: Top hat apodisation function and its corresponding instrument lineshape (sinc function).

The Hamming apodisation function takes the following form in interferogram space:

$$A(x) = (1 - 2a_0) - 2a_0 \cos\left(\frac{\pi x}{x_{max}}\right) \tag{8}$$

where $a_0 = 0.23$, x is the optical path difference (OPD) of the interferometer and x_{max} is the maximum optical path difference (MPD). When this apodisation function is applied to an interferogram, it suppresses the tails which contain information about the most detailed spectral structures, and hence apodisation reduces the effective spectral resolution of the spectra. An important reason for applying such an apodisation function is so that negative side-lobes in the instrument line shape (ILS) are suppressed, which can otherwise prevent the generation of fast radiative transfer coefficients. In comparison to other apodisation functions, the Hamming function has the benefit that a Hamming-apodised spectrum can be unapodised easily as its values are always non-zero for measured OPD values.

The Hamming function is shown in Figure 23 with its corresponding ILS. Compared to the top hat function, the side-lobes are substantially reduced in magnitude, but the central lobe is broader, hence when a raw spectrum is convolved with the Hamming ILS, the spectral resolution is reduced.

For a given maximum OPD, there is a limit to the resolution, as we have seen above. It is possible to gain a more intuitive interpretation of the resolution by considering the interferograms of two plane waves with the same amplitudes but slightly different frequencies:

$$\begin{aligned} I_1 &= \cos(2\pi\nu_1 x) \\ I_2 &= \cos(2\pi\nu_2 x) \end{aligned} \tag{9}$$

where ν_1 and ν_2 are the wavenumbers and x is the OPD. The combined interferogram is:

$$\begin{aligned} I_1 + I_2 &= \cos(2\pi\nu_1 x) + \cos(2\pi\nu_2 x) \\ &= 2\cos(\pi x(\nu_1+\nu_2))\cos(\pi x(\nu_1-\nu_2)) \end{aligned} \tag{10}$$

and in order to fully resolve this interferogram and uniquely identify the two frequencies, at least one entire cycle of the envelope function (frequency $\nu_1 - \nu_2$) must be sampled, which is consistent with the position of the first minimum of the sinc function.

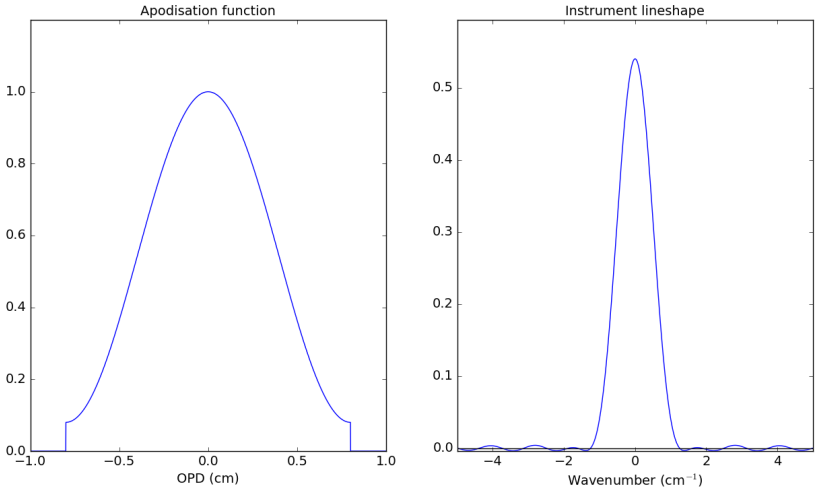

Figure 23: Hamming apodisation function and its corresponding instrument lineshape.

For GIIRS, the maximum OPD is 0.8cm and hence the spectral resolution is $0.625cm^{-1}$. If our samples in spectral space are separated by $0.625cm^{-1}$, a useful characteristic of the Hamming ILS shape can be exploited. The instrument lineshape with Nyquist sampling is given by the discrete Fourier transform of the Hamming function (Hamming [1989]):

$$I(\nu) = \frac{\sin 2\pi\nu x_{max}\cos\pi\nu}{\sin\pi\nu}\left[(1-2a_0) - 2a_0\left[\frac{\sin^2\pi\nu}{\sin^2\pi\nu - \sin^2\pi/(2x_{max})}\right]\right] \tag{11}$$

This ILS can be evaluated for relative channel indices, but L'Hôpital's rule is required to account for the singularities at $\nu = 0$ and $\nu = 1/(2x_{max})$:

$$\begin{aligned} I(\nu = -1/(2x_{max})) &= 2a_0 x_{max} \\ I(\nu = 0) &= 2(1-2a_0)x_{max} \\ I(\nu = 1/(2x_{max})) &= 2a_0 x_{max} \end{aligned} \tag{12}$$

and for all other multiples of $1/(2x_{max})$, the instrument lineshape is zero. Therefore, the apodised radiance for a given spectral channel is equal to a weighted sum of the unapodised radiance of that channel, and its two neighbouring channels. Normalising, to ensure the weights sum to unity, we have:

$$\begin{aligned} I_{norm,i-1} &= a_0 = 0.23 \\ I_{norm,i} &= (1-2a_0) = 0.54 \\ I_{norm_{i+1}} &= a_0 = 0.23 \end{aligned} \tag{13}$$

So, we have a simple means for apodising unapodised spectra very easily in spectral space, by means of a weighted running average using the coefficients in Equation 13. This was applied to a GIIRS spectrum, and it can be seen that the spectral structure is a better match to the simulations calculated with RTTOV, see Figure 24, c.f. Figure 21. Therefore, the subsequent observations considered here will all have had Hamming apodisation applied, although in the next section, unapodised observations are used in order to assess better the observations themselves rather than there being an influence from additional processing.

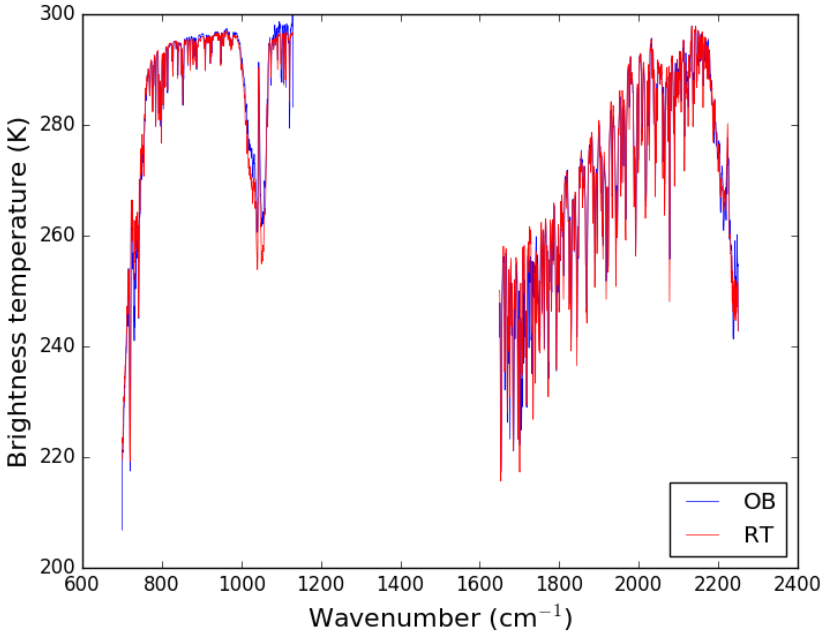

Figure 24: Observed and simulated GIIRS spectra, with Hamming apodisation applied to the observations, c.f. Figure 21.

4.3 Initial assessment of observations

Following the conversion to brightness temperature, the observations from individual scans have been plotted in order to examine whether any spatial features exist in the raw data. These are shown in Figure 25 for 00Z to 10Z and Figure 26 for 12Z to 22Z.

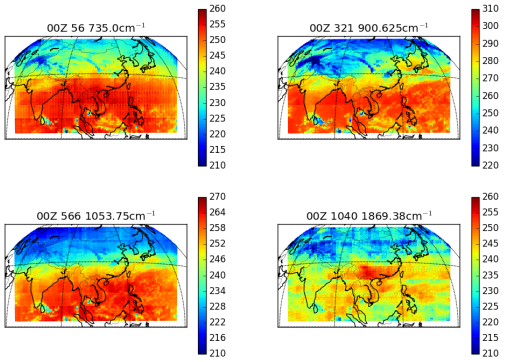

(a) 00Z

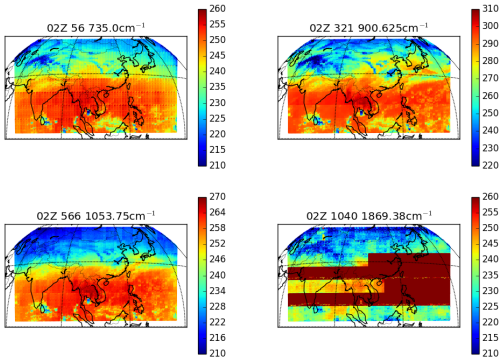

(b) 02Z

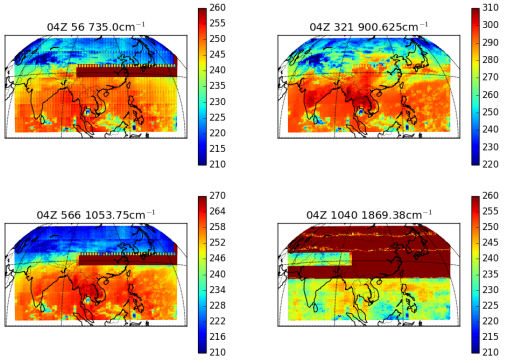

(c) 04Z

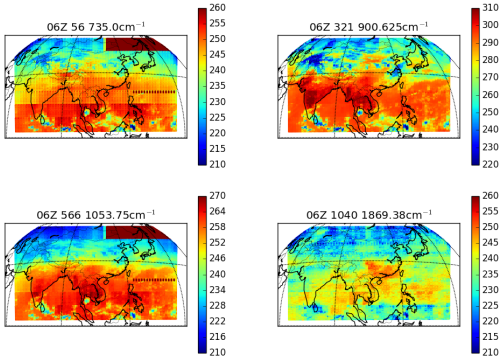

(d) 06Z

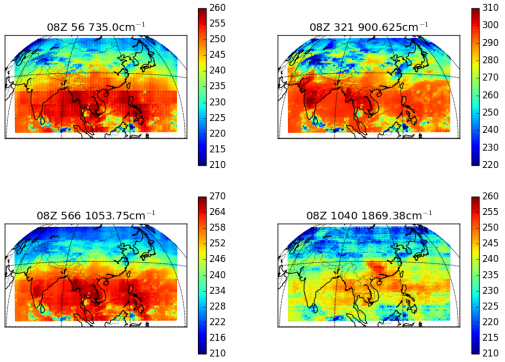

(e) 08Z

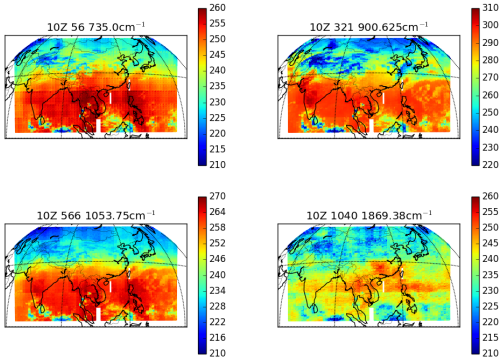

(f) 10Z

Figure 25: Unapodised GIIRS observed brightness temperatures for the first six 2-hour scans on 5/2/2019. Unapodised observations are shown for four channels that cover the key parts of the measured spectrum; $735cm^{-1}$ in the CO_2 band (top left), $900.625cm^{-1}$ which is surface sensitive (top right), $1053.75cm^{-1}$ in the ozone sounding band (bottom left) and $1869.375cm^{-1}$ in the water vapour band (bottom right).

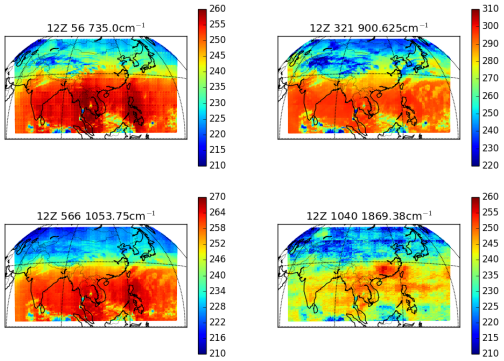

(a) 12Z

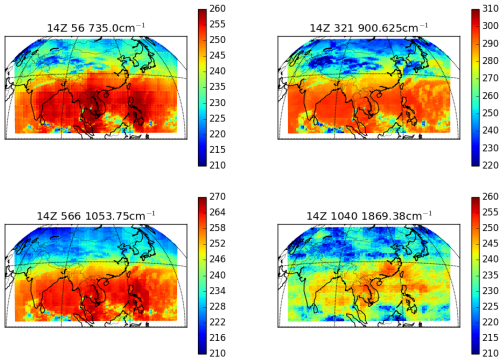

(b) 14Z

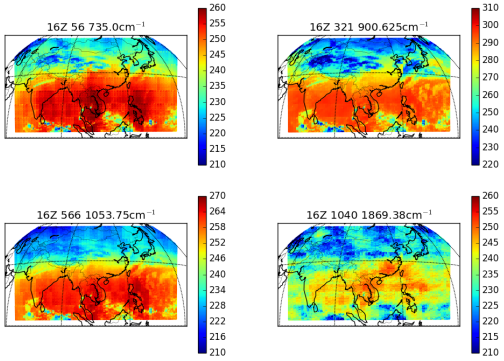

(c) 16Z

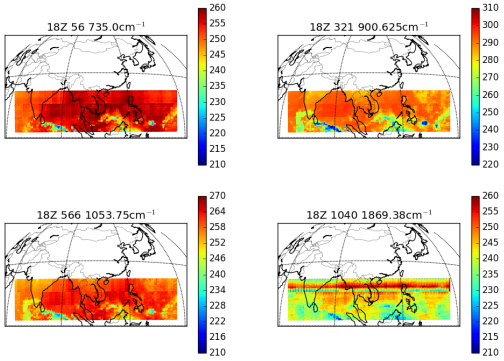

(d) 18Z

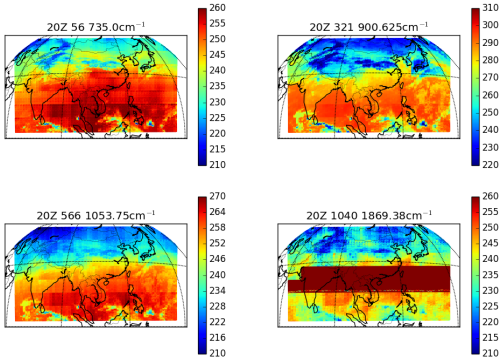

(e) 20Z

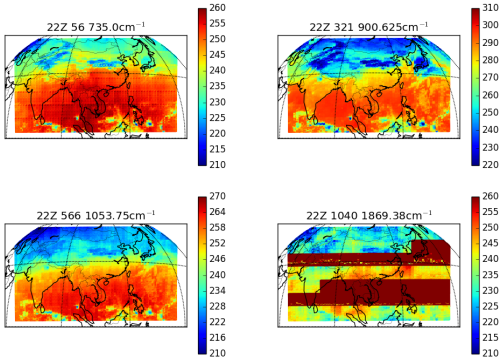

(f) 22Z

Figure 26: Unapodised GIIRS observed brightness temperatures for the final six 2-hour scans on 5/2/2019. Unapodised observations are shown for four channels that cover the key parts of the measured spectrum; $735cm^{-1}$ in the CO_2 band (top left), $900.625cm^{-1}$ which is surface sensitive (top right), $1053.75cm^{-1}$ in the ozone sounding band (bottom left) and $1869.375cm^{-1}$ in the water vapour band (bottom right).

Several features are apparent in these observations, some of which are cause for concern. Figures 27 to 36 show the geographic map plots of the apodised observations, the equivalent simulations and the differences for a number of channels. The clear-sky simulations were computed with RTTOV and the NWP SAF radiance simulator. The data are from 00Z on 3/2/2019 and the model field used as input was the operational ECMWF analysis valid at this time. Because cloud-affected observations are being compared with clear-sky simulations, large negative departures will be produced, so these cloud-affected areas should be ignored. The following sections will describe various features which are present in the data, with reference to these figures and those in Figures 25 and 26.

4.3.1 Horizontal banding

Initially, from Figures 25 and 26 we can see that there are extensive horizontal bands of very large brightness temperatures in the water vapour band at certain times (02Z, 04Z, 20Z and 22Z). At 04Z and 06Z there are similar features in the temperature sounding channel plotted. These bands appear to be due to all elements within the detector array being adversely affected for a number of dwells. Within the water vapour band, brightness temperatures of about 400K contribute to this banding. A quality control indicator provided in the HDF files appears to correlate with these contaminated bands (Fabien Carminati, private communication).

In the examples from 3/2/2019, such banding appears strongly throughout the mid-wave band, i.e. in Figures 33 to 36 where one complete band of FORs is adversely affected, as well as part of the row to the North. The large brightness temperature values strongly affect the mean values in the spectrum in a region from around $1670 cm^{-1}$ to $1770 cm^{-1}$. Within that range of wavenumbers, the observations are strongly affected by noise in general, see Figure 34, where any atmospheric signal in the observations is very difficult to discern. This band corresponds to a peak in the NEΔT, as can be seen in Figure 37.

Interestingly, the long-wave band also shows features which correspond to the horizontal banding in the mid-wave band. This is much less striking than in the long-wave band, and seems to correspond to a systematic difference in bias characteristics compared to neighbouring FORs, see Figures 27 to 30. In fact, in these affected regions, the observations in the long-wave band often match the simulations better, so it would be worthwhile to screen these data if possible to avoid having systematically-contaminated observations affecting the optimal behaviour of a bias correction scheme.

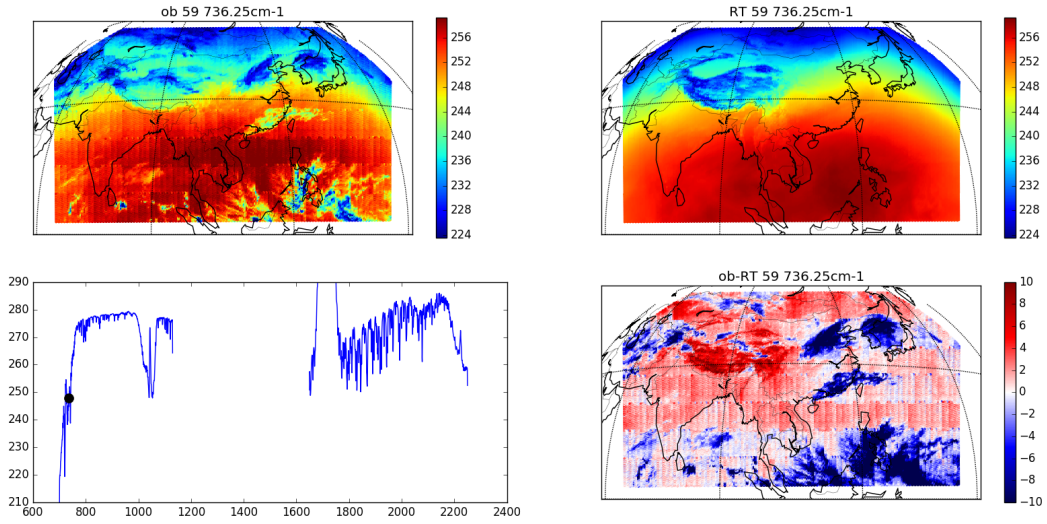

Figure 27: Comparisons of GIIRS observed brightness temperatures with simulations produced by RTTOV. Observations are from 00Z on 3/2/2019. Shown here are the observed, apodised, brightness temperatures from the 736.25cm^{-1} channel (top left), the equivalent clear-sky simulations (top right), the difference between the observations and simulations (bottom right) and the mean spectrum for this set of observations (bottom left), indicating with a black circle the location of the channel whose observations are being compared in the other panels.

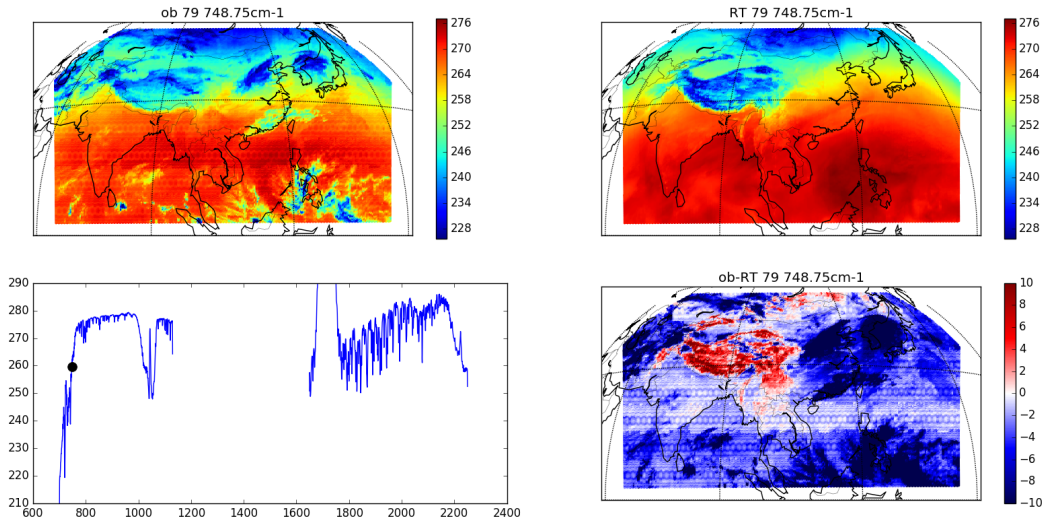

Figure 28: As Figure 27 but for the 748.75cm^{-1} channel.

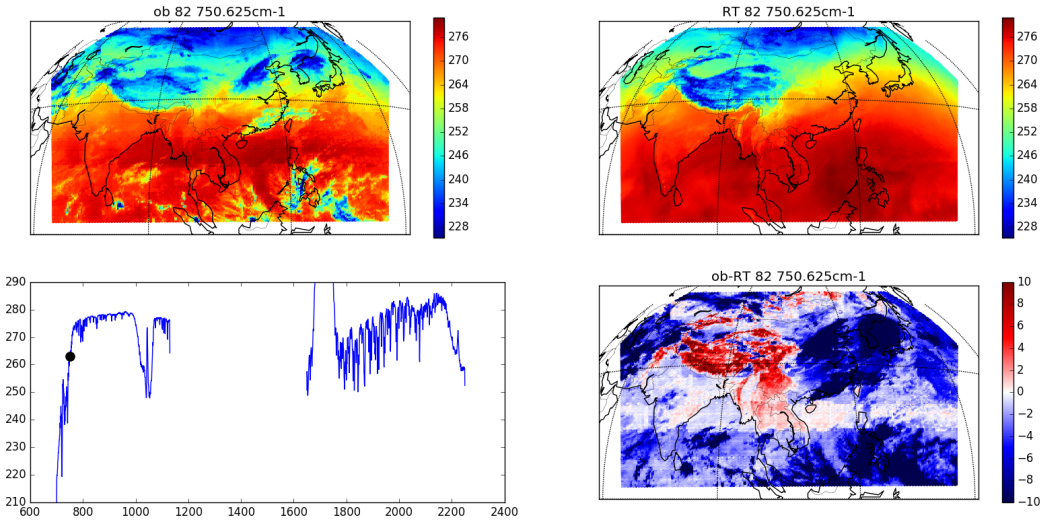

Figure 29: As Figure 27 but for the 750.625cm^{-1} channel.

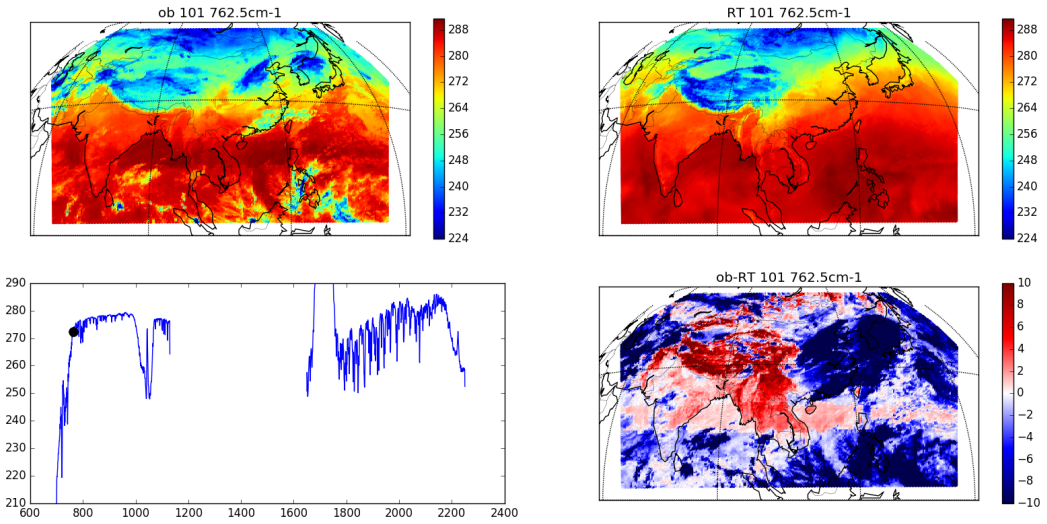

Figure 30: As Figure 27 but for the 762.5cm^{-1} channel.

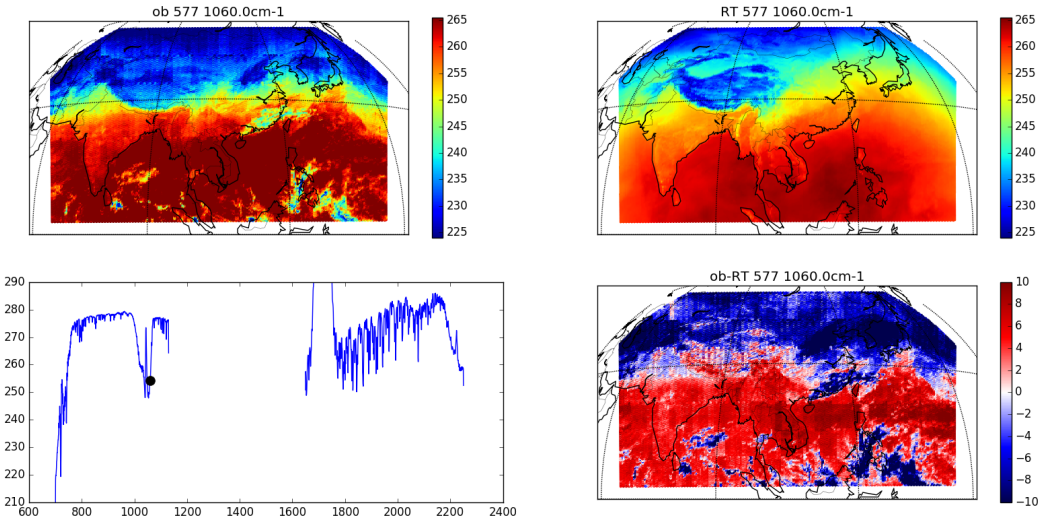

Figure 31: As Figure 27 but for the 1060cm^{-1} channel.

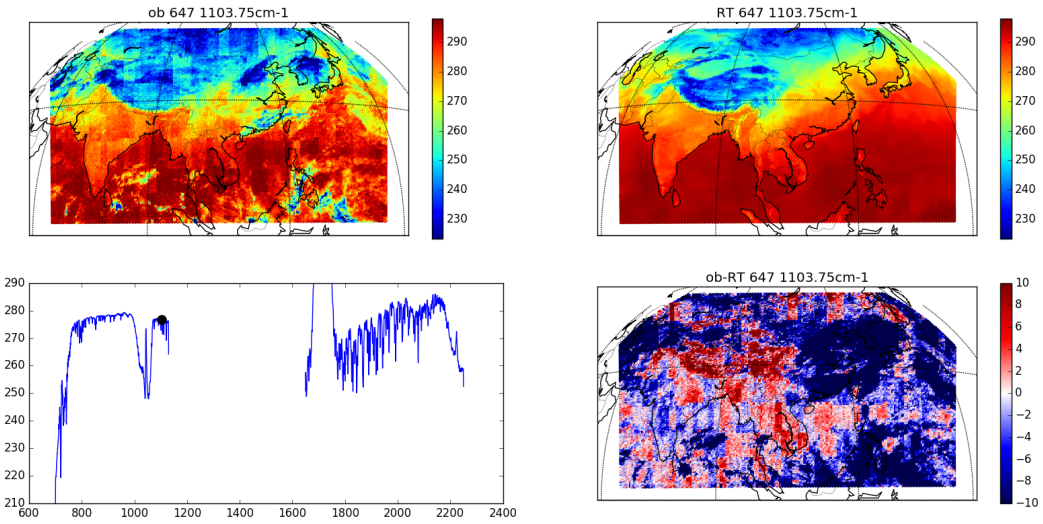

Figure 32: As Figure 27 but for the 1103.75cm^{-1} channel.

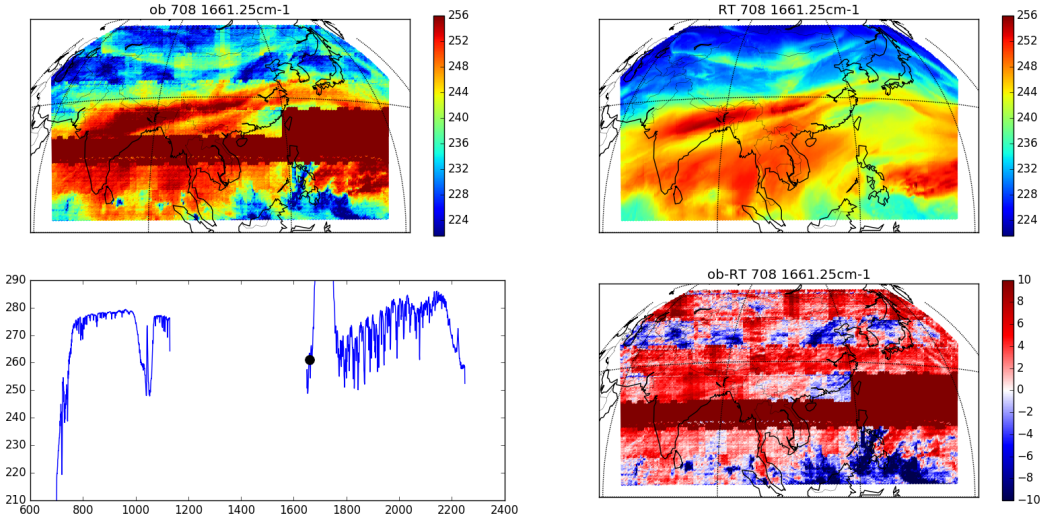

Figure 33: As Figure 27 but for the 1661.25cm^{-1} channel.

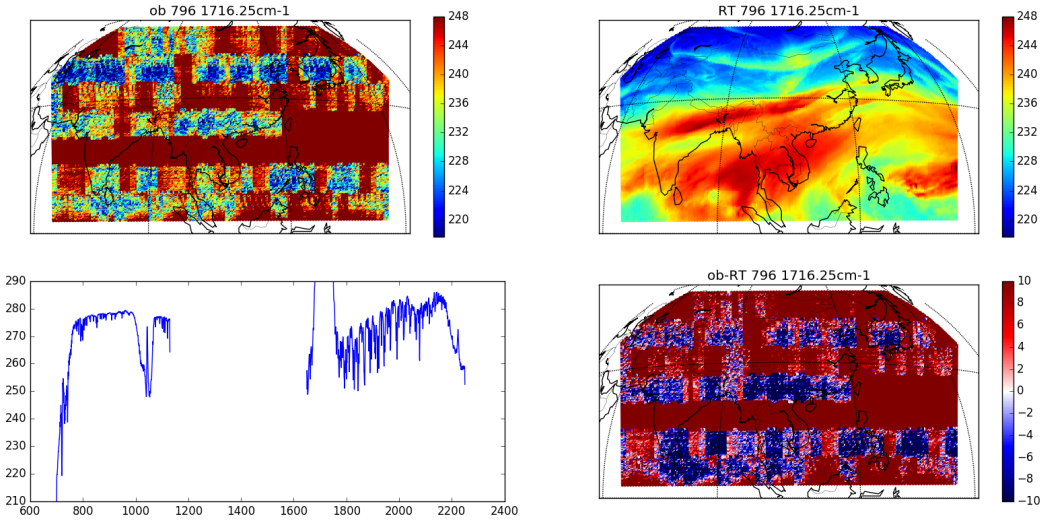

Figure 34: As Figure 27 but for the 1716.25cm^{-1} channel.

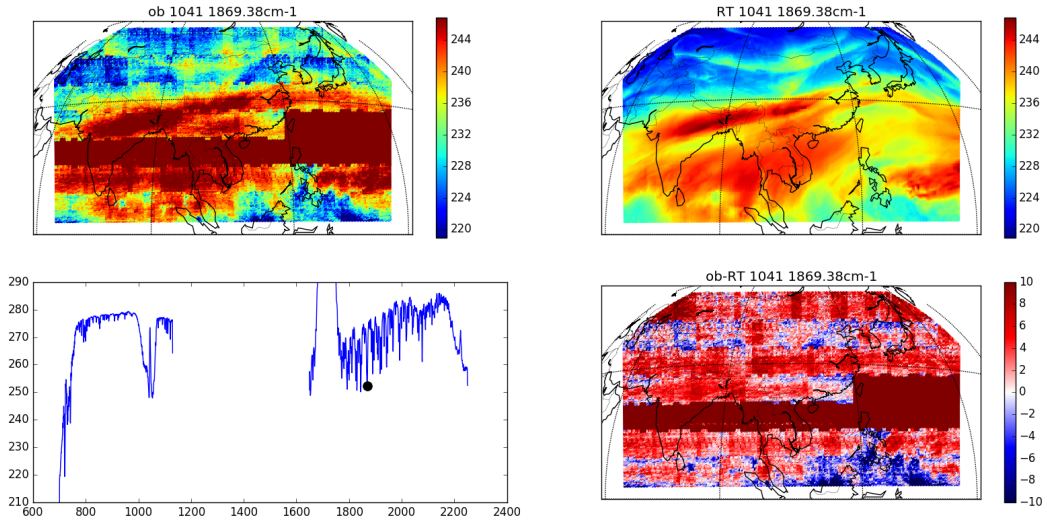

Figure 35: As Figure 27 but for the 1869.375cm^{-1} channel.

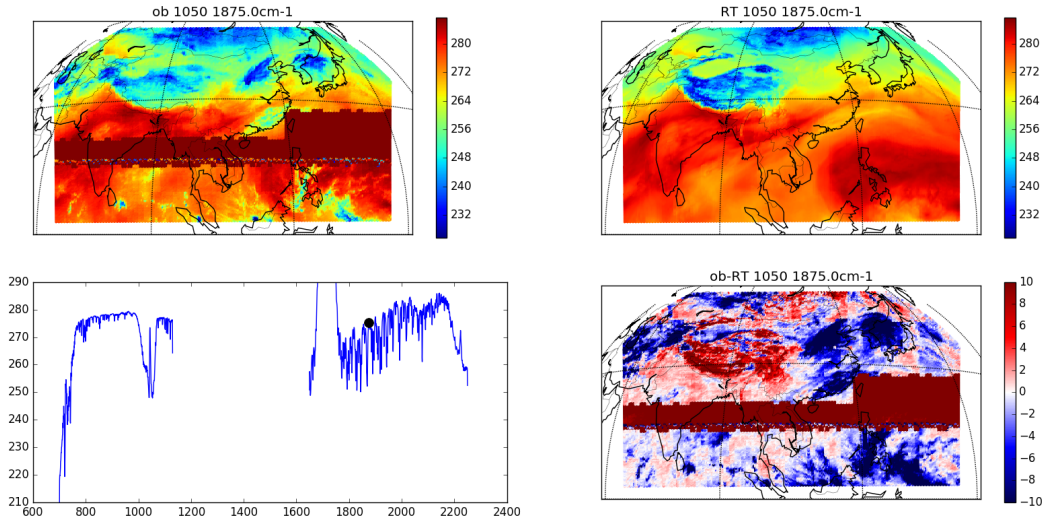

Figure 36: As Figure 27 but for the 1875cm^{-1} channel.

4.3.2 *Inter-dwell variation*

For certain channels, the observations exhibit discontinuities between a given FOR and its neighbours. Although apparent in the observations themselves, these signals are more prominent in the O-Bs. This variation can be seen strongly in the mid-wave band as a "patchwork" appearance in the observations and the observation-simulation differences (Figures 33 to 35). The inter-dwell variability is at its worst when the NEΔT is largest (c.f. Figure 37). Also, in the long-wave band, for channels where the NEΔT is large, there are apparent discontinuities between neighbouring dwells, as can be seen in Figure 32.

In the water vapour band, the magnitude of the inter-dwell variation is strongly dependent on whether or not the channel is aligned with an absorption line. For example, Figure 35 shows a channel directly on an absorption line and, ignoring the aforementioned band of unphysical data, the inter-dwell variability is clearly visible. A neighbouring channel that is off the line, however, does not have a significant inter-dwell variability; see Figure 36.

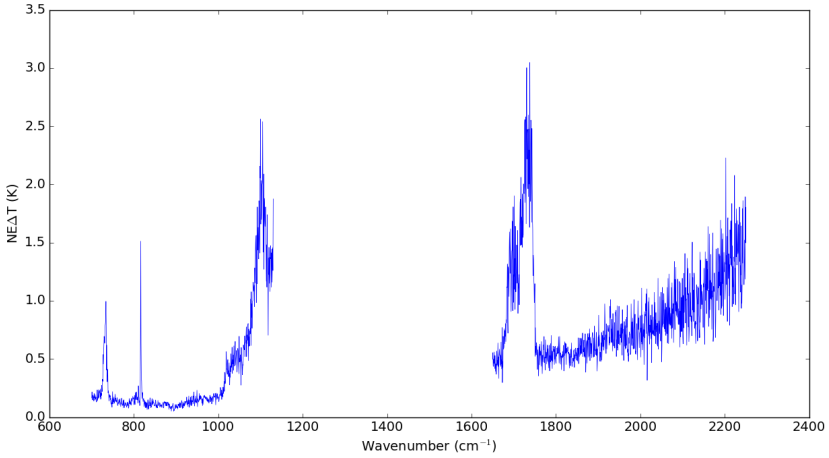

Figure 37: NEΔT for GIIRS, computed from the NEΔR provided with the observations for a reference temperature of 280K.

4.3.3 Intra-dwell variation

Within some dwells, there is the indication that systematic variations exist. In the mid-wave band this can be seen, for example in Figure 35 where the departures show that compared to the simulations, the observations are systematically warmer close to the centres of the dwells, and colder to the north and south. In the long-wave band, there appears to be an east-west variation within the dwells, as can be seen in Figure 27 where the observations are systematically warmer on the westernmost side of the dwells. This analysis is not exhaustive, but it suggests that care is required either to screen out observations that are affected adversely, or to provide a geometric bias correction, although such an approach is not desirable. If these characteristics are due to deficiencies in the pre-processing algorithms, these could potentially be improved.

4.3.4 Geometric features

In Figure 28, a systematic geometric pattern is present, which takes the form of circles in the centres of the dwells. This can be seen in both the observations and the departures, particularly in uniform areas. As well as the circles there also seems to be a fairly periodic pattern, repeating on a sub-FOR scale, particularly in the north-south direction. This is thought to be a spectral calibration issue (Hank Revercomb, private communication).

4.3.5 *Other comments*

As mentioned, some of these artefacts may be suppressed with developments in the pre-processing algorithms, but despite these current deficiencies, there are indications that there is useful information being provided by the observations. In particular, consider Figure 36 which shows a channel in the mid-wave band but that does not lie on an absorption line. We can ignore the large band of unphysical data and begin to identify observations which are contaminated by cloud. For example, there is convective cloud just to the east of Sri Lanka and more uniform cloud on the eastern parts of China and Korea, as well as in the vicinity of the Philippines. Another region to exclude is the high orography of the Tibetan plateau, where there is likely to be surface sensitivity from this channel. With these exclusions, the departures are much smaller, and take both positive and negative values and don't exhibit strong systematic features. In fact, the departures look similar in character to those of CSR observations, which gives confidence that GIIRS can provide useful information to an assimilation system.

4.4 Self co-locations

In assessing the inter- and intra-dwell variability, we can exploit the fact that the GIIRS FORs overlap slightly in the north-south direction. Hence, we get co-locations from these adjacent dwells, separated temporally by several minutes. These can be seen in Figure 38 where the eight southernmost pixels of one FOR are close matches to the northernmost eight pixels of the FOR to the south.

We can, therefore examine differences between the pixels where the FORs overlap. If the hypothesis is that there is a monotonic, systematic, north-south, intra-dwell bias pattern, then we would expect the differences at the co-location points to have a consistent sign everywhere.

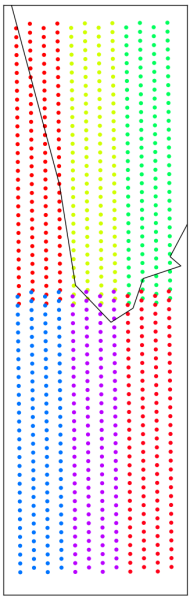

Figure 38: Geographic locations of GIIRS observations from six FORs, each coloured differently (only these six FORs are plotted here). The coastline visible is the southern tip of India, for scale.

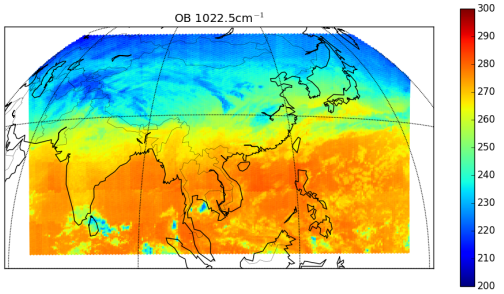

(a) Observations ($1022.5cm^{-1}$)

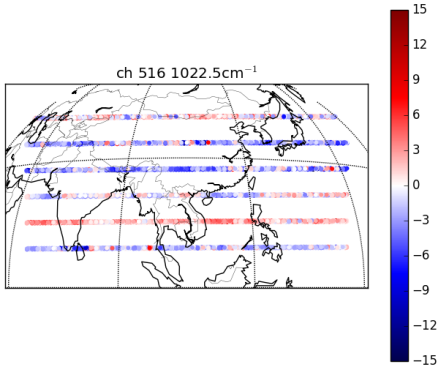

(b) Co-located differences ($1022.5cm^{-1}$)

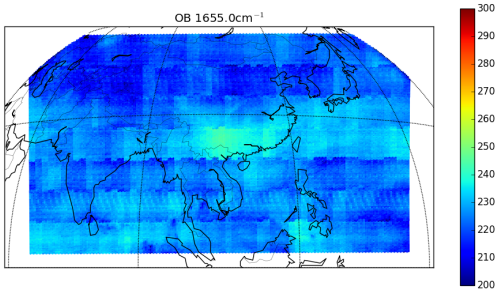

(c) Observations ($1655cm^{-1}$)

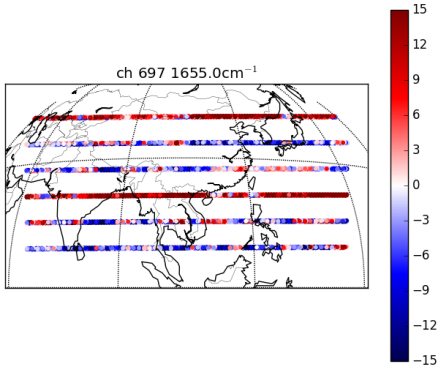

(d) Co-located differences ($1655cm^{-1}$)

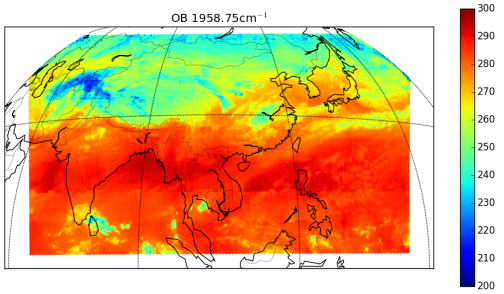

(e) Observations ($1958.75cm^{-1}$)

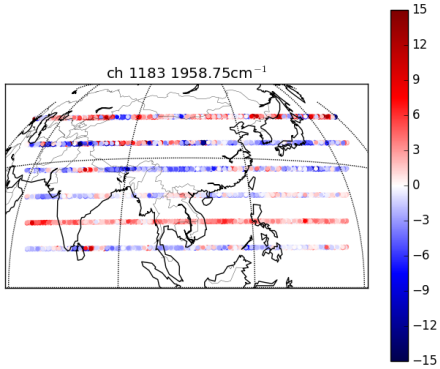

(f) Co-located differences ($1958.75cm^{-1}$)

Figure 39: Left hand panels show GIIRS observations for three channels and the right hand panels show the differences of the co-located observations from neighbouring north-south FORs, where the convention is that the observations from the FOR to the south are subtracted from those to the north in each pair.

Many channels are too noisy to display any systematic variations here, so the channels in Figure 39 are among those that show the strongest of any systematic variation. It seems that within each row of co-locations shown here, there is consistency in the sign of the difference, but this sign varies between the rows in a non-systematic way. This suggests that inter-dwell variation is dominating the signal in the co-locations. This may be a calibration effect that is row-specific. Although monotonic intra-FOR variation does not seem to be a dominant factor in the differences of these co-locations, it was seen in the previous section that the intra-channel variation was generally symmetric in the north-south direction within each FOR, so we would not expect to see any evidence of this by comparing only the extreme northernmost pixels in the FORs with the southernmost pixels in their neighbours to the north.

4.5 Spectral shift

In examining the characteristics of the GIIRS observations, a comparison against simulations in spectral space revealed that compared to a co-located IASI observation in a cloud-free region of the Indian ocean on 3/2/2019, the differences had much more spread for GIIRS than IASI. This can be seen in the lowest panel in Figure 40.

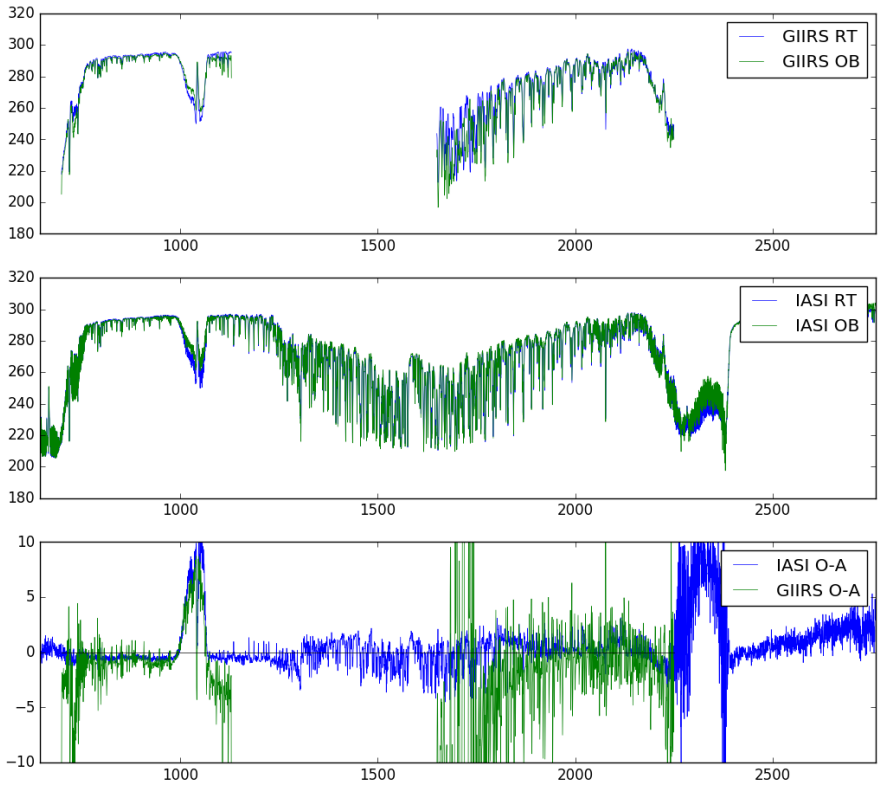

Figure 40: Top panel shows GIIRS observations and simulations for a cloud-free observation over the Indian ocean on 3/2/2019. The middle panel shows the same for a nearby IASI observation. The bottom panel shows the differences between observations and simulations for both GIIRS and IASI.

In the bottom panel, the lines for GIIRS appear to be much noisier than for IASI. On closer examination, it appears that the spikes in the differences are due to a spectral offset between the observations and the simulations. This can be seen in Figure 41 which shows zoomed regions for both the long-wave and mid-wave band.

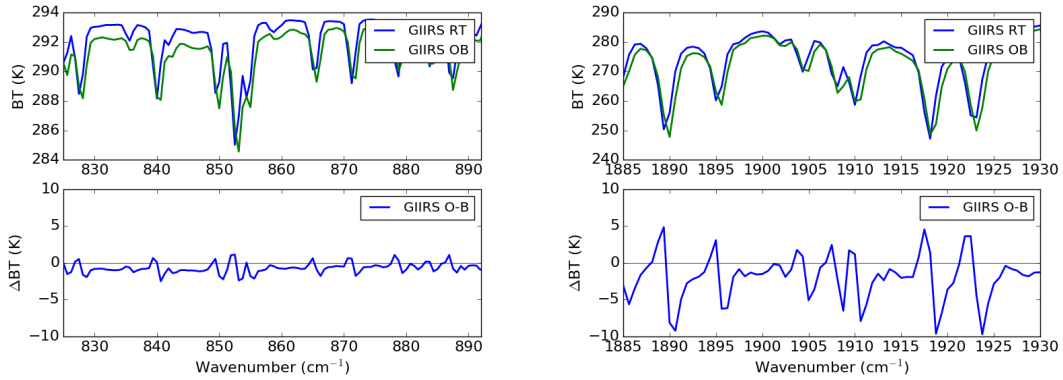

(a) Long-wave band

(b) Mid-wave band

Figure 41: A zoomed spectral region showing the same GIIRS observation as Figure 40, its corresponding simulation, and the difference between the observations and the simulations for the long-wave band (a) and the mid-wave band (b).

The absorption lines are present in both the observations and simulations but it will be seen that the observations are consistently shifted to slightly longer wavenumbers. The simulations were ruled out as the cause of the shift by comparing the positions of the absorption lines with IASI observations and IASI simulations, and these were consistent with the GIIRS simulations (results not shown here).

Initially, soon after launch, IASI on Metop-A was subject to a spectral shift, although much smaller than this one (Marco Matricardi, private communication). The magnitude of the shift for GIIRS is approximately $3.25 cm^{-1}$, so more than half of the channel separation. We can attempt to correct for this shift in a number of ways. Here, we follow a modified version of an approach involving using a multiplicative scaling factor to shift the wavenumber array (Claude Camy-Peyret, private communication), as would be required in the case of an incorrectly-specified laser wavelength.

Effectively, this approach results in a correction of the form:

$$\nu_{corr,i} = (1+\varepsilon)\nu_{raw,i}$$

This would have the effect of scaling the wavenumbers themselves, but we are restricted in that the RTTOV coefficients are specified for a particular set of channel-centre wavenumbers. For this reason, we must keep the wavenumbers the same and modify the observations in a way that is consistent with such a shift. In practice we do this here by interpolating the raw observations onto a very fine wavenumber grid, and then performing the required shift to the interpolated abscissa values. Then, the shifted observations can be interpolated back onto the original grid of wavenumbers. The scaling factor, ε, is diagnosed here empirically by finding the value for each spectral band that creates the best match to the simulated brightness temperatures. This is done by minimising the standard deviation of the differences (just for one spectrum initially, which is not ideal). The diagnosed values of ε are shown in Figure 42 where the minima indicate the optimum values of ε for the two bands.

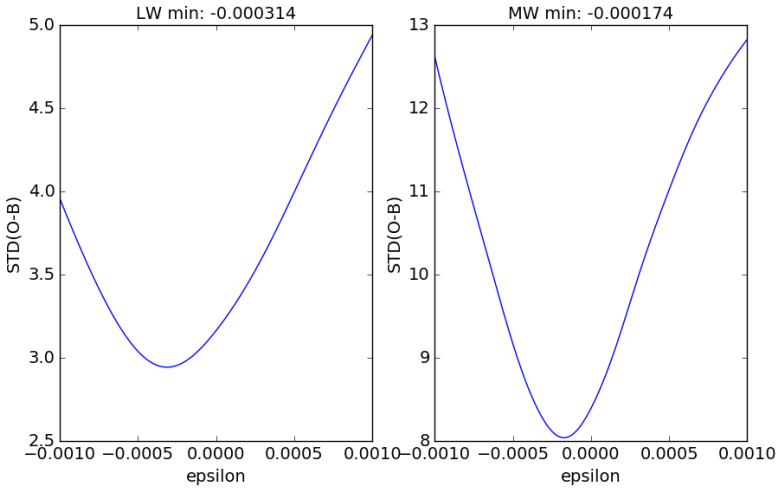

Figure 42: The standard deviation of the difference between the GIIRS observations and simulations for a range of wavenumber scaling values, ε. The optimum value of ε for each spectral band is the minimum of the curve and is shown in the titles of the two panels.

When these values of ε are applied to the observed spectra, the resulting comparison with the simulated brightness temperatures are improved as shown in Figure 43. Here, it can be seen that the absorption lines are better aligned when comparing the observations and the simulations, and in the differences, the sharp features are significantly reduced. It is worth noting that the requirement for performing this correction is not due to deficiencies in the way the apodisation is applied; the apodisation function is symmetrical (see Section 4.2).

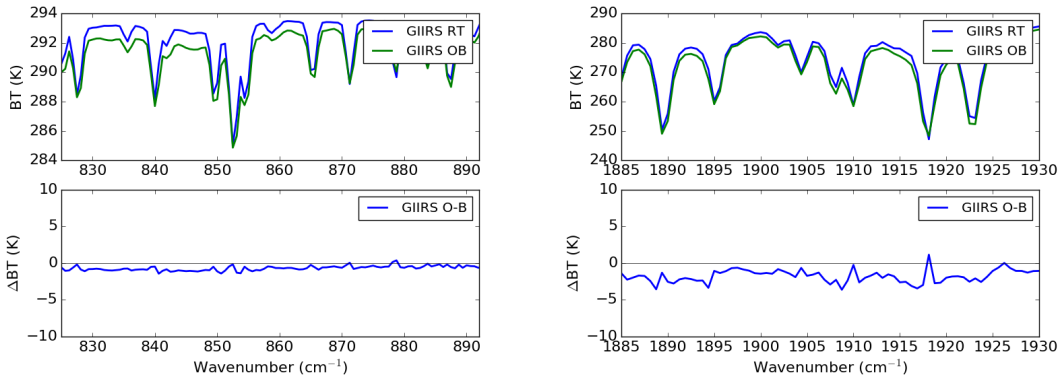

(a) Long-wave band

(b) Mid-wave band

Figure 43: As Figure 41, but after the spectral correction has been applied.

Looking at the entire spectrum as a whole, the impact of the shift is quite noticeable, and in the comparison of the GIIRS departures with those of IASI, the inter-channel variability is of a much more similar magnitude for large parts of the spectrum.

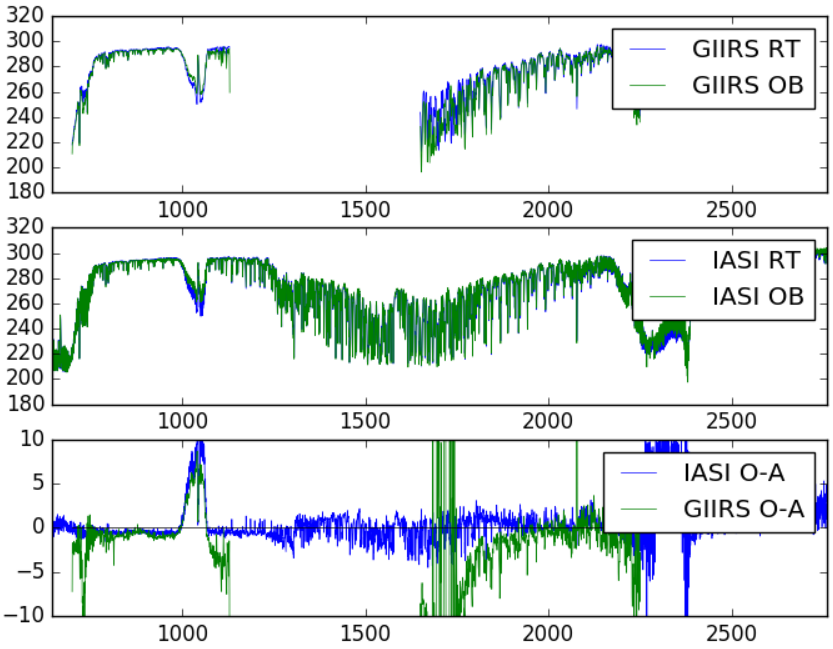

Figure 44: As Figure 40 but with the GIIRS observations having been spectrally shifted by the amounts diagnosed in Figure 42.

Here, only one detector has been investigated, but it is likely that the detectors may have different shifts, and these may vary in time. This will need to be assessed before the observations can be considered for assimilation experiments.

5 Future work

It is a priority to investigate further the GIIRS data with a view to running assimilation experiments. This will involve several steps. Initially, a channel selection will be required in order to withhold channels which are inherently noisy and have other systematic problems. Then, the data will need to be thinned spatially and/or temporally in order to avoid correlations between the observation errors across these dimensions. Observation errors will need to be prescribed and these should account for inter-channel error correlations by, for example, using the Desroziers method (Desroziers et al. [2005]). Quality control will also be an important consideration, in order to reject observations that are subject to systematic deficiencies. Cloud screening will also be a requirement, as well as providing a set of channel-dependent predictors which will be used to determine the coefficients in the variational bias correction scheme.

The preliminary considerations given to the varying of observation errors throughout the assimilation window will continue to be explored. Making use of geostationary observations is likely to prove particularly useful in assessing the impact of varying the observation errors as the temporal frequency of such observations is relatively high.

Also, the impact of observations with high temporal frequency will be investigated specifically for GOES-16 which is now providing full-disk scans every 10 minutes, which is unprecedented, and GOES-

17 is likely to provide similarly-sampled observations. Previous work has shown that as the temporal frequency of assimilated geostationary radiance is increased to hourly, the impact was not seen to saturate (Peubey and McNally [2009]), so it will be possible to extend this study to significantly higher temporal frequencies.

Acknowledgements

Many thanks to the following for their assistance and input:

Tony McNally, Peter Lean, Ioannis Mallas, Gabor Radnoti, John Hodkinson, Mohamed Dahoui, Elias Holm, Reima Eresmaa, Pete Weston, Marco Matricardi, Cristina Lupu, Niels Bormann, Tom King, Peter Keehn, Qiang Zhao, Sharon Nebuda, Haixia Liu, Andy Heidinger, Tim Schmit, Jim Jung, Ruth Taylor, Neill Bowler, Simon Elliott, Fabien Carminati, Roger Saunders, Xianjun Xiao, Qifeng Lu, Wei Han, Claude Camy-Peyret, Dave Tobin, Hank Revercomb, Fiona Smith and the EUMETSAT helpdesk.

References

N. Bormann, M. Bonavita, R. Dragani, R. Eresmaa, M. Matricardi, and A. McNally. Enhancing the impact of IASI observations through an updated observation error covariance matrix. *ECMWF Technical Memorandum 756*, July 2015.

C. P. Burrows. Assmilation of radiances from geostationary satellites: first year report. *EUMETSAT/ECMWF Fellowship Programme Research Report No. 47*, 2018.

C. Cardinali. Forecast sensitivity to observation (FSO) as a diagnostic tool. *ECMWF Technical Memoranda*, (599):26, 10 2009.

G. Desroziers, L. Berre, B. Chapnik, and P. Poli. Diagnosis of observation, background and analysis-error statistics in observation space. *Quarterly Journal of the Royal Meteorological Society*, 131(613): 3385–3396, 2005.

D. Di, J. Li, W. Han, W. Bai, C. Wu, and W. P. Menzel. Enhancing the fast radiative transfer model for FengYun-4 GIIRS by using local training profiles. *Journal of Geophysical Research: Atmospheres*, 123(22):12,583–12,596, 2018.

P. R. Griffith and J. A. de Haseth. *Fourier Transform Infrared Spectroscopy*. John Wiley & Sons inc., 1986.

R. W. Hamming. *Digital Filters*. Prentice Hall inc., 1989.

T. J. Hewison and J. Müller. Ice contamination of Meteosat/SEVIRI implied by intercalibration against Metop/IASI. *IEEE Transactions on Geoscience and Remote Sensing*, 51:1182–1186, 2013.

K. E. Howes, A. M. Fowler, and A. S. Lawless. Accounting for model error in strong-constraint 4D-Var data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 143(704):1227–1240, 2017.

R. H. Langland and N. L. Baker. Estimation of observation impact using the NRL atmospheric variational data assimilation adjoint system. *Tellus A*, 56(3):189–201, 2004.

J. Letertre-Danczak. The use of geostationary radiance observations at ECMWF and aerosol detection for hyper-spectral infrared sounders : 1st and 2nd year reports. *EUMETSAT/ECMWF Fellowship Programme Research Report No. 40*, 2016.

A. P. McNally. On the sensitivity of a 4D-Var analysis system to satellite observations located at different times within the assimilation window. *Quarterly Journal of the Royal Meteorological Society*, 145 (723):2806–2816, 2019.

C. Peubey and A. P. McNally. Characterization of the impact of geostationary clear-sky radiances on wind analyses in a 4D-Var context. *Quarterly Journal of the Royal Meteorological Society*, 135:1863 – 1876, 10 2009.