

EUMETSAT/ECMWF Fellowship Programme Research Report

52

Assimilation of radiance observations from geostationary satellites: third year report

Chris Burrows

March 2020



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

© Copyright 2020

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

Literary and scientific copyrights belong to ECMWF and are reserved in all countries. This publication is not to be reprinted or translated in whole or in part without the written permission of the Director-General. Appropriate non-commercial use will normally be granted under the condition that reference is made to ECMWF.

The information within this publication is given in good faith and considered to be true, but ECMWF accepts no liability for error or omission or for loss or damage arising from its use.

Executive summary

The assimilation of radiance data from Meteosat-8 and Meteosat-11 continues to be of value to the ECMWF system, along with other satellites in the geostationary ring. In particular, forecast sensitivity to observation impact (FSOI) results indicate that these observations are among the most beneficial infrared water vapour observations when normalised to produce the impact per observation, with the two Meteosat satellites giving the largest impact per observation in this category. Meteosat radiance observations with high zenith angles are now being exploited operationally in the ECMWF system, and we would encourage other data providers to disseminate observations with similarly large zenith angles.

In preparation for the assimilation of GOES-17 radiances, an investigation has taken place into the known instrument issue whereby at certain times of day close to the equinoxes the detector overheats. This is seen prominently in the first guess departure statistics, and the times at which the observations are degraded has been fed back to NOAA and may influence quality control decisions.

With the prospect of high temporal frequency data from both FCI and IRS on the MTG series of satellites, it is important to ensure that the assimilation system is prepared for receipt of these data. Although most clear sky radiance (CSR) products are provided hourly, the unofficial GOES-16 CSRs are provided every 10 minutes, each scan representing the full disk. This is a new opportunity, as observations have not been assimilated this frequently in global NWP. In order to extract the high temporal frequency information of the humidity field that will lead to improved wind tracing, the assimilation system has required modification to remove a level of time discretisation which would otherwise have prevented these observations being assimilated optimally. The results presented here indicate that increasing the temporal frequency of assimilated GOES-16 CSRs to one full disk every 20 minutes can improve short-range forecast skill in the ECMWF system. When observations are assimilated every 10 minutes, the short-range forecasts are degraded, but this can be mitigated by inflating the observation errors on account of the neglect of temporal error correlations.

ECMWF is actively preparing for the launch of MTG-IRS, which will be the first geostationary hyperspectral instrument covering Europe, and will be a step change in the remote sensing of this region. A new opportunity for this preparation has been the provision of observation data from the GIIRS instrument onboard the Chinese satellite FY-4A. Although this instrument has a lower specification than IRS, a great deal can be learned from these data in anticipation of IRS. As presented here, initially, the spectral calibration of GIIRS was unsatisfactory, but in collaboration with SSEC and CMA, the required shift was ascertained, using both ECMWF simulations and (at SSEC) co-located IASI and CrIS observations. A significantly improved spectral calibration has been implemented in the latest processing version of the GIIRS ground segment. Deficiencies have been discovered in the unofficial GIIRS RTTOV coefficients which we have used until recently, and it is shown that the official RTTOV coefficients look much better regarding consistency across the zenith angle range. Finally, some initial assimilation experiments have been performed using GIIRS data in the global assimilation system. The assimilation setup for GIIRS is deliberately basic for these initial experiments, and plans to augment various aspects of this have been described. Although the experiments have not run for a long time, encouraging signals are being seen in the verification of short-range forecasts against independent observations.

Contents

1	Overview of geostationary radiance usage and impact	3
2	GOES-17 detector issues	6
3	Assimilation of radiances with high temporal frequency	9
4	GIIRS	14
4.1	Apodisation	14
4.2	Spectral shift	15
4.3	Version 3 processing	18
4.4	Zenith angle dependence	20
4.5	Assimilation setup	22
4.5.1	Observation files	22
4.5.2	Pre-processing	22
4.5.3	Thinning	23
4.5.4	Channel selection	24
4.5.5	Observation errors	26
4.6	Assimilation results	26

1 Overview of geostationary radiance usage and impact

ECMWF currently assimilates “clear sky radiance” (CSR) and “all sky radiance” (ASR) products from several geostationary (GEO) satellites. The set of satellites currently used (February 2020) is shown in Table 1, along with some additional information.

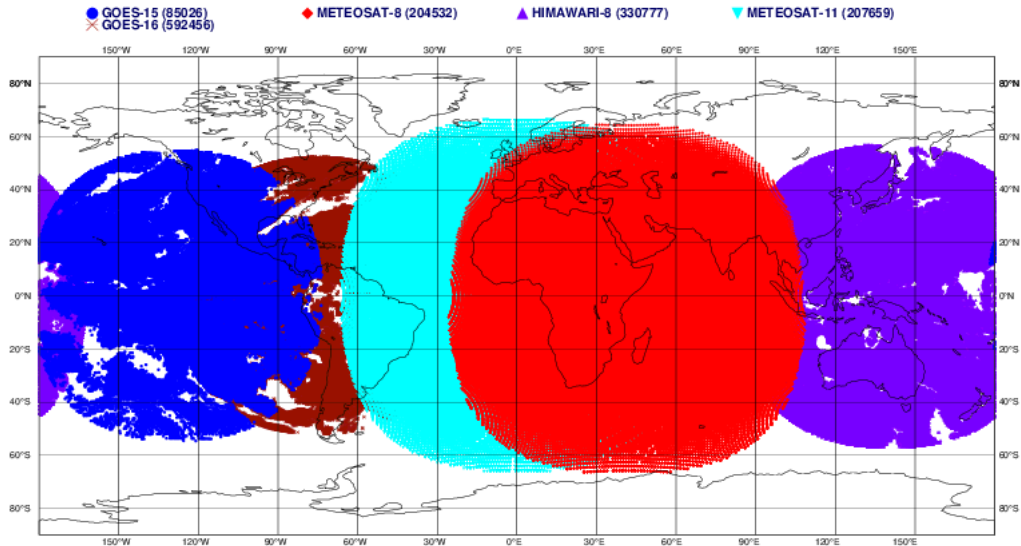
Satellite name	Longitude	Product assimilated	No. channels assimilated
Meteosat-11	0°E	ASR	2
Meteosat-8 (IODC)	41.5°E	ASR	2
Himawari-8	140.7°E	CSR	3
GOES-15	128°W	CSR	1
GOES-16	75.2°W	CSR	3

Table 1: GEO satellites whose radiances are currently being assimilated. Information valid for February 2020.

Spatially, the coverage of the GEO radiances is shown in Figure 1. Following the upgrade to the ECMWF system at cycle 46R1, these radiances are being used with zenith angles up to 74° instead of 60°. This was described previously Burrows [2018] and has particularly resulted in many more observations being used from Meteosat-11 and Meteosat-8 at the edges of their disks. The other satellites only have observations supplied up to smaller zenith angles. The increased usage of Meteosat-11 has led to significantly better coverage of the north Atlantic and northern Europe. An offline stream of GOES-16 data also includes data at these high zenith angles, and these should be used operationally in the future. Furthermore, other data providers should be encouraged to supply observations up to larger zenith angles.

ECMWF data coverage (all observations) - GEOSTATIONARY RADIANCES
26/02/2020 06

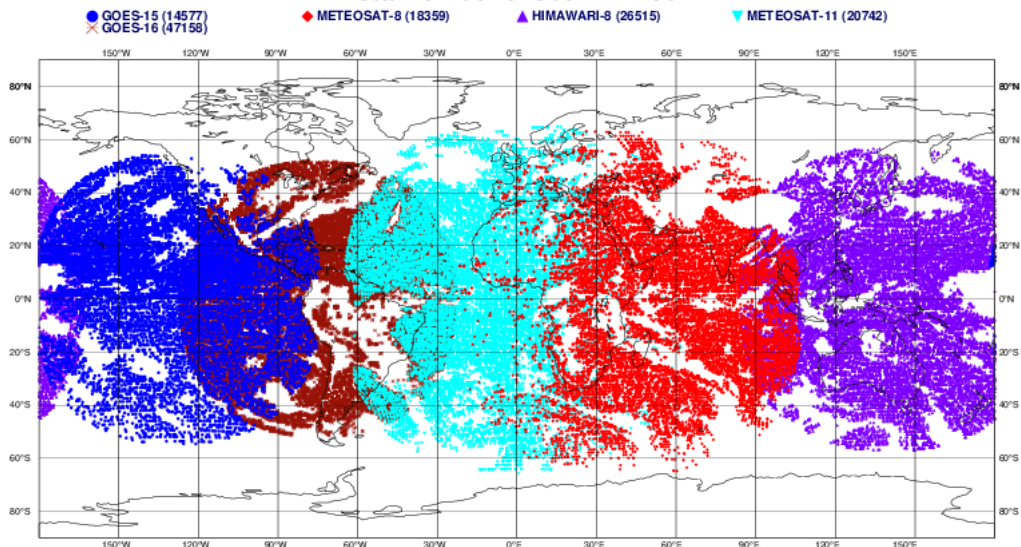
Total number of obs = 1420450



(a) All received

ECMWF data coverage (used observations) - GEOSTATIONARY RADIANCES
26/02/2020 06

Total number of obs = 127351



(b) Assimilated

Figure 1: Coverage of all (a) and assimilated (b) GEO radiance data from 03Z to 09Z on 26/2/20.

A proxy for the impact of observations on 24-hour forecasts can be estimated using the forecast sensitivity to observation impact (FSOI) method [Langland and Baker \[2004\]](#) which uses adjoint techniques to attribute the error reduction in the forecast to each individual observation which was assimilated in the analysis. Several caveats are required here and the FSOI cannot be expected to reproduce the results of observing system experiments (OSEs). Importantly, FSOI will not account for the effect whereby in

an OSE, the information content of the “denied” observation can be compensated by other parts of the observing system. Also, the forecast model used in FSOI is purely linear, which will become a poorer proxy of the nonlinear forecast at 24 hours. Also, the verifying analyses used in the technique are “own” analyses, and so correlations between the forecasts and their verifying analyses will affect any interpretation. Finally, the figure of merit in FSOI here is a dry energy norm which has been integrated to the full vertical extent of the model atmosphere. Such a scalar metric will favour the information content of some observations over others, for example, if a moist norm was used, it may be expected that humidity-sensitive satellite radiances would appear to provide a larger impact. Thanks to Alan Geer, Cristina Lupu and Bruce Ingleby for generating the following statistics from the operational ECMWF system.

The relative FSOI scores for all water-vapour-sensitive infrared radiance observations from the month of January 2020 are shown in Figure 2.


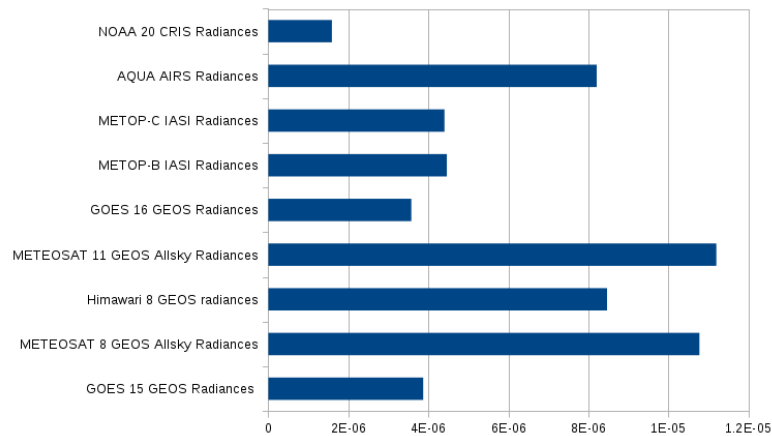
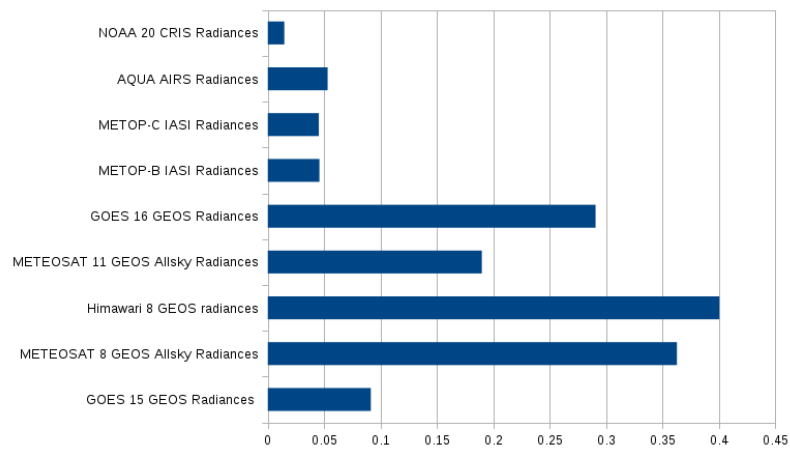


Figure 2: Relative FSOI for water-vapour sensitive satellite radiances in January 2020.

It can be seen that the largest impact from this subset of observations comes from the two IASI instruments (39 water vapour channels from each are assimilated). Broadly speaking, the impact from the various GEO satellites is consistent with the number of channels assimilated from each (c.f. Table 1), although note that during this period, Meteosat-11 underwent a decontamination procedure, so the relatively low impact can be attributed (at least in large part) to this. To mitigate sampling issues such as this, the relative impact can be normalised, either by the number of observations assimilated or the number of channels which are assimilated. The results using both normalisations are shown in Figure 3.



(a) Normalised by observation count



(b) Normalised by number of channels

Figure 3: Relative FSOI observations for water-vapour sensitive satellite radiances in January 2020, normalised by different quantities.

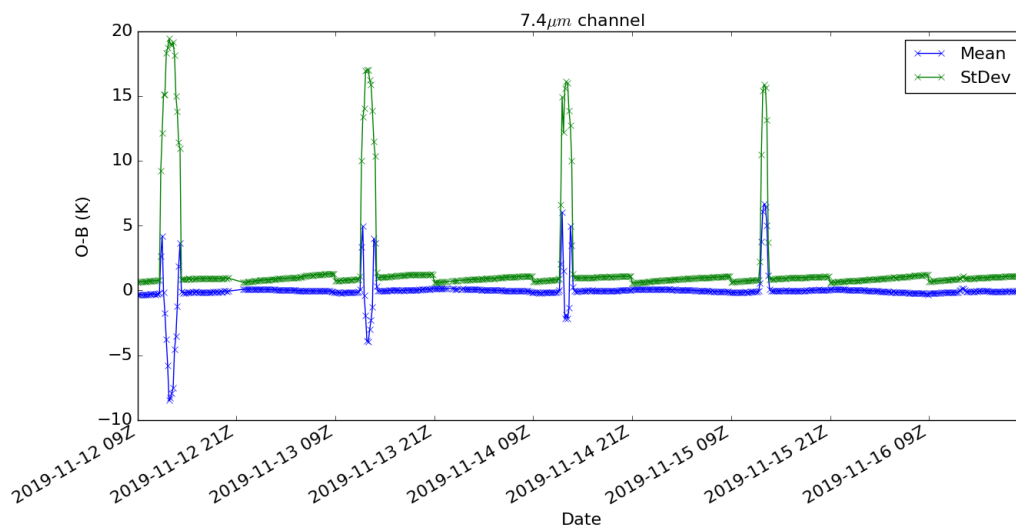
When the FSOI is normalised by the number of observations, the largest impact comes from the two Meteosats and Himawari-8. Normalising by the number of channels assimilated retains the limitation of the overall relative FSOI, in that Meteosat-11 appears down-weighted due to the decontamination outage. Still, the largest impact per channel comes from Himawari-8, Meteosat-8 and GOES-16. This is not necessarily surprising as the information provided by some of the hyperspectral channels may not project so well onto the energy norm.

2 GOES-17 detector issues

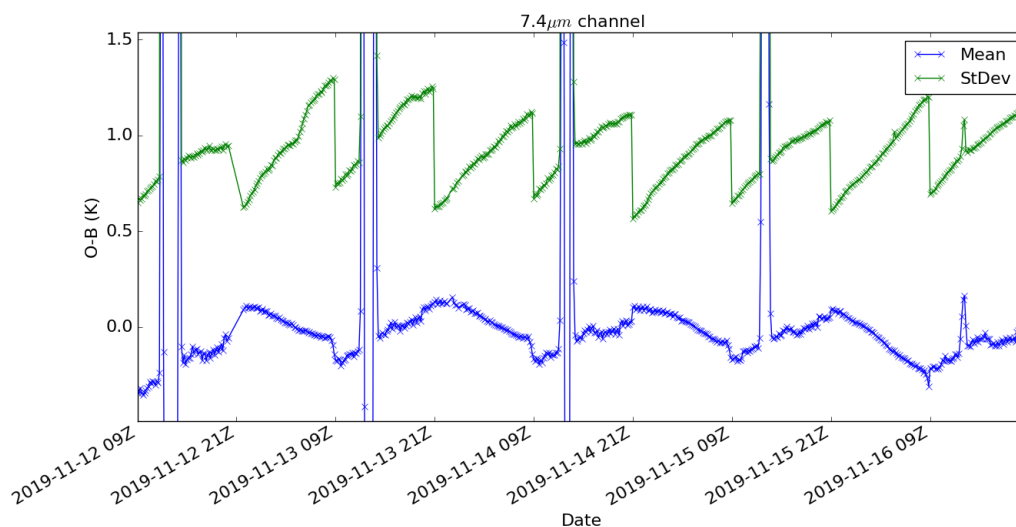
The Advanced Baseline Imager (ABI) onboard GOES-17 is nominally the same as the instrument carried on GOES-16. However, shortly after launch it was discovered that the infrared channels on GOES-17 are subject to a problem whereby the “loop heat pipe” is unable to keep the detector sufficiently cool during periods when the sun is in certain positions which correspond to times after local midnight when the focal plane temperature is at its maximum (Sharon Nebuda, private communication). As a result, it became clear that at certain times of day, radiances from certain channels would be unusable for assimilation.

These periods of degradation are longer close to the equinoxes, and the problem is not thought to be an issue at any time of day at the solstices.

A time series of first guess departure statistics from the $7.3\mu\text{m}$ channel of GOES-17 is shown in Figure 4. It can be seen that for the first few days, spikes occur at close to 13Z in both the mean and standard deviation.



(a)



(b)

Figure 4: Time series of first guess departure statistics for the $7.3\mu\text{m}$ channel of GOES-17. Shown are the full extent of the range (a) and a zoomed region showing the usual diurnal variation and variation associated with the 12-hour assimilation window (b). The sampling is every 10 minutes.

During this short period, both the magnitude and duration of the daily degradation events decreases monotonically. In fact, on the day following this (not shown), there is no discernible spike present in

the time series. The lessening of the severity of these periods is expected as the season passes further from the equinox. Putting the spikes aside for now, the remaining periodic variation in the mean and standard deviation is very similar to what is seen for GOES-16 (Figure 5). The variation in bias is likely to be due to a combination of biases in both the observations and the backgrounds. The sawtooth pattern in the standard deviation is a characteristic of the model error growth and is particularly noticeable for high temporal frequency observations such as radiances from geostationary satellites. The assimilation windows here are 12 hours long, and cover 09Z to 21Z and 21Z to 09Z. Throughout the assimilation window, the nonlinear forecast model is used to provide model fields at times close to those of the observations (fields are actually made available at every 30 minutes throughout the window). At the start of the window the forecast model has run for a shorter time, so the model error is smaller and hence the standard deviation of the first guess departures is relatively small. Conversely, at the end of the window, the model error is larger, which is reflected in larger standard deviations of first guess departures. Furthermore, the discontinuities in the means at the transitions between assimilation windows give an indication of a systematic diurnal bias in the background fields.

The precise times at which the degraded periods begin and end has been fed back to NOAA and if these times can be correlated with the detector temperature (which is not available to end-users in the BUFR files), a threshold can be applied as part of the pre-dissemination quality control for the CSRs. In the present study, only the three water vapour channels have been investigated for this short period, and the spikes are only present for the $7.3\mu\text{m}$ channel.

When performing comparisons with GOES-16 (which is unaffected by the local-midnight solar heating issue), it was noted that similar diurnal variation was present, of order several tenths of a Kelvin, but also small negative spikes were observed in the mean O-B, and these correspond to local midnight. This can be seen in Figure 5.

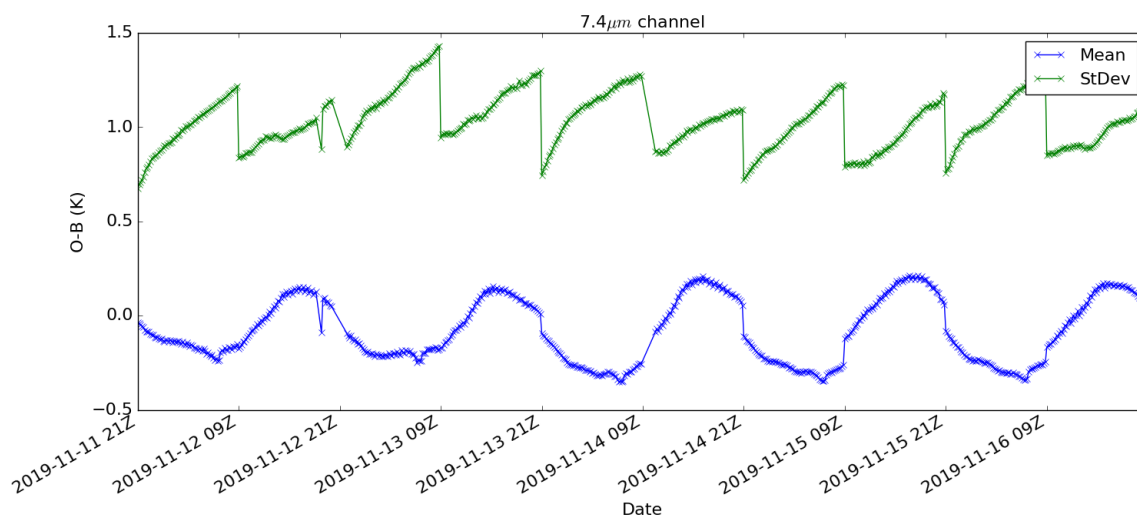


Figure 5: Time series of first guess departure statistics for the $7.3\mu\text{m}$ channel of GOES-16. Note the small negative spikes in the mean approximately corresponding to local midnight (05Z).

There may also be an indication of a similar effect for GOES-17 (Figure 4b), but local midnight is close to 09Z, so this is partly obscured by the jump between the assimilation windows.

3 Assimilation of radiances with high temporal frequency

Geostationary radiances have the unique benefit that the sampling is, in principle, much more frequent than is obtained from the typical coverage of polar orbiting satellites. As technology improves, the frequency of full disk scans increases, and FCI on Meteosat Third Generation will be able to produce full disk scans every 10 minutes. Currently it is typical for CSR/ASR products to be made available hourly, but recently, unofficial GOES-16 radiances have become available every 10 minutes. This provides both opportunities and challenges for data assimilation, and the ability to indirectly extract wind information in a 4D-Var system by measuring humidity features at high temporal resolution is likely to be the source of significant benefits to NWP. Furthermore, IRS will scan Europe every 30 minutes, but it is a hyper-spectral interferometer, so being able to extract temporal information will be even more important for this instrument as the effective vertical resolution is so much higher.

Previous work has shown that increasing the temporal usage of observations from one full disk scan every 12 hours, to one per hour, gave a significant improvement to the wind field Peubey and McNally [2009]. Now, it is possible to perform a similar analysis with sampling down to 10 minutes, for GOES-16.

It is important to explain an aspect of the ECMWF system that needs to be modified in order to perform this. 4D-Var minimises the following cost function J :

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

where \mathbf{x} is the state which is to be optimised (valid at time $T=0$), \mathbf{x}_b is the background state (also valid at $T=0$), \mathbf{B} is the background error covariance matrix, \mathbf{y} is the observation vector (with observations valid at various times throughout the observation window), $H()$ is the observation operator, $M()$ is the nonlinear forecast model and \mathbf{R} is the observation error covariance matrix. In order to evaluate the second term of the equation, the forecast model is run to times close to each observation. Although the model time step is currently 15 minutes in operations, the ECMWF 4D-Var has a second layer of temporal discretisation whereby the assimilation window is split into 30-minute “time slots”. All the observations that lie in a given time slot have their model equivalents calculated from a forecast field valid at the same time. This effectively limits the temporal information that can be extracted from observations, and is equivalent to the situation depicted in Figure 6 (although in practice the time slots are offset by 15 minutes).

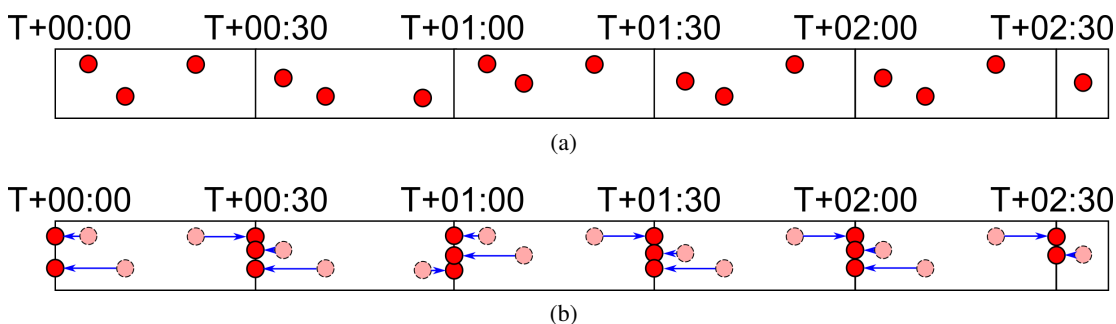


Figure 6: Observations distributed throughout part of the 12-hour assimilation window (a) and the timeslots the observations are assigned to (b).

It can be seen that even though the assimilated observations may be available at high temporal frequency, they will not be able to provide sub-30-minute temporal information to the analysis. Recent work by Elias Holm and Peter Lean has led to ability to relax these constraints, and reduce the length of the time

slots to the same as that of the time step. This has enabled a number of experiments to be run with effective time slots of 10 minutes. Following the example of [Peubey and McNally \[2009\]](#), the GOES-16 observations were sub-sampled in preparation for running three global assimilation experiments (set up with Elias Holm). These samples represented full disk scans every 10, 20 and 30 minutes. The 30-minute sample is equivalent to current operational usage for GOES-16 and so acts as a control against which to compare the other experiments.

In order to assess the impact of this change, we can examine the accuracy of the first guesses (i.e. short-range forecasts) relative to independent observations. This is shown in [Figure 7](#) for a selection of observations. For the geostationary radiances, the values for GOES-16 should be ignored as they are not independent.

These results are for a fairly short period, so may not be robust. Although some signals are mixed, others are clearer. Notably, the ATMS water vapour channels show a clear improvement (i.e. a relative reduction in standard deviation) when the observation frequency is increased to 20 minutes, but a degradation when this is extended to 10 minutes. This is quite consistent with the signal seen for the conventional humidity observations. Interestingly, the 10-minute data shows the best improvement in the fit to radiances from geostationary satellites with disks that overlap that of GOES-16 (i.e. GOES-15 and Meteosat-11). The signal for the winds is quite mixed, but in general it appears that 10-minute data is detrimental.

As stated above, this is the first time that observations with such high temporal information has been assimilated into the ECMWF system, and this indicates that even without significant tuning, there are signs that benefits can be seen in the humidity field. One aspect that has been neglected is the potential role of temporally-correlated observation errors. The temporal covariances of first guess departures for the $6.15\mu m$ channel have been calculated from 24 hours of GOES-16 observations which were assimilated in the 10-minute experiment. These are shown in [Figure 8](#) and include correlations of the observation errors themselves and the errors in the backgrounds which have been forward-modelled to the observation times.

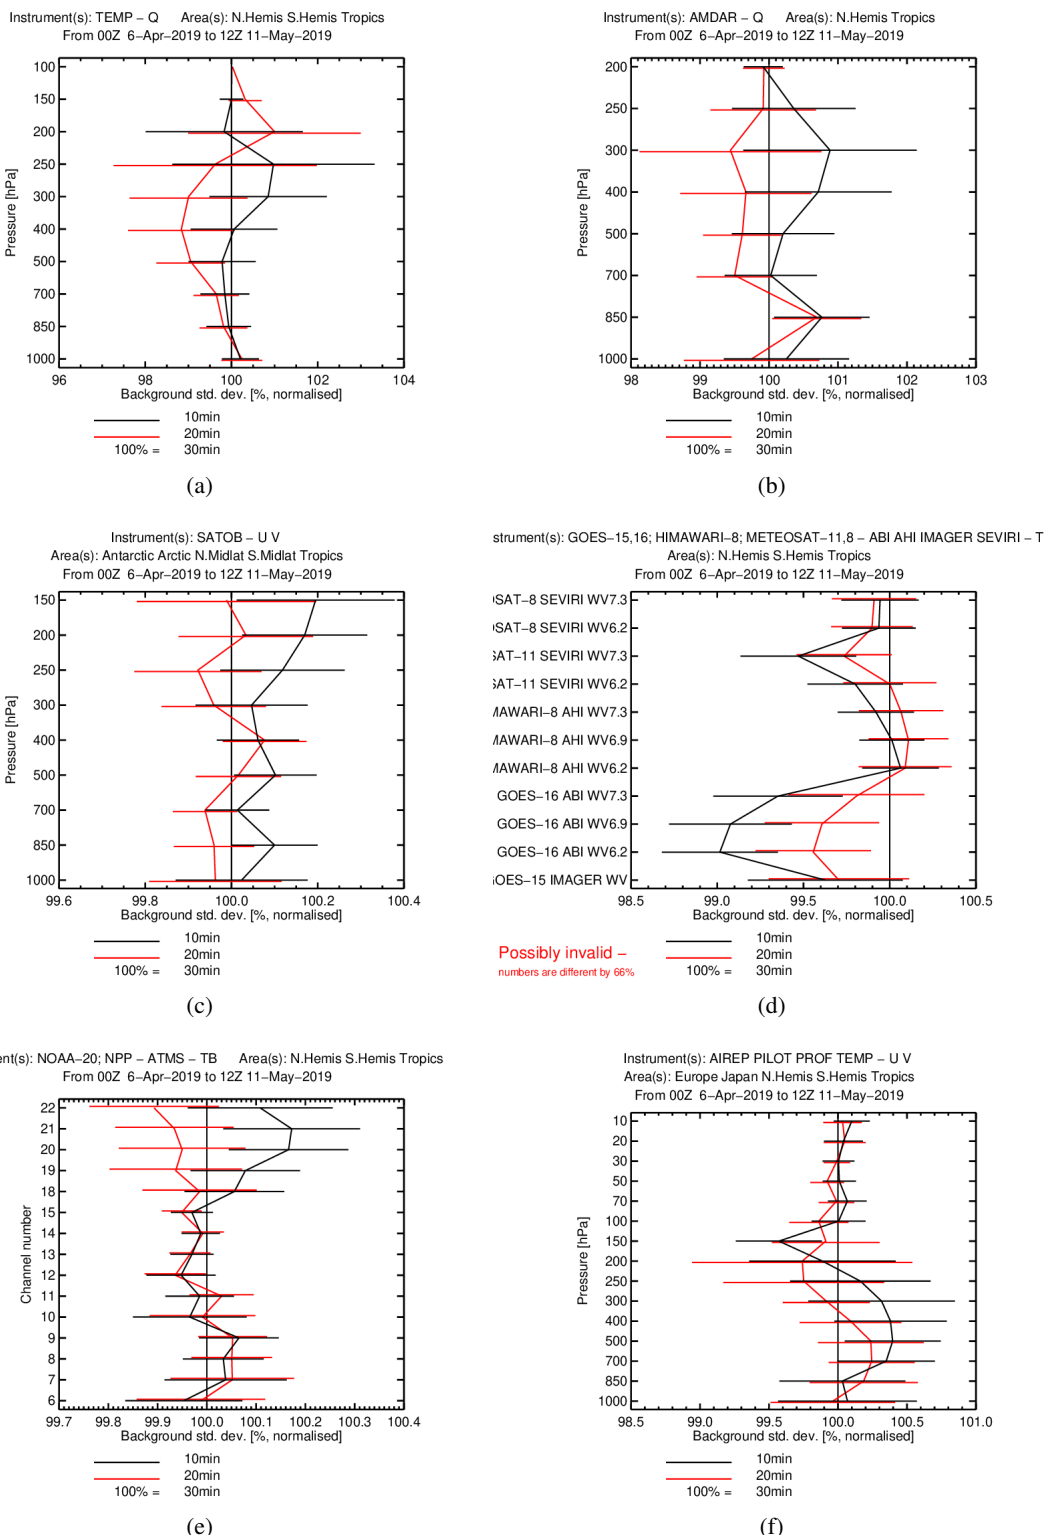


Figure 7: Relative change in global standard deviations of first guess departures for various observation types when varying the temporal sampling of assimilated GOES-16 CSR observations: radiosonde humidity (a), aircraft humidity (b), atmospheric motion vectors (c), geostationary radiances (d), ATMS (e) and *in situ* wind observations (f).




Figure 8: Temporal first guess departure covariances for the $6.15\mu\text{m}$ channel of GOES-16, and the number of match-ups for each ΔT . The $1/e$ value occurs at approximately 3.5 hours.

These temporal correlations are quite long given the size of the assimilation window, and even with hourly observations, the neglect of these correlations in the assimilation is likely to be detrimental. Accounting for these temporally correlated errors fully would be a very difficult task from a technical standpoint, but it is possible to mitigate the neglect of the correlations to a limited extent by inflating the observation errors. Indeed, the current empirically-determined “optimal” observation errors may have already accounted for this implicitly for 30-minute data. To assess this further, an additional set of experiments has been run using 10-minute observations and inflating the observation errors by factors of 1, 1.5, 2 and 4; again, using the 30-minute experiment as the control. Inter-channel correlations are neglected in all experiments for GOES-16 and the baseline errors are 2K. The scaling factors are applied to the standard deviations. The change in the first guess fits to independent observations is shown in Figure 9.

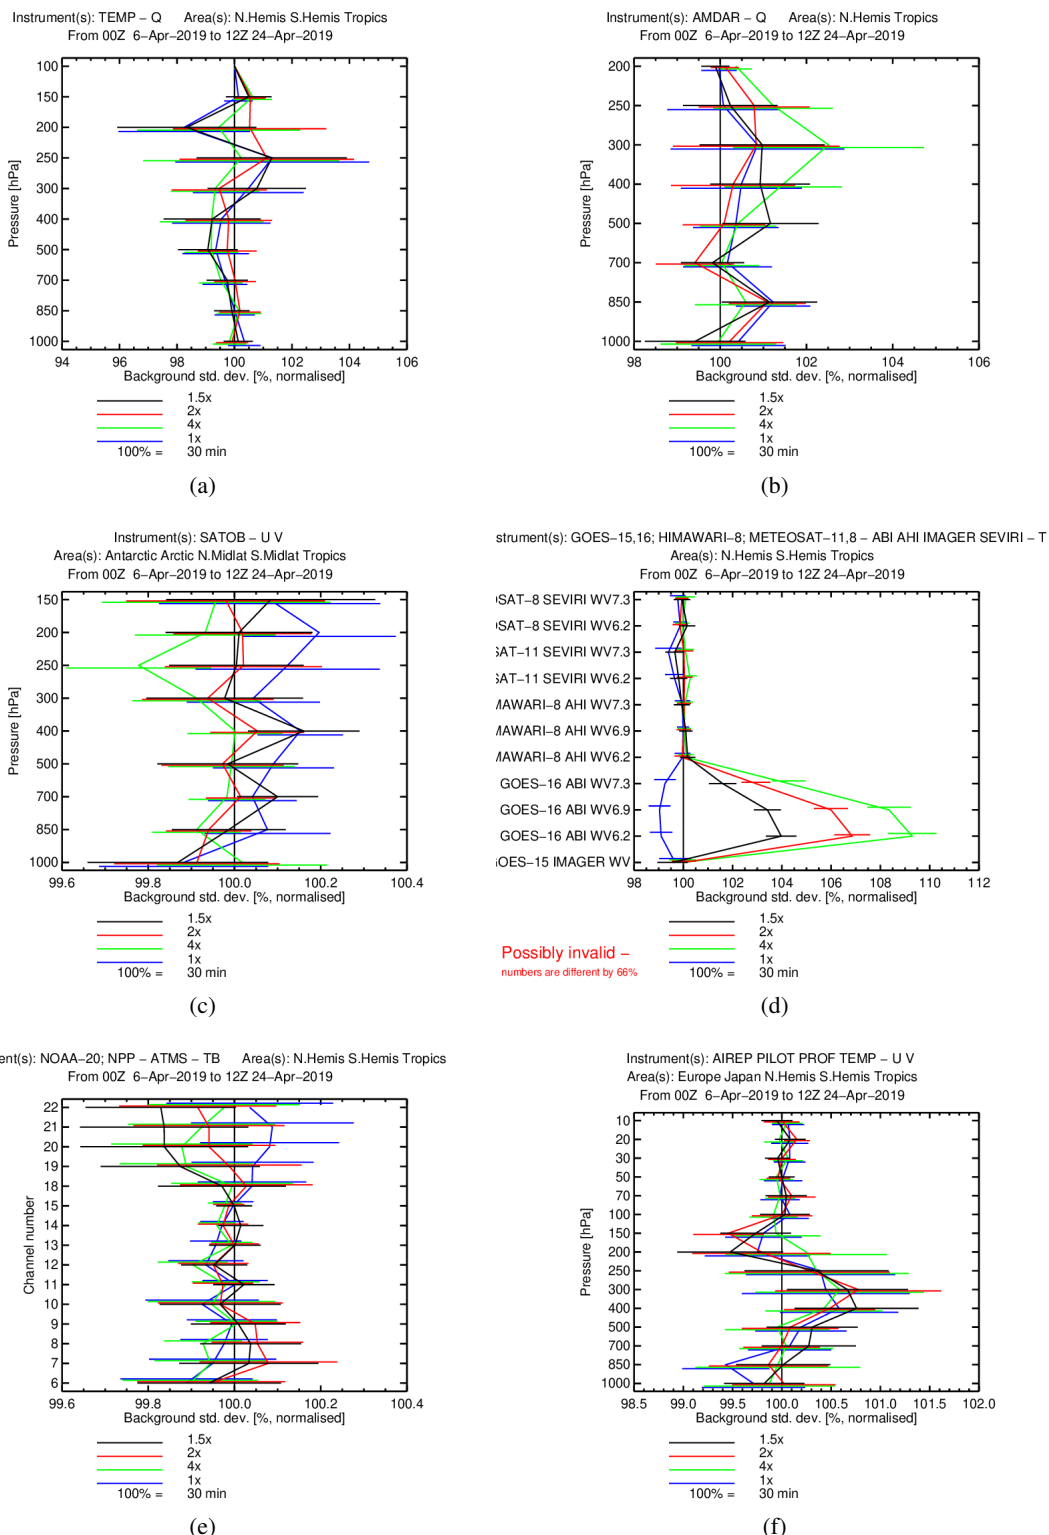


Figure 9: Relative change in global standard deviations of first guess departures for various observation types when inflating the observation errors for GOES-16 CSR observations. The experiments all use 10-minute data, but the control uses 30-minute observations. The independent observations used to perform the verification are: radiosonde humidity (a), aircraft humidity (b), atmospheric motion vectors (c), geostationary radiances (d), ATMS (e) and *in situ* wind observations (f).

These results are for a short period, so some caution is required in interpreting these results, but some signals are emerging already (again, ignore the results for GOES-16 itself). Firstly, the degradation that was previously seen against the ATMS water vapour channels has become an improvement compared to the 30-minute data when the errors are inflated, even with as small a factor as 1.5. Compared to the AMVs, the inflated observation errors seem to be reducing the degradation. For *in situ* measurements, the signals are less clear, and it would be worth extending these experiments to increase statistical significance. It would also be interesting to include inter-channel observation error covariances for GOES-16, as these have already been implemented for the other GEO satellites.

4 GIIRS

EUMETSAT plans to launch the instrument MTG-IRS in 2023, and as part of active preparations to make use of these observations at ECMWF, the Chinese instrument GIIRS onboard the geostationary satellite FY-4A can provide a proxy for the kind of data that IRS will produce. Although its spatial coverage, spectral range, temporal sampling and horizontal resolution are inferior to the specification of IRS, some useful experience can be gained from examining these data in an assimilation context. GIIRS continues to be developed and the specification of the follow-up missions FY-4B and FY-4C should have significantly improved capabilities.

GIIRS scans the domain shown in Figure 10 and note that the scanned region has been shifted to the west from the original domain (Burrows [2019]). This region is scanned once every two hours, and the details of the scanning schedule and focal plane array have been described previously Burrows [2019]. It is worth describing the spectral characteristics of the observations. Data are provided from two spectral bands: the long-wave (LW) band from 700cm^{-1} to 1130cm^{-1} and the mid-wave (MW) band from 1650cm^{-1} to 2250cm^{-1} . The spectral sampling is 0.625cm^{-1} — the same as the CrIS instrument.




Figure 10: Coverage of the GIIRS domain. Shown here are brightness temperatures from the 1993.75cm^{-1} channel. It takes 2 hours to scan this region.

4.1 Apodisation

During 2019, the data available from CMA via EUMETCAST Terrestrial underwent several changes of apodisation. Generally, the observations have been unapodised, but for a brief time in August 2019 the observations were apodised with an adapted Hamming function. The unofficial RTTOV coefficients (and

the recently-released official coefficients) assume that standard Hamming apodisation has been applied, so during the periods when no apodisation has been applied at CMA, this has been applied locally on receipt of the data using a 3-channel weighted average Hamming [1989]. Discussions with CMA suggest that they will apply standard Hamming apodisation if and when they revert to disseminating apodised data (Qiang Guo, private communication).

4.2 Spectral shift

As was previously reported Burrows [2019], a spectral shift of order 0.3cm^{-1} was determined by comparing the observed spectra with model simulations using RTTOV and ECMWF model fields. Work to assess and correct the shift has taken place in close collaboration with Bob Knuteson at SSEC and Qiang Guo at CMA, and their contributions are acknowledged.

Previously, the shift was computed by crudely interpolating the observations linearly onto a high resolution spectral grid, and then shifting the entire array by various steps in order to find the optimal fit between the shifted observations and the simulations Burrows [2019]. To a first order, this worked well but a more sophisticated method was suggested by Bob Knuteson which should eliminate spectral artefacts that may be produced from the simple interpolation method. The method involves performing a matrix multiplication on a vector of radiances for an entire spectral band. The multiplying square matrix is populated in order to apply the spectral shift, and is consistent with a multiplicative scaling of the wavenumber scale. This method is a computationally efficient approach as once the magnitude of the required correction has been determined, the values in the matrix can be stored as they will be the same for every subsequent application. The matrix multiplication method is equivalent to interpolating a shifted, highly oversampled spectrum generated by zero-filling the interferogram, and the equivalence of these methods has been determined (Bob Knuteson, private communication).

The matrix elements are given by:

$$L_{i,j} = \frac{1}{N} \frac{\sin(\pi(v_{sensor,i} - v_{user})/\delta v_{user,j})}{\sin((\pi/N)(v_{sensor,i} - v_{user,j})/\delta v_{sensor})}$$

where v_{sensor} is the wavenumber as measured by the instrument, v_{user} is the wavenumber onto which the spectrum will be interpolated, the δv quantities are the spectral samplings of the respective wavenumber grids and N is an arbitrarily large value of 2^{15} to represent the oversampling of the interferogram.

For GIIRS, two matrices must be computed, one for each band. Note that these transformations must be performed on radiances, not brightness temperatures. The wavenumber grid shift is performed as follows where \mathbf{S} is a vector of radiances,

$$\mathbf{S}_{user} = \mathbf{L}\mathbf{S}_{sensor}$$

and the magnitude of the multiplicative shift is given by the variable ‘‘PPM’’ (parts per million):

$$v_{user} = (1 + 10^{-6}PPM)v_{sensor}$$

In determining the size of the shift, a brute-force minimisation was carried out, with the aim of finding the shift that minimises the RMS difference between the shifted observations and a reference over a

specified wavenumber range for each spectral band. Here, simulations from ECMWF model fields were used, and at SSEC, simultaneous nadir overpasses (SNOs) of CrIS and IASI were used. The choice of RMS difference as the value to optimise is clear from Figure 11. The effect of having a shift present is such as to produce a large spike in the departures corresponding to each absorption line, which abruptly changes sign. The smaller the spectral shift, the smaller the RMS of ΔBT when integrated over a number of spectral lines.




Figure 11: Observations and ECMWF/RTTOV simulations for a small part of the GIIRS MW band (above) and the differences (below).

It is worth commenting that the magnitudes of the differences in Figure 11 render many of the channels unusable for NWP. Although the bias is systematic, bias correction schemes such as VarBC would not work in this instance because for a given channel, the observations and simulations will be sensitive to different parts of the atmospheric column. Therefore, it would be potentially damaging to apply bias corrections here using typical predictors; instead it is preferable to correct the shift directly.

Much of the activity related to the spectral shift determination was focussed on the Version 2 processing from CMA. The PPM values determined for both bands using model simulations and SNOs are shown in Figure 12 for each of the 128 detectors within the field of regard. Here, the PPM value corresponds to the value which produces minimum RMS differences summed over the ranges 740cm^{-1} to 760cm^{-1} and 2160cm^{-1} to 2180cm^{-1} for the LW and MW bands respectively (these ranges were chosen to match those used in the SSEC SNO comparisons). Entire spectra are rejected from the sample if the 900cm^{-1} window channel brightness temperature exceeds the range 290K to 310K, as an attempt to reduce cloud contamination in the sample.




Figure 12: Diagnosed PPM values for Version 2 of the CMA GIIRS processing for each detector in both the MW and LW bands. The values labelled “EC” were computed using simulations based on ECMWF model fields. The “CrIS” and “IASI” values were determined at SSEC using SNOs.

It can be seen that there is good agreement between the different methods, giving confidence that the variation between detectors is real and robust. The diagnosed PPM values for the LW band are approximately 400ppm and for the MW band they are approximately 250ppm. The effect of applying the optimal shift is shown in Figure 13, where, in cloud-free regions the differences between the observed and simulated brightness temperatures are greatly improved as a result of the shift, with the cold bias being largely removed.

After the shift has been corrected, the mean of the differences between observations and simulations is much closer to having zero in the cloud-free regions, and even by comparing the observed brightness temperatures qualitatively with the simulations, it can be seen that the agreement is much better. Interestingly, it can be seen from the position of the channel in the spectrum that after the shift is applied, the channel centre wavenumber is further from the absorption line centre, and hence for this channel, the weighting function will peak lower and be sensitive to warmer parts of the atmosphere, which is clearly an improvement compared to the simulations. It is clear that some systematic effects still persist across the field of regard, but their cause may be something other than the spectral calibration.

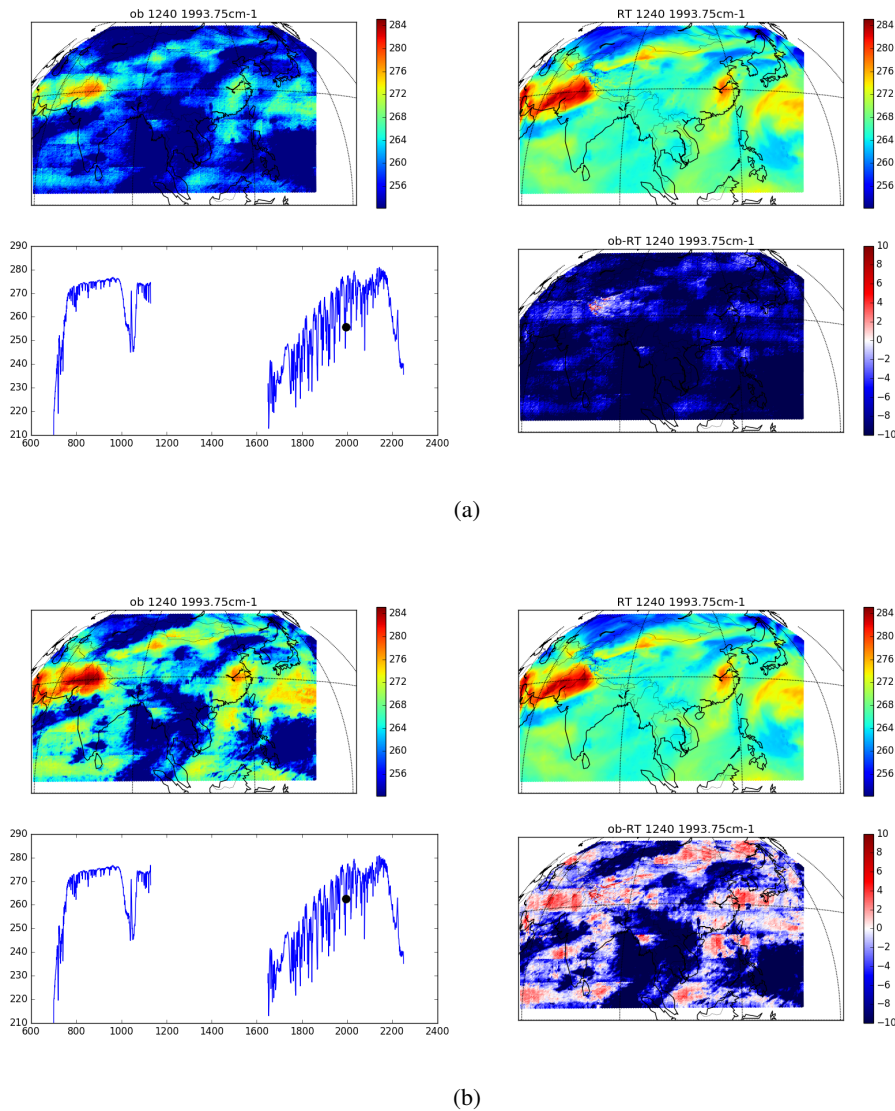


Figure 13: The effect of the spectral shift can be seen spatially, here for the 1993.75cm^{-1} channel. Each panel shows four plots: observed brightness temperatures (top left), simulated clear-sky brightness temperatures (top right), The observed spectrum with the location of the channel marked (bottom left) and the difference between the observed and simulated brightness temperatures (bottom right). The upper figure shows the Version-2 observations (having been apodised) and the lower figure shows the same, but with the diagnosed spectral shift having been applied.

4.3 Version 3 processing

In November 2019, CMA's GIIRS processing algorithms switched to Version 3 (V3), which included corrections for the spectral shift as determined by the comparisons above. The effect of this correction can be seen in Figure 14 which shows the diagnosed PPM values for each detector for V2 and V3.

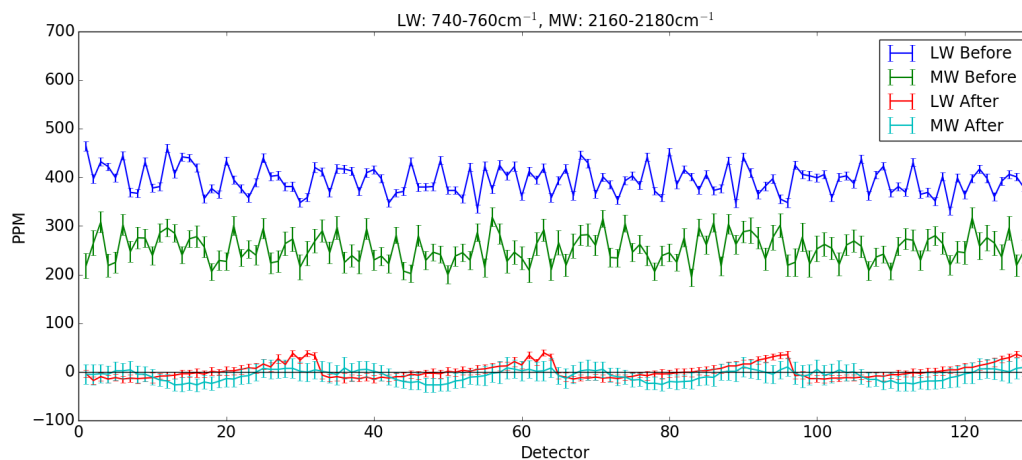


Figure 14: Diagnosed PPM values for Version 2 (“Before”) and Version 3 (“After”) of the CMA GIIRS processing for each detector in both the MW and LW bands, as computed using ECMWF simulations. The spectral ranges used to determine the PPM are shown in the plot title.

There is a significant improvement in the spectral shift as diagnosed by this method. A residual remains, which is acceptably small for NWP use, but further refinements may be made (Hank Revercomb, private communication). V3 does show a clear systematic variation between detectors, with periodicity of 32, i.e. the number of pixels in the north/south direction of the focal plane array. For each group of 32 pixels in the LW band, the variation is almost monotonic, but for the MW band it is more symmetrical. Results from EUMETSAT also show a systematic variation across the field of regard (Dorothee Coppens, private communication), and there is close agreement with results obtained by Wei Han (private communication).

4.4 Zenith angle dependence

During the assessment of the V3 GIIRS observations, a dependence on zenith angle was identified in the first guess departures by Wei Han (private communication). Following this, a similar analysis was performed using ECMWF simulations, and the results are shown in Figure 15 (thanks to Marco Matricardi).

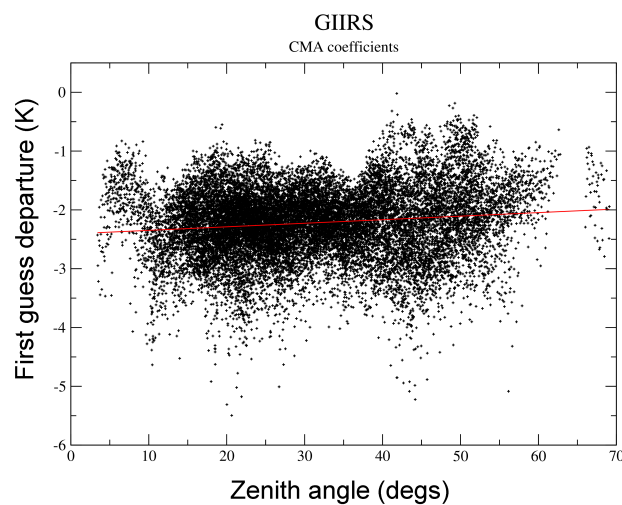


Figure 15: Scatter plot of GIIRS first guess departures from the 703.125cm^{-1} channel as a function of zenith angle. A linear regression is shown in red. To reduce the amount of cloud contamination, spectra are excluded for which the absolute first guess departure of the 900cm^{-1} window channel exceeds 7K . The GIIRS RTTOV coefficients from SSEC/CMA were used to produce the simulated brightness temperatures.

The regression shows a slope, suggesting that there is indeed a zenith-angle dependence on the first guess departures. This prompted further investigation into the origin of the dependence in order to assess any contribution to this bias which may come from the unofficial RTTOV coefficients obtained from SSEC/CMA (Di et al. [2018]). It was recognised that due to the identical channel spacing and assumed spectral response function of CrIS, the official CrIS RTTOV coefficients could be used to provide a good proxy of the official GIIRS coefficients that would be produced. Using the CrIS coefficients, the resulting scatter plot is shown in Figure 16 and it can be seen that although the envelope of the scatter plot shares many characteristics with Figure 15, the linear regression looks significantly more constant, thus it can be concluded that the zenith angle dependence is due to an aspect of the way in which the coefficients were generated.




Figure 16: Scatter plot of GIIRS first guess departures from the 703.125cm^{-1} channel as a function of zenith angle. A linear regression is shown in red. To reduce the amount of cloud contamination, spectra are excluded for which the absolute first guess departure of the 900cm^{-1} window channel exceeds 7K . The CrIS RTTOV coefficients generated by the NWP SAF were used to produce the simulated brightness temperatures.

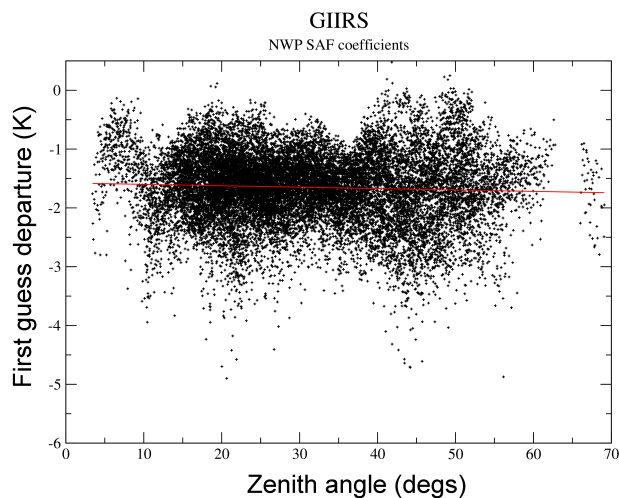


Figure 17: Scatter plot of GIIRS first guess departures from the 703.125cm^{-1} channel as a function of zenith angle. A linear regression is shown in red. To reduce the amount of cloud contamination, spectra are excluded for which the absolute first guess departure of the 900cm^{-1} window channel exceeds 7K . The GIIRS RTTOV coefficients generated by the NWP SAF were used to produce the simulated brightness temperatures.

More recently, the NWP SAF has produced official RTTOV coefficients for GIIRS, and these have been used to perform an additional comparison, as shown in Figure 17. Note the close agreement with the plot using CrIS coefficients.

From this analysis, it is clear that the official NWP SAF coefficients should be used in preference to the CMA/SSEC coefficients, for this channel at least.

4.5 Assimilation setup

For initial assimilation experiments with GIIRS observations, the setup has been kept as simple as possible, thus leaving scope for further refinement. This has allowed some initial progress to be made quickly, and provides a good starting point for ongoing improvements. In this section, the entire assimilation strategy is described, noting any deficiencies in the initial parameters, and plans to move towards a more optimal setup.

4.5.1 Observation files

The GIIRS observations are provided in HDF5 format, from CMA, and these are made available both on ftp, and via EUMETCast Terrestrial (thanks to Simon Elliott). Each file contains spectral radiances for one “dwell”, i.e. 128 pixels. Both LW and MW bands are included in the same file. Inspection of the files reveals that the pixels for the two bands are not at the exact same geographic locations, but are displaced by 20km near the sub-satellite point and 110km at the edge of the scanning domain (Fabien Carminati, private communication). This raises a technical question — should observations from the two bands be treated as two separate observations in order to best reflect their locations? There are two problems with this: firstly, this would mean that information from one band cannot easily be used to influence data usage in the other, which is important for cloud detection. Secondly, further analyses of first guess and analysis departure statistics can be better performed when all channels from a single field of view are combined, such as diagnosing observation inter-channel error correlations (and, indeed, applying these). If the two bands are handled separately, they will need to be re-combined for this, thus requiring co-location matching algorithms, or the addition of metadata to indicate the index of the parent spectrum. Initially, we treat the observations from both bands as being valid at the geographic locations of the LW observations. This choice was made on the basis that the LW band contains surface-sensitive channels, so misalignment near coastlines may be problematic if the locations were misplaced. Having said this, the MW band is sensitive to small scale humidity features where accurate geolocation is perhaps more important than the relatively horizontally-smooth features that the LW band temperature-sounding channels are sensitive to.

In the future, efforts should be made to allow QC such as cloud detection to be performed on entire spectra, but first ensuring that the model equivalents are calculated at the correct locations for the respective bands.

4.5.2 Pre-processing

An observation pre-processing system has been developed for GIIRS which reads in the HDF5 files, performs some initial processing and then writes the data in NetCDF format which can be used in ECMWF assimilation experiments via the “Experimental Observation Framework” (thanks to Peter Lean). The steps of the pre-processing system are described below:

1. The HDF5 files valid for a given assimilation window are identified.
2. Each file is read in.

3. Simple spatial thinning is currently applied at this stage (see Section 4.5.3).
4. Checks are performed to ensure the important arrays have the correct sizes.
5. Observations are checked to ensure they have valid lat/lons, and that the satellite zenith angle is between 0° and 74° . Observations not meeting these conditions are rejected.
6. Apodisation is applied if required (the Hamming function is applied in spectral space).
7. If required, a spectral shift can be applied at this stage (using the method described in Section 4.2), where ppm values can be specified for each detector in each band. This is not used for Version 3 and later of the GIIRS Level 1 processing algorithm.
8. Radiances are converted to brightness temperatures using the Planck function.
9. Arrays are defined for populating the NetCDF files, with congruence ensured (i.e. in some cases, metadata that appears once in the HDF5 file needs to be duplicated to be present for each radiance value).
10. The latitudes/longitudes from the LW band are used for both LW and MW bands.
11. The arrays are written to a NetCDF file (one per input HDF5 file), and the variable names are specifically chosen to map onto the ODB column names, which is an important aspect of integration with the “Experimental Observation Framework”.
12. All the NetCDF files valid for a particular assimilation window are tarred and made available to IFS experiments.

This process can be used for offline research testing, but is not intended for operational use. In due course, BUFR files will be required, and some discussions regarding a GIIRS BUFR sequence have taken place (Nigel Atkinson, private communication).

4.5.3 *Thinning*

Typically, when satellite observations are assimilated at ECMWF, they are thinned using a generic algorithm which retains at most one observation within boxes which have a specified angular size. For initial studies with GIIRS, a more simple thinning strategy has been chosen, which is to retain a subset of pixels from the focal plane array. In order to avoid issues with horizontally-correlated errors (i.e. in the observations themselves or through errors of representation), observations from only 3 out of 128 pixels are retained from the detector array. These were chosen to be the pixels numbered 6, 16 and 26 and thus occur in the second vertical column from the west. This sampling ensures relatively equal spacing in both north-south and east-west directions from the viewpoint of the satellite, but geographically, the observations will have the smallest separations close to nadir.




Figure 18: Observations are retained at these locations, which correspond to pixel numbers 6, 16 and 26 in the focal plane array. These provide uniform coverage from the viewpoint of the satellite.

More sophisticated thinning may be beneficial so that, for example, cloud-affected regions can be identified, but nearby clear scenes could still be used. With the simple approach used here, we are unable to optimally use the data we are presented with. Another aspect is that it may be prudent to exclude some pixels because of known, systematic issues, for example, those at the extreme top and bottom parts of the focal plane, so a retention of such a pre-processing capability may be worthwhile even if a more dynamic thinning strategy is also used subsequently.

On the subject of thinning, it should be stated that the data volumes from GIIRS are extremely large, so performing model simulations as part of the initial trajectory can become very costly. We have 1650 channels for 128 pixels, with 413 fields of regard being measured 11 times a day, which corresponds to nearly 1×10^9 brightness temperature observations per day, which is close to the total number of observations currently ingested into the IFS at ECMWF (Peter Lean, private communication). Looking further ahead, the assimilation of principal component scores or reconstructed radiances will help to allow more information to be assimilated whilst not exceeding constraints of computational cost (Matricardi et al. [2014]).

4.5.4 Channel selection

The initial channel selection used here was determined subjectively. The intention was to largely mirror the channel selection used for CrIS, although the MW band of CrIS covers the long-wave end of the water vapour band, and the MW band of GIIRS covered the short-wave end. Therefore, in the GIIRS MW band, a selection of 23 channels was made, spanning a range of weighting function peak heights, as seen in Figure 19, whilst avoiding those channels that lie very close to absorption lines as they are subject to high levels of noise (see Figure 21).




Figure 19: Temperature and humidity Jacobians for the selected GIIRS water vapour channels. The vertical axis is pressure in hPa.

The long-wave band of GIIRS starts at 700cm^{-1} , and the LW band of CrIS extends down to 650cm^{-1} , but where the bands of the two instruments overlap, the same channel selection was used (note that as mentioned in Section 4.4, the spectral response functions and channel centre frequencies of CrIS and GIIRS are the same for those spectral regions which are measured by both). Due to an error in the implementation, too many channels were selected to be used in the temperature-sounding and surface-sounding regions of the LW band. The total number of channels used in the LW and MW bands are 372 and 23 respectively. The selected channels are shown in Figure 20.




Figure 20: A typical GIIRS brightness temperature spectrum, indicating used channels in red and all others in black.

This process of identifying channels to use is somewhat subjective, so a more objective channel selection

will be considered in the future, such as optimising the degrees of freedom for signal (DFS, [Rodgers \[2000\]](#)). Such approaches are intended to find the channels that maximise the amount of information that the observations can impart to the analysis.

4.5.5 Observation errors

As with the other parameters used in these initial assimilation experiments, the observation errors have been specified to be simple. Here, we select a single fixed observation error which is the same for every channel and correlations between the errors of different channels are neglected. Further work should focus on diagnosing observation errors and their correlations using the Desroziers method [Desroziers et al. \[2005\]](#), but in order to perform this, it is first essential to have assimilated these observations in an experiment. Therefore, the archived analysis and first guess departures from these experiments can be used for this purpose.

In lieu of objectively-determined observation errors, we have here run assimilation experiments with a range of errors: 0.5K, 1K, 2K, 3K and 4K. The potential suitability of these values can be judged in reference to Figure 21, which shows the standard deviation of the first guess departure statistics for all channels of GIIRS. Firstly, the noisy region in the MW below 1800cm^{-1} is apparent, and should be ignored here, as there is no intention to use channels from this range. Close to absorption lines, the standard deviations are larger than off the lines, and generally, the MW band has larger standard deviation values than the LW band. Therefore, using a single observation error for both bands, as is done here, cannot be assumed valid across both bands, but by testing a range of values, we have some ability to gauge the response of different fields in the 4D-Var system to information coming from different parts of the spectrum. As mentioned in Section 20, too many channels have been used in the LW band, so inter-channel observation error correlations are likely to be non-negligible, and so with the limitation of using diagonal covariances here, the ‘optimal’ variances will be larger than they would be for a sparser channel selection.




Figure 21: Mean and standard deviation of first guess departures for GIIRS.

4.6 Assimilation results

Global GIIRS assimilation experiments have been run for 24 days, and some of the initial results are presented here. Although 24 days is a short time to draw firm conclusions about forecast impact, we can nevertheless make judgements based on short-range forecast fits to independent observations. We can also see the effect on forecast fields when we set the observations to be very small, which, although it will produce a degradation, will make it possible to see the locations and variables which can be incremented

via the adjoint and then propagated forward in time by the nonlinear forecast model. Such plots are shown in Figures 22-24 for 850hPa relative humidity, temperature and vector wind.

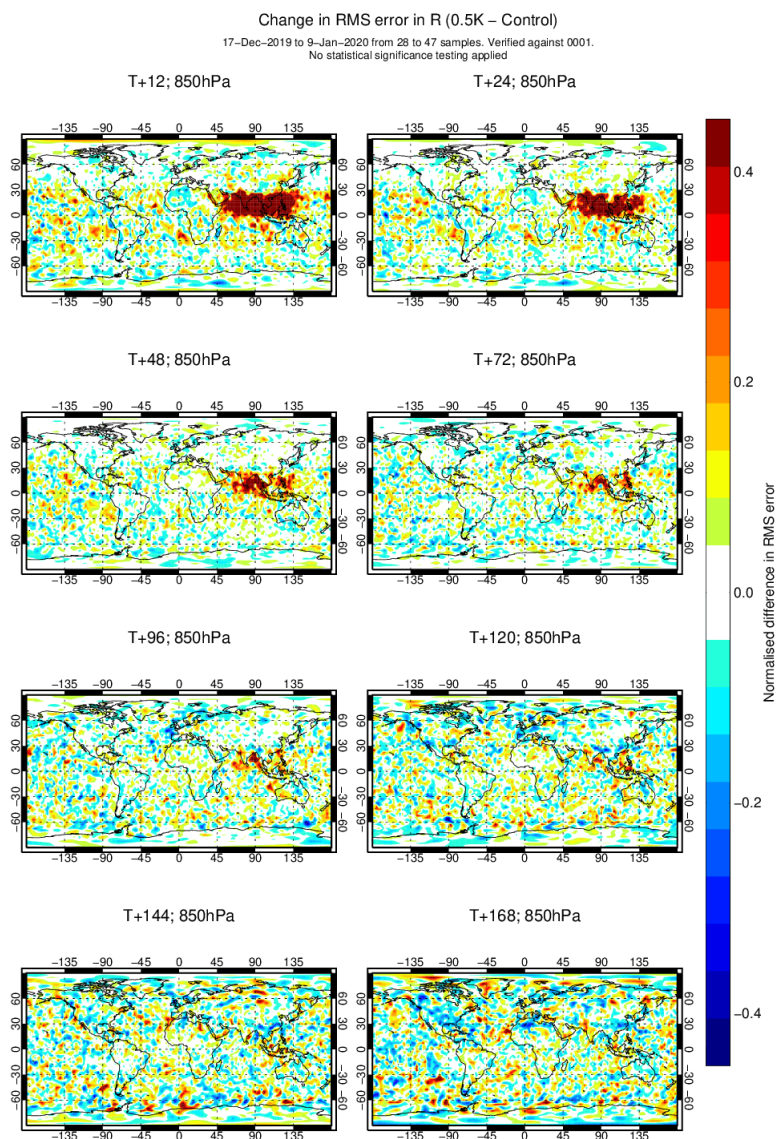


Figure 22: Change in RMS error of 850hPa relative humidity for different lead times, verified against operational analyses, when GIIRS is assimilated with 0.5K errors (i.e. very small in order to highlight the regions where positive impact can be anticipated when a more optimal assimilation strategy is adopted).

For relative humidity, the signal can be seen strongly over the GIIRS scanning domain, at short lead times, but a signal is present until at least day 3. It can be seen that the change in the humidity field is largest over ocean.

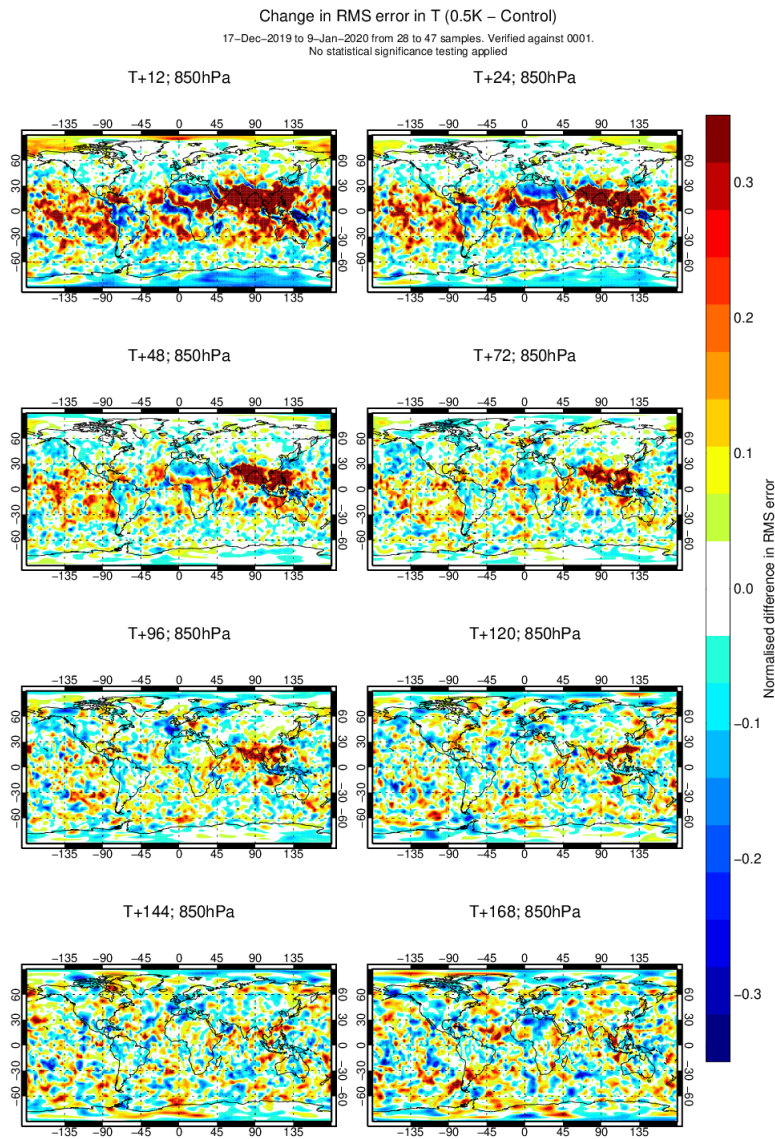


Figure 23: Change in RMS error of 850hPa temperature for different lead times, verified against operational analyses, when GIIRS is assimilated with 0.5K errors (i.e. very small in order to highlight the regions where positive impact can be anticipated when a more optimal assimilation strategy is adopted).

For temperature, a similar picture is seen, although at short lead times, the degradations are spread out more in the horizontal.

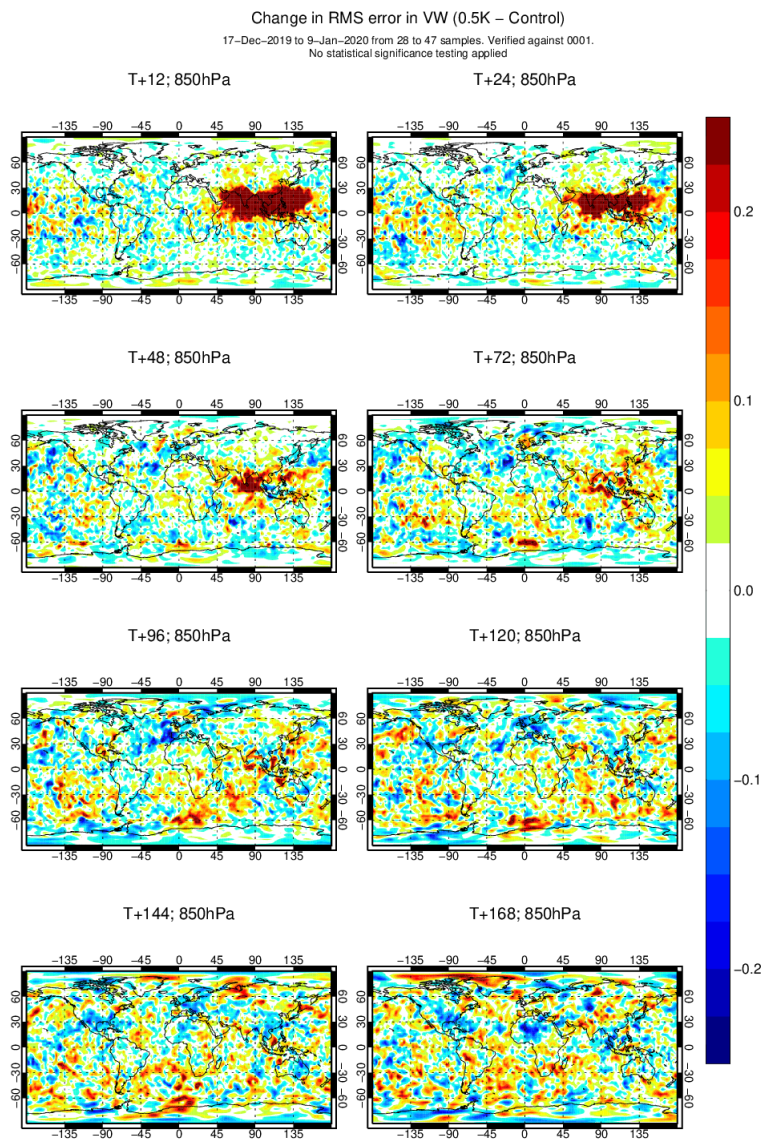
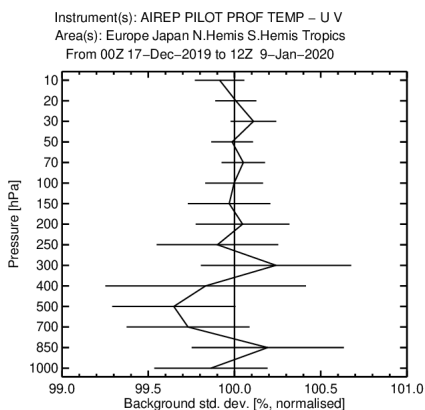


Figure 24: Change in RMS error of 850hPa vector wind for different lead times, verified against operational analyses, when GIIRS is assimilated with 0.5K errors (i.e. very small in order to highlight the regions where positive impact can be anticipated when a more optimal assimilation strategy is adopted).

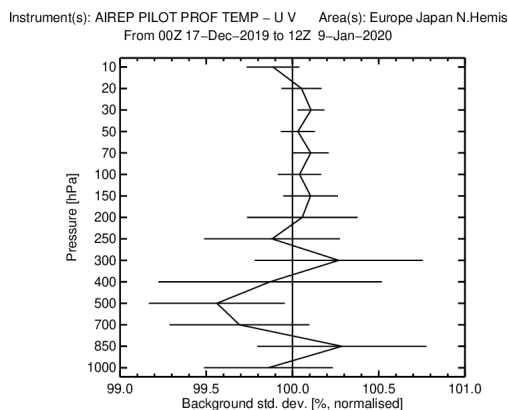
For vector wind, the geographic pattern of impact is similar to that of relative humidity and temperature, but the impact does not persist quite so long into the forecast.

To reiterate, although the impact is detrimental with such small observation errors, these figures are shown in order to demonstrate the aspects of the model where we can hope to achieve positive impact once the assimilation parameters have been optimised. Equivalent plots for larger observation errors show similar patterns but much weaker, and with 4K errors, the impact at lead times beyond 12 hours is neutral, although these statistics have been computed for a relatively short period, so signals could still emerge from the noise as the experiments are run for longer.

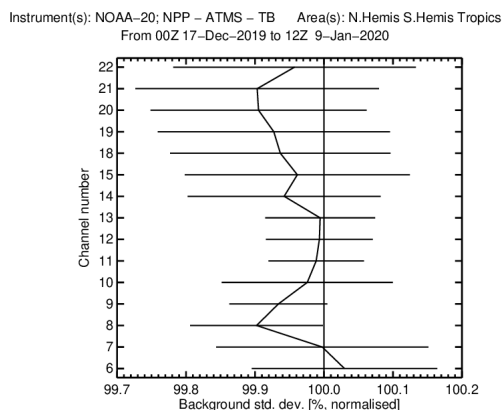
The analysis of first guess departure statistics of independent observations is a useful way to assess the impact on the short-range (12-hour) forecast. These are shown for a range of instruments in Figure 25, where the GIIRS observation errors were set to 4K, which, of the values tested, gave the most promising results.



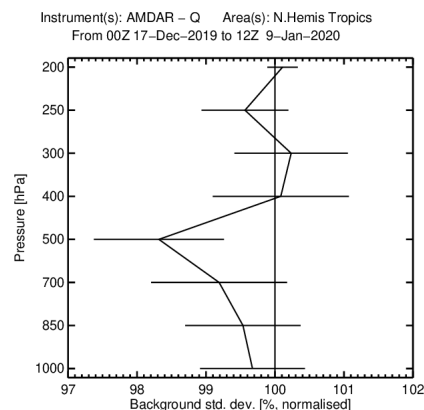
(a) *In situ* winds (global)



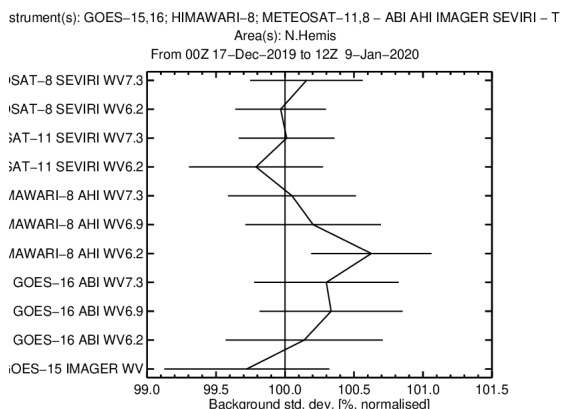
(b) *In situ* winds (NH)



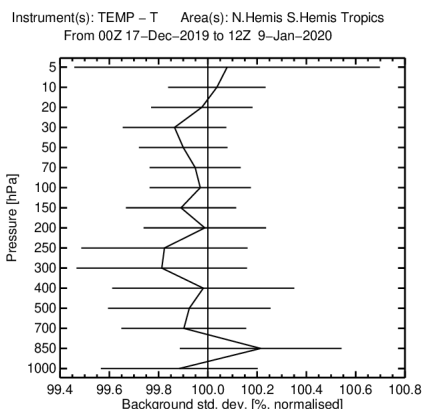
(c) ATMS global



(d) AMDAR humidity (NH/tropics)



(e) GEO radiances (NH)



(f) Radiosonde temperatures (global)

Figure 25: First guess departure statistics for the GIIRS assimilation experiment using 4K observation errors.

Initially, consider the *in situ* wind observations. These show an improvement at 500-700hPa globally

with borderline statistical significance, but restricting the verification to the northern hemisphere shows a similar profile of impact, but is statistically significant. The ability to improve the wind field by assimilating humidity-sensitive observations from temporally-frequent observations is thought to be one of the biggest benefits of assimilating GEO hyperspectral observations (Tony McNally, private communication). This suggests that observation errors of 4K are not far from their optimal value for the MW channels.

More directly-measured verification of humidity is provided by ATMS and the AMDAR humidity observations. The ATMS water vapour channels (18 and above) show a small improvement which, although not statistically significant, is consistent for all the channels. There is also some borderline improvement in channel 8 which is sensitive to upper-tropospheric temperature. The AMDAR humidity observations also show significant improvements at 500hPa.

Verifying against radiances from other geostationary satellites in the northern hemisphere shows a statistically significant degradation for the highest-peaking water vapour channel of Himawari-8, whose coverage includes the GIIRS scanning domain. This may be because the analysis and short-range forecast of tropospheric humidity in the Himawari-8 region was previously dominated by information coming from Himawari-8 itself, but with GIIRS, the analysis is able to be pulled closer to another observation type and hence further from Himawari-8. This is speculation, and will be investigated further.

Finally, radiosonde temperatures show a fairly neutral signal, but generally the points are to the left of the 100% line.

The period of this experiment is relatively short. From experience, signals seen at this length of experimentation time tend to persist, but this is not always the case, and signals can even change sign. This is, however, quite encouraging evidence that the observations are providing useful information to the assimilation system via the humidity field, and thence to the wind field. The 4K observation errors are clearly too large for the LW channels, so we are strongly down-weighting the surface and UTLS temperature-sounding channels. In fact, even with smaller observation errors, we do not see a positive signal in the lower stratosphere, which presumably is because with these errors we are over-weighting the WV channels, and hence degrading the analysis quite significantly at all heights.

Given the simplicity of the initial assimilation setup, these results are encouraging and should provide a solid starting point for further developments.

Acknowledgements

Many thanks to the following for their assistance and input:

Tony McNally, Reima Eresmaa, Marco Matricardi, Peter Lean, Elias Holm, Gabor Radnoti, Ioannis Malas, Qiang Guo, Simon Elliott, Dorothée Coppens, Bob Knuteson, Dave Tobin, Hank Revercomb, Wei Han, Fabien Carminati, Xianjun Xiao, Nigel Atkinson, James Hocking, Jérôme Vidot, Sharon Nebuda, Haixia Liu, Jim Jung, Tom King and Pete Keehn.

References

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

- C. P. Burrows. Assimilation of radiances from geostationary satellites: second year report. *EUMETSAT/ECMWF Fellowship Programme Research Report No. 51*, 2019.
- 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.
- R. W. Hamming. *Digital Filters*. Prentice Hall inc., 1989.
- 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.
- Marco Matricardi, A. McNally, and Niels Bormann. Principal component and reconstructed radiance based assimilation techniques. *ECMWF Annual Seminar proceedings*, 2014.
- 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.
- C. D. Rodgers. *Inverse Methods for Atmospheric Sounding*. World Scientific, 2000.