

EUMETSAT/ECMWF Fellowship Programme Research Report

RR68

Enhancing the exploitation of all-sky microwave sensors at ECMWF using inter-channel error correlations

Liam Steele, Niels Bormann, Alan Geer,
Marcin Chrust and David Duncan

December 2025

Series: EUMETSAT/ECMWF Fellowship Programme Research Reports

A full list of ECMWF Publications can be found on our web site under:

<http://www.ecmwf.int/en/publications/>

Contact: library@ecmwf.int

© Copyright 2025

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. The content of this document is available for use under a Creative Commons Attribution 4.0 International Public License.

See the terms at <https://creativecommons.org/licenses/by/4.0/>.

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.

Abstract

Microwave temperature and humidity sounding data are vital for improving numerical weather prediction skill, and are one of the most important observation types in the assimilation scheme used at ECMWF. In this report we present developments to the ECMWF Integrated Forecasting System (IFS) that allow for inter-channel error correlations and variational quality control to be used with all-sky observations from the Advanced Technology Microwave Sounder (ATMS), where cloudy or precipitating observations are not screened out. We discuss the creation of different error correlation matrices and assess their impact on the forecast.

We find that assimilating ATMS in all-sky conditions without accounting for inter-channel error correlations can lead to some improvements compared to clear-sky assimilation. Adding inter-channel error correlations results in some further slight improvements, but these are mainly limited to the southern hemisphere wind forecast. Using correlation matrices derived from background departure statistics gives a similar forecast performance as using matrices derived with the Desroziers diagnostic. The results have been achieved without specific reconditioning of the observation error covariance matrix, and introducing the error correlations does not lead to a significant increase in the number of iterations needed for 4D-Var to converge. Using an error correlation matrix derived from a reconditioned diagnosed covariance matrix has a neutral impact on forecast performance and the number of iterations needed.

The capabilities developed are highly relevant for future instruments with a larger number of cloud- and precipitation-sensitive channels, such as EUMETSAT's Micro-Wave Imager and Ice Cloud Imager on the Metop-SG satellites, ESA's Arctic Weather Satellite and the future EPS-Sterna microsatellite constellation, and Spire's Hyperspectral Microwave Sounder. The developments could also potentially realise greater benefit from instruments that are currently assimilated, as well as with future all-sky infrared assimilation.

Plain language summary

When producing weather forecasts we need to know the current state of the Earth's atmosphere and surface. To do this, ECMWF uses a technique called data assimilation, in which a large number of observations of the Earth system are combined with a computer model. This creates the initial conditions from which a forecast is performed.

In this report we develop and investigate a new capability to better take into account observations from the Advanced Technology Microwave Sounder (ATMS). The instrument provides observations in the microwave part of the spectrum, with different channels sensitive to temperature and humidity over different vertical layers of the atmosphere. When assimilating data from these different channels it is possible that some of the errors in the observations, or in the way they are represented in the assimilation system, are correlated between channels. For instance, two channels with similar sensitivity to clouds may be affected by similar errors of how clouds are represented in the forecast model or the radiative transfer. We assess the impact of taking these correlations into account, and find that they can lead to some improvements to the forecast up to around day 8, particularly for the wind in the southern hemisphere.

At present there are numerous instruments planned for launch that have large numbers of channels sensitive to clouds and precipitation. As such, the developments presented in this report could be useful in extracting as much information as possible from these instruments, resulting in improved weather forecasts.

1 Introduction

Microwave temperature and humidity sounding observations are one of the most important types of data in the assimilation scheme used at ECMWF (Bormann *et al.*, 2019). While humidity sounders are sensitive to water vapour, cloud and precipitation, temperature sounders are relatively insensitive to clouds, and provide important information about the atmosphere below. Microwave observations were initially only used in ‘clear-sky’ conditions, where cloudy or precipitating observations were screened out. However, developments in model physics, assimilation techniques and radiative transfer codes led to the use of all-sky assimilation, with all observations assimilated using the same scattering-capable radiative transfer model (Geer *et al.*, 2017, and references therein). The additional observations included in the all-sky framework come from meteorologically-active regions, and thus can have a larger impact on the forecast than those in clear-sky regions that are already covered by other observations.

Radiances from microwave imagers were the first to be assimilated in the all-sky framework, followed by microwave humidity sounders and then microwave temperature sounders. At ECMWF, the all-sky approach is now used to assimilate microwave-sounding channels on the Special Sensor Microwave - Imager/Sounder (SSMIS; Geer, 2013; Baordo and Geer, 2016), Microwave Humidity Sounder (MHS; Geer *et al.*, 2014), GPM Microwave Imager (GMI; Lean *et al.*, 2017), Micro-Wave Humidity Sounder-2 (MWHS-2; Lawrence *et al.*, 2018; Bormann *et al.*, 2021), and Advanced Microwave Sounding Unit-A (AMSU-A; Duncan *et al.*, 2022). At present, observations from the Advanced Technology Microwave Sounder (ATMS) instrument are still assimilated in the clear-sky framework, where cloudy and precipitating scenes are discarded.

Satellite data often display horizontally or vertically correlated errors (Janjić *et al.*, 2018), which can arise from representation errors (e.g. errors in the observation operator or errors in the location of clouds and precipitation, which are most prominent in humidity-sensitive channels) or instrument noise (which is most noticeable in temperature sounding channels that are insensitive to clouds and precipitation). As different channels sample different levels of the atmosphere, vertical error correlations are also called inter-channel error correlations. Accounting for inter-channel error correlations has led to significant positive forecast impacts when assimilating clear-sky data from the Infrared Atmospheric Sounding Interferometer (IASI; Weston *et al.*, 2014; Bormann *et al.*, 2016; Campbell *et al.*, 2017), Cross-track Infrared Sounder (CRIS; Eresmaa *et al.*, 2017), and geostationary instruments such as the Spinning Enhanced Visible Infra-Red Imager and Advanced Himawari Imager (SEVIRI and AHI; Burrows, 2018). Of particular note for the work in this report, Weston and Bormann (2018) found that assimilating ATMS in clear-sky conditions with inter-channel error correlations resulted in; (i) significant improvements to the short-range first guess fits against independent temperature, humidity and wind observations; and (ii) improvements of 0.5–1% in extra-tropical forecast scores out to day 5. More recently, several studies have experimented with accounting for inter-channel error correlations in all-sky assimilation, both for infrared (Feng and Pu, 2025; Geer, 2019) as well as microwave radiances (Ishibashi, 2024), with promising results. Combined, these studies suggest the possibility that all-sky assimilation of ATMS with inter-channel error correlations taken into account may provide even more benefits.

In ECMWF’s Integrated Forecasting System (IFS), observations assimilated in clear-sky conditions can have inter-channel error correlations accounted for. At present, this consists predominantly of infrared instruments on geostationary and low Earth orbit satellites; SEVIRI, AHI, ABI, CRIS, IASI and the Atmospheric Infra-Red Sounder (AIRS). ATMS is the only microwave sensor that is still assimilated in clear-sky conditions. For observations assimilated in all-sky conditions at ECMWF, variational quality control (VarQC; Andersson and Järvinen, 1999) is used to downweight or reject observations that are anomalous compared to other neighbouring observations, resulting in analysis departures becoming more

Gaussian. However, at present no VarQC takes place when inter-channel error correlations are accounted for. This is because VarQC assumes a diagonal observation error covariance matrix, which is not the case in the presence of correlations. For clear-sky assimilation, ignoring the quality control procedure is not much of a hindrance as problematic observations are already screened out prior to assimilation. However, for all-sky assimilation it is vital that quality control is performed (Geer and Bauer, 2011). For this reason, observations currently assimilated in all-sky conditions have VarQC active but inter-channel error correlations are ignored.

In this report we present developments to the IFS that allow for inter-channel error correlations and variational quality control to be used in tandem with all-sky observations. In general we follow the method used by Geer (2019), who perform VarQC in the eigen-space of the observation-error covariance matrix in their all-sky assimilation of IASI radiances with correlated errors in the IFS. We focus on the ATMS instrument, but the developments will be important for future instruments with increased numbers of channels sensitive to clouds and precipitation, such as:

1. ESA's Arctic Weather Satellite (which has been assimilated at ECMWF since July 2025), and on the future constellation of EPS-Sterna microsatellites, with channels at 183 and 325 GHz.
2. The combination of EUMETSAT's Micro-Wave Imager and Ice Cloud Imager on the Metop-SG satellites, with channels at 183, 325 and 448 GHz.
3. Spire's Hyperspectral Microwave Sounder, a demonstration concept on a single satellite, with several hundred channels around 183 GHz.

The report is structured as follows. In section 2 we give an overview of the ATMS instrument, then in section 3 we describe the steps needed to move ATMS from clear-sky to all-sky assimilation. In section 4 we describe the changes made to the IFS to allow inter-channel error correlations and variational quality control to be used with all-sky assimilation. In section 5 we discuss the creation of inter-channel error correlation matrices, and in section 6 we present the results of our experiments. Finally in section 7 we give our conclusions and discuss future plans.

2 Overview of the ATMS instrument

The ATMS instrument has been flown on NOAA's polar orbiting satellites, beginning with S-NPP in 2012, followed by NOAA-20 and NOAA-21 in 2018 and 2023 respectively (see Bormann *et al.*, 2013; Lawrence and Bormann, 2014; Weston *et al.*, 2017; Weston and Bormann, 2018; Bormann *et al.*, 2023, for details on the implementation and improvements to the assimilation of these instruments). These satellites all have a near-constant local time ascending node (LTAN) of around 13:30 at the equator. ATMS is a cross-track scanner, with a swath width of 2200 km and 96 fields of view (FOV). It combines temperature and humidity sounding channels which are similar to those on AMSU-A and MHS. At nadir, the spatial resolution is 75 km for channels 1–2, 32 km for channels 3–16 (temperature sounding) and 16 km for channels 17–22 (humidity sounding).

Information about the channels is given in Table 1. There are ten temperature sounding channels assimilated that are centred on the oxygen absorption band between 50–60 GHz, with similar channels found on AMSU-A. There are five humidity sounding channels assimilated that are centred on the strong water vapour absorption line at 183.31 GHz. Similar channels are also found on MWHS-2, MHS, SSMIS and GMI. As well as water vapour, the 183 GHz channels are sensitive to cloud and precipitation due to

| ATMS | | | Other instrument channel numbers | | | | |
|---------|-----------------------------|------|----------------------------------|--------|-----|-------|-----|
| Channel | Frequency (GHz) | A | AMSU-A | MWHS-2 | MHS | SSMIS | GMI |
| 1 | 23.8 | 0.50 | | | | | |
| 2 | 31.4 | 0.50 | | | | | |
| 3 | 50.3 | 0.15 | | | | | |
| 4 | 51.76 | 0.15 | | | | | |
| 5 | 52.8 | 0.35 | | | | | |
| 6 | 53.596 ± 0.115 | 0.20 | 5 | | | | |
| 7 | 54.4 | 0.10 | 6 | | | | |
| 8 | 54.94 | 0.05 | 7 | | | | |
| 9 | 55.5 | 0.05 | 8 | | | | |
| 10 | $f_0 = 57.290344$ | 0.05 | 9 | | | | |
| 11 | $f_0 \pm 0.217$ | 0.05 | 10 | | | | |
| 12 | $f_0 \pm 0.3222 \pm 0.048$ | 0.15 | 11 | | | | |
| 13 | $f_0 \pm 0.3222 \pm 0.022$ | 0.15 | 12 | | | | |
| 14 | $f_0 \pm 0.3222 \pm 0.010$ | 0.15 | 13 | | | | |
| 15 | $f_0 \pm 0.3222 \pm 0.0045$ | 0.05 | 14 | | | | |
| 16 | 88.2 | 0.50 | | | | | |
| 17 | 165.5 | 0.50 | | | | | |
| 18 | 183.31 ± 7.0 | 0.50 | | 15 | 5 | 9 | 13 |
| 19 | 183.31 ± 4.5 | 0.50 | | 14 | | | |
| 20 | 183.31 ± 3.0 | 0.30 | | 13 | 4 | 10 | 12 |
| 21 | 183.31 ± 1.8 | 0.30 | | 12 | | | |
| 22 | 183.31 ± 1.0 | 0.30 | | 11 | 3 | 11 | |

Table 1: ATMS channel information, along with the VarQC prior probabilities of gross error, A, adopted in the all-sky assimilation of ATMS when used without inter-channel error correlations (see Section 4.1 for further details). Also shown are the channel numbers of other instruments that share similar frequencies. Rows highlighted blue show channels not assimilated.

scattering from hydrometeors and absorption and emission by cloud liquid water (Geer *et al.*, 2017). The 183 GHz channels sample the atmosphere at different heights in the troposphere, while the 50–60 GHz channels sample the troposphere and stratosphere (Figure 1).

To achieve spatial resolution and noise characteristics in the IFS similar to AMSU-A, 3×3 averaging is performed on channels 3–22, where the neighbouring three scan positions and scan lines are averaged (Bormann *et al.*, 2013). Channels 9–15 are assimilated with no additional screening. For other channels which have some surface sensitivity, screening is applied to remove data over Antarctica (ch 6–8), over sea ice (ch 6 and 18–19), over snow (ch 18–19), polewards of $\pm 60^\circ$ over ocean (ch 18–19) and over high orography (ch 6–7 and 18–22). Variational bias correction (VarBC) is applied, with a globally constant predictor, four air mass predictors based on layer thicknesses, and third-order polynomial scan predictors. As of cycle 49R2, there is also an orbital bias predictor (Bormann *et al.*, 2023).

At ECMWF, the ATMS instrument has been operationally assimilated in the clear-sky framework since 26 September 2012 for S-NPP, 22 May 2018 for NOAA-20, and 23 August 2023 for NOAA-21. At present, few operational centres assimilate ATMS in all-sky conditions, likely because it was initially




Figure 1: Jacobians for (a) temperature and (b) humidity channels assimilated on ATMS, for a standard mid-latitude reference profile, normalised by the layer thickness in terms of log pressure. Both are with respect to perturbations in each layer, and the humidity Jacobians have been calculated with respect to a 10% increase in humidity. The dashed black line in panel (a) shows the pressure at the top of panel (b).

assimilated for clear-sky conditions and has not yet been moved to the all-sky framework (as is the case at ECMWF). [Zhu et al. \(2019\)](#) assimilated ocean-only ATMS all-sky radiances into the Global Forecast System (GFS), but found it had an overall neutral impact on forecast skill. [Wu et al. \(2019\)](#) tested the assimilation of ocean-only ATMS all-sky radiances into the regional-scale Hurricane Weather Research and Forecasting (HWRF) model to study hurricane Cristobal, noting an overall improved fit to observations, but with noticeable biases still evident.

In terms of the ATMS data quality, Figure 2 shows the mean and standard deviation of O–B for the assimilated temperature and humidity channels, with comparisons against other instruments with similar channels. This reveals information about observation biases and instrument noise. The data are from July 2022, and observations are only used for the Figure if they match the following criteria:

1. The observation is over ice-free oceans with $\text{abs}(\text{lat}) < 60^\circ$ and instrument zenith angle $< 60^\circ$.
2. The scattering index, SI (the difference between the brightness temperatures at 88.2 and 165.5 GHz) is < 3 K (for MHS and MWHS-2) or the liquid water path, LWP, is $< 0.3 \text{ kg m}^{-2}$ (for AMSU-A). See Section 3.1 for further details on the SI and LWP calculations.

These selection criteria ensure only ‘clear-sky’ data over oceans are used (so the results are not impacted by incorrect modelling of clouds or surface properties) and potential issues with observations at the outer scan positions of each swath are avoided.

In terms of biases (Figure 2a), both ATMS instruments generally have a negative bias for the temperature channels, similar to AMSU-A, but with a slightly larger magnitude. The bias in the humidity channels

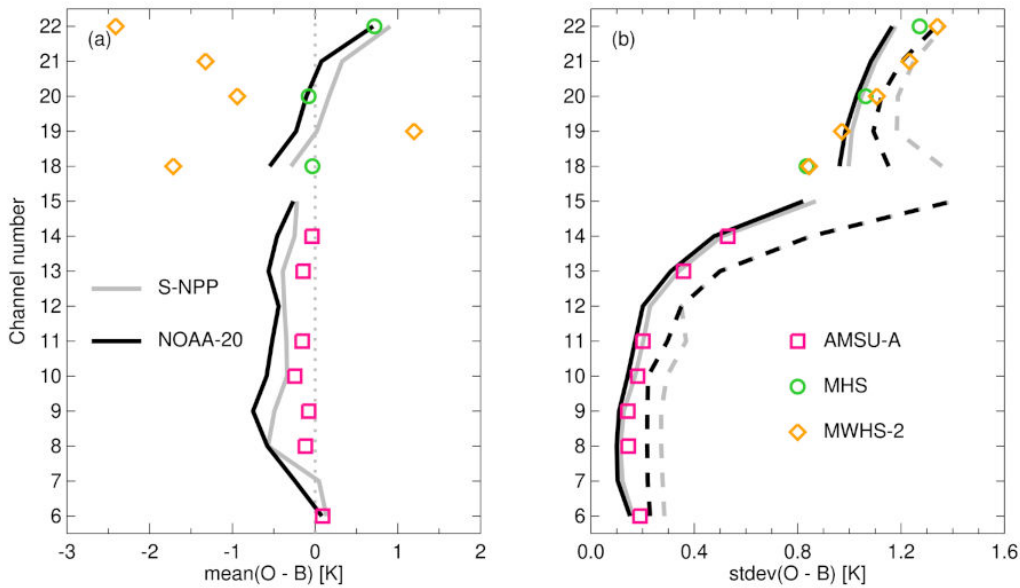


Figure 2: Background departure statistics for the assimilated channels of ATMS on S-NPP and NOAA-20, calculated from clear-sky data covering 1–31 July 2022. The mean $O-B$ (a) is shown before bias correction, and the standard deviation of $O-B$ (b) is shown after bias correction. Dashed lines in (b) show the assigned ATMS clear-sky observation error values. Coloured symbols show comparisons against similar channels of NOAA-15 AMSU-A, METOP-B MHS and FY-3C MWHS-2.

switches from negative to positive with increasing channel number, and is similar in magnitude to MHS. After bias correction (not shown) the mean $O-B$ s are similar for both ATMS instruments, and close to zero. In terms of noise, Figure 2b shows that ATMS on NOAA-20 has slightly lower noise than on S-NPP, and that after bias correction the noise is comparable to other similar instruments assimilated.

3 Moving ATMS from clear-sky to all-sky assimilation

When moving from clear-sky to all-sky assimilation, two important pieces of information are required. Firstly a variable error standard deviation needs provided, which depends on the level of ‘cloudiness’ of each observation. This is not required for clear-sky assimilation, where a constant error standard deviation can be used for each channel. Secondly, variational quality control is required, and so the prior probability of gross error needs to be supplied for each channel. The screening of observations can also differ between clear-sky and all-sky, and here we follow the methods used for other all-sky sensors (see Geer *et al.*, 2022).

3.1 Observation error modelling

As the IFS cannot precisely forecast the location and intensity of clouds and precipitation, the dominant source of random error in cloudy and precipitating situations is representativity, not instrument noise or radiative transfer inaccuracies (Geer and Bauer, 2011). For operational assimilation, these representativity errors are treated as observation errors, so an observation error model is required. Here we use the error model of Geer and Bauer (2011), which was originally developed for use with microwave imager

radiances. In this model, observation errors are given as a function of a cloud proxy, which is calculated from the liquid water path (LWP) or scattering index (SI) depending on the channel and surface type.

Channels 11–15 do not have a cloud-dependent error model as they are not sensitive to the troposphere. For temperature sounding channels 6–10 over oceans, the liquid water path (LWP) is used for the cloud proxy, with units of kg m^{-2} . This is based on [Grody *et al.* \(2001\)](#), but with updated coefficients; see [Geer *et al.* \(2012\)](#). The calculation is dependent on the zenith angle θ , and the brightness temperatures from channels 1 and 2 (TB_1 and TB_2) with frequencies of 23.8 GHz and 31.4 GHz:

$$\text{LWP} = \cos \theta [8.24 - (2.539 - 1.744 \cos \theta) \cos \theta + 0.754 \ln(285 - TB_1) - 2.265 \ln(285 - TB_2)]. \quad (1)$$

This calculation is performed using brightness temperatures from the observations as well as from the IFS short-range forecast (after bias correction), and these are combined to give the symmetric cloud predictor $C_{\text{sym}} = (\text{LWP}_{\text{obs}} + \text{LWP}_{\text{IFS}})/2$. The values from the forecast are obtained through spatial interpolation of model profiles (temperature, pressure, humidity, hydrometeors and a number of other variables) to observation locations, followed by a radiative transfer calculation using RTTOV-SCATT. Figure 3a shows a global map of the calculated LWP for one day of observations from NOAA-20 ATMS. Dark red regions with large LWP can be seen corresponding to frontal cloud features in the southern extratropics, the inter-tropical convergence zone (ITCZ), and thunderstorms in the western Pacific Ocean.

For temperature sounding channels 6–10 over land, and for humidity sounding channels over ocean and land, a scattering index is used for the cloud proxy. This is based on the difference between brightness temperatures at two different frequencies:

$$\text{SI} = TB_{v1} - TB_{v2}. \quad (2)$$

These frequencies are 23.8 GHz and 31.4 GHz (ATMS channels 1 and 2) for temperature sounding channels over land, and 88.2 GHz and 165.5 GHz (ATMS channels 16 and 17) for humidity sounding channels. For humidity sounding channels over oceans, where water vapour absorption can cause large differences between the brightness temperatures, a clear-sky term (calculated from the model background) is subtracted from Equation 2, which removes the effect of water vapour absorption:

$$\text{SI}_{\text{ocean}} = (TB_{v1} - TB_{v2}) - (TB_{v1}^{\text{clear}} - TB_{v2}^{\text{clear}}). \quad (3)$$

As with LWP, the SI calculations are performed using data from observations and the IFS, and combined to give the symmetric cloud predictor $C_{\text{sym}} = (\text{SI}_{\text{obs}} + \text{SI}_{\text{IFS}})/2$. Figure 3b shows a global map of the calculated SI over land and ocean for one day of observations from NOAA-20 ATMS. As with the LWP, large scattering indices highlight frontal features in the southern hemisphere, the ITCZ and thunderstorm activity in the Pacific. Over land, negative values of SI can be seen, particularly around coasts and close to inland lakes. These negative values appear because a clear-sky term is not subtracted from the SI over land like it is for the ocean. In future it may be worth performing this subtraction, given the relatively large amount of negative values over land.

The C_{sym} values for each observation are then used in conjunction with pre-determined clear and cloudy cloud proxies (C_{clr} , C_{cld}) and their associated errors (σ_{clr} , σ_{cld}) to determine the observation error, σ :

$$\sigma(C_{\text{sym}}) = \begin{cases} \sigma_{\text{clr}} & C_{\text{sym}} \leq C_{\text{clr}} \\ \sigma_{\text{clr}} + \gamma^\alpha (\sigma_{\text{cld}} - \sigma_{\text{clr}}) & C_{\text{clr}} < C_{\text{sym}} < C_{\text{cld}} \\ \sigma_{\text{cld}} & C_{\text{sym}} \geq C_{\text{cld}} \end{cases} \quad (4)$$

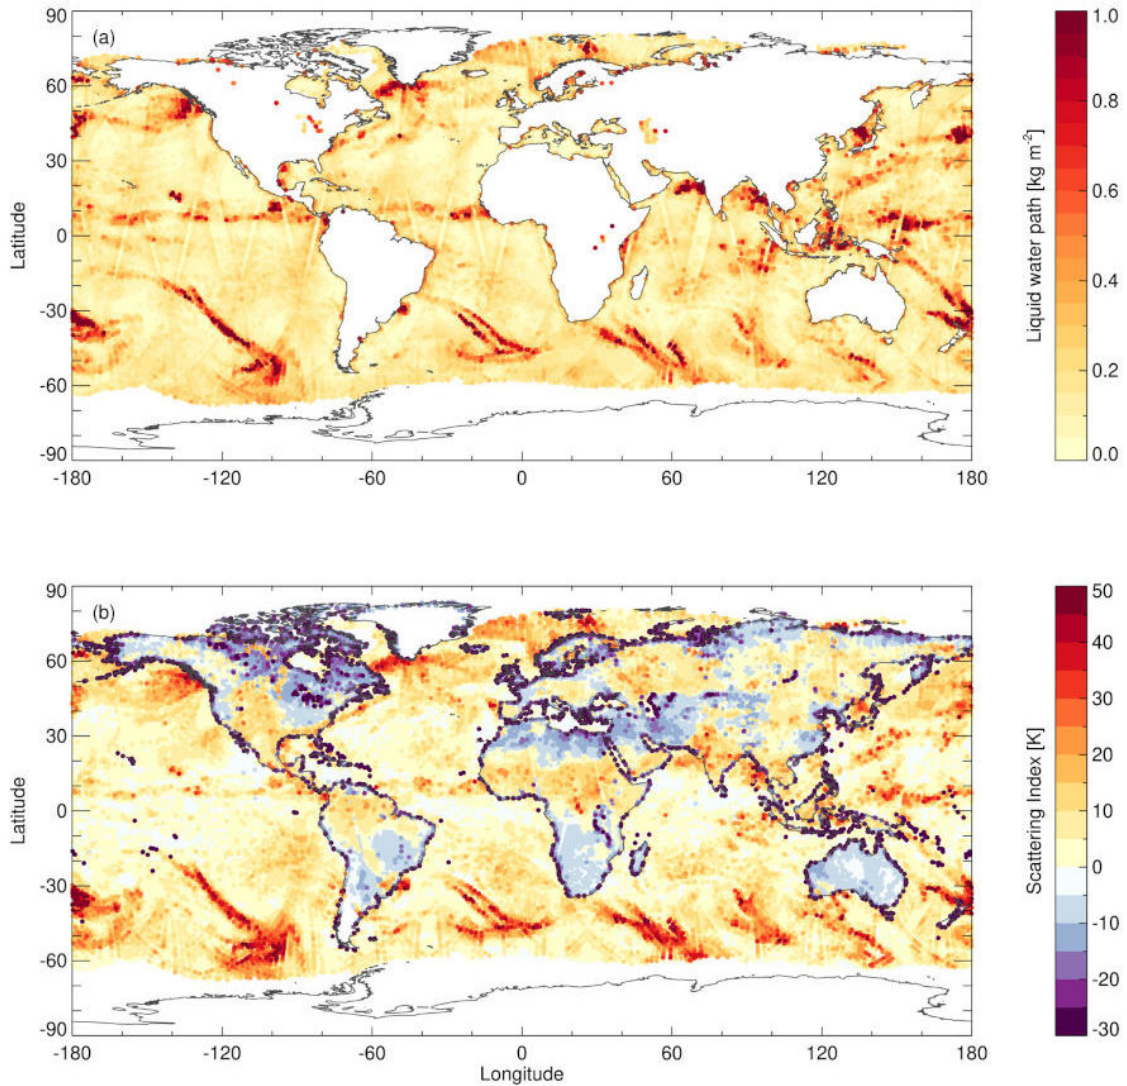


Figure 3: Maps of (a) liquid water path and (b) scattering index calculated for NOAA-20 ATMS observations on 15 July 2022.

where $\gamma = (C_{\text{sym}} - C_{\text{clr}})/(C_{\text{cld}} - C_{\text{clr}})$ and α can be equal to 1 or 2 for linear or quadratic interpolation between the boundary values.

Table 2 gives the values used in the all-sky observation error model for the assimilated channels of NOAA-20 ATMS, with S-NPP ATMS values being similar. They have been derived from background departure statistics using experiments in which ATMS has been actively assimilated via the all-sky assimilation route. Figure 4 shows examples of the standard deviation of O–B (black lines) from NOAA-20 ATMS as a function of C_{sym} for two channels (one temperature-sounding and one humidity-sounding) over ocean and land. The dashed pink line shows the corresponding observation errors from Equation 4. As can be seen, the error model begins to deviate from the standard deviation of O–B for larger C_{sym} values. However, as shown by the green line, the number of observations for large C_{sym} is low, and it is more important for the error model to fit the bulk of the observations, which are found at lower C_{sym} values.

| | Land | | | | Ocean | | | |
|---------|-----------------------|-----------------------|------------------|------------------|-----------------------|-----------------------|------------------|------------------|
| Channel | σ_{clr} | σ_{cld} | C_{clr} | C_{cld} | σ_{clr} | σ_{cld} | C_{clr} | C_{cld} |
| 6 | 0.26 | 2.55 | -2 | 14 | 0.21 | 1.70 | -0.1 | 0.8 |
| 7 | 0.16 | 0.90 | -2 | 17 | 0.14 | 0.54 | -0.05 | 0.8 |
| 8 | 0.15 | 0.45 | 0 | 16 | 0.14 | 0.25 | -0.05 | 1.0 |
| 9 | 0.14 | 0.25 | 0 | 15 | 0.14 | 0.19 | -0.05 | 0.2 |
| 10 | 0.17 | 0.22 | 0 | 20 | 0.17 | 0.20 | 0 | 0.8 |
| 11 | 0.23 | - | - | - | 0.23 | - | - | - |
| 12 | 0.34 | - | - | - | 0.34 | - | - | - |
| 13 | 0.50 | - | - | - | 0.50 | - | - | - |
| 14 | 0.81 | - | - | - | 0.81 | - | - | - |
| 15 | 1.40 | - | - | - | 1.40 | - | - | - |
| 18 | 2.3 | 35 | 0 | 22.4 | 2.1 | 27 | 0 | 30 |
| 19 | 2.0 | 28 | 0 | 22.4 | 1.8 | 20 | 0 | 32 |
| 20 | 1.8 | 20 | 0 | 22.5 | 1.6 | 16 | 0 | 34 |
| 21 | 1.7 | 13 | 0 | 22.5 | 1.6 | 11 | 0 | 36 |
| 22 | 1.8 | 8 | 0 | 22.5 | 1.7 | 8 | 0 | 37 |

Table 2: Parameters of the observation error model over land and ocean for assimilated channels of NOAA-20 ATMS (see equation 4). The σ values are in Kelvin, and the cloud proxies, C , are in kg m^{-2} (channels 6–10 over the ocean) and Kelvin (elsewhere). As channels 11–15 are not sensitive to the troposphere, the all-sky error model is not applicable, and so only clear-sky values are specified.

3.2 Variational quality control

In the 4D-Var cost function, errors in the observations and background are assumed to be Gaussian, which is not necessarily true. As such, variational quality control (VarQC; see [Andersson and Järvinen, 1999](#)) is performed so that the probability density functions (PDFs) of analysis departures become more Gaussian with fewer outliers. This is done by estimating, during the assimilation, the posterior probability that the error in the observations is drawn from a white-noise rather than a Gaussian distribution. In order to perform VarQC, a prior probability of gross error, A , is required, which specifies the probability that the error in the observation does not belong to a Gaussian distribution. For ATMS, the values for each channel are given in Table 1. The prior probability of gross error is larger for channels with Jacobians that peak closer to the surface, and for channels with more sensitivity to clouds and precipitation. The values of A were obtained empirically, in line with values for similar heritage channels.

Figure 5 shows examples of PDFs for two temperature-sounding channels (6 and 10) over ocean and land for clear and cloudy situations. Channel 6 is the lowest-peaking channel assimilated (see Figure 1), and is sensitive to cloud and precipitation. Channel 10, the highest-peaking channel for which an error model is supplied, peaks in the stratosphere and hence is little affected by clouds and precipitation. This difference is evident in how well the PDFs match Gaussian curves. Figure 5(b,d) shows that the channel 10 PDFs closely match a Gaussian within the range of around ± 1 K. For channel 6, the clear-sky PDFs also match a Gaussian within the range ± 1 K. However, for the cloudy sample the PDFs depart from being Gaussian, with sharp peaks and broad tails. Due to VarQC, the data in the wings which depart from Gaussian will receive less weight in the analysis, as can be seen by the posterior probability of gross error. VarQC and the calculation of weights is described in more detail in the next section.

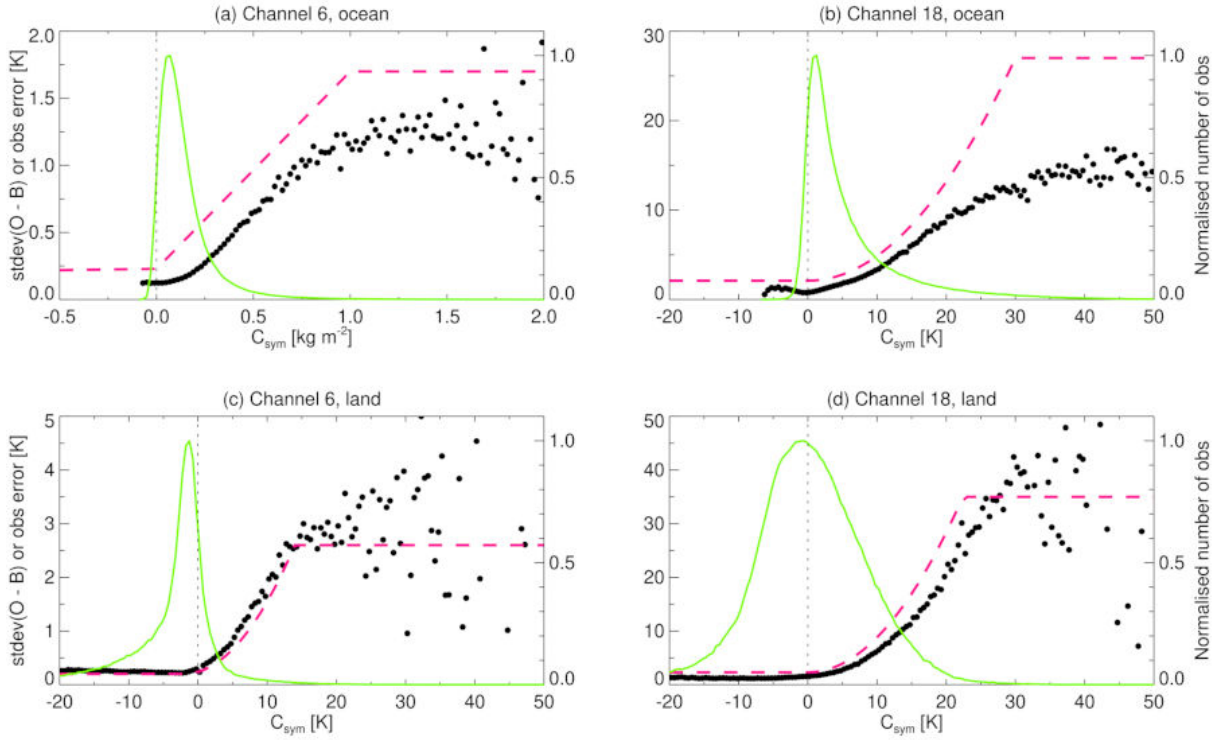


Figure 4: Standard deviation of O–B from NOAA-20 ATMS as a function of the symmetric cloud predictor (black dots) along with the observation error model (pink dashed line) and the normalised number of observations in each bin (green line). Values are shown for channels 6 and 18, over ocean and land. Based on used data from June 2022, with a C_{sym} bin size of (a) 0.02 kg m^{-2} , and (b–d) 0.5 K .

4 Correlated errors in the IFS

4.1 Current formulation

In the four-dimensional variational assimilation scheme (4D-Var) used at ECMWF, the best estimate of the model state, \mathbf{x} , is found by minimizing a cost function, J , consisting of terms relating to the background and observations. Here we are concerned with the observation term

$$J^o(\mathbf{x}) = \frac{1}{2} \mathbf{d}^T \mathbf{R}^{-1} \mathbf{d}, \quad (5)$$

where \mathbf{R} is the observation error covariance matrix and \mathbf{d} (with transpose \mathbf{d}^T) is the departure between the observations, \mathbf{y} , and the model's representation of the observations:

$$\mathbf{d} = \mathbf{y} - H(\mathbf{x}) = \mathbf{y}' - \mathbf{b} - H(\mathbf{x}), \quad (6)$$

where \mathbf{y}' are the raw observations, \mathbf{b} is a bias correction and $H()$ is the non-linear observation operator that maps from state space to observation space. For the purposes of this work \mathbf{R} has block diagonal structure, since we are not concerned with spatial error correlations. As such, the vectors \mathbf{y} and \mathbf{d} span only a single observation, with the number of elements equal to the number of levels/channels relevant to that observation. For ATMS, \mathbf{y} and \mathbf{d} contain up to 15 elements corresponding to channels 6–15 and 18–22. The complete J^o term is a summation over multiple block diagonal elements.

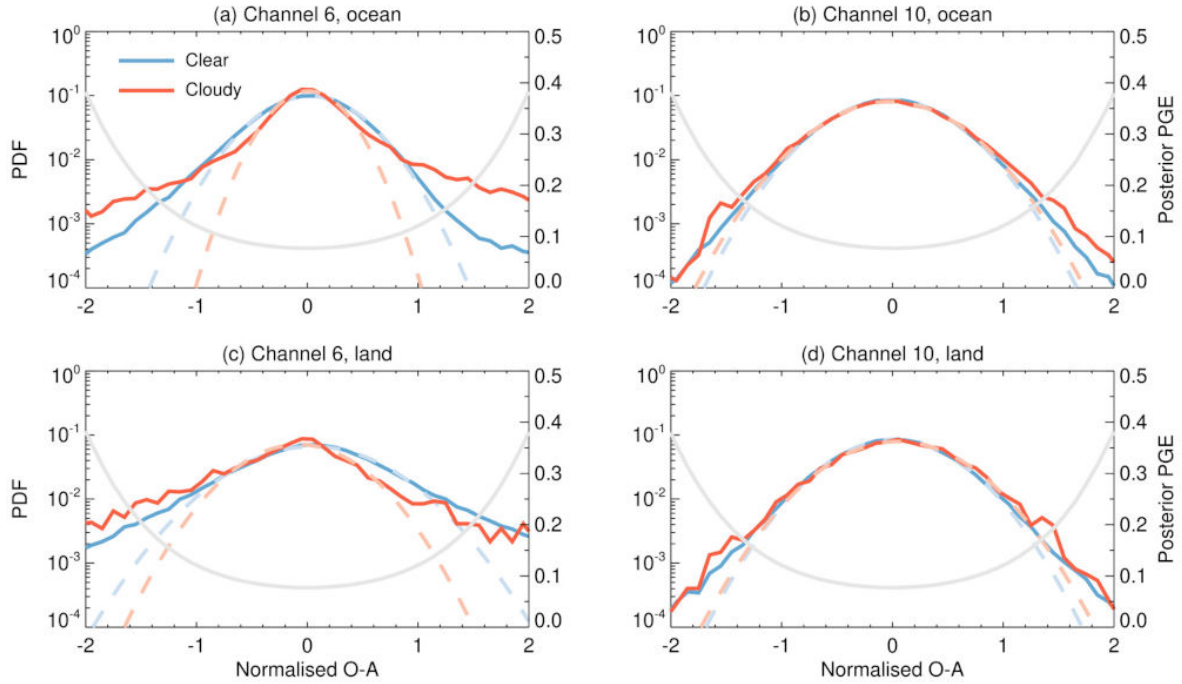


Figure 5: Each panel shows the probability density function (PDF) of normalized analysis departures (solid coloured lines), along with IDL's best fit Gaussian curve (dashed lines) and the posterior probability of gross error (grey line, right y-axis). Results are shown for channels 6 and 10, for clear and cloudy scenes over ocean and land. Cloudy is defined as a LWP > 0.3 kg m⁻² (over oceans), or a SI > 3 K (over land). Statistics are calculated from data covering 1-31 July 2022.

To allow for the application of variational quality control (VarQC), \mathbf{R}^{-1} in Equation 5 is replaced by $\tilde{\mathbf{R}}^{-1} = \mathbf{W}^{1/2} \mathbf{R}^{-1} \mathbf{W}^{1/2}$, where \mathbf{W} is a diagonal matrix of VarQC weights (Andersson and Järvinen, 1999). Thus, the weighted observation cost function is

$$J^{o,w}(\mathbf{x}) = \frac{1}{2} \mathbf{d}^T \tilde{\mathbf{R}}^{-1} \mathbf{d}. \quad (7)$$

Taking a correlation matrix, \mathbf{C} , which is symmetric and positive-definite, and a diagonal observation error variance matrix, \mathbf{V} , we can write $\mathbf{R} = \mathbf{V}^{1/2} \mathbf{C} \mathbf{V}^{1/2}$. Thus,

$$\tilde{\mathbf{R}}^{-1} = \mathbf{W}^{1/2} \mathbf{V}^{-1/2} \mathbf{C}^{-1} \mathbf{V}^{-1/2} \mathbf{W}^{1/2}. \quad (8)$$

For instruments where correlated errors are used (e.g. clear-sky ATMS) no VarQC takes place, so $\mathbf{W} = \mathbf{I}$, and $\tilde{\mathbf{R}}^{-1} = \mathbf{V}^{-1/2} \mathbf{C}^{-1} \mathbf{V}^{-1/2}$. In this case the correlation matrix is pre-computed and stored in a text file, and Cholesky decomposition is used to calculate $\tilde{\mathbf{R}}^{-1} \mathbf{d}$. Only one correlation matrix is needed per instrument, as the assimilation is for clear-sky conditions only.

For instruments where the observation errors are uncorrelated (e.g. $\mathbf{C}^{-1} = \mathbf{I}$), then $\tilde{\mathbf{R}}^{-1} = \mathbf{W}^{1/2} \mathbf{V}^{-1} \mathbf{W}^{1/2}$ is diagonal, and for each observation i it contains the inverse square of the observation error standard deviation multiplied by the VarQC weight: $w_i / (\sigma_i^o)^2$. This is the current procedure for all-sky microwave instruments. In order to calculate the VarQC weights for each observation, the unweighted observation

cost function in Equation 5 is used, which gives

$$J_i^o = \frac{1}{2} \left(\frac{d_i}{\sigma_i^o} \right)^2, \quad (9)$$

where d is the departure. The weight for each observation is then given by

$$w_i = 1 - P_i = 1 - \frac{\gamma_i}{\gamma_i + \exp(-J_i^o)}, \quad (10)$$

where P_i is the a-posteriori probability of gross error, and

$$\gamma_i = \frac{A_i \sqrt{2\pi}}{2D_i(1 - A_i)}, \quad (11)$$

with A_i the prior probability of gross error, and D_i defining the interval (centred at zero) over which there is a flat distribution. The value of D_i is set to 5 for radiance assimilation, while the value of A_i is specified per-channel for each instrument. For ATMS, the values of A_i are given in Table 1. The weights are then used in the weighted observation cost function (Equation 7).

4.2 Amended formulation

As noted in the Introduction, no variational quality control takes place when inter-channel error correlations are accounted for, due to the VarQC scheme assuming a diagonal observation error covariance matrix. As suggested by [Andersson and Järvinen \(1999\)](#), VarQC can be extended to include observation error correlations by diagonalising \mathbf{R} via an eigenvector decomposition. Here we perform the diagonalisation on the correlation matrix \mathbf{C} , giving $\mathbf{C} = \mathbf{E}\mathbf{\Lambda}\mathbf{E}^T$ and thus $\mathbf{C}^{-1} = \mathbf{E}\mathbf{\Lambda}^{-1}\mathbf{E}^T$, where $\mathbf{E} = [\mathbf{e}_1 \ \mathbf{e}_2 \ \dots \ \mathbf{e}_n]$ is a matrix of eigenvectors, \mathbf{e} , and $\mathbf{\Lambda}$ is a diagonal matrix of eigenvalues, λ . The inverse of the observation error covariance matrix can now be expressed as

$$\begin{aligned} \mathbf{R}^{-1} &= \mathbf{V}^{-1/2} \mathbf{C}^{-1} \mathbf{V}^{-1/2} \\ &= \mathbf{V}^{-1/2} \mathbf{E} \mathbf{\Lambda}^{-1/2} \mathbf{\Lambda}^{-1/2} \mathbf{E}^T \mathbf{V}^{-1/2}. \end{aligned} \quad (12)$$

If we take $\mathbf{R}^{-1/2} = \mathbf{\Lambda}^{-1/2} \mathbf{E}^T \mathbf{V}^{-1/2}$, then multiplication by the square root of the VarQC weights gives $\tilde{\mathbf{R}}^{-1/2} = \mathbf{W}^{1/2} \mathbf{\Lambda}^{-1/2} \mathbf{E}^T \mathbf{V}^{-1/2}$. Thus we obtain the amended form of the inverse of the weighted observation error covariance matrix

$$\begin{aligned} \tilde{\mathbf{R}}^{-1} &= \tilde{\mathbf{R}}^{-T/2} \tilde{\mathbf{R}}^{-1/2} \\ &= \mathbf{R}^{-T/2} \mathbf{W} \mathbf{R}^{-1/2} \\ &= \mathbf{V}^{-1/2} \mathbf{E} \mathbf{\Lambda}^{-1/2} \mathbf{W} \mathbf{\Lambda}^{-1/2} \mathbf{E}^T \mathbf{V}^{-1/2}, \end{aligned} \quad (13)$$

which can be compared with Equation 8. Calculation of the VarQC weights again requires the un-weighted observation cost function, which is now expressed as

$$J^o(\mathbf{x}) = \frac{1}{2} \mathbf{d}^T \mathbf{V}^{-1/2} \mathbf{E} \mathbf{\Lambda}^{-1} \mathbf{E}^T \mathbf{V}^{-1/2} \mathbf{d},$$

and can be simplified to

$$J_k^o = \frac{1}{2} \left(\frac{\mathbf{e}_k^T \mathbf{n}}{\lambda_k^{1/2}} \right)^2 \quad (14)$$

where \mathbf{e}_k is the k th eigenvector of \mathbf{E} , λ_k is the corresponding eigenvalue, and $\mathbf{n} = \mathbf{V}^{-1/2}\mathbf{d}$ is the normalised departure. As before, the values of J_k^o are used to calculate the VarQC weights as in Equations 10 and 11. While in the current formulation VarQC is applied to each departure, in the amended formulation VarQC is applied to each normalised eigendeparture $\mathbf{e}_k^T \mathbf{n}$. This amended formulation is broken up into the following four steps in the IFS, which are performed for each observation:

1. Calculate $\mathbf{R}^{-1/2}\mathbf{d}$, which gives the normalised departures projected onto an eigenbasis, termed normalised eigendepartures. The k th element of the vector $\mathbf{R}^{-1/2}\mathbf{d}$ is given by $\mathbf{e}_k^T \mathbf{n} / \lambda_k^{1/2}$.
2. Calculate the VarQC weights on an eigenbasis (Equation 10), using the results from step 1 to calculate J_k^o (Equation 14).
3. Multiply the results of step 1 by the weights from step 2, giving $\mathbf{W}\mathbf{R}^{-1/2}\mathbf{d}$.
4. Project the quality-controlled eigendepartures back onto the observation basis by applying the transform $\mathbf{R}^{-T/2} = \mathbf{V}^{-1/2}\mathbf{E}\mathbf{A}^{-1/2}$ to the results from step 3, giving $\tilde{\mathbf{R}}^{-1}\mathbf{d}$.

5 Creation of error correlation matrices

5.1 Methods for estimating error statistics

Several methods have been used in the past for estimating error statistics in order to determine \mathbf{R} . The Hollingsworth/Lönnerberg (HL) method (Hollingsworth and Lönnerberg, 1986) uses the expectation of the product of O–Bs (observation minus background, \mathbf{d}^b) to give

$$E[\mathbf{d}^b(\mathbf{d}^b)^T] = \mathbf{R} + \mathbf{H}\mathbf{B}\mathbf{H}^T, \quad (15)$$

where \mathbf{H} is the linearized observation operator and \mathbf{B} is the background error covariance. By assuming the background errors are spatially correlated while the observation errors are not, the contributions from \mathbf{R} and $\mathbf{H}\mathbf{B}\mathbf{H}^T$ can be separated.

Another method, proposed by Desroziers *et al.* (2005), uses the product of O–Bs and O–As (observation minus analysis, \mathbf{d}^a) to give

$$E[\mathbf{d}^a(\mathbf{d}^b)^T] = \mathbf{R}. \quad (16)$$

This method assumes the weights used by the assimilation, $\mathbf{H}\mathbf{B}\mathbf{H}^T(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})^{-1}$, are consistent with true error statistics.

Both the above methods can be used to estimate inter-channel error correlations by using \mathbf{d}^b and \mathbf{d}^a values from different channels. When creating the clear-sky correlation matrix for S-NPP ATMS, Weston and Bormann (2018) used a combination of the HL and Desroziers methods, in order to minimize the assumption about the true weights in the Desroziers method, which was violated due to the ATMS observation errors being inflated. For all-sky assimilation, Geer (2019) notes that the HL method cannot be used due to spatial correlations in the observation errors, while the Desroziers method can be used, but whenever it has been used in the past additional error inflation has been applied. For this reason, Geer (2019) opted to calculate \mathbf{R} based purely on O–B statistics.

In this report we use the O–B and Desroziers diagnostics to investigate the cloud-dependence of the error correlation matrices. The ultimate aim is to model the cloud-dependence in the error correlation structures through a suitably chosen look-up-table depending on the relevant cloud proxies. Such a

look-up-table approach has previously been used by [Ishibashi \(2024\)](#) for humidity-sensitive microwave radiances in all-sky assimilation. While [Ishibashi \(2024\)](#) employed a 1-dimensional look-up-table in a single cloud proxy, we use a 2-dimensional version, to capture the dependence on two different cloud proxies relating to the temperature and humidity sounding channels of ATMS. The approach is different from [Geer \(2019\)](#) who instead scaled the leading eigenvalues of the error covariance matrix depending on the cloud proxy. The [Geer \(2019\)](#) approach cannot be used here, as it is only applicable when the number of channels assimilated is fixed and as the assimilation of ATMS data requires the use of different cloud proxies for the temperature- and humidity-sensitive channels. Note that in the following we summarise the estimation of the error correlation matrices only. Throughout the work presented here, the modelling of observation error variances is kept the same as described in section 3.1 above.

5.2 Creating the error correlation matrices

To obtain the data for creating the correlation matrices, experiments were performed with ATMS assimilated in all-sky conditions with VarQC active, and with assigned diagonal observation errors as outlined in section 3.1. The ATMS data were actively assimilated, as O–A values were required for the Desroziers-derived correlations. For statistics based on just O–B values, the data could have been passively monitored. IFS cycle 49R1 was used for the experiments, with 137 vertical levels and 12-hour long window delayed cut-off assimilation cycles. The model resolution was set to Tco399 (triangular-cubic-octahedral truncation at wavenumber 399), giving a grid spacing of ~ 28 km, which is the typical resolution used for research. Unlike the operational clear-sky assimilation which uses 3×3 averaging (see Section 2), here we use superobbing onto Gaussian grid N80 (which has 80 grid points between each pole and the equator) corresponding to a resolution of ~ 125 km. This is close to the ~ 140 km thinning resolution used by the clear-sky 3×3 method. For operational assimilation at Tco1279 resolution a smaller superobbing scale may be preferred. However, at Tco399 resolution, additional testing with superobbing onto an N200 grid (~ 50 km resolution), and also with 3×3 averaging, showed that the results are relatively insensitive to the superobbing scale. The assimilation scheme uses three inner loops of increasing resolution, with the final loop resolution being TL255 (triangular-linear truncation at wavenumber 255), which is ~ 80 km. Background errors come from the operational ensemble of data assimilations (EDA) and represent a flow-dependent background error. For the radiative transfer calculations we used version 13 of RTTOV-SCATT (Radiative Transfer for TOVS microwave scattering package; see [Bauer et al., 2006](#); [Geer et al., 2017](#); [Geer, 2021](#); [Geer et al., 2021](#); [Geer et al., 2022](#)). Experiments were performed for three-month periods over summer 2022 (June, July and August) and winter 2022/2023 (December, January and February), but after investigation it was discovered that the correlations obtained differed very little between summer and winter, and so only the summer experiments were used.

Due to different cloud proxies being used by the temperature- and humidity-sensitive channels over land and ocean, different correlations matrices were required for these two surface types. To investigate how the correlations varied with the cloud proxy values, we binned the O–B and O–A data using bin sizes of 0.04 kg m^{-2} for the LWP and 2 K for the SI, and then inter-channel error correlations were calculated for each bin. As we continue to use the previously derived all-sky error model for the observation error standard deviations, the values diagnosed here for the observation error standard deviation are not retained. This is similar to the approach of [Campbell et al. \(2017\)](#). Figure 6 shows examples of how the correlation structure changes with SI and LWP over ocean and land between three pairs of channels; two temperature-sensitive channels (6 and 7), one temperature- and one humidity-sensitive channel (6 and 18) and two humidity-sensitive channels (18 and 22). Only results from NOAA-20 ATMS O–B data are shown, as the results from S-NPP ATMS and from Desroziers statistics are similar. In general the correlations increase as the SI and LWP increase, but it is not a linear relationship, with two-dimensional

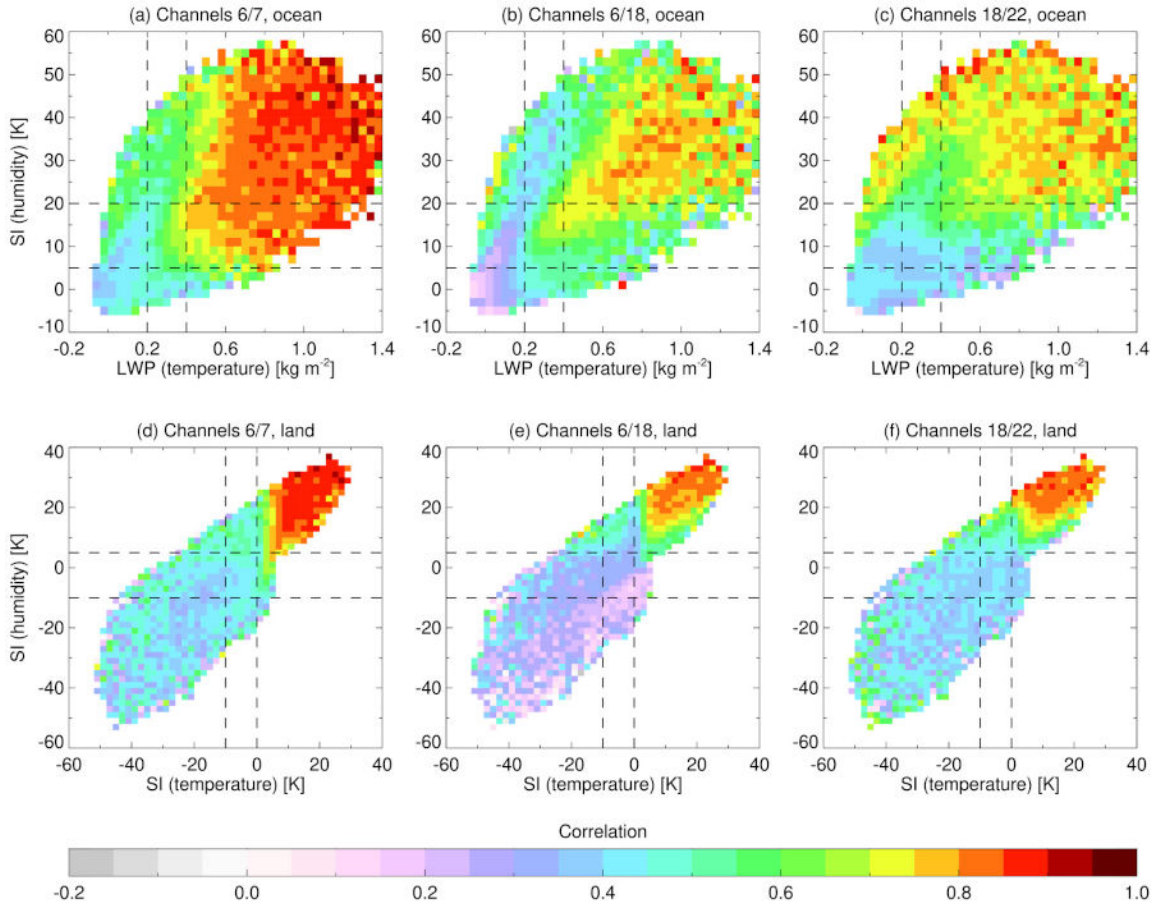


Figure 6: Error correlations between three pairs of channels over (a–c) ocean, and (d–f) land, as a function of cloud proxy. Statistics are calculated from NOAA-20 ATMS O–B data from a summer 2022 experiment (see text for details). Black dashed lines represent the cloud proxy boundaries for the chosen correlation matrices.

structure visible. For example, between channels 6 and 18 over the ocean (Figure 6b) there is a ‘tongue’ of lower correlations extending upwards and to the right of the plot. However, while these small bins capture the two-dimensional structure of the correlations, such fine binning leads to small sample sizes and hence potentially unreliable statistics. Such detailed statistics are also impractical to maintain.

As the larger SI and LWP values are unlikely to occur often, it is useful to look at the number of observations that fall into each bin. Figure 7 shows the number of observations in each of the bins used in Figure 6, normalised by the maximum number of observations present in the bins. In terms of the number of observations per 12-hour analysis cycle, the red coloured boxes in Figure 7 (which contain the most observations) correspond to ~ 8000 observations per cycle over the ocean and ~ 9000 over land. As such, a $\log_{10}(\text{normalised number of obs})$ value of 10^{-3} corresponds to around 8 and 9 observations per cycle over ocean and land respectively. Thus, there are very few observations in each analysis cycle with large LWP or SI values, and if correlation matrices were provided there are so few observations it is unlikely the analysis would be affected. As such, it was decided to use nine correlation matrices, which still captures a lot of the two-dimensional variation seen. The boundaries of the bins are shown by dashed black lines in Figures 6 and 7, and listed in Table 3. The choice of three cloud proxy bins is similar to that adopted by Ishibashi (2024), who used four cloud proxy bins to model error correlation matrices for microwave imager observations (diagnosed by the V and H polarization difference at 37 GHz), and three bins to model the correlation matrices for microwave humidity sounder observations.




Figure 7: The number of observations in the bins from Figure 6, normalised by the maximum number of observations in the bins and expressed in \log_{10} form. Black dashed lines represent the cloud proxy boundaries for the chosen correlation matrices.

| Channel sensitivity | Surface type | Low bin | Middle bin | High bin |
|---------------------|--------------|---------------------------|---|---------------------------|
| Temperature | Ocean | $< 0.2 \text{ kg m}^{-2}$ | $0.2 \text{ to } 0.4 \text{ kg m}^{-2}$ | $> 0.4 \text{ kg m}^{-2}$ |
| Temperature | Land | $< -10 \text{ K}$ | $-10 \text{ to } 0 \text{ K}$ | $> 0 \text{ K}$ |
| Humidity | Ocean | $< 5 \text{ K}$ | $5 \text{ to } 20 \text{ K}$ | $> 20 \text{ K}$ |
| Humidity | Land | $< -10 \text{ K}$ | $-10 \text{ to } 5 \text{ K}$ | $> 5 \text{ K}$ |

Table 3: Cloud proxy boundary values for the chosen correlation matrix bins.

Figure 8 shows three of the nine correlation matrices created for NOAA-20 ATMS over the ocean, with data from the lowest, middle and highest bins (see Figure caption for details). Correlations computed from O-B statistics (a–c) and the Desroziers diagnostic (d–f) are shown, along with the difference between them (g–i). Results from S-NPP ATMS, and for data over land, are not shown as they look similar, albeit with larger error correlations for the S-NPP temperature channels due to hardware differences. The largest inter-channel error correlations are generally found between the humidity-sensitive channels (the top right of each panel), though for the cloudiest bin there are also strong correlations between channels 6 and 7 (the bottom left of panel c). Notable correlations can also be found between the lowest two temperature channels (6–7) and the humidity channels (18–22). As the cloudiness increases from the low bin to the high bin, the correlations increase. When comparing the correlations derived using the O-B and Desroziers methods (panels g–i), three main differences are notable. First, the correlations between the humidity channels are reduced when using the Desroziers method, with the reduction being larger for less-cloudy conditions. Second, the correlations between the lowest two temperature channels and the humidity channels are reduced, again with a larger reduction in less-cloudy conditions. Third, the off-diagonal correlations between adjacent temperature channels are reduced. This reduction is similar in all bins, so appears not to be related to the cloudiness of the observations. Instead it may come from the fact that the Desroziers diagnostic aims to remove the contribution from background errors in the

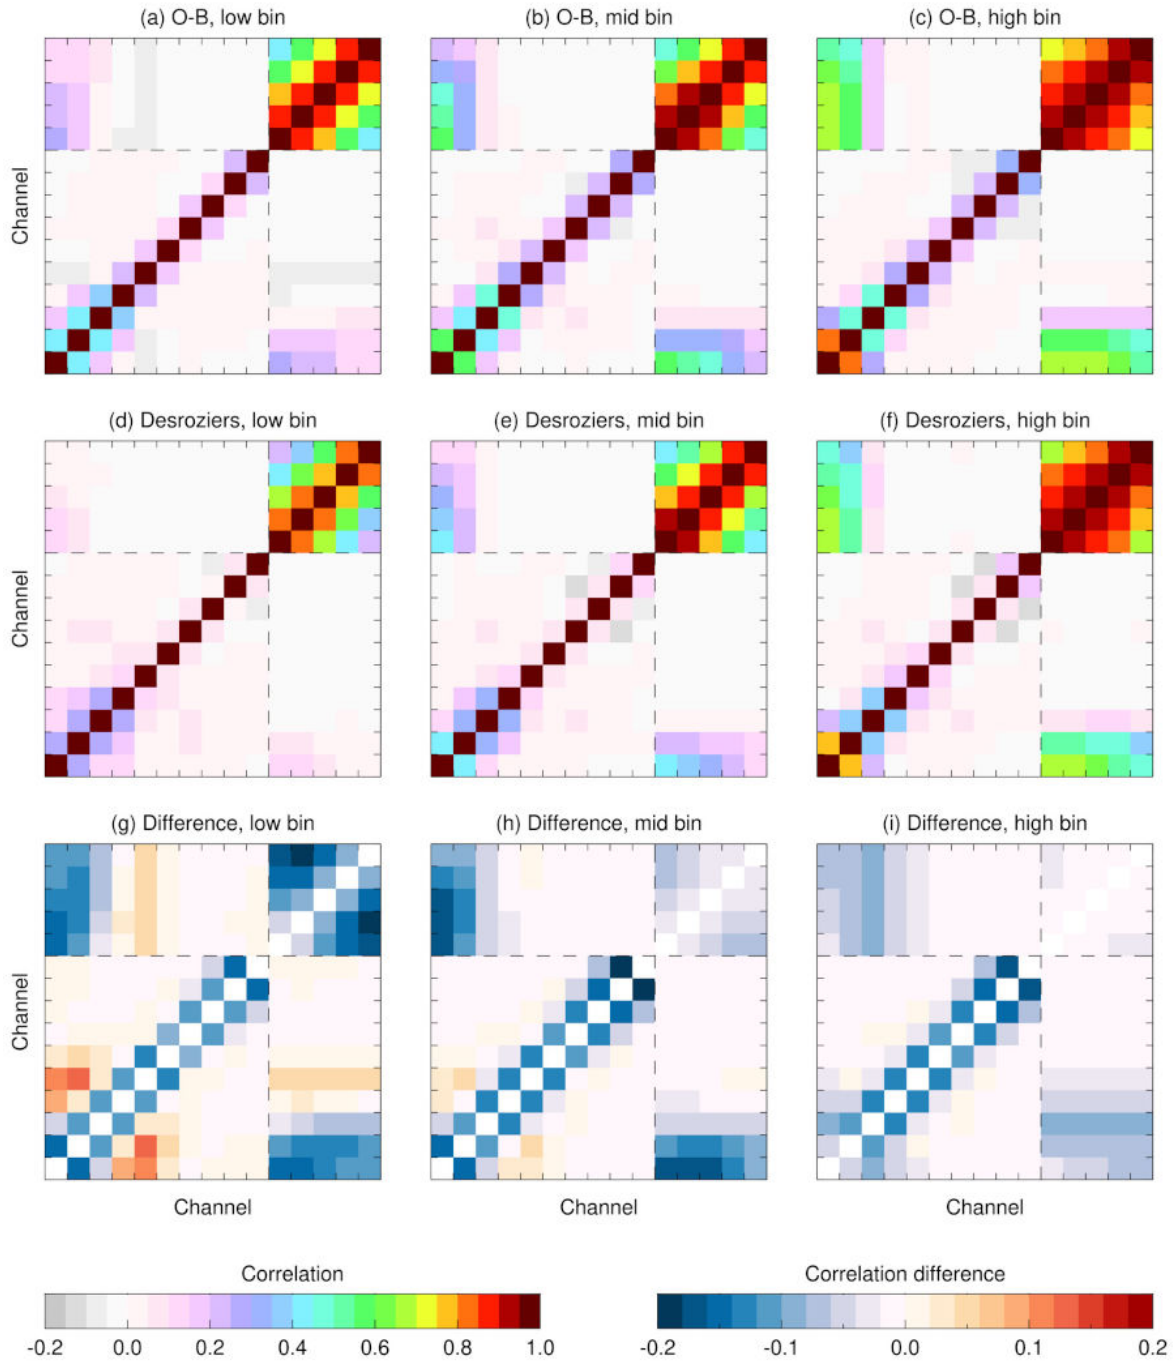


Figure 8: Correlation matrices from NOAA-20 ATMS data over the ocean, created using (a–c) O–B statistics, and (d–f) the Desroziers diagnostic, along with (g–i) the Desroziers – (O–B) correlation difference. Results are shown for the lowest bin ($LWP < 0.2 \text{ kg m}^{-2}$, $SI < 5 \text{ K}$), middle bin ($LWP = 0.2\text{--}0.4 \text{ kg m}^{-2}$, $SI = 5\text{--}20 \text{ K}$) and highest bin ($LWP > 0.4 \text{ kg m}^{-2}$, $SI > 20 \text{ K}$). Data are from a summer 2022 experiment, and the dashed lines separate the temperature- and humidity-sensitive channels.

observation departures, and adjacent temperature channels (which are sensitive to similar regions of the atmosphere) likely have larger correlated background errors.

5.3 Reconditioning

As found in previous studies, conditioning issues associated with \mathbf{R} can increase the number of iterations required for the minimisation to converge to a solution. To solve this problem, reconditioning of the covariance matrices can be performed, which is usually achieved by one of two methods; the minimum eigenvalue method or the ridge regression method (see [Tabcart *et al.*, 2020](#)). For both these methods, the required condition number of the covariance matrix, N_{req} , is specified, where the condition number, N , is the ratio of the maximum and minimum eigenvalues: $N = \lambda_{\text{max}}/\lambda_{\text{min}}$. In simple terms, the goal of reconditioning is to make the smallest and largest eigenvalues of the covariance matrix closer together. However, it is more desirable to alter the smaller eigenvalues, as altering the larger eigenvalues can lead to more significant alterations of the covariance matrix.

In the minimum eigenvalue method, a threshold eigenvalue is calculated via $\lambda_{\text{thresh}} = \lambda_{\text{max}}/N_{\text{req}}$, and any eigenvalue smaller than this is set to the threshold value. The covariance matrix is then reconstructed via reverse eigendecomposition using the original eigenvectors and new eigenvalues. In the ridge regression method, an increment is calculated via $\delta = (\lambda_{\text{max}} - \lambda_{\text{min}}N_{\text{req}})/(N_{\text{req}} - 1)$, and added to all diagonal elements of the covariance matrix, giving $\mathbf{R}_{\text{recond}} = \mathbf{R} + \delta\mathbf{I}$. The reconditioned correlation matrix can then be obtained from the reconditioned covariance matrix. [Campbell *et al.* \(2017\)](#) found the sensitivity to the choice of reconditioning method was small for their experiments using ATMS and IASI data, while [Weston *et al.* \(2014\)](#) used the ridge regression method for their IASI assimilation as it took fewer iterations to converge and the covariance matrix gave more realistic errors. [Bormann *et al.* \(2016\)](#) applied the minimum eigenvalue method in their IASI clear-sky assimilation work, but found comparatively little benefit in terms of the number of iterations or forecast improvement. [Geer \(2019\)](#) performed tests using the minimum eigenvalue method for the assimilation of IASI water vapour channels (but reconditioned the correlation matrix rather than the covariance matrix) and found that reconditioning improved the fits to ATMS data. The threshold eigenvalues used by [Geer \(2019\)](#) were directly specified (with values of $\lambda_{\text{thresh}} = 0.37$ and 1 tested), rather than calculated from a specified condition number.

The condition numbers of the covariance matrices range from ~ 3000 for the least cloudy bin $\sim 165,000$ for the cloudiest bin. As there is no clear consensus as to what the condition number should be, we decided to use a value of $N_{\text{req}} = 100$ for our experiments. [Weston *et al.* \(2014\)](#) used a value of $N_{\text{req}} = 67$ for their assimilation of IASI, which used 140 channels in clear-sky conditions and had an original condition number of ~ 2000 . We investigated both methods for reconditioning the covariance matrices, but little difference was seen in the resulting correlation matrices, possibly due to the large reduction in condition number. The main difference is that when using the ridge regression method the correlations between adjacent temperature channels are larger in the cloudier bins, but still small at < 0.1 . For this report we decided to use the minimum eigenvalue method.

The reconditioned covariance matrices are used to derive correlation matrices, and in 4D-Var these are combined with the σ^0 values from the error model to create covariance matrices. As such, the condition numbers of these ‘recreated’ covariance matrices will not equal 100, due to the different σ^0 values. Figure 9 shows the range of condition numbers in each of the nine bins, before and after reconditioning, for matrices derived using O–B statistics. (Results are almost identical when using the Desroziers diagnostic.) In general, the condition numbers of the ‘recreated’ covariance matrices after reconditioning are a factor of 3 smaller than the non-reconditioned matrices, and display more variation for the cloudier bins.




Figure 9: Box and whiskers plot showing the minimum, lower quartile, median, upper quartile and maximum condition numbers in each of the nine bins before (green) and after (purple) reconditioning. Data are from 5 days of a summer 2022 experiment using matrices derived with O–B statistics.

Figure 10 shows the differences between the correlation matrices derived from the original and reconditioned covariance matrices. The covariance matrices were calculated from O–B statistics, and results are shown for the low, mid and high bins. Apart from in the cloudiest bin, there is little difference in the correlations between the humidity channels before and after reconditioning. However, the correlations between the lowest-peaking temperature channels and the humidity channels are greatly reduced, especially in the cloudiest bin. Similarly, the off-diagonal correlations, representing correlations between adjacent channels, are reduced. So, while the strongest correlations between the humidity channels remain relatively unaltered, there is an overall reduction in correlations involving the temperature channels.

Figure 11 shows the impact of covariance matrix reconditioning on the eigenvalues of the correlation matrices over ocean and land, for clear and cloudy conditions (using the low and high bins). What is most obvious is that the reconditioning results in a large number of the eigenvalues being set close to 1, and that there is a larger increase in the trailing eigenvalues in the cloudiest conditions. Indeed, in the cloudiest bin the eigenvalues are practically identical in the reconditioned O–B and Desroziers matrices. In the clearest bin the reconditioned Desroziers correlation matrix has slightly larger trailing eigenvalues compared to the reconditioned O–B correlation matrix. However, even after reconditioning the trailing eigenvalues are still relatively small, at around 0.03 in the cloudiest bin and 0.06 in the clearest bin. These small values have the potential to impact the forecast through the amplification of vertical oscillations (see the next section for more details, and also Geer, 2019).

5.4 Eigenvectors and eigenjacobians

As discussed in Section 4.2, the departures associated with each channel are replaced by eigendepartures associated with each eigenvector in the amended formulation when including inter-channel correlated errors. To understand how these eigendepartures differ from normal background departures, Figure 12 shows all 15 eigenvectors from the O–B correlation matrices over the ocean from the low, mid and high bins discussed earlier. The eigenvectors are ordered from largest to smallest eigenvalue. For the

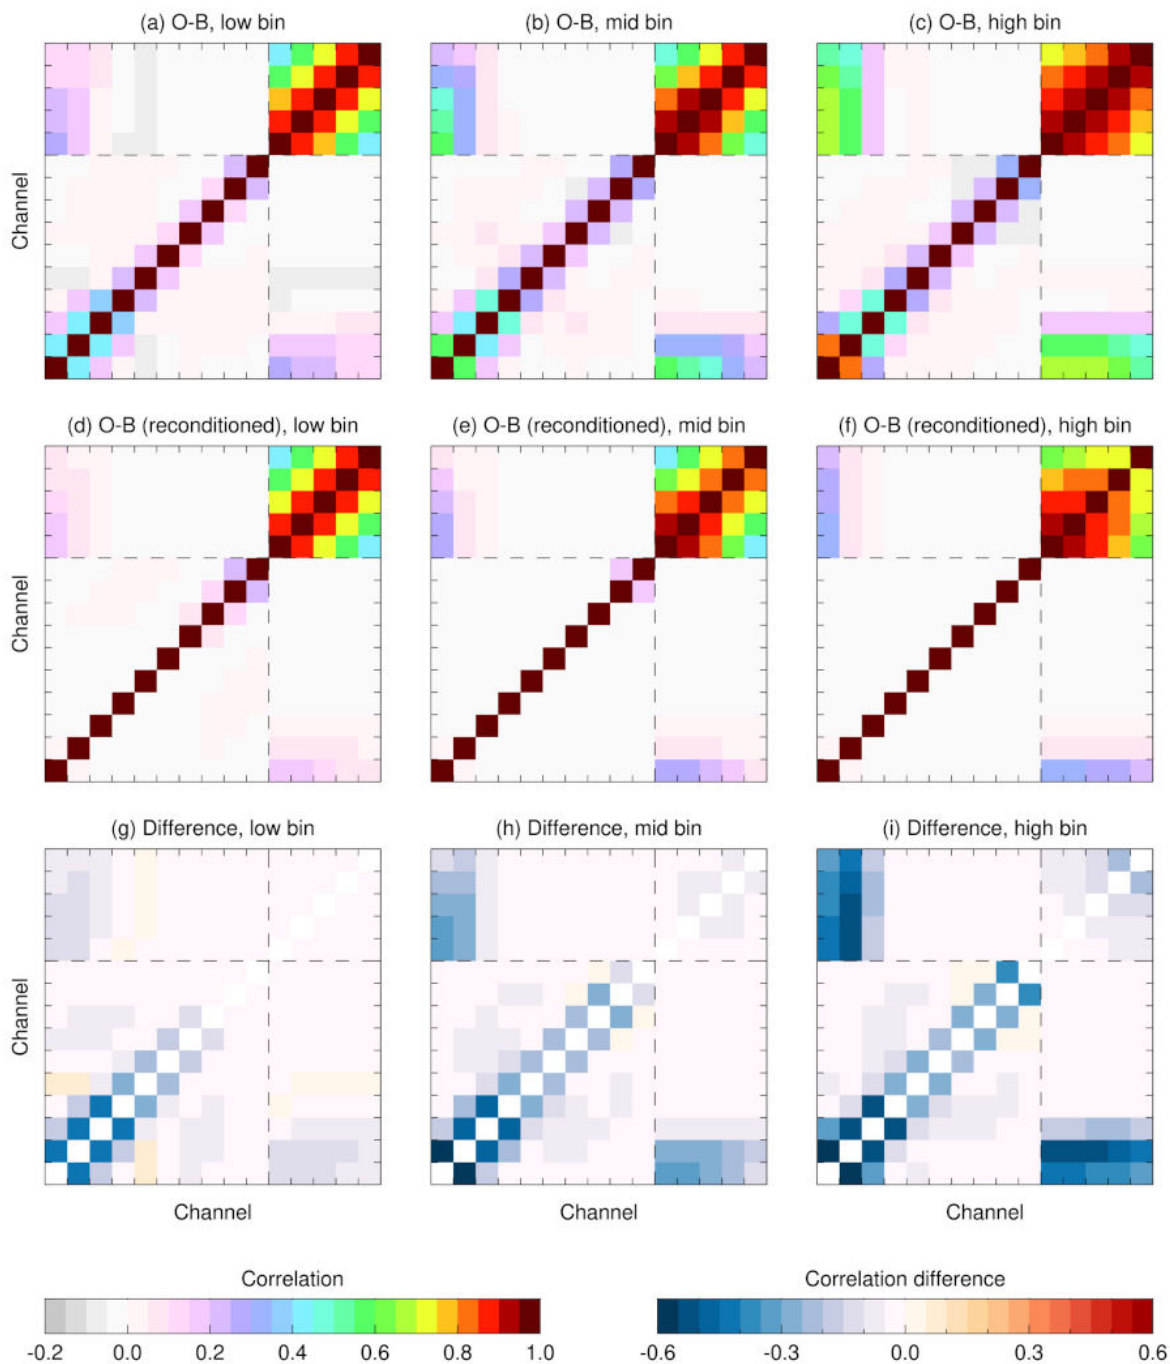


Figure 10: As Figure 8, but comparing correlation matrices derived with O-B statistics before and after the covariance matrices were reconditioned to a condition number of $N_{req} = 100$. Results are similar for matrices derived using the Desroziers diagnostic, and using data over land.

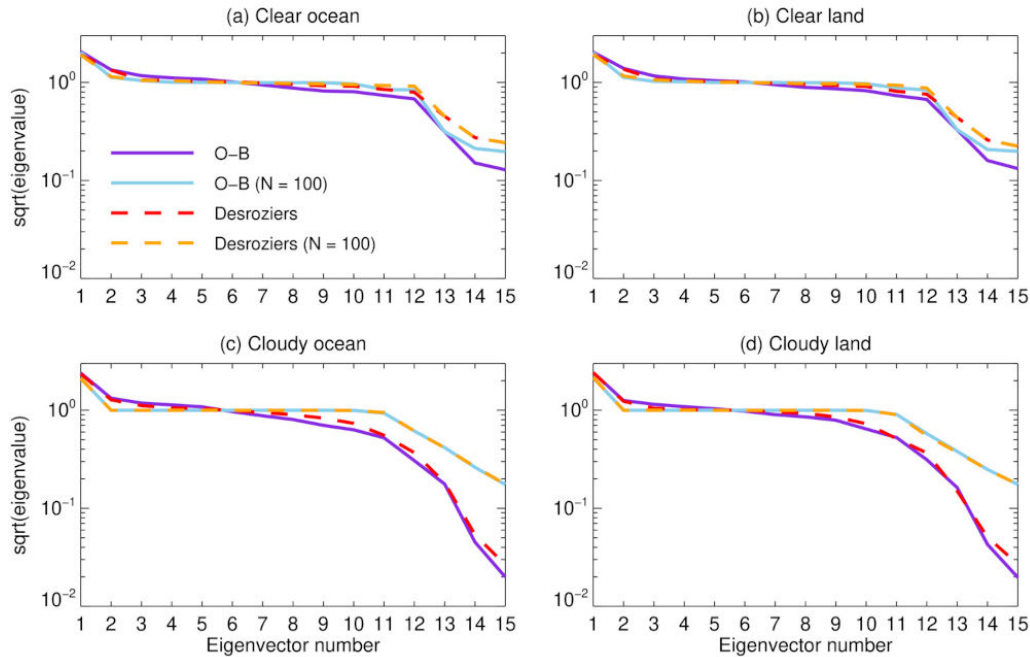


Figure 11: Square root of the eigenvalues of the error correlation matrix, as a function of eigenvector number, for the clearest and cloudiest bins (see Table 3) over ocean and land. Results are shown from correlation matrices derived using the O–B and Desroziers methods, before and after the covariance matrices were reconditioned to $N_{req} = 100$.

leading eigenvector (Figure 12a) there are roughly equal contributions to the eigendeparture from all five humidity channels (18–22), with slightly smaller contributions from the lower-peaking temperature channels (6–8). Thus, the eigendeparture associated with eigenvector 1 is comprised of background departures from the lower troposphere. When moving to eigenvector 2 (Figure 12b), there is now little contribution from the humidity channels, and it is mostly the temperature channels, particularly channels 8–11, which contribute to the eigendeparture. Thus this eigendeparture is mostly sensitive to background departures in the upper troposphere and lower stratosphere. Further eigenvectors begin to have more complicated combinations of channels, and greater differences between the structure for the three bins. For example, in eigenvector 3 the low bin (clearest conditions) has its largest contribution from channel 10, while the mid and high bins (more cloudy conditions) have their largest contributions from channel 14, which peaks higher in the stratosphere. For the three trailing eigenvectors (Figure 12m–o), there is little-to-no contribution from the temperature channels, with instead the contributions coming from increasing harmonics of the humidity channels.

Figure 13 compares a selection of six eigenvectors from the cloudiest bins over the ocean of four different correlation matrices: those derived from O–B and Desroziers statistics, before and after reconditioning. For eigenvector 1 there is little difference in the structure between the different correlation matrices, except that the reconditioned matrices have smaller contributions from the lowest three temperature channels. For the remaining eigenvectors the two non-reconditioned matrices have similar structures to each other, as do the two reconditioned matrices. However, there are differences in which channels the eigenvectors are sensitive to. For example, in eigenvector 2 the non-reconditioned matrices have most contributions from the lower-peaking temperature channels, while the reconditioned matrices have most contributions from the higher-peaking temperature channels. In eigenvector 9 the non-reconditioned matrices have most contributions from channels 10–15, while the reconditioned matrices are almost entirely using channel 8.

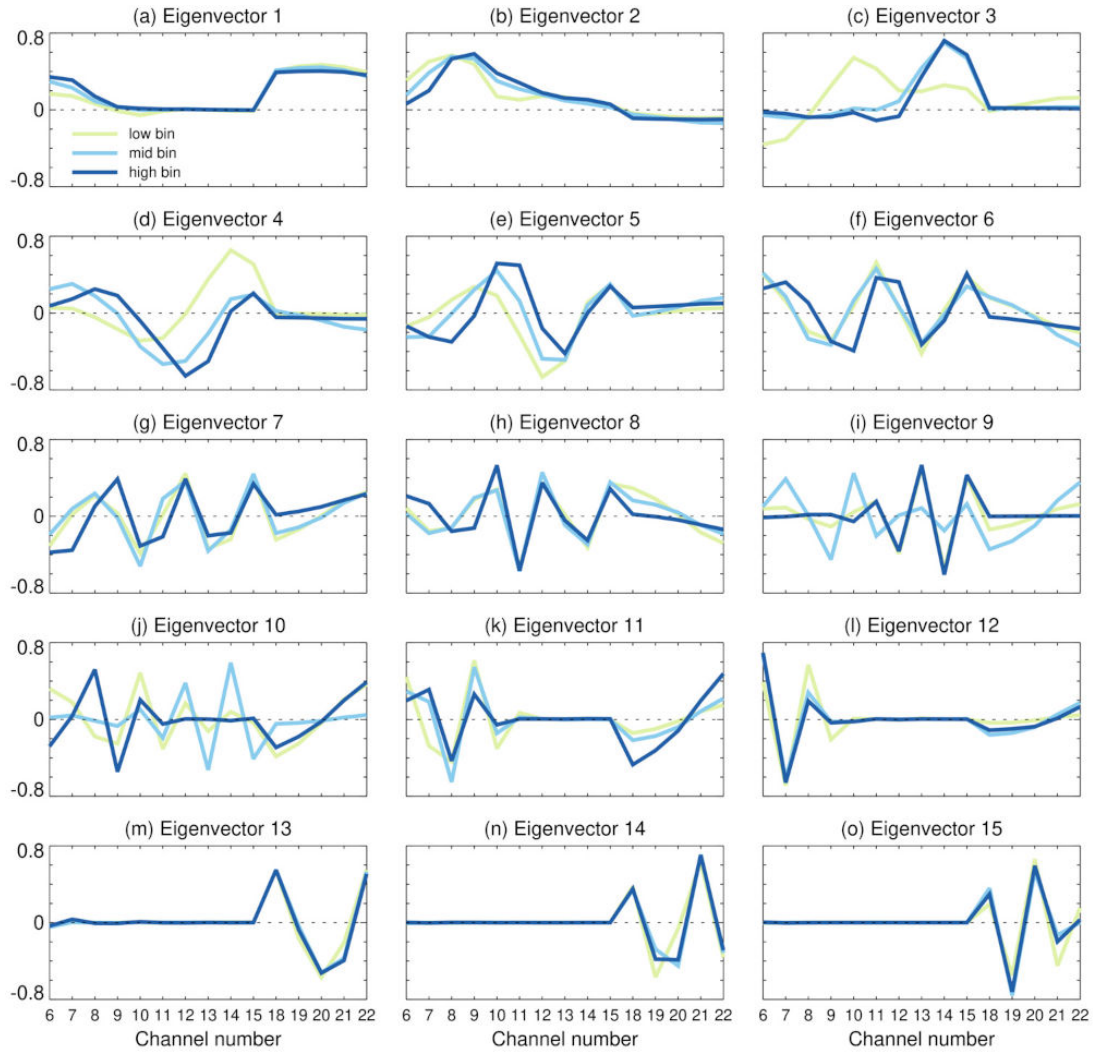


Figure 12: Eigenvectors from the three correlation matrices derived from O–B statistics shown in Figure 8.

Another useful way to think about the impact of using inter-channel error correlations is by studying the eigenjacobians associated with each eigenvector. From Equation 14, which gives the calculation of J_k^o for each observation, the gradient can be computed as

$$J_k^{o'} = \mathbf{H}^T \mathbf{e}_k \left(\frac{\mathbf{e}_k^T \mathbf{n}}{\lambda_k} \right). \quad (17)$$

Here, $\mathbf{H}^T \mathbf{e}_k$ is the eigenjacobian for each normalised eigendeparture $\mathbf{e}_k^T \mathbf{n}$. The sensitivities of J_k^o for each eigendeparture are given by $\mathbf{H}^T \mathbf{e}_k / \lambda_k^{0.5}$.

Figure 14(a,b) shows the temperature and humidity eigenjacobians from the cloudiest bin of a correlation matrix derived using O–B diagnostics for a standard mid-latitude reference profile. While the Jacobians shown in Figure 1 have peaks that progressively move higher in the atmosphere for subsequent channels, the eigenjacobian structure is more complicated. For the leading eigenvector, which as discussed earlier has contributions from all five humidity channels and the lower-peaking temperature channels, the temperature and humidity eigenjacobians have strong sensitivity to the troposphere. For eigenvector 2, there is some sensitivity of the temperature eigenjacobian to the troposphere, but most of the contribution comes from the stratosphere. As the eigenvector number increases, the eigenjacobians capture

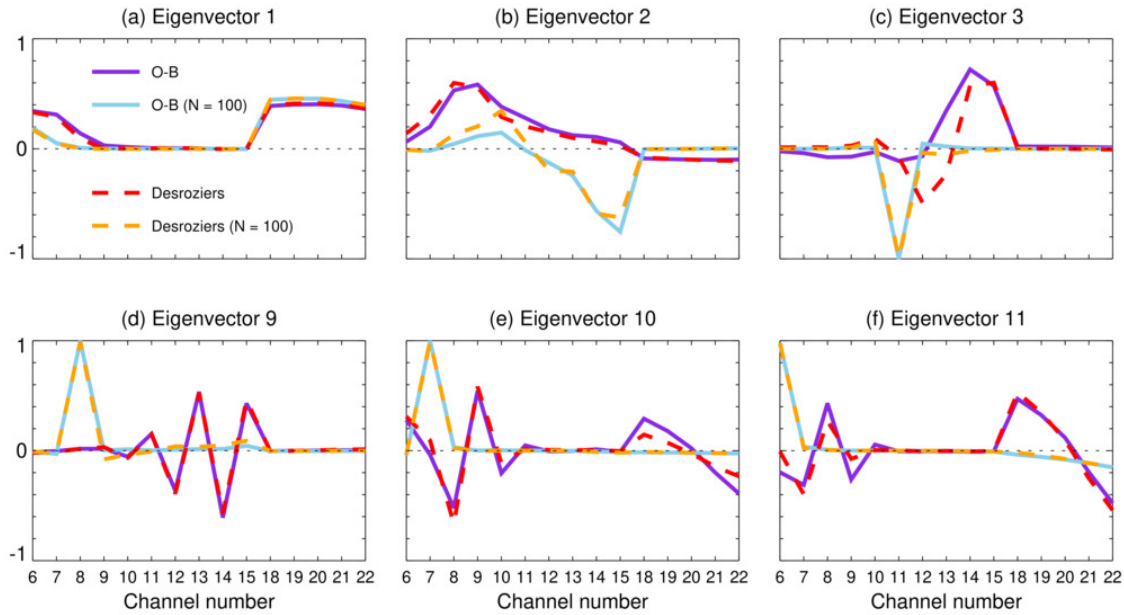


Figure 13: Comparison of six eigenvectors from the highest (cloudiest) bins of correlation matrices derived from O-B and Desroziers statistics, before and after reconditioning.

higher frequency vertical oscillations, though some have small amplitudes resulting in little sensitivity. As noted by Geer (2019) when studying eigenjacobians from IASI water vapour channels, the higher-order harmonic combinations of channels display strong sensitivity to vertically-oscillating structures in the atmospheric temperature profile. It is worth noting that these higher-order harmonic combinations of channels will be affected by inter-channel biases which may be poorly controlled.

Figure 14(c,d) shows the temperature and humidity sensitivities of J^0 for each eigendeparture, which is achieved by dividing each eigenjacobian by the square root of its eigenvalue. As the leading eigenvectors have eigenvalues > 1 (see Figure 11), this division results in a down-weighting of the broad vertical structures. Conversely, as the trailing eigenvectors can have eigenvalues $\ll 1$, the division by $\lambda_k^{0.5}$ amplifies the small-scale vertical oscillations. This is most noticeable for the temperature sensitivities shown in Figure 14c. Such an amplification might not be desirable if these oscillations are not capturing real features of the atmosphere that can be represented in the IFS or if the departures arise due to poorly controlled inter-channel biases. Or, if the statistics are reliable, the amplification may help to add vertical resolution if the error correlations are correctly taken into account.

Figure 14(e,f) again shows the temperature and humidity sensitivities of J^0 , but now for the reconditioned O-B matrix. In terms of the temperature sensitivity, the structure for the leading eigenvector remains the same as in the non-reconditioned matrix, and there are also some smaller structural differences for the subsequent eigenvectors. However, the biggest difference is the reduction in the amplification of the trailing eigenvectors, which may be desirable. A similar result can also be seen for the humidity sensitivity, although due to the reconditioning there is essentially no sensitivity to eigenvectors 2–8.

As shown in this section, the structure of the correlation matrix (which differs depending on the diagnostics used to create it and the ‘cloudiness’ it represents) affects the structure of the eigenvectors, and hence the eigendepartures used by VarQC. It also affects the vertical structure of the J^0 sensitivity, which can potentially negatively impact the analysis through the amplification of trailing eigenvectors. Thus, understanding the structures of the eigenvectors and eigenjacobians can help to interpret the analysis, and suggest pathways for further improvements.

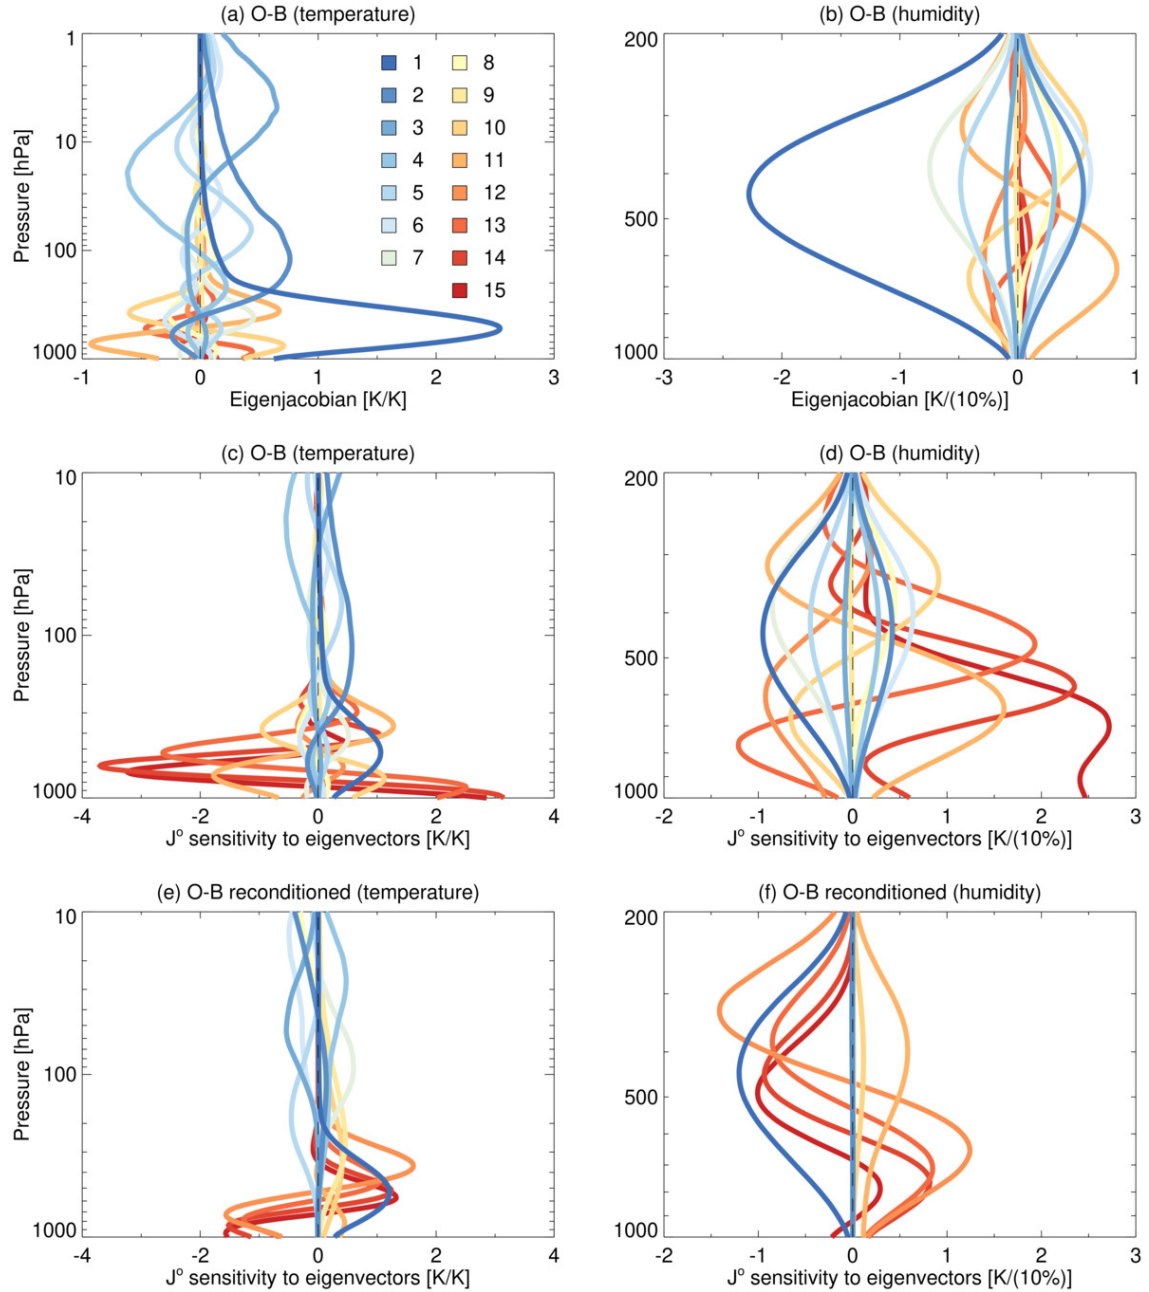


Figure 14: Temperature and humidity eigenjacobians from the cloudiest bin of a NOAA-20 ATMS correlation matrix derived using O–B diagnostics for a standard mid-latitude reference profile (a,b). Eigenvectors 1–14 are ordered from largest to smallest eigenvalue. Also shown is the sensitivity of the J^0 cost function by eigenvector from correlation matrices derived using O–B diagnostics before reconditioning (c,d) and after reconditioning (e,f). This highlights the role the eigenvalues play in amplifying sensitivities to smaller vertical wavelengths in the trailing eigenvectors, and reducing sensitivities to broad vertical features in the leading eigenvectors. Can be compared with the Jacobians shown in Figure 1.

6 Experimental setup and results

6.1 Experimental setup

We will now proceed with presenting results from assimilation experiments with ATMS in all-sky assimilation, investigating the impact of accounting for inter-channel error correlations through the methods described earlier. The setup for all the experiments discussed here is as described in Section 5.2, with only two main differences.

Firstly, the amended formulation for calculating $\tilde{\mathbf{R}}^{-1} \mathbf{d}$ is used, as described in Section 4.2, as inter-channel error correlations are now being included. The all-sky error model detailed in Section 3.1 is used to obtain the observation errors, rather than the values obtained while creating the correlation matrices, which is similar to the approach used by Campbell *et al.* (2017). Thus, while the correlation matrices differ between experiments, the errors assigned to each observation will be the same. Figure 15 shows comparisons between the errors derived when creating the different correlation matrices (σ_{matrix} , lines) and those from the all-sky error model (σ_{model} , circles). Note that σ_{matrix} is obtained by taking the square root of the diagonal elements of the covariance matrix. In general, $\sigma_{\text{model}} > \sigma_{\text{matrix}}$, particularly for the humidity-sensitive channels. However, for the reconditioned matrices as the cloud amount increases (panels b, c, e, f) $\sigma_{\text{matrix}} > \sigma_{\text{model}}$ for the temperature-sensitive channels, and shows little variation between channels. The reason for the increase in the error values is because reconditioning increases the magnitude of the diagonal elements of the covariance matrix.

Secondly, as we are using eigendecomposition, the prior probability of gross error (PPGE) values that have been specified for ATMS (and given in Table 1) are no longer valid. This is because these values are specified for the departures in each channel, and now we are concerned with eigendepartures for each eigenvector. The eigenvector structure is complicated (as seen from Figures 12 and 13), and can differ depending on the cloudiness of the bin and on the statistics used to derive the correlation matrix. This in turn means the eigendepartures can also be different depending on the correlation matrix, which makes specifying the PPGE difficult. One way to think about it is to consider which channels are being sampled by each eigenvector, what their PPGE is, and proportionally how much each channel contributes to the overall eigendeparture. This can be determined by taking the absolute value of an eigenvector, and normalising it so that its sum is equal to one. This can then be multiplied by the PPGE, and the sum gives a new PPGE for each eigenvector. This procedure was performed for all nine correlation matrices, and the range of results for each eigenvector are shown in Figure 16. The resulting PPGE values are largest for eigenvectors 1 and 13–15. This is because these eigenvectors sample the lowest-peaking temperature channels and the humidity channels (see Figure 12), which have the largest pre-existing values of PPGE (see Table 1). Eigenvectors 2–12 have generally lower PPGE values, ranging from < 0.1 up to almost 0.4, depending on the correlation matrix. This is because these eigenvectors sample the humidity channels less, with more weight instead given to stratospheric channels.

Another difficulty faced when specifying PPGE values for each eigenvector is that not all 15 ATMS channels are always assimilated. Due to various screening parameters, the number of channels assimilated varies from 5 to 15, as can be seen in Figure 17. The fewest assimilated channels occurs over the poles and over regions of high topography, while the most assimilated channels occurs over the ocean. When channels are missing, it is typically the lowest-peaking channels (6–9 and 18–22), with channels peaking in the stratosphere mostly present in all locations. When channels are missing the correlation matrix changes, and hence the eigenvectors and eigendepartures change. As such, the PPGE would also require changing. Faced with these difficulties, we made the pragmatic choice of using a $A = 0.25$ for all eigenvectors, based on Figure 16.

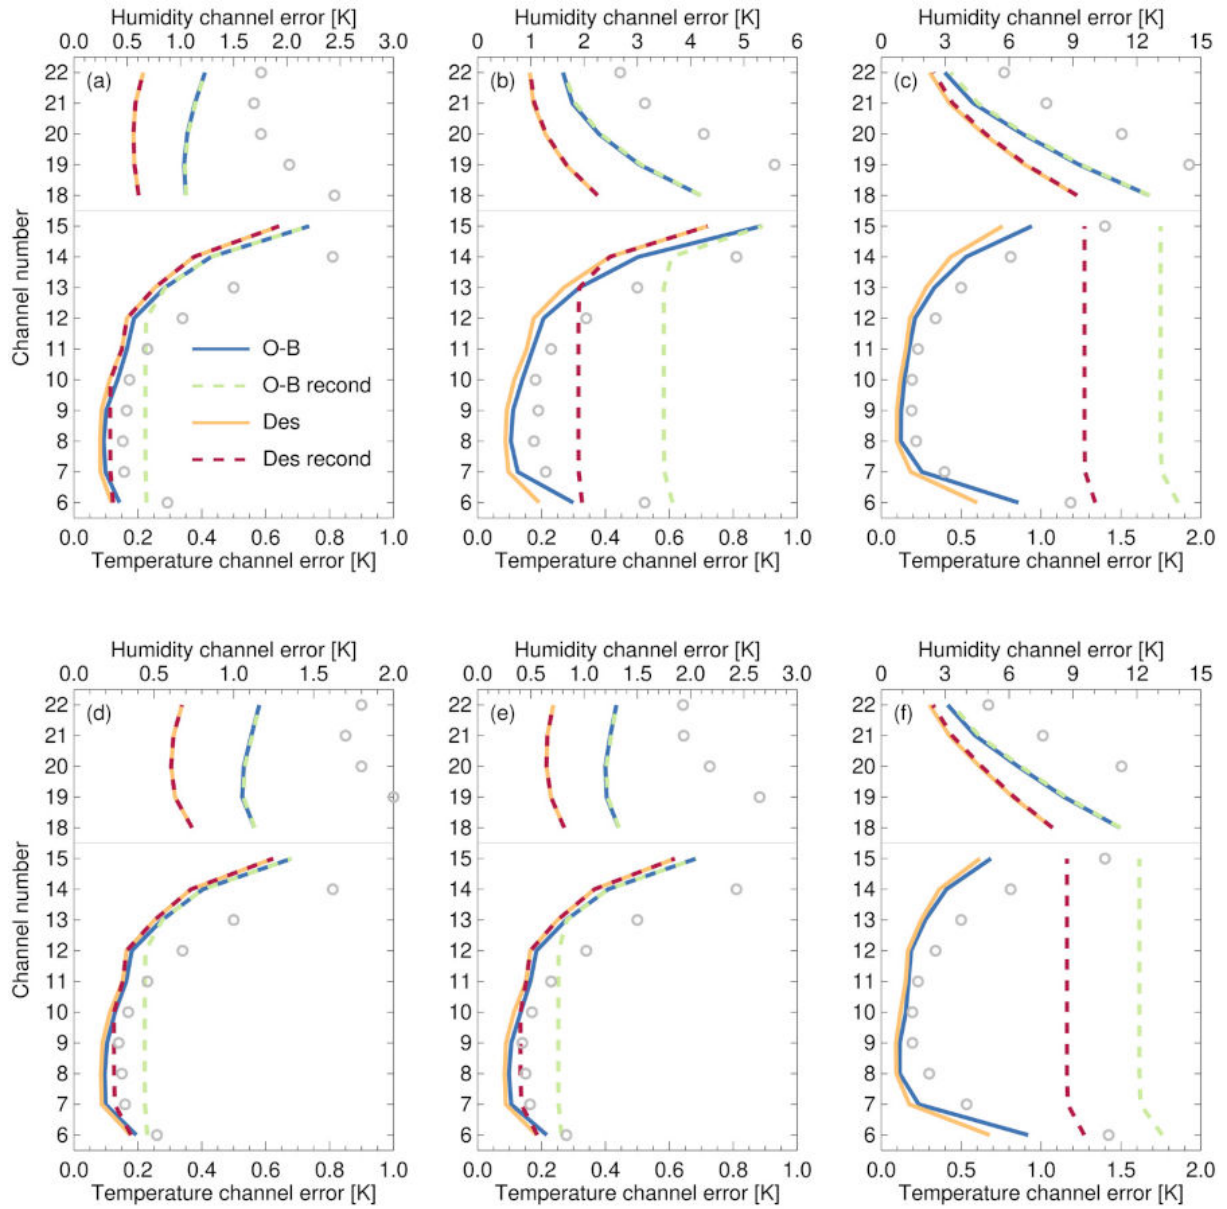


Figure 15: Errors derived from different covariance matrices (lines) compared to those from the all-sky error model (circles) for NOAA-20 ATMS. Results are shown for data over the ocean (a–c) and land (d–f). The left, middle and right columns correspond to results from the low bin ($LWP < 0.2 \text{ kg m}^{-2}$, $SI < 5 \text{ K}$), mid bin ($LWP = 0.2\text{--}0.4 \text{ kg m}^{-2}$, $SI = 5\text{--}20 \text{ K}$) and high bin ($LWP > 0.4 \text{ kg m}^{-2}$, $SI > 20 \text{ K}$) respectively, where the low bin is the least cloudy and the high bin is the most cloudy.




Figure 16: The variation of the prior probability of gross error, A , for each eigenvector, calculated from the set of correlation matrices used in the experiments. The value of A was determined by multiplying the absolute value of each normalised eigenvector by the prior probabilities currently in use at ECMWF (see text for details).




Figure 17: Number of NOAA-20 ATMS channels assimilated on 15 June 2022.

6.2 All-sky assimilation with and without inter-channel error correlations

We begin our examination of the results by looking at how the all-sky assimilation of ATMS compares to clear-sky assimilation (as is currently performed in operations), and whether including inter-channel error correlations derived from O–B statistics gives any improvement the forecast. Figure 18 shows the standard deviations of background departures from an experiment with ATMS data assimilated in clear-sky conditions, as well as from experiments assimilating ATMS in all-sky conditions with and without inter-channel error correlations accounted for. The standard deviations, σ , are normalised by those from a control experiment which excludes ATMS data, and are expressed as a percentage (i.e. $100 \cdot \sigma / \sigma_{\text{control}}$). These statistics were compiled from each of the individual 12 hour forecasts from the three-month periods over summer 2022 and winter 2022/2023. Values $< 100\%$ represent an improvement to the short-range forecast compared to the control, while values $> 100\%$ represent a forecast degradation. Such statistics are calculated against all assimilated instruments, but Figure 18 shows comparisons against six instruments, representing a range of different data types.

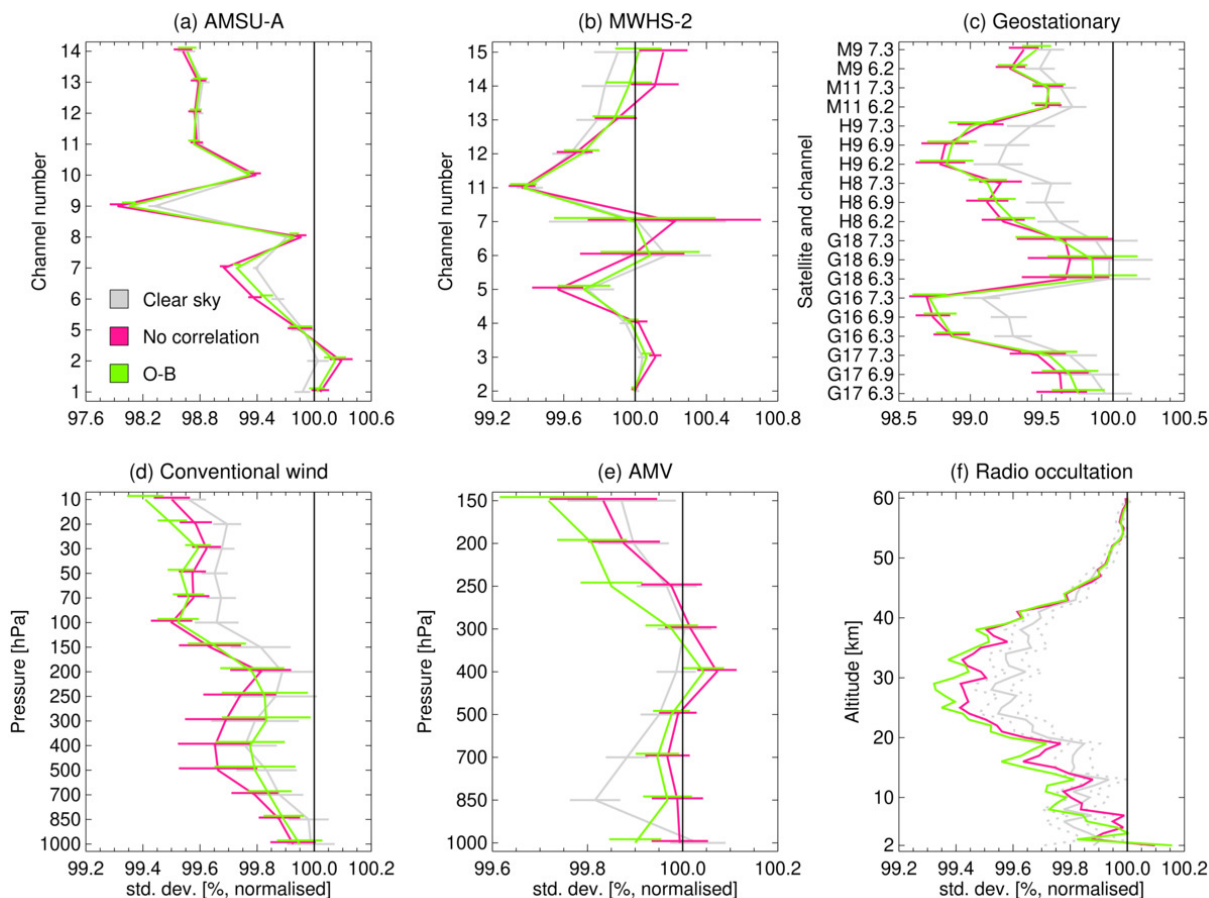


Figure 18: Standard deviations of background departures from experiments assimilating ATMS as clear sky (grey line) and as all-sky without inter-channel error correlations (pink line), and with correlations derived from O–B statistics (green line). Results are normalised by a control experiment with no ATMS data assimilated (black line at 100%). Each panel shows comparisons against a different observation type, with own analysis verification, using six months of data. The 95% confidence range is represented by (a–e) horizontal red lines, and (f) dashed lines for clarity, due to the large number of altitudes. In the y-axis for panel (c), M, H and G stand for Meteosat, Himawari and GOES satellites, with the numbers representing the channel frequency in microns.

The first thing of note is that both of the all-sky experiments mostly replicate the positive impact achieved when adding the assimilation of clear-sky ATMS, with normalised standard deviations generally below 100%. When comparing against geostationary water vapour channels (Figure 18c), conventional wind observations (Figure 18d) and radio occultation data (Figure 18f) the all-sky assimilations give consistently better results than the clear-sky assimilation. However, results are more mixed for comparisons against AMSU-A temperature-sensitive channels (Figure 18a), MWHS-2 temperature and humidity-sensitive channels (Figure 18b) and atmospheric motion vectors (AMV's; Figure 18e). In each of these cases there are improvements and degradations, depending on the channel or pressure level. The different behaviour seen when comparing against humidity-sensitive observations may be related to the fact that the MWHS-2 observations are assimilated in all-sky conditions, while the geostationary observations are assimilated in clear-sky conditions. For wind observations, the different behaviour seen may be due to a combination of the geographic distribution of observations (with conventional wind observations being more limited spatially) and the different number of observations at different altitudes. For example, at 700 hPa there are ~6 million AMV observations and ~45 million conventional wind observations.

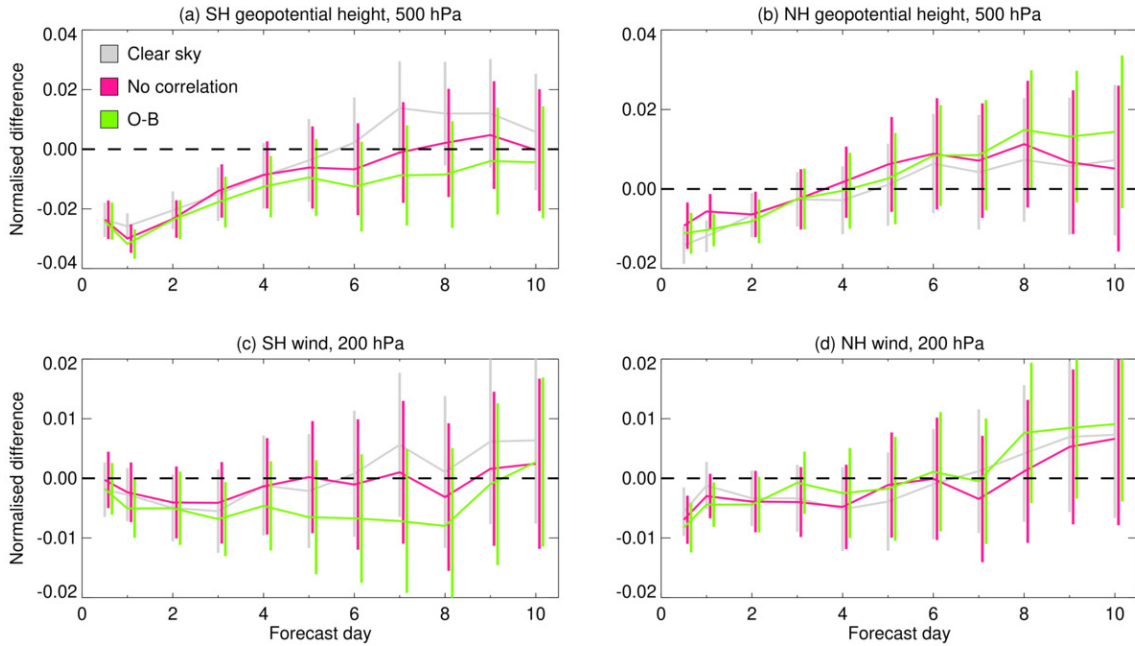


Figure 19: Normalised change in RMS error over 10 days for the same three experiments described in Figure 18, compared to the control with no ATMS data assimilated. Results are shown for both hemispheres separately, for (a,b) the 500 hPa geopotential height, and (c,d) the 200 hPa wind. Vertical lines show the 95% confidence range. Verification is against own analysis.

While the all-sky assimilations are generally giving a positive impact, is there any benefit to adding inter-channel error correlations? Comparing the experiments with and without error correlations (green and pink lines), the results are again mixed. For comparisons against AMVs and radio occultation data, including inter-channel error correlations gives slight short-range forecast improvements, while comparisons against other observation types suggest there is little difference when including the correlations. But what about the impact on the medium-range forecast? Figure 19 shows the normalised change in RMS error, ε , when comparing the clear-sky and all-sky assimilations against the control with no ATMS data assimilated, i.e. $(\varepsilon - \varepsilon_{\text{control}})/\varepsilon_{\text{control}}$. Results are shown for the 500 hPa geopotential height and the 200 hPa wind. In the southern hemisphere extratropics the results suggest that assimilating ATMS in all-sky conditions with inter-channel error correlations gives slightly better performance than assimilating in clear-sky or in all-sky without error correlations (Figure 19a,c), though none of the differences between the two experiments are statistically significant. In the northern hemisphere the results are more neutral, with no clear benefit being seen between any of the three experiments. The largest forecast improvements appearing in the southern hemisphere is something that has been noted before when adding microwave instruments with humidity channels (e.g. Lawrence *et al.*, 2018; Duncan and Bormann, 2020; Kan *et al.*, 2022; Steele *et al.*, 2023). This may be because the southern hemisphere has a lack of non-satellite observations, and so changes to the use of satellite data are more apparent here. To further highlight the slight improvement for wind forecasts when including inter-channel error correlations, Figure 20 shows zonal averages of the normalised difference in RMS wind error, but compared to a control assimilating ATMS data in all-sky conditions without error correlations. As can be seen, there are some regions of statistically-significant improvements of around 0.5–1%, and these extend to around day 8, but they do not cover a large area.

While the results presented so far are from the experiments using inter-channel error correlations derived from O–B statistics, experiments were also performed using correlation matrices derived from Desroziers

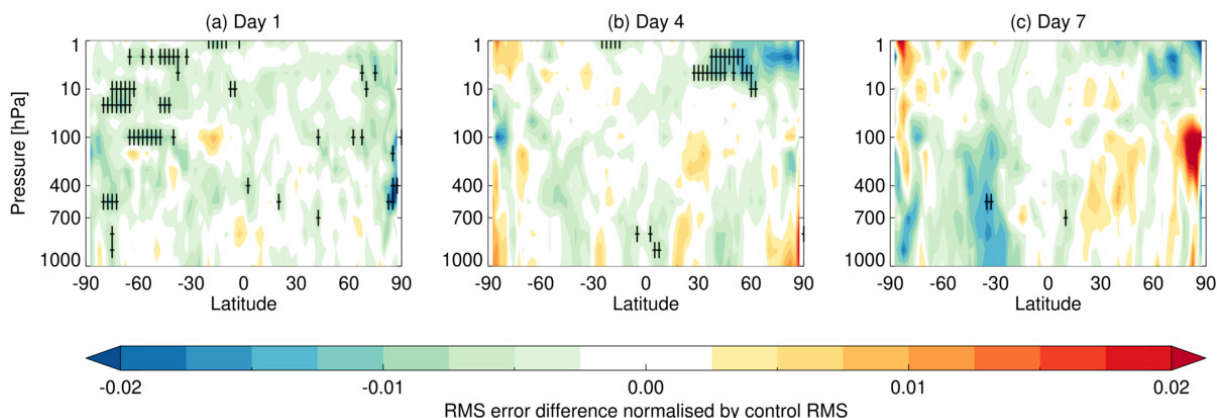


Figure 20: Zonal averages of the normalised difference in RMS wind error from an experiment assimilating ATMS in all-sky conditions with inter-channel error correlations from O–B statistics, compared to a control assimilating ATMS data in all-sky conditions without error correlations. Results are shown for forecast days 1, 4 and 7. Cross-hatching indicates the 95% confidence level and verification is against own analysis.

statistics. When looking at short-range forecasts there is little difference between the O–B and Desroziers correlation matrices, so the results are not shown. In terms of the medium range, Figure 21 shows the normalised change in RMS error when comparing the clear-sky and all-sky assimilations with different correlation matrices against a control with no ATMS data assimilated. In the southern hemisphere extra-tropics there is a signal that the correlation matrix derived from O–B statistics performs slightly better than that derived from Desroziers statistics, though the differences between the two are not statistically significant. Over the Northern Hemisphere, the result is neutral.

Considering the results presented here, there appear to be some benefits assimilating ATMS data in all-sky conditions, particularly in the southern hemisphere in the medium-range. When including inter-channel error correlations derived from O–B statistics there are mixed results for short-range forecasts, and slight improvements in the southern hemisphere in the medium range. However, the inclusion of inter-channel error correlations has less impact when assimilating ATMS in all-sky conditions than in clear-sky conditions. Here we only see slight improvements when comparing against other independent observations (Figure 18), while Weston and Bormann (2018) noted much more significant improvements in their clear-sky experiments. It is not clear why this should be the case. While the all-sky experiments use VarQC, the experiments with no inter-channel error correlations use PPGE values that have been extensively tuned, and are specified for each channel. Conversely, the experiments including inter-channel error correlations have a constant PPGE of 0.25. Perhaps tuning these values would allow for further forecast benefits.

There does not appear to be a benefit in using Desroziers statistics over O–B statistics to create the correlation matrices. This may be because we are using the observation error model rather than the errors obtained when creating the covariance matrices. If the errors diagnosed from the covariance matrices were used, there would likely be a larger difference, as the Desroziers errors are ~ 1.5 – 1.7 times smaller for the humidity channels (see Figure 15). Note however that in past studies the errors derived when using the Desroziers diagnostic have had to be inflated for successful assimilation (e.g. Weston *et al.*, 2014; Bormann *et al.*, 2016). As it is, the differences in the experiments performed here are due to the differences in the correlation matrices only. As seen in Figure 8, the Desroziers correlation matrices tend to have smaller correlations, with the differences compared to the O–B correlation matrices greatest in the least-cloudy bin. As the Desroziers statistics require background and analysis departures, they

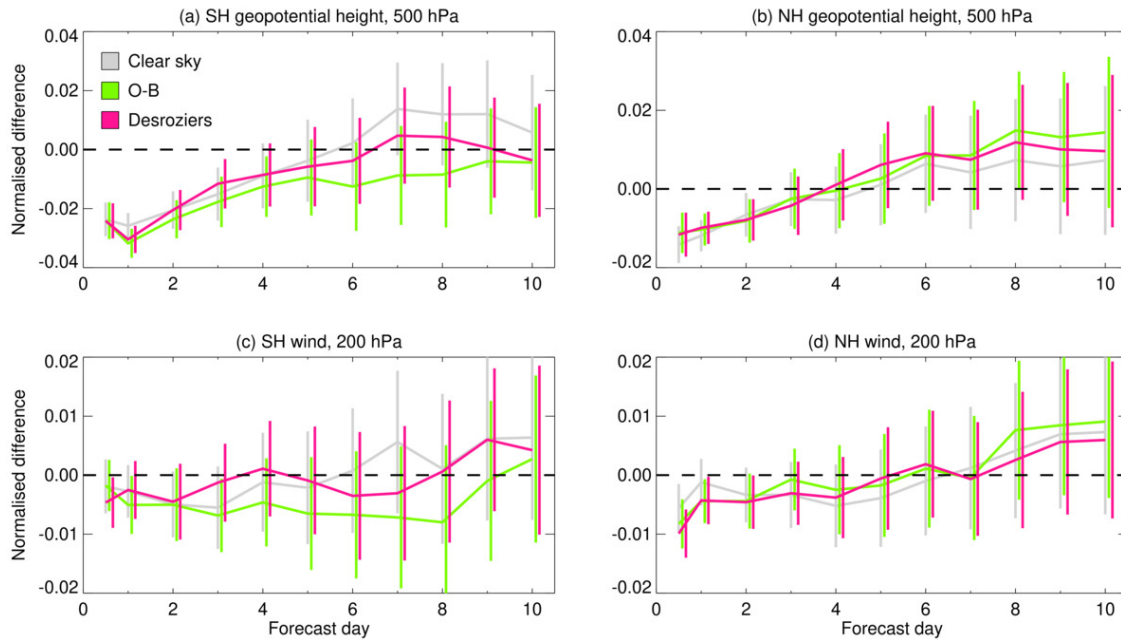


Figure 21: As Figure 19, but for experiments assimilating ATMS in clear-sky conditions (grey) and all-sky conditions with correlation matrices derived from O–B (green) and Desroziers (pink) statistics, compared to a control with no ATMS data assimilated.

require an experiment with the data actively assimilated. Thus, the correlation matrices derived from purely O–B statistics are simpler to produce and provide comparable results.

6.3 Effect of covariance matrix reconditioning

The goal of the experiments using the reconditioned correlation matrices is twofold: to see whether reconditioning helps speed up convergence in the minimisation, and whether the reduction of the correlations involving the temperature channels impacts the forecast. There is little difference between the results using the reconditioned O–B and Desroziers correlation matrices, so here we will just focus on the results using the reconditioned O–B statistics.

First we consider the effect of reconditioning on the 4D-Var convergence. Some previous studies using inter-channel error correlations have found that poor conditioning can substantially increase the number of iterations required for convergence (e.g. [Weston et al., 2014](#); [Campbell et al., 2017](#); [Geer, 2019](#)). In particular, when assimilating IASI data at the Met Office, [Weston et al. \(2014\)](#) noted that the number of iterations for convergence increased from 39 to 89, resulting in a large increase in the cost of the analysis. [Geer \(2019\)](#) found that using correlation matrices with adjustments to the trailing eigenvalues (a form of reconditioning) took ~ 30 iterations to converge, while using matrices without this adjustment took 34–35 iterations to converge. Our results are shown in Table 4. The lowest mean number of iterations is for the experiment with no ATMS data, while adding ATMS data in clear-sky conditions results in a slight increase to the number of iterations and condition number. When assimilating in all-sky conditions there is again a further slight increase in the number of iterations, with the experiment using the non-reconditioned O–B correlation matrix having the largest number. However, this is still only around 1 iteration more than when not assimilating ATMS data. This is similar to the results of [Bormann et al. \(2016\)](#), who also found only a small increase in the number of iterations.

| Experiment | Mean number of iterations | Median condition number |
|----------------------------------|---------------------------|-------------------------|
| No ATMS | 29.9 | 2670.2 |
| Clear-sky | 30.2 | 2726.0 |
| All-sky no correlation | 30.5 | 2771.5 |
| All-sky Desroziers reconditioned | 30.5 | 2799.3 |
| All-sky O-B reconditioned | 30.5 | 2860.9 |
| All-sky Desroziers | 30.5 | 2818.2 |
| All-sky O-B | 30.8 | 2785.0 |

Table 4: Statistics of the quality of the 4D-Var minimization in the experiments performed. Data are from the last inner-loop iteration, with the mean number of iterations and the median condition number given.

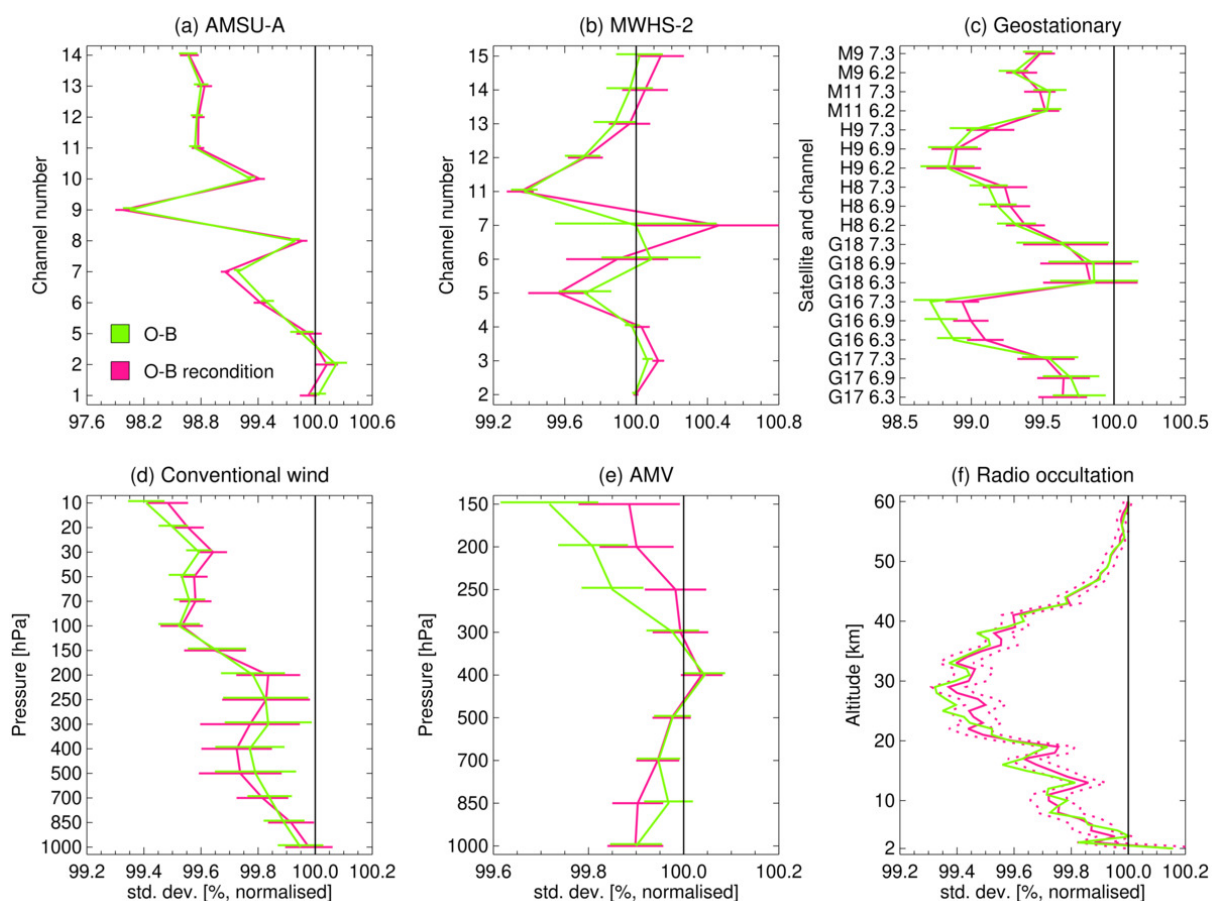


Figure 22: As Figure 18, but comparing experiments assimilating ATMS data in all-sky conditions using correlation matrices derived from O-B statistics before (green) and after (pink) reconditioning.

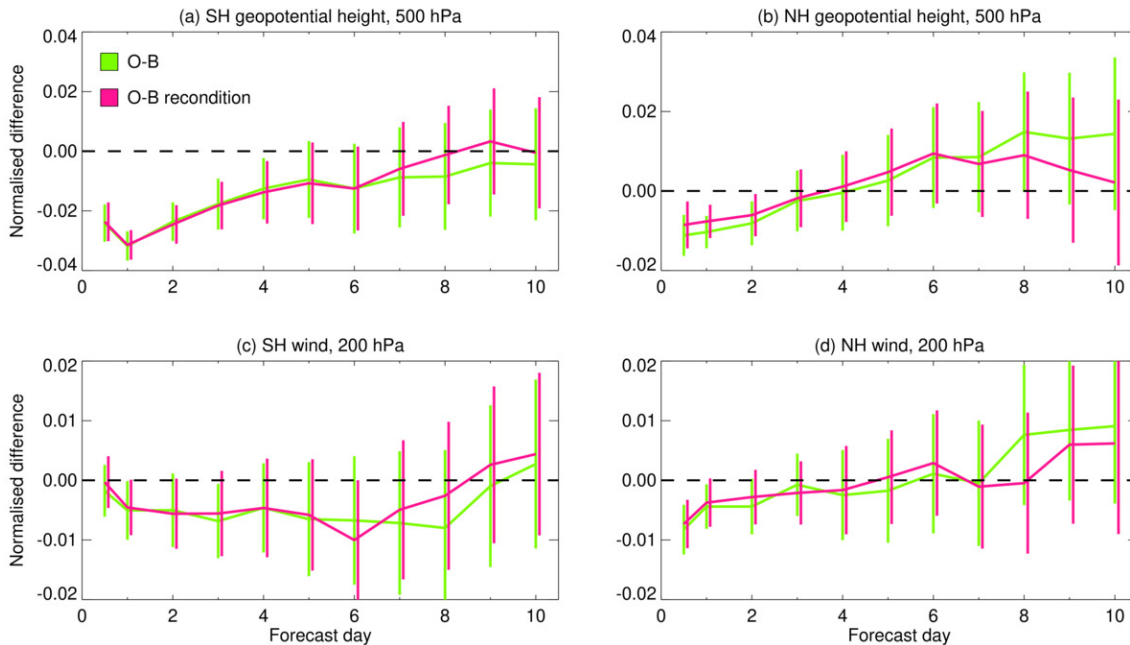


Figure 23: As Figure 21, but comparing experiments assimilating ATMS data in all-sky conditions using correlation matrices derived from O–B statistics before (green) and after (pink) reconditioning.

As the reconditioning does not seem to result in much change to the number of iterations or condition number, we now look at whether it improves the forecast. We begin by looking at results from short-range 12 hour forecasts in Figure 22, comparing the standard deviation of background departures for the same six instruments as earlier. In most cases there is little difference between the experiments, and in cases where some signs of statistically-significant differences appear (e.g. GOES-16 geostationary water vapour channels and AMVs above 300 hPa) it is the non-reconditioned matrix which performs better. Looking back at Figure 10, the biggest differences between the two correlation matrices is that the reconditioned matrices have reduced values for the off-diagonal correlations and for the correlations between the lowest-peaking temperature channels and the humidity channels. The Desroziers correlation matrices already have reduced off-diagonal correlations, and similar results to Figure 10 are seen in experiments using the reconditioned and non-reconditioned versions. Thus, it may be that the reduction in the correlations between the lowest-peaking temperature channels and the humidity channels resulting from the reconditioning is leading to the slightly worse performance.

Looking next at the medium-range impact, Figure 23 shows the normalised change in RMS error over 10 days for 500 hPa geopotential height and the 200 hPa wind. No statistically-significant difference is seen between the experiments using reconditioned and non-reconditioned correlation matrices. From the results presented in this section it does not appear that reconditioning is necessary for improving the 4D-Var minimisation, and while some benefits are seen in short-range forecasts using the non-reconditioned matrices, there is no statistically-significant impact in the medium-range.

7 Conclusions and future developments

In this report we described the changes made to the IFS to allow for all-sky assimilation with inter-channel error correlations and variational quality control. Our focus was on the Advanced Technology

Microwave Sounder (ATMS) instrument, but the amendments to the code should also allow other current and future instruments to make use of the new capabilities, including, for instance, MWI, ICI, and AWS/EPS-Sterna. We created correlation matrices using various different methods, and tested them to determine which performed the best. The results of the experiments can be summarized as follows:

- The all-sky experiments mostly replicate the positive impact achieved when adding the assimilation of clear-sky ATMS.
- When including inter-channel error correlations to the all-sky assimilations, further slight improvements are seen. While there is some evidence of improvements in short-range forecasts, it is in the medium range, up to around day 7 or 8, where the most statistically-significant improvements occur. The improvements are generally seen in the southern hemisphere, with results in the northern hemisphere being more neutral.
- When comparing all-sky assimilation using inter-channel error correlations derived from O–B and Desroziers statistics there is little difference in the results, with perhaps the matrices derived from O–B statistics performing slightly better. As the Desroziers statistics use background and analysis departures, they require an experiment with the data already actively assimilated. Thus, the correlation matrices derived from purely O–B statistics are simpler to produce, providing a practical reason for favouring this approach.
- Using inter-channel error correlation matrices derived from reconditioned diagnosed observation error covariances does not lead to any noticeable improvement to the forecast or the number of iterations required for 4D-Var minimisation to converge. The result may be specific to the particular reconditioning approach used, including the approach of leaving the error standard deviations unchanged, and further work on alternative approaches may be worthwhile.

Combined together, these results suggest there would be a benefit in moving the assimilation of ATMS from clear-sky conditions to all-sky. The impact of doing so ranges from neutral to positive in terms of forecast performance, and it would also standardise the treatment of microwave sensors at ECMWF. Including inter-channel error correlations on top of the all-sky assimilation could also give some further slight improvements, but this also adds increasing complexity to the assimilation scheme. For example, specifying the prior probability of gross error for each eigenvector is complicated, and the correlation matrices may need to be updated periodically due to changes in the model or the instrument performance. Thus, it may be prudent to await the assimilation of future instruments with increasing numbers of humidity-sensitive channels before making the inter-channel error correlation changes to the code operational. Additionally, there are still two areas that require investigation before the changes described here could be made operational. Firstly, the work presented here has investigated the impact of including two ATMS instruments (on the S-NPP and NOAA-20 satellites) but more recently a third instrument (on NOAA-21) has been added to the assimilation. Thus, experiments should be performed on more recent dates to assess the impact of having three ATMS instruments assimilated in all-sky conditions. Secondly, the changes made to the IFS to allow inter-channel error correlations to be used in all-sky conditions need to be tested further to assess their impact on the forecast sensitivity to observation impact (FSOI) and on the Ensemble of Data Assimilations (EDA) spread.

There are other aspects which could be investigated further in order to fine-tune the results presented here. As seen in Figure 6, the fact that ATMS uses two different cloud proxies (liquid water path, LWP, and scattering index, SI) for the temperature and humidity channels results in a two-dimensional correlation distribution between each pair of channels. This doesn't take the form of a simple increase in the strength

of the correlations as the LWP and SI increase, but shows specific ‘paths’ where correlations can increase rapidly or hardly increase at all. For the experiments performed in this report we simplified this two-dimensional structure by using nine bins, which captures some but not all of the observed structure. Additionally, using nine bins does not allow for interpolation between the bins, as those at the edge have no mid-point to use. Thus, it might be beneficial to investigate whether increasing the number of bins, and interpolating between the bins where possible, leads to forecast improvements due to capturing more of the correlation variation. The correlation matrices are required to be positive-definite for the eigendecomposition, and bilinear interpolation preserves this.

Another area of investigation could be the impact of the trailing eigenvalues on the forecast. As seen when looking at the J^0 sensitivity by eigenvector in Figure 14, the trailing eigenvalues can amplify small-scale vertical oscillations. This is not necessarily desirable for assimilation, as these eigenjacobians can map onto features that the assimilation scheme may not be able to correctly handle, such as gravity waves (see Geer, 2019). While reconditioning the covariance matrices before deriving the correlations does help to reduce this amplification, it also results in reductions in potentially important correlations, such as those between the lowest-peaking temperature and humidity channels (see Figure 10). Another option is not to recondition the covariance matrix, but to instead add a floor to the eigenvalues obtained from the correlation matrix during 4D-Var. As an example, Figure 24 shows the J^0 sensitivity for temperature and humidity without an eigenvalue floor (a,b) and with the eigenvalue floor set to 1 (c,d). The value of 1 was chosen as it is the value of most of the eigenvalues after reconditioning, as seen in Figure 11. The result is that the sensitivity to the leading eigenvectors is unaltered, while the amplification of the trailing eigenvectors is greatly reduced. When assimilating IASI infrared water vapour channels, (Geer, 2019) investigated using eigenvalue floors of 0.37 and 1, and found that both improved the results.

Finally, further investigations could be performed into the appropriate prior-probability of gross error (PPGE) values used by VarQC. In all the experiments in this report we used a constant value of 0.25. This is because of complications related to the different number of channels used in each observation (due to screening or broken channels; see Figure 17) which can impact the resulting eigendepartures. One possibility is to restrict the assimilation to observations where all channels are available, so that more detailed PPGE values can be specified. However, this reduces the geographical extent of the assimilated observations, the number of assimilated observations, and could result in problems in the future when channels fail. As noted by Geer (2019), for all-sky assimilation the assumption is that the gross error comes primarily from scenes where the analysis struggles to match the observed cloud or precipitation. It may therefore be beneficial to specify a PPGE for clear-sky eigendepartures, and scale this depending on the cloud proxy value from the temperature or humidity channel. However, this would require a more robust way of determining what PPGE values to use, and these may need updated in future as channels break or the noise increases.

Since some improvements are seen with ATMS, which only has a few channels with strong correlations, it is likely that accounting for inter-channel error correlations will be important for extracting the most benefit from future instruments with a larger number of cloud- and precipitation-sensitive channels, such as EUMETSAT’s Micro-Wave Imager and Ice Cloud Imager on the Metop-SG satellites, ESA’s Microwave Radiometer on the Arctic Weather Satellite, and on the future constellation of EPS-Sterna microsatellites (Lean and Bormann, 2024), and Spire’s Hyperspectral Microwave Sounder (Lean and Bormann, 2025). The technical developments presented here could also potentially help ECMWF gain more benefit from instruments that are currently operationally assimilated in all-sky conditions without inter-channel error correlations, as well as with future all-sky assimilation of infrared data.

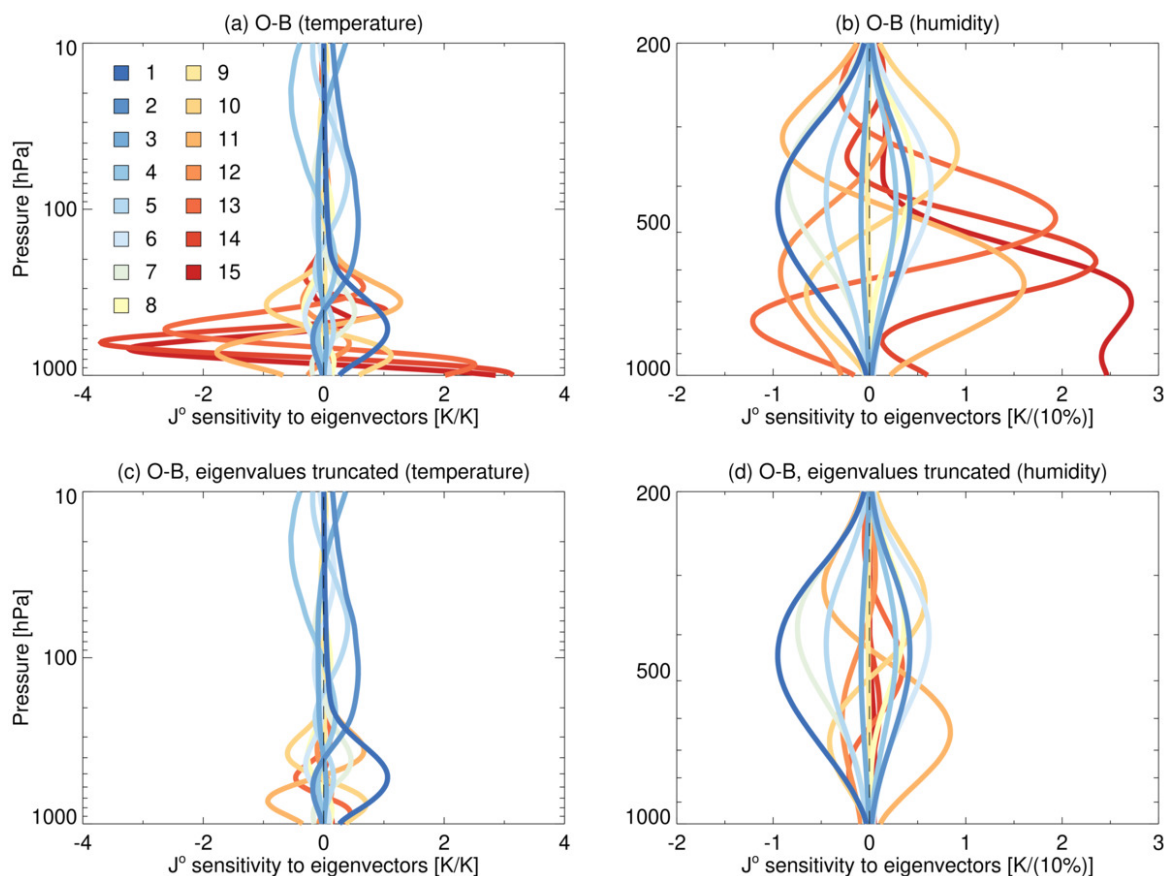


Figure 24: Temperature and humidity sensitivity of the J^0 cost function by eigenvector from correlation matrices derived using O–B diagnostics without eigenvalue truncation (a,b) and with truncation to a minimum value of $\lambda = 1$ (c,d). Eigenvectors 1–14 are ordered from largest to smallest eigenvalue, using the values before truncation.

Acknowledgements

Liam Steele is funded by the EUMETSAT Fellowship Programme. Thanks to Peter Lean for help with the experiments and analysis, and Peter Weston for helpful comments on an earlier version of this report.

References

- Andersson, E. and Järvinen, H. (1999). Variational quality control. *Quarterly Journal of the Royal Meteorological Society*, **125**(554), 697–722, doi:10.1002/qj.49712555416.
- Baordo, F. and Geer, A. J. (2016). Assimilation of SSMIS humidity-sounding channels in all-sky conditions over land using a dynamic emissivity retrieval. *Quarterly Journal of the Royal Meteorological Society*, **142**(700), 2854–2866, doi:10.1002/qj.2873.
- Bauer, P., Moreau, E., Chevallier, F. and O’Keeffe, U. (2006). Multiple-scattering microwave radiative transfer for data assimilation applications. *Quarterly Journal of the Royal Meteorological Society*, **132**(617), 1259–1281, doi:10.1256/qj.05.153.

- Bormann, N., Bonavita, M., Dragani, R., Eresmaa, R., Matricardi, M. and McNally, A. (2016). Enhancing the impact of IASI observations through an updated observation-error covariance matrix. *Quarterly Journal of the Royal Meteorological Society*, **142**(697), 1767–1780, doi:10.1002/qj.2774.
- Bormann, N., Duncan, D., English, S., Healy, S., Lonitz, K., Chen, K., Lawrence, H. and Lu, Q. (2021). Growing Operational Use of FY-3 Data in the ECMWF System. *Advances in Atmospheric Sciences*, **38**(8), 1285–1298, doi:10.1007/s00376-020-0207-3.
- Bormann, N., Fouilloux, A. and Bell, W. (2013). Evaluation and assimilation of ATMS data in the ECMWF system. *Journal of Geophysical Research (Atmospheres)*, **118**(23), 12,970–12,980, doi:10.1002/2013JD020325.
- Bormann, N., Lawrence, H. and Farnan, J. (2019). Global observing system experiments in the ECMWF assimilation system. *Technical Report 839*, ECMWF Technical Memoranda, doi:10.21957/sr184iyz, URL <https://www.ecmwf.int/node/18859>.
- Bormann, N., Magnusson, L., Duncan, D. and Dahoui, M. (2023). Characterisation and correction of orbital biases in AMSU-A and ATMS observations in the ECMWF system. doi:10.21957/d281dc221a.
- Burrows, C. (2018). Assimilation of radiance observations from geostationary satellites: First year report. *Technical Report 47*, EUMETSAT/ECMWF Fellowship Programme Research Report.
- Campbell, W. F., Satterfield, E. A., Ruston, B. and Baker, N. L. (2017). Accounting for Correlated Observation Error in a Dual-Formulation 4D Variational Data Assimilation System. *Monthly Weather Review*, **145**(3), 1019–1032, doi:10.1175/MWR-D-16-0240.1.
- Desroziers, G., Berre, L., Chapnik, B. and Poli, P. (2005). Diagnosis of observation, background and analysis-error statistics in observation space. *Quarterly Journal of the Royal Meteorological Society*, **131**(613), 3385–3396, doi:10.1256/qj.05.108.
- Duncan, D. and Bormann, N. (2020). On the addition of microwave sounders and NWP skill, including assessment of FY-3D sounders. *Technical Report 55*, EUMETSAT/ECMWF Fellowship Programme Research Report, doi:10.21957/iimhohmqt, URL <https://www.ecmwf.int/node/19760>.
- Duncan, D. I., Bormann, N., Geer, A. J. and Weston, P. (2022). Assimilation of AMSU-A in All-Sky Conditions. *Monthly Weather Review*, **150**(5), 1023–1041, doi:10.1175/MWR-D-21-0273.1.
- Eresmaa, R., Letertre-Danczak, J., Lupu, C., Bormann, N. and McNally, A. P. (2017). The assimilation of Cross-track Infrared Sounder radiances at ECMWF. *Quarterly Journal of the Royal Meteorological Society*, **143**(709), 3177–3188, doi:10.1002/qj.3171.
- Feng, C. and Pu, Z. (2025). All-Sky Assimilation of GOES-16 Water Vapor Channels in Consideration of Cloud-Dependent Interchannel Observation-Error Correlations. *Monthly Weather Review*, **153**(1), 23–47, doi:10.1175/MWR-D-24-0038.1.
- Geer, A. (2013). All-sky assimilation: better snow-scattering radiative transfer and addition of SSMIS humidity sounding channels. doi:10.21957/uxq5zlpk, URL <https://www.ecmwf.int/node/9502>.
- Geer, A., Baordo, F., Bormann, N. and English, S. (2014). All-sky assimilation of microwave humidity sounders. doi:10.21957/obosmx154, URL <https://www.ecmwf.int/node/9507>.
- Geer, A., Bauer, P. and English, S. (2012). *Assimilating AMSU-A temperature sounding channels in the presence of cloud and precipitation*. Shinfield Park, Reading, URL <https://www.ecmwf.int/node/9513>, also published as ECMWF Technical Memorandum 670.

- Geer, A., Lonitz, K., Duncan, D. and Bormann, N. (2022). Improved surface treatment for all-sky microwave observations. doi:10.21957/zi7q6hau, URL <https://www.ecmwf.int/node/20337>.
- Geer, A. J. (2019). Correlated observation error models for assimilating all-sky infrared radiances. *Atmospheric Measurement Techniques*, **12**(7), 3629–3657, doi:10.5194/amt-12-3629-2019.
- Geer, A. J. (2021). Physical characteristics of frozen hydrometeors inferred with parameter estimation. *Atmospheric Measurement Techniques*, **14**(8), 5369–5395, doi:10.5194/amt-14-5369-2021.
- Geer, A. J., Baordo, F., Bormann, N., Chambon, P., English, S. J., Kazumori, M., Lawrence, H., Lean, P., Lonitz, K. and Lupu, C. (2017). The growing impact of satellite observations sensitive to humidity, cloud and precipitation. *Quarterly Journal of the Royal Meteorological Society*, **143**(709), 3189–3206, doi:10.1002/qj.3172.
- Geer, A. J. and Bauer, P. (2011). Observation errors in all-sky data assimilation. *Quarterly Journal of the Royal Meteorological Society*, **137**(661), 2024–2037, doi:10.1002/qj.830.
- Geer, A. J., Bauer, P., Lonitz, K., Barlakas, V., Eriksson, P., Mendrok, J., Doherty, A., Hocking, J. and Chambon, P. (2021). Bulk hydrometeor optical properties for microwave and sub-millimetre radiative transfer in RTTOV-SCATT v13.0. *Geoscientific Model Development*, **14**(12), 7497–7526, doi:10.5194/gmd-14-7497-2021.
- Grody, N., Zhao, J., Ferraro, R., Weng, F. and Boers, R. (2001). Determination of precipitable water and cloud liquid water over oceans from the NOAA 15 advanced microwave sounding unit. *Journal of Geophysical Research: Atmospheres*, **106**(D3), 2943–2953, doi:10.1029/2000JD900616.
- Hollingsworth, A. and Lönnberg, P. (1986). The statistical structure of short-range forecast errors as determined from radiosonde data. Part I: The wind field. *Tellus Series A*, **38**(2), 111–136, doi:10.3402/tellusa.v38i2.11707.
- Ishibashi, T. (2024). Accurate Global Atmospheric State Analysis Using Objective Error Statistics Including Observation Error Dependence on Water Substance Field. *Earth and Space Science*, **11**(9), e2023EA003029, doi:10.1029/2023EA003029.
- Janjić, T., Bormann, N., Bocquet, M., Carton, J. A., Cohn, S. E., Dance, S. L., Losa, S. N., Nichols, N. K., Potthast, R., Waller, J. A. and Weston, P. (2018). On the representation error in data assimilation. *Quarterly Journal of the Royal Meteorological Society*, **144**(713), 1257–1278, doi:10.1002/qj.3130.
- Kan, W., Dong, P., Weng, F., Hu, H. and Dong, C. (2022). Impact of Fengyun-3E Microwave Temperature and Humidity Sounder Data on CMA Global Medium Range Weather Forecasts. *Remote Sensing*, **14**(19), 5014, doi:10.3390/rs14195014.
- Lawrence, H. and Bormann, N. (2014). *First year report: The impact of HIRS on ECMWF forecasts, adding ATMS data over land and sea ice and new observation errors for AMSU-A*. Shinfield Park, Reading, URL <https://www.ecmwf.int/node/10667>.
- Lawrence, H., Bormann, N., Geer, A. J., Lu, Q. and English, S. J. (2018). Evaluation and Assimilation of the Microwave Sounder MWHS-2 Onboard FY-3C in the ECMWF Numerical Weather Prediction System. *IEEE Transactions on Geoscience and Remote Sensing*, **56**(6), 3333–3349, doi:10.1109/TGRS.2018.2798292.
- Lean, K. and Bormann, N. (2024). *Evaluation of the EPS-Sterna 325 GHz channels in the Ensemble of Data Assimilations*. doi:10.21957/f53d05c057.

- Lean, K. and Bormann, N. (2025). *Evaluation of Hyperspectral MW for NWP: Simulation framework consolidation*. doi:10.21957/2752effd7a.
- Lean, P., Geer, A. and Lonitz, K. (2017). Assimilation of Global Precipitation Mission (GPM) Microwave Imager (GMI) in all-sky conditions. doi:10.21957/8orc7sn33, URL <https://www.ecmwf.int/node/17174>.
- Steele, L., Bormann, N. and Duncan, D. (2023). *Assimilating FY-3E MWHS-2 obs, and assessing all-sky humidity sounder thinning scales*. doi:10.21957/f42a9d9542, URL <https://www.ecmwf.int/node/17766>.
- Tabcart, J. M., Dance, S. L., Lawless, A. S., Nichols, N. K. and Waller, J. A. (2020). Improving the condition number of estimated covariance matrices. *Tellus Series A*, **72**, 1696646, doi:10.1080/16000870.2019.1696646.
- Weston, P. and Bormann, N. (2018). Enhancements to the assimilation of ATMS at ECMWF: Observation error update and addition of NOAA-20. *Technical Report 48*, EUMETSAT/ECMWF Fellowship Programme Research Report.
- Weston, P., Bormann, N., Geer, A. and Lawrence, H. (2017). *Harmonisation of the usage of microwave sounder data over land, coasts, sea ice and snow. First year report*. URL <https://www.ecmwf.int/node/17766>.
- Weston, P. P., Bell, W. and Eyre, J. R. (2014). Accounting for correlated error in the assimilation of high-resolution sounder data. *Quarterly Journal of the Royal Meteorological Society*, **140**(685), 2420–2429, doi:10.1002/qj.2306.
- Wu, T.-C., Zupanski, M., Grasso, L. D., Kummerow, C. D. and Boukabara, S.-A. (2019). All-Sky Radiance Assimilation of ATMS in HWRF: A Demonstration Study. *Monthly Weather Review*, **147**(1), 85–106, doi:10.1175/MWR-D-17-0337.1.
- Zhu, Y., Gayno, G., Purser, R. J., Su, X. and Yang, R. (2019). Expansion of the All-Sky Radiance Assimilation to ATMS at NCEP. *Monthly Weather Review*, **147**(7), 2603–2620, doi:10.1175/MWR-D-18-0228.1.