

Assimilation of Hyperspectral Infrared Observations with Optimal Spectral Sampling

Alan E. Lipton¹, Jean-Luc Moncet¹, Pan Liang¹,

Ross Hoffman^{1,2,3}, and Sid-Ahmed Boukabara³

¹Atmospheric and Environmental Research, Lexington, Massachusetts, USA

²Cooperative Institute for Climate and Satellites, University of Maryland, College

Park, College Park, Maryland

³NOAA/NESDIS/STAR, College Park, Maryland

Corresponding author:

Pan Liang

Atmospheric and Environmental Research

131 Hartwell Avenue

Lexington, MA, 02421-3126

Email: pliang@aer.com

Abstract

Current data assimilation (DA) systems use only a small fraction of the thousands of channels from hyperspectral infrared sounding instruments, including IASI, AIRS, and CrIS. An alternative that retains nearly the full hyperspectral information content at greatly reduced computation cost, while filtering instrument noise, is to assimilate the optimal spectral sampling (OSS) node radiances. The nodes are a set of optimally-selected spectral points, the node radiances are monochromatically calculated at the nodes using a fast forward model, and the nodes and weights are determined by a training procedure that ensures channel radiances can be accurately computed as a weighted sum of node radiances. For DA, we calculate the node observed radiances by inverting the OSS node-to-channel relationship in a least-squares sense. DA procedures can then be applied directly to the node radiances after some practical concerns are mitigated. The node-based approach has a quality-control advantage where emission from low clouds may contaminate channel radiances even in the stronger absorption bands because channels are wider than absorption lines, while there are uncontaminated monochromatic node radiances that have useful sensitivity to atmospheric column temperature and humidity closer to the cloud tops. Four preliminary, proof-of-concept global observing system experiments were conducted with the NOAA three-dimensional variational data assimilation system: a baseline that uses the 150 operational IASI channels; an experiment that uses a large portion (5014) of the IASI band 1 and 2 channels; and two experiments that assimilated 256 IASI-OSS node radiances with different specifications of observational errors. One of those two node-based experiments performed better than the other experiments for most metrics of forecast skill. The results of these preliminary

experiments are encouraging: the node-based approach provided positive impacts on the forecast for some metrics, regions, and lead times, despite the fact that there are aspects of the node-based processing that have not yet been optimized.

1. Introduction

Current hyperspectral infrared sounders have thousands of channels. The number of channels for such sensors is much higher than the potential degrees of freedom: i.e., the number of independent elements of information, in profiles of temperature, water vapor, and trace gases. The redundancy can, however, reduce the effect of instrument noise for the channel set as a whole. Variational data assimilation methods are capable of extracting the information content while simultaneously exploiting the noise reduction. It is impractical, however, to assimilate data from full channel sets due to the high computational costs of this approach, which involves both radiative transfer calculations and nonlinear minimization. Weather centers commonly reduce the volume of data from hyperspectral infrared sensors by presenting only a small fraction of the channels to the data assimilation system. Several methods of channel selection have been proposed (e.g., Collard 2007, and see further discussion in Sec. 3d). Typically, such methods attempt to retain information that is valuable for the forecast and eliminate channels that cannot be well modeled and hence cannot be properly assimilated. Inherently, however, there is some loss of profile information and the noise suppression effect of channel redundancy is lost. For example, at the National Centers for Environmental Prediction (NCEP), the Global Data Assimilation System (GDAS) operationally assimilates less than 2% of the channels observed by the hyperspectral Infrared Atmospheric Sounding Interferometer (IASI) sensors (Blumstein et al. 2004; Hilton et al. 2012). Weather centers have recently been experimenting with principal component (PC) approaches to compression of hyperspectral data in assimilation (Collard et al. 2010; Matricardi and McNally 2014; Lu and Zhang 2019). Here, we report on the implementation and testing of the use of optimal

spectral sampling (OSS) as an alternative to channel selection. This approach is shown to achieve a similar computational efficiency as channel selection yet retain the bulk of the information content of the full channel set.

The OSS method (Moncet et al. 2008; 2015) has been demonstrated to be a fast and accurate way to model radiometric observations and their Jacobians as linear combinations of monochromatic radiative transfer calculations (Saunders et al. 2007; Calbet et al. 2011). OSS models a set of channel radiances (\mathbf{y}) as

$$\mathbf{y} = \mathbf{W}\tilde{\mathbf{y}}, \quad (1)$$

where $\tilde{\mathbf{y}}$ represents the radiances at a set of monochromatic spectral points we call “nodes” and \mathbf{W} is a matrix of weights. OSS radiative transfer computations are thus monochromatic. It is useful to view \mathbf{W} as a projection operator from node space to channel space.¹ By differentiating $\mathbf{y} = \mathbf{W}\tilde{\mathbf{y}}$, it is clear that Jacobians (\mathbf{K}) are projected in the same way:

$$\mathbf{K} = \mathbf{W}\tilde{\mathbf{K}}. \quad (2)$$

The heart of OSS is the optimization process that determines the nodes and weights by minimizing the misfit between the OSS radiances and spectrally integrated radiances from a reference line-by-line radiative transfer model (LBLRTM; Clough et al. 1992; 2005) for sets of training profiles that span the range of conditions that occur globally, including anomalous conditions. A full description of the OSS training set properties and the OSS optimization method is provided by Moncet et al. (2008; 2015). The number of

¹ A physical interpretation for the weight values is provided in the Appendix of Moncet et al. (2015).

nodes required depends on the amount of independent information in the modeled channel set, including the number of variable gases modeled, and the required accuracy enforced during the optimization. We typically impose an accuracy requirement of 20% of the sensor noise standard deviation. For a typical accuracy requirement and modeling six variable gases (H₂O, CO₂, O₃, N₂O, CO, CH₄) only a few hundred OSS nodes are needed to compute the entire set of the more than 8000 IASI channels (Moncet et al. 2015). While it is typical for operational models to include six variable gases (Hocking et al. 2015; Chen et al. 2012), as was done in the present study, OSS can be trained to treat additional variable gases with a modest increase in the number of nodes (Moncet et al. 2008; 2015; Worden et al. 2019), such as to accommodate exceptional concentrations of SO₂ after a volcanic eruption.

Node-based assimilation operates on nodes (i.e., $\tilde{\mathbf{y}}$ and $\tilde{\mathbf{K}}$) instead of on channels (i.e., \mathbf{y} and \mathbf{K}), avoiding the transformation $\mathbf{K} = \mathbf{W}\tilde{\mathbf{K}}$, which can be computationally expensive when there are many channels and many atmospheric state variables: that is, when \mathbf{K} is large in both dimensions. Prior to assimilation, the observations are projected from channel space to node space by an inverse projection

$$\tilde{\mathbf{y}}^{\text{obs}} = \mathbf{A}\mathbf{y}^{\text{obs}}, \quad (3)$$

where \mathbf{A} is determined by least squares:

$$\mathbf{A} = (\mathbf{W}^T \mathbf{R}^{-1} \mathbf{W})^{-1} \mathbf{W}^T \mathbf{R}^{-1}, \quad (4)$$

with observation error covariance \mathbf{R} . This least-squares transformation that produces $\tilde{\mathbf{y}}^{\text{obs}}$ from \mathbf{y}^{obs} acts as a filter, both for sensor noise and for the effects of atmospheric constituents (e.g., trace gases) not included in the OSS training².

Assimilation in node space is mathematically equivalent to assimilation in channel space (as shown in APPENDIX A) if the observation error covariance in node space is defined as

$$\tilde{\mathbf{R}} = (\mathbf{W}^T \mathbf{R}^{-1} \mathbf{W})^{-1} \quad (5)$$

and the \mathbf{R} used in the variational cost function of the data assimilation (DA) system is the same as the \mathbf{R} used in (4). With practical applications, however, the \mathbf{R} used in the DA includes additional error terms (Sec. 3e). The existence of the inverse of $\mathbf{W}^T \mathbf{R}^{-1} \mathbf{W}$ is addressed in Appendix A. With this definition, $\tilde{\mathbf{R}}$ will have non-zero off-diagonal terms even if \mathbf{R} is diagonal, which is a potential disadvantage of node-based assimilation, since DA systems have often assumed that \mathbf{R} is diagonal; however, this assumption has been dropped in more modern DA systems and non-diagonal $\tilde{\mathbf{R}}$ is then no longer a disadvantage.

In our experiments, apodised IASI channel radiances were used and the nodes have been derived to best fit those apodised channel radiances, but an important motivation for

² Similarly, principal-component representations of radiances (Liu et al., 2006) apply filtering by retaining only a limited number of principal components: that is, by limiting to the number of independent pieces of information about the state of the atmosphere and surface one can extract from the spectral observations.

this work is the node-based method's suitability for assimilation of unapodised interferometric spectra from sensors such as IASI and the Cross-track Infrared Sounder (CrIS). In principle, assimilation of unapodised data will avoid the loss of spectral resolution and associated degrees of freedom (vertical profiling resolution) associated with apodisation. Projecting observations onto node space is the only currently known way to assimilate unapodised radiances while preserving the ability to filter out parts of the spectral observations corrupted by sources of atmospheric/surface variability not explicitly handled by forecast models or not well handled by radiative transfer models, without having to discard entire sets of observations. In particular, the monochromaticity of OSS nodes circumvents the spectral blending inherent in unapodised channel radiances (which have extensive side-lobes) and in principal-component (PC) representations of radiances (Liu et al. 2006; Collard et al. 2010; Matricardi and McNally 2014). Spectral blending refers to the fact that spectral signatures of any unmodeled (or poorly modeled) atmospheric components, such as trace gases or clouds, impact all PCs, making it impossible to isolate elements of a set of observations that are free of such impacts. In some DA applications it may be critical to avoiding spectral blending. Moncet et al. (2015; Sec. 5) covered the application of OSS to PCs and discussed other aspects of the relationship between PCRTM (Liu et al. 2006) and OSS.

2. Experimental design

We implemented a node-based method in the NCEP Gridpoint Statistical Interpolation (GSI) assimilation system in order to conduct experiments and analyze the performance of the node-based assimilation. For this demonstration, the node-based

version of GDAS was designed to be as consistent as possible with the original baseline version. In this work the (non-hybrid) three-dimensional variational (3D-Var) GDAS, rather than the newer hybrid four-dimensional ensemble variational (4DEnVar) GDAS (Kleist and Ide 2015a,b) was used to allow sufficient throughput with the available S4 computer system (Boukabara et al. 2016). In particular, we used GSI release 5.0.0, along with version 2.2.0 of the Community Radiative Transfer Model (CRTM), which has an option to use the OSS method for radiative transfer calculations, which we refer to as CRTM-OSS.

The CRTM-OSS software (as with other OSS forward model implementations) has a flexible design with respect to selections of nodes and weights. The only node-related information that must be supplied consists of: 1) a file containing monochromatic absorption coefficients for each relevant node, along with node-identifying information, and 2) a sensor-specific file containing weights assigned to each node for each channel the node is used to reconstruct, with indices to the node-specific absorption data in the other file. A single absorption file can be applied to more than one instrument and can cover a broad spectral interval. There are no limits on the extent of the spectral domain covered by OSS, provided that CRTM handles the Planck function, cloud, and surface property computations over that full domain.

The test case for our experiments was the two-month period June-July 2015. To facilitate implementation and interpretation, the baseline configuration assimilates only IASI brightness temperature and conventional observations, which include all scatterometer and cloud track winds as well as all various non-satellite observations that are used operationally. No other brightness temperature or global positioning system

radio occultation (GPS-RO) observations were included in the baseline. The motivation for this experimental design choice is that in research studies like this one, the control experiment is simplified to highlight the value of the new observations or method while making use of a limited sample.

We ran a series of experiments to test the impacts of expanding the use of IASI data from the operational set of selected channels (experiment: Baseline) to the full information content of IASI bands 1 and 2 in terms of channel (experiment: Full) and OSS node brightness temperatures (Table 1). When using OSS nodes, we tested two methods for representing the observation error (experiments: Node1 and Node2). A 7-day forecast was initiated once per day over the two-month experimental period. For analysis of results, we excluded the first two weeks, to ensure complete spin-up of the adaptive variational bias correction (VarBC, Zhu et al. 2014). The forecast error metrics are based on comparing each forecast to the subsequent analysis of that forecast's experiment that has the same valid time.

This experimental design reflects the fact that this is a proof-of-concept study, applying limited resources to stand up a node-based assimilation system and demonstrate its feasibility and see whether the preliminary performance assessments indicate that iterations of improvement to the system and its validation are warranted.

3. Practical aspects

a. Overview

To prepare for node-based experiments, there were several practical aspects of the GSI and associated data processing that were addressed as described in the following

subsections. Aside from these revisions, the GSI treats OSS node observations exactly as it treats channel observations. For example, the VarBC and the quality control (QC) processes operate without regard to whether the brightness temperatures are channel or node brightness temperatures.

Operationally, the 150-channel Baseline configuration includes channels from IASI bands 1 and 2, but not from band 3. OSS models the total 5420 IASI band 1 and band 2 channel set using only 404 nodes, based on an OSS training that had six variable gases (H_2O , CO_2 , O_3 , N_2O , CO , CH_4) and fit the reference model within 20% of channel noise (rms).

b. Applying CRTM-OSS to GSI

Our experiments required using CRTM-OSS as an alternative to the standard CRTM applied in GSI (hereafter CRTM-ODPS), which models channel radiances using a method called optical depth in pressure space (Chen et al. 2012). To prepare CRTM-OSS to replace CRTM-ODPS in the GSI, CRTM-OSS was modified as follows. First, the interface was revised to match the GSI identification and units of the variable gases, and to match the GSI use of brightness temperature instead of radiance. Second, an option was added to CRTM-OSS to provide $\tilde{\mathbf{y}}$ and $\tilde{\mathbf{K}}$ as outputs rather than handling them only as intermediate results in the computation of channel-space \mathbf{y} and \mathbf{K} . Third, since the GSI data QC process for IASI requires the layer optical depth, emissivity, and emissivity Jacobians, these quantities were also added as outputs.

For layer optical depths for nodes, revisions to CRTM-OSS merely took the optical depths that were treated as intermediate data and made them available as outputs. For

layer optical depths for channels, CRTM-OSS was revised so cumulative node transmittances (between each level and the top of the atmosphere) are projected to cumulative channel transmittances (similar to projection of radiances and Jacobians; Sec. 1). Then the channel layer optical depths are obtained as the logarithm of the ratio of the cumulative channel transmittances at the top and bottom of each layer. This method sometimes results in slightly negative channel layer optical depths because some OSS weights are negative and the accuracy of the OSS approximation is assured for radiances but not necessarily for layers that contribute little to the radiances. Any instances of negative channel layer optical depth are reset to zero.

For nodes, the inverse Planck function is used to transform radiances to brightness temperature. For channels, as an approximation, this transformation is applied to the channel center wavenumber.

c. IASI data preparation

The IASI data stream available at NCEP contains only 616 of the 8461 IASI channels. To get access to the full channel set, we obtained IASI L1c data from the Comprehensive Large Array-data Stewardship System (CLASS) and converted the data into Binary Universal Form for the Representation (BUFR) files.

In an off-line process, we projected IASI band 1 and 2 observations onto OSS nodes using (3) and (4). In this conversion, \mathbf{R} is a diagonal matrix of the IASI measurement error variances (D. Tobin, personal communication). This version of \mathbf{R} (\mathbf{R}_m hereafter) represents an assumption that the correlations of errors in the measurement system are negligible. The subscript m is used here to denote that \mathbf{R}_m represents measurement error

only, and does not represent other errors that are commonly included in estimates of radiometric observation error for assimilation. The monochromatic node radiances were converted to brightness temperatures using the inverse Planck function before writing the data to BUFR files. The measurement error covariance of the node radiances, $\tilde{\mathbf{R}}_m$, was computed by applying the OSS weights (5) to \mathbf{R}_m . The transformation of measurements from channels to OSS nodes introduced sparse but significant correlations (Fig. 1). To obtain the covariance in brightness temperature units we applied the linear approximation $\mathbf{D}\tilde{\mathbf{R}}_m\mathbf{D}$, where \mathbf{D} is a diagonal matrix of the derivative of the inverse Planck function at the nodes.

d. Static filtering

The baseline GSI assimilates 150 IASI channels. There are several criteria by which the IASI channel set was reduced to 150, the first of which is semi-subjective analysis based on information content (Collard 2007; Gambacorta and Barnett 2013). Additional subjective processes have been applied at NCEP, toward eliminating channels in parts of the spectrum where radiative transfer models perform relatively poorly (e.g., strong non-LTE effects) and channels where the radiative transfer is strongly nonlinear (including some water vapor channels) that have been found to interfere with convergence of the process that minimizes the variational cost function (J. Jung, personal communication). One mechanism to detect problem channels is the analysis of observation-minus-background (O–B) statistics.

To account for these practical considerations while still allowing experiments with node data and with a “full” IASI channel set, we first eliminated IASI band 3, thereby

avoiding non-LTE effects, and mimicking operational practice. We then eliminated channels and nodes for which the rms O–B was exceptionally high in relation to nearby channels. As a computationally economical way to obtain these O–B statistics, we configured our Baseline experiment to carry the full channel set and node data through the QC stage of processing, using the GSI mechanism for channels that are monitored but not assimilated. This filtering process, which is discussed in more detail in APPENDIX B, resulted in 5,014 channels and 265 nodes retained for assimilation. A shortcoming of this method is that it does not address the nonlinearity issue noted above. In future work, it may be feasible to include a nonlinearity criterion.

e. Observation error estimation

For brightness temperature assimilation in GSI, the prior estimate of observation error (PEOE) in each IASI channel represents a combination of measurement error, radiative transfer modeling error, and representativeness error. In the operational GSI, PEOEs are tabulated for 150 assimilated IASI channels and 315 monitored IASI channels (465 channels in total) in IASI bands 1 and 2. During QC, estimated observation errors (EOEs) at each field of view (FOV) are derived from the PEOEs. To complete the observation error estimation, GSI assumes inter-channel correlations are zero. To facilitate comparing our node-based results to the current GSI baseline, we retained this assumption in our experiments, while recognizing that inter-channel observation correlations for radiances, especially from hyperspectral sensors like IASI, should be included (Bormann et al. 2016; Campbell et al. 2017).

For our experiments, PEOEs were needed for the full channel set and for OSS node data. Considering that the PEOE data were produced through a partly subjective process (J. Jung, pers. comm.) that could not be simply replicated for the full channel set or for nodes, we used the operational PEOEs available for 465 channels to develop a linear regression prediction of the PEOE for a channel from the spatial pattern of radiance for that channel. We found that the spatial pattern could be adequately represented by the first 13 principal components (PCs) of the brightness temperature “maps” for one day (i.e., using PCs as predictors). Here, the PC transformation is based on the global *spatial*, location-to-location covariance matrix constructed from 465 channel brightness temperatures at all 6658 IASI FOVs for the chosen day. The regression equation was then applied to all channels and nodes for IASI bands 1 and 2. This approach essentially extrapolates PEOEs from the 465 channels, using spatial variation of a channel’s brightness temperature as a proxy for the channel’s radiometric response to varying atmospheric temperature and water vapor. Details of the PEOE extrapolation are given in APPENDIX C.

The error variance estimate for each specific measurement is adaptively calculated in the QC component of the GSI. The EOE values are initially set to the PEOE values. Then, at each step during QC, the EOE for a FOV may be multiplied by an inflation factor. At the end of the QC process, the squared EOE vector is taken to be the diagonal error variance matrix, denoted as \mathbf{V} in the case of channels and $\tilde{\mathbf{V}}$ in the case of nodes. When channel brightness temperatures are assimilated in the GSI, the observation error covariance matrix \mathbf{R} that is incorporated in the variational cost function is defined as $\mathbf{R} = \mathbf{V}$, thus

including estimates of the measurement error, radiative transfer modeling error, and representativeness error.

When node brightness temperatures are assimilated, the observation error covariance in the variational cost function is $\tilde{\mathbf{R}}$. If we were to define this as $\tilde{\mathbf{R}} = \tilde{\mathbf{R}}_m$, then the assimilation would neglect the adaptive estimates of modeling error and representativeness error. If, instead, we were to use $\tilde{\mathbf{R}} = \tilde{\mathbf{V}}$ (directly analogous to $\mathbf{R} = \mathbf{V}$ for channel assimilation), then the correlations associated with the channel-to-node transformation (Eq. 5) would be neglected. We tested two compromise options for specifying $\tilde{\mathbf{R}}$ while using the $\tilde{\mathbf{V}}$ that GSI produces when operating on nodes. In experiment Node1, $\tilde{\mathbf{R}}_m$ is scaled to obtain a matrix with the same diagonal as $\tilde{\mathbf{V}}$ and the same correlation as $\tilde{\mathbf{R}}_m$: $\tilde{\mathbf{R}} = (\tilde{\mathbf{V}})^{\frac{1}{2}} \tilde{\mathbf{C}}_m (\tilde{\mathbf{V}})^{\frac{1}{2}}$, where the correlation, $\tilde{\mathbf{C}}_m = \text{diag}(\tilde{\mathbf{R}}_m)^{-\frac{1}{2}} \tilde{\mathbf{R}}_m \text{diag}(\tilde{\mathbf{R}}_m)^{-\frac{1}{2}}$, is unaffected by conversion from radiance to brightness temperature (with the linearized conversion; Sec. 3c). In experiment Node2, $\tilde{\mathbf{R}}$ is equal to $\tilde{\mathbf{R}}_m$, except that the diagonal is set equal to the diagonal of $\tilde{\mathbf{V}}$. Note that in either case, the mathematical equivalence between channel-based and node-based assimilation is compromised by the fact that the version of \mathbf{R} used in (4) and in the definition of the GSI cost function are inherently different, because the product of (4) is used to compute an input to QC (namely $\tilde{\mathbf{y}}^{\text{obs}}$) and the $\tilde{\mathbf{R}}$ used in the cost function is an output of QC. Even if they were otherwise consistent, the QC adjustments to \mathbf{R} are inevitably different for channel-based and node-based assimilation.

f. Dynamic filtering

The QC module of the GSI operates as a dynamic filter to eliminate observations (individual channels or all channels in a FOV) and to adjust the estimated observation error variance (\mathbf{V} or $\tilde{\mathbf{V}}$) to account for factors that may affect RTM error. Factors by which channels are excluded or down-weighted for IASI include:

- IASI quality flags or consistency checks indicate unreliable data;
- high cloud fraction (per products from Advanced Very-High-Resolution Radiometer data);
- indicators of cloud impact on a per-channel basis;
- solar component of radiance is high;
- sensitivity to surface;
- sensitivity above top of the model; and
- high O–B difference.

These factors are applied to nodes in the same way as they are applied to channels, after accounting for practical issues discussed in Sec. 3b.

The test for high O–B difference compares the O–B to a channel-specific “gross check” threshold. For the operational GSI, those thresholds are provided in tabulated form along with the PEOE values for the same 465 channels. We inspected the relationship between the thresholds and the PEOE values for those channels and found that for channels $> 750 \text{ cm}^{-1}$ almost all thresholds were 6 K and for channels $< 750 \text{ cm}^{-1}$ the thresholds could be fairly well approximated by multiplying the PEOE values by 3 and capping the thresholds at 6 K. Thus, thresholds for the full channel set and for OSS

nodes were set to 6 K for channels $> 750 \text{ cm}^{-1}$ and the minimum of 6 K and 3 times the PEOE value otherwise.

We assessed the yields of data passing QC, and found that the yields were generally very consistent among the four experiments (Fig. 2). For some nodes, however, the yields were higher than those for channels at nearby spectral locations, and Sec. 4a explains how responses to cloud cover contribute to those differences in yields. Instances of nodes with anomalously low yields were attributed to the test for high O–B difference, and occurred with nodes having relatively low weights (and, hence, high noise) that were flagged by the static filtering process (Sec. 3d) to be excluded from assimilation.

g. Matrix inversion

The cost function minimization algorithm of the GSI requires the inverse of the observation error covariance matrix. For channel-based assimilation, this matrix is diagonal and inversion is computationally trivial. For node-based assimilation, a full matrix inversion step was invoked. To facilitate the ultimate transition of node-based assimilation into operations, we used the same matrix inversion software as is used by Bathmann et al. (2017) in their ongoing work at NCEP to adapt GSI to use non-diagonal observation error covariances.

4. Results and analysis

a. Sounding above clouds

With respect to making use of measurements above opaque clouds, we had found through prior 1D-Var experiments (Moncet et al. 2016) that node-based retrieval was

more robust than channel-based retrieval, because weighting functions (Jacobians) are more sharply peaked for monochromatic nodes than for channels. There are thus instances where upper-tropospheric information is available from nodes uncontaminated by low-level clouds while there are no such uncontaminated channels. We found that such circumstances occurred within our assimilation experiments. For example, Fig. 3 shows an instance where a node has peak temperature sensitivity near 520 mb and virtually no sensitivity below the 760-mb level, while the spectrally nearby channels with similar peak pressures have significant sensitivity down to almost 1000 mb. Fig. 4 shows many locations where the QC process filtered out these channels over low clouds while the QC passed (accepted) the node data for use in assimilation.

b. Forecast accuracy

The mean die-off curves for 500-hPa geopotential height anomaly correlations (AC) are similar for each of the four experiments (Fig. 5a), but the individual curves vary substantially as forecast time increases. The forecast impacts of alternatives to the Baseline are more significant at short forecast times, as is expected for data impact experiments, and this trend is more apparent in relative differences (Fig. 5c) than in absolute differences (Fig. 5b). The AC difference plots indicate significant improvements for Node1 at 24 h, for Node 2 at 120 and 144 h, marginal improvements for Node1 at 48 h and Node2 at 24 h, significant degradations for Full out to 48 h, and a marginal degradation for Full at 72 h.

In these figures, forecast impact Δ is defined as $z_c - z_e$, where z_c is the error metric for the control experiment (here, Baseline), and z_e is the error metric for the test

experiment (here, Node1, Node2, or Full). With this definition, forecast impact is positive for improvements and negative for degradations relative to the control experiment. The relative forecast impact δ is simply Δ (and its confidence intervals) divided by the sample mean of z_c . Since AC (Fig. 5) is a skill metric, we plot Δ and δ for $ACE=1-AC$. In the confidence results presented in Fig. 5 and following figures, no adjustment was made for the possibility of correlated errors.

The impact on 500-hPa geopotential height ACE forecasts is just one of many metrics that we examined. For a few key metrics, Fig. 6 has the δ -impact values and (in parentheses and color coded) the Student's t-statistic quantile (in %) for the impacts. In general, compared to Baseline, we see that impact of Node1 is positive, Node2 is neutral, and Full is negative. However, there is some variability among the relative performances of Baseline, Node1, and Node2. For example, Node2 performed best with respect to northern hemisphere AC in 4-day to 6-day forecasts, while Node1 performed best with respect to relative humidity.

To get a broad view of performance, we prepared summary assessment metrics (SAMs) following the approach of Hoffman et al. (2017; 2018). SAMs are averages of normalized assessment metrics. The normalization is via the empirical cumulative density function, which is determined for each category from the subset of all experiments and all verification times. The categories are all possible combinations of coordinates in five dimensions, as listed in Table 2. In practice, the normalization for all values z in a sample Z is given by the rank of z in Z minus $\frac{1}{2}$, all divided by the size of Z . Here z can be ACs or other skill scores or inverted RMSEs or other error metrics. Following the usual NCEP practice, ACs are calculated only for some combinations of level and variable. As a

result, the ACs are not calculated for RH (all levels), for 850 hPa height, and for 700 and 1000 hPa temperature and wind.

Fig. 7 shows the overall (global) SAM by experiment and the evolution of SAM as a function of forecast time. Overall, Node1 gives the best performance, followed by Baseline and Node2, which are approximately equivalent. Full performs much worse. The differences between experiments are greatest at the start of the forecast and decay thereafter. Fig. 8 shows how SAMs vary with level, domain, variable, and statistic. By level, Node1 is best at all levels. By domain, Node1 is best in the NHX and Tropics, but Baseline is best in the SHX. By variable, Node1 is best, but Baseline is equivalent to Node1 for RH. By statistic, Node1 is best for both AC and RMSE. Node2 performs better for AC than RMSE and the reverse holds for Baseline.

The finding that the assimilation with the “full” channel set (5,014 channels) performed worst (i.e., lowest ACs and highest errors) is likely due, in large part, to the assumption that the observation errors for channels are uncorrelated. This would have the effect of substantially overweighting the radiance observations compared to the other experiments. For example, Campbell et al. (2017) found that there were significant differences in forecast error (up to about 3%) depending on whether correlations were taken into account while assimilating 73 channels from IASI and 17 channels from the Advanced Technology Microwave Sounder (ATMS). Similarly, Bormann et al. (2016) reported substantial benefits of accounting for correlations, by various metrics, while assimilating data from 191 IASI channels. We would expect the impact of correlations to be much higher with 5,014 channels. There is currently an effort underway at NCEP to implement a capability to account for observation error correlations (Bathmann et al.

2017), as has been done at other forecast centers. The “full” channel set performance might also be degraded by shortcomings in the static filtering process (Sec. 3d), which may be passing channels with relatively high radiative transfer errors or otherwise not handled well by the cost function minimization process.

The finding that, for RH, the baseline did relatively well, may relate to a criterion of the channel reduction that was imposed on the baseline channel set but not on the node-based methods: elimination of instances with strongly nonlinear radiative transfer, including channels sensitive to water vapor, as mentioned in Sec. 3d. Ultimately, the best solution would be to advance the cost function minimization algorithm so that highly nonlinear channels or nodes do not degrade convergence. As an interim measure, a similar reduction criterion could be applied to nodes.

With regard to the fact that the baseline performed relatively well in the southern hemisphere, interpretation of the performance statistics is not straightforward because of a peculiarity of the analysis. In the southern hemisphere, a large majority of the assimilated data are from IASI, considering that no other infrared or microwave satellite sounder data were included. The forecast performance metrics are based on comparing an analysis of IASI data valid at one time with a forecast produced by starting from an analysis of IASI data at a prior valid time. In the hypothetical case where the QC applied to IASI data causes all channels to be excluded, the analysis step leaves the prior forecast unchanged and so forecast–analysis differences are zero. A method of using IASI data that introduces more information to the analysis can result in higher forecast–analysis differences, particularly if the forecast model has shortcomings that lead to systematic errors. In regions where the analysis is constrained by more diverse data sources, such as

the northern hemisphere, forecast–analysis differences would generally be more indicative of analysis errors and their growth over time.

A factor in the performance of Baseline relative to Full, Node1 and Node2 was that the channel filtering and the PEOE have been optimized (to some degree) for the 150-channel baseline but not yet for the other cases.

c. Computation time

To get insight into how the assimilation computation time is affected by using the node-based approach, we collected CPU time data from those portions of the GSI computational process that depend heavily on the radiative transfer model or otherwise depend on the number of channels. The timing results (Table 3) cover a batch of 10,748 IASI FOVs from one ± 3 -hr analysis time window. Subprogram `setuprad` runs the CRTM-OSS forward model, including radiative transfer for OSS nodes and, for channel-based assimilation, mapping of radiances and Jacobians from nodes to channels [by (1) and (2)]. Subprogram `setuprad` also runs QC on the observation data. It operates on assimilated and monitored channels. Subprogram `intrad` computes the gradient of the radiance observation terms, as part of the cost function minimization. The radiative transfer component of the `setuprad` processing was the same for all of these runs, because CRTM-OSS calculated the channel radiances at the same 404 nodes in each case. The extra time in `setuprad` for the channel-based runs included projecting radiances and Jacobians from nodes to channels, which depends on the numbers of nodes and channels and other factors (Eq. 3 of Moncet et al. 2015) and included performing QC on 465 or 5,420 channels versus 404 nodes.

The `intrad` time was much higher for the node-based runs than for the channel-based runs. The vast majority of the `intrad` time for the node-based runs was used for the inversion of the non-diagonal observation error covariance matrix needed for node-based assimilation. This attribution was confirmed in a test that treated this matrix as diagonal, and found the `intrad` time to be about the same as for the baseline (bottom row of Table 3). When the GSI implementation of correlated observation errors is complete, the matrix inversion distinction between channel-based and node-based approaches will be eliminated. While the time for the Full runs is inherently higher than for the node-based runs, the information content from the observations is virtually identical.

The computational burden of some operations of `setuprad` and `intrad` scale in relation to the counts of channels and nodes, but the timing is affected by numerous computational details and additional operations that are included in `setuprad`, `intrad`, and other components of the GSI system. Prior to operational implementation of node-based assimilation, a detailed computational profiling could be performed with the current GSI as a step toward optimizing the efficiency.

5. Conclusions

Our primary conclusions are the following.

- The GSI (3D-Var version) was successfully modified to assimilate hyperspectral sounder data that have been transformed from channels to OSS nodes, while doing all GSI processing in node space.

- The radiative transfer module was the only computational aspect of the GSI that inherently required modification to accommodate node-based assimilation. The observation error processing (including matrix inversion) required modification, but only as an interim measure until the implementation of GSI capabilities to treat observation error correlations is complete.
- Adaptation of the GSI to treating node-based data requires replacement of channel-specific tabulated error metrics with node-specific metrics. The methods we developed to essentially extrapolate channel-based metrics from the baseline channel set to the full channel set and to nodes were adequate to allow some forecast benefits of node-based assimilation to be shown.
- Performing GSI assimilation with less-exclusive channel selection (more channels than the baseline) and with node data requires methods to filter out channels or nodes that cannot be well handled by the GSI. The method we used, based on observation-minus-background difference statistics, was sufficiently effective for node-based assimilation to perform well relative to the baseline, but a method that can be objectively and consistently applied to small and large sets of channels and to nodes would be preferable.
- We tested two methods for representing the observation error covariances for nodes, which differed with respect to their assumptions about how node-projected measurement errors are combined with other observation errors. The method we called Node1 performed better than Node2 overall, but we do not consider that result to be definitive without broader testing, and there may be additional methods worthy of consideration.

- In areas with low clouds, the node-based approach was advantageous over the channel-based approach, with respect to allowing measurements characterizing the above-cloud atmosphere to be accepted (passed) by QC tests of cloud contamination, owing to the vertically narrower weighting functions for nodes than for channels.
- An experiment with node-based assimilation (Node1) performed significantly better than the channel-based baseline in overall summary assessment metrics and for most of the stratified metrics of forecast skill (variables, levels, regions) for which differences were statistically significant.
- An experiment that assimilated many more channels than the baseline (5014 versus 150) had worse forecast skill than the baseline, and we suspect that neglect of correlations of observation errors was predominantly responsible for the degradation.
- Node-based assimilation was advantageous over channel-based assimilation with respect to computation time for radiative transfer (by a factor of 4). The introduction of non-diagonal observation error covariances with the node-based approach caused the time for computing the gradient of the satellite observation terms to increase greatly, but this aspect of the distinction between the speed of channel and node-based approaches will be eliminated when the GSI transition to using non-diagonal covariances for channel observations is complete.

Overall, the results of our experiments are encouraging, from the standpoint that the node-based approach provided positive impacts on the forecast for some metrics, regions, and lead times, despite the fact that there are aspects of the node-based processing that

have not yet been optimized. Investigations of some of the performance issues discussed above would be promising subjects of follow-on work. In addition, it may be useful to perform sensitivity experiments to tease out the degree to which positive impacts of node-based assimilation are due to effective sampling of information in various parts of the spectrum (e.g. CO₂, H₂O, window regions) or improved QC yield above clouds and the surface.

Any follow-on studies should test performance over longer periods and should invest in data processing and verification aspects that avoid shortcomings of the proof-of-concept study presented here. In particular, self-analysis verification, as was used here, is well short of being ideal for data impact experiments, where the analysis error and short-term forecast error are crucial. Verification versus an independent analysis or versus high-quality observations, such as radiosonde or microwave radiances, would be desirable. In addition, data processing revisions should eliminate the dependence of node-based assimilation on adaptations from the partly subjective configuration data of the 150-channel baseline. Aspects of that dependence include static filtering that eliminates problematic channels or nodes, data representing the prior estimate of observation error in each channel or node, and thresholds applied in dynamic filtering of individual observations. To “level the playing field” that currently tilts in favor of the 150-channel baseline, fully objective methods should be developed for these aspects of processing and applied equally to node-based assimilation and a new channel-based baseline. Those new developments might include departing from the baseline GSI method of estimating observation error covariances to implement a diagnostic approach that includes off-diagonal terms.

Acknowledgements. The work described here was supported by the Joint Center for Satellite Data Assimilation through NOAA/NESDIS (award NA15NES4400005). We thank Kevin Garrett for assistance with navigating the GSI experimental system and interpreting results, Ling Liu for facilitating our work by providing a GSI test configuration and baseline data, and Jim Jung for providing valuable information about filtering, quality control, and other aspects of current operational assimilation of hyperspectral sounder data. The AVHRR Cloud Properties -NASA CDR used to provide cloud top data for this study was acquired from the NOAA National Centers for Environmental Information (<https://www.ncdc.noaa.gov/cdr/atmospheric/avhrr-cloud-properties-nasa>).

1 **APPENDIX A: Mathematical equivalence of assimilation with channels and OSS**
 2 **nodes**

3 The familiar variational assimilation problem consists of finding the minimum of the
 4 cost function

$$J = \mathbf{x}^T \mathbf{B}^{-1} \mathbf{x} + [\mathbf{y} - \mathbf{y}^{\text{obs}}]^T \mathbf{R}^{-1} [\mathbf{y} - \mathbf{y}^{\text{obs}}] \quad (6)$$

5 where \mathbf{x} is the atmospheric state vector, \mathbf{y}^{obs} is the vector of radiance observations, \mathbf{B} is
 6 the background error covariance, \mathbf{R} is the observation error covariance, and \mathbf{y} is the
 7 vector of radiances generated by a forward model operating on \mathbf{x} , which performs
 8 radiative transfer calculations. Considering that radiative transfer calculations are treated
 9 as one-dimensional problems, we can view this as a 1D-Var problem, where \mathbf{x} represents
 10 conditions along a single profile, for the purposes of considering the channel-node
 11 equivalence. This is a nonlinear problem, but the nonlinear solution is equivalent in the
 12 two modes (channel-based and node-based) if the linear solution is equivalent, so we can
 13 focus on a linear solution that represents one step in an iterative nonlinear solution.

14 For illustration, we work with a form of the iterative non-linear solution shown in
 15 Rodgers (2000), written in node space:

$$\hat{\mathbf{x}}_{i+1} = (\tilde{\mathbf{K}}_i^T \tilde{\mathbf{R}}^{-1} \tilde{\mathbf{K}}_i + \mathbf{B}^{-1})^{-1} (\tilde{\mathbf{K}}_i^T \tilde{\mathbf{R}}^{-1} \tilde{\mathbf{y}}^{\text{obs}} + \mathbf{B}^{-1} \mathbf{x}_i), \quad (7)$$

16 where i indicates the iteration, the tilde denotes node space and \mathbf{K} is a Jacobian matrix.
 17 Conversion to channel space uses the relationships (2) and (5) to show the equivalence
 18 $\tilde{\mathbf{K}}^T \tilde{\mathbf{R}}^{-1} \tilde{\mathbf{K}} = \mathbf{K}^T \mathbf{R}^{-1} \mathbf{K}$. Conversion of the term $\tilde{\mathbf{K}}^T \tilde{\mathbf{R}}^{-1} \tilde{\mathbf{y}}^{\text{obs}}$ can be done by substituting (5)
 19 into (4) and substituting that result into (3), yielding $\tilde{\mathbf{y}}^{\text{obs}} = \tilde{\mathbf{R}} \mathbf{W}^T \mathbf{R}^{-1} \mathbf{y}^{\text{obs}}$, so

20 $\tilde{\mathbf{K}}^T \tilde{\mathbf{R}}^{-1} \tilde{\mathbf{y}}^{\text{obs}} = \tilde{\mathbf{K}}^T \mathbf{W}^T \mathbf{R}^{-1} \mathbf{y}^{\text{obs}}$. Then applying (2) to replace $\tilde{\mathbf{K}}^T \mathbf{W}^T$ with \mathbf{K}^T yields
 21 $\tilde{\mathbf{K}}^T \tilde{\mathbf{R}}^{-1} \tilde{\mathbf{y}}^{\text{obs}} = \mathbf{K}^T \mathbf{R}^{-1} \mathbf{y}^{\text{obs}}$. Using these equivalences by substitution in (7) produces the
 22 same linear solution in channel space:

$$\hat{\mathbf{x}}_{i+1} = (\mathbf{K}_i^T \mathbf{R}^{-1} \mathbf{K}_i + \mathbf{B}^{-1})^{-1} (\mathbf{K}_i^T \mathbf{R}^{-1} \mathbf{y}^{\text{obs}} + \mathbf{B}^{-1} \mathbf{x}_i). \quad (8)$$

23 This derivation shows that the same result $\hat{\mathbf{x}}$ is obtained regardless of what is used
 24 for \mathbf{W} and for \mathbf{R} , as long as 1) the observation error covariances in channel space and
 25 node space are related by (5) and 2) \mathbf{W} and \mathbf{R} are used consistently throughout the
 26 computations. Specifically, the \mathbf{W} used in (1) and (2) in the forward model for the
 27 channel-space inversion (8) must be the same as is used when applying (4) and (5) to
 28 prepare terms for the node-space inversion (7), and the same \mathbf{R} must be used in (4) and
 29 (5) for the node-space solution and in (8) for the channel-space solution. However, in the
 30 experiments described in the main text, there are steps in the GSI that introduce
 31 inconsistencies between the \mathbf{R} used in (4) and (5) and the \mathbf{R} and $\tilde{\mathbf{R}}$ used in the respective
 32 channel-space and node-space inversions.

33 This demonstration of the equivalence between the channel-based and node-based
 34 solutions depends on (5) and, therefore, on the invertibility of $\mathbf{W}^T \mathbf{R}^{-1} \mathbf{W}$, which holds if
 35 $\mathbf{W}^T \mathbf{W}$ is invertible. The only condition under which $\mathbf{W}^T \mathbf{W}$ would not be invertible is
 36 when two or more columns of \mathbf{W} are linearly dependent, meaning that monochromatic
 37 radiances at some of the selected nodes could be expressed as linear combinations of
 38 radiances at other nodes. By design, the node selection process eliminates this possibility,
 39 as explained in detail by Moncet et al. (2008; 2015). In other words, if a matrix \mathbf{W} has
 40 been produced by the OSS training process, then $\mathbf{W}^T \mathbf{W}$ is invertible.

41

42 **APPENDIX B: Filtering based on observation-minus-background statistics**

43 We generated observation-minus-background (O–B) difference statistics for the full
44 IASI channel set (bands 1 and 2) and for OSS nodes as a by-product of a two-month
45 experiment that assimilated only the baseline 150 channels (Sec. 2). The O–B difference
46 statistics are affected by the GSI adaptive bias correction (Zhu et al. 2014). To verify the
47 bias correction was operating properly with CRTM-OSS, on the full channel set and on
48 OSS nodes, we compared these O–B difference statistics to the results from CRTM-
49 ODPS for the 465 channels in IASI bands 1 and 2, with and without bias correction. The
50 bias correction was as effective with CRTM-OSS as with CRTM-ODPS in terms of O–B
51 bias and rms (not shown)³. The O–B statistics were similar for channels and OSS nodes,
52 except for exceptionally high rms differences for some nodes (Fig. B1). We would expect
53 higher differences for nodes with relatively low OSS weights, for which the measurement
54 noise gets magnified in the transformation from channels to weights (5), and analysis
55 showed that the highest O–B rms generally occurred with nodes that had relatively high
56 measurement noise. Another contributor to O–B rms differences for nodes would be
57 background errors with vertical length scales resolvable by nodes but not by channels
58 (e.g., Fig. 3).

59 To apply the O–B results to static filtering of the full channel set, we identified
60 channels with exceptionally high O–B rms differences, relative to the overall spectral

³ The bias correction algorithm has the incidental effect of reducing O–B rms as well as bias.

61 trends. Specifically, we applied a Lee filter to the O–B rms differences to represent the
62 spectral trends. A Lee filter is a smoother of the form

$$h_i = \bar{k} + s(g_i - \bar{k}), \quad (9)$$

63 where g_i is the i th element of the input vector, h_i is the corresponding output, \bar{k} is the
64 mean of the window that centers on point i , and s is the weight function

$$s = \frac{\sigma_k^2}{\sigma_k^2 + \sigma^2}, \quad (10)$$

65 where σ^2 is the variance of the spectrum and σ_k^2 is the variance of the window. We chose
66 a window with 20 points on either side of point i , by subjective interpretation of plots
67 (e.g., Fig. B2a). We chose to exclude channels with O–B rms exceeding the Lee filter
68 output by greater than a threshold 0.15 K, except the threshold was increased to 0.3 K in
69 the range 1350–1750 cm^{-1} because the filter output in this range was affected by a
70 scattering of points with relatively low O–B rms and thus a 0.15-K threshold resulted in
71 excluding many channels. This process resulted in 5014 channels being selected and 406
72 channels excluded out of the total 5420 channels. Among the excluded channels are some
73 around the 667- cm^{-1} Q branch of CO_2 . For application to OSS nodes, we started by
74 interpolating the channel-based Lee filter output from channel centers to nodes,
75 considering that the nodes are too sparse and irregularly spaced for application of such a
76 filter. With nodes, a threshold of 0.5 K difference from the Lee filter output resulted in
77 139 of the total 404 nodes being identified as excluded and 265 being selected for
78 assimilation (Fig. B2).

79

80 **APPENDIX C: Extrapolation of PEOE to channels and nodes**

81 To run experiments with the full IASI channel set and with OSS nodes, we needed a
82 method to obtain PEOE values for nodes and for channels not normally treated by GSI.
83 To be as consistent as possible with the baseline GSI, we essentially extrapolated the
84 baseline PEOE data to the full channel set and to the OSS nodes⁴.

85 The approach we chose took account of the fact that, in spectral plots of PEOE data
86 provided in the baseline GSI configuration for IASI, it was apparent that there was some
87 relation to the atmospheric absorption spectrum. The absorption spectrum dictates how
88 brightness temperatures in a channel at any particular satellite FOV reflect the vertical
89 profile and surface (e.g., temperature, water vapor) at the FOV. With this in mind, we
90 created a dataset composed of simulated brightness temperatures for all 6658 globally
91 distributed IASI FOVs from 17 June 2015, including all 465 channels that are either
92 assimilated or monitored by the baseline GSI. We used brightness temperatures simulated
93 during the Baseline experiment (i.e., background atmosphere data processed through the
94 RTM) rather than using real measurements in this process, as a means to exclude
95 measurement error from this PEOE extrapolation process. We computed a matrix

⁴ Considering that the PEOE values include radiative transfer (RT) modeling error, this approach implicitly assumes that the RT errors are the same for OSS as for the ODPS method that is the standard approach with CRTM. Experiments have shown that these RT errors are lower for OSS than for ODPS. However, the PEOE values are empirically derived (tuned), so there is no quantification of the portion of the PEOE representing RT error, and thus we have no firm basis to adjust the PEOE for such model differences.

96 representing the covariance of the FOVs (a 6658×6658 matrix) based on the 465 spectral
97 samples, and then computed the eigenvectors of that covariance matrix. Because the
98 FOVs represent a single day, the eigenvectors make a coherent spatial pattern when
99 plotted as maps. For example (Fig. C1), the eigenvector 1 structure is dominated by
100 equator-pole contrast and the eigenvector 2 structure highlights the subtropics and the
101 south pole in contrast to other latitudes. We used this eigenvector matrix to transform
102 each of the 465 channel “maps” into principal components (PCs) and trained a linear
103 regression to “predict” the PEOE from varying subsets of the PCs. In this regression,
104 PEOEs were represented as standard deviations of brightness temperatures, as in the
105 original tabulation of the 465-channel PEOE data. We assessed the performance of this
106 regression approach with varying numbers of PCs as “predictors” (Fig. C2) and
107 concluded that 13 eigenvectors captured virtually all the predictive skill with no
108 indication of overfitting in the jack-knife results (i.e., no evidence of error rising with
109 increasing numbers of PCs).

110 To obtain the PEOE for the full channel set, we computed each channel’s brightness
111 temperature for each of the 6658 FOVs, transformed to PCs, and applied the regression
112 coefficients (Fig. C3a). As compared with the values provided by the baseline GSI
113 configuration data, the data produced by the regression are quite consistent throughout
114 the spectrum. As a point of reference, we computed the difference between the PEOE and
115 the O–B rms for each of the 465 channels for which tabulated PEOE data were available.
116 We found that this difference was higher than 0.1 K for 99% of the FOVs. For
117 application to the full channel set and to OSS nodes, we increased the PEOE, as needed,

118 to ensure that it exceeded the O–B rms for the respective channel by at least 0.1 K⁵. The
119 impact of adjusting PEOE values by this criterion was modest (Fig. C3b), and most of the
120 affected channels were in the 1400–1800 cm⁻¹ portion of the water vapor v2 band.

121 When the same approach was applied to OSS node brightness temperatures, the
122 results were generally consistent with the data for channels, but there were instances of
123 exceptionally high (~650 cm⁻¹) and low (~1300 cm⁻¹) PEOE values. The uneven results
124 are unsurprising when considering that the OSS node brightness temperatures are
125 monochromatic, and hence there are node value combinations (across the FOVs) that are
126 outside the range of anything that was included in the regression training data. We
127 mitigated this behavior for nodes whose nearest neighboring channel was among the 465
128 assimilated+monitored channels for which PEOE was provided by replacing the OSS
129 node regression results with data from that nearest channel (Fig. C4a). The adjustment
130 derived from O–B data had a more prominent effect on the PEOE for nodes (Fig. C4b)
131 than for channels. Those more prominent effects were, however, mostly limited to nodes
132 that were not among the 265 selected for assimilation.

133

⁵ It would be reasonable to expect PEOE to be lower than the O–B rms, since forecast error contributes to the O–B rms and they otherwise represent the same error sources. Nevertheless, we chose to retain this relationship between the specified PEOEs and the respective channel’s computed O–B rms as we extrapolated PEOEs.

REFERENCES

- 134
135 Bathmann, K., R. Todling, A. Collard, W. Gu, and J. Derber, 2017: Assimilating infrared
136 sounder observations with correlated errors, 15th JCSDA Tech. Review and Science
137 Workshop, NOAA Center for Weather and Climate Prediction, College Park, MD, 17–
138 19 May.
- 139 Blumstein, D., G. Chalon, T. Carlier, C. Buil, P. Hebert, T. Maciaszek, G. Ponce, T.
140 Phulpin, B. Tournier, D. Simeoni, P. Astruc, A. Clauss, G. Kayal, and R. Jegou, 2004:
141 IASI instrument: technical overview and measured performances. Proc. SPIE 5543,
142 Infrared Spaceborne Remote Sensing XII. <https://doi.org/10.1117/12.560907>.
- 143 Bormann, N., M. Bonavita, R. Dragani, R. Eresmaa, M. Matricardi, and A. McNally,
144 2016: Enhancing the impact of IASI observations through an updated observation-
145 error covariance matrix. *Q. J. R. Meteorol. Soc.*, **142**, 1767-1780,
146 <https://doi.org/10.1002/qj.2774>.
- 147 Boukabara, S.-A., et al., 2016: An O2R/R2O infrastructure for optimizing satellite data
148 utilization in NOAA numerical modeling systems. *Bull. Amer. Meteorol. Soc.*, **97**,
149 2359–2378, <https://doi.org/10.1175/BAMS-D-14-00188.1>.
- 150 Calbet, X., R. Kivi, S. Tjemkes, F. Montagner, and R. Stuhlmann, 2011: Matching
151 radiative transfer models and radiosonde data from the EPS/Metop Sodankylä
152 campaign to IASI measurements. *Atmos. Meas. Tech.*, **4**, 1177-1189,
153 <https://doi.org/10.5194/amt-4-1177-2011>.
- 154 Campbell, W. F., E. A. Satterfield, B. Ruston, and N. L. Baker, 2017: Accounting for
155 Correlated Observation Error in a Dual-Formulation 4D Variational Data Assimilation

156 System. *Mon. Wea. Rev.*, **145**, 1019–1032, <https://doi.org/10.1175/MWR-D-16->
157 0240.1.

158 Chen, Y., Y. Han, and F. Weng, 2012: Comparison of two transmittance algorithms in
159 the community radiative transfer model: Application to AVHRR, *J. Geophys. Res.*,
160 **117**, D06206, <https://doi.org/10.1029/2011JD016656>.

161 Clough, S. A. et al., 2005: Atmospheric radiative transfer modeling: a summary of the
162 AER codes. Short Communication, *J. Quant. Spectrosc. Radiat. Transfer*, **91**, 233-
163 244.

164 Clough, S. A., M. J. Iacono, and J.-L. Moncet, 1992: Line-by-line calculation of
165 atmospheric fluxes and cooling rates: Application to water vapor. *J. Geophys. Res.*,
166 **97**, 15,761-15,785.

167 Collard, A. D., 2007: Selection of IASI channels for use in numerical weather prediction,
168 *Q. J. R. Meteorol. Soc.*, **133**, 1977–1991, <https://doi.org/10.1002/qj.178>.

169 Collard, A. D., A. P. McNally, F. I. Hilton, S. B. Healy, and N. C. Atkinson, 2010: The
170 use of principal component analysis for the assimilation of high-resolution infrared
171 sounder observations for numerical weather prediction. *Quart. J. Roy. Meteorol. Soc.*,
172 **136**, 2038-2050, <https://doi.org/10.1002/qj.701>.

173 Gambacorta, A. and C. D. Barnet, 2013: Methodology and Information Content of the
174 NOAA NESDIS Operational Channel Selection for the Cross-Track Infrared Sounder
175 (CrIS), *IEEE Trans. Geosci. Remote Sens.*, **51**, 3207–3216,
176 <https://doi.org/10.1109/TGRS.2012.2220369>.

177 Hilton, F., R. Armante, T. August, C. Barnet, A. Bouchard, C. Camy-Peyret, V. Capelle,
178 L. Clarisse, C. Clerbaux, P. Coheur, A. Collard, C. Crevoisier, G. Dufour, D.

179 Edwards, F. Faijan, N. Fourrié, A. Gambacorta, M. Goldberg, V. Guidard, D.
180 Hurtmans, S. Illingworth, N. Jacquinet-Husson, T. Kerzenmacher, D. Klaes, L.
181 Lavanant, G. Masiello, M. Matricardi, A. McNally, S. Newman, E. Pavelin, S. Payan,
182 E. Péquignot, S. Peyridieu, T. Phulpin, J. Remedios, P. Schlüssel, C. Serio, L. Strow,
183 C. Stubenrauch, J. Taylor, D. Tobin, W. Wolf, and D. Zhou, 2012: Hyperspectral
184 Earth Observation from IASI: Five Years of Accomplishments. *Bull. Amer. Meteor.*
185 *Soc.*, **93**, 347–370, <https://doi.org/10.1175/BAMS-D-11-00027.1>.

186 Hocking, J., P. Rayer, D. Rundle, R. Saunders, M. Matricardi, A. Geer, P. Brunel, and J.
187 Vidot, 2015: RTTOV v11 users guide. NWP SAF Tech. Rep. NWPSAF-MO-UD-028,
188 139 pp. [Available online at
189 [https://www.nwpsaf.eu/site/download/documentation/rtm/docs_rttov11/users_guide_1](https://www.nwpsaf.eu/site/download/documentation/rtm/docs_rttov11/users_guide_1_1_v1.4.pdf)
190 [1_1_v1.4.pdf](https://www.nwpsaf.eu/site/download/documentation/rtm/docs_rttov11/users_guide_1_1_v1.4.pdf)]

191 Hoffman, R. N., S.-A. Boukabara, V. K. Kumar, K. Garrett, S. P. Casey, and R. Atlas,
192 2017: An Empirical Cumulative Density Function Approach to Defining Summary
193 NWP Forecast Assessment Metrics. *Mon. Wea. Rev.*, **145**, 1427–1435,
194 <https://doi.org/10.1175/MWR-D-16-0271.1>.

195 Hoffman, R. N., V. K. Kumar, S.-A. Boukabara, K. Ide, F. Yang, and R. Atlas, 2018:
196 Progress in forecast skill at three leading global operational NWP centers during 2015-
197 2017 as seen in Summary Assessment Metrics (SAMs). *Weather Forecast.*, **33**, 1661–
198 1679, <https://doi.org/10.1175/WAF-D-18-0117.1>.

199 Kleist, D. T. and K. Ide, 2015a: An OSSE-Based Evaluation of Hybrid Variational–
200 Ensemble Data Assimilation for the NCEP GFS. Part I: System Description and 3D-

201 Hybrid Results. *Mon. Wea. Rev.*, **143**, 433–451, <https://doi.org/10.1175/MWR-D-13->
202 00351.1.

203 Kleist, D. T. and K. Ide, 2015b: An OSSE-Based Evaluation of Hybrid Variational–
204 Ensemble Data Assimilation for the NCEP GFS. Part II: 4DEnVar and Hybrid
205 Variants. *Mon. Wea. Rev.*, **143**, 452–470, <https://doi.org/10.1175/MWR-D-13->
206 00350.1.

207 Liu, X., W. L. Smith, D. K. Zhou, and A. Larar, 2006: Principal component-based
208 radiative transfer model for hyperspectral sensors: theoretical concept. *Appl. Opt.*, **45**,
209 201–209.

210 Lu, Y. and F. Zhang, 2019: Toward ensemble assimilation of hyperspectral satellite
211 observations with data compression and dimension reduction using principal
212 component analysis. *Mon. Wea. Rev.*, **147**, 3505–3518, <https://doi.org/10.1175/MWR->
213 D-18-0454.1

214 Matricardi, M., and A. P. McNally, 2014: The direct assimilation of principal
215 components of IASI spectra in the ECMWF 4D-Var. *Quart. J. Roy. Meteor. Soc.*, **140**,
216 573–582, <https://doi.org/10.1002/qj.2156>.

217 Minnis, P. and co-authors, 2016: A Consistent Long-Term Cloud and Clear-Sky
218 Radiation Property Dataset from the Advanced Very High Resolution Radiometer
219 (AVHRR). Algorithm Theoretical Basis Document. AVHRR Cloud Properties –
220 NASA [CDRP-ATBD-0826].

221 Moncet, J.-L., G. Uymin, A. E. Lipton, H. E. Snell, 2008: Infrared radiance modeling by
222 optimal spectral sampling. *J. Atmos. Sci.*, **65**, 3917–3934,
223 <https://doi.org/10.1175/2008JAS2711.1>.

224 Moncet, J.-L., G. Uymin, P. Liang and A. E. Lipton, 2015: Fast and accurate radiative
225 transfer in the thermal regime by simultaneous optimal spectral sampling over all
226 channels. *J. Atmos. Sci.*, **72**, 2622–2641, <https://doi.org/10.1175/JAS-D-14-0190.1>.

227 Moncet, J.-L., G. Uymin, and A. E. Lipton, 2016: Evaluation of OSS Node-Based
228 Assimilation: Initial Assessment. 14th JCSDA Technical Review Meeting & Science
229 Workshop on Satellite Data Assimilation, May 31–June 2,
230 https://www.jcsda.noaa.gov/documents/meetings/wkshp2016/posters/ALipton_JCSD
231 [Awkshp2016poster.pdf](https://www.jcsda.noaa.gov/documents/meetings/wkshp2016/posters/ALipton_JCSD).

232 Rodgers, C. D., 2000: *Inverse Methods for Atmospheric Sounding: Theory and Practice*,
233 World Scientific, Singapore, 238 pp.

234 Saunders, R. W., P. Rayer, P. Brunel, A. von Engeln, N. Bormann, L. Strow, S. Hannon,
235 S. Heilliette, X. Liu, F. Miskolczi, Y. Han, G. Masiello, J.-L. Moncet, G. Uymin, V.
236 Sherlock and D. S. Turner, 2007: A comparison of radiative transfer models for
237 simulating Atmospheric Infrared Sounder (AIRS) radiances. *J. Geophys. Res.*, **112**,
238 D01S90, <https://doi.org/10.1029/2006JD007088>.

239 Worden, J. R., S. S. Kulawik, D. Fu, V. H. Payne, A. E. Lipton, I. Polonsky, Y. He, K.
240 Cady-Pereira, J.-L. Moncet, R. L. Herman, F. W. Irion, and K. W. Bowman, 2019:
241 Characterization and evaluation of AIRS-based estimates of the deuterium content of
242 water vapor. *Atmos. Meas. Techniques*, **12**, 2331–2339, <https://doi.org/10.5194/amt->
243 [12-2331-2019](https://doi.org/10.5194/amt-12-2331-2019).

244 Zhu, Y., J. Derber, A. Collard, D. Dee, R. Treadon, G. Gayno, and J. A. Jung, 2014:
245 Enhanced radiance bias correction in the National Centers for Environmental

- 246 Prediction's Gridpoint Statistical Interpolation data assimilation system. *Quart. J. Roy.*
247 *Meteorol. Soc.*, **140**, 1479–1492, <https://doi.org/10.1002/qj.2233>.

248 Table 1. Summary of assimilation experiments

Experiment	Assimilated	Obs error method
Baseline:	150 channels	
Full:	5014 channels	
Node1:	265 nodes	method 1
Node2:	265 nodes	method 2

249

250 Table 2. The dimensions over which SAMs were computed, and the specific coordinates
251 at which they were computed.

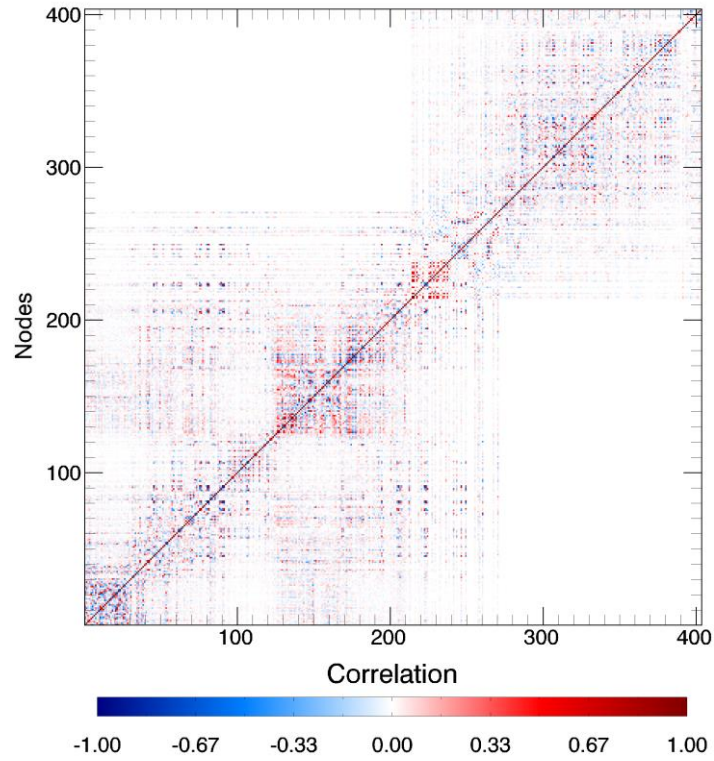
Dimension	Coordinates
forecast time	24, 48, 72, 96, 120, 144, 168 h
level	250, 500, 700, 850, 1000 hPa
domain	northern hemisphere extratropics (NHX), southern hemisphere extratropics (SHX), tropics
variable	geopotential height (Z), temperature (T), vector wind (V), relative, humidity (RH)
statistic	anomaly correlation (AC), root mean square error (RMSE)

252

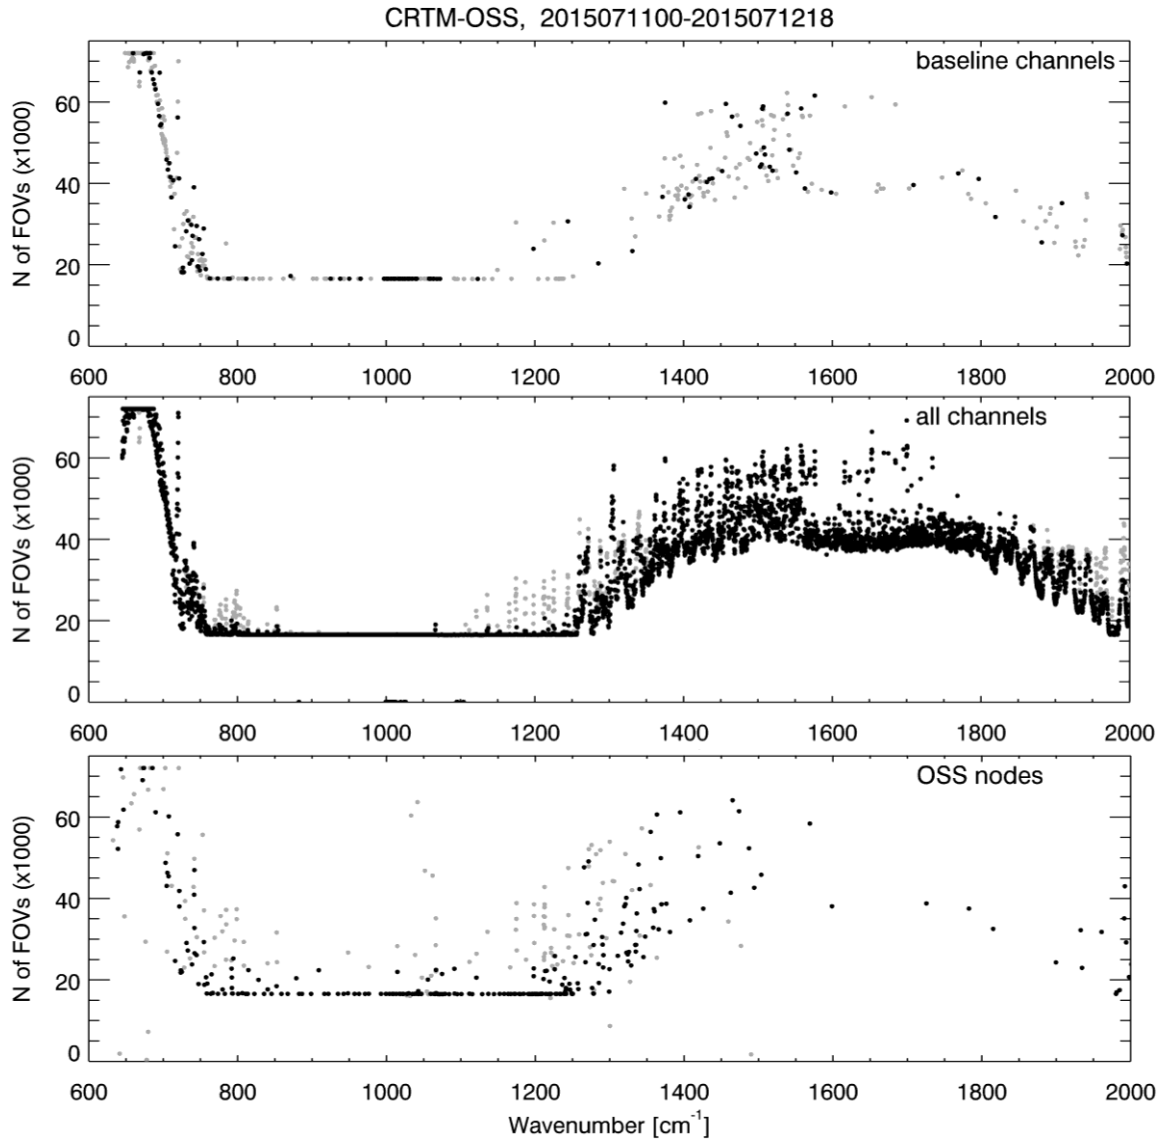
253 Table 3. Computation time per IASI field of view for subprocesses `setuprad` (3
 254 iterations, as in baseline GSI execution) and `intrad` (2 iterations), in absolute terms and
 255 relative to the 150-channel baseline assimilation, with counts of channels or nodes
 256 processed through radiative transfer, mapping Jacobians from nodes to channels, quality
 257 control, and assimilation.

	counts				setuprad		intrad	
	RT	map	QC	assim	time (s)	relative	time (s)	relative
Baseline: 150 channels	404	465	465	150	2.01	1.00	0.03	1.00
Full: 5014 channels	404	5420	5420	5014	11.05	5.49	0.55	18.28
Node1: 265 nodes	404	0	404	265	0.49	0.25	15.30	509.04
Node2: 265 nodes	404	0	404	265	0.50	0.25	15.34	510.32
Test: 265 nodes, uncorrelated	404	0	404	265	0.50	0.25	0.03	0.96

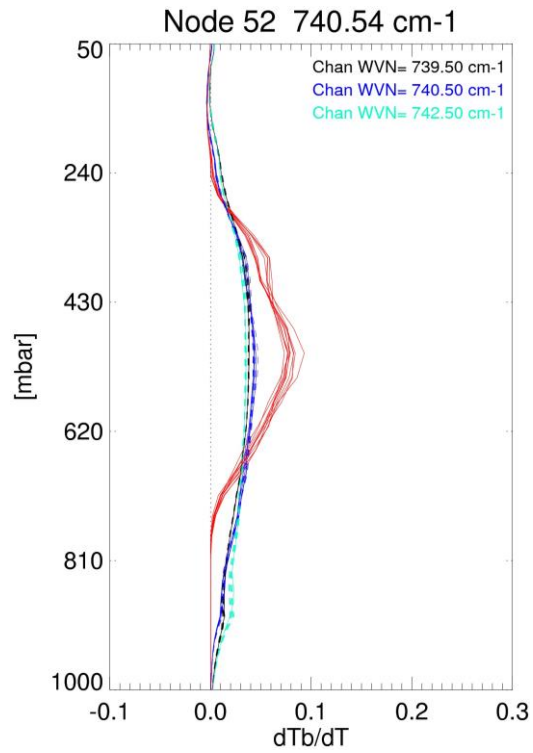
258



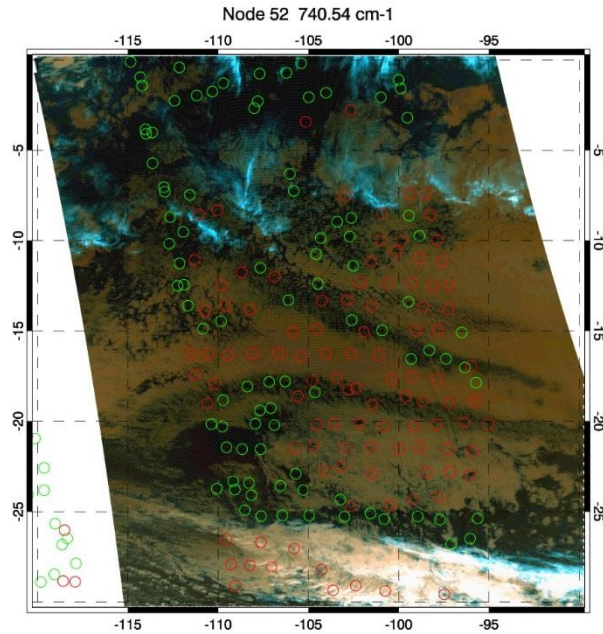
259
260 Fig. 1. Measurement error correlation after projection of all 5420 channels from IASI
261 bands 1 and 2 onto the 404 IASI OSS nodes.



262
 263 Fig. 2. The numbers of measurements passing all of the QC tests for the 465 baseline
 264 assimilated+monitored channels (top), the full channel set (middle), and OSS nodes
 265 (bottom), based on data from 11 July 2015 00 UTC – 12 July 2015 18 UTC. Black and
 266 gray markers indicate selected and excluded channels or nodes, respectively, according to
 267 the static filtering.



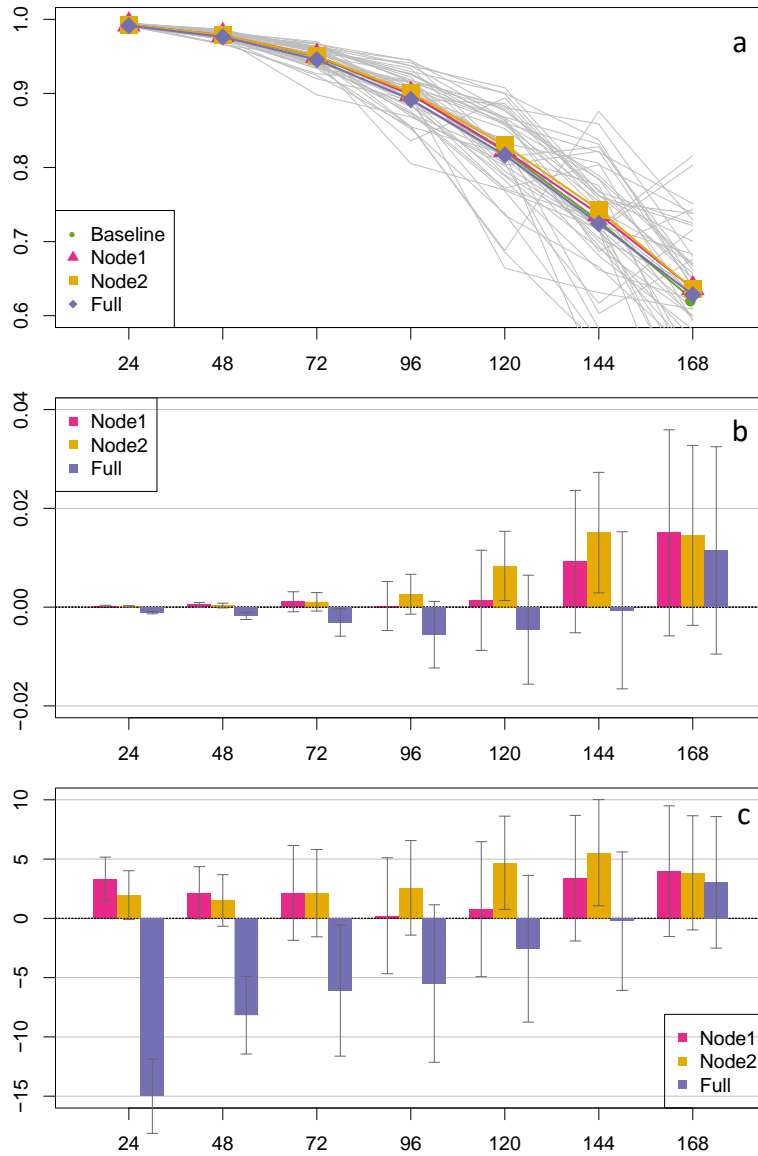
268
 269 Fig. 3. Temperature Jacobians for an OSS node at 740.54 cm^{-1} (red) and for three nearby
 270 channels (black, blue, and cyan) for which the Jacobians peak at about the same pressure
 271 level. The 11 curves are for observations within the subtropical South Pacific region
 272 shown in Fig. 4, on July 28, 2015.



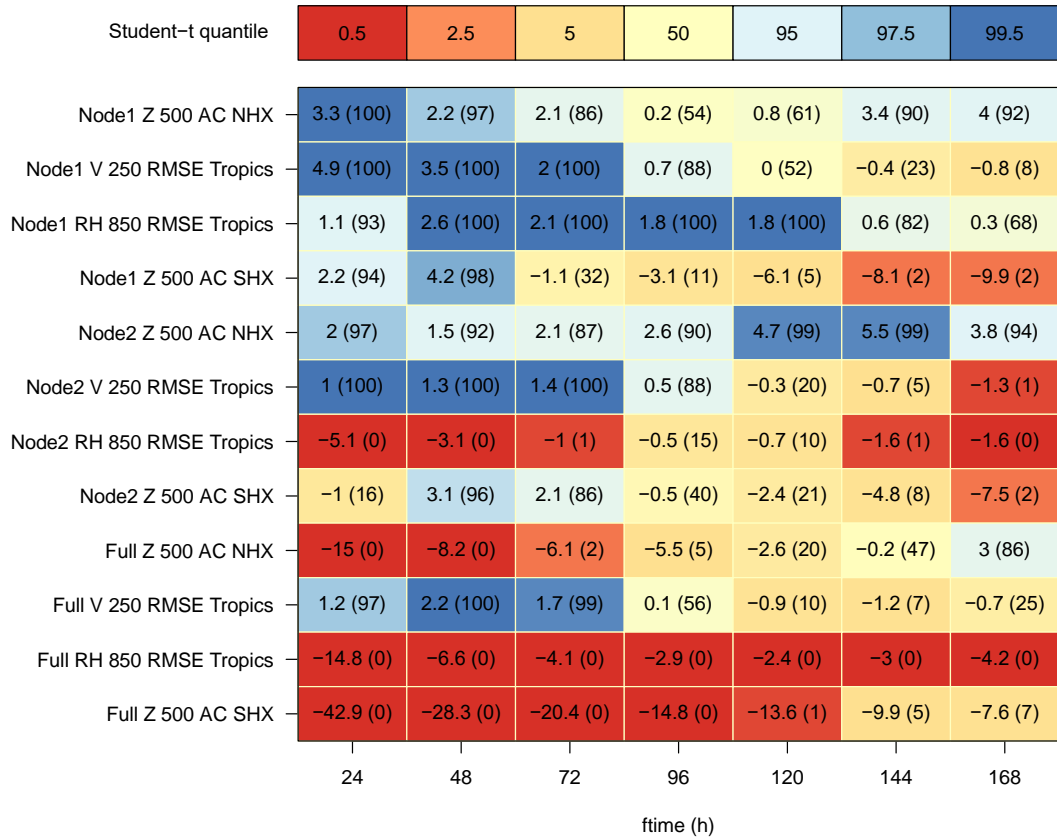
273
 274 Fig. 4. July 28, 2015 image composite of Advanced Very-High-Resolution Radiometer
 275 (AVHRR) channels centered near 3.7 μm (red), 10.8 μm (green), and 12.0 μm (blue),
 276 whereby high, thin clouds appear cyan, low clouds appear reddish brown, and thick
 277 clouds appear white. The clouds appearing reddish brown here had top pressures roughly
 278 in the range 820–880 mb, according to cloud property retrievals (Minnis et al. 2016).
 279 IASI observation sites where the QC process passed the node at 740.54 cm^{-1} and also
 280 passed the associated channels (see Fig. 3) are marked green, and sites where the node
 281 passed but the channels did not pass are marked red.

282

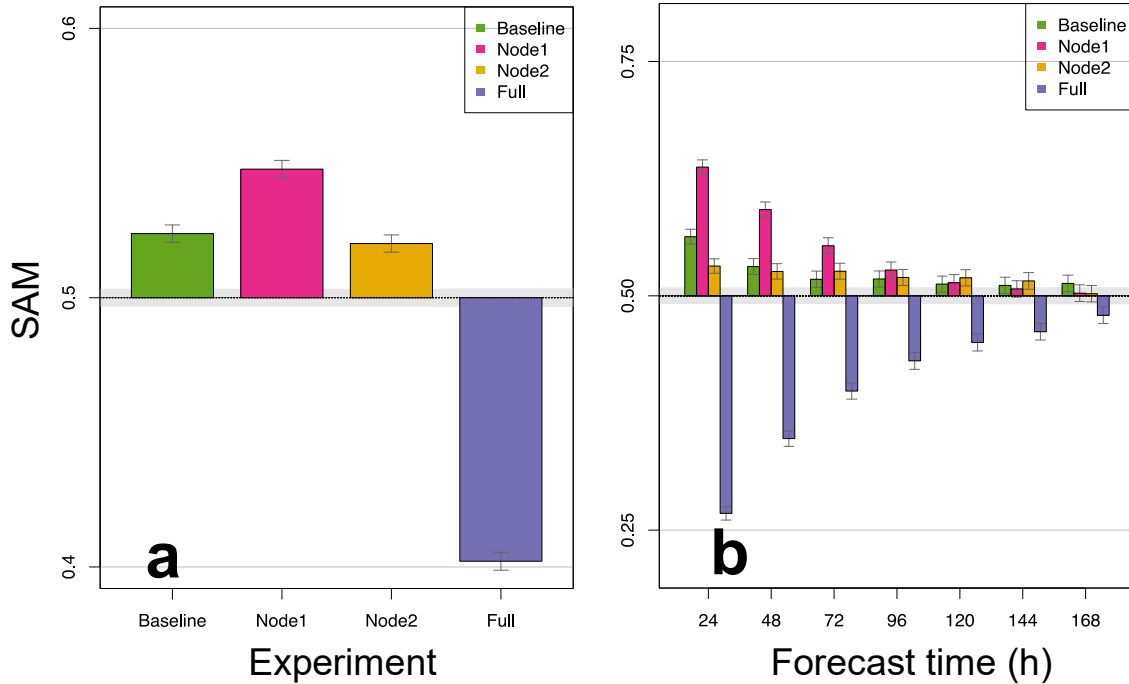
283



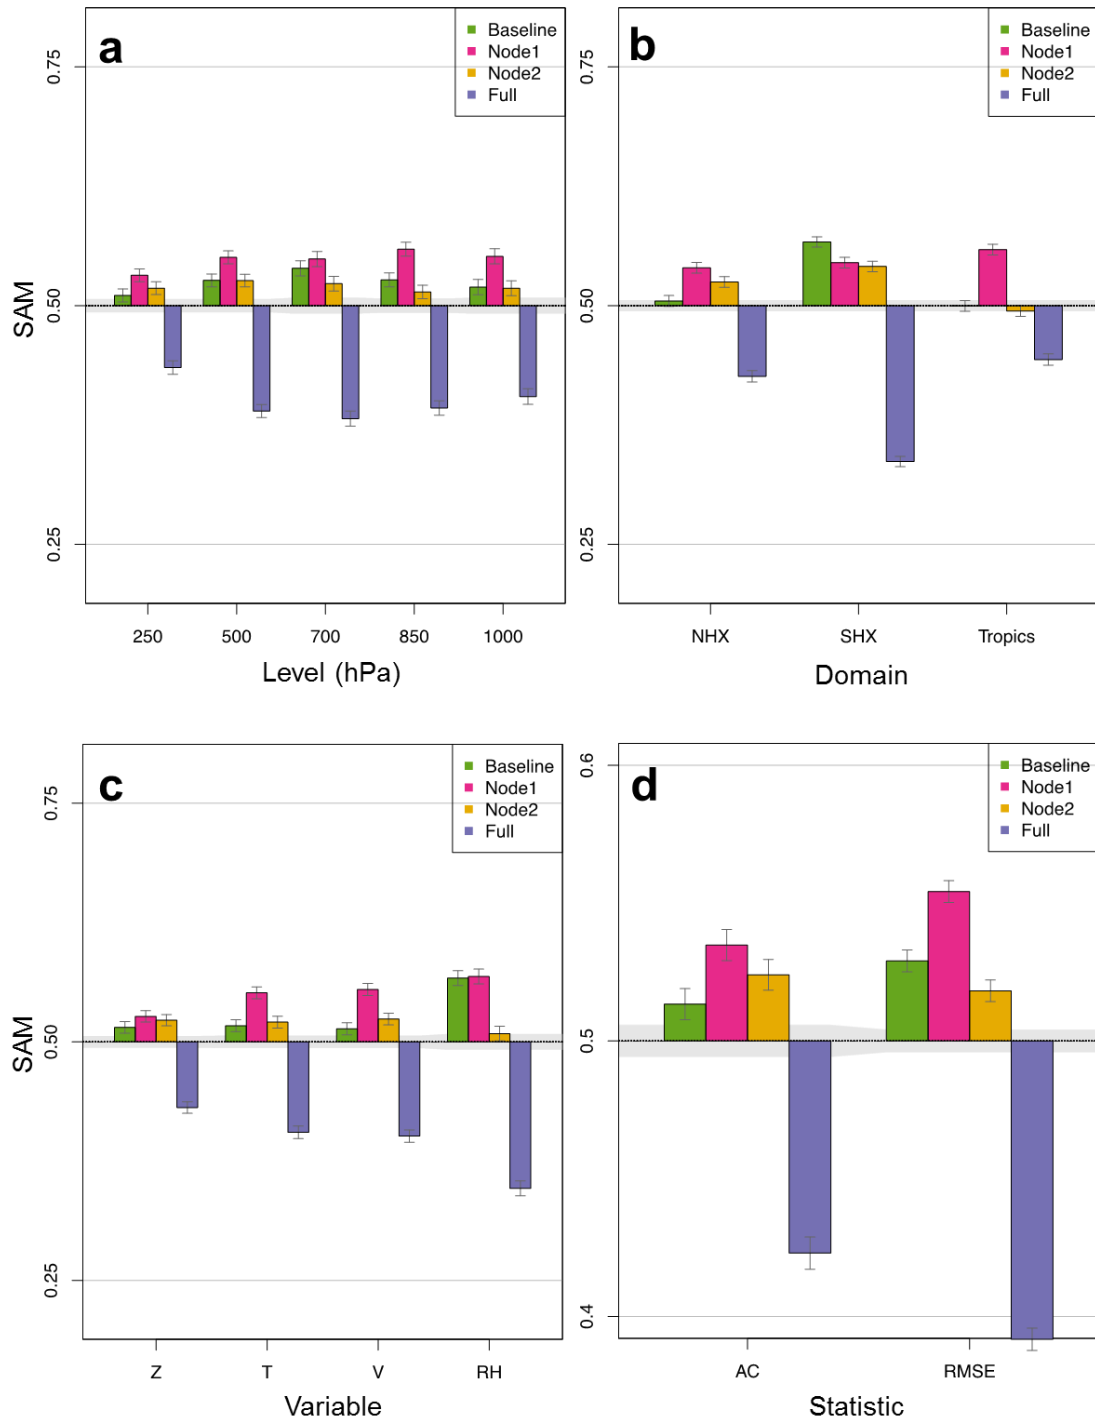
284
 285 Fig. 5. Anomaly correlations (ACs) for 500-mb geopotential height for the northern
 286 hemisphere extratropics (NHX) as a function of forecast time averaged over all
 287 verification times, in terms of a) AC for each of the four experiments (color coded), b)
 288 each experiment relative to Baseline (color bars), and c) each experiment relative to
 289 Baseline as a percentage. The grey curves in frame a represent all Baseline forecasts
 290 individually. Error bars are plotted at the 95% confidence interval in frames b and c.



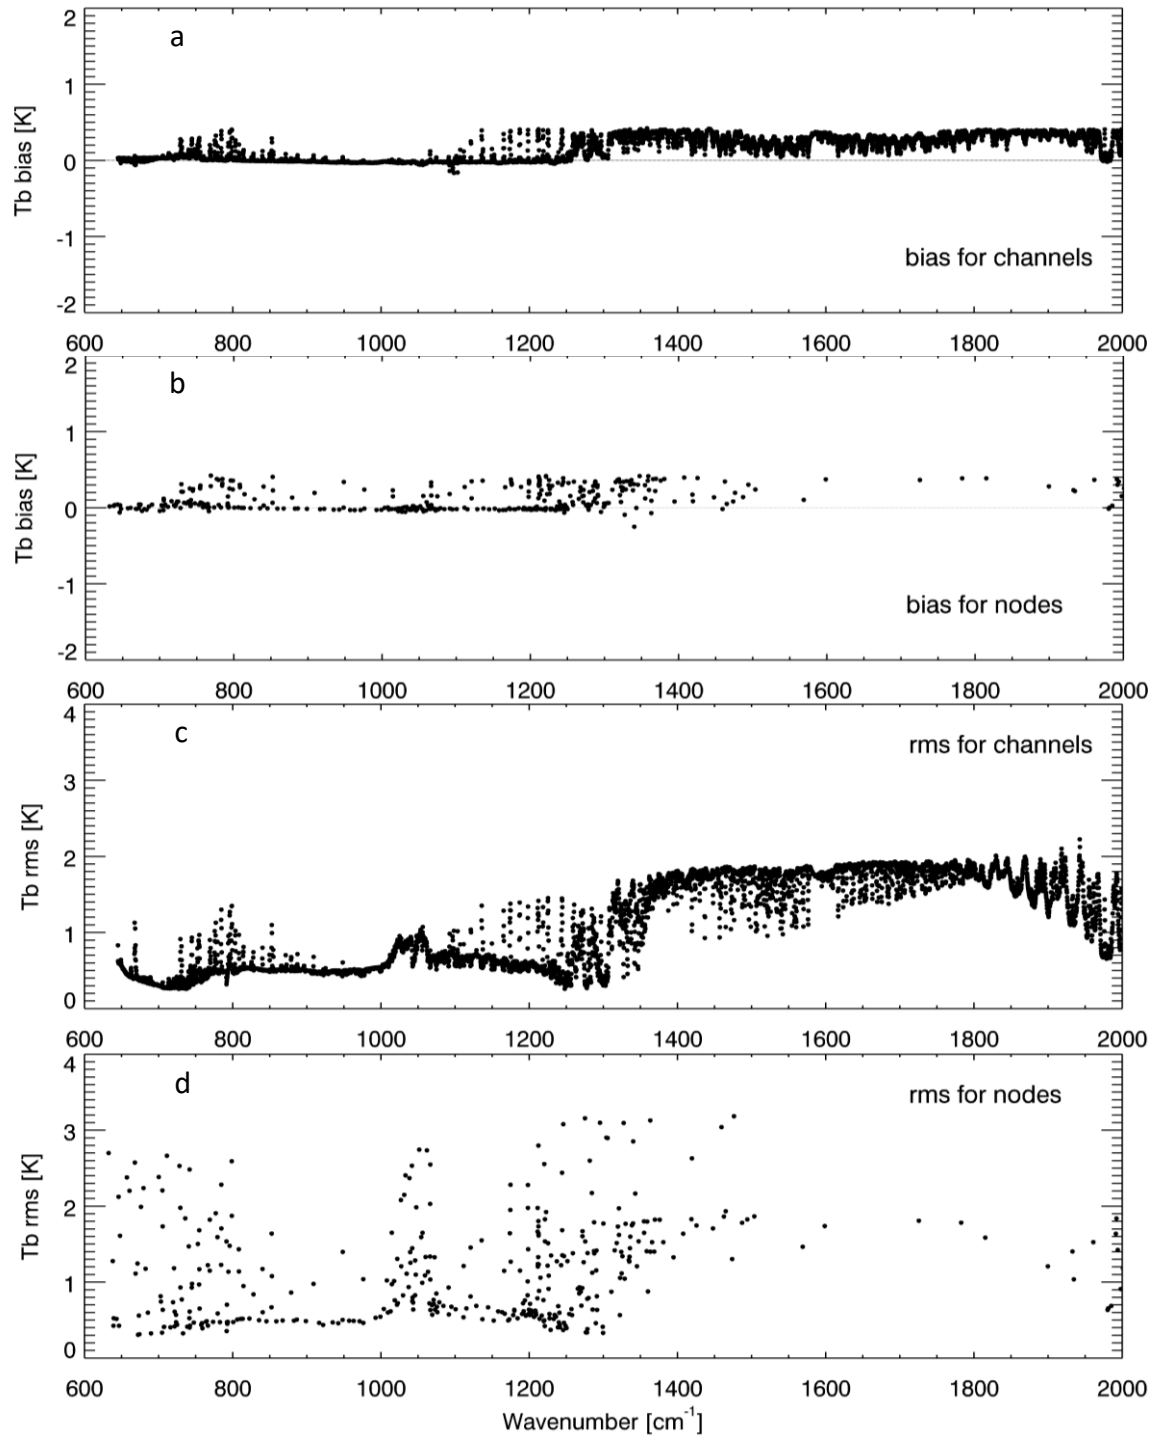
291
 292 Fig. 6. Impact relative to Baseline of selected error metrics as a percentage of the
 293 Baseline error. The Student-t quantile for each δ impact is printed in parentheses and
 294 color coded, with the shades of blue indicating improvement and the shades of red
 295 indicating degradation. In the axis label, symbols Z, V, and RH represent geopotential
 296 height, wind speed, and relative humidity, respectively, and the following numbers (250,
 297 850, 500) indicate the pressure level (hPa). NHX and SHX represent extra-tropical
 298 northern and southern hemispheres, respectively.
 299



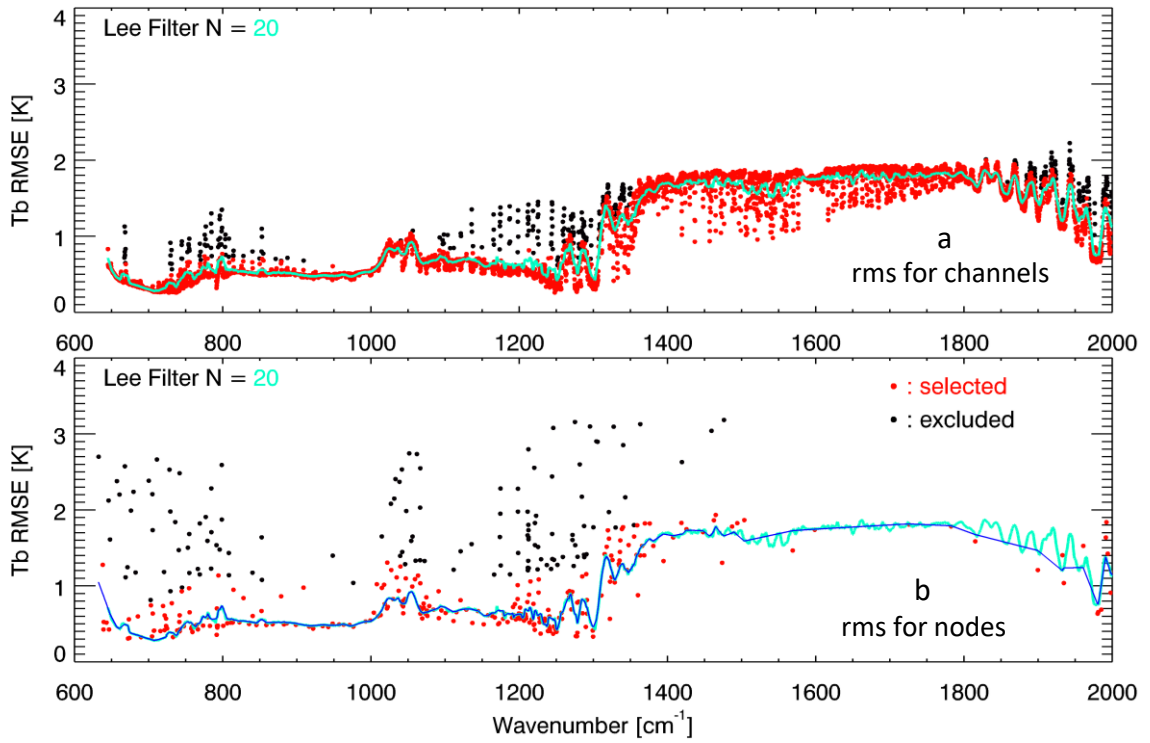
300
 301 Fig. 7. Summary assessment metrics (SAMs) a) for each experiment globally and b) for
 302 each experiment as a function of forecast time (h). Here, a 95% uncertainty band is
 303 shaded surrounding the SAM=0.5 line, under the null hypothesis that there is no
 304 difference between centers for this metric. The estimated uncertainty at the 95% level is
 305 indicated by small error bars at the ends of the color bars, which are anchored at the
 306 expected value (0.5).
 307



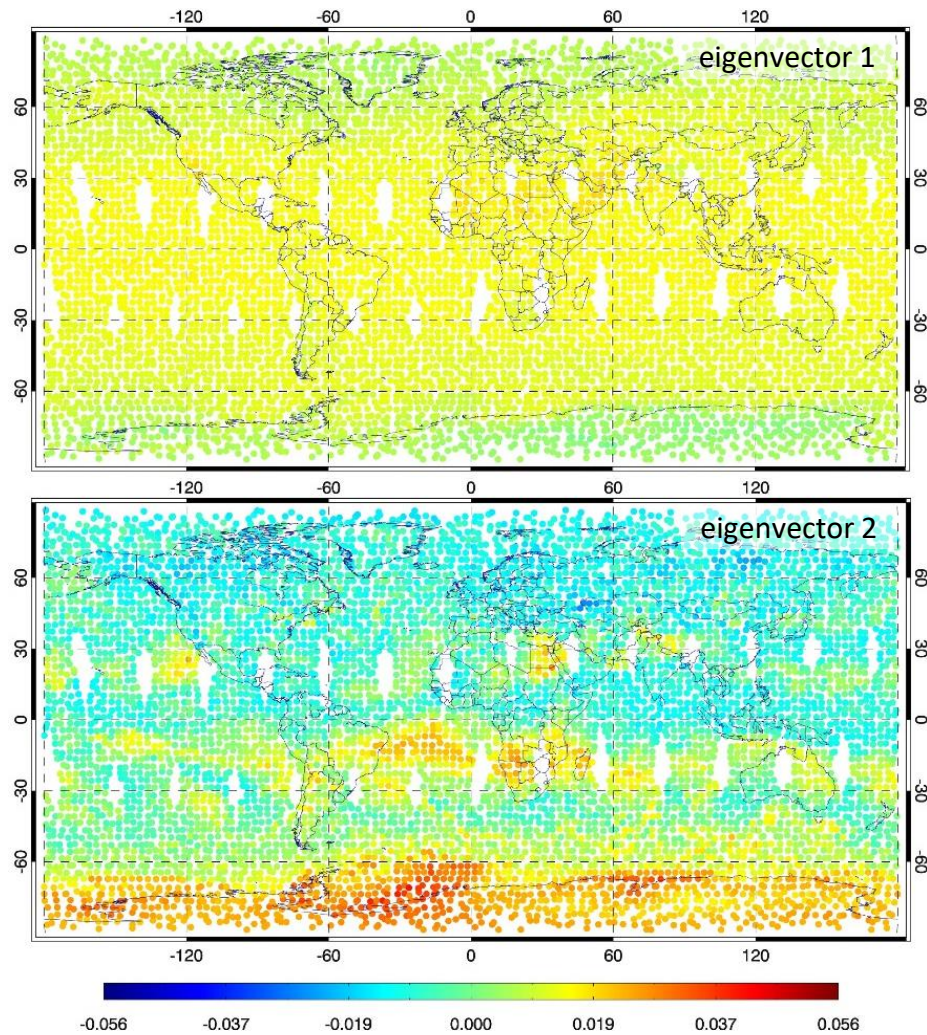
308
 309 Fig. 8. Variation of SAM by experiment and by a) pressure level, b) geographic domain,
 310 c) variable, and d) statistic. Colors and markers are as in Fig. 7, and the dimension
 311 abbreviations are as in Table 2.



312
 313 Fig. B1. O–B bias and rms for the full channel set (bands 1 and 2) and for OSS nodes,
 314 based on data from 11 July 2015 00 UTC – 12 July 2015 18 UTC.



315
 316 Fig. B2. O–B rms for a) channels as in Fig. B1c and b) nodes as in Fig. B1d. In both
 317 frames, the cyan curve marks the result of the Lee filter (as applied to data from
 318 channels) and the red and black indicate selected and excluded channels or nodes,
 319 respectively, according to the difference from the Lee filter result. The filter window size
 320 covered 20 points to each side of the smoothed point. For clarity, the channel Tb RMSE
 321 is plotted with smaller symbols.

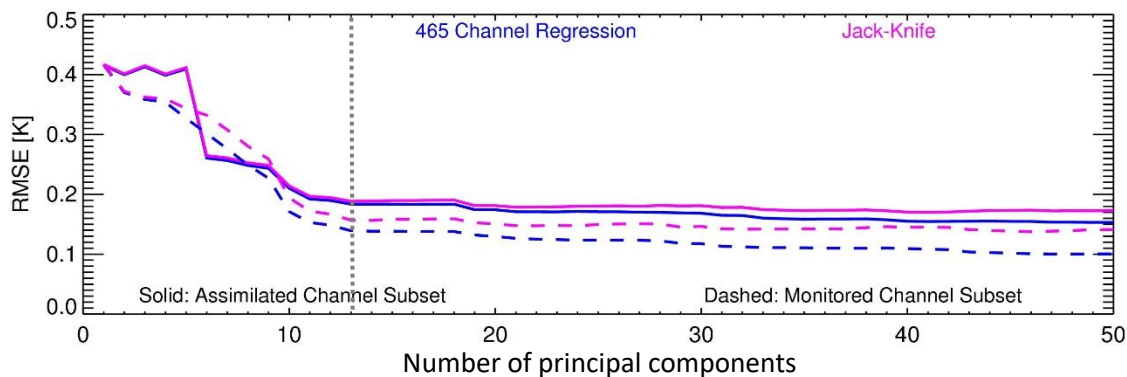


322


323 Fig. C1. Eigenvectors 1 (top) and 2 (bottom) of the FOV-to-FOV covariance matrix
 324 derived from 465 channel set of brightness temperature samples. The amplitude of each
 325 element of the eigenvector is plotted at the location of the corresponding FOV.

326 Amplitudes are normalized so the sum of squares is 1.

327

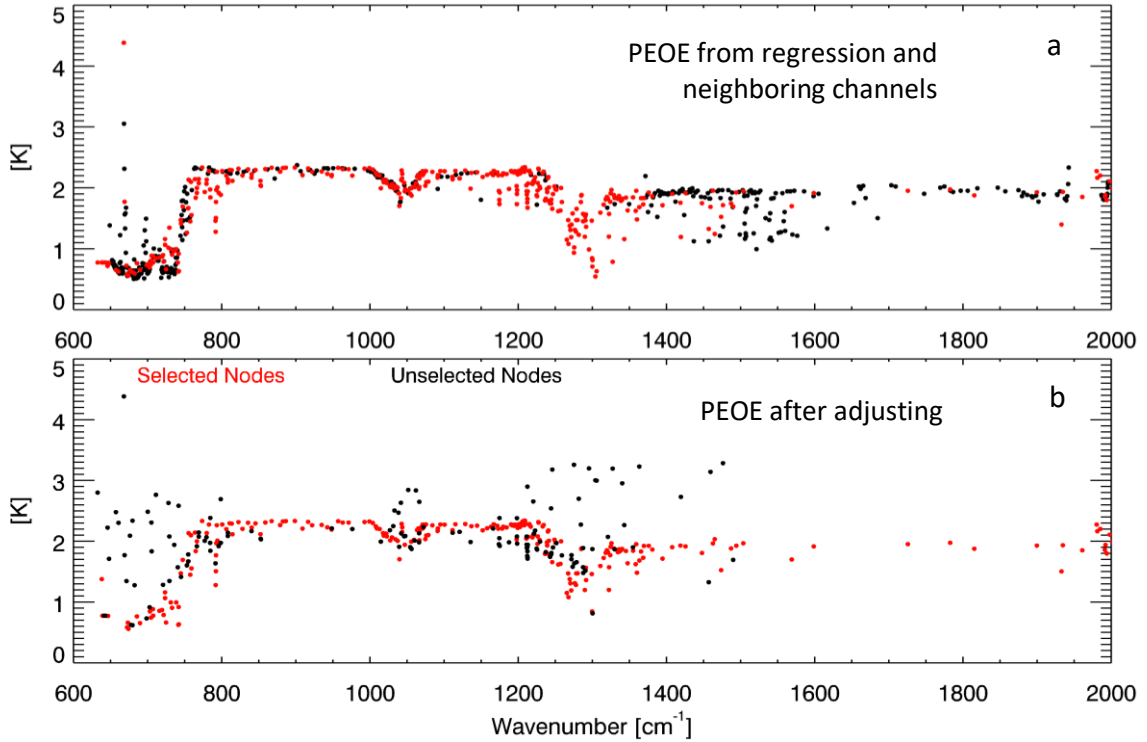


328
 329 Fig. C2. The regression fit rms error, comparing the predicted PEOE values against the
 330 actual baseline GSI values, for the 465 channels in the training set, plotted as a function
 331 of the number of PCs used as predictors. The blue curves represent the performance on
 332 the dependent (training) dataset and the magenta curves show the result of applying the
 333 standard jack-knife method to estimate the performance on independent data. Results are
 334 shown separately for the 150 assimilated channels (solid) and the 315 monitored channels
 335 (dashed). The vertical dotted line indicates the selected number of PCs (13).
 336



337
 338 Fig. C3. PEOE for IASI channels before (top) and after (bottom) adjusting relative to
 339 O–B rms. In the top frame (a), channels that were among the 465 channels designated for
 340 assimilation or monitoring in the baseline GSI are marked cyan and the other channels
 341 are marked red. For the 465-channel subset, the GSI baseline values are marked black. In
 342 the lower frame (b), the PEOE results from the 5014 channels selected for assimilation by
 343 our filtering approach (Sec. 3d) are marked red and the rest are black.

344



345 Fig. C4. PEOE for IASI OSS nodes before (a) and after (b) adjusting relative to O-B
 346 rms. The data in the upper frame were derived from a combination of values from
 347 neighboring channels and results from regression applied to node brightness
 348 temperatures, as explained in the text. For reference, PEOE data for the 465 channels
 349 designated for assimilation or monitoring in the baseline GSI are shown the same as in
 350 Fig. C3. In the lower frame, the PEOE results from the 265 nodes selected for
 351 assimilation by our filtering approach (Sec. 3d) are marked red and the rest are black.
 352