Met Office

Sensitivity of radio occultation-based dry temperature retrievals to upper-level information and its relevance to radiosonde bias corrections

8 November 2016

Chris Burrows (Met Office), Sean Healy (ECMWF)

Forecasting Research Technical Report 615

Contents

Abstract

Retrieving temperature values from radio occultation (RO) bending angles can be performed in the absence of water vapour. The so-called dry temperature retrieval requires a series of assumptions and *a priori* information. For example, the bending angle profile must be extrapolated above the highest level (typically 60km) in order to produce refractivity values, and from these refractivities an assumed temperature value must be used to calculate hydrostatically the pressure and thence temperature profiles. These assumptions and the fact that the retrieval steps involve vertical integrals result in information from high up propagating down to lower levels in the resulting dry temperatures. A linear version of the retrieval chain is used to assess the sensitivity to this upper information, and from these results it is shown that it is reasonable to truncate profiles of mean bending angle background departures when propagating these to temperature space, with the intention of using these to diagnose biases in radiosonde temperatures.

1 Introduction

When raw bending angle measurements are processed into dry temperature values, some additional prior information is required as integrals in the retrieval algorithm have their upper limit at infinity, but the bending angle profile typically extends to just 60km. The retrieved temperatures at all heights will be affected by these assumptions, so here we investigate the extent to which this occurs in the method (used in Tradowsky [2016], Tradowsky et al.) to reduce the undesirable influence of these arbitrary assumptions.

2 Retrieval chain

2.1 Abel transform

The first step to calculate dry temperatures is to calculate the refractive index, n, and refractivity, N, (where $N = 10^6(n-1)$) using the Abel transform:

$$n(x) = \exp\left(\frac{1}{\pi}\int_x^\infty \frac{\alpha(a)}{\sqrt{a^2-x^2}}da\right) \tag{1}$$

By noting that $a + x \simeq 2x$ and $\ln(n) \simeq n - 1$, the refractivity can be written as:

$$N(x) = \frac{10^{-6}}{\pi\sqrt{2x}}\int_x^\infty \frac{\alpha(a)}{\sqrt{a-x}}da \tag{2}$$

Following Syndergaard [1999], we discretise this integral by assuming linear variation of bending angle with impact parameter. This results in a linear equation of the form $\boldsymbol{N} = \boldsymbol{A}\boldsymbol{\alpha}$, where $\boldsymbol{\alpha}$ is a vector of bending angle values in a given profile, $\boldsymbol{A}$ is the matrix that defines the linear mapping that approximates Equation 2 and $\boldsymbol{N}$ is the resulting vector of retrieved refractivity values in the profile.

Note that the upper limit of the integral is ∞. The approach used in NRT processing is to blend the observed bending angles with climatological values to produce a smooth function $\alpha(a)$ which extends above the highest observation. In Tradowsky [2016], climatology was not used, and the uppermost bending angle value was used to initialise an exponentially decaying function from this height to infinity. This implies a contribution from this uppermost bending angle to all refractivity values below the uppermost bending angle level. The contribution to the refractivity at height x, due to the extrapolated bending angle profile is denoted as $N_{above}(x)$. This can be represented in terms of a Gaussian error function:

$$N_{above}(x) = \frac{10^{-6}}{\sqrt{x}}\sqrt{\frac{h}{2\pi}}\alpha(a_{top})\exp\left(\frac{a_{top}-x}{h}\right)\left[1-\mathrm{erf}\left(\sqrt{\frac{a_{top}-x}{h}}\right)\right] \tag{3}$$

The exponential scale height, h, is calculated for an isothermal atmosphere at 60km, using the co-located model temperature:

$$h = \frac{RT}{g} \tag{4}$$

Note that this is the same inverse scale height that represents the density field, and hence refractivity. This can be understood by taking Equation 7 of Healy and Thepaut [2006] at the limit where x_i is the impact parameter and x_{i+1} is ∞, with constant inverse scale height, k:

$$\alpha(a) = 10^{-6}\sqrt{2\pi a k}N(x=a) \tag{5}$$

Because $\sqrt{a}$ is almost constant due to the large, fixed offset of the Earth's radius of curvature, the height variation of α is dominated by the refractivity at the tangent height, and this varies exponentially due to hydrostatic balance of an ideal gas, given the isothermal constraint.

This extrapolation means that, in computing the refractivity, the bending due to the entire atmosphere above the highest observation is being approximated by two pieces of information: the highest bending angle observation and the co-located model background (i.e. 6-hour forecast) temperature at this height.

2.2 Dry pressure

Once the refractivity profile has been calculated, the pressure can be inferred, but more assumptions are required. In the absence of water vapour, the refractivity is proportional to the density, and the constant of proportionality is the coefficient of the dry term of (for example) the Smith and Weintraub expression times the gas constant; $\rho = N/cR$. The hydrostatic equation then becomes:

$$dP = -\frac{N(z)}{Rc}g\,dz \tag{6}$$

Assume exponentially varying refractivity in each layer:

$$N(z) = N_i \exp\left(-k_i \left(z - z_i\right)\right) \tag{7}$$

For each layer, the contribution to the dry pressure is, therefore:

$$\Delta P_i = \frac{g}{Rc} \frac{N_i - N_{i+1}}{\ln\left(\frac{N_i}{N_{i+1}}\right)} \left(z_{i+1} - z_i\right) \tag{8}$$

The pressure at the height of a given refractivity value is computed by summing the ΔP_i values for all layers above — downwards from the top. At this stage, the integral is indefinite, so a "known" pressure is required as a constant of integration. This value of pressure at the top is calculated from the top level refractivity (computed above) and the model temperature at that level (as used previously to calculate the exponential scale height):

$$P_{60km} = \frac{N_{60km} T_{60km}}{c} \tag{9}$$

The retrieved pressure at every level will contain information from the *a priori* temperature and the top level pressure, which itself is calculated using the same *a priori* temperature used to obtain the bending angle scale height for the extrapolation to infinity. The important difference with this step compared to the other levels is that the value at the uppermost level defines the entire atmospheric state above this height. Therefore, this level has a much larger impact on the retrieved pressures than any other levels.

2.3 Dry temperature

Once the pressure has been calculated for each level, the temperature is easily obtained using the refractivity values:

$$T(z) = c \frac{P(z)}{N(z)} \tag{10}$$

The entire chain from bending angles to dry temperature has been linearised. This allows the sensitivity of the resulting temperatures to their inputs to be investigated.

3 Jacobians

In this context, Jacobians are the matrices of partial derivatives of the outputs of the retrieval chain with respect to their inputs. These are not quite the same as the adjoints of the Jacobians of the linearised observation operator as slightly different assumptions are made in the retrieval. These allow the propagation of increments through a linear version of the retrieval chain (as is done in error propagation studies). Furthermore, the columns of the Jacobians can be plotted to see, for

example, how sensitive the resulting temperatures are to the input bending angles. In a simplified form, the Jacobian of the dry temperature with respect to bending angle is (using the chain rule):

$$\frac{\partial \boldsymbol{T}}{\partial \boldsymbol{\alpha}} = \left[\frac{\partial \boldsymbol{T}}{\partial \boldsymbol{P}}\frac{\partial \boldsymbol{P}}{\partial \boldsymbol{N}} + \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{N}}\right]\frac{\partial \boldsymbol{N}}{\partial \boldsymbol{\alpha}} \tag{11}$$

The tangent linear code can be made to produce Jacobians by providing a linearisation state and inputting columns of the identity matrix successively. For the Abel transform and dry pressure integrals, the separate Jacobians are upper triangular as the retrieval at a given height only depends on values above this height. When a Jacobian has been computed, its rows can be interpreted as the partial derivatives of the "output" at a single level with respect to the "inputs" on all levels.

The Jacobian matrix for refractivity is shown in Figure 1. This is obtained by coding the retrieval operator from Section 2.1, then coding its tangent linear and finally, inputting columns of the identity matrix in order to obtain the full Jacobian matrix.

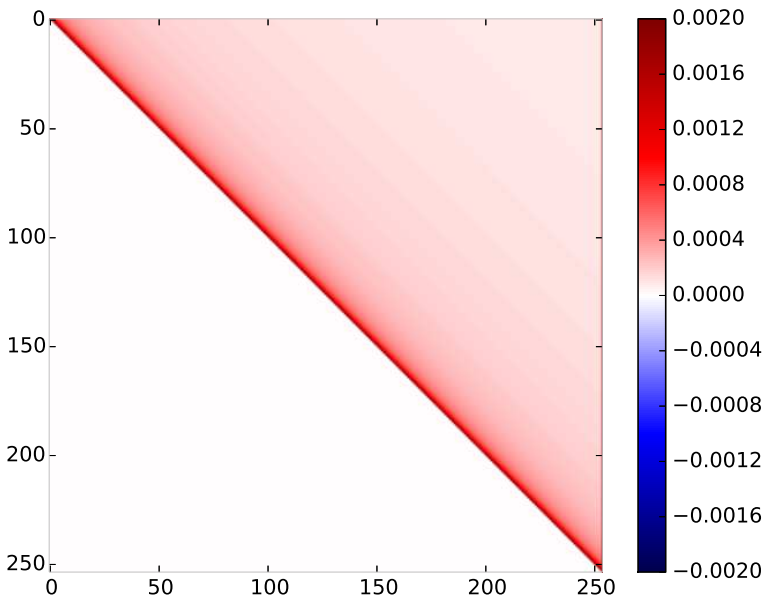

Figure 1: Jacobian matrix of $\partial N/\partial \alpha$ the axes are element numbers (zero is approximately 9 km and 250 is approximately 60km), with larger numbers representing higher levels; these levels are equally spaced from 9km to 60km. The scale is saturated, and note that the largest values occur in the final column. The colour scale is in units of N-units/rad.

The refractivity Jacobian is upper triangular and has large values close to the diagonal which reduce rapidly with height (with the exception of the final column), indicating that if the bending angle close to the tangent point is perturbed, the resulting refractivity will be altered much more than if a higher bending angle was perturbed. This is mostly a feature of the "kernel" of the Abel transform (i.e. the denominator $\sqrt{a^2 - x^2}$). The large values in the final column are due to the

uppermost bending angle being used to represent the entire atmosphere above, by means of the extrapolation described in Section 2.1.

The pressure Jacobian, with respect to bending angle, is shown in Figure 2 (computed the same way as Figure 1, with the addition of the equations in Section 2.2). This shows that the retrieved pressure at a given height is not significantly affected by those bending angle values close to that height — rather, the uppermost bending angles have a larger contribution. This is due to repeated integration over all refractivities, each of which contain information from all levels above.

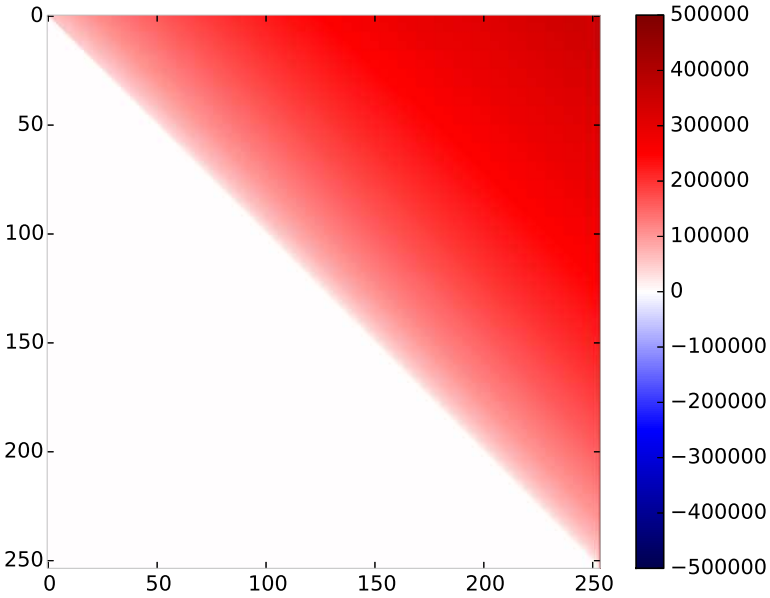

Figure 2: Jacobian matrix of $\partial P/\partial \alpha$ the axes are element numbers (zero is approximately 9 km and 250 is approximately 60km), with larger numbers representing higher levels; these levels are equally spaced from 9km to 60km. The scale is saturated, and note that the largest values occur in the final column. The colour scale is in units of Pa/rad.

All the elements of both the refractivity and pressure Jacobians are either zero or positive. This means that increased bending angles at any level will lead to increased refractivities and pressures. The temperature Jacobian is more complicated because Equation 10 has refractivity in the denominator, so negative elements may be present in the matrix. This temperature Jacobian is plotted in Figure 3 (computed the same way as Figure 2, with the addition of the equations in Section 2.3).

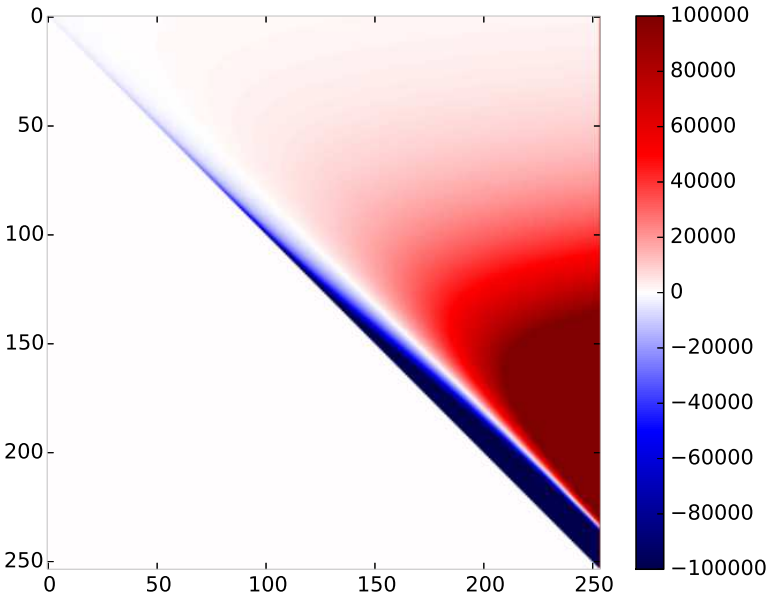

Figure 3: Jacobian matrix of $\partial T/\partial \alpha$ the axes are element numbers (zero is approximately 9 km and 250 is approximately 60km), with larger numbers representing higher levels; these levels are equally spaced from 9km to 60km. The scale is saturated, and note that the largest values occur in the final column. The colour scale is in units of K/rad.

Unlike bending angle, refractivity and pressure, the temperature does not vary exponentially with height, so it is worth noting that a unit change in α will be much more significant high up than low down as the bending angles vary by several orders of magnitude. This explains the behaviour close to the diagonal; at low levels, a unit change in bending angle at any other height has little impact on the temperature, but high up where the bending angles themselves are small, a unit change can produce a large impact in the temperatures. Close to the diagonal, the elements are negative. This arises from the term $\partial \mathbf{T}/\partial \boldsymbol{N}$ in Equation 11 being negative (N is in the denominator of Equation 10). The sharp spike near the diagonal is due to the kernel of the Abel transform. As height increases, the sign changes and the positive elements due to the pressure Jacobian dominate. These are larger for temperatures at higher levels; again, this is due to the relative magnitude of the bending angle at different heights, and the competing influence of the pressure and refractivity Jacobians. This matrix also has very large values in the final column due to repeated use of extrapolated information.

Line plots of a selection of rows from each of these three Jacobians is shown in Figure 4.

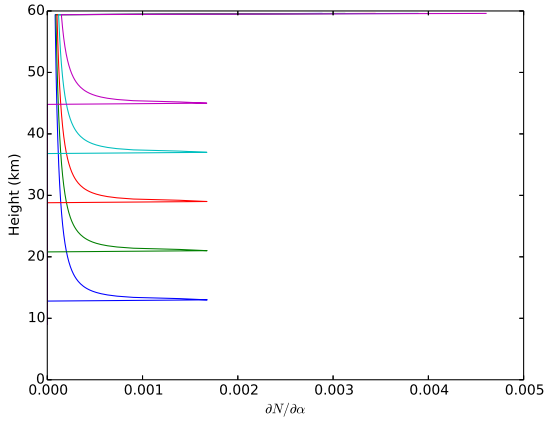

(a) $\frac{\partial N}{\partial \alpha}$

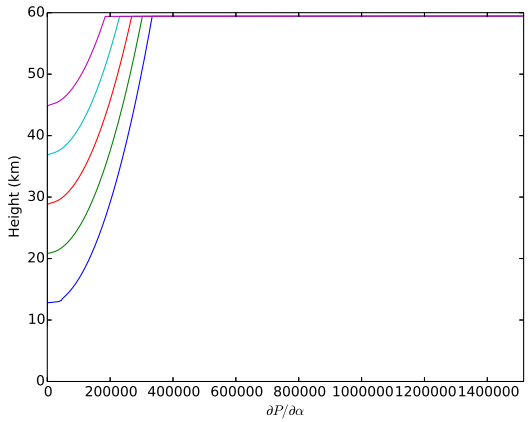

(b) $\frac{\partial P}{\partial \alpha}$

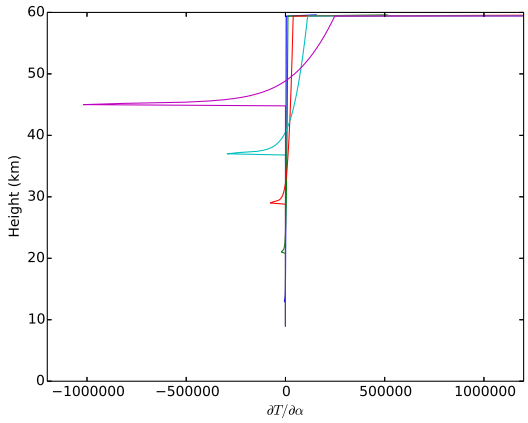

(c) $\frac{\partial T}{\partial \alpha}$

Figure 4: Selected rows from the three Jacobian matrices in Figures 1 to 3. The x-axis ranges of (a) and (b) have been restricted to show the most detail. In both cases the value at 60km reaches 1.4×10^7.

It can be seen that the temperatures at all heights are very sensitive to the top bending angle value, and also to those close to the top. The temperature values at low levels have this sensitivity to the upper bending angles, but this has a small influence on the retrieved temperatures in reality due to the magnitude of the bending angles being much greater low down.

4 Linear propagation of increments

For small bending angle increments, the Jacobian matrix $\partial T/\partial \alpha$ can be used to propagate these increments into temperature space. Let $\mathbf{A}$ be the linearised version of the retrieval chain A (i.e. the temperature Jacobian):

$$A(\boldsymbol{\alpha}_2) - A(\boldsymbol{\alpha}_1) \simeq \mathbf{A}(\boldsymbol{\alpha}_2 - \boldsymbol{\alpha}_1) \tag{12}$$

which can be applied as:

$$\delta\mathbf{T} = \mathbf{A}\delta\boldsymbol{\alpha} = \frac{\partial \mathbf{T}}{\partial \boldsymbol{\alpha}}\delta\boldsymbol{\alpha} \tag{13}$$

where bold $\mathbf{T}$ and $\boldsymbol{\alpha}$ indicate that the quantities are vectors. The Jacobian must be linearised about a bending angle profile, and it is reasonable to select either $\boldsymbol{\alpha}_1$ or $\boldsymbol{\alpha}_2$ for this purpose. Now, consider a temperature increment at a given level (i.e. a single element of $\delta\mathbf{T}$). In matrix form, this scalar is given as:

$$\begin{aligned} \delta T_i &= \frac{\partial T_i}{\partial \boldsymbol{\alpha}}\delta\boldsymbol{\alpha} = \begin{bmatrix} \frac{\partial T_i}{\partial \alpha_1} & \frac{\partial T_i}{\partial \alpha_2} & \cdots & \frac{\partial T_i}{\partial \alpha_n} \end{bmatrix} \begin{bmatrix} \delta\alpha_1 \\ \delta\alpha_2 \\ \vdots \\ \delta\alpha_n \end{bmatrix} \\ &= \sum_{k=1}^{n} \frac{\partial T_i}{\partial \alpha_k}\delta\alpha_k \end{aligned} \tag{14}$$

This summation weights the bending angle increments by the sensitivity, i.e. the Jacobian elements. This discrete, linear formulation allows a degree of control to be applied which cannot be done in the nonlinear case. We consider temperature increments based on bending angle mean background departures from the whole of 2015 in the vicinity of radiosonde station 10393 (Lindenberg, Germany). These mean bending angle departures are the mean values of the observed bending angles minus the forward-modelled equivalents from co-located short-range forecasts, with the observation being collected from within 500 km of the radiosonde station. This is calculated for temperatures at two heights; 15km and 25km. The elements of $\delta\boldsymbol{\alpha}$ and $\partial T/\partial\boldsymbol{\alpha}$ are plotted in Figure 5 along with the separate components of $(\partial T_i/\partial\boldsymbol{\alpha})\,\delta\boldsymbol{\alpha}$ which should be summed, as in Equation 14 to produce the temperature increment.

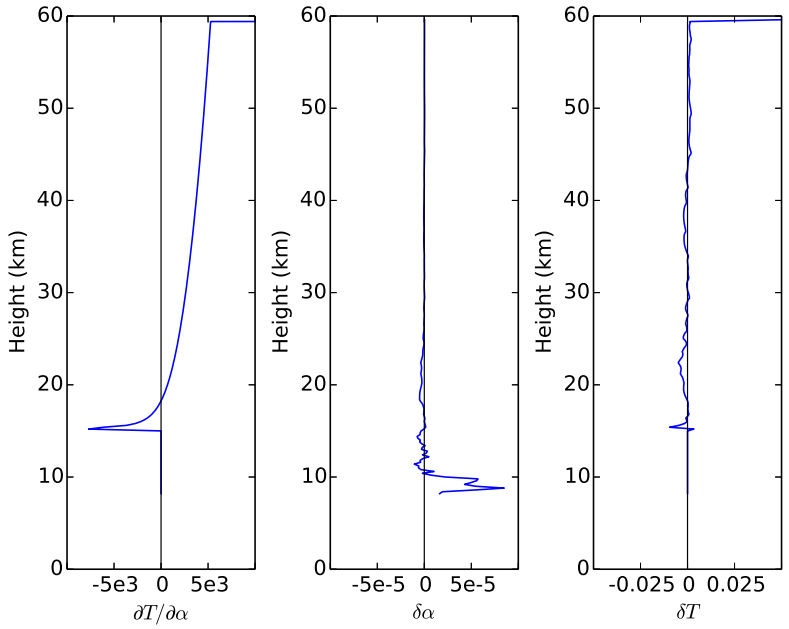

(a) 15km

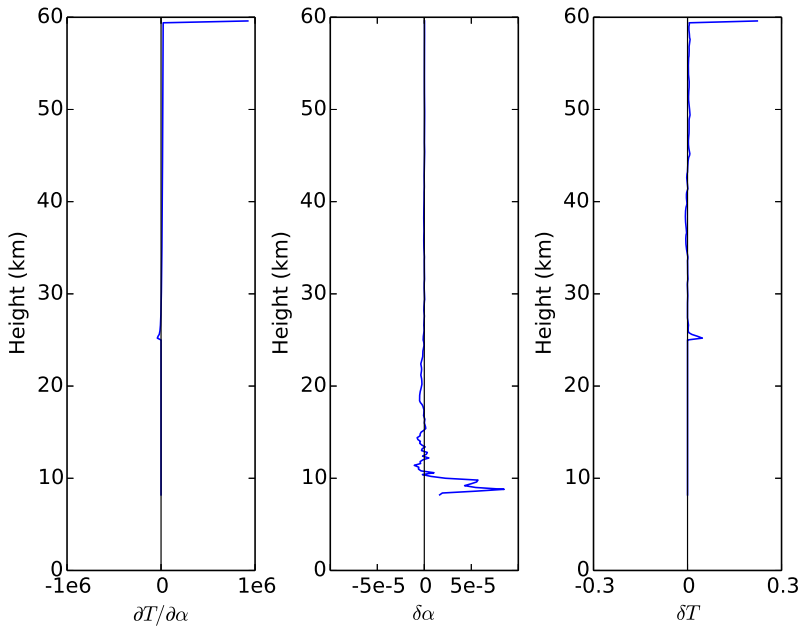

(b) 25km

Figure 5: Jacobian rows, bending angle increments and the resulting terms that constitute the temperature increments δT at 15km and 25km.

Figure 5 shows that the temperature increments are comprised of information at all levels above the height of the temperature being retrieved. For both heights shown here, there is a spike close to the chosen height which tails off rapidly, above which the values increase with height, with a large

spike at the uppermost level. It is instructive to view the cumulative (from below) version of this plot to understand how significant each additional term of the summation (Equation 14) is. This is shown in Figure 6.

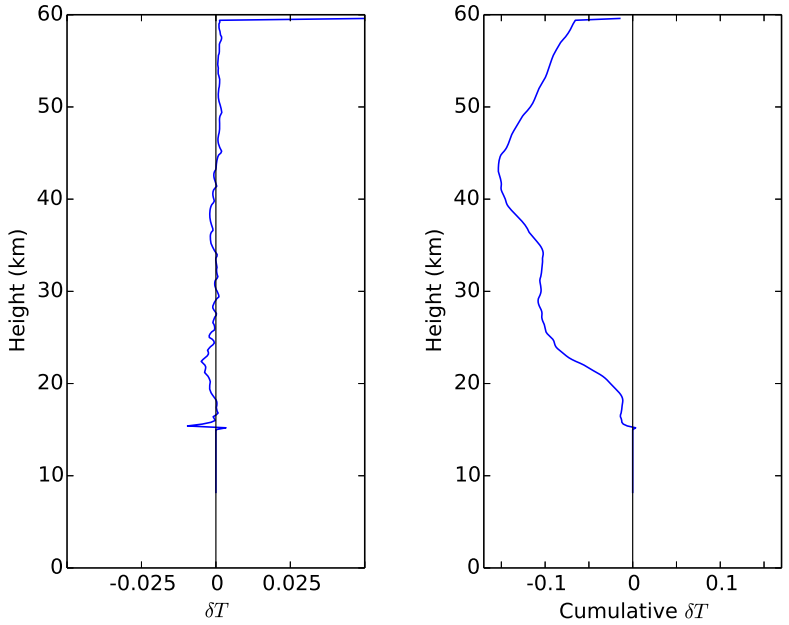

(a) 15km

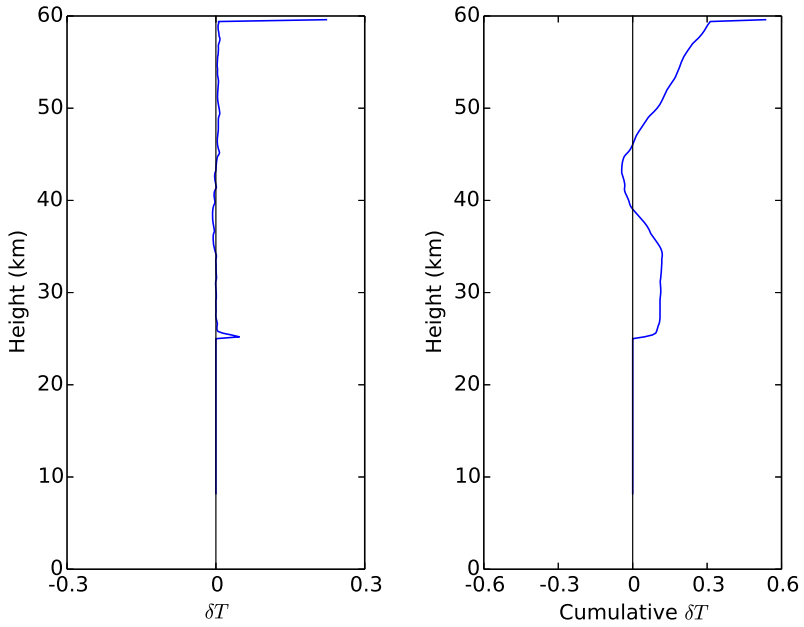

(b) 25km

Figure 6: Cumulative (from below) value of δT for two heights, summed from below, i.e. $\delta T_{cumul}(z') = \int_0^{z'} \delta T(z) dz$.

The “cumulative” plots show that as additional terms are added from higher bending angle increments, the retrieved temperature increments do not converge — they continue to change significantly, particularly when the term from the upper level is added. Note that in both cases, the variation in δT reduces temporarily at around 35km. Could it be that the contribution above this height is a combination of residual noise and *a priori* information, rather than information from the observations themselves? If so, it would seem reasonable to ignore the values above this height in the summation; this is equivalent to setting the rightmost columns in the Jacobian (Figure 3) to zero. It is known that the bending angle background departures above around 35km increase in magnitude and change sign. This is at least partly due to the model bias at these levels, and is different for Met Office and ECMWF model fields as seen in Figure 7. These values may be large enough to propagate through the retrieval and produce large contributions to the retrieved temperature increments. Similar plots have been produced from RO data in the vicinity of other radiosonde stations (not shown here) and although the levelling off at around 35km is present in several cases, others contain some variation, though this variation is significantly smaller than that above 40km. For this reason, it is proposed that in the calculation of δT, the bending angle departures above 35km are set to zero, which is equivalent to setting the corresponding columns of the Jacobian $\partial T/\partial \alpha$ to zero. Furthermore (as shown in Tradowsky [2016], Figure 4.3-4.5), the large random errors in the bending angle departures above 35km mean that the integrals in the retrieval chain introduce vertically correlated errors. By setting the bending angle departures above 35km to zero, the random errors below this height are reduced.

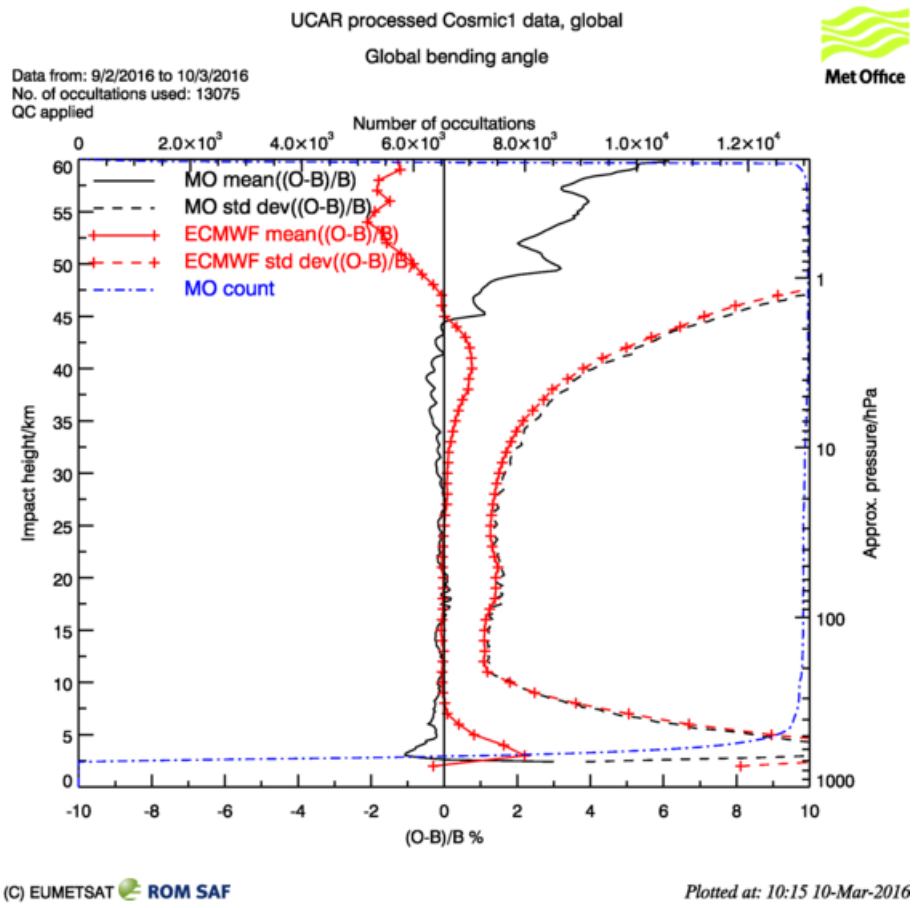

Figure 7: Bending angle departure statistics for a month of COSMIC-1 data, showing the difference in bias between the Met Office and ECMWF above around 30km. Note also the exponentially increasing standard deviation at these levels.

A different interpretation of the application of this cut-off may be reached (K. Lauritsen and S. Syndergaard, priv. comm.). That is, by setting the bending angle departures to zero, this implies replacing the observation values above this height with the background values. It is true that this would produce the same result, but so would replacing the background values with the observation values. The decision to apply the cut-off is based on our understanding that the *product* of the Jacobian and the bending angle departures is unduly large at these heights because of the reliance on a priori, noisy data and simplifications in the retrieval. Near-real-time refractivity values use statistical optimisation (i.e. the blending of observed and climatological values) to reduce these factors. In this study however, the use of climatology is undesirable, so applying a cut-off is a way to mitigate these factors in a controlled way.

5 Radiosonde bias correction

A new method of producing a bias correction for a given radiosonde station using RO data as an unbiased reference is described in Tradowsky [2016]. This involves subtracting mean RO background departures (in temperature space) in the vicinity of the radiosonde station from the radiosonde departures. The use of model fields mitigates the need for co-locations between RO profiles and radiosonde ascents, with the assumption that the background information in both sets of departures cancels to first order. This equation describes the concept, where primed quantities are observations and unprimed quantities are forward-modelled equivalents:

$$\overline{\mathbf{T}'_{RS} - \mathbf{T}'_{RO}} \simeq \overline{\mathbf{T}'_{RS} - \mathbf{T}_{RS}} - \overline{\mathbf{T}'_{RO} - \mathbf{T}_{RO}} \tag{15}$$

As shown above, the RO temperature departures used here are in fact the bending angle departures which have been operated on by the linearised version of the retrieval chain:

$$\overline{\mathbf{T}'_{RS} - \mathbf{T}'_{RO}} \simeq \overline{\mathbf{T}'_{RS} - \mathbf{T}_{RS}} - \mathbf{A}\overline{\boldsymbol{\alpha}' - \boldsymbol{\alpha}} \tag{16}$$

At a particular height, $z = z_{ob}$, the radiosonde bias is given by:

$$\begin{aligned} \overline{T'_{RS}(z = z_{ob}) - T'_{RO}(z = z_{ob})} \quad \simeq \quad & \overline{T'_{RS}(z = z_{ob}) - T_{RS}(z = z_{ob})} - \\ & \mathbf{A}_{z_{ob}} \overline{\alpha'(z = z_{ob} \rightarrow \infty) - \alpha(z = z_{ob} \rightarrow \infty)} \end{aligned} \tag{17}$$

The radiosonde temperature departures are straightforward to compute as the radiosonde temperature at a given height has an equivalent that can be interpolated from two adjacent model levels. The RO geometry means that for a given tangent point, the bending angle is influenced by the atmospheric state along the entire ray path between the satellites. Therefore, even before the linear retrieval is applied, the observed bending angle at a given tangent height contains information about the atmosphere at all levels above, and similarly, the forward-modelled equivalent contains model information above the tangent point. For the approximate equality above to hold, the mean backgrounds should cancel, but the RO backgrounds for a given height contain different model information from the radiosonde backgrounds. Fortunately, the bending angle forward model has a weighting function that peaks very sharply close to the tangent point, so to a good extent the forward-modelled bending angles are fairly localised. As an example of this, an offline test (results not shown here) demonstrates that a typical forward-modelled bending angle at 25 km has half of its contribution from the model levels between 25km and ~26.5 km, and has 95% of its contribution from the levels up to ~36.6 km.

As shown above, the linear retrieval has the potential to propagate significant model or observation information down from well above the tangent height. Applying a cutoff of the bending angle

departures also localises the implied RO temperature departures to match better the domain of the radiosonde observations; the balloons typically ascend to around 30km. As a quantitative example of the role biases have above the domain of interest for a global monthly mean of implied COSMIC-1 RO dry temperature departures, the difference between the value calculated with backgrounds from ECMWF and the Met Office is over 1K at 10hPa when no cutoff is applied (see Figure 8). Therefore, unless a cutoff is applied to the bending angle departures, the Abel transform and hydrostatic integral will transfer this bias to lower levels. If the cutoff is lower, this will localise the model bias that exists in the calculations, and the double-differencing will be better able to cancel the model bias which should be similar in both the radiosonde and RO departures. To reiterate; applying a cutoff mitigates the influence of any model bias that exists above the vertical range sampled by the radiosondes, and therefore should not contribute to the resulting bias corrections.

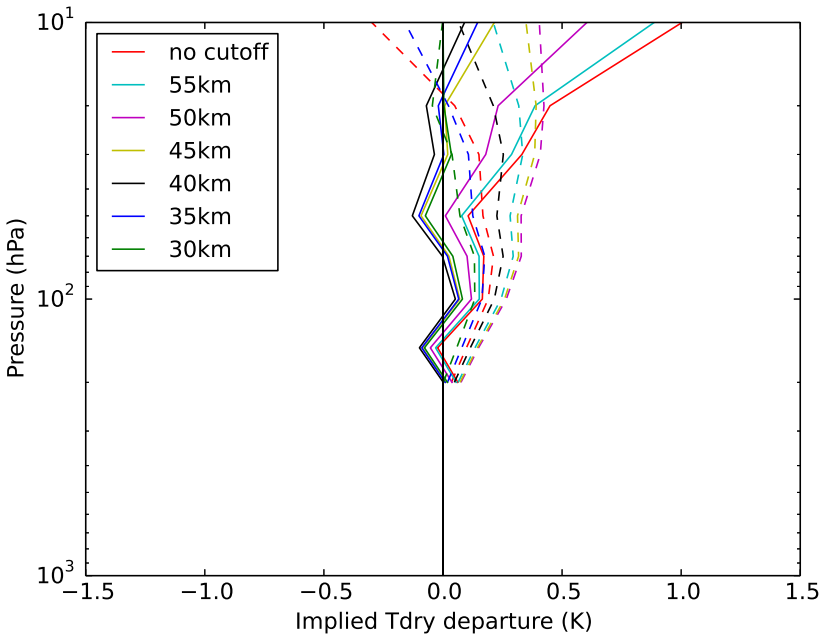

Figure 8: Global RO dry temperature departures from COSMIC-1 for one month (October 2016) calculated using backgrounds from the Met Office (solid) and ECMWF (dashed). These were calculated for a range of bending angle cutoff heights. The ECMWF background bending angle departures were provided in 1km bins, and as percentages — therefore, the mean observed bending angle from a comparable sample was used to obtain a proxy for the absolute ECMWF departures. Therefore, the values in this figure are not to be taken literally, but it does illustrate the importance of accounting for biases which may exist in the model above the vertical range sampled by the radiosondes.

6 Conclusions

The retrieval chain for radio occultation dry temperatures has been investigated to see the way in which information from upper-level bending angles can propagate down to dry temperatures at lower levels. This was achieved using a linearised version of a simplified retrieval algorithm, and Jacobian matrices have shown the sensitivity of the refractivity, pressure and temperature to the original bending angles. In particular, the uppermost bending angle has a large influence on retrieved temperatures at all levels below due to its use in extrapolating the bending angle profile above the highest level, and this value's repeated use in the integrals for refractivity and pressure. When background departures are propagated through the linear chain, the large, known, model biases above 40km become significant in the retrieved temperature departures. By decomposing the temperature increments at a given level into a summation of terms (Jacobian elements multiplied by the bending angle departures at all heights), a justification was provided for neglecting the bending angle increments above 35km, as these have an undue influence on the retrieved temperature departures.

The influence of upper information on the resulting temperature departures in the context of radiosonde bias correction has been discussed, and the locality of the Abel transform in the bending angle forward model was given as a reason for accepting the assumption that the backgrounds cancel when the difference of mean radiosonde and RO increments in the vicinity of a radiosonde station are calculated.

7 Acknowledgements

Thanks to Jordis Tradowsky, John Eyre and Stig Syndergaard for providing useful comments.

References

S. B. Healy and J.-N. Thepaut. Assimilation experiments with CHAMP GPS radio occultation measurements. *Quart. J. Roy. Meteorol. Soc.*, 132:605–623, 2006.

S. Syndergaard. Retrieval analyses and methodologies in atmospheric limb sounding using the GNSS radio occultation technique. Scientific Report 99–6, Danish Meteorological Institute, Copenhagen, 1999.

J. Tradowsky. Visiting Scientist Report 26: Characterisation of radiosonde temperature biases and errors using radio occultation measurements. SAF/ROM/DMI/REP/VS26/001, Version 1.1, 1 June 2016.

J. S. Tradowsky, C. P. Burrows, S. B. Healy, and J. Eyre. A new method to correct radiosonde temperature biases using radio occultation data (in review).

Met Office
FitzRoy Road, Exeter
Devon, EX1 3PB
UK

Tel: 0370 900 0100
Fax: 0370 900 5050
enquiries@metoffice.gov.uk
www.metoffice.gov.uk

Produced with the Met Office LaTeX template v2.0.0