A new floating model level scheme for the assimilation of boundary layer top inversions: the univariate assimilation of temperature.

Article

Published Version

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.

To link to this article DOI: http://dx.doi.org/10.1002/qj.955

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online
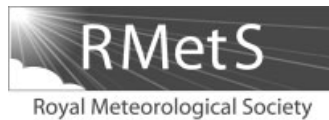
A new floating model level scheme for the assimilation of boundary-layer top inversions: the univariate assimilation of temperature

Alison Fowler,[a]* Ross Bannister[a] and John Eyre[b†]

[a]*Department of Meteorology, University of Reading, UK*
[b]*Met Office, Exeter, UK*

Correspondence to: A. Fowler, Department of Meteorology, University of Reading, Earley Gate, PO Box 243, Reading RG6 6BB, UK. E-mail: a.m.fowler@reading.ac.uk

[†]The contribution of this author was written in the course of his employment at the Met Office, UK and is published with the permission of the Controller of HMSO and the Queen's Printer for Scotland.

The assimilation of observations with a forecast is often heavily influenced by the description of the error covariances associated with the forecast. When a temperature inversion is present at the top of the boundary layer (BL), a significant part of the forecast error may be described as a vertical positional error (as opposed to amplitude error normally dealt with in data assimilation). In these cases, failing to account for positional error explicitly is shown to result in an analysis for which the inversion structure is erroneously weakened and degraded.

In this article, a new assimilation scheme is proposed to explicitly include the positional error associated with an inversion. This is done through the introduction of an extra control variable to allow position errors in the *a priori* to be treated simultaneously with the usual amplitude errors. This new scheme, referred to as the 'floating BL scheme', is applied to the one-dimensional (vertical) variational assimilation of temperature. The floating BL scheme is tested with a series of idealised experiments and with real data from radiosondes.

For each idealised experiment, the floating BL scheme gives an analysis which has the inversion structure and position in agreement with the truth, and outperforms the assimilation which accounts only for forecast amplitude error. When the floating BL scheme is used to assimilate a large sample of radiosonde data, its ability to give an analysis with an inversion height in better agreement with that observed is confirmed. However, it is found that the use of Gaussian statistics is an inappropriate description of the error statistics of the extra control variable. This problem is alleviated by incorporating a non-Gaussian description of the new control variable in the new scheme. Anticipated challenges in implementing the scheme operationally are discussed towards the end of the article. Copyright © 2011 Royal Meteorological Society and British Crown Copyright, the Met Office

1. Introduction

An important feature of a numerical weather forecast is the temperature inversion near the top of the boundary layer (BL). This is often present during anticyclonic conditions and acts to separate the moist BL air from the drier free-tropospheric air above. When stratocumulus cloud (Sc) is present in the BL, the cloud is capped by the inversion and prevented from rising into the free troposphere (Stull, 1988), making the height of the inversion an important aspect for cloud forecasts.

One of the factors that strongly influences the quality of a forecast of the BL capping inversion is its accurate representation in the initial conditions of the forecasting model. The height of the inversion, for instance, is frequently misrepresented by previous forecasts valid for the present time (referred to hereafter as the background) which can lead to fundamental errors in the inversion height and structure in a future forecast, even after data assimilation.

The importance of the inversion for diagnosing cloud means that an inaccurate inversion in the initial conditions may also lead to errors in the prediction of maximum and minimum temperatures, and hence the possibility of fog and ground frost, which are of particular public interest. An accurate inversion is also important for forecasting pollutant transport. The height can determine the extent to which the pollutants can rise (e.g. Stull, 1988) and the intensity of the inversion (the difference in potential temperature between the inversion top and bottom) is also directly proportional to its ability to inhibit the vertical movement of pollutants (Milionis and Davies, 1992), which is of interest for air quality forecasts.

In this article we do not consider cloud, but first look at the basic problem of how to allow for and reduce displacement errors of a quantity that has large vertical gradients (the BL temperature inversion). Part of this article shows that standard data assimilation techniques can fail to correct for BL inversion features that are at the wrong height in the background state, even when assimilated observations give sufficient information about the true inversion height. Standard variational data assimilation (Var) works by minimising the cost function

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}^{\text{b}})^{\text{T}}\mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^{\text{b}})$$
$$+ \frac{1}{2}\{\mathbf{y} - H(\mathbf{x})\}^{\text{T}}\mathbf{R}^{-1}\{\mathbf{y} - H(\mathbf{x})\}, \quad (1)$$

where $\mathbf{x}^{\text{b}}$ is the background vector, $\mathbf{y}$ is the observation vector and $\mathbf{B}$ and $\mathbf{R}$ are their respective error covariance matrices (square, symmetric and positive-definite by definition). The Ide *et al.* (1997) notation is used where possible. The value of the state vector, $\mathbf{x}$, at the minimum of J is hoped to be a 'best' estimate of the current state of the atmosphere. This is referred to as the analysis, $\mathbf{x}^{\text{a}}$, which we wish to compute. In this work, $\mathbf{x}$ represents a 1D vertical profile of potential temperature. The observation operator, H, allows $\mathbf{x}$ to be compared with $\mathbf{y}$ (here observations are also of potential temperature) and so J evaluates the disagreement of $\mathbf{x}$ with $\mathbf{x}^{\text{b}}$ (first term), and with $\mathbf{y}$ (second term) each weighted by the inverse of their respective error covariances. The minimum of J at $\mathbf{x} = \mathbf{x}^{\text{a}}$ can be shown to be the most probable state of the atmosphere, providing the background and observation errors are correctly defined, and are Gaussian and non-biased (Bouttier and Courtier, 1999). When H is linear, J is quadratic and the minimum can be calculated with practical algorithms. In reality, these criteria (Gaussianity, unbiased and linear operator) are rarely met and so the minimisation of (1) will be sub-optimal, although in practice the minimum may still give a useful analysis state.

A study of the characteristics of the background-error statistics for vertical temperature profiles was performed by Fowler *et al.* (2010). This study showed that the height of the inversion may be subject to a significant positional error. When there is a discrepancy between the height of the inversion in the observations and in the background, data assimilation can be problematic. As we show in this article, large vertical gradients present can mean that the observations and the equivalent background temperature values can differ well in excess of the error statistics specified by $\mathbf{B}$ and $\mathbf{R}$. This can make it difficult to find an estimate of the state that agrees with both. The standard 1D-Var framework (Eyre *et al.*, 1993) is otherwise successful at combining a background and observations, but in the case of a BL height discrepancy, it can lead to anomalously large analysis increments in the vicinity of the inversion, which are smoothed out by the otherwise correct B-matrix. Thus, in the attempt to correct for the height of the inversion, important information about the structure of the BL, including the inversion, can be washed out and anomalous spreading of information from the BL to the free atmosphere above can occur. An incorrect inversion is then likely to continue to propagate forward with the forecast model without the observations being able to bring the forecast back in line with the truth.

The assimilation of the BL-top inversion when a positional error is present was highlighted as one of the most important problems in numerical weather prediction at the UK Met Office (MO) at the end of 2006 when fog led to large disruptions at Heathrow during the Christmas period (Lorenc, 2007). In this instance, radiosonde observations were not able to improve the analysis of the inversion and so the fog was not accurately forecast. This will also adversely impact on the utility of the high-resolution models currently being developed at the MO.

Since then, little progress has been made on this specific problem, however the difficulty in assimilating features exhibiting positional errors has been highlighted in previous literature. Such features include fronts, thunderstorms, squall lines, hurricanes and precipitation, the height of the tropopause, and wild fires (e.g. Thiebaux *et al.*, 1990; Jones and MacPherson, 1997; Alexander *et al.*, 1998; Beezley and Mandel, 2008; Michel, 2010). A few attempts to include the positional error in the assimilation have been made. One such proposed technique developed by Alexander *et al.* (1998), for the improvement of forecasts of features associated with marine cyclones, is based on manually identifying corresponding features. A warping technique was then used to warp the entire field to match the observations. Mariano (1990) used contour analysis to calculate the average position of contours which was shown to avoid the weakening and smearing of geophysical features. Inclusion of position errors has also been used by Brewster (2003) for the assimilation of thunderstorms. Brewster (2003) defined a field of translation vectors in order to shift and distort the forecast field to match the observations, the aim being to minimise a displacement-based cost function and a penalty term based on the inverse of a second-order autoregressive term to prevent displacing too far from the original position.

Ravela *et al.* (2007) also developed a similar method making use of displacement vectors but without the autoregressive constraint. Because this method involves defining a field of displacement vectors for the entire field (the statistical error covariances for which are not known), it is very expensive and so the positional and amplitude adjustments were done separately. However, this will not find the optimal balance between the amplitude and positional error adjustment. Other work in this area has been performed by Hoffman *et al.* (1995), with later development by Nehrkorn *et al.* (2003). Within these two articles, a feature calibration and alignment technique was used in order to characterise forecast errors. Beezley and Mandel (2008) have also looked at the possibility of representing a positional error within the Ensemble Kalman filter with applications in wildfire modelling.

The idea of a change of co-ordinates (as proposed by Hoffman *et al.*, 1995; Brewster, 2003; Ravela *et al.*, 2007, and others) is used here to allow only the vertical grid to distort to allow a feature (the BL inversion) to move vertically to match the observations. This is done while preserving the underlying background-error covariance structure associated with the inversion. This is particularly useful when a state-dependent B-matrix is available that recognises the inversion in the background. Such a B-matrix may be obtained from studying the variability in an ensemble of forecasts of the BL temperature profile (Fowler *et al.*, 2010).

This paper follows on from Fowler *et al.* (2010), who showed that there is a pressing need to account for positional errors in the BL, and who suggested this method for the purposes of data assimilation. The structure of this article is as follows. In section 2, the new assimilation method is introduced. In section 3, results of experiments involving idealised cases are presented. In this idealised set-up, the truth, background and observations are generated in such a way that the background errors are consistent with the proposed amplitude plus positional error model. This comprises the basic assimilation tests (section 3.1), the sensitivities to certain parameters introduced within this scheme (section 3.2), an assessment of whether an equivalent conventional (amplitude-only) scheme could be modified to do the same job as the new scheme (section 3.3), and a study of the impact of the nonlinear observation operator (section 3.4). In section 4, the method is applied to a large sample of real cases, with discussion of statistics of the positional error of the background inversion (section 4) and the effect of modelling the positional error as non-Gaussian (sections 4.2 and 4.3). Finally in section 5 the key findings of this article are summarised and ideas for future work are given. Special consideration is given to the method of implementation in operational 3D- or 4D-Var.

2. Method

Hoffman *et al.* (1995) demonstrated that splitting the forecast errors into amplitude and displacement errors for certain features can give valuable information. It has been shown by Fowler *et al.* (2010) that a significant part of the error in the background inversion may be described by the vertical displacement of the inversion height. Therefore within the assimilation scheme it is not only the level-by-level temperatures (the amplitudes) that need to be varied

in order to give the 'best fit' (as in standard Var (1)), but also the height of the inversion in the background (the displacement).

Fowler *et al.* (2010) used an ensemble of background states to study the flow-dependence of the vertical B-matrix in the presence of strong and weak boundary-layer inversions (strong being defined as a significant rise in temperature through the inversion). They found that there is a sudden drop (to near zero) in temperature correlation between those levels below and above the inversion top when the inversion is strong (Figure 1 of their paper). In our new data assimilation method, which includes the correction of displacement errors, it is important for such correlations to be preserved with the inversion structure as it moves up or down, as the assimilation tries to correct its positional error. This conservation is achieved here by shifting the vertical positions of the underlying vertical levels themselves. Since the state vectors and the B-matrix are represented on these levels, changing level heights of some levels during the assimilation will not affect the background-error covariances between those levels.

One way of allowing the levels to shift in height is to assign one extra control variable per level, each describing how the level height is perturbed by the assimilation. However, this may lead to the possibility of interchanging levels or levels moving below the ground. In this work, only one new control variable is introduced, denoted a. This single variable, to be described in section 2.1, allows the coherent movement of a range of levels in the region of the inversion in such a way that the allowed shift fades to zero away from the inversion, ensuring levels near the ground remain fixed. Moving the levels in this way allows the inversion in the background to move without undesirable distortion. From now on, this scheme will be referred to as the 'floating BL scheme' (dealing with amplitude and position errors), as opposed to the 'standard scheme' (dealing with only amplitude errors).

2.1. A cost function described in terms of floating levels

In order to utilise a, a new cost function is written in terms of a new augmented control vector, $\widetilde{\mathbf{v}}_n$. In the following, a tilde indicates that the information in a vector or matrix is carried on levels that are 'floating' (that they are able to change their heights), and the subscript n indicates the Var iteration index.

$$
\begin{aligned}
J(\widetilde{\mathbf{v}}_n) = &\frac{1}{2}(\widetilde{\mathbf{v}}_n - \widetilde{\mathbf{v}}^{\mathrm{b}})^{\mathrm{T}}\widetilde{\mathbf{B}}_{\mathrm{v}}^{-1}(\widetilde{\mathbf{v}}_n - \widetilde{\mathbf{v}}^{\mathrm{b}}) \\
&+ \frac{1}{2}\{\mathbf{y} - H_n(\widetilde{\mathbf{v}}_n)\}^{\mathrm{T}}\mathbf{R}^{-1}\{\mathbf{y} - H_n(\widetilde{\mathbf{v}}_n)\},
\end{aligned}
\tag{2}
$$

where $\widetilde{\mathbf{v}}_n$ is the new control vector comprising the level-by-level values, $\widetilde{\mathbf{x}}_n$, and a_n:

$$
\widetilde{\mathbf{v}}_n = \begin{pmatrix} \widetilde{\mathbf{x}}_n \\ a_n \end{pmatrix}.
\tag{3}
$$

The quantities stored in the $\widetilde{\mathbf{x}}_n$ part of $\widetilde{\mathbf{v}}_n$ are potential temperature (θ) on each level. (Unlike temperature, θ is conserved as the levels move, assuming adiabatic conditions.) The *a priori* value of $\widetilde{\mathbf{v}}$ is

$$
\widetilde{\mathbf{v}}^{\mathrm{b}} = \begin{pmatrix} \widetilde{\mathbf{x}}^{\mathrm{b}} \\ a_0 \end{pmatrix} = \begin{pmatrix} \widetilde{\mathbf{x}}^{\mathrm{b}} \\ 0 \end{pmatrix},
\tag{4}
$$

where $\widetilde{\mathbf{x}}^b$ has the same numerical values as in $\mathbf{x}^b$, the ordinary background state on fixed levels. The background-error covariance matrix for $\widetilde{\mathbf{v}}^b$ is $\widetilde{\mathbf{B}}_v$ and is also described on the floating levels

$$\widetilde{\mathbf{B}}_v = \begin{pmatrix} \widetilde{\mathbf{B}} & 0 \\ 0 & \sigma_a^2 \end{pmatrix}, \tag{5}$$

where $\widetilde{\mathbf{B}}$ is numerically the same B-matrix as in (1). This contains the error covariances associated with θ on each model level–but now applied to floating rather than fixed levels. σ_a^2 is the variance of the extra variable. For convenience, it is assumed in (5) that the errors in $\widetilde{\mathbf{x}}^b$ and a are uncorrelated. In practice, it is difficult to separate out these errors in order to calculate what the correlations should be. Since amplitude errors are represented in terms of θ, it is reasonable to assume that these are less correlated with errors in a than for temperature if cross-correlations between errors in θ and a are dominated by adiabatic effects. Some attempt to calculate these correlation have been made using the Met Office Global and Regional Ensemble Prediction System (MOGREPS Bowler *et al.*, 2008) ensemble data. Due to difficulties in separating out the amplitude and positional errors and the strong state-dependence, these results are not included but they did suggest that the correlations may indeed be non-negligible (of the order of a half). This is not investigated further within this article and it is possible to argue that these correlations are small in comparison to those commonly neglected between control variables in operational schemes (Bannister *et al.*, 2008). However this assumption may be relaxed in (5) if required and may be beneficial when observations of the inversion are unavailable.

In order to minimise (2), both amplitude and position adjustments to the background are made simultaneously. If the background does not contain an inversion which is present in the observations, or vice versa, then no reduction in (2) will be achieved by varying a and so only amplitude adjustments will be made.

2.2. How a controls the level heights

Let the heights of the fixed model levels be z^i and the heights of the floating levels (for iteration n) be $\widetilde{z}_n^i$. These are related by

$$\widetilde{z}_n^i = z^i + a_n D(z^i). \tag{6}$$

Here $D(z^i)$ is a prescribed function of model level height that we call the 'displacement function'. This tells the assimilation which levels are allowed to move (and to what degree) and which are not. It is chosen such that it is unity within the centre of the background inversion and tends to zero away from the inversion. Then, only those levels around the inversion have the potential to move. This is the means by which a group of levels can be moved coherently with the single extra control variable, a. Within this framework the value of $a_n D(z^i)$ is interpreted as the adjustment of the height of model level i, and a_n is the adjustment at the level where the displacement function has value unity. Each level i now has a vertical position error equal to $\sigma_a D(z^i)$. Within this work, a mixed uniform and triangular function is used which has the desired characteristics described above; it is

given by

$$D(z) = \begin{cases} 1 + \frac{z - z_b}{r} & \text{for} \quad z_b - r < z \le z_b, \\ 1 & \text{for} \quad z_b < z < z_t, \\ 1 - \frac{z - z_t}{r} & \text{for} \quad z_t \le z < z_t + r, \\ 0 & \text{elsewhere}, \end{cases} \tag{7}$$

where z_t and z_b are the height of the top and bottom of the background inversion respectively and r defines the lengthscale in which $D(z)$ linear decreases from 1 to 0. Therefore for the depth of the inversion in the background, the displacement is equal to 1, ensuring that the inversion is not distorted in any way as it is moved vertically. The gradual linear decrease in the displacement function either side of the inversion helps to prevent levels from interchanging which could lead to discontinuities within the temperature profile.

The assimilation is performed in floating level space, but once the analysis, $\widetilde{\mathbf{v}}^a$, is calculated, values can be interpolated to the fixed model levels if necessary. As with any interpolation, this will introduce a degree of error dependent on the resolution of the model and the depth of the inversion. In general, if the resolution of the model has high enough resolution to capture the inversion, then the interpolation error should not be significantly large.

2.3. The observation operator

Point observations of potential temperature (at fixed heights) are considered here and so the observation operator, which computes the model's version of the observations, consists of interpolation. The introduction of a_n means that as J is minimised iteratively, $\widetilde{z}_n^i$ must be recalculated with each iteration (using (6)). This in turn alters the observation operator from iteration to iteration. Even with linear interpolation, H becomes a nonlinear function of elements of $\widetilde{\mathbf{v}}_n$, due to the presence of a (this issue is explored further in section 3.2).

The value of the jth element of H_n acting on $\widetilde{\mathbf{v}}_n$ can be expressed by linear interpolation as a weighted combination of the values of $\widetilde{\mathbf{v}}_n$ either side of the jth observation level:

$$H_n(\widetilde{\mathbf{v}}_n)^j = \left(\frac{\widetilde{z}_n^{II} - z_o^j}{\widetilde{z}_n^{II} - \widetilde{z}_n^I} \right) \widetilde{x}_n^I + \left(\frac{z_o^j - \widetilde{z}_n^I}{\widetilde{z}_n^{II} - \widetilde{z}_n^I} \right) \widetilde{x}_n^{II}, \tag{8}$$

where z_o^j is the height of the jth observation, and $\widetilde{x}_n^I$ and $\widetilde{x}_n^{II}$ are the values of $\widetilde{\mathbf{v}}_n$ either side of the jth observation level and $\widetilde{z}_n^I$ and $\widetilde{z}_n^{II}$ are their heights respectively. As a is varied, the floating heights $\widetilde{z}_n^I$ and $\widetilde{z}_n^{II}$ used to calculate $H_n(\widetilde{\mathbf{v}}_n)^j$ will change and so the levels needed may be different. This is illustrated in Figure 1 for the case when the assimilation has decided to move a floating level across an observation location between two iterations. In the $(n + 1)$th iteration, the level used to define $\widetilde{x}_{n+1}^{II}$ was previously used to define $\widetilde{x}_n^I$.

3. Results involving idealised data

3.1. An idealised example

Background and observation profiles are produced which are consistent with their specified error covariances. The
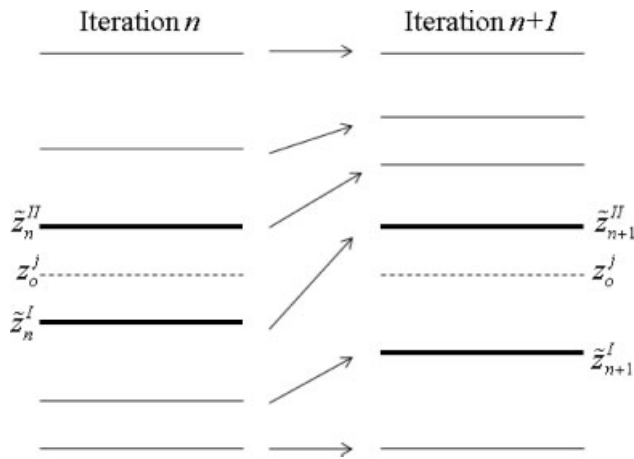
Figure 1. Schematic of how the floating model levels (full lines) can change between iterations respective to an observation level (dashed lines). The two floating levels that are used in the interpolation for the location of the observation are shown in bold. The interpolation formula is (8).

background and truth differ in two ways: firstly by noise drawn from the normal distribution with error covariance **B** on the model levels, and secondly by a vertical shift to simulate the error for a (see below). The observations are a subset of the truth with added noise drawn from the normal distribution with error covariance **R**. These profiles are then used to test the floating BL scheme.

In this example, the background gives a measure of potential temperature at level spacing of 40 m within the lowest 3 km of the atmosphere and then a spacing of 200 m between 3 and 8 km giving a total of 100 levels. Observations of potential temperature are made at the heights of every fifth model level.

The B-matrix is calculated from the spread in an ensemble (e.g. Bannister, 2008) about its mean using

$$B_{ij} = \frac{1}{M-1} \sum_{k=1}^{M} \left\{ (\mathbf{x}^b)^i_k - \left\langle (\mathbf{x}^b)^i \right\rangle \right\} \left\{ (\mathbf{x}^b)^j_k - \left\langle (\mathbf{x}^b)^j \right\rangle \right\}, \quad (9)$$

where i and j represent elements of the state vector, M is the number of ensemble members, k represents an ensemble member and $\left\langle (\mathbf{x}^b)^i \right\rangle$ is the average value of the ith element of $\mathbf{x}^b$ over all ensemble members. In this study, the ensemble is provided by MOGREPS. MOGREPS has only 20 vertical levels in the lowest 8 km of the Earth's atmosphere, so linear interpolation is used to give the higher resolution described above necessary for this study. The resulting B-matrix is valid for the mean of the ensemble which we take to be the background, $\mathbf{x}^b$. For this example the B-matrix used is plotted in Figure 1A of Fowler *et al.* (2010), which was there considered without any vertical shift. We consider this amplitude part of **B** close to optimal since it is derived explicitly from an ensemble. From this background, we calculate a state $\mathbf{x}^t$ which is considered the truth for this idealised case. $\mathbf{x}^t$ is calculated in three steps:

1. First, produce a profile, $\mathbf{x}^{t*}$, which removes error consistent with **B** on the fixed model levels.

$$\mathbf{x}^{t*} = \mathbf{x}^b - \mathbf{L}\Lambda^{\frac{1}{2}} \mathbf{e}_x, \quad (10)$$

where **L** is the matrix of eigenvectors of **B**, Λ is the diagonal matrix of eigenvalues and $\mathbf{e}_x$ is a vector of 100 random numbers, each taken from Gaussian distribution with mean zero and variance 1, $N(0, 1)$.

2. Second, shift the levels of $\mathbf{x}^{t*}$ according to the standard deviation of the error for a, σ_a, and a chosen displacement function, $D(z)$. This gives a new set of heights, $\widetilde{z}$

$$\widetilde{z}^i = z^i + \sigma_a e_z D(z^i), \quad (11)$$

where e_z is a scalar taken from $N(0, 1)$. σ_a here is chosen as 200 m (this is consistent with the error statistics of a based on real data to be discussed in section 3.4). The profile describing the truth generated in the first step, $\mathbf{x}^{t*}$, is now described on the shifted levels, $\widetilde{z}$.

3. Third, this profile is interpolated back to the original model levels, $\mathbf{z}$, to give $\mathbf{x}^t$.

This procedure gives background and truth states which contain a relative vertical positional error which is not thought to be unusual for a real data assimilation situation.

The imposed observation errors are chosen to be Gaussian, unbiased and uncorrelated in the vertical and so can be described solely by their variances. Observations can then be simulated from the truth by

$$\mathbf{y} = H(\mathbf{x}^t) + \mathbf{R}^{1/2} \mathbf{e}_y, \quad (12)$$

where $\mathbf{e}_y$ is an observation space vector of random numbers, each generated from $N(0, 1)$. In this case $H(\mathbf{x}^t)$ gives a vector of $\mathbf{x}^t$, but on every fifth model level.

An example of background and observation profiles generated in this way can be seen in Figure 2. Although potential temperature is used in the assimilation, temper-
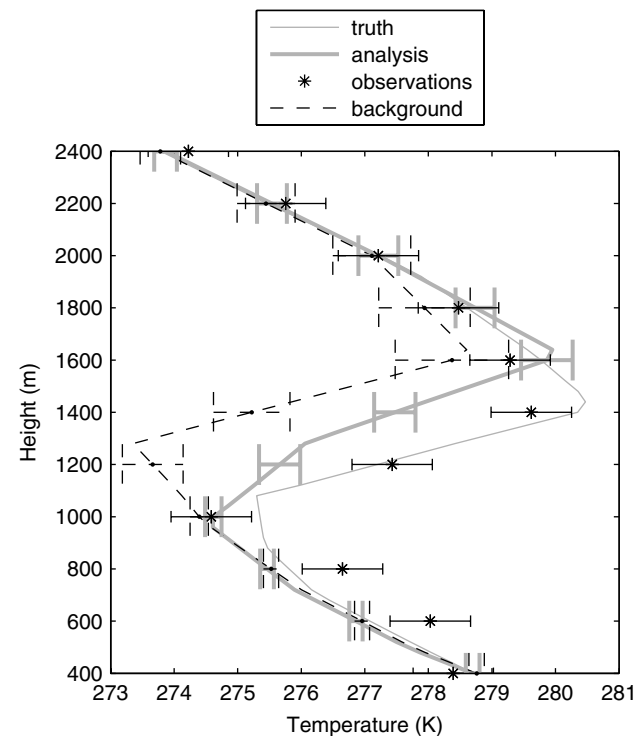

Figure 2. Illustration of how the standard 1D-Var scheme reacts to a disagreement between the background (dashed) and the observation (stars) temperature inversion height. The true temperature profile is shown by the thin grey line. The resulting analysis is shown in thick grey. Error bars denote the standard deviations (amplitude only).

ature is plotted to highlight the inversion structure. This shows a clear shift between the observations (stars) and the background (dashed) profiles.

The particular parameters chosen to define $D(z)$ in (7) in this example were $z_t = 1560$ and $z_b = 1280$, chosen to match the background temperature structure and $r = 460$. The random shift chosen in this experiment was $\sigma_a e_z = +240$ m. Due to the upward shift of the background inversion with respect to the truth, the background values are systematically too cool in the inversion region. This simulated shift is comparable to the difference seen between the inversion height as measured by radiosondes and a collocated forecast (section 3.4).

Performing the standard 1D-Var (1) for this case gives the analysed (thick grey) profile in Figure 2. The analysis does not sufficiently warm the background in the region of the inversion as the assumed background errors are too small (dashed error bars). The spreading of analysis increments is limited in this example due to the use of a state-dependent B-matrix which has very low correlation between levels in the inversion and elsewhere. If a static B-matrix were used, which is unaware of the inversion, the analysis increments would be spread vertically and the inversion would be weakened (Fowler *et al.*, 2010). In the case of Figure 2, the analysed inversion appears to be significantly broader in height than it should be. The resulting analysis has been unable to correct the height of the inversion or to retain the structure of the inversion present in the observation and the background profiles. It is also inconsistent with the error statistics (shown as error bars). This analysis could consequently alter the properties of any BL Sc diagnosed and subsequently forecast and would also impact on the evolution of the BL structure.

Now the new floating BL scheme is tested with the same background, observations and displacement function as above by minimising (2) instead of (1). The result is shown in Figure 3. In this analysis, a has been minimised to a value of -243 m, which is close to the magnitude random value 240 m used in (11), causing the background inversion structure to move down in order to agree with the observed structure. This shift downwards means that the displacement error has been reduced with the extra variable a and not by making large analysis increments to the amplitudes. This means that, unlike in the standard scheme, the spreading of the (large, in the case of the standard scheme) analysis increments has not washed out some of the inversion structure present in the background.

The movement of the levels, and the shape of the displacement function chosen are shown in Figure 4. In Figure 4(a), the black line shows the standard deviations of the errors of the background level heights which are found from $\sigma_a D(z^i)$ for level i. These are the vertical error bars shown in Figure 3 around the background. Similarly the grey line gives the standard deviations of the errors of the analysis level heights (derived from the inverse Hessian of the cost function). The presence of the vertical error bars in this scheme allows the background and observation profiles to agree with greater consistency than in the standard scheme, despite the height shift. This consistency is seen in the value of J at the analysis. If all assumptions about the errors are correct, we would expect J to equal half the number of observations (Bennett, 1992, 2002). In these experiments, there are 20 observations and so J has an expected value of 10. Using the standard scheme $J = 19.9$ and using the
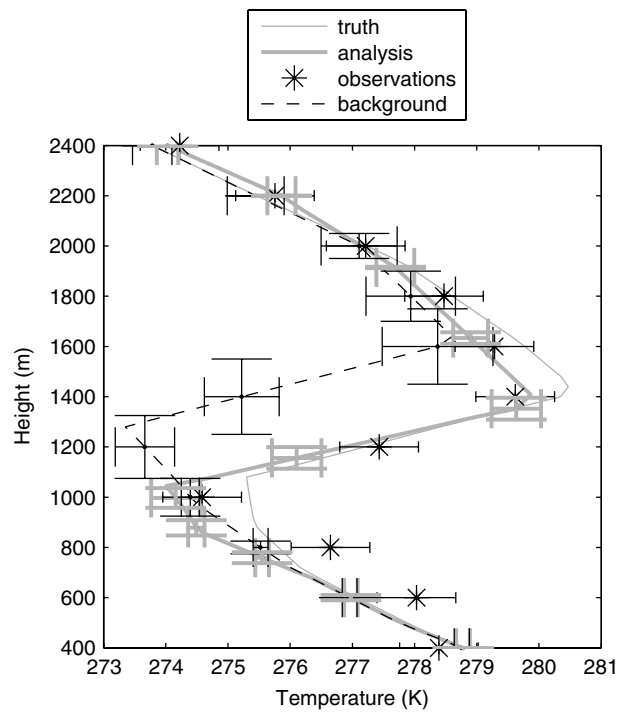

Figure 3. Analysis (thick grey) recalculated for the floating BL scheme. Observations (stars) and background (dashed) are the same as in Figure 2. This scheme has minimised the cost function by moving the background inversion down 243 m as well as varying the values of temperature on each model level. Error bars denote the standard deviations. For the background and analysis, these include height error standard deviations associated with the floating BL scheme.
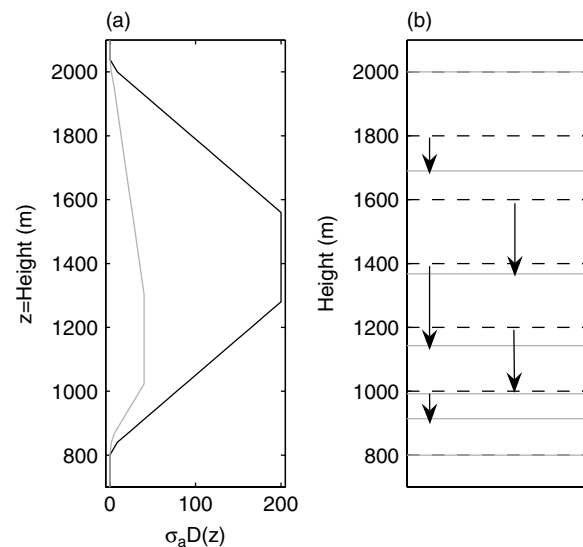

Figure 4. (a) shows the standard error associated with the heights of the model levels (black) and the heights of the final levels for the analysis (grey). Only the levels which are within the range of influence of the displacement function are associated with an error and so are free to move, as seen in (b). Since the error for the analysis is much smaller than the background, we have more confidence in the height of the inversion as diagnosed by the analysis temperature profile. In (b), the floating levels for the final iteration of Figure 3 (grey) have been plotted alongside the original model levels (dashed). Only one in five levels has been plotted for each field to allow for easy visual comparison.

floating BL scheme $J = 9.08$. This suggests that, not only has the floating BL scheme produced a better analysis than the standard scheme, it has done it in a way that is more consistent with the statistics.
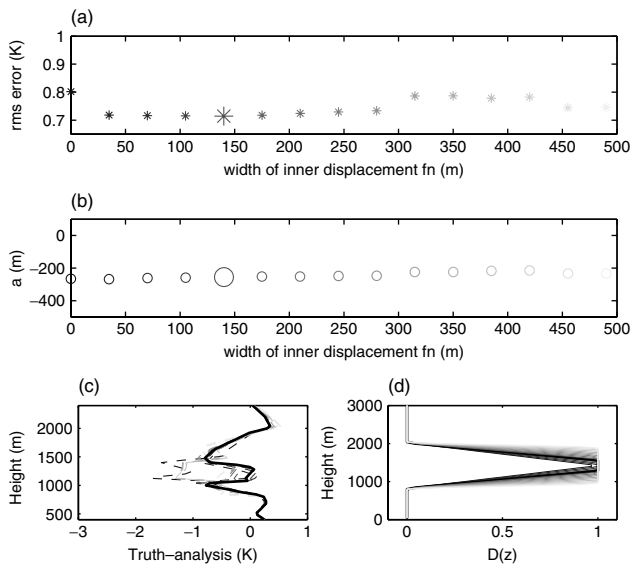
Figure 5. (a) The rms error as a function of the width of the region where the displacement function, D_{iw}, is unity. (b) a at the analysis as a function of D_{iw}. (c) The analysis minus truth temperature profiles for each value of D_{iw}. The black bold line gives the analysis which has used the true displacement function and the continuous black line is the true temperature profile. (d) $D(z)$ for each value of D_{iw}. The thick line is the 'true' displacement function corresponding to the larger marker in (a) and (b). In all panels, the line shade helps to distinguish the different experiments.

In these examples, low-resolution observations have been used to illustrate the standard and floating BL schemes to highlight the potential problems with the standard scheme. Access to high-resolution observations could indeed reduce some of the issues seen with the standard scheme, as the observations reinforce one another giving the analysis a greater trust in them. However the analysis will still be inconsistent with the given errors and prone to spreading the large increments vertically.

Throughout this article, in all cases, incremental Gauss–Newton minimisation of the cost function is used.

3.2. Sensitivity to $D(z)$ and σ_a^2

The choice of the displacement function, although chosen subjectively, is essential to the success of the floating BL scheme. In the idealised experiments, the true displacement function, given by (7), and positional error variance, σ_a^2, are known (and used), but in the case of real data these would be unknown. We can test the sensitivity of the scheme to changes in the displacement function and σ_a^2. The following four figures (Figures 5–8) show how the floating BL assimilation is sensitive to changes in aspects of the displacement function. Note that the background and observation states are the same as those used in Figures 2 and 3. Shown in each figure is (a) the root-mean-square (rms) error of the potential temperature analysis, (b) the analysed value of a, (c) the resulting truth minus analysis temperature profiles and (d) the displacement functions used. The four parameters of the displacement function that are varied (Figure 4) are:
the width of the region where the displacement function is unity, $D_{iw}(= z_t - z_b$ for the optimal $D(z)$) (Figure 5);
the overall width of the displacement function, $D_{tw}(= D_{iw} + 2r$ for the optimal $D(z)$, i.e. where it is non-zero) (Figure 6);
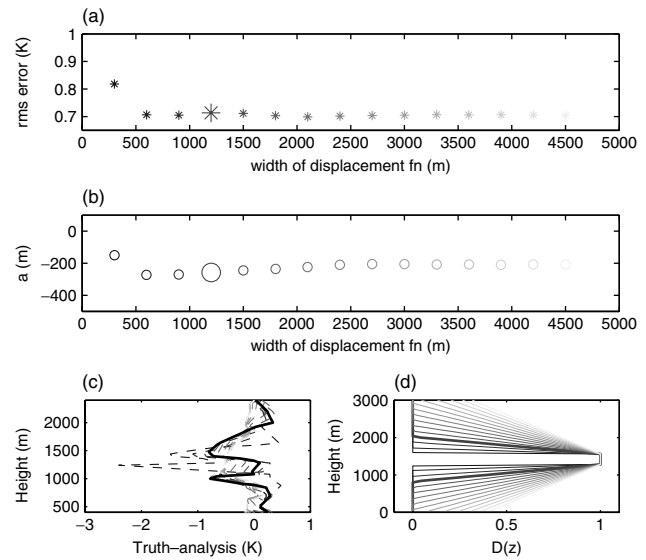

Figure 6. As Figure 5, but for the overall width of the displacement function, D_{tw}, greater than zero.
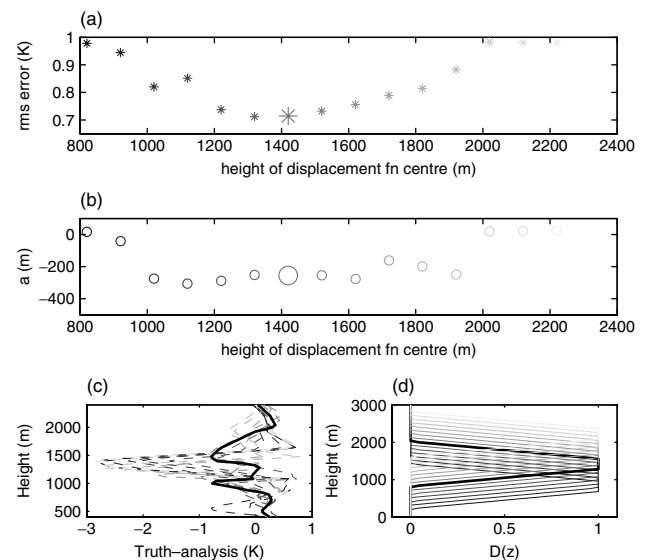

Figure 7. As Figure 5, but for the height of the centre of the displacement function, D_h.

the height of the centre of the displacement function, $D_h(= 0.5(z_t + z_b)$ for the optimal $D(z)$) (Figure 7);
and σ_a (Figure 8).
In each figure the other variables are kept constant.

From these figures, we see that the analysis is relatively insensitive to the exact shape of the displacement function. Figure 5 shows that the floating BL scheme works well even when the displacement function is triangular (D_{iw} is zero). In Figure 6 we see that the floating BL scheme is not as effective if $D(z)$ is uniform (i.e. $D_{tw} = D_{iw}$–the leftmost point in Figure 6(a, b)) where a is equal to only 130 m, approximately half the necessary shift. This also results in an increased rms error. This is not surprising as the range of levels that are allowed to move in this case is very narrow. Conversely in Figure 7, it is seen that the analysis is sensitive to the position of the displacement function. If the function centre is too high (> 2 km) or too low (< 1 km), then $a = 0$ in the analysis and, for the set of observations used, the floating BL scheme in effect reverts back to the standard
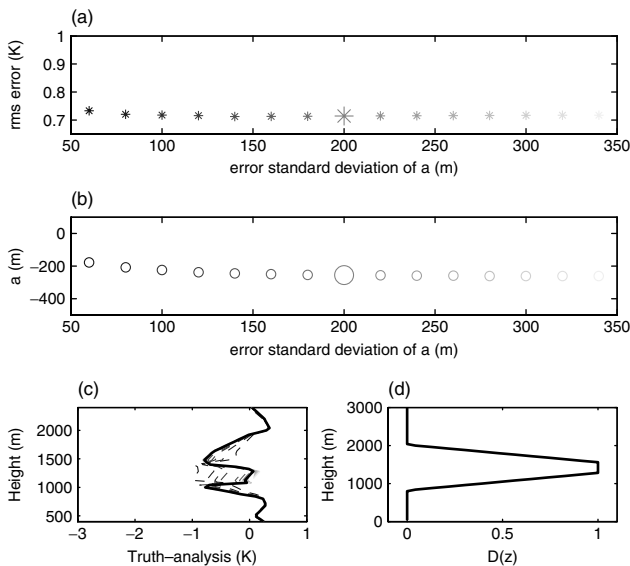
Figure 8. As Figure 5, but for the value of σ_a. $D(z)$ is equal to the true function for each experiment.

scheme and only amplitude adjustments are made. When no positional adjustments are made, the rms error is largest at approximately 1 K. In Figure 8, we see that the floating BL scheme is insensitive to the exact value of σ_a within the range tested.

In conclusion, it is important that the displacement function is centred on the inversion in the background but the width and shape of the displacement function need only be sensible. Therefore when real data are used in section 4, an adaptive displacement function is used which is calculated from the background temperature profile before the assimilation has begun. This ensures that the displacement function is positioned accurately. The mixed uniform and triangular function was seen to give a good description of the necessary displacement correlations for a wide range of non-idealised situations.

Only elevated inversions which correspond to the top of the BL are considered when assigning the displacement function. In the cases when no inversion is present in the background state, the displacement function is set to zero everywhere letting the assimilation effectively revert back to the standard amplitude-only scheme.

Potential sensitivities to the correlation length-scales in the amplitude part of **B** in the standard scheme are discussed in Fowler *et al.* (2010). However, the floating BL scheme has a decreased need to add large-amplitude increments to minimise the cost function, so it is reasonable to expect that the floating BL scheme will be less sensitive to these correlation length-scales. Since the amplitude part of **B** is derived explicitly from an ensemble, we consider it to be close to optimal in this work and so we do not investigate this here.

3.3. Inflated and implied background-error covariances

The extra degree of freedom of a allows the floating BL scheme to give an analysis which is capable of keeping the structure of the inversion seen in the background, but has allowed it to be moved to a height more in keeping with the observations. It may be thought that a similar result may be gained with the standard scheme but by

first inflating the background-error variances in the vicinity of the inversion. This will allow the observations to pull the background more effectively towards the truth. We have tested this idea by deriving a new B-matrix from the ensemble with the modification that each ensemble member is shifted using (11), each with a different random number e_z. In Fowler *et al.* (2010), this was shown to give a clear increase in the background-error variances within the inversion region. Using the B-matrix generated in this way with the standard scheme gives an analysis which is in good agreement with the observations, but with a loss of the BL inversion structure present in the background (not shown). This is undesirable in general as the background inversion structure is often accurate (even if it has the wrong height). This is a particularly significant drawback of this intermediate scheme when observations are noisy or are of a low vertical resolution.

A second method for producing amplitude-only errors which includes the positional error implicitly is to calculate the so-called implied background-error covariances. The matrix $\widetilde{\mathbf{B}}_v$ used in the floating BL scheme can be transformed to give these implied background-error covariances in the fixed level space by

$$\mathbf{B}_{\text{implied}} = \mathbf{Z}\widetilde{\mathbf{B}}_v\mathbf{Z}^{\mathrm{T}}, \qquad (13)$$

where $\mathbf{Z}$ represents the Jacobian of the interpolation from floating to fixed model space, and has been calculated using the value of a in the analysis of Figure 3. If the state vector $\widetilde{\mathbf{v}}$ has $N+1$ elements, $\widetilde{\mathbf{B}}_v$ is an $(N+1)\times(N+1)$ matrix and $\mathbf{Z}$ is an $N\times(N+1)$ matrix giving $\mathbf{B}_{\text{implied}}$ to be an $N\times N$ matrix. We propose that $\mathbf{B}_{\text{implied}}$ is the effective B-matrix that should be used in the standard scheme if the background inversion height were to be corrected. The implied background-error standard deviations (square roots of diagonal elements $\mathbf{B}_{\text{implied}}$) for the example given in Figure 3 are plotted in Figure 9 (compared to the standard deviations taken from the ordinary B-matrix used by the standard scheme). The result shows an increase of the standard deviations for $\mathbf{x}^b$ in the region between the background inversion and the true inversion. Hence this
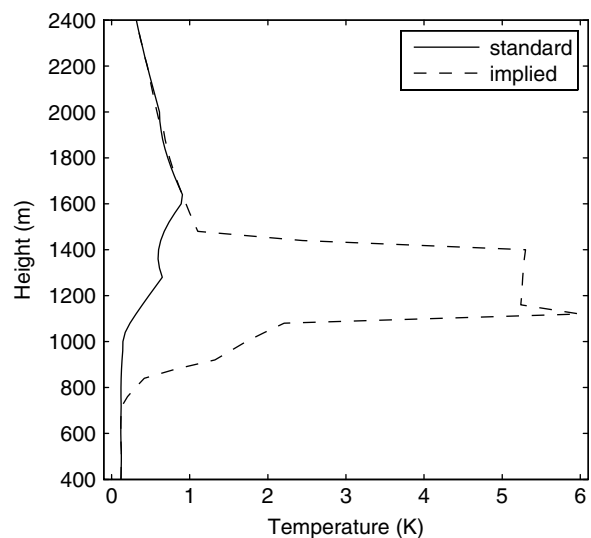

Figure 9. Background-error standard deviations on fixed levels as would be used in standard 1D-Var (continuous line) and the implied standard deviations derived from (13) (dashed line). The implied errors have been calculated using the value of a given by the analysis in Figure 3.
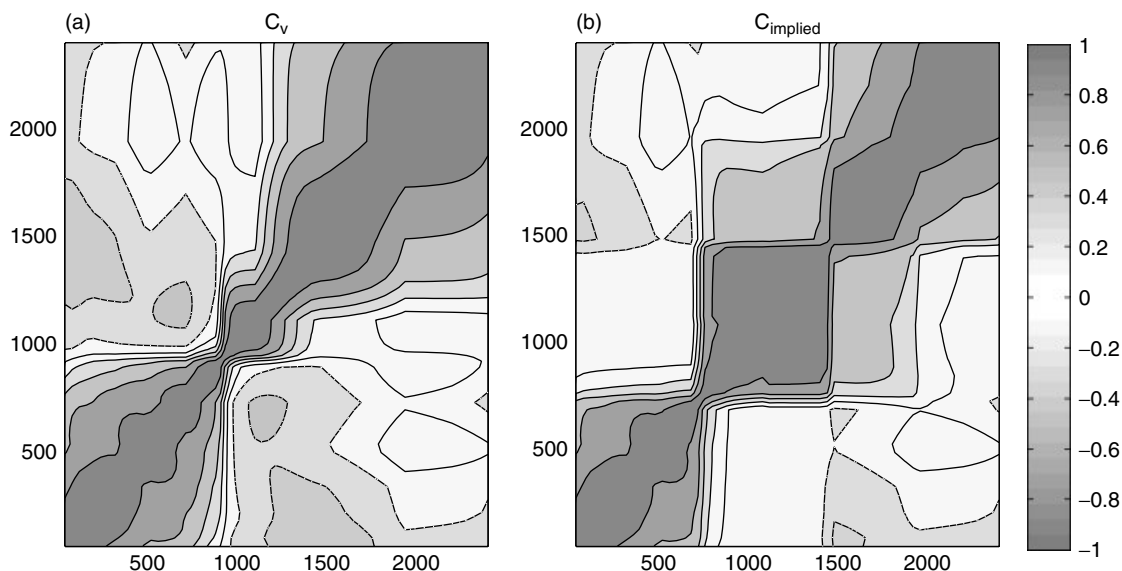
Figure 10. The background-error correlation matrix (a) used in the floating BL scheme calculated from the spread in an ensemble of forecasts, and (b) the correlation matrix implied on the original fixed model levels. Dashed contours represent negative correlations. The axes are given in metres.

effective error increase is dependent upon knowledge of the true inversion height which is intrinsic to the value of a at the analysis and hence differs from the first method discussed.

In Figure 10, the correlation structure for potential temperature errors derived from $\widetilde{\mathbf{B}}_v$ (Figure 10(a)) and from $\mathbf{B}_{\mathrm{implied}}$ (using the analysed a as in Figure 3) (Figure 10(b)) are shown. As $\widetilde{\mathbf{B}}_v$ is calculated from the spread in a set of ensembles (each valid at the same time), the correlations shown are state-dependent and it is possible to see the decoupling between the air above and below the BL for the position and structure of the inversion in this specific case (the BL top is at approximately 1 km). In the analysis of Figure 3, the assimilation has moved the inversion down by approximately 240 m and so the BL is shallower and the region of strong positive correlation seen in $\widetilde{\mathbf{B}}_v$ has shrunk in $\mathbf{B}_{\mathrm{implied}}$. An increase in the correlation length-scale can also be seen between 800 m and 2 km, caused by the displacement function.

We now test whether $\mathbf{B}_{\mathrm{implied}}$ is suitable for use as a B-matrix in standard 1D-Var. We wish to test this because the inflated standard deviations around the inversion will allow potentially large analysis increments, allowing the observations to have more relative weight. However, unlike in the first method, the correlations are now based on the necessary displacement provided by minimising the cost function for the floating BL scheme (2). The resulting analysis is poor, as shown in Figure 11. The strong correlation structure of $\mathbf{B}_{\mathrm{implied}}$ means that using this as the background-error covariance matrix in standard 1D-Var does not give a realistically smooth analysis. Thus, even with the larger background-error variances and adjusted correlations, the standard 1D-Var is still unable to effectively 'move' the background inversion whilst retaining its shape. It can be concluded that the extra variable, a, and the nonlinear observation operator are necessary to solve this problem.

This result is consistent with Lawson and Hansen (2005) and Ravela *et al.* (2007), who stated that the presence of a large positional error means the use of Gaussian amplitude errors only is no longer appropriate. Applying the results of
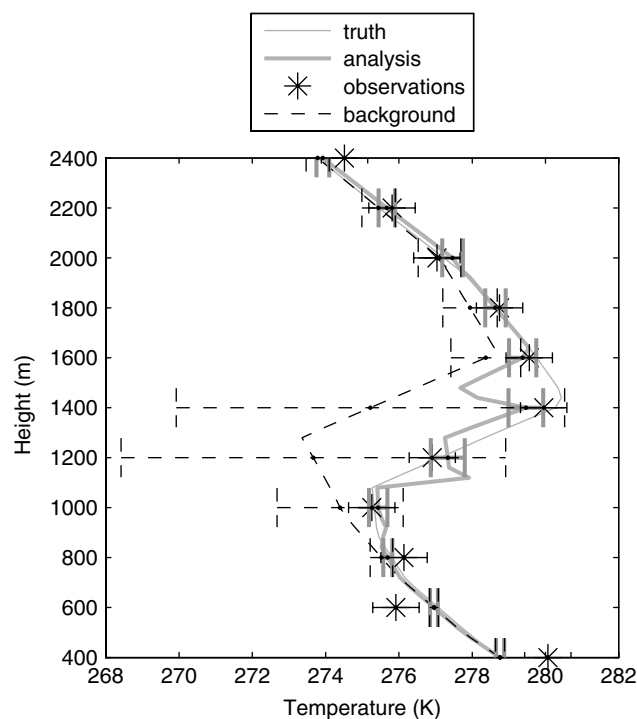
Figure 11. Assimilation using the standard scheme (as in Figure 2), but with the implied B-matrix given by (13). The cost function $J = 8.33$ at the analysis.

Lawson and Hansen (2005) to our problem, the total error, η, as assumed by the floating BL scheme may be written as

$$\eta(z^i) = x^{\mathrm{b}}(z^i - \eta_{\mathrm{D}}(z^i)) - x^{\mathrm{b}}(z^i) + \eta_{\mathrm{A}}(z^i). \quad (14)$$

Here z^i represents the original gridpoint height and $\eta_{\mathrm{D}}(z^i)$ represents the adjustment to this gridpoint given by $\eta_{\mathrm{inv}}D(z^i)$, where η_{inv} is the error in the inversion height itself. $\eta_{\mathrm{A}}(z^i)$ is the amplitude error which can then be added to the displaced background values to give the truth. Fields in (14) are written in function notation, rather than vector notation, to show the shift explicitly. Expansion of (14)

about $\eta_D = 0$ and $\eta_A = 0$ leads to the following expression for the total error

$$\eta(z^i) \approx \eta_A(z^i) - \eta_D(z^i)\frac{\partial x^b(z^i)}{\partial z} + \frac{\eta_D^2}{2}\frac{\partial^2 x^b(z^i)}{\partial z^2}, \qquad (15)$$

with the assumption that $\eta_{inv}\partial D(z^i)/\partial z \ll 1$ for simplicity. The first two moments of η can be expressed as

$$\langle \eta(z^i) \rangle \approx \frac{1}{2}\frac{\partial^2 x^b(z^i)}{\partial z^2}\langle \eta_D^2(z^i) \rangle, \qquad (16)$$

assuming amplitude and displacement errors are unbiased and

$$\langle \eta(z^i)\eta(z^j) \rangle \approx \mathbf{B}_{ij} + \frac{\partial x^b(z^i)}{\partial z}\langle \eta_D(z^i)\eta_D(z^j) \rangle \frac{\partial x^b(z^j)}{\partial z}, \qquad (17)$$

assuming the amplitude and displacement errors are independent as in (5). The covariance in (17) is that between levels i and j. Equations (16) and (17) show that in general a positional error can lead to:
a bias in the total error, violating one of the key assumptions in deriving the cost function;
an increase in the variances; and
a change in the correlation length-scales in the region where η_D is non-zero (Michel, 2010) i.e. in the region of the inversion as seen in Figures 9 and 10.

Lawson and Hansen (2005) showed that, for non-biased Gaussian amplitude and displacement errors which are independent, η is adequately Gaussian i.e. terms involving η_D^2 and above in (15) are small compared to first-order terms if

$$\frac{\eta_{inv}}{2L} \ll 1, \qquad (18)$$

where L is the typical length-scale of the feature. In our case, this is the depth of the inversion and is approximately 200 m. Therefore if η_{inv} is anything greater than about 40 m, say, the ratio (18) is not significantly less than one and so η cannot be described as Gaussian and hence impossible to represent in a Gaussian amplitude error framework. In section 4, we will see that the value of η_{inv} can be as large 600 m. This implies that any assimilation scheme based on Gaussian statistics, such as Var and the ensemble Kalman filter, would fail to coherently shift the inversion structure. An assessment of the ensemble Kalman filter when a positional error is present is given by Beezley and Mandel (2008).

There may be ways of avoiding the need to shift, e.g. assimilating observations more frequently to prevent the wrong inversion height in the forecasts, but this is not studied here because this is not a flexibility that is generally available operationally. A primitive assimilation technique, e.g. nudging, may be used instead to correct for large-amplitude errors without the anomalous smoothing effects seen in Figure 2, but nudging does not consider important uncertainty statistics, nor can it deal with other kinds of observation (such as those remotely sensed).

The result implied by (18) for the inversion is illustrated in Figure 12. In Figure 12(a), an ensemble of temperature profiles (truth plus different realisations of the full error–amplitude plus position (14)), are plotted. These errors are consistent with those used to simulate the idealised data in section 3.1, i.e. the same B-matrix and displacement function. At the levels marked, a histogram of the ensemble is given by the bars in Figure 12(b). The dashed lines represent a Gaussian distribution in which only amplitude errors are present. In the region of the inversion, these curves are far too narrow as only the statistics of the first term in (15) are represented. The solid lines represent a Gaussian distribution in which both the amplitude and positional error variances are present. Within this region, it is clear that a Gaussian approximation of the shape of the PDF is inaccurate. The distributions associated with the levels with the largest displacement error appear to be highly skewed as predicted by (18).

3.4.   Issues of nonlinearity

As already mentioned in section 2.1, the presence of a makes the observation operator, H, nonlinear. A linearised estimate of H is introduced for the Gauss–Newton minimisation of the cost function. Let $\mathbf{H}$ represent the matrix of partial differences of H with respect to the state vector $\widetilde{\mathbf{v}}$ evaluated at a current estimate of $\widetilde{\mathbf{v}}$. The linearisation error introduced when approximating perturbations in H by $\mathbf{H}\delta\widetilde{\mathbf{v}}$ can be quantified as

$$\text{linearisation error} = H(\widetilde{\mathbf{v}}_n) - \{H(\widetilde{\mathbf{v}}_{n-1}) + \mathbf{H}\delta\widetilde{\mathbf{v}}_n\}. \qquad (19)$$

When $D = 0$ there is no linearisation error. The magnitude of the linearisation error tends to zero as $\delta\widetilde{\mathbf{v}}$ tends to zero so, once the solution has converged to the minimum of J, the linearisation error will be small.

A nonlinear observation operator leads to a non-quadratic cost function, which can make it more difficult to find its minimum. In Figure 13, the cost function has been calculated for the set-up in Figure 3 as a function of a with the rest of the elements in $\widetilde{\mathbf{v}}$ equal to their background values. J_B (the background part of the cost function $(\widetilde{\mathbf{v}} - \widetilde{\mathbf{v}}^b)^T\mathbf{B}^{-1}(\widetilde{\mathbf{v}} - \widetilde{\mathbf{v}}^b)$) and J_O (the observation part of the cost function $\{\mathbf{y} - H(\widetilde{\mathbf{v}})\}^T\mathbf{R}^{-1}\{\mathbf{y} - H(\widetilde{\mathbf{v}})\}$) are also plotted separately.

The quadratic nature of J_B is unaffected by the introduction of a. The shape of J is dominated by the J_O contribution and its minimum is clearly seen to be shifted towards $a = -240$ m. This is the vertical shift between the inversion heights of the background and truth states in the set-up of Figure 3. Within this region of a-space, there are no other significant maxima or minima confirming that the floating BL scheme is working satisfactorily in this case.

It is possible for double minima to appear in J with the introduction of the floating BL scheme, especially when the resolution of the observations is not adequate to represent the true position of the inversion. In such cases, it is possible in the assimilation for the background inversion to be moved to the wrong position as a result of fitting to the observations. This can make it questionable whether the particular analysis calculated is truly the most probable state of the atmosphere. This can be demonstrated by assimilating just one observation (at 1160 m) as shown in Figure 14. In Figure 14(a) the truth (grey line), the background (dashed line) and the observation (star). The true and background profiles are the same as in Figures 2 and 3. The floating BL scheme now has a choice to move the inversion either up or down to fit the observation. J, J_B and J_O as a function of a have been plotted in Figure 14(b). The ambiguity in which way the background should be shifted introduces a
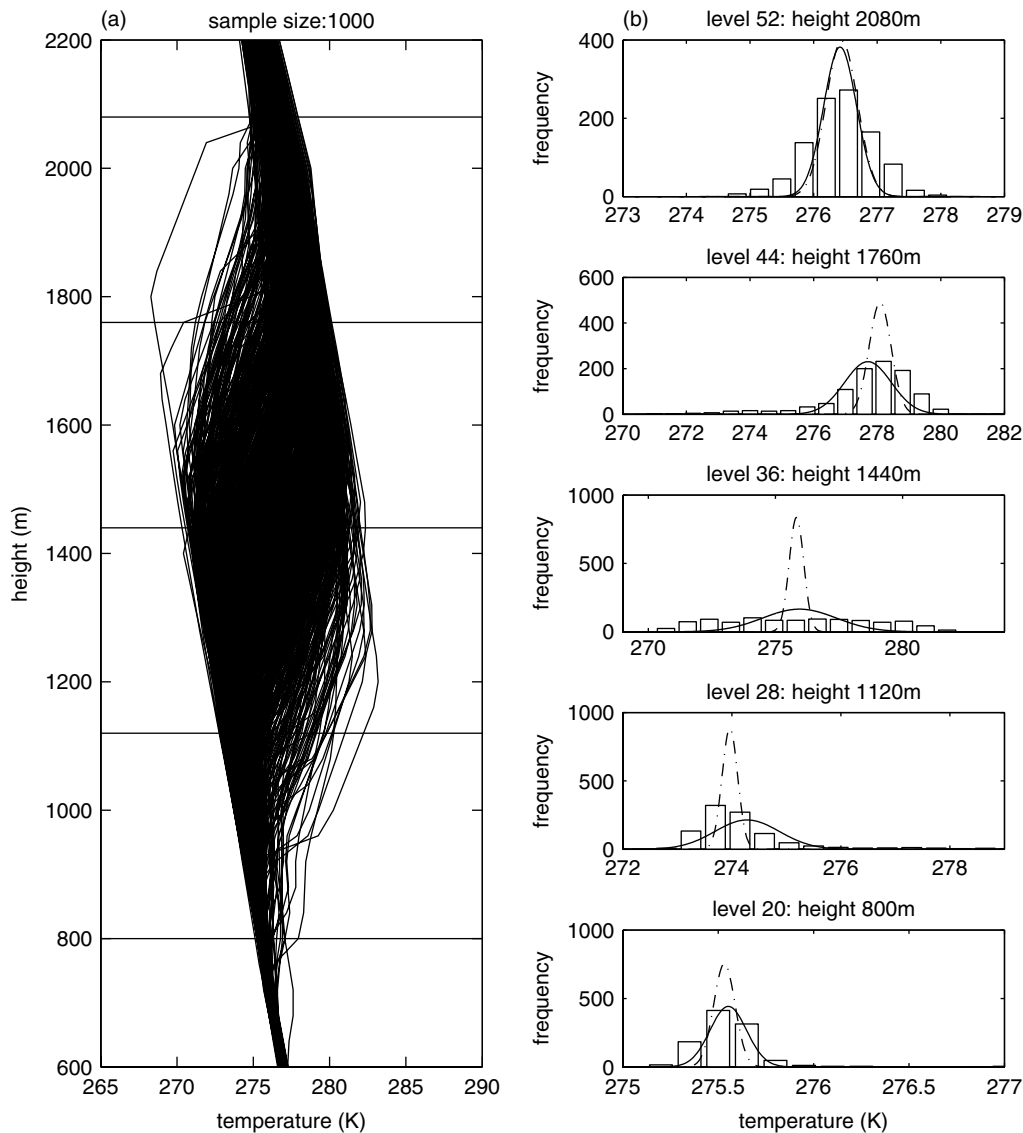
Figure 12. (a) An ensemble of temperature profiles generated by adding an error consistent with $\widetilde{\mathbf{B}}_v$ (i.e both amplitude and positional error). At each level marked, a bar plot of the distribution in temperature given by the sample is shown in (b), where dashed lines represent the assumed distribution given by $\mathbf{B}$ (i.e. amplitude error only), and solid lines give a Gaussian approximation to the sample.

local minimum at approximately +200 m in J, while the global minimum is still at $a = -240$ m. This could result in the scheme moving the levels in the wrong direction and giving a worse analysis than the standard 1D-Var scheme. In this example, the maximum value of J separating the two minima is at $a = -100$ m whilst the background value of a is 0 and hence the starting point of the minimisation is within the well of the (wrong) local minimum.

Although this is a contrived example, it does highlight the importance of an adequate observation density and/or a minimisation method that looks for the global minimum to remove ambiguities. In fact we find no occurrences of such ambiguities when assimilating real data (section 4.3).

4. Results involving real data

4.1.   The error statistics of the new control variable, a, derived from real observations

The standard deviation of a, σ_a, describes the likely range of distances that the inversion found in the background

may be allowed to move vertically in order to give a fit to the observed inversion. σ_a may then be thought of as the standard deviation of the error in the height of the background inversion. This can be estimated by comparing the heights of the inversion in a population of collocated actual forecasts and radiosonde observations.

The error in the height of the inversion in the background, η_{inv}, is the difference between the true height of the inversion, $z_{\text{inv}}^{\text{t}}$, and the background height of the inversion, $z_{\text{inv}}^{\text{b}}$:

$$\eta_{\text{inv}} = z_{\text{inv}}^{\text{t}} - z_{\text{inv}}^{\text{b}}. \qquad (20)$$

In these results, the inversion is calculated as the lowest level of an elevated temperature inversion. If we assume that the positional error is intrinsic to errors within the forecast model and so the cause of the positional error in the background data is not present within the observed data, we can substitute the true inversion height, $z_{\text{inv}}^{\text{t}}$, for the observed inversion height, $z_{\text{inv}}^{\text{y}}$. This gives an expression for η_{inv} in terms of knowable quantities:

$$\eta_{\text{inv}} = z_{\text{inv}}^{\text{y}} - z_{\text{inv}}^{\text{b}}. \qquad (21)$$
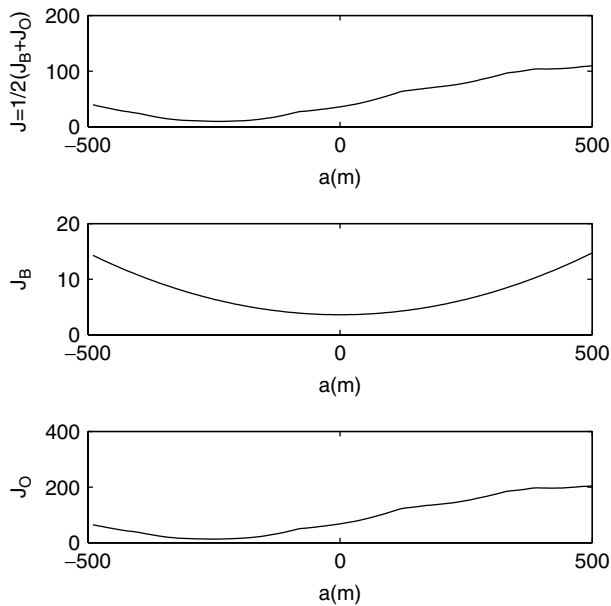
Figure 13. Cost functions plotted as a function of a: J, J_B and J_O.

However, in general, the observations themselves are subject to measurement errors, which could lead to an error in the observed inversion height. In (21), we have assumed that radiosondes measure the inversion height with error that is negligible compared to the resolution of the data. This is not easily testable objectively since we do not have co-located radiosondes to check against. However, we do believe this to be a fair assumption for measurements by radiosondes in which the resolution of the data is high compared to that of the model, and subject to only a small measurement error. For noisy or low-resolution observations, there is likely to be a significant source of error, $\eta_{\text{inv}}^{\text{y}}$:

$$\eta_{\text{inv}}^{\text{y}} = z_{\text{inv}}^{\text{t}} - z_{\text{inv}}^{\text{y}}. \qquad (22)$$

The difference then between the observed and background inversion heights, instead of (21) is

$$z_{\text{inv}}^{\text{y}} - z_{\text{inv}}^{\text{b}} = \eta_{\text{inv}} - \eta_{\text{inv}}^{\text{y}}, \qquad (23)$$

using definitions (20) and (22). The error statistics of (23) are

$$\left\langle (z_{\text{inv}}^{\text{y}} - z_{\text{inv}}^{\text{b}})^2 \right\rangle = \left\langle (\eta_{\text{inv}})^2 \right\rangle + \left\langle (\eta_{\text{inv}}^{\text{y}})^2 \right\rangle, \qquad (24)$$

where the background and observation errors are assumed to be uncorrelated. The left-hand side (LHS) of (24) is calculable from a given population of background states and observation profiles, and the first term on the RHS can be estimated separately by synthesising a population of such low-resolution or noisy observations profiles from a truth to derive the variance of $\eta_{\text{inv}}^{\text{y}}$. The second term on the RHS is σ_a^2 which can then be derived.

We assume that this procedure is not necessary in this work as the radiosonde observations used (see below) are accurate and of high resolution and so (21) is assumed to be applicable. This means that a sample of radiosonde and collocated background temperature profiles allow the variance of the error of the height of the background inversion to be calculated:

$$\sigma_a^2 = \left\langle \eta_{\text{inv}}^2 \right\rangle, \qquad (25)$$

where the triangular brackets represent an average over a suitable population when an inversion is present in both the background forecasts and the observations. Ideally, for data assimilation, the error distribution of any control variable should be Gaussian and unbiased so that the assumptions that the variational assimilation theory is based on are true. Figure 15 shows the distribution of η_{inv}. The background data are from the global Unified Model (UM; Cullen, 1993). The background and collocated observations have been taken from a population comprising 50 days of data at

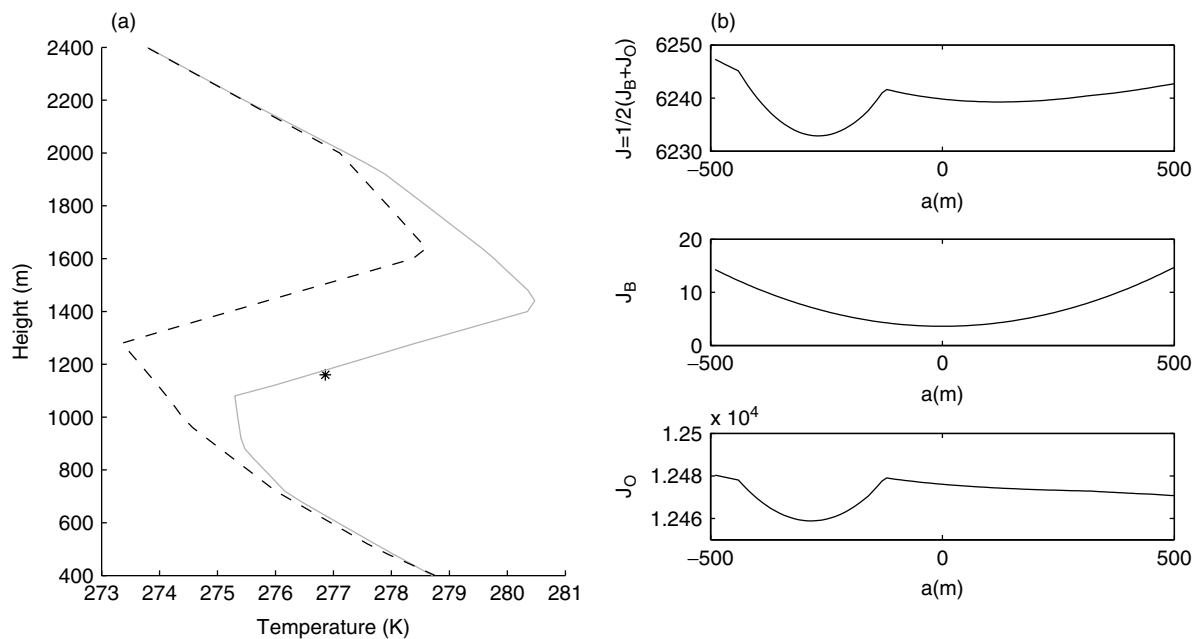**Figure 14.** (a) A single observation at 1160 m (star) measuring the centre of a true inversion (grey) which is 240 m lower than the inversion in the background (dashed). In order for the analysis to agree with the observations in the floating BL scheme, the background could be raised or lowered. (b) Components of the cost function as a function of a showing double minima in J at approximately -240 m and $+200$ m.
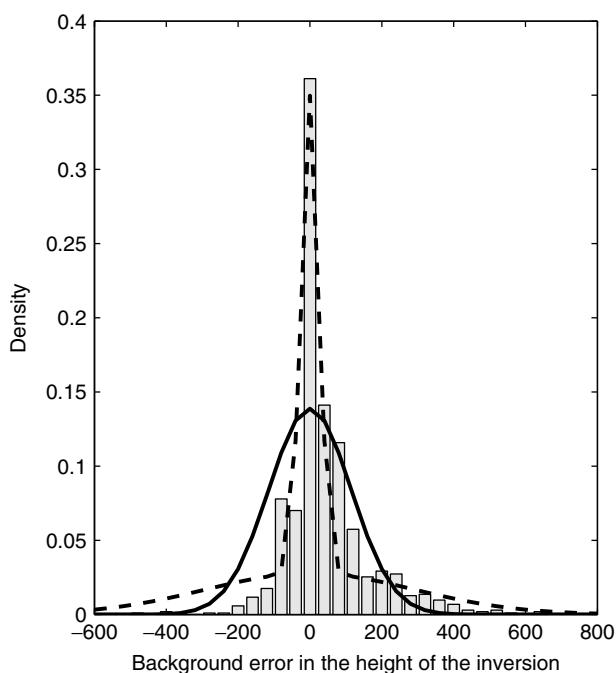
Figure 15. A histogram of observation minus background inversion height. This distribution is assumed to describe the error distribution of *a*. Examples of different distributions have also been plotted. The continuous curve is Gaussian with mean zero and standard deviation as calculated from the distribution. The dashed curve is an example of a sum of two Gaussian distributions which more accurately describes the shape of the distribution.

midday covering Europe over land from 24 December 2006 to 13 October 2008. During these 50 days an inversion was present in the background approximately 40% of the time, which gave us a sample of 1000 profiles to compare. The data have been interpolated to the same 100 level vertical resolution as that used for the idealised examples. The height error is seen to be significant, ranging from -400 m to $+600$ m, with a standard deviation of 115 m.

The distribution of the errors of *a* is non-Gaussian and asymmetric. The mean of the distribution is 42 m, so there is a slight bias towards the inversion in the background being too low. This may be due, for example, to too much large-scale subsidence in the model, erroneous front locations, or incorrect strength of mixing in the BL, cloud-top cooling, cloud-top entrainment or wind shear.

Modelling the distribution of *a* as Gaussian (continuous curve in Figure 15) underestimates the probability in the assimilation that the difference between the observation and background inversion heights is small (smaller than 40 m). This would mean that the floating scheme, when assimilating real data, would try to make a shift of the background more often than is necessary. That is, the scheme might try to shift the levels when the background inversion height is correct. However in cases when there is a large difference between the observation and background inversion heights, the floating BL scheme would be able to move the heights and give an improved analysis. This effect was seen in the preliminary assimilation results, given in Figure 8 of Fowler *et al.* (2010)

4.2. Modelling the non-Gaussian error statistics of the new control variable, a.

A better description of the prior error distribution of *a* could improve the floating BL scheme. There are two obvious

routes which could be taken. The first is the transformation from *a* to a new control variable to describe the movement of the background levels. This new control variable could be chosen such that its errors are better described by a Gaussian distribution. Work done by Hólm (2003) has shown that any variable, *g*, can be transformed to a variable, γ, with a Gaussian error distribution by finding a transform, *f*, such that the cumulative distribution of the probability of *g* less than a fixed value $\delta\phi$, $P(g \leq \delta\phi)$, equals the cumulative distribution of the probability of γ less than the transformed fixed value $f(\delta\phi)$, $P\{\gamma \leq f(\delta\phi)\}$. A second route is to reformulate the cost function for non-Gaussian errors and simply to choose a distribution which more accurately describes the errors seen. It is believed in this case that this is a simpler solution as we are dealing with only one variable which is taken to be uncorrelated with the rest of the state variables. The main drawback of the current Gaussian distribution for *a* is that it does not give high enough probability to the event of the background inversion height being in close agreement to that observed. Other important features of the distribution are the long tails representing the events when the background height is in a large disagreement with the observations. It is necessary to represent accurately these two characteristics of the error distribution in order for our floating BL scheme to accept these events and to be able to move the background accordingly. These criteria are satisfied empirically by a sum of two Gaussian curves (a two-component Gaussian mixture model). One Gaussian distribution with a large standard deviation in order to capture the long tails and a second with a small standard deviation to represent the sharp peak. The intention is to reduce the chance of the floating BL scheme moving the background inversion more often than is necessary. This is a similar approach to that of Andersson and Järvinen (1999) for their description of observation errors used for quality control purposes. This non-Gaussian distribution of $\eta_{\text{inv}} = a$ is described as P_a

$$P_a = (1 - A)P_{a_1} + AP_{a_2}, \qquad (26)$$

where P_{a_i} is a Gaussian distribution $(N(0, \sigma_{a_i}^2))$ and A is a constant between 0 and 1. An example of this is shown in Figure 15 (dashed curve), where $A = 0.5$, $P_{a_1} \sim N(0, 25^2 \text{m}^2)$ and $P_{a_2} \sim N(0, 300^2 \text{m}^2)$, chosen to fit the data in this sample. The parameters A, σ_{a_1} and σ_{a_2}, are found by trial and error to optimise the fit to the measured distribution. The sample used to make Figure 15 comprises data over a long period and so we assume it is applicable to different datasets. In an operational setting, it is expected these parameters would need to be calculated only once. The parameters, however, may be changed to fit a new population as long as it has characteristics of the sum of two Gaussians. By using (26), we account approximately for the sharp peak and long tails of the known error distribution but not its asymmetry, which we do not attempt to account for in this work.

The sum of two Gaussians is no longer Gaussian (as seen in Figure 15), and so using this distribution to represent the error distribution of *a* means the background part of the cost function is no longer quadratic. The background part of the cost function is now written explicitly in two parts as

shown below (cf. (2); the Appendix gives a derivation)

$$J_B = \frac{1}{2}(\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}^b)^T \widetilde{\mathbf{B}}_x^{-1} (\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}^b)$$
$$- \ln \left\{ \frac{1-A}{\sigma_{a_1}} \exp\left(\frac{a^2}{2\sigma_{a1}^2}\right) + \frac{A}{\sigma_{a_2}} \exp\left(-\frac{a^2}{2\sigma_{a2}^2}\right) \right\}. \quad (27)$$

The first term represents the (quadratic) amplitude description of the background profile (as in (1)) and the second term represents the (now non-quadratic) positional description of the background inversion. The separation is possible due to the assumption that the amplitude and positional errors are uncorrelated.

4.3. Testing the assimilation with real data

The performance of the 1D-Var schemes is now tested with a population of actual radiosonde data and UM forecasts with the following background state formulations:
(i) the floating BL scheme (fully Gaussian);
(ii) the floating BL scheme (non-Gaussian formulation described above); and
(iii) the standard scheme (Gaussian).
As discussed in section 3.2, in order for the floating BL scheme to be successful, the displacement function must be adaptive to the height of the inversion in the background. For this reason, before the assimilation of background and observation data is performed, the inversion top and bottom are derived from the background profile and the displacement function is then set to 1 within these limits. Outside of this region, the displacement function linearly tends to zero as in the idealised example (section 3.1).

A summary of assimilation results is now given in terms of frequency histograms of differences between the analysed and observed inversion heights (Figure 16). These plots show how effective the assimilation schemes are at shifting the background profiles to agree better with the inversion heights derived from the radiosondes. We assume that the observed inversion heights have negligible error (and so the plots may be interpreted as analysis inversion-height error distributions), and noting that the radiosonde-derived inversion heights are not assimilated themselves, these plots contain useful information for this study. The same population of backgrounds and observations is assimilated as that used to derive Figure 15. The objective is to produce an analysis-error distribution which is as tall and thin as possible about an analysis error of zero. In Figure 16(a), the analysis error distributions are compared between the Gaussian floating BL scheme (stems) and the non-Gaussian floating BL scheme (bars). The difference between the shapes of the two distributions is clear: the non-Gaussian formulation produces more analyses with a smaller error than the Gaussian formulation, and the standard deviation of the analysis height error has also been marginally reduced from 55 m to 50 m. This demonstrates the advantage of the non-Gaussian formulation.

In Figure 16(b), the analysis error distributions are compared between the standard scheme (stems) and the non-Gaussian floating BL scheme (bars). The two distributions have a similar shape with a relatively similar frequency of cases when the positional error is small (central bin). However in the standard scheme there are still extreme cases when the analysis has a very large positional error. The standard deviation of the analysis height error in the
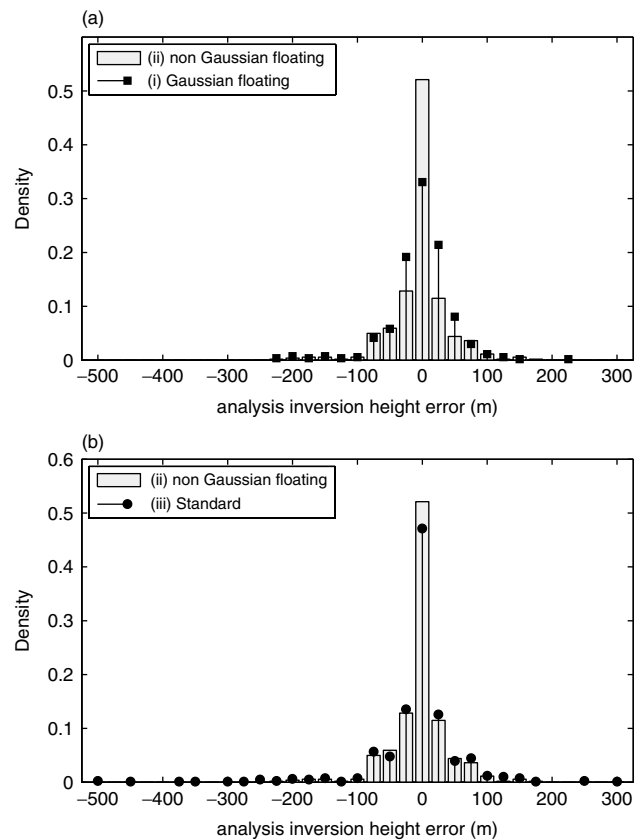

Figure 16. Inversion height analysis error distributions (a) with the floating BL scheme and the error in the background inversion height is approximated as Gaussian (stems) and as the sum of two Gaussian distributions (bars), and (b) with the standard scheme (stems) with the bars the same as in (a).

non-Gaussian floating BL scheme is significantly smaller than of the standard scheme (50 m compared to 75 m) and the non-Gaussian floating BL scheme performs better than even the standard scheme when the background inversion height is accurate (small analysis errors).

This diagnostic measures the improvement in the inversion height only. However, from what we have learnt about the idealised examples, we would also expect the floating BL scheme, with both the Gaussian and non-Gaussian positional error formulation, to give a better representation of the inversion structure as the background inversion structure is more easily maintained. We can compare the average temperature gradient of the inversion in the analysis using the three different methods to that of the background and observations. On average, the inversion in the analysis using the standard scheme is weaker than in the analyses using the floating BL formulations ($8.62 \times e^{-3}$ K m^{-1} compared to $9.87 \times e^{-3}$ K m^{-1} for the Gaussian floating BL scheme and $9.30 \times e^{-3}$ K m^{-1} for the non-Gaussian floating BL scheme). The stronger inversions seen in the floating BL scheme analyses are consistent with the average strengths seen in the observations ($9.70 \times e^{-3}$ K m^{-1}) and the backgrounds ($9.26 \times e^{-3}$ K m^{-1}).

As seen from the statistics plotted in Figure 16, the differences in the analyses using the three different methods can be subtle. It is perhaps easier to appreciate the differences between the three schemes for real data by looking at a few examples in detail. In Figure 17, three cases are shown where the model of the positional error used is important. In each case the background (thick grey line) and the observations
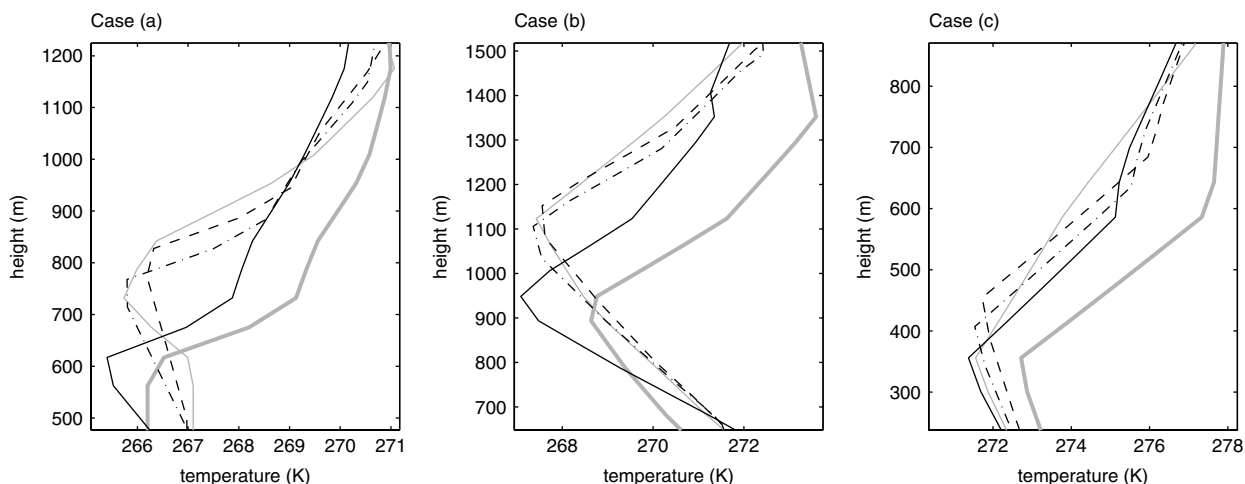
Figure 17. Examples of assimilations with real data using the three different methods. The background is given by the bold grey line and the observation profile is represented by the thinner grey line. The analyses profiles are given by the black lines; the solid uses the standard scheme, the dashed uses the floating BL scheme with a Gaussian representation of the positional error and the dashed-dot uses the floating BL scheme with the non-Gaussian representation of the positional error described in section 4.2. Dates for the three cases are given in Table I.

Table I. Comparison of the analysis error in the inversion height, η_{inv}, and the gradient of the inversion (inv grad) for the three assimilation methods for three cases using real radiosonde and background data. Cases a, b, and c correspond to the panels in Figure 17.

Case	Date	Location	Standard		Gaussian FBL			Non-Gaussian FBL		
			η_{inv}	inv grad	η_{inv}	inv grad	a	η_{inv}	inv grad	a
a	17 Dec 2007	Switzerland	−108	0.0066	105	0.0118	216	45	0.0098	157
b	13 Dec 2007	Switzerland	−163	0.0063	47	0.0133	212.8	0	0.0128	165.5
c	18 Dec 2007	Norway	0	0.0097	100	0.014	100	52	0.0107	52.8

(thin grey line) both described a strong capping inversion. The strength of these inversions have fairly good agreement between the background and observations albeit with slightly different structures, however in cases (a) and (b) there is a clear disagreement in the inversion height. In case (a), the background inversion is approximately 160 m too low and in case (b) the inversion is approximately 175 m too low. In case (c), there is not such a clear positional error and instead the background simply seems too cool in this region. For each case the background and observation data have been assimilated using the three different schemes. A summary of the assimilation results for these three cases with the three different formulations is given in Table I. As expected, in cases (a) and (b) the standard scheme (black solid) has been unable to give a good representation of the inversion height ($\eta_{inv} = -108$ m and -163 m respectively) or the temperature structure in this region ('inv grad' in Table I). In these two cases, the analyses given by the floating BL schemes perform better (given by the dashed line for the Gaussian formulation and dashed-dot line for the non-Gaussian formulation). These examples have been selected as it is seen that in each case the floating BL scheme with Gaussian formulation shifts the inversion in the background to a greater extent than the non-Gaussian formulation. In case (c), where there is no real positional error, the standard scheme has given a fairly good analysis. The floating BL scheme in this cases interprets the cool background relative to the observations as a positional error and so in the analysis has shifted the background inversion up unnecessarily. Using the non-Gaussian formulation means that this shift is slightly smaller and so the error in the inversion height is reduced.

5. Summary and future work

The floating BL scheme is a new method for performing 1D-Var which attempts to correct explicitly for vertical positional errors in the background as well as the usual errors dealt with in standard assimilation schemes (amplitude errors). The scheme works by introducing an additional control variable, *a*, which specifies the degree to which the heights of certain levels are allowed to be modified coherently by the assimilation. In this article this scheme has focused on the BL capping inversion, which is a sharp structure that is often forecast well, but at the wrong height. Compared to standard assimilation, this scheme gives a better representation of the errors associated with the background inversion by introducing an error in the height of the inversion.

It has been shown that, in an idealised case, the floating BL scheme gives an improved analysis, compared to the standard scheme, when there is disagreement in the vertical position of the inversion in the background and the observations, by making it possible explicitly to shift the whole of the background inversion structure. The floating BL scheme introduces nonlinearities into the observation operator and therefore implies a non-quadratic cost function. A non-quadratic cost function can lead to problems with uniqueness of the solution if the inversion is not well observed.

A study of the statistics of the new variable *a* has been conducted. This shows that the background errors of *a* are non-Gaussian and asymmetric which (strictly speaking) violates the assumptions on which Var is normally based. Performing the assimilation with the assumption that the

errors of a are Gaussian leads to the floating BL scheme shifting the background inversion more often than is necessary. We have been able to write the cost function with a more complicated description of the error statistics of a, due to a being a single variable with errors assumed to be uncorrelated with the errors in potential temperature. In this work, a two-component Gaussian mixture model was chosen to describe simultaneously the high probability that the positional error is small at the same time as accounting for the possibility that the positional error is large.

The improved representation of the temperature profile in the analysis given by the floating BL scheme should have a positive impact on the forecast of features directly affected by the height and strength of the inversion, e.g. the formation of BL clouds and the dispersion of pollutants. BL clouds are of particular interest due to their large impact on the propagation of radiation which in turn affects the BL dynamics. However, the diagnosis of clouds is dependent upon humidity as well as temperature. It is intended to extend the floating BL scheme to assimilate humidity as well as temperature. This will allow us to assess the impact that the floating BL scheme may have on the diagnosis of clouds.

Building upon the complexity of the floating BL scheme would also allow an assessment of the scheme's effect on future forecasts. Whether the scheme could lead to an improved forecast of the inversion or whether the forecast model would adjust back to the wrong inversion height is of particular interest. Part of the motivation for this work is based on the assumption that correcting wrong inversion heights and structures in analyses could significantly yield more accurate forecasts of such features and of BL cloud. This scheme, however, makes no attempt to correct the model itself which may contribute to an inaccurate inversion height. A study of the the current floating BL scheme's impact on the positional error in future forecasts would be enlightening.

There are a number of ways that this scheme may be implemented in an operational setting, e.g. within 3D- or 4D-Var. The anticipated research issues associated with this (e.g. those given below) would have to be resolved before the scheme could be used to generate 3D analyses so that realistic forecasts could be made to allow the scheme to be evaluated thoroughly.

In 3D-Var an extra control variable and a displacement function may have to be introduced for each horizontal location. The different meteorological conditions from one location to the next (resulting in the existence or not, and the different character, of inversions) is likely to lead to discontinuities in the 2D field of a values. Consequently, an error covariance matrix would have to be introduced to describe how perturbations in $a(\mathbf{r})$ co-vary with those in $a(\mathbf{r}')$, where $\mathbf{r}$ and $\mathbf{r}'$ are neighbouring horizontal positions. The nature of inversions changes with position and time and so these covariances are likely to be inhomogeneous, anisotropic and flow-dependent. Additionally, as the levels shift, all model variables (and not just temperature) would be carried with them. The effects of this would need careful evaluation to check that the other variables are not degraded as the heights of temperature inversions are corrected. Analyses produced on vertical grids that have been shifted by the floating BL scheme would need to be interpolated back to the fixed model grid before forecasts can be made.

In 4D-Var, other issues emerge in addition to the above. For instance, the extra field of control variables, $a(x, y)$, and the displacement functions may additionally become a function of time. Ignoring this variation over the 4D-Var time window may be problematic, e.g. due to the strong diurnal cycle of the boundary layer (Stull, 1988). Additionally, the perturbation forecast model and its adjoint used within incremental 4D-Var would have to operate on shifted levels. Not only would this require development effort, but may conceivably introduce unwanted imbalances during the propagation of field perturbations.

Appendix

The non-Gaussian background cost term for the floating BL scheme

The probability of the background, $\widetilde{\mathbf{v}}^{\mathrm{b}} = (\widetilde{\mathbf{x}}^{\mathrm{b}}, a = 0)^{\mathrm{T}}$, can be described as the weighted sum of two Gaussian distributions, P_1 and P_2:

$$P_{\mathrm{B}} = (1 - A)P_1 + AP_2, \tag{A.1}$$

where A is a constant $(0 \le A \le 1)$ and

$$P_i = \frac{1}{(2\pi)^{\frac{N+1}{2}} |\widetilde{\mathbf{B}}_{\mathrm{v}_i}|^{1/2}} \exp\left\{ -\frac{1}{2} (\widetilde{\mathbf{v}} - \widetilde{\mathbf{v}}^{\mathrm{b}})^{\mathrm{T}} \widetilde{\mathbf{B}}_{\mathrm{v}_i}^{-1} (\widetilde{\mathbf{v}} - \widetilde{\mathbf{v}}^{\mathrm{b}}) \right\}. \tag{A.2}$$

Here $i = 1, 2$ and N is the length of the background vector, $\widetilde{\mathbf{x}}^{\mathrm{b}}$. $\widetilde{\mathbf{B}}_{\mathrm{v}_i}$ is the background error covariance matrix for the background $\widetilde{\mathbf{v}}^{\mathrm{b}}$:

$$\widetilde{\mathbf{B}}_{\mathrm{v}_i} = \begin{pmatrix} \widetilde{\mathbf{B}}_{\mathrm{x}} & 0 \\ 0 & \sigma_{a_i}^2 \end{pmatrix}. \tag{A.3}$$

Since errors in a are taken to be uncorrelated with errors in $\widetilde{\mathbf{x}}^{\mathrm{b}}$, the part of the distribution concerned with a may be separated as follows (the distribution for amplitude errors in $\widetilde{\mathbf{x}}^{\mathrm{b}}$ remains Gaussian, while only the distribution for a becomes non-Gaussian):

$$P_i = \frac{1}{(2\pi)^{\frac{N+1}{2}} |\widetilde{\mathbf{B}}_{\mathrm{x}}|^{1/2}} \exp\left\{ -\frac{1}{2} (\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}^{\mathrm{b}})^{\mathrm{T}} \widetilde{\mathbf{B}}_{\mathrm{x}}^{-1} (\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}^{\mathrm{b}}) \right\}$$
$$\times \frac{1}{\sigma_{ai}} \exp\left\{ -\frac{a^2}{2\sigma_{ai}^2} \right\}. \tag{A.4}$$

Substituting (A.4) into (A.1) gives

$$P_{\mathrm{B}} = \frac{1}{(2\pi)^{\frac{N+1}{2}} |\widetilde{\mathbf{B}}_{\mathrm{x}}|^{1/2}} \exp\left\{ -\frac{1}{2} \left(\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}^{\mathrm{b}}\right)^{\mathrm{T}} \widetilde{\mathbf{B}}_{\mathrm{x}}^{-1} \left(\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}^{\mathrm{b}}\right) \right\}$$
$$\times \left\{ \frac{1 - A}{\sigma_{a_1}} \exp\left(-\frac{a^2}{2\sigma_{a_1}^2} \right) + \frac{A}{\sigma_{a_2}} \exp\left(-\frac{a^2}{2\sigma_{a_2}^2} \right) \right\}. \tag{A.5}$$

The background cost function is commonly defined as

$$J_{\text{B}} = -\ln P_{\text{B}} + c, \qquad (A.6)$$

where c is a constant which has no effect on the minimisation of the cost function and so can be added without altering the analysis. Substituting (A.5) into (A.6) gives

$$J_{\text{B}} = \frac{1}{2} \left(\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}^{\text{b}} \right)^{\text{T}} \widetilde{\mathbf{B}}_{\text{x}}^{-1} \left(\widetilde{\mathbf{x}} - \widetilde{\mathbf{x}}^{\text{b}} \right)$$
$$- \ln \left\{ \frac{1-A}{\sigma_{a_1}} \exp\left(-\frac{a^2}{2\sigma_{a_1}^2} \right) + \frac{A}{\sigma_{a_2}} \exp\left(-\frac{a^2}{2\sigma_{a_2}^2} \right) \right\},$$
$$(A.7)$$

where the constant c has been added:

$$c = \ln \left\{ \frac{1}{(2\pi)^{\frac{N+1}{2}} |\widetilde{\mathbf{B}}_{\text{x}}|^{1/2}} \right\}. \qquad (A.8)$$

The choice of A dictates the shape of the function. As A tends to 0, P_{B} tends to a Gaussian with $\sigma_a = \sigma_{a_1}$ and as A tends to 1, P_{B} tends to a Gaussian with $\sigma_a = \sigma_{a_2}$.

References

Alexander GD, Weinman JA, Schols JL. 1998. The use of digital warping of microwave integrated water vapour imagery to improve forecasts of marine extratropical cyclones. *Mon. Weather Rev.* **126**: 1469–1495.

Andersson E, Järvinen H. 1999. Variational quality control. *Q. J. R. Meteorol. Soc.* **125**: 697–722.

Bannister RN. 2008. A review of forecast-error covariance statistics in atmospheric variational data assimilation. I: Characteristics and measurements of forecast-error covariances. *Q. J. R. Meteorol. Soc.* **134**: 1951–1970.

Bannister RN, Katz D, Cullen MJP, Lawless AS, Nichols NK. 2008. Modelling of forecast errors in geophysical fluid flows. *Intl. J. Num. Methods Fluids* **56**: 1147–1153.

Beezley JD, Mandel J. 2008. Morphing ensemble Kalman filters. *Tellus* **60A**: 131–140.

Bennett AF. 1992. *Inverse modelling in ocean and atmosphere*. Cambridge University Press: Cambridge, UK.

Bennett AF. 2002. *Inverse methods in physical oceanography*. Cambridge University Press: Cambridge, UK.

Bouttier F, Courtier P. 1999. 'Data assimilation concepts and methods'. In *Notes from training course on data assimilation and use of satellite data.* ECMWF: Reading, UK. http://www.ecmwf.int/newsevents/training/rcourse_notes

Bowler NE, Arribas A, Mylne KR, Robertson KB, Beare SE. 2008. The MOGREPS short-range ensemble prediction system. *Q. J. R. Meteorol. Soc.* **134**: 703–722.

Brewster KA. 2003. Phase-correcting data assimilation and application to storm-scale numerical weather prediction. Part I: Method, description and simulation testing. *Mon. Weather Rev.* **131**: 480–492.

Cullen MJP. 1993. The unified forecast/climate model. *Meteorol. Mag.* **122**: 1993–1997.

Eyre JR, Kelly GA, McNally AP, Andersson E, Persson A. 1993. Assimilation of TOVS radiance information through one-dimensional variational analysis. *Q. J. R. Meteorol. Soc.* **119**: 1427–1463.

Fowler AM, Bannister RN, Eyre JR. 2010. Characterising the background errors for the boundary-layer capping inversion. *Austral. Meteorol. Oceanic J.* **59**: 17–24.

Hoffman RN, Liu Z, Louis J-F, Grassoti C. 1995. Distortion representation of forecast errors. *Mon. Weather Rev.* **123**: 2758–2770.

Hólm EV. 2003. 'Revision of the ECMWF humidity analysis: Construction of a Gaussian control variable'. In *Proceedings of the ECMWF/GEWEX Workshop on humidity analysis*, 8–11 July 2002. ECMWF: Reading, UK. http://www.ecmwf.int/publications/library/ecpublications/_pdf/workshop/2002/Humidity/holm.pdf

Ide K, Courtier P, Ghil M, Lorenc AC. 1997. Unified notation for data assimilation: Operational, sequential and variational. *J. Meteorol. Soc. Japan* **75**: 181–189.

Jones CD, MacPherson B. 1997. A latent heat nudging scheme for assimilation of precipitation data into an operation mesoscale model. *Meteorol. Appl.* 269–277.

Lawson WG, Hansen JA. 2005. Alignment error models and ensemble-based data assimilation. *Mon. Weather Rev.* **133**: 1687–1709.

Lorenc AC. 1986. Analysis methods for numerical weather prediction. *Q. J. R. Meteorol. Soc.* **112**: 1177–1194.

Lorenc AC. 2007. 'Ideas for adding flow-dependence to the Met Office VAR system'. In *Proceedings of Workshop on flow-dependent aspects of data assimilation*, 11–13 June 2007. ECMWF: Reading, UK. http://www.ecmwf.int/publications/library/do/references/list/14092007

Mariano AJ. 1990. Contour analysis: A new approach for melding geophysical fluids. *J. Atmos. Oceanic Technol.* **7**: 285–295.

Michel Y. 2010. Data assimilation of tropopause height using dry intrusion observations. *Mon. Weather Rev.* **138**: 101–122.

Milionis AE, Davies TD. 1992. A five-year climatology of elevated inversions at Hemsby (UK). *Int. J. Climatol.* **12**: 205–215.

Nehrkorn T, Hoffman RN, Grassotti C, Louis J-F. 2003. Feature calibration and alignment to represent forecast model errors: Empirical regularization. *Q. J. R. Meteorol. Soc.* **129**: 195–218.

Rabier F. 2005. Overview of global data assimilation developments in numerical weather prediction centres. *Q. J. R. Meteorol. Soc.* **131**: 3215–3233.

Ravela S, Emanuel KE, McLaughlin D. 2007. Data assimilation by field alignment. *Physica D* **230**: 127–145.

Stull RB. 1988. *An introduction to Boundary-Layer Meteorology*. Kluwer Academic Publishers: Dordrecht, Netherlands.

Thiebaux HJ, Julian PR, DiMego GJ. 1990. 'Areal versus collocation data quality control'. In *Proceedings of International Symposium on assimilation of observations in meteorology and oceanography*, Clermont-Ferrand, France. World Meteorological Organization: Geneva. 255–260.