Met Office

Meteorology Research and Development

Cloud-top properties from SEVIRI data using 1-D variational analysis

Technical Report No. 518

R.B.E.Taylor

email:nwp_publications@metoffice.gov.uk

Cloud-top properties from SEVIRI data using 1-D variational analysis

R.B.E.Taylor

Document Change History

First draft, Ruth Taylor. Comments recieved from Roger Saunders.	08/10/2007
Updated version, Ruth Taylor. Further comments recieved from Roger Saunders.	24/04/2008
Approved for publication by Roger Saunders.	16/05/2008
Submitted as Met R&D Technical Report.	17/06/2008

Abstract

The use of 1-D variational analysis (1D-Var) to obtain cloud-top height and a measure of cloud amount from SEVIRI radiances is described. A multiple field-of-view formulation is adopted, providing both computational savings and a useful constraint on retrievals. The cloud products are plausible but there are systematic differences to products derived from other sensors, and to products derived from SEVIRI data using other algorithms. There are also some aspects of the minimisation accuracy that are of concern, and addressing these would make easier a more extensive validation of the cloud properties against data from other sensors.

1 Introduction

Cloud products are routinely generated from satellite imagery at the UK Met Office for use in forecasting and nowcasting applications, and there is a continuing research effort aimed at ensuring they are of as high quality as possible. Variational analysis provides an objective method of deriving cloud products that are consistent with one another and with other atmospheric properties. This study applies one-dimensional variational analysis (1D-Var) to radiances observed in the infra-red channels of the Spinning Enhanced Visible and Infrared Imager (SEVIRI), on board the Meteosat Second Generation (MSG) series of geostationary satellites. Cloud-top pressure and a measure of cloud cover are retrieved through minimisation of a cost function, using background temperature and humidity information from model forecasts.

Cloudy 1D-Var is a computationally expensive technique, due both to the use of a radiative transfer model and to the highly non-linear effect of cloud on the forward-modelled radiances, which requires minimisation to be performed using an iterative technique. However, even at northern European latitudes, SEVIRI images have a relatively high resolution compared to that of the Met Office's current North Atlantic European (NAE) model, and a simulation study (Szyndel et al., 2004) has suggested that this disparity can be exploited via a multiple field-of-view formulation, in which one retrieval is carried out for a number of pixels, sharing a cloud-top pressure and atmospheric profile. In the simulation study, retrieval skill is improved, whilst use of a relatively simplistic cloud model minimises the number of radiative transfer calculations required.

Given the computational expense of 1D-Var, it is unlikely to prove a suitable method for routine generation of cloud products in the near future, but there may be cases where it can improve on the current operational cloud detection and measurement scheme. In addition, experience gained through this work may be relevant to the direct assimilation of SEVIRI radiances in the forecast model.

In this study, different configurations of 1D-Var are tested and the results compared against the Met Office's current operational cloud-characterisation scheme (developed for SEVIRI data) and against MODIS cloud products, which are completely independent. Data from the Cloud Profiling Radar (CPR) on the CloudSat platform and the Cloud-Aerosol Lidar with Orthogonal Polarization (CALIOP) aboard the Calipso satellite was not yet available at the time of the case studies used.

2 1D-Var Retrieval Scheme

Given a set of observations $\mathbf{y^o}$, an *a priori* or "background" estimate of the atmospheric profile $\mathbf{x_b}$ and a measure of the uncertainty in each, probabilistic arguments (Lorenc, 1986) identify the most likely value of the "true" atmospheric profile $\mathbf{x}$ as that which minimises the cost function

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x}-\mathbf{x_b})^{\mathbf{T}}\mathbf{B}^{-1}(\mathbf{x}-\mathbf{x_b}) + \frac{1}{2}\{\mathbf{y^o}-H(\mathbf{x})\}^{\mathbf{T}}(\mathbf{E}+\mathbf{F})^{-1}\{\mathbf{y^o}-H(\mathbf{x})\}, \qquad (1)$$

where $H(\mathbf{x})$ are observations simulated using a forward (radiative transfer) model and $\mathbf{B}$, $\mathbf{E}$ and $\mathbf{F}$ are error covariance matrices for $\mathbf{x_b}$, $\mathbf{y^o}$ and $H(\mathbf{x})$ respectively, assuming Gaussian statistics.

2.1 Cloud Model

For this study, forward modelling was performed using RTTOV 6.9 (Saunders *et al.*, 1999). The profile $\mathbf{x}$ includes cloud parameters. Cloud is represented as a single layer of geometrically thin,

Channel	Measurement error K	Forward modelling error K	Combined error R K
6.2μm	0.40	1.27	1.33
7.3μm	0.48	0.99	1.10
8.7μm	0.15	0.78	0.79
10.8μm	0.13	0.78	0.79
12.0μm	0.21	0.78	0.81
13.4μm	0.29	0.49	0.57

Table 1: Estimated standard deviation values used to form the diagonal elements of the error matrices **E** and **F**, together with the corresponding values obtained for the combined error **R**.

grey cloud, with the forward-modelled radiance for a partially cloudy scene being a linear combination of the clear-sky radiance $H^{\rm cl}(\mathbf{x}_a)$ and the radiance calculated for 100% coverage of opaque cloud at height $P_{\rm cl}$, $H^{\rm op}(P_{cl}, \mathbf{x}_a)$:

$$H(P_{cl}, N, \mathbf{x}_a) = (1 - N)H^{\rm cl}(\mathbf{x}_a) + NH^{\rm op}(P_{cl}, \mathbf{x}_a). \quad (2)$$

Here N is an "effective cloud amount" (ECA), representing fractional cloud cover and emissivity simultaneously, and $\mathbf{x}_a$ is the profile $\mathbf{x}$ excluding cloud parameters. We make the "grey cloud" assumption, in which emissivity and thus ECA are assumed to be independent of wavelength for the channels used in this study.

Forward modelling of radiances is performed on the 43 fixed pressure levels of the RTTOV model. Piece-wise linear interpolation is used to obtain $H^{\rm op}(P_{cl}, N, \mathbf{x}_a)$ for $P_{\rm cl}$ falling between two of these levels.

Note that brightness-temperature rather than radiance space was used for performing the minimisation described in this study, and "H" should be interpreted according to context.

2.2 Application to SEVIRI

Six SEVIRI channels were used in the retrieval: the 6.2μm and 7.3μm water vapour channels, the 8.7μm, 10.8μm and 12.0μm infra-red window channels and the 13.4μm CO_2 channel. Bias correction was not applied, because the radiance biases reported are much smaller than the cloud signals.

Both the observation error matrix **E** and the forward-modelled error matrix **F** were taken to be diagonal, with diagonal elements set as shown in Table 1. Observation errors were taken from Schmid (2004) and forward-modelling errors were estimated from experience with channels on other instruments that closely match the SEVIRI channels (P.N.Francis, personal communication). The combined observational- and forward-modelled error matrix **R** was obtained by adding the variances **E** and **F**.

The full retrieval vector contained cloud-top pressure and effective cloud fraction(s), temperature on the 43 levels of RTTOV, humidity on the lowest 28 of these levels (corresponding approximately to the troposphere), 2-m air temperature and humidity values, surface skin temperature and surface pressure. Background temperature, pressure and humidity information for this study were obtained from the Met Office's 38-level mesoscale NWP model (the limited area model in operational use at the time of the case studies) interpolated onto the 43 levels of RTTOV. Background error covariance (**B**) values for non-cloud parameters were taken from an ECMWF

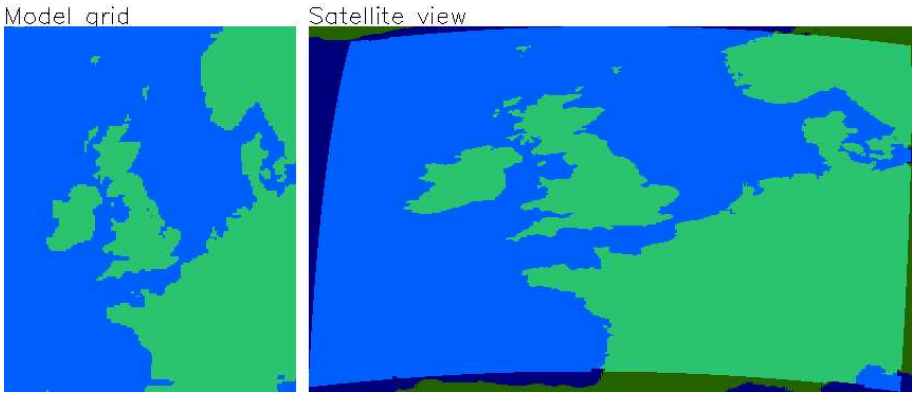

Figure 1: Mesoscale model area used for cloudy 1D-Var retrieval in this case study, displayed on the model (left) and satellite pixel (right) grids.

forecast error covariance matrix. Background values for the cloud parameters were estimated (crudely) using the background relative humidity profile, while the retrieval's sensitivity to these values was minimised by using very large values for the corresponding diagonal **B**-matrix elements. Other (off-diagonal) elements involving cloud parameters in **B** were set to zero, as no prior assumptions were being made concerning the dependence of cloud on other profile elements. Note that this treatment of the cloud background means that brightness temperatures obtained by forward-modelling cloudy background profiles are not expected to match observed brightness temperatures particularly well, so "O-B"s are not used as a diagnostic in this study.

For some tests the retrieval scheme used a reduced retrieval vector, excluding most or all of the temperature elements. Removal of temperatures reduces the length of the profile vector by a factor of around a half, reducing the time required for matrix manipulation and, in turn, overall computational expense; it can be justified if confidence in the background temperatures, as provided by the mesoscale model, is high. Allowing temperatures to vary in response to SEVIRI information may not then result in significant improvements to the cloud retrieval. The results of these tests are discussed in Section 3.6.1.

The domain of the mesoscale model, which is defined using latitudes and longitudes on a rotated-pole grid, is shown in Figure 1, together with its mapping onto a segment of the full-earth disk seen by SEVIRI (the "pixel grid"). Across the scene, the median number of satellite pixels per model gridbox is 7; the range of latitudes and hence of satellite zenith angles for this region is such that the mean number of pixels per gridbox is significantly larger for the southern portion of the region than towards the northern edge. Distribution of pixels across gridboxes is a task routinely performed by AUTOSAT for other applications, and is complicated by the requirement that, along coastlines, land (sea) pixels be matched to land (sea) gridboxes. In Figure 1, the satellite's land-sea mask includes a condition on the model's land-sea mask, accounting for the slightly ragged appearance of portions of the coastline.

For this study retrievals were carried out for sea points only, because of greater confidence in the surface characterisation.

2.3 Minimisation method

Finding the minimum of the cost function $J(\mathbf{x})$ was done using the implementation of the Marquardt-Levenberg method described in Szyndel *et al.* (2004), in which successive approximations to the

minimum are given by

$$\mathbf{x}_{n+1} = \mathbf{x}_n - (\nabla_{\mathbf{x}}\nabla_{\mathbf{x}}J(\mathbf{x}_n) + \gamma\mathbf{I})^{-1}(\nabla_{\mathbf{x}}J), \tag{3}$$

$\mathbf{I}$ being the identity matrix. For small values of the scalar γ this approximates Newton's method, while for large γ it is equivalent to the steepest-descent method with step size γ^{-1}.

A new profile $\mathbf{x}_{n+1}$ is accepted only if $J(\mathbf{x}_{n+1}) < J(\mathbf{x}_n)$; in that case γ is reduced and the iteration proceeds. Otherwise $\mathbf{x}_{n+1}$ is recalculated with an increased value of γ, as many times as required to produce a reduction in J. Minimisation is declared complete when the reduction in the cost function in one iteration is less than 1%, without any increase being made to γ. Retrievals requiring more than 12 iterations were abandoned in this study as "failures".

Note that if γ is increased within an iteration, the recalculation of $\mathbf{x}_{n+1}$ requires some extra computation (including a call to the radiative transfer model in order to re-test the cost function). The time taken for a retrieval is therefore not necessarily proportional to the number of iterations required.

First-guess values for the cloud parameters were obtained using the minimum-residual method (Eyre and Menzel, 1989) and brightness temperatures measured in the 8.7μm, 10.8μm, 12.0μm and 13.4μm channels. The minimum-residual technique is in effect a 1D-Var scheme with a very short profile vector, consisting only of the cloud parameters CTP and ECA; this means that best-fit forward-modelled radiances can be found without resort to numerical minimisation, so it is much cheaper than 1D-Var. The method uses the same NWP information as the 1D-Var retrievals as an input to the radiative transfer model, but it is not allowed to vary to fit radiances. Water vapour channels are excluded because of a suspicion that the information they provide might be too sensitive to the possibly inaccurate model background. Retrieval was carried out for all pixels, even those which are likely to be cloud-free, as retrieval performance in these circumstances is still of interest.

In the course of 1D-Var minimisation, variables which transgressed predefined physical limits were reset to that limit. An upper limit to cloud-top height was initially placed at 200 hPa but moved upwards to 100 hPa when it was found that using a limit of 200 hPa constrained a significant proportion of retrievals. Effective cloud amount was allowed to take only values between 0 and 1 for most of the retrieval configurations presented here. Other profile elements were effectively constrained by requirements on the inputs to the RTTOV model.

2.4 Error characterisation

Limited use was made in this study of the second derivative of the cost function, given by

$$\mathbf{J}''(\mathbf{x}) = \nabla_{\mathbf{x}}\nabla_{\mathbf{x}}J(\mathbf{x}) = \mathbf{B}^{-1} + \mathbf{H}^T(\mathbf{x})\mathbf{R}^{-1}\mathbf{H}(\mathbf{x}), \tag{4}$$

where $\mathbf{H}(\mathbf{x})$ is the Jacobian of $H(\mathbf{x})$, $\nabla_{\mathbf{x}}H(\mathbf{x})$. Assuming Gaussian statistics, $\nabla_{\mathbf{x}}\nabla_{\mathbf{x}}J$'s inverse $\mathbf{A}(\mathbf{x})$ gives the error covariance of the retrieved profile, henceforth referred to as "analysis error". $\mathbf{J}''(\mathbf{x})$ and $\mathbf{A}$ give a measure of the sharpness of the cost function minimum in the linear limit. Note that this is not the shape of the cost-function surface explored by the minimisation scheme, except in cases that are close to linearity, where $\mathbf{H}$ has only a weak dependence on $\mathbf{x}$.

Due to computational constraints, only the diagonal of the analysis error covariance matrix was output for this study. Profile vectors at intermediate steps in the minimisation path were output only for the diagnostic work described in Section 3.4, as were forward-model Jacobians for the cloud parameters.

2.5 Multiple Fields of View ("MFOV") formulation

The horizontal resolution of the mesoscale model is approximately 12km. SEVIRI resolution, dependent on the satellite's viewing angle, is of order 5km in the latitude range being used for this study. (Because of the relatively high latitude, it is also different in the East-West and North-South directions.) Thus several SEVIRI pixels will typically share a common vertical background profile.

In the multiple-FOV formulation (Szyndel *et al.* (2004)), a retrieval is performed for each (model) gridbox, rather than for each (satellite) pixel. The retrieval vector contains an effective cloud amount (ECA) for each satellite pixel falling within that gridbox and one cloud-top pressure (CTP) common to all such pixels. Any variation in radiance between the pixels associated with a given model gridbox can only be accounted for in the retrieval scheme by allowing ECA to vary between pixels. With the cloud model described in section 2.1, this approximation enables computational savings: by partial differentiation of Equation 2, we obtain

$$\frac{\partial H(P_{cl}, N, \mathbf{x}_a)}{\partial P_{cl}} = N \times \frac{\partial H^{\mathrm{op}}(P_{cl}, \mathbf{x}_a)}{\partial P_{cl}} \tag{5}$$

and

$$\frac{\partial H(P_{cl}, N, \mathbf{x}_a)}{\partial N} = -H^{\mathrm{cl}}(\mathbf{x}_a) + H^{\mathrm{op}}(P_{cl}, \mathbf{x}_a), \tag{6}$$

so that if P_{cl} and $\mathbf{x}_a$ are the same for all FOVs, Jacobians for the cloud parameters can, in principle, be obtained for all FOVs from a single calculation of $H^{\mathrm{op}}(P_{cl}, \mathbf{x}_a)$ and $\partial H^{\mathrm{op}}(P_{cl}, \mathbf{x}_a)/\partial P_{cl}$. Forward modelling of opaque radiances therefore needs to be carried out once per gridbox per iteration, rather than once per pixel per iteration.

Note that in RTTOV 6.9, $\partial H^{\mathrm{op}}(P_{cl}, \mathbf{x}_a)/\partial P_{\mathrm{cl}}$ is obtained for cloud-top pressure on RTTOV model levels, but is not interpolated for cloud-top pressures between RTTOV levels. This calculation is consistent with piece-wise linear interpolation of radiances between RTTOV levels, and is significant in the context of an apparent tendency for retrievals to have cloud-top pressure peaking on RTTOV levels (Section 3.4.2).

The MFOV approach means that the length of the retrieval vector depends on the number of pixels per gridbox N_{pix}. If the retrieval includes all temperatures, the retrieval vector has length $76 + N_{pix}$. Using the median value of N_{pix}, 7, gives a typical retrieval vector length of 83. The number of measurements contributing to the retrieval is given by the number of channels (6 for most of the configurations used for this study) multiplied by N_{pix}, giving 42 for the median case.

When determining first-guess values for the cloud parameters, CTP and ECA are calculated for each pixel, and the common first-guess CTP is set to that of the cloudiest FOV, where it is assumed that the minimum-residual technique is most accurate. The average value of this common first-guess CTP across the scene was found not to be significantly different to the average of the values for individual FOVs, that is, the cloudiest FOV was not found to have cloud that was systematically higher or lower than other FOVs in the retrieval.

Szyndel *et al.* (2004) describes a simulation study using the MFOV formulation. The cost-function surface for a profile including P_{cl} and N as cloud parameters typically contains a long "valley" connecting high cloud/low coverage retrievals and low cloud/high coverage retrievals, such that points in the valley are difficult to distinguish on the grounds of observed brightness temperature. The multiple-FOV formulation should help to constrain the retrievals. The simulation study concluded that:

- the skill of multiple-FOV retrievals increases with cloud height and with ECA, as Eyre (1989) had earlier demonstrated for single-FOV retrievals;

- in most scenarios, particularly for low cloud and for pixels with low ECAs in a mixed scene, multiple-FOV retrievals are more skilful than single-FOV retrievals;
- in many cases, this improvement in skill is seen even where "true" CTP varies between pixels within a gridbox.

This work applies the techniques used in Szyndel *et al.* (2004) to measured radiances.

2.6 Use of an additional cost function term

An additional, profile-dependent, term $J_{add}(\mathbf{x})$ can be used to change the shape of the cost-function surface and thus affect the path of a minimisation through profile space. As well as using $J_{add}(\mathbf{x})$ itself to assess convergence, extra terms must be added to both $\nabla_{\mathbf{x}} J$ and $\nabla_{\mathbf{x}} \nabla_{\mathbf{x}} J(\mathbf{x})$ in equation 3 above. An example is discussed in Section 3.4.1.

2.7 A variant of the Marquardt-Levenberg algorithm

In equation 3, it can be seen that if the various elements of $\mathbf{x}$ have different numerical magnitudes, the sizes of the elements of $\nabla_{\mathbf{x}} \nabla_{\mathbf{x}} J(\mathbf{x})$ relative to those of $\gamma \mathbf{I}$ are likely to be different for different elements of $\mathbf{x}$. In some implementations of the Marquardt-Levenberg algorithm, the term $\gamma \mathbf{I}$ is replaced by $\gamma \mathbf{D}$, where $\mathbf{D}$ is a diagonal scaling matrix (Rodgers (2000)). For this study, $\mathbf{D}$ was constructed using the diagonal elements of $\mathbf{A}^{-1}$, where $\mathbf{A}$ is the theoretical error defined in section 2.4. This choice of $\mathbf{D}$ involved no significant computational overhead. Section 3.4.1 describes some effects of making this substitution.

3 Results

Two cases are studied here, 0930Z, 30.11.2004 and 1200Z, 17.01.2006. The second corresponds to the date chosen for a EUMETSAT-sponsored cloud intercomparison study led by the Swedish Meteorological and Hydrological Institute, Norrköping, Sweden (Thoss, 2006; see also Section 3.1). Images derived directly from SEVIRI observations (an RGB image and brightness temperature measured in the $10.8 \mu m$ channel) are shown for the two cases in Figures 2 and 3.

Various configurations of 1D-Var were investigated for this report. For convenience, a summary of the configurations described in the following sections is given in Table 2, with labels which will be used to refer to different configurations in the subsequent text. Note that configurations can be divided into two groups: those with prefixes BAS, AXF, MLV, SFV, SHP and INH differ in their implementations of 1D-Var, but use the same set of input data, whilst those with prefixes P32, P33, N87 and OFG are all based on the control configuration but are supplied with different input information.

The "baseline" results for CTP and ECA are shown for the 2004 case (BAS2004) in Figure 4 and for the 2006 case (BAS2006) in Figure 6. The 1D-Var does not explicitly include a cloud mask and retrieval is performed for all pixels. "Cloud-free" pixels are shown to aid interpretation of the figures. The choice of a threshold of 10% to indicate "cloud-free" pixels is somewhat arbitrary and, as discussed in Section 3.2.1, is more helpful for BAS2004 than for BAS2006.

In the rest of this section we present a comparison between 1D-Var (BAS2004 and BAS2006) and other retrievals before discussing the impact of using different 1D-Var configurations or different sets of input data. It will be seen that altering some aspect of the 1D-Var retrieval generally

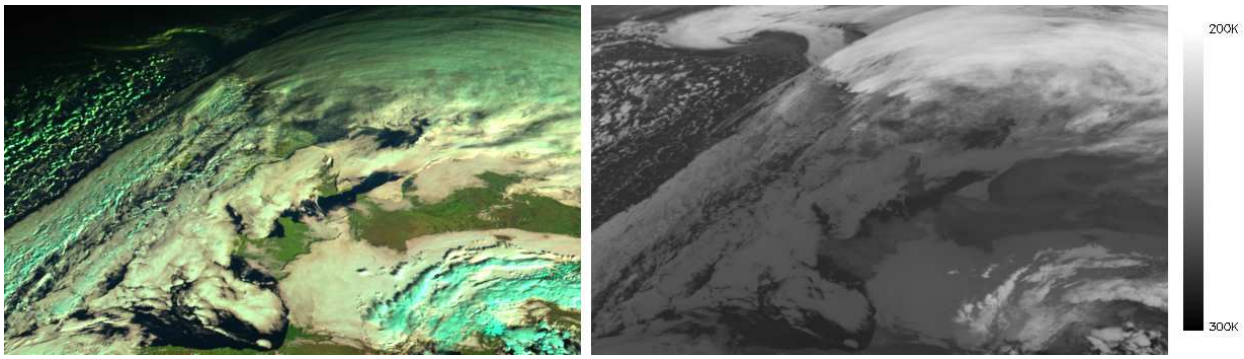

Figure 2: Left: False-colour image for 30.11.2004 0930Z, generated using SEVIRI's two (standard-resolution) visible channels, at 0.6μm and 0.8μm, and the near-infrared channel at 1.6μm. (None of these channels are used in the 1D-Var retrievals.) The cyan and beige colours give an indication of whether clouds consist of ice or liquid water respectively. Note that in this case, the northwestern portion of the image is compromised by the lack of sunlight, and shadows are apparent in all areas. Right: Infra-red brightness temperature, as observed in SEVIRI's 10.8μm channel, for the same case.

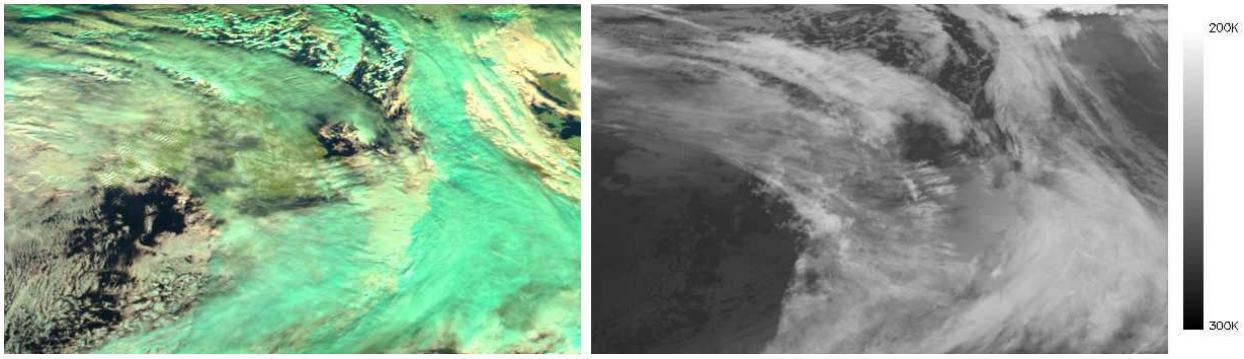

Figure 3: As Figure 2, for the 17.01.2006 1200Z case. Here, the RGB image is not affected by the lack of illumination at visible wavelengths.

Brief description	label for 2004 run	label for 2006 run
Control configuration (multiple FOV)	BAS2004	BAS2006
Additional cost function term	ACF2004	-
Variant Marquardt-Levenberg	MLV2004	-
Single-FOV test	SFV2004	-
True shared profile	SHP2004	-
Initial Jacobian test	INH2004	-
Fixing all temperatures	P322004	-
Fixing upper-air temperatures	P332004	-
Not using 8.7μm channel	N872004	-
Using AUTOSAT products as first guess	OFG2004	-

Table 2: Summary of 1D-Var configurations used, with labels.

has a rather small systematic effect on the scene as a whole, compared to the difference between 1D-Var and other products, although certain individual pixels may be affected dramatically by a change in retrieval configuration.

3.1 Sources of data for comparison

1D-Var retrievals are compared in detail in this report to results from the Met Office's current operational cloud scheme, derived from the same SEVIRI (and background) data using a different algorithm, and with MODIS cloud products, which are obtained by applying a different algorithm to different data. A brief comparison to cloud products contributed to the Norrköping cloud intercomparison workshop was also performed.

The Met Office's operational scheme (Saunders *et al.*, 2006) (henceforth referred to as the "AUTOSAT" scheme) incorporates a hierarchy of methods. The minimum-residual technique (Eyre and Menzel, 1989; see also Section 2.3) is applied first, using information from the 10.8-, 12.0- and 13.4μm SEVIRI channels. The same cloud model is used as for 1D-Var retrievals (see Section 2.1). If the minimum-residual scheme fails or produces a very large calculated error (as is usually the case for lower cloud), the stable-layer and profile-matching methods (Moseley, 2003) are attempted in turn. Of these methods, only the minimum-residual technique allowed fractional cloud cover when this study was performed. The other techniques use observations only from the 10.8μm SEVIRI channel.

Cloud-top pressures from the AUTOSAT scheme are shown in the top panels of Figures 5 and 7. The AUTOSAT processing for the 2006 case included substantial revisions to the processing used to select which of the three schemes in the hierarchy to apply to each pixel.

MODIS data were obtained from the closest available overpasses of the AQUA and TERRA satellites. The cloud-top pressure products (Menzel *et al.*, 2002) are included in Figures 5 and 7. In the 2004 case, a TERRA overpass occurred 90 minutes after the SEVIRI data were recorded, and a spatial misalignment affects "pixel-by-pixel" comparisons. In the 2006 case, TERRA and AQUA overpasses both occurred less than fifteen minutes after the SEVIRI image was recorded. Like the AUTOSAT products, the MODIS products also involve a choice of methods: "CO_2 slicing" and a simpler single-channel brightness temperature matching method, the Equivalent Blackbody Brightness Temperature (EBBT) method. Only the CO_2 slicing technique allows cloud fraction and emissivity to be less than unity and so there is a large proportion of pixels for which ECA = 1.

Both the AUTOSAT and MODIS cloud products include a cloud mask, determined before the retrieval of CTP and ECA. The schemes do not attempt retrieval for pixels judged to be cloud-free.

Data was contributed to the Norrköping cloud intercomparison workshop by several centres. The data is used here only to demonstrate the degree to which the 1D-Var scheme departs from any consensus between the other schemes for this one case, so individual schemes are not examined in detail (or identified) here. Five schemes from three other centres were included in the comparison, all using SEVIRI observations but different background information. Cloud masks were also submitted to the intercomparison.

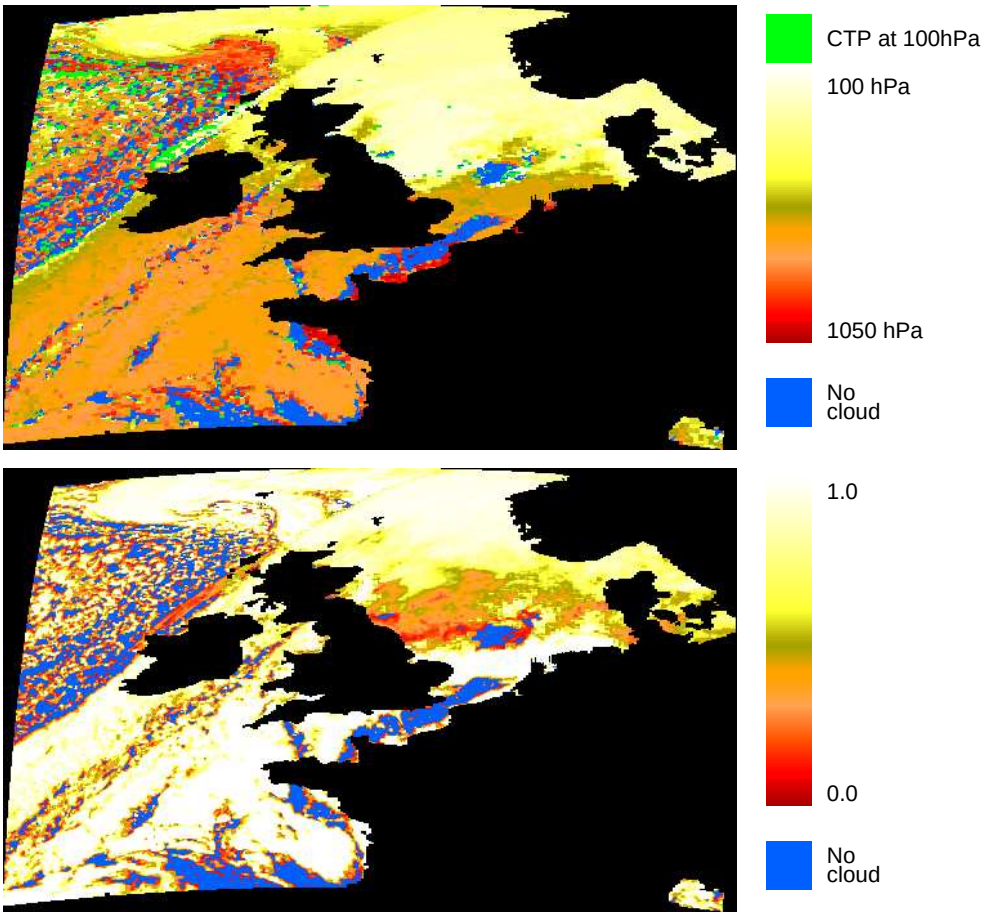

Figure 4: (Cloud-top pressure (CTP) (top) and effective cloud amount (bottom) retrieved from SEVIRI data for the 30.11.2004 0930Z case, using the "baseline" 1D-Var configuration. Retrieval was performed for all sea points; the areas shown in blue indicate pixels for which the retrieved cloud amount is less than 10%. Pixels shown in dark green in the CTP plot are those for which retrieved cloud is at the highest altitude allowed by the retrieval scheme, corresponding to a pressure of 100hPa.

3.2 Retrieved cloud properties: control configuration compared to other retrievals

3.2.1 Comparison to AUTOSAT products

In Figures 4 and 6, the 1D-Var retrievals are indicated as being cloud free where retrieved ECA is below 0.1. The current 1D-Var scheme does not include a separate cloud mask; if required, the AUTOSAT cloud mask could easily be applied before 1D-Var retrieval. The AUTOSAT cloud mask is derived using a sophisticated sequence of tests, and it is not surprising that putting a (somewhat arbitrary) threshold on 1D-Var ECA does not produce a particularly good comparison to the AUTOSAT cloud mask. Using a threshold of 0.1 for ECA for BAS2004 gives reasonably consistent cloud-free areas (12.0% of pixels being classified as cloud-free by AUTOSAT, compared to 11.4% of 1D-Var pixels, with 9.0% of pixels being cloud-free for both), but the same is not the case for BAS2006, where 9.9% of pixels are classified as cloud-free by AUTOSAT but only 2.9% by 1D-Var, with an overlap of 1.9%. Some of the 1D-Var cloud-free points, those appearing at the edges of cloud sheets, where radiance changes rapidly, may be due to the MFOV formulation accounting for radiance variation across a model gridbox with ECA (as it has to) rather than CTP.

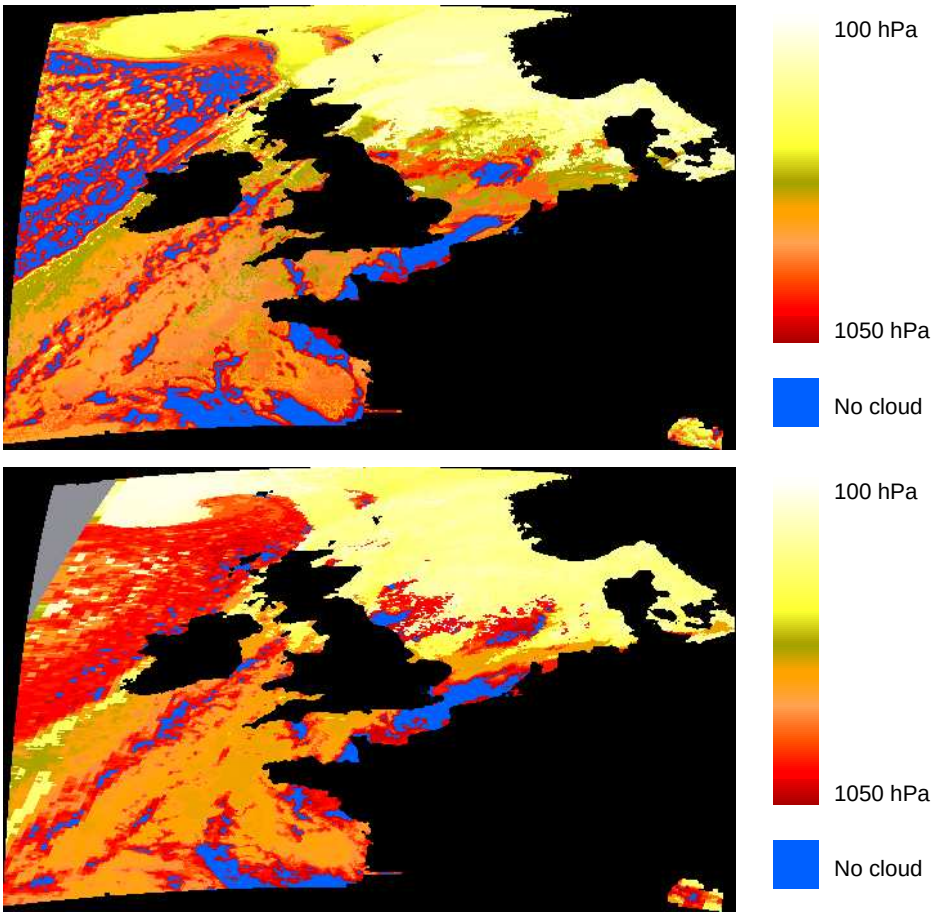

Figure 5: Cloud-top pressures obtained from SEVIRI via AUTOSAT (top) and from MODIS (bottom) for the 2004 case. In both cases products are obtained for all areas, land as well as sea, but land points have been masked out here for ease of comparison with 1D-Var results. MODIS cloud-top pressure is derived from the instrument on the TERRA platform, and is a composite of three images, recorded at 1055Z, 1100Z and 1105Z and "reprojected" onto SEVIRI's pixel grid simply by finding the MODIS pixel closest in location to each SEVIRI pixel. The relative resolutions of the datasets are such that, with this method, several SEVIRI pixels may derive their cloud information from the same MODIS datapoint, especially towards the edges of each swath. Grey areas are not covered by the MODIS swaths.

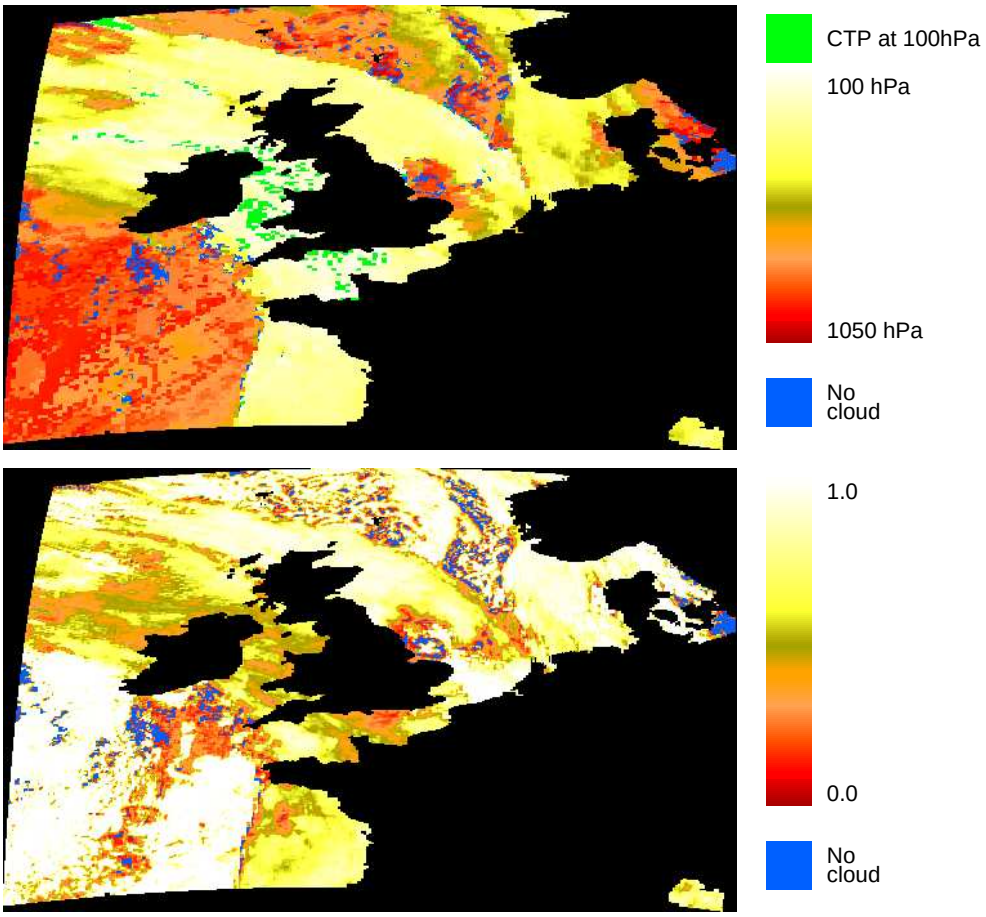

Figure 6: As Figure 4, for the 17.01.2006 1200Z case.

Most of the analysis in the rest of this report does not include a test on retrieved ECA to screen out pixels with very little cloud, or exclude the pixels which the AUTOSAT cloud mask indicates are cloud-free. Although there are good arguments for doing this, most analysis is performed in terms of model gridboxes, not pixels. Using model gridboxes effectively counts retrievals rather than pixels, and the area corresponding to a model gridbox varies less across the scene than the area corresponding to a pixel. To decide whether a model gridbox should be described as cloud free, it would be necessary to decide how to use information aggregated from individual pixels.

For BAS2004, AUTOSAT CTP tends to be more variable across the scene than 1D-Var CTP, and AUTOSAT ECA (not shown) tends to be less variable than 1D-Var ECA. This is apparent (in areas of lower cloud) in the western portion of the 2004 cloud field, where the AUTOSAT scheme gives a lot more of the cloud very low tops, particularly around the edges of bodies of cloud, but tends to characterise the cloud as full or nearly full. Where cloud is broken, as in the northwestern region of the image, the spatially-averaged cloud-top height can be significantly greater for 1D-Var than for AUTOSAT. The two methods appear to be using different parameters to account for variations in brightness temperature at the edges of clouds.

At the very edge of cloud, the restriction on variation in CTP on the scale of the model gridbox probably has some effect, but in general the CTP variations seen in the AUTOSAT fields across bodies of cloud have a larger spatial scale. The use of three different retrieval methods in the AUTOSAT scheme is probably more significant. Only one of these methods (the minimum-residual method) allows for partial cloud cover; the other, simpler, methods (stable-layers and

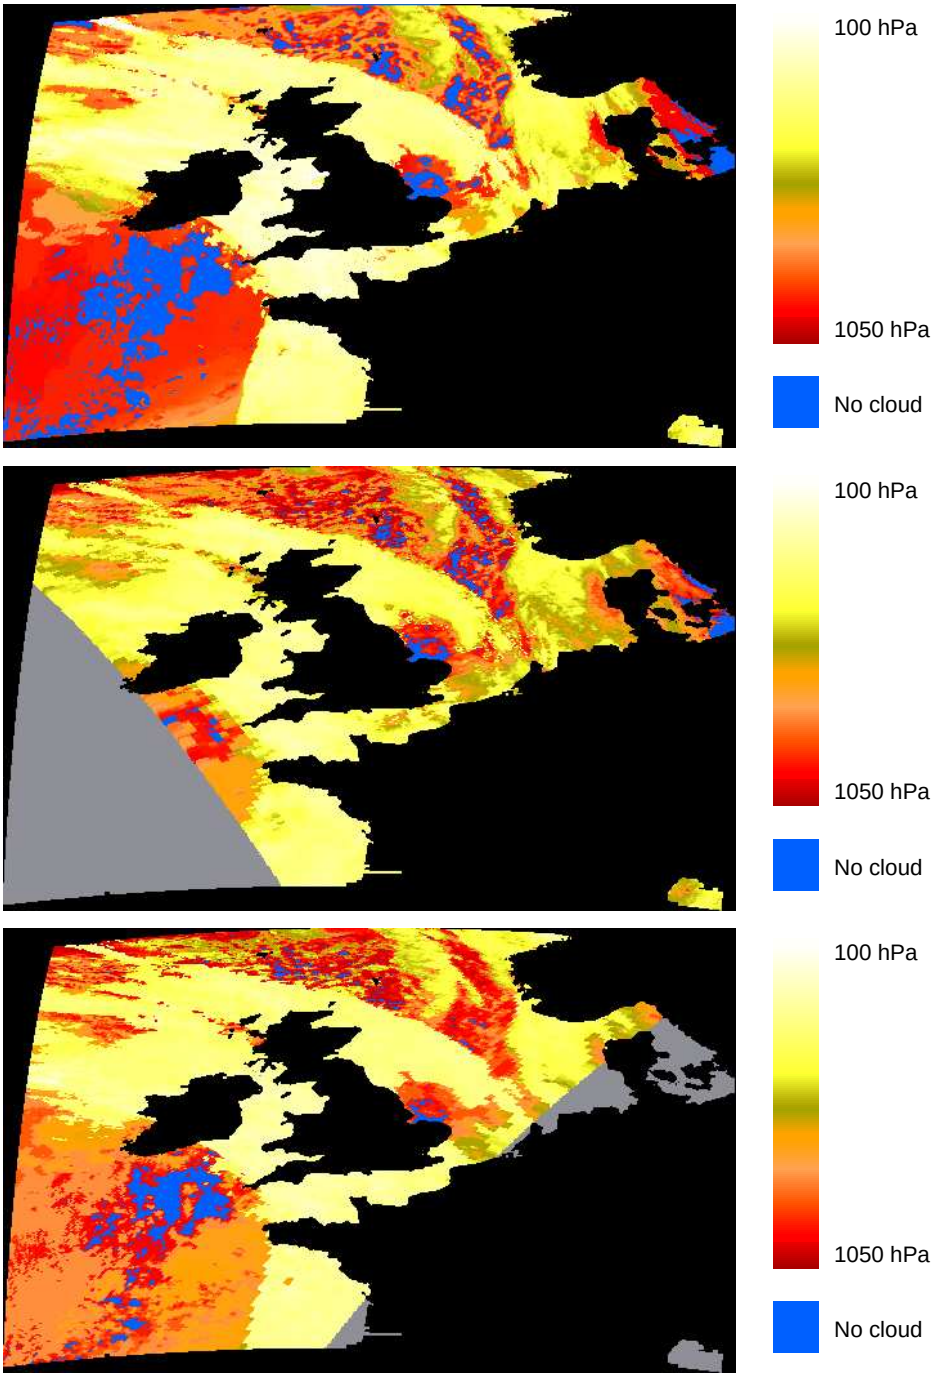

Figure 7: Cloud-top pressures obtained from SEVIRI via AUTOSAT (top), from MODIS AQUA (middle) and from MODIS TERRA (bottom) for the 2006 case. The AQUA data is a composite of two images, recorded at 1205Z and 1210Z, while the TERRA data is a composite of images recorded at 1210Z and 1215Z. See the caption to Figure 5 for further explanation.

profile matching), which assume an ECA of 1, tend to be used more frequently for lower, warmer cloud. However, even for bodies of solid cloud where both AUTOSAT and 1D-Var give ECAs equal or close to 1, such as that over the Bay of Biscay in BAS2004, 1D-Var places the cloud higher in the atmosphere. Less frequently, 1D-Var cloud is consistently lower, such over the North Sea off the Dutch and German coasts.

For BAS2006, low cloud is also higher in 1D-Var than in AUTOSAT retrievals. To a lesser extent the same is true for high cloud. The AUTOSAT cloud is rather more uniform than the 1D-Var cloud; some of the CTP variability shown by the 1D-Var cloud seems unlikely to be physical. Adjacent pixels in the AUTOSAT fields are much more likely to use the same retrieval method in the 2006 case than in the 2004 case, probably due to revisions to the AUTOSAT processing.

The false-colour images in Figures 2 and 3 indicate the presence of very thin cloud in both the cases studied. In the 2004 case there are examples of this over the North Sea off the northeast coast of England, and in a "spit" of cloud running between the northwestern coast of Ireland and the west of Scotland. In the 2006 case such cloud can be seen over the Irish Sea and the Channel. 1D-Var identifies this cloud as high and partial (with ECAs below $\sim$50%) in both cases. In the 2004 case the area of partial high cloud in the 1D-Var retrieval extends eastwards to the coast of Denmark, but neither the false-colour nor the infra-red image (also shown in Figure 2) indicate clearly whether this would be more correctly identified as lower full cloud, or whether there is multilayer cloud. The AUTOSAT retrieval identifies most of this area of cloud as lower full cloud, using the stable layers and profile matching methods. Where the minimum residual method is used, towards the Danish coast, some higher partial cloud is retrieved.

Where there is thin, high cloud in the 2006 case, the false-colour and infra-red images in Figure 3 both give a stronger indication of areas of multilayer cloud. AUTOSAT, which consistently uses the minimum residual method to retrieve thin high cloud, is in better agreement with 1D-Var here than in the 2004 case. 1D-Var shows slightly more variability in both the CTP and ECA fields. There is some indication that this variability is correlated with the presence of underlying cloud.

Two-dimensional histograms comparing CTPs retrieved using 1D-Var to those calculated by AUTOSAT for both the 2004 and 2006 cases (BAS2004 and BAS2006 respectively) are shown in the left-hand column of Figure 8. They confirm that the comparison between AUTOSAT and 1D-Var is much cleaner for BAS2006 than for BAS2004, reflecting the improvements to the AUTOSAT scheme, but both cases demonstrate the tendency of 1D-Var to raise the height of cloud compared to AUTOSAT, in some cases (where AUTOSAT cloud is low) very significantly. Some cloud is lowered by 1D-Var relative to AUTOSAT in the 2006 case, but rarely by more than $\sim$50 hPa.

The top row of Figure 9 illustrates the decomposition of the 2D-histogram by AUTOSAT processing method for the 2006 case. It confirms that AUTOSAT CTPs are clustered according to the method used: where the minimum residual method is successful, pixels are given a relatively low CTP, while the other two methods produce lower cloud. These methods show a lower degree of correlation with the 1D-Var results than the minimum residual, probably, as discussed above, because they only allow ECAs of 1.

Given that the minimum-residual method is effectively a "cut-down" 1D-Var, the relatively good agreement between these two methods is not surprising. There is a systematic difference between results from the two methods across the CTP range where the minimum-residual method is used, with 1D-Var CTPs being lower than the minimum-residual CTPs for the highest clouds, but higher than minimum-residual CTPs for lower clouds. 1D-Var differs from the minimum-residual processing in two ways: temperature and humidity profiles are allowed to vary, and

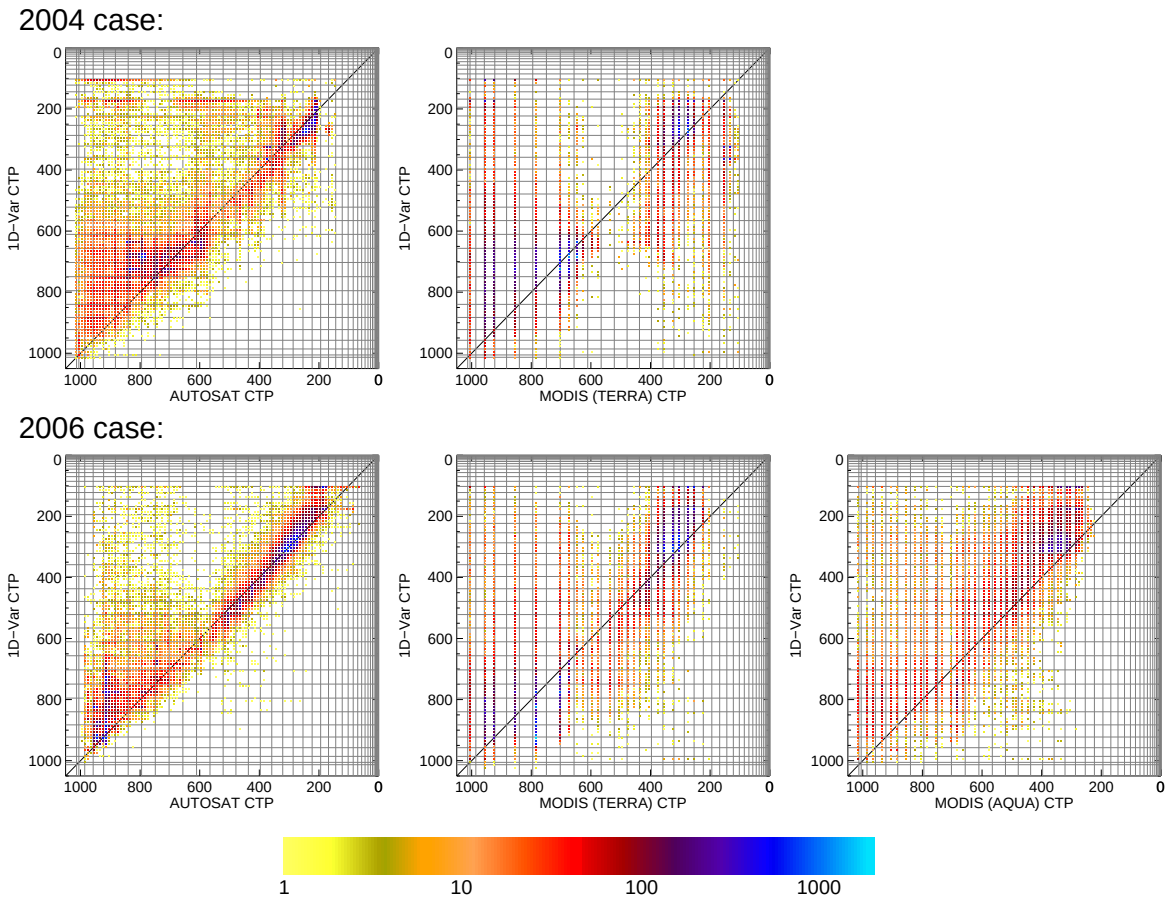

Figure 8: Two-dimensional histograms comparing CTPs from 1D-Var, AUTOSAT and MODIS for both the 2004 (top row) and 2006 (bottom row) cases. Each count represents a comparison between two satellite pixels. The bin size on each CTP axis (for this and all subsequent 2-D histograms with CTP on one or both axes) is 10 hPa. MODIS CTPs are assigned only to particular discrete values, hence the vertical "stripes". Grey lines indicate RTTOV levels.

additional channels are used. It is not clear if the systematic AUTOSAT/1D-Var differences can be attributed to one or other of these methods alone.

Horizontal streaks or stripes in histograms in Figure 8 are discussed in Section 3.4 below.

3.2.2 Comparison to MODIS products

In the 2004 case, the CTP field retrieved from TERRA MODIS data in general looks more similar to the AUTOSAT product than to the 1D-Var retrievals, although in the area of broken cloud in the western half of the image it produces even fewer cloud-free pixels than 1D-Var, replacing them with very low cloud. The MODIS CTP field has more pixel-to-pixel variability in areas of solid low- to mid-level cloud than the 1D-Var CTP field, but less than the AUTOSAT CTP. Areas of solid high cloud towards the western edge of the TERRA swath tend to be higher than in either the AUTOSAT or the 1D-Var retrievals. The 2006 case, like the 2004 case, shows that the MODIS scheme (with both TERRA and AQUA data) tends to diagnose fewer cloud-free points than AUTOSAT, though not as few as 1D-Var, which, as mentioned above, gives very few cloud-free points for this case. The fields derived from TERRA MODIS show less pixel-to-pixel CTP

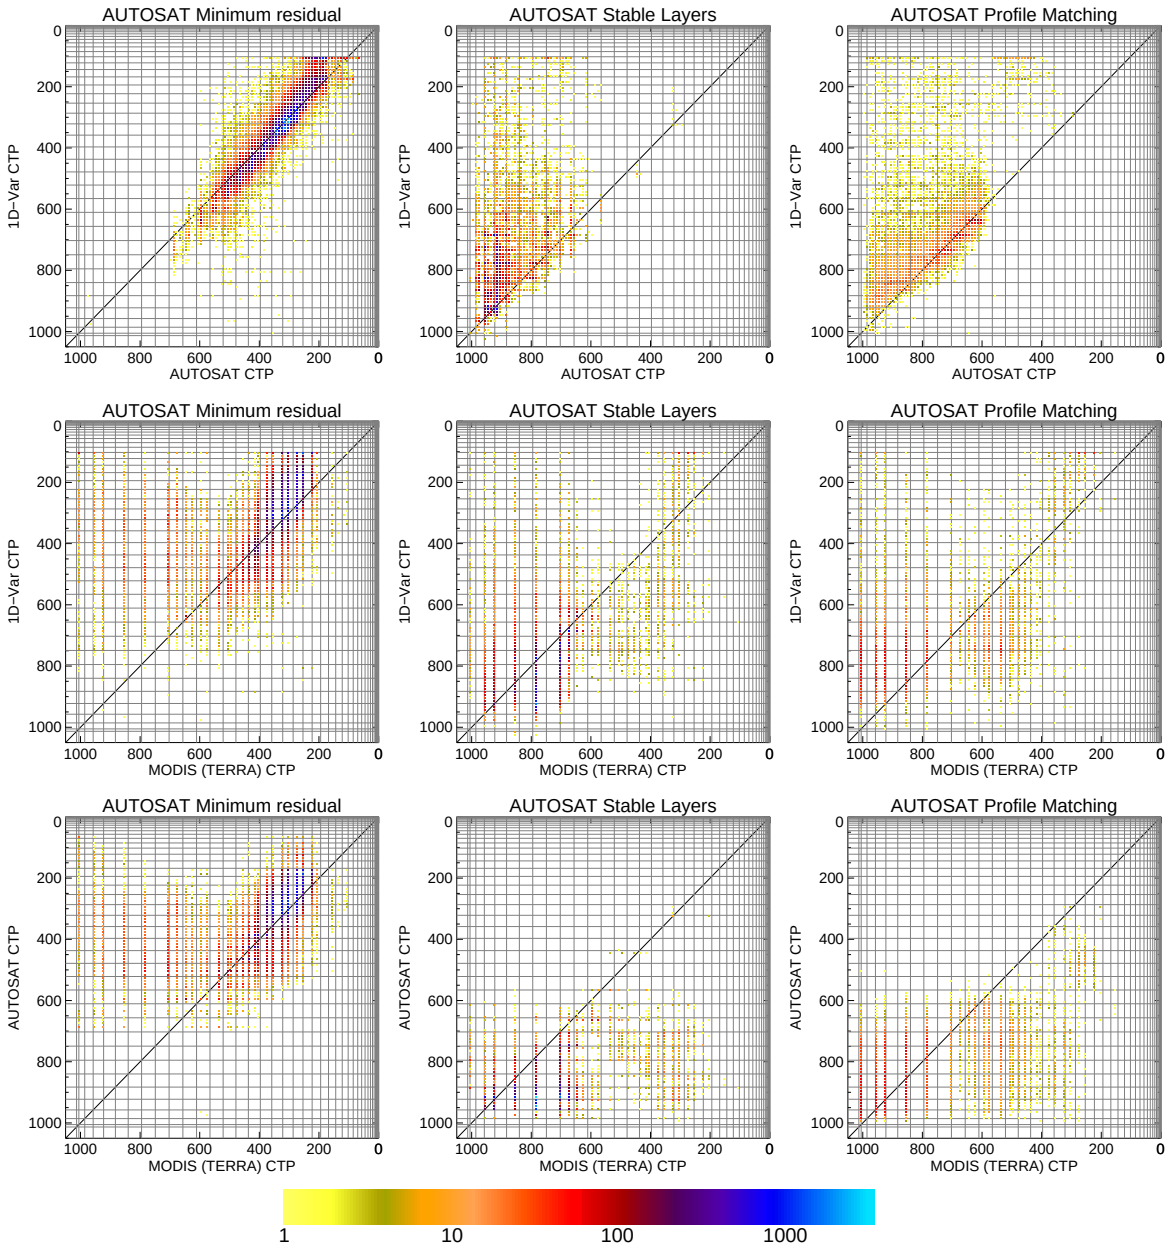

Figure 9: Two-dimensional histograms comparing CTPs from 1D-Var, AUTOSAT and MODIS for the 2006 case, distinguishing according to AUTOSAT method. (Grey lines indicate RTTOV levels.)

variability and more ECA variability than the AQUA MODIS fields, and in general the TERRA MODIS fields appear more similar to both AUTOSAT and 1D-Var than the AQUA MODIS fields. Higher cloud is slightly higher in both AUTOSAT and 1D-Var than in either the AQUA or TERRA MODIS products, although the TERRA MODIS field has more high cloud than either 1D-Var or AUTOSAT in the region of broken cloud along the northern edge of the image. The area of mid-level cloud in the southwestern portion of the TERRA MODIS field is more consistent with 1D-Var than with AUTOSAT, which returns significantly lower cloud, although this is an area where minimisation difficulties are common in 1D-Var (see Section 3.3 below).

MODIS's characterisation of what the false-colour and infra-red images suggest is thin, high cloud has more in common with AUTOSAT than with 1D-Var, in both the 2004 and 2006 cases.

In the 2004 cases, the area off the eastern coast of England shows a mixture of low full cloud and higher partial cloud in the MODIS retrieval, but discontinuities in the CTP and ECA fields suggest that MODIS CTP is strongly dependent on retrieval method, which in turn is frequently different for adjacent fields of view. In the 2006 case, in common with AUTOSAT, the thin cloud over the Irish sea has CTPs more consistent with 1D-Var.

Point-by-point comparisons of MODIS and 1D-Var CTPs are shown in the histograms in the second and third columns of Figure 8. The 1D-Var CTPs compare less cleanly to MODIS than to AUTOSAT, particularly for the 2004 case where, as noted above, a time mismatch contributes to scatter in the histogram. For both BAS2004 and BAS2006 (and for both MODIS sensors in the 2006 case) 1D-Var CTPs for high cloud are higher than the MODIS CTPs. The signal for low cloud is less clear, with the TERRA MODIS sensor indicating that MODIS low cloud is lower than 1D-Var cloud in the 2004 case but higher in the 2006 case (although there is a significant "tail" of 1D-Var CTPs to low values in both cases). For the 2006 case, a lot of the lower cloud is excluded from the AQUA MODIS analysis because of the position of the swath. The plots confirm that AQUA MODIS compares less well to 1D-Var than does TERRA MODIS.

The second and third rows of Figure 9 show comparisons between 1D-Var and TERRA MODIS CTPs and between AUTOSAT and TERRA MODIS CTPs respectively for the 2006 case, decomposed by AUTOSAT processing method as for the 1D-Var/AUTOSAT comparison in the top row. The two MODIS processing methods also partition CTPs, with almost all retrieved CTPs less than 650 hPa being produced by CO_2-slicing, and almost all retrieved CTPs greater than 650 hPa being derived using the EBBT method. While use of minimum residual processing by AUTOSAT usually coincides with MODIS's use of CO_2-slicing, the correspondence is not perfect, so some of the "minimum residual" pixels have MODIS's EBBT processing applied. Similarly, there are significant numbers of pixels for which AUTOSAT uses the stable layers or profile matching methods but MODIS uses CO_2-slicing. These points where methods do not coincide generally have much greater disagreement between CTPs in the AUTOSAT/MODIS comparison, suggesting that choice of retrieval method may be as important a source of uncertainty in CTP as the method itself.

For points where AUTOSAT uses minimum residual processing, both 1D-Var and AUTOSAT clouds are high compared to MODIS clouds, with the difference being slightly less for AUTOSAT. For the "stable-layers" points, 1D-Var CTPs show better agreement with MODIS CTPs than the AUTOSAT CTPs do, with AUTOSAT's cloud being on average lower than MODIS or 1D-Var. 1D-Var does place some cloud significantly higher than MODIS, but where MODIS uses CO_2-slicing, there is some agreement for this higher cloud also. For pixels where AUTOSAT applies profile matching and MODIS the EBBT method, agreement between MODIS and AUTOSAT is better than between MODIS and 1D-Var, with 1D-Var cloud being too high; where AUTOSAT applies profile matching but MODIS uses CO_2-slicing, most 1D-Var cloud is rather lower than the MODIS cloud.

3.2.3 Comparison with Norrköping cloud intercomparison workshop data

Histograms of cloud-top pressure for retrieval schemes submitted to the Norrköping intercomparison workshop are shown in Figure 10. Note that most schemes processed a large fraction of the full-earth disc and performed retrieval over land as well as sea; here only data for the pixels processed in the 1D-Var retrieval are considered. Cloud-free pixels were also considered in constructing the histograms, each dataset being masked using its own cloud mask.

All schemes suggest a double-peaked distribution of CTPs. The 1D-Var scheme shows good

agreement with most other schemes over the higher-altitude peak. There are two features which make it significantly different to other schemes, for this case at least. Firstly, its low-cloud peak is somewhat broader than those of most other schemes, with a shoulder towards higher cloud; this feature was noted in the comparison of cloud spatial distributions between 1D-Var and AUTOSAT retrievals in Section 3.2.1. Secondly, 1D-Var gives fewer "cloud-free" pixels than any other scheme, again already noted in the comparison to AUTOSAT results.

3.3 Some retrieval characteristics (control configuration)

In this section we examine some diagnostics of the behaviour of the minimisation scheme: number of iterations required for convergence, cost function at convergence, analysis error covariances and the brightness temperatures modelled by RTTOV for the retrieved profiles.

A histogram showing convergence rates for both BAS2004 and BAS2006 is given in Figure 11. The mean number of iterations required for convergence in BAS2004 is 5.2, which is smaller than that for BAS2006, 5.7. The latter peaks at a lower value but has a larger tail towards higher iteration numbers. BAS2004 also shows a smaller number of retrievals failing to converge within 12 iterations: 2.8%, as opposed to 5.2% for BAS2006. The spatial distribution of convergence failures for the two cases is shown in Figure 12. From the false-colour and infra-red images (Figures 2 and 3), the convergence failures appear to be associated preferentially with mid-level full cloud in both the 2004 and 2006 cases, although a few occur in regions of very high cloud. The clustering of convergence failures, particularly in BAS2006, suggests that there are conditions under which minimisation is particularly difficult; these could be physical (such as a structure in the background temperature profile that produces a complex cost-function surface) or associated with the design of the retrieval scheme (for example, the multiple-FOV configuration finding it

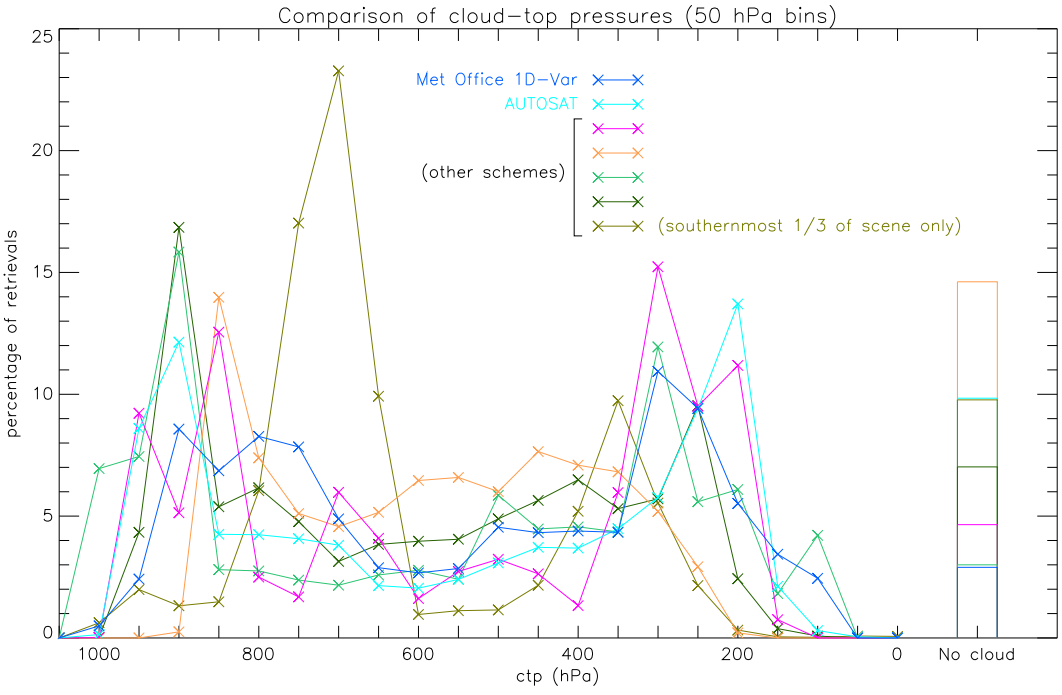

Figure 10: Histograms of CTP for a subset of retrieval schemes contributing to the Norrköping inter-comparison study. A relatively broad bin-size of 50hPa is used.

difficult to reconcile widely varying brightness temperatures using a common CTP).

Histograms of the values of the cost function $J(\mathbf{x})$ at the end of the retrieval are shown for both the 2004 and 2006 cases in Figure 13. The plot also shows the effect of excluding retrievals which have failed to converge within the limit of 12 iterations. The mean value of J for BAS2004 (5.22 for all retrievals and 5.08 for retrievals that converge within 12 iterations) is significantly greater than that for BAS2006 (3.47 for all retrievals and 3.38 for retrievals that converge within 12 iterations). There is not an obvious correlation between the number of iterations required for convergence and the final value of J - to demonstrate this, it would be necessary to take into account the likely dependence of the expected final value of J on the number of FOVS in the retrieval. Neither convergence failures nor high-J retrievals have been excluded from subsequent analyses in this report, though in an operational context they could be disregarded and a fall-back method used to obtain cloud properties.

Two of the diagonal elements of the analysis error matrix discussed in Section 2.4, those for cloud-top pressure and effective cloud amount, are shown as scatterplots in Figure 14. Given the non-linear nature of the retrieval, these should be interpreted with caution, but some trends are apparent. Errors in both quantities are generally lower for clouds retrieved higher in the atmosphere, indicating that, like other methods, 1D-Var is more skilful for high cloud, where thermal contrast between the surface and cloud top is greater than it is for low cloud. For the highest cloud (above approximately 300 hPa) there is a suggestion that error in CTP begins to increase with height for pressures above approximately 300 hPa, and a significant number of retrievals have large CTP analysis errors. The latter tend to have retrieved CTP either close to one of the levels of the RTTOV model (also plotted) or at the lower limit of pressure allowed by the scheme, 100 hPa. Clustering on RTTOV levels is not observed in the 2006 case; rather there is a suggestion (more clearly visible in the 1D-histogram of cloud-top pressures plotted in Figure 17) that retrieval on RTTOV levels is inhibited. These situations will be discussed in Sections 3.4.2 and 3.4.3 respectively. ECA errors for these retrievals do not appear to depart from the general pattern, which is for effective cloud amount error to decrease with increasing ECA, at least for

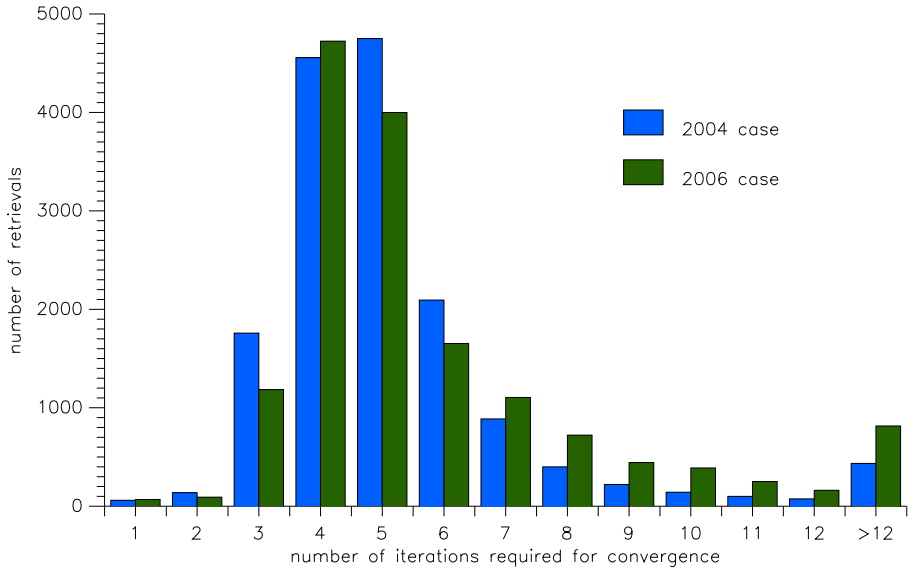

Figure 11: Histogram of the number of iterations required for the retrieval to converge, for both 2004 and 2006 cases.

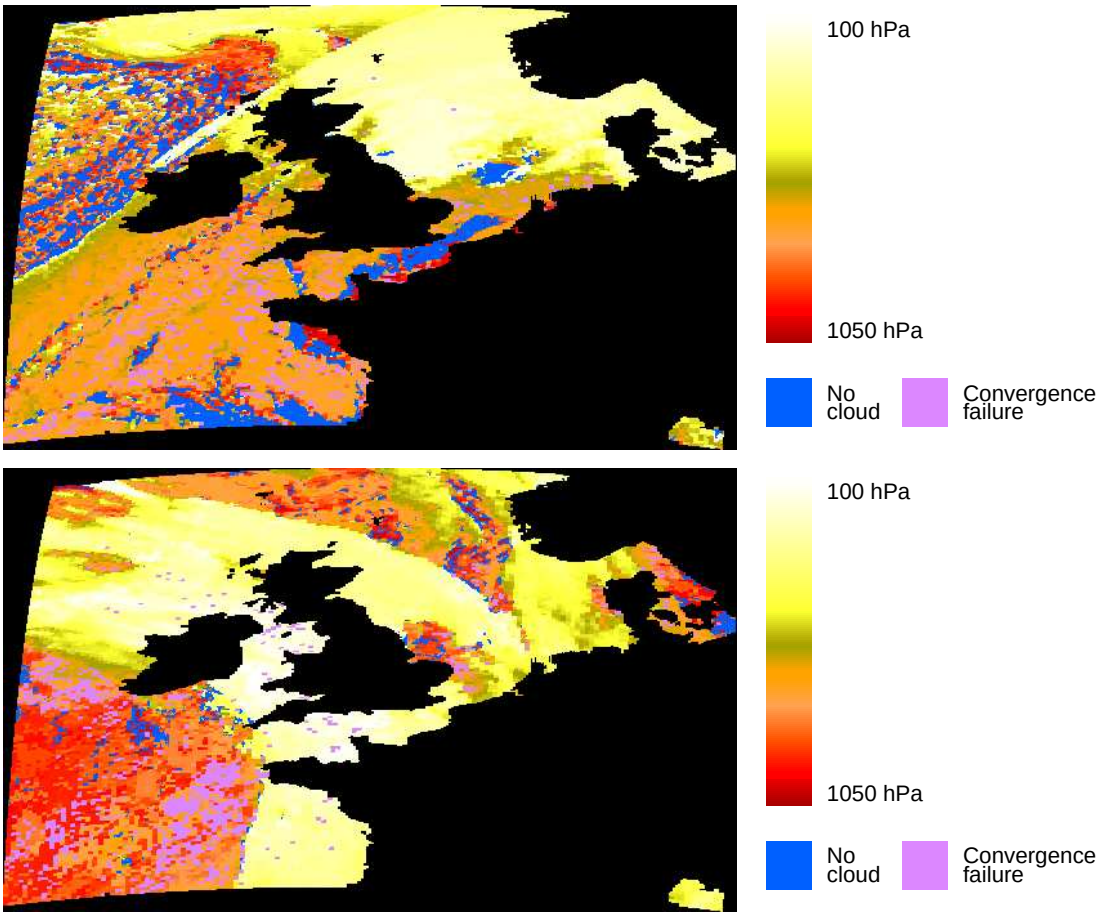

Figure 12: Spatial distribution of retrievals for which convergence fails, for the 2004 (top) and 2006 (bottom) cases.

higher cloud. (For lower cloud this signal is lost, or possibly reversed.)

The use of final cost function value to shade points in the plots of CTP analysis error in Figure 14 illustrates that retrievals with smaller values for the theoretical analysis error are not necessarily those for which the final value of the cost function is small, that is, those for which a good fit of profile and brightness temperatures is achieved. A tendency for retrievals of CTP for low- to mid-level cloud to cluster around the midpoints between RTTOV levels, and for these retrievals to have relatively high values of J, is evident in both the 2004 and 2006 cases and will be addressed in Section 3.4.1.

Histograms of the brightness temperature fits for each channel in are shown for BAS2004 in Figure 15 and for BAS2006 in Figure 16. Two sets of retrievals are considered in each case: the full scene, and "clear-sky" retrievals, selected conservatively to consist only of pixels which the AUTOSAT cloud mask tests judge to be clear and which have (1D-Var) ECA below 0.1. The plots also show the forward-modelled brightness temperatures obtained from the minimum-residual first-guess profiles, for the same set of pixels. The plots of summary statistics show the mean (retrieved−observed) temperature difference for each channel, with bars indicating one standard deviation on either side.

The summary statistics plots confirm that 1D-Var retrieval gives better fits to the observed brightness temperatures than the first guess, in most cases in terms of both the mean and the width of the distribution. For the window channels the improvement over the first guess is small,

and the standard deviation of the 1D-Var fits for "all-sky" cases - of order 1K to 1.5K - is large compared to the forward-modelling and measurement errors, given in Table 1. (By contrast, "clear sky" 1D-Var fits for these channels produce much narrower histograms, even though skin temperature is only loosely constrained by the background error covariance matrix $\mathbf{B}$ — the value of the diagonal element representing skin temperature variance in the $\mathbf{B}$-matrix is $(1.64\text{K})^2$.) These features suggest either deficiencies in the minimisation where cloud is present, and/or that the cloud model being used is over-simplistic. A mechanism by which the minimisation might be producing tails to high positive brightness temperature differences, visible in the histograms for the window channels, is discussed in Section 3.4.1 below.

All-sky fits for the water vapour channels, on the other hand, show 1D-Var giving a significant improvement on the first guess; this is not surprising, as these channels are not used in the minimum-residual fits. The spread of retrieved brightness temperature fits is smaller than the forward-modelling and measurement errors suggest. Except at high levels, these channels are less sensitive to cloud than the window channels, and so might be expected to be less affected by the difficulties associated with cloud suggested above (although clear-sky fits for these channels are still better than the all-sky fits). It may even be that fitting these channels well is too "cheap" for the minimisation scheme, indicating that the water vapour elements in $\mathbf{B}$ are too large.

3.4 Problems with the retrieval scheme

The qualitative comparison of CTP and ECA fields with those obtained from other methods and sensors, together with examination of retrieval diagnostics, suggests that the 1D-Var scheme gives plausible results in many cases. However, the following features were noted in Section 3.3 above, and are confirmed by the histograms of retrieved cloud-top pressure shown in Figure 17:

- for mid-level clouds (CTPs above ∼600 hPa), in both cases, relatively broad peaks in the distributions between the levels used in the RTTOV radiative transfer model;

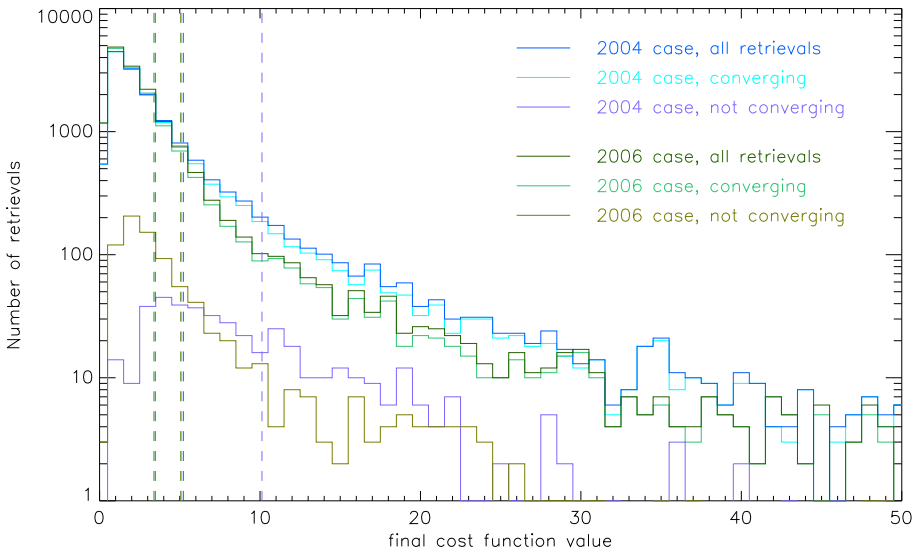

Figure 13: Histograms of final J-values, including and excluding retrievals which do not converge within 12 iterations. The vertical lines indicate the mean J in each case.

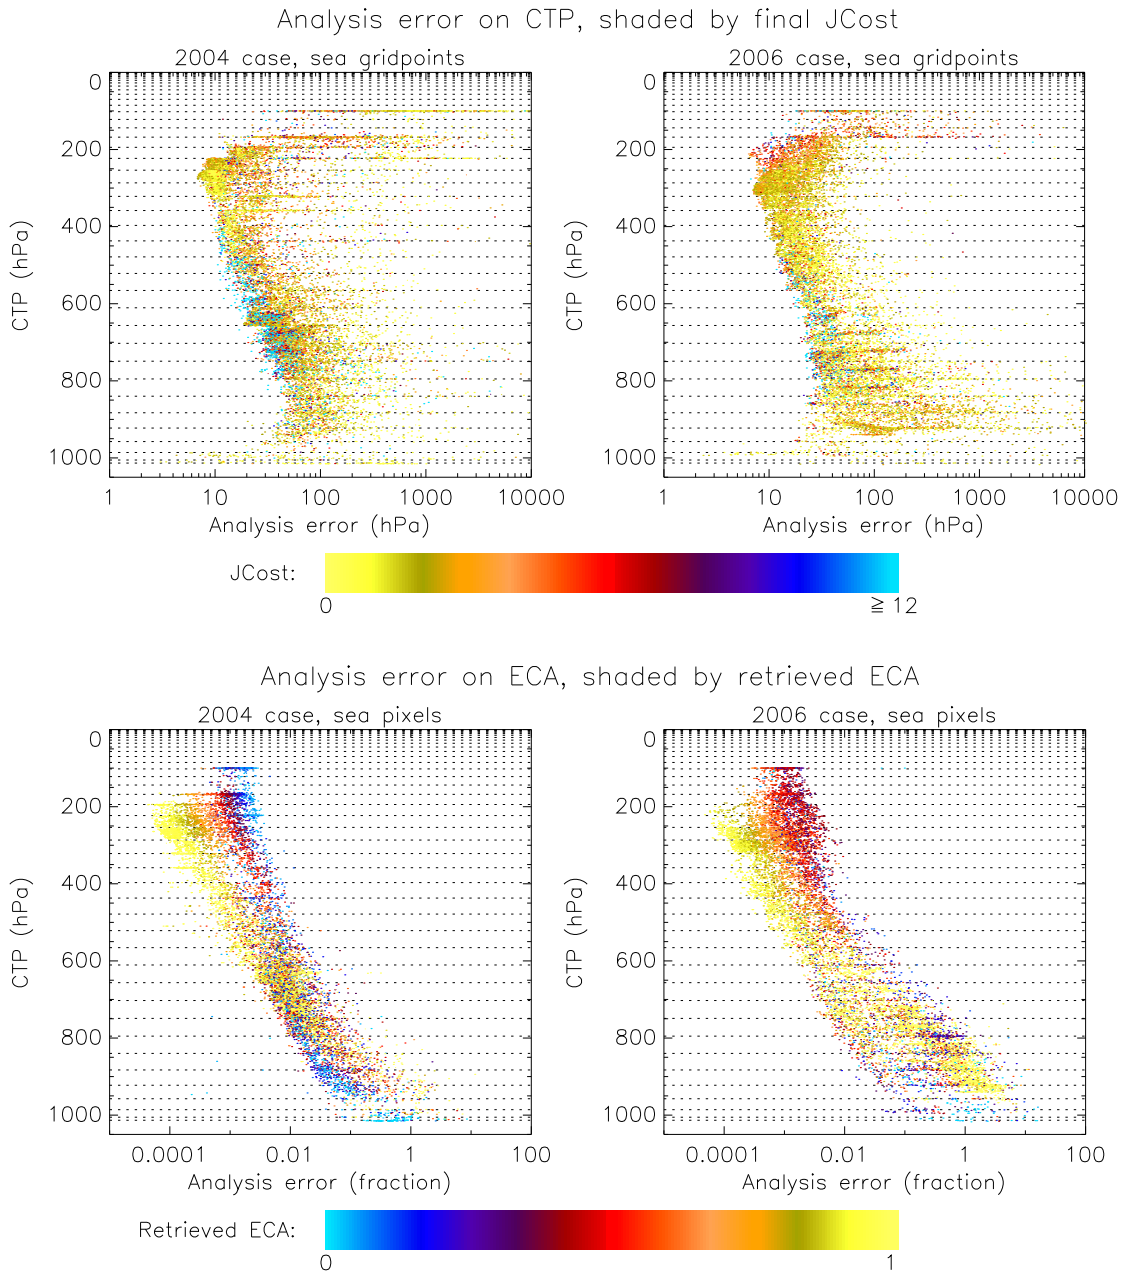

Figure 14: Scatterplots of analysis error on retrieved cloud-top pressure (top) and retrieved effective cloud amount (bottom), for both the 2004 and 2006 cases. The vertical axis is retrieved cloud-top pressure in all plots, and horizontal dotted lines indicate the 43 levels of the radiative transfer model RTTOV 6.9. CTP errors are shaded according to the value of the cost function at the end of minimisation, whilst ECA errors are shaded using the value of ECA. Note that a CTP analysis error has been plotted for each model gridbox, whilst ECA errors have been plotted for every eighth pixel in the pixel grid. Pixel thinning is done only for clarity.

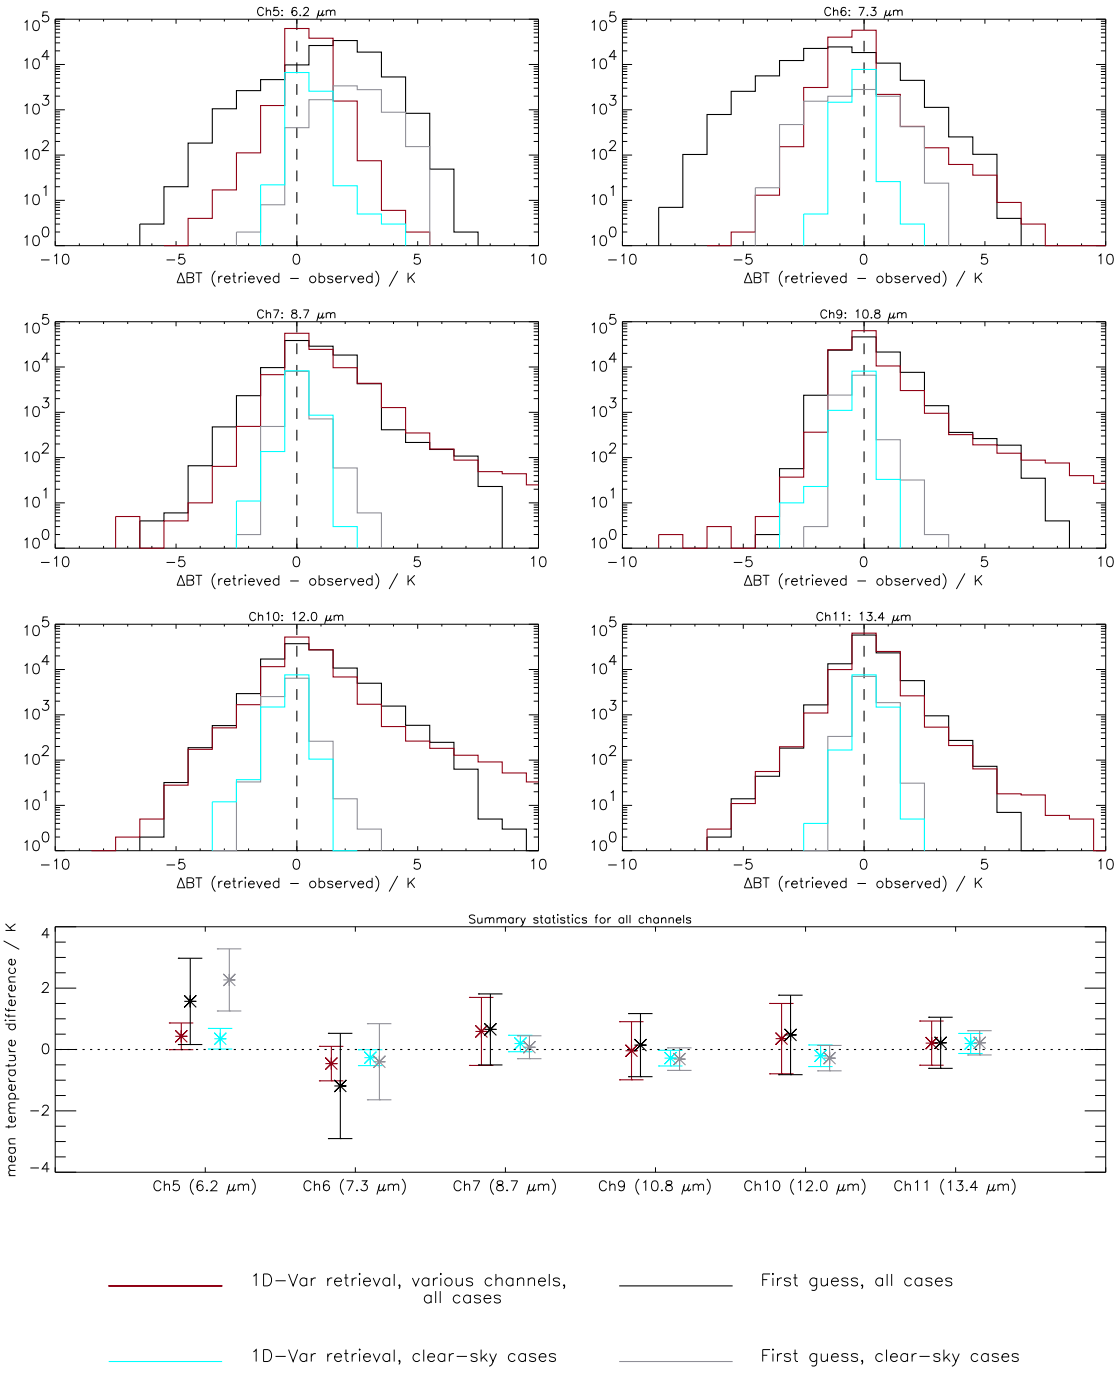

Figure 15: Histograms showing the fit of retrieved to observed brightness temperatures for each channel in the retrieval, for BAS2004. "Clear sky" retrievals (see text for definition) are isolated and overplotted. "Clear sky" 1D-Var fits are effectively independent retrievals of surface (skin) temperature in the window channels and of upper-tropospheric humidity in the water-vapour channels. The "summary statistics" plot shows the mean retrieved-observed temperature differences, with the bars representing one standard deviation of the distribution on either side.

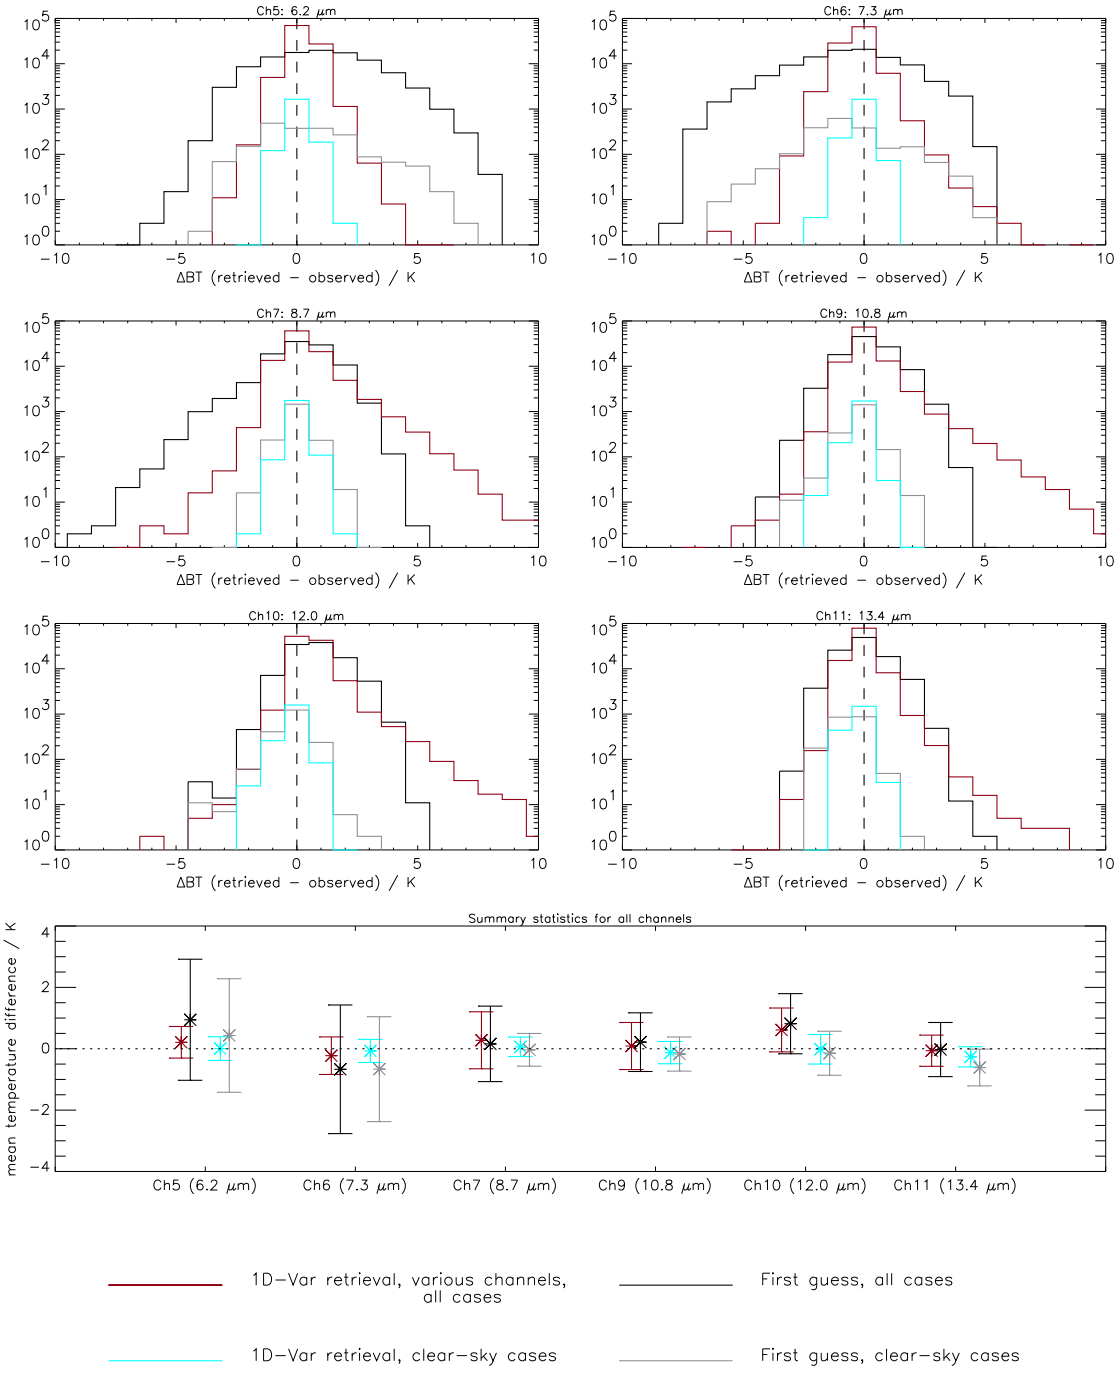

Figure 16: As for Figure 15, but using data from the 2006 case.

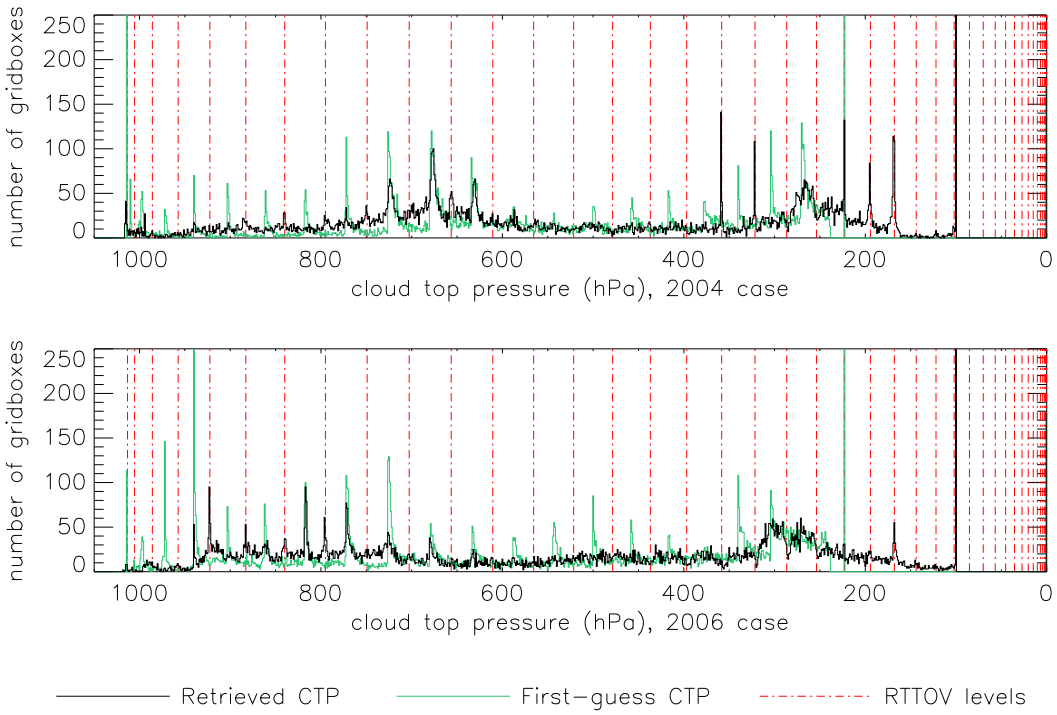

Figure 17: Histograms of retrieved and first-guess cloud-top pressure, for both the 2004 and 2006 cases. The vertical scale does not cover the full heights of the peaks at 100 hPa, which contain 476 points in the 2004 case and 265 points in the 2006 case. (15,784 retrievals are performed in total.)

- for higher clouds (with CTPs below ∼400 hPa), clustering on or very close to RTTOV levels, in the 2004 case, and a possibility that in the 2006 case retrievals with CTPs close to RTTOV levels are being prevented;
- in both cases, a significant amount of cloud placed at 100 hPa, the highest altitude allowed by the scheme.

The characteristics and possible causes of these features are discussed, in turn, below. Variants of the 1D-Var scheme that might address the problems were tested where feasible.

3.4.1 Peaks between RTTOV levels

Figure 14 suggests that retrievals with CTP clustering between RTTOV levels tend to be associated with high values of the 1D-Var cost function J, indicating poor fitting of brightness temperatures and/or profile information. Convergence rates for these retrievals (not shown) are not particularly poor. Since the minimisation scheme tests for convergence using the change in value of the cost function between iterations, rather than its absolute value, the convergence rate suggests that the retrievals may be being accepted by the scheme as well-behaved because profiles are not changing significantly from iteration to iteration. The retrievals also tend to produce full cloud cover. It is suggested that the deficiency of the 1D-Var scheme here lies in its failure to explore solutions that depart sufficiently, in CTP, from the first guess. Whilst clustered CTPs in the first guess[1]

[1]Clustering in the first-guess values arises from the implementation of the minimum-residual technique used for this study. Like 1D-Var, the minimum-residual technique requires the minimum of a cost function to be found.

(evident over a wide range of cloud-top pressures in Figure 17) are unsatisfactory, and improving the first guess is important for reasons of efficiency, a retrieval scheme should not be constrained to its first guess. (In this respect the imperfect first guess provides a useful diagnostic.)

Premature convergence is thought to be associated with the limits on ECA of 0 and 1 imposed in the 1D-Var retrieval code. In the course of minimisation, profile increments that would take ECA above its physical limit of 1 may be proposed by equation 3, and changes proposed in other profile elements will be consistent with this. However if the value of the ECA is then reset to 1 by the 1D-Var code, changes in other profile elements are not adjusted consistently. This accounts for lack of improvement in fitting brightness temperatures, and corresponding high values of J; it also means that once ECA has reached 1 the profile, and hence J, may not change significantly from iteration to iteration. The convergence criteria can thus be satisfied, and minimisation stopped, before an optimum solution has been found.

Two variants of the scheme were tested, using the 2004 data, in an attempt to force the 1D-Var scheme to explore more profiles during minimisation. In the first (ACF2004), the restriction on N was relaxed to limits of -0.25 and 1.25, and an additional cost function term J_{add} (as described in Section 2.6 above) was introduced for profiles with $N > 1$ and $N < 0$. By penalising unphysical cloud, the J_{add} term tends to reduce the proposed increments to ECA and allows greater changes in other profile elements, such as CTP. The functional form of the additional cost function term was that suggested by Szyndel *et al.* (2004), after Phalippou (1996):

$$J_{add} = \begin{cases} -CN^3 & N < 0 \\ 0 & 0 \leq N \leq 1 \\ C(N-1)^3 & N > 1 \end{cases} , \tag{7}$$

where C is a constant, this being the simplest positive definite function continuous up to the second derivative. In a given retrieval, a J_{add} contribution was made for each FOV with $N > 1$ or $N < 0$, using a value of 5000 for C for all FOVs. In this study only one value of C (chosen using typical values of the cost function) was tested. There is scope to tune both the value of C and its dependence on the number of FOVs in the retrieval.

Unphysical cloud amounts were common in the 2004 case: approximately 13.4% of retrievals had a final ECA of greater than 1, and 1.5% a final ECA below 0. Apart from obvious problems of interpretation, unphysical cloud amounts can cause retrievals to fail. If $N > 1$ or $N < 0$ then one of the two terms on the right-hand side of Equation 2 is negative, and depending on the values of $H^{\mathrm{cl}}(\mathbf{x})$ and $H^{\mathrm{op}}(P_{cl}, \mathbf{x})$, may produce an unphysical negative value for the radiance $H(P_{cl}, N, \mathbf{x})$. Where this occurred (in fewer than 0.1% of cases) retrieval was abandoned.

A second approach was to use the variant of the Marquardt-Levenberg method described in Section 2.7. This test (MLV2004) was motivated by the observation that with CTP expressed in hPa, the numerical difference between first-guess and retrieved CTP is typically (in non-problematic retrievals) larger than that for other elements, by a factor of at least ten, whilst the change in the ECA element is often smaller than that in other elements. Substituting an appropriate $\mathbf{D}$ for $\mathbf{I}$ should make it "cheaper" to explore changes in CTP in order to produce better fits to the observed radiances, and so reduce the tendency noted above to propose changes in ECA that would lead to

Here, the cost function is calculated for cloud-top at the set of pressures given by the RTTOV levels, and the CTP at which the cost function is minimised is estimated from the discrete values using quadratic interpolation between three points: the level where cost is a minimum and the two levels immediately above and below. The minimum-residual cost function is often found to have an asymmetric form, varying much more quickly with pressure on one side of the minimum than on the other, which results in the quadratic interpolation locating the minimum close to the mid-point between levels.

unphysical cloud amounts. (It is likely that re-scaling CTPs into different units, say of 100 hPa, would have a similar effect, but in this case many more changes to the code and its inputs would be required.) The initial value for the parameter γ was adjusted by a factor reflecting typical values for the trace of $\mathbf{D}$.

The results of using both variants are illustrated in Figures 18 and 19. Both variant methods are effective in smoothing away the peaks between the RTTOV levels (Figure 19(a)), although new peaks do appear on RTTOV levels. Both methods raise the mean height of low- and mid-level cloud slightly. ACF2004 shows very little effect on higher cloud, whilst MLV2004 affects the spread of the CTP-CTP comparison but does not produce a significant mean difference. MLV2004 retrievals appear slightly more "unstable", in that there is a larger number of cases in which cloud-top is placed at the upper limit of 100 hPa allowed by the minimisation (behaviour which is discussed below) and is in general more likely to produce large alterations (several hundred hPa) in CTP.

The fact that both variant methods result in lower cloud-top pressures, on average, than the control, in the range of concern, together with lower values of J at convergence, suggests that the retrieval scheme is better able to fit the observed radiances using higher cloud, and is consistent with the idea that poor fits are produced by profiles for which simulated brightness temperatures are too high. Statistics of the brightness temperature fits (Figure 20) confirm this, with colder mean fitted brightness temperature in the 8.7-, 10.8- and 12.0μm channels for both variant methods, and the variances of the fits reduced. (It is possible that the 10.8μm brightness temperature is now systematically being fitted too cold.) In histograms of brightness temperature fits (not shown), the asymmetry evident for the 8.7-, 10.8- and 12.0μm channels in Figure 15 is reduced, although a tail to high values is still apparent. Tests in which information from the 8.7μm channel is excluded from retrievals, discussed in Section 3.5 below, suggest that channel choice may be another factor exacerbating the problem.

ACF2004 improves fits more effectively than MLV2004, probably because it allows unphysical cloud amounts. Mean convergence rates were 5.8 iterations for ACF2004 and 6.0 for MLV2004 (as compared to 5.2 for BAS2004). Timing tests were not performed, but it was noted when using the additional cost function that there was a reduction in the number of cases where increments had to be recalculated because γ was increased by the minimisation scheme. The higher iteration count may be partially offset by this reduction.

Neither method was tested in combination with any of the other retrieval configurations discussed in this report.

3.4.2 Peaks/troughs on RTTOV levels

The tendency for CTPs to be "attracted to" RTTOV model levels in the course of minimisation in BAS2004 appears to be associated with the occurrence of cloud close to a local minimum in the vertical temperature profile, and is thought to be due to discontinuous CTP Jacobians $\partial H^{\mathrm{op}}(\mathbf{x})/\partial P_{\mathrm{cl}}$, noted in Section 2.5, changing sign on either side of an RTTOV model level. Considering retrievals for which $\partial H^{\mathrm{op}}(\mathbf{x})/\partial P_{\mathrm{cl}}$ for the initial profile is somewhere negative in the 2004 case (approximately 2.0% of all pixels) 29% of these result in retrieved CTPs within 0.5 hPa of an RTTOV level, as compared to 6.3% of all pixels. For the 2004 case the tropopause is observed to be as low as 400 hPa in places, and the distribution of retrievals that return CTP on model levels (not shown) exhibits spatial clustering, implying some dependence on the background fields.

It can be shown that in the general form of $\nabla_{\mathbf{x}} J$ (obtained from equation 1),

$$\nabla_{\mathbf{x}} J = \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x_b}) - \nabla_{\mathbf{x}} H(\mathbf{x}) \mathbf{R}^{-1}\{\mathbf{y^o} - H(\mathbf{x})\}, \tag{8}$$

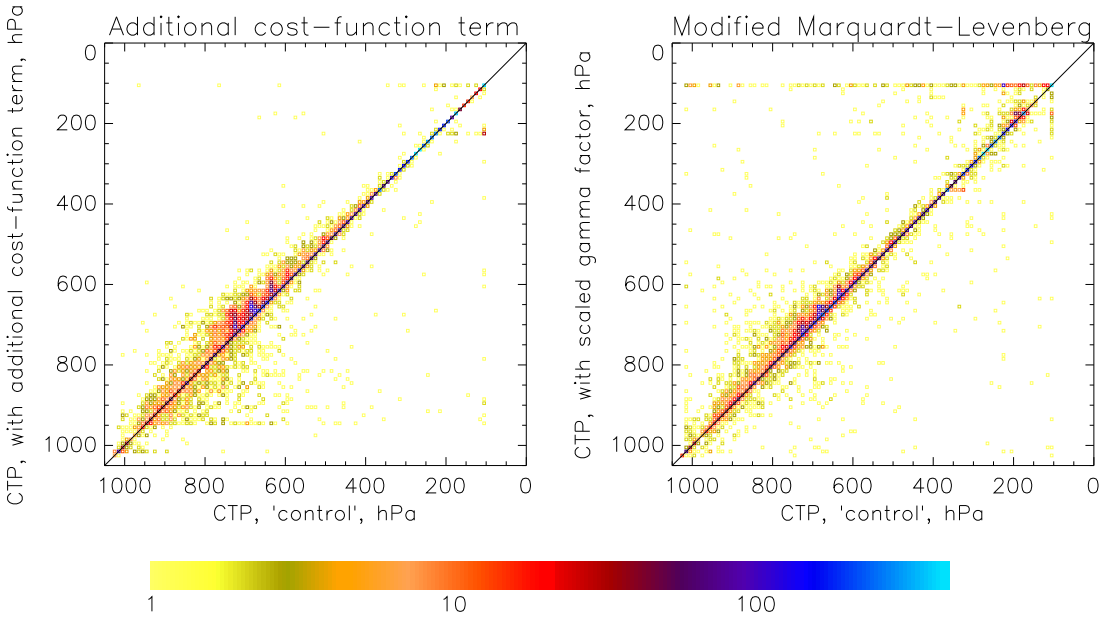

Figure 18: Two-dimensional histograms comparing retrieved cloud-top pressures in each of the variant retrieval methods with those from the control configuration, gridbox-by-gridbox.

a **B**-matrix with the form described in Section 2.2 gives a CTP element $\partial J/\partial P_{\rm cl}$ which can be reduced to the following expression:

$$\frac{\partial J}{\partial P_{\rm cl}} = -\sum_{\rm FOVs}\sum_{\rm channels} \frac{\partial H_{ch}(\mathbf{x})}{\partial P_{\rm cl}} R_{ch}^{-1}(y^o - H_{ch}(\mathbf{x})) \qquad (9)$$

where H_{ch} is a forward-modelled brightness temperature specific to a particular channel and R_{ch} is the appropriate element of the diagonal error matrix **R**. In cases where $\partial H(\mathbf{x})/\partial P_{\rm cl}$ is uniformly positive, the sign of the cost function gradient is anticorrelated with the sign of the term $(y^o - H(\mathbf{x}))$, so that if modelled brightness temperatures are too warm a reduction in CTP will be suggested, and vice versa. However if the sign of $\partial H(\mathbf{x})/\partial P_{\rm cl}$ depends on $P_{\rm cl}$ the situation is less straightforward. If modelled brightness temperatures are consistently too warm $(y^o - H(\mathbf{x}) < 0)$ the sign of the cost function gradient will instead be determined by the sign of $\partial H(\mathbf{x})/\partial P_{\rm cl}$. If $\partial H(\mathbf{x})/\partial P_{\rm cl}$ is positive on the low-pressure side of an RTTOV pressure level and negative on the high-pressure side, the local cost function minimum will coincide with that level. This provides a possible mechanism for preferential convergence on RTTOV levels in the 2004 case, and suggests that, as long as there is not a secondary minimum in the cost-function surface, the error in the retrieved cloud-top pressure may not be greater than half the pressure difference between RTTOV levels. (This argument is not complete, as each element in the increment $(\mathbf{x}_{n+1} - \mathbf{x}_n)$ in equation 3 depends on all elements of $\nabla_{\mathbf{x}} J$ through the matrix $(\nabla_{\mathbf{x}}\nabla_{\mathbf{x}} J(\mathbf{x}_n) + \gamma \mathbf{I})^{-1}$. However this matrix is not usually expected to affect the sign of the increment.) Average final cost-function values for these retrievals are if anything lower than for the scene as a whole, indicating that the fits $(y^o - H(\mathbf{x}))$ are not particularly poor. It may also be that the sign of $(y^o - H(\mathbf{x}))$ is different for different channels and different fields of view.

It is possible that for the 2006 case, where troughs are seen, the term $(y^o - H(\mathbf{x}))$ in equation 9

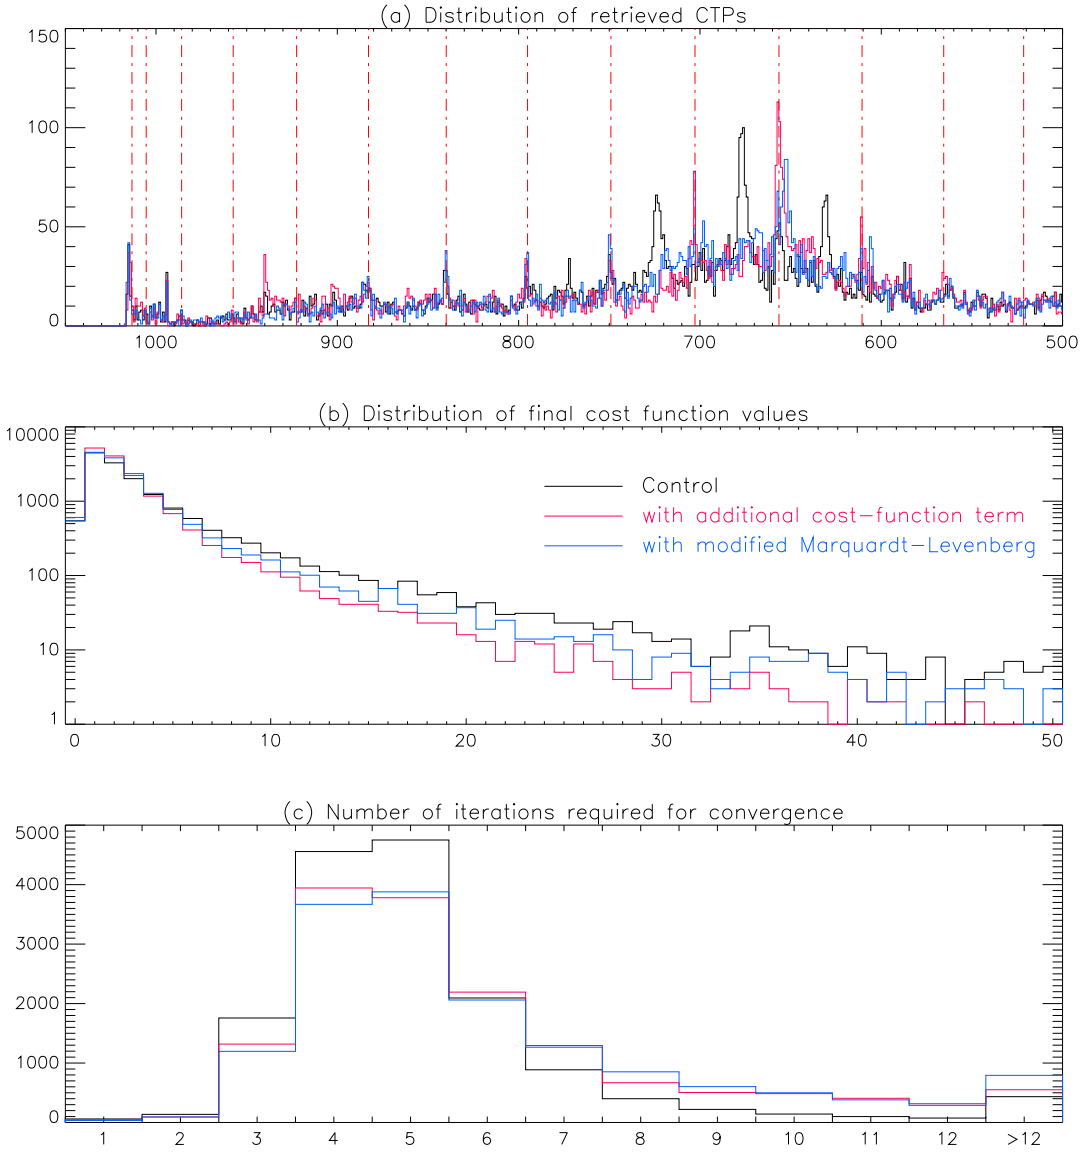

Figure 19: One dimensional histograms of (a) retrieved cloud-top pressure, (b) final cost-function value and (c) number of iterations required for convergence in the control configuration and each of the two variant retrieval methods. In each histogram one count represents one model gridbox.

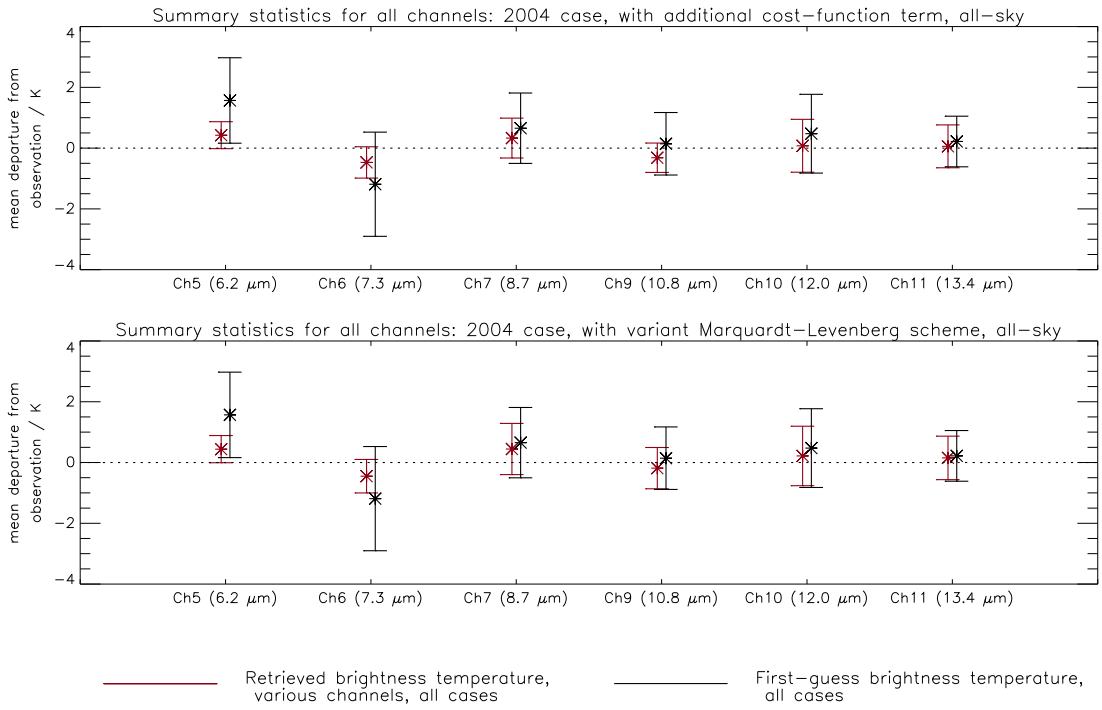

Figure 20: Summary statistics for brightness temperature fits (retrieved − observed) for each channel, for the ACF2004 and MLV2004 tests. The first-guess temperatures are the same as for BAS2004, that is, the same as in Figure 15 and 16.

is consistently negative and so if $\partial H(\mathbf{x})/\partial P_{\rm cl}$ changes sign on either side of an RTTOV level that level becomes a local cost-function maximum. However the affected retrievals are not easily isolated as in the 2004 case.

As already noted in Section 3.3 above, some instances of very large CTP analysis error are seen for the "peaking" retrievals in BAS2004. Analysis error is obtained from the inverse of the matrix $\nabla_{\mathbf{x}}\nabla_{\mathbf{x}}J$, and large values of analysis error reflect cost function surfaces which curve slowly and may thus not have well-defined minima. This is consistent with the observation that these retrievals are more likely to involve a nett increase in the Marquardt-Levenberg parameter γ in the course of minimisation, that is, increments have relatively large components from the steepest-descent method. The affected retrievals are not particularly slow to converge.

Another observation that may be relevant is that if CTP is retrieved on a given level, the temperature retrieved at that level can be much lower than the background value — 2-4K, compared to a typical background error of less than 1.5K. Such temperature changes (spread to other levels by the off-diagonal elements of $\mathbf{B}$) could tend to deepen any minimum in the CTP-dependence of J.

Neither of the variant minimisation schemes tested in ACF2004 and MLV2004 eliminate the problem, although there are some significant effects, such as the removal of the peak at the 222.940 hPa RTTOV level in MLV2004 (not shown). Figure 19 indicates that the variant schemes cause (relatively small) peaks on RTTOV levels below 600 hPa to appear or grow. These cases have not been investigated.

3.4.3 Peak at 100 hPa

It is not clear why a significant number of retrievals (such as those highlighted in Figures 4 and 6) place cloud-top at the upper limit allowed by the scheme, at 100 hPa. The average effective cloud amount for these retrievals in the 2004 case is 11%, and they have a relatively low average final J value (2.84), indicating that poor fitting of brightness temperatures (specifically, being unable to cool scenes sufficiently) is not a problem. The false-colour and infra-red images (Figures 2 and 3) and retrieved CTP in surrounding pixels suggest that most of this cloud is probably high (and thin) anyway, particularly in the 2006 case. In BAS2004 a substantial fraction of the affected retrievals also occur on the (northern) edges of broken cloud, as in the northwestern portion of the region. However it seems unlikely that rapid variation in radiance across FOVS in a retrieval is a contributory factor, as a similar distribution is seen when an independent retrieval is performed for each FOV (see Section 3.5.1 below). Large analysis errors are observed (Figure 14), and in the 2004 data at least there seems to be an element of instability in the retrievals, in that varying some aspect of the retrieval can cause retrieved cloud-top pressure to change from any height to 100 hPa, or vice versa. This is apparent for example when comparing BAS2004 and MLV2004 (Figure 18) as well as in later comparisons in this report (such as Figure 25). It may be that the cost function does not include a sufficiently restrictive tropopause, so that once the scheme has inserted cloud above the tropopause it tends to rise to the upper limit, rather than fall.

3.5 Sensitivity to configuration

3.5.1 Single vs. multiple fields of view formulations

The cloud parameters obtained when an independent retrieval is performed for every pixel (SFV2004) are shown in Figure 21. Comparing these to the cloud parameters retrieved using the multiple-FOV configuration (BAS2004, Figure 4), the distributions can be seen to be qualitatively very similar (more similar to one another than either are to the AUTOSAT or MODIS products considered earlier). A greater amount of detail is apparent in the map of retrieved single-FOV CTP, as might be expected, as the resolution of the CTP field has effectively been increased. In particular, cloud edges are more sharply defined, in terms of CTP, than in the multiple-FOV retrievals, showing an even greater contrast to the other (AUTOSAT and MODIS) retrievals. A significant number of pixels have very high cloud in SFV2004, associated especially with broken and isolated cloud. As in the multiple-FOV retrievals, these appear to occur preferentially on the north/northwest edges of cloud. The percentage of pixels identified as being clear (having an ECA of less than 0.1) is slightly higher in SFV2004 (13.2%) than in BAS2004 (11.4%).

The mean convergence rate for SFV2004, 4.5 iterations, is faster than the rate for BAS2004 (5.2); this is consistent with results reported by Szyndel *et al.* (2004). The failure rate for single-FOV retrievals (that is, the percentage of retrievals that have not met the convergence criteria after 12 iterations) is low, 0.6%. Comparing Figure 21 (in which convergence failures are indicated) to Figure 12 for the 2004 case shows that the distribution of convergence failures differs between the single- and multiple-FOV configurations. A large fraction of convergence failures in BAS2004 (and BAS2006) occur in regions of mid-level full cloud coverage to the southwest and southeast of the UK, whilst there is a cluster of single-FOV failures in the region of higher, partial cloud in the North Sea, where the false-colour and infra-red images (Figure 2) suggest the presence of thin (semi-transparent) cloud. It is possible that multiple-FOV retrievals fail more frequently in the south of the region because the 1D-Var scheme has difficulty reconciling radiances from a larger number of FOVs with a set of common background parameters: as noted in Section 2.2, due to the

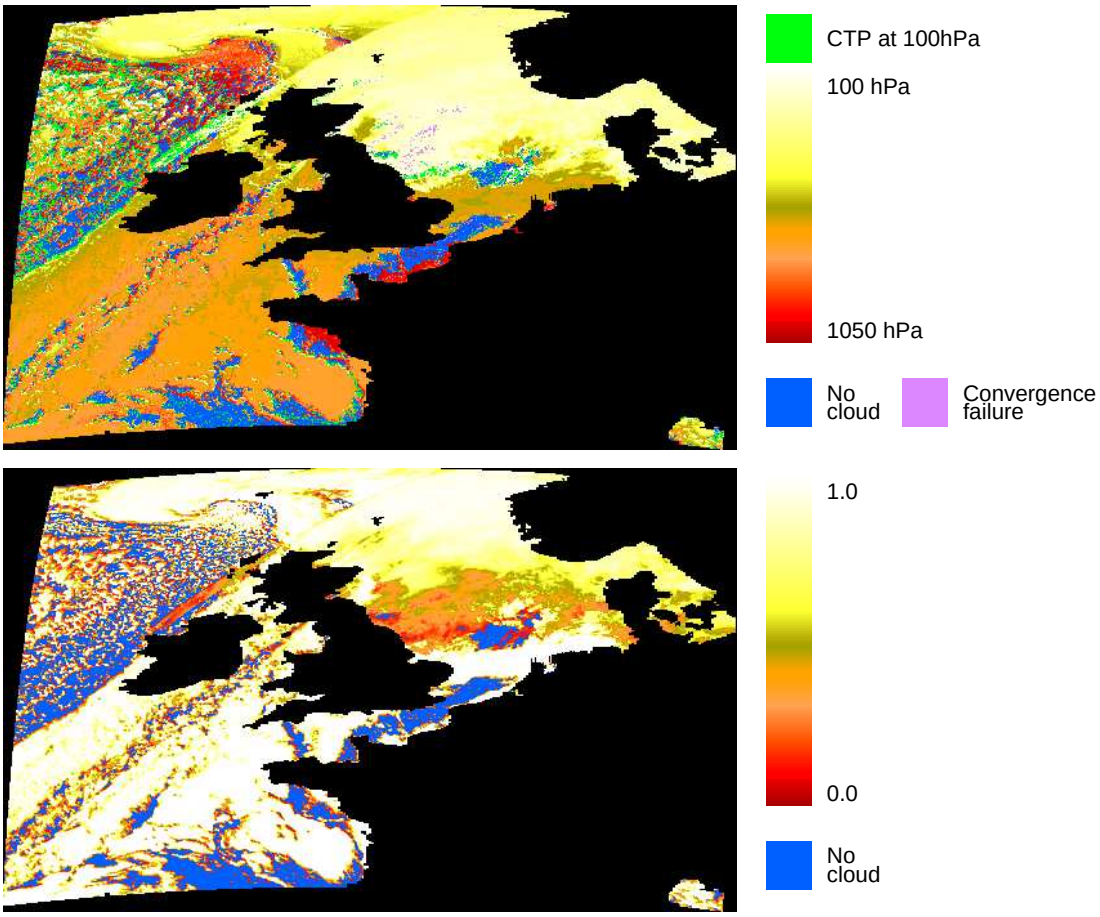

Figure 21: Maps of retrieved CTP and cloud fraction retrieved with single fields of view.

viewing geometry, the average number of gridboxes per pixel decreases quite rapidly with latitude across the region. However the infra-red images do not indicate particularly high pixel-to-pixel brightness-temperature variability for these areas, and it is not clear why multiple-FOV failures are associated more commonly with mid-level cloud. The successful convergence of single-FOV retrievals means it is unlikely that the background profiles make convergence difficult. In the case of thin, high cloud, where there are fewer failures in BAS2004 than in SFV2004, the multiple-FOV configuration may be providing a useful constraint.

One-dimensional histograms of cloud-top pressure (not shown) for SFV2004 contain the other non-physical features of peaks on and between RTTOV levels, if anything to a greater extent than for the multiple-FOV retrievals.

Figure 22, which shows the two-dimensional distribution of single-FOV vs multiple-FOV cloud-top pressure in histogram form, indicates that on average, retrieved cloud is slightly lower in BAS2004 than in SFV2004. Where cloud-top height differs by a large amount, it is more common for cloud to be lowered than raised by the multiple-FOV scheme. In particular, there are a significant number of pixels for which cloud in SFV2004 is very high — above ∼250 hPa — which are given a wide range of CTP values in BAS2004. The common cloud-top pressure requirement in the multiple-FOV configuration is probably constraining the retrievals, effectively imposing a spatial smoothing and confirming the suggestion in Figure 21 that some of the pixels given very high CTP by the single-FOV retrievals are isolated "noise", and possibly problematic retrievals. BAS2004 does raise some cloud, in particular cloud which is retrieved below ∼950 hPa by SFV2004,

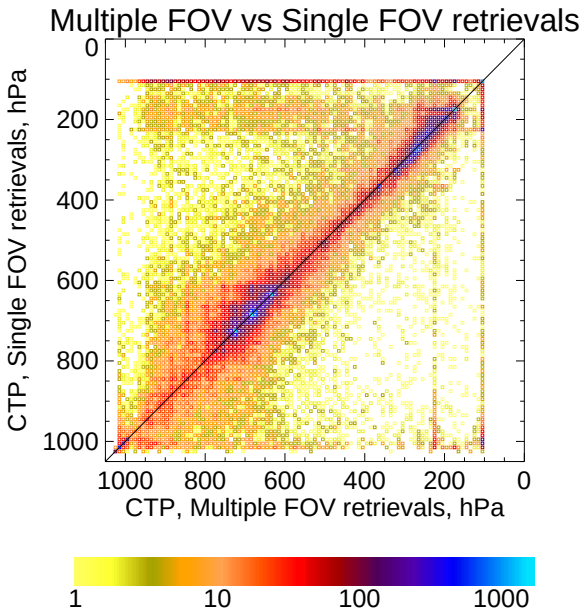

Figure 22: Two-dimensional histogram comparing CTPs retrieved in the single- and multiple fields of view configurations, on a pixel-by-pixel basis.

which may also be due to the common CTP requirement. There is also a significant (though smaller) number of pixels which have common retrieved CTP either at the 100 hPa limit imposed by the minimisation scheme, or close to the RTTOV level at 222.94 hPa (which is both the upper limit in the minimum-residual scheme used for the first guess, and one of the RTTOV peaks on which peaks in histograms of retrieved CTP tend to occur).

Figure 23 shows CTPs retrieved with the single- and multiple FOV configurations, plotted against their respective first-guess values. The single-FOV retrievals tend to show more fidelity to their first guesses, while the multiple-FOV scheme is more likely to lower cloud significantly from the first-guess value, particularly if this value is at the upper limit of the first-guess scheme. (Note that the first-guess values used in these plots are not the same — BAS2004 uses a common value for all FOVs in a retrieval, derived from the cloudiest first-guess in a retrieval. The high first-guess clouds are thus less likely to be isolated "noisy" pixels but smoother regions of high cloud.) In contrast, some of the very low cloud in the multiple-FOV first guess is raised by the 1D-Var retrieval.

The "edge effects" seen in the distribution of very cloud, in both the single- and multiple-FOV configurations, may be associated with the high latitude via a satellite zenith angle effect, but the mechanism is not clear. (Solar zenith angle effects, apparent in Figure 2, should not be apparent, as none of the channels used is sensitive to solar wavelengths; in any case the angle of the apparent "shadowing" is more consistent with satellite than a solar effect.) This tendency is also seen in the first-guess (minimum-residual) maps (not shown), although the lower pressure limit in the minimum-residual scheme is 222.94 hPa, rather than 100 hPa. The high cloud could be either common to both schemes, or inherited by the 1D-Var retrievals from their first guesses.

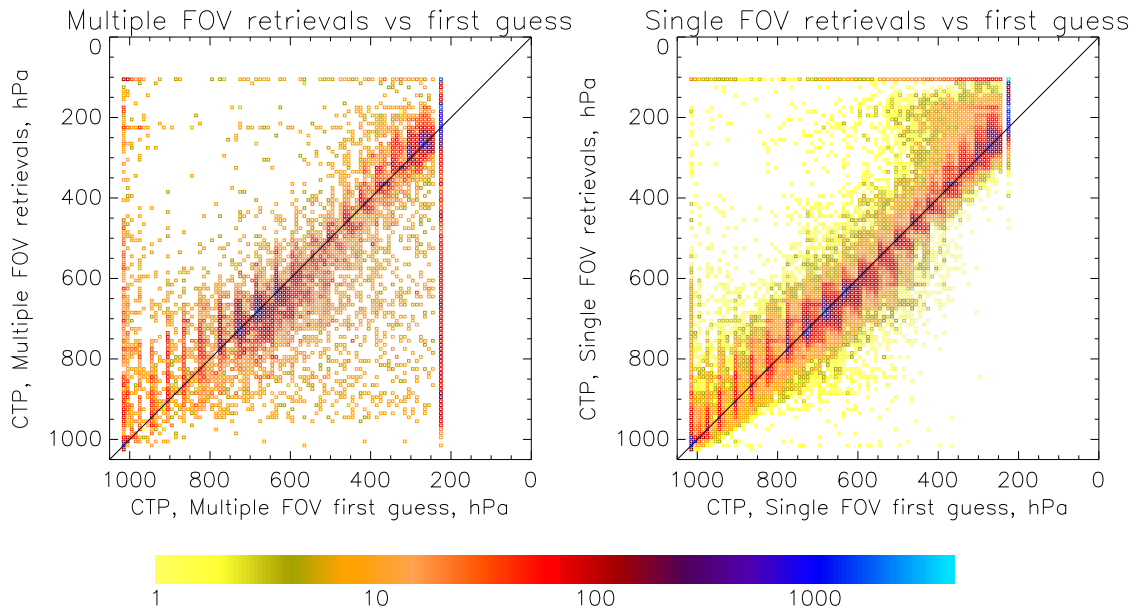

Figure 23: Two-dimensional histograms comparing CTPs retrieved in the single- and multiple fields of view configurations to their respective first-guess values, on a pixel-by-pixel basis. (In the multiple-FOV histogram, the first-guess and retrieved CTP are both common to all FOVs. Hence the very small amount of yellow, where only one or two pixels are counted, in this plot.)

3.5.2 Optimisation of the multiple-FOV formulation

The implementation of the multiple fields of view formulation used for this study did not in fact employ a single calculation of $H^{\rm op}(P_{cl}, \mathbf{x}_a)$ and $\partial H^{\rm op}(P_{cl}, \mathbf{x}_a)/\partial P_{cl}$ for all FOVs common to a model gridbox, as described in Section 2.5. (Common values of P_{cl} and $\mathbf{x}_a$ were used.) In principle results using a common calculation should be identical, but a test (SHP2004) showed that retrieved CTPs were not identical to the control case. Compared to the other tests presented the differences were small, and were probably due to the small differences observed in calculated values of $H^{\rm op}(P_{cl}, \mathbf{x}_a)$ and $\partial H^{\rm op}(P_{cl}, \mathbf{x}_a)/\partial P_{cl}$ between FOVS. Using common calculations of $H^{\rm op}(P_{cl}, \mathbf{x}_a)$ and $\partial H^{\rm op}(P_{cl}, \mathbf{x}_a)/\partial P_{cl}$ is expected to give considerable computational saving, but reliable timing comparisons were not available due to difficulties in controlling for external factors affecting the machine being used, and factors such as the level of diagnostic output. The mean number of iterations required for convergence increased slightly, from 5.2 in BAS2004 to 5.3 in SHP2004.

When setting common first-guess values for the cloud parameters, it was found that setting a common cloud-top pressure Jacobian (INH2004) (in addition to a common cloud-top pressure) gave significant savings: the mean number of iterations was reduced to 4.5, from 5.2 in BAS2004, and the number of recalculations due to increases in γ was slightly lower. There was a slight rise (4%) in number of CTPs reaching 100hPa, but otherwise no significant effect on retrieved CTP (not shown) or the distribution of final J values. This adjustment was not adopted in any of the other tests in this study, but given that it is made for good physical reasons, it should be considered.

3.6 Sensitivity to Input Information

3.6.1 Reducing profile size

As discussed in Section 2.2 above, fixing temperatures may allow computational savings, without significantly affecting retrievals. This is rather difficult to test without examining all retrieved quantities simultaneously, as temperature increments are correlated with increments to other elements of the profile vector by off-diagonal elements of $\mathbf{B}$.

2-D histograms showing the effect of reducing the length of profiles on retrieved CTP, ECA and T_{skin} are shown in Figure 24. The left-hand column shows the effect of fixing upper-air and 2-m temperatures but leaving T_{skin} free to vary (P332004). For lower cloud, with CTPs above about 500 hPa, cloud is on average slightly lower when temperatures are fixed (although with sufficient spread that some cloud is raised), while the majority of retrievals have very similar values of ECA and T_{skin}. Lowering cloud tends to increase modelled brightness temperatures; as the distributions of retrieved brightness temperature (not shown) are statistically very similar to those for the full retrieval vector, other changes in the profile are probably acting to cool the modelled brightness temperatures. However it is difficult to verify this argument without quantifying the sensitivity of modelled brightness temperature to all retrieved quantities. The opposite effect is seen for higher cloud (CTPs below 400 hPa) — retrieved cloud is slightly higher and so colder when temperatures are fixed, cloud amount is unchanged and T_{skin} is slightly colder, further cooling the scene. Other changes in the profile are probably acting to warm retrieved brightness temperatures and so preserve brightness-temperature fits.

The right-hand column of Figure 24 compares retrievals where T_{skin} is free but all other temperatures fixed (P332004) to retrievals where all temperatures including T_{skin} are fixed (P322004). CTP is only affected significantly for high cloud (CTPs below ∼400 hPa), where it is slightly raised by allowing T_{skin} to vary. For these high clouds, the overall distribution of T_{skin} increments has a significant tail towards negative values, with a reduction in ECA which probably has a compensating effect on retrieved brightness temperature. (Again, the effect on other profile elements is not examined; these may have at least as important an effect on retrieved brightness temperature.) For lower cloud, CTP is largely unaffected, but there is a suggestion that increased ECA may be compensating for increased T_{skin}. Clustering of CTPs on RTTOV levels and at 100 hPa is clearly visible in the P322004 T_{skin}-CTP plot, through tails to large positive values of T_{skin} difference. The absence of such tails in the P332004 T_{skin}-CTP plot suggests that where skin temperature does show big variations, it is consistent between different configurations, although the CTP-CTP plots indicates that the set of pixels affected is not identical.

Retrievals using the cut-down retrieval vector were observed to be quicker, but by a factor of less than two, which included a slight reduction in the number of iterations required for convergence (a mean of 4.8 for fixing upper-air and 2m temperatures, and 4.9 for fixing all temperatures). This probably indicates that, in this study, forward modelling was a more significant expense than linear algebra.

Excluding temperatures from the retrieval means that the typical length of the retrieval vector n is reduced from 83 to 39 or 40. Since the number of measurements m does not change, a typical "reduced" retrieval has $m > n$, as opposed to $m < n$ for the full-length retrieval vector. There is an alternative formulation for the matrix algebra used by 1D-Var (Rodgers, 2000) which is more efficient for cases where the measurement vector is longer than the retrieval vector, but this can only be easily applied with Gauss-Newton minimisation, rather than the Marquardt-Levenberg method used for this study.

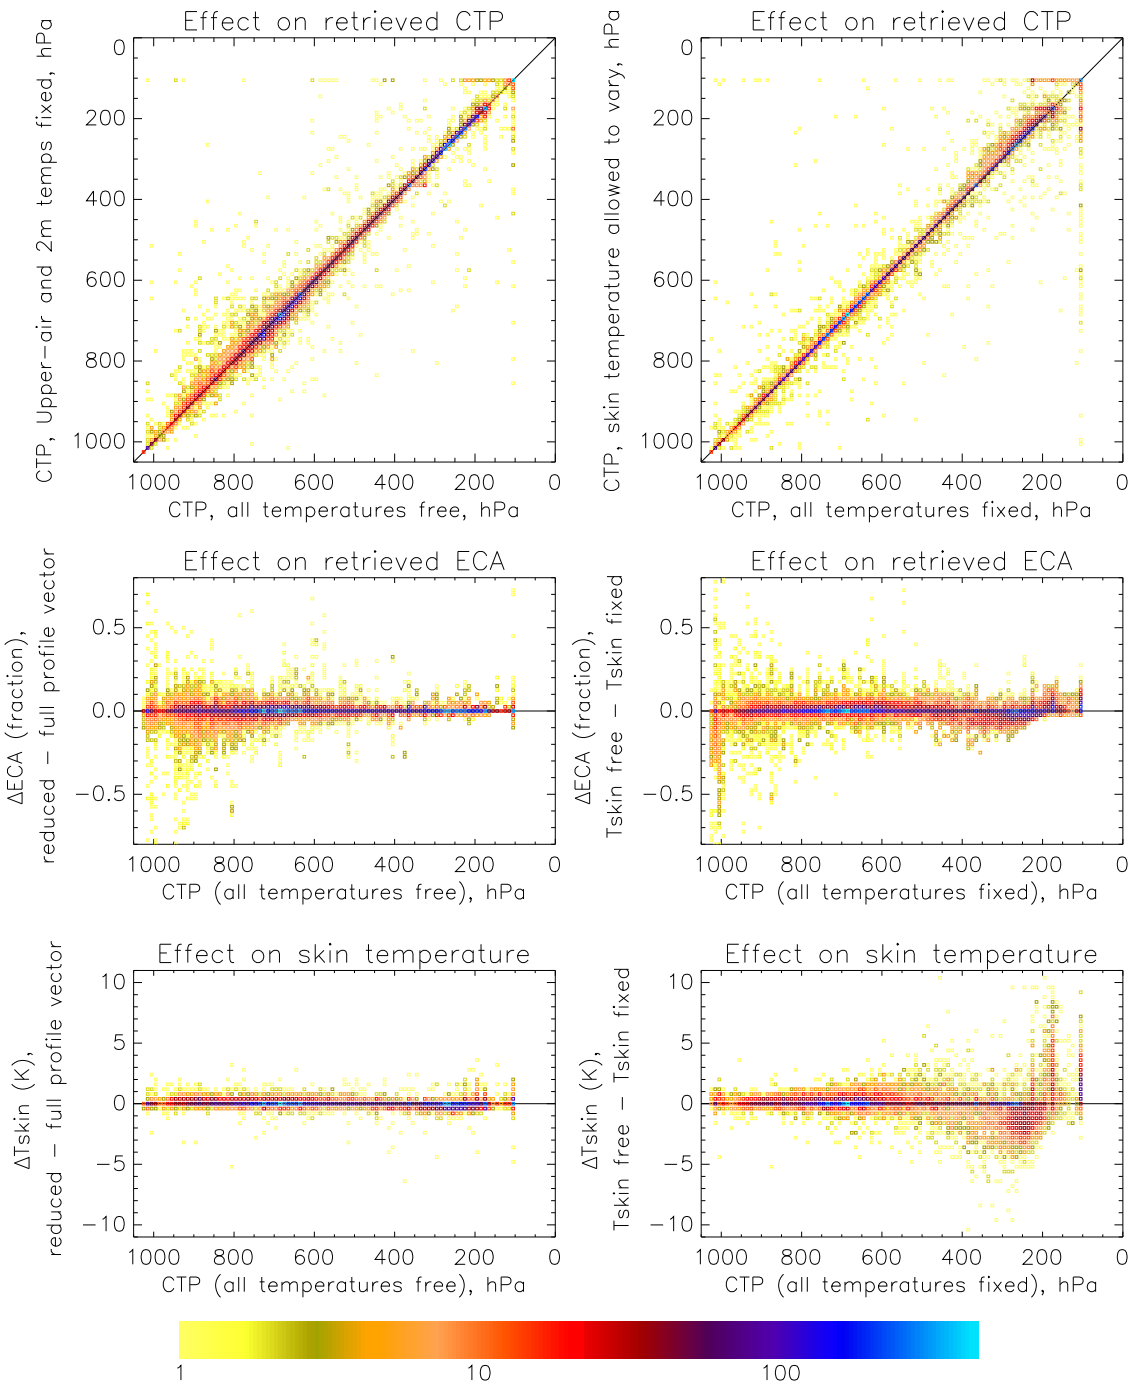

Figure 24: Two-dimensional histograms showing the effect of removing temperature elements from the profile vector on retrieved CTP, ECA and skin temperature T_{skin}. The left-hand column compares retrievals with and without upper-air and 2m air temperatures (T_{skin} being allowed to vary in both), while the right-hand column compares retrievals where all temperatures are fixed with retrievals in which T_{skin} alone is allowed to vary. In the CTP-CTP (top row) and T_{skin}-CTP (bottom row) histograms, each count represents one model gridbox, but in the ECA-CTP histogram (middle row) each count represents a SEVIRI pixel. The bin size on the ECA axis is 0.025, whilst on the T_{skin} axis it is 0.4K.

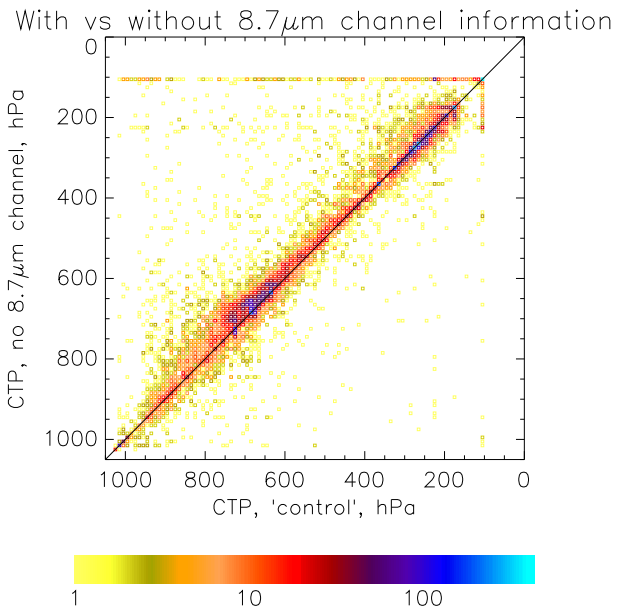

Figure 25: Two-dimensional histogram comparing CTPs retrieved with and without information from the 8.7μm channel, on a gridbox-by-gridbox basis.

3.6.2 Excluding the 8.7μm channel

A comparison of the cloud-top pressures obtained in retrievals with and without brightness temperatures from the 8.7μm channel (BAS2004 *vs* N872004) is shown in Figure 25. The 8.7μm channel is the channel for which the grey cloud approximation — the assumption that emissivity can be taken to be constant across all channels used in the retrieval — is least accurate, particularly for ice cloud: the imaginary part of the refractive index of ice is lower for the 8.7μm channel than for other channels (e.g. Warren, 1984) so absorbtivity and hence emissivity are lower. Forward-modelled brightness temperatures for a given ECA then have a warm bias, as the modelled radiance contribution from opaque cloud is larger than would be measured. This suggests that including observations in the 8.7μm channel in retrievals would tend to cool scenes, and any effect on cloud would be to raise it (lowering CTP) and/or extend its coverage.

Cloud-top pressures are however on average higher in BAS2004, which includes the 8.7μm channel, than in N872004, particularly in the range 500 to 900 hPa, where the effect is similar in magnitude to that of the methods tested for reduction of the peaks between RTTOV levels (Section 3.4.1). The one-dimensional histogram of CTP (not shown) shows that the peaks between RTTOV levels are considerably smaller in N872004. (The first guesses used in the two cases were not identical, as the 8.7μm channel information was also excluded from the first guess.) This suggests that the more important argument against using the 8.7μm channel is the exacerbation of problems in the minimisation (possibly because a warm bias in RTTOV calculations in the 8.7μm channel leads to difficulties reconciling it with other channels), rather than a systematic bias in retrieved CTP.

The one-dimensional histogram of CTP also demonstrates that excluding the 8.7μm channel does not prevent clustering of retrievals on RTTOV levels. "Scatter" in the CTP-CTP plot

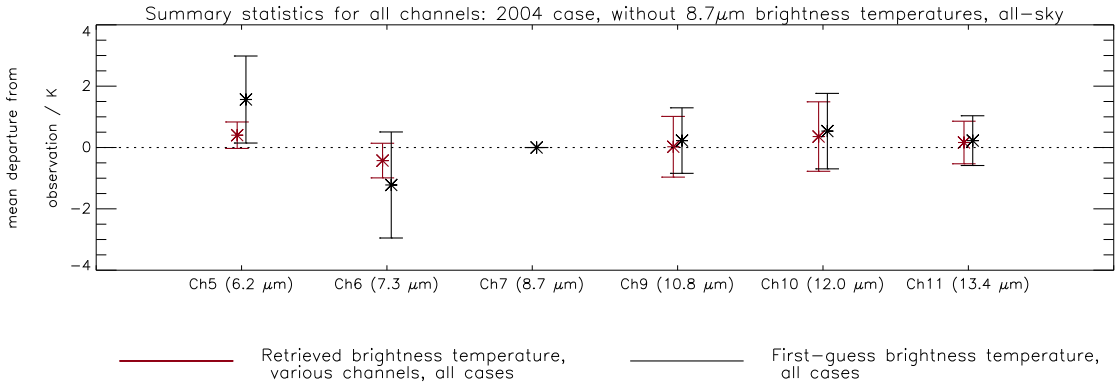

Figure 26: Summary statistics for brightness temperature fits (retrieved − observed) for N872004, where the 8.7μm channel is excluded from the retrieval.

(Figure 25) is relatively high (compared to other configuration tests), with CTP in some cases being significantly affected by exclusion of the 8.7μm brightness temperature. In particular, the number of retrievals where cloud is placed at 100 hPa is larger in N872004 than in BAS2004. The constraint provided by the 8.7μm brightness temperature might thus have positive as well as negative effects.

Summary statistics for the brightness temperature fits, shown in Figure 26, are similar to those in the control configuration (Figure 15). The statistics are similar to those for BAS2004 - both (first-guess − observed) and (retrieved − observed) temperature distributions have a larger variance in N872004 than in BAS2004, and the retrieval improves the first-guess fits to a similar extent, when compared to the effect of using one of the variant minimisations (Figure 20).

It is difficult to recommend inclusion or exclusion of the 8.7μm channel based on the tests presented here. A clearer signal of any systematic bias might be apparent using one of the variant minimisations discussed in Section 3.4.1 as a baseline. It might also be worthwhile to compare the effect of taking out either the 10.8- or 12.0μm channels (the other channels for which the variant minimisations significantly improve brightness-temperature fits) to see if the 8.7μm channel is particularly significant in the context of the problem of peaks between RTTOV levels.

3.6.3 Using AUTOSAT products as a first guess

Retrievals were performed for the 2004 case using AUTOSAT products (Section 3.1) to provide first-guess CTP and ECA (OFG2004), and comparisons between retrievals with different first guesses are given in Figure 27. In an ideal minimisation, retrieved state should be independent of first guess. Comparison of retrieved CTPs (shown on the left-hand side of Figure 27) shows that the 1D-Var scheme used here is reasonably robust to first guess; the comparison of retrieved CTPs is a lot cleaner than the comparison of first-guess CTPs (shown on the right).

Comparing the 1D-Var CTPs in BAS2004 and OFG2004 to their own first-guess CTPs (not shown) shows a tendency for 1D-Var to raise low cloud from both first guesses, with a greater effect when using the (lower) AUTOSAT first guess. There appears to be only a slight tendency for cloud retrieved in OFG2004 to be lower than in BAS2004 (complicated by some evident overfidelity to the first guess, as identified in Section 3.4.2 - note that there are also peaks in the distribution of AUTOSAT CTPs, giving rise to a grid pattern in the histograms, but these appear on, rather than between, RTTOV levels). This is consistent with the tendency noted in Section 3.2.1 for

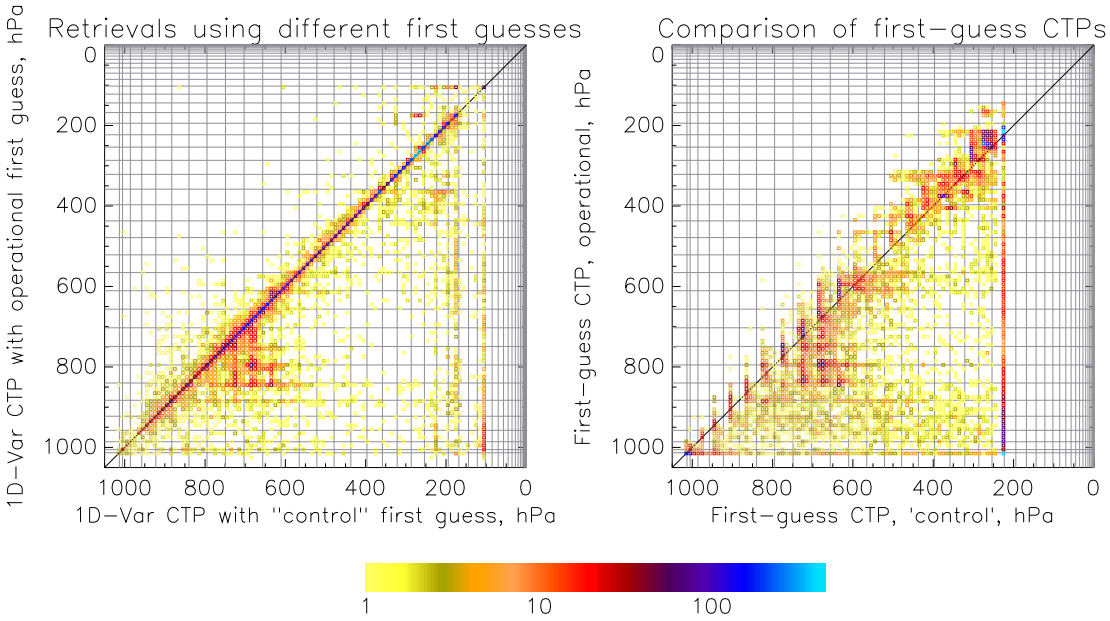

Figure 27: Comparison of retrieved CTPs using two different sets of first-guess parameters: 1D-Var vs. 1D-Var (left) and first guess vs. first guess (right).

low cloud in BAS2004 to be higher than that diagnosed by AUTOSAT, and suggests that this is unlikely to be due to the use of first-guess cloud that is too high. The number of high-cloud pixels in BAS2004's first guess - apparent in the right-hand plot in Figure 27 - does however seem to be problematic; there is less very high cloud retrieved in OFG2004 (0.8% of gridboxes) than in BAS2004 (3.1%), and the comparison of BAS2004's 1D-Var CTP with its first guess suggests that most of the very high 1D-Var cloud derives from cloud at the upper limit in the BAS2004 first-guess scheme. The high first-guess cloud can be spread downwards to all levels, which does indicate some robustness in the 1D-Var scheme.

The convergence rate in OFG2004 is slower, a mean of 5.9 iterations being required. than in BAS2004 (5.2 iterations). This is probably because the profile changes less in the course of minimisation when the 1D-Var scheme's own first guess is used. Although the 1D-Var scheme's first guess is likely to have deficiencies, the apparent robustness of the 1D-Var retrieval means it is unlikely to be worth coding a more accurate, and probably itself more expensive, first guess. The exception to this would be some effort to reduce the amount of very high cloud in the first guess, as this appears to exacerbate the problem of very high cloud in the retrievals.

4 Summary and recommendations for future work

The cloud products produced by the 1D-Var retrieval seem plausible, but show some consistent differences to both AUTOSAT and MODIS products. Notably, these include a tendency to place low cloud at higher altitudes than either AUTOSAT or MODIS (and higher than other schemes contributing to the Norrköping cloud intercomparison workshop), and to describe the edges of cloud in terms of a variation in cloud amount rather than cloud-top pressure. Comparison with

visible imagery suggests that the 1D-Var scheme may give a better description of thin, high cloud than other retrievals, but this is less consistent and is complicated by the possibility of underlying cloud. Other observations, particularly from other types of instrument, such as the cloud-profiling lidar CALIOP (unavailable for the cases presented) would be very useful in resolving these differences.

In terms of accuracy, the 1D-Var scheme has some specific problems:

- CTPs peaking between RTTOV levels, probably due to the use of a hard limit on cloud amount in the numerics of the retrieval. This can be reduced by adopting one of the two variant minimisations described in this report, which is recommended (slower convergence rates notwithstanding). Use of an additional cost-function term produces some retrievals in which cloud amount is unphysical; it also introduces extra parameters whose values must be decided. Modifying the Marquardt-Levenberg scheme is arguably a more objective method but gives smaller improvements to brightness-temperature fits.

- CTPs peaking on RTTOV levels. It is recommended that more smoothly interpolated cloud Jacobians - available in more recent versions of RTTOV - be tested in the retrieval scheme.

- A CTP peak at 100hPa. This was unexplained. It is possible that there some physical constraint is not adequately represented in the retrieval scheme. This also seems to be rather sensitive to the exact configuration and input data being used.

Points affected by the first two of these problems cannot easily be identified, so comparisons with other products may be complicated by errors.

Although not tested in an operational environment, it is likely that the scheme is too slow for operational use in its current configuration; processing the scene used for this study took several hours. There may be scope for applying the scheme selectively, if particular cases where it might be more accurate could be identified - it was noted, when comparing products, that where a scheme involves more than one method, the choice of method can be a significant source of uncertainty. Convergence in the minimisation scheme itself generally seems acceptable - using multiple fields of view gives slightly slower convergence, a result consistent with the simulation study by Szyndel *et al.* (2004). There are indications that slow convergence and convergence failures may be more common in particular circumstances, but these were not investigated. Some problematic aspects of the retrieval, such as the tendency for cloud to be placed at very high altitudes, appeared to be reduced by the constraint provided by the multiple fields of view formulation. It also appeared to reduce sensitivity to the first-guess cloud parameters used by the scheme. There are a couple of adjustments that could be made to the MFOV retrieval code (ensuring it uses common values for both forward-modelled cloudy radiance and cloud-top pressure Jacobian for all pixels associated with a model gridbox) to improve the economy of the scheme.

Fixing temperature elements in the background profile was not found to have a very large effect on the retrieved cloud products, although an accurate estimate of the computational savings was not obtained. Skin temperature appeared to be rather sensitive to fixing temperatures, but this may be reduced if minimisation problems are addressed. The question of whether or not to use the 8.7μm channel in the retrieval is less clear, and again, is probably best tested in conjunction with an improved minimisation scheme. It is probably not worth using extra computational resources to produce more accurate first-guess cloud parameters, as the scheme seems reasonably robust to its first guess; the exception would be to reduce the amount of very high cloud in the first guess, which in turn would reduce the amount of problematic high cloud in the retrievals.

5 References

Eyre, J.R., 1989. Inversion of cloudy satellite sounding radiances by nonlinear optimal interpolation. I: Theory and simulation for TOVS. QJRMS, **115**, 1001–1026.

Eyre, J.R. and Menzel, W.P., 1989. Retrieval of cloud parameters from satellite sounder data: a simulation study. J. Appl. Meteorol., **28**, 267–275.

Lorenc, A.C., Analysis methods for numerical weather prediction. QJRMS, **112**, 1177–1194.

Menzel, W.P., B.A. Baum, K.I. Strabala and R.A. Frey, 2002. Cloud top properties and cloud phase algorithm theoretical basis document, version 6. MODIS ATBD Reference Number: ATBD-MOD-04 (NASA-GSFC, 2002).

Moseley, S., 2003. Changes to the Nimrod Cloud Top Height Diagnosis. Met Office Forecasting Research Technical Report No. 424 [2] .

Phalippou, L., 1996. Variational retrieval of humidity profile, wind speed and cloud liquid water path with the SSM/I: Potential for numerical weather prediction, *Q. J. R. Meteorol. Soc.,* **122**, 327–355

Rodgers, C.D., 2000. Inverse Methods for Atmospheric Sounding: Theory and Practice. World Scientific, Singapore.

Saunders, R.W., M. Matricardi and P. Brunel, 1999. An Improved Fast Radiative Transfer Model for Assimilation of Satellite Radiance Observations. QJRMS, **125**, 1407–1425.

Saunders, R.W. *et al.*, 2006. The exploitation of Meteosat Second Generation data in the Met Office. Proceedings of the 2006 EUMETSAT Meteorological Satellite Conference, Helsinki, Finland.

Schmid, J. The SEVIRI Instrument. Proceedings of the 2000 EUMETSAT Meteorological Satellite Data Users' Conference, Bologna, Italy, May 29-June 2, 2000, pp. 23-32

Szyndel, M.D.E., A.D. Collard and J.R. Eyre, 2004. A simulation study of 1D variational cloud retrieval with infrared satellite data from multiple fields of view. QJRMS, **130**, 1489–1503.

Thoss, A. Report on the Cloud Parameter Retrieval Workshop 17-19 May 2006, Norrköping. [3]

Warren, Stephen G., 1984. Optical constants of ice from the ultraviolet to the microwave. Applied Optics, **23**, 1206–1225.

[2] `http://www.metoffice.gov.uk/research/nwp/publications/papers/technical_reports`

[3] Available from `ftp://ftp.eumetsat.int/pub/CWS/CWS_general`