Numerical Weather Prediction

Improvements to the Nimrod Visibility Analysis/Forecast System

Forecasting Research Technical Report No. 217

Bruce Wright

email: nwp_publications@metoffice.com

Forecasting Research
Technical Report No. 217

Improvements to the Nimrod Visibility Analysis/Forecast System

by

Bruce Wright

July 1997

Forecasting Research
Meteorological Office
London Road
Bracknell
Berkshire
RG12 2SZ
United Kingdom

Abstract

An objective visibility analysis/forecast system (VAFS) was developed as part of Nimrod, the United Kingdom Meteorological Office (UKMO) very-short-range forecasting system. The development was carried out in a number of stages, and after testing on a range of case studies, the VAFS was implemented on Nimrod. In the case studies, which were, naturally enough, mainly fog cases, the analysis performed well. However, monitoring of the day-to-day running of the system, coupled with long-term verification statistics and provisional results from an NMC Trial, have indicated some deficiencies with the current scheme. This has led to a major revision of the way in which the visibility is diagnosed from the prognostic variables (liquid water temperature and total water content), a new scheme for the use of satellite data, the addition of a fog probability diagnostic and various other smaller changes.

To obtain a more realistic assessment of the performance of the new scheme against the current one, a parallel visibility trial (PVT) was run on NIMDEV (the Nimrod backup machine), using the same data as the current scheme. Real-time monitoring of the results of the PVT, led to the discovery of problems with the revised scheme, and subsequent revisions were introduced. The lack of extensive fog through much of the PVT period, especially towards the end, was a disappointment, but the trial did yield useful information both on the current and the revised scheme. In light of this, further tests were carried out on an extended visibility case study (EVCS), covering seven days of foggy weather in the middle of the PVT.

The most significant change, using a fixed probability based total water value rather than the integral over a distribution of states for the diagnosis of the visibility, appears to have a small positive impact, if an adjustment is applied to the Mesoscale Model (MM) total water to remove the bias perceived by the new scheme. Although the improvement is small overall, the distribution of states, as displayed through the fog probability statistics, appears to be better. Subjective verification confirms that visibilities are slightly higher in the new scheme, but there is little overall difference in quality. The improved satellite scheme, which essential spreads surface observation information, provides significantly better analyses where satellite data are available. However, further benefits could potentially be gained by using the surface observations in a similar way when no satellite data are available.

Improvements to the Nimrod Visibility Analysis/Forecast System

1. Introduction

An objective visibility analysis/forecast system (VAFS) has developed as part of Nimrod, the United Kingdom Meteorological Office (UKMO) very-short-range forecasting system (Golding, 1995). The development was carried out in a number of stages (Wright, 1995, 1996a and 1996b, and Wright and Thomas, 1995). The final system was tested on a range of case studies, with test periods up to 18 hours, with the results presented in Wright and Thomas (1996). In mid-September 1996, the VAFS, together with a verification package was implemented on Nimrod. In following month, various problems were discovered and addressed, but the system was essentially stable from the end of October 1996 onwards. Through February 1997 into the beginning of March the system was trialled in the National Meteorological Centre (NMC), being compared directly with forecaster predictions for Heathrow and Gatwick.

In the case studies, which were, naturally enough, mainly fog cases, the analysis performed well. However, monitoring of the day-to-day running of the system, coupled with long-term verification statistics and provisional results from the NMC Trial, have indicated some deficiencies with the current scheme. This has led to a major revision of the way in which the visibility is diagnosed from the prognostic variables (liquid water temperature and total water content), a new scheme for the use of satellite data, the addition of a fog probability diagnostic and various other smaller changes. To obtain a more realistic assessment of the performance of the new scheme against the current one, a parallel visibility trial (PVT) was run on NIMDEV (the Nimrod backup machine), using the same data as the current scheme. Real-time monitoring of the results of the PVT, led to the discovery of problems with the revised scheme, and subsequent revisions which were introduced. The lack of extensive fog through much of the PVT period, especially towards the end, was a disappointment, but the trial did yield useful information both on the current scheme and the proposed revisions. In light of this, further tests were carried out on an extended visibility case study (EVCS), covering seven days of foggy weather in the middle of the PVT. This paper discusses the details of the new scheme, presents some results from the PVT and EVCS and suggests some further development work.

2. The new scheme

The current version of the VAFS is run hourly on Nimrod. Routine monitoring has highlighted a number of problems. The new scheme, which has been tested in the PVT and EVCS, differs from the current scheme in three distinct ways, each of which addresses a known shortfall of the current system; these are considered in the following three sections, with reference to the current scheme. A further diagnostic, fog probability, has been added; this is described in section 2.4.

2.1. Improved use of satellite data

In the satellite scheme, pixels are initially classified as either clear, cloud-contaminated or fog/low cloud using the infrared and visible (‘fog’ for AVHRR) imagery and the UK Mesoscale Model (MM) surface temperature; this step remains unchanged in the new scheme (see Wright and Thomas (1996)). Cloud-contaminated pixels are not used in either scheme, but clear and fog/low cloud pixels are ‘calibrated’ using the surface observations.

In the current scheme, clear pixels are assumed to be completely fog-free and the analysis is weighted towards values of T_L and q_t corresponding to a ‘clear’ relative humidity ($0.75 + 0.25\ RH_{crit}$) (see section 2.3.1 for an explanation of RH_{crit}); clear pixels are not set for AVHRR data . Potentially foggy pixels are assigned a visibility taken from the nearest observation within the ‘fog’; some allowances

are made for non-contiguous fog and use of foggy observations not perfectly co-located with satellite-derived fog, and the 'spread' of observational information is limited to land points up to a maximum of twenty grid-points away from the observation. Those points which have a visibility less than 1 km are treated as foggy, with the analysis values of T_L and q_t being weighted towards a relative humidity of 1.

In the new scheme, the visibility information in the surface observations is used more extensively, and temperature is used to recover T_L and q_t, to avoid the need to weight towards limits, which may not be appropriate in a particular instance. The observed visibilities and temperatures are spread independently for clear and potentially foggy areas; the logarithm of visibility is used. The values from the different observations within range (20 grid-points) are combined using a one over radius cubed weighting; this was found to give an acceptable 'satellite visibility product'. As with the current scheme, data are not used over the sea. The final visibility and temperature fields are converted to T_L and q_t values, using the same approach as for the surface observations, and are analysed in the same way as the other data. This approach leads to an analysis with a much closer areal fit to surface observations, but with significant boundaries being defined by the satellite data.

2.2. Limiting of model aerosol values

In the MM, visibility observations are assimilated through the aerosol mass mixing ratio. This can result in very large values, where the observed low visibilities are due to errors in the moisture field, but are (incorrectly) incorporated through changes in the aerosol field. To avoid the subsequent over-prediction of fog in the VAFS when the MM aerosol values are used, the values are limited to lie in the range 10^{-10} to 10^{-7} $kgkg^{-1}$ when they are read in.

2.3. Change to the diagnosis of visibility

The current scheme over-predicts fog, as perceived viewing the 'mean' visibility product. This is a trait shared by the UKMO Mesoscale Model (MM), which is not surprising as the same scheme is used by both to diagnoses mean visibility. To address this problem, it was decided to change the visibility diagnosis to use a 'median' cloud water value, rather than a 'mean' value, which should produce a 'later', but sharper, transition into thick fog.

2.3.1. Moisture variables

The prognostic variables are the liquid water temperature, T_L, and the total water content, q_t, which are defined by equations (1) and (2).

$$T_L = T - \frac{Lq_L}{C_p} \qquad (1)$$

$$q_t = q + q_L \qquad (2)$$

where T is the temperature, q is the specific humidity, q_L is the liquid water content, L is the latent heat of condensation and C_p is the specific heat capacity at constant pressure. Strictly speaking, equations (1) and (2) should include terms representing the contribution from the ice phase, but as ice fog is very rarely observed in the UK, the ice phase has been neglected throughout. Although T_L and q_t characterise the temperature and moisture environment within a grid-box, they do not say anything about the distribution of the moisture between water vapour and condensed water; further assumptions have to be made to carry out this moisture partitioning.

In the current scheme, which is also that used by the Unified Model (UM), a triangular distribution of moisture states about the saturated specific humidity is assumed (see Figure 1). The width of the

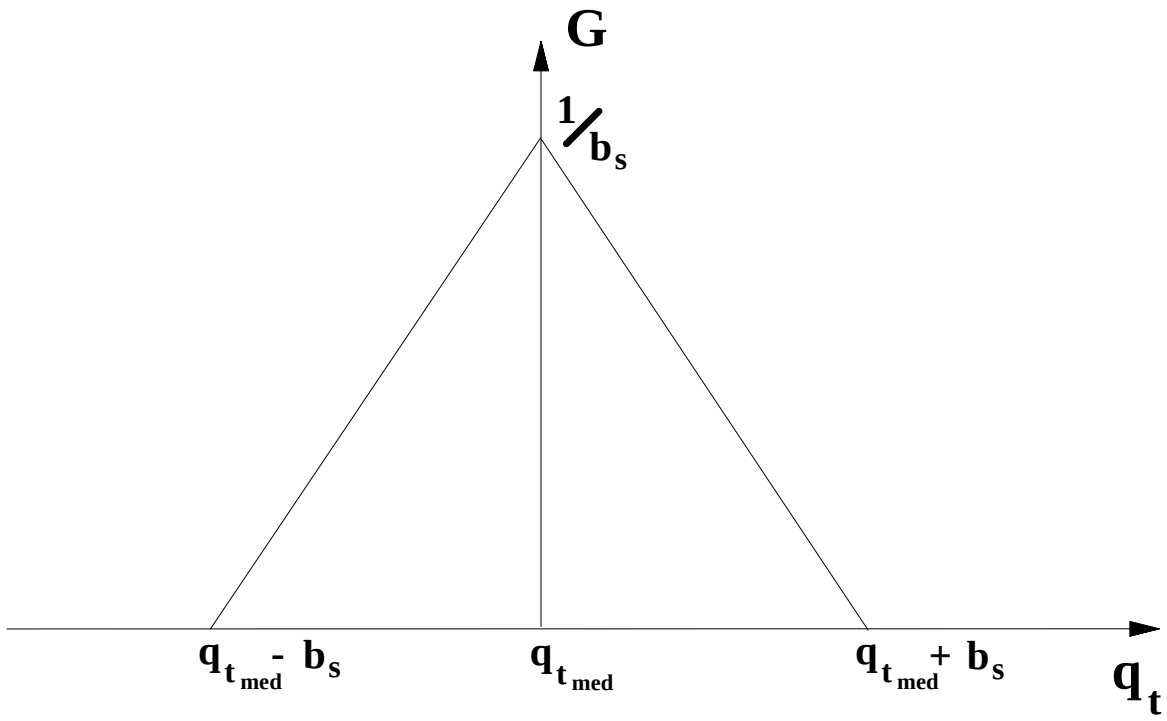

Figure 1 Distribution of states of total water within a grid-box.

distribution is characterised in terms of a relative humidity value, RH_{crit}, above which saturated states exist. Thus, the half-width of the moisture distribution expressed as a water content, b_s, is:

$$b_s = (1 - RH_{crit}) q_s(T) \quad \textbf{(3)}$$

where $q_s(T)$ is the saturated specific humidity at temperature T; this can be derived using a lookup table given T, and approximated using a Taylor expansion given T_L using equations (1) and (2). The cloud water content, q_L, is calculated by integrating over the distribution of states. Temperature and humidity are calculated using equations (1) and (2). The idea of this scheme is represent a variety of different states that may exist within the grid-box; it allows liquid water to exist, with a relative humidity of less than 1.0 (any value above RH_{crit}, in fact). The final quantities calculated can be thought of as grid-box mean values.

The new scheme does not partition the total water into humidity and liquid water before passing it to visibility diagnostic routines. This partitioning is internal (and integral) to those routines; broadly, it assumes that moisture in excess of saturation is liquid water, but does allow some super-saturation. This internal partitioning allows a consistent treatment through droplet activation (see section 2.3.2). The diagnostic routines (visibility, fog probability, etc) all require total water, temperature, pressure and aerosol. The latter two are taken unchanged from the MM. A temperature (and dew point, relative humidity, etc) are derived from the T_L and q_t using the UM scheme with an RH_{crit} value of 1.0; this is effectively taking the median grid-box state, but is fairly close to the mean value (current scheme) for these variables. For the total water value, the new scheme assumes the same triangular distribution of states, but modifies the total water value in light of the distribution, rather than integrating over all states to obtain a liquid water value. For example, the median (50% probability) visibility would be the value diagnosed from the q_t value with 50% of the area of the triangle to the right of it in Figure 1 (i.e. the unchanged prognostic value of q_t); the 30% probability visibility would be the value diagnosed from the q_t value with 30% of the area of the triangle to the right of it in Figure 1 (i.e. slightly larger than the prognostic value of q_t). Although, this change has little effect on the temperature, dew point, etc, it significantly effects the visibility, which drops very rapidly as liquid water is formed. Thus, in the new scheme visibility is far more sensitive to changes in humidity close to saturation, especially in clean air (low aerosol content).

2.3.2. Visibility

Visibility can be related to a fog droplet radius, r, and number density, N, using the following equation (based on Koschmeider, 1924):

$$Vis = \frac{-\ln\epsilon}{Nr^2\beta_o + \beta_{air}} \quad \textbf{(4)}$$

where ε is the liminal contrast (0.2), and the scattering coefficient, β_o, is given by:

$$\beta_o = \pi Q \eta \quad \textbf{(5)}$$

where Q and η represent the scattering efficiency and the effect of the aerosol particle size distribution respectively, with Q × η taking the value 1.5. β_{air} is a term to represent clear air scattering such that the maximum visibility is limited to the clear air visibility (100 km):

$$Vis_{air} = \frac{-\ln\epsilon}{\beta_{air}} \quad \textbf{(6)}$$

The relationship between visibility and droplet radius remains unchanged in the new scheme, but the derivation of the droplet radius from the total water ratio is changed.

With respect to visibility, the atmosphere can be divided in to two distinct regimes, depending on whether or not the fog droplets are activated. Unactivated droplets are essentially aerosol particles with a small amount of liquid water, and their mean droplet radius, r, can be related to the relative humidity, RH, by the equilibrium equation (Pruppacher and Klett, 1978):

$$RH = \exp\left(\frac{A_o}{r} - \frac{B_o}{\left(\frac{r}{r_d}\right)^3 - 1}\right) \quad \textbf{(7)}$$

where r_d is the mean volume radius of the dry aerosol, A_o is a constant (1.2×10^{-9} m) and B_o is the Activation Parameter (0.5). Activated fog droplets grow rapidly, and the cloud water content, q_L, can be related to the mean droplet radius, r, by geometric consideration:

$$q_L = \frac{4}{3}\pi (r^3 - r_d^3)\rho_w N \quad \textbf{(8)}$$

where ρ_w is the density of water (1000 kgm^{-3}) and r_d is the dry particle radius. As is apparent from the equations (7) and (8) a number density, N, and a dry particle radius, r_d, are required. These are provided by the MM aerosol field, which is in the form of a mass mixing ratio. To recover N and r_d requires an assumption to be made about the aerosol particle size distribution; namely that r_d varies as a power of the aerosol mass concentration, m, i.e.:

$$r_d = r_o \left(\frac{m}{m_o}\right)^p \quad \textbf{(9)}$$

where r_o is the radius of a standard aerosol particle (0.16×10^{-6} m), p is the power used to represent the variation in aerosol particle size with mass-loading ($^1/_6$), and m_o, the standard mass mixing ratio of the aerosol, is given by:

$$m_o = \frac{4}{3}\pi r_o^3 \left(\frac{\rho}{\rho_a}\right) N_o \quad \textbf{(10)}$$

where N_o is the standard number density of the aerosol (5.0×10^8 m^{-3}), ρ is the density of the aerosol (1700 kgm^{-3}) and ρ_a is the density of air (1 kgm^{-3}). Thus, the aerosol number density, N, is given by the equation:

$$N = N_o \left(\frac{m}{m_o} \right)^{1-3p} \qquad \textbf{(11)}$$

The number density of the aerosol is assumed to be equivalent to that of the fog droplets for use in both equations (4) and equation (8); this assumes that all fog droplets are activated in the later case.

Figure 2 shows the variation of relative humidity according to equation (7); the panel on the right has an expanded humidity scale and a compressed radius scale to illustrate the super-saturated section of the curve, where activation of the droplets occurs. In the current scheme, the first term in equation (7) is dropped to allow an analytic solution for r to be calculated:

$$r = r_d \left(1 - \frac{B_o}{\log(RH)} \right)^{\frac{1}{3}} \qquad \textbf{(12)}$$

This modified curve asymptotes towards saturation with increasing droplet radius size, but does not become super-saturated; its behaviour resembles that of the section of the full curve shown in the left panel of Figure 2. This means that as the relative humidity approaches 1.0, the droplet radius tends to infinity. Thus, to ensure sensible values of droplet radius are generated, the relative humidity is limited to 0.999. This is not really a problem, as significant cloud water is present before reaching the relative humidity limit, and the radius is calculated using equation (8), as soon as this method produces a larger droplet radius (lower visibility). Physically, this change of equations corresponds to the activation of droplets. However, in this simplified form of the equations it really just represents a rather messy change in the diagnosis method.

For the new scheme, which uses total water, q_t, directly, liquid water is assumed not to be present until the relative humidity exceeds 1.0 (the air is supersaturated); the relative humidity is defined using the equation:

$$RH = \frac{q_t}{q_s} \qquad \textbf{(13)}$$

To cope with this, it is necessary to solve the full form of equation (7), which accounts for the presence of unactivated droplets in a supersaturated air. This is illustrated in the right panel of Figure 2; droplet radius will grow slowly as the relative humidity is increased, until a activation radius is reached, and then the droplets will grow very rapidly, resulting in a drop in relative humidity. This approach still implies a change in the method of solution for r from equation (7) (in terms of relative humidity) to equation (8) (in terms of liquid water). It is not difficult to calculate the activation radius; neglecting the -1 in the denominator (r is much greater than r_d close to saturation), and differentiating with respect to r yields:

$$r_{act} = r_d \sqrt{\frac{3B_o}{A}} \qquad \textbf{(14)}$$

This can be substituted into equation (7) to obtain the activation humidity. It can, then, be assumed that if the relative humidity is greater than the activation relative humidity, that the droplet radius is calculated using a re-arranged form of equation (8), assuming that all water in excess of saturation is liquid water. However, in practice it is necessary to go both ways (from total water to droplet radius and vice versa), and it turns out to be difficult to cope with the transition at the droplet activation point consistently. For this reason (and because it results in a single simpler scheme), the two methods of deriving the droplet radius have been combined. Thus, the droplet radius, r, is obtained by the solving the equation:

$$q_t = RH(r)\, q_s(T) + q_L(r) \qquad \textbf{(15)}$$

where RH(r) and $q_L(r)$ are given by equations (7) and (8) respectively. In practice, this has to be solved using Newton-Raphson (see Appendix A). Although, the functions are rather messy, there is

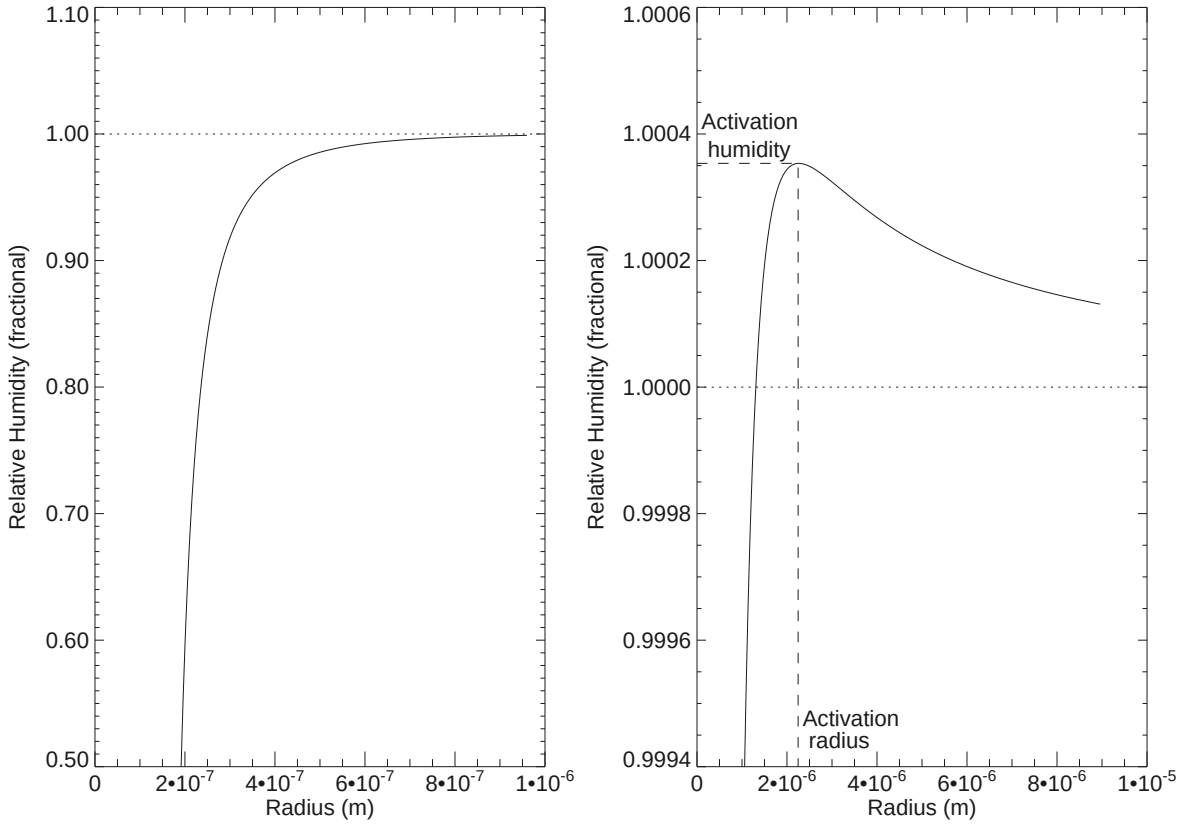

Figure 2: The variation of relative humidity against fog droplet radius given the droplet growth equation.

no problem in obtaining a solution, as the curve described by equation (15) is continuous and single-valued. Figure 3 shows the variation of the two terms in equation (15) and their sum against droplet radius. Equation (4) is used to diagnose the visibility from the calculated droplet radius.

2.3.3. Recovering the prognostic variables from the surface observations

Surface observations of temperature, T, dew point T_d, visibility, Vis, and pressure, p, are taken, but the prognostic variables are liquid water temperature, T_L, and total water mixing ratio, q_t, so these quantities have to be derived from the observed values and the MM aerosol, m. If the visibility is greater than 10 km or no visibility has been reported, then temperature and dew point are used to derive the prognostic variables. Otherwise, Vis and T are use.

If Vis > 10 km (or no visibility is reported), then a critical temperature, T_{crit}, which corresponds to the gridbox median temperature at which condensation starts occurring is calculated; it is the dew point for a relative humidity of RH_{crit}. Of course with the trial scheme, RH_{crit} = 1.0, so T_{crit} = T_d, but this is not the case with the current scheme. If T_d > T_{crit}, then no liquid water is present and the prognostic variables are given by:

$$T_L = T$$
$$q_t = q_s(T_d) \qquad \textbf{(16)}$$

If $T_d \leq T_{crit}$, then T_L and q_t are set to -1.0 times the values set in equation (16), to flag their values to be used as limits. If T_d > T, then the prognostic variables are set to missing data.

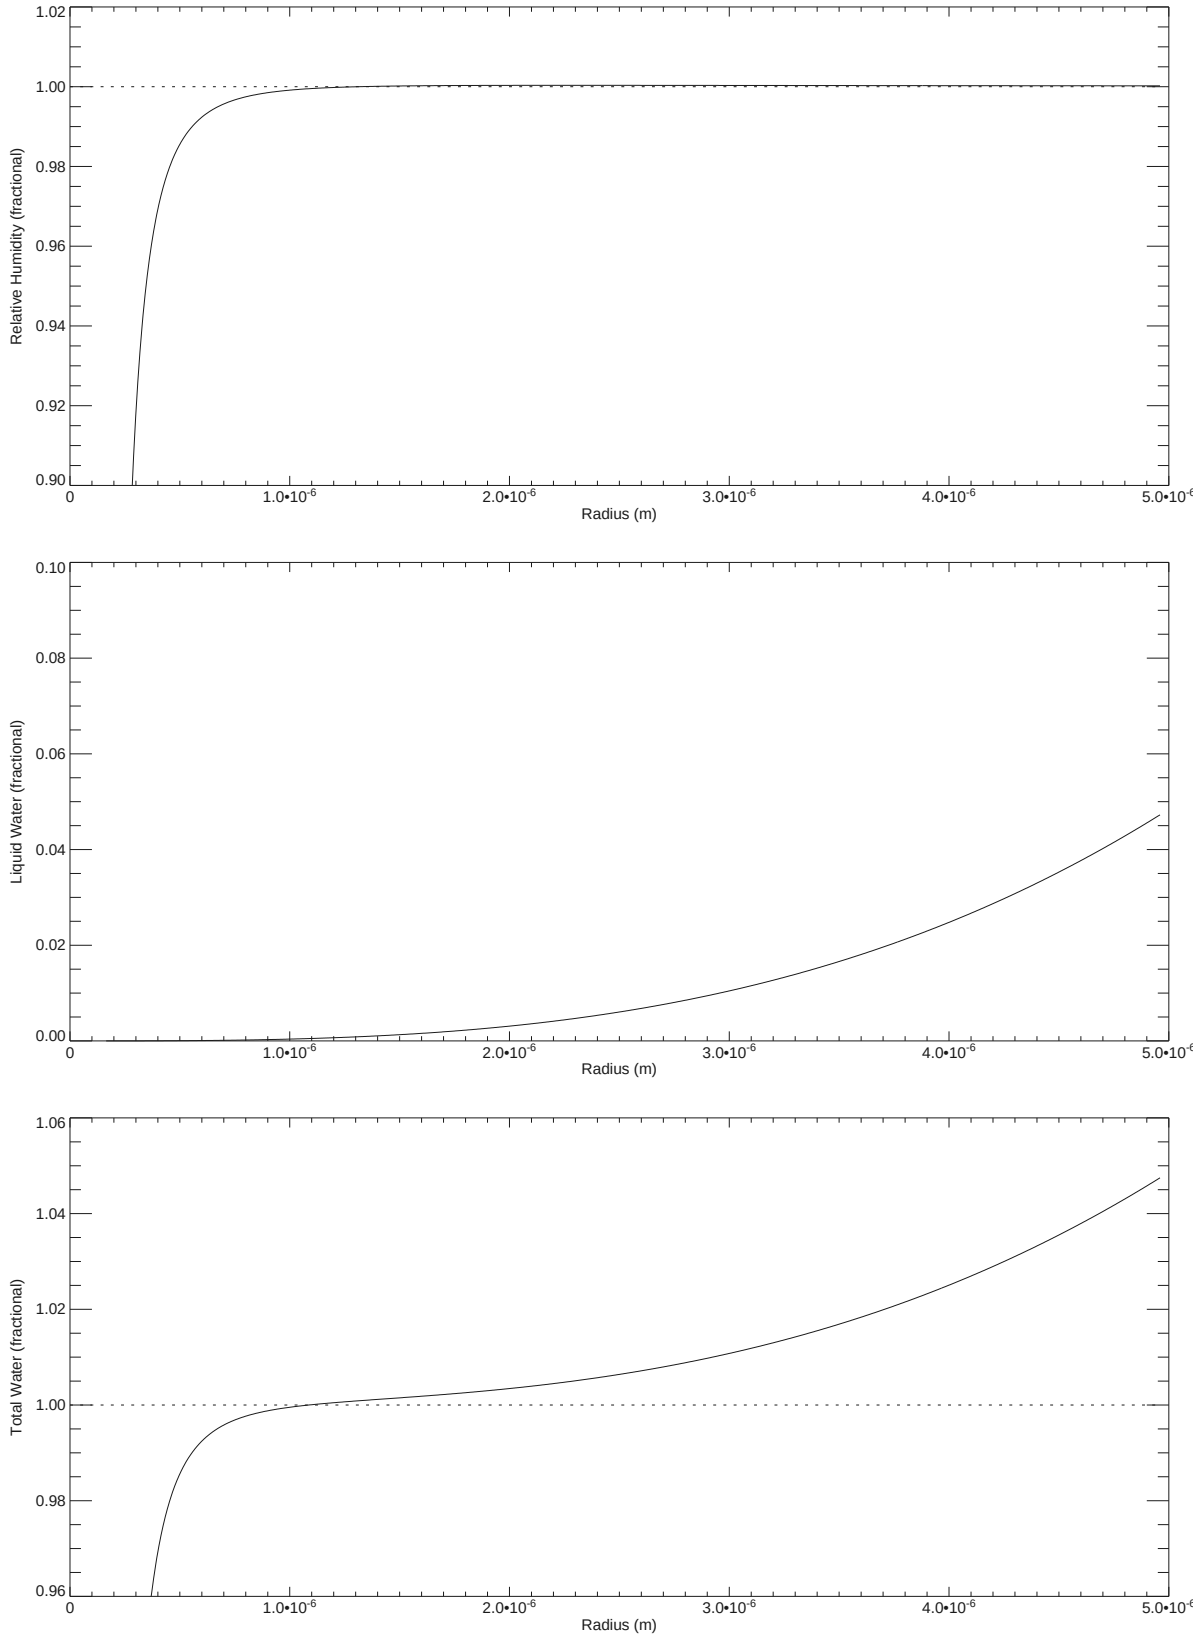

Figure 3: Droplet radius of (top) relative humidity, (middle) liquid water normalised by $q_s(T)$ and (bottom) the sum of the two.

If Vis ≤ 10 km, then the inverse of the visibility diagnosis (described earlier in this section) is used to recover T_L and q_t. A re-arrangement of equation (4) is used to recover the droplet radius from the

visibility:

$$r = \sqrt{\frac{\ln\epsilon}{N \beta_o}\left(\frac{1}{Vis} - \frac{1}{Vis_{air}}\right)} \qquad \textbf{(17)}$$

The details of the recovery of T_L and q_t from r for the current scheme are given in Wright & Thomas (1996). For the new scheme, if this droplet radius is less than the dry particle radius, then it is not possible to obtain values of the prognostic variables from the visibility, so the temperature and dew point are used as previously described. Otherwise, a median value of q_t is recovered using equations (7), (8) and (2). T_L is recovered from the inverse of equation (1), with the liquid water, q_L, equal to 0 if $q_t < q_s(T)$, or given by equation (18) otherwise.

$$q_L = q_t - q_s(T) \qquad \textbf{(18)}$$

The median value of total water, q_{tmed} is adjusted using the appropriate probability (see section 2.3.1), as follows. If the probability, prob, lies in the range: $0.0 \leq prob < 0.5$, then:

$$q_t = q_{t_{med}} + \left(1 - \sqrt{2\ prob}\ b_s\right) \qquad \textbf{(19)}$$

If $0.5 \leq prob \leq 1$, then:

$$q_t = q_{t_{med}} - \left(\sqrt{2\left(prob - \frac{1}{2}\right)}\ b_s\right) \qquad \textbf{(20)}$$

The final value if limited to be greater than 0.0001 $kgkg^{-1}$.

2.4. Fog probability

Fog probability is defined as the probability of the grid-square having a visibility below a given visibility threshold (typically 200 m, 1000 m or 5000 m). As the diagnosis procedure considers a distribution of states over the grid-square, the fog probability can equally well be considered as a fog fraction, the fraction of the grid-square having a visibility below the given threshold. The scheme requires the total water content, q_t, the temperature, T, the surface pressure and aerosol mass mixing ratio (MM values), and the threshold visibility, Vis_{thresh}. Firstly, a threshold total water content, $q_{t\ thresh}$, is diagnosed from Vis_{thresh} and T. Comparison between this threshold total water and the actual total water content is used to determine the cloud fraction.

The code for calculating a total water threshold from the visibility threshold is a simplified version of that used to obtain the values of T_L and q_t from the observed temperature, dew point and visibility; it uses only visibility and temperature, and does not calculate liquid water temperature values. Also, where the droplet radius calculated using equation (17) is smaller than the dry radius, the q_t is set to the lower limit value, 0.0001 $kgkg^{-1}$. Although, the visibility threshold is a scalar value, an array of total water threshold values are produced, as calculation also involves screen temperature, and screen level pressure and level 1 aerosol values (both provided by the MM).

A triangular distribution of moisture states about the median q_t value, with a width b_s, is assumed (see section 2.3.1). The fog probability for a particular threshold if found by calculating the fraction of states in the distribution with a total water value, which exceeds the threshold total water value. The easiest way to calculate this fraction is geometrically. There are four different situations:

1. $q_{t\ thresh} \leq q_t - b_s$: in which all states in the distribution exceed the total water threshold, and the fog fraction is 1.0

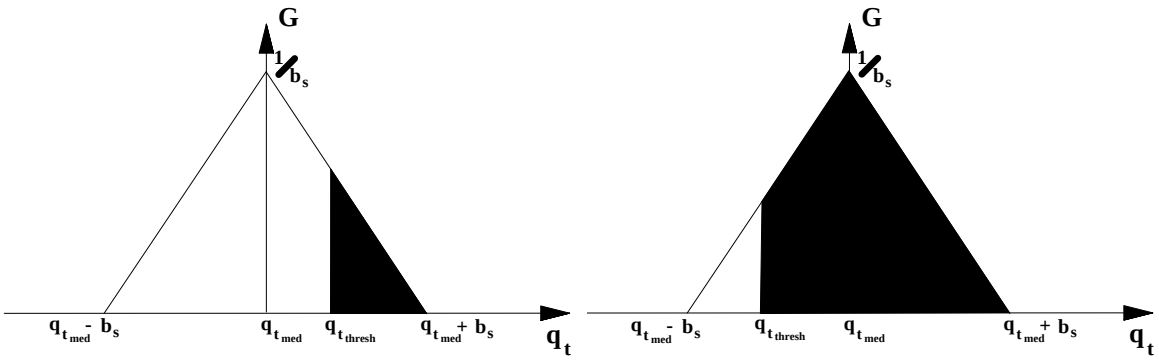

Figure 4 Two examples of the distribution of total water about the median. The shaded region denote states corresponding to a visibility below the prescribed threshold.

2. $q_t - b_s < q_{t\ thresh} \leq q_t$: in which case the fog fraction lies in the range 0.5→1.0; the situation is illustrated by Figure 4 (left panel). The fog fraction, FF, id given by:

$$FF = 1 - \frac{1}{2}\left(\frac{q_{t_{thresh}} - (q_t - b_s)}{b_s}\right)^2 \qquad \textbf{(21)}$$

3. $q_t < q_{t\ thresh} \leq q_t + b_s$: in which case the fog fraction lies in the range 0.0→0.5; the situation is illustrated by Figure 4 (right panel). The fog fraction, FF, id given by:

$$FF = \frac{1}{2}\left(\frac{q_t + b_s - q_{t_{thresh}}}{b_s}\right)^2 \qquad \textbf{(22)}$$

4. $q_{t\ thresh} > q_t + b_s$: in which no state in the distribution exceeds the total water threshold, and the fog fraction is 0.0

3. Verification

The operational verification scheme produces comparative statistics hourly against surface observations for the analysis, merged forecast, extrapolation forecast, persistence and the mesoscale model; these can be amalgamated to produce the following statistics for the period of interest:

- Area-based hit rate (HR), false alarm rate (FAR) and critical success index (CSI) for 1 km and 200 m thresholds for visibility.
- Site-specific HR, FAR and CSI for 1 km and 200 m thresholds for visibility.
- Mean factor error (MFE), root mean square factor error (RMSFE) and percentage within 30% or 200 m (P30) for the visibility.
- Mean error (ME), root mean square error (RMSE) and percentage within 2°C (P2) for the temperature and dewpoint.
- Site-specific HR, FAR and CSI for 1 km and 200 m thresholds for 'worst' visibility.

The statistics produced against the analysis for the various forecast components are:

- HR, FAR and CSI for 1 km and 200 m thresholds using visibility fields smoothed with a 25 km half-width.
- Site-specific HR, FAR and CSI for 1 km and 200 m thresholds for visibility.
- MFE, RMSFE and P30 for the visibility.
- ME, RMSE and P2 for the temperature and dewpoint.

For the PVT only, distributions of percentages observations less than the specified threshold corresponding to different fog probabilities are produced for the thresholds: 5000 m, 1000 m and 200 m.

As the quality of the analysis is one of the things in question, the verification results will concentrate on the comparison against observations.

4. The parallel trial

A parallel visibility trial (PVT) was considered a good way to test the new scheme, as it allowed direct comparison with the operational scheme over a large number of cases, with all the standard verification.

4.1. Details

The PVT was intended to run for about three weeks. In actual fact, it ran for just over nine weeks: it was started on the 11th February 1997 and was stopped on 15th April 1997 (to make way for testing of the cloud scheme). There were a number of reasons for this long testing period:

- The lack of fog in the early (and late) stages of the PVT.
- The usefulness of having a full near-real-time comparison as a test bed.
- The ability to make and test changes thoroughly and quickly.
- Because it was available.

Two changes were made to the new scheme which was originally being tested, so the PVT can be divided into three distinct regimes:

11/2/97 - 5/3/97	Median visibility with two distinct regimes for activated and non-activated fog droplets. Prognostic variables are derived from observations using approximate inverse of the visibility scheme.
1200 UTC 5/3/97 - 0800 UTC 14/3/97	Median visibility scheme changed to have a single regime including the activated and non-activated fog droplets. Prognostic variables are derived from observations using exact inverse of the visibility scheme. Fog probabilities at 5000 m, 1000 m and 200 m generated for the analysis and merged forecasts.
1200 UTC 14/3/97 - 0700 UTC 15/4/97	30% probability visibility replaces median visibility. Derivation of prognostic variables from observations changed to be consistent. Although no change made to the fog probabilities, the other changes have a knock-on effect.

The first period served to highlight problems in the use of a full scheme with two distinct regimes for activated and non-activated droplets; namely diagnosing the transition point and consistently carrying out the inverse procedure to extract the observational information. However, these problems mean that, beyond that, the period is of little use for verification.

The second period, although the shortest, was by far the foggiest, and as such makes a very interesting study period. However, day-to-day monitoring of the analysis in the PVT suggested that the use of the median (50% probability) visibility led to the underprediction of fog. Hence, the introduction of the 30% probability visibility for the final period.

The third and final period of the PVT had virtually no foggy cases in it, despite being the longest period (32 days), which has tended to negate the usefulness of the period, especially for the fog verification parameters.

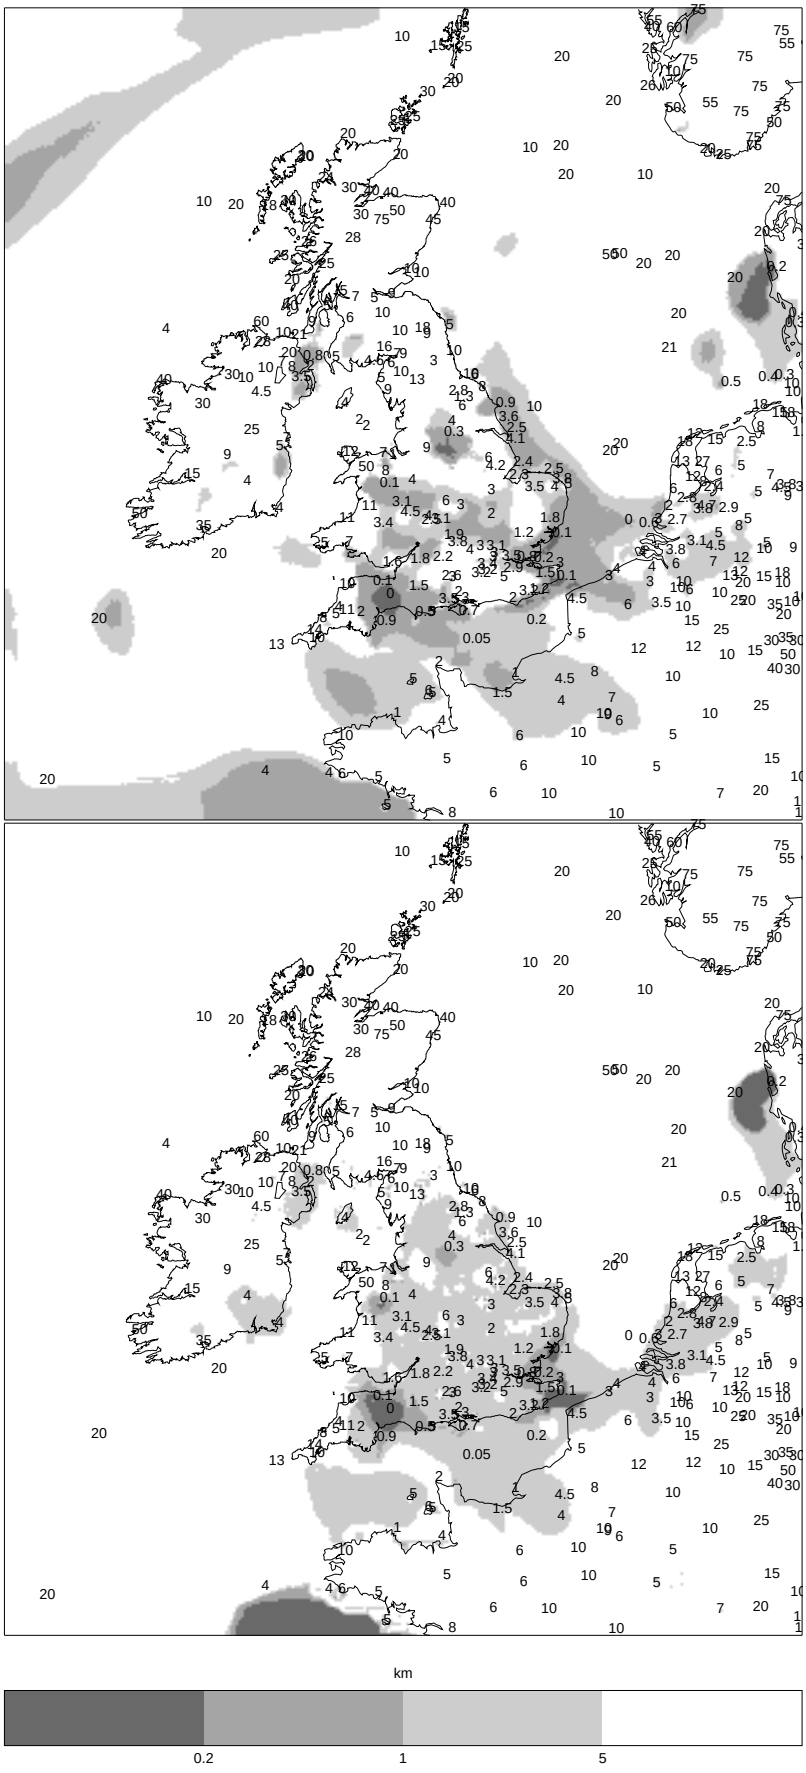

Figure 5 Visibility analysis for the current (top) and new (bottom) schemes with observations overlaid for 12 UTC 9/3/97.

4.2. Results

Despite the problems with PVT, it did provide some useful results. The assessment will tend to concentrate on the second period of the PVT, as this uses the significant changes to the scheme and includes a lot of foggy occasions.

The new satellite scheme significantly improved the analyses where satellite data were available. Figure 5 shows the analyses from both schemes for 12 UTC 9/3/97, which is a good example, as the UK is predominantly cloud-free, and so the satellite data are used over most of the land areas. The new scheme provides a much better areal fit to the observations, especially in the range 1–5 km. The current analysis has too much fog over southern England, where the visibilities were poor rather than foggy, and too high visibilities over much of France.

It is very difficult to isolate the impact of the aerosol limiting, although early tests with it in place helped remove the persistent, erroneous areas of thick fog over the continent, which are often produced by the current scheme.

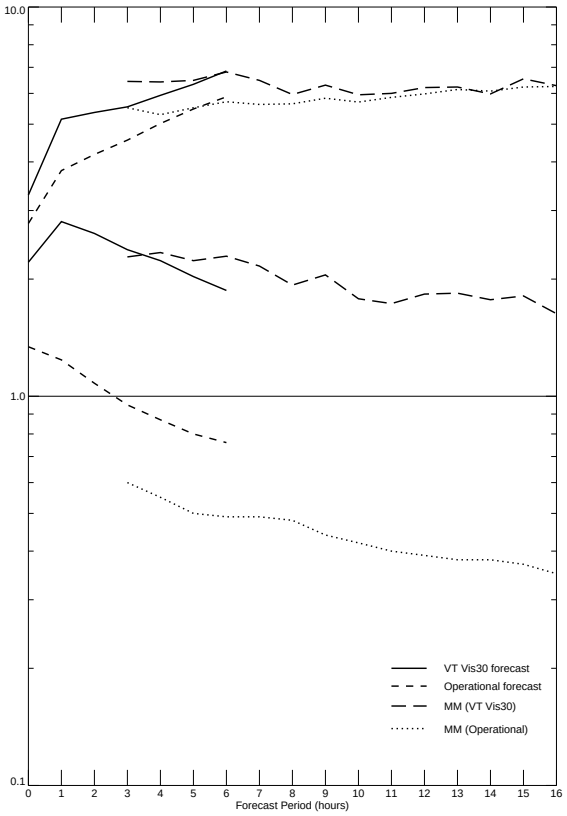

Figure 6 PVT visibility MFE and RMSFE for the MM and merged forecast for current scheme and Vis30 scheme.

The main change in the new scheme was, of course, the move to the use of a fixed probability visibility diagnostic (originally a median, then 30%), as apposed to a mean visibility. The results of this change were somewhat disappointing. The use of a median produced visibilities which were far too high, although it did provide improvements in the fitting of the thick fog (it should be noted, that an error discovered at a later stage did worsen this bias). Hence, the move to the 30% probability visibility (Vis30), which still appeared to show a significant bias in the visibilities, both in the forecast visibility and the diagnosed MM visibility (Figure 6), with the diagnosed MM visibility being particularly bad. The current scheme shows an opposite bias, producing visibilities which are too low. However, there was virtually no fog in this period, so the results may not be comprehensive. For this reason it was not possible to obtain any useful information from looking at HR, FAR and CSI at 1 km or 200 m.

5. The extended case study

The lack of fog in the later half of the PVT had made it impossible to fully compare the Vis30 scheme with the current scheme, so an extended visibility case study (EVCS) covering seven days in the foggy period in the early part of the PVT was setup for this purpose. This seemed to offer the ideal way to consider a large number of analyses and forecasts for verification, whilst ensuring a large number of fog events. It also allowed direct comparison of further modifications to the scheme.

5.1. The details

The EVCS was set up to run for seven days from 1200 UTC 9th March 1997 to 1200 UTC 16th March 1997, with a 6 hour ‘spin-up’ period before this. The system was setup to mimic the operational running as closely as possible, with MM data becoming available at T+180 minutes, and full verification. The main output was the verification files, although the hourly analyses were saved

for later reference and data to re-start the analysis-forecast cycle was saved every 12 hours. The full 174 hour run took about 30 hours elapsed time on the HP.

5.2. Results

The current operational scheme was run on the EVCS to act as a control. The 30% probability visibility scheme (Vis30) was run as a first comparison. The results were broadly similar to those for the PVT (see Figure 6), except the bias in the Vis30 run was slightly reduced. Comparison of the skill scores shows that at the 200 m threshold the Vis30 run performs better than the current scheme. At the 1 km threshold, the CSI is comparable for the site-specific assessment, but the area-based score is significantly worse, and the HR is worse for both. These results tend to confirm that using a fixed probability visibility does not give low enough visibilities at around 1 km, because of the MM bias in moisture as perceived by the new scheme.

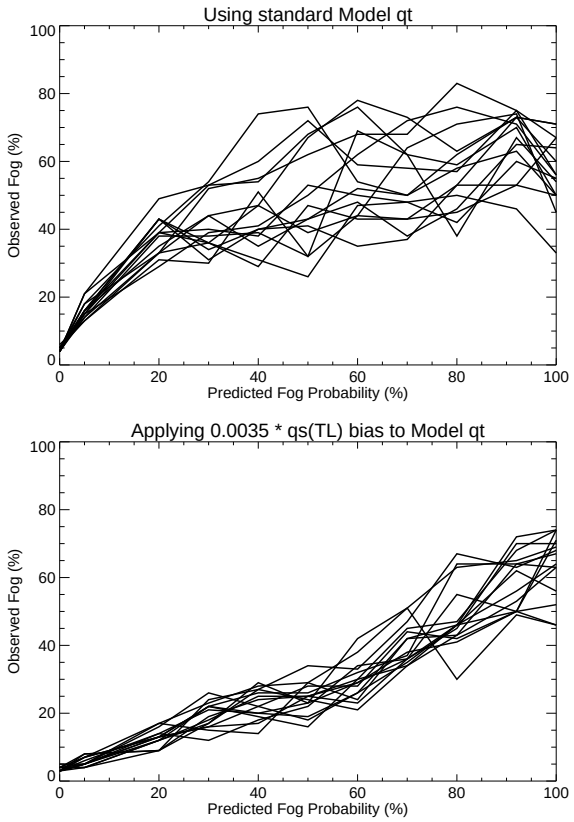

Figure 7 Percentage of foggy observations against MM fog probability for a 1 km visibility threshold for current (top) and Bias0035 (bottom) schemes.

Based on this result, it was decided to re-examine the median visibility (50% probability), after applying a correction to the MM total water content to remove the bias. The correction was applied as a scaling of the saturated specific humidity, so that the adjusted total water value, $q_{t\,adj}$, was given by:

$$q_{t_{adj}} = q_t + bias\, q_s(T_L) \qquad \textbf{(23)}$$

where q_t is the original MM total water content, $q_s(T_L)$ is the saturated specific humidity at the liquid water temperature and bias is a scaling factor. The MM visibility MFE was used to tune the value of bias; a value of 0.0035 appeared to give the best fit. In terms of the threshold statistics, the 0.0035 bias corrected median visibility (Bias0035) from the MM appeared to be marginally better then the current scheme overall, and significantly better then the MM Vis30 results. Figure 7 shows the percentage of observation with a visibility less than 1 km corresponding to the different diagnosed MM fog probabilities (the different lines represent different forecast times). An ideal solution would give a 45° line from 0% on the left to 100% on the right. Neither the current scheme (top) or the Bias0035 scheme (bottom) reach 100%, which indicates the failure of the MM to predict all the foggy events. However, the Bias0035 scheme produces a better grouped set of lines which a much more linear variation; this implies that the correction to the MM total water is improving the distribution of visibility states. Similar results are obtained for thresholds of 200 m and 5 km.

A full run of the Bias0035 scheme was carried out. Figure 8 shows the site-specific and area-based CSI scores for 1 km and 200 m threshold visibilities (the crosses mark the target scores at different times). In general, the Bias0035 scheme shows an improvement over the current scheme. For the site-specific verification the improvement is small, but for the area-based verification it is quite marked (except at T+6 hours for the 1 km threshold). Figure 9 shows the visibility RMSFE and MFE scores. The RMSFE for the forecasts are very similar, but the Bias0035 scheme probably has a better MFE, although still too high at all but T+6 hours. Figure 10 shows an example of a typical set of analyses and forecasts from the two schemes, valid at 0900 UTC 10th March 1997. The Bias0035 analysis

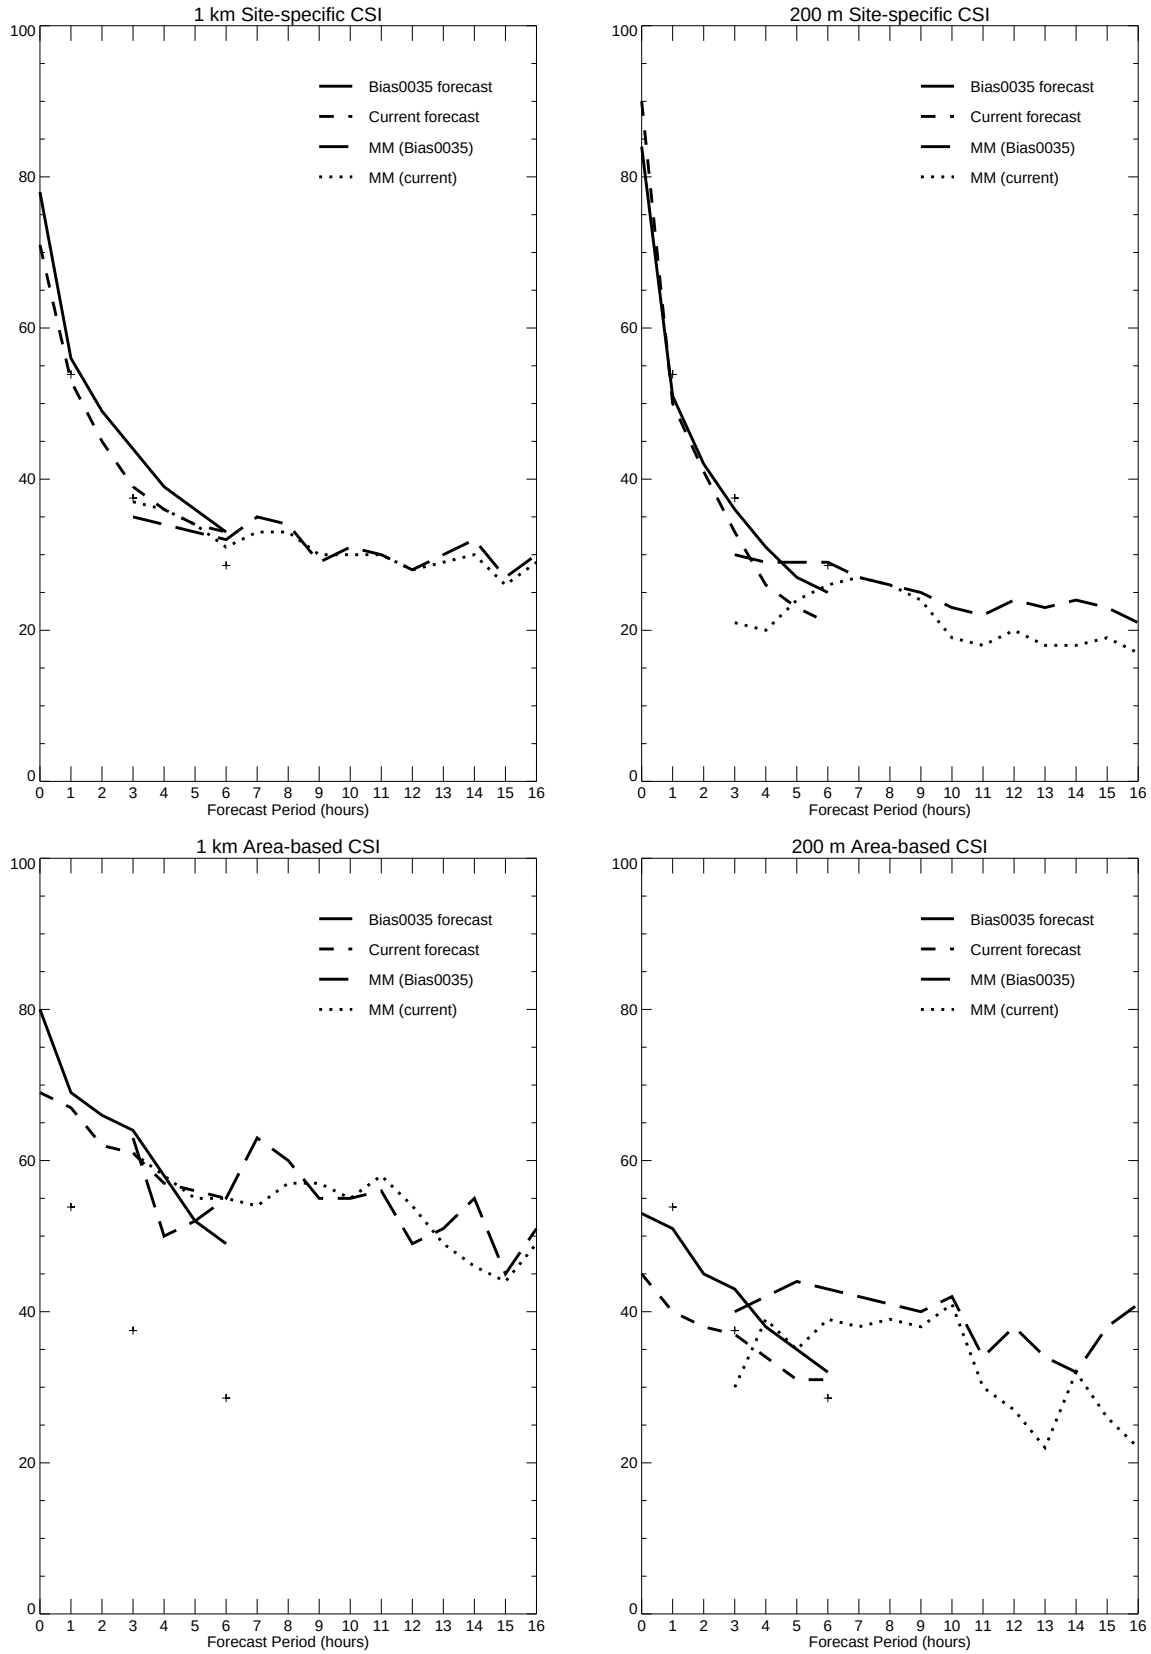

Figure 8 EVCS Visibility CSI scores for the MM and merged forecast for the current and Bias0035 scheme (right=1 km, left=200 m, top=site-specific and bottom=area-based).

(bottom right) is much better than the current analysis (bottom left), as it does not have the spurious fog over southern Scotland; this, of course, is primarily due to the use of the new satellite scheme. The

Bias0035 T+1 forecast (satellite data used in the analysis) also lacks the spurious fog present over southern Scotland in the current T+1 forecast, but it has too little fog in northeastern England. At T+3, the Bias0035 forecast has slightly less fog, but the forecasts are broadly similar. Overall, it is difficult to choose between the two schemes.

6. Conclusions

The improved satellite scheme provides significantly better analyses where satellite data are available. However, it is probably wrong to consider this as solely the impact of the satellite data, as it is the areal usage of the surface observations, which provides much of the benefit; the satellite data are simply providing an appropriate way of spreading the observational information. Hence, further benefits could potentially be gained by using the surface observations in a similar way when no satellite data are available. However, this then begs the questions: do we need to use the surface observations independently at all? Is a variational scheme the most appropriate method of combining the data?

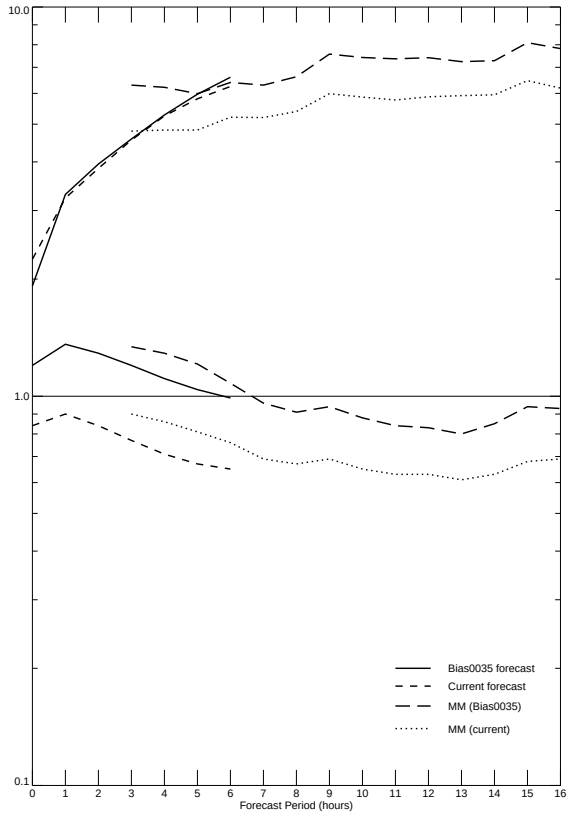

Figure 9 EVCS visibility MFE and RMSFE for the MM and merged forecast for the current and Bias0035 schemes.

The impact of the limiting of the aerosol concentrations is harder to judge independently of the modifications to the diagnosis of the visibility. However, it looks as though the extensive areas of erroneous fog over the continent, which usually coincide with high aerosol concentrations, are absent when the limiting is used.

The most significant change, using a fixed probability based total water value rather than the integral over a distribution of states for the diagnosis of the visibility, appeared to have a small positive impact, if an adjustment is applied to the MM total water to remove the bias perceived by the new scheme. Using an adjustment of $0.035 \times q_s(T_L)$ for the MM q_t value allows a median visibility (50% probability) to be used successfully. Although the improvements are small overall, the distribution of states, as displayed through the fog probability statistics, appears to be better for both the MM and the forecast. Subjective verification confirms that visibilities are slightly higher in the new scheme, but there is little overall difference in quality. As mentioned earlier the analyses which use satellite data are significantly better.

The overall conclusion, is that the move to a median visibility, produces a more satisfactory distribution of visibility states, provided a correction is applied to the MM total water content, but the performance of the scheme is only marginally better then the current scheme. The use of the satellite scheme described here to spread all observations, regardless of whether or not satellite data are available, could produce further improvements to the analyses, and might allow the removal of the variational scheme, which would give a substantial saving in computer time. A scheme for treating hill fog has been developed (not used here) and might provide better performance in both the analysis and the forecast. Further improvements to the forecast are desirable, but would require a re-assessment of the extrapolation forecast technique and the merging technique.

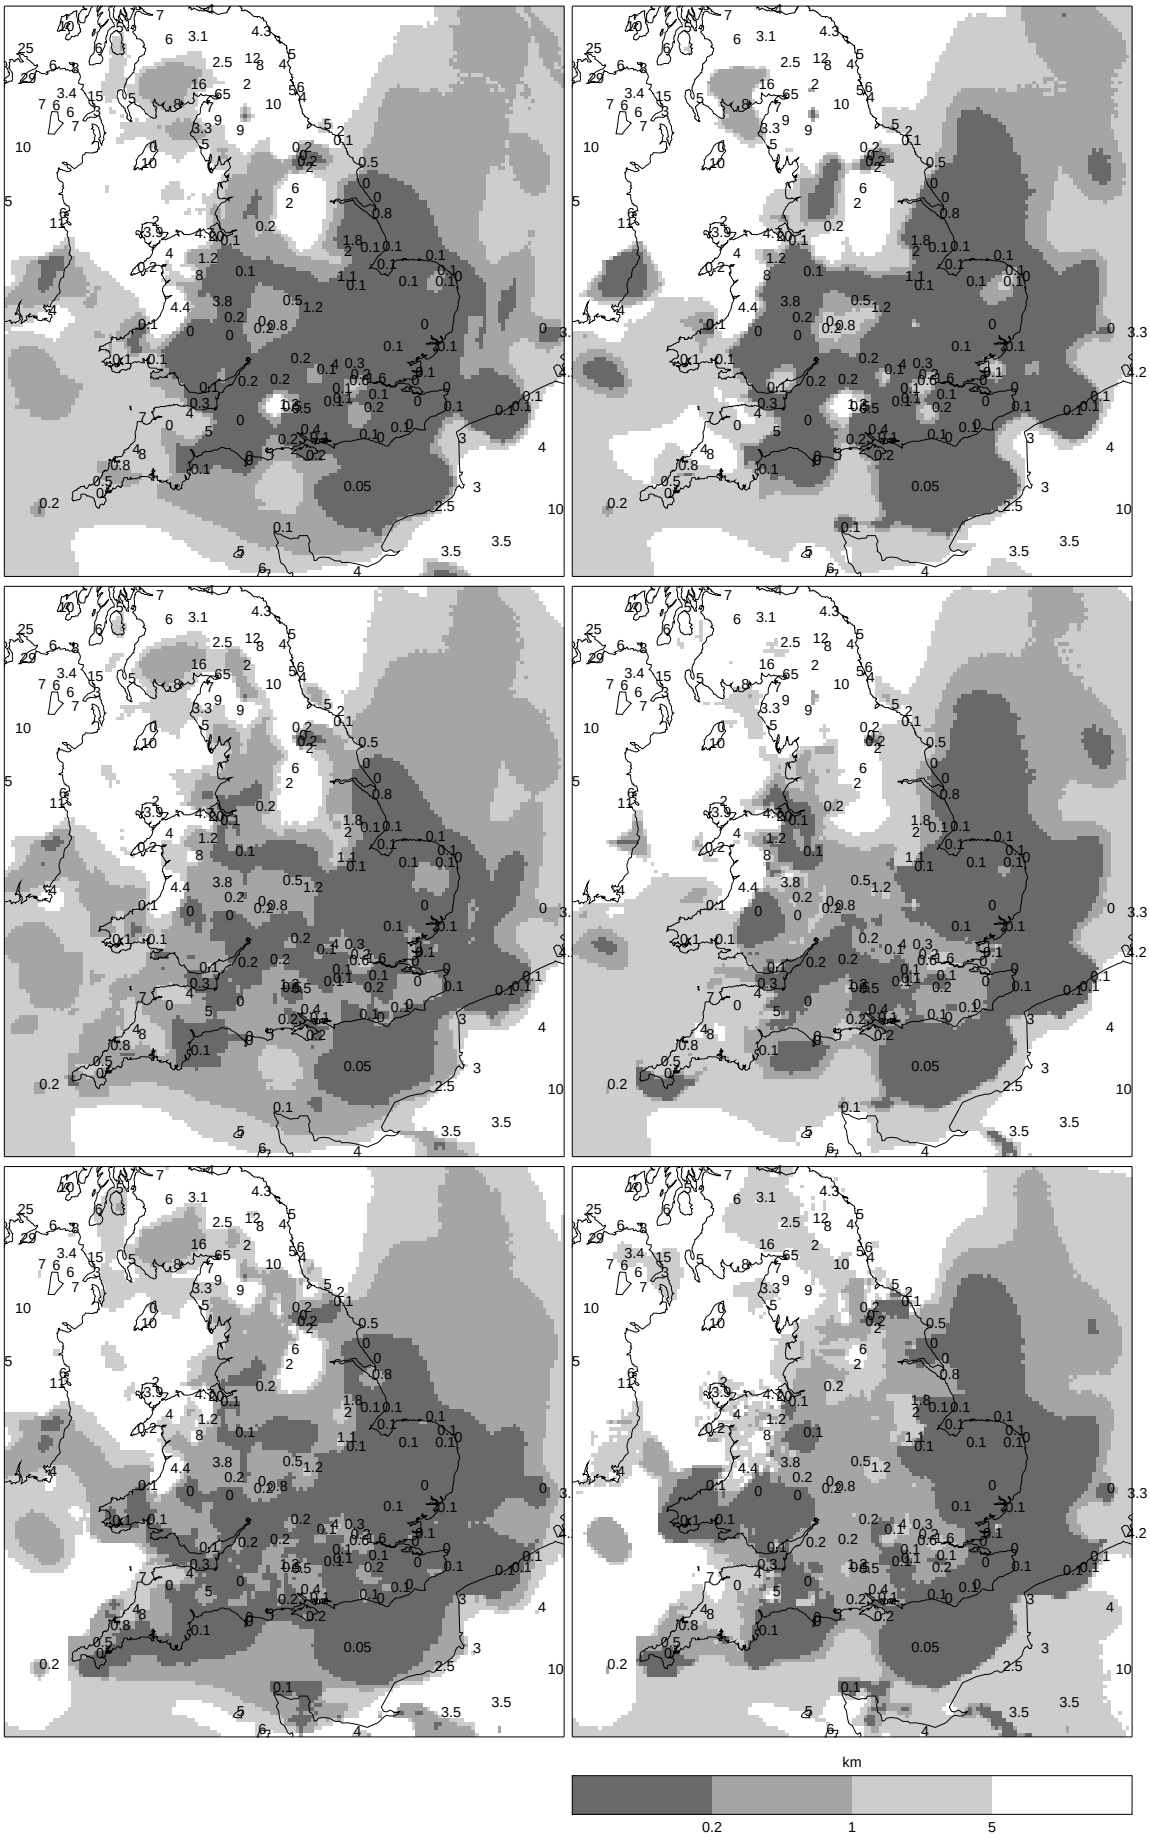

Figure 10 Comparison of T+3 (top), T+1 (middle) and analysis (bottom) visibilities for current (left) and Bias0035 (right) schemes for 0900 UTC 10/3/97.

Appendix A. Solution for fog droplet radius

The fog droplet radius, r, is calculated from the total water content, q_t, by solving for r_* in the equation:

$$q_t(r_*) = RH(r_*)\,q_s(T) + q_L(r_*) \qquad (24)$$

where:

$$RH(r_*) = \exp\left(\frac{A}{r_*} - \frac{B_o}{r_*^3 - 1}\right) \qquad (25)$$

and:

$$q_L(r_*) = \tfrac{4}{3}\pi r_d^3 (r_*^3 - 1)\,\rho_w N \qquad (26)$$

where:

$$A = \frac{A_o}{r_d} \qquad and \qquad r_* = \frac{r}{r_d} \qquad (27)$$

So, and $q_s(T)$ is calculated as a function of T, which is given and all the other constants are defined in the text. This can not be solved analytically, so the solution is found using Newton-Raphson, i.e. a function $f(r_*)$ is defined as:

$$f(r_*) = q_t - RH(r_*)\,q_s(T) - q_L(r_*) \qquad (28)$$

and it is this which is solved iteratively using the Newton-Raphson approximation to obtain successively better predictions for r_*:

$$r_{*_{n-1}} = r_{*_n} - \frac{f(r_{*_n})}{f'(r_{*_n})} \qquad (29)$$

The derivation of the function, $f(r_*)$, is:

$$f'(r_*) = -q_s(T)\,RH'(r_*) - q_L'(r_*) \qquad (30)$$

where relative humidity derivatives is given by:

$$RH'(r_*) = \left(-\frac{A}{r_*^2} + \frac{3B_o r_*^2}{\left(r_*^3 - 1\right)^2}\right) RH(r_*) \qquad (31)$$

and the liquid water derivative is given by:

$$q_L'(r_*) = 4\pi\rho_w N r_d^3 r_*^3 \qquad (32)$$

A first-guess value for r_* is calculated from the relative humidity, after it has been limited to lie in the range 0.001 to 0.999, using equation (12), in there-arranged form:

$$r_* = \left(1 - \frac{B_o}{\log(RH)}\right)^{\frac{1}{3}} \qquad (33)$$

The Newton-Raphson solution is iterated until the change in r_* is less than 0.001 or 20 iterations have been carried out. Equation (27)b is used to recover the droplet radius. To ensure convergence, the values of r_* are limited to lie in the range 1.0 to 1000.0 after each iteration, and the maximum change in the value of r_* allowed by one iteration is a factor of 4.

References

Golding, B. W. (1995), 'An integrated system for very short range forecasting'. *Proc. Int. Conf. Meteorol. and Hydrol. Technology and Its Management,* Meteohytec 21, Geneva, **WMO/TD No. 672,** 114–117.

Koschmeider, H. (1924), 'Therie der horizontalen sichtweite'. *Beitr. Phys. Atmosph.*, **12**, 33–53, 171–181.

Pruppacher, H. R. and Klett, J. D. (1978), *Microphysics of clouds and precipitation*, D. Reidel Publishing Company, Dordrecht, Holland.

Wright, B. J. (1995), 'Proposals for the stage 2 development of the Nimrod visibility/fog analysis'. M*et. Office NDG Work. Pap.*, **No. 98**.

Wright, B. J. (1996a), 'The stage three development of the Nimrod visibility analysis'. M*et. Office NDG Work. Pap.*, **No. 107**.

Wright, B. J. (1996b), 'The development and testing of the stage two Nimrod visibility analysis'. M*et. Office NDG Work. Pap.*, **No. 108**.

Wright, B. J. and Thomas, N. (1995), 'Proposals for a visibility/fog analysis and forecast scheme for Nimrod'. M*et. Office NDG Work. Pap.*, **No. 93**.

Wright, B. J. and Thomas, N. (1996), 'An objective visibility and very-short-range forecasting system'. M*et. Office FR Tech. Rep..*, **No. 200**.