

A Variational Scheme for Retrieving Rainfall Rate and Hail Reflectivity Fraction from Polarization Radar

ROBIN J. HOGAN

Department of Meteorology, University of Reading, Reading, United Kingdom

(Manuscript received 30 June 2006, in final form 17 January 2007)

ABSTRACT

Polarization radar offers the promise of much more accurate rainfall-rate R estimates than are possible from radar reflectivity factor Z alone, not only by better characterization of the drop size distribution, but also by more reliable correction for attenuation and the identification of hail. However, practical attempts to implement retrieval algorithms have been hampered by the difficulty in coping with the inherent noise in the polarization parameters. In this paper, a variational retrieval scheme is described that overcomes these problems by employing a forward model for differential reflectivity Z_{dr} and differential phase shift ϕ_{dp} and iteratively refining the coefficient a in the relationship $Z = aR^b$ such that the difference between the forward model and the measurements is minimized in a least squares sense. Two methods are used to ensure that a varies smoothly in both range and azimuth. In range, a is represented by a set of cubic-spline basis functions; in azimuth, the retrieval at one ray is used as a constraint on the next. The result of this smoothing is that the retrieval is tolerant of random errors in Z_{dr} of up to 1 dB and in ϕ_{dp} of up to 5°. Correction for attenuation is achieved simply and effectively by including its effects in the forward model. If hail is present then the forward model is unable to match the observations of Z_{dr} and ϕ_{dp} simultaneously. This enables a first pass of the retrieval to be used to identify the radar pixels that contain hail, followed by a second pass in which the fraction of the Z in those gates that is due to hail is retrieved, this time with the scheme being able to forward-model both Z_{dr} and ϕ_{dp} accurately. The scheme is tested on S-band radar data from southern England in cases of rain, spherical hail, oblate hail, and mixtures of rain and hail. It is found to be robust and stable, even in the presence of differential phase shift on backscatter.

1. Introduction

Accurate rainfall estimates are essential for short-range flood forecasts, but with a conventional single-polarization radar one must use a simple relationship of the form

$$Z_h = aR^b, \quad (1)$$

where Z_h is radar reflectivity factor at horizontal polarization, R is rainfall rate, and a and b are constants that must be assumed. Such relationships have been found to be accurate to no better than a factor of 2 because of inherent variations in the size distribution. Conventional radars are also unable to distinguish hail from heavy rain and are very difficult to correct for attenu-

ation because gate-by-gate methods are inherently unstable (Hitschfeld and Bordan 1954).

In principle, dual-polarization radar can overcome all these problems. The crucial size-distribution information is obtained by exploiting the fact that raindrops become increasingly oblate with size and fall with their major axes aligned in the horizontal plane. Differential reflectivity (expressed in decibels) is defined as the ratio of reflectivity factor measured ($\text{mm}^6 \text{m}^{-3}$) at horizontal polarization Z_h and vertical polarization Z_v :

$$Z_{dr} = 10 \log_{10}(Z_h/Z_v). \quad (2)$$

It increases with mean drop size (Seliga and Bringi 1976) and hence can be used to refine rain-rate estimates. Differential phase shift ϕ_{dp} is a measure of the lag of the phase of the horizontally polarized radiation with respect to the vertically polarized radiation in the presence of oblate drops. Because it is a propagation effect, the usual approach is to make use of its gradient with range, known as the (one way) specific differential phase shift $K_{dp} = 0.5d\phi_{dp}/dr$. Here Z_{dr} and K_{dp} provide

Corresponding author address: Robin J. Hogan, Department of Meteorology, Earley Gate, P.O. Box 243, Reading RG6 6BB, United Kingdom.

E-mail: r.j.hogan@reading.ac.uk

similar information on the raindrop size distribution, and the potential for hail detection arises because when hail is present the two sources of information are contradictory (Smyth et al. 1999). Attenuation correction is possible by making use of ϕ_{dp} (e.g., Testud et al. 2000), or differential attenuation leading to negative Z_{dr} at the far side of heavy rain (Smyth and Illingworth 1998a).

In practice, the inherent noise in Z_{dr} and especially K_{dp} presents problems when these techniques are implemented operationally (Illingworth 2003). Because K_{dp} is the derivative of an already noisy field, negative values occur (Ryzhkov and Zrnić 1996) and it is impossible to use at each radar range gate; instead one must use the average value over several kilometers (Blackman and Illingworth 1997). A similar approach is necessary with Z_{dr} (Illingworth and Thompson 2005). Such methods can result in sharp changes in retrieved rain rate at the interface between averaging regions. Furthermore, it is difficult to design conventional algorithms to make use of Z_{dr} and ϕ_{dp} simultaneously in all situations, and therefore operationally one must usually choose the most appropriate algorithm depending on the conditions [e.g., using the criteria of Ryzhkov et al. (2005)].

In this paper, a variational method is applied for the first time to the problem of retrieving rainfall rate from polarization radar. This approach (also known as “optimal estimation theory”; Rodgers 2000) is ubiquitous in satellite retrievals but has only recently been applied to radar (e.g., Austin and Stephens 2001; Löhnert et al. 2004). By an explicit treatment of errors, a seamless transition can be achieved between each of the various rain/hail/attenuation regimes that previously one would have to have treated by a separate algorithm. Moreover, it is straightforward to include attenuation in the forward model without the instabilities discussed by Hitschfeld and Bordan (1954), and by forward modeling ϕ_{dp} , we avoid the error associated with taking the derivative to get K_{dp} . The various regimes that the new scheme is designed to approximate are now summarized.

- 1) In very light rain, the drops are spherical so both Z_{dr} and K_{dp} are zero. In this case the polarization variables contain no useful information and one must assume values of a and b in (1) that are suitable for light rain. This is achieved by the use of an a priori.
- 2) In light to moderate rain, Z_{dr} increases above 0 dB while K_{dp} remains very close to 0° km^{-1} . Here Z_{dr} provides information on a .
- 3) As the rain becomes heavier, K_{dp} increases above 0° km^{-1} and also provides information on a , for example, as a path average (Testud et al. 2000). Unlike

other techniques, the new variational method then uses the known errors in Z_{dr} and ϕ_{dp} to weight the information provided by these two variables appropriately.

- 4) When attenuation becomes significant, information to correct for it is first obtained from ϕ_{dp} (Holt 1988; Bringi et al. 1990; Testud et al. 2000).
- 5) Stronger attenuation results also in differential attenuation, whereby Z_h is attenuated more strongly than Z_v such that Z_{dr} becomes negative at the far side of a region of heavy rain. It was argued by Smyth and Illingworth (1998a) and Illingworth (2003) that the total attenuation is more closely related to differential attenuation than to ϕ_{dp} so that in principle more accurate attenuation correction can be performed if this negative Z_{dr} region is well sampled (Smyth and Illingworth 1998a). However, the relationship is likely to break down in the presence of wet hail.
- 6) Hailstones are usually close to spherical and tumble as they fall, and so their intrinsic Z_{dr} and K_{dp} are close to zero. Horizontally aligned, oblate hail is also known to occur (e.g., Smyth et al. 1999) but has a considerably lower Z_{dr} than rain of the same reflectivity and still has a K_{dp} of close to zero. Hence the combination of Z , Z_{dr} , and K_{dp} can be used to identify the presence of hail. When rain and hail coexist, the rain rate can be estimated from K_{dp} alone (e.g., Sachidananda and Zrnić 1987; Ryzhkov and Zrnić 1996). The new variational method also retrieves the fraction of the reflectivity that is due to hail.

In section 2 the various steps of the retrieval scheme are described, followed by the forward model in section 3. Then in section 4, the performance of the scheme on S-band data taken by the Chilbolton radar in southern England during three case studies is described.

2. Retrieval method

a. Overview

The sequence of operations performed in the retrieval is illustrated in Fig. 1. It is assumed that the measurements have first been processed to subtract instrument noise, to remove clutter, and to identify the pixels where a valid meteorological signal has been detected. This is best achieved using the polarization parameters (e.g., Gourley et al. 2007). Pixels within and above the melting layer should be identified and removed, using, for example, the temperature field from an operational forecast model (Kitchen et al. 1994; Mittermaier and Illingworth 2003). Although the scheme remains stable when applied to ice, the retrieved surface rain rate will obviously be incorrect because of the

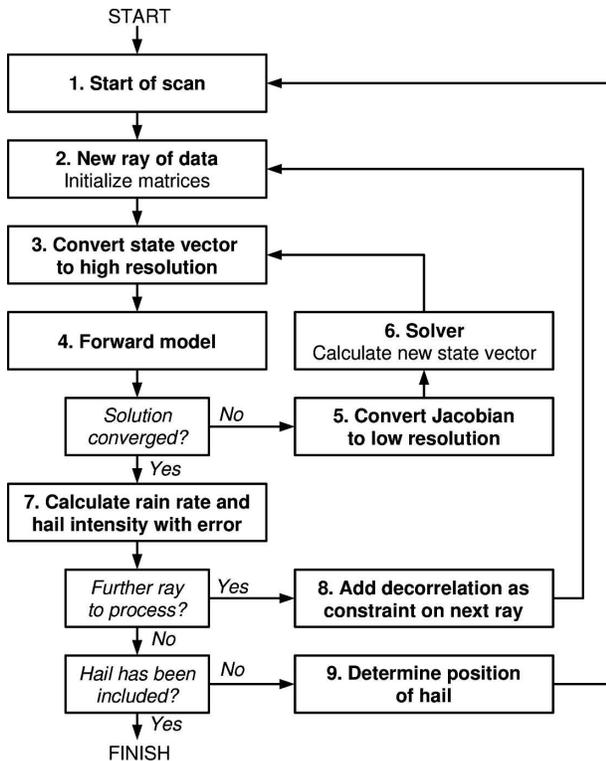


FIG. 1. Flowchart that shows the sequence of operations performed by the retrieval scheme. The detailed description of each step is given in section 2.

inappropriate assumptions regarding the nature of the scatterers. In this situation, an alternative method to estimate surface rain rate should be used, such as one based on the vertical profile of reflectivity (Kitchen et al. 1994; Smyth and Illingworth 1998b).

First, the processing of a single ray of data containing m pixels with a detectable meteorological signal is described. Consideration of the full scan is given in section 2e. In a variational scheme one must decide what variables to use to describe the rain; these variables will be retrieved and constitute the state vector \mathbf{x} . An obvious choice would be rain rate R and some measure of the mean drop size at each radar pixel along a ray (hail is considered in section 2g). The observations at each pixel along the ray (e.g., Z_h , Z_{dr} , and ϕ_{dp}) are designated as the observation vector \mathbf{y} . The forward model $H(\mathbf{x})$ would then use a first guess of \mathbf{x} to predict the observations at each gate, with the difference between $H(\mathbf{x})$ and \mathbf{y} being used to refine \mathbf{x} such that a better fit with the observations is achieved, in a least squares sense. This process would be repeated until convergence.

However, for computational efficiency it is desirable to minimize the number of elements in \mathbf{x} and \mathbf{y} . Fur-

thermore, the noise in the Z_{dr} and ϕ_{dp} measurements is such that it is not realistic to retrieve rain rate independently at each gate; rather, the information from the polarization variables must be spatially smoothed in some way (e.g., Illingworth and Thompson 2005). We therefore choose to retrieve the coefficient a in (1) and represent it by a reduced set of n basis functions (typically $n \approx m/10$) such that smooth variation in range is guaranteed. If the resolution at which a can vary is still reasonably high then the value of b can be specified; a value of 1.5 or 1.6 is typically used. A simplified version of the scheme has been written in which both a and b were allowed to vary (Furness 2005), but the differences in retrieved rain rate were only around 5%. It should be stressed that a smoothly varying a field does not mean that retrieved rain rate is forced to be smooth, because the fine structure in the Z_h field will feed through directly to R .

It was shown by Bringi and Chandrasekar (2001) that if the size distribution is represented by a gamma distribution with a fixed value for the normalized number concentration parameter N_w , then this results in $b = 1.5$, with the a coefficient being inversely proportional to the square root of N_w . Therefore, the assumption that a varies smoothly with range is similar to the assumption by Testud et al. (2000) that normalized number concentration parameter is constant or varies smoothly with range.

The state vector for a single ray is therefore

$$\mathbf{x} = \begin{pmatrix} \ln a_1 \\ \vdots \\ \ln a_n \end{pmatrix}. \quad (3)$$

By using the logarithm of a , we avoid the unphysical possibility of retrieving negative a . It also turns out that the forward model for Z_{dr} and ϕ_{dp} is less nonlinear when formulated in terms of $\ln a$, resulting in more rapid convergence. After the scheme has converged, rain rate is calculated at high resolution from Z_h using (1). We are effectively assuming that, in relative terms, the error in Z_h is much less than the errors in Z_{dr} and ϕ_{dp} so that the retrieval should be forced to be exactly consistent with Z_h , and hence Z_h can be omitted from the observation vector. Thus \mathbf{y} is given by

$$\mathbf{y} = \begin{pmatrix} Z_{dr,1} \\ \vdots \\ Z_{dr,m} \\ \phi_{dp,1} \\ \vdots \\ \phi_{dp,m} \end{pmatrix}. \quad (4)$$

b. Variational formulation

The essence of the technique is to minimize a cost function J :

$$2J = \sum_{i=1}^m \frac{(Z_{\text{dr},i} - Z'_{\text{dr},i})^2}{\sigma_{Z_{\text{dr}}}^2} + \frac{(\phi_{\text{dp},i} - \phi'_{\text{dp},i})^2}{\sigma_{\phi_{\text{dp}}}^2} + \sum_{i=1}^n \frac{(x_i - x_i^a)^2}{\sigma_{x^a}^2}. \quad (5)$$

The first summation in (5) represents the deviation of the observations Z_{dr} and ϕ_{dp} from the values predicted by the forward model Z'_{dr} and ϕ'_{dp} . The forward model is described in section 3.

The second summation represents the deviation of the elements of the state vector from some a priori estimate x^a (referred to as the ‘‘background’’ in data assimilation). This term is necessary to ensure that, in the presence of very light rain when very little information on a is contained in the measurements, a tends to some predefined value, typically $200 \text{ mm}^6 \text{ m}^{-3} (\text{mm h}^{-1})^{-1.5}$ for $b = 1.5$. The terms $\sigma_{Z_{\text{dr}}}$ and $\sigma_{\phi_{\text{dp}}}$ are the root-mean-square observational errors, and σ_{x^a} is the error in the a priori estimate.

In matrix notation, the cost function may be written as

$$2J = \delta\mathbf{y}^T \mathbf{R}^{-1} \delta\mathbf{y} + (\mathbf{x} - \mathbf{x}^a)^T \mathbf{B}^{-1} (\mathbf{x} - \mathbf{x}^a), \quad (6)$$

where $\delta\mathbf{y} = \mathbf{y} - H(\mathbf{x})$ and \mathbf{R} and \mathbf{B} are the error covariance matrices of the observations and the a priori, respectively. In this application we assume that \mathbf{R} is diagonal; that is, that the errors in the observations are not spatially correlated. By contrast, the off-diagonal components of \mathbf{B} [not represented in (5)] play an important role in smoothing the retrieval in range, as described in section 2d.

The cost function cannot be minimized in one step because of the presence of the nonlinear forward-model operator $H(\mathbf{x})$, and so we employ the Gauss–Newton method (Rodgers 2000) in which a linearized version of the cost function is minimized iteratively. At iteration k we have an estimate of the state vector \mathbf{x}_k and the corresponding forward-model estimate of the observations $H(\mathbf{x}_k)$. The linearized cost function J_L is obtained by replacing $H(\mathbf{x})$ in (6) by $H(\mathbf{x}_k) + \mathbf{H} \times (\mathbf{x} - \mathbf{x}_k)$, where \mathbf{H} is the Jacobian, a matrix containing the partial derivative of each observation with respect to each element of the state vector. In this case it is a $2m \times n$ matrix given by

$$\mathbf{H} = \begin{pmatrix} \partial Z'_{\text{dr},1}/\partial \ln a_1 & \cdots & \partial Z'_{\text{dr},1}/\partial \ln a_n \\ \vdots & \ddots & \vdots \\ \partial Z'_{\text{dr},m}/\partial \ln a_1 & \cdots & \partial Z'_{\text{dr},m}/\partial \ln a_n \\ \partial \phi'_{\text{dp},1}/\partial \ln a_1 & \cdots & \partial \phi'_{\text{dp},1}/\partial \ln a_n \\ \vdots & \ddots & \vdots \\ \partial \phi'_{\text{dp},m}/\partial \ln a_1 & \cdots & \partial \phi'_{\text{dp},m}/\partial \ln a_n \end{pmatrix}, \quad (7)$$

and is calculated at the same time as the forward model, as described in section 3c. By setting the derivative of J_L with respect to each element of \mathbf{x} to zero and rearranging, an expression for the state vector at the minimum of J_L is obtained:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{A}^{-1} [\mathbf{H}^T \mathbf{R}^{-1} \delta\mathbf{y} - \mathbf{B}^{-1} (\mathbf{x}_k - \mathbf{x}^a)], \quad (8)$$

where the symmetric matrix \mathbf{A} is known as the Hessian and is given by

$$\mathbf{A} = \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} + \mathbf{B}^{-1}. \quad (9)$$

For efficiency, \mathbf{A} is not inverted but rather is kept on the left-hand side of (8) and the matrix problem is solved by Cholesky decomposition. This step is represented in Fig. 1 by box 6.

For the first iteration the a priori may be used for \mathbf{x}_1 , or the values from the previous ray or a previous radar scan may alternatively be used. The entire loop (boxes 3–6 in Fig. 1) is repeated iteratively (with the forward model and \mathbf{H} recalculated each time), until the solution is judged to have converged satisfactorily. A χ^2 convergence test may be used, although in practice the problem is close enough to linear that only around four iterations are needed.

c. Use of cubic-spline basis functions for smoothing in range

The forward model described in section 3 works on the radar range grid; therefore, within the iterative loop, the state vector \mathbf{x} containing the n coefficients of the basis functions must first be converted to m values of $\ln a$ on the radar grid: $\hat{\mathbf{x}}$. This is achieved using an $m \times n$ matrix \mathbf{W} containing the precomputed weights of each basis function:

$$\hat{\mathbf{x}} = \mathbf{W}\mathbf{x}. \quad (10)$$

This operation is shown by box 3 in Fig. 1. In principle, any set of basis functions may be used, such as a reduced set of Fourier modes. However, the use of local basis functions is preferable because it allows smoothing to be implemented in the azimuthal direction (see section 2e), and also the fact that \mathbf{W} is then sparse

makes the scheme more efficient. We use cubic-spline basis functions, which result in the retrieved $\ln a$ being continuous in itself and its first and second derivatives. The calculation of \mathbf{W} is described in the appendix.

The Jacobian that is output from the forward model $\hat{\mathbf{H}}$ is also on the radar range grid; that is, it consists of the partial derivative of each observation with respect to the value of $\ln a$ at each radar gate, rather than with respect to each basis-function coefficient as in (7). It is converted to a form that can be used in (8) and (9) by

$$\mathbf{H} = \hat{\mathbf{H}}\mathbf{W}. \quad (11)$$

This step is depicted by box 5 in Fig. 1. It should be noted that as one increases the number of basis functions n then not only does this slow down the solution of (8) but also the calculation of (11), which is generally the slowest part of the entire scheme (even using code optimized to exploit the sparseness of \mathbf{W}). It is still much faster than not using basis functions at all and having one member of the state vector for every radar pixel.

d. Use of a priori error covariances for smoothing in range

The use of a cubic spline enforces smoothness on local scales but does not spread information beyond the domain of an individual basis function. In practice it might be desirable to spread information on the value of $\ln a$ obtained from Z_{dr} or ϕ_{dp} into adjacent regions of lower rainfall where the observations contain little information on $\ln a$. In the case of a diagonal a priori error covariance matrix \mathbf{B} , the value of $\ln a$ in these low-rainfall regions would revert rapidly to the a priori value. Off-diagonal elements in \mathbf{B} state that the difference between the actual value of $\ln a$ and the a priori value is spatially correlated and hence has the effect of a smoothing in range.

We represent the diagonal elements of \mathbf{B} by a constant value $B_{i,i} = 1.0$, which is equivalent to saying that the error in rain rate retrieved by assuming constant a and b in a Z - R relationship is approximately a factor of 2. If it is assumed that the correlation coefficient between two basis-function coefficients centered at ranges r_i and r_j decreases as an inverse exponential with the

separation distance, then the off-diagonal elements of \mathbf{B} are given by

$$B_{i,j} = B_{i,i} \exp(-|r_j - r_i|/r_0), \quad (12)$$

where r_0 is the decorrelation distance. Hence \mathbf{B} is a symmetric Toeplitz matrix. The problem of deciding what value to use for r_0 is common to many applications in data assimilation (Daley 1991), and in the case of radar the choice will depend on the error in the observations.

It should be stressed that, unlike the use of splines in section 2c, the use of error covariances does not *force* the solution to be smooth. Rather, it adds a *tendency* for smoothness, which could be overridden by accurate observations if they showed the true solution not to be smooth.

e. Smoothing in azimuth

In principle, smoothing in azimuth could be achieved by treating the scan in its entirety and using two-dimensional basis functions to ensure smoothness in both range and azimuth. In practice, the size of the matrices would be unmanageably large; therefore, a different method is used in the azimuthal direction. Two passes of the algorithm are performed through the rays in a scan.

In the forward pass, the retrieval at ray j is constrained by the solution from the previous ray, $j - 1$. This is followed by a backward pass in which the retrieval at ray j is constrained by both the forward-pass solution at ray $j - 1$ and the backward-pass solution at ray $j + 1$. This process is exactly analogous to a Kalman smoother, except that we are performing the smoothing in azimuth rather than in time.

In the forward pass the solution at ray $j - 1$ is denoted as \mathbf{x}_{j-1} and its error covariance is denoted as \mathbf{S}_{j-1} . Note that \mathbf{S}_{j-1} is simply the inverse of the Hessian at the final iteration, \mathbf{A}_{j-1}^{-1} . Before it can be used as a constraint on the next ray, the additional error due to spatial decorrelation between rays $j - 1$ and j must be considered. This we represent by an error covariance matrix $\mathbf{D}_{j-1,j}$ that is added to \mathbf{S}_{j-1} (represented by box 8 in Fig. 1). The constraint is added to the retrieval at ray j (except for the first ray) simply by adding an extra term to (8) to obtain

$$\mathbf{x}_{j,k+1} = \mathbf{x}_{j,k} + \mathbf{A}^{-1}[\mathbf{H}^T \mathbf{R}^{-1} \delta \mathbf{y} - \mathbf{B}^{-1}(\mathbf{x}_{j,k} - \mathbf{x}^a) - (\mathbf{S}_{j-1} + \mathbf{D}_{j-1,j})^{-1}(\mathbf{x}_{j,k} - \mathbf{x}_{j-1})]. \quad (13)$$

Likewise the Hessian becomes

$$\mathbf{A} = \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} + \mathbf{B}^{-1} + (\mathbf{S}_{j-1} + \mathbf{D}_{j-1,j})^{-1}. \quad (14)$$

In the backward pass, similar terms are added to both (13) and (14) involving the backward-pass solution at ray $j + 1$, \mathbf{x}_{j+1} , and its error covariance, $\mathbf{S}_{j+1} + \mathbf{D}_{j,j+1}$.

The forward-pass solutions from ray $j - 1$ must still be included, which carries the modest overhead of keeping the forward-pass error covariances for each ray in memory.

The $\mathbf{D}_{j-1,j}$ is modeled as a diagonal matrix with the error variances on the diagonal being proportional to the separation distance of the centers of the basis functions between ray $j - 1$ and ray j , and similarly for $\mathbf{D}_{j,j+1}$. The constant of proportionality is chosen to achieve comparable levels of smoothing in range and azimuth.

In an operational context, not only could one use data from adjacent rays as a constraint, but also the retrievals from the ray at the same azimuth in the previous scan. In this case the formalism introduced in (13) and (14) would be used but with subscript $j - 1$ referring to the ray from the previous scan, rather than the previous ray in the same scan. Likewise, matrix $\mathbf{D}_{j-1,j}$ would represent the error covariance due to temporal rather than spatial decorrelation. If the retrieved rain rate were to be reported immediately then the backward pass would not be used.

f. Correction of attenuation

As described in section 3, the correction for attenuation is achieved within the forward model by using $\ln a$ and Z_h at a particular gate to estimate the associated attenuation, and then using it to correct Z_h at all subsequent gates. This is essentially the Hitschfeld and Bordan (1954) method, except that the size-distribution information in $\ln a$ enables a more accurate correction to be performed than is possible using Z_h alone. Nonetheless, the method is still potentially inaccurate because small errors in the radar calibration or the scattering model can be amplified from one gate to the next. This problem is overcome by the iterative nature of the variational approach. If the first guess of $\ln a$ is too low then it will lead to an overestimate of the attenuation for a given measured Z_h , and hence the correction applied to Z_h will be increasingly overestimated at the gates beyond. This will lead to the forward model overestimating both Z_{dr} and ϕ_{dp} at these gates. When compared with the observed values, the scheme will know that it needs to increase $\ln a$ at the earlier gates to achieve a better fit to the observations, and the subsequent iterations will converge on a retrieval of $\ln a$ that is consistent with them. In practice it is necessary to cap the total attenuation at around 20 dB because some of the iterations may have a profile of $\ln a$ that leads to the attenuation correction being so unstable that the total attenuation predicted exceeds the maximum floating-point value that can be held in the computer.

By requiring the forward model to be consistent with

ϕ_{dp} at the far end of the ray, the correction for attenuation will be similar to that obtained by other methods using ϕ_{dp} (e.g., Holt 1988; Brangi et al. 1990). However, it was stressed by Jameson (1992) and Smyth and Illingworth (1998a) that attenuation is dependent primarily on the imaginary part of the refractive index, whereas Z_h , Z_{dr} , and ϕ_{dp} depend primarily on the real part. Additional information on total attenuation is available from the differential attenuation that exhibits itself as a region of negative Z_{dr} at the back of regions of heavy rain and which does depend on the imaginary part (Smyth and Illingworth 1998a). In principle, the variational scheme can use this information to assist in the correction for attenuation because the combination of variables that correctly predicts any negative Z_{dr} will have to have correctly predicted the attenuation of Z_h and Z_v separately.

g. Retrieval when hail is present

When hail is present in a ray, the scheme described so far is unable to find a solution for $\ln a$ that, when used in the forward model, can closely predict both Z_{dr} and ϕ_{dp} . This property enables a first pass of the algorithm to be used to detect the presence of hail followed by a second to retrieve the contribution that hail makes to the measured reflectivity factor at each pixel, recognizing the fact that there is likely to be a continuum among rain, a rain-hail mixture, and hail. These two passes constitute the outer loop shown in Fig. 1.

The first pass is carried out with the error assigned to Z_{dr} [denoted $\sigma_{Z_{dr}}$ in (5)] increased by a factor of 10. This ensures that the solution is consistent with the measured ϕ_{dp} , and if hail is present it appears as an overestimate of Z_{dr} by the forward model. We then impose the ad hoc criteria that if Z_h is greater than 35 dBZ (after correction for attenuation) and Z_{dr} is overestimated by more than 1.5 dB then hail is present in a pixel (note that other variables could also be used such as the copolar correlation coefficient ρ_{hv} and the linear depolarization ratio). In the second pass, the error in Z_{dr} is returned to its original value and terms are added to the state vector representing the fraction f of the reflectivity factor that is due to hail at each pixel identified by the criteria above. In this way the retrieval is able to find the combination of $\ln a$ and f (forced to lie between 0 and 1) that enables both Z_{dr} and ϕ_{dp} to match the observations.

This strategy for identifying hail was inspired by Smyth et al. (1999), who used Z_h and Z_{dr} to predict ϕ_{dp} along a ray. If it deviated significantly from the observed ϕ_{dp} then hail was deemed to be present in the ray, although its exact location was uncertain. Here, Z_h and ϕ_{dp} are used to predict Z_{dr} , and differences with the

observed Z_{dr} at individual pixels allow the location of the hail to be pinpointed precisely.

The ad hoc criteria given above for the detection of hail enable the forward model to simultaneously match the Z_{dr} and ϕ_{dp} observations, but there may well be further work to be done to tune these parameters for different radars. Nonetheless, it is intended that the use of low Z_h and Z_{dr} thresholds prevents the scheme from being excessively sensitive to the precise thresholds chosen; if the hail flag is triggered incorrectly then the second pass of the scheme should retrieve a vanishingly small hail fraction f .

The situation is made more complicated by the fact that regions of dry tumbling hail with zero Z_{dr} are often surrounded by regions of oblate water-coated hail with higher Z_{dr} . According to Smyth et al. (1999), oblate hail still has a lower Z_{dr} than rain with the same Z_h and ϕ_{dp} and so should be detectable as hail by this method even if its Z_{dr} is difficult to characterize. Differential phase shift on backscatter, particularly at shorter radar wavelengths, makes the K_{dp} of hail similarly difficult to model. For these reasons, even if the location of the hail has been identified correctly, the retrieved f may not be accurate. Hence, rather than reporting the retrieved radar reflectivity Z_h^{hail} due to hail, it may be more appropriate to use a more categorical determination such as, for instance, issuing a hail warning flag when Z_h^{hail} exceeds some threshold. These difficulties are best elucidated in the context of the results of applying the scheme to real hailstorms, and so the discussion is continued in section 4e where the scheme is applied to a case with oblate hail. The difficulty in forward modeling the attenuation and differential attenuation of hail leads to ambiguities in the rain retrievals beyond and is discussed in section 4d.

Because hail occurs in very localized regions, it is not appropriate to represent f by basis functions as with rain. However, there is still a need to impose some kind of smoothness with range whenever hail appears in adjacent gates, because otherwise the retrieved hail structure will attempt to fit the noise in the Z_{dr} and ϕ_{dp} measurements. Smoothness is achieved by adding an additional term J_f to the cost function that penalizes the second derivative of f with range. For a sequence of contiguous hail gates from i_{min} to i_{max} it is given by

$$J_f = \lambda \sum_{i=i_{min}}^{i_{max}} (f_{i-1} - 2f_i + f_{i+1})^2, \quad (15)$$

where λ is a tunable smoothing coefficient. In addition to requiring the solution for f to be smooth, it is also desirable for it to tend toward zero at each end of the sequence. In essence, we wish (15) to behave as if there

were points at $i_{min} - 1$ and $i_{max} + 1$ at which f is always zero. The additional term in the cost function is implemented by creating a ‘‘Twomey’’ matrix \mathbf{T} (Rodgers 2000), adding it to the right-hand side of (9) and subtracting $\mathbf{T}\mathbf{x}$ from the terms in square brackets in (8). The matrix \mathbf{T} contains zeros for elements corresponding to the $\ln a$ part of \mathbf{x} , but for each section of \mathbf{x} containing contiguous f values, it is pentadiagonal. Thus if elements i to $i + 4$ of \mathbf{x} contain the hail fractions for five contiguous gates, then the appropriate subregion of \mathbf{T} would contain

$$\mathbf{T}\{i, \dots, i + 4; i, \dots, i + 4\} = \lambda \begin{bmatrix} 5 & -4 & 1 & 0 & 0 \\ -4 & 6 & -4 & 1 & 0 \\ 1 & -4 & 6 & -4 & 1 \\ 0 & 1 & -4 & 6 & -4 \\ 0 & 0 & 1 & -4 & 5 \end{bmatrix}. \quad (16)$$

This is the same as the Twomey matrix described by Rodgers (2000), except that the outermost rows and columns have been excluded, which neatly enforces the constraint that f should tend toward zero at each end of the sequence. At each iteration, it is necessary to check that each value of f lies between 0 and 1 and if necessary correct it.

With the inclusion of hail and attenuation, the formula for deriving rain rate from the observed linear reflectivity Z_h and the retrieved variables a and f is

$$R = [(1 - f)10^{0.1A_h}Z_h/a]^{1/b}, \quad (17)$$

where A_h is the total two-way attenuation at horizontal polarization in decibels.

h. Calculation of error in final rain-rate field

After the solution has converged, the Hessian matrix \mathbf{A} can be used to estimate the error in the retrieved variables. The first step is to subtract any smoothing terms that had been added [i.e., use the Hessian defined by (9) rather than (14)]. This is because the extra constraints act like an a priori with a low error, so have the effect of artificially reducing the inferred error in the solution. When this has been done, \mathbf{A}^{-1} constitutes the error covariance matrix of the solution in \mathbf{x} . To derive the error in the high-resolution rain-rate field, we first determine the error covariance of the high resolution $\ln a$ field, which is given by

$$\mathbf{S}_{\ln a} = \mathbf{W}\mathbf{A}_{\ln a}^{-1}\mathbf{W}^T, \quad (18)$$

where here $\mathbf{A}_{\ln a}^{-1}$ is taken to be just the part of \mathbf{A}^{-1} that describes the error of the $\ln a$ basis-function coefficients

(i.e., removing any columns and rows associated with hail fraction). The diagonal elements of $\mathbf{S}_{\text{ln}a}$ we denote as $\sigma_{\text{ln}a}^2$ [note that (18) is calculated much more efficiently if only the diagonals are required]. For simplicity, it is assumed that the errors in the various terms in (17) are independent, which enables the error in retrieved rain rate at each gate to be written as

$$\frac{\sigma_R}{R} \approx \sigma_{\text{ln}R} = \frac{1}{b} \left[\left(\frac{\ln 10}{10} \right)^2 (\sigma_{Z_h, \text{dB}}^2 + \sigma_{A_h}^2) \sigma_{\text{ln}a}^2 + \frac{\sigma_f^2}{(1-f)^2} \right]^{1/2}, \quad (19)$$

where $\sigma_{Z_h, \text{dB}}$ is the random error in Z_h in decibels, which has been neglected until now. The error variance in hail fraction σ_f^2 is simply taken from the appropriate diagonal element of \mathbf{A}^{-1} . For simplicity, the error in attenuation σ_{A_h} is simply taken to be 25% of the mean value, that is, $\sigma_{A_h} = A_h/4$.

Note that (19) has neglected errors in the forward model, such as those that are due to deviations of the shape of the size distribution from the one assumed, and uncertainties in the relationship between drop size and axial ratio. The sensitivity of the retrieval scheme to systematic (i.e., calibration) errors in Z_{dr} and ϕ_{dp} is also not represented in (19) but is similar to that for other algorithms as described in detail by Bringi and Chandrasekar (2001).

3. Forward model

In this section the details of the forward model for Z_{dr} and ϕ_{dp} are described. The forward model encapsulates our understanding of how the properties of rain and hail affect these variables. Readers who are familiar with the interpretation of these variables may wish to skip some or all of this section.

a. Lookup tables for rain

The scattering properties of oblate raindrops have been calculated using the T-matrix method (Waterman 1969) for equivalent-volume drop diameters between 0.1 and 10 mm. The Andsager et al. (1999) relationship for drop axial ratio as a function of diameter has been used. However, above 4.5 mm this relationship is not constrained by observations and predicts unrealistically low axial ratios, and therefore here we change seamlessly to the Goddard et al. (1995) shapes. The temperature-dependent refractive index of liquid water is calculated following Liebe et al. (1989).

The task of the forward model is to predict Z'_{dr} and ϕ'_{dp} from the observed Z_h (assumed to be exactly correct) and the elements of the state vector, in particular

$\text{ln}a$. First, the simple case of no attenuation or hail is described. Lookup tables are constructed that relate Z_{dr} and the ratio K_{dp}/Z_h to the ratio Z_h/R . These have been calculated from gamma size distributions with a range of median diameters. At this stage the normalization of the distributions is arbitrary because we are just relating ratios of moments of the size distribution, and so the normalization factor [denoted N_w by Bringi and Chandrasekar (2001)] is eliminated. The results are shown in Fig. 2 for temperatures of 0° and 20°C, and for gamma distributions with shape parameters of $\mu = 0$ and $\mu = 5$. Based on the evidence of Tokay and Short (1996), Wilson et al. (1997), and Illingworth and Johnson (1999), a value of 5 is used later in this paper, but it is trivial to change in the algorithm. It can be seen that for these variables the value of μ leads to a greater uncertainty than does temperature.

By rearranging (1), $\ln(Z_h/R)$ is calculated at each gate from Z_h and $\text{ln}a$. The lookup tables then provide Z_{dr} and K_{dp}/Z_h at each gate from $\ln(Z_h/R)$. Differential phase shift ϕ'_{dp} is calculated in the forward model simply by multiplying K_{dp}/Z_h by Z_h and integrating in range.

In the presence of attenuation the forward model must estimate α_h and α_v , the one-way specific attenuations at the two polarizations (dB km^{-1}). Lookup tables for α_h/Z_h and α_v/Z_h are constructed in the same way as for K_{dp}/Z_h , and Fig. 3a shows α_h/Z_h versus Z_h/R . It can be seen that temperature has a much greater role than in Fig. 2, because attenuation depends primarily on the imaginary part of the refractive index, which is strongly temperature dependent, whereas Z_{dr} and K_{dp} depend on the real part, which is not. Figure 3b shows the corresponding plot for differential attenuation $(\alpha_h - \alpha_v)/Z_h$.

b. Forward model in the presence of attenuation and hail

The lookup tables calculated in the previous section are used in the forward model as follows. Each range gate is considered in turn. At gate i we have the measured linear reflectivity factor $Z_{h,i}$ and two variables from the state vector, $\text{ln}a_i$, and f_i . We also have the total two-way attenuations at the two polarizations $A_{h,i}$ and $A_{v,i}$, in decibels. At the first gate these are both zero. First, the *unattenuated* linear reflectivity factors of the rain and hail are calculated using

$$Z_{h,i}^{\text{rain}} = (1 - f_i) 10^{0.1 A_{h,i}} Z_{h,i} \quad \text{and} \quad (20)$$

$$Z_{h,i}^{\text{hail}} = f_i 10^{0.1 A_{h,i}} Z_{h,i}, \quad (21)$$

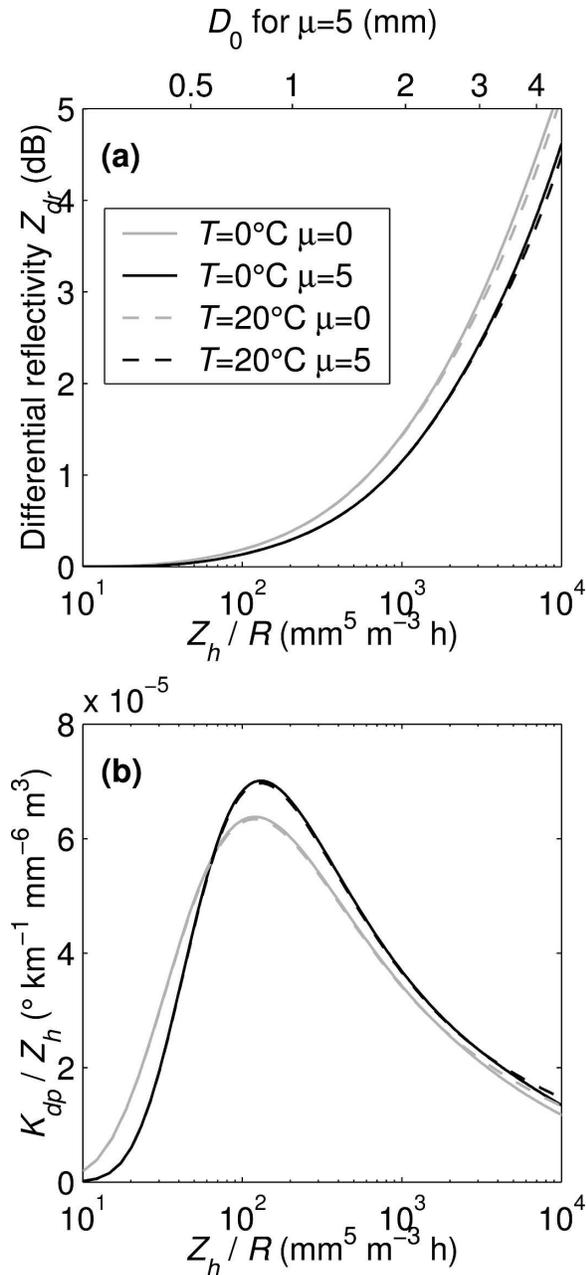


FIG. 2. (a) Differential reflectivity Z_{dr} vs the ratio of reflectivity factor to rain rate Z_h/R for two values of temperature and two gamma-distribution shape parameters μ . The corresponding median volumetric diameter D_0 for $\mu = 5$ is shown on the upper axis. (b) The ratio of one-way specific differential phase shift to reflectivity factor K_{dp}/Z_h vs Z_h/R . The calculations have been performed at S band (3 GHz) using the T-matrix method.

Next, by rearranging (1) we obtain $\ln(Z_{h,i}^{\text{rain}}/R_i)$ from $Z_{h,i}^{\text{rain}}$ and $\ln a_i$. This is used in the lookup table described in the previous section (and plotted in Fig. 2a) to derive the unattenuated differential reflectivity for rain $Z_{dr,i}^{\text{rain}}$ (dB). The measured value is affected by hail, which is

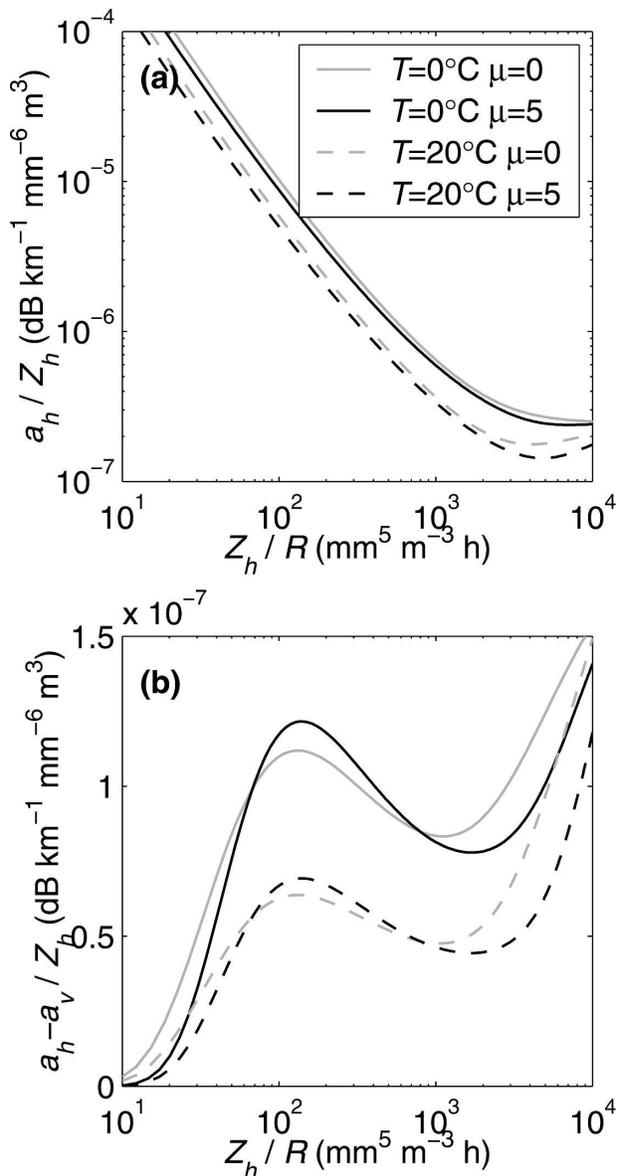


FIG. 3. (a) The ratio of one-way specific attenuation at horizontal polarization to reflectivity factor α_h/Z_h vs Z_h/R for two values of temperature and two gamma-distribution shape parameters μ . (b) The ratio of one-way specific differential attenuation to reflectivity factor $(\alpha_h - \alpha_v)/Z_h$ vs Z_h/R . The calculations have been performed at S band (3 GHz) using the T-matrix method.

assigned an intrinsic differential reflectivity of $Z_{dr,i}^{\text{hail}}$ (usually assumed to be 0 dB), and differential attenuation. Thus the forward-model estimate of differential reflectivity is given by

$$Z'_{dr,i} = Z_{dr,i}^{\text{all}} - A_{h,i} + A_{v,i}, \tag{22}$$

where it can be shown that the unattenuated differential reflectivity of the rain-hail mixture is

$$Z_{\text{dr},i}^{\text{all}} = -10 \log_{10} \left[f_i 10^{-0.1 Z_{\text{dr},i}^{\text{hail}}} + (1 - f_i) 10^{-0.1 Z_{\text{dr},i}^{\text{rain}}} \right]. \quad (23)$$

We now consider the propagation variables. The lookup tables are used to calculate the ratios K_{dp}/Z_h , α_h/Z_h , and α_v/Z_h for rain from $\ln(Z_{h,i}^{\text{rain}}/R_i)$. These ratios are then multiplied by $Z_{h,i}^{\text{rain}}$ itself to recover the actual values $K_{\text{dp},i}^{\text{rain}}$, $\alpha_{h,i}^{\text{rain}}$, and $\alpha_{v,i}^{\text{rain}}$. The forward-model estimate of two-way differential phase shift is then incremented by the contributions from rain and hail:

$$\phi'_{\text{dp},i+1} = \phi'_{\text{dp},i} + 2\Delta r (K_{\text{dp},i}^{\text{rain}} + K_{\text{dp},i}^{\text{hail}}), \quad (24)$$

where Δr is the range-gate spacing. The differential phase shift at the first gate is taken to be zero. Note that for spherical or tumbling hail the one-way specific differential phase shift $K_{\text{dp},i}^{\text{hail}}$ would be zero, but if other information were available on the fall mode then it could be represented as a function of $Z_{h,i}^{\text{hail}}$. The total two-way attenuations are calculated in the same way:

$$A_{h,i+1} = A_{h,i} + 2\Delta r (\alpha_{h,i}^{\text{rain}} + \alpha_{h,i}^{\text{hail}}) \quad \text{and} \quad (25)$$

$$A_{v,i+1} = A_{v,i} + 2\Delta r (\alpha_{v,i}^{\text{rain}} + \alpha_{v,i}^{\text{hail}}). \quad (26)$$

Again, the intrinsic attenuation of the hail, $\alpha_{h,i}^{\text{hail}}$ and $\alpha_{v,i}^{\text{hail}}$, could be parameterized as a function of $Z_{h,i}^{\text{hail}}$, but it is often reasonable to assume hail attenuation to be small relative to rain attenuation. Now that the total attenuation at the next gate has been estimated, we proceed to that gate and repeat the procedure, thus obtaining Z_{dr} and ϕ_{dp} at each gate.

No attempt is made to forward-model effects such as differential phase shift on backscatter, which can add a random offset of either sign to ϕ_{dp} in the cores of hailstorms. Because K_{dp} is not used directly, the retrieval should be relatively insensitive to this effect, with ϕ_{dp} at the far end of the storm still providing a reliable path constraint. However, to avoid the possibility of this phenomenon degrading the retrieval, the error in the measured ϕ_{dp} could be increased just in those hail-containing pixels that are likely to be affected.

c. The Jacobian

The Jacobian matrix is found by calculating the partial derivatives of the expressions in the previous section with respect to each element of the state vector. Its elements are defined in (7) for the simplified case in which f is excluded from the state vector.

We start by considering the derivatives of the total attenuations, which will be required in the formulation of the derivatives of the measured variables. The simplest case is the dependence of A_h at one gate on $\ln a$ at

the gate immediately before. Taking the partial derivative of (25) with respect to $\ln a_i$, we obtain

$$\frac{\partial A_{h,i+1}}{\partial \ln a_i} = 2\Delta r \frac{\partial \alpha_{h,i}^{\text{rain}}}{\partial \ln(Z_{h,i}^{\text{rain}}/R_i)} \frac{\partial \ln(Z_{h,i}^{\text{rain}}/R_i)}{\partial \ln a_i}. \quad (27)$$

Note that $A_{h,i}$ in (25) is dependent only on gates earlier than i and hence not on $\ln a_i$. From (20) it is seen that $Z_{h,i}^{\text{rain}}$ is independent of $\ln a_i$, and so $\partial \ln(Z_{h,i}^{\text{rain}}/R)/\partial \ln a_i = 1/b$, and (27) may be rewritten as

$$\frac{\partial A_{h,i+1}}{\partial \ln a_i} = 2\Delta r \frac{Z_{h,i}^{\text{rain}}}{b} \left[\frac{\partial (\alpha_h/Z_h)}{\partial \ln(Z_h/R)} \right]_i^{\text{rain}}. \quad (28)$$

The term in square brackets is taken directly from the appropriate lookup table and is simply the slope of the graph in Fig. 3. For the derivative of A_h at gate j with respect to $\ln a$ at a gate i much earlier in the ray, things are somewhat more complicated. Not only does a change in $\ln a_i$ imply a change in the attenuation due to gate i (and hence the total attenuation A_h at all subsequent gates), it also changes the degree to which the reflectivity of intervening gates needs to be corrected for attenuation and, in turn, changes their implied contribution to the attenuation at the gate j . This complication may be treated by considering each gate sequentially; the derivative of (25) may be written as follows using the chain rule (omitting the hail term for simplicity):

$$\frac{\partial A_{h,j+1}}{\partial \ln a_i} = \frac{\partial A_{h,j}}{\partial \ln a_i} \left(1 + 2\Delta r \frac{\partial \alpha_{h,j}^{\text{rain}}}{\partial Z_{h,j}^{\text{rain}}} \frac{\partial Z_{h,j}^{\text{rain}}}{\partial A_{h,j}} \right). \quad (29)$$

The derivative of attenuation-corrected linear reflectivity with respect to the two-way attenuation in decibels is $\partial Z_{h,j}^{\text{rain}}/\partial A_{h,j} = 0.1 \ln(10) Z_{h,j}^{\text{rain}}$. The other derivative in the parentheses in (29) is given by

$$\frac{\partial \alpha_{h,j}^{\text{rain}}}{\partial Z_{h,j}^{\text{rain}}} = \frac{\alpha_{h,j}^{\text{rain}}}{Z_{h,j}^{\text{rain}}} + \left(1 - \frac{1}{b} \right) \left[\frac{\partial (\alpha_h/Z_h)}{\partial \ln(Z_h/R)} \right]_j^{\text{rain}}. \quad (30)$$

Thus it is possible to proceed from one gate to the next using the derivative from the previous gate.

The derivation of the derivatives for A_v and ϕ'_{dp} proceeds in a very similar fashion. In the case of ϕ'_{dp} , the adjacent-gate derivative [the equivalent of (28)] is

$$\frac{\partial \phi'_{\text{dp},i+1}}{\partial \ln a_i} = 2\Delta r \frac{Z_{h,i}^{\text{rain}}}{b} \left[\frac{\partial (K_{\text{dp}}/Z_h)}{\partial \ln(Z_h/R)} \right]_i^{\text{rain}}, \quad (31)$$

where the term in square brackets is the gradient of the appropriate line in Fig. 2b. For nonadjacent gates we have (still omitting the hail term)

$$\frac{\partial \phi'_{dp,j+1}}{\partial \ln a_i} = \frac{\partial \phi'_{dp,j}}{\partial \ln a_i} + 0.2 \ln(10) \Delta r \frac{Z_{h,i}^{\text{rain}}}{b} \frac{\partial A_{h,j}}{\partial \ln a_i} \left\{ \frac{K_{dp}}{Z_{h,j}^{\text{rain}}} + \left(1 - \frac{1}{b} \right) \left[\frac{\partial (K_{dp}/Z_h)}{\partial \ln(Z_h/R)} \right]_i^{\text{rain}} \right\}, \quad (32)$$

which is the equivalent of (29) and (30).

For differential reflectivity we use (22) and (23) to first obtain the derivative of Z_{dr} at a pixel with respect to $\ln a$ at the same pixel:

$$\frac{\partial Z'_{dr,i}}{\partial \ln a_i} = \frac{1}{b} \left[\frac{\partial Z_{dr}}{\partial \ln(Z_h/R)} \right]_i^{\text{rain}} (1 - f_i) 10^{0.1(Z_{dr,i}^{\text{all}} - Z_{dr,i}^{\text{rain}})}, \quad (33)$$

where the term in square brackets is the gradient of the appropriate line in Fig. 2a. For the derivative of $Z'_{dr,j}$ with respect to $\ln a_i$ at an earlier gate we first apply the chain rule to $Z_{dr,j}^{\text{all}}$:

$$\frac{\partial Z_{dr,j}^{\text{all}}}{\partial \ln a_i} = \frac{\partial Z_{dr,j}^{\text{all}}}{\partial Z_{dr,j}^{\text{rain}}} \frac{\partial Z_{dr,j}^{\text{rain}}}{\partial \ln(Z_{h,j}^{\text{rain}}/R_j)} \frac{\partial \ln(Z_{h,j}^{\text{rain}}/R_j)}{\partial Z_{h,j}^{\text{rain}}} \frac{\partial Z_{h,j}^{\text{rain}}}{\partial A_{h,j}} \frac{\partial A_{h,j}}{\partial \ln a_i}. \quad (34)$$

Evaluating each of these derivatives and including the attenuation terms in (22), we obtain

$$\frac{\partial Z'_{dr,j}}{\partial \ln a_i} = 0.1 \ln(10) (1 - f_j) 10^{0.1(Z_{dr,j}^{\text{all}} - Z_{dr,j}^{\text{rain}})} \left(1 - \frac{1}{b} \right) \times \left[\frac{\partial Z_{dr}}{\partial \ln(Z_h/R)} \right]_j^{\text{rain}} \frac{\partial A_{h,j}}{\partial \ln a_i} - \frac{\partial A_{h,j}}{\partial \ln a_i} + \frac{\partial A_{v,j}}{\partial \ln a_i}. \quad (35)$$

The calculation of the derivatives of the observations with respect to f proceeds in a similar way (not shown).

4. Results

In this section, three case studies are used to demonstrate the application of the technique.

a. Demonstration of the retrieval of rain rate

Figures 4a–c show Z_h , Z_{dr} , and ϕ_{dp} observed by the Chilbolton 3-GHz radar for a case in which no hail was present. The radar has a beamwidth of 0.28° and a range resolution of 300 m. It has been calibrated using the method of Goddard et al. (1994), which uses the redundancy of Z_{dr} and ϕ_{dp} in moderate to heavy rain. The scans have been truncated at a range of 80 km because of contamination by the melting layer beyond this distance. Ground clutter and other nonprecipita-

tion echoes have been removed by rejecting all pixels with Z_h less than 0 dBZ or with linear depolarization ratio larger than -20 dB. The algorithm has been run using basis functions with centers spaced 3 km apart in range and a decorrelation length of 5 km. The coefficient b in (1) is fixed at 1.5, and retrieved coefficient a is assigned an a priori value of $200 \text{ mm}^6 \text{ m}^{-3} (\text{mm h}^{-1})^{-1.5}$ and an error standard deviation in the a priori value of $\ln a$ of 1.0 (i.e., $+172\%/ -63\%$, which translates to an error in rain rate of a factor of 2). At large signal-to-noise ratios we assume root-mean-square random errors in Z_{dr} and ϕ_{dp} of $\sigma_{Z_{dr}} = 0.2$ dB and $\sigma_{\phi_{dp}} = 3^\circ$; these values are close to the theoretical prediction for the Chilbolton radar (Li et al. 1994) and are generally what is found in the data. The errors are assumed to increase at low signal-to-noise ratios. It is unfortunate that ρ_{hv} was not available from Chilbolton at this time to perform a more rigorous error calculation. Note that operational radars with both higher scan rates than Chilbolton (which is only 1° s^{-1}) and less perfect antennas will tend to have higher errors, even at large signal-to-noise ratio.

Figure 4d shows the rain rate retrieved by the algorithm, and Figs. 4e,f depict the corresponding forward-modeled fields of Z_{dr} and ϕ_{dp} at the final iteration for each ray. It can be seen that the algorithm has converged to a solution in which Z_{dr} and ϕ_{dp} are very well modeled, but without reproducing the random noise in these observations, as would have occurred had a simple $R(Z_h, Z_{dr})$ relationship been applied at every pixel.

Figure 5a shows the corresponding retrieval of the a coefficient, which the algorithm has ensured varies smoothly on a scale of 3–5 km in both directions. Note that the forward model defines these values at all locations, even where no signal was present in the observations. An important step will be to test the retrieved values of a using distrometers, although that is beyond the scope of this paper. Figure 5b depicts $\sigma_{\ln a}$, the retrieval error in $\ln a$, as described in section 2h. It should be stressed that this is only the retrieval error associated with random errors in the measurements of Z_{dr} and ϕ_{dp} ; errors in the forward model (specifically due to size distributions and drop shapes) have not been included. It can be seen to be around 15% through most of the region but is higher within 20 km of the radar because of the reduced number of observations here. It is also higher far from the radar in regions of low signal-to-noise ratio because of the larger error ascribed to Z_{dr} in these regions. Furthermore, at low rain rate the drops are nearly spherical and so the information on drop size contained in Z_{dr} is low (i.e., $\partial Z_{dr}/\partial a \approx 0$) and

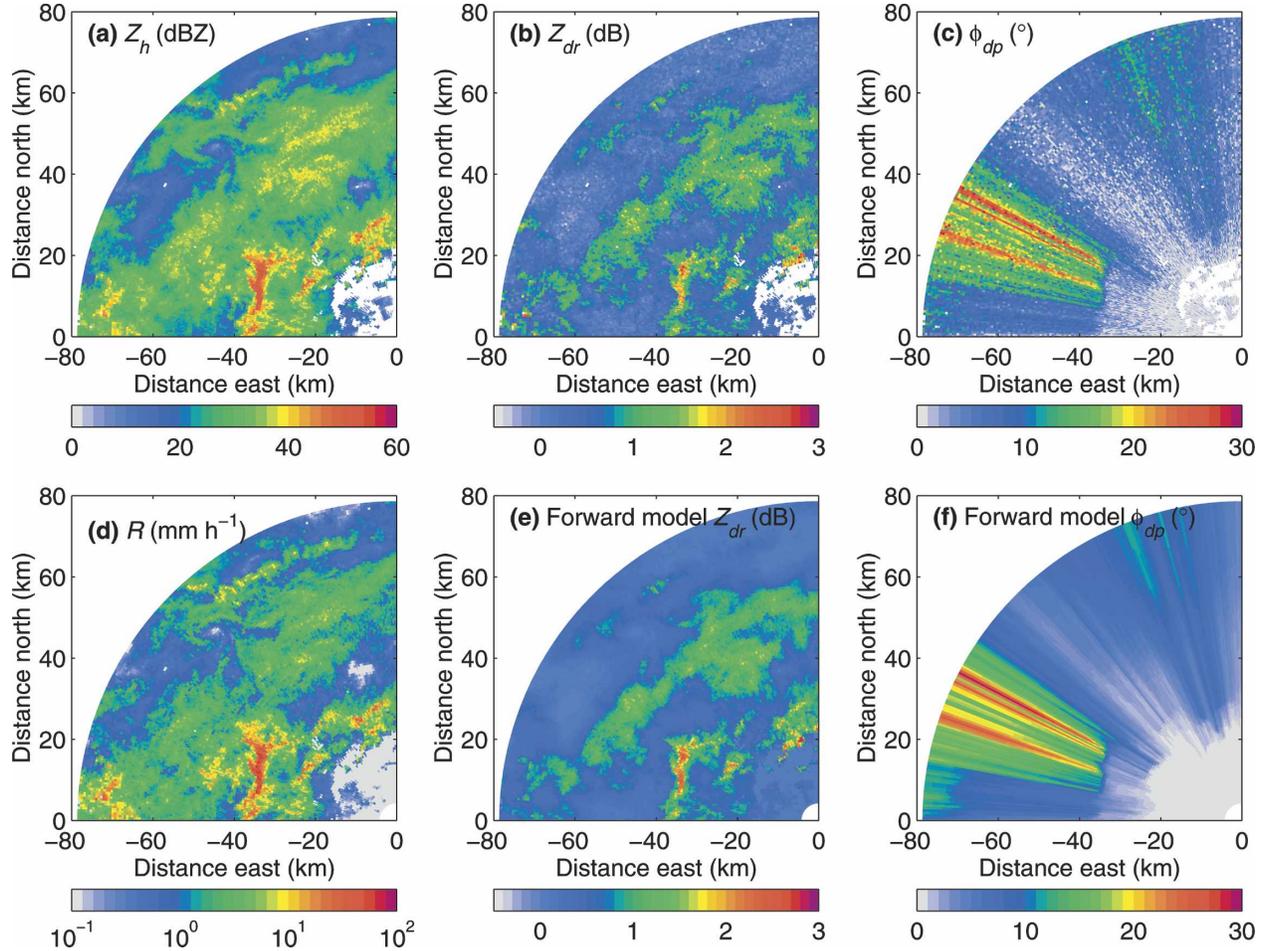


FIG. 4. Radar observations and retrieval from a 1.0° -elevation scan by the Chilbolton 3-GHz radar at 0955 UTC 18 Aug 2000: (a) observed radar reflectivity factor, (b) observed differential reflectivity, (c) observed two-way differential phase shift, (d) retrieved rain rate, (e) differential reflectivity predicted by the forward model at the final iteration of the retrieval, and (f) forward-modeled differential phase shift at the final iteration of the retrieval.

the error in retrieved $\ln a$ tends toward the a priori error of 1.0.

A single ray from this case is shown in Fig. 6. Figures 6a–c depict the observations and, in the case of Z_{dr} and ϕ_{dp} , the corresponding forward-modeled values. The fact that the large Z_{dr} and differential phase shift at a range of 36 km can be accurately modeled with a single value of a indicates clearly that the associated high radar reflectivity factor is due entirely to rain and there is no hail present. Figure 6d shows retrieved coefficient a , again illustrating the increased error in regions with few observations or low signal-to-noise ratio. The corresponding rain rate is shown in Fig. 6e and can be seen to be different by up to a factor of 2 from the rain rate calculated using the fixed a priori estimate of a and to be outside the expected error bounds. Note that the retrieval error has been calculated using $\sigma_{Z_h, dB} = 0$ dB in

(19) to highlight the error due to the retrieval method rather than errors in the final scaling with Z_h .

b. Sensitivity tests

In this section we investigate the dependence of the retrievals on random errors in the observations of Z_{dr} and ϕ_{dp} , as well as the relative importance of these two variables to the rainfall estimate. To investigate the role of Z_{dr} in the retrieval, Figs. 7a,b depict the same as Fig. 5 but with only Z_{dr} being used in the retrieval, that is, with ϕ_{dp} removed from the observation vector in (4). The results are virtually identical, indicating that most of the information on a is being taken from Z_{dr} . This is due to the lower relative error of Z_{dr} as measured by the Chilbolton radar, coupled with the fact that ϕ_{dp} only provides information in regions of moderate to

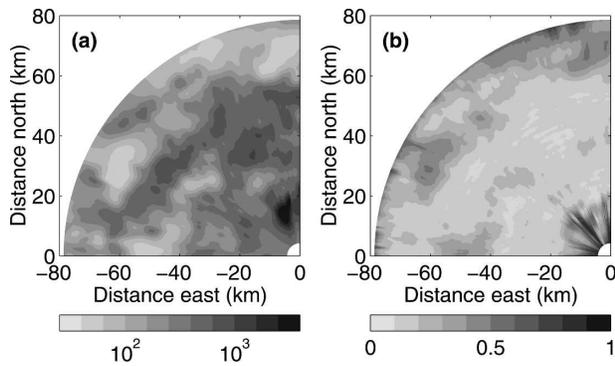


FIG. 5. Retrieval output for the scan shown in Fig. 4: (a) coefficient a in $Z = aR^b$ [$\text{mm}^6 \text{m}^{-3} (\text{mm h}^{-1})^{-1.5}$] and (b) the retrieval error in $\ln a$ assuming a perfect forward model.

high rain rate where there is a significant gradient in differential phase shift.

To investigate the effect of increased observational error on the retrieval, Gaussian-distributed random noise has been added to the Z_{dr} measurements with a standard deviation of 1 dB, and the retrieval has been rerun. The results are shown in Figs. 7c,d. The increased error results in less weight being applied to the Z_{dr} measurements relative to the a priori and the smoothing terms, with the result that the field is smoother and closer to the a priori value, particularly in the low-reflectivity regions farther from the radar. Nonetheless, the broad region of elevated a in the center of the scan is still present, but with increased retrieval error (Fig. 7d). The increased tendency for the solution to approach the a priori may be countered simply by increasing the error assigned to the a priori. Figure 7e shows the retrieved a field for the same noisy Z_{dr} field, but increasing the error in the a priori value of $\ln a$ from 1.0 to 2.5. The results are surprisingly similar to the original retrieval in Fig. 7a that had only the intrinsic Z_{dr} noise of ± 0.2 dB. This indicates that noisier Z_{dr} measurements by a radar with a much shorter dwell time than Chilbolton could still be used for improving rain-rate estimates, providing that the appropriate smoothing of the inferred a coefficient is applied. Note that errors in Z_{dr} of ± 1 dB can only be tolerated if they are truly random, as in this example; there is still a need for Z_{dr} to be unbiased, that is, to be calibrated to better than 0.2 dB.

We next consider the case in which only ϕ_{dp} is used in the retrieval. The top panels of Fig. 8 show the same as Fig. 5 but with Z_{dr} removed from the observation vector. Because of the larger relative error of ϕ_{dp} , the retrieved a is smoother and closer to the a priori than in Fig. 5, although the general pattern is similar to that shown in Fig. 7c, indicating that ϕ_{dp} and Z_{dr} are pro-

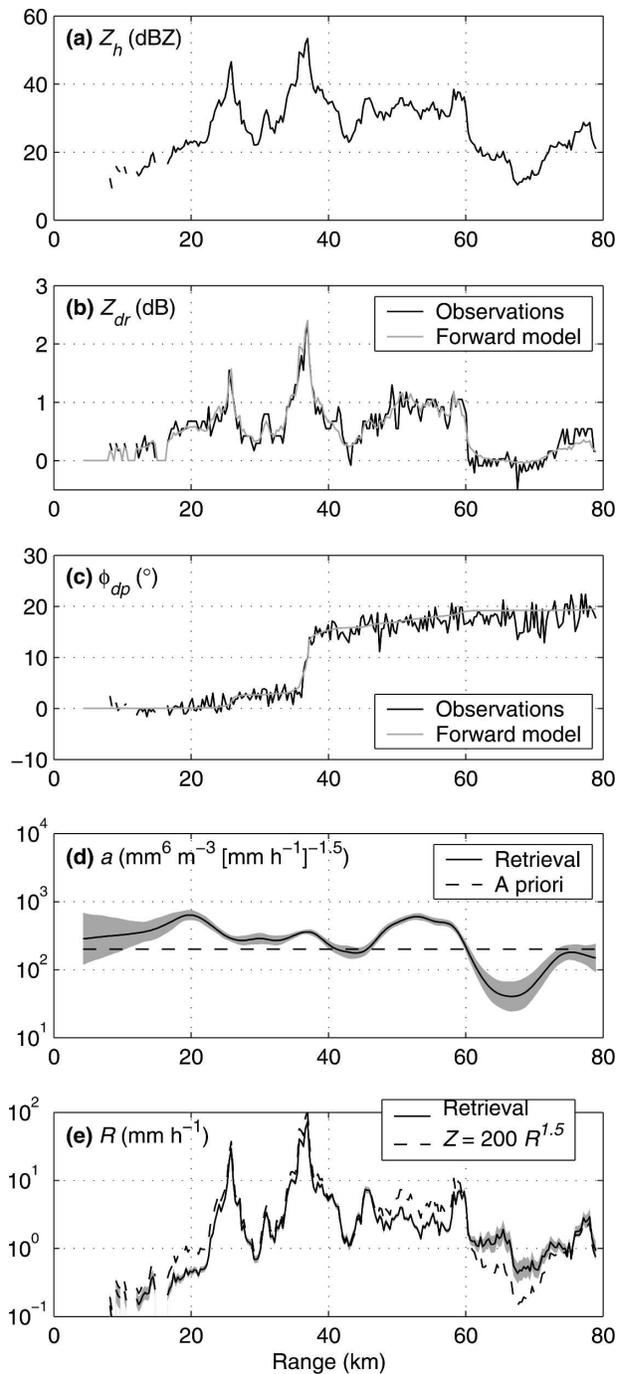


FIG. 6. Observed and retrieved quantities for the 293° -azimuth ray from the case shown in Figs. 4 and 5: (a) observed radar reflectivity factor; (b) differential reflectivity from the observations and the forward model at the final iteration; (c) differential phase shift from the observations and the forward model at the final iteration; (d) retrieved coefficient a in $Z = aR^b$ with the gray band indicating the retrieval error assuming a perfect forward model, together with the a priori value of $200 \text{ mm}^6 \text{ m}^{-3} (\text{mm h}^{-1})^{-1.5}$; and (e) rain rate from the retrieval with the gray band indicating the retrieval error (assuming a perfect forward model and no observational error in Z_h) and using a simple empirical relationship of the form $Z_h = 200R^{1.5}$.

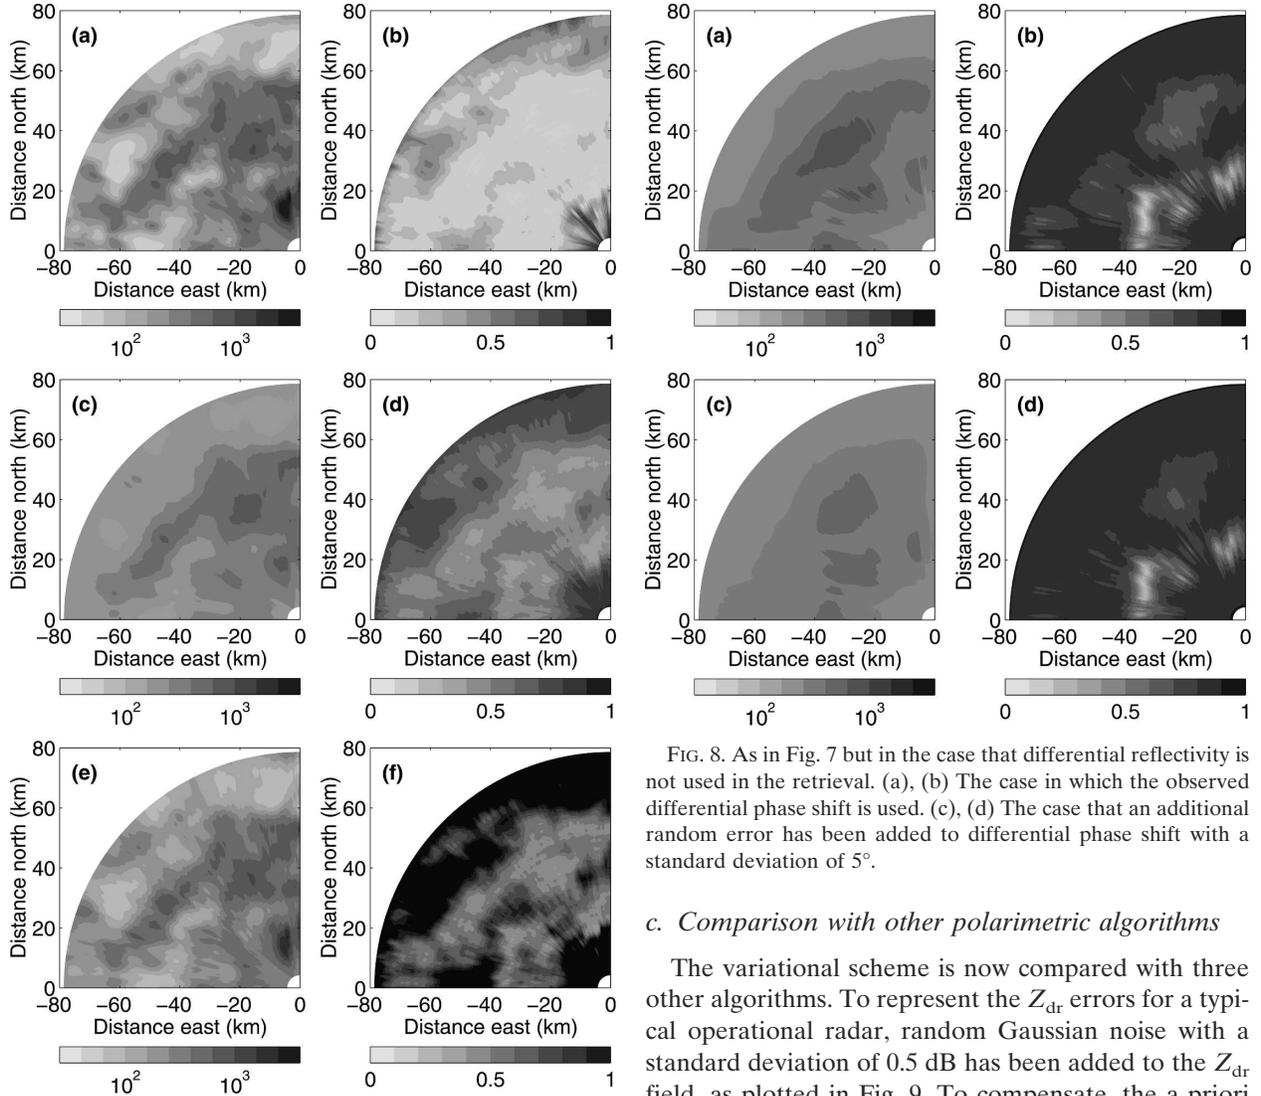


FIG. 7. Retrieval for the scan shown in Fig. 4, but in the case that differential phase shift is not used in the retrieval: (a) coefficient a in $Z_h = aR^b$ [$\text{mm}^6 \text{m}^{-3} (\text{mm h}^{-1})^{-1.5}$] and (b) the retrieval error in $\ln a$. (c), (d) As in (a) and (b), respectively, but in the case that an additional random error has been added to differential reflectivity with a standard deviation of 1 dB. (e), (f) As in (c) and (d), respectively, but with the error in the a priori estimate of $\ln a$ being increased to 2.5.

viding the same information on drop size in this case. The retrieval error in Fig. 8b shows values significantly less than the a priori error only in regions where there is a significant gradient in ϕ_{dp} . This is in contrast to Z_{dr} , which provides information on drop size in much weaker rain rates. Figures 8c,d show the results of a retrieval carried out after adding 5° Gaussian noise to the ϕ_{dp} measurements and keeping the error in the a priori constant. As in Figs. 7c,d, the a coefficient can be seen to be even more constrained by the a priori value.

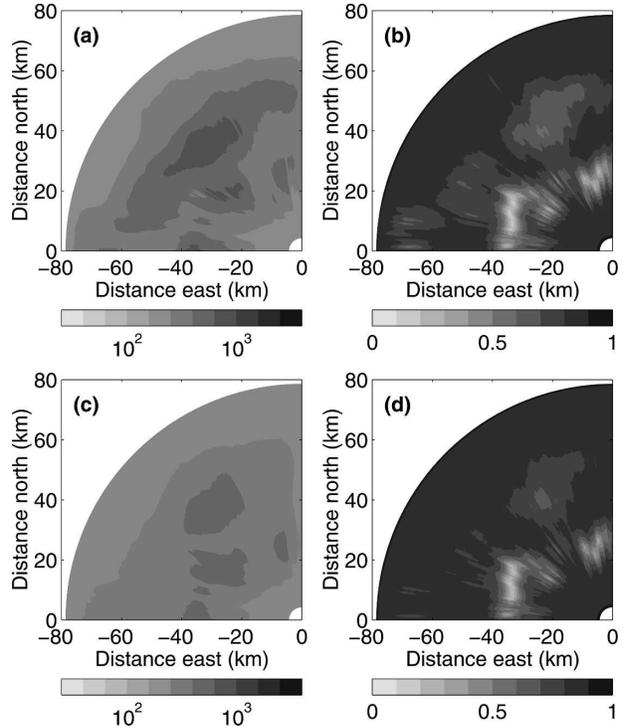


FIG. 8. As in Fig. 7 but in the case that differential reflectivity is not used in the retrieval. (a), (b) The case in which the observed differential phase shift is used. (c), (d) The case that an additional random error has been added to differential phase shift with a standard deviation of 5° .

c. Comparison with other polarimetric algorithms

The variational scheme is now compared with three other algorithms. To represent the Z_{dr} errors for a typical operational radar, random Gaussian noise with a standard deviation of 0.5 dB has been added to the Z_{dr} field, as plotted in Fig. 9. To compensate, the a priori error in $\ln a$ has been increased to 1.7. Comparison of the black line in Figs. 6d and 9b confirms the point made in the previous section that the variational scheme is very robust to random errors in Z_{dr} .

The first of the other algorithms considered is the “ZPHI” method of Testud et al. (2000), which may be implemented as a limiting case of the variational scheme. In essence, a single value of a is retrieved for each ray such that when combined with Z_h it correctly predicts the ϕ_{dp} at the far end of the ray. It can be seen in Fig. 9b that the constant retrieved value of a [labeled $R(Z, \phi_{dp})$ in the figure] matches the variational retrieval at a range of 37 km, because this is where most of the differential phase shift occurred. Elsewhere the two retrievals diverge. This algorithm can also be implemented by dividing ϕ_{dp} into a number of segments along the ray and retrieving a separate value of a for each one. To some extent this is akin to the ϕ_{dp} -only

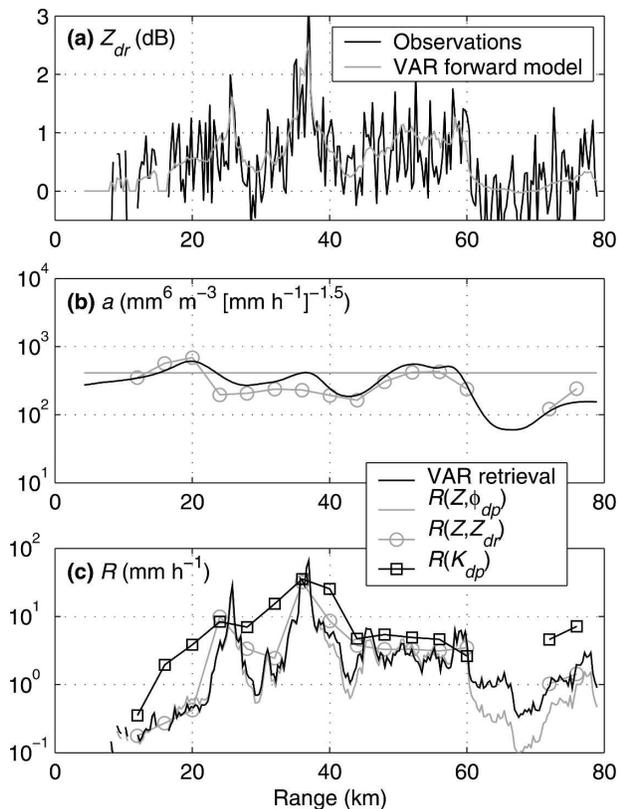


FIG. 9. (a) Observed differential reflectivity with an additional random error of 0.5 dB added, and forward-model values from the variational retrieval (denoted “VAR”) for the same ray of data as shown in Fig. 6; (b) comparison of a retrieved by the variational retrieval, an algorithm using just Z_h and ϕ_{dp} , and an algorithm using just Z_h and Z_{dr} ; (c) comparison of rain rate retrieved by the three algorithms shown in (b), plus an algorithm using K_{dp} alone.

scheme illustrated in Fig. 8a, but without the smoothing of a in the azimuthal direction.

The second algorithm utilizes Z_{dr} and is denoted $R(Z, Z_{dr})$ in Fig. 9. Because Z_{dr} is too noisy to use at each gate, both Z_h and Z_{dr} have been averaged over regions measuring 4 km in range and 2° in azimuth. The lookup table shown in Fig. 2a is then used to estimate rain rate and a . In general, the results are reasonably similar to the variational scheme, except that fine structure in the rain field cannot be resolved and no retrievals are possible between 60 and 70 km where Z_{dr} is slightly negative.

The final algorithm utilizes K_{dp} alone, which has the advantage that it is insensitive to calibration errors that can beset both Z_h and Z_{dr} . However, ϕ_{dp} is a noisy field and so has been averaged over the same $4 \text{ km} \times 2^\circ$ regions used above, before calculating the derivative with range. The relationship of Bringi and Chandrasekar (2001) is used to calculate rain rate: $R(K_{dp}) = 50.7K_{dp}^{0.866}$, where one-way K_{dp} has units of degrees per

kilometer and rain rate is in millimeters per hour. It can be seen in Fig. 9c that the remaining noise in K_{dp} makes the retrieval of low rain rates inaccurate, and sometimes impossible when K_{dp} is negative. There is a trade-off between resolution and accuracy for this method; higher resolutions were attempted, but it was found that they resulted in many more negative K_{dp} regions where rain rate could not be retrieved and, in general, a greater deviation from the other methods in Fig. 9c. In an operational environment, an adaptive window is required if the heavy rain is to be resolved.

Despite the differences in the detailed structure, each of the methods agrees within approximately a factor of 2 when the path-integrated rain rate is considered. Note that both the K_{dp} -only method and the ZPHI method would expect more accuracy at shorter wavelengths in moderate rain because of the increased differential phase shift for a given rain rate, but at higher rain rates this is countered by the increase in attenuation.

d. Demonstration of the retrieval of hail intensity

To demonstrate the ability of the algorithm to retrieve hail, Figs. 10 and 11 show the results of its application to a much more challenging case with differential phase shift of up to 80° and differential attenuation of up to 1 dB. The gray regions in Figs. 10b,e show the negative Z_{dr} due to differential attenuation by heavy rain closer to the radar. The settings are as in the previous case except for the additional loop to retrieve hail (as described in section 2g).

The retrieved a and its error are shown in Figs. 11a,b. As in the previous case, the error tends to increase toward the edge of the regions of rain. The fraction of the radar reflectivity factor that is due to hail is shown in Fig. 11c and appears in the form of isolated patches of hail that are several kilometers across. To indicate the workings of the hail retrieval, it is illuminating to consider the retrievals for a single ray through the most intense part of the hailstorm, shown in Fig. 12. In the first pass of the algorithm to locate the hail, the scan is assumed to consist only of rain and the error on Z_{dr} is multiplied by 10. This way, the initial retrieval is mostly constrained by ϕ_{dp} , so that differences between the information provided by Z_{dr} and ϕ_{dp} (an indicator of the presence of hail; see Smyth et al. 1999) may be highlighted. The observed peak in reflectivity factor of 59 dBZ at a range of 64 km (see Fig. 12a) is associated with no significant increase in Z_{dr} or gradient in ϕ_{dp} . In an attempt to make this consistent with ϕ_{dp} , the first pass of the algorithm increases a (see the gray dashed line in Fig. 12d) because larger drops result in a lower specific differential phase shift for a given radar reflectivity factor. Even then, the forward model cannot accurately

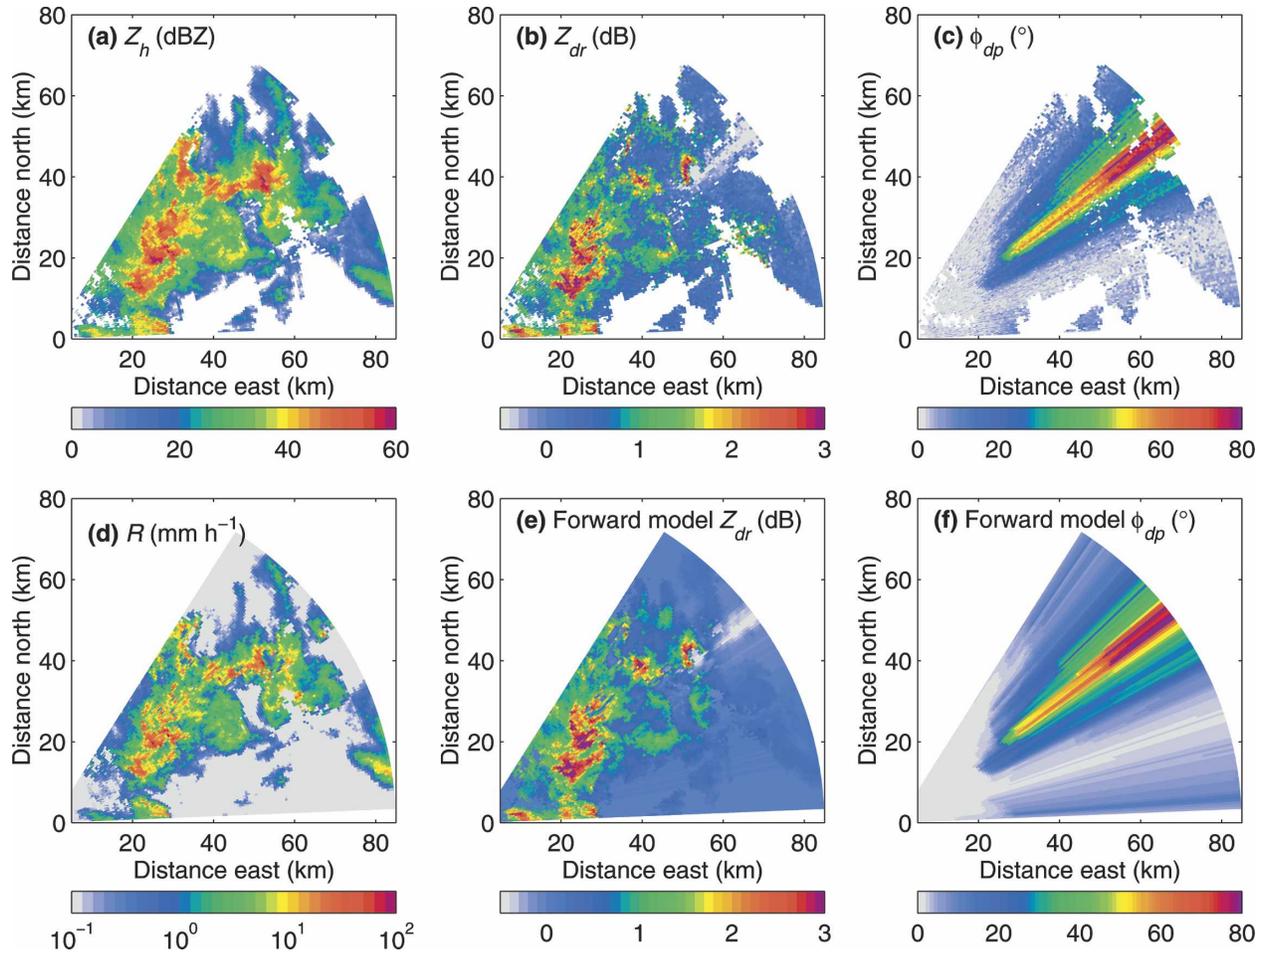


FIG. 10. As in Fig. 4, but for a 1.5° -elevation scan by the Chilbolton 3-GHz radar at 1722 UTC 19 May 1999. Note the change in color scale for ϕ_{dp} .

match ϕ_{dp} in this region (see the gray dashed line in Fig. 12c). When the forward-modeled Z_{dr} from the first retrieval (gray dashed line in Fig. 12b) is compared with the observations, the difference of 6 dB is found to exceed the threshold of 1.5 dB, and so hail is deemed to be present.

Next, the assumed observational error in Z_{dr} is returned to its original value and a second pass of the algorithm is carried out, in which hail fraction f is retrieved at the pixels where the Z_{dr} threshold was exceeded. The solid gray lines in Figs. 12b,c indicate that now it is possible for the retrieval to be consistent with both Z_{dr} and ϕ_{dp} . The gray-filled regions in Fig. 12a show the retrieved contribution of hail to Z_h , highlighting an intense hail shaft centered on a range of 64 km, but with weaker hail regions earlier in the ray. With the dominance of hail rather than rain at 64 km, the algorithm is no longer forced to retrieve a very high a coefficient in an attempt to reconcile Z_h and ϕ_{dp} , and a

lower value is retrieved (see Fig. 12d). Figure 12e shows that when the hail contribution to Z_h is removed, the inferred rain rate is more than an order of magnitude lower than that obtained by blind application of an empirical Z - R relationship. This is crucially important if false alarms are to be avoided in the forecasting of flash floods. Nonetheless, it should be noted that the rain rate is not very accurate when hail is present because of the large uncertainty in the fraction of the reflectivity factor that is due to rain. This is reflected by the increased error reported in a at this point [because of the final term in square brackets in (19)]. When heavy rain coexists with hail (as reported by the algorithm between 30 and 40 km) then ϕ_{dp} still provides a reasonable constraint on rain rate and the rain part of the retrieval will be more accurate.

The performance of the scheme in the presence of heavy rain and hail is encouraging, but there is a need to examine its reliability in the rain regions that lie in

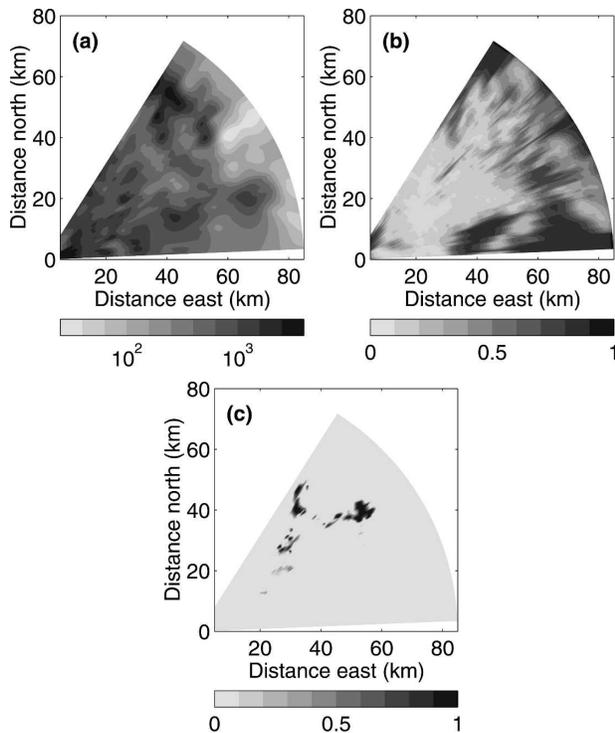


FIG. 11. Retrieval output for the scan shown in Fig. 10: (a) coefficient a in $Z_h = aR^b$ [$\text{mm}^6 \text{m}^{-3} (\text{mm h}^{-1})^{-1.5}$], (b) the retrieval error in $\ln a$ assuming a perfect forward model, and (c) the fraction f of radar reflectivity factor that is due to hail.

the shadow of intense hail shafts. The heavy rain between 20 and 40 km in Fig. 12 caused appreciable attenuation. The forward model includes a prediction for attenuation of the horizontally and vertically polarized radiation that is able to predict correctly the observed negative Z_{dr} in Figs. 10b and 12b. However, Figs. 11a and 12d reveal that the coefficient a in the shadow of the largest region of hail is lower than anywhere else in the scan. The most likely reason for this is not that the drops are unusually small in this region, but rather that the forward model has underestimated the differential attenuation that occurred earlier in the ray. It is possible that there is an error in the forward modeling of differential attenuation by heavy rain, but it is more likely that there were nonspherical hailstones present that caused additional differential attenuation. There is clearly a need to investigate how to represent this effect. One approach could be to retrieve an additional variable describing the asphericity of the hail, but that is beyond the scope of this paper.

e. Retrieval in the presence of oblate hail and differential phase shift on backscatter

Last, the scheme is tested on an isolated thunderstorm observed on 10 July 1995, which was found in the

detailed study of Smyth et al. (1999) to contain oblate hail that led to significant differential phase shift on backscatter δ . This case therefore provides a stringent test of the stability of any hail retrieval method. A horizontal scan through this storm was shown in Fig. 7 of Smyth et al. (1999) and so is not reproduced here. A ray through the base of the storm is shown in Fig. 13 and has been split into the same four regions as they used.

Region 1 was stated by Smyth et al. (1999) to be composed predominantly of a small concentration of large raindrops, and indeed that is what has been retrieved at the leading edge of the storm. In region 2 there was a mixture of rain and tumbling hail, and the forward-modeled values of Z_{dr} and ϕ_{dp} are able to simulate accurately the observations at this location. Region 4 contained lighter rain, and the near-constant ϕ_{dp} in this region provides a useful path constraint for the scheme in its estimate of the rain rate through the core of the storm.

A more complex picture emerges in region 3 (and to a lesser extent between regions 1 and 2), where substantial δ reduced the measured ϕ_{dp} by as much as 20° below what would be expected from propagation effects alone. This can be explained by the presence of large oblate wet hailstones with a degree of alignment (i.e., not falling with their major axis always in the horizontal plane), confirmed by Smyth et al. (1999) who collected a number of disc-shaped hailstones from this storm that had a typical aspect ratio of 0.6 and typical maximum dimension of 2 cm. Because the forward model is unable to represent nonspherical hail, the retrieval scheme has diagnosed the presence of a mixture of raindrops and spherical hail in this location, enabling it to fit the observed Z_{dr} but not the observed ϕ_{dp} . Some inaccuracy should therefore be expected in the weighting between rain and hail, because oblate hail would be expected to contribute itself to positive Z_{dr} . However, the forward-modeled ϕ_{dp} does correctly predict the propagation component and is very similar to the curve that was “fitted by eye” by Smyth et al. (1999). This implies that the estimated mean rain rates across the storm should still be reasonable.

Note that some of the anomalous ϕ_{dp} values may have been caused by increased statistical errors due to the reduced copolar correlation, rather than purely a δ effect, but ρ_{hv} was unfortunately not measured during this case. However, the nature of the anomalous ϕ_{dp} region when viewed in a horizontal scan was to be systematically negative, not random as would be the case if it were simply statistical error. In any case, the retrieval scheme is robustly stable whatever the mechanism.

There is some evidence that this configuration of a

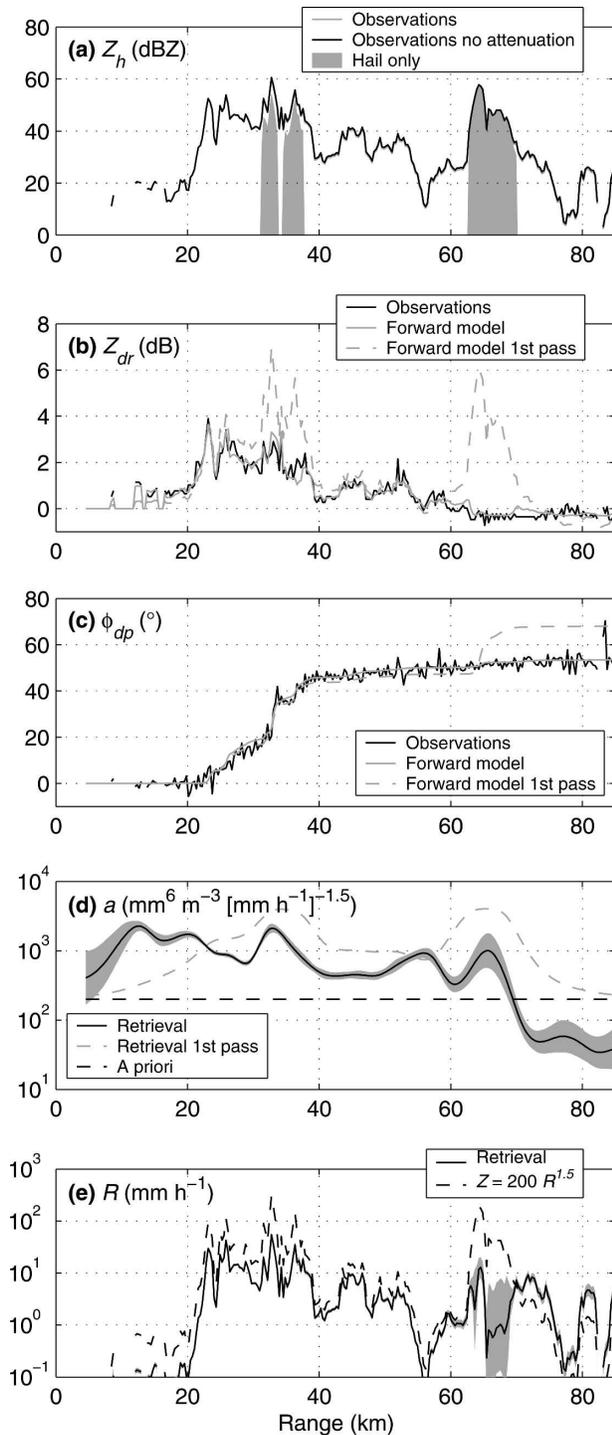


FIG. 12. Observed and retrieved quantities for the 54°-azimuth ray from the case shown in Figs. 10 and 11: (a) radar reflectivity factor as observed, observed but with correction for attenuation, and from the hail component only; (b) differential reflectivity from the observations, the forward model, and the forward model in the first pass of the algorithm to identify the location of the hail; (c) two-way differential phase shift from the observations, the forward model, and the forward model in the first pass; (d) retrieved coefficient a in $Z_h = aR^b$ with the gray band indicating the

core of dry tumbling hailstones being encircled by a region of water-coated oblate hail also occurred in the previous case; Fig. 10b shows an arc of high Z_{dr} immediately adjacent to the largest region of hail identified in Fig. 11c. However, there is no evidence of δ in the observed ϕ_{dp} , and so it is equally likely that this arc was simply heavy rain.

Hence we conclude that the scheme is stable in the presence of oblate hail and differential phase shift on backscatter and correctly identifies the location of hail, even if its reflectivity fraction is not accurate. Note that any algorithm that explicitly calculated K_{dp} by taking the derivative of the observations in Fig. 13c would have severe problems in interpreting the negative values followed by large positive values in the vicinity of region 3.

5. Conclusions

For the first time a variational method has been used to tackle the problem of retrieving rain rate and hail intensity from polarization radar. By an explicit treatment of the observational errors, the scheme is able to make optimum use of the complementary but sometimes conflicting information from Z_{dr} , ϕ_{dp} , and any a priori information we have on the size distribution. Horizontal smoothing of the information provided by the polarimetric variables enables random errors in Z_{dr} of up to 1 dB to be tolerated, and the use of a forward model appears to solve the problem of how to handle negative K_{dp} and Z_{dr} . Attenuation correction is included in the forward model in such a way that it is consistent with both ϕ_{dp} and the measured differential attenuation; hence it will be at least as accurate as methods that use only one of these variables. The algorithm is fast enough that it may be applied in real time to operational radar data, where derived rain-rate fields are required within a matter of minutes of the event. It should be stressed that other formulations of the problem are possible within the variational framework, for example, use of a different choice for the variables in the state vector or a different method for spatial smoothing.

The retrievals are still sensitive to problems that affect existing (nonvariational) polarimetric algorithms,

←

retrieval error assuming a perfect forward model, together with the value retrieved in the first pass, and the a priori value of $200 \text{ mm}^6 \text{ m}^{-3} (\text{mm h}^{-1})^{-1.5}$; and (e) rain rate for both the full retrieval (with the gray band indicating the error) and assuming a simple empirical relationship of the form $Z_h = 200R^{1.5}$.

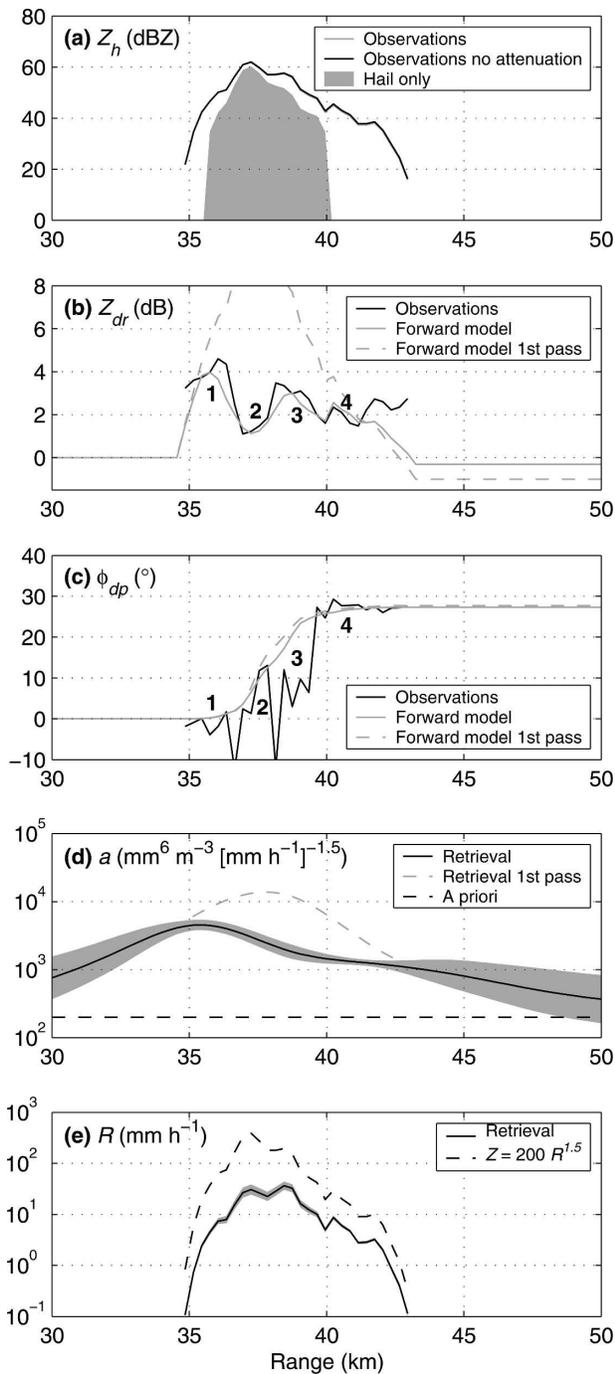


FIG. 13. As in Fig. 12 but for a ray at an azimuth of 29° , elevation 0.7° , through an isolated storm at 1707 UTC 10 Jul 1995. This ray was also analyzed in Fig. 12 of Smyth et al. (1999), although note that they averaged four adjacent rays whereas only a single ray is shown here. The numbered regions 1–4 in (b) and (c) are discussed in the text.

such as calibration of Z_{dr} , beam obscuration, and clutter rejection. Nonetheless, it should be possible to alleviate these problems within a variational framework. As an example, the retrieval in low-rain-rate regions has been found to be very sensitive to the calibration of Z_{dr} . A simple solution would be to manipulate the elements of the observational error covariance matrix \mathbf{R} [in (8) and (9)], either by including the appropriate off-diagonal terms that arise from a systematic error in all Z_{dr} values or by simply increasing the error assigned to the Z_{dr} measurements at low values of Z . This would then cause the retrieved coefficient a to tend toward the a priori value at low rain rate. A more ambitious approach would be to devise a variational calibration scheme that would be run on a routine basis to correct Z_h and Z_{dr} ; in essence one would formulate a scheme in which the Z_h and Z_{dr} calibration factors were included as retrieved variables in the state vector.

The next step is clearly to validate the retrieved drop-size information against distrometer data. The performance of the scheme is dependent on the relative weight given to the observations, the a priori and the various smoothness constraints, so if the scheme is to be applied operationally then observational error needs to be calculated rigorously for that particular radar, in particular to include its dependence on signal-to-noise ratio and whether ground clutter is likely to be present, making use of ρ_{hv} where available. In intense convection it would be desirable to increase the error in ϕ_{dp} due to the possibility of backscatter differential phase, and in Z_h due to the possibility of non-Rayleigh scattering by large hydrometeors. Of particular interest in an operational context is the application of the scheme at C band (typically around 5.6 GHz), which is commonly used in Europe and Japan but where attenuation and differential phase shift are much larger than for the S-band data used in this paper.

Acknowledgments. I thank Anthony Illingworth, Gemma Furness, Nancy Nichols, Sarah Dance, and Robert Thompson for useful discussions. The Chilbolton 3-GHz radar data were provided by the Rutherford Appleton Laboratory. The T-matrix code was written and provided by S. Vivekanandan.

APPENDIX

The Basis-Function Matrix

Here we outline how matrix \mathbf{W} in section 2c is defined. In (10), \mathbf{x} represents the amplitudes (or “control points”) of the set of n basis functions contained in the columns of \mathbf{W} . The rows of \mathbf{W} represent the m points of

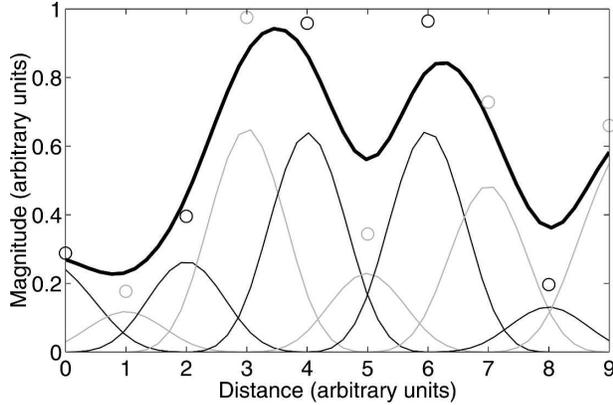


FIG. A1. Illustration of the use of cubic splines to produce a smooth, high-resolution function $\hat{\mathbf{x}}$ (thick black line) from a small number of control points \mathbf{x} (circles). The individual weighted basis functions are shown by the thin black and gray lines and sum to equal $\hat{\mathbf{x}}$.

the high-resolution grid at which each basis function has been evaluated and at which we require the values of $\hat{\mathbf{x}}$. Cubic splines have the property that the resulting $\hat{\mathbf{x}}$ will be continuous in itself and its first and second derivatives but will not necessarily pass through the control points (Prenter 1975). We consider element j of $\hat{\mathbf{x}}$, denoted \hat{x}_j , that lies between control points i and $i + 1$ of \mathbf{x} . The order of the spline determines how many control points on which \hat{x}_j will depend. For a cubic spline, \hat{x}_j depends on the four nearest control points, that is, x_{i-1} , x_i , x_{i+1} , and x_{i+2} . Defining u as a variable between 0 and 1 that represents the distance that point j lies between the control points i and $i + 1$, the four nonzero elements of row j of \mathbf{W} are given by

$$\mathbf{W}\{j; i - 1, \dots, i + 2\} = \frac{1}{6} \begin{pmatrix} (1 - u)^3 \\ 4 - 6u^2 + 3u^3 \\ 1 + 3u + 3u^2 - 3u^3 \\ u^3 \end{pmatrix}^T. \quad (\text{A1})$$

Each row of \mathbf{W} is built up in the same way for each point j on the high-resolution grid. At the edge of the domain, the basis functions are constructed such that $\hat{\mathbf{x}}$ is produced as if there were a repeated control point beyond each end of the domain. Figure A1 demonstrates the fitting of a smooth curve using this method.

REFERENCES

- Andsager, K., K. V. Beard, and N. F. Laird, 1999: Laboratory measurements of axis ratios for large raindrops. *J. Atmos. Sci.*, **56**, 2673–2683.
- Austin, R. T., and G. L. Stephens, 2001: Retrieval of stratus cloud microphysical parameters using millimeter-wave radar and visible optical depth in preparation for CloudSat. 1. Algorithm formulation. *J. Geophys. Res.*, **106**, 28 233–28 242.
- Blackman, T. M., and A. J. Illingworth, 1997: Examining the lower limit of K_{dp} rainrate estimation including a case study at S-band. Preprints, *28th Conf. on Radar Meteorology*, Austin, TX, Amer. Meteor. Soc., 117–118.
- Bringi, V. N., and V. Chandrasekar, 2001: *Polarimetric Doppler Weather Radar: Principles and Applications*. Cambridge University Press, 636 pp.
- , —, N. Balakrishnan, and D. S. Zrnić, 1990: An examination of propagation effects in rainfall on radar measurements at microwave frequencies. *J. Atmos. Oceanic Technol.*, **7**, 829–840.
- Daley, R., 1991: *Atmospheric Data Analysis*. Cambridge University Press, 457 pp.
- Furness, G., 2005: Using optimal estimation theory for improved rainfall rates from polarization radar. M.Sc. dissertation, Dept. of Mathematics, University of Reading, 66 pp.
- Goddard, J. W. F., J. Tan, and M. Thurai, 1994: Technique for calibration of meteorological radars using differential phase. *Electron. Lett.*, **30**, 166–167.
- , K. L. Morgan, A. J. Illingworth, and H. Sauvageot, 1995: Dual wavelength polarization measurements in precipitation using the CAMRa and Rabelais radars. Preprints, *27th Int. Conf. on Radar Meteorology*, Vail, CO, Amer. Meteor. Soc., 196–198.
- Gourley, J. J., P. Tabary, and J. Parent du Chatelet, 2007: A fuzzy logic algorithm for the separation of precipitating from non-precipitating echoes using polarimetric radar observations. *J. Atmos. Oceanic Technol.*, **24**, 1439–1451.
- Hitschfeld, W., and J. Bordan, 1954: Errors inherent in the radar measurement of rainfall at attenuating wavelengths. *J. Meteor.*, **11**, 58–67.
- Holt, A. R., 1988: Extraction of differential propagation phase from data from S-band circularly polarised radar. *Electron. Lett.*, **24**, 1241–1242.
- Illingworth, A. J., 2003: Improved precipitation rates and data quality by using polarimetric measurements. *Weather Radar: Principles and Advanced Applications*, P. Meischner, Ed., Springer-Verlag, 130–166.
- , and M. P. Johnson, 1999: The role of raindrop shape and size spectra in deriving rainfall rates using polarisation radar. Preprints, *29th Int. Conf. on Radar Meteorology*, Montreal, QC, Canada, Amer. Meteor. Soc., 301–304.
- , and R. J. Thompson, 2005: The estimation of moderate rain rates with operational polarisation radar. Preprints, *32d Conf. on Radar Meteorology*, Albuquerque, NM, Amer. Meteor. Soc., CD-ROM, P9R.1.
- Jameson, A. R., 1992: The effect of temperature on attenuation-correction schemes in rain using polarization propagation differential phase shift. *J. Appl. Meteor.*, **31**, 1106–1118.
- Kitchen, M., R. Brown, and A. G. Davies, 1994: Real-time correction of weather radar data for the effects of bright band range and orographic growth in widespread precipitation. *Quart. J. Roy. Meteor. Soc.*, **120**, 1231–1254.
- Li, H., A. J. Illingworth, and J. Eastment, 1994: A simple method of Dopplerizing a pulsed magnetron radar. *Microwave J.*, **37**, 226–236.
- Liebe, H. J., T. Manabe, and G. A. Hufford, 1989: Millimeter-wave attenuation and delay rates due to fog/cloud conditions. *IEEE Trans. Antennas Propag.*, **37**, 1617–1623.
- Löhnert, U., S. Crewell, and C. Simmer, 2004: An integrated ap-

- proach toward retrieving physically consistent profiles of temperature, humidity, and cloud liquid water. *J. Appl. Meteor.*, **43**, 1295–1307.
- Mittermaier, M. P., and A. J. Illingworth, 2003: Comparison of model-derived and radar-observed freezing level heights: Implications for vertical reflectivity profile correction schemes. *Quart. J. Roy. Meteor. Soc.*, **129**, 83–96.
- Prenter, P. M., 1975: *Splines and Variational Theory*. John Wiley and Sons, 323 pp.
- Rodgers, C. D., 2000: *Inverse Methods for Atmospheric Sounding: Theory and Practice*. World Scientific, 238 pp.
- Ryzhkov, A. V., and D. S. Zrnić, 1996: Advantages of rain measurements using specific differential phase. *J. Atmos. Oceanic Technol.*, **13**, 454–464.
- , T. J. Schuur, D. W. Burgess, P. L. Heinselm, S. E. Giangrande, and D. S. Zrnić, 2005: The joint polarization experiment: Polarimetric rainfall measurements and hydrometeor classification. *Bull. Amer. Meteor. Soc.*, **86**, 809–824.
- Sachidananda, M., and D. S. Zrnić, 1987: Rain rate estimates from differential polarization measurements. *J. Atmos. Oceanic Technol.*, **4**, 588–598.
- Seliga, T. A., and V. N. Bringi, 1976: Potential use of radar differential reflectivity measurements at orthogonal polarizations for measuring precipitation. *J. Appl. Meteor.*, **15**, 69–76.
- Smyth, T. J., and A. J. Illingworth, 1998a: Correction for attenuation of radar reflectivity using polarization data. *Quart. J. Roy. Meteor. Soc.*, **124**, 2393–2415.
- , and —, 1998b: Radar estimates of rainfall rates at the ground in bright band and non-bright band events. *Quart. J. Roy. Meteor. Soc.*, **124**, 2417–2434.
- , T. M. Blackman, and A. J. Illingworth, 1999: Observations of oblate hail using dual-polarisation radar and implications for hail-detection schemes. *Quart. J. Roy. Meteor. Soc.*, **125**, 993–1016.
- Testud, J., E. L. Bouar, E. Obligis, and M. Ali-Mehenni, 2000: The rain profiling algorithm applied to polarimetric weather radar. *J. Atmos. Oceanic Technol.*, **17**, 332–356.
- Tokay, A., and D. A. Short, 1996: Evidence from tropical rain-drop spectra of the origin of rain from stratiform versus convective clouds. *J. Appl. Meteor.*, **35**, 355–371.
- Waterman, P., 1969: Scattering by dielectric obstacles. *Alta Freq.*, **38**, 348–352.
- Wilson, D. R., A. J. Illingworth, and T. M. Blackman, 1997: Differential Doppler velocity: A radar parameter for characterizing hydrometeor size distributions. *J. Appl. Meteor.*, **36**, 649–663.