

A Variational Method to Retrieve the Extinction Profile in Liquid Clouds Using Multiple-Field-of-View Lidar

NICOLA L. POUNDER AND ROBIN J. HOGAN

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

TAMÁS VÁRNAI

Joint Center for Earth Systems Technology, University of Maryland, Baltimore County, Baltimore, Maryland

ALESSANDRO BATTAGLIA

Earth Observation Science, Department of Physics and Astronomy, University of Leicester, Leicester, United Kingdom

ROBERT F. CAHALAN

Climate and Radiation Laboratory, NASA Goddard Space Flight Center, Greenbelt, Maryland

(Manuscript received 5 November 2010, in final form 13 September 2011)

ABSTRACT

Liquid clouds play a profound role in the global radiation budget, but it is difficult to retrieve their vertical profile remotely. Ordinary narrow-field-of-view (FOV) lidars receive a strong return from such clouds, but the information is limited to the first few optical depths. Wide-angle multiple-FOV lidars can isolate radiation that is scattered multiple times before returning to the instrument, often penetrating much deeper into the cloud than does the single-scattered signal. These returns potentially contain information on the vertical profile of the extinction coefficient but are challenging to interpret because of the lack of a fast radiative transfer model for simulating them. This paper describes a variational algorithm that incorporates a fast forward model that is based on the time-dependent two-stream approximation, and its adjoint. Application of the algorithm to simulated data from a hypothetical airborne three-FOV lidar with a maximum footprint width of 600 m suggests that this approach should be able to retrieve the extinction structure down to an optical depth of around 6 and a total optical depth up to at least 35, depending on the maximum lidar FOV. The convergence behavior of Gauss–Newton and quasi-Newton optimization schemes are compared. Results are then presented from an application of the algorithm to observations of stratocumulus by the eight-FOV airborne Cloud Thickness from Off-Beam Lidar Returns (THOR) lidar. It is demonstrated how the averaging kernel can be used to diagnose the effective vertical resolution of the retrieved profile and, therefore, the depth to which information on the vertical structure can be recovered. This work enables more rigorous exploitation of returns from spaceborne lidar and radar that are subject to multiple scattering than was previously possible.

1. Introduction

Boundary layer clouds play an important role in the global radiation budget and yet remain one of the largest uncertainties in climate models (e.g., Randall et al. 2007) as well as being an important source of error in weather

forecasting (Martin et al. 2000). Satellite remote sensing of clouds is necessary to obtain global cloud observations. Vertical cloud profiles are very difficult to obtain observationally yet can be used to quantify subadiabatic behavior and therefore to study the role of entrainment and boundary layer parameterization.

Spaceborne lidar measurements of clouds are affected by multiple scattering of the lidar signals (Flesia and Schwendemann 1995), as are cloud radar measurements of deep convective clouds (Battaglia et al. 2010). The direct lidar return consists of a single scattering event, and

Corresponding author address: Nicola Ponder, Dept. of Meteorology, University of Reading, Earley Gate, P.O. Box 243, Reading RG6 6BB, United Kingdom.
E-mail: n.l.pounder@reading.ac.uk

the return delay is linearly related to the vertical height in the cloud where the scattering occurred. Multiply scattered returns consist of radiation that may have undergone many scattering events before being returned to the lidar receiver, often at an angle to the incoming lidar beam. The extra distance traveled between scattering events means the relationship between return delay and cloud height is no longer linear.

Multiply scattered returns potentially contain much information about cloud structure and optical depth—in particular, if observed in multiple fields of view (FOV)—but are very challenging to interpret. Bissonnette et al. (2005) successfully retrieved profiles of cloud extinction coefficient using multiple-FOV lidar. Their algorithm uses a small-angle diffusion approximation but does not include wide-angle multiply scattered returns (by “wide angle” we mean scattering from the part of the Mie phase function that is outside the narrow forward lobe that is typically only a degree or two wide) and so is limited to receivers whose footprint diameter is smaller than the scattering mean free path. To utilize the small-angle scattering while excluding wide-angle multiply scattered returns, Bissonnette et al. used a multiple-FOV lidar that detects returns sequentially from narrow fields of view with different widths. In this paper we consider lidars that receive wide-angle returns, in addition to direct and small-angle returns, in wide-angle multiple-FOV receivers that can operate simultaneously. In this paper, multiple-FOV lidar refers to these wide-angle multiple-FOV receivers and not to the narrow-angle lidar of Bissonnette et al. (2005). Davis et al. (1999) and Polonsky and Davis (2004) proposed an approach to interpret multiply scattered returns on the basis of diffusion theory and applied it to observations by a ground-based instrument (Polonsky et al. 2005). Although convenient analytical expressions were obtained, a limitation is that the shape of the extinction profile needs to be specified although the main variables of interest, cloud base and optical depth, were retrieved. Moreover, diffusion theory is only strictly valid after many scattering events, and so practically this approach is limited to interpretation of the tail of the backscatter profile, although this limitation has been partially overcome by further developments of Davis (2008). Cahalan et al. (2005) interpreted multiply scattered returns from an airborne lidar with a multiple-FOV receiver by comparison with a library of profiles that were precalculated using a Monte Carlo model, but the retrieval will always be limited by the scope of the library.

In this paper, another approach is taken: a variational algorithm (an approach also known as optimal estimation theory; Rodgers 2000) is developed in which a first guess of the extinction profile is iteratively refined on the

basis of its ability to forward model the observations. This is facilitated by a combination of the models of Hogan (2008) and Hogan and Battaglia (2008), which can estimate multiply scattered returns approximately six orders of magnitude faster than can Monte Carlo techniques and which, therefore, can be run multiple times within an iterative retrieval algorithm. The retrieved extinction profile is then the one that best forward models the returns from all available fields of view in a least squares sense. So that the retrieved profiles do not reproduce noise in the measurements, Twomey–Tikhonov regularization is employed, with the degree of smoothness optimized by performing an L-curve analysis (Hansen 1992).

Section 2 describes the retrieval method and details how the fast forward model and additional constraints may be included in the variational retrieval scheme. Section 3 studies the behavior of the retrieval method with synthetic measurements and examines the ability of the method to retrieve the vertical structure of the extinction coefficient and the total cloud optical depth. We also describe the use of averaging kernels to quantify the effective spatial resolution. The retrieval method is applied to data from the Cloud Thickness from Off-Beam Lidar Returns (THOR) instrument (Cahalan et al. 2005) in section 4. Section 5 provides a brief summary and outlook for the wider applications of this approach.

2. Retrieval method

a. Overview

The retrieval obtains a one-dimensional profile of the visible extinction coefficient α_v (at the wavelength of the lidar: c/ν) from observed profiles of apparent backscatter β at one or more different fields of view. Extinction coefficient is useful because it is directly related to optical depth. It is related to apparent backscatter using the lidar equation in the following form:

$$\beta(r) = \hat{\beta}(r) \exp\left[-2\int_0^r \alpha_v(r') dr'\right] + \beta_{\text{MS}}(r, \rho, \rho_l), \quad (1)$$

where $\hat{\beta}(r)$ is the true, unattenuated, lidar backscatter coefficient at range r and is proportional to α through the extinction-to-backscatter ratio S :

$$\hat{\beta} = \alpha_v/S. \quad (2)$$

The first term on the right of (1) describes the apparent backscatter in the absence of multiple scattering. The second term includes the contribution to apparent backscatter from multiple scattering and is dependent on the receiver FOV ρ and the lidar beam divergence ρ_l .

We use a variational retrieval scheme (Rodgers 2000) to obtain a best estimate of α at each range gate. The best estimate is obtained by minimizing a cost function,

$$J = J_{\text{obs}} + J_{\text{prior}} + J_{\text{constraint}}, \quad (3)$$

that is the sum of cost functions for observations (J_{obs}), prior constraints (J_{prior}), and additional constraints ($J_{\text{constraint}}$).

The observation part of the cost function penalizes the squared difference between the real observations β and the predicted observations β' as forward modeled from the estimated profile of extinction coefficient; that is (for N observations),

$$J_{\text{obs}} = \frac{1}{2} \sum_{i=1}^N \frac{(\beta_i - \beta'_i)^2}{\sigma_{\beta_i}^2}, \quad (4)$$

where σ_{β_i} is the standard error in β_i , which typically consists of the contribution from observational error (e.g., due to photon-counting noise) and from forward-model error. Note that the summation in (4) is over all lidar fields of view.

Generally applicable prior profiles of extinction coefficient for cloud are not available; Miles et al. (2000) observed a typical maximum extinction coefficient of 0.08 m^{-1} in stratocumulus clouds, however, and therefore we can expect $0 < \alpha_i < 0.08 \text{ m}^{-1}$. We add a Gaussian prior constraint centered at $\alpha_i^{(p)} = 0 \text{ m}^{-1}$ with a width of $\sigma_{(p),i} = 0.08 \text{ m}^{-1}$. This prior does not prevent unphysical values of $\alpha_i = 0 \text{ m}^{-1}$, and so in this algorithm we prevent negative values by forcing $\alpha_i = 0 \text{ m}^{-1}$ wherever the minimization algorithm tries to make it negative. The lack of a positivity constraint needs to be taken into account when estimating the uncertainties on the retrieved profile, and this point is discussed in section 2e. An alternative to this approach would be to use a prior distribution that excluded negative values of extinction coefficient, such as a lognormal distribution. This could be implemented by formulating the cost function in terms of the natural logarithm of the extinction coefficient.

In regions with no, or poor, observations, the prior pulls the profile to the clear-sky solution, which is a sensible assumption in the absence of information. In regions where there are observations, the prior constraint will be relatively weak and the retrieval will be dominated by the observations. The contribution of the prior to the cost function is

$$J_{\text{prior}} = \frac{1}{2} \sum_{i=1}^M \frac{[\alpha_i - \alpha_i^{(p)}]^2}{\sigma_{(p),i}^2}, \quad (5)$$

where there are M parameters to be retrieved.

Additional constraints on the retrieved state vector can be applied as an additive term in the cost function $J_{\text{constraint}}$. We use the Twomey–Tikhonov smoothness constraint introduced below in section 2b.

The cost function can be conveniently written in matrix notation as

$$J = 0.5[\mathbf{y} - H(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - H(\mathbf{x})] + 0.5[\mathbf{x} - \mathbf{x}_{(p)}]^T \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}_{(p)}] + J_{\text{constraint}}, \quad (6)$$

where \mathbf{x} is the state vector, a vector of the α_i values to be retrieved, and $\mathbf{x}_{(p)}$ is a vector of the $\alpha_i^{(p)}$ values of the prior. In addition, \mathbf{y} is the observation vector, a vector of the β_i values for all fields of view; $H(\mathbf{x})$ is the forward-model operator outlined in section 2d; and \mathbf{R} and \mathbf{B} are the error covariance matrices of the observations and the prior, respectively.

b. Smoothness constraint

Lidar measurements can be noisy, which can contaminate the retrieved extinction-coefficient profile. In their variational radar–lidar ice-cloud retrieval, Delanoë and Hogan (2008) reduced the impact of noise by including a smoothness constraint on the retrieved extinction-coefficient profile by penalizing the second derivative of the α_v profile. The constraint is (Rodgers 2000)

$$J_{\text{constraint}} = 0.5\lambda \mathbf{x}^T \mathbf{T} \mathbf{x}, \quad (7)$$

where \mathbf{T} is a Twomey–Tikhonov matrix whose elements, for example for a state vector of length 6, are

$$\mathbf{T} = \begin{pmatrix} 1 & -2 & 1 & 0 & 0 & 0 \\ -2 & 5 & -4 & 1 & 0 & 0 \\ 1 & -4 & 6 & -4 & 1 & 0 \\ 0 & 1 & -4 & 6 & -4 & 1 \\ 0 & 0 & 1 & -4 & 5 & -2 \\ 0 & 0 & 0 & 1 & -2 & 1 \end{pmatrix} \quad (8)$$

and similarly for other sizes. The constraint is weighted relative to the observations and prior information by the constant λ . Section 3f discusses how λ may be chosen.

c. Minimizing the cost function

An optimal estimate of the extinction-coefficient profile is obtained by minimizing the cost function (6) starting from an initial guess of the state vector, \mathbf{x}_1 . There are several methods that can be used to minimize

a function; we consider the Gauss–Newton and quasi-Newton methods.

The iterative Gauss–Newton method (e.g., Rodgers 2000) has been applied by a number of authors in the formulation of radar and lidar retrievals of cloud properties (e.g., Austin and Stephens 2001; Löhnert et al. 2004; Hogan 2007; Delanoë and Hogan 2008). In this approach, the forward model is linearized by making the approximation

$$H(\mathbf{x}) \simeq H(\mathbf{x}_k) + \mathbf{H}\Delta\mathbf{x}, \quad (9)$$

where $\Delta\mathbf{x} = \mathbf{x} - \mathbf{x}_k$, \mathbf{x}_k is the estimated state vector at iteration k and $H(\mathbf{x}_k)$ is the corresponding forward-modeled estimate of the observations. Here, $\mathbf{H} = \partial\mathbf{y}/\partial\mathbf{x}$ is the *Jacobian* matrix: the rate of change of each forward-modeled observation with respect to each element of the state vector. Matrix \mathbf{H} is recalculated each time the forward model is called. At each iteration of the algorithm, the new estimate of the state vector, \mathbf{x}_{k+1} , is taken to lie at the minimum of the *linearized* cost function, J_L , that is obtained by substituting (9) into (6). At this minimum, $\nabla_{\Delta\mathbf{x}}J_L = 0$, which may be rearranged to obtain

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{A}^{-1}\nabla_{\Delta\mathbf{x}=0}J_L, \quad (10)$$

where

$$\begin{aligned} \nabla_{\Delta\mathbf{x}=0}J_L = & -\mathbf{H}^T\mathbf{R}^{-1}[\mathbf{y} - H(\mathbf{x})] \\ & + \mathbf{B}^{-1}[\mathbf{x}_k - \mathbf{x}_{(p)}] + \lambda\mathbf{T}\mathbf{x}_k \end{aligned} \quad (11)$$

is a vector containing the gradient of the full cost function with respect to each element of $\Delta\mathbf{x}$ at $\Delta\mathbf{x} = 0$, and the symmetric *Hessian* matrix is given by

$$\mathbf{A} = \nabla_{\Delta\mathbf{x}=0}^2J_L = \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H} + \mathbf{B}^{-1} + \lambda\mathbf{T}. \quad (12)$$

In an operational scheme this process would be iterated until convergence as determined by a χ^2 test. As this paper is a proof of concept, we perform a fixed number of iterations and confirm that convergence has been achieved by observing only small changes in the value of the cost function for the final iterations. The convergence behavior is studied in section 3g.

A benefit of the Gauss–Newton method is that a good estimate of the error covariance matrix of the retrieved variables \mathbf{S}_x is given by the inverse of the Hessian matrix after the final iteration: $\mathbf{S}_x = \mathbf{A}^{-1}$ (Rodgers 2000). Moreover, the fact that the curvature of the cost function is available, in the form of (12), ensures that convergence is very rapid; indeed, if the full forward model is perfectly linear then the minimum can be found in one iteration of (10).

The main disadvantage of the Gauss–Newton method is that the Jacobian of the forward-modeled observations is required. If the state and observation vectors each have N elements, then \mathbf{H} is an $N \times N$ matrix and the computational cost to fill it is proportional to at least the square of N [i.e., $O(N^2)$]. As described in section 2d, the forward model we use for estimating the wide-angle multiply scattered returns already has a computational cost of $O(N^2)$, and the nature of this algorithm means that it is not possible to calculate the Jacobian more efficiently than $O(N^3)$. This is too slow for operational usage—for example, from a multiple-FOV spaceborne lidar for which we would want to process each individual profile in less than 1 s.

A solution is offered by atmospheric data assimilation systems, which cannot use the Gauss–Newton method because N is so large that the cost of computing and storing \mathbf{H} and \mathbf{A} is excessive. Most of these systems minimize the cost function using only information on its gradient and, therefore, do not calculate the Hessian. The *quasi-Newton* family of methods performs iterations using (10) but uses an approximation for \mathbf{A}^{-1} . We use the limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) method (Nocedal 1980; Liu and Nocedal 1989), which uses the cost-function gradients from a limited number of the most recent iterations to reconstruct an estimate of the curvature of the cost function. This approach was found by Gilbert and Lemaréchal (1989) to be superior to several of its competitors for large-scale problems, and their implementation of L-BFGS is currently used in the data assimilation system of the European Centre for Medium-Range Weather Forecasts. It appears from (11) that calculating $\nabla_{\Delta\mathbf{x}=0}J_L$ requires \mathbf{H} to be calculated first, which is expensive. This can be avoided by using the *adjoint method*, in which the vector $\nabla_{\Delta\mathbf{x}=0}J_{\text{obs}}$ is calculated from the gradient of the cost function with respect to each forward-modeled observation, $\nabla_{H(\mathbf{x})}J_{\text{obs}} = \mathbf{R}^{-1}[\mathbf{y} - H(\mathbf{x})]$ (also a vector), without requiring the intermediate matrix \mathbf{H} . This is achieved by coding the adjoint of the forward model (e.g., Giering and Kaminski 1998), which is typically slower to compute by approximately a factor of 3 than the original forward model but is much faster than the additional order of N in computational cost associated with computing the full Jacobian. The adjoint code calculates $\nabla_{\Delta\mathbf{x}=0}J_{\text{obs}}$ from $\nabla_{H(\mathbf{x})}J_{\text{obs}} = \mathbf{R}^{-1}[\mathbf{y} - H(\mathbf{x})]$ by performing what can be thought of as a “time reversed” sequence of linearized forms of the operations in the original forward model. The other terms in (11) take much less time to compute.

Because the L-BFGS method uses an approximation to \mathbf{A}^{-1} , more iterations are required to reach convergence than for the Gauss–Newton method. The difference in the number of iterations is typically less than the factor of approximately $N/3$ between the costs of each iteration of

the two methods, however, and therefore, for large N , L-BFGS can be much faster than Gauss–Newton to reach a solution. The difference in the number of iterations required depends on both N and the nonlinearity of the problem. In section 3g, the convergence rates are compared for retrievals using multiply scattered returns. The approximate nature of \mathbf{A}^{-1} unfortunately means that it is less accurate as an estimate of the error covariance matrix of the solution. It is also a little tricky to calculate because it is not held explicitly by the L-BFGS algorithm (see Fisher and Courtier 1995). An alternative for calculating the error covariance matrix is to perform a single iteration of the Gauss–Newton method after the L-BFGS method has converged, yielding a more accurate \mathbf{A}^{-1} .

d. Forward model and adjoint

The forward model we use for lidar returns that are subject to multiple scattering consists of the sum of the output from two fast algorithms: the photon variance-covariance (PVC) method of Hogan (2006, 2008) for the single- and small-angle-scattering contribution, and the time-dependent two-stream (TDTS) method of Hogan and Battaglia (2008) for the wide-angle contribution. The PVC method was used for lidar scattering in the combined radar–lidar algorithm of Delanoë and Hogan (2008), which employs the Gauss–Newton method to minimize the cost function. This algorithm has been applied to satellite observations of ice clouds (Delanoë and Hogan 2010), for which wide-angle scattering may be safely neglected. In wide-FOV lidar observations of stratocumulus clouds, radiation that has undergone wide-angle scattering can dominate the returned signal, necessitating the use of the TDTS method. This method involves integrating a pair of coupled partial differential equations forward in time, and its computational cost is $O(N^2)$. Hogan and Battaglia (2008) reported that the TDTS method applied to a profile of $N = 100$ points took approximately 16 ms to compute on a 1-GHz Intel Corporation processor, with the PVC method being much faster. For $N = 50$, this reduces by a factor of 4.

As outlined in section 2c, it appears not to be possible to formulate an exact Jacobian model with a cost of less than $O(N^3)$ (i.e., approximately 1.6 s per iteration for $N = 100$ and 0.2 s for $N = 50$ on a 1-GHz Intel processor). Therefore, we have coded the adjoints of both the PVC and TDTS methods so that the L-BFGS method may be applied. The adjoint for the TDTS method is exact, but that for the PVC method is approximate; it is the adjoint equivalent to the Jacobian calculation for the simple small-angle multiple-scattering model of Platt (1973) as described by Hogan (2008). Because most of the information in the retrieval comes from wide-angle scattering, the L-BFGS algorithm is still able to converge rapidly with this approximate adjoint.

e. Calculating optical depth and its error

The total optical depth down to range gate m can be calculated from the retrieved extinction-coefficient profile as

$$\delta_m = \sum_{i=1}^m \alpha_i \Delta z \quad (13)$$

$$= \mathbf{w}\mathbf{x}, \quad (14)$$

where Δz is the range-gate spacing and the row vector $\mathbf{w} = \Delta z[1, 1, \dots, 1]$ is of length m . The error variance S_{δ_m} of the optical depth to range gate m may naively be calculated as

$$S_{\delta_m} = \mathbf{w}^T \mathbf{S}_{\mathbf{x}}^{(m)} \mathbf{w}, \quad (15)$$

where $\mathbf{S}_{\mathbf{x}}^{(m)}$ is a matrix containing the first $m \times m$ elements of the full covariance matrix $\mathbf{S}_{\mathbf{x}}$. This provides a reasonable estimate of the positive uncertainty on the optical depth, but more thought is required for the negative uncertainty, which is overestimated because the prior does not include a positivity constraint. Consider a retrieved extinction-coefficient profile for an optically thick cloud in which the lidar has been completely attenuated. Near cloud top, the retrieval is dominated by the observations and the error on the retrieved extinction is dominated by the observation errors. Toward cloud base, there is no information from observations and the retrieval and error are dominated by the prior and the large prior uncertainty. The optical depth and associated error at each range gate can be calculated using (14) and (15). At each range gate we can consider $\delta_m - S_{\delta_m}$ to be a minimum bound on the total optical depth to that range gate; as the retrieved extinction-coefficient error becomes large toward cloud base, however, so does S_{δ_m} . In some cases the estimated “minimum bounds” at each range gate increase as we descend into the cloud until the retrieval uncertainties become large, at which point the minimum bounds decrease again toward cloud base. We cannot have less knowledge about the minimum extent of cloud optical depth as we descend into the cloud, and therefore we use the maximum minimum bound δ_{\min} to recalculate the negative uncertainty on the retrieved optical depth below that point as $S_{\delta_m}^- = \delta_m - \delta_{\min}$. This results in asymmetric errors.

3. Studies with synthetic data

a. Retrieving a triangular extinction-coefficient profile

The general behavior of the retrieval algorithm has been studied using synthetic measurements. The expected observations (in the absence of noise) for a chosen α_r profile are simulated using the forward model described

above. Random, normally distributed fluctuations are added to the simulated observations to represent instrument noise, with a variance that is characteristic of a photon counter with a sensitivity comparable to the THOR receiver to be described in section 4.

Two extinction-coefficient profiles are used: a triangular profile (cloud top at 1600 m and cloud base at 705 m) and a sinusoidal profile (cloud top at 1600 m and cloud base at 400 m). The triangular profile is chosen to represent an adiabatic cloud, and the sinusoidal profile is chosen to test the sensitivity of the method to a highly structured cloud. The synthetic data use 540-nm lidar with a 325- μ m beam divergence at the $1/e$ level. The lidar is at an altitude of 7980 m. The lidar receiver has up to three fields of view: a central, circular FOV with a footprint of 10 m at ground (1.25-mrad full-width FOV) and two, concentric, annular fields of view whose outer limits encompass footprints of 100 and 600 m at ground (12.53- and 75.19-mrad full-width fields of view, respectively). In each case the receiver has a top-hat pattern.

Figure 1 shows an example of a retrieval using a triangular extinction-coefficient profile with a total optical depth of 40. Triangular profiles of liquid water content are commonly observed (e.g., Slingo et al. 1982), indicating an extinction coefficient with an approximately triangular profile as well. For this retrieval, $\lambda = 10^5$ was chosen using the method described in section 3f below. We use scattering properties that are suitable for liquid droplets: an asymmetry factor of 0.85, single-scattering albedo of 1, lidar ratio $S = 18.5$ sr (Pinnick et al. 1983; O'Connor et al. 2004), and droplet equivalent-area radius of 10 μ m (required by the PVC method for calculating the width of the forward-scattering lobe). All are kept constant with height and are the same for simulating the synthetic data and for the retrieval. The lidar range-gate spacing is 30 m. Figure 1a shows the true extinction-coefficient profile and retrieved profiles using the three-FOV receiver and using the central field alone. The error bars are the square root of the diagonal of the retrieval error covariance matrix \mathbf{S}_x . The simulated observations are shown in Fig. 1b along with the corresponding values forward modeled from the extinction profile retrieved using all three fields of view. The retrieval using only a narrow FOV underestimates α_p because the signal is rapidly attenuated. The retrieved total optical depth using only the narrow FOV is $11.4^{+16}_{-3.2}$.

Using all three fields of view, the retrieved total optical depth is $35.4^{+13}_{-2.7}$, in much better agreement with the true optical depth. The retrieved extinction profile is underestimated near cloud top and base although the slope is generally well reproduced and the extinction goes to zero near cloud base. The retrieved extinction coefficient is accurate to 15% down to 1000 m (35 optical depths).

Above this height the retrieval is dominated by the observations, but below it the constraint provided by the observations is weaker and the prior becomes relatively more important. Information about the relative contributions of parts of the cost function is contained in the averaging kernel matrix described in section 3b below. In section 3d we shall see that the total optical depth of this example is about five optical depths greater than can be retrieved using this receiver configuration, and so we expect the profile to be underestimated in this case.

The synthetic data and the retrieval algorithm use the same forward model. These studies demonstrate how the retrieval behaves with a perfect forward model. To study the effect of our forward model on the retrieved extinction profile and total optical depth, we perform the retrieval on simulated observations, for the same extinction profile, generated using the Monte Carlo simulator of Battaglia et al. (2006), which is the same as that used by Hogan and Battaglia (2008) to test the forward model. The retrieved profile is shown in Fig. 1a. The retrieved total optical depth is $32.3^{+13}_{-2.7}$, which agrees with the total optical depth retrieved from the idealized synthetic data. The extinction coefficient in the lower part of the cloud is well retrieved, but the retrieval does not reproduce cloud top as well. This appears to be because the forward model does not perfectly model the backscatter peak at the transition between the parts of the profile for which narrow- and wide-angle scattering are dominant. Nonetheless, the algorithm is good for retrieving the total optical depth and the extinction-coefficient gradient toward cloud base. The studies with the synthetic data illustrate the potential of a variational retrieval scheme for liquid clouds and that the current forward model can still be improved.

b. The information content of a retrieval

The averaging kernel matrix $\mathbf{W} (= \partial \mathbf{x}_{\text{retrieved}} / \partial \mathbf{x}_{\text{truth}})$ describes the way the observing system smooths the profile. It is given by Rodgers (2000) as

$$\mathbf{W} = (\nabla^2 J)^{-1} [\mathbf{y} - H(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - H(\mathbf{x})]. \quad (16)$$

Because the state vector represents a profile, the rows, \mathbf{a}_i^T of \mathbf{W} are averaging kernels or smoothing functions, one for each point in the extinction-coefficient profile. If the inverse method were perfect, \mathbf{W} would be a unit matrix. In reality, the averaging kernels are functions peaked at their associated range gate with a half-width that is a measure of the spatial resolution of the observing system. The area of the averaging kernel, calculated as $\mathbf{a}_i^T \mathbf{u}$, where \mathbf{u} is a column vector of unit elements, can be considered to be a rough measure of the fraction of the retrieval, at gate i , that comes from the observations rather than from the prior or additional constraints.

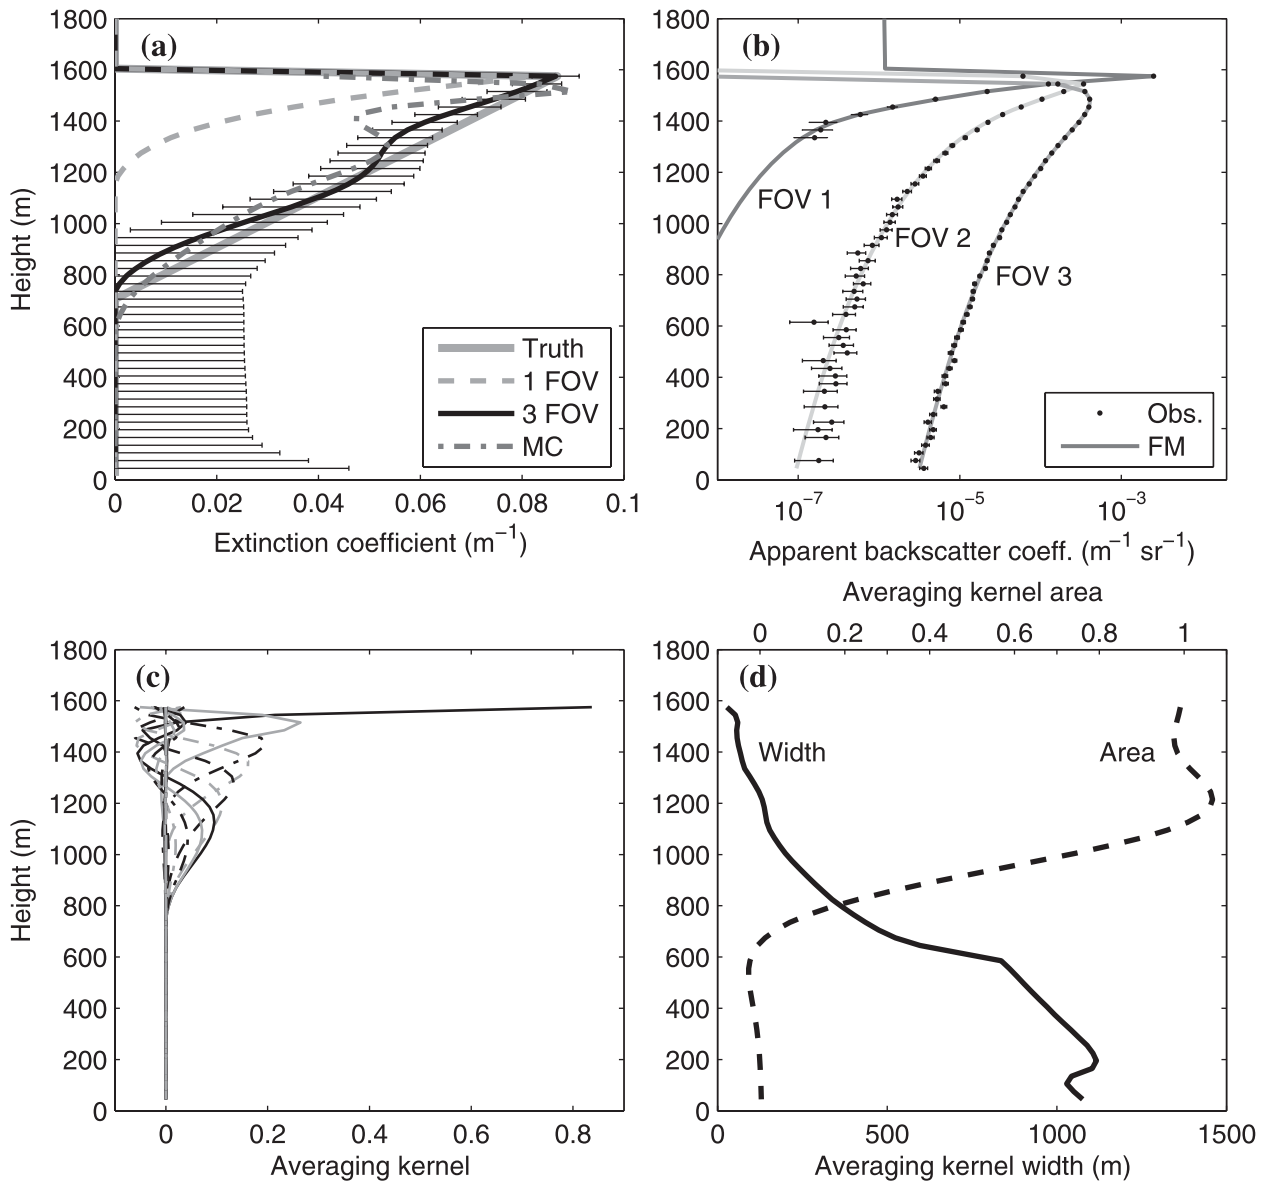


FIG. 1. Example results from a retrieval of an idealized triangular extinction-coefficient profile. (a) The true extinction-coefficient profile (“Truth”) and profiles retrieved from synthetic observations using two different receiver configurations: a single-FOV receiver with a 10-m footprint (1 FOV) and a three-FOV receiver with a 10-m central FOV and two concentric annular fields of view encompassing 100- and 600-m footprints (3 FOV). Also shown is the profile retrieved from Monte Carlo generated observations using the three-FOV receiver (MC). (b) The observed (Obs.) apparent backscatter coefficients for each of the fields of view (points with error bars) and the forward-modeled (FM) observations (lines) for the extinction-coefficient profile retrieved from the synthetic observations using all three fields of view: FOV 1 (central FOV), FOV 2, and FOV 3 (widest FOV). (c) The averaging kernels for the retrieval of the synthetic observations that used the three-FOV receiver. For clarity, only every third kernel is plotted. (d) The area (top scale) and width (bottom scale) of each averaging kernel.

Figure 1c shows some of the averaging kernels for the three-FOV retrieval. For clarity, only every third kernel is shown. The first kernel is strongly peaked at the first range gate, which indicates very good spatial resolution at cloud top. Further into the cloud, the spatial resolution worsens and the kernels broaden. The width and area of the kernels are shown in Fig. 1d. The kernel width is approximated as

$$\sigma_{a,i} = \left[\frac{\sum_{j=0}^M W_{ij}(z_i - z_j)^2}{\sum_{j=0}^M W_{ij}} \right]^{1/2}. \quad (17)$$

The kernel width is approximately equal to the range-gate spacing for the first kernel, broadening to 2 times




FIG. 2. Performance of the retrieval algorithm for an extinction-coefficient profile with sinusoidal structure. (a) As in Fig. 1a, but without the MC line. (b) As in Fig. 1b. (c) As in Fig. 1d.

that, around 60–70 m, down to a height of about 1300 m, where the kernel width increases to about 100 m, before rising rapidly below a height of about 1100 m as the retrieval loses spatial resolution. In physical terms, the process of multiple scattering means that information on vertical structure is lost as more scattering events have taken place. The dip near ground is an artifact of the truncated kernels at the base of the retrieval. The areas of the kernels are approximately 1 from 1600 m down to about 1100 m, and therefore the retrieval is dominated by the observations over this range. From 1100 m, the areas smoothly drop to near zero at about 700 m. In this region the retrieval is only weakly determined by the observations and has a significant contribution from the prior, which is why the extinction coefficient is underestimated in this region. The smoothing constraint has had the effect of smoothing the information from the observations across that range. In a real case in which the true extinction coefficient is not known, the averaging kernels show which

regions of the retrieval are only weakly determined by observations and are therefore less trustworthy.

c. Retrieval of a sinusoidal extinction profile

Figure 2 shows another simulated retrieval for a cloud with exaggerated vertical structure to determine to what extent the information on vertical structure is smoothed out after many scattering events. The total optical depth of this profile is 14.4, and the optical depth between troughs is about 2.4. The results of retrievals using two different receiver configurations are shown. As with the triangular profile, when using only the narrowest FOV it is possible to locate the cloud top but it is not possible to retrieve the extinction-coefficient profile with any accuracy. Using the three-FOV receiver, the structure of the profile can be retrieved to about 6 optical depths (three peaks). Below this, although the structure is no longer retrieved, it is still possible to constrain the total optical depth and the extinction coefficient goes to zero near true




FIG. 3. Retrieved optical depth for synthetic, triangular extinction-coefficient profiles with different optical depths. Each retrieval is performed for three different lidar receiver configurations: a single-FOV receiver with a 10-m footprint, a single-FOV receiver with a 100-m footprint, and a three-FOV receiver with 10-, 100-, and 600-m footprints for each FOV. The shaded regions indicate the central 60% of the retrieved optical depths for 100 retrievals of independent observations including instrument noise (not shown for retrievals using the 100-m footprint receiver). The lines with error bars are the retrieved optical depths for observations with no instrument noise. The insert contains an enlarged view for 0–10 optical depths.

cloud base. The retrieved total optical depth is $13.3^{+8.3}_{-1.3}$. The true optical depth down to a height of 1005 m, where structure is retrieved, is 7.2, and the retrieved optical depth to this height is $7.1^{+6.4}_{-0.2}$.

d. Optical depth

Figure 1 showed an example retrieval for one extinction profile. We have repeated the retrieval for a number of triangular extinction-coefficient profiles, with different optical depths, for each of the lidar receiver configurations described in section 3a. For each profile, the cloud-top height was 1600 m and the gradient da/dz was 10^{-4} m^{-2} . The physical thickness and peak extinction coefficient were varied to change the total cloud optical depth. For each extinction profile we simulated one set of observations without instrument noise and 100 sets with instrument noise. For the simulated observations that do not include instrument noise, we do assign a measurement error that is consistent with what we would expect for noisy observations. Figure 3 shows the retrieved total optical depth as a function of input total optical depth. The lines with error bars are the retrievals for observations without noise, and the shaded regions indicate the central 60% of retrieved optical depths for the observations with noise.

For each of the receiver configurations, the true optical depth for the idealized, noise-free, observations is well retrieved for small optical depths and then, after some point, the retrieved optical depth levels off. Using a single FOV with a 10-m footprint, the optical depth can be retrieved up to about 2 with a negative error of about 0.3 before this “saturation” effect occurs. At this point, the positive error has saturated and is about 14. This confirms the common “rule of thumb” that a narrow-FOV lidar contains useful information only in the first two or three optical depths of a cloud. Use of a wider FOV with a 100-m footprint [comparable to *Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations (CALIPSO)*; Winker et al. 2004] allows retrieval of optical depth up to about 25. The three-FOV receiver can retrieve optical depths up to about 35 before the retrieved optical depth begins to level off. Instrument noise increases the retrieved optical depth in retrievals using the three-FOV receiver for optical depths up to about 20. The maximum optical depth retrievable by this method will depend on the fields of view of the lidar and the forward model. In principle, increasing the width of the widest FOV will increase the maximum total optical depth that can be retrieved as long as the assumptions in the forward model are still valid.




FIG. 4. The sensitivity of the retrieved optical depth to the assumed parameters. (a) Retrieved optical depth as a function of true optical depth for 10 values of r_e between 5 and 15 μm . (b) As in (a), but for five values of g between 0.845 and 0.867.

The standard deviation of the retrieved optical depths for observations with instrument noise is indicative of the statistical uncertainty on the measurements. The spread in the retrieved optical depth due to this noise is significantly reduced when using the three-FOV receiver as compared with the narrow-FOV receiver alone. The statistical spread agrees well with the negative error bars on the noise-free retrievals at small optical depths where the retrieval is dominated by statistical uncertainty. The positive errors in optical depth include the uncertainty in the retrieval that is due to the lack of observations once the lidar is completely attenuated. The magnitude of this uncertainty is determined by the uncertainty on the prior, and the positive errors are large where they become dominated by the prior uncertainty. Above approximately 40 optical depths, the peak extinction coefficient has become unphysically large and inconsistent with the prior uncertainty of 0.08 m^{-1} , and therefore we do not expect the positive errors on the retrieved optical depth to incorporate the true optical depth because these are unphysical clouds.

e. Sensitivity to input parameters

In the studies described above, the droplet effective radius r_e and scattering asymmetry parameter g that are assumed by the retrieval have been chosen to match those used to generate the simulated observations. In this section we demonstrate the sensitivity of the retrieval algorithm to the choice of these parameters. Varying r_e has

two effects on the retrieval. First, g is dependent on r_e , affecting the depth to which wide-angle multiply scattered photons penetrate into the cloud. Second, the width of the forward lobe of the phase function varies inversely with r_e , affecting the degree of small-angle multiple scattering in the first few optical depths in the cloud. We separate the two effects by varying g independent of the value of r_e in the forward lobe. Figure 4a shows the retrieved optical depth as a function of true optical depth for the instrument-noise-free synthetic data in Fig. 3 but for 10 input values of r_e in the forward lobe between 5 and 15 μm ($r_e = 10 \mu\text{m}$ was used to simulate the observations). The spread of the retrieved optical depth for the three-FOV receiver and the 10-m-footprint receiver are small relative to the statistical uncertainties on the retrieval in Fig. 3, and therefore the uncertainty on the retrieved optical depth due to assuming r_e to calculate the width of the forward lobe of the phase function is negligible.

Figure 4b is the same as Fig. 4a except that r_e is fixed at 10 μm and five different input values of g , between 0.845 and 0.867, are used in the retrievals ($g = 0.850$ was used to simulate the data). That range corresponds to the change in g when varying r_e between 5 and 15 μm . The effect of varying g is small for the narrow FOV, which is dominated by single scattering. Varying g has a considerable effect on the retrieval for the three-FOV receiver. At a true total optical depth of 35 (the limit of this receiver's ability to retrieve optical depth), the retrieved optical

depth varies by ~ 5 as g is varied between 0.845 and 0.867. The uncertainty associated with assuming g is reduced as the true optical depth decreases.

f. Optimizing the smoothness constraints

Section 2b introduced a smoothness constraint to the retrieval. The constraint is weighted relative to the rest of the cost function by a constant λ that should be optimized to smooth out contributions from noise without smoothing away true structure. Hansen (1992) proposed an L-curve analysis as a computational tool for choosing this regularization parameter. The so-called L-curve is a plot of $\mathbf{x}^T \mathbf{T} \mathbf{x}$ ($= 2\lambda^{-1} J_{\text{constraint}}$), a measure of the sum of the squared curvature of the solution, against the norm of the residual of the forward model to observations ($= 2J_{\text{obs}}$) for different values of λ . Figure 5 shows an example L-curve for retrievals of a sinusoidal α_ν profile using a three-FOV receiver (i.e., the black curve in Fig. 2a). The L-curve demonstrates the balance between information coming from the observations and the constraint. In the top-left vertical part of the curve ($\lambda < 10^4$), λ is small and the retrieval is dominated by the observations. Large values of $J_{\text{constraint}}$ indicate that the retrieval is noisy. Increasing λ in this region increases the smoothness of the solution while having little impact on the residual from observations. The smoothness of the solution has been increased by removing fluctuations in the retrieved solution that are due to noise in the observations. Beyond $\lambda = 10^5$, the curve becomes less steep. In this region, increasing λ pulls the forward-modeled solution away from the observations, increasing the residual, while having little further impact on the smoothness of the solution. As λ is increased beyond $\lambda = 10^6$, the L-curve becomes vertical again. Here, the retrieval is dominated by the constraint and the smoothness of the retrieved solution is increased by removing the true structure. In this particular example there are two scales of structure: 1) small-scale structure due to instrument noise that is not in the true extinction-coefficient profile and that varies on the scale of the spacing between neighboring points in the profile and 2) the true, large-scale structure of the original sine curve. The optimal choice for λ is at the heel of the L-curve around $\lambda = 10^5$. At this point the noise has been reduced in the retrieved solution without significantly increasing the residual of the solution.

g. Convergence of different minimization methods

Section 2c introduced two methods for minimizing the cost function: the Gauss–Newton and L-BFGS methods. We stated that the Gauss–Newton method is quick to converge for a linear system but requires the Jacobian matrix of the forward model to be calculated, which is computationally expensive. In contrast, the L-BFGS




FIG. 5. Example L-curve for retrievals of the sinusoidal extinction profile in Fig. 2a using the Twomey–Tikhonov smoothness constraint. The retrieval is repeated, applying a different weighting λ to the smoothness constraint. The optimum choice for λ lies in the heel of the curve.

method does not require the Jacobian but may require more iterations to reach convergence.

Figure 6 compares the convergence rates of the two methods when retrieving the triangular extinction profile in Fig. 1. The starting point for the minimization was $\ln(\alpha_i) = -4$ (with α_i in reciprocal meters) for both methods, and for these studies the cost function was formulated in terms of the natural logarithm of extinction coefficient and apparent backscatter. The retrieved extinction-coefficient profiles agreed within errors. The minimization was halted when the change in the cost function between iterations was less than 10^{-4} . This happened after 15 Gauss–Newton iterations but required 135 quasi-Newton iterations: a factor-of-9 difference. As discussed in section 2c, each quasi-Newton iteration is around $N/3$ times as fast as a Gauss–Newton iteration for observation and state vectors containing N elements. The retrieval in Fig. 6 has a 53-element state vector and a 135-element observation vector, and therefore each quasi-Newton iteration is at least 18 times as fast as a Gauss–Newton iteration, making the quasi-Newton convergence faster than the Gauss–Newton option despite the extra iterations required.

4. Application to THOR data

The method has been applied to airborne 540-nm lidar measurements collected over Oklahoma on 25 March 2002




FIG. 6. Convergence rates for the Gauss-Newton and L-BFGS minimization methods for the triangular profile in Fig. 1. The minimization was halted when the change in the cost function was less than 10^{-4} .

by the THOR instrument (Cahalan et al. 2005). During the flight, the region was covered by low-level stratus cloud plus a layer of highly variable but optically thin cirrus clouds at about 5.5 km above the ground. The retrieval has been designed for a single layer of liquid water cloud and therefore has been tested on profiles from THOR taken in regions with very thin cirrus. The THOR receiver has eight fields of view: one circular, central FOV and seven concentric annular rings. The receiver fields of view are summarized in Table 1.

Figure 7b shows a sample THOR observation of apparent backscatter. The altitude of the lidar was 8046 m for these observations. At this altitude the footprint of the central FOV at ground level is 7 m while the widest FOV encompasses a footprint of 859 m. The central FOV has an intense signal at cloud top that is quickly attenuated. The signals from the wider fields of view are weaker and broader because they originate from the multiply scattered radiation coming from deep within the cloud. The observed apparent backscatter, above 1200 m, in the second and third fields of view is likely caused by cross talk with the first FOV. The raw THOR observations are photon counts and must be calibrated and converted to apparent backscatter. The central FOV is calibrated by matching the molecular return from above cloud to the Rayleigh scattering predicted from a nearby radiosonde ascent. The signals from the other fields of view are calibrated relative to the central FOV and each other using the approach of

TABLE 1. Fields of view of the THOR receiver channels. The widest FOV is divided into three segments (channels 8–10) that are merged before being used in the retrieval.

Channel	FOV (full angle; mrad)
1	0.000–0.840
2	1.029–1.681
3	1.681–3.361
4	3.361–6.723
5	6.723–13.40
6	13.40–26.72
7	26.72–53.40
8–10	53.40–106.7

Cahalan et al. (2005). The raw observations include random background noise the amplitude of which is independent of apparent height. The mean magnitude and standard deviation of the background are estimated from the apparent measured signal at the range gates below the ground. To remove the background, this mean is subtracted from the signal at each range gate, and observations within 4 standard deviations of 0 are removed. The observation errors are estimated from the raw photon counts by assuming Poisson statistics.

The smoothness constraint λ used in the retrieval of the THOR data was optimized using an L-curve analysis as shown in Fig. 8. The THOR observations have structure on several length scales. Reducing λ below 10^3 has no effect on the retrieval. Increasing λ to 10^5 increases the smoothness of the retrieved extinction profile with only a small effect on the residual of the observations, and this is chosen as the optimal value. The L-curve does not level out, because increasing λ continues to remove structure on larger and larger length scales.

Figure 7a shows the retrieved extinction-coefficient profile and associated errors. The total retrieved optical depth is $16.4^{+9.1}_{-0.6}$. Cahalan et al. (2005) retrieved physical cloud thickness, but not optical depth, for this profile using THOR's three outermost fields of view. They obtained a thickness of 560 ± 20 m. Using the narrowest FOV to infer cloud top, this implies that cloud base is at 528 m, which is consistent with where our retrieved extinction coefficient approaches zero.

The widths and areas of the averaging kernels are summarized in Fig. 7c. The areas indicate that the retrieval is dominated by the observations down to an altitude of about 500 m. The averaging-kernel areas show the resolution of the retrieval gradually increasing to about 120 m, at an altitude of 540 m, before increasing more rapidly as the prior and smoothness constraints become relatively more important.

The apparent backscatter coefficient profiles as forward modeled from the retrieved extinction profile are shown in Fig. 7. The observations and forward model




FIG. 7. Observations of apparent backscatter and retrieval of extinction coefficient for a profile observed by THOR. (a) The extinction-coefficient profiles retrieved using observations from all fields of view (all FOV; black line) and retrieved excluding observations from the narrow FOV (gray line). (b) The observed apparent backscatter coefficient from each of the eight fields of view (points with error bars) and the apparent backscatter forward modeled from the retrieved extinction profile in (a) (FOV 1 is the central FOV, and FOV 8 the widest). (c) The area (top scale) and width (bottom scale) of each averaging kernel.

agree well for the tails of the wide fields of view; the forward model overestimates the peak of the narrowest FOV by approximately a factor of 4, however. There is evidence that this is due to a nonlinearity in the response of the first FOV of THOR: this channel has been calibrated by matching the molecular scattering above the stratocumulus with what would be expected from the air density, but the results are not consistent with the calibration method of O'Connor et al. (2004) that instead uses the signal at the upper end of the dynamic range of the receiver. O'Connor et al. showed that the vertically integrated attenuated backscatter observed in an optically thick liquid water cloud by a lidar with an FOV narrow enough that it only detects single-scattered returns is $1/(2S) \approx 0.027 \text{ sr}^{-1}$ (assuming a lidar ratio $S = 18.5 \text{ sr}$). We

expect the first FOV of the THOR lidar to also detect photons that have been forward scattered by the narrow forward lobe in the Mie phase function, resulting in an expected integrated backscatter of $1/S \approx 0.054 \text{ sr}^{-1}$. This is almost exactly the value calculated from the forward-modeled backscatter for the first FOV. The corresponding value for the THOR measurements in this channel is only 0.019 sr^{-1} , however. Recalibration of this channel by multiplying by 2.8 would yield the expected integrated backscatter but would be inconsistent with the molecular calibration. Therefore, it seems likely that the signal measured in this channel has suffered from saturation.

It is possible that this saturation has affected the accuracy of the retrieved extinction profile. To investigate




FIG. 8. The L-curve for retrievals, using the Twomey–Tikhonov smoothness constraint, of a profile observed by the THOR experiment.

this, we repeated the retrieval but excluded measurements from the first FOV. The result is shown by the gray line in Fig. 7a. The retrieved extinction profile is now much smoother near cloud top, and the trough is no longer present. The gradient of the tail is unchanged, as expected, because information about cloud base can only come from the wide fields of view. The apparent backscatter forward modeled from the retrieved extinction profile agrees well with the observations at the peak as well as the tail for all but the widest FOV, where the forward model has a peak that is too low and too large. The retrieved total optical depth is $18.2^{+9.6}_{-0.6}$, only 1.8 optical depths larger than was retrieved when using all fields of view. The incompatible narrow FOV affects the shape of the retrieved profile, but the wide fields of view are able to constrain the total optical depth despite this influence. This is because the wide fields of view tightly constrain the total optical depth and the gradient at cloud base while the narrow FOV constrains cloud top. The trough at around 1000 m is an artifact of the retrieval trying to reconcile these two constraints when the narrow FOV is inconsistent with the others.

Discrepancies between the narrow and wide fields of view could also be introduced by the presence of cirrus and also by horizontal inhomogeneities in the stratocumulus cloud. High-altitude cirrus will predominately attenuate the direct return but not the multiply scattered return. The cirrus will also scatter parts of the downwelling lidar pulse into a wider cone, increasing the returns in

the wider fields of view. Although there was some high cirrus in this profile, it was optically thin and cannot explain the majority of the observed differences. We have also assumed that the stratus cloud is horizontally homogeneous. The profile used here was averaged over 500 lidar pulses. With accounting for the aircraft speed, the central FOV sampled a $7\text{ m} \times 70\text{ m}$ area whereas the widest FOV samples over a 1-km square area, and therefore sampling errors could cause differences between the fields of view. In this case, however, saturation of the narrow FOV is likely the dominant effect.

5. Conclusions and discussion

A variational retrieval scheme has been described that retrieves extinction-coefficient profiles in liquid water clouds from wide-FOV lidar measurements. This scheme incorporates the fast multiple-scattering models of Hogan (2008) and Hogan and Battaglia (2008). Using this technique, we have demonstrated that it is possible to retrieve cloud vertical structure down to around 6 optical depths with a perfect forward model and to retrieve total cloud optical depth up to about 35 optical depths, as compared with only 2 optical depths when using a conventional narrow-FOV lidar. Larger optical depths than 35 are likely to be retrievable for lidars the receivers of which have a wider FOV than is considered in this paper. The retrieval has been tested on Monte Carlo and idealized synthetic data with realistic instrument noise, as well as on real observations by the airborne THOR instrument.

It has been necessary to make several assumptions so that the retrieval problem is tractable. The forward model assumes that the cloud is horizontally homogeneous on the scale of the footprint of the receiver with the largest FOV, making the algorithm more applicable to stratiform liquid clouds than to cumulus. We have also assumed values for the lidar ratio S , single-scattering albedo; and asymmetry factor. In practice these are not significant weaknesses; S is only relevant for the narrowest FOV and in any case is approximately constant for the range of droplet size distributions in stratocumulus. Likewise, the single-scattering albedo, important for calculating the multiply scattered returns, is approximately constant over the droplet size range of interest. The asymmetry factor can have a significant effect for multiply scattered returns in optically thick clouds, although this is also approximately constant over the droplet size range of interest.

We hope that the ability to routinely and rapidly interpret returns from multiple-FOV lidar in liquid clouds encourages the development of more such instruments. In particular, this technology is perfectly suited to a satellite platform, which would enable global measurements

of the extinction profile of liquid clouds for the first time. The dependence of spaceborne lidar returns on FOV is highlighted by comparing the profiles from the 1994 Lidar In-Space Technology Experiment (LITE; Winker et al. 1996), with those from *CALIPSO*. LITE, with its FOV of up to 1 km, frequently exhibited returns with so much multiple scattering that they appeared to originate from below the surface (Miller and Stephens 1999), whereas the effect in *CALIPSO*, with its 90-m FOV, is only apparent on close inspection of the data. Idealized retrievals, such as those in Figs. 1 and 2, can be used to optimize the design of lidar receivers. In particular, a study of total retrieved optical depth, as in section 3d, could shed light on the optimum receiver footprint, and the effect of increasing the number of fields of view could be tested in retrievals with exaggerated vertical structure, such as in section 3c. Such studies are beyond the scope of this paper.

We have concentrated on retrievals of cloud total optical depth and vertical cloud structure, but cloud geometric thickness is also of interest. For cloud retrievals for which the total optical thickness can be retrieved by the algorithm, the height of cloud base can be estimated as the height at which the retrieved extinction coefficient goes to zero. It is not possible to retrieve confidently the position of the cloud base with the current algorithm in very optically thick clouds. This could be achieved by including in the cost function physically based constraints (e.g., including cloud adiabaticity), however, and we intend to include such a constraint in the future.

The technique is most powerful when applied to lidars equipped with multiple-FOV receivers, but we have demonstrated that useful extra information on cloud optical depth is available even with a single-, wide-FOV receiver. Both the spaceborne *CALIPSO* lidar and the *CloudSat* radar are affected by multiple scattering in optically thick media (liquid clouds in the case of *CALIPSO* and deep convective clouds in the case of *CloudSat*), but since only a single FOV is available there is a limit to what can be retrieved unambiguously from a single instrument. We are therefore currently developing a comprehensive variational cloud and precipitation retrieval that exploits multiple instruments simultaneously; in this scheme, the additional constraints necessary to interpret multiple-scattered radar or lidar signals come not from extra fields of view but from other instruments such as radiometers. This is a natural extension to the radar, lidar, and radiometer retrieval scheme of Delanoë and Hogan (2008) in which small-angle lidar multiple scattering was accounted for by the Hogan (2006) forward model.

Acknowledgments. This work benefited from the support of NERC's National Centre for Atmospheric Science

(Grant NE/H003894/1) and the European Space Agency (Contract 22442/09/NL/CT). Marta Janiskova and Dong Huang are thanked for useful discussions on adjoint coding and L-curve analysis, respectively.

REFERENCES

- 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.
- Battaglia, A., M. O. Ajewole, and C. Simmer, 2006: Evaluation of radar multiple scattering effects from a GPM perspective. Part I: Model description and validation. *J. Appl. Meteor. Climatol.*, **45**, 1634–1647.
- , S. Tanelli, S. Kobayashi, D. Z. R. J. Hogan, and C. Simmer, 2010: Multiple-scattering in radar systems: A review. *J. Quant. Spectrosc. Radiat. Transfer*, **111**, 917–947.
- Bissonnette, L. R., G. Roy, and N. Roy, 2005: Multiple-scattering-based lidar retrieval: Method and results of cloud probings. *Appl. Opt.*, **44**, 5565–5581.
- Cahalan, R. F., M. McGill, J. Kolansinski, T. Várnai, and K. Yetzer, 2005: THOR—Cloud thickness from off-beam lidar returns. *J. Atmos. Oceanic Technol.*, **22**, 605–627.
- Davis, A. B., 2008: Multiple-scattering lidar from both sides of the clouds: Addressing internal structure. *J. Geophys. Res.*, **113**, D14S10, doi:10.1029/2007JD009666.
- , R. F. Cahalan, J. D. Spinhirne, M. J. McGill, and S. P. Love, 1999: Off-beam lidar: An emerging technique in cloud remote sensing based on radiative Green-function theory in the diffusion domain. *Phys. Chem. Earth*, **24B**, 177–185; Erratum, **24B**, 757–765.
- Delanoë, J., and R. J. Hogan, 2008: A variational scheme for retrieving ice cloud properties from combined radar, lidar, and infrared radiometer. *J. Geophys. Res.*, **113**, D07204, doi:10.1029/2007JD009000.
- , and —, 2010: Combined *CloudSat*–*CALIPSO*–*MODIS* retrievals of the properties of ice clouds. *J. Geophys. Res.*, **115**, D00H29, doi:10.1029/2009JD012346.
- Fisher, M., and P. Courtier, 1995: Estimating the covariance matrices of analysis and forecast error in variational data assimilation. ECMWF Tech. Memo. 220, 28 pp.
- Flesia, C., and P. Schwendimann, 1995: Special section on Multiple Scattering in Lidar Experiments (MUSCLE). *Appl. Phys.*, **B60**, 315–362.
- Giering, R., and T. Kaminski, 1998: Recipes for adjoint code construction. *ACM Trans. Math. Softw.*, **24**, 437–474.
- Gilbert, J. C., and C. Lemaréchal, 1989: Some numerical experiments with variable-storage quasi-Newton algorithms. *Math. Programming*, **45**, 407–435.
- Hansen, P. C., 1992: Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Rev.*, **34**, 561–580.
- Hogan, R. J., 2006: Fast approximate calculation of multiply scattered lidar returns. *Appl. Opt.*, **45**, 5984–5992.
- , 2007: A variational scheme for retrieving rainfall rate and hail intensity from polarization radar. *J. Appl. Meteor. Climatol.*, **46**, 1544–1564.
- , 2008: Fast lidar and radar multiple-scattering models. Part I: Small-angle scattering using the photon variance–covariance method. *J. Atmos. Sci.*, **65**, 3621–3625.
- , and A. Battaglia, 2008: Fast lidar and radar multiple-scattering models. Part II: Wide-angle scattering using the time-dependent two-stream approximation. *J. Atmos. Sci.*, **65**, 3636–3651.

- Liu, D. C., and J. Nocedal, 1989: On the limited memory method for large scale optimization. *Math. Programming*, **45B**, 503–528.
- Löhnert, U., S. Crewell, and C. Simmer, 2004: An integrated approach toward retrieving physically consistent profiles of temperature, humidity, and cloud liquid water. *J. Appl. Meteor.*, **43**, 1295–1307.
- Martin, G. M., M. R. Bush, A. R. Brown, A. P. Lock, and R. N. B. Smith, 2000: A new boundary layer mixing scheme. Part II: Tests in climate and mesoscale models. *Mon. Wea. Rev.*, **128**, 3200–3217.
- Miles, N. L., J. Verlinde, and E. E. Clothiaux, 2000: Cloud droplet size distributions in low-level stratiform clouds. *J. Atmos. Sci.*, **57**, 295–311.
- Miller, S. D., and G. L. Stephens, 1999: Multiple scattering effects in the lidar pulse stretching problem. *J. Geophys. Res.*, **104**, 22 205–22 219.
- Nocedal, J., 1980: Updating quasi-Newton matrices with limited storage. *Math. Comput.*, **35**, 773–782.
- O'Connor, E. J., A. J. Illingworth, and R. J. Hogan, 2004: A technique for autocalibration of cloud lidar. *J. Atmos. Oceanic Technol.*, **21**, 777–786.
- Pinnick, R., S. Jennings, P. Chlek, C. Ham, and W. Grandy Jr., 1983: Backscatter and extinction in water clouds. *J. Geophys. Res.*, **88**, 6787–6796.
- Platt, C. M. R., 1973: Lidar and radiometric observations of cirrus clouds. *J. Atmos. Sci.*, **30**, 1191–1204.
- Polonsky, I. N., and A. B. Davis, 2004: Lateral photon transport in dense scattering and weakly-absorbing media of finite thickness: Asymptotic analysis of the space-time green function. *J. Opt. Soc. Amer.*, **21A**, 1018–1025.
- , S. P. Love, and A. B. Davis, 2005: Wide-angle imaging lidar deployment at the ARM Southern Great Plains site: Intercomparison of cloud property retrievals. *J. Atmos. Oceanic Technol.*, **22**, 628–648.
- Randall, D. A., and Coauthors, 2007: Climate models and their evaluation. *Climate Change 2007: The Physical Science Basis*, S. Solomon et al., Eds., Cambridge University Press, 589–662.
- Rodgers, C. D., 2000: *Inverse Methods for Atmospheric Sounding: Theory and Practice*. World Scientific, 238 pp.
- Slingo, A., S. Nicholls, and J. Schmetz, 1982: Aircraft observations of marine stratocumulus during JASIN. *Quart. J. Roy. Meteor. Soc.*, **108**, 833–856.
- Winker, D. M., R. H. Couch, and M. P. McCormick, 1996: An overview of LITE: NASA's Lidar In-Space Technology Experiment. *Proc. IEEE*, **84**, 164–180.
- , W. H. Hunt, and C. A. Hostetler, 2004: Status and performance of the CALIOP lidar. *Laser Radar Techniques for Atmospheric Sensing*, U. N. Singh, Ed., International Society for Optical Engineering (SPIE Proceedings, Vol. 5575), 8–15.