A variational method to retrieve the extinction profile in liquid clouds using multiple field-of-view lidar

Nicola L. Pounder[1] and Robin J. Hogan[1]

Tamás Várnai[2,3]

Alessandro Battaglia[4]

Robert F. Cahalan[3]

[1]*Department of Meteorology, University of Reading, Reading, UK*

[2]*University of Maryland Baltimore County, Joint Center for Earth System Technology*

[3]*NASA GSFC Climate and Radiation Laboratory (Code 613)*

[4]*Department of Physics and Astronomy, University of Leicester, Leicester, UK*

Prepared for the *Journal of Applied Meteorology and Climatology*

October, 2011

Popular summary

While liquid clouds play a very important role in the global radiation budget, it's been very difficult to remotely determine their internal cloud structure. Ordinary lidar instruments (similar to radars but using visible light pulses) receive strong signals from such clouds, but the information is limited to a thin layer near the cloud boundary. Multiple field-of-view (FOV) lidars offer some new hope as they are able to isolate photons that were scattered many times by cloud droplets and penetrated deep into a cloud before returning to the instrument. Their data contains new information on cloud structure, although the lack of fast simulation methods made it challenging to interpret the observations. This paper describes a fast new technique that can simulate multiple-FOV lidar signals and can even estimate the way the signals would change in response to changes in cloud properties—an ability that allows quick refinements in our initial guesses of cloud structure. Results for a hypothetical airborne three-FOV lidar suggest that this approach can help determine cloud structure for a deeper layer in clouds, and can reliably determine

the optical thickness of even fairly thick liquid clouds. The algorithm is also applied to stratocumulus observations by the 8-FOV airborne "THOR" lidar. These tests demonstrate that the new method can determine the depth to which a lidar provides useful information on vertical cloud structure. This work opens the way to exploit data from spaceborne lidar and radar more rigorously than has been possible up to now.

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, UK

TAMÁS VÁRNAI

Joint Centre 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, UK

ROBERT F. CAHALAN

Laboratory for Atmospheres, NASA Goddard Space Flight Center, Greenbelt, Maryland

* *Corresponding author address:* Nicola Pounder, Department of Meteorology, University of Reading, Reading, UK.

E-mail: n.l.pounder@reading.ac.uk

1. Introduction

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). Vertical cloud profiles can be used to quantify sub-adiabatic behavior and therefore to study the role of entrainment and boundary layer parameterization. A better understanding of cloud properties is important for both climate modeling and weather forecasting. Satellite remote sensing of clouds is necessary to obtain global cloud observations.

Space-borne lidar measurements of clouds are affected by multiple scattering of the lidar signals (Flesia and Schwendimann 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 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 a lot of information about cloud structure and optical depth, particularly if observed in multiple fields of view, but they are very challenging to interpret. Bissonnette et al. (2005) successfully retrieved cloud extinction coefficient profiles using multiple field of view lidar. Their algorithm uses a small-angle diffusion approximation but does not include wide angle multiply scattered returns. To utilize the small angle scattering while excluding wide angle multiply scattered returns Bissonnette et al. use multiple field-of-view lidar that detect returns sequentially from narrow fields of

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 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 THOR (Cloud Thickness from Off-beam Lidar Returns) 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 visible extinction coefficient, α_ν (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

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

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

$$\hat{\beta} = \alpha_\nu / S. \tag{2}$$

lack of positivity constraint needs to be taken into account when estimating the uncertainties on the retrieved profile and this is discussed in section 2e. An improvement to this approach would use a prior distribution that excluded negative values of extinction coefficient, such as a log-normal distribution.

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

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

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 in section 2b.

The cost function can be conveniently written in matrix notation as

$$\begin{aligned} J = &\frac{1}{2} \left[\mathbf{y} - H(\mathbf{x})\right]^T \mathsf{R}^{-1} \left[\mathbf{y} - H(\mathbf{x})\right] \\ &+ \frac{1}{2} (\mathbf{x} - \mathbf{x}_{(\text{p})})^T \mathsf{B}^{-1} (\mathbf{x} - \mathbf{x}_{(\text{p})}) + J_{\text{constraint}}, \end{aligned} \tag{6}$$

where $\mathbf{x}$ is the state vector, a vector of the α_i values to be retrieved, and $\mathbf{x}_{(\text{p})}$ is a vector of the $\alpha_i^{(\text{p})}$ values of the prior. $\mathbf{y}$ 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 R and B are the error covariance matrices of the observations and the prior respectively.

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

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

where $\Delta\mathbf{x} = \mathbf{x} - \mathbf{x}_k$, $\mathbf{x}_k$ is the estimated state vector at iteration k, and $H(\mathbf{x}_k)$ the corresponding forward-modeled estimate of the observations. $\mathsf{H} = \partial\mathbf{y}/\partial\mathbf{x}$ is the *Jacobian* matrix: the rate of change of each forward-modeled observation with respect to each element of the state vector. It 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) in to (6). At this minimum, $\nabla_{\Delta\mathbf{x}} J_\mathrm{L} = 0$, which may be rearranged to obtain

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \mathsf{A}^{-1}\nabla_{\Delta\mathbf{x}=0} J_\mathrm{L}, \tag{10}$$

where

$$\nabla_{\Delta\mathbf{x}=0} J_\mathrm{L} = -\mathsf{H}^T\mathsf{R}^{-1}\left[\mathbf{y} - H(\mathbf{x})\right] + \mathsf{B}^{-1}(\mathbf{x}_k - \mathbf{x}_{(\mathrm{p})}) + \lambda\mathsf{T}\mathbf{x}_k \tag{11}$$

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

$$\mathsf{A} = \nabla^2_{\Delta\mathbf{x}=0} J = \mathsf{H}^T\mathsf{R}^{-1}\mathsf{H} + \mathsf{B}^{-1} + \lambda\mathsf{T}. \tag{12}$$

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

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 $\nabla_{\Delta\mathbf{x}=0}J$ requires H to be calculated first, which is expensive. However, 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 $\nabla_{H(\mathbf{x})}J_{\text{obs}} = \mathsf{R}^{-1}[\mathbf{y} - H(\mathbf{x})]$ (also a vector), without requiring the intermediate matrix H. This is achieved by coding the adjoint of the forward model (e.g. Giering and Kaminski 1998), which is typically around three times slower to compute than the original forward model, but much faster than the additional order of N in computational cost associated with computing the full Jacobian. The other terms in (11) are much more rapid to compute.

As the L-BFGS method uses an approximation to A^{-1}, more iterations are required to reach convergence than for the Gauss-Newton method. However, the difference in the number of iterations is typically less than the factor of around $N/3$ between the costs of each iteration of the two methods, so, 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 non-linearity of the problem. In section 3g, the convergence rates are compared for retrievals using multiply scattered returns. Unfortunately, the approximate nature of A^{-1} 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 since it is not held explicitly by the

$N = 50$ on a 1-GHz Intel processor). Therefore, we have coded the adjoints of both the PVC and TDTS methods in order that the L-BFGS method may be applied. While the adjoint for the TDTS method is exact, 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), described by Hogan (2008). Since 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

$$\delta_m = \sum_{i=1}^{m} \alpha_i \Delta z \tag{13}$$

$$= \mathbf{w}\mathbf{x}, \tag{14}$$

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

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

where $\mathsf{S}_{\mathbf{x}}^{(m)}$ is a matrix containing the first $m \times m$ elements of the full covariance matrix $\mathsf{S}_{\mathbf{x}}$. 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

scribed 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 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 μrad bean divergence at the $1/e^2$ level. The lidar is at an altitude of 7980 m. The lidar receiver has up to three fields of view: a central, circular field of view with a footprint of 10 m at ground (1.25 mrad full width field of view) and two, concentric, annular fields of view whose outer limits encompass footprints of 100 m and 600 m at ground (12.53 mrad 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. We use scattering properties 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 μm (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 field-of-view receiver and using the central field alone. The error bars are the square root of the diagonal of the retrieval error covariance matrix $\mathsf{S}_\mathbf{x}$. The simulated

potential of a variational retrieval scheme for liquid clouds and the current forward model can still be improved.

b. The information content of a retrieval

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

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

As the state vector represents a profile, the rows $\mathbf{a}_i^T$ of W are averaging kernels or smoothing functions, one for each point in the extinction coefficient profile. If the inverse method were perfect, 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 vector of unit elements, can be considered a rough measure of the fraction of the retrieval, at gate i, that comes from the observations rather than the prior or additional constraints.

Figure 1c shows some of the averaging kernels for the three field-of-view 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 figure 1d. The kernel width is approximated as

$$\text{width}_\text{i} = \left(\frac{\sum_{j=0}^{M} W_{ij} (z_i - z_j)^2}{\sum_{j=0}^{M} W_{ij}} \right)^{\frac{1}{2}}. \tag{17}$$

The kernel width is approximately equal to the range gate spacing for the first kernel,

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. The retrieved total optical depth is $13.7^{+8.3}_{-1.7}$. 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 $7.28^{+5.4}_{-0.31}$.

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 $\mathrm{d}\alpha/\mathrm{d}z = 10^4$ 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. Although the simulated observations do not include instrument noise we do assign a measurement error 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 field of view 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

e. Sensitivity to input parameters

In the studies described above the droplet effective radius (r_{eff}) and scattering asymmetry parameter (g) assumed by the retrieval have been chosen to match those used to simulate the data. In this section we demonstrate the sensitivity of the retrieval algorithm to the choice of these parameters, separating the effect of varying r_{eff} on the retrieval due to changes in width of the forward lobe of the phase function from its effect due to changes in g. 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 ten input values of r_{eff} between 5 and 15 μm ($r_{\text{eff}} = 10$ μm was used to simulate the data). The spread of the retrieved optical depth for the three-FOV receiver and the 10 m footprint receiver are small compared to the statistical uncertainties on the retrieval in Fig. 3, so the uncertainty on the retrieved optical depth due to assuming r_{eff} to calculate the width of forward lobe of the phase function is negligible.

Figure 4b is the same as Fig. 4a except seven different input values of g, between 0.82 and 0.88, are used in the retrievals ($g = 0.85$ was used to simulate the data) and r_{eff} is fixed at 10 μm. The effect of varying g is small for the narrow field of view, which is dominated by single scattering. Varying g has a significant effect on the retrieval for the three-FOV receiver. At a true total optical depth of 35 (the limit of this receivers ability to retrieve optical depth) The retrieved optical depth varies by ~ 10 as g is varied between 0.82 and 0.88. The uncertainty associated with assuming g reduces as the true optical depth reduces. This is a fairly conservative estimate of the error due to assuming g although it has a significant affect.

the scale of the spacing between neighboring points in the profile, and 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 of the forward model to be calculated, which is computationally expensive. In contrast, the L-BFGS 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 figure 1. The starting point for the minimisation was $\log(\alpha_i) = -4$ for both methods and for these studies the cost function was formulated in terms of the 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 faster than 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 so each quasi-Newton iteration is at least 18 times faster than a Gauss-Newton iteration, making the quasi-Newton

the central field of view and each other using the approach of Cahalan et al. (2005). The raw observations include random background noise whose amplitude is independent of apparent height. The mean magnitude and standard deviation of the background is 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 four standard deviations of zero are removed. The observation errors are estimated from the raw photon counts assuming Poisson statistics.

The smoothness constraint, λ used in the retrieval of the THOR data was optimized using an L-curve analysis 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, as 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 field of view to infer cloud top this implies 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. 7. 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

profile is now much smoother near cloud top and the trough is no longer present. The gradient of the tail is unchanged as expected as 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 field of view 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 retrieved using all fields of view. The incompatible narrow FOV affects the shape of the retrieved profile but the wide FOV are able to constrain the total optical depth despite this. The wide fields of view tightly constrain the total optical depth and the gradient at cloud base and the narrow field of view constrains cloud top. The trough at around 1000 m is an artifact of the retrieval trying to reconcile these two constraints when the narrow field of view 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. While 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 the stratus cloud is horizontally homogeneous. The profile used here was averaged over 500 lidar pulses. Accounting for the aircraft speed the central field of view sampled a 7 m by 70 m area while the widest field of view samples over a 1 km square area so sampling errors could cause differences between the fields of view. However, in this case, saturation of the narrow field of view is the dominant effect.

of-view lidar in liquid clouds encourages the development of more of such instruments. In particular, this technology is perfectly suited to a satellite platform, which would enable measurements to be made of the extinction profile of liquid clouds globally for the first time. The dependence of spaceborne lidar returns on field of view 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 field of view 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), while the effect in CALIPSO, with its 90-m field of view, is only apparent on close inspection of the data. Idealized retrievals, such as those in figures 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 affect of increasing the number of fields of view could be seen 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 geometrical thickness is also of interest. For cloud retrievals where the total optical thickness can be retrieved by the algorithm, the height of cloud base can be estimated as the height where the retrieved extinction coefficient goes to zero. It is not possible to retrieve the position of cloud base with the current algorithm in very optically thick clouds; however, this could be achieved by including physically based constraints, such as the expected cloud adiabaticity, in the cost function and we intend to include such a constraint in the future.

The technique is most powerful when applied to lidars equipped with multiple field-

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., S. Tanelli, S. Kobayashi, D. Z. R. J. Hogan, and C. Simmer, 2010: Multiple-scattering in radar systems: a review. *J. Quant. Spectrosc. and Radiat. Transf.*, **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 offbeam lidar returns. *J. Atmos. Oceanic Technol.*, **22**, 605–627.

Davis, A. B., 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 B*, **24**, 177–185; Erratum 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.

Delanoë, J. and R. J. Hogan, 2010: Combined Cloudsat-CALIPSO-MODIS retrievals of the properties of ice clouds. *J. Geophys. Res.*, **115**, D00H29, doi:10.1029/2009JD012346.

Liu, D. C. and J. Nocedal, 1989: On the limited memory method for large scale optimization. *Mathematical Programming B*, **45**, 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.

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. *Mathematics of Computation*, **88**, 6787–1983.

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, 1983: Backscatter and extinction in water clouds. *J. Geophys. Res.*, **30 (C11)**, 1191–1204.

Platt, C. M. R., 1973: Lidar and radiometric observations of cirrus clouds. *J. Atmos. Sci.*, **30**, 1191–1204.

Polonsky, I. N., 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.

List of Tables

List of Figures

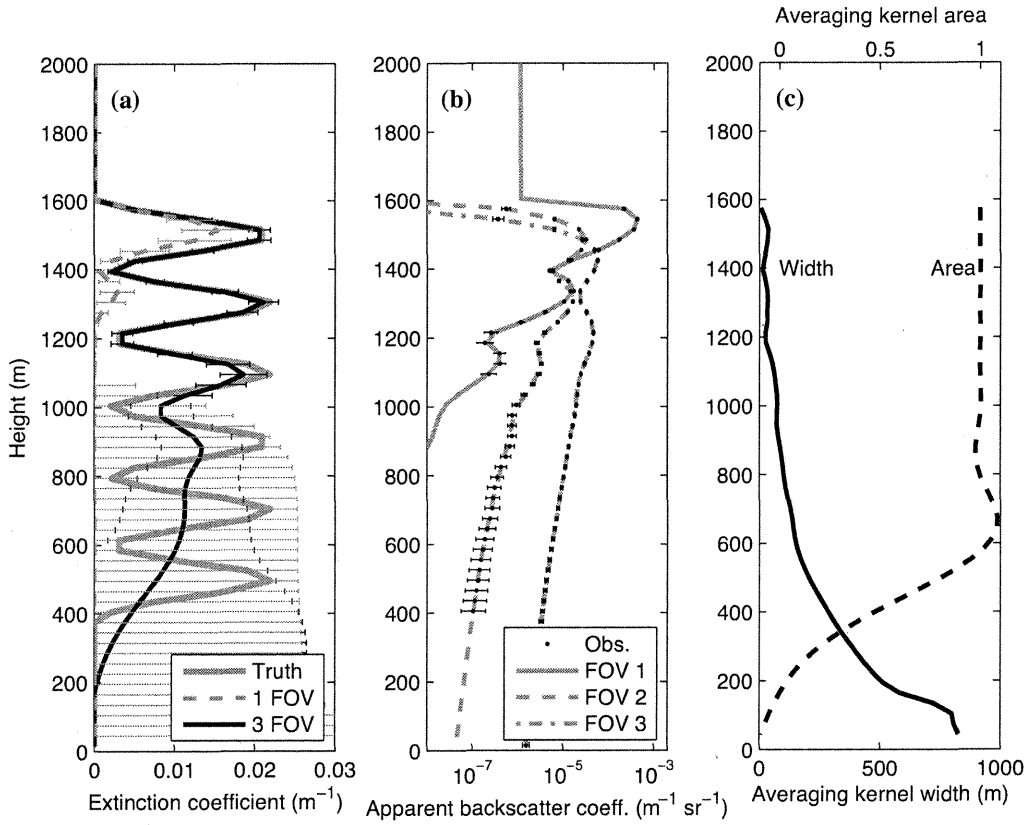

FIG. 2. Performance of the retrieval algorithm for an extinction coefficient profile with sinusoidal structure. (a) As figure 1a except without the MC line. (b) As figure 1b. (c) As figure 1d.

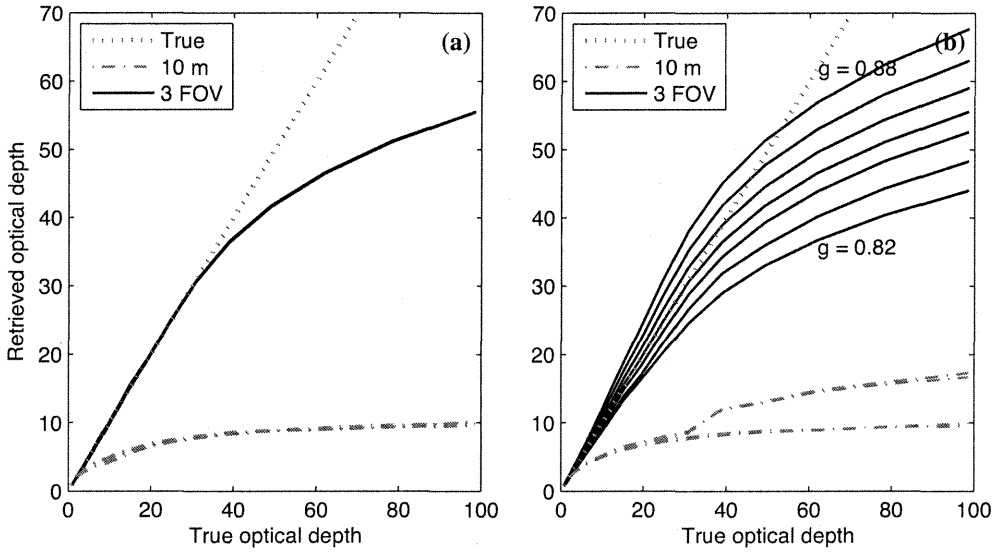

FIG. 4. The sensitivity of the retried optical depth to the assumed parameters. a)Retrieved optical depth as a function of true optical depth for ten values of r_eff between 5 and 15 μm. (b) As (a) but for seven values of g between 0.82 and 0.88.

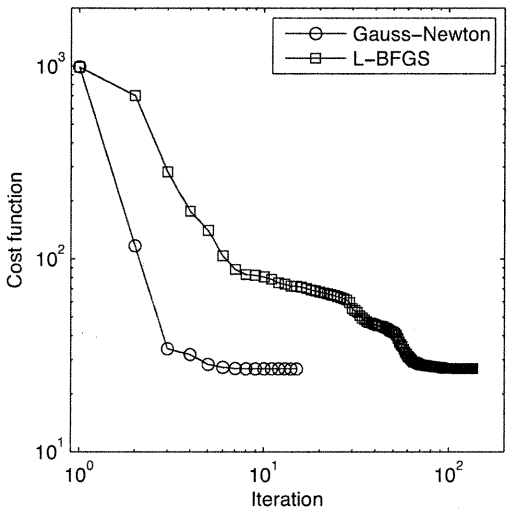

FIG. 6. Convergence rates for Gauss-Newton and L-BFGS minimization methods for the triangular profile in Fig. 1. The minimization was halted when the change in the cost function was less that 10^{-4}.

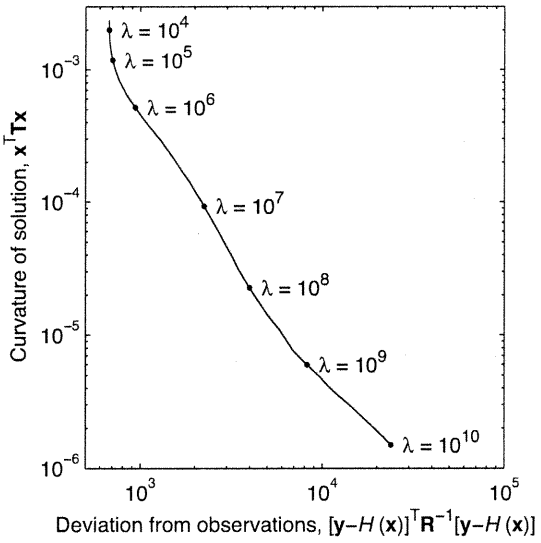

FIG. 8. L-curve for retrievals, using the Twomey-Tikhonov smoothness constraint, of a profile observed by the THOR experiment.