Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

A practical method to remove *a priori* information from lidar optimal estimation method retrievals

Ali Jalali[1], Shannon Hicks-Jalali[1], Robert J. Sica[1,2], Alexander Haefele[2,1], and Thomas von Clarmann[3]

[1]Department of Physics and Astronomy, The University of Western Ontario, London, Canada

[2]Federal Office of Meteorology and Climatology, MeteoSwiss, Payerne, Switzerland

[3]Forschungszentrum Karlsruhe, Institut fur Meteorologie und Klimaforschung, Karlsruhe, Germany

Correspondence to: sica@uwo.ca

Abstract. Lidar retrievals of atmospheric temperature and water vapour mixing ratio profiles using the Optimal Estimation Method (OEM) typically use a retrieval grid whose number of points is larger than the number of pieces of independent information obtainable from the measurements. Consequently, retrieved geophysical quantities contain some information from their *a priori*, which can affect the results in the higher altitudes of the temperature and water vapour profiles due to decreasing signal-to-noise ratios. The extent of this influence can be estimated using the retrieval's averaging kernels. The removal of formal *a priori* information from the retrieved profiles in the regions of prevailing *a priori* effects is desirable, particularly when these greatest heights are of interest for scientific studies. We demonstrate here that removal of *a priori* information from OEM retrievals is possible by transforming the retrieval from a fine grid to a coarser grid such that the averaging kernel is close to unity at each grid point. In this case, setting the *a priori* term in the OEM retrieval equation to zero minimizes the effect of the *a priori* for the coarse grid retrieval. We demonstrate the improvements gained by this technique for the case of a large power-aperture Rayleigh scatter lidar nighttime temperature retrieval and for a Raman scatter lidar water vapor mixing ratio retrieval during both day and night.

1 Introduction

Rodgers (2011) introduced an Optimal Estimation Method (OEM) based on information theory to use for atmospheric remote sensing retrievals. The OEM has primarily been used in passive remote sensing (Rodgers, 1976; Cunnold et al., 1989; Boersma et al., 2004) and it was not until recently that the OEM was applied to lidar measurements to retrieve atmospheric aerosol properties, temperature, and water vapor profiles (Povey et al., 2014; Sica and Haefele, 2015, 2016). OEM is advantageous for lidar work not only because desired geophysical quantities are retrieved (e.g. temperature, water vapour mixing ratio, etc.)

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

but also because it produces averaging kernels and a full uncertainty budget on a profile-by-profile basis. The averaging kernel matrix is a diagnostic tool that indicates the degree to which the retrieval is determined by the lidar measurements or by the retrieval *a priori* values.

Lidars have high temporal and spatial resolution compared to passive remote sensing instruments, coupled with high signal-to-noise (SNR) ratio measurements over much of their dynamic range, and thus have averaging kernels close to unity for the majority of their retrievals with a grid spacing much finer than passive instruments. At most retrieval altitudes, the majority of the information comes from the lidar measurements. However, near the top of the lidar retrieval range, and in other regions where the SNR is low, the *a priori* contribution to the retrieval increases and consequently the amount of information from the measurement decreases. The *a priori* influence at the top of the retrieval should be considered when comparing OEM lidar measurements, particularly if different *a priori* profiles are used.

An estimate of the *a priori*'s contribution to the retrieval, otherwise known as the "measurement response", can be calculated by taking the sum of the averaging kernel functions. The measurement response is calculated by multiplying the averaging kernel matrix, $\mathbf{A}$, with a unit vector, $\mathbf{u}$, which we will refer to henceforth as $\mathbf{Au}$.

An example of the *a priori*'s influence is shown in Fig. 1 of Jalali et al. (2018). Jalali et al. (2018) used more than 500 nights of measurements from the Purple Crow Lidar in London, Ontario between 1994 and 2013 to calculate the OEM temperature climatology. The cutoff height used for the climatology was the altitude at which the measurement response equaled 0.9, or where the retrieval is roughly comprised of 90% measurements and 10% *a priori* information. In order to see the influence of the *a priori* on the temperature retrieval, temperature profiles from two different models, CIRA 86 and the US Standard Model (Committee on Extension to the Standard Atmosphere, 1976), were chosen to use as *a priori* temperatures. Temperatures were retrieved using both *a priori* profiles, and the differences between the two were compared at the altitudes where $\mathbf{Au} = 0.9$ and $\mathbf{Au} = 0.99$. The distribution of the influence of the *a priori* at these altitudes for the entire climatology was also calculated (Fig. 1). However, the temperature *a priori*'s effect is always one or two degrees smaller than the random uncertainties at these altitudes.

The mean value of the histogram at the altitude where $\mathbf{Au} = 0.99$ is $0.53 \pm 1.29\,\mathrm{K}$ and the mean at $\mathbf{Au} = 0.9$ increases to $0.96 \pm 3.25\,\mathrm{K}$. There is a positive bias in both histograms which is due to the fact that the monthly CIRA-86 temperature profiles are consistently warmer than the yearly US Standard Model profile. The effect of the *a priori* increases as the values of $\mathbf{Au}$ decrease. Also, all values in the histogram are within two sigma of the statistical uncertainty of the PCL climatology.

As Rodgers (2011) suggested, it is important to pick the most accurate *a priori* for the retrieval. We used the CIRA-86 and US Standard Model to investigate the influence of the choice in *a priori* more clearly, as the differences between these two model temperatures profiles is large. If *a priori* profile values from the CIRA-72 and CIRA-86 models had been chosen for comparison, the mean values on the histogram would have been much smaller.

Several solutions regarding reducing the *a priori*'s influence on the retrieval have been suggested by Vincent et al. (2015), Ceccherini et al. (2009), von Clarmann and Grabowski (2007) and Joiner and Silva (1998). Their method to minimize the effect of the *a priori* was based on transforming a regularized to a maximum likelihood retrieval by moving from a fine grid to a coarser grid. Our work applies the methodology of von Clarmann and Grabowski (2007) to a Rayleigh lidar OEM temperature

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

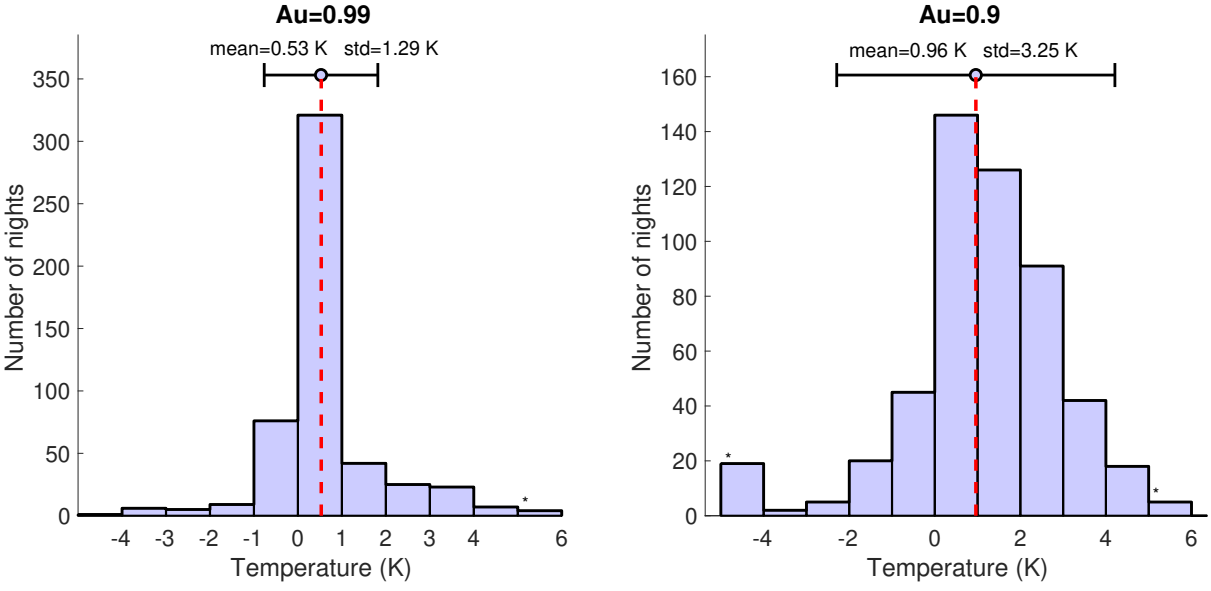

Figure 1. Distribution of the differences in temperatures retrieved at the altitudes where the sum of the averaging kernels ($\mathbf{Au}$) is 0.99(a) and 0.9(b) using two *a priori* temperature profiles - the US Standard Model and CIRA-86 for over 500 nights as detailed in Jalali et al. (2018). The red dashed line shows the mean. For each case, the difference in temperatures is always smaller than the statistical uncertainty at the same altitude.

retrieval and a Raman lidar OEM water vapour retrieval. The method uses a grid transformation on the retrieved temperature and water vapor lidar profiles to remove the *a priori* temperature and water vapour contribution. The transformation is applied in such a way that each final grid point carries roughly one degree of freedom. Then, the retrieved profiles are calculated on the coarse grid by re-running the OEM in a way that the effect of the *a priori* constraint is minimized.

We have used two lidars in this study, whose specifications are discussed in more detail in Section 2. Section 3 summarizes some basic foundation material of the OEM which will be referenced throughout the paper. Section 4 discusses the *a priori* removal methodology with a simple example. The method is then applied in Section 5 for three cases: Raman water vapour daytime, Raman water vapour nighttime, and Rayleigh nightly temperature retrievals. Section 6 discusses the differences between our practical application and the method in von Clarmann and Grabowski (2007) and some of the method's advantages. Sections 7 and 8 are the Summary and Conclusions respectively.

2 Description of the lidar systems

Two lidars were used in this study, the RAman Lidar for Meteorological Observation (RALMO) in Payerne, Switzerland and the Purple Crow Lidar in London, Ontario. RALMO was used for the water vapour daytime and nighttime retrievals and the PCL was used for the Rayleigh temperature retrievals.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

2.1 RALMO

RALMO is located at the MeteoSwiss research station in Payerne, Switzerland (46.81° N, 6.94° E, 491 m a.s.l.). RALMO was built at the École Polytechnique Fédérale de Lausanne (EPFL) and was designed as an operational lidar for model validation and climatological research. RALMO uses a 355 nm wavelength laser operating at 30 Hz with a nominal power of 300 mJ. Measurements are made in one-minute intervals with an altitude resolution of 3.75 m. A typical 30 min water vapour profile will extend to 10 - 12 km at night and 4 to 5 km during the day. Detailed specifications for the RALMO can be found in Dinoev et al. (2013) and Brocard et al. (2013). The water vapour retrieval for daytime and nighttime followed the same procedure as described in Sica and Haefele (2016), with the exception that we now retrieve the overlap, which is no longer a model parameter. Only raw (uncorrected) photocount measurements are used for the water vapour retrievals. The lidar input measurements are 30 min profiles beginning at the same time as the coincident radiosonde launch from the Payerne station. The US Standard Model water vapour profile is used as the water vapour *a priori* input for both daytime and nighttime retrievals.

2.2 Purple Crow Lidar

The Purple Crow Lidar (PCL) is located at the Environmental Science Western Field Station (43.07° N, 81.33° W, 275 m a.s.l.) near The University of Western Ontario in London, Canada. The PCL uses a 532 nm wavelength Nd: YAG laser with 1000 mJ per pulse power at 30 Hz. The PCL is comprised of two Rayleigh channels, a High Level Rayleigh (HLR) channel whose high gain detector is useful from between 40 to 110 km and a Low Level Rayleigh (LLR) low gain channel which is nearly linear due to the use of a neutral density filter above 25 km. Returns from below 25 km are blocked by a mechanical chopper which controls the firing of the laser. The backscattered photons are collected by a 2.65 m diameter liquid mercury mirror. The temporal and spatial resolution of the PCL is one minute, or 1800 laser shots, and 7.5 m, respectively. The details of the PCL OEM Rayleigh temperature retrieval are discussed in Sica and Haefele (2015) and its application to the PCL data set in Jalali et al. (2018). The PCL OEM temperature profiles are created using nightly integrated HLR and LLR measurements and typically reach up to 100 km. The *a priori* temperature profiles are the CIRA86 (Fleming et al., 1988) and US Standard Model temperatures.

3 OEM theoretical background

The Optimal Estimation Method (OEM) is an inverse method based on Bayesian statistics which calculates the maximum *a posteriori* solution by minimizing a cost function involving both the fit residual and the difference between the result and the *a priori* information. The measured signal $\mathbf{y}$ can be represented as:

$$\mathbf{y} = \mathbf{F}(\mathbf{x}, \mathbf{b}) + \epsilon, \tag{1}$$

where $\mathbf{y}$ is the measurement vector which includes measurement noise (ϵ), $\mathbf{F}$ is the forward model, $\mathbf{x}$ is for the state or retrieval vector and $\mathbf{b}$ is a vector including all model parameters which are included in the forward model but not retrieved. Note that all

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

vectors and matrices will be in bold font, but vectors will be written in lower case and matrices will be capitalized in the same format as Rodgers (2011).

The OEM assumes Gaussian probability density functions (PDFs) to maximize the *a posteriori* probability of the atmospheric state, given the value of the measurements ($P(\mathbf{x}|\mathbf{y})$) and choice of *a priori*:

$$P(\mathbf{x}|\mathbf{y}) = \frac{P(\mathbf{y}|\mathbf{x})P(\mathbf{x})}{P(\mathbf{y})}. \tag{2}$$

The possible values of measurements and solutions are distributed by the PDFs $P(\mathbf{y})$ and $P(\mathbf{x})$ respectively, and $P(\mathbf{y}|\mathbf{x})$ is the probability of the measurement given the atmospheric state $\mathbf{x}$. The solution can be optimized in a number of ways depending on the goal of the observer. The method implemented by Sica and Haefele (2015) and Jalali et al. (2018) picks the most likely state as the solution by minimizing a cost function. The cost function (Eq. 3) is a weighted least squares regression with a regularization term comprised from measurements and *a priori* components.

$$Cost = (\frac{1}{2}(\mathbf{y} - \mathbf{F}(\mathbf{x},\mathbf{b}))^T {\mathbf{S}_\epsilon}^{-1}(\mathbf{y} - \mathbf{F}(\mathbf{x},\mathbf{b})) + \frac{1}{2}(\mathbf{x} - \mathbf{x_a})^T {\mathbf{S_a}}^{-1}(\mathbf{x} - \mathbf{x_a}), \tag{3}$$

where $\mathbf{x_a}$ is the *a priori* value for the retrieval vector $\mathbf{x}$ and $\mathbf{S_a}$ is the corresponding covariance matrix. The first term in the cost function is the weighted least squares minimization problem, or the fit residuals. Minimizing the cost function produces the retrieval solution ($\hat{\boldsymbol{x}}$), where the solution is then the maximum *a posteriori* solution based on the PDFs and is given by

$$\hat{\mathbf{x}} = \mathbf{x_a} + (\mathbf{K}^T {\boldsymbol{S}_\epsilon}^{-1}\mathbf{K} + {\mathbf{S_a}}^{-1})^{-1}\mathbf{K}^T {\mathbf{S}_\epsilon}^{-1}(\mathbf{y} - \mathbf{F}(\mathbf{x_a})) = \mathbf{x_a} + \mathbf{G}(\mathbf{y} - \mathbf{F}(\mathbf{x_a})), \tag{4}$$

where $\mathbf{K}$ refers to the Jacobian matrix, $\mathbf{G}$ is the gain matrix, and $\mathbf{S}_\epsilon$ is the covariance matrix of the error measurements. The gain matrix describes the sensitivity of the retrieval to the observations:

$$\mathbf{G} = \frac{\partial \hat{\mathbf{x}}}{\partial \mathbf{y}} = (\mathbf{K}^T {\mathbf{S}_\epsilon}^{-1}\mathbf{K} + {\mathbf{S_a}}^{-1})^{-1}\mathbf{K}^T {\mathbf{S}_\epsilon}^{-1}. \tag{5}$$

One of the advantages of the OEM is that in addition to obtaining a retrieval/solution vector, the method also provides diagnostic tools and a full uncertainty budget. The primary diagnostic tool is the averaging kernel matrix ($\mathbf{A}$) which represents the sensitivity of the retrieved state to the true state. At each retrieval grid point (level or altitude), the averaging kernel shows the sensitivity of the retrieval to the measurement. The full width at half-maximum of the averaging kernel at each altitude represents the vertical resolution. Equation. 4 can be rewritten based on the averaging kernel as

$$\hat{\mathbf{x}} = \mathbf{A}(\mathbf{x} - \mathbf{x_a}) + \mathbf{G}\epsilon. \tag{6}$$

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

Equation 6 shows that if the $\mathbf{A}$ is unity at each altitude, the retrieval is sensitive only to the measurements, with no contribution from the *a priori*. Wherever $\mathbf{A}$ is less than unity, the *a priori* is contributing to the retrieval and the extent of its contribution can be estimated using the measurement response. The averaging kernel also provides a means of calculating the number of degrees of freedom (dgf) in the retrieval by evaluating the trace of $\mathbf{A}$,

$$dgf = Tr(\mathbf{A}). \tag{7}$$

Ideally, the contribution of the *a priori* is zero at all levels, and dgf equals the number of levels of the retrieved water vapour or temperature profiles.

4 Methodology

Our objective in this study is to find a practical method to remove the *a priori* information from the retrieval vector. We have based our work upon the methodology of von Clarmann and Grabowski (2007), and have developed a quick and straightforward method to remove the *a priori* from the lidar retrieval. von Clarmann and Grabowski (2007) proposed removing the effect of the *a priori* by using an information-centered grid approach. Each level of the retrieval on the information-centered grid contains one degree of freedom and therefore, the number of degrees of freedom of the signal is the same as the number of the retrieval levels. In this condition, the formal *a priori* information can be removed without de-stabilizing the retrieval.

To create an information-centered grid that contains close to one degree of freedom per level requires the averaging kernel of the fine grid retrieval. For a lidar, this is either the raw measurement spacing or a grid found by integrating some number of raw measurements into larger bins. Therefore, the first step is to run the OEM retrieval following the same procedures as in Sica and Haefele (2015) or Sica and Haefele (2016). This produces a temperature or water vapor retrieval along with their respective averaging kernel matrices and uncertainty budgets on the "fine grid" or first retrieval grid. For RALMO water vapour retrievals, the fine grid altitude resolution is 100 m and 50 m resolution for the daytime and nighttime retrievals respectively, and 1024 m for the PCL Rayleigh temperature retrieval. The fine grid averaging kernel contains the information regarding the degrees of freedom of the retrieval along the diagonal elements of the matrix (see Sect. 3). The cumulative trace of the averaging kernel is the total degrees of freedom of the retrieval (Eq. 7).

To illustrate the method, we will give a simple example with the fine grid levels, diagonal components of the averaging kernel matrix, and the cumulative trace of the averaging kernel, as shown in Table 1.

We then use the triangular representation from von Clarmann and Grabowski (2007) to create the information-centered grid using the fine grid averaging kernel. First, the cumulative trace of the averaging kernel matrix is used to determine the amount of information needed for each grid point on the coarse grid using Eq. 8:

$$dgf_c = \frac{dgf}{int(dgf) - 1} \approx 1, \tag{8}$$

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

Fine Grid Levels	Diagonal elements of **A**	Cumulative Trace of **A**
1	1	1
2	1	2
3	1	3
4	1	4
5	0.9	4.9
6	0.8	5.7
7	0.7	6.4
8	0.6	7.0
9	0.5	7.5
10	0.4	7.9
11	0.2	8.1
12	0.1	8.2

Table 1. A simple example for demonstrating the averaging kernel matrix's role in finding the coarse grid which resembles the typical structure of a lidar temperature retrieval averaging kernel. The first column is the retrieval level and for lidar OEM retrievals is typically an altitude. The second column is the elements along the diagonal of the averaging kernel matrix **A**. The third column is the cumulative trace of **A**, where the last value determines the number of degrees of freedom per grid point for the coarse grid using Eq. 8.

where dgf_c refers to the degrees of freedom per level on the coarse grid, dgf is the cumulative trace of the fine grid averaging kernel matrix (Eq. 7), and $int(dgf)$ is the integer value of dgf. The degrees of freedom per grid point is determined by dividing the total degrees of freedom by one less than the integer value of the total. For example, if the total degrees of freedom of the retrieval is 8.2, then the degrees of freedom per grid point is 8.2/(8-1) = 1.1 degrees of freedom per grid point. In the triangular representation the information is spread over $dgf-1$ grid points because the first and last points remain the same as those in the fine grid. It is then necessary to interpolate the fine grid to the points where the diagonal elements are equal to the appropriate degrees of freedom to create the coarse grid. As each grid point contains equal amounts of degrees of freedom, the grid points are distributed irregularly. The final levels which are used in the coarse grid are shown in Fig. 2. In this case, we now have coarse grid points at 1, 2.2, 3.4, 4.6, 6.1, 8, and 12. As the sensitivity of the averaging kernel decreases the number of points used in the coarse grid from the fine grid increases.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

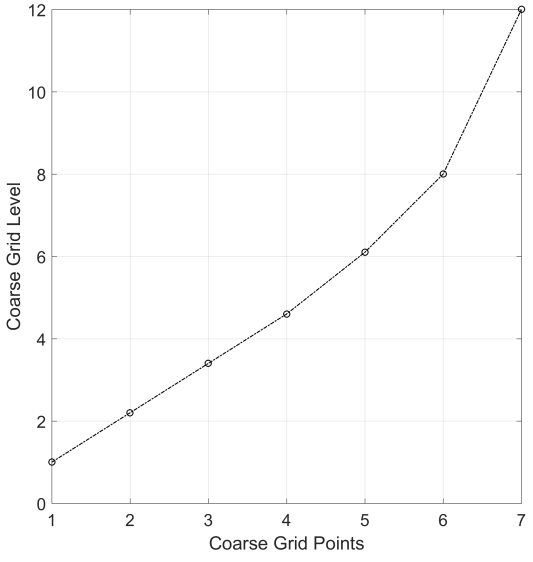

Figure 2. The coarse grid levels are shown for the example case as a function of the cumulative trace of the averaging kernel matrix. The total degrees of freedom for the retrieval is 8.2, which is spread over the entire retrieval grid such that each point has roughly one degree of freedom. As the SNR of the measurements decreases, more fine grid points are used in the coarse grid, therefore the distance between points generally increases with altitude.

The resulting coarse grid is then used as the retrieval grid for a second iteration of the OEM calculation. However, before running the retrieval again it is necessary to remove the regularization term, or equivalently, the inverse of the *a priori* covariance matrix ($\mathbf{S_a^{-1}}$) in Eq. 4. If $\mathbf{S_a^{-1}}$ is set to zero, the optimal estimation becomes the unconstrained weighted least squares solution (von Clarmann and Grabowski, 2007). We set $\mathbf{S_a^{-1}}$ to be close to zero by choosing a large uncertainty for the *a priori* covariance matrix.The OEM is then run using the new coarse retrieval grid, and the effect of the *a priori* is minimal due to minimizing the regularization term. The coarse grid averaging kernel now equals one at all levels.

5 Results

We now apply our information-centered approach, using the triangular representation from von Clarmann and Grabowski (2007), to lidar OEM retrievals in order to minimize the effect of the *a priori*. We will examine the method's effectiveness with RALMO daytime and nighttime water vapor retrievals, as well as with a PCL Rayleigh temperature retrieval. This method is also applicable in general, and can be applied to other lidar retrievals. First, we will discuss the results from the triangular representation and the creation of the "coarse grid" and how it is used as the new retrieval grid. Then we will discuss its effect on the retrieval, vertical resolution, uncertainty budgets, and averaging kernel.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

5.1 Daytime RALMO water vapor *a priori* removal

RALMO water vapour mixing ratio retrievals typically extend up to 4 or 5 km due to the large daytime background signal. This daytime water vapor retrieval on 22 January 2013 at 1200 UT (Fig. 5) shows the large impact our method has on low signal-to-noise ratios, which occur during the daytime due to the high solar background or in dry layers. The example retrieval is a 30 min integration obtained in conjunction with a Vaisala RS92 radiosonde launch from the Payerne station. The input data grid for this case was binned to 50 m to remove numerical features in the retrieval due to the high background noise levels.

The daytime case fine grid averaging kernels (Fig. 3a) quickly drop below 1 after 2 km due to a dry layer. The measurement response is shown by the red line and drops below the altitude at which $\mathbf{Au} = 0.9$ at 2.7 km, or the last altitude at which we consider the retrieval to have significant influence from the *a priori*. The coarse grid averaging kernels (Fig. 3b), by design, are all equal to 1 as discussed in Sect. 4 and reach up to 10 km. While the coarse grid insures that each altitude has 1 degree of freedom, we do not necessarily consider the entire retrieval valid, which will be discussed further below. The vertical resolution of each point on the fine and coarse retrieval grids is shown in Fig. 4. In this case, the fine grid averaging kernels are never exactly 1, therefore they have some *a priori* contribution which is why the vertical resolution is generally higher in the coarse grid retrieval. The second to last point in the coarse retrieval grid has a vertical resolution of over 600 m. The coarse grid points which have incorporated more fine grid points have a lower vertical resolution than others (between 2.8 and 10 km).

The daytime water vapor fine and coarse grid retrievals are shown in Fig. 5a and Fig. 5b respectively. The fine and coarse grid retrievals are the same up to 2.5 km, at which point the coarse grid retrieval (in red) begins to more closely follow the path of the radiosonde and the traditional profile (dotted blue) and not the fine grid retrieval (black). The coarse grid retrieval agrees with the radiosonde until 4.5 km. At 4.8 km the statistical uncertainty is above 100%, and the last two points are above 80% statistical uncertainty; therefore, the retrieval is no longer viable. All valid points are below the red dotted line. The large peaks in the fine grid retrieval above 5 km show features that are numerical and not physical. If we consider the last valid point to be 4.5 km with a statistical uncertainty of 27%, the *a priori* removal method extends the valid altitude range of the daytime OEM retrievals by 2 km.

The three main components of the uncertainty budget are shown in Fig. 5b. The fine grid statistical and air density uncertainties increase with altitude due to decreasing SNR of the return photocounts and then decrease as the retrieval falls back to the *a priori* as the signal goes to zero. The coarse grid statistical uncertainties and the uncertainty due to air density continue to increase with altitude, instead of falling back to zero, on the coarse grid because the *a priori* has been removed. The *a priori* covariance matrix is now zero in Eq. 5. The *a priori* covariance also acts as a constraint on the uncertainty by minimizing the gain, which acts as an scaling factor for the uncertainties (Rodgers, 2011). When the *a priori* covariance is removed, the solution space is no longer constrained and the coarse grid uncertainties increase compared to the fine grid uncertainties. The calibration uncertainty also increases, but now remains constant at all altitudes with the exception of the last point, as it is no longer influenced by the *a priori* constraint. The second-to-last point in the statistical uncertainty has a mixing ratio uncertainty of 100% due to the lack of signal above 4.5 km. Therefore, the meaningful section of the coarse grid retrieval is 4.5 km and below. The maximum uncertainty is 46% statistical uncertainty at 3.8 km, where the water vapour signal is very small due to

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

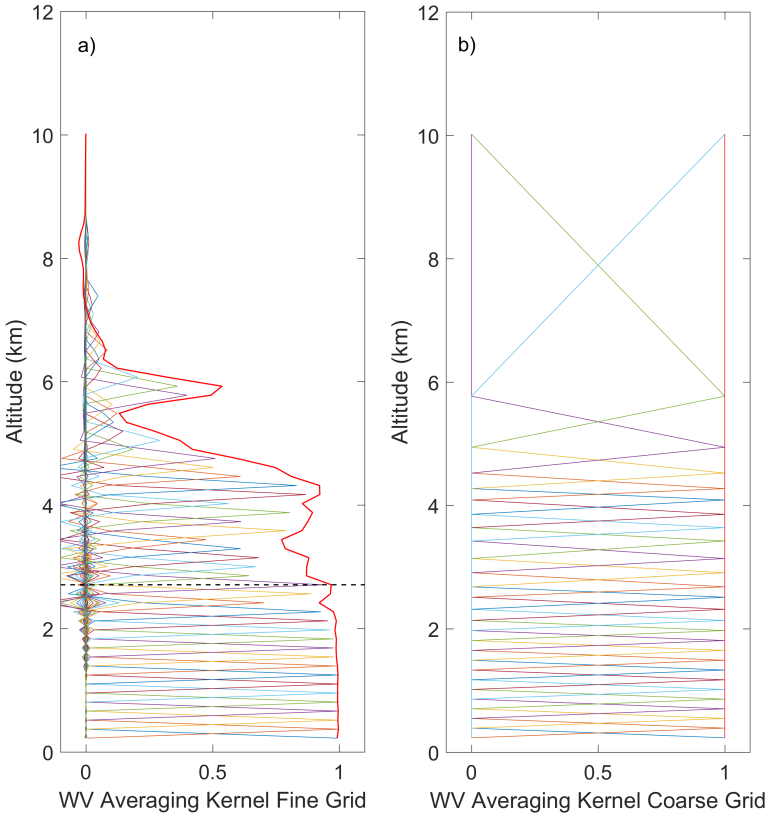

Figure 3. The clear daytime water vapour averaging kernel matrix for 22 January 2013 at 1200UT (a) on the fine grid and (b) on the coarse grid. Every other averaging kernel has been plotted for clarity. a)The measurement response $\mathbf{Au}$, or the sum of the averaging kernel rows, is the red solid line. The horizontal dashed line is the height at which the measurement response is equal to 0.9 and is the line above which we would consider there to be large influence from the *a priori*. b) The coarse grid averaging kernels all equal 1 and reach up to the last retrieval altitude at 10 km.

the presence of a dry layer at that altitude. While the *a priori* removal technique increases the maximum retrieval altitude, in addition to removing the contribution from the *a priori* profile, it will increase the statistical uncertainty of the retrieval as well.

Finally, we compare the fine and coarse grid retrievals with the radiosonde profile in Fig. 6. Below 2.3 km, where the fine grid *a priori* influence is below 10%, both of the retrievals are equivalent and show the same difference with respect to the radiosonde. However, between 2.3 and 4.8 km, where the fine grid retrieval has above 10% *a priori* information, the coarse grid retrieval more closely matches the radiosonde, and reduces the difference by 10 to 50% depending on the altitude. Above 4.8 km the coarse grid and fine grid are no longer valid due to the lack of signal at those altitudes. The coarse grid method improves the retrieval in regions where the *a priori* has significant influence and performs as we would expect.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

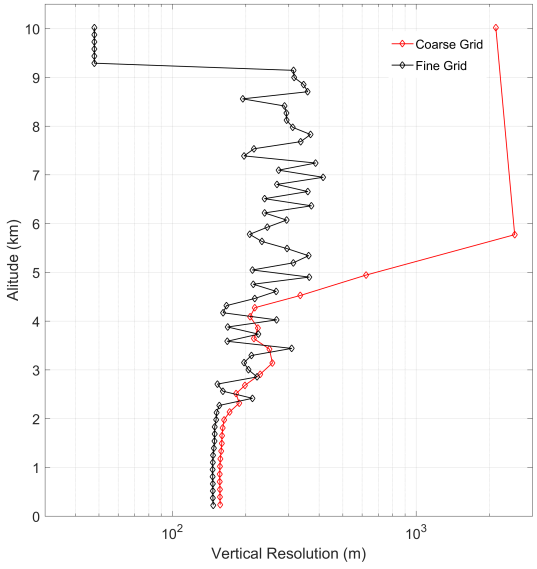

Figure 4. The vertical resolution profile on 22 January 2013 1200 UT. The vertical resolution will decrease on the coarse grid as the points are used to reach one degree of freedom. The last two points have vertical resolutions of several hundred meters and are not used in the retrieval.

5.2 Nighttime RALMO water vapour *a priori* removal

The nighttime retrieval uses a 30-minute integration on 24 April 2013 0000 UT which coincides with the time of the radiosonde launch. The fine retrieval grid for the RALMO water vapor retrieval is 50 m.

The averaging kernel matrix for the fine and coarse grid retrievals is shown in Fig. 7a and Fig. 7b, respectively. The altitude where $\mathbf{Au} = 0.9$ for the fine grid retrieval is at 9.1 km, which is typical for a 30 min nighttime measurement. The coarse grid averaging kernels all equal 1, with the second to last altitude at 11 km.

Unlike the daytime case, the nighttime vertical resolution between the fine and coarse grid retrievals is very close up to 5 km where they begin to diverge (Fig. 8). This is because the nighttime averaging kernels are very close to 1 until 5 km, therefore the fine and coarse grid are the same or very close to each other. As the *a priori* enters the signal, more points from the fine grid are used to create the coarse grid, resulting in larger coarse grid averaging kernels and decreasing the vertical resolution. Figure 9 shows the final water vapor retrievals on the fine and coarse grid, as well as a GRUAN Vaisala RS92 radiosonde profile and the traditional method results for comparison. Both fine and coarse grid profiles agree past the 0.9 cutoff and up to 9 km at which point the coarse grid retrieval diverges from both the fine grid retrieval and the radiosonde. We do see small differences in dry layers where the signal level is lower, however, the differences are inside the total uncertainty. The last four points in the retrieval are shown in dashed lines because we do not consider them to be meaningful points as their total uncertainties are 70% or larger.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

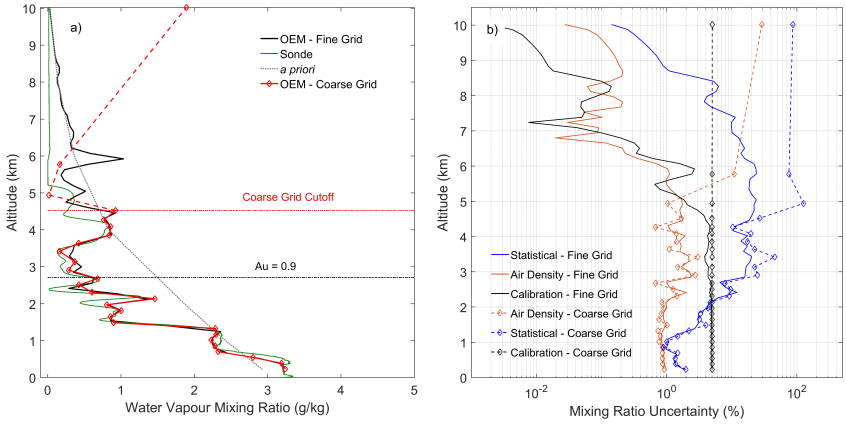

Figure 5. a) The retrieved daytime water vapour profile for 22 January 2013 1200 UT. The fine grid retrieval is in black and includes the *a priori* information. The coarse grid retrieval is in red and the *a priori*(grey) has been removed. The radiosonde is shown in green. The points which we do not consider meaningful because their uncertainties are larger than 80% in the retrieval are shown in dashed red lines. The coarse grid retrieval increases the last valid point by 2 km (red dashed line) and now more closely resembles the radiosonde above the original cutoff altitude of 2.7 km (black dashed line). b)The three primary contributors to the uncertainty budget on January 22 2013 1200 UT are shown for comparison: the statistical uncertainty, the uncertainty due to the calibration constant, and the uncertainty due to air density. The solid lines are the uncertainties from the fine grid retrieval, and the dashed lines are from the coarse grid retrieval. The *a priori* begins influencing the profile above 2 km where the uncertainty increases.

The uncertainties for the nighttime retrievals are shown in Fig. 9b. Similarly to the daytime retrievals, we have shown the top three uncertainty contributors for comparison. Below 5 km the uncertainties are the same, as there is no influence from the *a priori*. However, above 5 km the uncertainties begin to increase due to the removal. The statistical uncertainty increases to almost 100% uncertainty at the second-to-last point due to the lack of signal above 11 km.The mixing ratio uncertainty due to the calibration uncertainty is now constant with altitude, which we would intuitively expect and contributes roughly 5% uncertainty to the mixing ratio measurements. The uncertainty due to air density increases by a maximum of 0.2% at the second-to-last point. We would consider anything above 9.7 km to be invalid since points above that height have a total uncertainty of 60% or higher. The last valid point has a total uncertainty of 52% at 9.7 km. Therefore, the *a priori* removal technique increases the maximum valid altitude of the retrieval by 600 m.

The fine and coarse grid retrievals do not change very much with respect to each other until 9.1 km where the averaging kernels begin to drop off significantly. They both produce similar differences with the radiosonde (Fig. 10) until 9.7 km where the coarse grid retrieval shows a larger difference with the radiosonde. However, above this altitude we would no longer consider the coarse grid retrieval viable since its total uncertainties at that height are above 70%.

Using the *a priori* removal technique for nighttime retrievals may be helpful when trying to improve water vapour measurements of the Upper Troposphere and Lower Stratosphere (UTLS) region. However, in this case, because the nighttime measurements have large SNRs and a rapid change from high to low signal values, we do not see as large of a difference between the coarse and fine grid retrievals as we do in the daytime retrievals.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

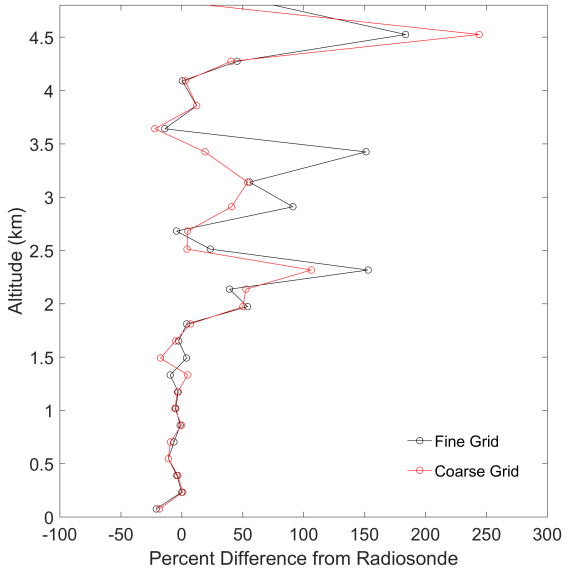

Figure 6. The percent difference between the radiosonde and the fine and coarse grid retrievals on 22 January 2013 1200 UT. The two retrievals are the same below 2 km, where the fine grid retrieval has less than 10% of *a priori* contribution. However, the coarse grid retrieval is closer to the radiosonde above 2 km and decreases the percent difference between the fine grid and the radiosonde by up to 50% in regions where the *a priori* contributes to more than 10% to the fine grid retrieval. Above 4.5 km the statistical uncertainties are too large to consider the retrieval meaningful.

5.3 Purple Crow Lidar Rayleigh temperature *a priori* removal

We picked a sample night, 12 May 2012, from the Rayleigh temperature climatology in Jalali et al. (2018) to illustrate the *a priori* removal procedure for a Rayleigh temperature retrieval. The original OEM retrieval fine grid was 1024 m, and the *a priori* temperatures were taken from the CIRA-86 model. The details regarding the OEM retrieval are discussed in Sica and Haefele (2015) and its applied result to the climatology is discussed in Jalali et al. (2018).

The averaging kernels for the fine grid and coarse grid retrievals are shown in Fig. 11a and Fig. 11b. The red line is the measurement response or the estimate of the averaging kernel's sensitivity to the measurements. The height at which the measurement response equals 0.9 was chosen as a "cutoff" height in Jalali et al. (2018), which is shown in Fig. 11a with a dashed line. After applying the *a priori* removal, the averaging kernel on the coarse grid is equal to 1 at each point. Fig. 11b shows that at the coarse grid points, the averaging kernel is completely sensitive to the measurements and therefore there is no *a priori* contribution.

The vertical resolution for both grids is similar up to 85 km altitude (Fig. 12). Above this height the coarse grid incorporates more points from the fine grid, and thus, the vertical resolution increases. The values of the vertical resolution (Fig. 12) of the two highest points for the coarse grid are 10 km at 100 km and 8 km at 110 km. However, the corresponding total uncertainties at these altitudes is above 100% and 60%, therefore we do not consider them to contribute to the retrieval.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

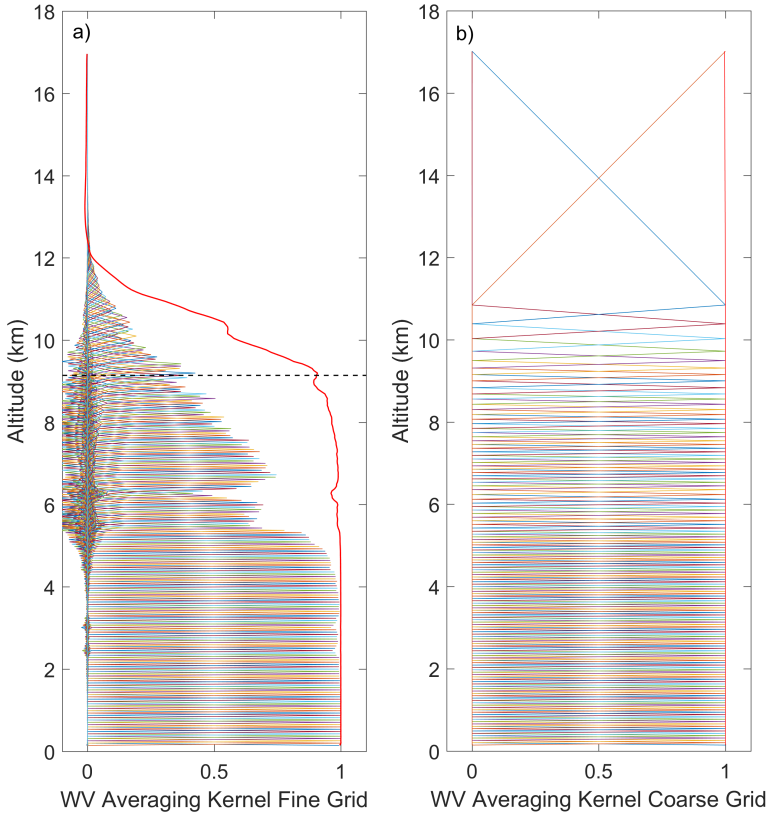

Figure 7. The averaging kernel matrix for the nighttime water vapour retrieval on 24 April 2013 0000 UT. a)The fine grid retrieval with a maximum altitude of 9.1 km (black dashed line). The measurement response is shown in red. b) The coarse grid retrieval, where each averaging kernel is 1 for all altitudes.

Figure 13a shows the temperature retrieval calculated using the OEM, with and without the *a priori*. The two retrievals are identical up to 88 km. After 88 km the coarse grid retrieval differs from the fine grid retrieval and provides 4 more levels to the retrieval. The last 2 levels are shown with dashed lines in Fig. 13a and are points that we would not consider in the retrieval due to their large uncertainties. The last valid point shown in Figure 13a is around 100 km, where the corresponding statistical uncertainty and systematic uncertainties due to the tie-on pressure and ozone cross section are 15, 9 and 2.3 K, respectively (Fig. 13b). Therefore, the last valid point of the retrieved temperature on the fine grid is within the total uncertainty of the coarse grid and the final retrieval altitude increases by 4 km.

A consequence of applying this method is that the uncertainties in the retrieval increase where the coarse grid is not equal to the fine grid. Figure 13b shows the statistical uncertainty on the fine and coarse grid, as well as two of the largest systematic uncertainties, including the uncertainty in the retrieved temperature due to the tie-on pressure and ozone cross section. The most sensitive uncertainty parameter is the statistical uncertainty, which changes from 13 K to 20 K at 98 km. The details of

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

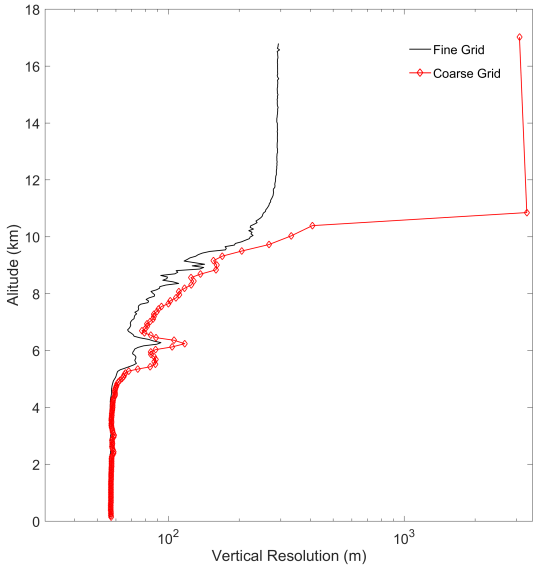

Figure 8. The vertical resolution for April 24 2013 0000 UT. The vertical resolution on the coarse grid retrieval decreases as more points are added to ensure that each bin has one degree of freedom. The coarse grid resolution is shown in red and each point is marked. The fine grid has points every 50 m therefore they are not shown individually.

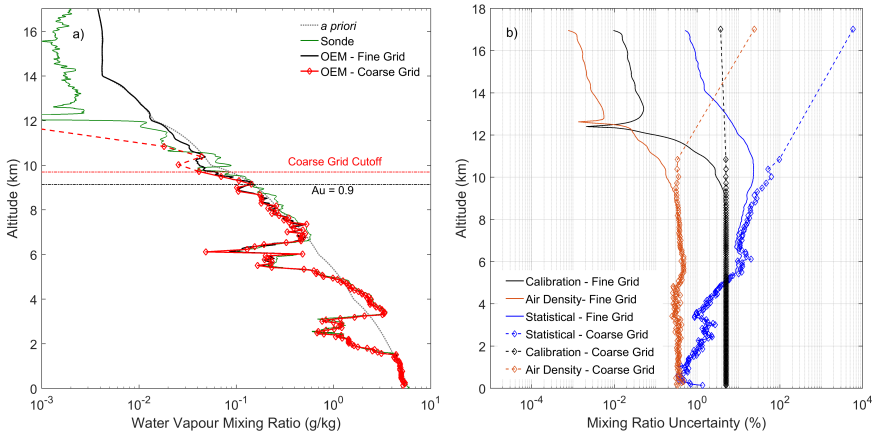

Figure 9. a)The water vapour retrieval for 24 April 2013 0000 UT. The fine grid retrieval is in black, the coarse grid retrieval is in red. In general, both OEM retrievals on the coarse and fine grid, and the radiosonde agree until the original cutoff altitude at 9.1 km (dashed black line). The dashed red lines above 9.7 km show the points we do not consider meaningful due to their large uncertainties. Therefore, the *a priori* removal technique increases the last altitude bin by 600 m. The method is limited by the lack of water vapour in the upper troposphere which causes a large and rapid drop in signal. b)The three largest uncertainty components are compared here on the fine and coarse grid. The drawback of the *a priori* removal technique is that while you gain in altitude, you increase the uncertainty. At 9.7 km the statistical uncertainty is 52%, which is where we no longer consider the rest of the retrieval to be viable.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

Atmospheric Measurement Techniques Discussions

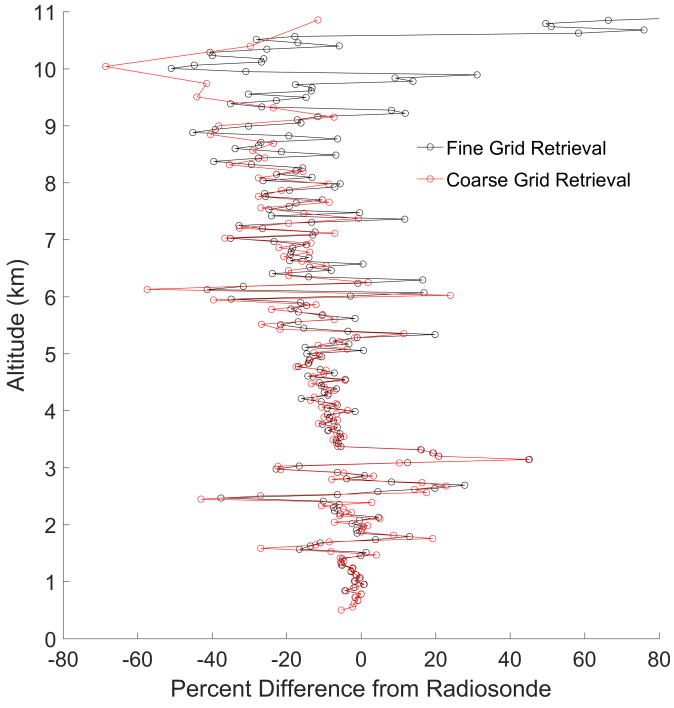

Figure 10. The percent difference from the radiosonde for both the fine and coarse grid retrievals. Both show similar differences with the radiosonde and the last valid height is 9.7 km.

the systematic uncertainties on the fine grid are discussed in Sica and Haefele (2015) and Jalali et al. (2018). The systematic uncertainties increase after *a priori* removal due to the gain matrix (Eq. 5) increasing after the regularization term is removed. In general, all uncertainties on the coarse grid (Fig. 13b) increase at higher altitudes, where contribution from the *a priori* starts. The increasing of the random uncertainties at the highest altitudes is due to decreasing photocounts from the exponential decrease in air density.

To illustrate that the *a priori* is in fact being removed, we compared the temperature retrievals using two very different *a priori* temperature profiles, one calculated by CIRA-86 and one calculated by the US Standard Model (Fig. 14). The difference between the two temperatures on the fine grid retrieval is shown by the black curve and is about 2 K at the 0.9 cutoff line, within the statistical uncertainty. The difference increases rapidly above that height. The same temperature difference after the *a priori* is removed is shown in red and is on the order of zero at all altitudes.

6 Discussion

We have developed a method to remove the influence of the *a priori* temperature and water vapour profiles on the retrieval based on the method discussed in von Clarmann and Grabowski (2007). von Clarmann and Grabowski (2007) presented a method to

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

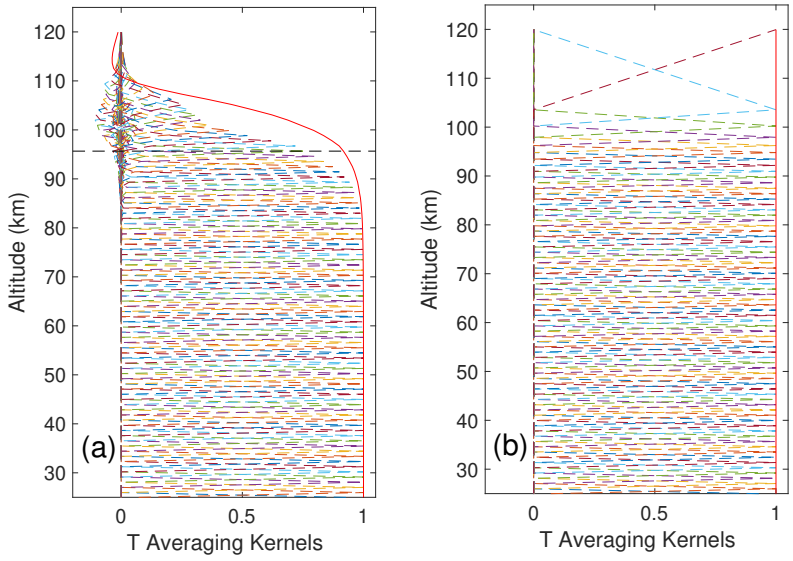

Figure 11. The PCL averaging kernels for the temperature retrieval on 12 May 2012 on the fine grid (a) and on the coarse grid (b). The $\mathbf{Au} = 0.9$ cutoff height on the fine grid is shown by the black horizontal dashed line at 97 km. The red lines on the edges of the averaging kernels are the measurement response. The coarse grid extends the temperature upwards by 4 km.

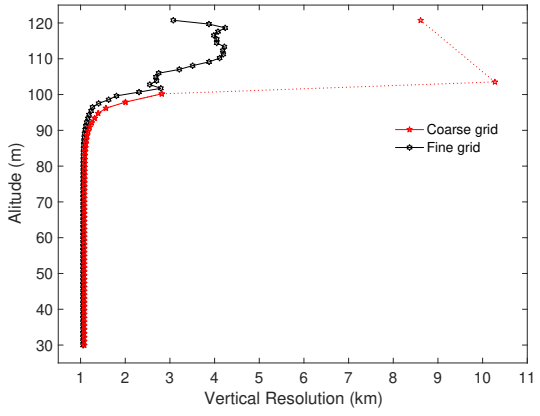

Figure 12. The PCL vertical resolution for 12 May 2012 on the fine and coarse grid. The vertical resolution is similar up to 85 km on both grids. Above this height the vertical resolution decreases until it is 10 km in resolution above 100 km altitude (dotted red line). We consider 100 km to be the highest meaningful point on the coarse grid due to large uncertainties above that height.

re-regularize the retrieval on a coarser grid after performing the OEM retrieval on a fine grid, effectively removing the *a priori* information from the retrieval. The method transformed the retrieved data from the “fine grid", or the initial retrieval grid, to a coarser grid in a way that the averaging kernel matrix on the coarse grid equals unity (that is, no *a priori* effect) at each vertical grid point.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

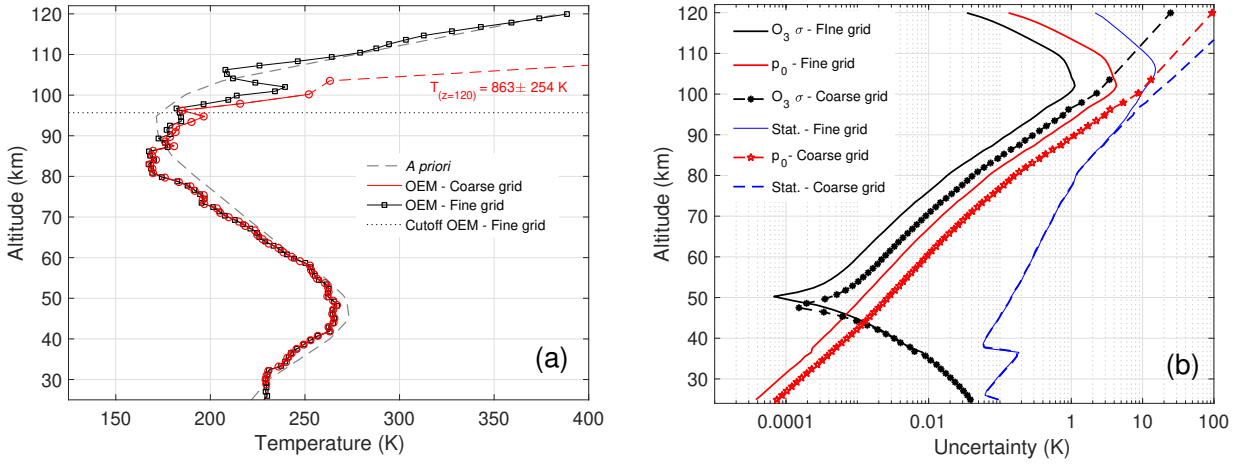

Figure 13. (a) PCL temperature retrieval for the fine and coarse grids on 12 May 2012. The temperature and its uncertainty for the last coarse grid point has a large value and it is not shown. (b) the statistical and systematic uncertainties due to the tie-on pressure and ozone cross section for the PCL temperature retrieval. The other systematic uncertainty terms included in our retrieval are not shown.

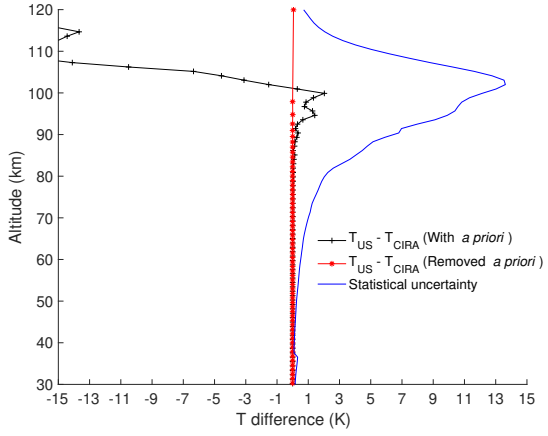

Figure 14. PCL temperature difference between the OEM retrieved temperature profiles using values from the US Standard Atmosphere and CIRA-86 as the *a priori*.

von Clarmann and Grabowski (2007) presented two approaches, a “staircase" representation, and a “triangular" representation, to transform the retrieval from the fine to the coarse grid. The cumulative trace of **A** shows the total degrees of freedom of the retrieval. In these representations, the cumulative trace of the averaging kernel matrix **A** as a function of altitude is calculated and is then interpolated to the coarse grid based on the centered information approach. As each space contains only one degree of freedom, the spaces are distributed irregularly. The staircase representation is not a realistic representation of the atmosphere (von Clarmann and Grabowski, 2007), therefore we use the triangular representation here to create the coarse grid.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

In the triangular representation, the highest and lowest level of the coarse grid are considered to be the same as the fine grid and the rest of the grid points are distributed such that each element of $\mathbf{A}$ at each level is close to one.

Our method differs from von Clarmann and Grabowski (2007) in that we do not re-regularize the retrieval to remove the *a priori*. Instead, after the initial retrieval, we remove the regularization term from the retrieval and re-run the retrieval using the triangular coarse grid. Both methods are equally effective, however, this method is more of a brute-force technique but easier to practically implement.

We have shown how the *a priori* removal method works for three sample retrievals: water vapor during both daytime and nighttime, and a nighttime Rayleigh temperature. The *a priori* removal technique is most useful when the SNR is low, particularly in dry layers. The method increased the retrieval altitude by roughly 2 km which is highly beneficial for meteorological studies that rely on accurate tropospheric measurements. The nighttime water vapor retrieval was provided for contrast to illustrate how the *a priori* removal technique does not provide significantly more information when the signal level falls off rapidly. However, the method did improve the difference between the radiosonde and the retrieval in the dry layer regions. The signal in the upper troposphere is significantly weaker and as such there is no measurement information left with which to increase the retrieval altitudes.

For Rayleigh temperature retrievals, we used measurements from the PCL in London, Ontario. Jalali et al. (2018) suggested that the 0.9 level be used as the valid cut-off height. In the case of the PCL, we see that the second to last point on the coarse grid has a vertical resolution not much larger than the fine grid retrieval (Fig. 12) and is very close to the same height; therefore, the 0.9 measurement response value seems to be a conservative choice for a valid cutoff. We also showed that the effect of the *a priori* is removed completely in the Rayleigh temperature retrieval when we compared the differences in the retrieved temperature using the values from CIRA-86 and from the US Standard Model as the *a priori* profiles (Fig. 14). The presented method provides us with higher altitudes for the retrieved temperature profiles. Additionally, where the retrieved temperature profile in the coarse grid is the same as it is for the fine grid, we can be confident the temperature retrieval has a negligible contribution from the chosen *a priori* temperature profile.

An advantage of our method is that the entire coarse grid profile is valid, in the sense that the regularization term does not contribute to the retrieval. In regions where the SNR is low or the averaging kernel is significantly less than 1, the *a priori* removal method improves the validity of the retrieval. However, the systematic and statistical uncertainties in the retrieval increase due to the removal of the inverse of the *a priori* covariance matrix from the gain equation (Eq. 5). The vertical resolution of the profile also increases as a consequence of the method. While the *a priori* removal gives us more confidence in the retrieval, we may not consider the entire profile meaningful due to high uncertainties. Hence, the last few points with unity averaging kernel value on the coarse grid may not be recognized as valid retrieval levels.

7 Summary

We have developed a practical and robust method which removes the effect of *a priori* information in lidar OEM retrievals. The method utilizes an information-centered coarse grid which is derived using the averaging kernels from the initial "fine

grid" retrieval. The resulting coarse grid is then used, in addition with setting the inverse of the *a priori* covariance matrix to zero, to create the final retrieval without any *a priori* information. The method has little computational cost; the OEM retrieval is extremely fast even on a laptop computer, so having to do the retrieval twice for each profile is not critical. We illustrated the method using a simple example in Sect. 4 and demonstrated the removal method using the water vapour signal from the RALMO and the Rayleigh temperature signal from the PCL. We summarize the results from both of these examples as follows:

1. Figure 1b) shows that 90% of the nights in the temperature climatology from Jalali et al. (2018) had less than a 5 K influence from the *a priori* temperature profiles at the $\mathbf{Au} = 0.9$ cutoff height. Additionally, in all cases the *a priori* temperature influence was less than the statistical uncertainty, as was illustrated in Fig. 6 in Jalali et al. (2018). Although small, the *a priori* temperature profile does contribute to the retrieved temperature in regions where the measurement response is smaller than 1.

2. The daytime water vapor *a priori* removal showed better agreement between the OEM coarse grid retrieval and radiosonde by up to 50% in regions where the fine grid measurement response was smaller than 0.9. Using this method helped to increase the altitude range of the daytime and nighttime water vapor retrievals by up to 2 km and 600 m, respectively. The nighttime water vapor *a priori* removal does not show large differences with the fine grid retrieval, but does show differences in dry regions with low signal. The difference in improvement between the daytime and nighttime retrievals is due to the difference in SNR of the RALMO measurements.

3. Applying the method to the PCL temperature retrieval showed a gain in information above the $\mathbf{Au} = 0.9$ cutoff height by 2 km, validating the choice of $\mathbf{Au} = 0.9$ for a cutoff made in Jalali et al. (2018) to form their climatology up to an altitude where tie-on pressure effects were minimal. The temperatures below the cutoff height were the same.

4. In all cases, the vertical resolution of the OEM retrieval decreases after *a priori* removal.

5. The systematic uncertainties after *a priori* removal increase roughly by a factor of 2, but remain on the same order of magnitude as before the *a priori* removal. The values of the systematic uncertainties also remain significantly smaller than the statistical uncertainties.

6. The temperature difference between the PCL retrieved temperature profiles using two different *a priori* profiles were used to show the efficiency of the *a priori* removal method. The temperature difference before removal around the 0.9 cutoff height was more than 2 K, however, this value was zero for the entire range after *a priori* removal.

Atmos. Meas. Tech. Discuss., https://doi.org/10.5194/amt-2018-347
Manuscript under review for journal Atmos. Meas. Tech.
Discussion started: 23 October 2018

8 Conclusions

A question that often arises from our colleagues when introduced to the OEM is: what is the effect of the *a priori* on the retrieval? This effect has been explored in detail for satellite-based and passive ground-based instruments, but not for the new area of applying OEM to active-sensing measurements such as lidar. Lidars are high resolution instruments with significant amounts of information available from their measurements, as evidenced by the retrieval averaging kernels. The OEM helps to illustrate the robustness of the lidar data products with the advantage of providing diagnostic tools, such as the averaging kernel and a full uncertainty budget.

Newcomers to the OEM may find this *a priori* removal technique helpful for checking their *a priori*'s influence on the retrieval and in determining the appropriate *a priori*. It is most useful for lidar measurements with low signal to noise and a slow transition from regions of high signal to low signal. The method is less effective when signal strength changes rapidly, such as when the nighttime water vapour measurements quickly enter the dry upper troposphere or lower stratosphere.

Another advantage of this method is that the same coarse grid for a typical night can be used for multiple lidar retrievals. In some cases, the coarse grid will not be optimal but still reasonable. With this consistent grid choice, the altitude resolution of a multi-year time series will be consistent. Varying information content of the individual measurements will lead to error bars of different size. The coarse grid allows time series analysis or trend analysis for single altitudes without problems caused by varying altitude resolution.

In the future, this method will be applied to the entire 10 years of RALMO measurements to retrieve the water vapour day time and nighttime measurements and create a water vapour climatology. We anticipate that this technique will increase the altitude of the daytime water vapour retrievals by several kilometers. It is also our hope that this method may provide statistically significant measurements in the UTLS region. Finally, the RALMO water vapour climatology will be used to find trends.

Author contributions. Ali Jalali was responsible for developing the *a priori* removal method and code as well as manuscript preparation. He also applied the method to the Rayleigh Temperature analysis.This work is the second component of his doctoral thesis. Shannon Hicks-Jalali applied the removal method to the water vapour daytime and nighttime analyses as well as helped with manuscript preparation. This will serve as a component of her doctoral thesis. Robert J. Sica was responsible for supervising the doctoral theses and contributed to manuscript preparation. Alexander Haefele contributed to manuscript preparation and scientific discussions. Thomas Von Clarmann significantly contributed to the scientific discussions which resulted in this paper, helped develop the method based on his original work, and contributed to manuscript preparation.

Acknowledgements. This project has been funded in part by Discovery Grants and a CREATE Training Program in Arctic Atmospheric Science (PI K. Strong) from the National Science and Engineering Research Council of Canada and awards from the Canadian Foundation

for Climate and Atmospheric Science. We would like to thank the Federal Office of Meteorology and Climatology, MeteoSwiss, for its support of this project and providing the water vapour lidar measurements.

References

Boersma, K. F., Eskes, H. J., and Brinksma, E. J.: Error analysis for tropospheric NO2 retrieval from space, J. Geophys. Res., 109, 2004.

Brocard, E., Philipona, R., Haefele, A., Romanens, G., Mueller, A., Ruffieux, D., Simeonov, V., and Calpini, B.: Raman Lidar for Meteorological Observations , RALMO – Part 2 : Validation of water vapor measurements, Atmos. Meas. Tech., 6, 1347–1358, 2013.

Ceccherini, S., Raspollini, P., and Carli, B.: Optimal use of the information provided by indirect measurements of atmospheric vertical profiles, Opt. Express, 17, 4944–4958, 2009.

Committee on Extension to the Standard Atmosphere: U.S. standard atmosphere, US Government Printing Office, pp. 1–227, doi:NASA-TM-X-74335, NOAA-S/T-76-1562, 1976.

Cunnold, D. M., Chu, W., R.A.Barnes, McCormick, M. P., and Veiga, R. E.: Validation of SAGE II ozone measurements, J. Geophys. Res., 94, 8447–8460, 1989.

Dinoev, T., Simeonov, V., Arshinov, Y., Bobrovnikov, S., Ristori, P., Calpini, B., Parlange, M., and Van Den Bergh, H.: Raman Lidar for Meteorological Observations , RALMO – Part 1 : Instrument description, Atmos. Meas. Tech., 6, 1329–1346, 2013.

Fleming, E. L., Chandra, S., Shoeberl, M. R., and Barnett, J. J.: Monthly Mean Global Climatology of Temperature, Wind, Geopotential Height and Pressure for 0-120 km, NASA Tech. Memo., NASA TM100697, 85 pp, 1988.

Jalali, A., Sica, R. J., and Haefele, A.: A middle latitude Rayleigh-scatter lidar temperature climatology determined using an optimal estimation method, Atmos. Meas. Tech. Discussions, 2018, 2018.

Joiner, J. and Silva, A. D.: Efficient methods to assimilate remotely sensed data based on information content, Q. J. Roy. Meteor. Soc., 124, 1669–1694, 1998.

Povey, A. C., Grainger, R. G., Peters, D. M., and Agnew, J. L.: Retrieval of aerosol backscatter, extinction, and lidar ratio from Raman lidar with optimal estimation, Atmos. Meas. Tech., 7, 757–776, 2014.

Rodgers, C. D.: Retrieval of atmospheric temperature and composition from remote measurements of thermal radiation, Rev. Geophys., 14, 609–624, 1976.

Rodgers, C. D.: Inverse Methods for Atmospheric Sounding: Theory and Practice, vol. 2, World Scientific, Hackensack, NJ, USA, 2011.

Sica, R. J. and Haefele, A.: Retrieval of temperature from a multiple-channel Rayleigh-scatter lidar using an optimal estimation method, Appl. Optics, 54, 1872–1889, 2015.

Sica, R. J. and Haefele, A.: Retrieval of water vapor mixing ratio from a multiple channel Raman-scatter lidar using an optimal estimation method, Appl. Optics, 55, 763–777, 2016.

Vincent, R. A., Dudhia, A., and Ventress, L. J.: Vertical level selection for temperature and trace gas profile retrievals using IASI, Atmos. Meas. Tech., 8, 2359–2369, 2015.

von Clarmann, T. and Grabowski, U.: Elimination of hidden a priori information from remotely sensed profile data, Atmos. Chem. Phys., 7, 397–408, 2007.