



Time series inversion of spectra from ground-based radiometers

O. M. Christensen and P. Eriksson

Department of Earth and Space Sciences, Chalmers University of Technology, Gothenburg, Sweden

Correspondence to: O. M. Christensen (ole.m.christensen@chalmers.se)

Received: 3 December 2012 – Published in Atmos. Meas. Tech. Discuss.: 12 February 2013

Revised: 7 May 2013 – Accepted: 29 May 2013 – Published: 4 July 2013

Abstract. Retrieving time series of atmospheric constituents from ground-based spectrometers often requires different temporal averaging depending on the altitude region in focus. This can lead to several datasets existing for one instrument, which complicates validation and comparisons between instruments. This paper puts forth a possible solution by incorporating the temporal domain into the maximum a posteriori (MAP) retrieval algorithm. The state vector is increased to include measurements spanning a time period, and the temporal correlations between the true atmospheric states are explicitly specified in the a priori uncertainty matrix. This allows the MAP method to effectively select the best temporal smoothing for each altitude, removing the need for several datasets to cover different altitudes.

The method is compared to traditional averaging of spectra using a simulated retrieval of water vapour in the mesosphere. The simulations show that the method offers a significant advantage compared to the traditional method, extending the sensitivity an additional 10 km upwards without reducing the temporal resolution at lower altitudes. The method is also tested on the Onsala Space Observatory (OSO) water vapour microwave radiometer confirming the advantages found in the simulation. Additionally, it is shown how the method can interpolate data in time and provide diagnostic values to evaluate the interpolated data.

1 Introduction

The maximum a priori method (Rodgers, 2000), or optimal estimation, is a commonly used method for retrieving atmospheric properties from spectroscopic measurements. Such measurements will always contain thermal noise, which negatively impacts the retrieval process. In order to overcome this noise, the measured atmospheric spectra are averaged

over time. This averaging increases the signal to noise ratio, but reduces the temporal resolution of the measurements. The strength of the measured spectral lines, the required accuracy, and the sensitivity of the instrument determine the need for temporal averaging. Hence, depending on which atmospheric phenomena and altitudes that are investigated in any single study, a specific compromise must be made between temporal resolution and noise reduction.

The use of different temporal resolutions is exemplified by, for example, looking at microwave spectrometers measuring water vapour in the middle atmosphere. One example is the the *Wasserdampf- und Spurengasmessungen in der Atmosphäre mit Mikrowellen* (WASPAM) instrument, which has used a 6 h averaging time in a case study of sudden stratospheric warming (Seele and Hartogh, 2000) as well as a 24 h averaging time to study the annual variation of water vapour around the mesopause (Seele and Hartogh, 1999). Different averaging times are used since retrievals for higher altitudes require lower thermal noise and thus longer integration times.

The ratio between the required averaging times for high and low altitudes will in large part be determined by the altitude range of the instrument. As newer instruments, such as those described in Nedoluha et al. (2011) and Bleisch et al. (2011), offer the possibility to increase this range, the differences in averaging time needed for the upper- and lowermost altitudes will increase.

Although the use of different averaging times in itself poses no problems, it does complicate the validation and cross-comparison of instruments. An example is the water vapour radiometer at the Onsala Space Observatory (OSO) (Forkman et al., 2003). This instrument has mainly used one-day spectra in studies of atmospheric dynamics (Forkman et al., 2005; Scheiben et al., 2012), whereas only a scheme with varying averaging time depending on tropospheric

opacity has been cross-compared with other instruments (Haefele et al., 2009).

One way to circumvent the problem of multiple averaging times for a single dataset is to incorporate the averaging of spectra directly in the retrieval process. To achieve this, several measurements are simultaneously inverted using temporal correlation data of the quantity to be retrieved. Earlier attempts to use the temporal information in retrievals have resorted to recursive filtering (Askne and Westwater, 1986), whereas the method presented in this paper uses a non-recursive approach based on the maximum a posteriori method (Rodgers, 2000).

The study is described using the following structure. Section 2 introduces the retrieval theory, terminology and the time series inversion technique. In Sect. 3 we apply the time series inversion method on a simulated instrument to show the advantages of the method. Section 4 investigates the practical use of the method, and Sect. 5 discusses the computational requirements. The conclusion is given in Sect. 6.

2 Retrieval methodology

2.1 Terminology

In passive atmospheric remote sensing, properties of the atmosphere are determined by analysing the radiation emitted from, and passing through, the atmosphere. This analysis is called a retrieval, or inversion, and is done by solving an inverse problem. The relationship between the measured radiation, \mathbf{y} , and the atmospheric properties is described by a forward model, $\mathbf{y} = \mathbf{F}(\mathbf{x})$, where \mathbf{x} , denoted as the state vector, contains the variables to be retrieved. These can include atmospheric variables at different altitudes as well as instrument variables. If the forward model is locally linear, the measurement can be expressed as

$$\mathbf{y} = \mathbf{F}(\mathbf{x}_0) + \mathbf{K}(\mathbf{x} - \mathbf{x}_0) + \epsilon, \quad (1)$$

where \mathbf{x}_0 is the state vector used for linearising the forward model, \mathbf{K} is the Jacobian, or weighting function, matrix, defined as $\partial \mathbf{y} / \partial \mathbf{x}$, and ϵ represents errors in the measurement.

The inverse problem involves finding \mathbf{x} for a given \mathbf{y} . In remote sensing, inverse problems can be ill-posed. This means that several atmospheric states can give rise to the same measurement. To obtain sensible results, the inversion must then be constrained through some regularisation algorithm. The regularisation in the time series inversion method is based on the maximum a posteriori (MAP) method, also called optimal estimation. It uses statistical properties of the measurements and the atmosphere to constrain the solutions (Rodgers, 2000).

Assuming Gaussian statistics, the relationship between a set of stochastic variables is described by a covariance matrix, \mathbf{S} , in which each element, $\mathbf{S}_{i,j}$, is the covariance between variables i and j . If the covariance of the a priori state is

given by the matrix \mathbf{S}_a and the measurement uncertainties are specified by \mathbf{S}_ϵ , a cost function can be expressed as

$$\chi^2 = (\mathbf{x} - \mathbf{x}_a)^T \mathbf{S}_a^{-1} (\mathbf{x} - \mathbf{x}_a) + (\mathbf{y} - \mathbf{F}(\mathbf{x}))^T \mathbf{S}_\epsilon^{-1} (\mathbf{y} - \mathbf{F}(\mathbf{x})), \quad (2)$$

where \mathbf{x} is the true state vector and \mathbf{x}_a is the a priori state vector, which indicates a “best guess” atmosphere. The state minimising this function is the maximum of the a posteriori probability density function and is given by

$$\hat{\mathbf{x}} = \mathbf{x}_a + \mathbf{G}(\mathbf{y} - \mathbf{F}(\mathbf{x}_a)), \quad (3)$$

where \mathbf{G} , the gain-matrix, is

$$\mathbf{G} = \partial \mathbf{y} / \partial \mathbf{x} = \left(\mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{K} + \mathbf{S}_a^{-1} \right)^{-1} \mathbf{K}^T \mathbf{S}_\epsilon^{-1}. \quad (4)$$

An alternative way of expressing Eq. (3) is by introducing the averaging kernel (AVK) matrix, $\mathbf{A} = \partial \hat{\mathbf{x}} / \partial \mathbf{x} = \mathbf{G} \mathbf{K}$. Combining Eqs. (1) and (3) gives

$$\hat{\mathbf{x}} = \mathbf{x}_a + \mathbf{A}(\mathbf{x} - \mathbf{x}_a) + \mathbf{G}\epsilon. \quad (5)$$

This shows that a retrieved value, in the absence of noise, is the sum of the corresponding a priori value plus the change from the a priori state convolved with the matching row of the AVK matrix. The sum of a row in the AVK matrix is a measure of the retrieval’s sensitivity to changes in the state vector and is called measurement response (Baron et al., 2002), or measurement sensitivity.

2.2 Time series inversion

For ground-based instruments using single spectrum inversions, the measurement vector, \mathbf{y} , usually holds the brightness temperature for each channel of the instrument. Thus, for an instrument with m channels, the length of \mathbf{y} is m , and similarly \mathbf{x} contains values of the state variables at the time of the measurement. The time series inversion method proposed here expands both \mathbf{y} and \mathbf{x} to include measurements from N different times. This increases the length of \mathbf{y} to $m \cdot N$, and the length of \mathbf{x} to $n \cdot N$, where n is the length of the single spectrum state vector. In a similar fashion, the Jacobian matrix, \mathbf{K} , becomes a block diagonal matrix with block elements equal to the Jacobian matrix for each measurement, giving

$$\mathbf{y} = \begin{pmatrix} y_1^1 \\ \vdots \\ y_m^1 \\ y_1^2 \\ \vdots \\ y_m^2 \\ y_1^3 \\ \vdots \\ y_m^N \end{pmatrix}, \quad \mathbf{K} = \begin{pmatrix} \mathbf{K}^1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \mathbf{K}^N \end{pmatrix}, \quad \text{and} \quad \mathbf{x} = \begin{pmatrix} x_1^1 \\ \vdots \\ x_n^1 \\ x_1^2 \\ \vdots \\ x_n^2 \\ x_1^3 \\ \vdots \\ x_n^N \end{pmatrix},$$

where the upper index denotes measurement number.

Since the state and measurement vectors cover several measurements, the covariance matrices in Eq. (2) must be adjusted. Assuming that the measurement uncertainties are uncorrelated between the measurements (i.e. only thermal noise) and uncorrelated between the channels, \mathbf{S}_ϵ simply becomes a diagonal matrix. The uncertainties in the a priori state are, however, *not* uncorrelated between the times of the measurements. To incorporate this temporal correlation, even the non-diagonal blocks in \mathbf{S}_a must have non-zero values. Thus, the covariance matrices become

$$\mathbf{S}_\epsilon = \begin{pmatrix} \mathbf{S}_\epsilon^1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \mathbf{S}_\epsilon^N \end{pmatrix}, \text{ and } \mathbf{S}_a = \begin{pmatrix} \mathbf{S}_a^{1,1} & \mathbf{S}_a^{1,2} & \dots & \mathbf{S}_a^{1,N} \\ \mathbf{S}_a^{2,1} & \mathbf{S}_a^{2,2} & \dots & \mathbf{S}_a^{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{S}_a^{N,1} & \mathbf{S}_a^{N,2} & \dots & \mathbf{S}_a^{N,N} \end{pmatrix},$$

where $\mathbf{S}_a^{i,j}$ is the a priori covariance matrix corresponding to measurement i and j , and \mathbf{S}_ϵ^i is the noise covariance matrix from measurement i .

Introducing correlation between measurements in \mathbf{S}_a results in an averaging kernel matrix which contains non-zero elements in its off-diagonal blocks. The resulting matrix,

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}^{1,1} & \mathbf{A}^{1,2} & \dots & \mathbf{A}^{1,N} \\ \mathbf{A}^{2,1} & \mathbf{A}^{2,2} & \dots & \mathbf{A}^{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}^{1,N} & \mathbf{A}^{2,N} & \dots & \mathbf{A}^{N,N} \end{pmatrix},$$

will contain information about how much smoothing occurs with respect to both altitude and time. Temporal smoothing occurs because the MAP method uses information from several measurements to retrieve a single profile. The amount of smoothing will partly depend on the setup of the a priori uncertainty matrix, which will be described in detail in the next section.

2.3 Specification of the a priori covariance matrix

The specification of the uncertainty matrices is central to MAP. Since the state vector of the time series inversions has been extended to include several measurements, both the temporal and vertical correlation of the state variables need to be specified in \mathbf{S}_a , and it is through the latter correlation that the MAP method can take into account results from adjacent measurement times when retrieving a profile.

The correlation between state variables can be conveniently described with a correlation function combined with a correlation length. We will use an exponential function to describe correlations. This means that for state vector variables related in altitude (e.g. species concentration), the correlation coefficient between the variable at altitude z_k and z_p for measurement i will be given by $\rho_z(\mathbf{x}_k^i, \mathbf{x}_p^i) = \exp(-|z_k - z_p|/l_c)$, where l_c is denoted as the correlation length, meaning the length at which the correlation has dropped to $1/e$.

In a similar fashion the correlation of a variable in time can be represented by $\rho_t(\mathbf{x}_k^i, \mathbf{x}_k^j) = \exp(-|t^i - t^j|/t_c)$, where t^i (t^j) is the time of measurement i (j) and t_c is the temporal correlation length. Assuming that the correlation is independent in the two dimensions (separable), the total correlation is calculated as the product of ρ_z and ρ_t . In this study we also assume that the a priori covariance is stationary, so that it can be described by the same matrix at all measurement times.

The a priori covariance matrix of the investigated variable should represent both the uncertainty arising from natural variability and the uncertainty of the a priori mean value (Eriksson, 2000). The latter uncertainty arises from a limited knowledge of the atmospheric mean state at that particular altitude and time. These two terms have quite different correlations. In general, the error in the mean value of the state variable should be characterised by longer vertical correlation length and a smaller standard deviation than the natural variability. An additional, and more important, feature for inverting the time series is that errors in the mean will be correlated over long temporal periods compared to the natural variability. If, for example, the assumed a priori value for the concentration of one atmospheric species is too high with respect to the true mean at one time, it is likely that it will remain too high for a considerable time (weeks, months, etc.) thereafter.

We will use the retrieval of water vapour in the mesosphere as an example of the time series inversion method. Figure 1 shows covariance matrices used to describe the a priori water vapour profile in the retrievals. In Fig. 1a and b, the vertical and temporal covariances of the volume mixing ratio of water vapour at 60 km are shown. The dashed green curve represents the natural variability of water vapour, which is given a standard deviation of 50 %, a vertical correlation length of 4 km, and a temporal correlation length of 12 h. It should be noted that since information about the temporal correlation of water vapour at these altitudes is limited, these covariance matrices are created in a somewhat “ad-hoc” fashion. As such, the covariance matrices should be viewed as illustrative examples rather than a perfect representation of the atmospheric variability.

As mentioned earlier, the correlation lengths of the uncertainty in the a priori mean are different from the natural variability. The red dot-dashed curve in Fig. 1 shows properties of a covariance matrix set to represent the uncertainty in the a priori mean. The matrix has a standard deviation of 20 % and a correlation length of 8 km in altitude and 7 days in time. By adding both the natural variability and the a priori mean uncertainty, the complete covariance matrix, described by the solid red line (NatMean), is obtained. A selected number of elements from the complete \mathbf{S}_a matrix are shown in Fig. 1c. The block structure, explained in Sect. 2.1, is indicated by the green and black squares. The diagonal block (green square) represents the covariance within a measurement time, whereas the off-diagonal block (black square)

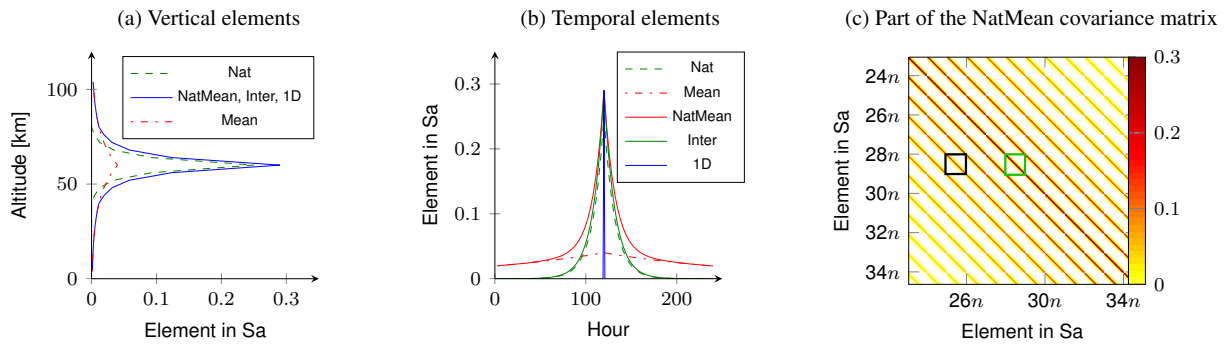


Fig. 1. The covariance matrices used for the a priori information on water vapour. **(a)** depicts the temporal and **(b)** the vertical elements of the covariance matrix. The matrices represent natural variability (Nat), uncertainty in a priori mean (Mean), the sum of natural variability and uncertainty in mean (NatMean), single spectrum inversions (1-D) and an intermediate covariance matrix (Inter). The structure of the NatMean a priori covariance matrix is shown in **(c)**. A diagonal matrix block is highlighted by the green square and an off-diagonal block is highlighted by the black square. The axis labels are multiples of n , i.e. the length of the single spectrum state vector.

represents the variance scaled with the correlation between measurement times.

The result from the time series retrievals will depend on the temporal correlation used. To investigate this, a second covariance matrix is created which assumes that the entire a priori uncertainty has a temporal correlation of only 12 h, but with the same standard deviation (54 %) and vertical correlation length as the total covariance matrix. This matrix is described by the solid-green line (Inter) in Fig. 1b. For comparison, inversions are also performed with zero correlation in time (blue curve) to mimic single spectrum (1-D) retrievals. Though these retrievals could be done on each spectrum separately, it was chosen, for comparison purposes, to perform the retrievals simultaneously using the same formalism as the time series inversions. This is achieved by using a block diagonal a priori covariance matrix in the retrievals.

A retrieval using the traditional method of averaging spectra is also performed by doing a 48 h running mean over the simulated spectra. However, the expected variance of a 48 h mean is different from that expected in a 3 h mean, so to correctly specify the covariance of these inversions, the NatMean covariance matrix is projected onto a 48 h grid following Rodgers (2000, Ch. 10.3.1.1). This results in an a priori standard deviation of 36 %, which is a decrease from the 54 % for the 3 h measurements.

3 Theoretical test case

3.1 The simulation and retrievals

In order to test the time series inversion method, a model scenario is set up. The simulation and retrievals are done with the atmospheric radiative transfer simulator (ARTS v.2.0) and the retrieval toolkit Qpack (Eriksson et al., 2005, 2011). The simulated instrument is designed to mimic the 22 GHz radiometer currently operating at OSO, but some

simplifications are made to illustrate the more general use of this method. Most notably, the bandwidth is increased from 20 MHz to 1 GHz and the noise temperature is reduced from 170 K to around 100 K. The instrument back end is simulated using 83 channels, each 25 kHz wide and unevenly distributed across the bandwidth of the instrument. At the line-centre, a distance of 25 kHz between the channels is used. This is increased further away from the centre, reaching 100 MHz at the band edges. The calibration used is a beam switching method.

Spectra from ground-based radiometers are often corrected for tropospheric loss before retrievals are performed. To model this, the simulated instrument is located above the troposphere (15 km) and the thermal noise level is doubled. This represents a tropospheric transmission of 0.5, which, together with the loss of observational time due to the beam switching, leads to an effective noise temperature of 400 K, which is used to specify S_{ϵ} . Additionally, the thermal noise in the system is left uncorrelated between the channels, and the integration time is set to 3 h. Note that no thermal noise is actually added to the simulated spectra, but only used to specify the covariance matrix, i.e. the spectra themselves are noise free.

The simulated atmosphere is created by extracting temperature and water vapour profiles for 25 February from the Mass Spectrometer and Incoherent Scatter (MSIS) (Hedin, 1991) temperature database and a climatology based on retrieved water vapour over OSO from the Microwave Limb Sounder on the Aura satellite (Aura-MLS). The spectroscopic parameters for the water vapour line at 22 GHz are taken from the JPL-catalogue (Pickett et al., 1998) (line strength and position) and HITRAN 2004 database (Rothman et al., 2005) (broadening parameters).

The retrieval of water vapour is done on an altitude grid ranging from 4 km to 104 km with a grid resolution of 4 km. The covariance matrices used are the same as specified in

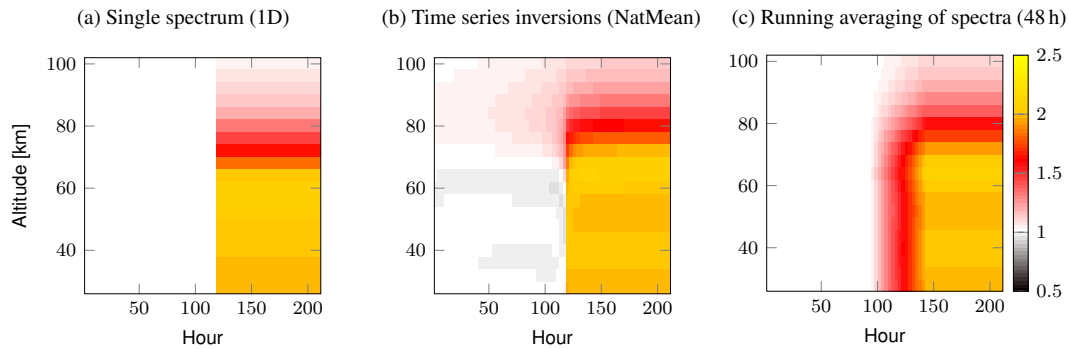


Fig. 2. Retrieved concentrations of water vapour, relative to a priori, from the simulated retrievals. The true water vapour concentration is doubled (i.e. set to two) at the 120th hour. **(a)** is from the single spectrum (1-D) inversions, **(b)** from the retrievals using the time series inversion method with the NatMean covariance matrix, and **(c)** is from the retrievals using spectra averaged over 48 h.

Sect. 2.3, with S_e being a pure diagonal matrix and S_a having a correlation in both altitude and time.

3.2 Response to a sudden doubling of H₂O

The time series inversion method combines information from measurements at several times, this makes the temporal characteristics of the retrieved profiles of particular interest. To investigate these temporal characteristics, we run a test scenario in which water vapour in the atmosphere is kept constant, equal to the a priori, until the 120th hour and then instantaneously increased to the double of a priori at the 121st hour.

The results of the retrievals are shown in Fig. 2. The results are shown as water vapour volume mixing ratio relative to the a priori volume mixing ratio. The single spectrum inversions (Fig. 2a) have no errors before the increase. This is because no noise was added to the spectra, and the true atmosphere is equal to the a priori. Afterwards, the retrievals will only change at certain altitudes determined by the measurement response. At the highest levels, the retrieved value remains 1, i.e. equal to a priori, due to the lack of measurement response, whereas at lower altitudes the retrieved values reflect the true atmosphere.

The conventional way to increase the measurement response at high altitudes is by averaging spectra in order to reduce the influence of thermal noise. Figure 2c shows results from the 48 h averaging of spectra. An improvement around 70 km can be seen for measurements later than 24 h after the sudden increase of water vapour in the atmosphere. This however, comes at the cost of smoothing out the step increase in time. Figure 2b shows that this smoothing can be avoided (for low altitudes) by applying the time series inversion method. Just as with the traditional averaging, the response around 70 km improves after the increase compared to the single spectrum inversions. However, the high temporal resolution is maintained at lower altitudes, and the abrupt change can clearly be seen in the retrievals. Thus, by inverting the entire time series simultaneously, temporal resolution

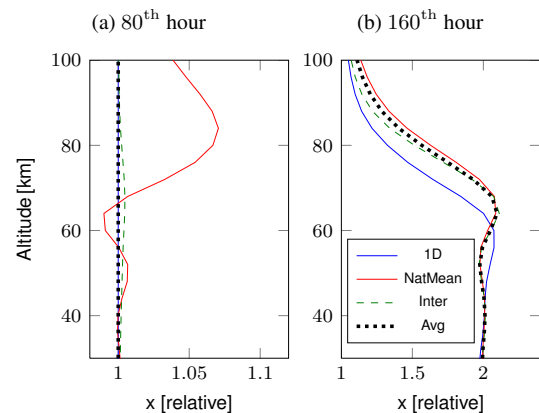


Fig. 3. Retrieved water vapour profiles (relative to the a priori) from the simulated retrievals. The different curves represent the single spectrum inversions (blue), the time series inversions using the NatMean (red line) and the intermediate covariance matrix (dashed green line), and the inversions using the averaged spectra (black dashed line). **(a)** is the retrieved profile at the 80th hour, when the true profile is equal to the a priori. **(b)** is the retrieved profile at the 160th hour, after the step-increase, when the true profile is double that of the a priori.

can be maintained at the lower altitudes while the sensitivity at higher altitudes is increased.

The increase in sensitivity is seen more clearly in Fig. 3b, which shows the profiles at the 160th hour. Both the traditional method of averaging spectra (black dashed line) and the time series inversions (red and dashed green lines) show a significant improvement above 60 km. The long temporal correlation in the a priori mean uncertainty does, however, lead to a large temporal smoothing at high altitudes seen by the increased water vapour above 70 km in the red line in Fig. 3a. The time series inversion method also leads to some oscillatory patterns shown by the negative values around 60 km in Fig. 3a.

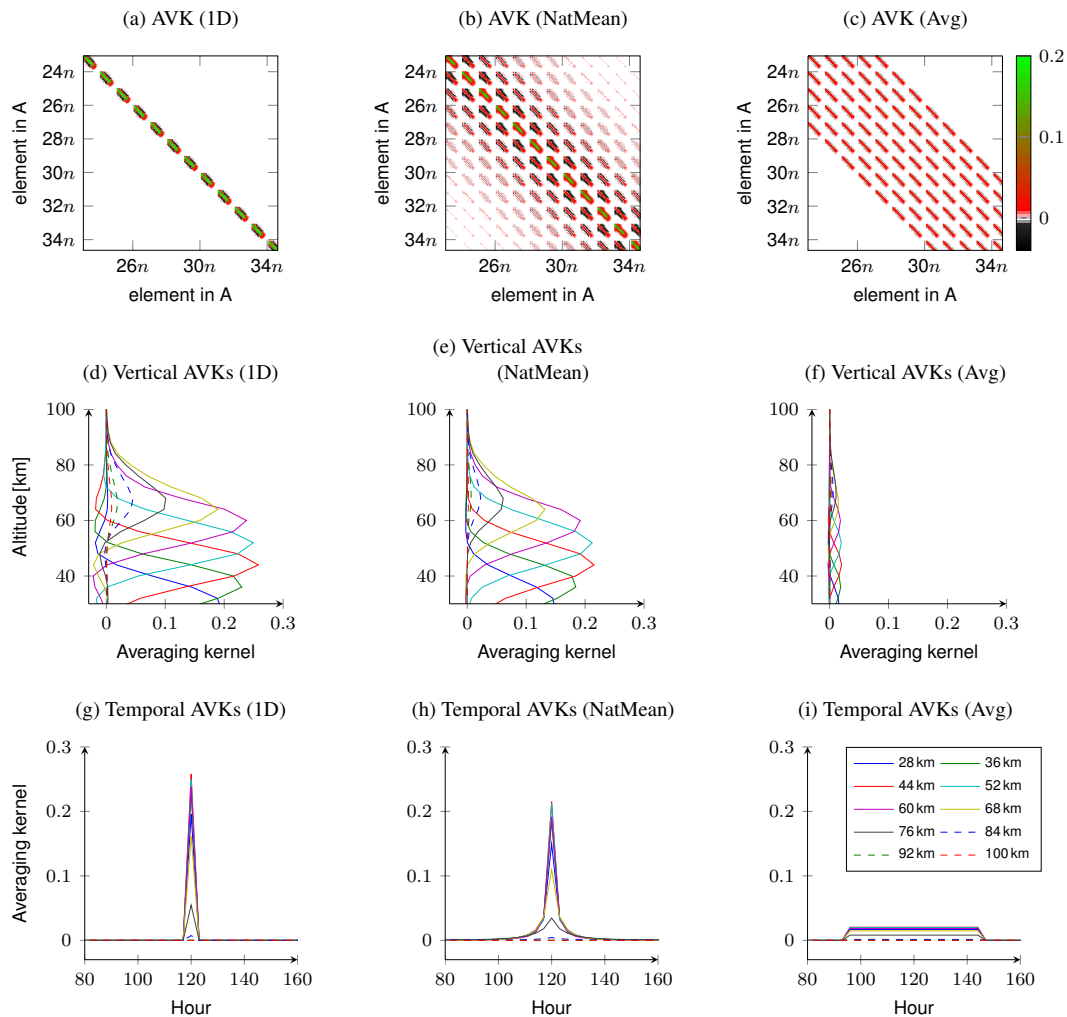


Fig. 4. Selected elements of the averaging kernels for the theoretical test case. The first row (a–c) shows the structure of the AVK matrices of single spectrum retrieval (left panel), time series retrieval with the NatMean a priori uncertainty matrix (centre panel) and 48 h averaging of spectra (right panel). The axis labels are multiples of n , i.e. the length of a single spectrum state vector. The second row (d–f) shows vertical AVKs for the respective cases and the third row (g–i) temporal AVKs for the respective cases. For the vertical and temporal averaging kernels, each of the curves corresponds to different altitudes.

3.3 Retrieval diagnostics

The temporal and vertical resolution of the inversions can be explored further by analysing the AVK matrices. Selected elements of the AVK matrices are shown in Fig. 4. The single spectrum inversions (Fig. 4a) give an AVK matrix which is completely diagonal with respect to time (at 3 h time resolution), meaning that the matrix is a block diagonal matrix where the non-zero elements are confined to elements no further away from the diagonal than the number of elements in the single measurement state vector. This is due to the fact that the single spectrum retrievals are performed as 2-D retrievals with no correlation in time in the \mathbf{S}_a matrix. If correlation between the days is introduced (Fig. 4b), the blocks adjacent to the diagonal block become non-zero and fall off exponentially from the diagonal. This implies

that a smoothing occurs in the temporal dimension, as already shown in Sect. 3.2. For direct averaging of the spectra (Fig. 4c), the AVK elements are constant across all blocks inside the averaging time, albeit reduced with a factor of roughly 1/16 compared to the single spectrum inversion to account for the averaging.

The averaging kernels are the rows of the AVK matrix. For clarity, it is convenient to focus on some particular elements of the rows. The first are the elements which correspond to the n columns around the diagonal, where n is the number of elements in the single spectrum state vector. These represent the vertical averaging kernel for each altitude. These kernels are seen in the second row of Fig. 4. The vertical averaging kernels for the three inversions are quite similar. Most notable is the reduction of the values that occurs for the 48 h

averaging, but this is compensated by the kernels spanning a larger number of measurements. Time series inversion produces vertical averaging kernels that have smaller negative values for elements within the same measurement compared to single spectrum inversions, but it should be noted that elements corresponding to different altitudes and different times can be negative (black areas in Fig. 4b).

The temporal averaging kernels describe the smoothing in time and are given by the elements corresponding to the same altitude for different times. These are shown in the third row of Fig. 4. The single spectrum inversions are simulated by performing the retrievals with zero correlation in time, and thus they have Dirac delta function kernels. The time series inversions have averaging kernels showing how the retrieval takes values from adjacent measurements into account. The lower values at the wings show how the inversions put diminishing weight measurements further away. For the averaging of spectra, the temporal AVKs have a constant value over the averaging time and zero elsewhere.

The full width at half maximum (FWHM) of the AVKs in the different dimensions can be used to roughly describe the resolution of the inversion in those dimensions. Figure 5a shows the FWHM of the temporal AVKs. The single spectrum inversions (blue curve) and the inversions averaging over spectra (black dashed curve) have a constant temporal resolution across all the altitudes corresponding to their respective averaging times. The red and the dashed green curves show the FWHM from the time series inversions. These inversions have a temporal FWHM which varies with altitude. At lower altitudes the FWHM is close to 3 h, i.e. the same as the single spectrum inversions. At higher altitudes the AVKs become wider indicating a reduction of temporal resolution as more information from adjacent measurements are used in the retrievals. Also, the larger a priori correlation between days of the NatMean matrix (red curve) compared to the intermediate (green curve) matrix results in wider AVKs. It is this wider averaging time at higher altitudes that allows the time series inversions to extend the retrievals higher than the single spectrum inversions.

There is some limitation in using only the FWHM to describe the resolution, as it does not take into account the full shape of the AVKs. The FWHM will have a different meaning for different shapes. For example, the FWHM in the temporal dimension of 48 h averaged retrievals will define where the averaging is cut off. For the exponentially shaped temporal AVKs of the time series inversions, however, the retrievals can have significant contributions from measurements beyond the FWHM. This explains why the time series inversion shows more temporal smoothing above 80 km in Fig. 2, yet, it has a smaller temporal FWHM in Fig. 5a.

Figure 5b shows the FWHM of the vertical AVKs. The FWHM is more or less the same for the inversions except for the 48 h averaging over spectra where the reduced noise in the measurements results in a better vertical resolution. Once again, some care should be taken when comparing the




Fig. 5. Properties of the simulated retrievals at the the 120th hour. (a) is the FWHM of the temporal AVKs, (b) is the FWHM of the vertical AVKs, (c) is the measurement response from the retrievals, and (d) is the retrieval noise from the retrievals in units relative to the a priori concentration. The different curves represent the single spectrum inversions (blue), the time series inversions using the Nat-Mean (red) and the intermediate covariance matrix (dashed green), and the inversions using the averaged spectra (black dashed).

FWHM from the different inversions. In particular, the negative lobes seen in the 1-D inversions will not be accounted for, and thus, AVKs with weaker lobes, like those from the time series inversions, will have a larger FWHM, though this mainly comes from the removal of the lobes and not a decrease in vertical resolution.

The measurement response corresponds, as anticipated, to the observed changes when x is doubled (Fig. 3b). The measurement response for the different inversions shows once again that the time series inversion method enables the retrieval of atmospheric values up to roughly the same altitude as the traditional averaging over spectra, actually exceeding the traditional averaging when using the NatMean covariance matrix.

Retrieval noise describes the error in the retrieved profiles from thermal noise, and is calculated as $\mathbf{G}\mathbf{S}_\epsilon\mathbf{G}^T$ (Rodgers, 2000). Figure 5d shows the square root of the diagonal elements of the retrieval noise matrix. As anticipated, the reduction of thermal noise in the measurements from the traditional averaging method (black dashed line) will result in a lower retrieval noise compared to the single spectrum inversions (blue line). This is the result of both a reduction in the thermal noise due to averaging and the change from

using a different a priori uncertainty. The retrieval noise in the time series inversions ends up a bit below the single spectrum inversions. This shows that some error reduction is achieved with the time series method, but that the main improvement it offers is the increased measurement response at high altitudes.

In addition to the diagonal elements, the retrieval noise covariance matrix will have non-diagonal elements arising from temporal and vertical correlations. This means that the retrieval noise for the time series inversions has a correlation in time even though the underlying thermal noise is uncorrelated in time. This correlation is introduced through the temporal correlation in the a priori constraints. The FWHM of this correlation (not shown) can be different from that of the AVKs. For the time series inversions (NatMean), it is roughly 15 h up to around 70 km, above this it increases and reaches 50 h at 85 km. For the intermediate a priori covariance matrix inversions, the temporal FWHM of the retrieval noise stays at around 15 h for all altitudes.

4 Test using a real instrument

To illustrate the practical use of the time series inversion method, we invert atmospheric spectra measured from the water vapour radiometer at OSO. When using real measurements, instrument-related issues might degrade the efficiency of the retrieval, or introduce biases, which complicates the retrievals and error analysis.

4.1 OSO radiometer

The radiometer used to test the time series inversion method is placed at OSO (57.4° N, 12° E). It measures water vapour at 22.235 GHz with a resolution of 25 KHz and bandwidth of 20 MHz. The system has an uncooled High-electron-mobility transistor (HEMT) front end and uses a 800 channel autocorrelator back end. Receiver temperature is estimated to 170 K and the calibration is done by a hot-cold calibration and beam switching. The spectra are also corrected for tropospheric absorption. Each spectrum consists of 5 min measurements averaged together into six 3 h intervals for each day, 0–3, 4–7, 8–11, 12–15, 16–19 and 20–23. The averaging is done so that measurements with lower noise values have a larger weight in the average. This is done since we assume that the thermal noise in the measurements is uncorrelated to the volume mixing ratio of water vapour during the 3 h interval. The thermal noise of each 3 h spectrum is determined separately by fitting a 3rd order polynomial to one of the line-wings and calculating the standard deviation of the residual. For a complete description of the instrument and retrieval parameters, see Forkman et al. (2003) and Haefele et al. (2009).

When performing the time series inversions over all days and all channels, the dimension of the matrices in Eqs. (3), (4) and (5) become quite large. To reduce the size of the

matrices, only a subsample of the channels is selected, as in the theoretical test case. In the line-centre, all channels are used, but at the line-wings, the channel separation is increased gradually, reaching 620 kHz at the far ends. In total, 83 of the 800 channels are used. Furthermore, the retrievals are performed in 30-day intervals. To minimise edge effects each interval has a 10-day overlap, which allows for 5 days on each end of the retrieval intervals to be removed. These intervals are then combined to create the complete time series. Further discussion regarding the computational demands can be found in Sect. 5.

Just as in the theoretical test case, the retrievals are performed on each 3 h spectrum separately, with the time series inversion method, and a 48 h moving average of spectra. However, since the thermal noise in the measurements varies with time, the simple averaging is replaced with a weighted average giving measurements with lower noise more weight. This is done to simplify the comparison with the time series method. In addition, since some gaps exist in the measurements, some 48 h averages have fewer measurements than the nominal 12.

The inversions are set up as described in Sect. 3, except that since the noise level varies with time, the thermal noise in each measurement must be estimated from each corresponding spectrum rather than having a constant noise level as in Sect. 3.3. Additionally, an instrumental baseline (5th order polynomial) is added to the state vector and retrieved. This polynomial fit is given a priori uncertainties from 10 K (0th order) to 2 K (5th order), and without correlation in time. Since the atmosphere over OSO changes over time, the a priori is also set to vary according to the climatologies (temperature and water vapour) in Sect. 3, rather than have a constant value.

4.2 Dealing with measurement gaps

The OSO time series has periods where no measurement data could be recorded. These periods are mainly caused by rain. Data gaps create additional problems when handling measurement series. A time interpolation using neighbouring retrieved data requires the user to select an interpolation strategy (nearest, linear, spline, etc.), as well as to make subjective judgements on the validity of these interpolated values based on experience and knowledge of the atmospheric variables measured.

The time series inversion method provides an elegant solution to this problem. To obtain values at the gaps, \mathbf{x} is expanded to cover the times where measurements are lacking. This increases the size of \mathbf{x} to $n \cdot (N + N')$, where N' is the number of missing measurements, and N the number of times with measurements. n and later m are defined as in Sect. 2.2. The expansion of \mathbf{x} allows the MAP algorithm to retrieve the missing values, maintaining a consistent inversion methodology over the complete time period.




Fig. 6. H₂O retrievals from the summer of 2005 obtained by the OSO radiometer. The results are relative to the a priori concentration and are shown for 76 km (a), 64 km (b) and 52 km (c). The different curves represent the single spectrum inversions (blue), the time series inversions using the NatMean (red) and the intermediate covariance matrix (dashed green), and the inversions using the averaged spectra (black dashed).

Since the size of \mathbf{x} is increased, the number of columns in \mathbf{K} must increase correspondingly. This gives \mathbf{K} a size of $N \cdot m \times (N + N')$. The elements in the N' extra columns will be zero as no measured spectra exists for this time in y . Physically this is equivalent to only using a virtual measurement of the a priori atmosphere at the time of the data gap. However, this does not mean that only a priori information is used for the retrieval of the corresponding state. Since \mathbf{S}_a contains information about the temporal correlation of the atmosphere, the MAP method will automatically use information from a neighbouring measurement to “optimally” estimate \mathbf{x} at the time of the measurement gap.

For an “interpolated” value, the amount of information taken into account from nearby measurements is given by the corresponding measurement response. The measurement response will depend on the a priori uncertainty matrix used and the amount of noise in the adjacent measurements. For the single spectrum inversions, the measurement response will be zero at the interpolated values, whereas for the time series inversions it will increase with increasing a priori temporal correlation. The measurement response will provide a value on which the validity of the interpolated value can be determined. This value is based on the underlying

statistical properties of the retrievals, and thus a consistent selection scheme can be applied, for example, by only using data points above a certain measurement response threshold (e.g. 0.8). It should, however, be noted that for further data analysis (e.g. trend estimation), the exact influence that the a priori and AVKs will have on the analysis must be considered, just as with data from any other retrieval methods based on MAP.

4.3 Result from time series inversions

Retrievals from the OSO instrument were done for the entire measurement period (2002–2012). For comparison of the different inversions methods, an example period from the end of April to the end of June 2005 was selected for further study as this period offers a long set of continuous measurements, with few measurement gaps. The results of the retrievals at three different altitudes are shown in Fig. 6 in units relative to the a priori (i.e. 1 = a priori concentration). These results include estimated values where data gaps occur (13 of 368 times), interpolated using the method discussed in Sect. 4.2.




Fig. 7. 1σ of the thermal noise from the OSO radiometer from the summer of 2005, estimated as described in the text. The measurements are performed with 3 h integration time. The circles show the three days selected for further AVK analysis.

By comparing the single spectrum retrievals (blue curve) to the retrieval of 48 h averaged spectra (black dashed curve), the effect of the averaging can be seen. The variability is reduced from 1σ of ~ 0.22 in the single spectrum inversions to 1σ of ~ 0.15 in the averaged ones, with the averaged spectra having a longer temporal correlation. This correlation has two causes. The first one is the temporal correlation of the atmospheric changes over the instrument, the other cause is the thermal noise which, as discussed earlier, will also have a correlation in time when averaging is performed. In addition to the change in variability, the averaged spectra show a clearer deviation from the a priori at higher altitudes. This shows the effect of the increased measurement response at these altitudes.

The time series inversions (red and dashed green curves) show an increased measurement response at higher altitudes similar to the 48 h averaged inversions, and the variation is correlated over several days. The more longer-term averages, over a couple of days, seem to follow the averaged inversions. At 76 km the measured mean volume mixing ratio over the entire period is actually lower than the a priori mean volume mixing ratio. This illustrates why it is important to include the uncertainty in the a priori mean in the inversions. Without this part, the measurement response would be lower and the inversions would not reveal this information. At lower altitudes (52 and 64 km), the time series inversions preserve many of the short-term variations, indicating a high temporal resolution at these altitudes.

It is hard to distinguish whether the short-term variations are a result of noise in the instrument or natural variance in the atmosphere. However, the main point of these retrievals is not to determine the true water vapour concentration in the atmosphere, but rather to show that the time series inversions produce comparable results to the single spectrum inversions at lower altitudes while producing results more similar to those of the averaged inversions higher up.

4.4 Averaging kernels

Since the thermal noise in the real measurements varies with time, the AVKs vary as well. Thus, to study some typical AVKs, three dates are selected for further inspection. Figure 7 shows the magnitude of the thermal noise in each of the measurements from the time series in Fig. 6, and the selected

measurements are marked by the three circles. The first measurement (10 May 005, red circle) is from a measurement with a low noise value. The second (10 May 2005, green circle) and third measurements (10 May 2005, blue circle) are separated by only four hours and have a high and intermediate thermal noise value, respectively.

The measurement response of the three measurements is shown in the top row of Fig. 8. The measurement response of the low-noise measurement is similar to the theoretical test case above 60 km. Below 60 km the measurement response starts declining due to the fitting of the instrumental baseline polynomials. The similarity above 60 km is not surprising considering that the noise of the measurement is 0.043 K, which resembles the noise in the test case of 0.037 K. The two other cases, however, have much higher noise than the test case with values of 0.16 and 0.07 K. This results in a very low measurement response for single spectrum inversions, but the time series inversions and the averaged spectra still have a good measurement response between 55 and 75 km.

The amount of information taken from each measurement is given by the temporal averaging kernels shown in the second row of Fig. 8. However, unlike the averaging kernels from the theoretical test case, the temporal averaging kernels are not smoothly exponentially declining (see Fig. 4), but vary depending on the weight placed on each adjacent measurement. This variation comes from the fact that the MAP method puts a lower weight on noisier measurements. In fact, the temporal averaging kernels from the measurement with intermediate noise (Fig. 8f) show that the inversion gives little weight to the noisy measurement from 3 h earlier. In the high-noise case (Fig. 8e), it can be seen that the adjacent measurement is actually weighted more than the central measurement, as the AVK has its peak displaced from the centre.

The last row of Fig. 8 shows the FWHM of the temporal AVKs. For the low-noise measurement (Fig. 8g), the FWHM is similar to the theoretical test case above 60 km having a minimum width at around 60 km before increasing in width with increasing altitude. For the measurements with higher noise, however, the irregular shape of the temporal AVKs means that the interpretation of the FWHM is not as straightforward as in the theoretical case. The maximum might not

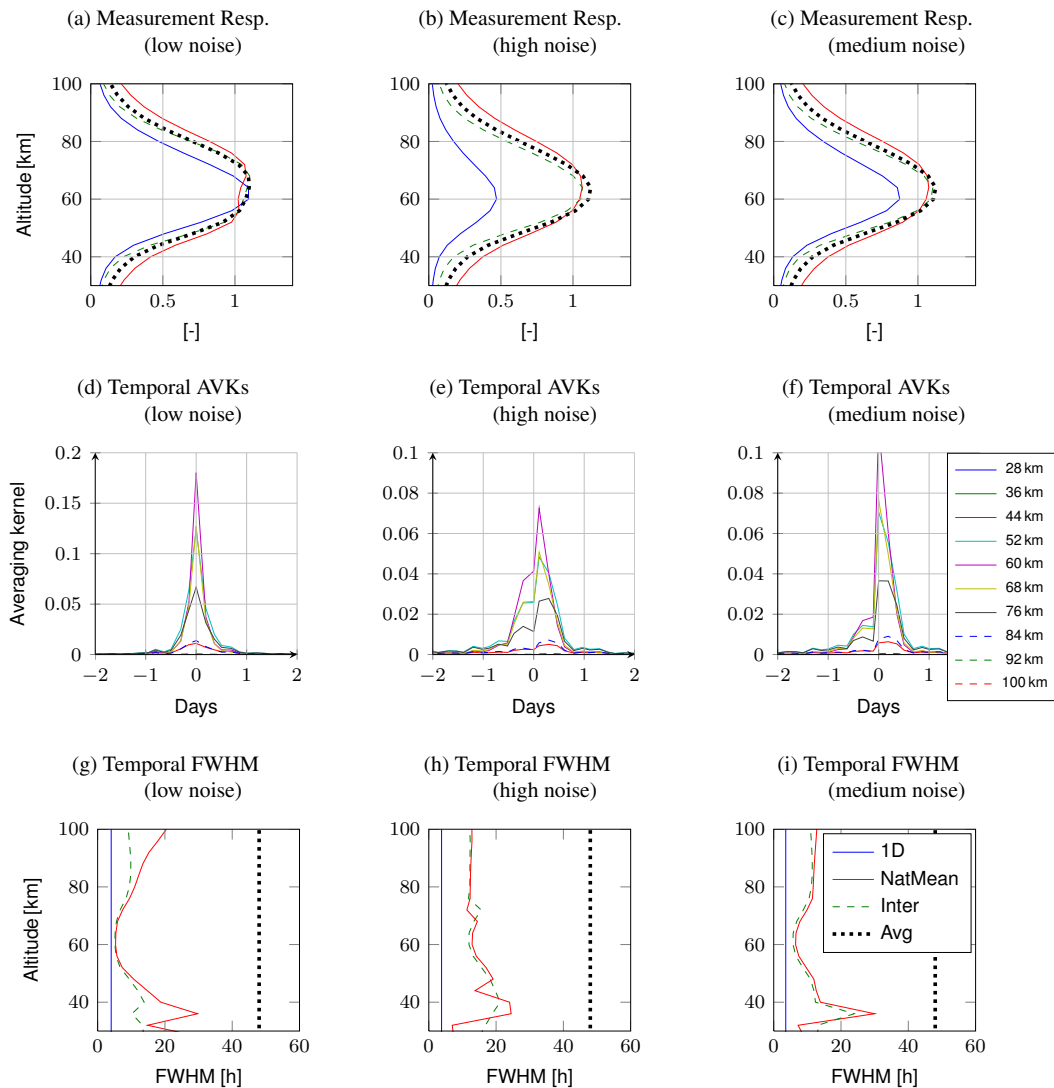


Fig. 8. Properties of the AVKs from three selected measurements from the OSO radiometer. The left column shows the low-noise measurement (10 May 2005), the centre column shows the high-noise measurement (21 May 2005), and the right column shows the intermediate noise measurement (21 May 2005). The first row (a–c) depicts the measurement response of the respective measurements, the second row (d–f) temporal averaging kernels of different altitudes, and the third row (g–i) the FWHM of the temporal AVKs. The different curves in the first and last rows represent the single spectrum inversions (blue), the time series inversions using the NatMean (red) and the intermediate covariance matrix (dashed green), and the inversions using the averaged spectra (black dashed).

be centred at zero, and the position of the half-value point might even be ambiguous. As a result the FWHM of the temporal AVKs from these measurements (Fig. 8h and i) differs quite a lot from the theoretical test case, especially for the high-noise case where it fluctuates at lower altitudes.

5 Computational demands

Though there are several advantages of expanding the inversions into the temporal dimension, a drawback is the increased computational demand. The computational demand includes both larger memory usage and the increased number

of CPU operations required for the matrix operations. This paper will not go into detail on optimising the efficiency of the retrievals, but a discussion of the major issues is required.

Depending on the retrieval setup, either \mathbf{K} , \mathbf{S}_ϵ or \mathbf{S}_a will have the largest memory demand. All three matrices tend to be diagonal heavy (i.e. highest values around the diagonal), thus considerable memory can be saved storing them as sparse matrices. For \mathbf{S}_a , this could require some cut-off value for the covariance, as the exponential correlation theoretically never reaches zero.

Considering CPU cycles, the linear algebra can be optimised for either $n < m$ (n-form) or $n > m$ (m-form) (Rodgers,

2000), where n and m (and later N) are defined as in Sect. 2.2. For the retrievals in this paper $n < m$, and the most demanding operations are the calculation of $\mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{K}$, which scales as $N^3 n^2 m$, and \mathbf{S}_a^{-1} and $(\mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{K} + \mathbf{S}_a^{-1})^{-1}$, which both scale as $N^3 n^3$.

The available computational power limits the size of N , n , and m . In the retrievals from the OSO radiometer, N is limited to measurements from 30 consecutive days ($N \sim 180$). To limit m , we use a simple solution of only selecting a 83-channel subset of the 800 channels in the spectrometer. Other methods make it possible to take advantage of all the channels while keeping down the size of m . These might be as straightforward as binning channels at the line wing, i.e. averaging the channels together to reduce the noise, or more advanced data reduction methods based on eigenvector expansions (e.g. Eriksson et al., 2002).

In addition to reducing the size of the matrices, the algebra itself can be optimised. In particular, an explicit inversion of $(\mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{K} + \mathbf{S}_a^{-1})$ can be avoided by instead solving Eq. (3) using methods such as Cholesky decomposition (Livesey et al., 2006) or the iterative bi-conjugate gradient method (Reburn et al., 2000).

Data reduction algorithms and optimisation of the linear algebra might indeed improve the practical use of the time series inversion methods, but a thorough discussion of such optimisation is beyond the scope of this paper as it will be highly dependent on the specific retrieval setup and needs.

6 Discussion and conclusion

This paper presents a method for inverting time series data from ground-based instruments by extending the retrieval method into the temporal dimension. This is done by directly specifying the correlation of the atmosphere in time to achieve “optimal averaging” at all altitudes. The implications and analysis of the temporal averaging kernels are discussed thoroughly in the paper, including their importance and limitations in describing the temporal resolutions of the retrievals.

To investigate the effect of using different temporal correlations, the time series inversions are performed with two different a priori matrices: one modelled to represent a realistic a priori uncertainty (NatMean), and one intermediate matrix with shorter temporal correlation. Interestingly enough, in both the simulated retrievals (Fig. 5c) and the practical example (Fig. 8a–c), the retrieval using the intermediate covariance matrix shows almost the same increase in measurement response between 60 and 80 km as the retrievals using the realistic covariance matrix. This is confirmed as the difference in the retrieved data from the OSO radiometer between the two matrices (Fig. 6) is minuscule below 80 km.

The similar increase in measurement response for both matrices shows that the major improvement of the time series inversions comes from the basic step of extending the

inversions into the temporal dimension rather than to specify the covariance matrix in detail. This is important for the practical use of the method since it means that the method can be applied to cases where the temporal correlation is unknown, or hard to specify. In these cases care should however be taken, especially if the atmospheric variable has a systematic, e.g. diurnal, variation in time. Such variations might be hard to specify using a normal distribution described by a simple covariance matrix, and as such might be unsuited for MAP retrievals.

The demonstration of the method using real data retrieves 10 yr of water vapour data from the OSO radiometer. During this test case the computation time for the time series method was less than one order of magnitude larger than for single spectrum inversions. This shows that the computational demand, though increased, is not insurmountable. For large-scale retrievals, however, further optimisation might be advantageous, in particular, the data reduction method for reducing the size of the measurement vector can easily be improved.

The practical demonstration also shows how the time series method can be used to interpolate data to times where no measurements are performed. The interpolation is carried out directly during the retrieval. It is based on the same underlying a priori statistics of the atmosphere, and it automatically takes into account the quality of the nearby measurements. By using the measurement response, a selection of the valid interpolated values can be made. This selection is consistent with the retrieval, and removes the need for ad hoc, post-processing selection algorithms to fill data gaps.

Some earlier studies have also used the temporal dimension in the retrievals; in particular, the Aura-MLS retrieval of “noisy” products (Livesey et al., 2006) is similar to the time series method suggested here. The difference is that, whereas the MLS retrievals invert all spectra simultaneously into *one* mean profile over the entire time period, the time series inversions are based on the MAP method and determine the optimal averaging period at each point and produce a complete time series.

The advantages of using the time series inversion technique will depend on the instrument and species studied. This paper has focused on water vapour retrieval from a microwave radiometer, but the method will similarly benefit instruments retrieving other species such as O_3 , or using other methods, such as Fourier transform infrared spectroscopy. Another useful application of the method is the retrieval of several species, or atmospheric variables requiring different averaging times.

An additional, interesting aspect of the approach is the possibility to also consider time correlations of instrument variables. For example, the practical test case used here included the retrieval of polynomial coefficients to describe the “baseline ripple”. The a priori variability of these coefficients are set to be the same, independent of integration time, and uncorrelated between measurements. However, if

the temporal correlation of baseline changes is determined, it can be incorporated in time series inversion. This would result in an extension of the measurement response downwards compared to retrievals using a single spectrum or 48 h averaged spectra.

The time series inversion technique offers several advantages over traditional averaging. In particular it offers a way to produce a single consistent dataset from retrievals that span a wide set of altitudes, optimising the temporal resolution at each altitude. This removes the need for multiple datasets for variables requiring different integration times. However, as the method increases the dimensions of the retrieval and resulting averaging kernels, it also increases the complexity for any end user using the data.

Acknowledgements. We would like to thank Peter Forkman for providing measurement data from the OSO water vapour radiometer. We would also like to thank Marston Johnston and the anonymous reviewers for valiantly proofreading the manuscript.

Edited by: K. Strong

References

- Askne, J. and Westwater, E.: A review of ground-based remote sensing of temperature and moisture by passive microwave radiometers, *IEEE T. Geosci. Remote*, GE-24, 340–352, 1986.
- Baron, P., Ricaud, P., d. Noë, J. L., Eriksson, P., Merino, F., and Murtagh, D.: Studies for the Odin sub-millimetre radiometer: Retrieval methodology, *Can. J. Phys.*, 80, 341–356, 2002.
- Bleisch, R., Kämpfer, N., and Haefele, A.: Retrieval of tropospheric water vapour by using spectra of a 22 GHz radiometer, *Atmos. Meas. Tech.*, 4, 1891–1903, doi:10.5194/amt-4-1891-2011, 2011.
- Eriksson, P.: Analysis and comparison of two linear regularization methods for passive atmospheric observations, *J. Geophys. Res.*, 105, 18157–18167, doi:10.1029/2000JD900172, 2000.
- Eriksson, P., Jiménez, C., Bühler, S., and Murtagh, D.: A Hotelling transformation approach for rapid inversion of atmospheric spectra, *J. Quant. Spectrosc. Ra.*, 73, 529–543, doi:10.1016/S0022-4073(01)00175-3, 2002.
- Eriksson, P., Jiménez, C., and Buehler, S. A.: Qpack, a general tool for instrument simulation and retrieval work, *J. Quant. Spectrosc. Radiat. Transfer*, 91, 47–64, doi:10.1016/j.jqsrt.2004.05.050, 2005.
- Eriksson, P., Buehler, S., Davis, C., Emde, C., and Lemke, O.: ARTS, the atmospheric radiative transfer simulator, version 2, *J. Quant. Spectrosc. Ra.*, 112, 1551–1558, doi:10.1016/j.jqsrt.2011.03.001, 2011.
- Forkman, P., Eriksson, P., and Winnberg, A.: The 22 GHz radio-astronomy receiver at Onsala Space Observatory, *J. Quant. Spectrosc. Ra.*, 77, 23–42, doi:10.1016/S0022-4073(02)00073-0, 2003.
- Forkman, P., Eriksson, P., and Murtagh, D.: Observing the vertical branch of the mesospheric circulation at lat N60° using ground based measurements of CO and H₂O, *J. Geophys. Res.*, 110, D05107, doi:10.1029/2004JD004916, 2005.
- Haefele, A., Wachter, E. D., Hocke, K., Kämpfer, N., Nedoluha, G. E., Gomez, R. M., Eriksson, P., Forkman, P., Lambert, A., and Schwartz, M. J.: Validation of ground-based microwave radiometers at 22 GHz for stratospheric and mesospheric water vapor, *J. Geophys. Res.*, 114, D23305, doi:10.1029/2009JD011997, 2009.
- Hedin, A. E.: Extension of the MSIS thermosphere model into the middle and lower atmosphere, *J. Geophys. Res.*, 96, 1159–1172, doi:10.1029/90JA02125, 1991.
- Livesey, N., Van Snyder, W., Read, W., and Wagner, P.: Retrieval algorithms for the EOS Microwave limb sounder (MLS), *IEEE T. Geosci. Remote*, 44, 1144–1155, doi:10.1109/TGRS.2006.872327, 2006.
- Nedoluha, G., Gomez, R., Hicks, B., Helmboldt, J., Bevilacqua, R., and Lambert, A.: Ground-based microwave measurements of water vapor from the midstratosphere to the mesosphere, *J. Geophys. Res.*, 116, D02309, doi:10.1029/2010JD014728, 2011.
- Pickett, H., Poynter, R., Cohen, E., Delitsky, M., Pearson, J., and Muller, H.: Submillimeter, millimeter, and microwave spectral line catalog, *J. Quant. Spectrosc. Ra.*, 60, 883–890, 1998.
- Reburn, W. J., Siddans, R., Kerridge, B., Bühler, S., von Engeln, A., Eriksson, P., Kuhn-Sander, T., Künzi, K., and Verdes, C.: Critical assessments in millimetre-wave atmospheric limb sounding, Tech. rep., ESTEC Contract No. 13348-98-NL-GD, European Space Research and Technology Centre, 2000.
- Rodgers, C.: Inverse methods for atmospheric sounding: Theory and practice, World Scientific, Singapore, 2000.
- Rothman, L. S., Jacquemart, D., Barbe, A., Chris Benner, D., Birk, M., Brown, L. R., Carleer, M. R., Chackerian Jr., C., Chance, K., Coudert, L. H., Dana, V., Devi, V. M., Flaud, J.-M., Gamache, R. R., Goldman, A., Hartmann, J.-M., Jucks, K. W., Maki, A. G., Mandin, J.-Y., Massie, S. T., Orphal, J., Perrin, A., Rinsland, C. P., Smith, M. A. H., Tennyson, J., Tolchenov, R. N., Toth, R. A., Vander Auwera, J., Varanasi, P., and Wagner, G.: The HITRAN 2004 molecular spectroscopic database, *J. Quant. Spectrosc. Ra.*, 96, 139–204, 2005.
- Scheiben, D., Straub, C., Hocke, K., Forkman, P., and Kämpfer, N.: Observations of middle atmospheric H₂O and O₃ during the 2010 major sudden stratospheric warming by a network of microwave radiometers, *Atmos. Chem. Phys.*, 12, 7753–7765, doi:10.5194/acp-12-7753-2012, 2012.
- Seele, C. and Hartogh, P.: Water vapor of the polar middle atmosphere: Annual variation and summer mesosphere conditions as observed by ground-based microwave spectroscopy, *Geophys. Res. Lett.*, 26, 1517–1520, doi:10.1029/1999GL900315, 1999.
- Seele, C. and Hartogh, P.: A case study on middle atmospheric water vapor transport during the February 1998 stratospheric warming, *Geophys. Res. Lett.*, 27, 3309–3312, doi:10.1029/2000GL011616, 2000.