



Improving retrieval quality for airborne limb sounders by horizontal regularisation

J. Ungermann^{1,2}

¹Institute of Energy and Climate Research – Stratosphere (IEK-7), Research centre Jülich GmbH, Jülich, Germany

²National Center for Atmospheric Research, Boulder, Colorado, USA

Correspondence to: J. Ungermann (jorn@ucar.edu)

Received: 18 June 2012 – Published in Atmos. Meas. Tech. Discuss.: 13 September 2012

Revised: 20 November 2012 – Accepted: 3 December 2012 – Published: 2 January 2013

Abstract. Modern airborne infrared limb sounders are capable of measuring profiles so fast that neighbouring profiles are very similar to one another. This can be exploited by retrieving whole 2-D cross-sections instead of simple 1-D profiles.

This paper presents algorithms that are able to perform such a large-scale retrieval and that efficiently produce typical diagnostic quantities. The characteristics and capabilities of the proposed method are analysed and demonstrated in a detailed case study using a series of profiles that were measured by CRISTA-NF (Cryogenic Infrared Spectrometers and Telescope for the Atmosphere–New Frontiers).

It is shown that cross-section retrievals can either reduce noise-induced artefacts or produce finer vertical structures while maintaining the same image noise level. Further, it is discussed how the presented methodology can also be applied to improve the retrievals for other instrument types including current satellite-borne nadir-sounders and near-future satellite-borne limb sounders.

1 Introduction

Two-dimensional (2-D) cross-sections of trace gas volume mixing ratios derived from limb sounder measurements taken from aircraft are a valuable tool for examining mixing processes in the upper troposphere/lower stratosphere. This region is of special importance, as it significantly influences radiative forcing (e.g. Forster and Shine, 1997); it also provides the boundary between tropospheric and stratospheric air, so its dynamic properties largely determine the exchange between these layers (e.g. Gettelman et al., 2011). To examine

such dynamic structures, one would like to have a highly resolved 3-D image of the atmosphere including information on temperature, pressure, and trace gases of different lifetimes. While there are promising instruments in development that may deliver such highly resolved volumes (e.g. Friedl-Vallon et al., 2006; Ungermann et al., 2010b), current limb imagers can only provide 2-D cross-sections.

Using high-flying aircraft (such as the Russian M55 Geophysica or the German research plane HALO) as carriers, limb sounders are typically mounted to be looking towards the side. This enables them to have an excellent resolution in flight direction (limited mostly by acquisition time), a good resolution in the vertical direction (limited by their vertical field of view; typically on the order of several hundred meters or more) and a much worse resolution perpendicular to the flight path (on the order of several hundred km; e.g. Ungermann et al., 2012a). The latter is obviously a problem, if one tries to resolve atmospheric structures of smaller extent. However, modern (chemical) forecasts are often reliable enough to enable the aircraft to cross long-drawn filamentary structures orthogonally. Under such circumstances, be they achieved by meticulous flight planning or pure chance, the averaging along the line of sight becomes less important than the vertical and horizontal resolution.

Especially the vertical resolution is a key factor when examining very thin layers in the atmosphere such as the tropopause inversion layer or tropopause folds with an extent of only 1 km or less (e.g. Birner, 2006; Hegglin et al., 2009). A vertical resolution of 500 m, already quite respectable for a limb sounder, generates only one to two vertical data points for such structures and may not reproduce thinner layers at all.

This paper describes a retrieval scheme that is able to produce 2-D cross-sections with both better vertical resolution and generally fewer noise-induced artefacts. Such cross-sections can be more easily analysed due to the consistency between neighbouring profiles and lower smoothing in the vertical directions. Obviously, the lower limit for the improvement is given by the vertical sampling, which needs to be fine enough to sample structures of interest. This paper introduces the concept of cross-section retrievals, where a whole cross-section consisting of many individual profiles is jointly retrieved in the usual non-linear fashion. This is in principle similar to the concept of tomographic 2-D retrievals of satellite limb sounder data, which were first produced by Carlotti et al. (2001) and Steck et al. (2005) for the Michelson Interferometer for Passive Atmospheric Sounding (MIPAS) and by Livesey and Read (2000) for the Microwave Limb Sounder. However, the cross-section retrievals presented here do lack the tomographic component, as consecutively measured radiance profiles do measure completely different air masses. Also, the intent is different, as tomographic satellite retrievals are primarily used to improve along line-of-sight resolution, while the cross-section retrieval presented here shall instead improve the vertical resolution of the retrieval result at the cost of a slightly worse along-flight track resolution.

The retrieval of trace gases or other quantities from infrared nadir sounder or limb sounder measurements is inherently an ill-posed problem, implying that no unique solution may exist and that small measurement errors may cause large differences in retrieval values. One way to obtain a unique and physically reasonable solution is regularisation. Regularisation approximates the original ill-posed problem by a similar well-posed problem. This is typically accomplished by adding constraints to the solution. From here on, the term horizontal refers solely to the along-flight track direction. By introducing horizontal regularisation, i.e. constraints on the horizontal variation of retrieved values, the influence of stochastic error sources becomes reduced. Further, this paper introduces algorithms to provide an associated set of typical diagnostic information following the linear framework of Rodgers (2000). Last, the feasibility of the approach is demonstrated upon analysing measurements of the Cryogenic Infrared Spectrometers and Telescope for the Atmosphere–New Frontiers (CRISTA-NF) taken during the RECONCILE (reconciliation of essential process parameters for an enhanced predictability of Arctic stratospheric ozone loss and its climate interactions) campaign. The effect of horizontal regularisation on the retrieval quality, the impact of measurement noise, and the achievable horizontal resolution is examined and discussed. It is demonstrated that introducing horizontal regularisation allows lowering the vertical regularisation strength in order to reproduce finer features than otherwise possible. Last, this more complicated method is compared against the weighted a posteriori averaging of conventionally retrieved 1-D profiles.

2 Cross-section retrieval

The retrieval is the process of deriving atmospheric quantities such as temperature or trace gas volume mixing ratios from remote sensing measurements. In practise, a discrete representation of the atmospheric state $\mathbf{x} \in \mathbb{R}^n$ is modified under certain constraints until the fit between the real measurements $\mathbf{y} \in \mathbb{R}^m$ and simulated measurements $\mathbf{F}(\mathbf{x})$ produced by a forward model $\mathbf{F} : \mathbb{R}^n \mapsto \mathbb{R}^m$ is deemed good enough according to some quality criteria. Thereby \mathbf{F} is a sufficiently accurate forward model capable of simulating the physics of radiative transport and instrument principle. If the forward model is linear, a solution can be directly derived while non-linear models require an iterative procedure. As the retrieval problem is ill-posed, the addition of some kind of a priori information is usually required to generate physically reasonable results. This is formalised by introducing a cost function $J : \mathbb{R}^n \mapsto \mathbb{R}$:

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

This specific cost function is a quadratic form, which has the advantage that its minimum can be found rather efficiently. It consists of two distinct terms. The first term penalises differences between the measurements and the radiances provided by the forward model weighted by the measurement error covariance matrix $\mathbf{S}_\epsilon \in \mathbb{R}^{m \times m}$. This matrix describes the available knowledge about measurement errors, thereby assuming that they are Gaussian distributed. The second term penalises differences between the derived atmospheric state and so-called a priori knowledge, which is usually taken from a climatology, from models, or from ad hoc assumptions about smoothness of atmospheric quantities. Thereby \mathbf{x}_a is the expected atmospheric state and $\mathbf{S}_a \in \mathbb{R}^{n \times n}$ is a covariance matrix thereof codifying correlations between atmospheric quantities.

A state vector minimising the cost function J can then be found, e.g. using an iterative truncated quasi-Newton method (often within less than 10 iterations) (e.g. Nocedal and Wright, 2006):

$$\mathbf{x}_{i+1} = \mathbf{x}_i - \left(\mathbf{S}_a^{-1} + \mathbf{F}'(\mathbf{x}_i)^T \mathbf{S}_\epsilon^{-1} \mathbf{F}'(\mathbf{x}_i) + \lambda_i \mathbf{I}_n \right)^{-1} \cdot \left(\mathbf{S}_a^{-1} (\mathbf{x}_i - \mathbf{x}_a) + \mathbf{F}'(\mathbf{x}_i)^T \mathbf{S}_\epsilon^{-1} (\mathbf{F}(\mathbf{x}_i) - \mathbf{y}) \right). \quad (2)$$

The Levenberg–Marquardt parameter λ_i is chosen in each iteration to decrease the cost function and should go to zero in the limit. The matrix \mathbf{I}_n is the identity matrix in $\mathbb{R}^{n \times n}$, and $\mathbf{F}'(\mathbf{x}_i) \in \mathbb{R}^{m \times n}$ is the Jacobian matrix of \mathbf{F} evaluated at \mathbf{x}_i .

This mathematical framework is identical to the one presented by Ungermann et al. (2012a) but fits also well to the approach of cross-section retrievals. The basic difference to conventional profile retrievals is just the exchange of the 1-D representations of the atmospheric state with a 2-D state vector and the addition of horizontal correlation or smoothness conditions to the regularisation.

The remainder of this section discusses technical algorithmic and implementation details necessary to perform cross-section retrievals in practise. First, a specific regularisation is presented that is suitable for large-scale retrievals. Second, it is discussed how to practically derive the effects of various error sources on retrieval results for large-scale retrievals. The last subsection presents a novel stochastic approach to diagnostics very suitable for large-scale retrievals as its computational cost does increase only with the square instead of the third power of the number of unknowns.

2.1 Regularisation

For large-scale retrievals, Tikhonov regularisation (e.g. Tikhonov and Arsenin, 1977) is very convenient, if not required, as it allows directly defining the needed precision matrix (i.e. the inverse of the covariance matrix) $\mathbf{S}_a^{-1} \in \mathbb{R}^{n \times n}$. In so far, the regularisation of cross-section retrievals differs from that of 1-D retrievals only by the addition of the last term:

$$\mathbf{S}_a^{-1} = (\alpha_0)^2 \mathbf{L}_0^T \mathbf{L}_0 + (\alpha_1^v)^2 \mathbf{L}_1^{vT} \mathbf{L}_1^v + (\alpha_1^h)^2 \mathbf{L}_1^{hT} \mathbf{L}_1^h. \quad (3)$$

The parameters α_0 , α_1^v , and α_1^h allow weighing the different regularisation terms differently. A good starting point is to set all to 1, even though this paper chooses a smaller α_0 (see Table 2).

Having a sparse precision matrix makes it very inexpensive to set the matrix up and to multiply it with vectors. This allows the efficient use of iterative algorithms such as the conjugate gradient method for solving linear equation systems. However, this poses severe restrictions on which covariance matrices are admissible for the problem and prevents the use of a pure optimal estimation approach that usually employs arbitrary covariance matrices often derived from 3-D models.

As compromise, the Tikhonov regularisation used in this paper is chosen to approximate the precision matrix of an optimal estimation covariance matrix employing the autoregressive model to fill the covariances (see Appendix A). The parametrisation used here follows closely the one described by Ungermann et al. (2012a) with the notable exception of the added term for horizontal regularisation. To summarise the setup briefly, $\mathbf{L}_0 \in \mathbb{R}^{n \times n}$ is a diagonal matrix, with climatological standard deviations on the diagonal. The \mathbf{L}_1^v and \mathbf{L}_1^h matrices pose constraints on the first-order derivative in vertical and horizontal direction, scaled with quantity-specific scaling factors c_q^v and c_q^h for \mathbf{L}_1^v and \mathbf{L}_1^h , respectively (that means that different rows of the matrices are scaled differently).

2.2 Linearised diagnostics

The diagnostics used in this work follow the linearised diagnostics described by Rodgers (2000). The key point of

this section is how these well-known diagnostics may be derived in a memory conserving and numerically stable way required for large-scale retrievals. It thereby expands the previous work that described solely the calculation of the noise error (Ungermann et al., 2010b) to the more complicated estimation of systematic errors induced by background gases and uncertainties in spectral line data.

The underlying idea is to insert the \mathbf{x}_f minimising the cost function as fix point into Eq. (2) and to replace the vector of measurements \mathbf{y} by a first-order Taylor approximation around the true atmospheric state $\mathbf{x}_t \in \mathbb{R}^n$. This gives a linear relation between the retrieval result \mathbf{x}_f , the (unknown) true atmospheric state \mathbf{x}_t and the measurement errors $\boldsymbol{\epsilon} \in \mathbb{R}^m$:

$$\mathbf{x}_f = \mathbf{A} \mathbf{x}_t + (\mathbf{I} - \mathbf{A}) \mathbf{x}_a + \mathbf{G} \boldsymbol{\epsilon}. \quad (4)$$

Thereby, the matrix $\mathbf{G} \in \mathbb{R}^{n \times m}$ is the so-called gain or contribution matrix mapping differences in the measurements onto differences in the retrieval result. It is defined as

$$\mathbf{G} = \left(\mathbf{S}_a^{-1} + \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} \mathbf{F}'(\mathbf{x}_f) \right)^{-1} \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1}. \quad (5)$$

The matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with

$$\mathbf{A} = \mathbf{G} \mathbf{F}'(\mathbf{x}_f) \quad (6)$$

is the averaging kernel matrix that describes the influence of the true atmospheric state on the retrieval result.

One can analyse the averaging kernel matrix to determine how the applied regularisation influences the retrieval result and to quantify the spatial resolution of the result. The gain matrix is also instrumental in estimating the effect of various errors on the retrieval result. With $\mathbf{S} \in \mathbb{R}^{m \times m}$ being the covariance matrix of an error in the measurements, the covariance matrix $\mathbf{G} \mathbf{S} \mathbf{G}^T$ describes the resulting uncertainty in the retrieval result. Such an error covariance matrix \mathbf{S} can be generated (if only in an ad-hoc manner) for most errors affecting the retrieval, e.g. for the uncertainty of non-retrieved background gases or uncertainties with respect to spectral line strengths (see below). Obviously, many error sources are not of Gaussian nature and can thus only be approximated to a certain degree by this approach.

Providing the diagnostic information for a cross-section retrieval is slightly more complicated than for 1-D retrievals, as the involved matrices are of a size that is neither easy to keep in memory without resorting to sparse storage techniques, nor straightforwardly invertible. The inversion is thereby both computationally demanding and, even worse, numerically unstable (scaling the diagonal of the matrix to one and using an LU-decomposition as done by Livesey et al. (2006) works mostly well for the presented type of problem up to a certain size of $\approx 40\,000$ unknowns).

A different approach that avoids all matrix inversions or factorisations drastically reduces the amount of required memory at the cost of additional computation. The basic idea is to calculate only a single row of \mathbf{A} and \mathbf{G} at a time by

solving a linear equation system and to derive all relevant diagnostic quantities for the associated atmospheric quantity from these matrix rows. Doing so for all points in turn is trivially to parallelise on supercomputers.

To calculate the i -th row \mathbf{g}_i of \mathbf{G} , the i -th row \mathbf{s}_i of $(\mathbf{S}_a^{-1} + \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} \mathbf{F}'(\mathbf{x}_f))^{-1}$ is calculated by solving the equation system $(\mathbf{S}_a^{-1} + \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} \mathbf{F}'(\mathbf{x}_f)) \mathbf{s}_i = \mathbf{e}_i$ with \mathbf{e}_i being the i -th unity vector. Naturally, iterative methods are used for this to avoid the calculation of any matrix-matrix product or the generation of dense matrices. The resulting vector \mathbf{s}_i is then multiplied with the appropriate matrices to generate the i -th row \mathbf{g}_i of \mathbf{G} and the i -th row \mathbf{a}_i of \mathbf{A} .

Using this, one can directly calculate the variance of influence of measurement noise on the retrieval result by evaluating $\mathbf{g}_i^T \mathbf{S}_\epsilon \mathbf{g}_i$.

For deriving the influence of unknown background gases, one first needs to assemble a covariance matrix \mathbf{S}_{gas} describing the state of this gas. This can then be multiplied from both sides with the Jacobian matrix $\mathbf{F}'_{\text{gas}}(\mathbf{x}_f)$ of the forward model with respect to this gas to derive a covariance matrix for the influence of this gas on the (simulated) measurements. Assembling a full covariance matrix \mathbf{S}_{gas} by, e.g. using the auto-regressive approach (i.e. $(\mathbf{S}_{\text{gas}})_{i,j} = \sigma_i \sigma_j e^{(v_i - v_j)(c_{\text{gas}}^v)^{-1}} e^{(h_i - h_j)(c_{\text{gas}}^h)^{-1}}$ with v_i and h_i being vertical and horizontal coordinates of data point i) and some horizontal and vertical correlation lengths is too costly both with respect to memory consumption and computation time. In this case, it is much more efficient to analytically define the (sparse) precision matrix $\mathbf{S}_{\text{gas}}^{-1}$ instead as described in Appendix A. Using the precision matrix instead of the covariance matrix requires the solution of another equation system to generate the desired result. The standard deviation σ_i^{gas} of the influence of the gas on the retrieved quantity i can then be deduced as

$$(\sigma_i^{\text{gas}})^2 = \mathbf{g}_i^T \mathbf{F}'_{\text{gas}}(\mathbf{x}_f)^T \mathbf{z}, \text{ with } \mathbf{S}_{\text{gas}}^{-1} \mathbf{z} = \mathbf{F}'_{\text{gas}}(\mathbf{x}_f) \mathbf{g}_i. \quad (7)$$

Solving this linear equation system can be done very efficiently by employing iterative methods as only a couple of diagonals of the precision matrix $\mathbf{S}_{\text{gas}}^{-1}$ are non-zero. The effort is therefore typically negligible compared to the effort to calculate the vector \mathbf{g}_i itself.

The influence of uncertainties of spectral line data of one specific trace gas can be estimated in a similar fashion. Under the assumption that the error is fully correlated, the corresponding covariance matrix $\mathbf{S}_{\text{spec,gas}}$ is fully determined by its standard deviations $s_{\text{spec,gas}}$ ($\mathbf{S}_{\text{spec,gas}} = s_{\text{spec,gas}} \mathbf{s}_{\text{spec,gas}}^T$), which are set to the relative spectral error times the local gas volume mixing ratios (this assumes that the relative spectral error is identical over all used integrated microwindows (IMWs); should this not be the case, it is straightforward to calculate these individually for each IMW, but more complicated in notation and therefore omitted here). Again, it is too costly to assemble the matrix $\mathbf{S}_{\text{spec,gas}}$, especially as

the equation is easily simplified. Thus, the standard deviation $\sigma_i^{\text{spec,gas}}$ induced by uncertainties in spectral characterisation of a gas can be calculated as

$$\begin{aligned} (\sigma_i^{\text{spec,gas}})^2 &= \mathbf{g}_i^T \mathbf{F}'_{\text{gas}}(\mathbf{x}_f)^T (s_{\text{spec,gas}} \mathbf{s}_{\text{spec,gas}}^T) \mathbf{F}'_{\text{gas}}(\mathbf{x}_f) \mathbf{g}_i \\ &= (s_{\text{spec,gas}}^T \mathbf{F}'_{\text{gas}}(\mathbf{x}_f) \mathbf{g}_i)^2. \end{aligned} \quad (8)$$

In this fashion, all relevant diagnostic quantities can be calculated using only sparse matrices and iterative algorithms. The resulting quantities are numerically identical to those derived by straightforward matrix inversion, if feasible. The sole disadvantage of this approach is that the effort to calculate the diagnostics for all retrieved quantities quickly surpasses the effort for the retrieval itself. A more efficient way to deduce error bars and to estimate the influence of individual error terms on the full set of retrieved quantities is presented in the next subsection.

2.3 Monte Carlo diagnostics

Deriving the diagnostics for each retrieved quantity requires on the order of $O(n^3)$ operations, as for each of the n retrieved elements of the atmospheric state vector, a linear n -by- n equation system must be solved. This already assumes that a fixed number of iterations much smaller than n produces a sufficiently accurate solution in the iterative solver.

One way of reducing the computational cost is to use stochastic algorithms. Such algorithms can also be used to produce a solution with associated (non-Gaussian) distribution, or to estimate the errors in a non-linear way (e.g. Tarantola, 2004); here, only a simple linear error estimate is desired.

The previous section introduced a number of error effects and associated covariance matrices. Using these Gaussian characterisations of systematic and stochastic errors influencing the measurements, it is easy to generate random errors according to the given distributions by multiplying a root of the covariance matrix with a vector of uncorrelated Gaussian random variables with standard deviations of one. Appendix B gives an efficient method to construct the inverse of such roots. These error terms can then be added to the actual measurements and be fed into a linear retrieval to produce perturbed solutions. The differences between the actual retrieval result and the outcomes of the linear retrievals with disturbed inputs can then be analysed to derive the error bars. As the involved distributions are assumed to be Gaussian and as the retrieval is linearised for this purpose, the resulting distributions are also Gaussian, simplifying the mathematics for computing all estimates significantly. Depending on the required confidence level for the error bars, a couple of solved n -by- n linear equation systems deliver already usable estimates.

Let $\boldsymbol{\epsilon}_i^{\text{mc}} \in \mathbb{R}^m$, $1 \leq i \leq N$, be N error vectors generated according to a multivariate Gaussian error distribution, e.g. of the noise error. Usually, one is interested especially in

the total error (the sum over all individual errors), for which an error vector according to the joint distribution of all errors may be inexpensively generated. Given this set of error vectors, a linearised retrieval generates a set of associated solutions \mathbf{x}_i^{mc} :

$$\mathbf{x}_i^{\text{mc}} = \mathbf{x}_f - \left(\mathbf{S}_a^{-1} + \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} \mathbf{F}'(\mathbf{x}_f) \right)^{-1} \cdot \left(\mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} \boldsymbol{\epsilon}_i^{\text{mc}} \right). \quad (9)$$

The retrieval is linearised to generate comparable results to the conventional error calculation that is also based on linear approximation.

The standard deviation of these Monte Carlo solutions delivers the desired error estimate. According to Kenney and Keeping (1951), the standard deviation can be estimated as

$$\sigma^{\text{mc}} = \frac{1}{c_4(N)} \sqrt{\frac{1}{N} \sum_{i=1}^N \left(\mathbf{x}_i^{\text{mc}} - \frac{1}{N} \sum_{j=1}^N \mathbf{x}_j^{\text{mc}} \right)^2}. \quad (10)$$

Thereby, the correction factor

$$c_4(N) = \sqrt{\frac{2}{N-1}} \frac{\Gamma\left(\frac{N}{2}\right)}{\Gamma\left(\frac{N-1}{2}\right)} \quad (11)$$

makes this estimator unbiased also for small N . The typically used correction factor of dividing by $N-1$ instead of N within the square root of Eq. (10) removes only the bias for the estimate for the variance, which causes noticeable discrepancies between deterministically calculated and stochastically estimated standard deviations of retrieval results for very small values of N . The error of the estimate decreases with the number of samples. Kenney and Keeping (1951) also provide an estimate for the standard deviation of the standard deviation estimate as

$$\text{std}(\sigma^{\text{mc}}) = \sigma^{\text{mc}} \sqrt{\frac{1}{N} \left(N - 1 - \left(\frac{2\Gamma\left(\frac{N}{2}\right)}{\Gamma\left(\frac{N-1}{2}\right)} \right)^2 \right)}. \quad (12)$$

Please note that the paper distinguishes here and in the following between the estimate of the standard deviation of the retrieval result (giving the error of the retrieval result) and the standard deviation of the estimate itself (giving the error of the error estimate).

Equation (12) shows that exemplary 32 random samples give an estimate of the standard deviation with a $1-\sigma$ relative error of about 12.4%. Thus one will almost certainly encounter outliers if many points are diagnosed. A quadrupling of the number of samples roughly halves the relative error. In our implementation of the algorithm and for the example presented later on, it is possible to calculate roughly 128 linear solutions in the same time that is required for the

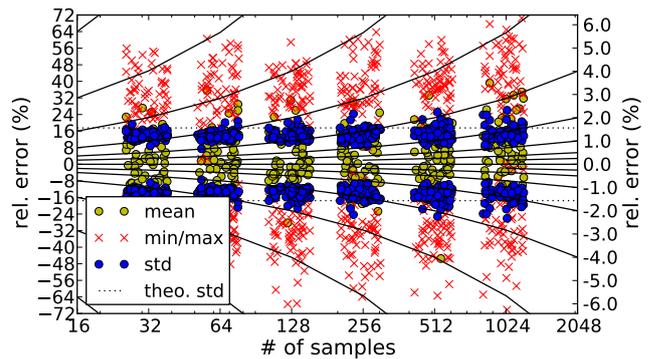


Fig. 1. Accuracy for estimation of errors for varying samples sizes. Shown are the means of the relative error of the standard deviation estimate (yellow circles), the real standard deviation of the estimates generated by comparison of the estimates with deterministically calculated errors (blue circles) and minimal and maximal relative errors (red crosses). The dotted horizontal line is the theoretical standard deviation of the standard deviation estimate. Please note the varying vertical scale with solid lines indicating identical relative errors. The samples sizes are always powers of two.

non-linear retrieval itself, delivering a precision of about 6%. Please note that this is the precision of the estimated standard deviation. For example, if the estimated standard deviation points to a 10% uncertainty for a certain retrieved trace gas (not uncommon for the CRISTA-NF instrument), a relative error of 6% translates into a $1-\sigma$ uncertainty of the retrieval result ranging from 9.4 to 10.6%. This uncertainty in the error is well within the range of other uncertainties introduced by instrument characterisation or spectral line data.

If one is only interested in a worst-case estimate, for example to determine the relevance of different error sources, it is advantageous to use a very small number of samples to generate a coarse estimate of the standard deviation and add to this one (or multiple) standard deviations of the estimate to derive a reliable upper bound for the error.

To give an example of a typical spread, Fig. 1 shows the converging behaviour for a real diagnosis of the total error for a range of sampling sizes. The quality of the estimates follows the expected theoretical forecast, which makes a separate calculation superfluous in practise. As expected for a large number of estimates and a Gaussian distribution, the minimum and maximum errors are rather large compared to a single standard deviation. For a small number of just 32 samples, relative errors of the standard deviation estimate of up to 50% are not uncommon.

This method lends itself well to parallelisation. Each Monte Carlo solution may be generated in parallel to the others on an arbitrary number of cores. It is also possible to refine the estimate later on. One may combine the sample mean and sample variance of different sample runs to improve upon the estimate at a later time, when more processing

power may be available or the need for a more accurate estimate has arisen.

As the number of samples for a desired precision is independent of the number of unknowns, this algorithm has an effective cost of only $O(n^2)$, albeit with a large constant coefficient. Due to the large amount of involved unknowns, this algorithm is much more efficient than deterministic methods that require $O(n^3)$ operations. Thus, this method provides an efficient way to provide diagnostic information for large-scale retrievals.

3 Case study

This section presents a case study that examines the differences between cross-section retrievals and conventional 1-D retrievals.

The RECONCILE campaign took place from January to March 2010. The high-flying Russian research aircraft M55-Geophysica performed 12 flights in this time frame with a variety of in situ and remote sensing instruments. One of the instruments aboard was CRISTA-NF, an airborne infrared limb sounder. The data taken by CRISTA-NF in this campaign are rather unique in having at the same time a high frequency of taken profiles (one profile every ≈ 15 km) and a high vertical sampling (≈ 250 m). Its measurements during one flight of the campaign shall function as an example. The one-dimensional retrieval and validation of this dataset was described by Ungermann et al. (2012a). The impact of stochastic errors on the measurements is rather high, first from actual measurement noise introduced by the detectors and associated components, second from uncertainties in the vertical pointing of individual spectra (and even spectral samples within one spectrum). The second flight of 2 March 2010 has been selected as example as it offers the most vertical and horizontal structure.

3.1 Instrument

CRISTA-NF is an air-borne infrared passive limb sounder. It measures emissions of thermally excited atmospheric trace gases in the mid-infrared region (4 to 15 μm). The CRISTA-NF viewing direction is perpendicular to the flight direction of the M55-Geophysica. Spectra are scanned from the flight altitude down to ≈ 5 km in vertical steps of ≈ 250 m using a Herschel telescope with a tiltable mirror. A detailed description of the optical system and the cryostat of CRISTA-NF is given by Kullmann et al. (2004). The optical system of CRISTA-NF has been re-purposed from the CRISTA experiment that was launched by the Space Shuttle on the Shuttle Pallet Satellite SPAS in November 1994 (STS 66) and August 1997 (STS 85) (Offermann et al., 1999; Grossmann et al., 2002). Obtaining a single spectrum takes only 1.2 s due to the low temperature of the detectors and the whole optical system of the instrument.

Table 1. A list of employed integrated microwindows and their spectral range.

IMW	range (cm^{-1})	IMW	range (cm^{-1})
0	777.5–778.5	7	808.0–809.0
1	784.0–785.0	8	810.0–813.0
2	787.0–790.0	9	820.5–821.5
3	791.5–793.0	10	831.0–832.0
4	794.1–795.0	11	846.0–847.0
5	795.5–796.5	12	863.0–866.0
6	796.6–797.5		

While the pointing stability of the instrument had been improved upon compared to previous campaigns, its pointing stability is still only on the order of 0.02° , which corresponds to an uncertainty in the tangent point altitude of ≈ 125 m vertically 10 km below flight level. This corresponds roughly to a relative error in measured radiation of $\approx 1\%$.

3.2 Retrieval processor

The Jülich Rapid Spectral Simulation Code Version 2 (JURASSIC2) is used as forward model and retrieval processor. The Python/C++ based JURASSIC2 and its predecessor were used in several experiments and studies (e.g. Hoffmann et al., 2008). It uses pre-calculated tables of spectrally averaged optical paths to avoid costly line-by-line calculations. Combining the radiances derived from the emissivity growth approximation method (EGA; e.g. Weinreb and Neuendorfer, 1973; Gordley and Russell, 1981) and the Curtis–Godson approximation (CGA; Curtis, 1952; Godson, 1953) using a simple regression scheme (Francis et al., 2006), the standard deviation between full line-by-line calculations and the approximate scheme used by JURASSIC2 is less than 0.22% (Ungermann et al., 2012a). This is less than the assumed stochastic error of 1% and much less than assumed uncertainties in spectral line data and parametrisation schemes.

3.3 Retrieval setup

The discussed retrieval derives the primary retrieval targets CCl_4 , CFC-11, H_2O , HNO_3 , O_3 , and ClONO_2 and the secondary retrieval targets aerosol, temperature, PAN, CFC-113, and HCFC-22 from 13 integrated microwindows (see Table 1). The primary distinction between primary and secondary retrieval targets is that the signal of the primary targets is strong enough to deliver reliable results also in the presence of systematic errors in background gases. The retrieval grid sampling distance is 250 m below 20 km, 1 km between 20 and 30 km, and 2 km above. All targets are derived between 0 and 25 km except for the trace gases HNO_3 , O_3 , and ClONO_2 that are subject to an increase in volume mixing ratio above flight altitude. These gases are therefore derived up to an altitude of 65 km. In total, this gives a state

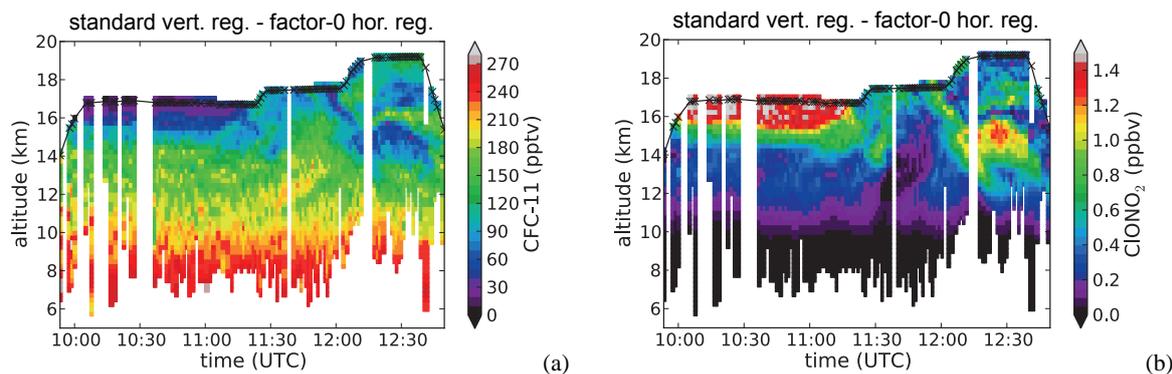


Fig. 2. Two cross-sections of CFC-11 (a) and ClONO₂ (b) volume mixing ratios for the baseline regularisation (standard vertical regularisation and no horizontal regularisation). The straight black line indicates the altitude of the instrument. Crosses on this line indicate the mean position of the instrument during a profile measurement.

Table 2. Parameters employed for the regularisation: Tikhonov strength and correlation lengths.

Parameter	Value	Parameter	Value
α_0	$\sqrt{0.1}$	c_{ClONO_2}	4 km
α_1^v, α_1^h	1	$c_{\text{H}_2\text{O}}$	12 km
$c_{\text{temperature}}$	2 km	$c_{\text{HCFC-22}}$	32 km
c_{aerosol}	640 km	c_{HNO_3}	4 km
c_{CCl_4}	1 km	c_{O_3}	6 km
$c_{\text{CFC-11}}$	0.3 km	c_{PAN}	8 km
$c_{\text{CFC-113}}$	2 km		

vector \mathbf{x} with 93 870 entries and a measurement vector \mathbf{y} with 73 660 entries. The regularisation is parametrised according to Table 2. A priori information and initial guess are chosen to be identical. For temperature, pressure and water vapour, ECMWF ERA-Interim data are used. A zero profile is used for PAN. For the remaining targets appropriate profiles are taken from the climatology assembled by Remedios et al. (2007). All standard deviations used for the regularisation are also taken from this source except for temperature (for which a fixed standard deviation of 1 K was chosen) and PAN (for which the climatology of Glatthor et al., 2007 is used). Additional information about the 1-D retrieval and a validation of the results against measurements by other instruments is given by Ungermann et al. (2012a).

3.4 Effect of horizontal regularisation

This section quantifies the effect of horizontal regularisation on a cross-section retrieval. Besides the obvious horizontal smoothing of volume mixing ratios, it needs to be examined if one can in turn reduce vertical regularisation to improve the vertical resolution at the cost of horizontal resolution and how much the image noise can be reduced without degrading the horizontal resolution of the cross-section too much. Image noise refers here to the artefacts in retrieved volume

mixing ratios that are induced by measurement noise and similar stochastic error sources in the radiances.

The presented case study derives multiple gases from the measurements, but, to conserve space, only the volume mixing ratios of CFC-11 and ClONO₂ will be discussed, as the other trace gases behave quite similarly. The choice of CFC-11 is disadvantageous for cross-section retrievals, as its signal is strong and the image noise is low to begin with and therefore the possible gain is small compared to other retrieval gases. In contrast, ClONO₂ has a weaker signal and should thus profit more. Thus, the two trace gases represent a worst-case and a best-case scenario for the proposed algorithm and thereby allow a better quantification of expected benefits for other scenarios.

The baseline retrieval is a cross-section retrieval that is identical to a conventional series of 1-D retrievals as no horizontal regularisation has been added. The CFC-11 and ClONO₂ volume mixing ratios produced by the baseline retrieval are depicted in Fig. 2. These pseudo-colour plots show each retrieved trace gas volume mixing ratio value as a rectangle. This cross-section exhibits a lot of detail and fine structure. Please note that most structures are horizontal or diagonal and that neighbouring profiles are often very similar. This suggests that consecutive profiles taken at this horizontal sampling are not fully independent from one another. It is likely that many of the inconsistencies between neighbouring profiles stem from random errors in the given elevation angle knowledge that affects lower altitudes more than those close to the instrument.

The CFC-11 cross-section contains image noise in several regions, which is caused by the given measurement noise combined with a rather weak vertical regularisation; but this parametrisation allows a good reproduction of the filament of increased CFC-11 volume mixing ratio at 15.5 km on the right, which is situated between two air masses of smaller volume mixing ratios. The vertical extent seems to be ≈ 500 m, which corresponds to two samples of the employed grid. On the left-hand side of the figure, two very faint

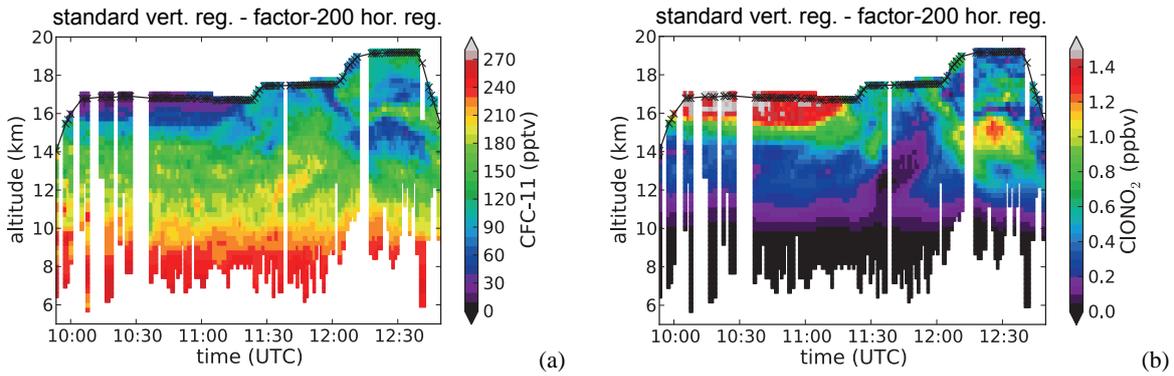


Fig. 3. Two cross-sections of CFC-11 (a) and ClONO₂ (b) volume mixing ratios with standard vertical regularisation and factor-200 horizontal regularisation.

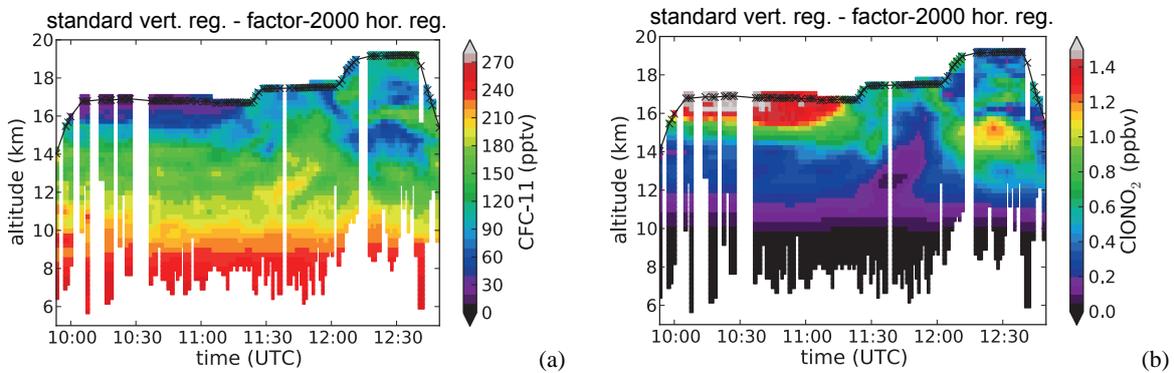


Fig. 4. Two cross-sections of CFC-11 (a) and ClONO₂ (b) volume mixing ratios with standard vertical regularisation and factor-2000 horizontal regularisation.

horizontal filaments might be present at 13 km and 14 km. Also, a large artefact is present at 9.5 to 11 km at 10:35 UTC where no measurements were taken due to aircraft movements; this problem is revealed by the decreased information content in this area (not depicted).

The ClONO₂ cross-section is generally similar in quality to the CFC-11 cross-section but less resolved. Only at lower altitudes, the ClONO₂ volume mixing ratios contain less image noise, mostly because the measured signal becomes rather weak and more a priori knowledge enters the retrieval result, which is reflected in a sharp drop off in measurement contribution (not depicted) below 12 km.

A natural starting point for the horizontal regularisation strength would be an approximate scale difference between vertical and horizontal length for large meso- or synoptic-scale structures in the atmosphere, i.e. ≈ 200 . This strength is referred to as factor-200 horizontal regularisation within this paper. For the vertical regularisation strength, initially the values of Table 2 are used. The result for using thus $c_q^h = 200 c_q^v = 200 c_q$ for all quantities except aerosol gives Fig. 3 (it was found to be beneficial to always keep $c_{\text{aerosol}}^h = 100 c_{\text{aerosol}}^v$ as the vertical regularisation strength for aerosol was very strong to begin with). Starting with the CFC-11

cross-section, the differences are subtle, but it is visually already more pleasing than the baseline setup (see Sect. 3.5 for a quantitative discussion). With this regularisation strength, the largest effect is on individual pixels, where closely neighbouring profiles have differing values, i.e. mostly outliers presumably caused by spikes in the measurements and small-scale oscillations. Some of these outliers in the thin vertical filaments have been reduced; especially the region between 10 and 12 km has become more uniform. The artefact at 10:35 UTC has become smaller as the missing information starts being filled in from the right side (the left profile is spatially even farther away than suggested by the temporal proximity due to a turn of the aircraft and therefore exerts only a minimal influence). Also the ClONO₂ cross-section looks improved. No major features were smoothed over. Several features even look better. For example, the inclined outflow of increased ClONO₂ at 11:50 UTC at 12 km is now consistent over all neighbouring profiles. This seems correct as indicated by retrieved HNO₃ volume mixing ratios, which have a stronger signal and show the same behaviour without any horizontal regularisation (not depicted).

A stronger horizontal regularisation is depicted in Fig. 4, where the horizontal strength has been set to be 2000 times

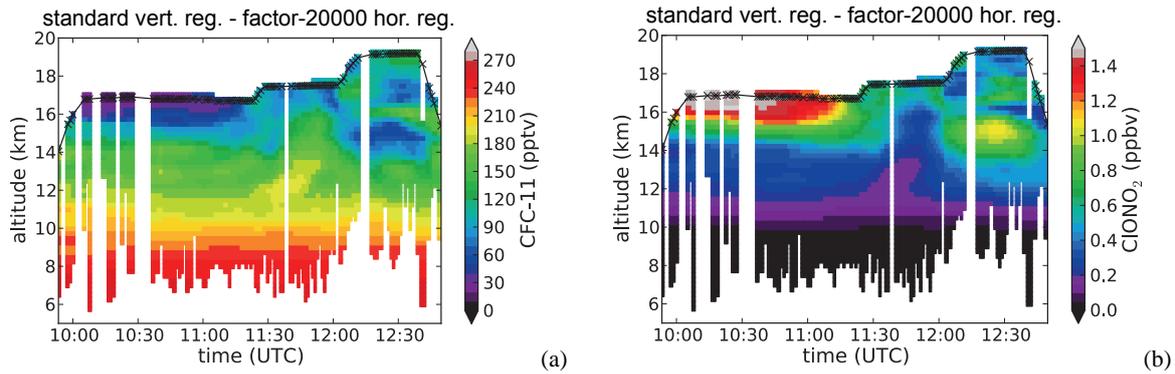


Fig. 5. Two cross-sections of CFC-11 (a) and ClONO₂ (b) volume mixing ratios with standard vertical regularisation and factor-20 000 horizontal regularisation.

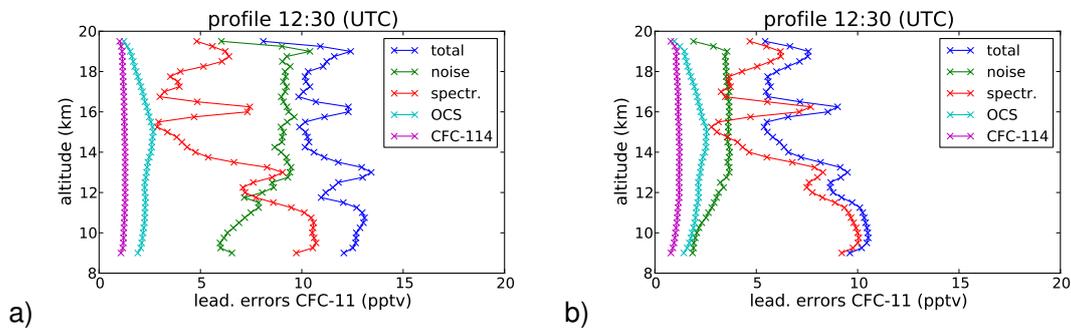


Fig. 6. Leading errors in CFC-11 volume mixing ratio for a representative profile (12:30 UTC) and the baseline regularisation in panel (a) and a factor-2000 horizontal regularisation in panel (b). “spectr.” notes the error induced by uncertainties in CFC-11 spectral data. “OCS” and “CFC-114” are the systematic errors induced by the uncertainty in volume mixing ratios of these background gases.

the vertical one (factor-2000). This result is devoid of obvious defects as they have been smoothed over. But strong horizontal gradients are still reproduced, e.g. the small intrusion of air with low CFC-11 volume mixing ratios at 11:30 UTC and 14 to 16 km. But the low volume mixing ratios within this filament are slightly increased due to smoothing. Also the diagonal structure at 16 km and 12:00 UTC shows larger CFC-11 volume mixing ratios than in the setup without horizontal regularisation. However, while certain details have been lost, it certainly reproduces the major features well. The weaker ClONO₂ volume mixing ratios behave similarly. Most features are still present, but increased volume mixing ratios of small filaments have been reduced and vice versa.

For comparison, a setup with an obviously too strong horizontal regularisation is shown. Using a factor-20 000 horizontal regularisation gives results as shown in Fig. 5. Previously clear-cut horizontal and diagonal structures have been distorted by this setup. Even though the vertical regularisation is unmodified, the vertical resolution has also deteriorated, which results in obvious vertical smoothing. Depending on the retrieved quantity and the measurement noise level, such a strength might still be useful, as structures with strong horizontal gradients are preserved even here.

3.5 Quantification of noise error reduction

After presenting the visual differences of the various regularisation strengths, this section focuses on the impact of the horizontal regularisation on the more objective diagnostic noise error estimates.

As the errors are quite similar from one profile to another, it is sufficient to discuss the errors of a single profile. Figure 6 shows the errors associated with the profile measured at 12:30 UTC for the baseline retrieval setup (panel a) and factor-2000 horizontal regularisation (panel b). The total error is produced by taking the Euclidean norm of a vector consisting of all individual error terms. Only the most significant error terms are depicted in descending order. In Fig. 6a, the error according to stochastic noise (which was assumed to be 1 %) is the leading error down to 12 km. Only from there on, the error induced by spectral uncertainties in CFC-11 becomes dominant (this error was set to 5 % in an ad hoc manner based on HITRAN (Rothman et al., 2009) error quantification). The corresponding plot for factor-2000 regularisation is given in Fig. 6b. It is shown that the noise error was reduced by a factor of ≈ 3 , leaving the total error being dominated by systematic errors implying an improved

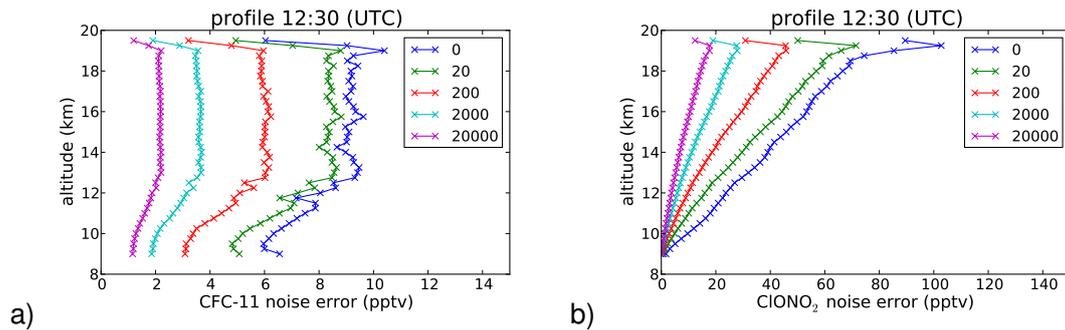


Fig. 7. Error induced by measurement noise for CFC-11 (a) and ClONO₂ (b), respectively, for one representative profile (12:30 UTC) and varying horizontal regularisation strengths. The given number indicates the horizontal regularisation multiplier compared to baseline vertical regularisation strength.

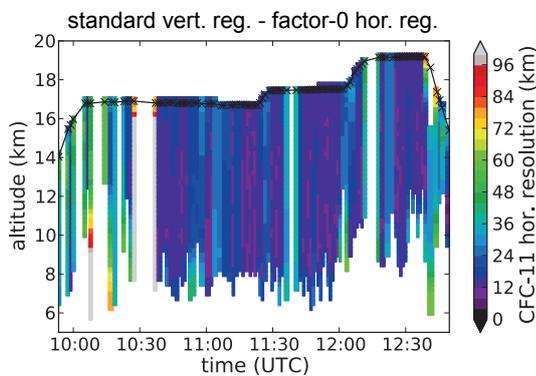


Fig. 8. A cross-section of CFC-11 horizontal resolution (FWHM) for the baseline setup.

precision with similar accuracy of the result. The other errors are largely unaffected.

Figure 7 shows the reduction of noise error for CFC-11 (panel a) and ClONO₂ (panel b) for several horizontal regularisation strength. The effect of measurement noise on CFC-11 volume mixing ratios is depicted in Fig. 7a. Introducing horizontal regularisation in a natural manner reduced the noise error to 70 % of the original level. Using a factor-2000 horizontal regularisation reduces the noise error further down to 30 % of the original level. Increasing the horizontal regularisation strength further gives only diminishing returns, especially as the impact of measurement noise is reduced to the same order as most systematic or quasi-random error sources such as influence of unknown background gases, so that the gain in precision is insignificant compared to the given accuracy. For ClONO₂, the factor-200 horizontal regularisation already halves the effect of measurement noise. As its signal is weaker, it profits more from additional regularisation.

3.6 Achievable horizontal resolution

Adding horizontal regularisation obviously worsens the achievable horizontal resolution. However, it may be that the resulting resolution is still sufficient for the scientific purposes or that the given reduction of image noise is seen as more advantageous.

Figure 8 shows the horizontal resolution as derived from the full width at half maximum (FWHM) of the averaging kernel in horizontal direction. The width is determined by linear interpolation between data points, which causes the step-like structure due to the irregular horizontal grid size. Vertical variations are partly caused by differing azimuthal directions of the detector, which causes the lower tangent points of affected neighbouring profiles to be farther apart than upper ones (the profile taken at $\approx 10:10$ UTC is the most prominent example of this effect). The large value around 10:30 UTC is caused not only by the sparse measurement density, but also by a turn of the aircraft implying that the horizontal distance of the two seemingly neighbouring profiles is much larger than suggested by the temporal proximity. A regular horizontal grid would remove several of these artefacts, but on the other hand also introduce many points with purely interpolated information.

The impact of the naturally dimensioned factor-200 horizontal regularisation is very small as shown in Fig. 9. The horizontal resolution for CFC-11 is nearly unchanged compared to applying no horizontal regularisation: the resolution is in both cases typically the maximum of ≈ 16 km and the resolution of the baseline setup. This is only slightly larger than the typical minimum distance between profiles of ≈ 15 km. Several individual points have a noticeably worsened horizontal resolution; these are typically points of profiles where there is a gap in the vertically measured spectra. For these cases, the retrieval prefers to interpolate the mixing ratios from neighbouring profiles rather than doing so from measurements farther above or below, especially if the neighbouring profile is close by. A second kind of artefact arises in the profile at 10:35 UTC, where the horizontal resolution

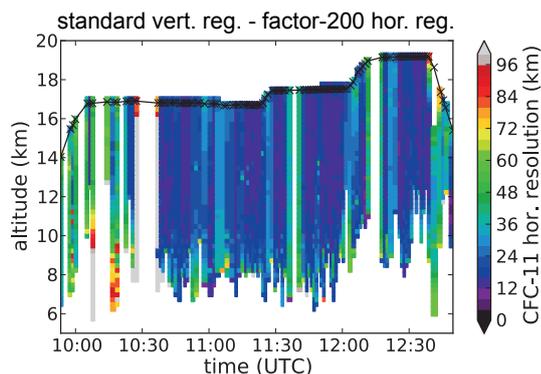


Fig. 9. A cross-section of CFC-11 horizontal resolution (FWHM) for the factor-200 horizontal regularisation strength.

seemingly was much improved at 9 km altitude. Here, an especially large gap in taken spectra is present and the retrieval extrapolates from the profile to the right, which was taken in close spatial proximity. The averaging kernel matrix indicates that the information is spatially well localised, but centred on the succeeding profile. This dislocation can also be visualised by plotting the distance between the maximum of the averaging kernel matrix row and the location of the retrieved point; however for the given case study, such a problem arises only for the 10:35 UTC profile.

Figure 10 shows the results of factor-2000 horizontal regularisation strength. Here, the horizontal resolution is obviously worse than for the baseline case, especially in the lower altitudes. But where a sufficient measurement density is given, e.g. around 11:45 UTC, the horizontal resolution is still ≈ 25 km, which is for many long-living trace gases quite sufficient.

3.7 Trade-off of vertical against horizontal resolution

It was established that adding horizontal regularisation reduces the impact of measurement noise at the expense of worse horizontal resolution. One might therefore use this achieved stability to reduce the strength of vertical regularisation to, possibly, achieve less vertical smoothing.

Reducing the vertical regularisation strength by a factor of ten and using a factor-200 horizontal regularisation (related to the original vertical regularisation strength) gives similar noise error figures as the baseline retrieval. Therefore, they should give qualitatively comparable results, but with different trade-offs. In Fig. 11, the corresponding CFC-11 and ClONO₂ volume mixing ratios are depicted. The figure seems to contain a similar amount of image noise as Fig. 2 but with a different distribution of errors. As neighbouring profiles are slightly more consistent, it is slightly more pleasing to the eye, but objectively with respect to error bars (not depicted) of comparable quality. Especially the lower altitudes around 10 km are quite improved. Some very small structures like the extension of air with low CFC-11

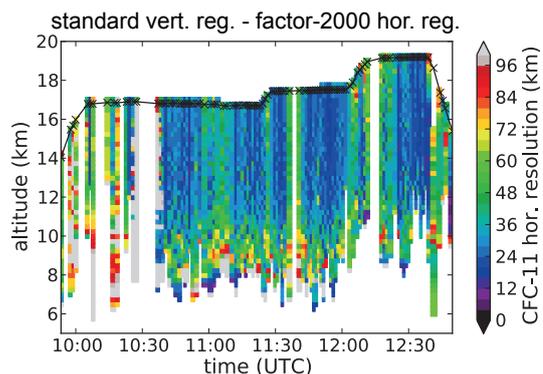


Fig. 10. A cross-section of CFC-11 horizontal resolution (FWHM) for the factor-2000 horizontal regularisation strength.

volume mixing ratios at 11:15 UTC and 15.5 km altitude seem slightly more consistent.

The horizontal resolution is similar to the setups discussed in Sect. 3.6, which is slightly reduced compared to the baseline case without horizontal regularisation, i.e. on the order of 15 km (i.e. equal to the profile sampling distance). The difference in vertical resolution is shown in Fig. 12. The vertical resolution of CFC-11 is already quite close to the achievable optimum constrained by the observation geometry, so its vertical resolution cannot improve significantly. But the other major retrieval targets benefit noticeably by up to 30 % at lower altitudes. Also water vapour benefits greatly by more than halving the vertical resolution around 16 km.

CFC-11 might be a bad example here, as the vertical regularisation was small to begin with and is effectively turned off altogether by the reduction by factor 10. The largest improvement in resolution of a primary target is therefore achieved in ClONO₂. Figure 11b reproduces very well the horizontal filaments at 12 and 14 km altitude, which are difficult to discern in Fig. 2b due to the effects of measurement noise and uncertainty in elevation angle. These filaments coincide with well-defined filaments of elevated HNO₃ volume mixing ratio (not depicted). Also the low mixing ratios in the filament at 15.5 km in the later part of the flight around 12:30 UTC are more pronounced. At only one place, there may be a negative effect of the horizontal smoothing. At 11:35 UTC in Fig. 2a, there seems to be a vertically continuous mass of low ClONO₂ volume mixing ratios from 8 to 14 km. In Fig. 11b, this faint structure has become less pronounced. It is however unclear whether this is an artefact of the regularisation or the proper structure, as only some correlated retrieved trace gases exhibit a similar structure in the baseline retrieval.

Concluding, this technique allows adjusting the characteristics of the retrieved volume mixing ratios to the scientific question and the atmospheric situation. Especially for examining the upper troposphere and lower stratosphere and the

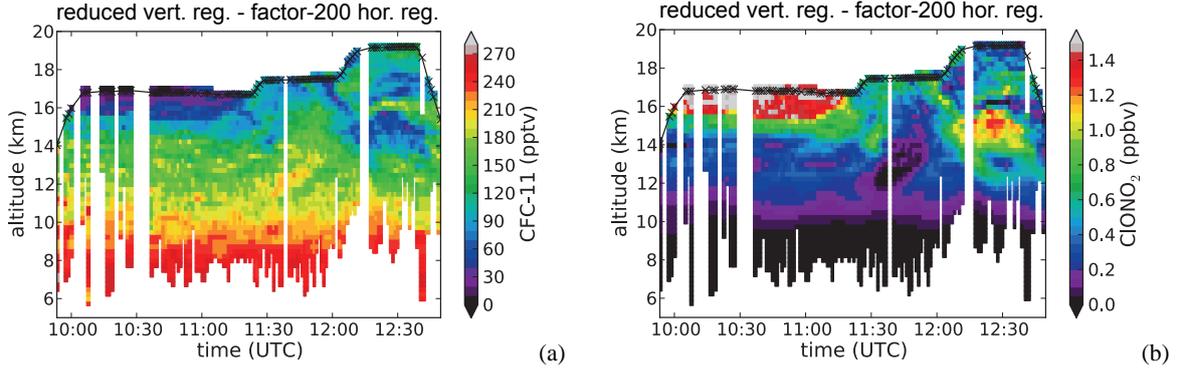


Fig. 11. Two cross-sections of CFC-11 (a) and ClONO₂ (b) volume mixing ratios for a reduced vertical regularisation strength and factor-200 horizontal regularisation.

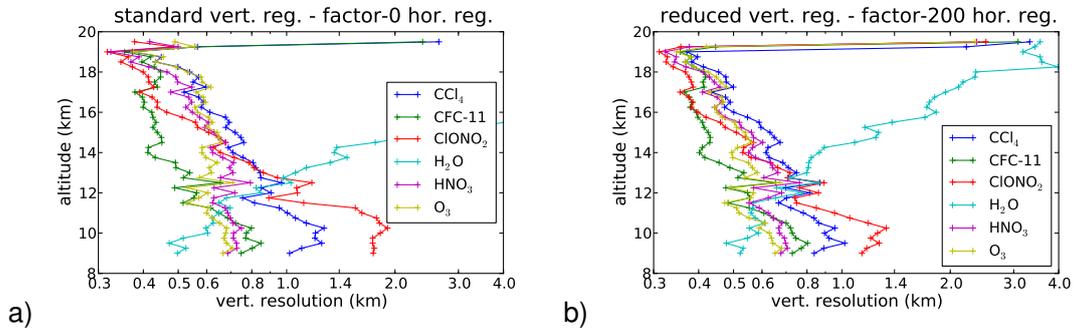


Fig. 12. Vertical resolution of selected retrieval targets for a representative profile (12:30 UTC) for the baseline setup in panel (a) and a reduced vertical regularisation strength and factor-200 horizontal regularisation in panel (b)

tropopause inversion layer, which is of very narrow, this technique might reproduce otherwise invisible details.

3.8 Comparison to smoothing of 1-D retrievals

The basic idea behind the cross-section retrieval is to exploit the similarity of neighbouring profiles. This can be exploited also by simple weighted moving averages. This section presents two such less costly alternatives to a full cross-section retrieval, operating solely on the results of a series of individual 1-D retrievals, and compares these to the cross-section retrieval.

In this section, x_i refers to the 1-D result of the i -th measured profile and S_i is the associated covariance matrix. Then the simple moving average for the i -th profile x_i^{avg} is defined as

$$x_i^{\text{avg}} = \sum_j w_{i,j} x_j. \quad (13)$$

The weights $w_{i,j} \in \mathbb{R}$ are chosen according to a Laplacian distribution (the probability density function is $\sqrt{2}\sigma^{-1} \exp(-\sqrt{2}|x - \mu|\sigma^{-1})$ with mean μ and standard deviation σ) with a given FWHM, thereby effectively prescribing a fixed horizontal resolution. The Laplacian

distribution was chosen because of its obvious relationship to the auto-regressive model and Tikhonov regularisation. In this way, the averaging should produce similar results to a cross-section retrieval given the FWHM of the Laplacian distribution is similar to the horizontal resolution of the cross-section retrieval. However, there are many other choices for window functions with different trade-offs, which cannot be fully explored here.

A second, more complex smoothing option is to weigh the profiles not only by a simple scalar but, in addition, with the inverse of their covariance matrix. This corresponds to a maximum likelihood estimator and is mathematically similar to a linear cross-section retrieval linearised at the state given by the assembly of the 1-D solutions.

The maximum likelihood estimator for the i -th profile x_i^{ml} is defined as

$$x_i^{\text{ml}} = \left(\sum_j w_{i,j} S_j^{-1} \right)^{-1} \sum_j w_{i,j} S_j^{-1} x_j. \quad (14)$$

The weights $w_{i,j}$ are again defined based on a Laplacian probability distribution with a mean of zero and a standard deviation corresponding to a FWHM of 30 km. This correlates roughly with the cross-section retrieval results for a

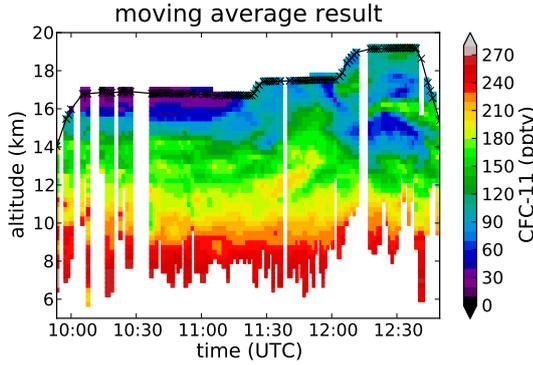


Fig. 13. Cross-section of CFC-11 volume mixing ratios derived by horizontally averaging 1-D retrieval results using a simple Laplacian distribution with 30 km FWHM to generate the weights.

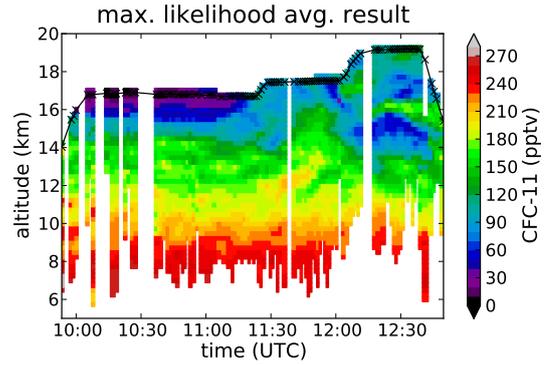


Fig. 14. Cross-section of CFC-11 volume mixing ratios derived by horizontally averaging 1-D retrieval results using both a simple Laplacian distribution with 30 km FWHM and the retrieval result covariance matrices to generate the weights.

factor-2000 horizontal regularisation. This stronger smoothing was chosen, as the effect of these methods and thereby also their differences become more apparent for stronger smoothing. For the maximum likelihood estimator, the actual width of the weighting functions is naturally also influenced by the involved covariance matrices, should they differ significantly from each other.

The horizontal distance between profiles is taken to be the average distance of the lower 20 km of the profiles. Other values for the FWHM have also been tested, and it was asserted that the presented strength delivers the result most similar to the factor-2000 horizontal regularisation. A key difference of cross-section retrievals is that the strength of horizontal smoothing is also influenced by the signal-to-noise ratio of the measurements and the vertical structure of the standard deviations of the climatology used for a priori, which leads to a stronger smoothing for lower altitudes than for higher altitudes. For either moving average, this effect might be incorporated to some extent by altitude-dependent weights, i.e. a diagonal matrix with altitude-dependent weights on the diagonal.

The result in CFC-11 volume mixing ratios for the simple weighted moving average is shown in Fig. 13. The volume mixing ratios contain obviously much less image noise than those in the baseline retrieval shown in Fig. 2. In so far this method is an effective means to simply reduce the impact of measurement noise. The result looks overall similar to the corresponding Fig. 4; the only obvious discrepancies can be found around obvious artefacts of the baseline retrieval. The profile with missing observations at 10:35 UTC still presents too low CFC-11 volume mixing ratios at lower altitudes. Interestingly, also applying stronger smoothing by Laplace distributions with FWHMs up to 200 km does not fully remedy this problem. In addition, there are some individual points with increased CFC-11 volume mixing ratio around 10:00 UTC, which are likely induced by uncertainties in elevation angle due to aircraft movements during the ascent. Neighbouring profiles fit much better to one

another with horizontal regularisation compared to this simple weighted moving average. For this method, diagnostic information can simply be assembled by weighted averaging of the available error covariance matrices of the individual profiles. It is also straightforward to generate a 2-D averaging kernel matrix.

The more complicated maximum likelihood estimator delivers very similar results presented in Fig. 14. The only notable difference to Fig. 13 is in the profile at 10:35 UTC, which is slightly better re-constructed. However, some values in the range between 9 and 11 km do not have an increased standard deviation as one might expect, even though their value seems obviously biased. As a consequence, the weighted average does not fill these values from the surrounding profiles. The way that this method combines the available profiles is so close to the method of the actual cross-section retrievals that an even closer agreement between this result and the cross-section retrieval results would be expected. However, the cross-section retrieval is a non-linear process, while this recombination of profiles is inherently linear. In this way, the cross-section retrieval is more potent.

With respect to providing diagnostic information, the covariance matrix \mathbf{S}_i^{ml} of the retrieval result of each profile \mathbf{x}_i^{ml} follows from Eq. (14) and can be calculated by the following formula:

$$\mathbf{S}_i^{\text{ml}} = \left(\sum_j w_{i,j} \mathbf{S}_j^{-1} \right)^{-2} \sum_j w_{i,j}^2 \mathbf{S}_j^{-1}. \quad (15)$$

It is straightforward to derive further diagnostic quantities like gain and averaging kernel matrix from this. When applying this on the presented numerical example, the calculation of \mathbf{S}_i^{ml} became numerically unstable even after applying typical and more elaborate countermeasures, causing several negative variances to appear in the matrix. This is surprising as such issues arise typically neither in the diagnostics of conventional 1-D profiles with matrices of similar size nor

in the diagnostics of the full cross-section retrieval involving much larger matrices. However, a closer analysis reveals that the matrices being multiplied and inverted in Eq. (15) are not too well conditioned to begin with (about $\approx 10^5$ for the summed matrices after symmetrically scaling the diagonal to 1) and both the matrix-matrix multiplication and the matrix-inversion are rather sensitive to high condition numbers: for the given matrices, the combination of the squaring and inversion in Eq. (15) introduces sufficient error to effectively remove up to ≈ 33 bit of information, which is more than is usually employed for storing the individual covariance matrices.

A different question is whether such smoothing methods can also be used to reduce the vertical regularisation strength. Initial tests with the discussed test case showed that the increase in oscillations in the 1-D profiles due to vertical decreased regularisation cannot be fully compensated by the a posteriori horizontal smoothing (not depicted). Here, the stabilising influence of horizontal regularisation during the retrieval is more critical as the chosen vertical regularisation strength is insufficient to guarantee numerical stability and well-posedness.

Combining individual profiles in a weighted manner can produce some of the benefits of the cross-section retrieval, even though the results are slightly worse as the weighting is only a linear process compared to the non-linear cross-section retrieval. Moving averages have the advantage of being able to provide diagnostic information in a computationally less costly way at the cost of a result of reduced retrieval quality.

4 Discussion and conclusions

This paper proposed and demonstrated non-linear 2-D cross-section retrievals of trace gases including vertical and horizontal regularisation from airborne infrared limb sounder measurements. As test case, the second measurement flight on 2 March 2010 during the RECONCILE campaign was used (Ungermann et al., 2012a).

An algorithm that is able to retrieve large-scale 2-D cross-sections was presented, and it was shown how associated diagnostic information can be produced on common hardware. Stochastic algorithms can provide the error estimates for various error sources and all retrieved quantities with a cost of only $O(n^2)$ compared to $O(n^3)$ for deterministic methods. These algorithms were then applied to a real-life test case to demonstrate the feasibility and quantify the benefits of the cross-section retrieval.

First, the effect of adding horizontal regularisation of varying strength was examined. Using 200 times the vertical regularisation strength as horizontal strength (factor-200 horizontal regularisation) produced visibly more uniform images without noticeable degradation of features. Increasing the horizontal regularisation strength further produces ever more

uniform images that still retain all evident structures until the picture becomes over-smoothed around factor-20 000 horizontal regularisation.

After establishing a suitable range for the horizontal regularisation strength, its impact on the retrieval error was examined. A detailed analysis showed that all error sources remained unaffected with the exception of a reduced noise error. The factor-200 horizontal regularisation reduced the noise error for the CFC-11 retrieval down to 70 % of the original level, while a more aggressive strength could reduce it down to about 30 % of the original level. For CIONO₂, factor-200 horizontal regularisation could reduce the image noise even by up to 50 %.

Introducing horizontal regularisation affects obviously the along-flight path resolution. Factor-200 horizontal regularisation imposes a lower limit of ≈ 17 km for well-resolved trace gases on this horizontal resolution, which is just about 2 km worse than the minimal profile distance. The more the horizontal regularisation strength is increased from there on, the worse the horizontal resolution becomes. For factor-2000 horizontal regularisation, the imposed minimum mean horizontal resolution is already more than 30 km.

Using horizontal regularisation, it is possible to reduce the vertical regularisation strength below the strength necessary for successful 1-D retrieval given a certain fixed noise error target. In the presented example, the vertical resolution of CIONO₂ could be improved by up to 30 % without increasing the image noise level. This enhanced the visibility of thin filaments and generally reduced the vertical smoothing. Choosing a good combination of vertical and horizontal regularisation thereby improves the reproduction of boundaries between air masses and thereby the examination of, for example, mixing processes.

The last section compared the presented approach to a simpler weighted averaging of one-dimensionally retrieved profiles. A simple linear recombination using a Laplace function with a set FWHM and an average that also includes the covariance matrices of the results was examined. Both approaches reduce the impact of measurement noise in the same way as the cross-section retrieval. However, they were both less able to reduce artefacts caused by missing observations.

By using cross-section retrievals, it is possible to produce a better representation of the true atmospheric state by exploiting the high measurement density of modern instruments and the self-similarity of the atmosphere. Especially for trace gases with weak signature, the technique reduces the image noise significantly without noticeable degradation of the horizontal resolution. The better reproduction of thin vertical layers is important for the analysis of mixing processes in the upper troposphere/lower stratosphere. The algorithm is therefore currently used to process further CRISTA-NF data (Ungermann et al., 2012b) and initial (non-tomographic) measurements by GLORIA (Gimballed Limb Observer for Radiance Imaging of the Atmosphere; see Ungermann et al.,

2010b). This technique might also be used to more reliably derive constant and slowly varying instrument parameters, which cannot be determined from pre- or post-flight calibration. The technique should also be applicable to older airborne sounders with sparser sampling such as MIPAS-STR (e.g. Woiwode et al., 2012), albeit with less resulting image enhancements.

In contrast to the title of this paper, which was chosen mostly due to the discussed use-case, it is straightforward to extend the presented technique to retrievals for current satellite-borne nadir-sounders, as these instruments also measure closely spaced profiles (e.g. both the Atmospheric Infrared Sounder (Aumann et al., 2003) and the Infrared Atmospheric Sounding Interferometer (Clerbaux et al., 2009) have a (sub)pixel size of 12 to 13.5 km; the operational retrievals for both instruments combine sub-pixels to reduce measurement noise, which could be achieved with less cost of horizontal resolution with the proposed algorithm). As the viewing geometry is different, the technique needs to be slightly adapted. Instead of cross-sections, 3-D cubes would be retrieved with added horizontal regularisation in both along-flight path and across-flight path direction. This only changes the state vector representation and the setup of the regularisation matrix. While the resulting problem size would be noticeably larger than in the presented cross-section retrieval, it is not larger than tomographic retrieval problems already treated by Ungermann et al. (2011). The technique might be even more beneficial when applied to nadir-sounders due to the greater ill-posedness of the retrieval compared to limb sounder retrievals. Examining and quantifying this in detail deserves further study.

It is also straightforward to extend the presented technique to proposed near-future satellite limb imagers (Riese et al., 2005; ESA, 2012). Such a rearward-looking instrument uses a 2-D detector to acquire multiple profiles simultaneously. Due to the high measurement speed, consecutive images overlap in the sense that they measure largely the same air mass. Assembling the measured profiles into several 2-D swaths parallel to the flight path allows for the use of 2-D tomographic retrieval techniques (e.g. Carlotti et al., 2001; Ungermann et al., 2010a) to achieve an excellent resolution in all three dimensions. The proposed technique is capable of evaluating all 2-D swaths together in a single 3-D tomographic retrieval. By exploiting the similarity between neighbouring swaths, the described technique would stabilise the inherently more ill-posed tomographic retrieval problem and reduce the image noise level while possibly also improving the resolution. This might enhance the scientific capabilities of the limb imager, for example with respect to gravity wave detection (see Preusse et al., 2009).

Appendix A

Analytical inverse of covariance matrices defined by an auto-regressive model

The auto-regressive model is a simple way to construct a covariance matrix from available variances and (ad hoc) assumptions about correlation length. While more common for the one-dimensional case, it can be generalised to an arbitrary number of dimensions. For simplicity's sake, only a single type of quantity and a two-dimensional cross-section setup is assumed for the discussion:

$$(\mathbf{S})_{i,j} = \sigma_i \sigma_j e^{-\frac{|v_i - v_j|}{c_v}} e^{-\frac{|h_i - h_j|}{c_h}}. \quad (\text{A1})$$

Thereby v_i , and h_i represent the vertical and horizontal coordinate of data point i . If the grid is rectilinear, one can define two matrices, which each remove all exponential terms of \mathbf{S} belonging to one spatial dimension. Each such matrix is symmetric, and only three of its diagonals are non-zero. The product of the two matrices composes the desired precision matrix \mathbf{S}^{-1} . This is a straightforward extension of the analytical inverse for 1-D problems described by Steck and von Clarmann (2001). With an efficient implementation, this inverse can be set up in linear time, thereby consuming only a negligible amount of computation time and memory.

The two-dimensional, rectilinear grid is defined by two vectors, $\mathbf{g}^v \in \mathbb{R}^{n_v}$ and $\mathbf{g}^h \in \mathbb{R}^{n_h}$, that consist of the ascending coordinates. The elements of the state space vector \mathbf{x} are ordered also according to the coordinates, varying fastest in vertical and slowest in horizontal coordinate. Last, two maps of m_v and m_h are needed that map the index i of an element of the state vector \mathbf{x} to the corresponding index of the appropriate coordinate in the vectors \mathbf{g}^v and \mathbf{g}^h , respectively. Here, the convention is used that g_i^v denotes the i -th element of \mathbf{g}^v .

Let the helper expression $\beta^v(i, j)$ be $\exp(-(g_i^v - g_j^v) \cdot c_v^{-1})$, and $b_{i,j}^v$, the element in the i -th row and j -th column of the matrix \mathbf{B}^v . Then the matrix \mathbf{B}^v is defined as

$$b_{i,j}^v = \begin{cases} \frac{1}{1 - \beta^v(m_v(i) + 1, m_v(i))^2} & \text{if } i = j \text{ and } m_v(i) = 0 \\ \frac{1 - \beta^v(m_v(i) + 1, m_v(i) - 1)^2}{(1 - \beta^v(m_v(i), m_v(i) - 1)^2)(1 - \beta^v(m_v(i) + 1, m_v(i))^2)} & \text{if } i = j \text{ and } 0 < m_v(i) < n_v - 1 \\ \frac{1}{1 - \beta^v(m_v(i), m_v(i) - 1)^2} & \text{if } i = j \text{ and } m_v(i) = n_v - 1 \\ -\frac{\beta^v(m_v(i) + 1, m_v(i))}{1 - \beta^v(m_v(i) + 1, m_v(i))^2} & \text{if } m_v(i) = m_v(j) - 1, m_v(i) < n_v - 1 \text{ and } m_h(i) = m_h(j) \\ -\frac{\beta^v(m_v(i), m_v(i) - 1)}{1 - \beta^v(m_v(i), m_v(i) - 1)^2} & \text{if } m_v(i) = m_v(j) + 1, m_v(i) > 0 \text{ and } m_h(i) = m_h(j) \\ 0 & \text{else.} \end{cases} \quad (\text{A2})$$

The first three terms describe the main diagonal; the following two describe the two side-diagonals. In the case of multi-target retrievals, an additional condition must ensure that the entry is zero in case i and j describes different entities. It is straightforward to likewise define the matrix \mathbf{B}^h . To handle the standard deviations σ_i in \mathbf{S} , an element-wise division of only \mathbf{B}^v by the matrix $(\sigma_i \cdot \sigma_j)_{i,j}$ is sufficient. Then $\mathbf{B}^v \cdot \mathbf{B}^h = \mathbf{S}^{-1}$. The resulting precision matrix has only up to 9 non-zero entries per row.

Appendix B

Analytical root of the precision matrix defined by an auto-regressive model

With the sparse precision matrix \mathbf{S}^{-1} constructed as described above, one can in turn apply the Cholesky–Banachiewicz algorithm (Banachiewicz, 1938) to efficiently construct a Cholesky decomposition, which delivers $\mathbf{S}^{-\frac{1}{2}}$. As shorthand, the individual elements of the inverse and the root are identified by $\mathbf{S}^{-1} = (a_{i,j})_{i,j}$ and $\mathbf{S}^{-\frac{1}{2}} = (l_{i,j})_{i,j}$. Then, the root can be constructed row by row as

$$l_{i,j} = \begin{cases} \sqrt{a_{i,i} - \sum_{k=1}^{i-1} l_{i,k}^2}, & \text{for } i = j, \\ l_{j,j}^{-1} \left(a_{i,j} - \sum_{k=1}^{j-1} l_{i,k} l_{j,k} \right), & \text{for } i > j, \\ 0, & \text{for } i < j. \end{cases} \quad (\text{B1})$$

Using this definition, the root can be row-wise constructed. However, the already assigned elements are accessed column-wise during the construction of further elements. To handle this efficiently, the root should be stored in a row-major compressed matrix storage with an additional data structure containing sorted indices of the non-zero elements of each column. The additional data structure is only required to speed up the construction of the root and can be discarded afterwards. For precision matrices constructed as described above, the non-zero elements of the root are a subset of the non-zero elements of the precision matrix itself. When exploiting this, the amount of time required for the construction of the root is directly proportional to the number of rows and is thereby negligible compared to the total cost for a retrieval.

Acknowledgements. J. Ungermann is supported by the Deutsche Forschungsgemeinschaft (GZ UN 311/1-1 and HALO SPP/1294 KA2324/1-2) and the visitor's program of the Atmospheric Chemistry Division at the National Center for Atmospheric Research (NCAR). NCAR is funded by the National Science Foundation. The RECONCILE project is funded under the European Commission Seventh Framework Programme (FP7) under grant number

RECONCILE-226365-FP7-ENV-2008-1. The author thanks the RECONCILE coordination and flight planning teams, MDB and Enviscope for the successful implementation of and support during the Arctic field campaign. The author especially thanks the CRISTA-NF team for the great work before, during, and after the RECONCILE campaign and for the provision of the measurements used in the case study. The European Centre for Medium-Range Weather Forecasts (ECMWF) is acknowledged for meteorological data support.

The service charges for this open access publication have been covered by a Research Centre of the Helmholtz Association.

Edited by: U. Friess

References

- Aumann, H. H., Chahine, M. T., Gautier, C., Goldberg, M. D., Kalnay, E., McMillin, L. M., Revercomb, H., Rosenkranz, P. W., Smith, W. L., Staelin, D. H., Strow, L. L., and Susskind, J.: AIRS/AMSU/HSB on the Aqua Mission: Design, Science Objective, Data Products, and Processing Systems, *IEEE T. Geosci. Remote*, 41, 253–264, doi:10.1109/TGRS.2002.808356, 2003.
- Banachiewicz, T.: Methode de Resolution Numerique des Equations Lineaires, Du Calcul des Determinants et des Inverses et de Reduction des Formes Quadratiques, *Bull. Ac. Pol. Sci. A*, 393–404, 1938.
- Birner, T.: Fine-scale structure of the extratropical tropopause region, *J. Geophys. Res.*, 111, D04104, doi:10.1029/2005JD006301, 2006.
- Carlotti, M., Dinelli, B. M., Raspollini, P., and Ridolfi, M.: Geo-fit approach to the analysis of limb-scanning satellite measurements, *Appl. Optics*, 40, 1872–1885, doi:10.1364/AO.40.001872, 2001.
- Clerbaux, C., Boynard, A., Clarisse, L., George, M., Hadji-Lazaro, J., Herbin, H., Hurtmans, D., Pommier, M., Razavi, A., Turquety, S., Wespes, C., and Coheur, P.-F.: Monitoring of atmospheric composition using the thermal infrared IASI/MetOp sounder, *Atmos. Chem. Phys.*, 9, 6041–6054, doi:10.5194/acp-9-6041-2009, 2009.
- Curtis, A. R.: Discussion of “A statistical model for water vapour absorption” by R. M. Goody, *Q. J. Roy. Meteorol. Soc.*, 78, 638–640, 1952.
- ESA: Report for Mission Selection: PREMIER, vol. SP-1324/3, ESA Communication Production Office, Noordwijk, The Netherlands, 2012.
- Forster, P. M. D. F. and Shine, K. P.: Radiative forcing and temperature trends from stratospheric ozone changes, *J. Geophys. Res.*, 102, 10841–10855, doi:10.1029/96JD03510, 1997.
- Francis, G. L., Edwards, D. P., Lambert, A., Halvorson, C. M., Lee-Taylor, J. M., and Gille, J. C.: Forward modeling and radiative transfer for the NASA EOS-Aura high resolution dynamics limb sounder (HIRDLS) instrument, *J. Geophys. Res.*, 111, D13301, doi:10.1029/2005JD006270, 2006.
- Friedl-Vallon, F., Riese, M., Maucher, G., Lengel, A., Hase, F., Preusse, P., and Spang, R.: Instrument concept and preliminary performance analysis of GLORIA, *Adv. Space Res.*, 37, 2287–2291, doi:10.1016/j.asr.2005.07.075, 2006.

- Gettelman, A., Hoor, P., Pan, L. L., Randell, W. J., Hegglin, M. I., and Birner, T.: The extra tropical upper troposphere and lower stratosphere, *Rev. Geophys.*, 49, RG3003, doi:10.1029/2011RG000355, 2011.
- Glatthor, N., von Clarmann, T., Fischer, H., Funke, B., Grabowski, U., Höpfner, M., Kellmann, S., Kiefer, M., Linden, A., Milz, M., Steck, T., and Stiller, G. P.: Global peroxyacetyl nitrate (PAN) retrieval in the upper troposphere from limb emission spectra of the Michelson Interferometer for Passive Atmospheric Sounding (MIPAS), *Atmos. Chem. Phys.*, 7, 2775–2787, doi:10.5194/acp-7-2775-2007, 2007.
- Godson, W. L.: The evaluation of infra-red radiative fluxes due to atmospheric water vapour, *Q. J. Roy. Meteorol. Soc.*, 79, 367–379, 1953.
- Gordley, L. L. and Russell, J. M.: Rapid inversion of limb radiance data using an emissivity growth approximation, *Appl. Optics*, 20, 807–813, doi:10.1364/AO.20.000807, 1981.
- Grossmann, K. U., Offermann, D., Gusev, O., Oberheide, J., Riese, M., and Spang, R.: The CRISTA-2 mission, *J. Geophys. Res.*, 107, 8173, doi:10.1029/2001JD000667, 2002.
- Hegglin, M. I., Boone, C. D., Manney, G. L., and Walker, K. A.: A global view of the extratropical tropopause transition layer from atmospheric chemistry experiment Fourier transform spectrometer O₃, H₂O, and CO, *J. Geophys. Res.*, 114, D00B11, doi:10.1029/2008JD009984, 2009.
- Hoffmann, L., Kaufmann, M., Spang, R., Müller, R., Remedios, J. J., Moore, D. P., Volk, C. M., von Clarmann, T., and Riese, M.: Envisat MIPAS measurements of CFC-11: retrieval, validation, and climatology, *Atmos. Chem. Phys.*, 8, 3671–3688, doi:10.5194/acp-8-3671-2008, 2008.
- Kenney, J. F. and Keeping, E. S.: *The Distribution of the Standard Deviation*, vol. 2, 2nd Edn., Van Nostrand, 1951.
- Kullmann, A., Riese, M., Olschewski, F., Stroh, F., and Grossmann, K. U.: Cryogenic infrared spectrometers and telescopes for the atmosphere – new frontiers, *Proc. SPIE*, 5570, 423–432, 2004.
- Livesey, N. J. and Read, W. G.: Direct Retrieval of Line-of-Sight Atmospheric Structure from Limb Sounding Observations, *Geophys. Res. Lett.*, 27, 891–894, doi:10.1029/1999GL010964, 2000.
- 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.
- Nocedal, J. and Wright, S. J.: *Numerical Optimization*, Springer Series in Operations Research and Financial Engineering, 2nd Edn., Springer, New York, 2006.
- Offermann, D., Grossmann, K.-U., Barthol, P., Knieling, P., Riese, M., and Trant, R.: Cryogenic infrared spectrometers and telescopes for the atmosphere (CRISTA) experiment and middle atmosphere variability, *J. Geophys. Res.*, 104, 16311–16325, doi:10.1029/1998JD100047, 1999.
- Preusse, P., Schroeder, S., Hoffmann, L., Ern, M., Friedl-Vallon, F., Ungermann, J., Oelhaf, H., Fischer, H., and Riese, M.: New perspectives on gravity wave remote sensing by spaceborne infrared limb imaging, *Atmos. Meas. Tech.*, 2, 299–311, doi:10.5194/amt-2-299-2009, 2009.
- Remedios, J. J., Leigh, R. J., Waterfall, A. M., Moore, D. P., Sembhi, H., Parkes, I., Greenhough, J., Chipperfield, M.P., and Hauglustaine, D.: MIPAS reference atmospheres and comparisons to V4.61/V4.62 MIPAS level 2 geophysical data sets, *Atmos. Chem. Phys. Discuss.*, 7, 9973–10017, doi:10.5194/acpd-7-9973-2007, 2007.
- Riese, M., Friedl-Vallon, F., Spang, R., Preusse, P., Schiller, C., Hoffmann, L., Konopka, P., Oelhaf, H., von Clarmann, T., and Höpfner, M.: Global limb radiance imager for the atmosphere (GLORIA): scientific objectives, *Adv. Space Res.*, 36, 989–995, doi:10.1016/j.asr.2005.04.115, 2005.
- Rodgers, C. D.: *Inverse Methods for Atmospheric Sounding: Theory and Practice*, vol. 2 of Series on Atmospheric, Oceanic and Planetary Physics, World Scientific, Singapore, 2000.
- Rothman, L., Gordon, I., Barbe, A., Benner, D., Bernath, P., Birk, M., Boudon, V., Brown, L., Campargue, A., Champion, J.-P., Chance, K., Coudert, L., Dana, V., Devi, V., Fally, S., Flaud, J.-M., Gamache, R., Goldman, A., Jacquemart, D., Kleiner, I., Lacome, N., Lafferty, W., Mandin, J.-Y., Massie, S., Mikhailenko, S., Miller, C., Moazzen-Ahmadi, N., Naumenko, O., Nikitin, A., Orphal, J., Perevalov, V., Perrin, A., Predoi-Cross, A., Rinsland, C., Rotger, M., Simeckov, M., Smith, M., Sung, K., Tashkun, S., Tennyson, J., Toth, R., Vandaele, A., and Auwera, J. V.: The HITRAN 2008 molecular spectroscopic database, *J. Quant. Spectrosc. Ra.*, 110, 533–572, doi:10.1016/j.jqsrt.2009.02.013, 2009.
- Steck, T. and von Clarmann, T.: Constrained profile retrieval applied to the observation mode of the Michelson interferometer for passive atmospheric sounding, *Appl. Optics*, 40, 3559–3571, doi:10.1364/AO.40.003559, 2001.
- Steck, T., Höpfner, M., von Clarmann, T., and Grabowski, U.: Tomographic retrieval of atmospheric parameters from infrared limb emission observations, *Appl. Optics*, 44, 3291–3301, doi:10.1364/AO.44.003291, 2005.
- Tarantola, A.: *Inverse Problem Theory*, Society for Industrial and Applied Mathematics, Philadelphia, 2004.
- Tikhonov, A. N. and Arsenin, V. Y.: *Solutions of ill-posed problems*, Winston, Washington DC, USA, 1977.
- Ungermann, J., Hoffmann, L., Preusse, P., Kaufmann, M., and Riese, M.: Tomographic retrieval approach for mesoscale gravity wave observations by the PREMIER Infrared Limb-Sounder, *Atmos. Meas. Tech.*, 3, 339–354, doi:10.5194/amt-3-339-2010, 2010a.
- Ungermann, J., Kaufmann, M., Hoffmann, L., Preusse, P., Oelhaf, H., Friedl-Vallon, F., and Riese, M.: Towards a 3-D tomographic retrieval for the air-borne limb-imager GLORIA, *Atmos. Meas. Tech.*, 3, 1647–1665, doi:10.5194/amt-3-1647-2010, 2010b.
- Ungermann, J., Blank, J., Lotz, J., Leppkes, K., Hoffmann, L., Guggenmoser, T., Kaufmann, M., Preusse, P., Naumann, U., and Riese, M.: A 3-D tomographic retrieval approach with advection compensation for the air-borne limb-imager GLORIA, *Atmos. Meas. Tech.*, 4, 2509–2529, doi:10.5194/amt-4-2509-2011, 2011.
- Ungermann, J., Kalicinsky, C., Olschewski, F., Knieling, P., Hoffmann, L., Blank, J., Woiwode, W., Oelhaf, H., Hösen, E., Volk, C. M., Ulanovsky, A., Ravegnani, F., Weigel, K., Stroh, F., and Riese, M.: CRISTA-NF measurements with unprecedented vertical resolution during the RECONCILE aircraft campaign, *Atmos. Meas. Tech.*, 5, 1173–1191, doi:10.5194/amt-5-1173-2012, 2012a.

- Ungermann, J., Pan, L., Kalicinsky, C., Olschewski, F., Knieling, P., Blank, J., Guggenmoser, T., Stroh, F., Hoffmann, L., and Riese, M.: Filamentary structure in chemical tracer distributions near the subtropical jet following a wave-breaking event, *Atmos. Chem. Phys. Discuss.*, submitted, 2012b.
- Weinreb, M. P. and Neuendorffer, A. C.: Method to apply homogeneous-path transmittance models to inhomogeneous atmospheres, *J. Atmos. Sci.*, 30, 662–666, doi:10.1175/1520-0469(1973)030<0662:MTAHPT>2.0.CO;2, 1973.
- Woiwode, W., Oelhaf, H., Gulde, T., Piesch, C., Maucher, G., Ebersoldt, A., Keim, C., Höpfner, M., Khaykin, S., Ravegnani, F., Ulanovsky, A. E., Volk, C. M., Hösen, E., Dörnbrack, A., Ungermann, J., Kalicinsky, C., and Orphal, J.: MIPAS-STR measurements in the Arctic UTLS in winter/spring 2010: instrument characterization, retrieval and validation, *Atmos. Meas. Tech.*, 5, 1205–1228, doi:10.5194/amt-5-1205-2012, 2012.