

A 3-D tomographic retrieval approach with advection compensation for the air-borne limb-imager GLORIA

J. Ungermann¹, J. Blank¹, J. Lotz², K. Leppkes², L. Hoffmann³, T. Guggenmoser¹, M. Kaufmann¹, P. Preusse¹, U. Naumann², and M. Riese¹

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

²LuFG Informatik 12: Software and Tools for Computational Engineering, RWTH Aachen, Germany

³Jülich Supercomputing Centre, Research Centre Jülich GmbH, Jülich, Germany

Received: 6 June 2011 – Published in Atmos. Meas. Tech. Discuss.: 17 June 2011

Revised: 4 October 2011 – Accepted: 14 November 2011 – Published: 24 November 2011

Abstract. Infrared limb sounding from aircraft can provide 2-D curtains of multiple trace gas species. However, conventional limb sounders view perpendicular to the aircraft axis and are unable to resolve the observed airmass along their line-of-sight. GLORIA (Gimballed Limb Observer for Radiance Imaging of the Atmosphere) is a new remote sensing instrument that is able to adjust its horizontal view angle with respect to the aircraft flight direction from 45° to 135°. This will allow for tomographic measurements of mesoscale structures for a wide variety of atmospheric constituents.

Many flights of the GLORIA instrument will not follow closed curves that allow measuring an airmass from all directions. Consequently, it is examined by means of simulations, what spatial resolution can be expected under ideal conditions from tomographic evaluation of measurements made during a straight flight. It is demonstrated that the achievable horizontal resolution in the line-of-sight direction could be reduced from over 200 km to around 70 km compared to conventional retrievals and that the tomographic retrieval is also more robust against horizontal gradients in retrieved quantities in this direction. In a second step, it is shown that the incorporation of channels exhibiting different optical depth can further enhance the spatial resolution of 3-D retrievals enabling the exploitation of spectral samples usually not used for limb sounding due to their opacity.

A second problem for tomographic retrievals is that advection, which can be neglected for conventional retrievals, plays an important role for the time-scales involved in a tomographic measurement flight. This paper presents a method to diagnose the effect of a time-varying atmosphere on a 3-D

retrieval and demonstrates an effective way to compensate for effects of advection by incorporating wind-fields from meteorological datasets as a priori information.

1 Introduction

The upper troposphere/lower stratosphere (UTLS) is a region of importance for the radiative forcing and thereby for understanding climate change (e.g. Forster and Shine, 1997). Its composition and structure determines the exchange between the tropospheric and stratospheric air (e.g. Gettelman et al., 2011). When examining small-scale dynamic structures in this region, there is an observational gap between in situ instruments and air-borne or satellite-borne remote sensing instruments: in situ instruments provide measurements with a very good resolution but lack with respect to coverage, whereas remote sensing instruments provide good coverage albeit at the cost of comparably worse spatial resolution along at least one spatial dimension.

Small-scale dynamic structures arise frequently in the atmosphere. They are not well understood, as they cannot yet be properly observed by current in situ, limb, or nadir observations. Quite often, model calculations are the only way to examine the underlying small-scale processes. One example are mesoscale structures like filaments and tropopause folds in the upper troposphere/lower stratosphere, which are crucial for the UTLS composition and variability (Konopka et al., 2009).

Limb-sounding measures incident radiation from the limb of the Earth's atmosphere. In the case of infrared limb-sounding, this radiation is emitted by ro-vibrationally excited molecules along the ray path or line-of-sight (LOS) of the



Correspondence to: J. Ungermann
(j.ungermann@fz-juelich.de)

instrument. The point of the LOS closest to the surface is called the tangent point and the shortest distance between the tangent point and the surface of the Earth is called the tangent altitude. As the atmosphere is densest at the tangent point, the part of the LOS around this location generally contributes most to the measured radiation if absorption can be neglected. Limb-sounders typically take measurements at different tangent altitudes to collect information about the vertical structure of the atmosphere.

While limb-sounders offer a sufficient vertical resolution to observe small-scale vertical structures in the atmosphere, they require horizontal homogeneity along the LOS for undistorted non-tomographic retrievals, which is only given if the usually long-drawn atmospheric structure is properly aligned with the LOS of the instrument. As numerical simulations for dynamic structures are currently not accurate enough to reliably predict the exact location and sometimes even existence of such structures, one cannot ensure the proper alignment of the measuring instrument with the structure at hand. Consequently, remote sensing instruments with high resolution along all three spatial axes are required to properly examine such structures.

One way to improve the achievable resolution in the desired way is evaluating limb-sounder measurements with tomographic methods that combine multiple views of the same volume from different angles to produce a spatially resolved reconstruction of the examined object (e.g. Natterer, 2001). First steps towards tomographic evaluation of satellite measurement data have already been taken by combining the measurements of multiple consecutive vertical profiles to recreate a 2-D slice of atmosphere with presumably better horizontal resolution. Such retrievals for limb-sounding measurements were first explored by a variety of authors for different purposes and instruments (e.g. Solomon et al., 1984). Practical implementations for the large-scale retrieval of atmospheric constituents from satellite measurements 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 et al. (2006) for the Microwave Limb Sounder.

The 3-D tomographic evaluation of measurements made by the new air-borne GLORIA instrument (Gimballed Limb Observer for Radiance Imaging of the Atmosphere; Riese et al., 2005; Friedl-Vallon et al., 2006; Ungermann et al., 2010b) is a new conceptual approach for atmospheric limb-sounding. Figure 1 shows the tomographic measurement principle for GLORIA, which relies on its high measurement speed and its capability to pan the instrument to measure the same volume from multiple angles. The 3-D tomographic retrieval is thereby not a simple extension of 2-D tomography: first, the observation geometries differ as satellites usually are forward- or backward-looking and GLORIA looks sideways; second, the carrier speeds and thereby the involved time-scales are quite different. For instance, a satellite-borne instrument moves fast enough so that the atmospheric state




Fig. 1. Tomographic flight paths for GLORIA. A circular flight path is schematically shown in (a), while a linear flight path is depicted in (b). In practise, the measurement density is much denser than depicted here and the instrument is panned also for the circular flight depicted in (a).

does not change significantly between measurements with overlapping LOSs, whereas performing an air-borne tomographic measurement flight as shown in Fig. 1a may take hours. Further, retrieving a 3-D volume from measurements taken during a typical linear flight path of its air-borne carrier as shown in Fig. 1b poses a notoriously difficult limited-angle problem (e.g. Natterer, 2001, pages 144–151), as the atmospheric volume cannot be measured from all directions. Lastly, large parts of the LOSs of GLORIA measurements run through parts of the atmospheric volume, which cannot be viewed from multiple angles for practical matters, and therefore cannot be resolved with good quality.

This paper extends the studies presented in (Ungermann et al., 2010b). It proceeds to briefly review the GLORIA instrument and the inversion techniques required to perform tomographic retrievals and diagnose the results. The second part describes numerical studies assuming rather ideal conditions to demonstrate the principal feasibility of the approach and to derive first performance characteristics. Tomographic measurements obtained during a linear flight are simulated, evaluated and the results are compared to a conventional 1-D retrieval to identify the benefits of the more complex method. Tomography opens the very interesting opportunity to exploit channels of different optical depth, similar to how nadir sounders derive vertical profile information. In the last part, the effect of advection on tomographic measurements flights following a circular flight path is examined. A method for compensating advection by including wind-fields from meteorological datasets as a priori information is proposed and analysed.

2 The GLORIA infrared limb imager

GLORIA combines a Fourier transform spectrometer in the infrared spectral region with a two-dimensional detector array, thereby allowing to take 16 384 spectra simultaneously covering a wavelength region between 770 cm^{-1} and 1400 cm^{-1} (Friedl-Vallon et al., 2006). It is mounted to a cardanic frame that stabilises the instrument-LOS against

movements of the carrying aircraft and allows to point the instrument to different azimuth angles. With 0° being the flight direction, the instrument can be panned from 45° to 135° . However, the field-of-view might be slightly restricted by obstructing parts of the plane. The spectral resolution can be adapted from 0.1 cm^{-1} to 1.25 cm^{-1} , whereby the coarsest resolution allows to take a full set of 16 384 spectra every three seconds. The following numerical studies are based on this mode.

3 JURASSIC2 retrieval processor

The Jülich Rapid Spectral Simulation Code Version 2 (JURASSIC2) retrieval processor combines a fast radiative transfer forward model with a suit of inversion techniques. The Python/C++ based JURASSIC2 and its predecessor were used in the evaluation of numerous experiments (e.g. Hoffmann et al., 2008, 2009; Eckermann et al., 2009; Hoffmann and Alexander, 2009; Weigel et al., 2010; Ungermann et al., 2010b). The forward model consists of a core module implementing the emissivity growth approximation (EGA; e.g. Weinreb and Neuendorffer, 1973; Gordley and Russell, 1981) and Curtis-Godson approximation (CGA; Curtis, 1952; Godson, 1953) that are based on pre-calculated spectrally averaged values of emissivity stored in lookup tables that account already for the instrument line shape. The emissivity lookup tables are generated by the Reference Forward Model (Dudhia et al., 2002). This approach is about a factor 1000 faster than conventional line-by-line calculations at the cost of a slight loss of accuracy (Gordley and Russell, 1981).

The forward model $\mathbf{F} : \mathbb{R}^n \mapsto \mathbb{R}^m$ is then used to identify an atmospheric state that agrees with measurements made by the GLORIA instrument. This is implemented by minimising a cost function $J : \mathbb{R}^n \mapsto \mathbb{R}$,

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

containing two terms: the one on the left represents the fit of simulated measurements for a given discretised atmospheric state $\mathbf{x} \in \mathbb{R}^n$ to the actual measurements $\mathbf{y} \in \mathbb{R}^m$ and a second one to the right representing additional a priori information about the solution. The matrix $\mathbf{S}_\epsilon^{-1} \in \mathbb{R}^{m \times m}$ is the inverse of a covariance matrix of the measurement error, $\mathbf{x}_a \in \mathbb{R}^n$ represents an a priori state of the atmosphere, e.g. climatological means, and $\mathbf{S}_a^{-1} \in \mathbb{R}^{n \times n}$ may be either the inverse of a climatological covariance matrix or a Tikhonov regularisation matrix codifying, e.g. smoothness conditions. In this way, the original, ill-posed problem of inverting \mathbf{F} is approximated by a similar, well-posed problem. The added constraint stabilises the inversion at the cost of a bias that depends on the type and strength of the constraint. Different discretisations of the atmospheric state are supported by our forward model, most being realised as rectilinear 1-D, 2-D,

3-D, or 4-D (three spatial and one temporal dimension) grids with linear interpolation between the grid points.

The minimum of the cost function is searched for using a truncated Quasi-Newton minimiser employing the conjugate gradient algorithm for the involved linear equation systems:

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

where λ_i is a Levenberg-Marquardt parameter and \mathbf{I}_n should be exchanged with a scaling matrix if \mathbf{x} contains elements of different magnitude.

For the retrieval simulations presented in this paper, the regularisation matrix \mathbf{S}_a^{-1} is composed of four matrices \mathbf{L}_0 , \mathbf{L}_1^x , \mathbf{L}_1^y , and $\mathbf{L}_1^z \in \mathbb{R}^{n \times n}$ corresponding to regularisation of the absolute value and the first order derivative in the three spatial dimensions:

$$\mathbf{S}_a^{-1} = (\alpha_0)^2 \mathbf{L}_0^T \mathbf{L}_0 + (\alpha_1^x)^2 \mathbf{L}_1^x \mathbf{L}_1^x + (\alpha_1^y)^2 \mathbf{L}_1^y \mathbf{L}_1^y + (\alpha_1^z)^2 \mathbf{L}_1^z \mathbf{L}_1^z \quad (3)$$

with the tuning parameters α_0 , α_1^x , α_1^y , and $\alpha_1^z \in \mathbb{R}$. The matrix \mathbf{L}_0 is a diagonal matrix. The \mathbf{L}_1 matrices are simple Tikhonov regularisation matrices of first order (Ungermann, 2011). This approach can also be extended by further terms to implement the regularisation needed for deriving instrument parameters.

3.1 Adjoint Jacobian computation

Each column of the Jacobian matrix $\mathbf{F}'(\mathbf{x})$ required in the inversion can be approximately calculated by the finite difference of $(\mathbf{F}(\mathbf{x} + \mathbf{h}) - \mathbf{F}(\mathbf{x})) / \|\mathbf{h}\|_2$, where $\mathbf{h} \in \mathbb{R}^n$ perturbs only one element of \mathbf{x} . The Jacobian matrix $\mathbf{F}'(\mathbf{x})$ of \mathbf{F} for typical tomographic problems is rather sparse, as a single measurement is influenced only by a small fraction of the volume defined by \mathbf{x} . This sparsity can be exploited to speed up the calculation of the Jacobian matrix with finite differences by calculating only non-zero elements (Ungermann et al., 2010a), making this approach feasible even for tomographic problems. Another, more common, approach is implementing an analytical derivative of the model, which is computationally very efficient but requires elaborate, manual adjustments to the model.

A more efficient way is to employ algorithmic differentiation (e.g. Giering and Kaminski, 1998; Griewank and Walther, 2008) for the calculation of $\mathbf{F}'(\mathbf{x})$. Algorithmic differentiation exploits the chain rule and partial derivatives to calculate the exact directional derivative (to machine accuracy) of a function implemented in a programming language during its execution (Lotz, 2010). A brief introduction into the topic of algorithmic differentiation and its applications in JURASSIC2 is given in Appendix A. Due to the specific structure of our forward model, the adjoint method is even more efficient than typical for such a function: the computational effort for calculating the Jacobian matrix $\mathbf{F}'(\mathbf{x})$ is only

a small constant factor times the effort for evaluating \mathbf{F} instead of the factor $n + 1$ required for pure finite differences.

3.2 Diagnostics including missing dimensions

An important diagnostic entity is the averaging kernel matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ that maps the true atmospheric state $\mathbf{x}_t \in \mathbb{R}^n$ onto the retrieval result:

$$\mathbf{A} = \left(\mathbf{S}_a^{-1} + \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} \mathbf{F}'(\mathbf{x}_f) \right)^{-1} \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} \mathbf{F}'(\mathbf{x}_f). \quad (4)$$

It consists of the matrix product between the gain matrix $\mathbf{G} \in \mathbb{R}^{n \times m}$ and the Jacobian matrix evaluated at the solution \mathbf{x}_f with \mathbf{G} 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 averaging kernel matrix can be analysed to derive useful quantities like measurement contribution (indicating the influence of a priori information of zeroth order) or resolution. The simplest resolution measure is the inverse of the diagonal entries of \mathbf{A} (Purser and Huang, 1993). This resolution measure indicates the spread of information in relation to the sampling grid, but lacks a directional component. To derive a directed measure, each element of a row of \mathbf{A} may be placed at the location that it maps to the solution to derive the full-width at half-max (FWHM) in all directions. A second measure employed in this paper is the diameter of the smallest sphere containing all elements larger than the half-maximum (FWHM sphere in the following). Due to the difference in scale between horizontal and vertical dimensions, the extent of the sphere is entirely determined by the horizontal spread of the elements of the averaging kernel matrix row. Thus, the diameter can be interpreted as a measure for the horizontal resolution. The displacement between the centre of the sphere and the location of the retrieved atmospheric data point is also used in the following as vertical and horizontal dislocation.

When trying to analyse the effects of gradients of atmospheric quantities along the LOS of conventional 1-D retrievals or when trying to quantify the effect of a time-varying atmosphere on a static 3-D retrieval, it is necessary to embed the retrieval in a context of higher dimensionality: the solution \mathbf{x}_f is a fix-point of Eq. (2), implying

$$\left(\mathbf{S}_a^{-1} + \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} \mathbf{F}'(\mathbf{x}_f) \right)^{-1} \cdot \left(\mathbf{S}_a^{-1} \mathbf{x}_f + \mathbf{F}'(\mathbf{x}_f)^T \mathbf{S}_\epsilon^{-1} (\mathbf{F}(\mathbf{x}_f) - \mathbf{y}) \right) = (\mathbf{I} - \mathbf{A}) \mathbf{x}_a. \quad (6)$$

The vector of measurements \mathbf{y} can be approximated by a first order Taylor expansion of the forward model \mathbf{F} and the vector of measurement errors $\boldsymbol{\epsilon} \in \mathbb{R}^n$. In this case however, employing a different model enables the handling of a more accurate representation of the true atmospheric state, e.g. a 2-D atmosphere instead of a 1-D one for a conventional retrieval. The fix point \mathbf{x}_f can then be expressed in this representation

of higher dimensions by keeping its values constant along the newly added dimension. The forward model and vectors employing the additional dimension are marked by a tilde:

$$\mathbf{y} \approx \tilde{\mathbf{F}}(\tilde{\mathbf{x}}_f) + \tilde{\mathbf{F}}'(\tilde{\mathbf{x}}_f)(\tilde{\mathbf{x}}_t - \tilde{\mathbf{x}}_f) + \boldsymbol{\epsilon} \quad (7)$$

As $\tilde{\mathbf{x}}_f$ does not vary along the added dimension, it is obvious that $\tilde{\mathbf{F}}(\tilde{\mathbf{x}}_f) = \mathbf{F}(\mathbf{x}_f)$ and $\tilde{\mathbf{F}}'(\tilde{\mathbf{x}}_f)\tilde{\mathbf{x}}_f = \mathbf{F}'(\mathbf{x}_f)\mathbf{x}_f$ hold. It follows that

$$\mathbf{x}_f = \mathbf{G}\tilde{\mathbf{F}}'(\tilde{\mathbf{x}}_f)\tilde{\mathbf{x}}_t + (\mathbf{I} - \mathbf{A})\mathbf{x}_a + \mathbf{G}\boldsymbol{\epsilon}, \quad (8)$$

which resembles closely the usual form of linearised diagnostics (Rodgers, 2000). The main difference is the exchange of the true state \mathbf{x}_t and the Jacobian matrix $\mathbf{F}'(\mathbf{x}_f)$ by representations including an additional dimension. This defines a new kind of averaging kernel matrix

$$\tilde{\mathbf{A}} = \mathbf{G}\tilde{\mathbf{F}}'(\tilde{\mathbf{x}}_f) \quad (9)$$

that maps the true atmospheric state expressed in a higher dimensional representation onto the retrieval result. This matrix can be analysed in the usual way to derive a horizontal resolution for conventional 1-D retrievals. This is equivalent to a full retrieval including the additional dimension but with a regularising constraint that enforces the solution to remain constant along this dimension as von Clarmann et al. (2009) used to analyse the horizontal resolution of MIPAS. However, the presented approach is computationally much less involved, as it requires only a single evaluation of the higher-dimensional Jacobian matrix and some matrix-matrix multiplications in contrast to a full non-linear retrieval. Please note, that it is straightforward to extend this approach to adding more than a single dimension, if required.

This method should be used to compare high resolution 3-D model data with retrieval results, as this method incorporates the effects of horizontal gradients on the retrieval. The definition and application of higher-dimensional averaging kernel matrices is rather important, as conventional 1-D averaging kernel matrices disregard the effect of horizontal gradients and it is unclear whether discrepancies between the folded model data and the retrieval result stem from flaws in the model simulations, flaws in the retrieval, or from horizontal gradients. This definition also solves the problem, of where to sample horizontally the 3-D model atmosphere for comparison.

As the effort for the Jacobian calculation is independent of the size of the atmospheric representation, if an adjoint model is used and the ray-tracing step length remains constant, the effort to calculate $\tilde{\mathbf{F}}'(\tilde{\mathbf{x}}_f)$ is nearly identical with the effort for calculating $\mathbf{F}'(\mathbf{x}_f)$.

4 Simulation setup

To retain comparability, the simulation setup remains largely unchanged compared to the setup used by Ungermann et al. (2010b). The atmospheric situation used in the numerical

studies is based on a run of GEM-AQ recreating the synoptic situation during the European heat wave of July 2006 (Struzewska and Kaminski, 2008). Within this data set, there are many ozone filaments and streamers with an increased ozone concentration of more than 50 % in comparison to ambient air as shown in Fig. 2.

It is assumed that two horizontal lines of the detector array are co-added to generate a single, less noisy spectrum, of which only a single spectral sample, an ozone peak located around 778.5 cm^{-1} is used, which was also exploited by Weigel et al. (2010). This leaves 64 measurements at different tangent heights with a constant angular elevation distance of 0.0625° per image starting at 0.73° pointing upwards and going down from there on. This corresponds to a spatial distance of tangent points of 130 m 2 km below the aircraft, a distance of 280 m 5 km below the aircraft, and a distance of 400 m 10 km below the aircraft. The vertical FOV of two rows is assumed to have a full width at half maximum (FWHM) of 0.08° . Thereby, the FOV of two vertically neighbouring measurements overlaps meaningfully. The horizontal FOV is 4° . One image is taken every three seconds, increasing the azimuth angle by four degrees between taking successive images and returning to 45° to start again when no further increase is possible. Noise is modelled by a stochastic additive offset component and a stochastic multiplicative gain component conservatively chosen according to currently available instrument specification (standard deviation of offset: $0.1875 \times 10^{-5}\text{ W}/(\text{m}^2\text{ sr cm}^{-1})$, standard deviation of gain: 0.1 percent). The simulated measurements are generated by the same radiative transfer model that is also used in the retrieval, which usually provides optimistic results, as the measurements are not affected by systematic errors.

The retrieval shall derive ozone in the vertical region from 4 km to 20 km over the full horizontal extent from synthetic measurements. The inversion is simplified by assuming perfect knowledge of the top-column above 20 km and all entities except ozone in the retrieved area. This simplification allows to concentrate on a single target and to avoid the complexities of multi-target retrievals for the time being. For this study, we assume clear-sky conditions and disregard systematic errors. We choose ozone as test target, as it is a comparatively difficult trace gas to retrieve for an airborne limb-sounder because of the large contribution of radiance from the atmosphere above the instrument. While we assume the top-column of ozone to be known in the presented simulations, the tomographic retrievals produce essentially the same results below 15 km altitude for an unknown ozone top-column. We tried different scenarios from GEM-AQ and also CLAMS (Groß et al., 2005) datasets giving similar results as presented below. The presented scenario is basically a worst-case scenario due to large horizontal gradients at basically all altitude levels combined with a dominant wind speed orthogonal to the filament extent.




Fig. 2. The atmospheric situation. This figure depicts the ozone concentration of the true atmospheric state.

As stated before, we use Tikhonov regularisation of zeroth and first order. Ozone mid-latitude values and standard deviations of Remedios et al. (2007) are used for the a priori vector \mathbf{x}_a and regularisation matrix \mathbf{L}_0 . As initial guess \mathbf{x}_0 , mean ozone values for polar conditions are taken from the same source. The Tikhonov regularisation is parametrised according to Table 1. Using these parameters, the measurement contribution is always very close to 1 within the regions depicted in plots.

5 Linear flight path

In this section, 3-D tomographic retrievals are discussed that use synthetic GLORIA measurements made while following a linear flight track. As the GLORIA instrument will generally share the carrier with other instruments subject to different mission requirements, flying closed curves might not always be possible. The simplest shape usually present within a flight track is a straight line, which is examined here.

In this numerical simulation, the plane starts its flight at 6.5° E , 43.5° N and continues from there along the latitudinal circle until 4.0° W maintaining a constant altitude of 15 km. This flight track of 840 km length enables $\approx 600\text{ km}$ of air to be viewed from azimuthal angles differing up to 90° . The flight track has been placed so that the volume covered by tangent points at 12 km altitude is located at the boundary of the ozone filament. This generates the most challenging scenario because this positioning gives a steep gradient in across flight track direction. Figure 3 shows the tangent point coverage. With a sampling rate of one image consisting of 64 individual observations every three seconds, one full azimuth scan takes 66 s. With the given flight track length, the flight allows for 54 full azimuthal scans with 1 408 observations each. Altogether, these are 76 480 radiance values. To account for the effect of the FOV, 458 880 radiative transfer calculations are employed. The atmosphere is comprised

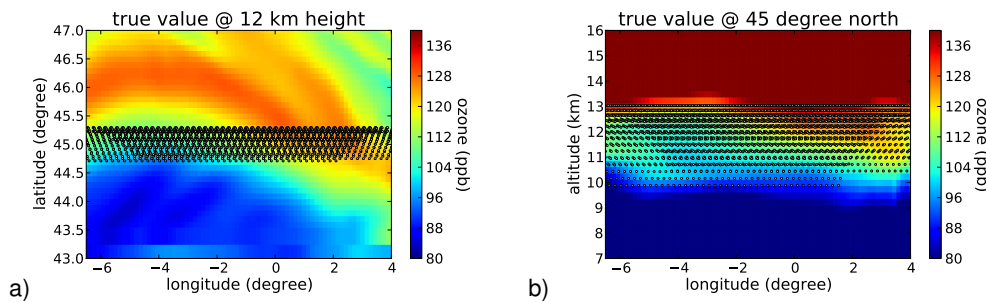


Fig. 3. The atmospheric situation and measurement geometry of the linear flight track numerical study. Panel (a) shows the true atmospheric state and the location of the tangent points at 12 km altitude. Panel (b) depicts the same for a vertical cutting plane at 45° N. The location of tangent points close to the depicted volume elements is indicated by small white circles.

Table 1. Regularisation strength for 3-D retrieval experiments.

Parameter	Value
α_0	0.1
α_1^x	0.8 km ppb ⁻¹
α_1^y	0.8 km ppb ⁻¹
α_1^z	4×10^{-4} km ppb ⁻¹

of volume elements, which extend $11 \text{ km} \times 11 \text{ km} \times 0.25 \text{ km}$ in the depicted regions. In Fig. 4, the sampling and sensitivity of the atmosphere with respect to the simulated measurements is depicted. This cutting plane is representative for the full volume except for the beginning and end of the flight, where less overlapping measurements are available. Panel (a) shows the number of measurements influenced by each volume element. The number of measurements traversing a volume element decreases with increasing distance from the instrument as the measurements vertically diverge. The number of ≈ 300 measurements northwards of 45° N per volume element is roughly consistent with each measurement passing through $\approx 3 \times 7$ (horizontal \times vertical) volume elements in 200 km distance due to FOV; in addition, ≈ 15 images are taken with different azimuth angles while traversing the 11 km side length of each volume element. As also minor contributions count as influence here, panel (b) shows the sum over the absolute values of the corresponding column of the Jacobian matrix. High values are a necessary, but not sufficient, precondition for a good reconstruction quality. The highest sensitivity is given in the measured airmass between the aircraft and the volume of high tangent point density, as the employed channel is not fully transparent.

5.1 Conventional 1-D retrieval for reference

In this section, the achievable horizontal resolution for a conventional 1-D retrieval performed under the same idealised

conditions as the tomographic one is analysed using the technique described in Sect. 3.2. It was already shown by Ungermann et al. (2010b) that over- and under-estimations of ozone in this atmospheric situation in the order of $\pm 20\%$ are possible for 1-D retrievals depending on the alignment between gradients in ozone and the LOSs.

One image pointing straight north is used as input for a 1-D retrieval with the same regularisation parameters as described in Table 1 (naturally without the horizontal smoothing component). This vertical regularisation strength is too weak, but the aim of this analysis is to derive a lower limit for the horizontal resolution, which benefits from a weak regularisation and is therefore favourable for the 1-D retrieval. For diagnosis, the 1-D atmosphere is extended with a gridding of 10 km spacing along the LOS and used to calculate a 2-D/1-D averaging kernel matrix $\tilde{\mathbf{A}}$. This matrix can be used to calculate the horizontal resolution of the retrieval. For this representative 1-D retrieval, the horizontal resolution within the tangent plane is close to 100 km, which corresponds to the length that the LOS of a measurement spends in the given altitude range. However, the diameter of the FWHM sphere (see Sect. 3.2) is much larger lying between ≈ 200 and 250 km, as volume elements above the tangent altitude also contribute meaningfully to the solution.

5.2 Analysis of the linear flight track

A 3-D tomographic retrieval of a simulated linear flight delivers the result depicted in Fig. 5. The relative error (which here and in the following is the difference of retrieval result and true state divided by the true state) stays below 5 % in the region covered by tangent points. The covered region is qualitatively and quantitatively well reproduced. Values further to the north and south stem largely from smoothing and extrapolation from the values in the central region. For this reason, should be disregarded or at least used with caution. It should be emphasised that the steep gradients in north-south direction in ozone do not cause large deviations in the retrieval result in the region covered by tangent points as would be expected for conventional 1-D retrievals. This is a first

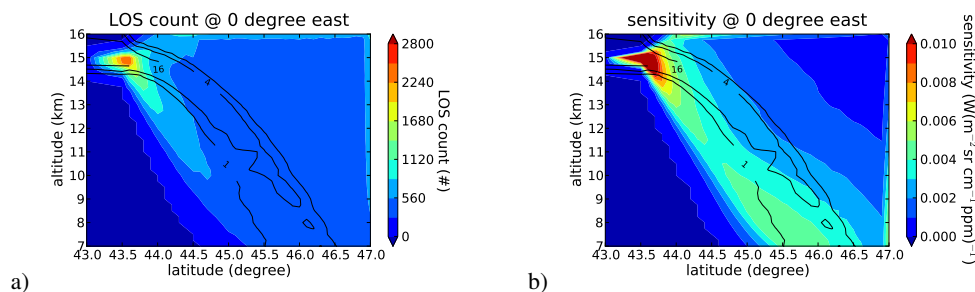


Fig. 4. The spatial distribution of measurements. Panel (a) shows for each volume element of the cutting plane the number of measurements sampling it. In (b), the sum over the absolute values of the columns of the Jacobian matrix is depicted. The number of tangent points within the depicted volume elements is overlaid as a contour plot.

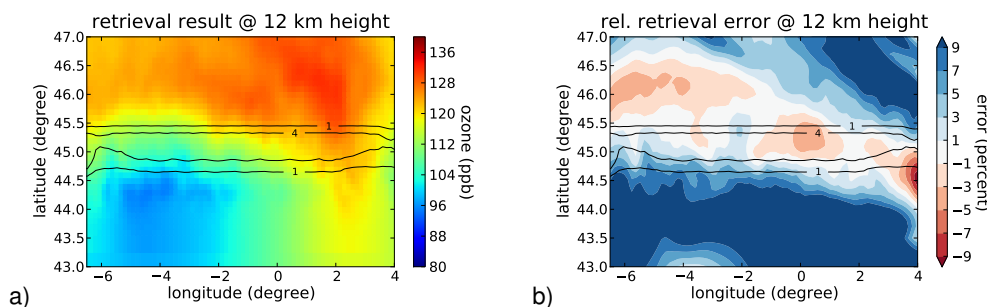


Fig. 5. The retrieval result for a linear flight track. Panel (a) shows the retrieved ozone concentrations. The relative error of the retrieval is plotted in (b). The number of tangent points within the depicted volume elements is overlaid as a contour plot.

indication of the robustness of tomographic retrievals against gradients along the LOS. Figure 6 shows the error induced by noise added to the simulated measurements and the relative error on a vertical cutting plane in north-south direction. The noise error in panel (a) is ≈ 1.5 ppb in the volume with high tangent point density. This corresponds to a relative error of $\approx 1.5\%$, which is roughly in the same order of magnitude as the total relative error of the retrieval result in the same region. Below the aircraft and behind the volume with high tangent point density, the noise error becomes small as the result in this volume is more determined by smoothing due to the regularisation constraint and less by individual measurements. Panel (b) shows how the relative error behaves at different altitudes and is representative for the central portion of the flight. The volume below the aircraft, being devoid of measurements, expectedly exhibits the largest relative error. The volume with high tangent point density shows relative errors which are largely smaller than 5 %, similar to the behaviour at 12 km. Behind the tangent points, there is also an area with relatively low errors, however with decreasing altitude and increasing distance from the aircraft, the quality drops rapidly. Due to the opacity of the air, the measurements are not very sensitive to these regions.

Figure 7 shows the horizontal resolution in longitude and latitude directions within the tangent plane. The achievable resolution in along-track direction in panel (a) is in the or-

der of 30 km and is largely determined by the strength of the horizontal regularisation in this direction. The resolution in across-track direction in panel (b) is about 50 km, which is much smaller than the 100 km that are usual for 1-D retrievals. Vertically, this retrieval exhibits a resolution of ≈ 300 m in the volume covered by tangent points.

Determining the resolution reliably is difficult for this setup as the distribution of significant non-zero elements within the rows of the averaging kernel matrix is rather irregular. In the centre of the region covered by tangent points, the information is well localised with a noticeable smoothing in meridional direction and calculating the FWHM within the tangent plane delivers reasonable results. However, north- and southwards of the well resolved area, the result is strongly influenced by both the true atmospheric state around the closest tangent points and the volume around the retrieved value. With some loss of information, this can be reliably condensed to the FWHM sphere diameter resolution shown in Fig. 8.

The diameter shown in panel (a) is between 30 and 60 km within the region covered by tangent points and about 50 km further south towards the instrument location as the employed channel is not fully transparent in this altitude range. This is a clear improvement over the FWHM sphere diameter for the corresponding pure 1-D retrieval of 200 km. The resolution is naturally best in front of the instrument, where the

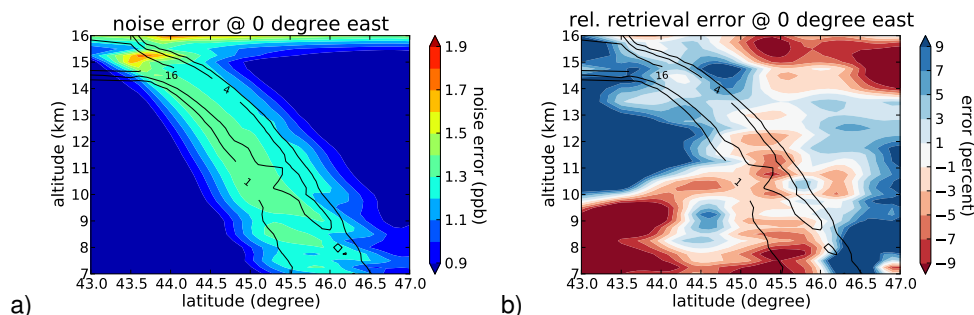


Fig. 6. The retrieval result for a linear flight track. Panel (a) shows the error induced by noise according to the linearised diagnostics. The relative error of the retrieval is plotted in (b). The number of tangent points within the depicted volume elements is overlaid as a contour plot.

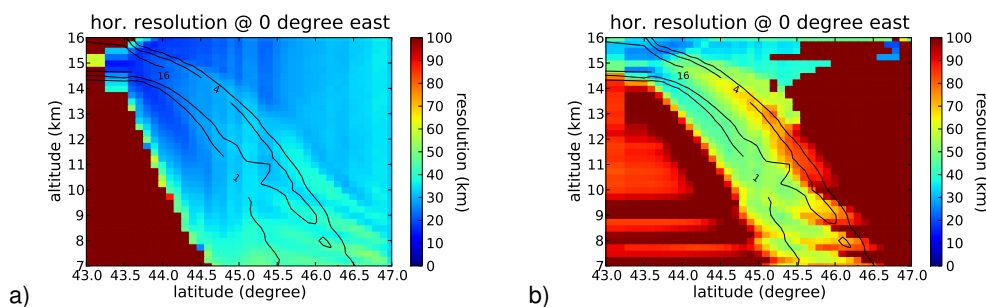


Fig. 7. The horizontal resolution for a linear flight track. Panel (a) gives the horizontal resolution in along track direction (zonal), while (b) gives the horizontal resolution in across track direction (meridional). The number of tangent points within the depicted volume elements is overlaid as a contour plot.

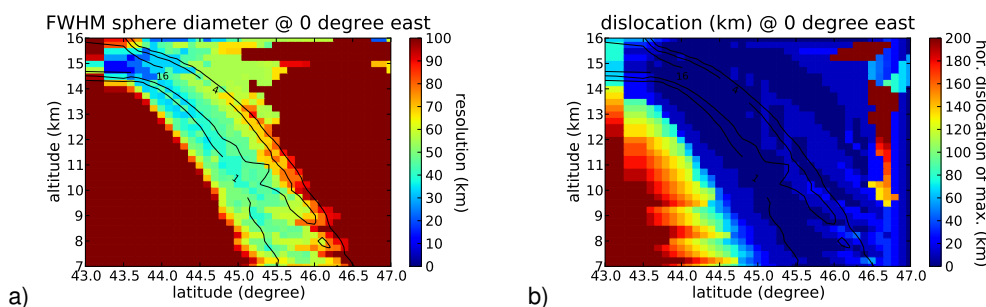


Fig. 8. The resolution and dislocation of the linear flight track. Panel (a) shows the FWHM sphere diameter. Panel (b) shows the distance between the shown data element and the element with the largest contribution in the associated row of the averaging kernel matrix. The number of tangent points within the depicted volume elements is overlaid as a contour plot.

measurement density is highest, and decreases with increasing distance. Please note that in contrast to 1-D retrievals, the two horizontal resolution measures shown in Figs. 7b and 8a generate similar values. The volume of true atmosphere that is averaged for generating one ozone concentration of the tomographic retrieval result is more concentrated than that for the 1-D retrieval. Thus, the tomographic retrieval is more robust against gradients along the LOS. Panel (b) shows that the dislocation is small within the well-resolved region. This indicates that there is sufficient measurement information to

reconstruct a spatial average of the true atmospheric state in this region. Outside this region however, there is an increase in dislocation, which indicates that the result there is increasingly derived by extrapolation from the well-resolved region.

The given setup can provide a good 3-D picture of the atmosphere, albeit only in a limited volume approximately defined by the volume covered by tangent points. Increasing the instrument altitude will significantly enlarge the well-resolved volume within the UTLS region and make this measurement mode more useful. Within this well resolved

volume, a vertical resolution of ≈ 300 m, an along track resolution of ≈ 30 km, and an across track resolution of ≈ 50 km seems feasible.

6 Linear flight path using channels of different optical depth

For tomographic retrievals it is possible to employ channels of different optical depth to increase the spatial resolution. This is similar to nadir sounding, which uses different channels to resolve the profile in the vertical direction. To examine the potential of this technique for 3-D tomography, the setup of Sect. 5 is augmented by three additional channels at 1020 cm^{-1} , 1043.75 cm^{-1} and 1055 cm^{-1} for this experiment only, which exhibit different optical depths. The choice has been rather arbitrary and we are currently looking into ways of choosing an optimal, minimal set of channels for this purpose under realistic constraints for multi-target retrievals. Figure 9 shows the transmissivity for the original and the three added channels at different tangent altitudes. Due to their decreased optical depth, the added channels are to varying degrees more sensitive to the atmosphere in front of the tangent points, where we expect the largest improvement.

The relative error of the retrieval result using these added channels can be seen in Fig. 10a. The retrieval result shows a better agreement with the true state of the atmosphere in a larger area. Especially the low concentrations southwards are better reproduced, as the errors remain acceptably low as far south as 44° N , which gives an additional well-resolved 50 km compared to Fig. 5. The volume towards the north seems not to profit as much as the volume closer to the measuring instrument. For example, the volume to the upper right in panel (a) remains severely overestimated. Also other altitudes than depicted benefit from the added channels: in the whole volume with high tangent point density and also for 50 km towards the aircraft, the relative error is now reduced by about 2 % down to largely 3 %.

Looking at the horizontal resolution in Fig. 10b, one can see a drastic improvement. Panel (b) shows that the volume with a FWHM sphere diameter below 100 km has been significantly increased. The resolution in meridional direction becomes as low as 25 km between the instrument and the tangent point locations and remains below 50 km for another ≈ 60 km. The shape of the averaging kernel matrix rows is better localised than previously, as is also indicated by the dislocation (not depicted), which remains close to zero inside the actually measured volume. Even the resolution in along-track direction (latitudinal) could be improved from the previous ≈ 30 km down to ≈ 20 km (not depicted). The vertical resolution could not be further improved and remains at ≈ 300 m.

Employing channels of different optical depth seems to improve the retrieval result considerably for the linear flight




Fig. 9. The transmissivity of the four employed channels plotted at varying tangent altitudes. For measurements pointing upwards, the tangent altitude behind the measuring instrument has been mirrored at the instrument altitude at 15 km.

path. Some of the information missing due to the limited amount of measured angles can be reconstructed from the added measurements. This opens the possibility to employ also for limb-sounding a host of channels that become optical thick in combination with the usual transparent ones. The horizontal resolution (FWHM sphere diameter) of the test retrieval in across-track direction could be improved from 30 to 60 km in a rather confined region down to 25 km to 50 km in a much larger volume encompassing nearly 300 km in across-track direction.

7 Circular flight path with advection

This study assumed a static atmosphere so far. Now the effect of a time-varying atmosphere on 3-D tomographic retrievals is analysed. While the time required to acquire the measurements is usually negligible for a conventional 1-D retrieval, the atmosphere might change significantly during the time required to fly a tomographic measurement flight, e.g. a full circle with 400 km diameter. The effect becomes the larger, the longer the timespan between two measurements of the same air mass becomes. Thus, a closed, circular flight path is a worst-case scenario, as basically all measurements are affected by the same enclosed air mass.

The aircraft is assumed to fly a circle with ≈ 400 km diameter over the north of France maintaining constantly 15 km altitude, with the centre of the circle being located at 0° E , 46° N . For this setup, 1775 images are taken, summing up to 113 600 individual measurements. This requires the casting of 681 600 individual pencilbeams to capture the effect of horizontal and vertical FOV. A volume element in the central, depicted region of the simulated volume is again $11\text{ km} \times 11\text{ km} \times 0.25\text{ km}$. Except for the changes described here and in the following, the setup of Sect. 5 is used for the following numerical experiments.

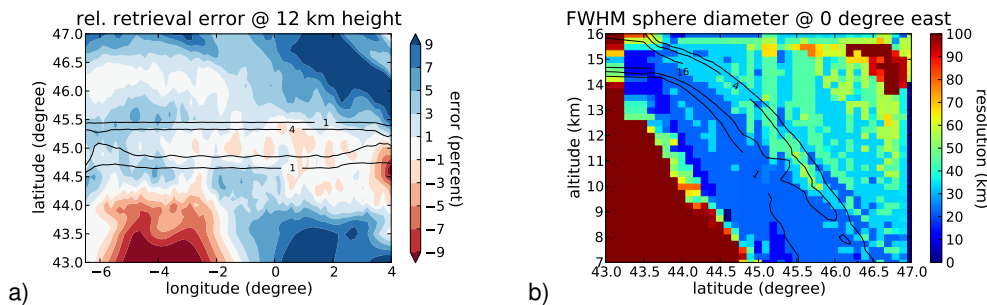


Fig. 10. The relative error of the retrieval and resolution for a linear flight track employing four channels instead of only one. Panel (a) shows the relative error of the retrieval (cf. Fig. 5b). The FWHM sphere diameter is shown in (b) (cf. Fig. 8a). The number of tangent points within the depicted volume elements is overlaid as a contour plot.

Section 7.1 describes a countermeasure against the smoothing in time by incorporating a priori information about advection into the retrieval. Then, the effect of advection on a static 3-D retrieval is examined in detail before the effectiveness of the countermeasure is analysed in Sect. 7.3.

7.1 Incorporating advection

For most species in the UTLS region, transport by advection is much more important than mixing and chemical processes, which play only a minor role for the typical time-scale of a tomographic measurement. To properly simulate the effect of advection, the so far static atmospheric state needs to be altered into a time-varying one. Compared to the high contrast in concentration of many trace gas species (e.g. water vapour), mesoscale winds are rather smooth and comparatively well known. It is therefore advantageous to exploit this a priori information about advection, instead of setting up a full 4-D retrieval with an additional regularisation component for time.

The European Centre for Medium-range Weather Forecast (ECMWF) maintains an operational assimilation and weather forecasting system providing real-time analysis and forecasting of atmospheric composition. In addition, regular re-analyses of the Earth-system are performed. Amongst other information, ECMWF also provides global wind fields, which are wind speeds in meridional, zonal and vertical direction.

For the following numerical studies, the ECMWF ERA40 interim data set (Simmons et al., 2006) is used. It is interpolated onto a grid with 0.25° in longitudinal and latitudinal direction and 250 m in vertical direction. It is thereby horizontally 2 to 2.5 times as coarse as the retrieval grid. Within the original ECMWF data, the vertical wind speed is given as pressure change (“omega”) in Pa s^{-1} . Using the supplied pressure values, this value is converted to a vertical velocity in m s^{-1} . A trajectory model operating directly on pressure levels might be more exact, though. ECMWF ERA40 interim global data sets from 12:00 UTC (Universal Time, Coordinated) and 18:00 UTC of 3 July 2006 were used. Wind

speeds at times in between are linearly interpolated from the two supplied datasets. It is assumed that the first simulated image is taken exactly at 12:00 UTC. Pressure changes were small enough during the measurement period to neglect their effect on air temperature. However, more complex meteorological situations might require the implementation of heating and cooling due to pressure changes.

A simple model uses the wind fields to calculate backward trajectories. The intent is to relate any part of the air volume during the circular flight backward in time to the position it had when the first measurement was taken. The trajectory model uses the wind speeds in the three spatial directions \mathbf{v} as map from \mathbb{R}^4 (time, longitude, latitude and altitude) to \mathbb{R}^3 given in m s^{-1} . Let $\mathbf{x}(t)$ describe the 3-D location of an air parcel at time t . The advection can then be captured by the differential equation

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{v}(t, \mathbf{x}(t)) \quad (10)$$

This equation can be numerically solved backward from a given time t' towards the start of the measurements at t_0 .

A simple fourth-order Runge-Kutta method with fixed step size (e.g. Press et al., 2007) was employed to calculate the trajectories in this study. Varying the step size in the Runge-Kutta scheme, it was found that a step of 10 s gives reliable results, meaning that reducing it further did not provide noticeably different trajectories.

As it is too time-consuming to calculate a backward-trajectory for each access to an atmospheric data point during the retrieval, trajectory lookup tables with pre-calculated trajectories are generated for all grid points of the atmosphere and regular time intervals. During the retrieval, the dis-location of the data point due to advection is derived by linear interpolation in these tables and the atmospheric composition is derived via a 3-D interpolation at the thus discovered original location. In effect, the 4-D state of the retrieval atmosphere is thereby reduced to a simple 3-D state at the time of the first measurement. The effect of the time grid employed for the trajectory lookup tables was estimated and a 300 s grid is used for the following numerical studies.

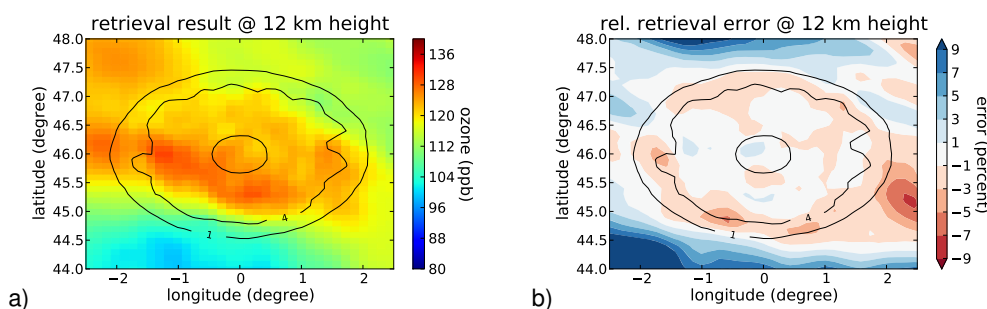


Fig. 11. The retrieval result of a tomographic retrieval using a circular flight path. Panel (a) shows the retrieved ozone values at 12 km height while (b) depicts the relative error in percent compared to the true values. The number of tangent points within the depicted volume elements is overlaid as a contour plot.

7.2 4-D analysis of a circular flight

Before proceeding with numerical studies including advection, it is insightful to first perform a temporal analysis of a static 3-D retrieval, as one can use this to approximate the effect of a time-varying atmosphere on a tomographic retrieval by using the technique described in Sect. 3.2. This method has the advantage, that it does not rely on a priori information about advection.

Figure 2a shows the atmospheric state provided by the GEM-AQ simulation at 12 km altitude and 12:00 UTC and is used as true state for a static 3-D retrieval discussed in this subsection. Using synthetic measurements with added noise generated by assuming this static 3-D atmosphere, a cutting plane through the 3-D retrieval result in 12 km altitude is shown in Fig. 11. Panel (a) shows the absolute values of retrieved ozone concentrations x_f , while panel (b) shows the relative error in relation to the true atmospheric state x_t . The area covered with tangent points (marked as a contour surface indicating the tangent point density per grid point) is well reproduced with an relative error of mostly less than $\pm 3\%$.

To analyse the effect of a time-varying atmosphere on the retrieval result, the 3-D atmosphere is expanded by a fourth dimension representing time. Time is thereby discretised in 150 s steps as a compromise between exactness and memory requirements. The 4-D atmosphere is filled with the 3-D result of the static 3-D retrieval, which is by definition homogeneous in time. It is straightforward to calculate the Jacobian matrix with respect to the 4-D atmosphere and multiply the 3-D gain matrix \mathbf{G} of the baseline setup with the 4-D Jacobian matrix to acquire a 4-D/3-D averaging kernel matrix $\tilde{\mathbf{A}}$ (mapping the 4-D true atmospheric state onto the 3-D retrieval result) that gives quantitative insights into the averaging over time.

Starting with a data element located in the centre of the volume, Fig. 12 shows cutting planes through one row of the 3-D (above) and 4-D/3-D (below) averaging kernel matrices. The upper row depicts cutting planes orthogonal to the three

major axes of one 3-D averaging kernel matrix row, while the lower row of panels essentially blows up the central column of the panel above it, adding time as a new dimension. In panels (a)–(c), the non-zero part is well localised in space around the retrieved data element at 12 km altitude in the central profile. However, panels (d)–(f) reveal that with respect to localisation in time, it is averaged rather evenly over the whole retrieval time. In addition, one can see that the small “bulges” in panel (d) and panel (e) wax and wane according to the position of the GLORIA instrument. The contribution from each single image is necessarily drawn along its LOS and thereby causes these bulges. The 4-D analysis shows that each measurement contributes an average of a not very well localised volume. Only the combination of the contribution of many such measurements increases the weight of the relevant element and lessens the weight of the surrounding elements.

Proceeding westwards, Fig. 13 shows cutting planes for a data element close to the border of the area covered with tangent points at 1.87° W, 46° N, and 12 km altitude. While still comparably well localised, some parabolic artefacts can faintly be seen in panels (a) and (b). These are located close to the points of maximum sensitivity of the LOS of the contributing measurements or at regions, where many contributing measurements are sensitive. Such artefacts arise, when only a few measurements are available that have their maximum sensitivity in the relevant space.

The smoothing in time looks quite different from the smoothing given for the central element in Fig. 12, as it exhibits strong maxima coinciding with the points in time where measurements are taken, the tangent points of which are close to the retrieved data point. Except for the central column of the encircle volume, this usually occurs twice during the flight: once while the instrument is pointing forward, and once while the instrument is pointing backward. This effect is rather typical for all data elements not located on the central profile. Depending on the distance to the centre, these maxima are more or less pronounced.

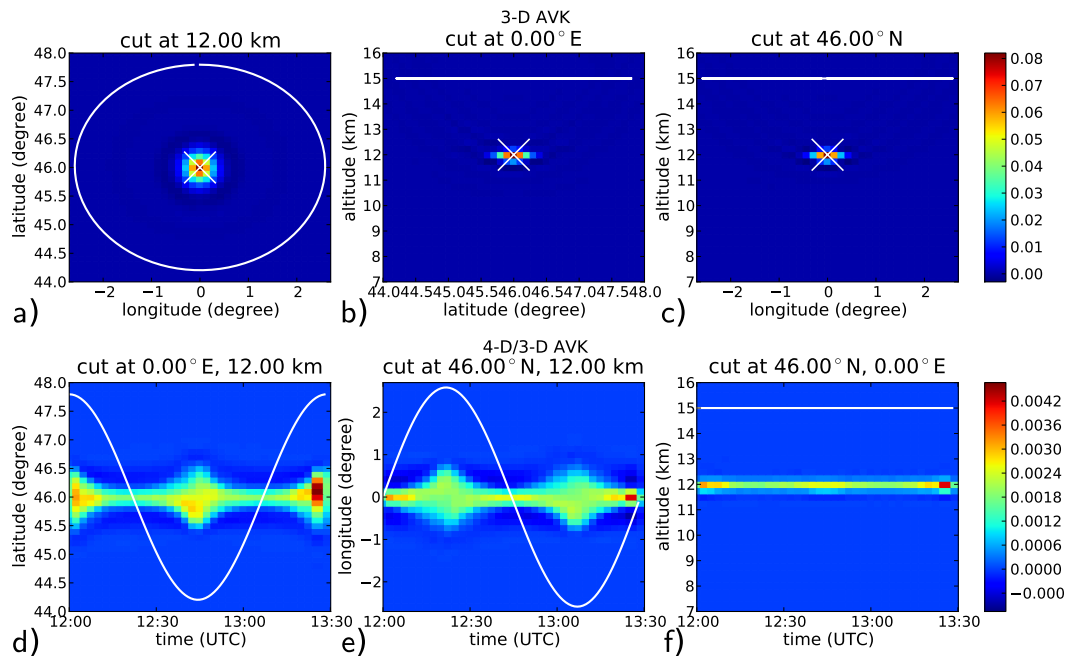


Fig. 12. A comparison between one row of the 3-D averaging kernel matrix and its 4-D/3-D equivalent for a central element. In the upper three panels, three cutting planes through a row of the 3-D averaging kernel matrix are depicted, all belonging to the element at 0° E, 46° N, and 12 km altitude. The projection of the flight track onto the cutting plane is drawn as a white line. The location of the data element is marked by a white cross. For the same data element, the three lower panels depict three cutting planes through the 4-D space described by the row of the 4-D/3-D averaging kernel matrix.

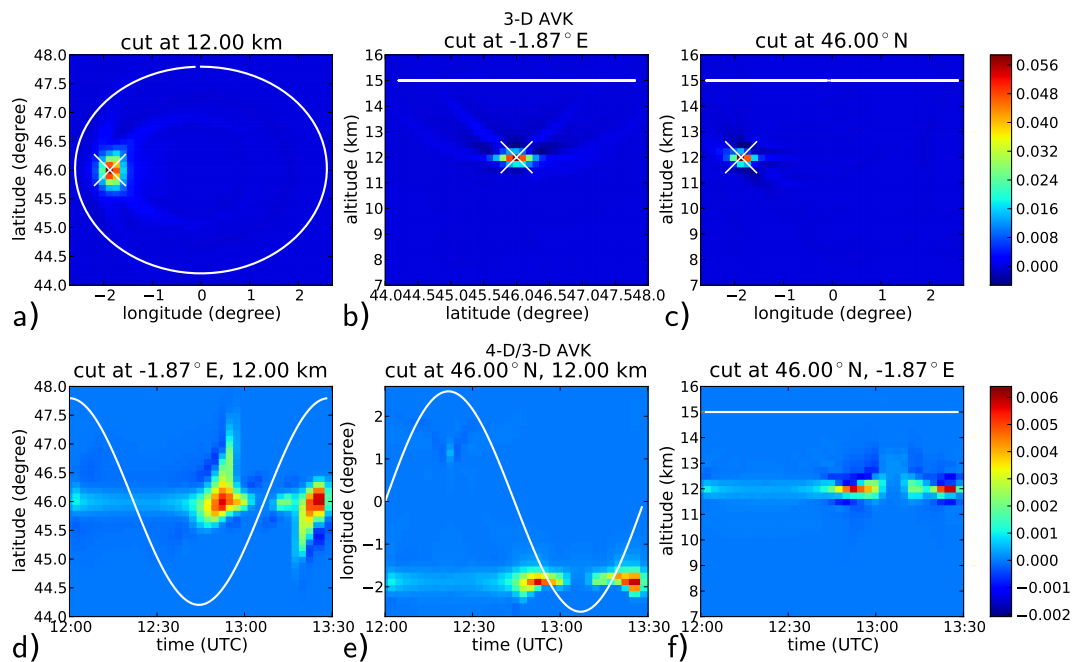


Fig. 13. A comparison between one row of the 3-D averaging kernel matrix and its 4-D/3-D equivalent for a border element. In the upper row of panels three cutting planes through a row of the 3-D averaging kernel matrix are depicted, all belonging to the element at 1.87° W, 46° N, and 12 km altitude. For the same element, the lower row of panels depicts three cutting planes through the 4-D space described by the row of the 4-D/3-D averaging kernel matrix.

Concluding, the smoothing in time emphasises those times when tangent points of measurements are close to a volume element. In addition, the methodology of the 4-D/3-D averaging kernel matrix allows to reliably compare the results of “simple” 3-D tomographic retrievals with corresponding 4-D simulations of atmospheric models or data assimilation systems.

7.3 Trajectory enhanced retrievals

After analysing the theoretical impact of a time-varying atmosphere on the retrieval result, this section analyses the effect of advection on 3-D tomographic retrievals by numerical simulations and offers means for compensation. Three different numerical studies are now described. In all three, the trajectory model uses the unmodified ECMWF wind fields to generate synthetic GLORIA measurements. The impact of the wind on the atmospheric state can be seen in Fig. 2, which shows the atmospheric situation during the first measurement, and Fig. 14, which shows the atmospheric situation during the last measurement of the simulated flight. The barbs in Fig. 14 indicate horizontal wind speeds ranging from 10 m s^{-1} to 20 m s^{-1} . The wind moves the filament $\approx 100 \text{ km}$ towards north-north-east. Additional visible changes are caused by vertical advection. This scenario is basically a worst-case test, as the wind blows nearly parallel to the steepest gradient observed in the ozone field. This causes a maximal displacement of the ozone filament and is therefore well suited to determine the feasibility of the method under worst-case circumstances.

Three different retrieval experiments are executed:

1. The first experiment described in Sect. 7.3.1 employs the trajectory model for generating the simulated measurements as well as for the retrieval. Thereby, it tries to reconstruct the atmospheric state at the time of the first measurement. This setup shall demonstrate how well the retrieval can compensate for advection with perfect knowledge about the wind.
2. The second experiment described in Sect. 7.3.2 uses the trajectory model to generate the simulated measurements but does not use it for the retrieval. Instead, the conventional 3-D retrieval from the previous sections is used. This setup shall examine if a retrieval without compensation is subject to convergence issues and how much the retrieval result is affected by uncompensated advection.
3. The third experiment described in Sect. 7.3.3 uses the trajectory model for generating the simulated measurements and for the retrieval. However, it uses a different, degraded wind field consisting of horizontally homogeneous, averaged wind speeds for the retrieval. This setup shall examine the effect of a potentially erroneous wind field on the retrieval result, just as the ECMWF data may also not reflect the true wind speeds.




Fig. 14. The temporal change caused by advection. This figure depicts the scene at the end of the measurement period. In addition, horizontal wind speed at 12 km altitude has been added using barbs, whereby half a line indicates 5 m s^{-1} and a full line 10 m s^{-1} horizontal wind speed.

Before proceeding to the numerical experiments, the distortion of the spatial distribution of measurements by advection is examined. Figure 15 depicts a vertical cut in north-south direction through the atmospheric state at 12:00. The colouring indicates, how often measurements pass through each volume element taking into account their displacement by advection. One can see clearly that airmasses towards the south are measured more often than those towards the north, especially at lower altitudes. With the prevalent direction of wind, this is expected. Still, the distribution of measurements is quite even so that a good reconstruction should be possible, except for the volume northwards of 47° N and below $\approx 10 \text{ km}$ altitude, where only few measurements pass through the volume elements.

7.3.1 Perfect a priori knowledge

The first experiment takes advection into account for generating the measurements and also uses perfect knowledge about the advection in the retrieval process. As the advection moves air from southwards of the flight track into the circle described by the instrument, it is expected that these air masses are better resolved, possibly at the cost of reduced resolution for air lying towards the north within the circle compared to the retrieval experiment devoid of advection in Fig. 11.

The results of such a retrieval, where the forward calculation and the retrieval were both calculated with the same wind fields is presented in Fig. 16. Comparing Fig. 16b to Fig. 11b, the southern boundary of the ozone filament is much better reproduced. The low ozone concentrations towards the south are quantitatively well given with an relative error below $\pm 2\%$. This is quite an improvement compared to Fig. 11b of the static setup, where southward of 44.5° N the error surpasses 9% . In contrast, outside the northern half of




Fig. 15. The spatial distribution of measurements. The number of measurements sampling each volume element is shown.

the circle described by the instrument, the errors are slightly increased with a slight tendency to overestimate true ozone concentrations. In Fig. 17, a vertical cut through the volume is depicted to show the behaviour at other altitudes. Panel (a) shows the error induced by the noise added to the measurements according to the linearised diagnostics. The error is large, where many tangent points are located and small, where no measurements are present. The smoothing influence of the a priori constraint reduces the influence of individual measurements and thereby the influence of noise. The vertical regularisation constraint is not very strong, so the influence of noise increases towards the top column. Panel (b) shows the relative error of the retrieval result compared to the true values. The reconstruction is best around 12 km altitude as this plane coincides well with the volume containing the most tangent points of measurements. But the reconstruction quality is still very good for several kilometres above and below. Only below ≈ 9 km, the relative error surpasses 5 % in a larger volume towards the north. This is expected, because this volume is quickly blown away without having been measured by a noticeable number of images as depicted in Fig. 15.

The horizontal resolution (FWHM sphere diameter) is shown in Fig. 18a. The well resolved area with a resolution of ≈ 30 km is non-circular and enlarged towards the south, the direction from where the wind blows. The same can be said for the vertical resolution in Fig. 18b. Except for the different distribution of well-resolved elements, the achievable resolution seems to be rather unchanged compared to simulations without advection. The more irregular shape of the rows of the averaging kernel matrix causes the resolution values to be slightly noisier, but the minimum value is nearly unchanged. Seemingly, the GLORIA instrument takes a sufficient amount of measurements to resolve the additional volume of air passing through the circle described by the instrument without a significant degradation of the achievable resolution.

7.3.2 No a priori knowledge

While the previous paragraph discussed the best case, perfect knowledge about advection, the worst case is presented in this section: no knowledge about advection. Consequently, this experiment uses measurements that were generated using the ECMWF wind field, but (wrongly) assumes a static atmosphere during the retrieval.

Slightly surprisingly, this setup did converge to a physically meaningful atmospheric state within the usual number of iterations. This result is shown in Fig. 19a. It looks quite different from the result of the run incorporating perfect wind speed knowledge in Fig. 16. The filament is moved towards the north-north-west, the same direction the wind blows to. Correspondingly, the relative error plot (comparing with the state at 12:00 UTC) in Fig. 16b shows large errors in the south-south-east of more than 10 %. Vertical cuts are not depicted, as the noise error behaves virtually identical to Fig. 17 and as the relative error is similar to 12 km also at other altitudes.

As discussed in Sect. 7.2, the 3-D retrieval without trajectory model compensation averages the true atmospheric state over time. So comparing the result with the state at 12:00 UTC might not be the best frame of reference. Instead a uniform average of the true atmosphere over time might be better suited as is shown in Fig. 20. The averaged atmosphere more closely resembles the retrieval result in Fig. 19a within the volume covered by tangent points and also towards the south, even though Sect. 7.2 showed that the average is non-uniform. This is also true for other altitudes down to ≈ 9 km (not depicted). This implies that the 3-D retrieval without any compensation of advection still generates reliable results, even though they provide only a time-averaged state.

The usual resolution measures cannot be meaningfully employed here, as they ignore advection. One could perform a 4-D analysis like in Sect. 7.2, but a more interesting measure would be the spatial instead of the temporal smoothing. Under the assumption that the only change in atmospheric state is caused by advection, it is possible to express the temporal smoothing of the 4-D/3-D averaging kernel matrix in terms of spatial smoothing. The temporal dimension of this matrix can be reduced using the trajectory model by projecting each element to the location it was located at the beginning of the measurement period. The temporal smoothing of the retrieval is thereby effectively translated into a spatial smoothing. The FWHM sphere diameter is determined and presented in the left panel of Fig. 21. The resolution varies from ≈ 40 km up to ≈ 140 km within the usually well resolved region. The first image is taken in the north following the circle in clockwise direction. The worst resolution is given around the start and end point of the instrument's flight path. Here, images from the beginning and end of the flight are combined, which are taken approximately 90 min apart. This involves a maximum displacement of measured air mass and consequently a bad spatial resolution. On the

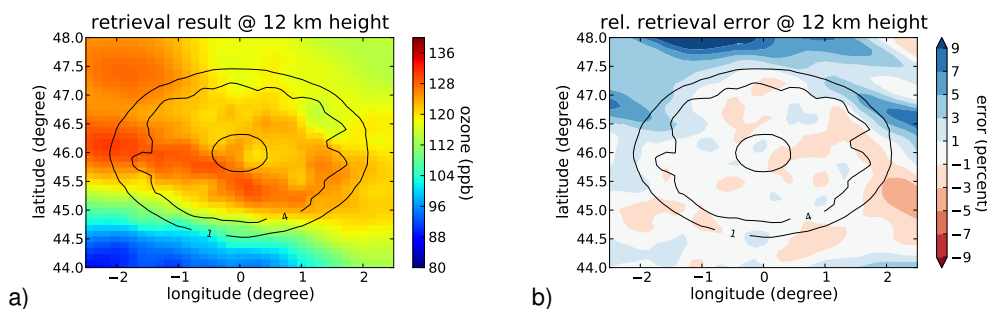


Fig. 16. The retrieval result of a run with advection and perfect a priori knowledge thereof. Panel (a) shows the retrieved ozone values at 12 km height while (b) depicts the relative error in percent compared to the true values. The number of tangent points within the depicted volume elements is overlaid as a contour plot.

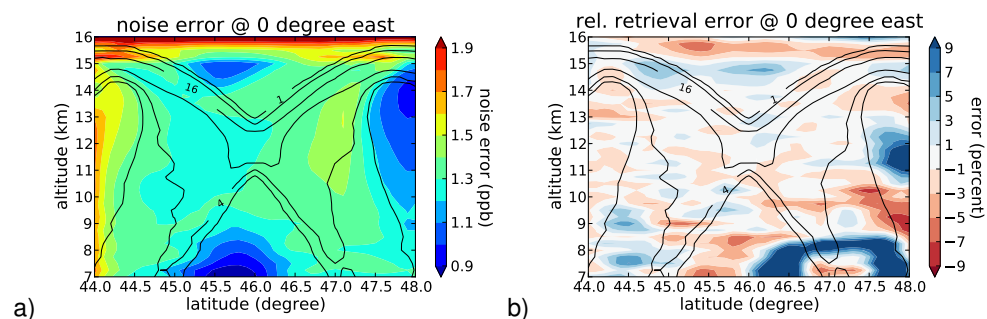


Fig. 17. The retrieval result of a run with advection and perfect a priori knowledge thereof. Panel (a) shows the error induced by noise according to the linearised diagnostics at 0° E while (b) depicts the relative error in percent compared to the true values. The number of tangent points within the depicted volume elements is overlaid as a contour plot.

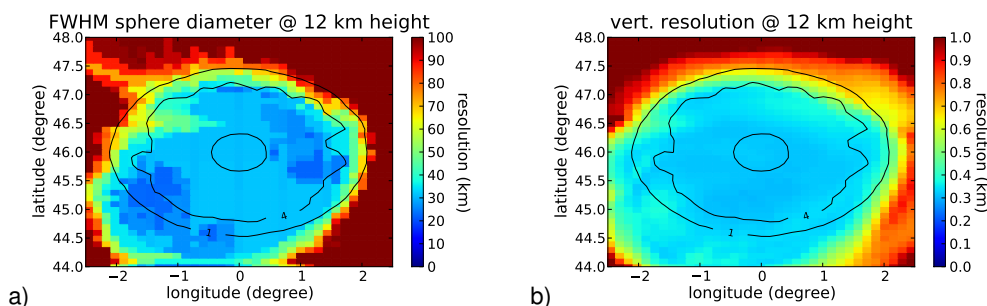


Fig. 18. Horizontal and vertical resolution of a retrieval with perfect wind speed knowledge. Panel (a) shows the FWHM sphere diameter and (b) shows the vertical resolution. The number of tangent points within the volume element is overlaid as a contour plot.

other hand, in the southern part of the circle, the forward and backward looking measurements are taken only 30 minutes apart, leading to a much smaller spatial smoothing.

Panel (b) of Fig. 21 shows the distance between the largest element of the row of the averaging kernel matrix and the data element this row represents. This can be interpreted as a measure for the dislocation of the retrieved element with respect to its location at 12:00 UTC. A value of 0 km means that the retrieved element was at its given location at 12:00 UTC, while a value of 100 km means that the retrieved

air parcel was about 100 km away at 12:00 UTC. The right panel shows how the parts of the volume measured first are barely displaced. But the displacement grows larger as more time passes. The nearly linear increase in displacement indicates that the airmass being currently in front of the instrument is effectively measured.

Combining both figures, one can state that except for the northern part of the volume, where measurements from the very beginning and end of the flight track are combined, the air located between the tangent point and the instrument at

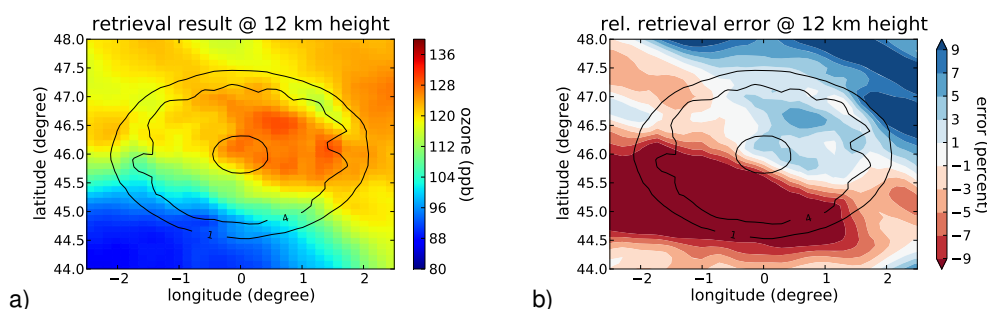


Fig. 19. The retrieval result of a run with advection and no compensation during the retrieval. Panel (a) shows the retrieved ozone values at 12 km height while (b) depicts the relative error in percent compared to the true values. The number of tangent points within the depicted volume elements is overlaid as a contour plot.

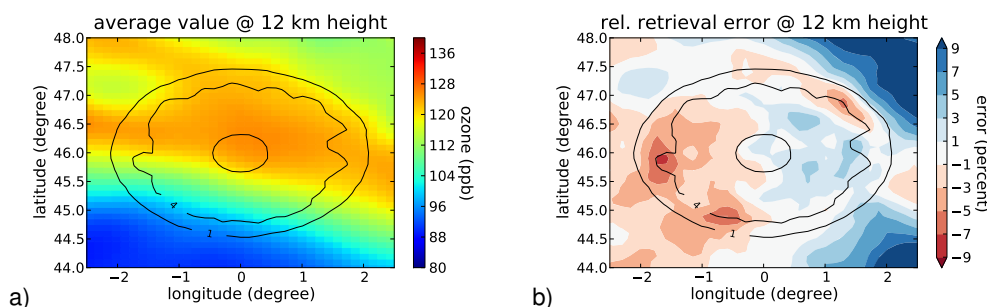


Fig. 20. The atmospheric state averaged over time in panel (a) and the relative error of the retrieval result without wind compensation in relation to the averaged state in (b). The number of tangent points within the volume element is overlaid as a contour plot.

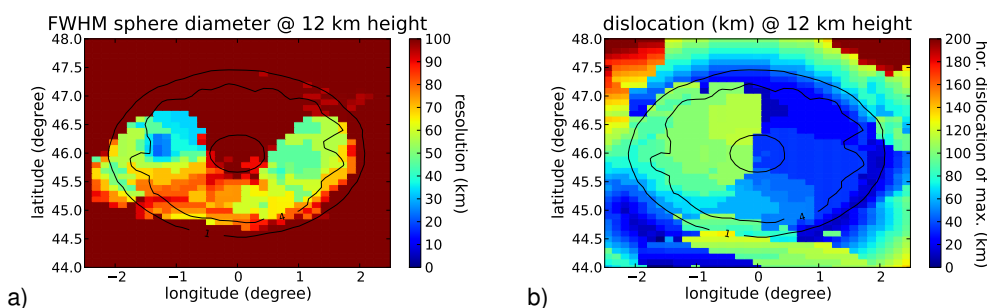


Fig. 21. The resolution of the retrieval without a priori wind speed information related to the atmospheric state prevalent during the first measurement. Panel (a) shows the FWHM sphere diameter. Panel (b) shows the distance between the shown data element and the element with the largest contribution in the associated row of the averaging kernel matrix. The number of tangent points within the depicted volume elements is overlaid as a contour plot.

the time of measurement is retrieved. While the resolution in the centre of the volume is quite bad due to the large averaging in time (and correspondingly in space), the outer parts of the volume represent mostly the state of affairs at the time when the instrument is closest.

7.3.3 Imperfect a priori knowledge

The third numerical study tries to address the problem of imperfect wind speed knowledge. Just as for the previous two experiments, the trajectory model is employed to generate

the simulated measurements. But for the retrieval, an averaged, horizontally homogeneous wind field is used.

The inaccuracies of the provided wind speeds and their effect onto trajectory models were examined by Pickering et al. (1994, 1996) in the southern Atlantic. He found that the ratio of average absolute difference between dropsonde measurements and ECMWF winds to mean of absolute values was up to $0.87 \text{ m s}^{-1} / \text{m s}^{-1}$. Problematic were especially regions with large horizontal or vertical wind shears, which may not be perfectly located by the generating model. With respect

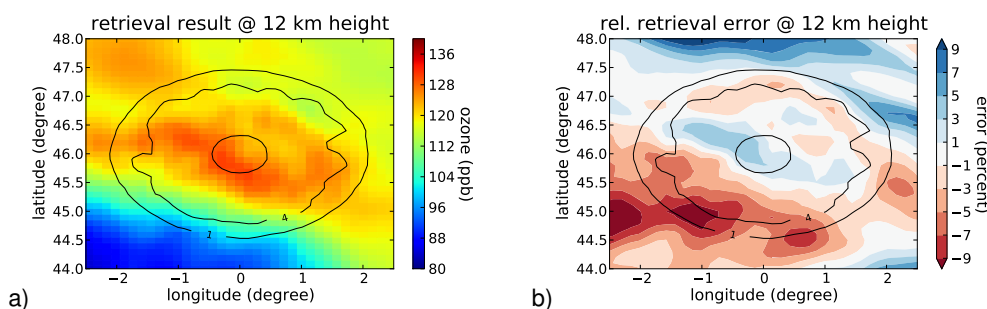


Fig. 22. The retrieval result of a run with advection and an averaged a priori wind speed knowledge. Panel (a) shows the retrieved ozone values at 12 km height while (b) depicts the relative error in percent compared to the true values. The number of tangent points within depicted volume elements is overlaid as a contour plot.

to deriving trajectories from this input data, he suggests to perform a cluster of runs with different closely spaced starting points. Should the resulting trajectories perform similarly, confidence can be placed in the results. While the ECMWF wind speed data has naturally been quite improved upon in the past 15 yr, one should not rely solely on the data provided by ECMWF or similar institutions. By combining GLORIA measurements with simultaneous dropsonde wind speed measurements around the volume of interest, it would be possible to validate the provided model wind speeds. One might even augment the ECMWF data using data assimilation techniques and the measured wind speeds.

While the wind fields provided by ECMWF may be subject to errors, on average, they provide a good approximation to reality. For example in a more recent analysis, Rickard and Lunnon (2001) examined the accuracy of wind predictions by the Met Office with measurements by commercial aircraft and found that the discrepancy between the predicted wind speeds and the measured ones averaged over the flight duration is only $\approx 4\%$.

Consequently, in this case study, the wind speeds used for the retrieval are generated by using horizontally homogeneous wind fields, whereby the value at each vertical level was generated by horizontally averaging the wind speed from 8° W to 8° E and from 37° N to 52° N. In 12 km altitude this gives a constant wind speed of 10 m s^{-1} in zonal and 8 m s^{-1} in meridional direction. The averaged meridional wind speed is thereby quite larger than the true wind speed (see Fig. 14) in the centre of the circle described by the instrument.

Comparing this averaged speed with the true speed used for generating the simulated measurements in Fig. 14, it becomes evident that the averaged meridional speed is quite a bit too large. Within the centre of the circle described by the instrument the true meridional wind speed is actually close to zero.

The result of this experiment is presented in Fig. 22. The result is an improvement over the simulation run without a priori knowledge of the wind field: the filament is placed closer to the true location and the error plot in Fig. 22b shows

errors below 8% . This is also true for the relative error at other altitudes. Except for the airmass to the bottom north, the error stays below 9% in the altitude range between 9 and 15 km within the volume covered by tangent points.

The same resolution and displacement figures as for the experiment without a priori wind speed knowledge are shown in Fig. 23 for comparison. The worst horizontal resolution within the volume covered by tangent points could be reduced from ≈ 140 km down to only ≈ 70 km. The part of the volume with a resolution of ≈ 40 km is greatly increased and encompasses nearly the full area covered by tangent points except for the middle and north-western part. As before, the volume closest to the start and end of the flight track is still most affected. The right panel of Fig. 23 shows the displacement of the retrieved data points. Again, the further the flight progresses, the larger the displacement becomes, even though the maximum displacement has been roughly halved.

This suggests that even inaccurate information about the prevalent wind can improve the retrieval result as long as it captures the main features of the flow. However, this needs to be examined on a case by case basis. It is evident that the quality of knowledge about the advection during the flight is essential for the achievable quality of the retrieval result. Even though plausible results can be generated without any wind speed a priori information, a maximum horizontal resolution of down to 30 km can only be achieved using accurate wind fields.

8 Conclusions

This paper extended the work of Ungermann et al. (2010b) by addressing the important issues of non-closed flight-paths and the effect of advection.

The forward model was supplemented by an adjoint model to calculate the Jacobian matrix. This significantly reduces the computation time as the cost for calculating the Jacobian matrix is now only a constant multiple of a single execution of the forward model, which is extremely fast due to the employed emissivity growth approximation. This allows

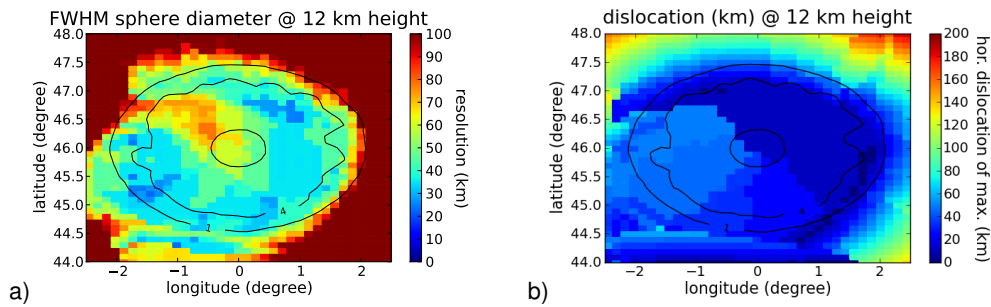


Fig. 23. The resolution of the retrieval run with advection and an averaged a priori knowledge. Panel (a) shows the FWHM sphere diameter. Panel (b) shows the distance between the shown data element and the element with the largest contribution in the associated row of the averaging kernel matrix. The number of tangent points within the depicted volume elements is overlaid as a contour plot.




Fig. 24. A comparison of computation time for the calculation of one Jacobian matrix using different methods. Three different problems representing single-target 1-D, 2-D and 3-D problems are used as example. The method “FD” uses finite differences without any optimisation. The method “FD+TRK” exploits the sparsity of the Jacobian matrix to calculate only the non-zero entries with finite differences. The method “ADM” uses algorithmic differentiation. No parallelisation was employed.

to cheaply extend the dimensions employed in the diagnostics, that is from 1-D to 2-D for conventional retrievals and from 3-D to 4-D for tomographic retrievals. This method can be used to correctly fold higher dimensional model data with an averaging kernel matrix of similar dimensionality to compare the model data with retrieval results.

It was examined by means of simulations what performance can be expected from the tomographic evaluation of linear flight paths in combination with a panning instrument swivelling the LOS between 45° and 135° under ideal circumstances. It was found that using only measurements made from a linear flight track in an extremely difficult situation with strong gradients in the exemplary target quantity ozone, the tomographic retrieval approach can reconstruct the true atmosphere in a limited volume roughly defined by the tangent point locations. The achievable horizontal resolution (FWHM sphere diameter, i.e. the diameter of the

smallest sphere containing all elements of an averaging kernel matrix row larger than half the largest value) was in the order of 30 to 60 km. Further, within the well-resolved volume, no indication of over- or underestimation of the target quantity compared to the true model data was found, which might have been caused by the strong gradient of the target quantity along the LOS. In contrast, a conventional 1-D retrieval executed under the same idealised circumstances exhibits a larger horizontal resolution of about 200 km (FWHM sphere diameter), has a much worse vertical resolution (if a proper regularisation strength to suppress oscillations in the result is employed) and is prone to bias due to the gradients along the LOS.

In a further step, the effect of using additional channels with a different optical depth was examined. It was shown that this further improves the retrieval quality of the simulations especially between the measuring instrument and the tangent point locations and overall improves the horizontal resolution. This opens the opportunity to derive information from opaque channels usually not used for limb retrievals and thereby further increase the spatial resolution of the retrieval.

The last part of this paper considered the effect of advection on tomographic retrievals using circular flight paths. Retrievals using data taken from linear flight paths are not affected as much by advection as those employing circular flight paths due to the shorter time frame between the measurements of the same volume. Thus, the circular flight path has been chosen as a worst-case scenario for examining the effect of advection. An approach was presented that uses wind speed information taken from global models such as ECMWF to correct for the advection during the retrieval. When the wind speed information used in the retrieval is of similar quality as the true wind speeds, the retrieval result is of similar quality as a retrieval for a static atmosphere. A linear analysis showed how the static 3-D retrieval produces a weighted average of the true 4-D atmosphere. This weighted average could be approximated by an uniform average with some success, but this disregards that the retrieval favours the atmospheric state at times, when measurements

highly sensitive to that volume are taken. Lastly, it was shown that also flawed wind speed information could improve the retrieval results to some extent. In this case the horizontal resolution (FWHM sphere diameter) still stays well below the typical horizontal resolution for conventional 1-D retrievals.

In conclusion, the GLORIA instrument should be able to derive highly resolved 3-D volumes of trace gas abundances in the UTLS, enabling a better understanding of meso-scale processes and structures in this region. The numerical studies show a vertical resolution in the order of 300 m and a horizontal resolution in the order of 30 km, for both linear flight paths and circular flight paths in the presence of advection.

Appendix A

Algorithmic differentiation and adjoints

In this section, it is shortly described how the Jacobian matrix of the JURASSIC2 forward model can be efficiently calculated using algorithmic differentiation techniques. A comprehensive description of the algorithms employed in the algorithmic differentiation and their implementation can be found in the technical report of Leppkes and Naumann (2011). The application of these tools to JURASSIC2 is described by Lotz et al. (2011).

For many scientific numerical applications, the tangent-linear model (TLM) or the adjoint model (ADM) are of interest. The TLM $\mathbf{M}_{\text{TLM}}^x : \mathbb{R}^n \mapsto \mathbb{R}^m$ of a real-valued function $\mathbf{F} : \mathbb{R}^n \mapsto \mathbb{R}^m$ at location $\mathbf{x} \in \mathbb{R}^n$ is

$$\mathbf{M}_{\text{TLM}}^x(\mathbf{y}) = \mathbf{F}'(\mathbf{x})\mathbf{y}, \quad \mathbf{y} \in \mathbb{R}^n, \quad (\text{A1})$$

and the ADM $\mathbf{M}_{\text{ADM}}^x : \mathbb{R}^m \mapsto \mathbb{R}^n$ is

$$\mathbf{M}_{\text{ADM}}^x(\mathbf{y}) = \mathbf{F}^T(\mathbf{x})\mathbf{y}, \quad \mathbf{y} \in \mathbb{R}^m. \quad (\text{A2})$$

Efficient implementations of these functions do not explicitly use the Jacobian matrix. However with either function available, it is trivial to assemble the Jacobian matrix itself by using Cartesian basis vectors for \mathbf{y} , which computes a row or column of the Jacobian matrix, respectively. If f is a function mapping to only a single parameter, the ADM generates the full Jacobian matrix of f , i.e. the gradient.

Typical applications for these functions are sensitivity analyses in the context of meteorological or oceanographic models. The Jacobian matrix shows the sensitivity of output parameters of the model to all input parameters. Consequently, the TLM can be used to examine the effect of a disturbance of one input parameter (i.e. to calculate a column of the Jacobian matrix), while the ADM is used to analyse the cause of anomalies in the result (i.e. to calculate a row of the Jacobian matrix) (Giering and Kaminski, 1998).

Quite often it is possible to manually provide efficient implementations of the TLM or ADM for a given task, but this

manual approach requires elaborate adjustment of those programs whenever the main function \mathbf{F} is modified. A different approach is to use algorithmic differentiation. JURASSIC2 employs a method that requires only small changes to the source code to automatically generate functions that implement the TLM or ADM based on C++ operator overloading. The normally used floating point types are exchanged by a more complex data type that records or “tapes” the data flow of the forward model during its execution (Leppkes and Naumann, 2011). This tape can be used afterwards to evaluate the TLM or ADM in an interpretation step. Naturally, the construction of the tape slows down the execution of \mathbf{F} and consumes additional memory per taped operation, yielding problems when dealing with large functions \mathbf{F} . To overcome this issue, so called checkpointing can be applied to stay within the available memory bound.

For example, the cost function J defined in Eq. (1) has only a single output parameter, so a single execution of its ADM computes the full gradient \mathbf{J}' . Using our overloading approach, this requires only a single execution of an instrumented J , which automatically constructs and interprets the tape. Depending on the code to be differentiated, the theoretical minimum for the cost of this operation is 2 to 5 times the cost of the execution of an unmodified J . A factor of 5 can often be achieved in practise (e.g. Griewank and Walther, 2008, pages 83–88). This makes this approach much more efficient for large vectors \mathbf{x} of length $n \gg 5$ than the method of finite differences, which requires $n + 1$ separate executions of the cost function J .

The efficient calculation of Jacobian matrices of the JURASSIC2 forward model \mathbf{F} is a direct application of an ADM at the appropriate place. Splitting the forward model \mathbf{F} into an instrument model \mathbf{H} and the computation of individual pencilbeams \mathbf{P} gives

$$\mathbf{F}(\mathbf{x}) = \mathbf{H}(\mathbf{P}(\mathbf{x})), \quad (\text{A3})$$

which implies for the Jacobian matrix \mathbf{F}' according to the chain rule:

$$\mathbf{F}'(\mathbf{x}) = \mathbf{H}'(\mathbf{P}(\mathbf{x})) \cdot \mathbf{P}'(\mathbf{x}). \quad (\text{A4})$$

Ignoring for simplicity's sake the instrument parameters gain, offset and elevation angle, the instrument model \mathbf{H} models the FOV and can be represented by a matrix \mathbf{H} . Thus, Eq. (A4) simplifies to

$$\mathbf{F}'(\mathbf{x}) = \mathbf{H} \cdot \mathbf{P}'(\mathbf{x}). \quad (\text{A5})$$

The function \mathbf{P} consists of the combined pencilbeam calculations. As the pencilbeams are mutually independent from one another, the function \mathbf{P} can be split into individual functions $p_i : \mathbb{R}^n \mapsto \mathbb{R}$, each mapping the full atmospheric state onto a single radiance value. The p_i are again functions with only a single output parameter, so an ADM associated with p_i can directly calculate the gradient $\mathbf{p}'_i(\mathbf{x})$. Evaluating a single pencilbeam in this manner is such a small problem that

the memory consumption of the taping process is negligible. The gradients $\mathbf{p}'(\mathbf{x})$ are then assembled to construct the Jacobian matrix $\mathbf{P}'(\mathbf{x})$, which in turn can be multiplied with \mathbf{H} to calculate $\mathbf{F}'(\mathbf{x})$. For efficiency, $\mathbf{P}'(\mathbf{x})$ is not actually constructed, but each gradient is individually multiplied with \mathbf{H} to construct $\mathbf{F}'(\mathbf{x})$ on the fly. This enables the assembly of the full Jacobian matrix \mathbf{F}' with effectively a single instrumented evaluation of the forward model \mathbf{F} .

The implementation is slightly more complicated than presented to include derivatives with respect to instrument parameters like elevation angle or more complicated constraints, e.g. with respect to hydrostatic equilibrium.

The efficiency of calculating the Jacobian with finite difference (“FD”), finite differences with tracking (“FD+TRK”; calculating only non-zero elements of the Jacobian matrix) and algorithmic differentiation (“AD”) is compared in Fig. 24. The time required to calculate the Jacobian matrix is depicted for three different use-cases and three different methods. On the left is a very simple 1-D test case. Here, half of the Jacobian matrix consists of zeros, so the computation with tracking brings a speed-up of about two compared to pure finite differences (“FD”). Even for this small test case, the adjoint method of computation (“ADM”) is much faster. For the larger 2-D and 3-D test cases, the speed-up is even more pronounced. While finite differences with tracking is ≈ 100 times faster than pure finite differences, the algorithmic differentiation provides another factor of ≈ 100 , being therefore $\approx 10\,000$ times faster than pure finite differences for the two examined tomographic problems on the right. As the exemplary calculations show, the time needed to construct the full Jacobian matrix \mathbf{F}' takes about five times as long as a single evaluation of the forward model \mathbf{F} for the tomographic test cases, thereby being close to the theoretical optimum.

Acknowledgements. We sincerely thank A. Dudhia, Uni. Oxf., for providing us with the reference forward model (RFM). We also thank the many unmentioned members of the GLORIA team involved in designing and building this instrument.

Edited by: K. Strong

References

- 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.
- Curtis, A. R.: Discussion of “A statistical model for water vapour absorption” by: R. M. Goody, *Q. J. Roy. Meteor. Soc.*, 78, 638–640, 1952.
- Dudhia, A., Morris, P. E., and Wells, R. J.: Fast monochromatic radiative transfer calculations for limb sounding, *J. Quant. Spectrosc. Ra.*, 74, 745–756, doi:10.1016/S0022-4073(01)00285-0, 2002.
- Eckermann, S. D., Hoffmann, L., Höpfner, M., Wu, D. L., and Alexander, M. J.: Antarctic NAT PSC belt of June 2003: observational validation of the mountain wave seeding hypothesis, *Geophys. Res. Lett.*, 36, L02807, doi:10.1029/2008GL036629, 2009.
- 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.
- 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., Heglin, M. I., and Birner, T.: The extra tropical upper troposphere and lower stratosphere, *Rev. Geophys.*, 49, RG3003, doi:10.1029/2011RG000355, 2011.
- Giering, R. and Kaminski, T.: Recipes for adjoint code construction, *ACM T. Math. Software*, 24, 437–474, doi:10.1145/293686.293695, 1998.
- Godson, W. L.: The evaluation of infra-red radiative fluxes due to atmospheric water vapour, *Q. J. Roy. Meteor. 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.
- Griewank, A. and Walther, A.: Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation, no. 105 in Other Titles in Applied Mathematics, SIAM, Philadelphia, PA, 2nd edn., 2008.
- Groß, J.-U. Günther, G., Müller, R., Konopka, P., Bausch, S., Schlager, H., Voigt, C. Volk, C. M., and Toon, G. C.: Simulation of denitrification and ozone loss for the arctic winter 2002/2003, *Atmos. Chem. Phys.*, 5, 1437–1448, doi:10.5194/acp-5-1437-2005, 2005.
- Hoffmann, L. and Alexander, M. J.: Retrieval of stratospheric temperatures from Atmospheric Infrared Sounder radiance measurements for gravity wave studies, *J. Geophys. Res.*, 114, D07105, doi:10.1029/2008JD011241, 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.
- Hoffmann, L., Weigel, K., Spang, R., Schroeder, S., Arndt, K., Lehmann, C., Kaufmann, M., Ern, M., Preusse, P., Stroh, F., and Riese, M.: CRISTA-NF measurements of water vapor during the SCOUT-O₃ Tropical Aircraft Campaign, *Adv. Space Res.*, 43, 74–81, doi:10.1016/j.asr.2008.03.018, 2009.
- Konopka, P., Groß, J.-U., Plöger, F., and Müller, R.: Annual cycle of horizontal in-mixing into the lower tropical stratosphere, *J. Geophys. Res.*, 114, 0148–0227, doi:10.1029/2009JD011955, 2009.
- Lane, T. P., Reeder, M. J., and Clark, T. L.: Numerical modeling of gravity wave generation by deep tropical convection, *J. Atmos. Sci.*, 58, 1249–1274, doi:10.1175/1520-0469(2001)058<1249:NMOGWG>2.0.CO;2, 2001.
- Leppkes, K. and Naumann, U.: DCO/C – Efficient Derivative Code by Overloading in C++, Tech. Rep. AIB-2011-06, RWTH Aachen, 2011.

- 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.
- Lotz, J.: Analyse adjungierten-basierter Optimierungsverfahren unter Verwendung algorithmischen Differenzierens, Master's thesis, RWTH Aachen, 2010.
- Lotz, J., Naumann, U., and Ungermann, J.: Efficient Discrete Adjoint Computation in a Spectral Simulation Code, Tech. Rep. AIB-2011-05, RWTH Aachen and Forschungszentrum Jülich, 2011.
- Natterer, F.: The Mathematics of Computerized Tomography, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001.
- Piani, C., Durran, D., Alexander, M. J., and Holton, J. R.: A numerical study of three-dimensional gravity waves triggered by deep tropical convection and their role in the dynamics of the QBO, *J. Atmos. Sci.*, 57, 3689–3702, doi:10.1175/1520-0469(2000)057<3689:ANSOTD>2.0.CO;2, 2000.
- Pickering, K. E., Thompson, A. M., McNamara, D. P., and Schoeberl, M. R.: An intercomparison of isentropic trajectories over the South Atlantic, *Mon. Weather Rev.*, 122, 864–879, doi:10.1175/1520-0493(1994)122<0864:AIOITO>2.0.CO;2, 1994.
- Pickering, K. E., Thompson, A. M., McNamara, D. P., Schoeberl, M. R., Fuelberg, H. F., Loring Jr., R. O., Watson, M. V., Fakhruzzaman, K., and Bachmeier, A. S.: TRACE a trajectory intercomparison – 1. Effects of different input analyses, *J. Geophys. Res.*, 101, 23909–23925, doi:10.1029/95JD03297, 1996.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P.: Numerical Recipes, The Art of Scientific Computing, vol. 1, Cambridge University Press, 3. edn., 2007.
- Purser, R. J. and Huang, H.-L.: Estimating Effective Data Density in a Satellite Retrieval or an Objective Analysis, *J. Appl. Met.*, 32, 1092–1107, doi:10.1175/1520-0450(1993)032<1092:EEDDIA>2.0.CO;2, 1993.
- 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.
- Rickard, G. J. and Lunnon, R. W.: The Met Office upper air winds: prediction and verification in the context of commercial aviation data, *Meteorol. Appl.*, 8, 351–360, doi:10.1017/S1350482701003115, 2001.
- 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, World Scientific, 2000.
- Simmons, A., Uppala, S., Dee, D., and Kobayashi, S.: ERA-interim: new ECMWF reanalysis products from 1989 onwards, *ECMWF Newsletter*, 110, 25–35, 2006.
- Solomon, S. C., Hays, P. B., and Abreu, V. J.: Tomographic inversion of satellite photometry, *Appl. Optica*, 23, 3409–3414, doi:10.1364/AO.23.003409, 1984.
- 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.
- Struzewska, J. and Kaminski, J. W.: Formation and transport of photooxidants over Europe during the July 2006 heat wave – observations and GEM-AQ model simulations, *Atmos. Chem. Phys.*, 8, 721–736, doi:10.5194/acp-8-721-2008, 2008.
- 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.: Tomographic reconstruction of atmospheric volumes from infrared limb-imager measurements, PhD thesis, Wuppertal University, 2011.
- von Clarmann, T., De Clercq, C., Ridolfi, M., Höpfner, M., and Lambert, J.-C.: The horizontal resolution of MIPAS, *Atmos. Meas. Tech.*, 2, 47–54, doi:10.5194/amt-2-47-2009, 2009.
- Weigel, K., Riese, M., Hoffmann, L., Hofer, S., Kalicinsky, C., Knieling, P., Olschewski, F., Preusse, P., Spang, R., Stroh, F., and Volk, C. M.: CRISTA-NF measurements during the AMMA-SCOUT-O3 aircraft campaign, *Atmos. Meas. Tech.*, 3, 1437–1455, doi:10.5194/amt-3-1437-2010, 2010.
- 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:MTAHT>2.0.CO;2, 1973.