



Initial Assessment of the Measurement and Retrieval Performance of the Upper- Atmospheric Terahertz Limb- Sounder LOCUS

D Gerber, R Siddans, W Feng, T Kovacs

September 2018

©2018 Science and Technology Facilities Council



This work is licensed under a [Creative Commons Attribution 4.0 Unported License](https://creativecommons.org/licenses/by/4.0/).

Enquiries concerning this report should be addressed to:

RAL Library
STFC Rutherford Appleton Laboratory
Harwell Oxford
Didcot
OX11 0QX

Tel: +44(0)1235 445384
Fax: +44(0)1235 446403
email: libraryral@stfc.ac.uk

Science and Technology Facilities Council reports are available online at: <http://epubs.stfc.ac.uk>

ISSN 1358-6254

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigations.

Initial Assessment of the Measurement and Retrieval Performance of the Upper-Atmospheric
Terahertz Limb-Sounder LOCUS

Daniel Gerber ^{1,2}, Richard Siddans ², Wuhu Feng ³, Tamas Kovacs ³

¹ Millimetre-wave Technology Group, RAL Space, Harwell Campus, Didcot OX11 0QX, UK

(Contact email: daniel.gerber@stfc.ac.uk)

² Remote Sensing Group, RAL Space, Harwell Campus, Didcot OX11 0QX, UK

³ School of Chemistry, University of Leeds, Woodhouse Lane, Leeds LS2 9JT, UK

Author Note

This is the Final Report of the CEOI funded pathfinder project “LOCUS Preparatory Activities for EE-9”. The aim of the study was to increase the Science Readiness Level of the ESA Earth Explorer 9 candidate mission LOCUS, a THz limb-sounder for the upper atmosphere. This is achieved by performing linear retrieval simulations to assess the retrieval errors of the mission. Crucially, this includes a first ever retrieval assessment of the key mesospheric – lower thermospheric species atomic oxygen from limb remote sensing measurements.

Table of Contents

Abstract.....	5
Initial Assessment of the Measurement and Retrieval Performance of the Upper- Atmospheric Terahertz Limb-Sounder LOCUS	6
Scope and Methodology	7
Scope.....	7
Methodology	8
Tools	8
Description of Tools and Data Sources used in the Study.....	10
Description of the Atmospheric Modelling	10
Model Description, and Improvements over previous versions.....	10
Model Results.....	11
Description of the Forward Model	18
The Reference Forward Model	19
Modification to the RFM Source Code for this Study.....	19
RFM Configuration	20
Description of the Retrieval Processor	21
Algorithm Theoretical Basis Description	22
Implementation.....	24
Flow Chart.....	24
Software Design	26
User Manual	27

Study Results.....	30
Climatological Reference Profiles	30
Atomic Oxygen (O).....	31
Nitric Oxide (NO).....	31
Nitrosylium (NO+).....	31
Hydroxyl (OH).....	31
Carbon Monoxide (CO).....	32
Water Vapour (H ₂ O)	32
Ozone (O ₃)	32
Hydroperoxyl (HO ₂).....	32
Remote Sensing of Atomic Oxygen	36
Measurement Frequencies	36
Atmospheric Distribution.....	38
Opacity Profile from Self-extinction	39
Atmospheric Gradient and Variability	43
Simulated Measurements.....	44
Sideband Configuration	44
Measurement Errors	45
Simulated Spectra.....	46
Band 1 – 4.75THz.....	47
Band 2 – 3.55THz.....	48

Band 3 – 1.15THz.....	49
Band 4 – 0.77THz.....	50
Single Species Retrievals.....	51
Measurement Vectors.....	52
Atmospheric Weighting Functions	57
Retrieval Precision.....	68
Vertical Resolution	72
Impact of Atmospheric Variability	82
Impact of Spectral Resolution.....	85
Nitrosylium Detectability	86
Signal to Noise Ratio.....	87
Lower State Energy	88
Results of Nitrosylium Detectability Study	89
Enhanced D-layer Nitrosylium.....	90
Mission Requirements Table.....	93
Conclusions.....	96
Acknowledgments.....	97
References.....	98

Abstract

This is the Final Report of the CEOI funded pathfinder project “LOCUS Preparatory Activities for EE-9”. The aim of the study was to increase the Science Readiness Level of the ESA Earth Explorer 9 candidate mission LOCUS, a THz limb-sounder for the upper atmosphere. This is achieved by performing linear retrieval simulations to assess the retrieval errors of the mission. Crucially, this includes a first ever retrieval assessment of the key Mesosphere – Lower Thermosphere species atomic oxygen from limb remote sensing measurements. The study shows that it is possible to retrieve the vertical abundance profiles of all the LOCUS targets species, and quantifies the expected measurement precision, the achievable vertical resolution, and the altitude ranges where a retrieval is possible. The core finding is that it is possible to extract some abundance information for atomic oxygen below attitudes of 120km, where the line centre starts to become opaque. We believe that this information comes from self-absorption of upper mesospheric atomic oxygen against the hot, Thermospheric background. This means that abundance profile retrievals of atomic oxygen are possible in an altitude range that overlaps with infrared heat-flux measurements by instrument like SABER. This vindicates the LOCUS science objectives of providing a first 3D abundance map of atomic oxygen from remote sensing measurements, and the quantification of CO₂-O quenching rates from collocated THz and infrared measurements.

Keywords: LOCUS, THz Limb Sounder, Retrieval Study, CEOI, ESA Earth Explorer

Initial Assessment of the Measurement and Retrieval Performance of the Upper-Atmospheric Terahertz Limb-Sounder LOCUS

The purpose of this study is to provide a first assessment of the retrieval performance that can be expected from a THz limb sounding satellite instrument for the upper atmosphere, specifically the Mesosphere – Lower Thermosphere (MLT) from about 50km – 150km. This work is of particular significance, because many of the trace gases that are of scientific interest in the MLT region have only scarcely been measured, or not at all. In the case where they have been measured, the data processing is under continuous academic development and has to be classified as ‘experimental’ [Hedin 1991]. For the most important MLT species, atomic oxygen (O) [Gray 2010, Plane 2003], the only existing measurements of the emission line taken from the Space Shuttle [Grossmann 2000, 1997] are not of sufficient quality to allow for a direct retrieval of the vertical abundance. For O, this study will therefore have to establish for the first time, whether direct retrievals of O abundance profiles are feasible at all, and if yes, what the expected precision as a function of altitude will be. The baseline instrument concept is that of the Low Cost Upper Atmospheric Sounder LOCUS. LOCUS is a mission concept for a passive heterodyne radiometer at THz frequencies (i.e. the far infrared region of the spectrum). Its distinctive feature is the use of Quantum Cascade Laser (QCL) local oscillators (LO) to pump Schottky diode mixers. This has never been attempted before, and will provide a significant reduction in complexity, weight and cost over conventional Los (i.e. gas lasers in the case of the MLS 2.5 THz channel). Furthermore, the concept allows heterodyne receivers to be built in the so-called THz-Gap, a frequency interval from ~1THz to 10THz, where no conventional LO sources exist.

LOCUS is a candidate mission for ESA’s Earth Explorer 9 Call. For a Phase-A proposal to be successful, the measurement performance of the instrument has to be

demonstrated. Because of the novelty of the concept, and the current lack of remote sensing observation in the higher THz region, there is no existing evidence from previous missions. We therefore have to rely on retrieval simulations to provide this evidence.

Work performed in this study presents the first retrieval study for a THz remote sensing instrument in the 1THz to 5THz region. The aim is to provide evidence that the detection of key Upper Atmospheric species is indeed possible, and to quantify the expected retrieval precision, vertical range, and vertical and spectral resolutions. This includes the development of a retrieval processor from scratch, which will be the basis for future science studies, and a starting point for a future L2 prototype processor.

Scope and Methodology

Scope

The scope of the study is to calculate estimated retrieval error profiles for the LOCUS targets species under realistic conditions of atmospheric composition, observation geometry and measurements errors. The measurement errors are the random errors resulting from the noise-performance of the receivers (aka. thermal noise, or NEdT). In this first iteration, we exclude systematic errors from the instrument model (i.e. field-of-view knowledge, pointing knowledge, etc.) although the processor developed in this study will be the basis for this next iteration in retrieval simulations. Retrieval errors shall be calculated for the following geophysical observables: O, OH, O₃, NO, NO⁺, CO, H₂O, HO₂, and Temperature. The simulated measurements shall be representative of the LOCUS instrument as defined in the ESA IOD Study [SSTL-IOD 2014].

The following statement of work summarises the scope of this study:

- Provision of 4D model fields of relevant trace gases in the Mesosphere – Lower Thermosphere (MLT) region of the Earth's atmosphere

- Software-Engineering of a radiative transfer model to simulate LOCUS measurements based on atmospheric profiles from the model database from the previous task
- Software-Engineering of an Optimal Estimation Method (OEM) processor for the retrieval of simulated measurements
- Definition of case scenarios for retrieval simulations to allow a basic assessment of the capacity to obtain geophysical observables from THz observations (i.e. quantification of retrieval precisions for each geophysical observable, maximum achievable vertical resolution and vertical intervals for which a retrieval is possible).
- Assess the potential to measure Nitrosylium (NO^+) in the atmosphere
- Performing of retrieval simulations, visualization, summary of findings

Methodology

To estimate the retrieval errors, a Zero-order retrieval is performed, in which the weighting functions and average kernels of an Optimal Estimation Method retrieval are used to perform a linear mapping of the known measurements errors onto unknown (i.e. estimated) abundance profile errors. This analysis is performed for a number of different input scenarios, i.e. different frequency resolutions, different vertical retrieval grid spacings, and different atmospheric scenarios (i.e. day/night for diurnal species) in order to establish the vertical range and resolution under which a given species can be observed, both at its peak and at its weakest abundances.

Tools

The tools required for the study are:

- An atmospheric model, to calculate the distribution profiles of all atmospheric species needed to perform radiative transfer calculations. We are using a modified version of

the WACCM-X model with included D-layer chemistry, as described in the section “Description of the Atmospheric Modelling”

- A radiative transfer model to simulate measurements along a given line of sight. The simulated measurements, collated for an entire atmospheric scan, are called the measurement-vector. We are using the Reference Forward Model (RFM) as described in section “Description of the Forward Model”
- A retrieval processor to perform an inversion from the simulated measurements-vector back onto the state space of atmospheric products, aka state-vectors. We are using an Optimal Estimation Method (OEM) inversion as described in the section “Description of the Retrieval Processor”

Description of Tools and Data Sources used in the Study

Description of the Atmospheric Modelling

Model Description, and Improvements over previous versions

For this project, we have done a few model experiments using the Whole Atmosphere Community Climate Model either without or with thermosphere and ionosphere extension (termed as WACCM or WACCM-X respectively) based on the framework from the fully coupled global climate model Community Earth System Model (CESM version 1, e.g., Hurrell et al., 2013). WACCM is a comprehensive numerical model from the Earth's surface to the lower thermosphere (~140 km) [Garcia et al., 2007; Marsh et al., 2013] while WACCM-X has a model top boundary between 500 and 700 km [Liu et al, 2018]. Both models use the same detailed physical processes as described in the Community Atmosphere Model, version 4 (CAM4) (Neale et al., 2012) and have the fully interactive chemistry described in Mills et al. (2016). The current configurations for WACCM and WACCM-X are based on a finite volume dynamical core (Lin, 2004) for the tracer advection. The horizontal resolution is 1.9° latitude x 2.5° longitude. For WACCM we used a specific dynamics (SD) version of WACCM discussed in Feng et al. 2017, which is nudged with NASA's Modern-Era Retrospective Analysis for Research and Applications (MERRA) below 60 km [Lamarque et al., 2012]. The horizontal resolution is 1.9° latitude × 2.5° longitude, WACCM has 88 vertical model layers giving a height resolution of ~3.5 km in the MLT. The vertical model layers and the vertical resolution for WACCM-X is 81 levels which is identical as Liu et al. [2010] and it is a “free running” model simulation which still produce the reasonable climatology in the model. We ran the model for year 2000 condition for 16 years for WACCM-X, and WACCM ran from 2002 to the end of 2014. Note that we have updated some important WACCM and

WACCM-X inputs, for example Solar forcing (F10.7, Ap, ion-pair ionization production rate, SSI etc.). We have also incorporated a detailed D-region chemistry above 50 km in WACCM [Kovacs et al., 2016]. We have saved the LOCUS key target variables (O, Ozone, NO, OH, H₂O, HO₂, NO⁺, O₂⁺, electron density, temperature etc).

Model Results

The key and basic question is how the models WACCM and WACCM-X perform in the mesosphere and lower thermosphere (MLT) which is interested region for this project. Here we only show some examples of the climatological fields from model simulations compared with available measurements. Here we just focus on the comparison with SABER measurement [Russell, 1999]. Note that the seasonal, diurnal variabilities and long term trend have not shown here.

The Sounding of the Atmosphere using Broadband Emission Radiometry (SABER) instrument aboard on the TIMED satellite has performed measurements of atmosphere temperature and density (e.g., ozone) from near the tropopause into the thermosphere (20-120 km) since 2002 [Russell, 1999]. The latitude coverage on a given day extends from 53° latitude in one hemisphere to 83 ° in the other hemisphere [Zhang et al., 2006]. The SABER temperature and ozone data were provided by Anne Smith at NCAR (version 2.0, Smith et al., 2010).

Figure 1 compares the annual mean temperature profiles from the upper stratosphere to lower thermosphere (50-110 km) from SABER (v2.0) satellite measurements (averaged from 2002-2015) and long term WACCM (2001-2014) and WACCM-X (2000-2015) simulations.

Remsberg et al. [2008] indicated that SABER temperature has 2-10 K random errors for mesosphere and lower thermosphere region and has systematic errors of ± 2 K in the lower mesosphere, and 2-4 K lower in the middle mesosphere. It has been widely used to validate other satellite observations and model in the middle/upper atmosphere [Xu et al., 2006]. compared the monthly mean midlatitude nighttime temperature profiles from sodium lidar and SABER measurements and showed that there are good agreement in observed temperature at 85-100 km. French and Mulligan [2010] compared temperature derived from OH(6-2) emission at Davis station with SABER and showed that the SABER has a warmer temperature trend than OH temperature. Generally, both WACCM and WACCM-X simulate the observed temperature structures in upper stratosphere to lower thermosphere region (i.e., temperature decreases with increasing altitude from stratopause to mesopause) and have reasonable stratopause and mesopause as well as equator to pole temperature gradient. It is noticeable that WACCM has better MLT than WACCM-X though both models are free running above 60 km (note WACCM is only nudged below 50-60 km by MERRA-2 reanalysis).

Figure 2 shows the climatological zonally averaged mean O density profiles as a function of latitude from SABER, WACCM and WACCM-X. Atomic oxygen is the most abundance chemical species and drives nearly all the chemistry in the mesosphere lower thermosphere (MLT) region [Plane et al., 2015]. SABER measurement shows the O is produced in the lower thermosphere which has a peak of 10^{12} cm^{-3} around 100 km and decreases with altitude mainly due to the photochemistry and dynamics. Both WACCM and WACCM-X predicts the reasonable peak O concentration height between 90 and 100 km but underestimate the observed O peak density. WACCM and WACCM-X use the same

chemistry/dynamical and physical modules in the MLT region, the only difference is that the WACCM is nudged by reanalysis below 50-60 km and using a different Prandtl number (Pr) which is the ratio of momentum flux to heat flux in the gravity wave parameters [Feng et al., 2013]. WACCM uses $Pr=2$ while WACCM-X uses $Pr=4$. Therefore, the differences in the simulated O abundances can be partly attribute the vertical diffusive transport used in the model.

Figure 3 compares the climatological zonally averaged mean O_3 density profiles as a function of latitude between SABER observations and model simulations from WACCM and WACCM-X. Ozone is formed from molecular oxygen in the atmosphere, which is mainly controlled by chemistry and transport processes. There is the secondary ozone peak layer in the mesosphere region around 83-93 km which has a similar volume mixing ratio (~ 10 ppmv) comparative with stratosphere ozone mixing ratio (not shown). However, the concentration is much smaller seen in the SABER measurement in Figure 3 because very low pressure in the upper atmosphere. The uncertainty of observed ozone in the MLT is quite large (e.g, Huang et al., 2008). Huang et al. [2008] mentioned that the systematic errors of the SABER O_3 in the mesosphere can be as large as 40 %. Smith et al. [2008] compared night-time MLT ozone from SABER and GOMOS measurements and they found that the SABER MLT ozone is lower than from GOMOS data above 85 km (about 20% lower at 95 km). However, there are few other measurements for comparison and validation of MLT ozone above 90 km [Smith et al., 2008]. WACCM and WACCM-X have similar amount of ozone concentration in the lower mesosphere compared with SABER data, however, the models significant underestimate the observed mesospheric ozone concentration (similar as O in Figure 2) because the modelled downward transport from the thermosphere appears to be too slow [Huang et al., 2015]. Need to mention that both model simulates quite well the observed ratio

of nighttime to daytime O_3 (not shown) which indicates the chemistry in the model is still well valid.

Figure 4 shows the climatology of modelled NO concentration (70-140 km). There is substantial source of NO in the lower thermosphere due to the reactions involving electronically excited atoms which are produced from highly exothermic ion-molecule reactions [Plane et al., 2015]. Both model have the peak NO concentration between 100-100 km which arises from energetic electron precipitation [Kovacs et al., 2015]) and has strong downward transport in the high latitude due to the winter polar vortex.

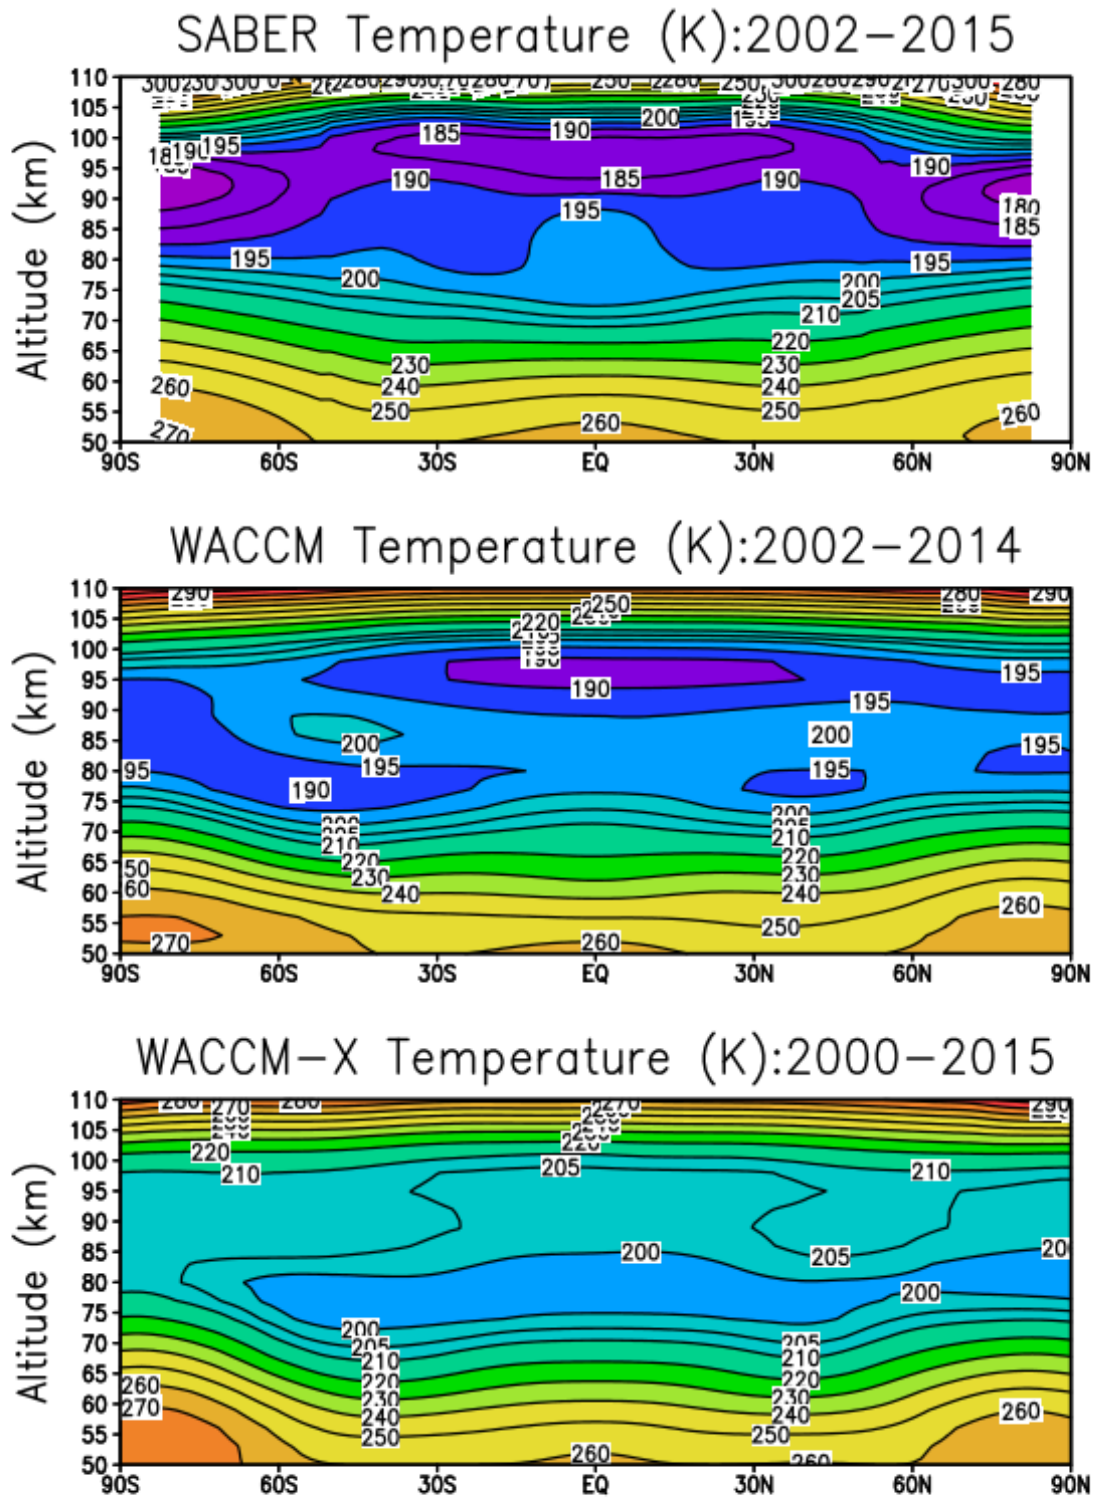


Figure 1: Climatological annual mean temperature profiles from 50–110 km as a function of latitude from SABER satellite measurement and long term simulations of WACCM and WACCM-X

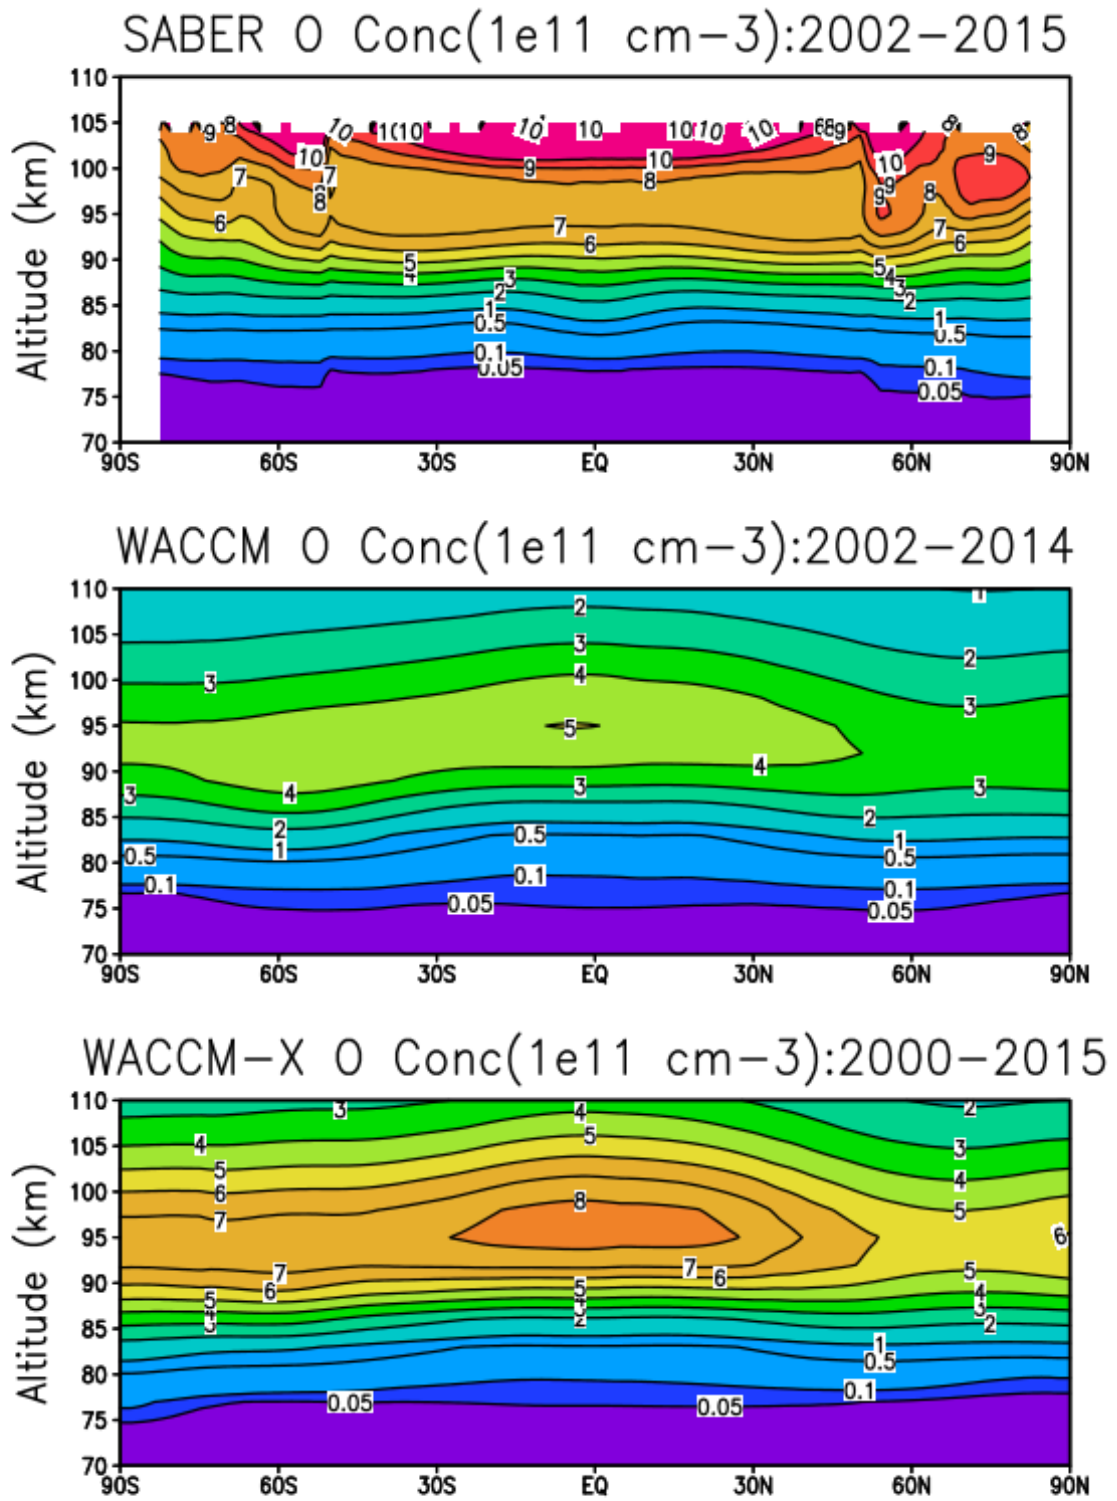


Figure 2: Similar as Figure 1, but for O concentration from 70 to 110 km

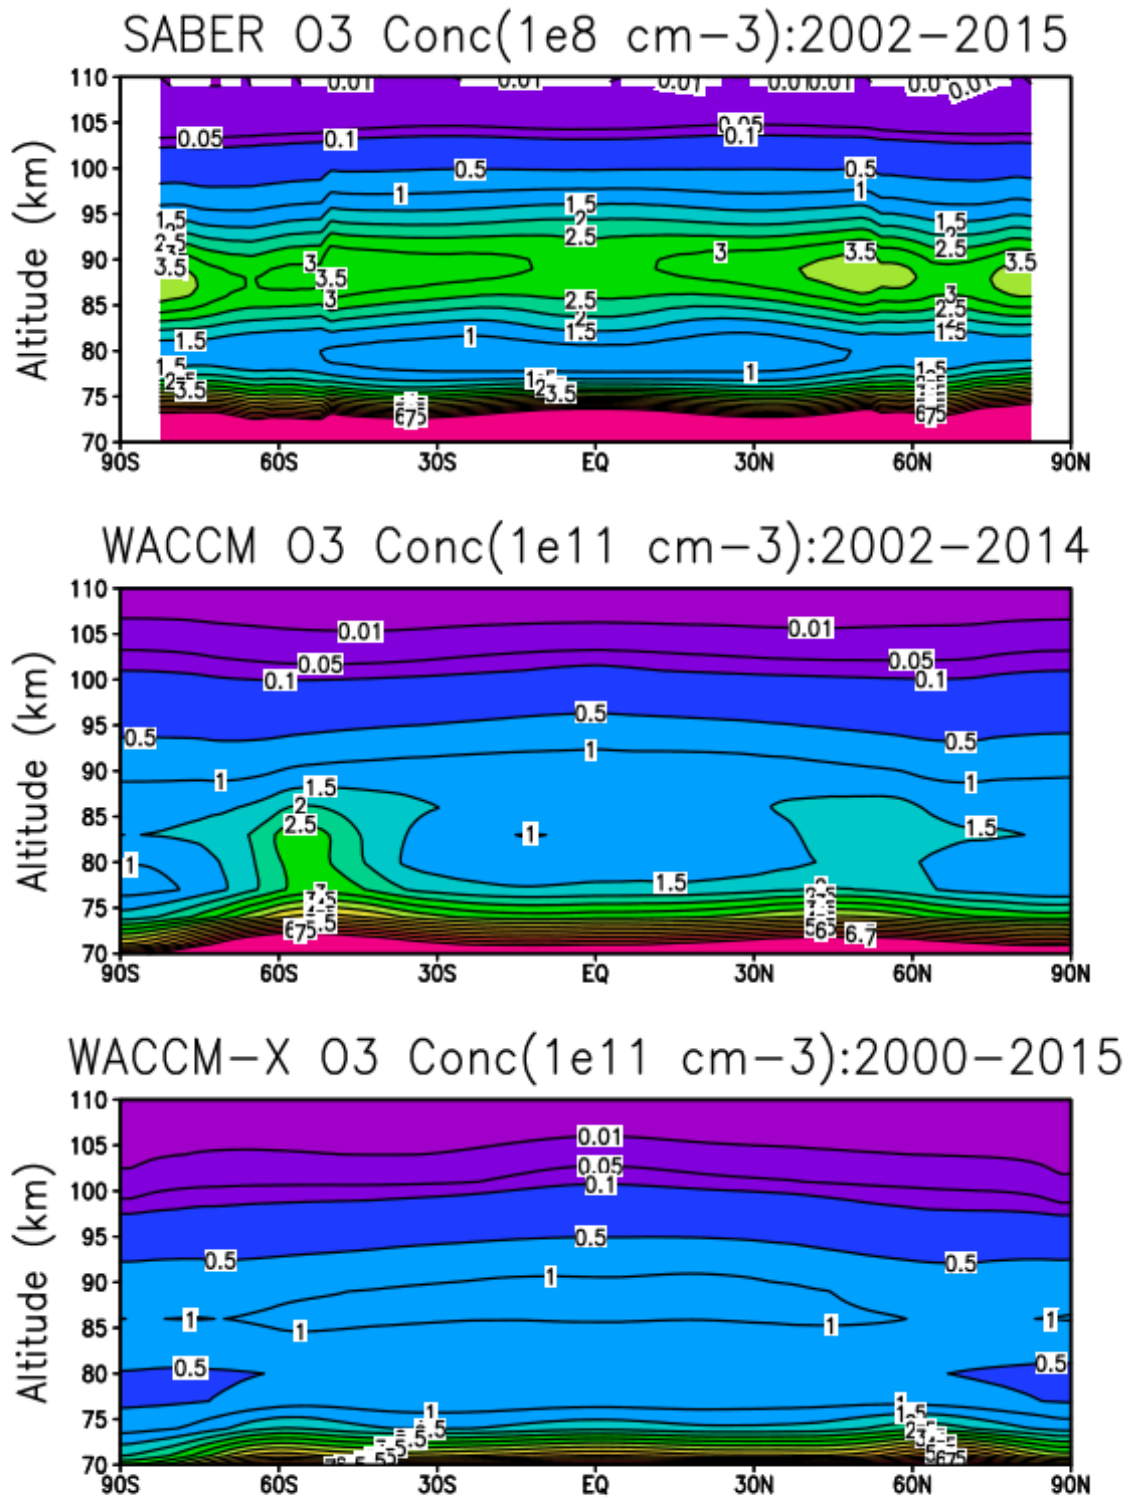


Figure 3: Similar as Figure 2, but for ozone concentration

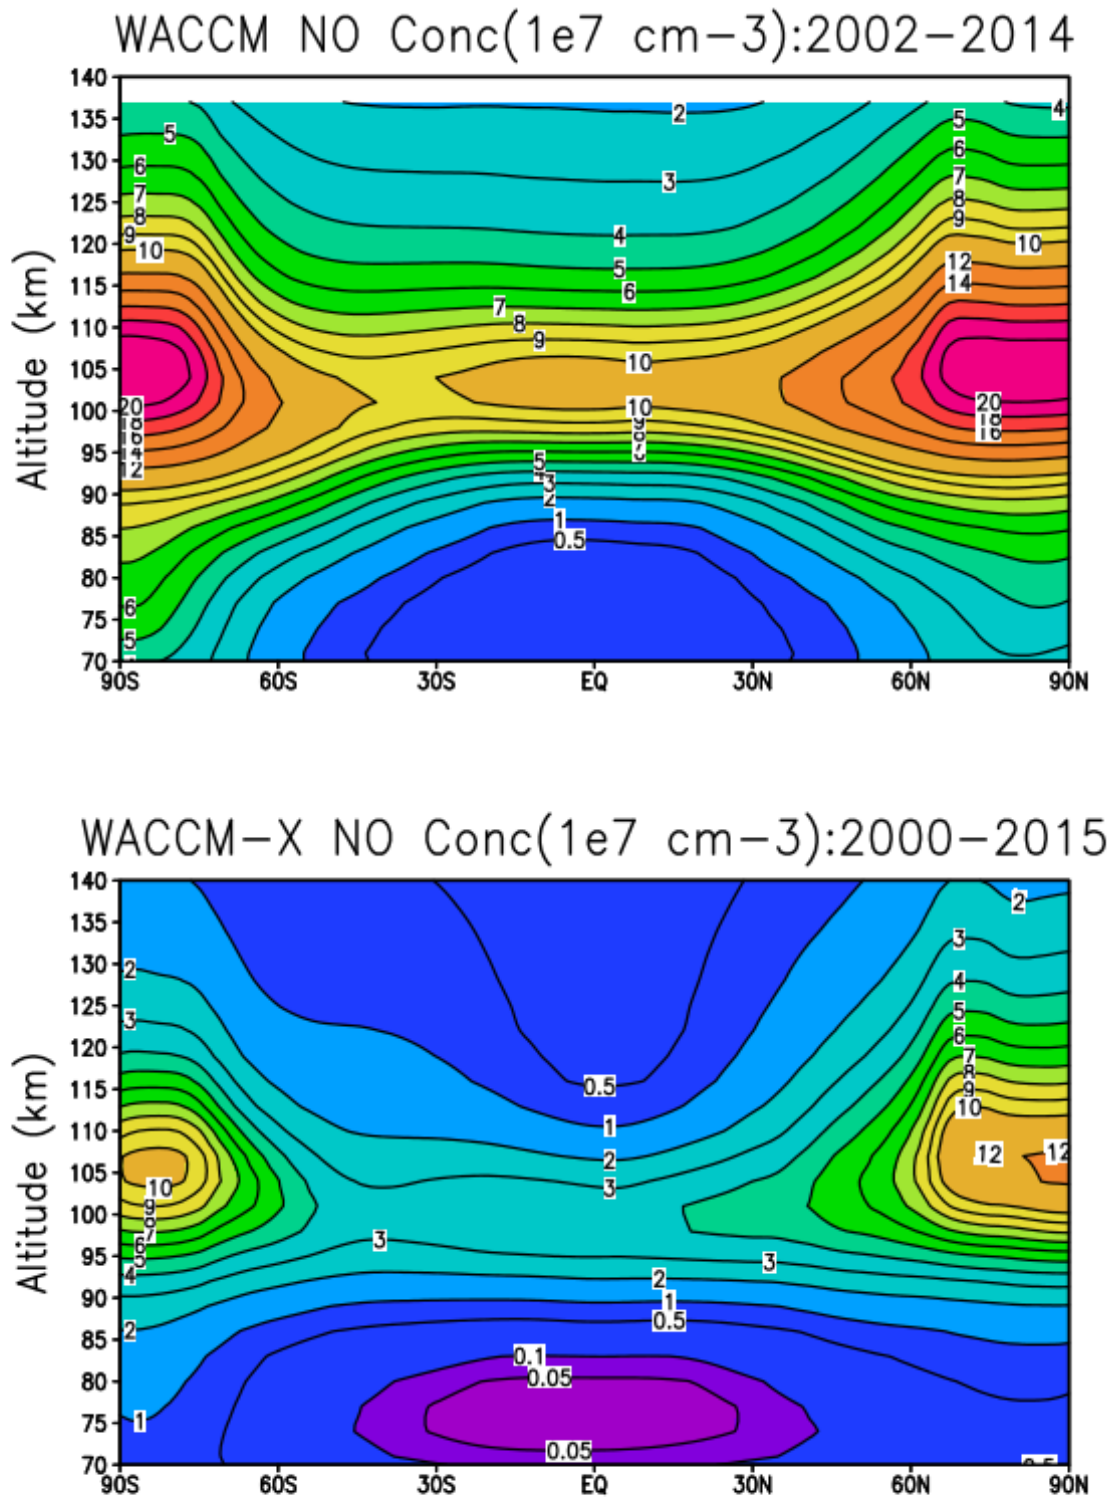


Figure 4: Similar as Figure 2, but for NO concentration from 70 to 140 km

Description of the Forward Model

The radiative transfer model for LOCUS has to fulfill a number of selection criteria. It has to support all of the target species, it has to work at THz frequencies, and it has to support

altitudes up to the top of the thermosphere at ~600km. It should also have the option to include effects of Non Local Thermal Equilibrium (NLTE). NLTE effects are a big factor for infrared emissions at MLT altitudes. While first research indicates that they are probably not crucial for THz lines, the feature is desirable to have for future, more in-depth studies.

We have decided to use the Reference Forward Model (RFM) for LOCUS. The RFM has been developed by our colleague Anu Dudhia at Oxford University specifically for the operational data processing of the MIPAS instrument [Dudhia 2015]. MIPAS is an atmospheric limb-sounder with an upper atmospheric observation mode, making it very similar to LOCUS. For this reason, the RFM turns out to meet all our requirements. The close working relation with the RFM developer is also of benefit. Because of the scientific novelty, radiative transfer codes have rarely been thoroughly tested at THz frequencies, and for exotic molecules like atomic oxygen. On occasions, we've encountered unexpected behavior of the code when trying out new things, which Anu was always very quick to fix.

The Reference Forward Model

The RFM is a line-by-line radiative transfer code written in FORTRAN. It supports a large number of molecules and observation geometries. It is heavily and easily customisable through the use of ASCII driver files, both for the measurement setup, as well as for the definition for the atmosphere. This makes batch processing of multiple scenarios easy, since these configuration files can easily be generated by a computer script. The spectral data is supplied in HITRAN format (binary format for speed).

Modification to the RFM Source Code for this Study

Because RFM is a FORTRAN code, the memory allocation is pre-set at compile time. Because LOCUS has an extensive scan, with many tangent points, as well as a high spectral

resolution, the resulting measurements vector is very large; too large for the default memory allocation it turned out. We therefore had to recompile RFM from source code, with an increased memory allocation. This was possible through the use of a powerful computing cluster at RAL, with computing nodes of up to 130GB of memory.

The RFM base version for the LOCUS simulations was **v4.35**. The modified version for LOCUS is identified as **v4.35_RAL_01**. The source code modifications applicable to v4.35_RAL_01 are listed in Table 1.

Table 1: List of source code modification in RFM v.4.35_RAL_01

Routine	Line	Parameter	V4.35	V4.35_RAL_01
rfmsiz.inc	54	MAXCLC	15,000	60,000
rfmsiz.inc	63	MAXFIN	2,016	30,000
rfmsiz.inc	89	MAXJAC	300	459
rfmsiz.inc	114	MAXOUT	3,000	30,000
rfmsiz.inc	119	MAXPTH	120,000	240,000
rfmsiz.inc	134	MAXTAN	4,000	20,000
rfmsiz.inc	151	MAXWID	5	135

RFM Configuration

The following is an example RFM driver file that illustrates a typical forward model run. Some aspects of the configuration will of course change for different scenarios, but the overall concept remains similar. E.g. this example is for a full run of all species, on an optimized frequency grid with a maximum spectral resolution of 1MHZ and limb-scan steps of 2km altitude. The RFM is asked calculate Jacobians/Weighting Functions for all products (faster than calculating them via perturbations). The spectral data file is based on HITRAN2012, but condensed down to include just the spectral lines and wavelength ranges applicable to LOCUS (again, to dramatically speed up the processing). The RFM atmosphere is generated automatically by the retrieval processor from a database of WACCMX model data.

```

! RFM Driver Table for LOCUS Retrieval Study
! Created by LOCUS_MK_RFM_DRV on Fri Mar 31 16:14:14 2017
*HDR
  LOCUS CEOI Prep EE9 2016 Daniel.Gerber@stfc.ac.uk
*FLG
  BBT GHZ JAC LUN SHP
*SPC
  ALL
~/Data/projects/locus/rfm/spc/LocusReducedFreqsALL1MHz0P1K.ghz
*GAS
  O OH HO2 NO NO+ CO H2O O3 O2
*ATM
  rfm_locus.atm
*TAN
  50 52 54 56 58 60 62 64 66 68 70 72 74 76 78 80 82 84 86 88
  90 92 94 96 98 100 102 104 106 108 110 112 114 116 118 120 122
  124 126 128 130 132 134 136 138 140 142 144 146 148 150
*HIT
  ~/Data/projects/locus/rfm/hit/HITRAN2012LOCUS.hit
*JAC
  TEM Jacobians.tan
  O Jacobians.tan
  OH Jacobians.tan
  HO2 Jacobians.tan
  NO Jacobians.tan
  NO+ Jacobians.tan
  CO Jacobians.tan
  H2O Jacobians.tan
  O3 Jacobians.tan
*SHP
  VOI *
*BBT
  locus_*.bbt
*END

```

Description of the Retrieval Processor

The retrieval processor uses an Optimal Estimation Method (OEM) to solve the inversion problem of finding the atmospheric state that most likely lead to a given measurement. A measurement being the collective measured atmospheric spectra at all altitudes of our limb scan. The inversion problem is either under-constrained, over-constrained, or both, so a regularisation mechanism is needed. We use the widely used Levenberg-Marquardt regularization. The OEM and Levenberg-Marquardt regularization is

extensively described in Clive Rodger's seminal book on Inverse Methods in Atmospheric Sounding [Rodgers 2000].

Algorithm Theoretical Basis Description

The use of OEM in geosciences, especially atmospheric remote sensing, is extensively described in Clive's book, and the literature. We just give a summary here as a reference to the output products of the retrieval processor.

The initial remote sensing challenge is that we have a measured spectrum that is linked to an underlying atmospheric composition, which we are trying to establish. The function that links the atmospheric composition, or atmospheric state, to the measured radiation spectrum is modelled by the radiative transfer equations, which is reasonably well known – within a certain error bar. In matrix notation, this situation is described as follows:

$$A \vec{x} = \vec{y}$$

where x is the atmospheric state vector, y is the measurement vector, and A is the function we are trying to invert. The essential concept is to transform the matrix, A , into a conditional probability and the variables, x and y into probability distributions by assuming Gaussian statistics and using empirically-determined covariance matrices. The measurement error covariance matrix S_Y is reasonably well known, because instrument builders have a good understanding of the measurement performance of their instrument.

Because of the assumption of Gaussian statistics, and Gaussian error distributions for the measurement vector and the state vector, the probability function for the most likely measurement from a given atmospheric state, as well as the most likely atmospheric state linked to a given measurement, are following a Gaussian distribution. The atmospheric state

can further be regularised by introducing an *a priori* atmospheric state X_a together with a *a priori* error covariance matrix S_{Xa} that both form a Gaussian distribution. With this, the state error covariance matrix S_X follows as

$$S_X = (A^T S_Y^{-1} A + S_{Xa}^{-1})^{-1}$$

and the profile of the retrieved atmospheric state follows as

$$\hat{x} = \hat{x}_a + S_X A^T S_Y^{-1} (y - A \hat{x}_a)$$

The retrieval error is profile is then simply calculated from the trace of the state error covariance matrix, the value of which denote the variance of the retrieval precision.

Typically, with optimal estimation, in addition to the vector of retrieved quantities, one extra matrix is returned along with the covariance matrix. This is sometimes called the resolution matrix or the averaging kernel and is calculated as follows:

$$K = (A^T S_Y^{-1} A + S_{Xa}^{-1})^{-1} A^T S_Y^{-1} A$$

This tells us, for a given element of the retrieved vector, how much of the other elements of the vector are mixed in. In the case of a retrieval of profile information, it typical indicates the altitude resolution for a given altitude.

Implementation

The retrieval code is written in IDL. The core routine for the OEM retrieval is the generic IDL routine LQ_RET.PRO developed in the Remote Sensing Group at RAL Space (Author: Richard Siddans). The core radiative transfer model is RFM, which is a binary FORTRAN module. RFM performs a single radiate transfer calculation, so in order to simulate a LOCUS measurement over a full atmospheric limb-scan and in double-sideband mode, RFM has to be called several times, and the resulting spectra have to be collated. For this purpose, an IDL wrapper was written to configure RFM, run a single instance in a sandbox environment, collect the resulting spectral simulations and collate them into a LOCUS measurements, which is a collated vector of double-sideband spectra at each scan altitude. The prior information, and first guess profiles are read from the WACCMX model field database.

Flow Chart

The flow chart of the retrieval processor is shown in Figure 5. A state vector structure is defined from a template, and populated with atmospheric data from a WACCMX model field, either the full database, or the reduced file with climatological max/mean/min reference profiles. The state vector structure contains the state vector proper (the description of the atmospheric state, i.e. the collated vector of all atmospheric profiles), as well as meta-data (i.e. the prior errors of the atmospheric state parameters, as well as parameters names, units, matrix dimensions based on the number of scan altitudes, and start/stop indices to identify individual products inside of the collated state vector).

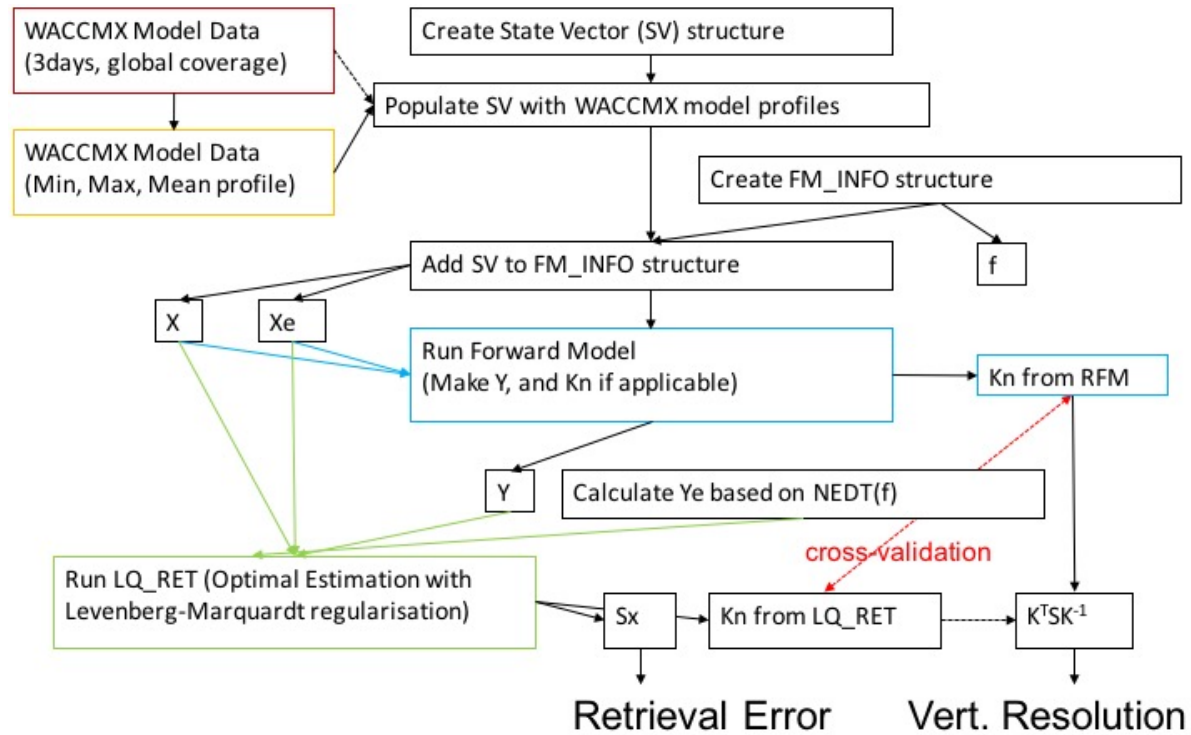


Figure 5: Flow chart of the LOCUS retrieval processor

The state vector structure is combined with the forward model configuration structure (FM_INFO), which hold meta-information required to configure the radiate transfer model RFM. This includes the list of retrieved species, the frequency grid, the spectral line data file to be used, the scan grid, the retrieval grid, integration times, filed-of-view definitions, spectral response definitions, and a full description of the atmospheric state (pressure, temperature, z-heights, and abundance profiles of all species).

From this, the forward model is run and the initial simulated measurement Y based on the state vector X is calculated. These are picked up by the OEM processor, which runs LQ_RET to find the retrieved atmospheric state X' most compliant with the prior state X and its associated variability X_e , as well as with the measurement Y and its associated measurement error Y_e .

In this initial study, we are performing linear error mapping to calculate the estimated retrieval errors based on our measurement errors. That means we are only running a single iteration of the OEM code. The retrieved state X' is therefore identical to the first guess state

X , but the retrieval diagnostics returned by LQ_RET allow us to quantify the expected retrieval performance. From the diagonal elements of the error covariance matrix S_e the estimated retrieval error can be calculated, and from the Averaging Kernel K^tSK^{-1} the vertical resolution can be derived.

Software Design

The following IDL routines make up the core of the LOCUS retrieval and forward model processor. They add up to over 3,500 lines of code. The LOCUS retrieval code is run under version control using Subversion (SVN), with a repository held on the servers of the British Atmospheric Data Centre (BADC).

Subversion URL: http://proj.badc.rl.ac.uk/svn/rsg_code/idl/dgerber/trunk/locus/ret

SVN Revision Number: 5661

Table 2: List of the core IDL routines that make up the LOCUS retrieval processor, and a description of their purpose

Procedure	Description
<code>def_locus_sv_add_indices.pro</code>	Add indices to help find and extract the records of individual geophysical data products in the state vector.
<code>def_locus_sv.pro</code>	Define the state vector structure format
<code>def_locus_svstruct.pro</code>	Define the format of individual data records in the state vector structure
<code>locus_fmatmos_ndens2vmr.pro</code>	Convert a state vector that was defined in number density to volume mixing ratios (RFM needs all input profiles defined in units of ppmv)
<code>locus_fm_info2idtag.pro</code>	Create short but unique ID string from all configuration options
<code>locus_merge_rfm_outfiles.pro</code>	Combine the outputs of individual RFM runs into a single measurement file
<code>locus_mk_fm_info.pro</code>	Create the FM_INFO structure to define the forward model run
<code>locus_mk_rfm_atm.pro</code>	Create the atmosphere file in the format required by RFM
<code>locus_mk_rfm_drv.pro</code>	Create the driver file for RFM
<code>locus_mk_je.pro</code>	Define the measurement error vector
<code>locus_pl_ret.pro</code>	Plot the results of a single retrieval

<code>locus_waccmx2sv.pro</code>	Populate the state vector with atmospheric profiles extracted from the WACCMX model field
<code>locus_waccmx2zg.pro</code>	Set the altitude grid that corresponds to the respective WACCMX atmospheric profiles
<code>lq_ret.pro</code>	Perform an OEM retrieval with Levenberg-Marquardt regularisation
<code>lq_sx.pro</code>	Compute the error covariance matrix S_x
<code>ret_locus.pro</code>	Perform a LOCUS retrieval for a given scenario
<code>rfm_locus.pro</code>	Run RFM for a given scenario
<code>run_ret_locus.pro</code>	Run a number of LOCUS retrievals for different scenarios

In addition to that a multitude of secondary tools were written for visualization, testing, data handling and manipulation (i.e. interpolation of profiles to sparse grids, optimized frequency grids, etc.), the analysis of ancillary data and general debugging.

User Manual

The retrieval processor is configured by editing run-time parameters in the wrapper script `run_ret_locus.pro`. The supported options are listed in Table 3.

Table 3: Run-time configuration options of the LOCUS retrieval processor

Parameter	Options	Description
<code>filename</code>	Filename	Name of WACCMX model fields for first guess and a priori profiles
<code>longitudes</code>	Vector of integers	Longitudes of reference profiles from WACCMX
<code>latitudes</code>	Vector of integers	Latitudes of reference profiles from WACCMX
<code>hours</code>	Vector of numbers	Time stamps of reference profiles from WACCMX (or when using climatological profiles: 0=Max, 1=Mean, 2=Min)
<code>int</code>	Vector of floats	Integration times in seconds
<code>ils</code>	Vector of floats	Instrument line shape (half-width of Gaussian spectral response in MHz). This is overridden by the use of an explicit <code>grid_file</code> (see below)
<code>fov</code>	Vector of floats	Field of view (half-power width of Gaussian antenna function) at the tangent point in km
<code>spec_ranges</code>	‘ALL’ ‘IOD’ ‘LAD’	All frequency bands (incl. optional 2 THz) Original 4 bands from the ESA IOD Study (excl. optional 2 THz) Only the 2 THz band
<code>grid_file</code>	Filename	Filename of an RFM grid file

scan_grid	'scan_grid_0	50km - 150km in 2km steps (Default)
	2'	
	'scan_grid_0	50km - 150km in 5km steps
	5'	
ret_grid	'scan_grid_1	50km - 150km in 10km steps (Debug)
	0'	
	Vector of floats	Custom scan grid points
	'ret_grid_2'	[indgen(51)*2+50,153,155,160,170,180,200,250,300]
	'ret_grid_5'	[indgen(21)*5+50,160,180,200,250,300]
	'ret_grid_10'	[indgen(11)*10+50,180,200,250,300]
hitran_file	'ret_grid_de	[indgen(6)*20+50]
	bug'	
	Filename	Name of spectral data file
	KSK	Binary keyword
KN	Binary keyword	Save averaging kernels in output file
		If set, forward model is expected to return weighting functions Kn
Gas_unit	'VMR'	Atmospheric profiles provided in volume mixing ratios
	'DENSITY'	Atmospheric profiles provided in number densities

There are a few more configuration options and default settings hardcoded into parts of the processor routines. I.e. the a priori errors are defined directly in the routine `locus_wacmx2sv.pro`, where the a priori profiles from WACCMX are put into the state vector, and the state vector errors are defined. The a priori errors can be set for each product individually. They consist of a relative and an absolute term (i.e. a multiplicative scaling factor (percentage of the absolute magnitude of the product), and a fixed offset (to prevent Zero errors at altitudes where the product is Zero/undefined).

The perturbations for the calculations of weighting functions Kn directly by the retrieval routine `LQ_RET.PRO` are set in the procedure `RET_LOCUS.PRO`. The code tries to be clever about defining the perturbations in order to avoid values that are either too large, or too small. The default perturbation is 1% of the value at every altitude (which is also the RFM default perturbation), but it's never smaller than 1% of the smallest non-zero value, i.e.

$$fp = 0.01$$

$$ap = \text{MIN}(x_ap[\text{WHERE}(x_ap \text{ GT } 0.0)]) * fp$$

where fp is the fractional perturbation, ap is the absolute perturbation, and x_ap is the a priori profile in the state vector structure. The same formalism is also used in the definition of a priori errors.

Study Results

Climatological Reference Profiles

University of Leeds have provided us with a WACCMX model field of global atmospheric profiles up to 500km of altitude, sampled at 3 hour intervals over a 3-day period starting at 2000-01-01 00:00:00. The assumption is that this will cover all diurnal variations, as well as latitudinal gradients over both summer (SH) and winter season (NH). The retrieval processor is set up that it can extract atmospheric profiles at any given latitude, longitude and time stamp and compile a reference atmosphere from that. This atmosphere serves as a priori information for the OEM retrieval. It's also used to compute the simulated measurements from the forward model (RFM).

The filename of the WACCMX model file, saved in NetCDF data format, is: `locus_waccmx_pressure_3days.nc`. The NetCDF file is 1.7GB in size, which makes accessing it slow. This is why, especially in this initial phase of setting up and debugging the retrieval processor, we are using three representative climatological reference cases as a basis for our atmosphere. There are the minimal, maximal, and average abundances for each species. We have extracted them from the full WACCMX NetCDF file, and saved them in the same format, so that the handling of the input profiles is 100% transparent to the code. This results in a new WACCMX model file called `locus_waccmx_pressure_3days_maxmin.nc`, which is identical in all but the temporal dimension to the original WACCMX file. Instead of the original 24 time steps, we retain but 3 steps and use them to sort the three statistical scenarios:

- Time index 0: Maximum profiles of all products
- Time index 1: Mean profiles of all products
- Time index 2: Minimum profiles of all products

Atomic Oxygen (O)

Atomic oxygen is created by the photo-dissociation of molecular oxygen at altitudes above ~90km under intense solar UV radiation. It reacts back to O₂ at night below ~120km, but above that the chemical lifetime is long enough that it prevails. At altitudes above 120km O quickly becomes the dominant atmospheric species in terms of mixing ratio. O has a strong diurnal variability in the MLT region, as seen in Figure 6 or Figure 7.

Nitric Oxide (NO)

Nitric oxide is formed in the upper atmosphere through chemical reactions prompted by percolating electrically charged particles from the solar wind. It is highly variable as a result of the variability of solar activity, and also because it is itself chemically active. Below ~100km the NO abundance quickly degrades, as seen in in Figure 6 or Figure 7.

Nitrosylium (NO⁺)

Nitrosylium is not one of the original WACAMMX species. However, University of Leeds (courtesy of Tamas Kovacs) have provided us with a few sample profiles to cover the expected variability of NO⁺ in the ionosphere (E- and F-layers) at altitudes above 100km. The maximum scenario is representative of the extreme solar proton event in the year 2003. This profile was also used in our Nitrosylium detectability study.

Hydroxyl (OH)

Hydroxyl exists in the stratosphere and mesosphere, e.g. as a product of methane oxidation. It is an extremely reactive radical, and there has a large variability. It does however phase out rapidly at altitudes above 80km, and at 100km it all but disappears, as shown in in Figure 6 or Figure 7.

Carbon Monoxide (CO)

Carbon monoxide is a stable species, and one of the few molecules that is still present even at very high altitudes. Its abundance peak is at 150km (in units of volume mixing ratio), as seen in in Figure 6 or Figure 7. This makes it a good tracer product to study dynamical processes.

Water Vapour (H₂O)

Water vapour is mainly constrained to the troposphere, where it has a massive abundance gradient, covering several orders of magnitude. Most water freezes out at the tropopause, but then it has a second peak in the stratosphere at around 60km, where water can be produced from methane oxidation. Ultimately though it is photo-dissociated in the mesosphere. The water vapour abundance profiles in in Figure 6 and Figure 7 are plotted on a semi-logarithmic scale because of the large gradient.

Ozone (O₃)

Ozone has a distinct abundance peak in the stratosphere (in the well-known ozone layer), but it is still present at higher altitudes, where it is involved in many chemical reactions, not least with atomic oxygen. It peaks in variability at ~80km altitude, which is where it starts to be broken up by solar UV radiation to form O. In the MLT region its variability is less pronounced, as seen in Figure 6 or Figure 7.

Hydroperoxyl (HO₂)

Hydroperoxyl, not unlike hydroxyl, is confined to altitudes below ~80km, above which it quickly disappears. Furthermore, it's absolute abundance is significantly lower than

most of the other species, which results in a comparatively weak spectral emission line.

Existing satellites missions (i.e. Aura MLS) only measure zonal monthly means of HO₂.

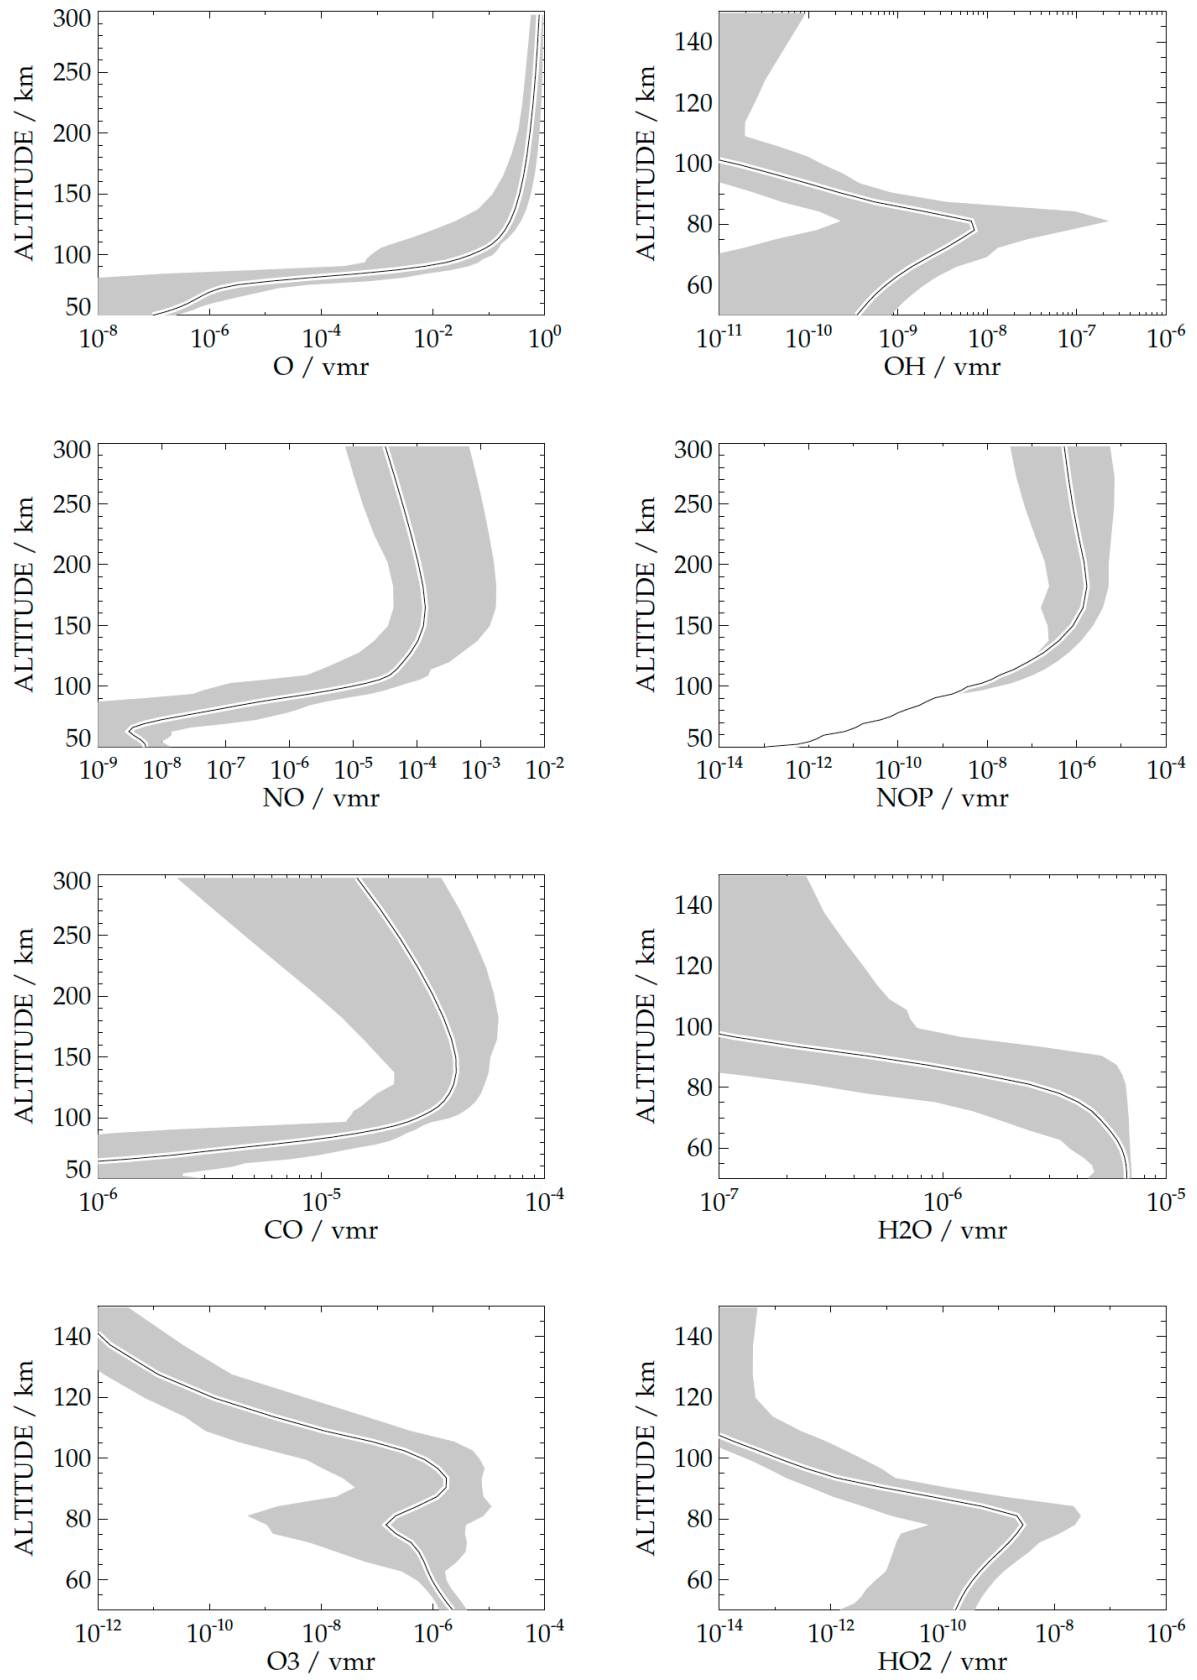


Figure 6: Climatology of atmospheric profiles in units of their volume mixing ratio

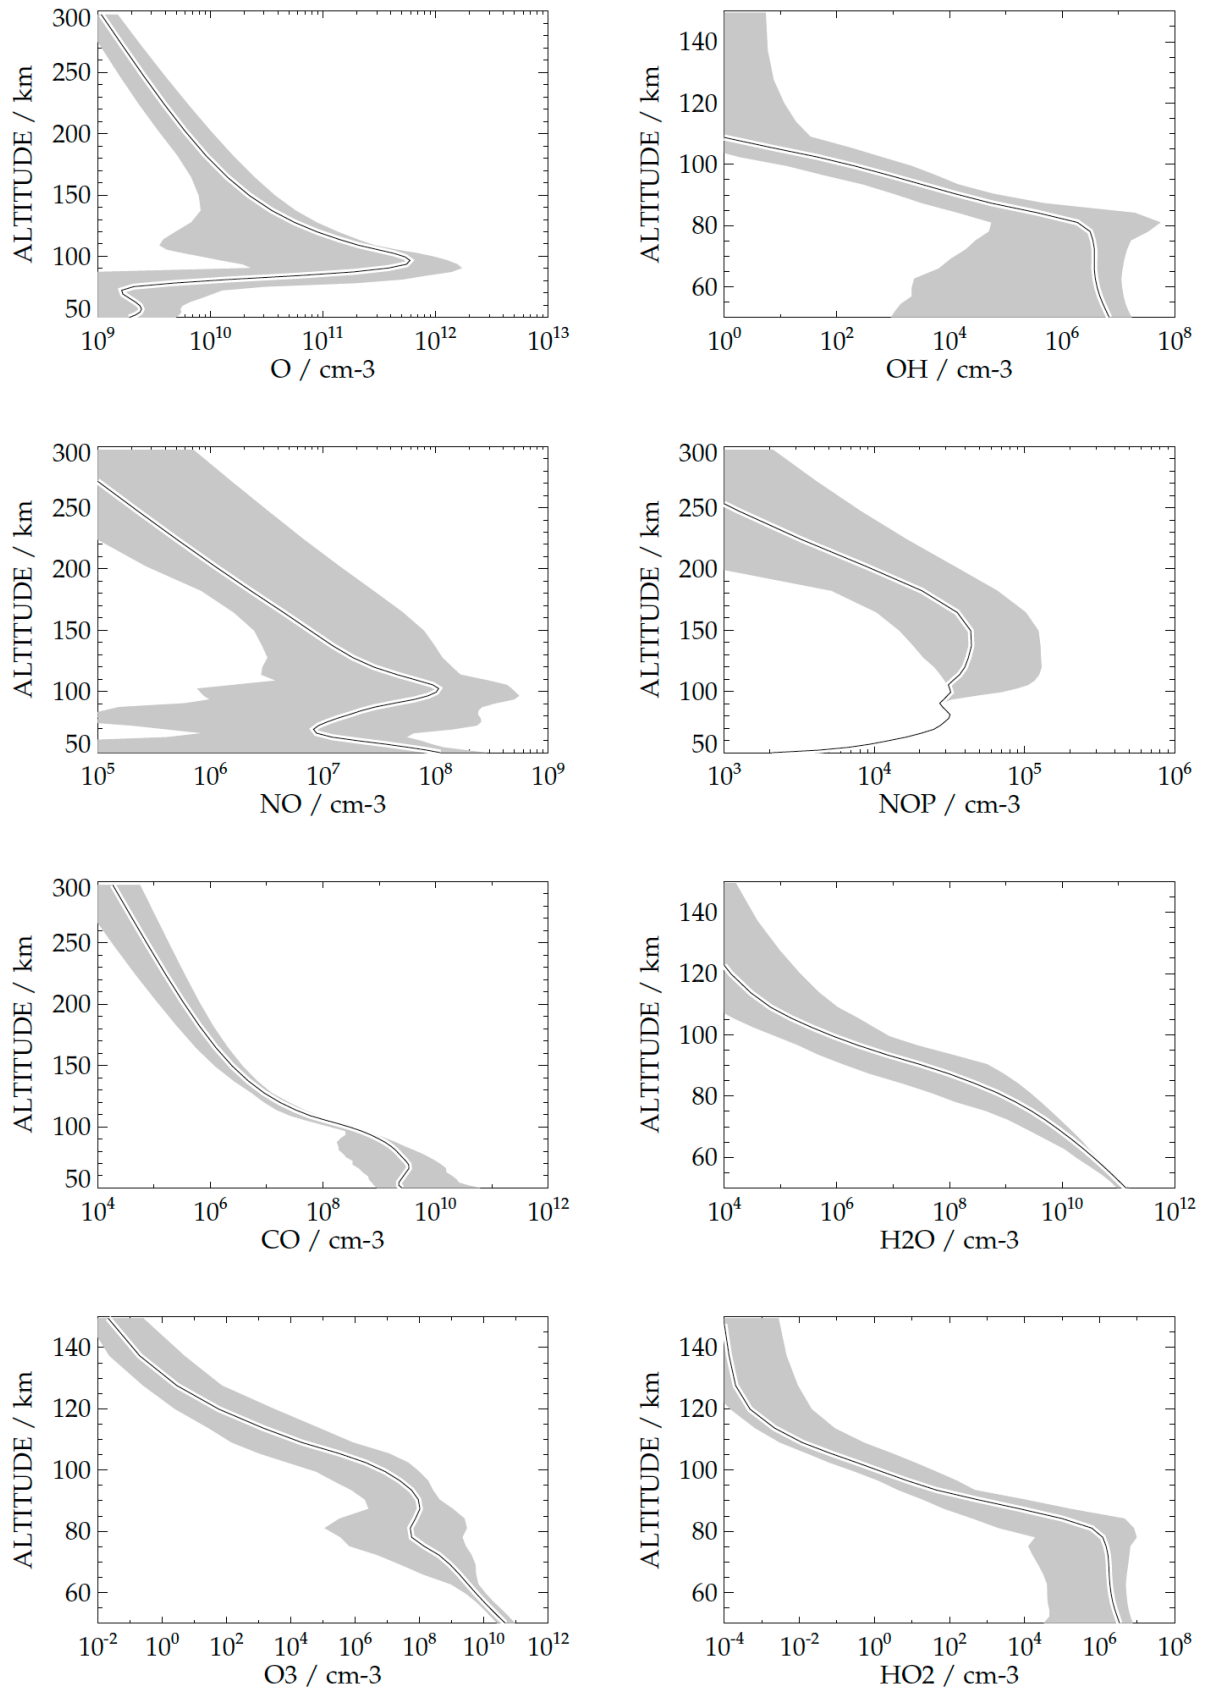


Figure 7: Climatology of atmospheric profiles in units of their number density (number of molecules per cm^{-3})

Remote Sensing of Atomic Oxygen

Atomic oxygen, despite being the major component of the atmosphere above ~120km, is not a conventional target species for atmospheric remote sensing. It exhibits several features, which make it particularly challenging to measure at a distant.

Measurement Frequencies

The first one is the fact that as a symmetric, single atom it doesn't feature any of the typical vibrational transitions that are exploited by infrared remote sensing instruments, or rotational transitions that are exploited by microwave instruments. Its electromagnetic emission lines are from two internal electrical transitions, listed in Table 4. Both of them are in the THz (and supra-THz) range, where technical difficulties imply that we are currently facing an absolute lack of heterodyne (aka. high spectral resolution) remote sensing detectors. This is referred to as the "THZ-Gap" [Moloney 2011] (Figure 8). For LOCUS, the technical aspect is addressed by the use of novel Quantum Cascade Laser diodes as Local Oscillators for the heterodyne frequency down-conversion.

Table 4: The only two electromagnetic emission lines from atomic oxygen

Transition	HITRAN Line Strength Parameter	Wave Number	Wavelength	Frequency
$O(^3P_1) \rightarrow O(^3P_2)$	1.117e-21	158.302980 cm-3	63 μ m	4.745THz
$O(^3P_0) \rightarrow O(^3P_1)$	9.577e-23	68.716470 cm-3	145 μ m	2.060THz




Figure 8: The THz-Gap is a frequency window ca. between 1-10THz, where current technology is not able to provide local oscillator sources powerful enough to pump a heterodyne frequency down-conversion mixer (which is the scientific requirements to retrieve vertical abundance profiles of a species from subtle features in their spectral line-shapes)

The only remote sensing measurement of an atomic oxygen line today was the measurements of the 63 μ m line with the far-infrared instrument CRISTA from the space shuttle. While ground breaking, the spectrometer resolution and signal-to-noise ratios did not allow a direct retrieval of the abundance profile of atomic oxygen. However, the CRISTA team have inferred the global distribution of atomic oxygen from the limb-extinction profile of the line centre down to altitudes of 120km, where the line centre becomes opaque. A comparison of this global map from two short Space Shuttle flights showed large discrepancies with MSIS model simulations (Figure 9).

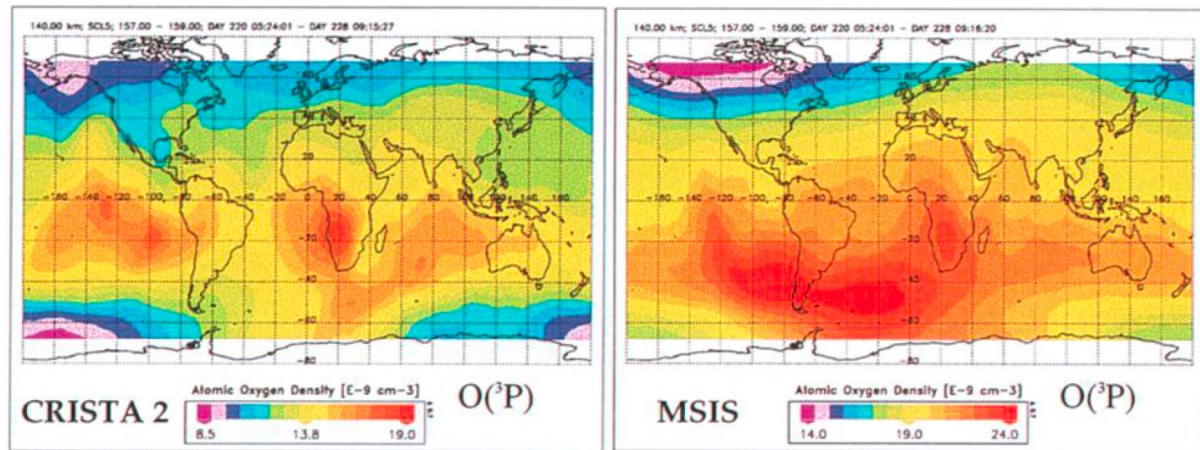


Figure 9: First detection of atomic oxygen from a remote sensing technique (far-infrared detectors CRISTA on the Space Shuttle, Grossmann et al, 2000) and comparison with the current state of knowledge from the MSIS model

Atmospheric Distribution

The second challenge is set by the vertical distribution of atomic oxygen. O is one of the very few species whose absolute abundance (in units of volume mixing ratio) increased with altitude, to reach its peak at the top of the atmosphere. Conventionally, atmospheric trace gases tend to phase out with altitude. For atmospheric limb-sounding, the technique that yields much improved vertical resolution over a nadir looking satellite instrument – this means that the majority of the gas is not necessarily found at the tangent point, but rather at higher altitudes along the line of sight. This is illustrated in Figure 10.




Figure 10: Locations of strongest emission regions along the line of sight for conventional species (left), and atomic oxygen(right) which - quite uniquely - is most prevalent at the top of the atmosphere

The main issue here is that the temperature also increased rapidly with altitude in the Thermosphere. Because the strength of the resonant emission from internal energy transition is a function of temperature, this means that the total spectral line measured along a line of sight across all of the upper atmosphere will be dominated by the contributions from the hot Thermosphere. This also means that the line centre becomes opaque as the satellite instrument scans down into the atmosphere, and the atmospheric path becomes longer. This generally starts to happen at ~120km of altitude. However – as discussed in the next section – there is the potential to extract information from lower layers by self-extinction in the line wing against the hot thermospheric background signal.

Opacity Profile from Self-extinction

The spectral line opacity for the two atomic oxygen emission lines as a function of altitude is shown in Figure 11. The stronger 4.7THz transition is shown to the left, and the weaker 2THz transition to the right. Spectral line plots show the observed Brightness Temperature (black lines) and the corresponding atmospheric transmittance (red lines). The transmittances show that the line centres become opaque below ~120km for the 4.7THz line, and below ~100km for the weaker 2THz line. The lower opacity threshold of the 2THz transition is consistent with the weaker spectral line strength at that frequency, which to some extent mitigates the impact of the high thermospheric temperatures.

However, looking at the atmospheric weighting functions (i.e. the Jacobians $K_n = \delta Y / \delta X$ that quantify how much the measurement vector Y changes at a given limb view n for perturbations of the atmospheric state profile X) we see that information is still present at altitudes below the respective line-centre extinction thresholds. This information comes from the non-opaque wings of the atomic oxygen spectral lines. In the case of the 4.7THz transition particularly, this weighting function information takes negative values (magenta

areas in the colour contour plots in Figure 11), meaning that an increased atomic oxygen abundance at these low altitudes leads to a reduction in the observed Brightness

Temperatures. This can only be explained by an increased self-absorption against the hot thermospheric background signal.

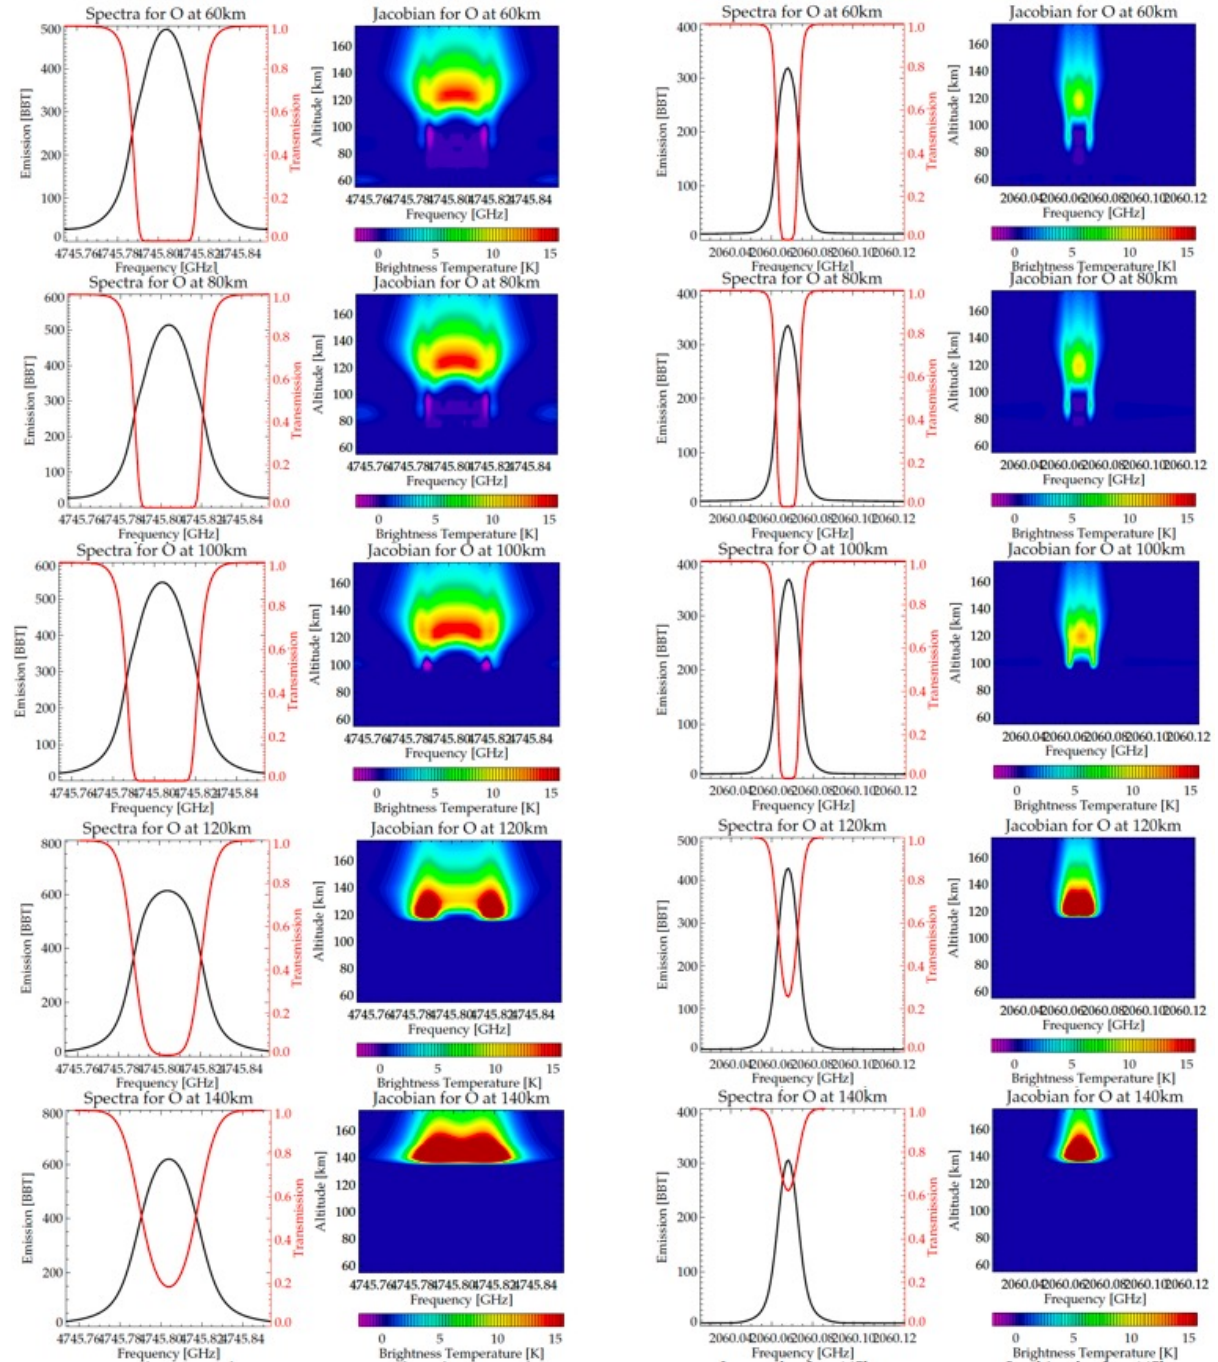


Figure 11: Limb-opacity profiles of atomic oxygen lines at 4.7THz (left two columns) and 2THz (right two columns). The spectral line plots show the observed Brightness Temperature (black line), and the atmospheric transmittance (red line). The line centre becomes opaque - seen in the red transmittance curves - at altitudes below 120km for the 4.7THz line, and

altitudes below 100km for the 2THz line. The colour contour plots show the atmospheric weighting functions for an observation along the line of sight at the respective tangent point altitude. They indicate that information on the abundance of atomic oxygen is present in the line wings of the spectral lines even at altitudes below the line-centre opacity thresholds. For the 4.7THz emission line in particular, this information is expressed in negative weighting functions (magenta), meaning that an increased abundance of atomic oxygen at low altitudes lead to a lower observed Brightness Temperatures, which can only be explained by an increased self-absorption against the hot thermospheric background signal.

The conclusion from this is that – despite the reversed abundance profile of atomic oxygen, with a dominant contribution from the hotter Thermosphere – it is possible to retrieve abundance information from atmospheric altitudes below the line centre opacity cut-off, so long as the spectral resolution is sufficient to resolve the artefacts of self-absorption in the shape of the spectral line wings. These are further illustrated in Figure 12, which traces the evolution of the shape of the wing of the spectral line when scanning to lower tangent point altitudes. It is clear from the Brightness Temperature spectra that these exhibit a sort of a ‘kink’ in their wings at the transition from the opaque centre-band and the transparent wing regions. As expected, this artefact moves down the line wings as the centre becomes more and more opaque. The magnitude of it, which can reach up to 20% of the top of the atmosphere spectral line at that frequency, means that if the spectral resolution of the measurement is high enough to resolve it, then it should be possible to extract altitude information of atomic oxygen even from altitudes below the line-centre opacity cut-off.

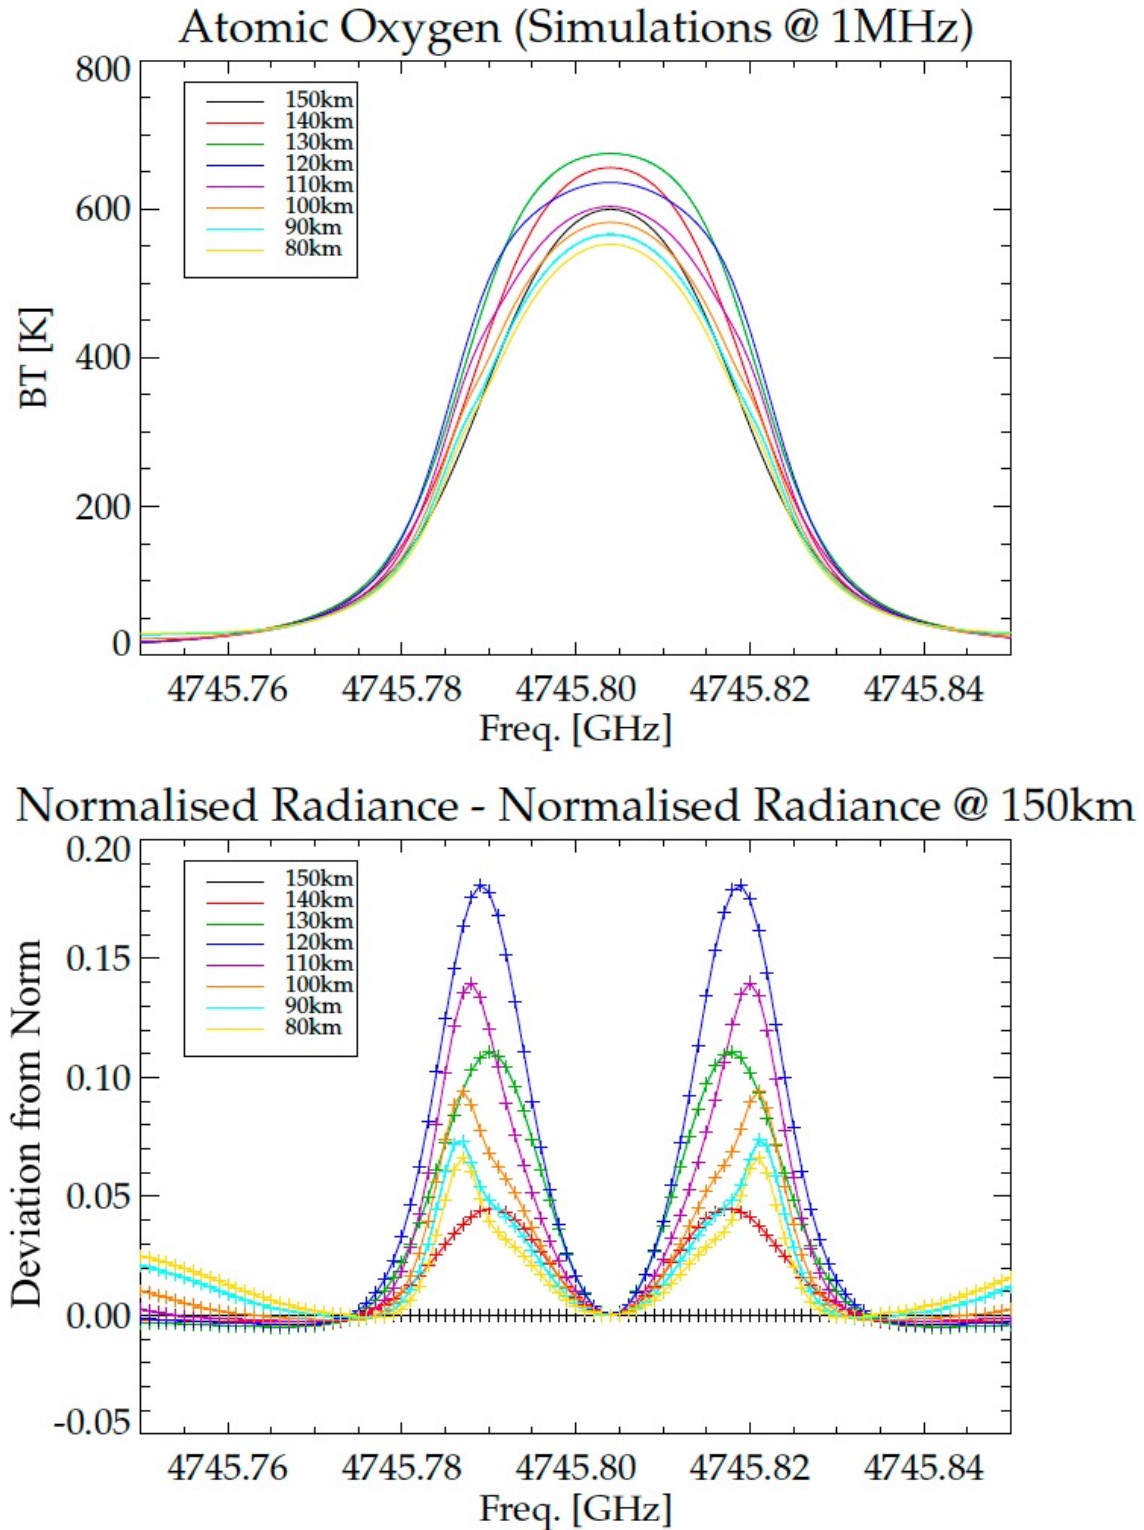


Figure 12: Artefacts in the shape of the spectral line shape of the atomic oxygen line in the altitude interval where self-extinction against the hot thermospheric background signal is observed. At 150km the spectral line displays a smooth Voigt line-shape profile. However, as the centre of the spectral line becomes opaque, a distinctive 'step', or 'bulge' appears in the shape of the line wing. The bottom panel shows the deviation of the line shapes at lower limb altitudes from the Voigt line at the top of the altitude. The deviation from Voigt reaches magnitudes of up to 20%, and moves away from the line-centre (i.e. down the wing) as the line-of-sight moves to lower altitudes (because a wider part of the line centre becomes opaque the lower we look).

Atmospheric Gradient and Variability

Finally, another aspect which makes remote sensing of atomic oxygen difficult, is the strong vertical gradient and diurnal variability of the species in the narrow altitude interval between 60km – 90km, where molecular oxygen starts to be phot-dissociated during day-time, but recombines instantly during nighttime. This is illustrated in Figure 13. This leads to some very distinct, sharp demarcation lines in the vertical regimes of atomic oxygen. The first one is the strong gradient region in volume mixing ratio between 80-100km, the other the high-variability band from 60-90km. This means that this region is likely to drive the requirements for vertical resolution. It's also possible that this could manifest in height-dependent artefacts in the retrieval process, so it's worth keeping in mind.

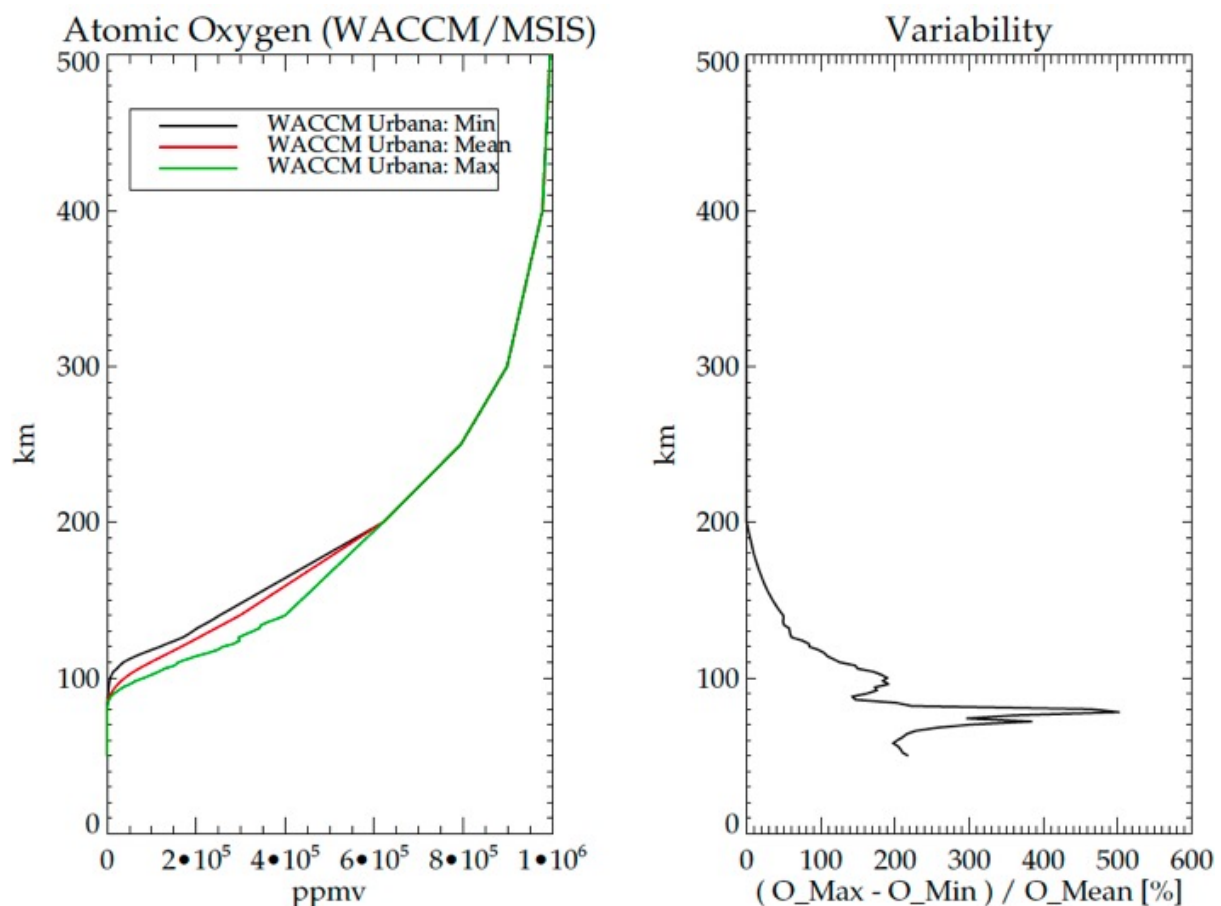


Figure 13: Variability of atomic oxygen vs. altitude. The largest relative variability is in the 60-90km region, where it can span 5 orders of magnitude of the local abundance

Simulated Measurements

Sideband Configuration

LOCUS is a double-sideband receiver. At the low pressures, prevalent in the upper atmosphere, spectral lines are very narrow, and therefore spectral confusion between lines in the image and mirror sidebands is highly unlikely. The spectral line width of the 4.7THz atomic oxygen line is shown in Figure 14. Pressure broadening becomes irrelevant at ~80km, above where the line width is determined by the local temperature via Doppler broadening. For a short interval between 80-110km the line width is more or less constant in balance, but beyond that additional altitude information is impressed on the spectral line from the respective Temperature profile.

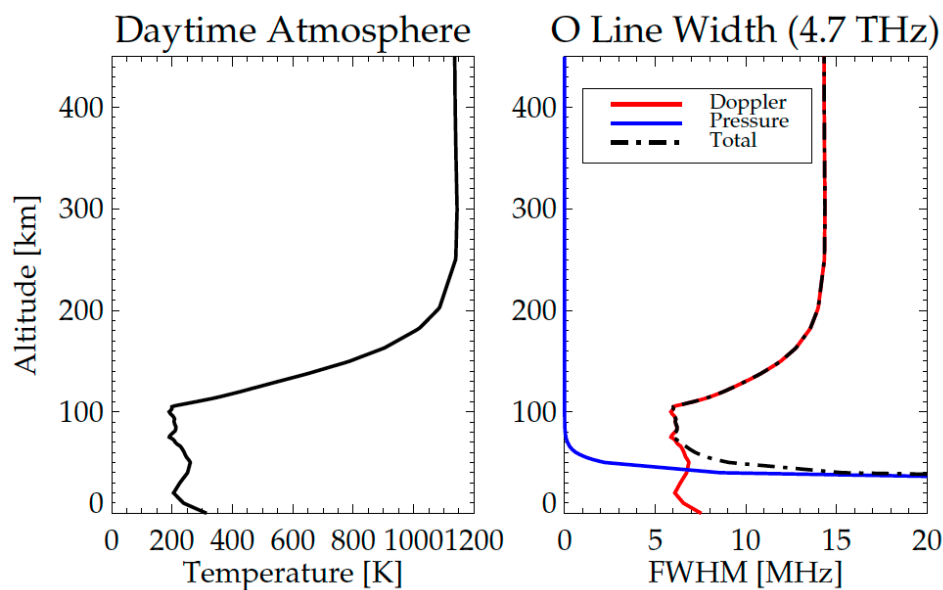


Figure 14: Spectral line-width of the 4.6THz atomic oxygen line. Pressure broadening becomes insignificant at 80km, above where the local temperature determines the line-width through Doppler broadening

The individual side-band contributions are shown in the consecutive paragraphs, together with the resulting Double Side-Band (DSB) spectra, both with and without noise. In some cases, DSB operation is exploited to collect several target lines in a single DSB spectrum (i.e. the two OH clusters in Band 2, or the multiple lines in Band 3). Because the spectral lines are spectrally disentangled, a single target retrieval need only consider a narrow frequency window that fully includes the spectral line in question. We have exploited this in

the single target retrieval simulations for computational efficiency. For multi target retrievals, or for retrievals of temperature, where the information content is spread over all molecules, the full measurement vector needs to be used.

Measurement Errors

For this initial phase of retrieval simulations, we study the impact of random (thermal) measurements errors only. This means that the errors on the measured spectral radiance (or Brightness Temperature in our case) is defined by the noise performance of the receivers alone. Systematic errors (aka. instrument errors, like pointing biases, calibration offsets, or uncertainties in the knowledge of the antenna pattern, or the spectral response function of the spectrometer) will be examined at a later stage. Where possible, we have however already included the ability to study these in the retrieval processor, e.g. by including the functionality to apply field-of-view convolutions, or instrument line-shape convolutions.

The noise performance of the detectors follows the radiometric formula:

$$NE\Delta T = \frac{T_{SYS}}{\sqrt{B \cdot \tau}}$$

NE ΔT is the noise equivalent Brightness Temperature, T_{SYS} is the system noise temperature, τ is the integration time, and b is the detection bandwidth (or in this case the spectral resolution).

The numbers for T_{SYS} are taken from existing radiometers (at the lower frequencies, where these exist), and extrapolated from published figures for the supra-THz channels at 4.7THz and 3.5THz. These extrapolations, based on the state of the art at the time of the ESA IOD study in 2014, are conservative. Are potential future implementation using the then

current technologies will likely improve on these figures. The values used for the retrieval simulations are listed in Table 5.

Table 5: Measurement errors derived from the receiver noise performance of the LOCUS channels (based on a 3 second integration time at 1 MHz spectral resolution)

Frequency	T_{sys} (DSB)	NEΔT (3 sec, 1 MHz)
4.7THz	80,000K	46K
3.5THz	20,000K	12K
2.1THz	20,000K	12K
1.1THz	7,000K	4K
0.8THz	2,000K	<2K

Simulated Spectra

Simulations of atmospheric spectra have been performed at every limb-view from 50km to 150km at 2km intervals. Because of the individual distribution profiles of each species, some molecules have stronger lines at low altitudes, whereas others (mainly atomic oxygen, NO and CO) are stronger at high altitudes. In the following sections, we have picked the 70km limb-view for illustrative purposes. It's a good compromise altitude at which all species have discernible spectral lines.

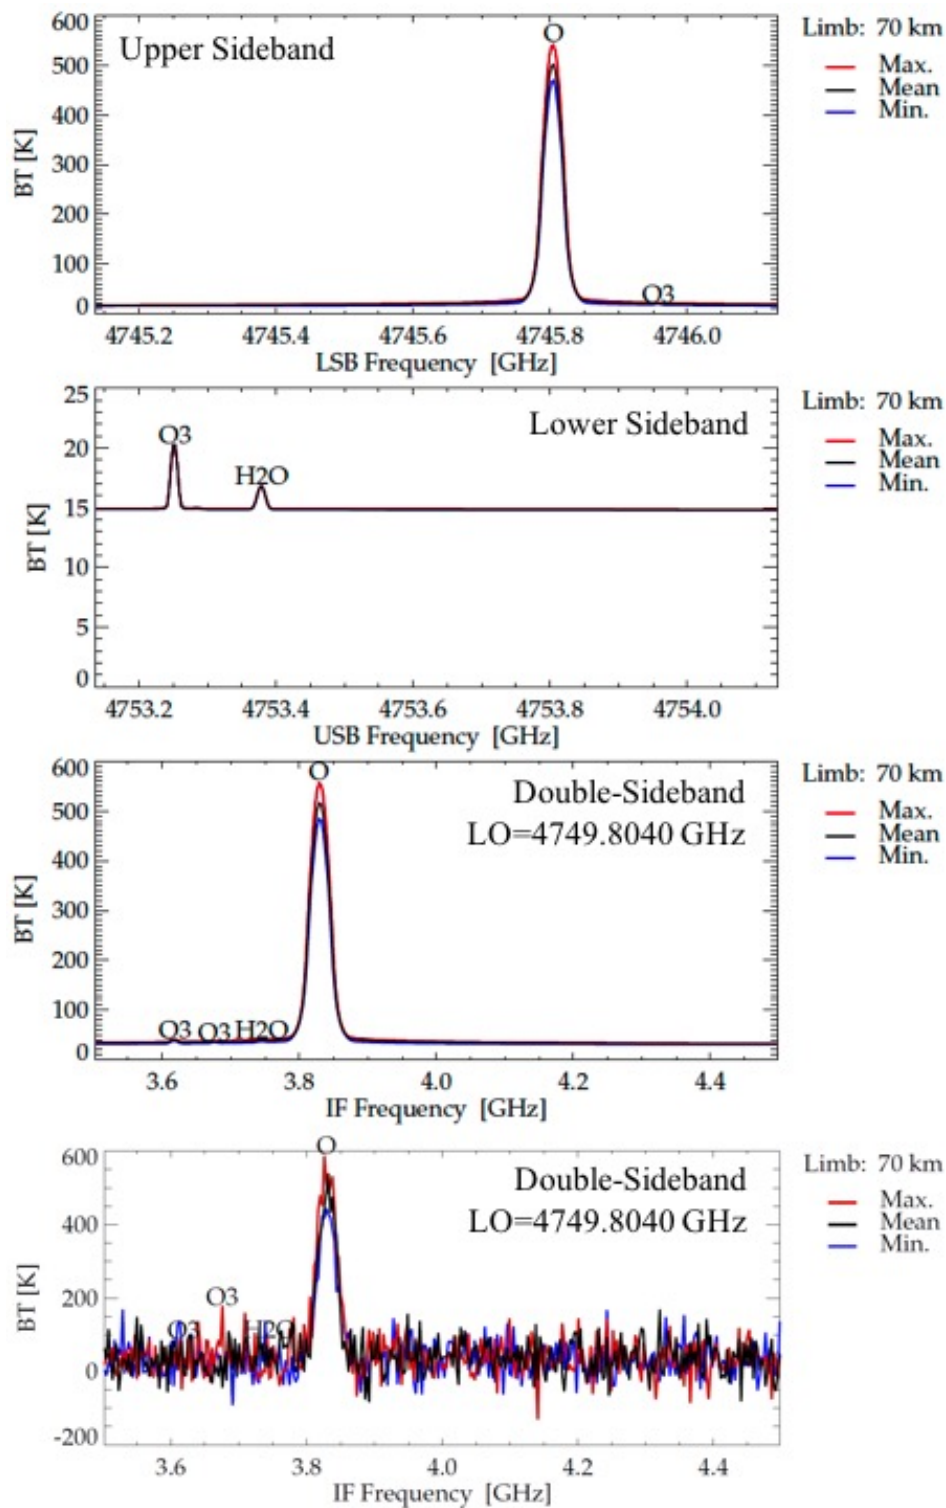
Band 1 – 4.75THz

Figure 15: Simulated atmospheric spectra for the 70km tangent view in the 4.7THz Band. The main target species here is atomic oxygen, located in the upper heterodyne sideband.

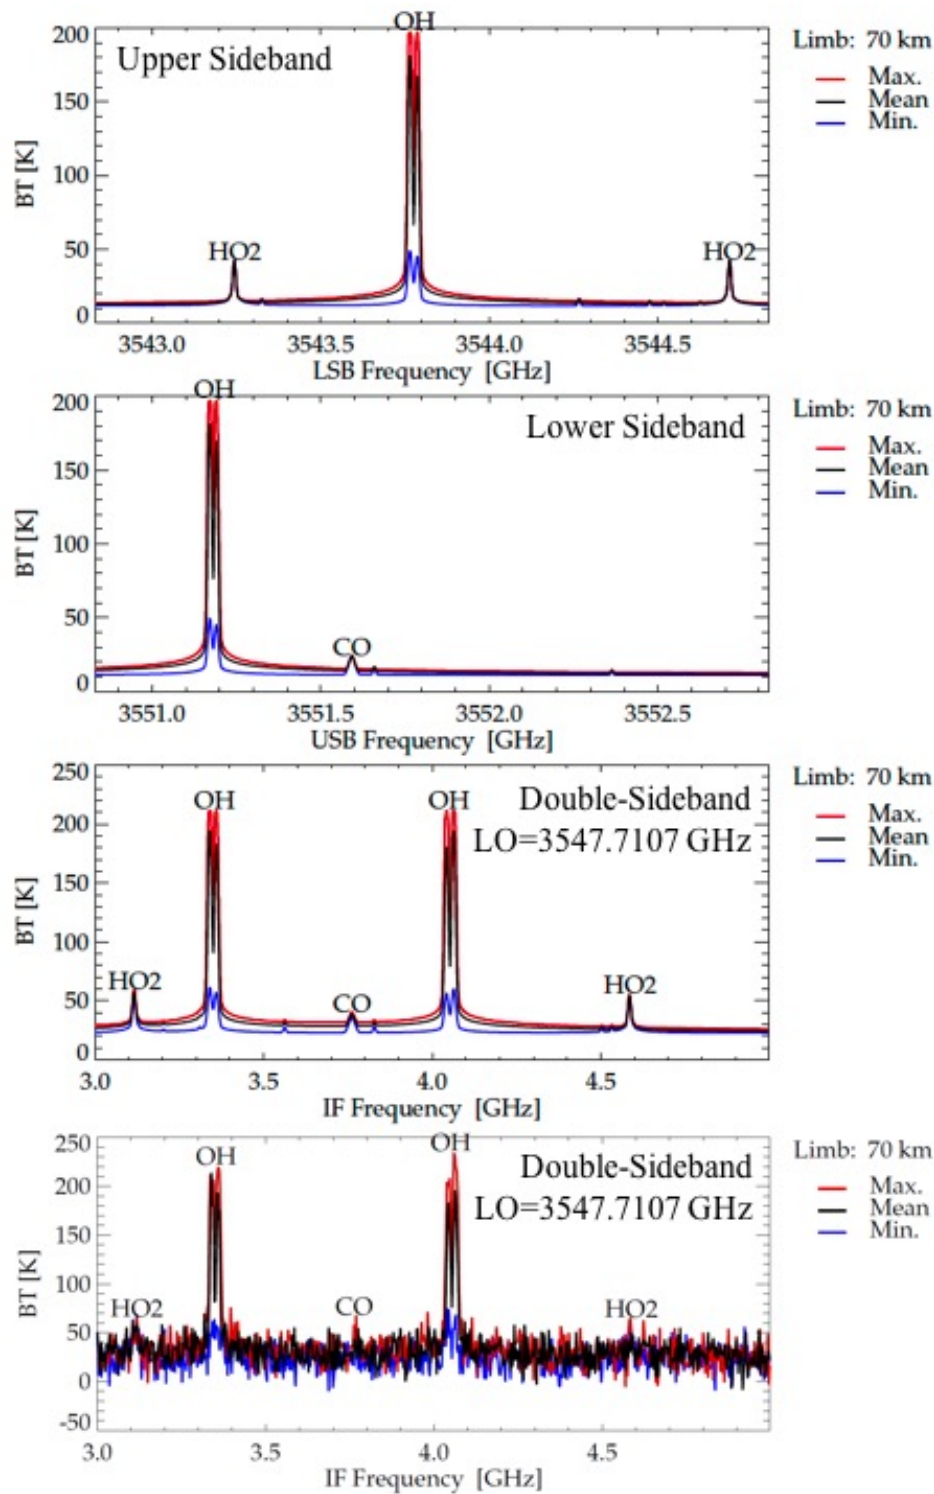
Band 2 – 3.55THz

Figure 16: Simulated atmospheric spectra for the 70km tangent view in the 3.5THz Band. The main target species here is hydroxyl (OH), with spectral line triplets in both of the heterodyne sideband. The secondary target Hydroperoxyl (HO2) has much weaker lines, so they are carefully placed in the gaps between OH by appropriate selection of the LO frequency. The CO line in the lower sideband will grow stronger at higher altitudes, offering an alternative CO measurements to the primary one in Band 3.

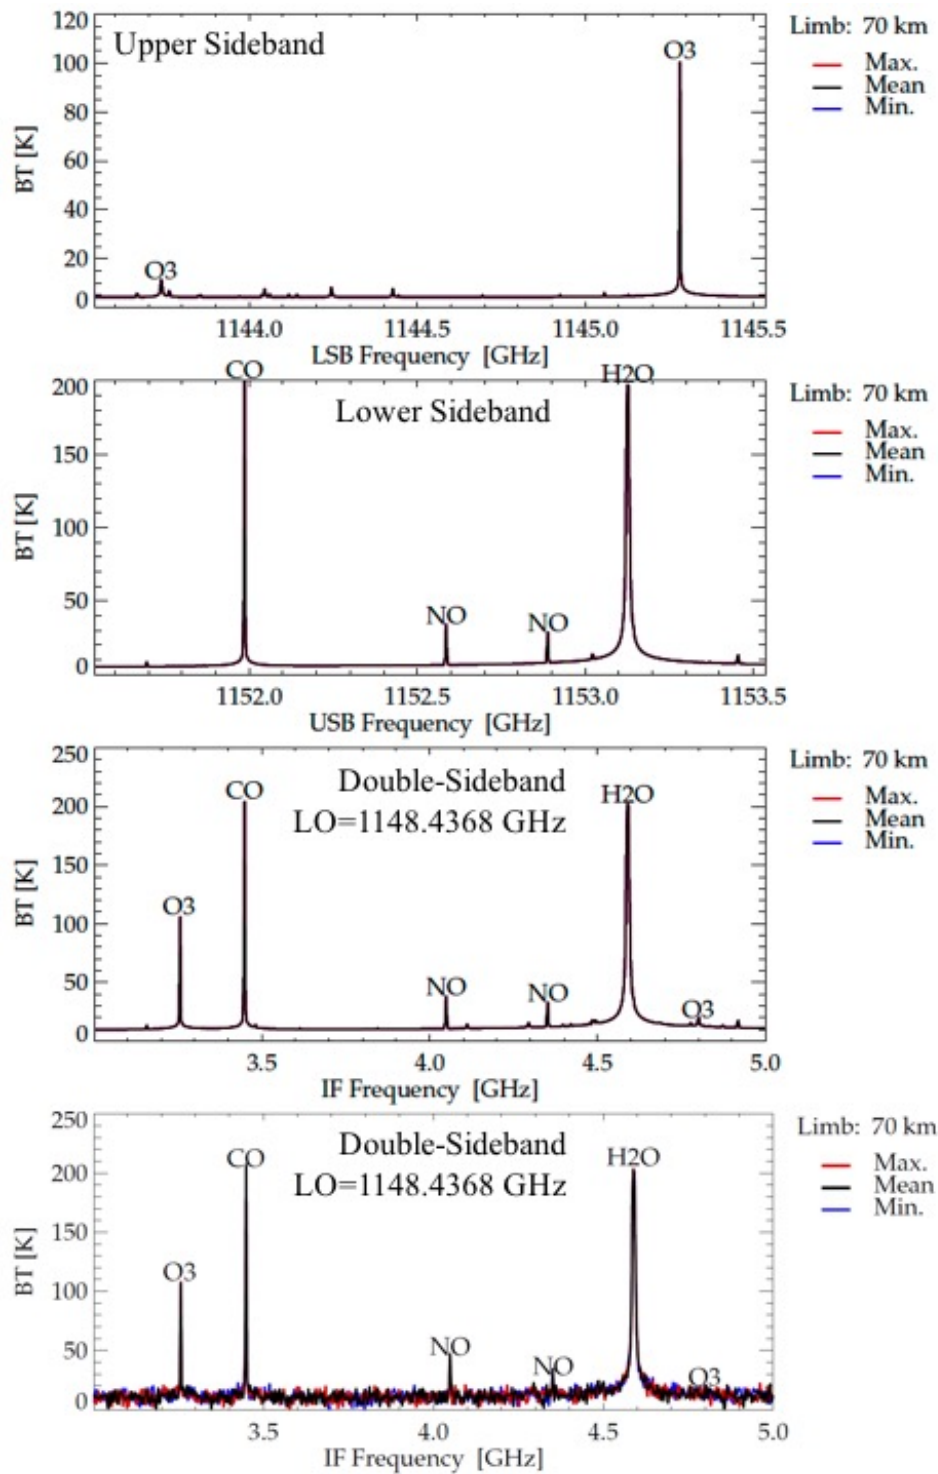
Band 3 – 1.15THz

Figure 17: Simulated atmospheric spectra for the 70km tangent view in the 1.1THz Band. The main target species here is nitric oxide (NO), with spectral lines in the lower heterodyne sideband. Other important target species in the lower sideband are CO and H₂O, while the upper sideband contains a strong O₃ line. The NO lines will get much stronger with altitude (as the NO abundance increased), where the possible interference from H₂O becomes negligible. At altitudes below the MLT, the H₂O line becomes pressure broadened and overpowering, but at

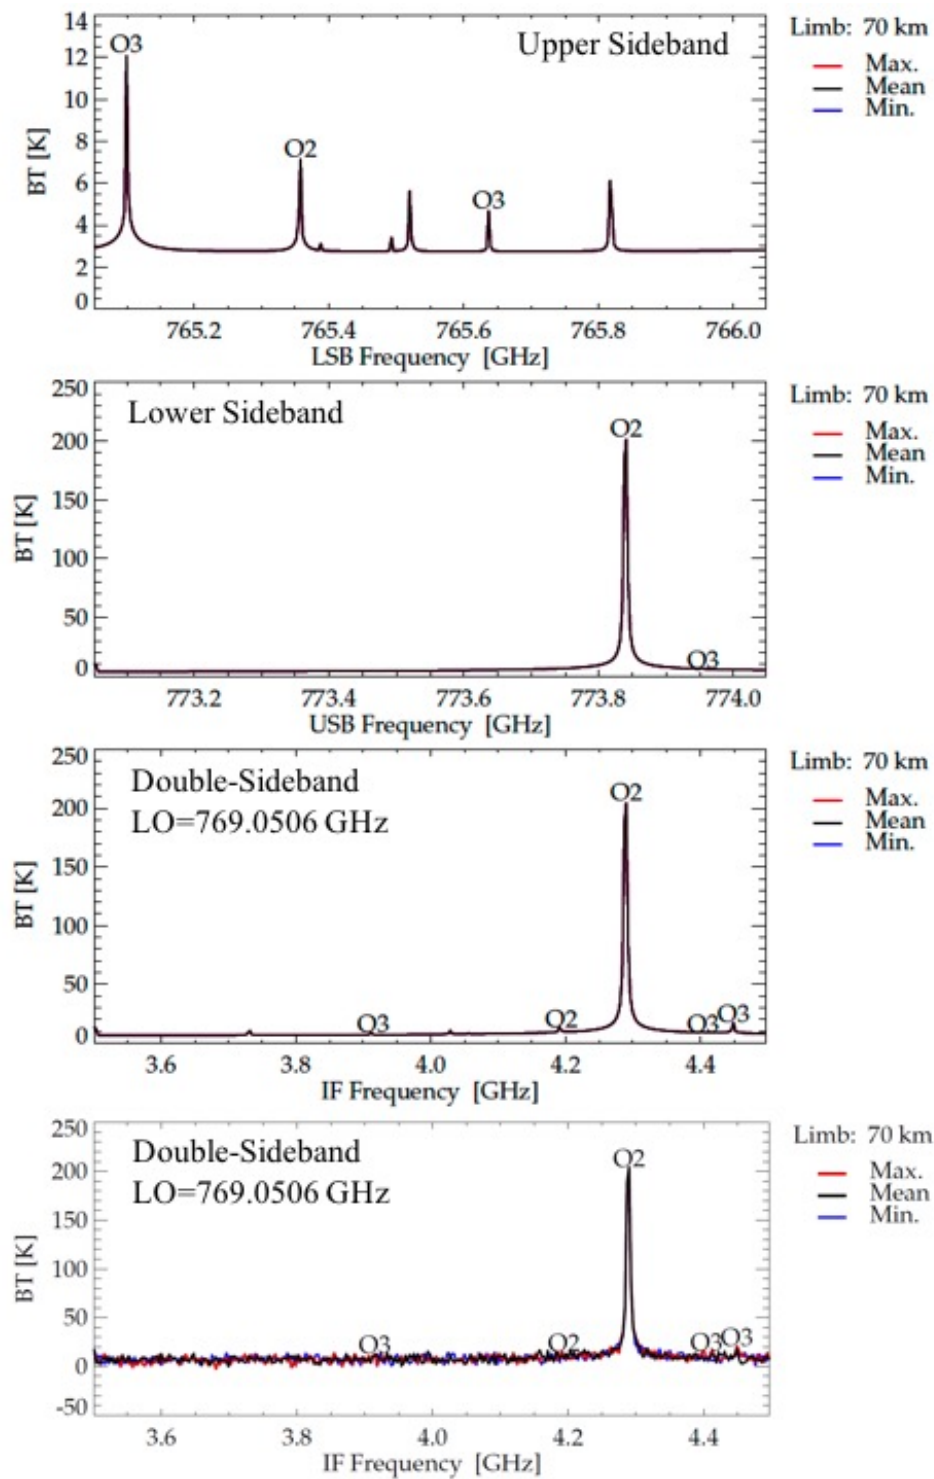
Band 4 – 0.77THz

Figure 18: Simulated atmospheric spectra for the 70km tangent view in the 770GHz Band. The main target species here is molecular oxygen (O₂) in the lower heterodyne sideband. Molecular oxygen is the source of information for temperature in the Mesosphere. A few ozone lines, particularly one in the Upper sideband, complement the primary ozone measurement in Band 3.

Single Species Retrievals

The fundamental retrieval simulations have been performed individually and independently for each product in the atmospheric state vector. This means that there are no cross-correlations between multiple retrieved products, which is a further complexity that will be explored in later experiments. Initially, we want to assess the best possible retrieval if we solely focus on a single species. For the single target retrievals, we limit the measurement vector to a micro-window that fully contains the spectral lines of the molecule under test, including its line-wing, but excluding any frequencies outside of the spectral line. For molecules with a single spectral line, this window is located around the line centre. For molecules with line tuples or triplets, it is located at the frequency centre-of-gravity of all spectral lines. Where a molecule has several spectral lines, or line tuples, that are a significant bandwidth apart, only a single micro-window is considered. In these cases, the predicted retrieval performances can be improved by a factor of $1/\sqrt{N}$ if additional N lines are included in the measurement. The frequency grid of the micro-window is monochromatic, with different frequency resolutions being tested in order to quantify their impact on the retrieval precision. The selection of spectral lines in the single target retrieval simulations is listed in Table 6.

Table 6: List of spectral lines included in the retrieval simulations, and potential additional lines in any of the LOCUS spectral bands whose inclusion in the measurement vector would statistically improve the retrieval results

Product	No. of lines in retrieval simulations	No. of additional lines in LOCUS bands	Reference
O – 4.5THz	1 line (4745.8039GHz)	N/A	Figure 19
O – 2.0THz	1 line (2060.0679GHz)	N/A	Figure 20
OH	1 line triplet (3543.7794GHz)	1 line triplet (3551.1852GHz)	Figure 21
NO	1 line tuple (1152.5862GHz)	1 line tuple (1152.8880GHz)	Figure 22

		2 weak tuples in Band 4 (773.02588GHz / 773.05064GHz)	
NO ⁺	1 line (1191.26GHz)	N/A	Figure 23
CO	1 line (1151.9854GHz)	1 line in Band 2 (3551.5923)	Figure 24
		1 CO isotope line in Band 4 (771.18386GHz)	
H ₂ O	1 line (1153.1268GHz)	N/A	Figure 25
O ₃	1 line (1145.2821GHz)	4 lines (1143.7382GHz, 1153.4545GHz, 765.10033GHz, 773.95319GHz)	Figure 26
HO ₂	1 multiplet (3544.7130GHz)	1 multiplet (3550.0478GHz)	Figure 27

Measurement Vectors

The measurements vectors and associated measurement errors for the single-target retrieval simulations are shown in the following plots. The measurement errors are the noise equivalent Brightness Temperatures (NE Δ T) of the detectors based on the receiver noise performance (Table 5). A spectrum is shown for every tangent altitude of the measurement scan sequence, which ranges from 50km to 150km at 2km intervals.




Figure 19: Measurement vector for the single target retrieval of O at 4.7THz. The amplitude of the spectral line increased with tangent altitude due to the inversion in the abundance profile of O (i.e. there is substantially more O at higher altitudes)




Figure 20: Measurement vector for the single target retrieval of O at 2.0THz. The amplitude of the spectral line increased with tangent altitude due to the inversion in the abundance profile of O (i.e. there is substantially more O at higher altitudes). The 2 THz line is weaker and narrower than the 4.7THz line.

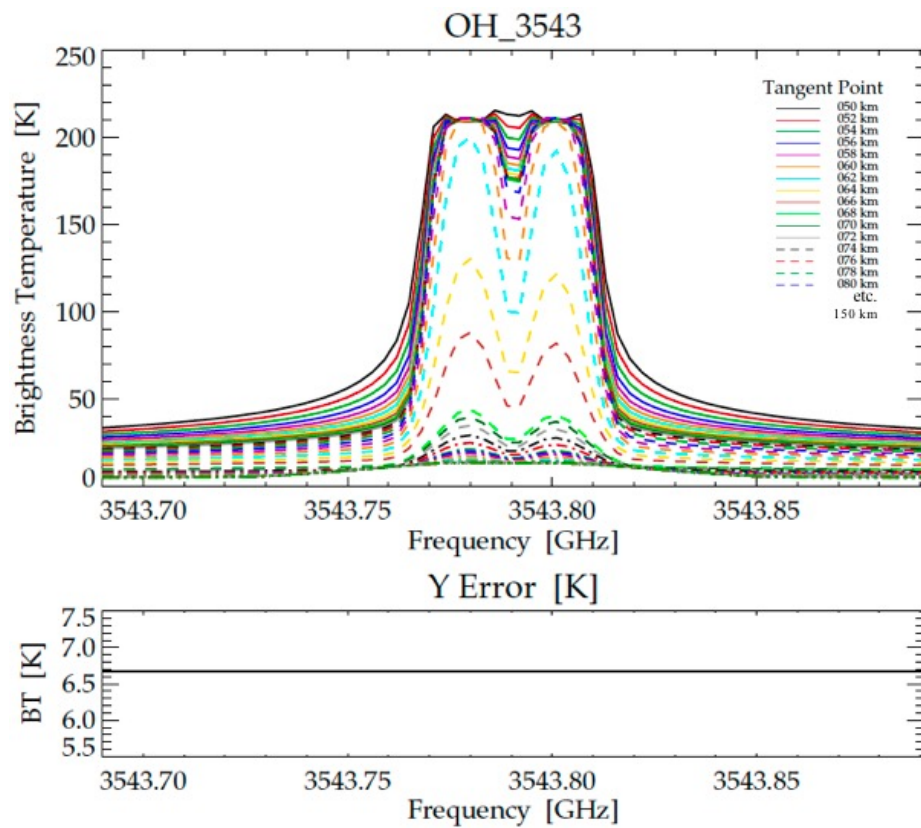


Figure 21: Measurement vector for the single target retrieval of OH at 3.543 THz. There is a symmetric OH feature 7.4 GHz beyond this window.




Figure 22: Measurement vector for the single target retrieval of NO. There are another 2 NO lines 300 MHz beyond this window.

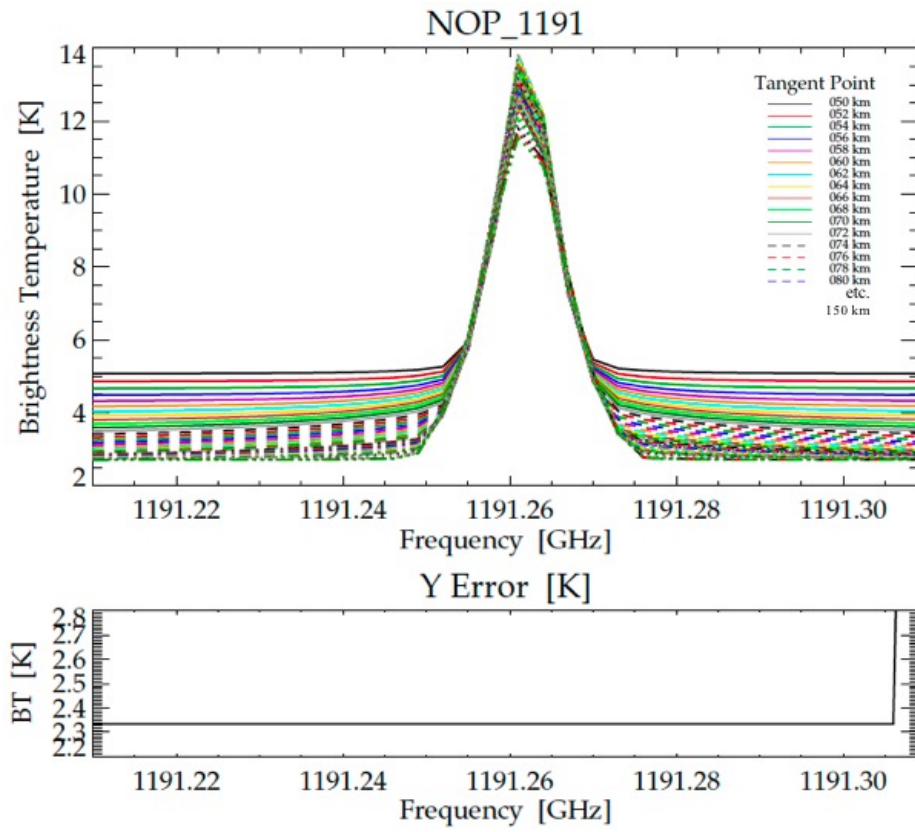


Figure 23: Measurement vector for the single target retrieval of NO^+ . Note that the NO^+ is considerably weaker than NO . This is due to the much-reduced abundance of the ion,




Figure 24: Measurement vector for the single target retrieval of CO

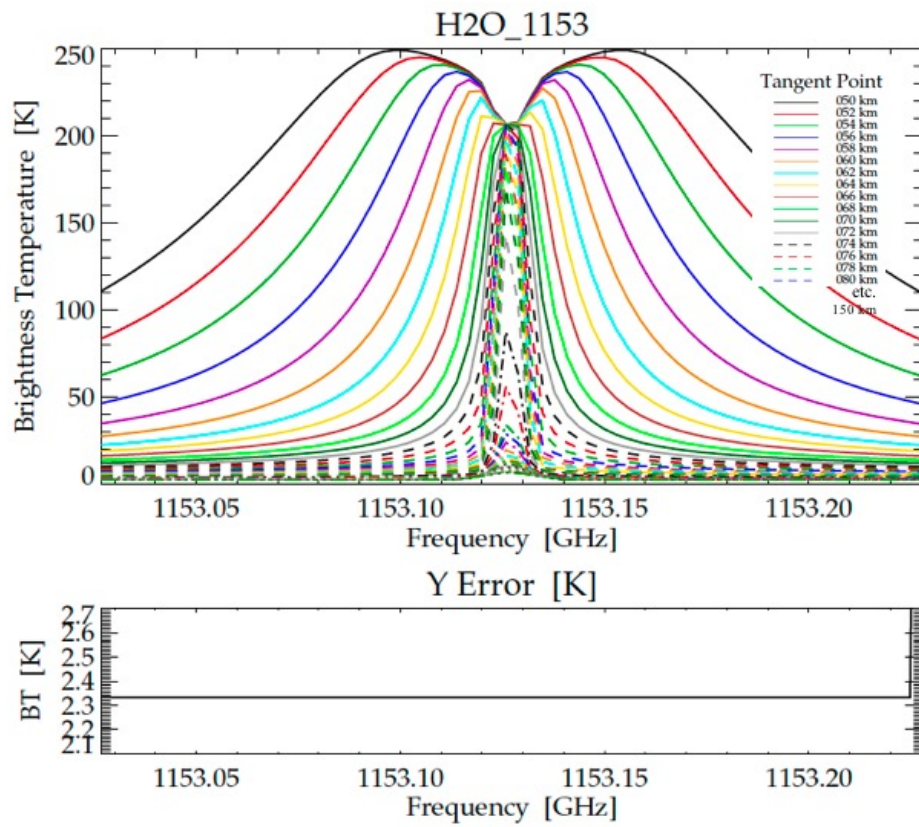


Figure 25: Measurement vector for the single target retrieval of H_2O




Figure 26: Measurement vector for the single target retrieval of O_3

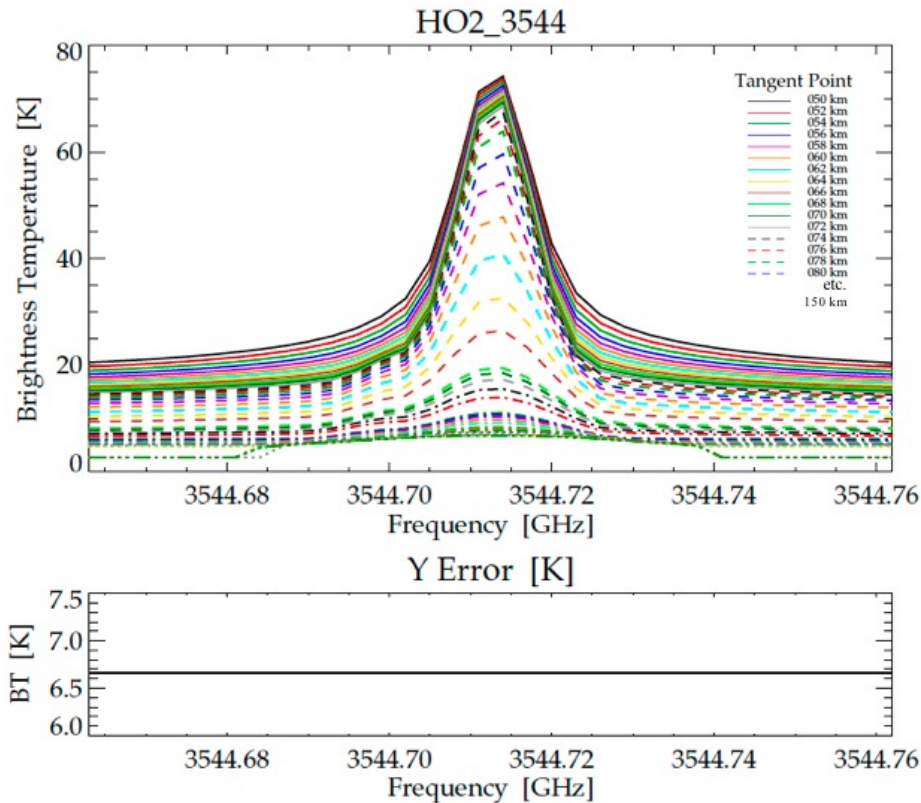


Figure 27: Measurement vector for the single target retrieval of HO₂. Note that the HO₂ spectral line is significantly weaker than those of other gases.

Atmospheric Weighting Functions

The atmospheric weighting functions (or Jacobians) show the response of the measurement to perturbations in the atmospheric state at various altitudes. This is dependent of the abundance profile of the target species, but also the overall opacity of the atmosphere. If e.g. the spectral line from a certain species becomes opaque – either because there is too much of the molecule in the path for the signal to pass through to the observer, or because another species blacks out the electromagnetic spectrum at that frequency – then changes in the distribution of said species along the line-of-sight would not be detectable as a change in Brightness Temperature at the observer. The weighting functions are an important component in the calculation of the retrieval diagnostics. In their own right, they offer an indication of the altitude range over which a given species can be observed in the atmosphere.

We compute weighting functions for every view angle (i.e. observed tangent point) in the limb-scan. In the base configuration, that would be every 2km. At every tangent point, the atmospheric profile is perturbed at every point of the atmospheric model grid in turn. The resulting colour contour plots show how brighter (in units of Brightness Temperature) the measured spectrum becomes, when there is more gas in the atmosphere. Because of the characteristic geometry of atmospheric limb-sounding, with the line of sight piercing the atmospheric layers tangentially, like a needle piercing the shells of an onion, there is never any response from below the respective tangent altitude. This is because the line of sight doesn't intersect with any layers below the tangent layer. Normally, the strongest response is from the tangent layer itself, because it contributes the longest segment to the line of sight. Layer in the atmosphere above are still sampled, but at decreasingly steep angles and decreasing relevance. Atomic oxygen is an exception, because of the large amount of O present in the Thermosphere, so for O the strongest contribution to the weighting functions comes from the upper atmospheric layers.

The following figures show a selection of weighting functions at representative tangent point altitudes of 80km to 150km in 10km steps. These were done for a 10km retrieval grid, which means that perturbations are triangular functions with a 10km half-width.

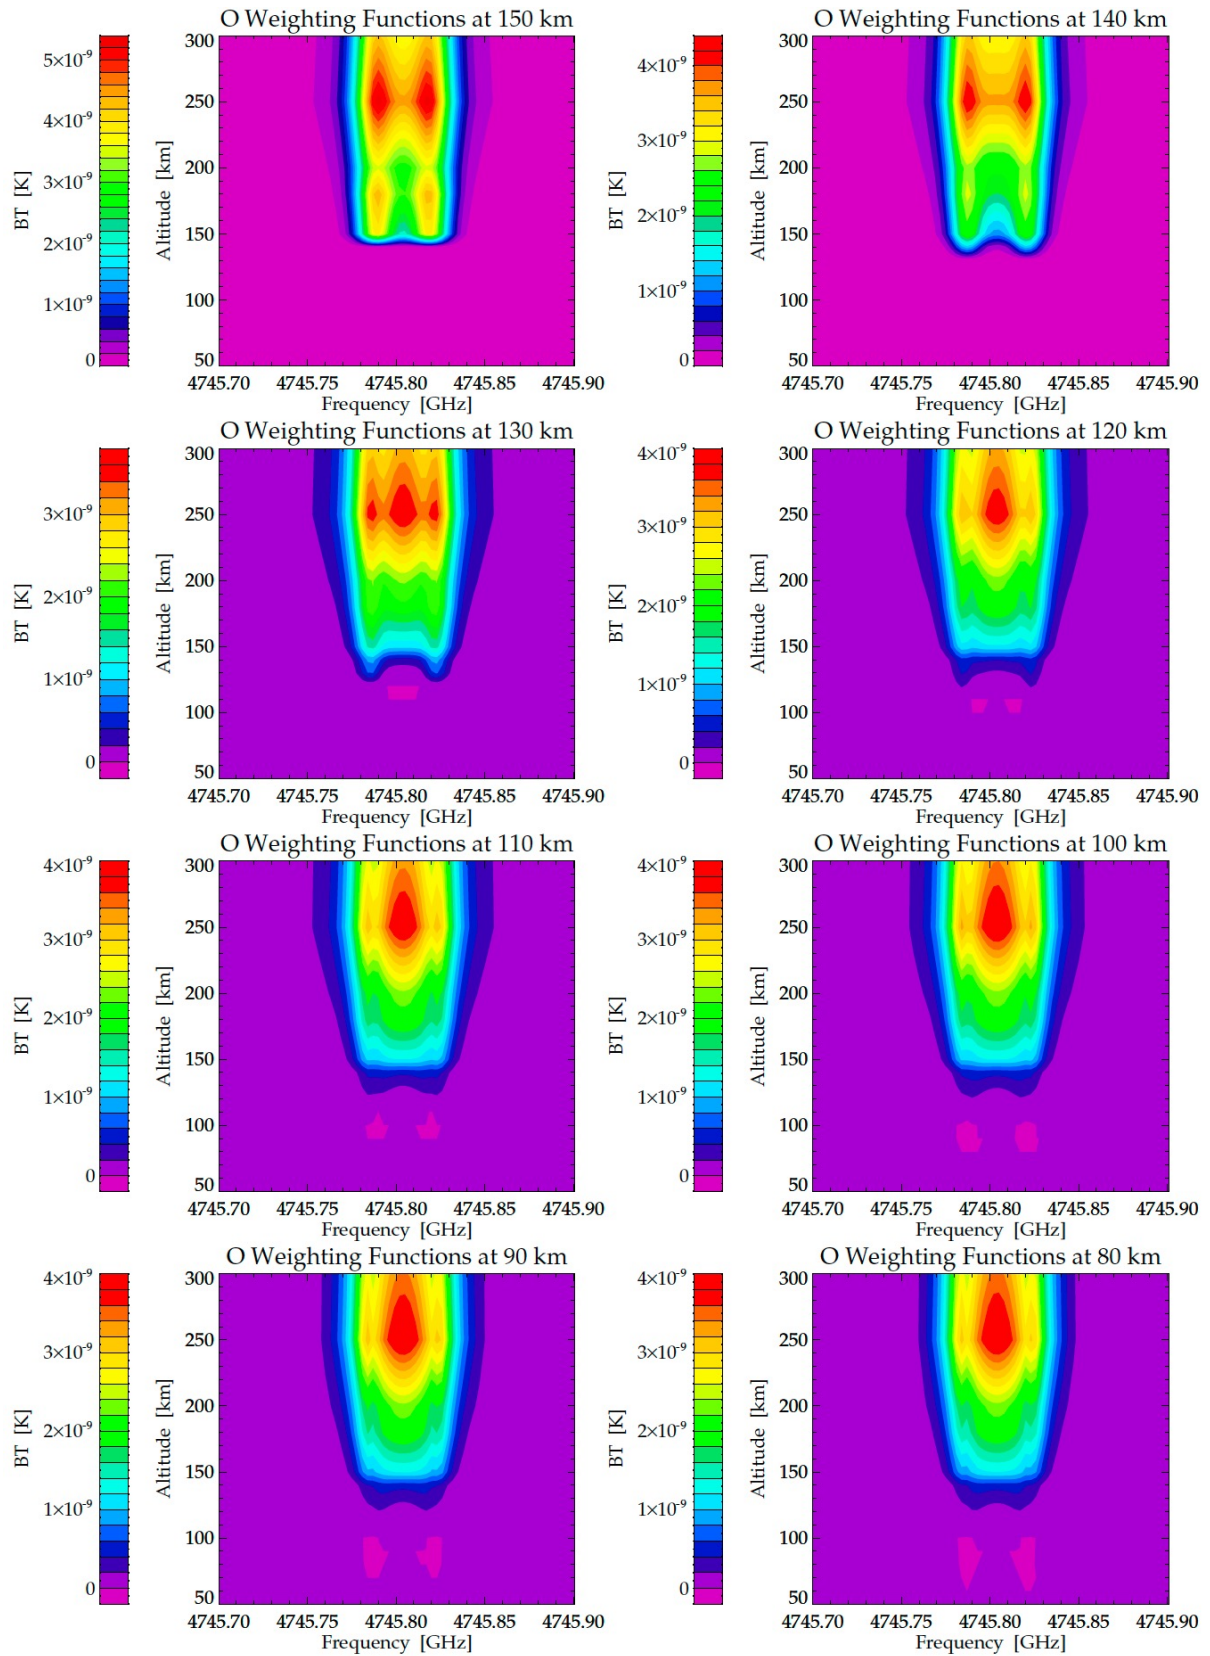


Figure 28: Atmospheric weighting functions (Kn) for O at 4.7 THz

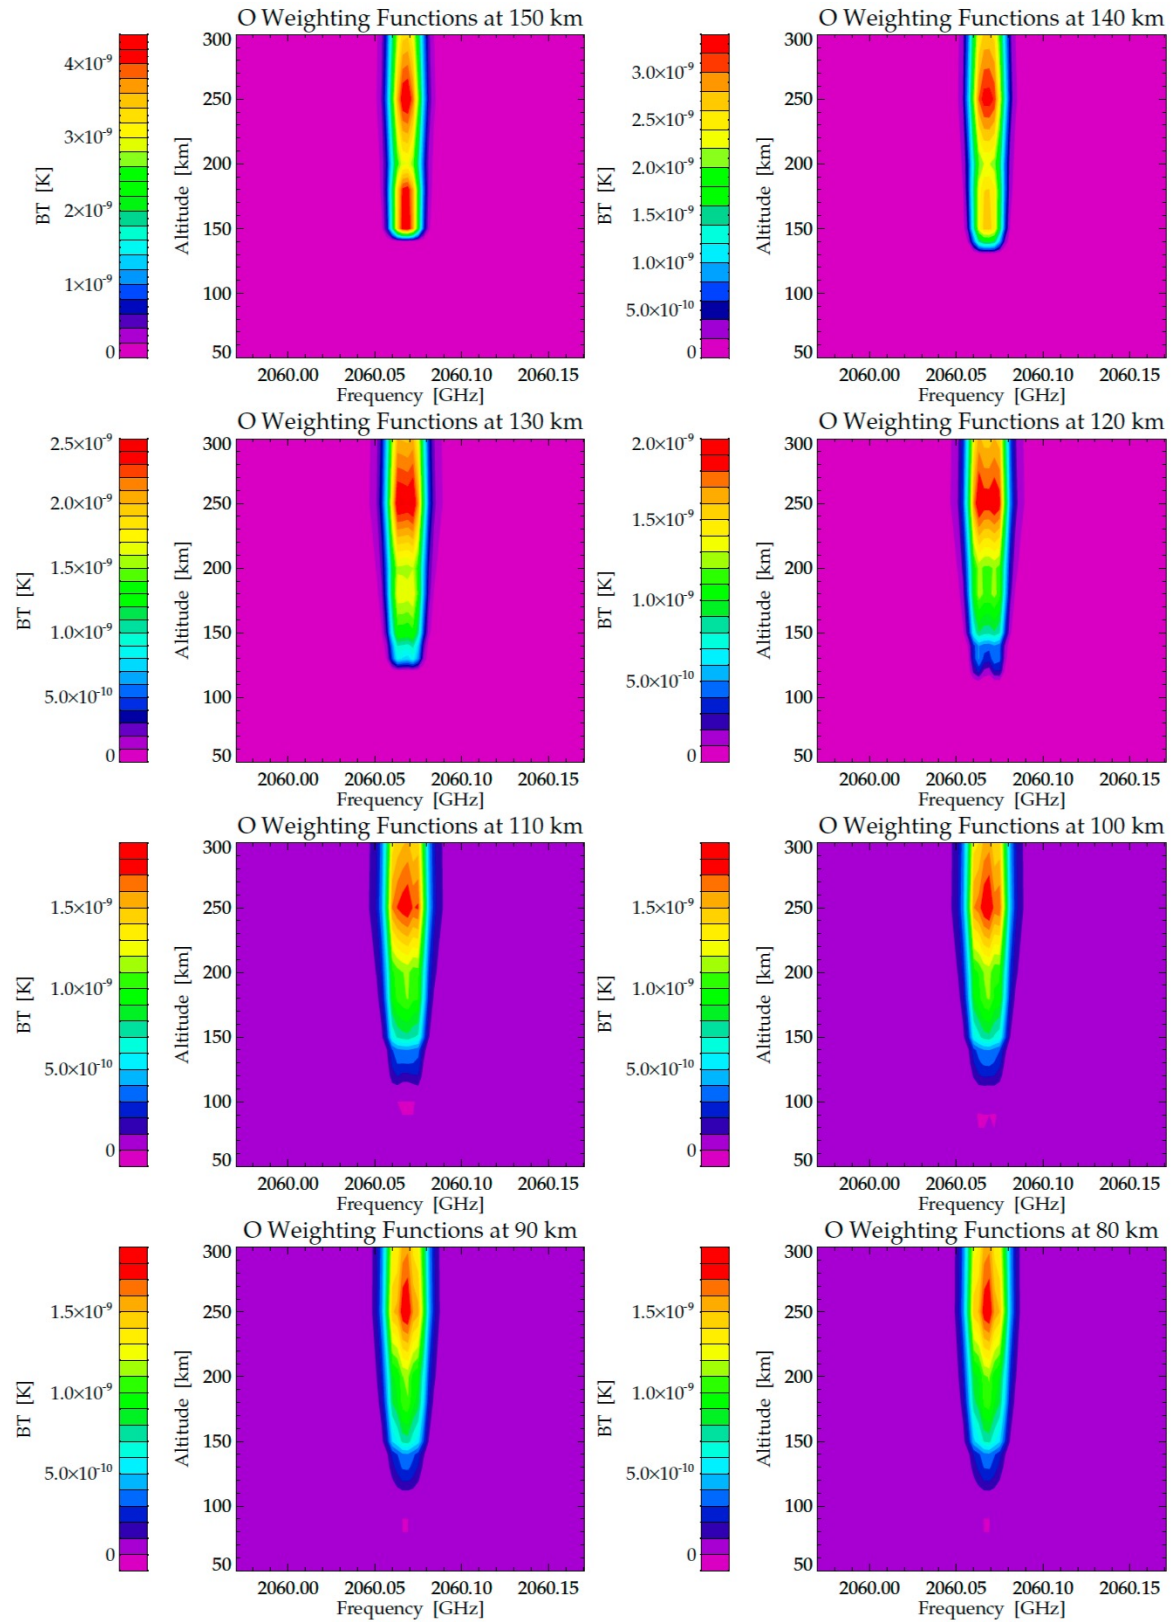


Figure 29: Atmospheric weighting functions (Kn) for O at 2.0THz

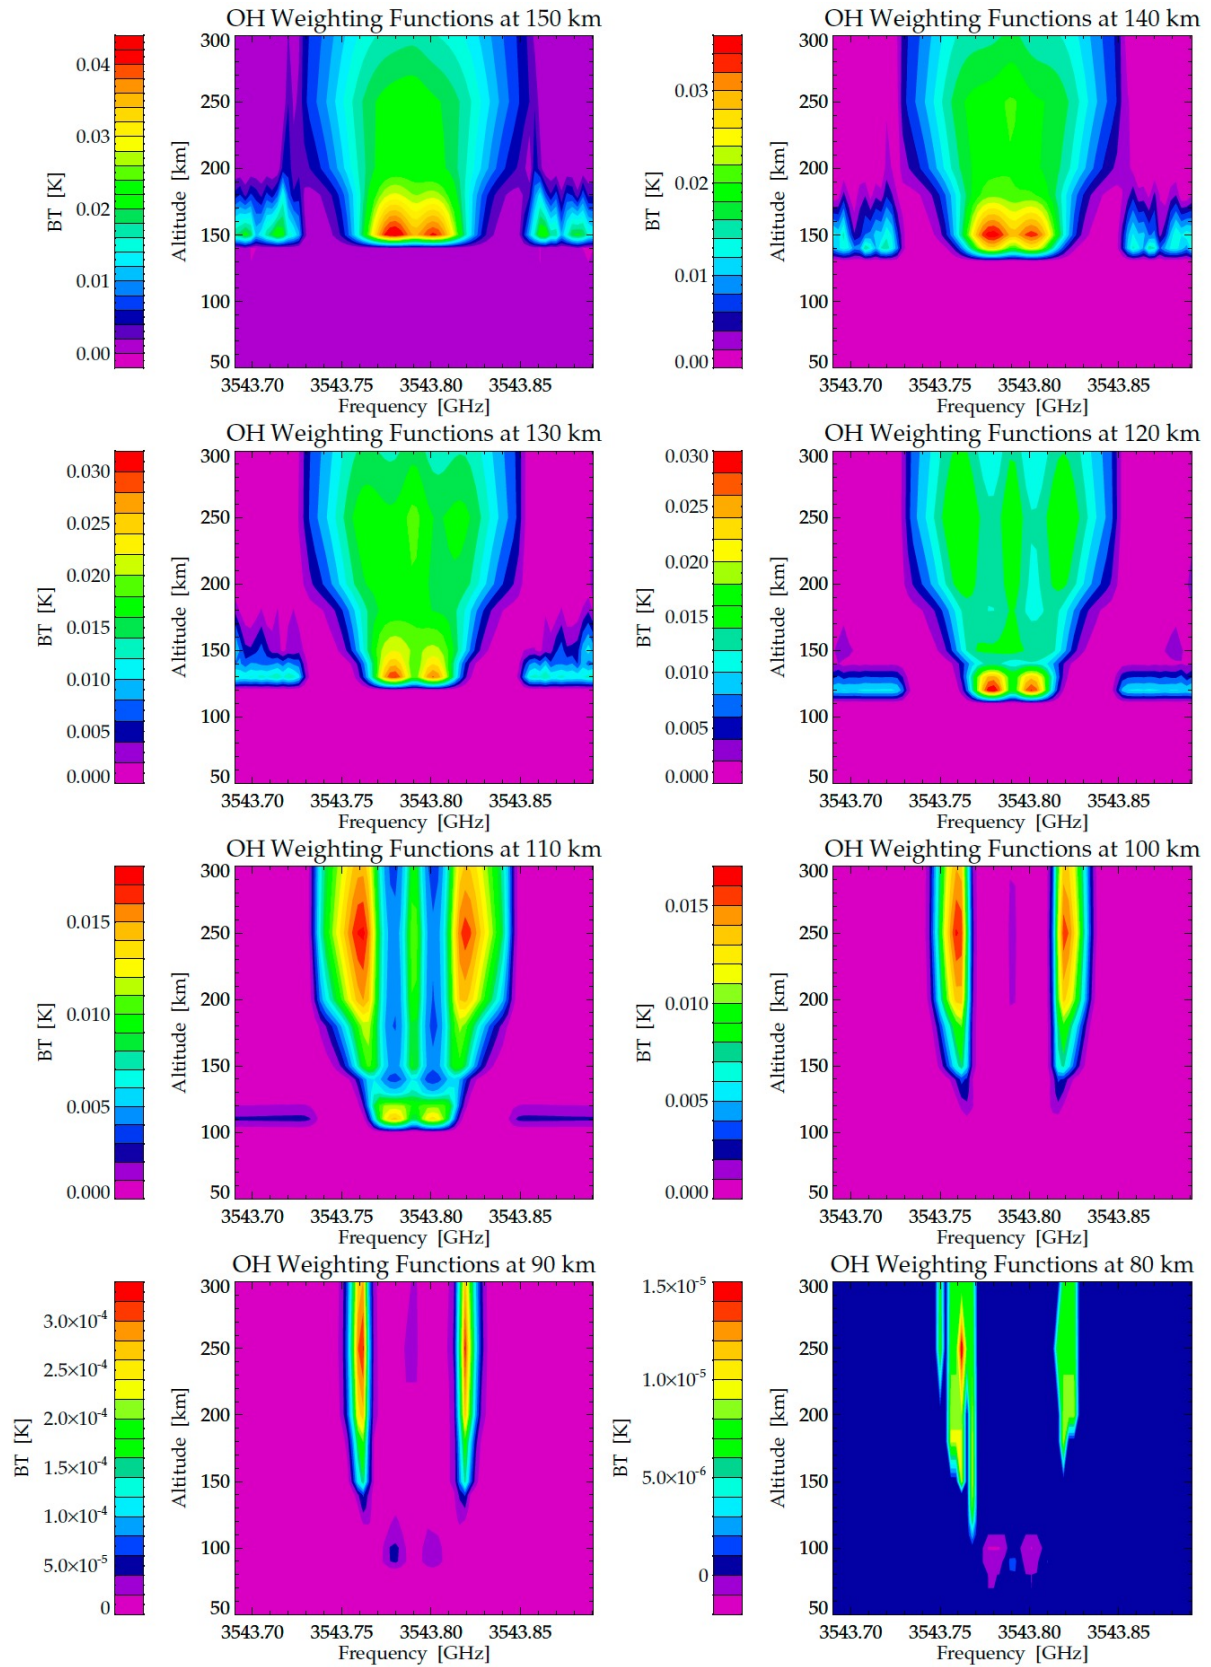


Figure 30: Atmospheric weighting functions (Kn) for OH. For maximal abundances (as in this example) the OH line centres become opaque at low altitudes, but information is still present in the line wings.

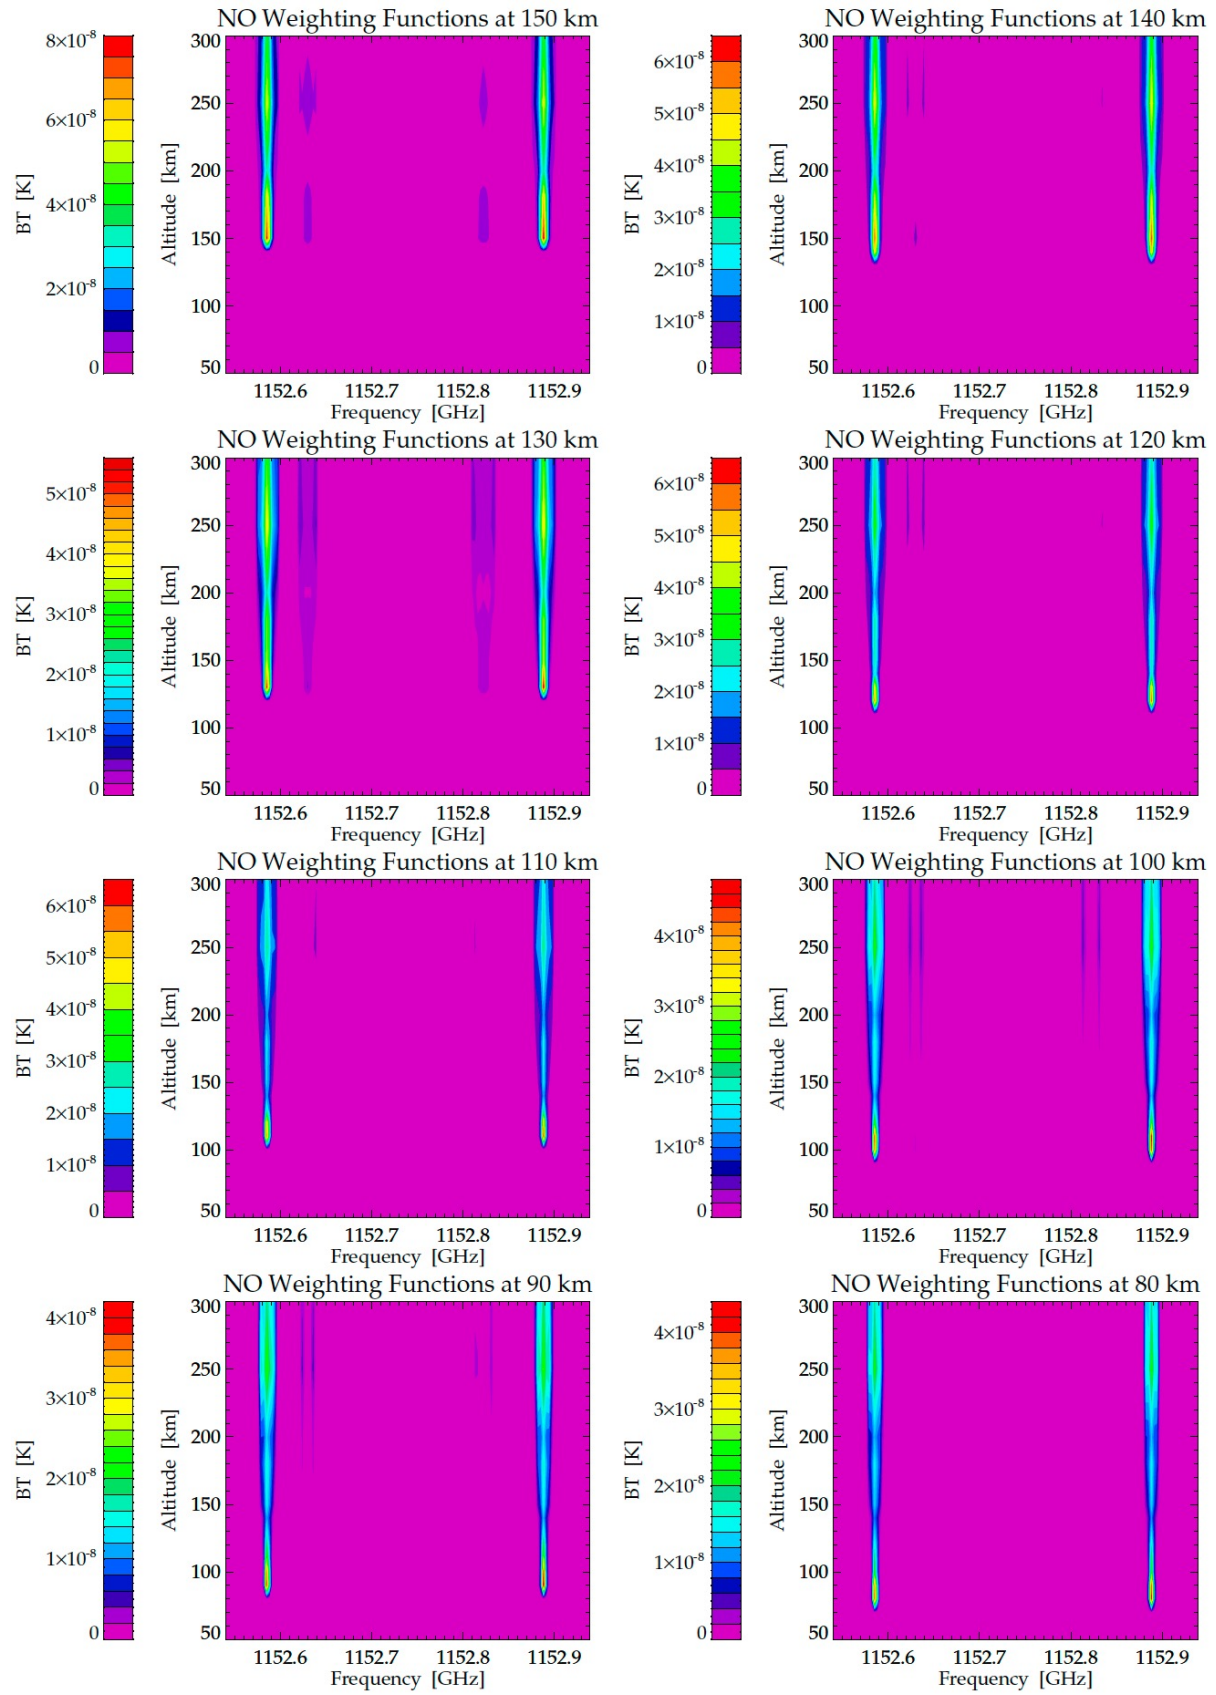


Figure 31: Atmospheric weighting functions (Kn) for NO. Contributions from minor NO lines in the vicinity of the main features are visible between 100km and 140km.

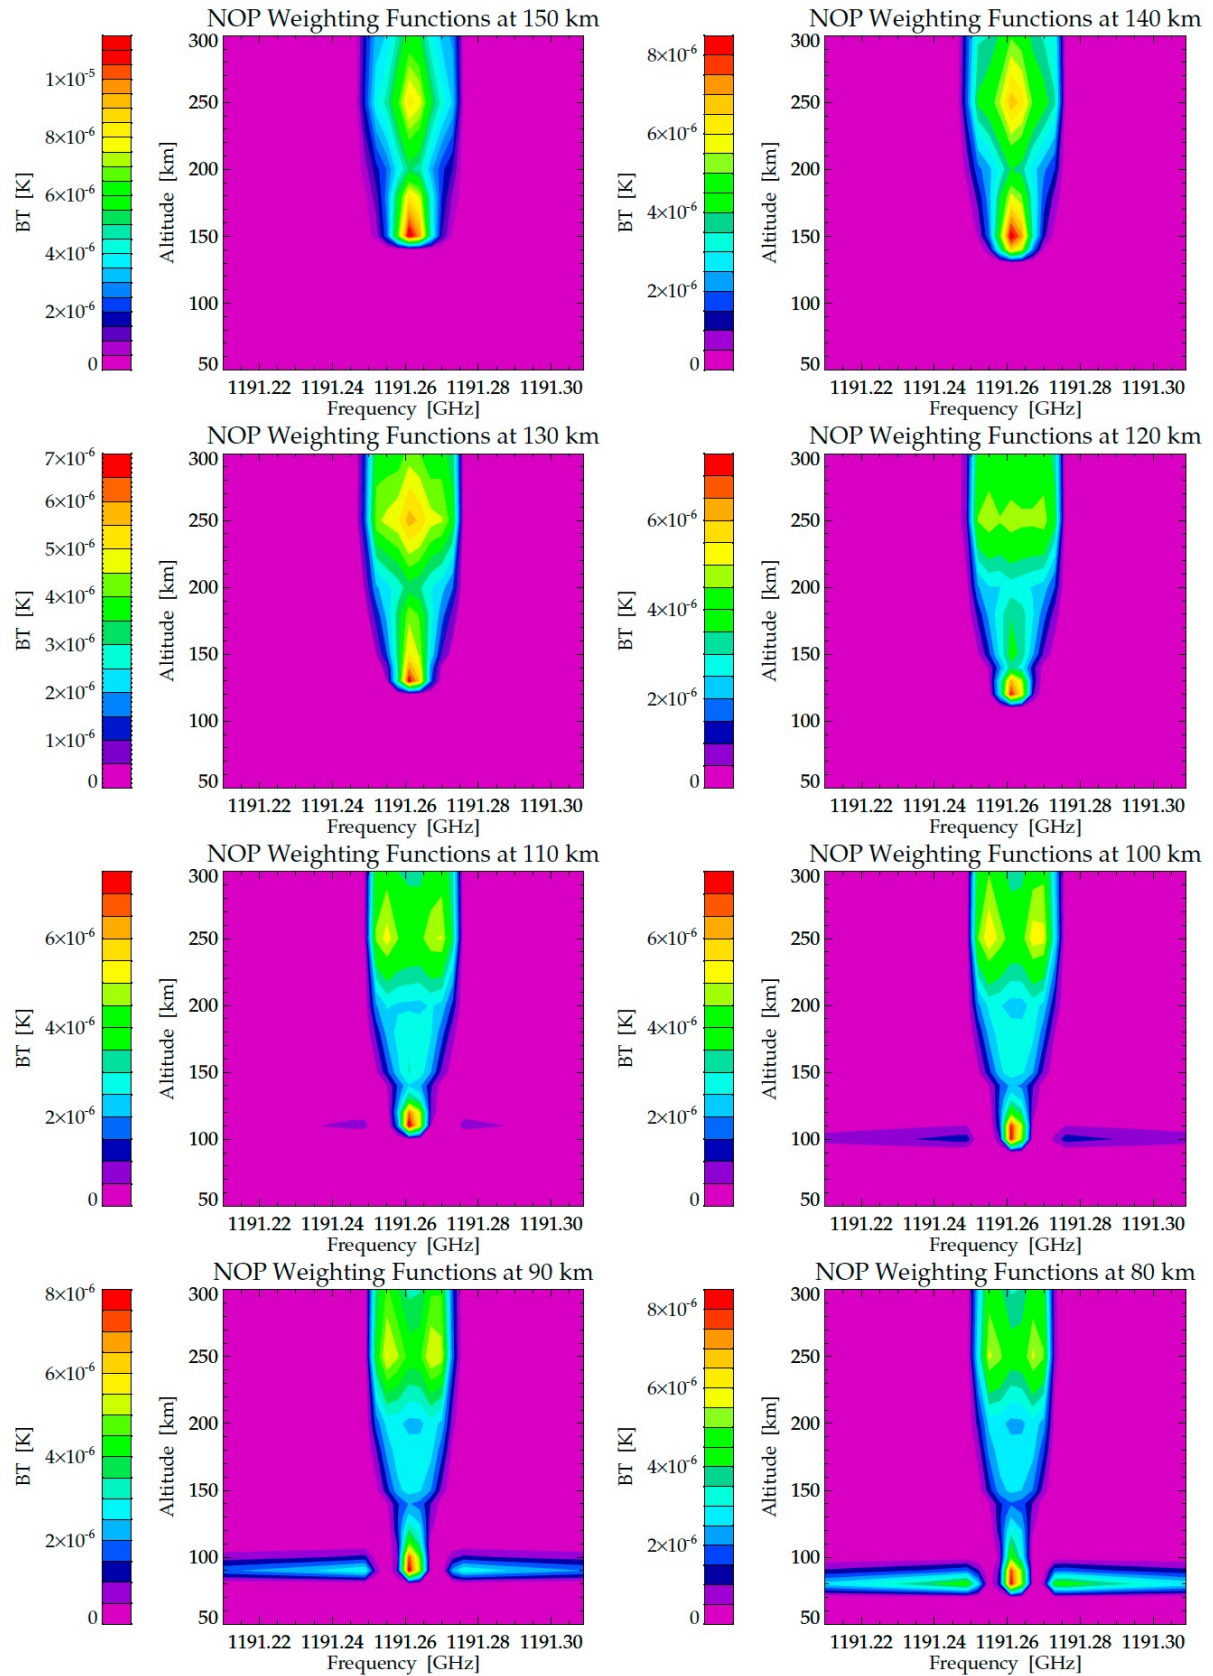


Figure 32: Atmospheric weighting functions (Kn) for NO+

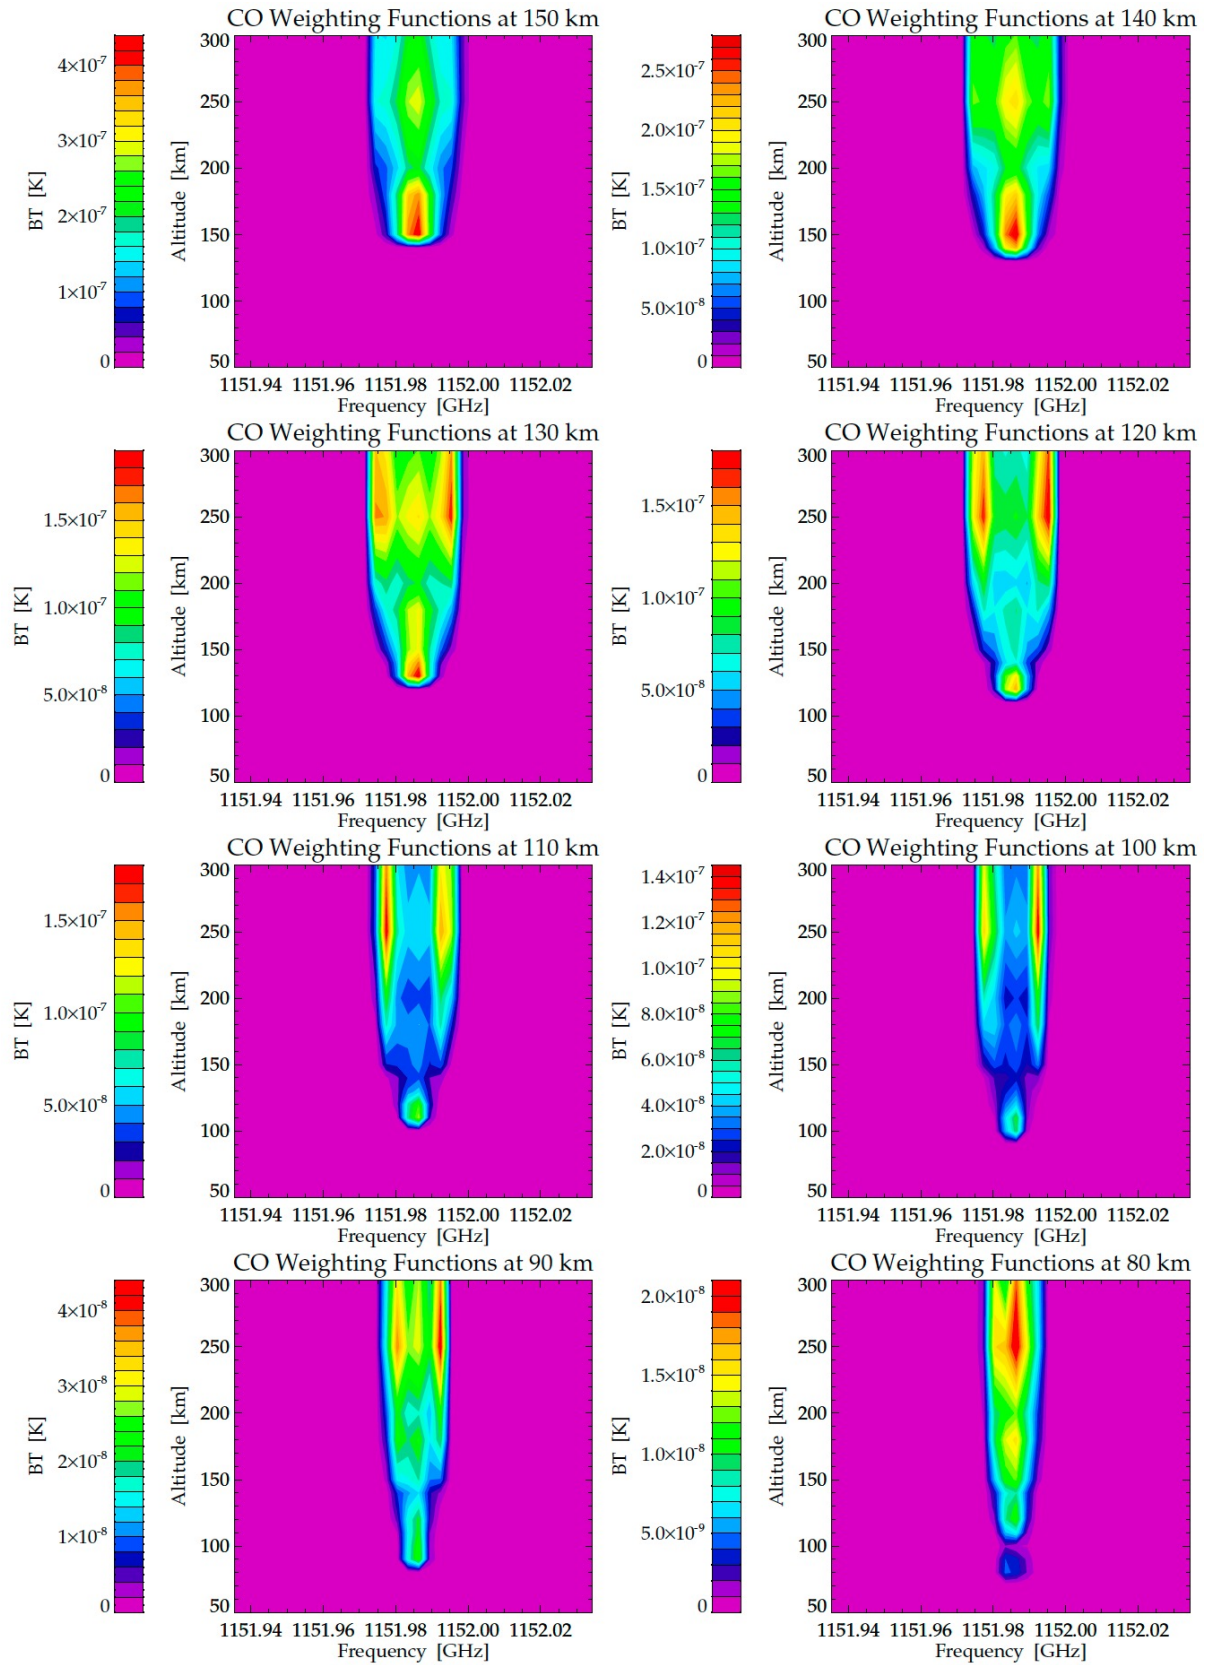


Figure 33: Atmospheric weighting functions (Kn) for CO

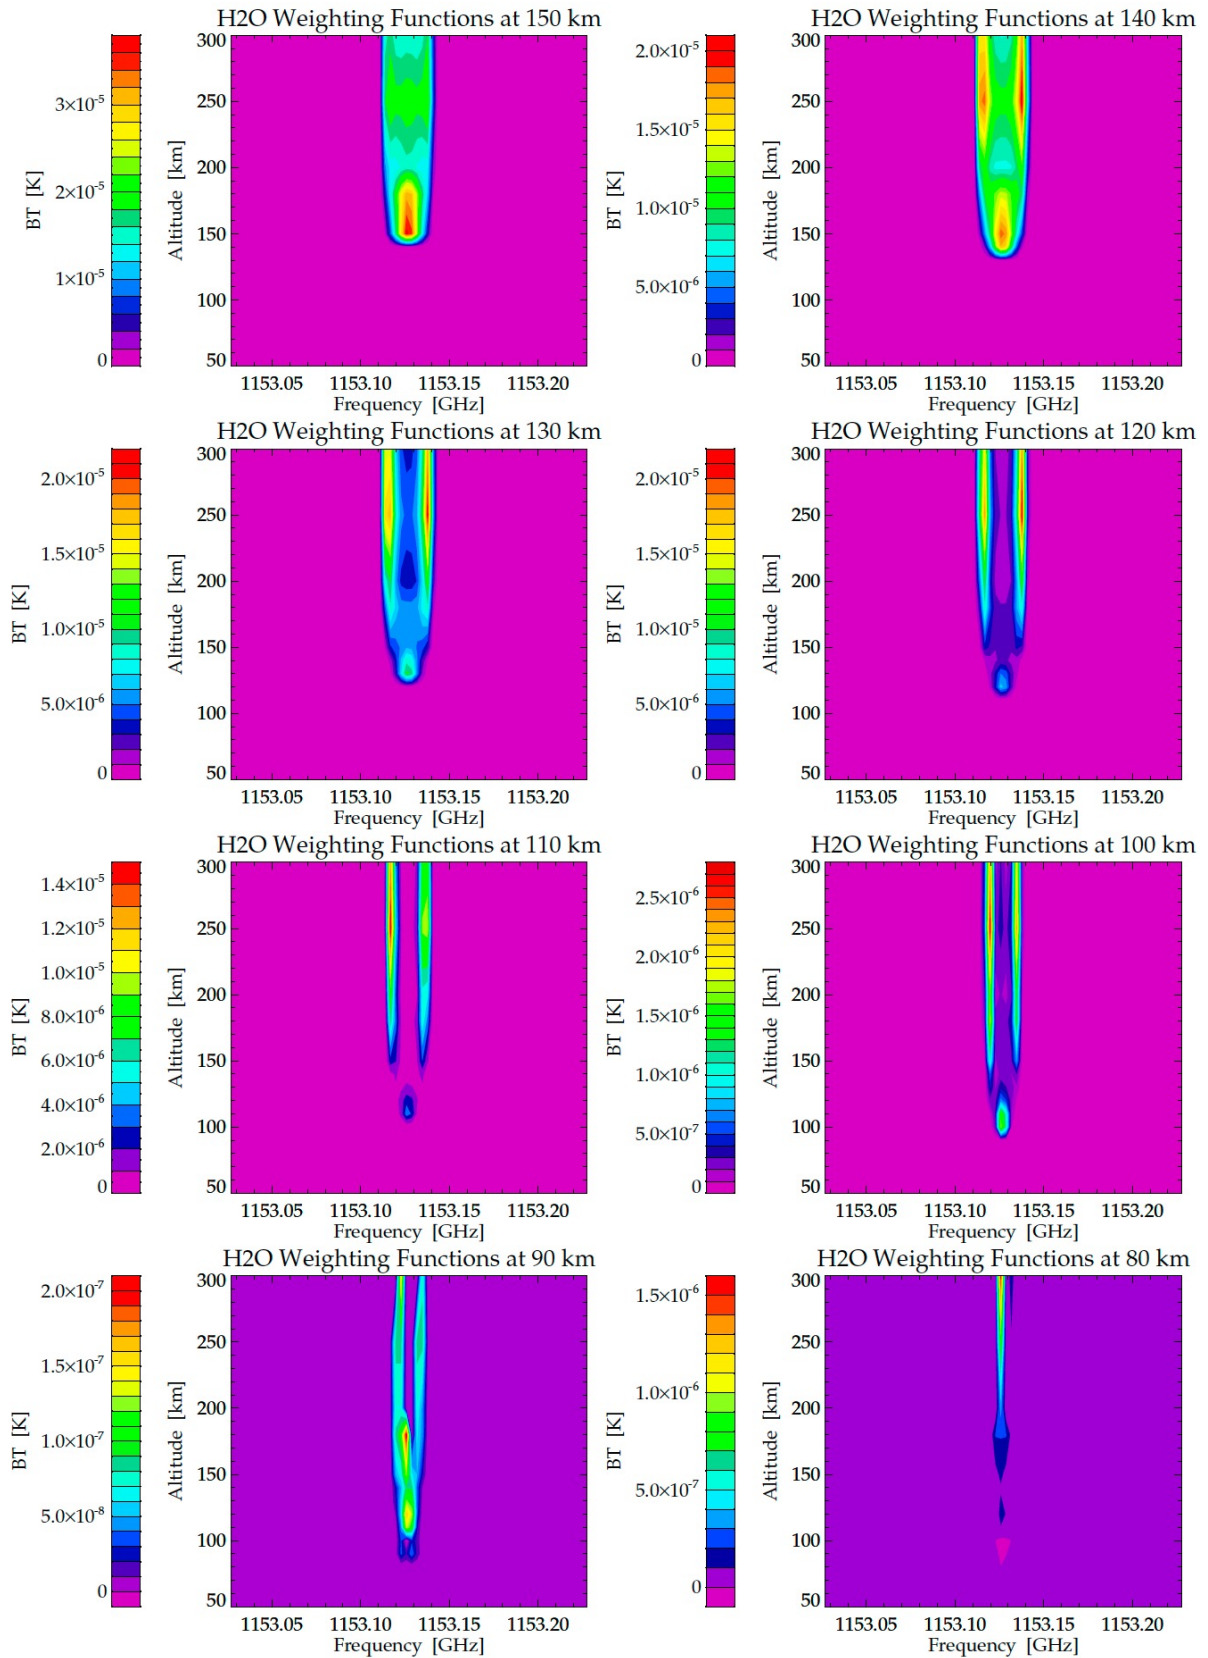


Figure 34: Atmospheric weighting functions (Kn) for H₂O

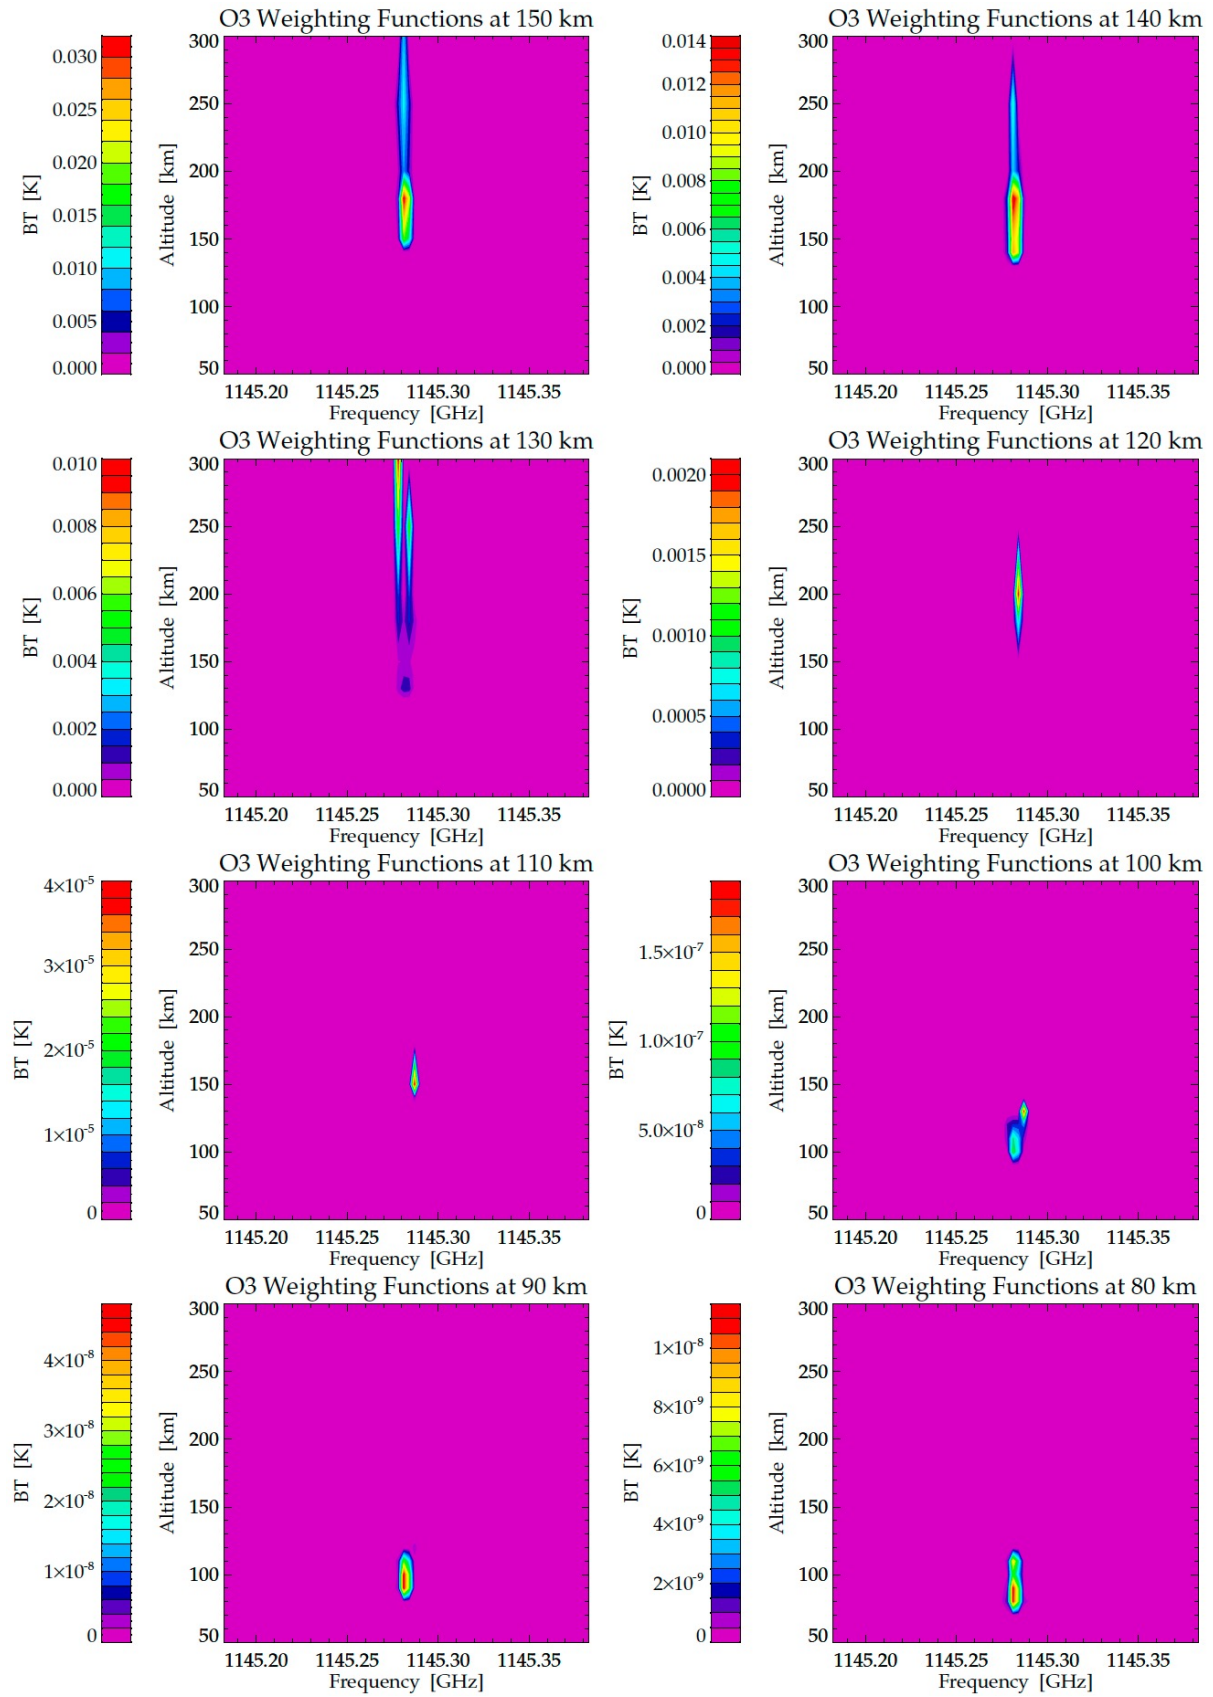


Figure 35: Atmospheric weighting functions (Kn) for O₃

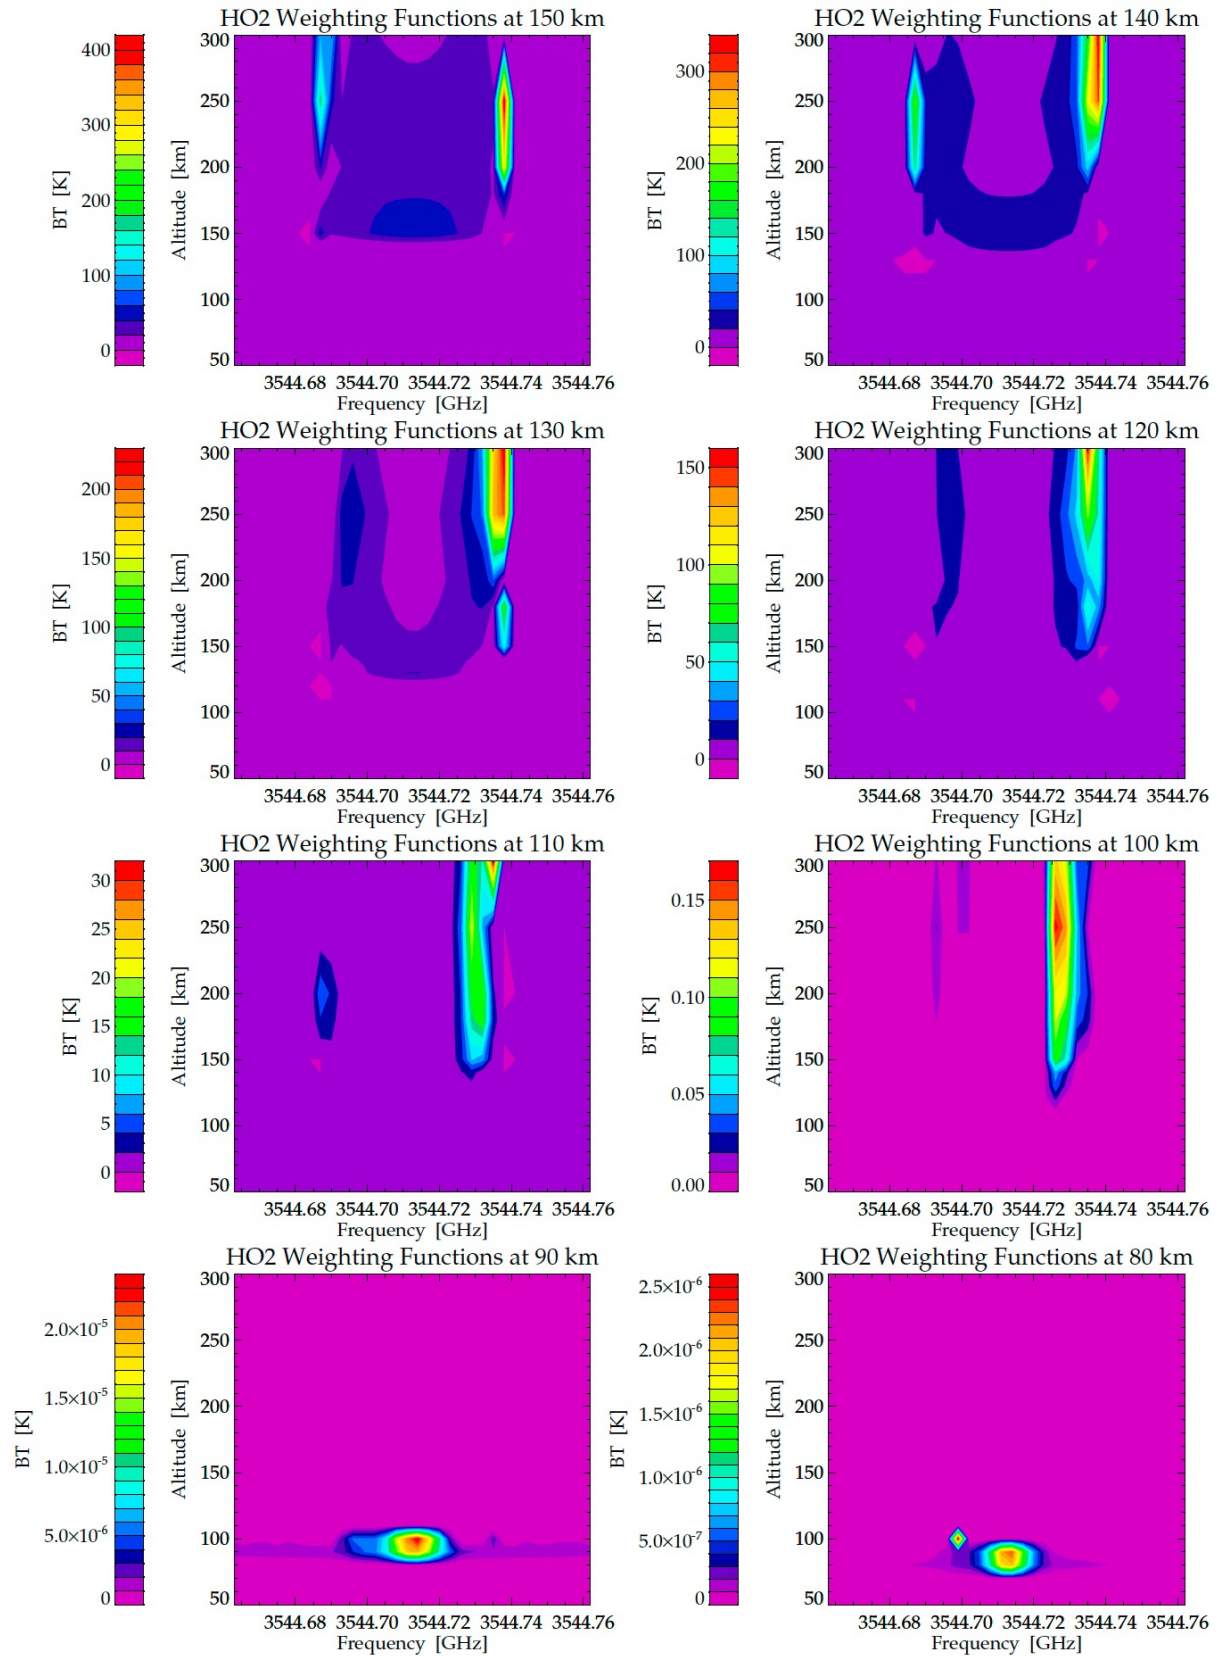


Figure 36: Atmospheric weighting functions (Kn) for HO₂. The main signal is from the line centre at low altitudes. The large Kn values at high altitudes are artefacts from a perturbation at altitudes where there is normally no HO₂, so a hypothetical presence would produce a strong impact.

Retrieval Precision

The retrieval precision is expressed as the errors of the retrieved state vector profile. The errors of the state vector profile are quantified in the state vector error covariance matrix S_X as described in the ATBD in a previous section. The trace of the state error covariance matrix is the variance of a corresponding expectation values from a Gaussian statistic. The corresponding expectation value is the retrieved profile of the atmospheric state. The retrieval errors x_e are therefore calculated as:

$$x_e = \sqrt{\text{TRACE}(K^T S_x K^{-1})}$$

The following plots (Figure 37 to Figure 45) show the results of a Zero-order iteration of the OEM retrieval run. The first guess, and also the a priori profile, are the input profiles from the WACCMX model climatology. The ‘retrieved’ profile in this non-iterative run will be unchanged from the first guess, but the retrieval diagnostics allow the calculation of retrieval errors and averaging kernels (resolution matrix).

The results are for a retrieval of a single measurement with a 3 second integration time every 2km over an atmospheric limb-scan from 50km to 150km. There are no measurements above 150km, although we have included some retrieval points at altitudes above 150km, because the line-of-sight passes through upper atmospheric layers, and these therefore contribute some information to the measurements at lower tangent points. In all following plots, we have extrapolated the retrieval precision statistically to the scenario of averaging over 100 individual three second measurements. With an estimated monthly observation frequency of ~33 measurements in a 5° Lat/Lon bin, this is roughly the equivalent of a seasonal mean over a ground pixel, or a 6h long orbit segment.

The retrieval grid spacing for the consecutive plots is a conservative 10km, applied uniformly to all species for direct comparison. The best achievable vertical resolution is studied in the next section.

A flat a priori error profile was chosen so as to not impress any vertical shape from the a priori error to the retrieval results. Similarly, a very large value for the a priori error was chosen, so as to not constrain the retrieval algorithm, and to guarantee that any information content that the algorithm identifies is attributed to the measurement, and not the regularisation. The a priori error profile is the red dashed line in all plots.

A direct interpretation of the vertical range of a successful retrieval is the intersection point of the retrieval error profile line with the retrieved state profile, which corresponds to a relative retrieval error of 100%. Anywhere where the retrieval error is lower than the absolute retrieved value, the retrieval effort is valid in principle.

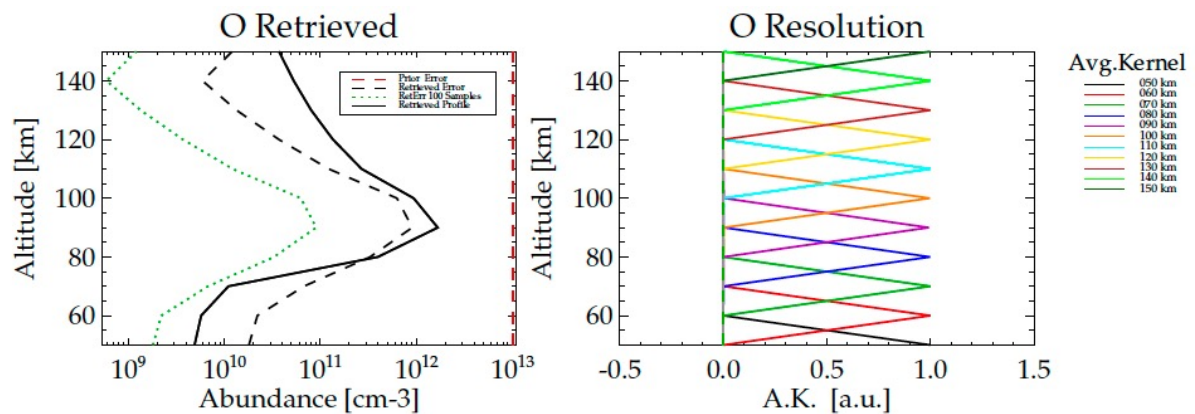


Figure 37: Retrieval error profile (dashed line) for O at 4.7THz. The retrieval errors for a single measurement are lower than the total abundance profile (solid black line) from ~90km upwards. The averaging kernels (right panel) are perfectly uncorrelated for a 10km retrieval grid spacing.

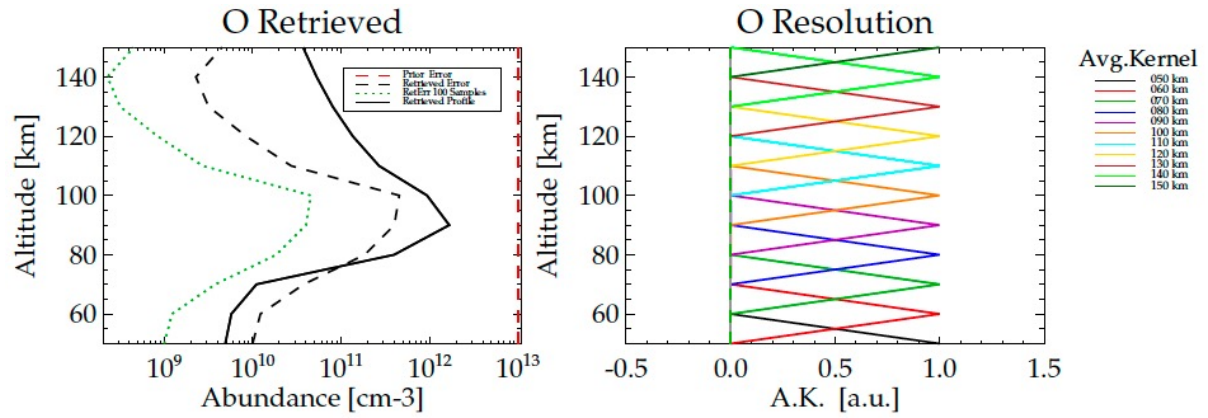


Figure 38: Retrieval error profile (dashed line) for O at 2THz. The retrieval errors for a single measurement are lower than the total abundance profile (solid black line) from ~80km upwards. The averaging kernels (right panel) are perfectly uncorrelated for a 10km retrieval grid spacing.

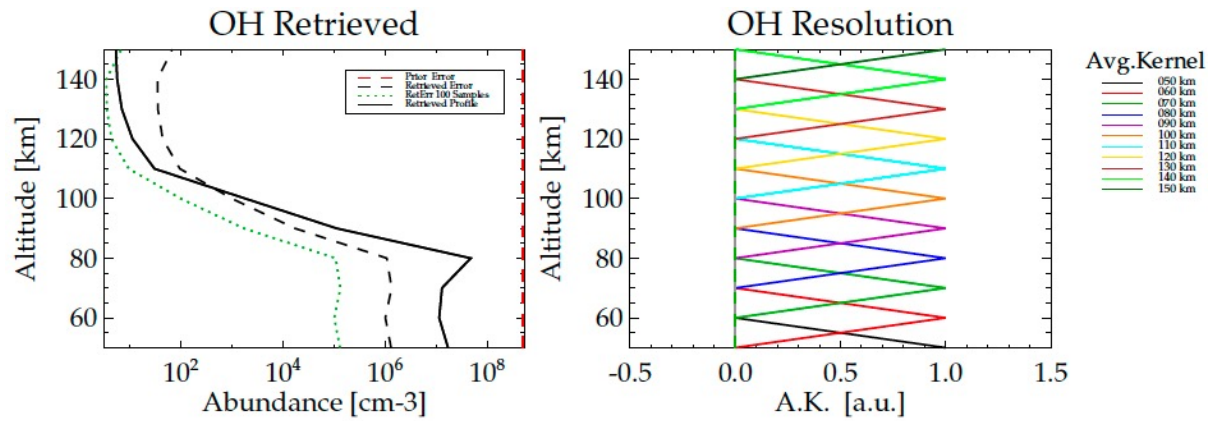


Figure 39: Retrieval error profile (dashed line) for OH. The retrieval errors for a single measurement are lower than the total abundance profile (solid black line) up to ~100km. The averaging kernels (right panel) are perfectly uncorrelated for a 10km retrieval grid spacing.

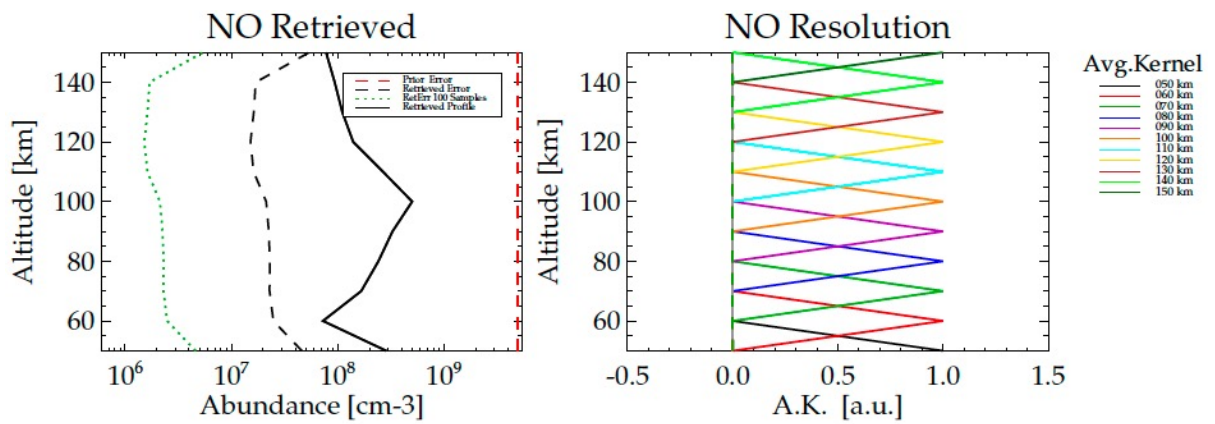


Figure 40: Retrieval error profile (dashed line) for NO. The retrieval errors for a single measurement are lower than the total abundance profile (solid black line) for the full scan range. The averaging kernels (right panel) are perfectly uncorrelated for a 10km retrieval grid spacing.




Figure 41: Retrieval error profile (dashed line) for Nitrosylium. The retrieval errors for a single measurement are never lower than the total abundance profile (solid black line), meaning that multiple measurements will have to be averaged together (i.e. green dotted line for 100 samples). The averaging kernels (right panel) are perfectly uncorrelated for a 10km retrieval grid spacing.




Figure 42: Retrieval error profile (dashed line) for CO. The retrieval errors for a single measurement are lower than the total abundance profile (solid black line) up to ~140km. The averaging kernels (right panel) are perfectly uncorrelated for a 10km retrieval grid spacing.




Figure 43: Retrieval error profile (dashed line) for H₂O. The retrieval errors for a single measurement are lower than the total abundance profile (solid black line) up to ~120km. The averaging kernels (right panel) are perfectly uncorrelated for a 10km retrieval

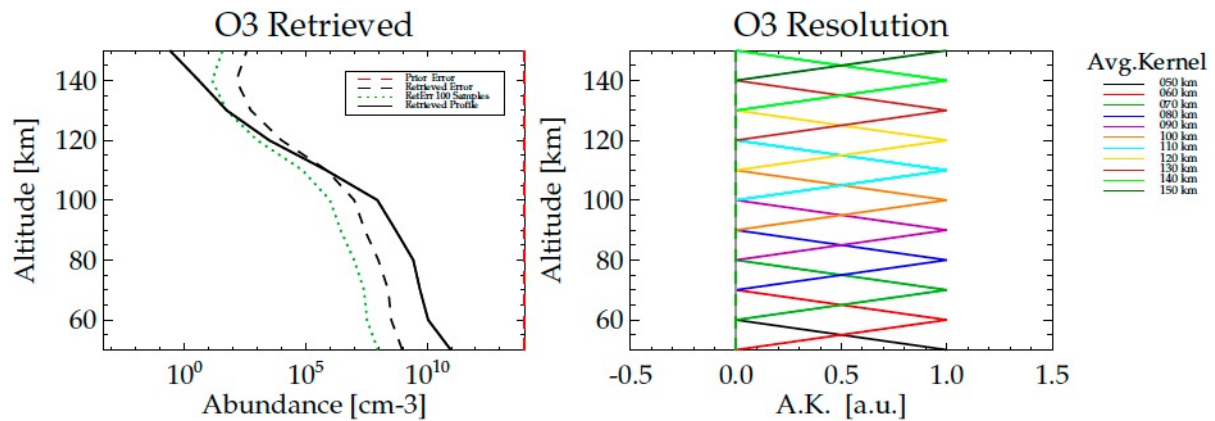


Figure 44: Retrieval error profile (dashed line) for O_3 . The retrieval errors for a single measurement are lower than the total abundance profile (solid black line) up to ~ 110 km. The averaging kernels (right panel) are perfectly uncorrelated for a 10 km retrieval grid spacing.




Figure 45: Retrieval error profile (dashed line) for HO_2 . The retrieval errors for a single measurement are lower than the total abundance profile (solid black line) up to ~ 90 km. The averaging kernels (right panel) are perfectly uncorrelated for a 10 km retrieval grid spacing.

Vertical Resolution

The scientific motivation is to push the vertical resolution to the maximum. The point to which this can be done is revealed by the resolution matrix, aka the averaging kernels. Ideally, the averaging kernels at each altitude should be perfectly triangular, and peak at a normalized value of one (as displayed by the plots in the previous section). This means that the retrieved information at any altitude is originating from this single altitude layer only. If the averaging kernels peak at less than one, then they are wider than the vertical layer, which

means that information of the neighboring layer is impressed on the retrieved value at this given altitude. This introduces an additional error, often called the smoothing error. The best approach then, is to start with a moderate vertical resolution (i.e. 10km in our example in the previous section), and then gradually increase the resolution, while keeping an eye on the averaging kernels. If the envelope of these is starting to diverge from unity, then we have reached the maximum achievable vertical resolution.

We have performed retrieval simulation for three different retrieval grids with vertical spacing's of 10km, 5km and 2km. 2km is equal to the spacing of the measurements in our atmospheric limb-scan, so there is little information to be gained from going even narrower. The relative retrieval error profiles, together with the averaging kernels, are shown in the following plots. We've highlighted what we consider to be the optimal of the three retrieval grids in bold letters.

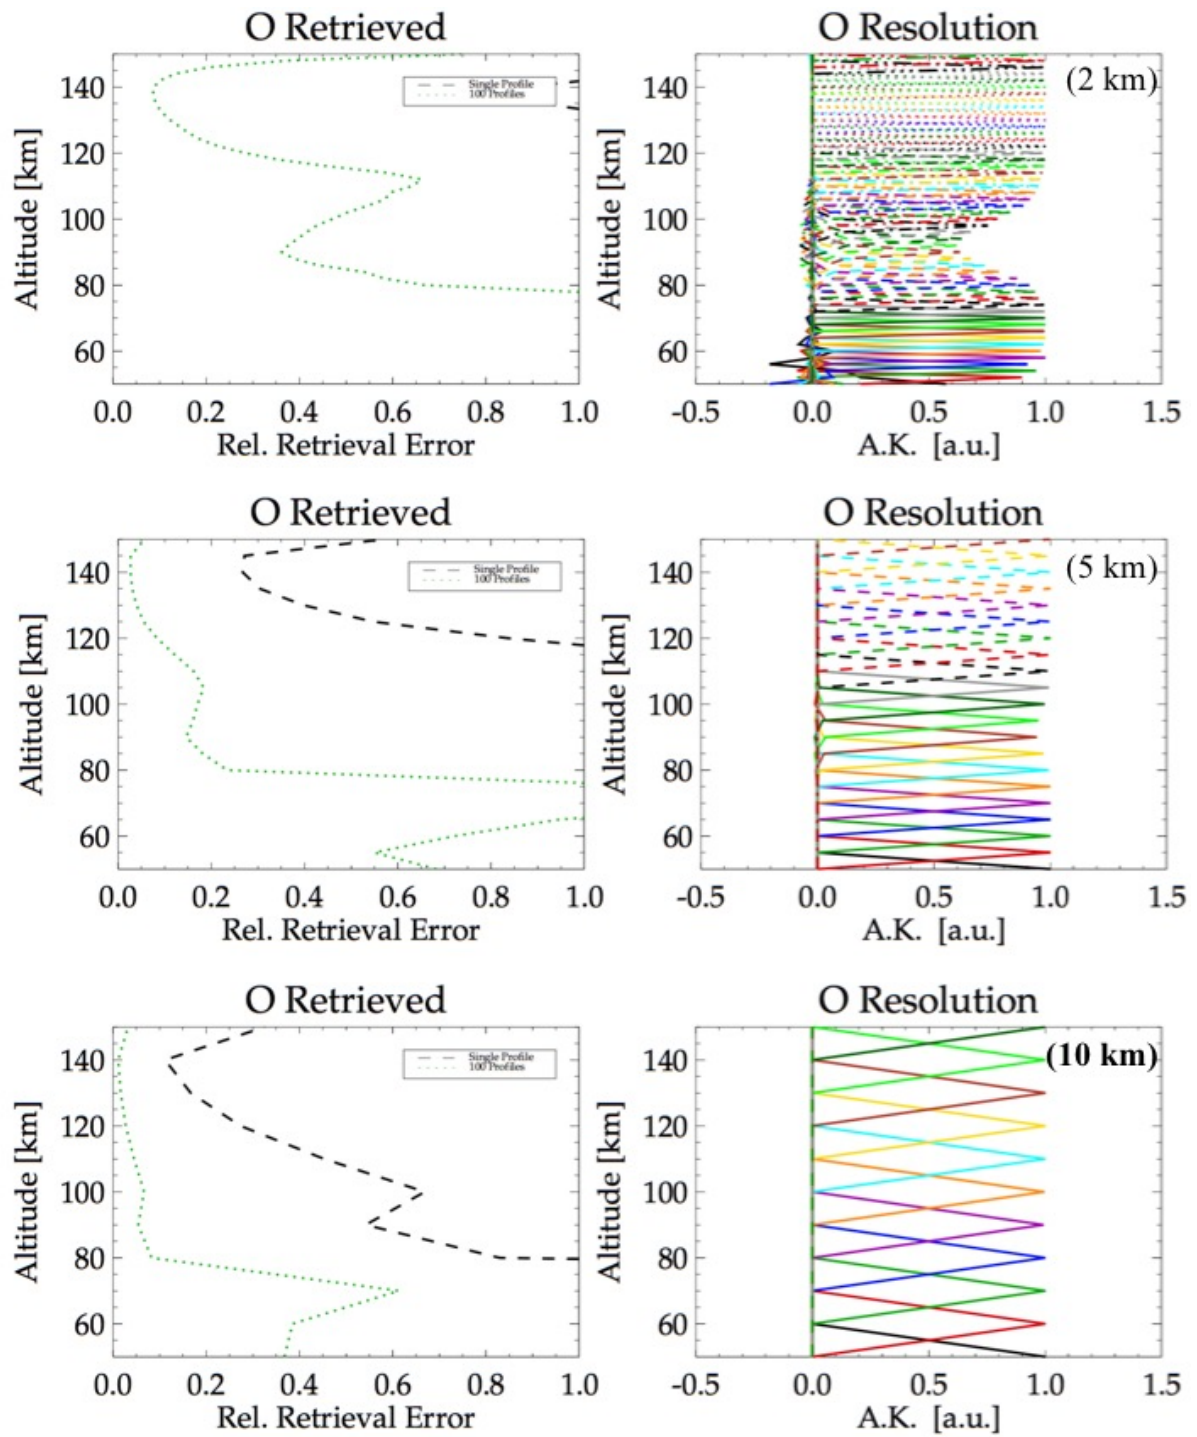


Figure 46: O retrieval errors for different vertical retrieval grids. While the averaging kernels still look halfway decent at 5km vertical resolution, the retrieval errors at low altitudes are already significantly degraded. Therefore, it is best to retrieve O at 10km resolution, at least from single measurements (i.e. without averaging over multiple samples). At 2km resolution the averaging kernels first start to break down, expectedly in the altitude region of the steep gradient and high variability.

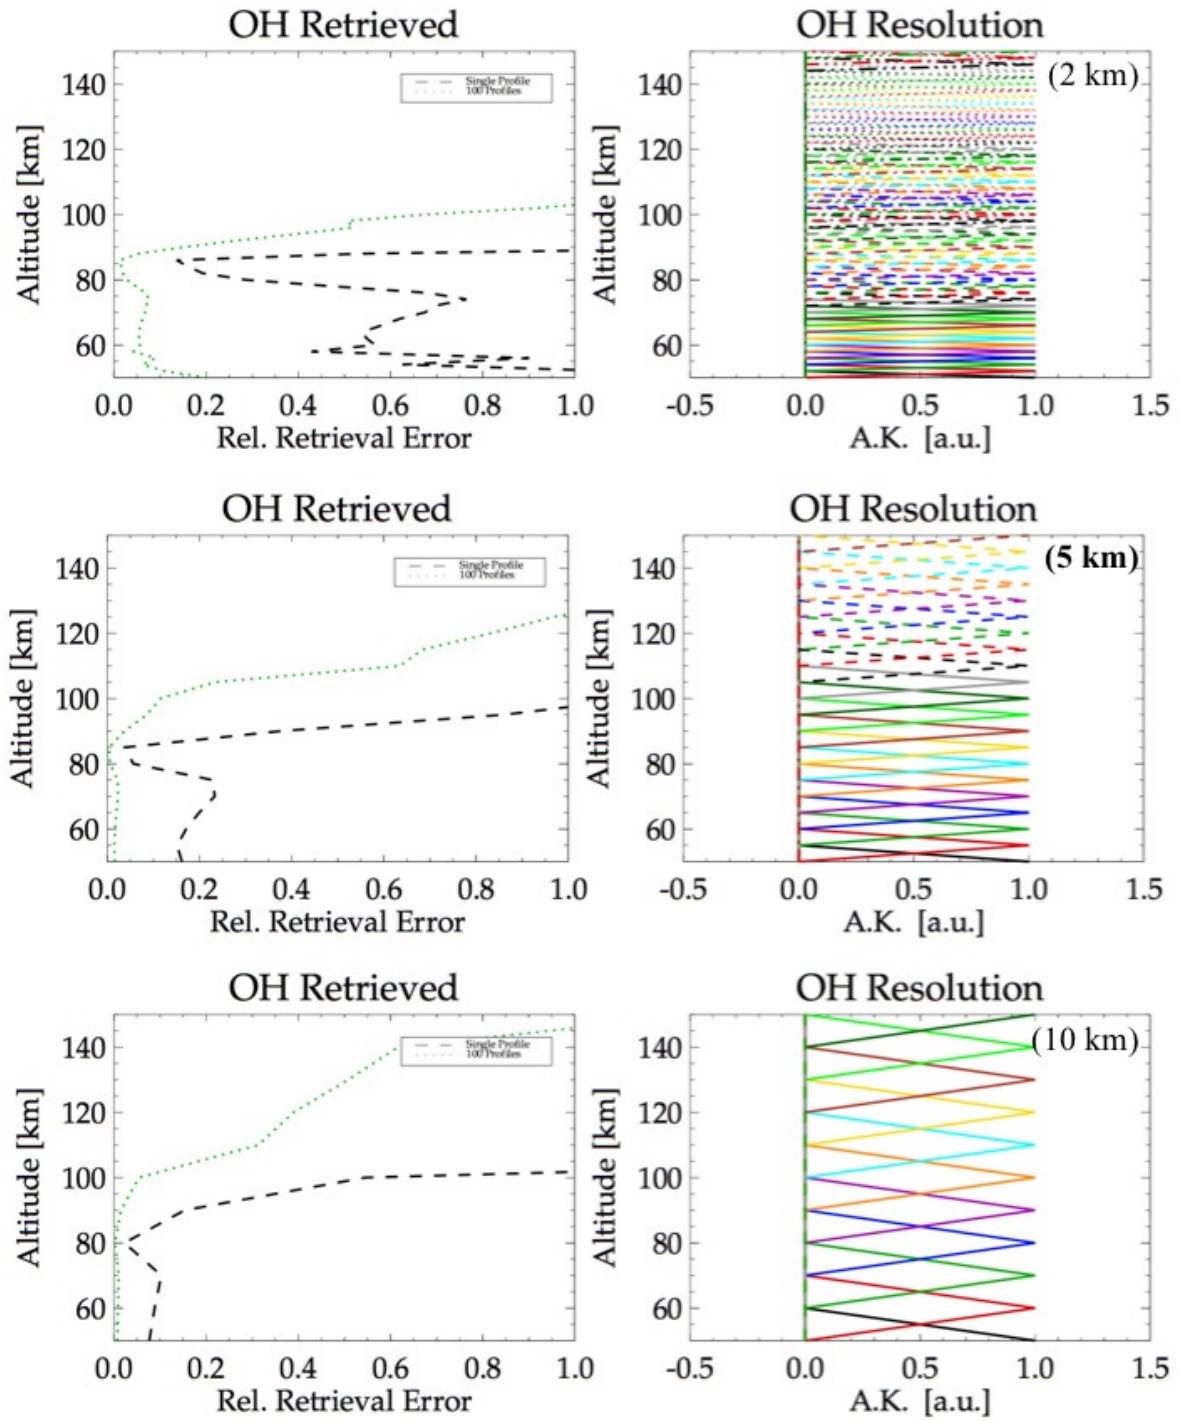
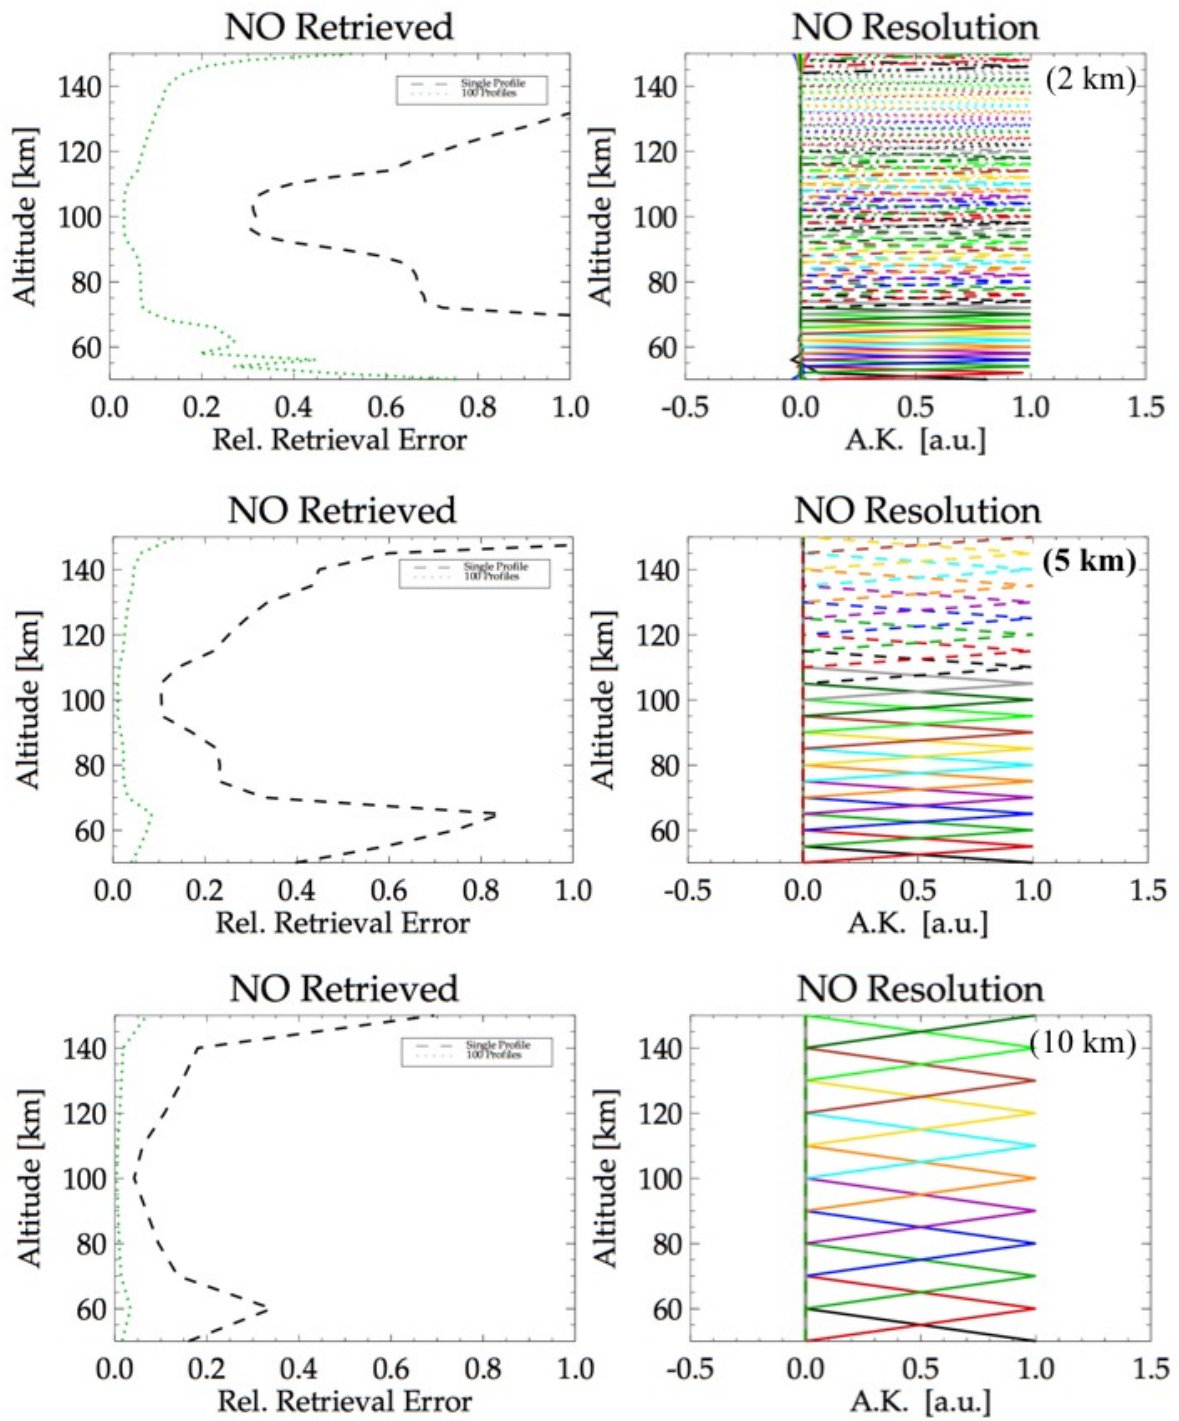


Figure 47: OH retrieval errors for different vertical retrieval grids. For OH, there is little penalty in pushing the vertical resolution down to 5km. The averaging kernels at 2km resolution are still well separated, but the retrieval errors at the lowest altitudes start to increase. It's not worth pushing the retrieval of a single measurement to a resolution finer than 5km.



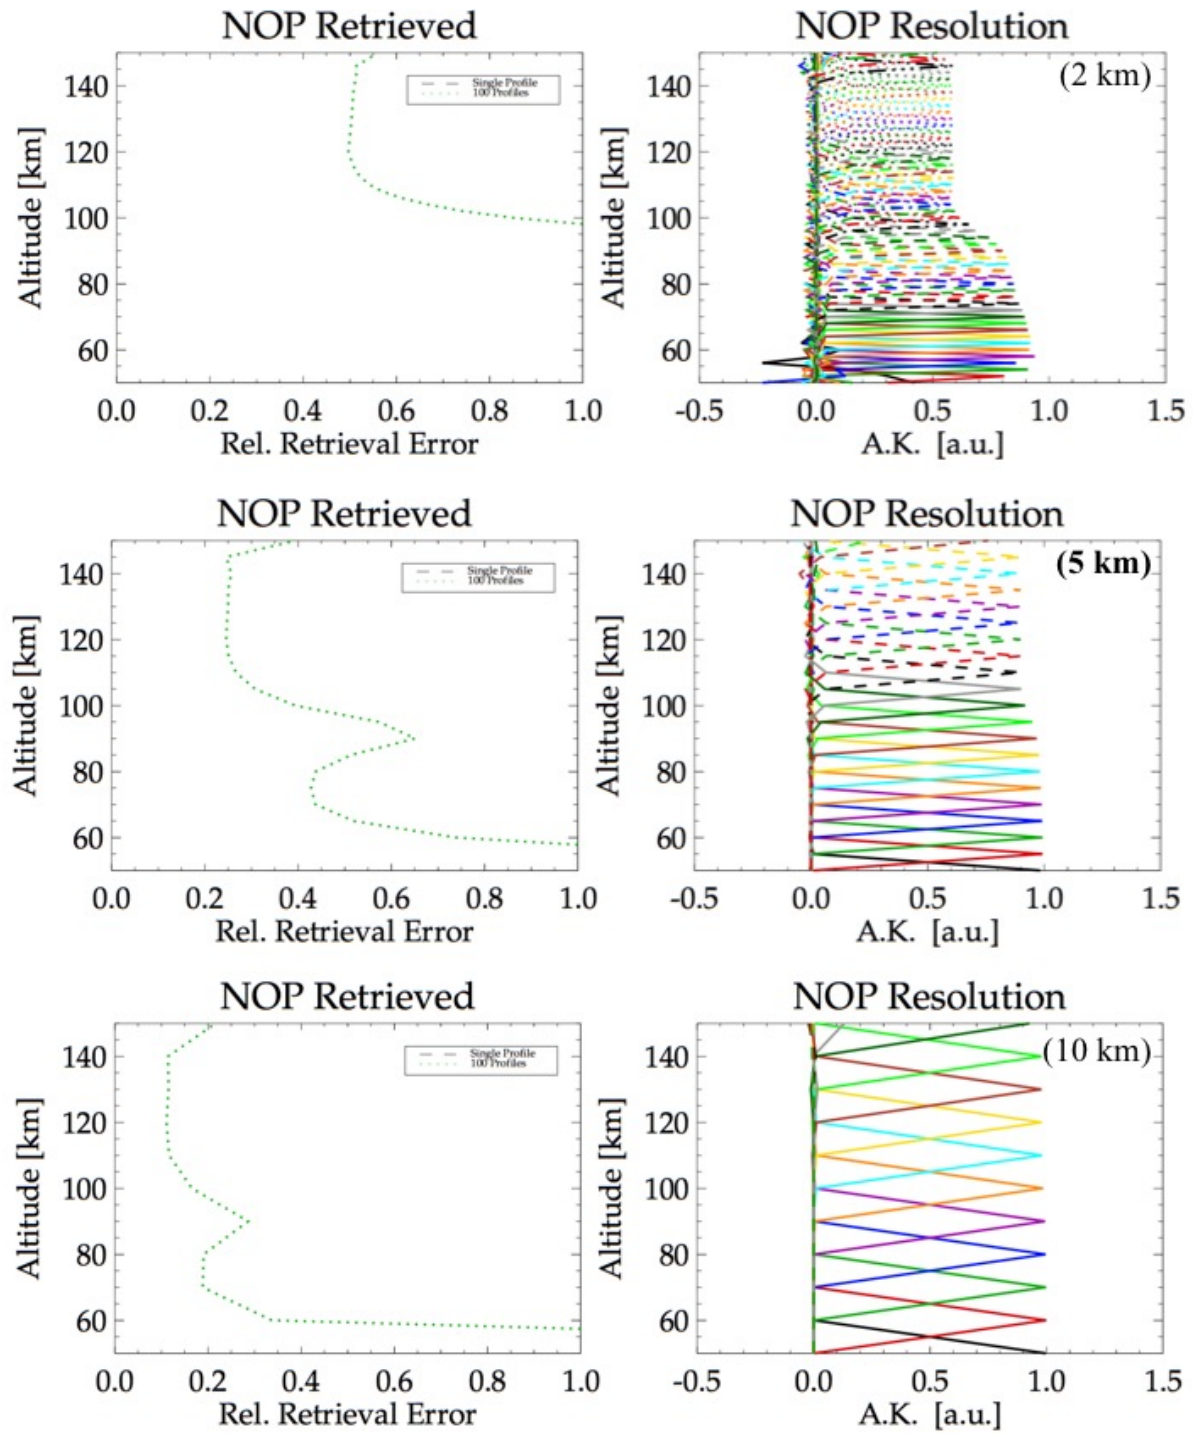


Figure 49: NO⁺ retrieval errors for different vertical retrieval grids. Because the spectral line of NO⁺ is so weak, it's not possible to retrieve a single measurement. That was already confirmed by the retrieval error plot in Figure 41. Some sort of global averaging will have to be applied anyway. The relative retrieval error plots here are therefore not indicative of an optimal retrieval grid spacing. The averaging kernels seem to indicate that an improved vertical resolution of 5 km is feasible, but this will have to pass the test of reality.

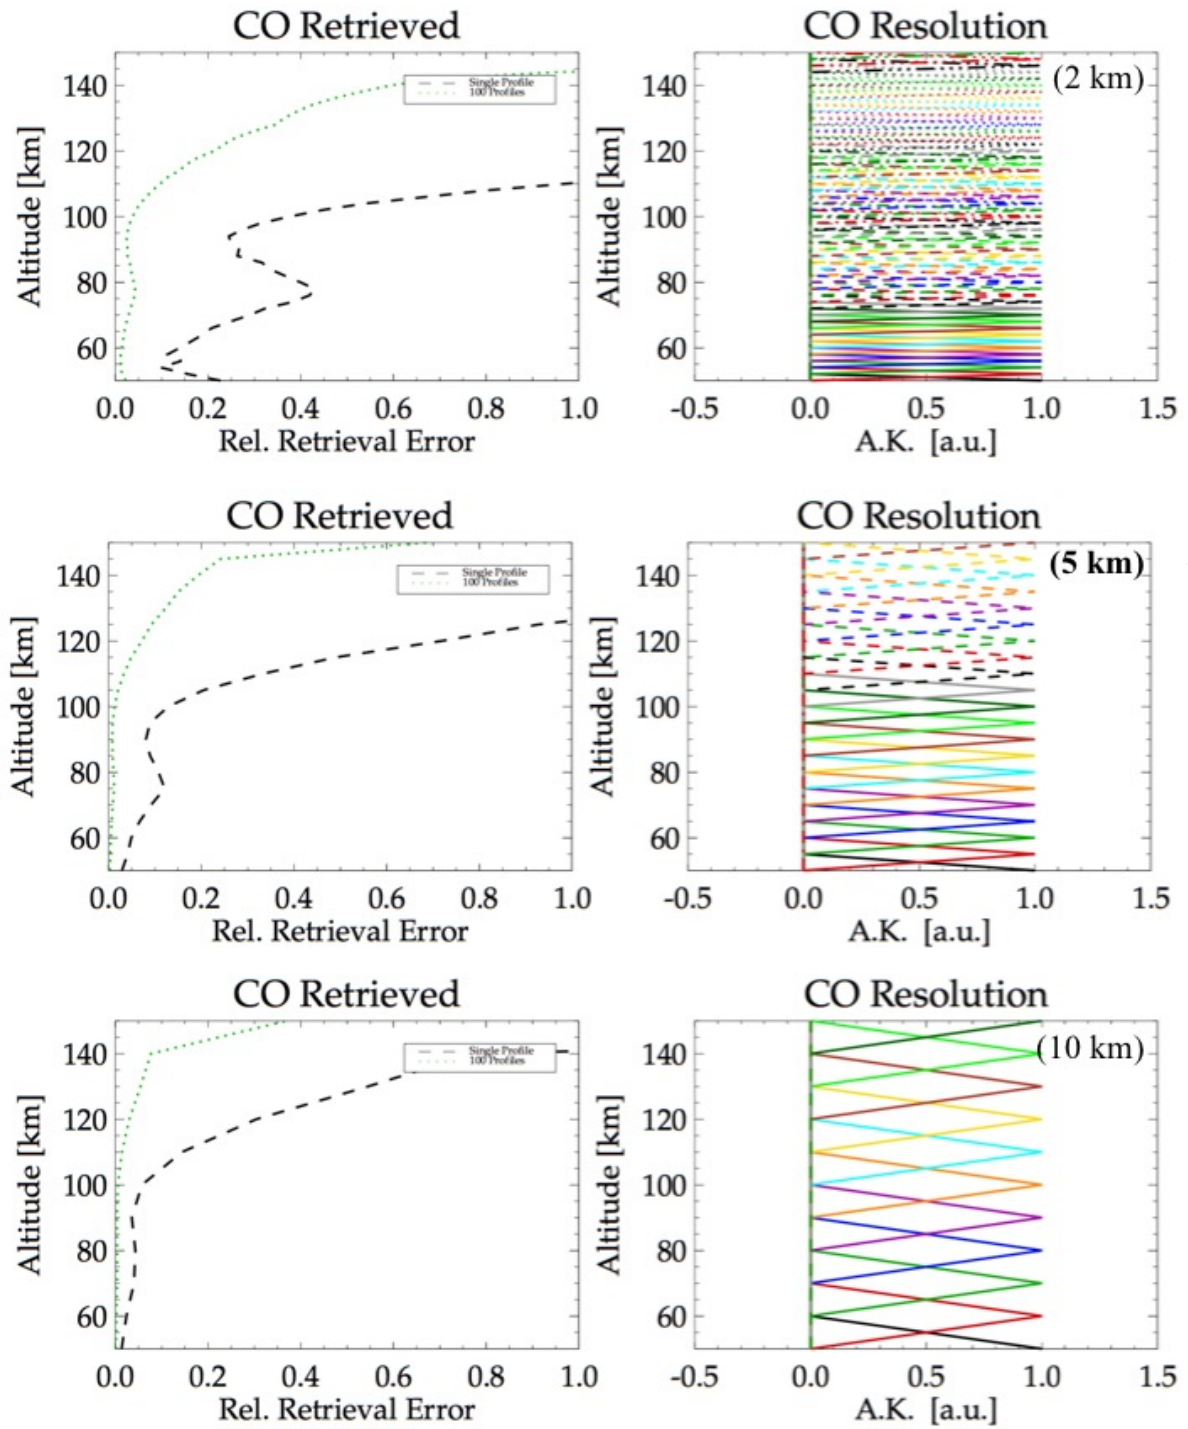


Figure 50: CO retrieval errors for different vertical retrieval grids. For CO, there is little penalty in pushing the vertical resolution down to 5km. The averaging kernels at 2km resolution are still well separated, but the retrieval errors at the lowest altitudes start to increase. It's not worth pushing the retrieval of a single measurement to a resolution finer than 5km.

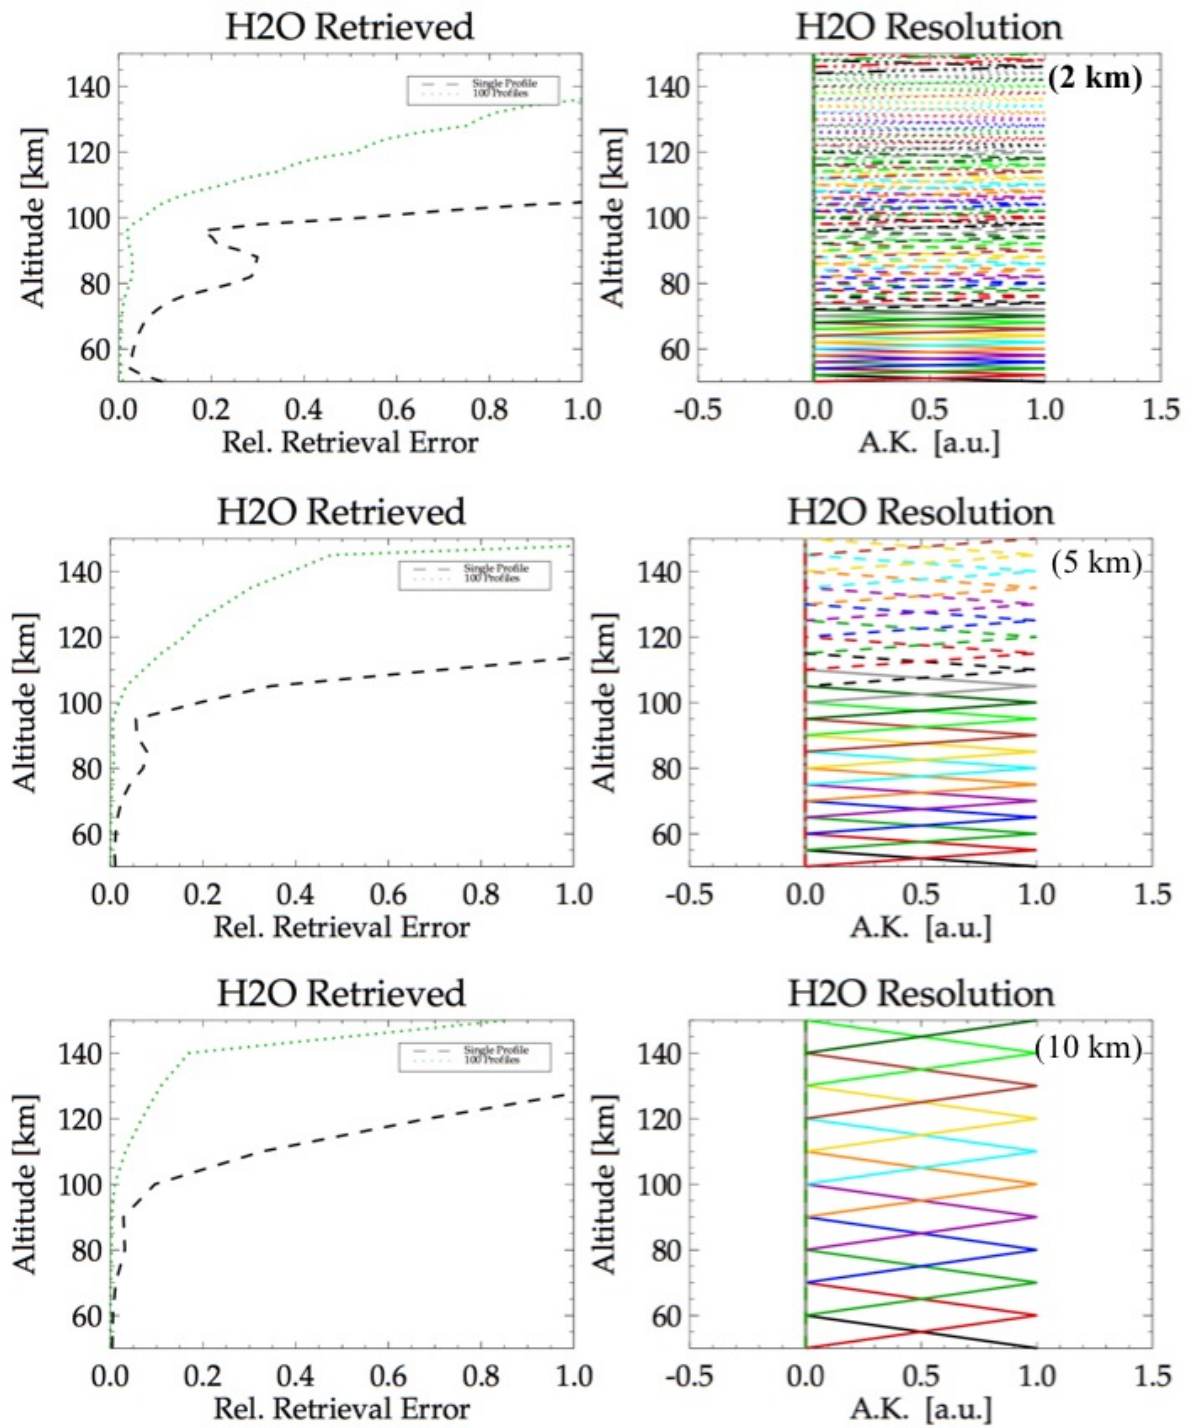


Figure 51: H₂O retrieval errors for different vertical retrieval grids. For H₂O, it seems acceptable to push the retrieval all the way to the maximum theoretical possible vertical resolution of 2km. This is no doubt thanks to the fact that water vapour is a strong spectral emitter. It's also a nice finding, because humidity (together with temperature) is a fundamental descriptor of the atmospheric state, and it's desirable to know it with a high level of detail.

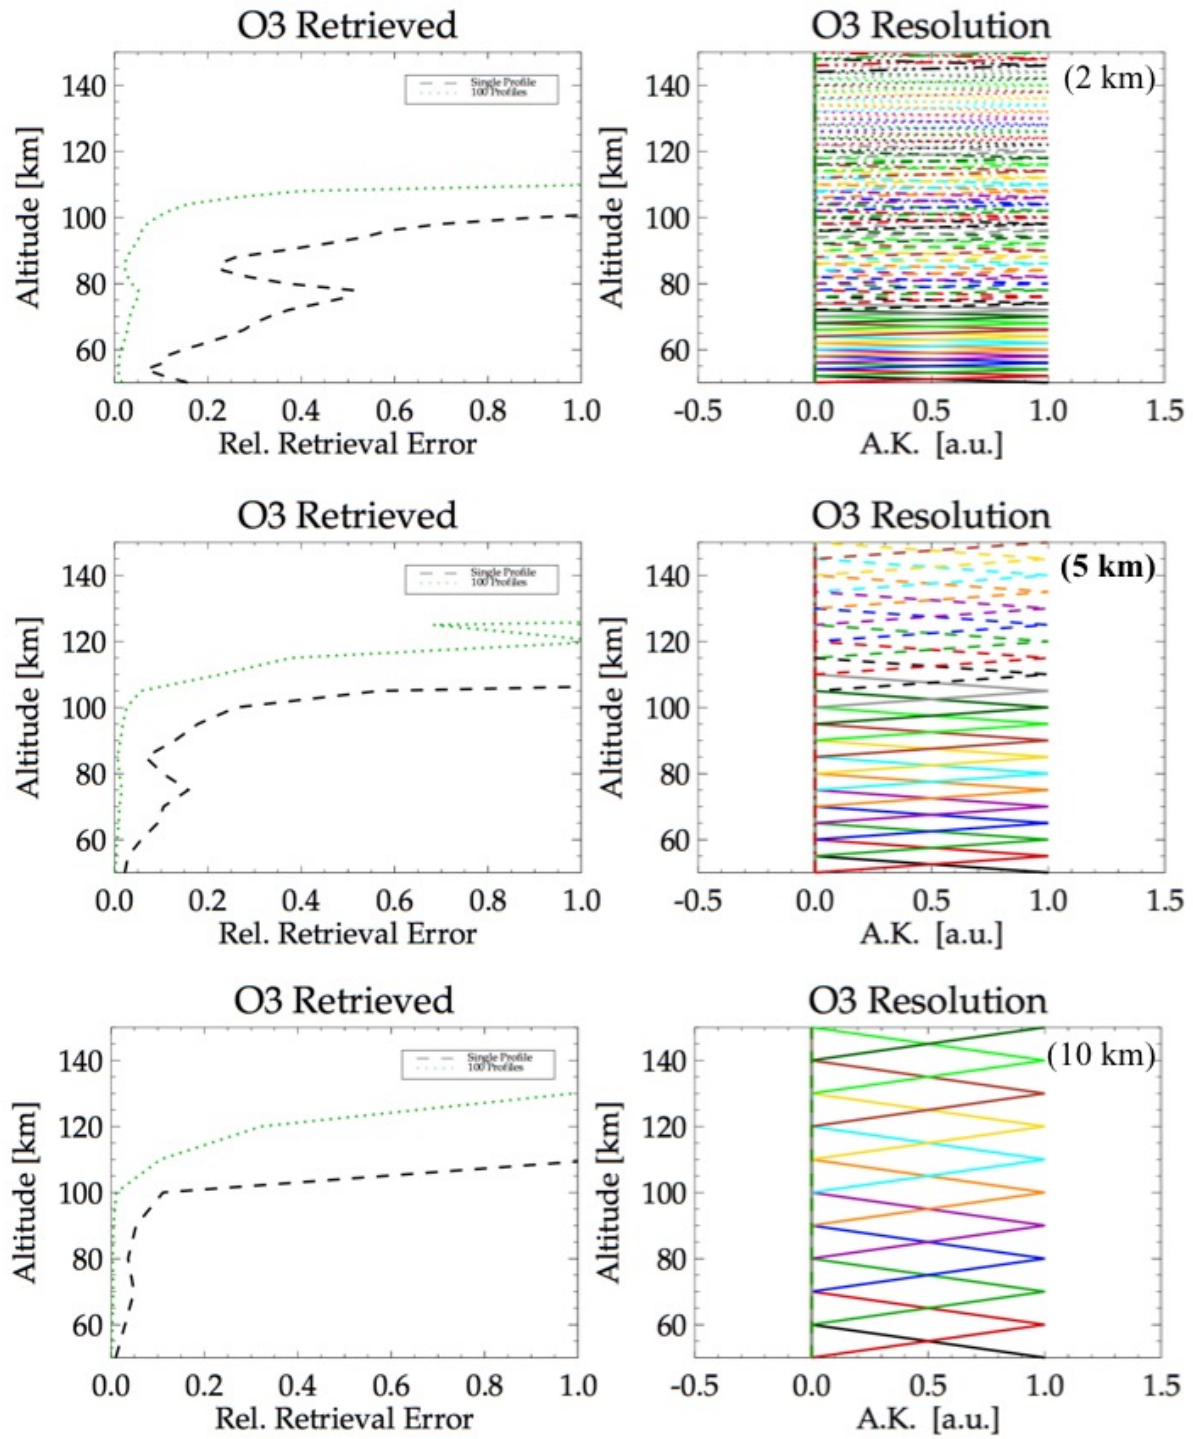


Figure 52: O_3 retrieval errors for different vertical retrieval grids. For O_3 , there is little penalty in pushing the vertical resolution down to 5km. The averaging kernels at 2km resolution are still well separated, but the retrieval errors at the lowest altitudes start to increase. It's not worth pushing the retrieval of a single measurement to a resolution finer than 5km.

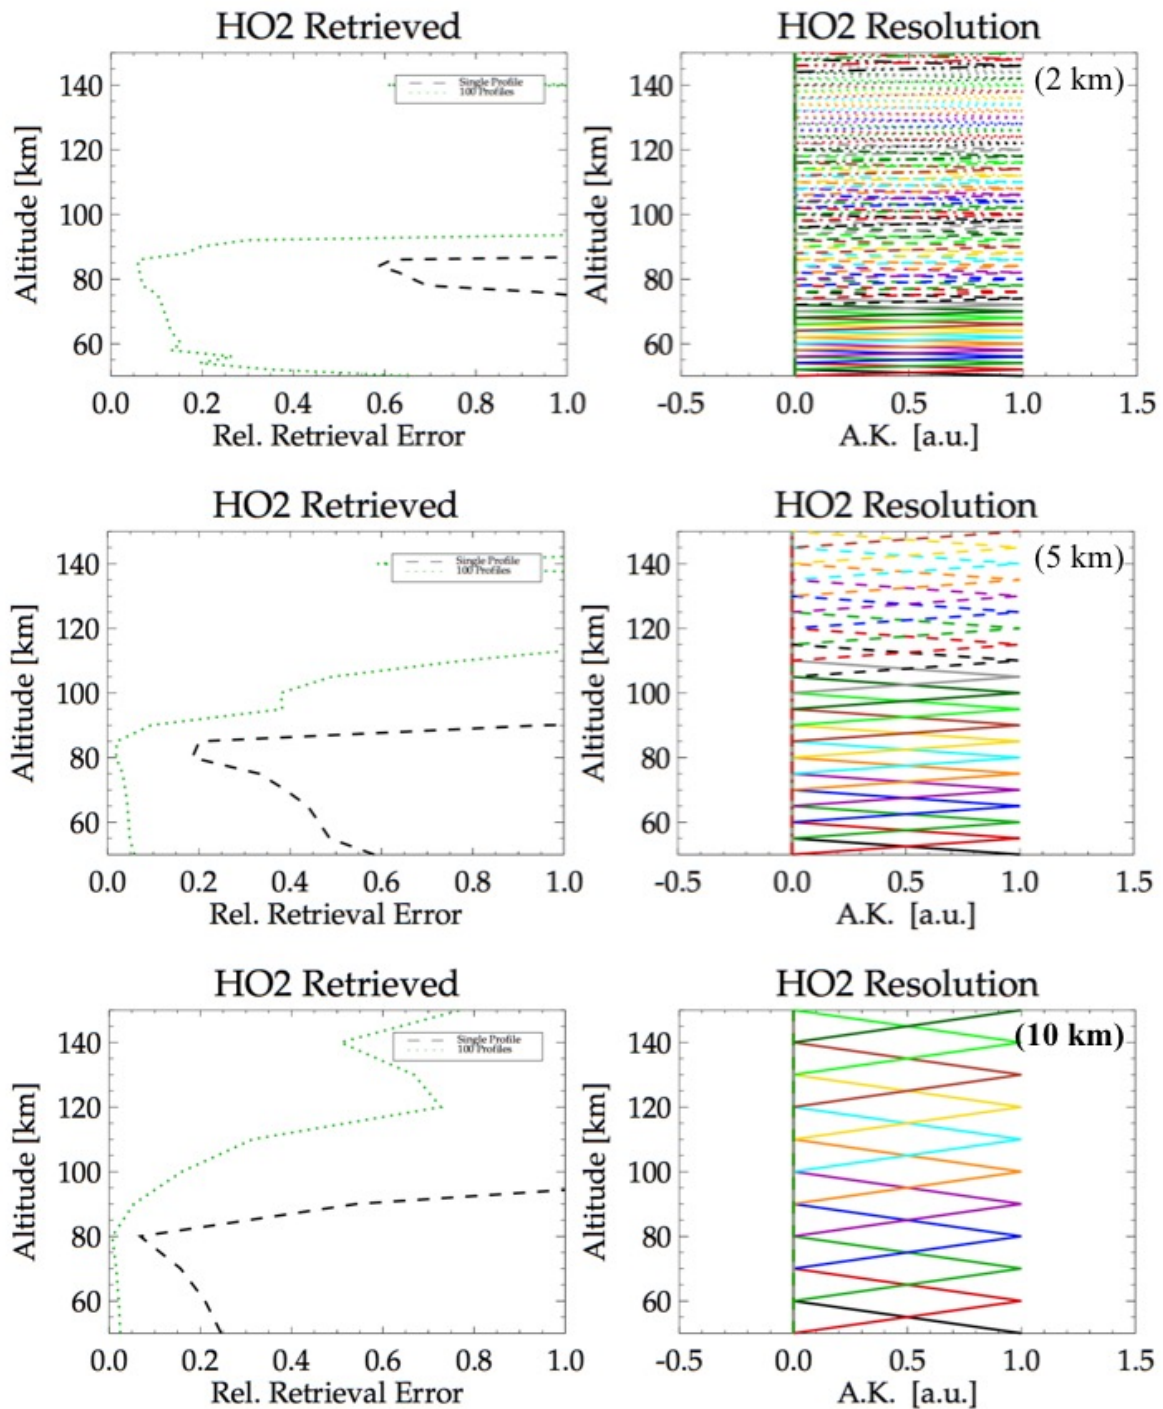


Figure 53: HO₂ retrieval errors for different vertical retrieval grids. For HO₂, although the averaging kernels look well behaved, the retrieval errors at the lowest altitudes start to increase quite significantly when increasing the vertical resolution even down to 5 km. HO₂ also features a comparatively weak spectral line, so the assumption is that for this species an averaging will be required. The only current instrument to measure HO₂ from limb-sounding, Aura MLS, only produces monthly zonal means for its distribution. For a single measurement, it doesn't seem worth pushing the retrieval to a vertical resolution finer than 5 km.

Impact of Atmospheric Variability

In the previous sections, we have established the physical limitations of the retrieval process, i.e. find out what the lowest possible retrieval errors and the best possible vertical resolutions are for the ideal atmospheric scenario (i.e. maximum abundance profile of the single species that is to be retrieved, combined with a fixed, average atmosphere). It is interesting to examine how the retrieval results change for each species for the three extreme climatological scenarios of maximum, mean, and minimum abundances. This is illustrated in the following figures. For some species, i.e. those with a strong diurnal variation, these differences can be extreme. In some cases, the species under test can all but disappear during the bottom part of its chemical lifecycle, in which case a detection will become impossible. Often though – e.g. in the case of atomic oxygen – this is restricted to a certain altitude interval only.

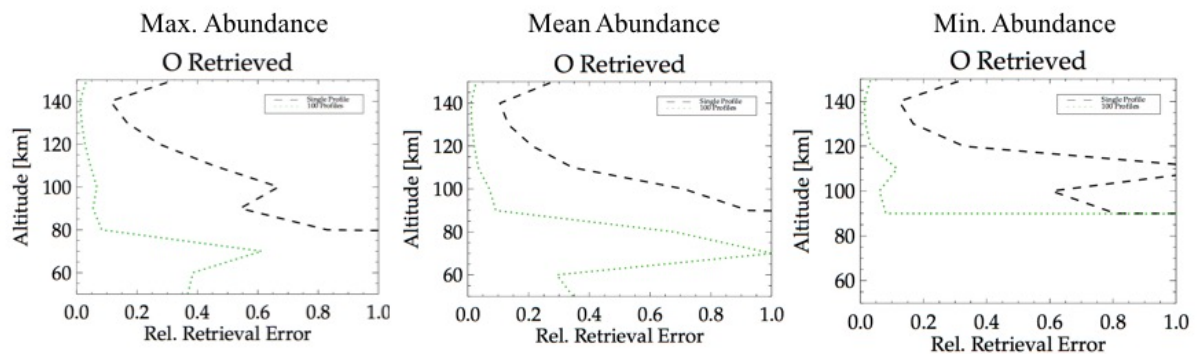


Figure 54: Climatological variability of O retrievals show the expected diurnal variation due to photo-chemistry in the 80-100km altitude range. The average profile is roughly halfway between the extremes, implying that both extremes are roughly equally likely. Above 120km the O concentration is always detectable.




Figure 55: Climatological variability of OH retrievals show that OH can be completely depleted at times, and be pretty much undetectable from a single measurement. The average profile however is still very much detectable, even though it phases out more quickly at the top of the OH presence in the atmosphere at ~90km. It's not clear why for the average profile only, the retrieval errors pick up a signal again at the topmost layer at 150km; this is most likely an artefact from a cross-talk between averaging kernels at remote altitudes.




Figure 56: Climatological variability of NO retrievals. Similar to OH, the depleted NO abundance is so low, that the detectability becomes marginal from single measurements. For the average abundance profile the detection is still feasible, but the altitude range over which the detection is feasible becomes narrow around the peak at 100km.




Figure 57: Climatological variability of NO+ retrievals is moot to discuss at the level of a single measurement, because the spectral line is too weak. We have dedicated a whole section to the detectability of Nitrosylium by looking at the signal to noise ratio of the spectral signature under various scenarios of atmospheric enhancements.




Figure 58: Climatological variability of CO retrievals. As expected from a classical tracer species, the lifetime of CO in the atmosphere is so long, that there is no significant difference between the Max and Min cases. This is reflected in the retrieval errors for climatological extremes.




Figure 59: Climatological variability of H_2O retrievals show that H_2O really only varies significantly at the top of its presence in the atmosphere, i.e. above the Mesosphere. Mesospheric water vapour is always detectable from a single measurement.

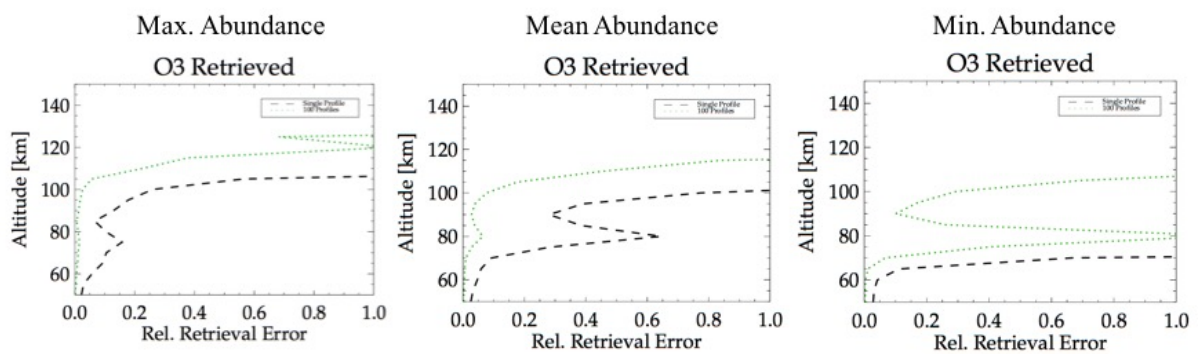


Figure 60: Climatological variability of O_3 retrievals testify the chemical reactivity of ozone as a species. O_3 is permanently detectable from a single measurement in the lower Mesosphere only (and below). In the upper Mesosphere and above, detectability varies dramatically with the atmospheric state. In the best case, it can be detected well into the Thermosphere at 100km. For an average case this might still hold true, but with visibly increased retrieval errors in the upper Mesosphere.




Figure 61: Climatological variability of HO_2 retrievals show that it will be difficult to measure HO_2 at the lowest possible abundances, but that the situation looks fairly positive for average abundance. This implies that a total depletion of HO_2 in the atmosphere is not the norm. As with OH , we don't know where the delta-peak artifacts at the top of the scan range come from, but suspect these are from cross-talk in averaging kernels between remote altitude layers.

Impact of Spectral Resolution

The baseline for the spectral resolution of LOCUS is 3MHz, with a target of 1MHz. The LOCUS spectrometer are digital Fourier transform spectrometers, meaning that their spectral resolution can be freely selected – within boundaries set by the Field-Programmable Gate Array (FPGA) – in a trade off against total instantaneous bandwidth. We have therefore studied the impact of increasing the spectral resolution from 3MHz to 1MHz (and possibly beyond) on the retrieval performance. This is illustrated in Figure 62 for the most important LOCUS species.

The finding is that the impact of spectral resolution is highly varied. NO seems to benefit quite strongly from it, and over a wide altitude range. The higher spectral resolution of 1MHz over 3MHz leads to a reduction in retrieval errors by ~20%, all else equal. For atomic oxygen at 2THz a benefit of similar magnitude is observed, however limited to a very narrow altitude interval at 100km (i.e. the altitude of the strongest abundance gradient and day-night variability). Interestingly, the primary atomic oxygen line at 4.7THz does not seem to be nearly as sensitive to spectral resolution than its weaker counter-part at 2THz. This must be because the 4.7THz is not only stronger, but also significantly wider, which seems to mitigate the need for extra fine spectral resolution.




Figure 62: Improvement in retrieval errors from increasing the spectral resolution from 3MHz (baseline) to 1MHz (target). This is expressed in the fraction of the retrieval error profile at 3MHz, divided by the retrieval error profile at 1MHz. The impact is limited to certain species and altitude ranges. For NO and O at 2THz a 20% improvement in retrieval error can be expected, but for O this is limited to a very narrow range around 100km (the interval of the strongest gradient). Interestingly, O at 4.7THz does not respond nearly as strongly to an increase in spectral resolution (only about 5% improvement is to be gained). This must be because the stronger spectral line at 4.7THz is also much wider, making spectral resolution less relevant.

Nitrosylium Detectability

Nitric oxide is one of the main scientific targets of the LOCUS mission. Its composition is driven by auroral forcing, and therefore indicative of solar wind activity. A recent paper by Chernicharo et al. [Chernicharo 2014] reports to have measured the sub-millimeter emission lines of the Nitric oxide ion Nitrosylium in a laboratory experiment. The measurements range from 238.4GHz up the THz frequencies at 953.2GHz. This has prompted us to test, if a THz radiometer that measures NO, could also detect its ion NO^+ .

Nitrosylium is a component of the Ionosphere, the layers of ionized particles that reflect long-wavelength radio signals, and makes radio communication over long distances possible. Nitrosylium is present in the lower ionospheric layers, the D- and F-layers. It's

particularly important in the lowest ionospheric layer, the D-layer between 60-90km.

Nitrosylium in the D-layer can get dramatically enhanced during extreme solar proton event, which can lead to a total telecommunication breakdown. If Nitrosylium was detectable from space, this would add significant scientific merit to the LOCUS mission.

We have therefore added NO^+ to the state vector of the retrieval processor, in order to assess the retrievability based on various atmospheric profiles. From a re-run of the WACCMX model we've compiled climatological (Max/Mean/Min) profiles in the same manner as for all the other species.

Signal to Noise Ratio

To quantify the detectability, we were looking at the signal to background noise ratio of the spectral emission line. We define the signal to noise ratio as the ratio between the line amplitude and the thermal measurement noise at that frequency ($\text{NE}\Delta T$). The line amplitude is the difference between the Brightness Temperature at the line centre, minus the Brightness Temperature of the background level (i.e. the continuum):

$$SNR = \frac{BT_{Peak} - BT_{Continuum}}{\text{NE}\Delta T}$$

We have performed simulations for every NO^+ line in the HITRAN catalogue, and all limb-views. The noise figure was obtained by linear interpolation of the frequency dependent $\text{NE}\Delta T$ values listed in Table 5. One surprising finding was that, instead of weakening with frequency – as do the spectral line strength parameters in the HITRAN catalogue - the Nitrosylium spectral line peak strengths were highest at frequencies of $\sim 3\text{THz}$.

Lower State Energy

A careful examination of this behavior lead us to realise, that this was because HITRAN line strengths are listed for a standard temperature of 296K, but that the temperatures in the MLT region are substantially higher ($\sim 700\text{K}$). This has an impact on the lower energy state of the molecular transition, the value of which has a string temperature dependence. When correcting for MLT temperatures, the peak in the line strength parameters moves to higher THz frequencies, as the radiative transfer model predicts. This behavior is shown in Figure 63

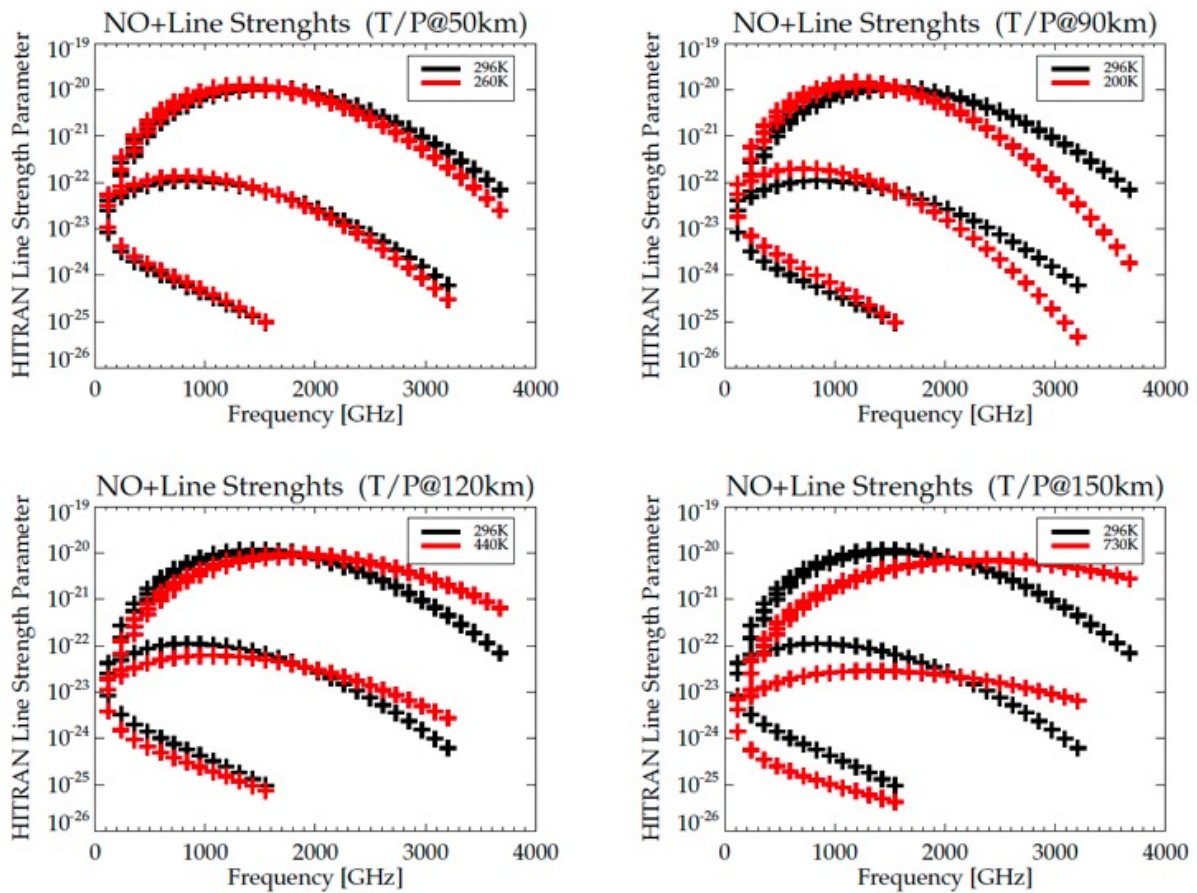


Figure 63: Correction of the spectral line strength parameters of Nitrosylum for the MLT altitudes. The shift to higher frequencies is rooted in the temperature dependence of the lower state energies of the rotational molecular energy transitions. This represent an inherit benefit to observe species with a peak abundance profile in the MLT region (where average temperatures are around 700K) at THz frequencies.

Results of Nitrosylium Detectability Study

The results of our simulations are shown in Figure 64 and Figure 65. For the best-case scenario of a maximally enhanced NO⁺ profile in Figure 64, the expected signal to noise ratio is exceeding 3. Furthermore, the peak signal to noise is observed at frequencies around 1THz. Of these, the transition at 1191GHz stands out by being very close to the existing LOCUS Band 3 with a centre frequency of 1148GHz. This would make it an attractive option for inclusion in the LOCUS science objectives, with minimal changes to the band definition. A signal to noise ratio of 3 should warrant a detection, especially when integrating over multiple measurements. This is still true for SNR 2 under average NO⁺ abundances, whereas under a depleted NO⁺ scenario the spectral line peak drops to the order of the noise level (as seen in Figure 65).




Figure 64: Detectability of Nitrosylium under a best-case scenario (maximal abundance profile, with enhanced D-layer). The best signal to noise ratio of just above 3 is obtained at frequencies around 1THz. The 1,191GHz line is close to the existing LOCUS Band 3 at 1,148GHz, which makes this a particularly attractive option.

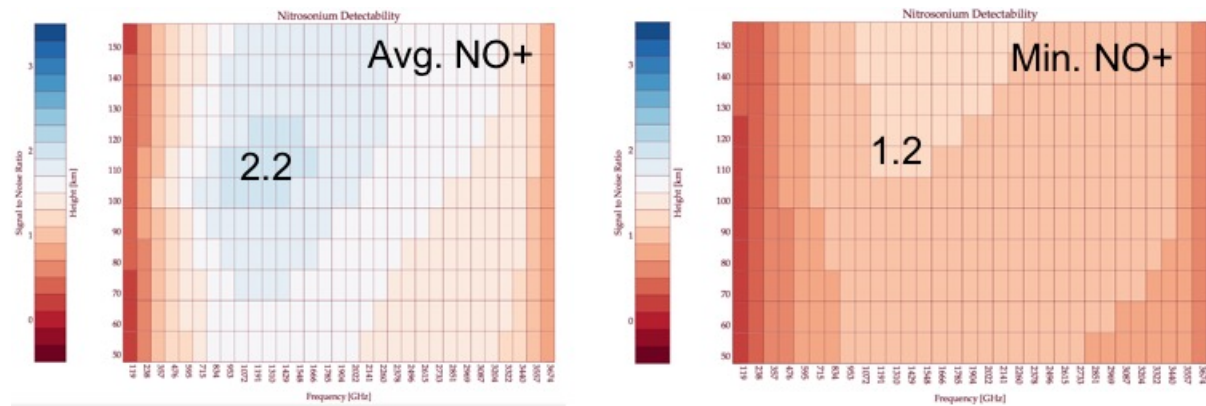


Figure 65: Detectability of Nitrosylium for a mean or minimal abundance profile respectively. While an average NO^+ abundance is borderline detectable with a signal to noise ratio of 2, especially when integrating over multiple measurements, a depleted NO^+ state means that the line signal drops to the order of the noise level.

Enhanced D-layer Nitrosylium

Because the detrimental impact of enhanced Nitrosylium is radio communication (i.e. the societal impact) is strongest in the D-layer, we have tested if it would be possible to detect an enhanced D-layer abundance. For this, a separate NO^+ abundance profile was modelled, representing the strong solar proton event of October 2003. This is shown in Figure 66. The enhanced profile was modelled with a different version of the WACCM model, which includes ionospheric D-layer chemistry.




Figure 66: Enhanced D-layer abundance profile of Nitrosylium, modelled for the extreme solar proton event in 2003

The top of the atmosphere for the D-layer WACCM model stops at 140km though, so for use in our retrieval simulations we had to combine the enhanced D-layer with the climatological WACCMX profiles, that span all the way up to a 500km top of the atmosphere. The procedure for this is illustrated in Figure 67. The dashed profiles represent the climatological Max/Mean/Min profiles of NO^+ from the WACCMX model field. The solid lines are modifications of the individual climatological profiles that feature the enhanced D-layer abundance.




Figure 67: Addition of the enhanced D-layer abundance to the climatological Nitrosylium profiles from our WACCMX model field

We've calculated simulated spectra of all these 6 scenarios, the three climatological scenarios without enhanced D-layer NO^+ , and the three climatological scenarios with enhanced NO^+ . The residual spectra (enhanced D-layer minus no D-layer) are shown in Figure 68. They show that an enhanced D-layer concentration could be detected (residual signature is larger than 5K in Brightness Temperature) if the Nitrosylium concentration in the F-layer above is low. It's not clear though if this is a realistic atmospheric scenario, as a strong solar proton event like the one in 2003 is likely to also enhance the F-layer of the ionosphere. However, we have already shown above that an enhanced F-layer Nitrosylium abundance is detectable even without the D-layer contribution.




Figure 68: Residual signature of an enhanced D-layer concentration of Nitrosylum as a function of three different climatological F-layer concentrations. An enhanced D-layer chemistry can reasonably be expected to be detectable for low F-layer abundances only, although we haven't established if this is a realistic atmospheric scenario.

Mission Requirements Table

Based on the findings from retrieval simulations described in the previous sections, we have compiled a Mission Requirements table, describing the conditions under which the LOCUS target species can be retrieved, and which what precision.

The altitude range over which a product can be retrieved is taken as the range where the retrieval errors are usefully low (i.e. below 40% for most species) to make a retrieval viable. For species with strong diurnal variations, detectability is defined for an average to enhanced abundance scenario, acknowledging the fact that for a very variable species a minimum abundance can be equivalent to complete chemical or photo-chemical depletion, which is case not worth including. Where ranges of retrieval errors. For some species like atomic oxygen, for which the exact atmospheric distribution is barely known, and which features a vast natural variability of 300%-500%, even retrieval errors of 100% are scientifically viable, so they are included in the table. Where multiple numbers of retrieval errors are given, this implies that the retrieval performance differs significantly over the retrieval interval, but that retrievals at the larger error range are still considered useful.

The vertical resolution is taken from the analysis of the averaging kernel matrix, in conjunction with the relative retrieval error profiles (as discussed in the corresponding section).

This table was included in our ESA Earth Explorer 9 proposal in the Mission Requirements Document (MRD). Table 7 is the full table of all results from this current study, discerning between the retrieval results from single, and from averaged measurements.

Table 7: Full mission requirements table from retrieval simulations from this study. Results are listed separately for retrievals from a single measurement (1 atmospheric lib-scan), as well as an average of 100 measurements. This accounts for the fact that some of the weaker spectral lines will require some sort of statistical averaging before they can be retrieved with sufficient precision.

Variable	Data Product	Altitude Range	Estimated Precision	Vertical Sampling
O	Single Profile [ppmv]	90km ¹ – 150km	60% – 20%	10km
	100 Averaged Profiles [ppmv]	80km – 150km	20% – 20%	5km
OH	Single Profile [ppmv]	50km – 90km	20% – 20%	5km
	100 Averaged Profiles [ppmv]	50km – 110km	10% – 20%	5km
HO₂	Single Profile [ppmv]	50km – 85km	20% – 20%	10km
	100 Averaged Profiles [ppmv]	50km – 100km	10% – 40%	5km
H₂O	Single Profile [ppmv]	50km – 100km	10% – 20%	2km
	100 Averaged Profiles [ppmv]	50km – 120km	10% – 40%	2km
NO	Single Profile [ppmv]	50km – 150km	20%-60%	5km
	100 Averaged Profiles [ppmv]	50km – 150km	10%-10%	5km
NO+	Single Profile [ppmv] ²	N/A	N/A	N/A
	100 Averaged Profiles [ppmv]	70km – 150km	40%-20%	5km
CO	Single Profile [ppmv]	50km – 100km	10%-20%	5km
	100 Averaged Profiles [ppmv]	50km – 150km	10%-20%	5km

¹ The lower limit of 90km is for enhanced abundances. For average abundances, profiles are retrievable from 110km-150km from a single measurement, or from 80km-150km when integrating over 100 samples

² NO+ retrievals are not viable from just a single measurement. Multiple measurements will have to be integrated in order to reduce the measurement error

O₃	Single Profile [ppmv]	50km – 100km	10%-40%	5km
	100 Averaged Profiles [ppmv]	50km – 120km	10%-40%	5km

Table 8 is the condensed version thereof, as presented in the executive summary of the EE-9 proposal. It combines the expected retrieval performance of the THz instrument – as established in this study – with the measurement performance of the infra-red instrument. The latter is, by design, a replica of the SABER performance.

Table 8: Simplified mission requirements table as presented in the executive summary of the EE-9 proposal, combining the expected retrieval performance from the THz instrument with the retrieval performance of the infra-red instrument (which per design replicates the SABER measurement performance).

Geophysical Parameters

Product Name	Height Range	Vertic. Resol.	Retrieval Precision	Scientific Rationale
O	90-150	10km	60%-20%	Thermal Balance/Chemistry/Dynamics
OH	50-90	5km	20%	Chemistry
NO	50-150	5km	20%-60%	Thermal Balance/Space Weather
CO	50-100	5km	10%-20%	Mesospheric Tracer
H₂O	50-100	2km	10%-20%	Climate/NLC Formation
O₃	50-100	5km	10%-40%	Climate
HO₂	50-90	10km	20%	Chemistry
Temp. *	50-110	2 km	2K-10K	Thermal Balance/Climate/Dynamics
Cooling Rates	From CO ₂ , NO and O ₃ infra-red channels analogous to SABER / TIMED mission performance. (*) Temperature retrieval performance from SABER			

Conclusions

We have modified and used the WACCMX model to provide global model fields of all the LOCUS target species at altitudes from 40km up to 500km. The model data cover a three-day period starting at 1 Jan 2000, with a high temporal resolution of 3 hours. We extracted climatological reference profiles (Max/Mean/Min) from this database to perform retrieval simulations.

We have written a retrieval processor for LOCUS in the interactive data language (IDL), using an Optimal Estimation Method (OEM) retrieval scheme with Levenberg-Marquardt regularisation. The processor can use any fix-point in the WACCMX model field, or the equivalent climatology, as first guess and/or prior profiles. Using this retrieval code, we have performed Zero-order iteration retrievals to compute the retrieval diagnostics parameters, i.e. the state error covariance matrix, and the averaging kernels (vertical resolution matrix). From these, the retrieval precision (aka retrieval errors) were extracted. We have used flat a priori error profiles of large magnitude in order not to constrain the retrieval, and to minimize information content arising from the prior knowledge. The measurement errors are defined as the noise performance of the individual receivers.

An assessment of retrieval errors for multiple scenarios of retrieval grid spacings, atmospheric abundances, and spectral resolution has been performed. From the analysis of these results we have compiled a table that quantifies the altitude ranges over which each species can be retrieved, the best possible vertical resolution that can realistically be achieved (without introducing a smoothing error), and ultimately the expected retrieval precisions from a single measurement, or from an average over 100 measurement samples).

The resulting table forms the core of the Mission Requirement Document in our ESA Earth Explorer 9 Phase-A proposal. The results quantify the link between geophysical parameters (atmospheric composition) to measurements (Level2 data) that defines SRL 4.

Acknowledgments

Work under this study was funded by the Centre for Earth Observation Instrumentation (CEOI) under contracts RP10G0327E14 (LOCUS Preparatory Activities for EE-9), and RP10G0435D03 (LOCUS Preparatory Activities for EE-10). The Centre for Earth Observation Instrumentation was established in 2007 in order to realise an excellent, internationally competitive national Earth observation instrument and technology research and development programme in the UK. The CEOI is funded by the UK Space Agency with parallel technology investment from industry. Its key aim is to develop UK innovative technologies to observe Earth from space through the teaming of scientists and industrialists.

References

- Cernicharo, J., et al., (2014), Tentative Detection of the Nitrosylium Ion in Space, *The Astrophysical Journal*, 795:40 (6pp), doi:10.1088/0004-637X/795/1/40.
- Dudhia, A. (2015), Reference Forward Model, <http://eodg.atm.ox.ac.uk/RFM>, Release v4.34, 1st Oct 2015.
- ESA-SRL (2015), Scientific Readiness Levels (SRL) Handbook, Reference: EOP-SM/2776, Issue: 1, Revision: 1, 17/06/2015
- Feng, W., D. R. Marsh, M. P. Chipperfield, D. Janches, J. Hoffner, F. Yi, and J. M. C. Plane (2013), A global atmospheric model of meteoric iron, *J. Geophys. Res.- Atmos.*, 118, 9456--9474, doi:10.1002/jgrd.50708.
- Feng, Wuhu, Bernd Kaifler, Daniel R. Marsh, Josef Hoffner, Ulf-Peter Hoppe, Bifford P. Williams and John M. C. Plane, Impacts of a sudden stratospheric warming on the mesospheric metal layers, *J. Atmos. Sol.-Terr. Phys.*, <http://dx.doi.org/10.1016/j.jastp.2017.02.004>, 2017.
- French, W. J. R. and Mulligan, F. J.: Stability of temperatures from TIMED/SABER v1.07 (2002-2009) and Aura/MLS v2.2 (2004-2009) compared with OH(6-2) temperatures observed at Davis Station, Antarctica, *Atmos. Chem. Phys.*, 10, 11439-11446, doi:10.5194/acp-10-11439-2010, 2010.
- Garcia, R. R., D. Marsh, D. E. Kinnison, B. Boville, and F. Sassi, Simulations of secular trends in the middle atmosphere, 1950-2003, *J. Geophys. Res.*, 112, D09301, doi:10.1029/2006JD007485, 2007.
- Gerber, D., et al. (2016), LOCUS – Linking Observations of Climate, the Upper-atmosphere, and Space-weather, Phase-A Proposal for the 9th ESA Earth Explorer Mission, ESA Ref.: CEE9/017, 2016

- Gray, L. J., et al. (2010), Solar influences on climate, *Rev. Geophys.*, 48, RG4001, doi:10.1029/2009RG000282
- Grossmann, K. U. (1997), The fine structure emission of thermospheric atomic oxygen. *Adv. Space Res.*, 19, 595-598.
- Grossmann, K. U. (2000), A global measurement of lower thermospheric atomic oxygen densities, *Geophysical Research Letters*, 27(9), 1387-1390.
- Hedin, A. E. (1991), *Extension of the MSIS Thermospheric Model into the Middle and Lower Atmosphere*, *J. Geophys. Res.* 96, 1159.
- Huang, W., X. Chu, C. S. Gardner, J. D. Carrillo-Sanchez, W. Feng, J. M. C. Plane, and D. Nesvorny (2015), Measurements of the vertical fluxes of atomic Fe and Na at the mesopause: Implications for the velocity of cosmic dust entering the atmosphere, *Geophys. Res. Lett.*, 42, doi:10.1002/2014GL062390.
- Huang, F. T., H. G. Mayr, C. A. Reber, J. M. Russell III, M. G. Mlynczak, and J. G. Mengel, Ozone quasi-biennial oscillations (QBO), semiannual oscillations (SAO), and correlations with temperature in the mesosphere, lower thermosphere, and stratosphere, based on measurements from SABER on TIMED and MLS on UARS, *J. Geophys. Res.*, 113, A01316, doi:10.1029/2007JA012634, 2008.
- Hurrell, J.W., M.M. Holland, P.R. Gent, S. Ghan, J.E. Kay, P.J. Kushner, J. Lamarque, W.G. Large, D. Lawrence, K. Lindsay, W.H. Lipscomb, M.C. Long, N. Mahowald, D.R. Marsh, R.B. Neale, P. Rasch, S. Vavrus, M. Vertenstein, D. Bader, W.D. Collins, J.J. Hack, J. Kiehl, and S. Marshall, 2013: The Community Earth System Model: A Framework for Collaborative Research. *Bull. Amer. Meteor. Soc.*, **94**, 1339–1360, <https://doi.org/10.1175/BAMS-D-12-00121.1>

- Lauritzen, P. H., Bacmeister, J. T., Callaghan, P. F., and Taylor, M. A.: NCAR global model topography generation software for unstructured grids, *Geosci. Model Dev.*, 8, 1-12, doi:10.5194/gmd-8-1-2015, 2015.
- Lin, S.-J., A “vertically-Lagrangian” finite-volume dynamical core for global atmospheric models, *Mon. Wea. Rev.*, 132, 2293-2307, 2004.
- Liu, H.-L., Bardeen, C. G., Foster, B. T., Lauritzen, P., Liu, J., Lu, G., Wang, W. (2018). Development and validation of the Whole Atmosphere Community Climate Model with thermosphere and ionosphere extension (WACCM-X 2.0). *Journal of Advances in Modeling Earth Systems*, 10, 381–402. <https://doi.org/10.1002/2017MS001232>
- Liu, H.-L., Foster, B. T., Hagan, M. E., McInerney, J. M., Maute, A., Qian, L., et al. (2010). Thermosphere extension of the Whole Atmosphere Community Climate Model. *Journal of Geophysical Research: Space Physics*, **115**, A12302. <https://doi.org/10.1029/2010JA015586>
- Marsh, D. R., M.J. Mills, D.E. Kinnison, J.-F. Lamarque, N. Calvo, and L. M. Polvani, Climate change from 1850 to 2005 simulated in CESM1(WACCM), 73727391, *Journal of Climate*, 26(19), doi:10.1175/JCLI-D- 12-00558.1, 2013.
- Mills, M. J., A. Schmidt, R. Easter, S. Solomon, D. E. Kinnison, S. J. Ghan, R. R. III Neely, D. R. Marsh, A. Conley, C. G. Bardeen, et al. (2016), Global volcanic aerosol properties derived from emissions, 1990–2014, using CESM1(WACCM), *J. Geophys. Res. Atmos.*, 121, 2332–2348, doi:10.1002/2015JD024290.
- Mills, M. J., Richter, J. H., Tilmes, S., Kravitz, B., MacMartin, D. G., Glanville, A. A., Kinnison, D. E. (2017). Radiative and chemical response to interactive stratospheric sulfate aerosols in fully coupled CESM1(WACCM). *Journal of Geophysical Research: Atmospheres*, 122, 13,061–13,078. <https://doi.org/10.1002/2017JD027006>

- Moloney, J. V., et al. (2014), Compact, high-power, room-temperature, narrow-line terahertz source, 15 March 2011, SPIE Newsroom. DOI: 10.1117/2.1201102.003523
- Neale, R.B., C.C. Chen, A. Gettelman and Coauthors, 2012: Description of the NCAR Community Atmosphere Model (CAM 5.0). NCAR Tech. Note NCAR-TN-486+STR, 274 pp.
- Plane, J. M. C. (2003), Atmospheric chemistry of meteoric metals, *Chem. Rev.*, 103(12), 4963-4984.
- Remsberg, E. E., et al., Assessment of the quality of the Version 1.07 temperature-versus-pressure profiles of the middle atmosphere from TIMED/SABER, *J. Geophys. Res.*, 113, D17101, doi:10.1029/2008JD010013, 2008.
- Russell, J. M., III, M. G. Mlynczak, L. L. Gordley, J. Tansock, and R. Esplin, An overview of the SABER experiment and preliminary calibration results, paper presented at the 44th Annual Meeting, Int. Soc. for Opt. Eng., Denver, Colo., 18--23 July 1999.
- Smith et al., Satellite observations of high nighttime ozone at the equatorial mesopause, *J. Geophys. Res.*, 113, D17312, doi:10.1029/2008JD010066, 2008.
- Smith, A. K., D. R. Marsh, M. G. Mlynczak, and J. C. Mast, Temporal variations of atomic oxygen in the upper mesosphere from SABER, *J. Geophys. Res.*, 115, D18309, doi:10.1029/2009JD013434, 2010.
- SSTL-IOD (2014), Requirements and Concepts for IOD Missions for Breakthrough Concepts and Approaches ("Integrated THz Mission for Atmospheric Sounding"), Surrey Satellite Technology Ltd. (SSTL), Final Report of ESA Contract 4000109287/13/NL/MV, SSTL Ref.: 0237735, Rev.: 001.011, 13/08/2014.
- Swinyard, B., Ellison, B., Gerber, D., Plane, J., Feng, W., Nocerino, T., Navarathinam, N. (2013), Low Cost Upper-atmosphere Sounder - LOCUS CEOI Study Report, Final Report of the UKSA/CEOI study.

- Rodgers, C.D. (2000), *Inverse Methods for Atmospheric Sounding: Theory and Practice*, World Scientific.
- Xu, J., C. Y. She, W. Yuan, C. Mertens, M. Mlynczak, and J. Russell, Comparison between the temperature measurements by TIMED/SABER and lidar in the midlatitude, *J. Geophys. Res.*, 111, A10S09, doi:10.1029/2005JA011439, 2006.
- Xu, J., H. Gao, A. K. Smith, and Y. Zhu, Using TIMED/SABER nightglow observations to investigate hydroxyl emission mechanisms in the mesopause region, *J. Geophys. Res.*, 117, D02301, doi:10.1029/2011JD016342, 2012.
- Zhang, X., J. M. Forbes, M. E. Hagan, J. M. Russell III, S. E. Palo, C. J. Mertens, and M. G. Mlynczak, Monthly tidal temperatures 20-120 km from TIMED/SABER, *J. Geophys. Res.*, 111, A10S08, doi:10.1029/2005JA011504, 2006.