




Contents lists available at ScienceDirect

Journal of Quantitative Spectroscopy & Radiative Transfer

journal homepage: www.elsevier.com/locate/jqsrt

Fast radiative transfer using monochromatic look-up tables

R. Anthony Vincent, Anu Dudhia

Atmospheric, Oceanic and Planetary Physics, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, UK

ARTICLE INFO

Article history:

Received 12 February 2016

Received in revised form

11 April 2016

Accepted 11 April 2016

Keywords:

Radiative transfer

Line-by-line models

Absorption cross section

Look-up tables

ABSTRACT

Line-by-line (LBL) methods of numerically solving the equations of radiative transfer can be inhibitingly slow. Operational trace gas retrieval schemes generally require much faster output than current LBL radiative transfer models can achieve. One option to speed up computation is to precalculate absorption cross sections for each absorbing gas on a fixed grid and interpolate. This work presents a general method for creating, compressing, and validating a set of individual look-up tables (LUTs) for the 11 most abundant trace gases to use the Reference Forward Model (RFM) to simulate radiances observed by the Infrared Atmospheric Sounding Interferometer (IASI) at a more operational pace. These LUTs allow the RFM to generate radiances more than 20 times faster than LBL mode and were rigorously validated for 80 different atmospheric scenarios chosen to represent variability indicative of Earth's atmosphere. More than 99% of all IASI simulated spectral channels had LUT interpolation errors of brightness temperature less than 0.02 K, several factors below the IASI noise level. Including a reduced spectral grid for radiative transfer speed up the computation by another factor of six at the expense of approximately doubling interpolation errors, still factors below IASI noise. Furthermore, a simple spectral compression scheme based upon linear interpolation is presented, which reduced the total LUT file size from 120 Gbytes to 5.6 Gbytes; a compression to just 4.4% of the original. These LUTs are openly available for use by the scientific community, whether using the RFM or to be incorporated into any forward model.

Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The retrieval of atmospheric temperature and composition from infrared measurements is usually an iterative process requiring repeated calculations of the radiative transfer equation to simulate the observations. In its simplest form the radiance L reaching the satellite at wavenumber ν from viewing geometry (e.g., tangent height, scan angle) z can be expressed as the sum of atmospheric and background terms:

$$L(\nu, z) = \int_{\tau_0}^1 B(\tau) d\tau + B_0 \tau_0, \quad (1)$$

E-mail addresses: tony.vincent@physics.ox.ac.uk (R. Anthony Vincent), anu.dudhia@physics.ox.ac.uk (A. Dudhia).

<http://dx.doi.org/10.1016/j.jqsrt.2016.04.011>

0022-4073/Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

where τ is the transmittance along the line-of-sight from the satellite ($\tau = 1$) to the remote boundary of the atmosphere ($\tau = \tau_0$), B is the Planck function along this path, and B_0 is the Planck function of the background at the far side of or beyond the atmospheric path.

The measurement itself, R_{ij} , for nominal wavenumber ν_i and viewing geometry z_j , is modelled as a convolution of this radiance with the appropriate instrument functions

$$R_{ij} = \iint L(\nu, z) \Phi(z - z_j) \Psi(\nu - \nu_i) d\nu dz. \quad (2)$$

Φ_j represents the instrument field-of-view (usually omitted with nadir-viewing) and Ψ_i represents the spectral response (the instrument line shape for Fourier transform spectrometers, or channel response for radiometers).

The three main geometries for infrared remote sensing can be distinguished by the relative importance of the

background term in the radiative transfer equation; for solar occultation $B_0\tau_0$ is dominant, for limb-viewing it is negligible and for nadir-viewing it is comparable to the atmospheric term. However, in all cases, since B is only a function of wavenumber and temperature, information on the composition and pressure (as well as temperature) is derived from the transmittance.

Calculating transmittance is usually the time-consuming part of the retrieval process. The most accurate method is to use a line-by-line (LBL) model. However, this is slow and only feasible in real-time data processing for solar occultation measurements, which typically have a maximum of 28 profile acquisitions per day.

Nadir-viewing instruments, such as those commonly used for operational weather forecasting, provide measurements at typically 10^5 – 10^6 locations per day and this presents a greater challenge. With (spectrally broad) measurements from filter radiometers, the usual method is to apply the radiative transfer to pre-computed spectrally integrated quantities, effectively:

$$R_i = \int \bar{B} d\bar{\tau} + \bar{B}_0\bar{\tau}_0 \quad (3)$$

where

$$\bar{\tau} = \int \tau(\nu)\Psi(\nu - \nu_i) d\nu \quad (4)$$

and similarly for \bar{B} [1,2]. This differs fundamentally from the monochromatic approach in that Beer's Law does not hold for spectrally integrated transmittances. In other words, the net transmittance of a composite path is no longer simply the product of the transmittances of the component paths when dealing with averaged quantities. This can be readily appreciated if one considers the net transmittance through two consecutive identical cells, each opaque over the lower half of the spectral range and transparent in the upper half. The spectrally averaged transmittance of each cell is 0.5, but the transmittance of the two cells combined is, clearly, also 0.5, rather than the product $0.5 \times 0.5 = 0.25$. To overcome this it is necessary to parametrise spectrally averaged transmittances not only as a function of the 'natural' quantities such as pressure and temperature, but also a number of ad hoc 'predictors' such as geometry and contaminants, relating to each channel and generally obtained by statistical regression from a large sample (with all the risks that entails when dealing with anomalous atmospheric cases) [3].

The recent trend has been towards replacing broad band filter radiometers, such as the High-resolution Infrared Radiation Sounder (HIRS; order of 10 channels of widths ~ 10 's of cm^{-1}), with Fourier transform or grating spectrometers (order of 1000 channels with resolution $\sim 1 \text{ cm}^{-1}$). Extending the spectrally averaged approach requires treating each spectral sample as an independent channel with its own set of tailored predictors. While this remains the 'fast' solution, it becomes unwieldy and, noting the steady increase in computing speeds, it is anticipated that a monochromatic approach will gradually be adopted even for the operational nadir-viewing instruments.

Significant time savings can be made by using the line-by-line model to pre-compute monochromatic transmittances

$\tau(\nu)$ (or, more usually, a related quantity such as absorption cross section), over a variety of atmospheric conditions, and storing these in a 'look-up table' (LUT) for use in the retrieval forward model. Such an approach has, for example, been used for the Michelson Interferometer for Passive Atmospheric Sounding (MIPAS) [4] limb-sounding instrument and the nadir-viewing Infrared Atmospheric Sounding Interferometer (IASI) [5]. The use of monochromatic LUTs of absorption cross sections is not new [6]. However, in many cases the LUTs themselves are closely tied to their application and details of their generation, testing, and optimisation are not widely published.

The aim of this paper is to review the criteria required to construct such LUTs and the various methods chosen for their implementation. Simulations of IASI radiances are used as a test case. Furthermore, a detailed description is given of the construction and testing of LUTs which can be implemented in the Reference Forward Model (RFM) for radiative transfer [7] as a direct replacement for the line-by-line calculation, which therefore allows comparisons of the speed/accuracy trade-off.

Finally, given that these LUTs require a relatively simple file format (even if the files are large) it is hoped that this paper might encourage such LUTs to be regarded as independent databases, decoupled from their originating line-by-line models and any particular instrumental application.

2. Calculating transmittance

It will be assumed that a generic atmospheric radiative transfer model (RTM) represents the transmittance of a line-of-sight path from the satellite through an inhomogeneous atmosphere as the product of the monochromatic transmittances of the component path segments,

$$\tau = \prod_l \tau_l. \quad (5)$$

Each segment (l) is defined, for example, by the intersection of the line-of-sight with the internal vertical grid on which the atmospheric profile is represented. In turn, the transmittance of each segment is related to the optical depth (χ_g) of each absorbing species (g) through

$$\tau_l = \exp\left(-\sum_g \chi_{gl}\right). \quad (6)$$

The optical depths themselves are computed from

$$\chi_{gl} = k_{gl}u_{gl}, \quad (7)$$

where k is the absorption cross section (e.g., units of $\text{m}^2 \text{ mol}^{-1}$) and u is the integrated absorber amount within the segment (inverse units of k). As an example, Fig. 1 shows the contribution to optical depth for a water vapour spectral line in a tropical atmospheric scenario for each 100 hPa tropospheric layer. Notice that the optical depth shape changes depending upon the physical atmospheric conditions of that layer. The total optical depth is then the sum of the optical depths from each layer.

The absorption cross section itself, as well as its spectral dependence, is also a function of pressure, temperature and, sometimes, absorber concentration if the absorbing




Fig. 1. Optical depths of a water vapour spectral line are shown for tropospheric layers of 100 hPa for a typical tropical atmosphere. Each optical depth curve is colour coded to the approximate atmospheric temperature of that layer. These optical depths were calculated with the RFM using the 2012 HITRAN database. All spectroscopic parameters listed in HITRAN were used, including those accounting for H₂O self-broadening and asymmetric pressure shifts. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

molecule has a permanent electric dipole moment. Typically k would be evaluated for absorber weighted mean values of these quantities within the segment, i.e., approximating the optical depth of the inhomogeneous path segment as being the same as an equivalent homogeneous path characterised by these mean quantities. The Curtis–Godson (CG) approximation [8–10] is commonly used to set an equivalent pressure, \bar{p}_g , for gas g as

$$\bar{p}_g = \frac{\int_z^{z+\Delta z} p(z) \rho_g(z) dz}{\int_z^{z+\Delta z} \rho_g(z) dz} \quad (8)$$

and an equivalent temperature as

$$\bar{T}_g = \frac{\int_z^{z+\Delta z} T(z) \rho_g(z) dz}{\int_z^{z+\Delta z} \rho_g(z) dz}. \quad (9)$$

Each absorbing gas has layer values distinct from those defined by the atmospheric levels. This potentially complicates the process of setting LUT grids, as each gas could have its own CG pressure and temperature grid.

Fig. 2 shows the RFM results of the CG approximation using the same water vapour line and tropical atmosphere from Fig. 1. The top plot shows that using just one homogeneous layer with the CG pressure and temperature yields a surprisingly accurate optical depth spectral shape compared to a 100 layer inhomogeneous atmosphere. The bottom plot residuals also show the effect empirical H₂O pressure shifts and self-broadening [11–13] have upon the accuracy of the CG approximation. Interestingly, including self-broadening from the High Resolution Transmission (HITRAN) database [14] in the optical depth calculation leads to less residual error for the one layer example. Caution is needed, however, to avoid concluding that since the optical depth residual error is $\approx 3\%$ that only one or a few layers are necessary. The spectral line considered in Fig. 2 is largely Lorentzian. For lineshapes [15] falling in the Voigt regime, the resulting radiance error is often significant

compared to instrument noise for anything less than at least 30 layers [16, chapter 5].

Finally, the difference between line-by-line and LUT calculations is in the evaluation of k : For a line-by-line model, k is evaluated as the sum of contributions of all local transitions appropriately adjusted for the CG mean values; while for a LUT model, k is interpolated from pre-tabulated axis values.

3. Previous approaches

In principle, precomputing LUTs for a particular molecule, is straightforward; use a line-by-line model to calculate $k(\nu, p, T, e)$ (where p is pressure, T is temperature and e is partial pressure for water vapour) at a sufficient range and density of axis points that atmospheric radiance calculations with interpolated LUTs are of acceptable accuracy compared to the equivalent uninterpolated line-by-line calculation. However, there are two immediate practical problems.

First, in order to meet the ‘monochromatic’ criterion that atmospheric spectral features can be adequately resolved (set by the Doppler width of atmospheric lines), a wavenumber spacing of 0.001 cm^{-1} is generally required for nadir-viewing and 0.0005 cm^{-1} for limb-viewing. Assuming of the order of 1000 points to represent the p, T, e domains, any LUT with a uniform spectral axis will require 10^6 real numbers just to represent the absorption cross section of a single absorber in a 1 cm^{-1} interval.

The second problem is defining the required accuracy. It is quite simple to define axis spacing such that k can be reconstructed with a given relative or absolute accuracy. The problem is that the monochromatic radiance reaching the satellite is not linearly related to k , and that is without even considering the effects of other absorbers or spectral convolution.

Scott and Chedin [6] were among the first to bypass LBL processing time by creating a set of LUTs known as the Automated Atmospheric Absorption Atlas for use with HIRS and the Infrared Interferometer Spectrometer (IRIS). Transmittance values were predetermined over a range of plausible pressure, temperature, and H₂O partial pressure scenarios and stored in LUTs at each spectral increment. By interpolating transmittance from the LUTs for the various atmospheric conditions, radiance calculation speeds were increased by a factor of 15–40. Parametrisation error was found to be negligible for the desired applications when compared to instrument noise.

The base steps to create a faster monochromatic RTM based on LUTs are summarised in the paper by Amato et al. [17] describing the σ -IASI RTM designed specifically for atmospheric retrievals using IASI. In this model, individual optical depths for H₂O, CO₂, O₃, N₂O, CH₄, N₂, and O₂ were tabulated from the Line-By-Line Radiative Transfer Model (LBLRTM) [18] and thinned in size using methods to be discussed in Section 4.4. Grid thinning reduced the final LUTs from 27.7 Gbytes to 726 Mbytes. However, self-broadening effects of H₂O line shapes due to variable partial pressures were not included, resulting in radiance error approaching IASI noise levels in spectral regions




Fig. 2. Top: Optical depth for a H_2O spectral line from a tropical atmosphere, the sum of layer contributions from Fig. 1. The solid curve shows the optical depth for a 100 layer inhomogeneous atmosphere and the dashed line shows optical depth for an equivalent one layer homogeneous atmosphere using the CG approximation. Bottom: The difference between the top curves (residual) is shown for three separately modelled conditions, which are classical ideal gas pressure broadening, H_2O self-broadening, and self-broadening including the H_2O collisional pressure shift. The reference spectra for the green and black curves are not plotted in the top diagram as they are indistinguishable by eye. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

associated with H_2O absorption for the six analysed atmospheric scenarios.

A similar process was repeated when creating the forward model for the Fast Optimal Retrieval on Layers for IASI (FORLI) [19], which is based upon previously defined LUTs and includes a partial pressure dimension for H_2O . As of publication, there were three versions of this scheme specific to retrievals of O_3 , HNO_3 , and CO . Since the spectral ranges considered in FORLI were small ($< 150 \text{ cm}^{-1}$), spectral compression of the LUTs was not needed. Absorption cross-sections at a given pressure and temperature interpolated from the LUTs were claimed to match the true values to less than 0.1% error. However, validation in radiance space was not mentioned, in particular for comparing radiances over a host of atmospheric conditions to those produced by a full LBL model. Validation of radiances generated from faster LUTs is a crucial step, because forward model errors inevitably become artefacts of estimated quantities [20].

Optical depth LUTs were also created for retrievals using the Atmospheric Infrared Sounder (AIRS) in much the same fashion, but Strow et al. [21] implemented a singular value decomposition (SVD) scheme to compress the ν dimension and make the LUT size more manageable. They showed that optical depth interpolation can occur in the compressed space, which sped up transmittance calculations by couple orders of magnitude while keeping interpolation errors

below a desired threshold in radiance. This compression method is further discussed in Section 4.4.

Finally, specific LUTs for MIPAS were generated for each microwindow used in the retrievals [22]. The procedure was to start with a high-density tabulation of $k(\nu, p, T)$ for each absorbing species (the maximum microwindow width of 3 cm^{-1} makes this feasible). These LUTs were then progressively sub-sampled in the p, T domain, compressed using SVD and finally sub-sampled in the ν domain, at each stage comparing the simulated radiances with those of a line-by-line model to ensure the error budget remained within 10% of the instrument noise.

4. Selecting look-up table axes

This paper outlines the process of generating, compressing, and validating LUTs to model nadir-viewing spectral radiances simulating those observed by IASI [5]; a polar orbiting Fourier transform spectrometer measuring in the thermal infrared. These LUTs are neither specific to the RTM that created them nor the instrument in consideration, but were designed for easy portability into any forward model to simulate instruments with similar spectral ranges and noise levels. Work presented here was performed using the Reference Forward Model to numerically simulate IASI radiances and to create the LUTs mentioned. The RFM is a line-by-line model originally written to simulate observations from the limb-viewing MIPAS. Based upon the GENLN2 model [23], the RFM has since been adapted to numerous viewing geometries, including nadir, and is now a general purpose RTM intended to be flexible, accurate, and easy to use. One possible use of the RFM is to simulate IASI observed spectral radiances. Configuration steps necessary to simulate IASI with the RFM are documented in Ventress [16, chapter 5].

Before creating LUTs, one must first decide upon the grid points along each dimension, i.e., temperature, pressure, and spectral wavenumber. For water vapour, there is also a partial pressure axis. Absorption cross sections, $k(\nu, p, T, e)$, must be tabulated at sufficient resolution in all four domains to avoid interpolation error. As seen in Eqs. (1) and (6), interpolation errors in k map non-linearly into errors in radiance space. Given a model atmosphere at defined pressure levels, it is somewhat trivial to create a LUT at the exact CG pressure and temperature values and reproduce highly consistent model radiances. The difficulty arises when attempting to create LUTs that encompass all possible atmospheric conditions and produce simulated radiances to a desired accuracy. As currently implemented, no extrapolation is performed for points lying outside the tabulated domain. Instead, the value at the closest boundary is used. As such, the LUTs were iteratively tested against an ensemble of atmospheres while constructing these grids by systematically comparing the changes in modelled spectral radiance to those produced in full LBL mode. Brightness temperature differences less than 0.02 K were desired, which is accuracy criteria approximately an order of magnitude less than IASI instrument noise levels in much of the spectral range.

The 80 atmospheres created by Matricardi [24] to parametrise the regression coefficients for RTTOV were used as the validation cases. Matricardi [24] took care when selecting the atmospheres so that the ensemble represents the statistical variability of Earth to include extreme seasonal events. Temperature and gas profiles are represented on a fixed grid of 101 pressure levels with 1 km spacing ranging from 1100 to 0.005 hPa. Profiles of temperature, H_2O , and O_3 were generated by the European Centre for Medium-Range Weather Forecasts (ECMWF) cycle30R2 forecasting system. All other trace gas profiles were sampled from the Global and regional Earthsystem Monitoring using Satellite and in situ data (GEMS) network [25]. Skin temperature was modelled to be 10 K greater than atmospheric surface temperature in order to enhance sensitivity of weak spectral lines and the pressure dependence of all lines in the lower atmosphere, which might otherwise be invisible in the absence of any temperature difference. Additionally a straight nadir viewing geometry was assumed. The objective is to meet the accuracy criteria over the entire 80 atmosphere ensemble while maintaining manageable file sizes with appreciable gains in computation speed.

4.1. Temperature

Creating robust LUTs requires that the ranges of temperature, pressure, and partial pressure bound all possible atmospheric scenarios. Fig. 3 shows the possible range of temperature for each pressure level represented in the 80 RTTOV atmosphere ensemble. The profile area within the solid black lines shows the absolute maximum and minimum ensemble temperature values, which is the space that needs to be sampled and not the total area represented in grey shading. Therefore, a temperature grid was used relative to the midpoints between maximum and minimum profile temperatures as opposed to a regular grid populating the entire space between the absolute minimum and maximum temperatures. Relative temperature spacing uses

approximately two thirds the memory as compared to sampling the entire grey area with the same sampling density. For reference, the dots in Fig. 3 show the temperature and pressure grid points for water vapour.

Absorption cross sections were calculated for all trace gases included in the ensemble atmospheres and are listed in Table 1. All relative temperature grids extend from ± 50 K from the midpoint temperature profile. Temperature sampling increments are displayed in the second row of Table 1 for each trace gas. These were determined by first starting with an overly dense sampling and increasing the temperature spacing width until the difference in radiance compared to full LBL mode became appreciable relative to the IASI instrument noise level.

While temperature affects both the relative line strengths and line shapes of an absorber, absorption cross sections smoothly vary with temperature for a given layer. As an example, Fig. 4 shows water vapour contour plots of absorption cross section for two different wavenumbers. For any given pressure level notice that absorption cross section changes with low curvature as a function of temperature. Therefore, Strow et al. [21] parametrised the temperature axis by fitting a 5th order polynomial through the optical depth values, while Amato et al. [17] used a 2nd order polynomial, thus reducing the temperature grid by a factor of 2–3. For the LUTs created and analysed in this study, simple linear interpolation between temperature grid points was used. While linear interpolation requires a denser temperature sampling and larger file sizes, the LUTs are portable and rapidly decompressed to specified grids.

4.2. Pressure

Pressure primarily affects absorption cross sections by changing the spectral line shape functions due to collisional broadening. Line strengths are pressure dependent through collisional excitation, but this effect is minor in comparison and generally neglected in spectroscopic databases. When deciding upon the pressure axis for a particular gas, spacings




Fig. 3. Left: the mean temperature profile from the 80 RTTOV atmosphere ensemble is shown as the solid blue line with lighter blue area representing the standard deviation of temperature profiles about the mean. The thin black lines depict the absolute minimum and maximum temperature values from the ensemble at those pressure levels. The blue dots help visualize the pressure and temperature grid points used in the H_2O LUT. Right: this plot shows the standard deviations of the CG pressure levels specific to H_2O , CO_2 , and O_3 for the 80 atmosphere ensemble. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

Table 1

Tabulated properties of the final and compressed gas specific LUTs. The first row shows the sample grid spacings of temperature where absorption cross sections are precalculated. The second row shows the absorption threshold from Eq. (11) used to thin the LUTs over the fine spectral grid. The third row displays the compressed size of the LUTs on hard disk and the fourth row shows the percent of the compressed size compared to the uncompressed LUTs.

LUT parameter	H ₂ O	CO ₂	O ₃	N ₂ O	CH ₄	HNO ₃	CO	NO ₂	NO	OCS	SO ₂	Total
ΔT (K)	8.3	4.5	5.5	11	11	11	11	11	11	11	11	
h	10^{-4}	10^{-5}	10^{-5}	10^{-5}	10^{-5}	10^{-7}	10^{-7}	10^{-7}	10^{-7}	10^{-7}	10^{-7}	
Size (Mbytes)	2298	1738	881	135	101	359	127	147	15	57	18	5623
Compression (%)	7.8	9.1	5.7	1.7	1.3	4.6	1.6	1.9	.19	.73	.23	4.4




Fig. 4. Contour plots of absorption cross section at line wing (top) and line centre (bottom) for water vapour using a mid-latitude type VMR. Pressure is expressed linearly in the top plot and logarithmically in the bottom.

can be incremented linearly with pressure, by the logarithm of pressure, or set with irregular spacings according to the resulting interpolation error. Fig. 4 shows that the dependence of absorption cross sections upon pressure is not monotonic and is unlikely to be parametrised well by a low order polynomial or other simplistic functions. It may be that pressure is best interpolated in a piecewise manner. Since the H₂O LUT requires an extra e dimension, it is significantly larger in size than the other gas LUTs. Therefore, a tailored pressure grid was created for the H₂O LUT to minimise

interpolation error while reducing the number of grid points to keep the file size relatively low. During the modelling of radiative transfer, absorption cross sections are linearly interpolated in the RFM along the pressure axis according to the logarithm of the given LUT pressure.

Reduced axis points for pressure were selected by starting with the 101 CG pressure levels and removing levels individually while checking the change in radiance over a small spectral range around 1500 cm^{-1} . This pressure level removal process was repeated until the changes in brightness temperature approximately met the 0.02 K accuracy criteria. Finally, the 32 pressure levels shown in Fig. 3 were selected specific to H₂O.

The right plot in Fig. 3 shows the standard deviation of CG pressure from the 80 atmosphere ensemble. These standard deviations may be plotted as y-axis errorbars for the left figure, but are too small to be visualized if not displayed separately. Variability of CG pressure from one atmosphere to another is quite small when the vertical profile is approximated with many layers. In fact, the maximum standard deviation for H₂O is $\approx 0.3\%$ at 100 hPa. Therefore, any atmosphere from the ensemble may safely be used to generate absorption cross sections without considering variations in CG pressure specific to the absorbing gas.

While an irregular pressure grid was used to create the H₂O LUT, the 101 level pressure grid native to the RTTOV ensemble modified by the CG approximation was used for the 10 remaining trace gases. Without a partial pressure axis increasing the size of the LUT, more pressure levels can be included while maintaining a manageable size on hard disk. A similar tailoring process was attempted for CO₂ and O₃, however it was found that far fewer pressure levels could be removed throughout the profile before interpolation error became noticeable.

4.3. Partial pressure for water vapour

Partial pressure axis values (units of hPa) for H₂O were selected in a manner similar to its pressure grid; where an overly dense e grid was reduced one point at a time, linearly interpolated between the remaining points, and verified by monitoring the change in radiance. The mean volume mixing ratios (VMRs) from the ensemble with extreme midpoint temperatures were used as the reference atmosphere when creating absorption cross sections. The axis points are scaling factors relative to the ensemble mean VMR and were selected to be 10, 35, 60, 80, 120, 200, 345, 600, and 1000%. Continuum effects are the extension and cumulation of self-broadening far from line centre and

are not included in the water vapour LUT. These effects could be included but, with its relatively low computation cost plus rapid variation as a function of p , T , and e (requiring a high tabulation density), it is more efficient to exclude continuum from the LUTs and treat it separately. This is the default behaviour of the RFM.

Partial pressure axes were not included for the remaining gases as spectral line shapes are only weakly dependent upon e for molecules without permanent dipole fields. Since the optical depth for a homogeneous layer is directly proportional to the density of the absorbing gas, these LUTs should provide accurate results regardless of e assuming the gas densities are properly adjusted when evaluating Eq. (6). To verify this posit, LUTs for CO₂ and O₃ at different partial pressures were created and found to yield negligibly different radiances as computed with the RFM.

4.4. Spectral compression

Accurately reproducing the narrowest of spectral features detectable from nadir sounders, i.e., Doppler broadened lines in the upper stratosphere, require modelling the spectral grid at approximately 1000 points per wavenumber. As atmospheric sounding instruments like IASI span thousands of wavenumbers, LUT size can push upwards of tens to hundreds of gigabytes on hard disk space. This can be unwieldy to share over network connections and the input/output operations alone may offset the potential gain in speed. Considering that each LUT originally contained over 2 million spectral points, the majority of LUT compression is achieved along the spectral axis. The fundamental question should compression be determined by considering absolute or relative errors? A few examples of previous spectral compression techniques may help to guide the decision.

Substantial reductions can be achieved simply by taking a threshold of optical depth or absorption [6,17]. For each gas and layer, every spectral point that is optically thin beyond some threshold is removed from storage and reconstructed with either a small predetermined number or zero. Scott and Chedin [6] used an absorption threshold value of 0.005. Alternatively, Amato et al. [17] used 10^{-4} , which removed 88.5% of all spectral points from storage, compressing the LUTs to 11.5% of the original.

Sparks [26] developed a technique to analytically select a reduced spectral grid by considering the change in line shape from removing spectral points and replacing the region with three-point Lagrangian interpolation. One benefit to this method is that the difference in absorption cross sections due to interpolating from a coarse to fine spectral grid need not be explicitly evaluated during grid selection, making it highly efficient. However, if the reduced spectral grid is fixed, then the cost to systematically raster through each fine grid point while checking interpolation error is only paid once. Furthermore, Kuntz and Höpfner [27] pointed out that this spectral sampling is proportional to the line strengths and number density of the absorbing gas. In other words, the densest spectral sampling occurs near the centre of the strongest line for the most abundant gas, so errors are analysed in the

absolute sense. One potential drawback to quantifying interpolation error in absolute terms is that a spectral grid constructed for a highly varying trace gas may produce unexpectedly large errors if that gas is significantly enhanced compared to the referenced amount.

Showing the limits of compressibility, Strow et al. [21] developed a truncated SVD method to compress their LUTs to 1.4% of their original size. However, LUTs compressed in this manner need to be expanded to their original dimensionality with the left and right singular vectors in order to extract the required absorption information. Thus, the truncated SVD method itself may not increase speed and likely add further processing time for the benefit of smaller LUTs. Interestingly, this problem was overcome by interpolating within the compressed space and regaining the speed increase. Additionally, Strow et al. [21] found that the dynamic ranges for the optical depths of CO₂ and H₂O were so great that large optical depth values in the lower atmosphere dominated the reconstruction. Therefore, they reduced the range to compress by taking the 4th root of the optical depths and then applied SVD truncation, which required far fewer singular vectors to represent. The 4th root was decided upon by trial and error, however this SVD method seems to suggest that a relative compression metric performs better by treating all atmospheric layers with equal importance.

The key point to the method employed in this work is to use each pressure level in the LUT to define an atmospheric layer and then set the threshold criteria as the error in absorption for each layer independently. The spectral range computed extends continuously from 644 to 2761 cm⁻¹ to cover the entire IASI range plus an extra wavenumber on either end to allow for apodization [28]. A study conducted by Ventress [16] showed that a convolution kernel width of 1 cm⁻¹ was adequate to represent a 0.5 cm⁻¹ full width at half maximum Gaussian apodization function. Spectral increments were set to spacings of $\Delta\nu = 0.001$ cm⁻¹ to capture all fine structure necessary for nadir-viewing scenarios. The total size of the 11 combined LUTs at this spacing was 124 Gbytes, which is unwieldy for storage and sharing.

This method to remove spectral grid points is numerical in nature. For an individual gas, at the l th CG level the difference in absorption by removing a spectral point, k_i , and replacing that with a linear interpolation from neighbouring points,

$$k'_i = \left(\frac{k_{i+1} - k_{i-1}}{\nu_{i+1} - \nu_{i-1}} \right) \nu_i + \frac{\nu_{i+1} k_{i-1} - \nu_{i-1} k_{i+1}}{\nu_{i+1} - \nu_{i-1}}, \quad (10)$$

is evaluated and compared to a threshold value:

$$\left| \exp(-k'_{igl} u_{gl}) - \exp(-k_{igl} u_{gl}) \right| < h_g, \quad (11)$$

where h is the set threshold. If Eq. (11) is true for *all* CG levels, then the spectral point is removed from the LUT and the thinning process proceeds to the next point with the updated grid. Absorption threshold values used for each gas are displayed in the second row of Table 1.

An example of the thinned spectral grid (blue stars) is shown in Fig. 5 for the same H₂O line from Figs. 1 and 2. Using a threshold value of $h = 10^{-4}$ reduces the number of




Fig. 5. This plot shows the absorption spectrum for an isolated H₂O line in a tropical scenario with self-broadening and continuum effects included. The irregular grid produced by the spectral thinning method described in Section 4.4 ($h = 10^{-4}$) is annotated by the blue stars. The IASI ILS is overlain in green at each quarter wavenumber for reference. The irregular grid considering convolution with an ILS is shown as the red circles. Absorption residual errors (right axis) from linearly interpolating the irregular grids to the fine grid are overlain as the solid blue and red lines. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

spectral grid points to just 3.7% that of the original, while the residual interpolation error in total absorption is less than 0.2%. However, there is an opportunity to remove even more spectral points if one considers that instruments such as IASI have much broader instrument line shape (ILS) functions than the irregular grid spacings shown. The apodized ILS for IASI is a convolution of a sinc and Gaussian function with a full width at half maximum (FWHM) of 0.5 cm^{-1} [28], as highlighted in Fig. 5 for reference. Therefore, positive and negative error contributions separated by less than $\approx 0.5 \text{ cm}^{-1}$ may be offsetting in the final apodized spectrum. To account for this, the thinning method was rerun with Eq. (11) evaluated after convolving the absorption cross sections over a local spectral cell with a simple triangular apodization function having a full base width of 0.5 cm^{-1} . This absorption residual is also shown in Fig. 5 where irregular grid points are further reduced to 1.6% of the original. Notice that the increase in positive and negative interpolation error on the fine grid will roughly cancel out once convolved with the IASI ILS and sampled on a grid of 0.25 cm^{-1} .

Compressed LUT file sizes are shown in the third row of Table 1 with the percent relative to their uncompressed size shown in the fourth row. The total size on hard disk for all 11 trace gas LUTs is 5.62 Gbytes, which is just 4.4% the size of the uncompressed LUTs. This is a significant compression of the originals, and can now be readily shared with the scientific community over the internet as opposed to transferring physical hard drives. It is important to note that interpolation error for this method is evaluated relative to each layer and treats all atmospheric levels with equal importance. This is surely a conservative approach and there are alternative methods available to further compress the spectral axis. However, if the lower limit using SVD truncation is approximately 1.4%, then the conservative approach may be worth the extra size if the

interpolation errors are small over a wide variety of atmospheric conditions.

Finally, care was taken to ensure that the minor trace gas LUTs were not overly compressed. Some of the weaker spectral lines for the less abundant gases may produce optical depth shapes that are essentially flat and entirely removed from the LUT. This will not be an issue for naturally occurring levels of that gas as represented in the GEMS database, but may lead to significant modelling errors for extreme and anomalous events such as forest fires and volcanic eruptions. Therefore, a far stricter threshold ($h = 10^{-7}$) was used for the last six minor trace gases listed in Table 1 to force retention of additional spectral grid points in the event of anomalously high mixing ratios. This was verified for SO₂ levels from a volcanic scenario following the validation method described in Section 6.

5. Reducing the radiative transfer spectral grid

Spectral compression reduces each LUT onto an irregularly thinned spectral grid. In order to calculate Eq. (1), the LUT spectral axis is interpolated to the regularly spaced fine grid, in this case $\Delta\nu = 0.001 \text{ cm}^{-1}$. However, there is another opportunity to speed up the model by reducing the fine spectral grid used to calculate radiative transfer. The simplest option is to take the spectral points from the individual compressed LUTs and combine them together to form the union. Unfortunately, the resulting interpolation error from the union of LUT grid points for IASI purposes is unacceptably high. This is likely due to the effect apodization has on locally mixing spectral features from separate gases. Therefore, another methodology must be implemented.

Rather than thinning the fine spectral grid in absorption cross section or absorption space, as in Section 4.4, the common grid of radiative transfer is thinned by monitoring interpolation errors according to the blackbody brightness temperature (BBT) spectrum. The RTTOV atmosphere ensemble was used to generate 80 BBT spectra using all the listed gases in the atmospheric profiles. Then the error from removing a spectral point and replacing that with a linear interpolation from neighbouring points is calculated similar to Eq. (10), but using BBT instead of k . Threshold criterion for this method is set to

$$|\text{BBT}_i - \text{BBT}'_i| < 0.01 \text{ K.} \quad (12)$$

If Eq. (12) holds true for all 80 spectra, then the i th spectral point is removed. The resulting irregular grid for radiative transfer was reduced to 12.5% of its regularly spaced original size. Note that spectral radiance could be used in Eq. (12) instead. However, BBT is chosen because it is a common representation of the observation and it conveniently allows accuracy criteria to be expressed as a single value applicable across the whole IASI range, unlike radiance which would require a spectral dependence.

Evaluating Eq. (1) on the fine spectral grid is the top level calculation. Therefore, reductions in this grid translate proportionately to reductions in computation time. If the radiative transfer grid is reduced by a factor of ten,

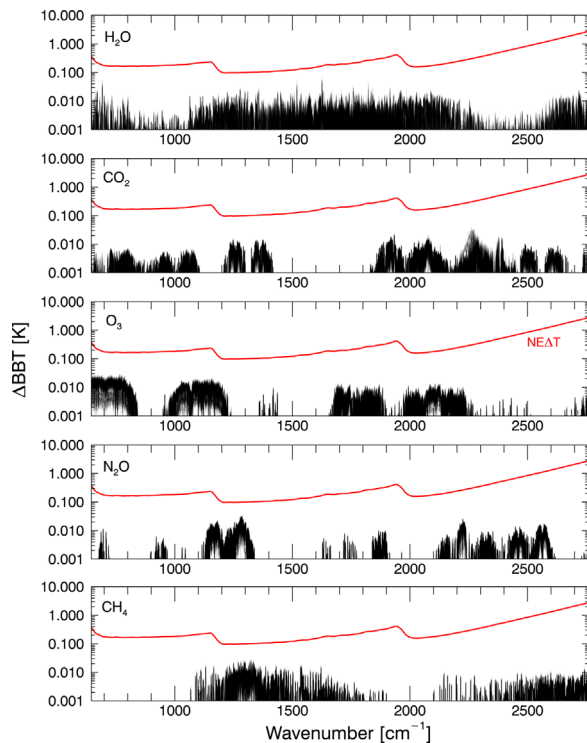


Fig. 6. The magnitude of differences in brightness temperature from using LUTs to simulate IASI spectra versus full LBL mode with the RFM (black lines). The plots are specific to the individual brightness temperatures of H_2O , CO_2 , O_3 , N_2O , and CH_4 . Each plot shows 80 ΔBBT lines over-plotted, one for each atmosphere from the RTTOV ensemble. The IASI NEAT noise level at 280 K is shown in red for reference. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)




Fig. 7. The magnitude of differences in brightness temperature from using LUTs to simulate IASI spectra versus full LBL mode with the RFM (black lines). All 11 gases were included in the radiative transfer calculations. In this figure, 80 ΔBBT lines are over-plotted, one for each atmosphere from the RTTOV ensemble. The IASI NEAT noise level at 280 K is shown in red for reference. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

then one could expect computation speed to increase by approximately that same factor. However, this reduction is not free. The implicit cost of a reduced grid is further dependence upon the atmospheric ensemble and the explicit cost is a slight increase in interpolation error.




Fig. 8. Same as in Fig. 7, but also utilising the reduced fine spectral grid from Section 5.

6. Look-up table performance

The desired outcome of using LUTs of absorption cross sections is to produce simulated radiances significantly faster than LBL methods with a negligible increase in error. Indeed, the computational time required to simulate an entire IASI spectrum using the 11 trace gases listed in Table 1 decreased by an average factor of 22 compared to LBL results. Including the reduced spectral grid from Section 5 further decreased computation time by an additional factor of six. Total computation time is 130 times faster than LBL mode on average. The RFM now computes full-range IASI radiances (8461 apodized channels) for the 11 combined gases in approximately 10 s on a standard desktop computer. Undoubtedly, there are numerous faster RTMs available to simulate IASI. However, the objective is to maintain uncompromising accuracy while becoming more operationally plausible.

Fig. 6 shows the individual gas differences in BBT (ΔBBT) of the LUT results compared to the LBL results for the first five trace gases. Also, the noise equivalent differential temperature (NEAT) of IASI is displayed for reference [5]. There are actually 80 separate lines of ΔBBT shown in these plots, one for each atmosphere from the RTTOV ensemble. Notice that all atmospheres studied produced ΔBBT values well below the IASI noise level. The remaining six trace gases are not shown because all of their NEAT values were below 0.002 K. Upon close inspection of the H_2O plot, there appear to be a few spikes where ΔBBT approaches the IASI noise level. However, after analysis these spikes are purely artificial and result from selecting different methods of interpolating (polynomial or inverse polynomial) continuum effects when using the LUTs versus LBL mode, where one can argue equally well for either interpolation method.

Finally, all 11 gases were combined and calculated with the RFM to simulate IASI radiances over the 80 atmospheres. The resulting ΔBBT values from using LUTs versus full LBL are shown in Fig. 7. Over 99% of the spectral channels have ΔBBT errors due to LUT interpolation less than 0.02 K for all atmospheric scenarios. As evident, many of these values are an order of magnitude or less than IASI noise. Additionally, if the LUTs were not appropriately gridded for any of the 80 atmospheres, then this would be apparent in Fig. 7. Similarly, Fig. 8

shows the resulting errors when including both LUTs and the reduced fine spectral grid from Section 5. As a result, over 99% of all spectral channels have Δ BBT errors less than 0.04 K. Since the interpolation error from using these LUTs is several factors less than the IASI NE Δ T, it is unlikely that the decrease in accuracy will be noticeable for any part of the spectrum over the vast majority of possible Earth atmospheres.

7. Conclusions

Look-up tables of absorption cross sections were created for the 11 most abundant trace gases along with an irregularly reduced spectral grid that allow the RFM to simulate IASI radiances approximately 130 times faster on average than its traditional LBL mode. These LUTs were spectrally compressed according to a simple linear interpolation scheme that accounts for offsetting positive and negative errors during apodization with the IASI ILS function. Thus, reducing total LUT file size from 120 Gbytes to 5.6 Gbytes, a compression to 4.4% of the original. Furthermore, the resulting LUT error, as measured in brightness temperature, was rigorously validated for the 80 different atmospheres found in the RTTOV ensemble. More than 99% of all IASI simulated spectral channels had a LUT interpolation error less than 0.02 K and a combined error including the reduced spectral grid less than 0.04 K, several factors below the IASI noise level.

As mentioned, several RTMs already incorporate LUTs or other “fast” methods in their numerical schemes. Perhaps the novelty in this work is the attention paid to error in the resulting radiances over a wide variety of atmospheric conditions with the flexibility afforded by creating a LUT separately for each individual trace gas. By carefully selecting the range and spacings of the LUT grids, it should be quite difficult to devise an atmospheric Earth scenario that results in large radiance errors due to LUT interpolation.

These LUTs of absorption cross sections were created with the intent of using the RFM as a forward model for IASI in a more operational context. However, these LUTs are not specific to IASI and may be used with the RFM in any nadir-viewing scenario from 645 cm^{-1} to 2760 cm^{-1} where the instrument noise and spectral spacings are roughly greater than or equal to that of IASI. In fact, these LUTs could potentially be incorporated into any existing or future RTM as the absorption cross section points were designed for simple linear interpolation within the LUT grids.

Finally, aside from presenting a generalised method of creating, compressing, and validating LUTs, these are also freely available to the scientific community. The 11 gases listed in Table 1 are just a starting point. Other trace gas species including ammonia, CFCs, and several organic compounds are also detectable by IASI. Instructions on how to use the RFM to generate LUTs of absorption cross sections are posted on the RFM website along with the source code for spectral compression and several other helper functions for manipulating the LUTs. Each gas specific LUT takes approximately two days to generate per computation node from creation to compression if covering the entire IASI range. This process is much faster for smaller spectral intervals. Further information on how to

download the latest version of the RFM and the presented LUTs can be found by browsing to www.atm.ox.ac.uk/RFM.

Acknowledgements

Portions of this work were funded by the United Kingdom National Centre for Earth Observation and the United States Air Force. The views expressed in this article are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense or the U.S. Government.

References

- [1] Goody R, West R, Chen L, Crisp D. The correlated-k method for radiation calculations in nonhomogeneous atmospheres. *J Quant Spectrosc Radiat Transf* 1989;42(6):539–50.
- [2] Lacis AA, Oinas V. A description of the correlated k distribution method for modeling nongray gaseous absorption, thermal emission, and multiple scattering in vertically inhomogeneous atmospheres. *J Geophys Res* 1991;96(15):9027–64.
- [3] Moncet J-L, Uymin G, Lipton AE, Snell HE. Infrared radiance modeling by optimal spectral sampling. *J Atmos Sci* 2008;65(12):3917–34. <http://dx.doi.org/10.1175/2008JAS2711.1> ISSN 0022-4928.
- [4] Fischer H, Birk M, Blom C, Carli B, Carlotti M, von Clarmann T, et al. MIPAS: an instrument for atmospheric and climate research. *Atmos Chem Phys* 2008;8(8):2151–88 URL (<http://www.atmos-chem-phys.net/8/2151/2008/>).
- [5] Hilton F, Armante R, August T, Barnet C, Bouchard A, Camy-Peyret C, et al. Hyperspectral earth observation from IASI: five years of accomplishments. *Bull Am Meteorol Soc* 2011;93(3):347–70. <http://dx.doi.org/10.1175/BAMS-D-11-00027.1> ISSN 0003-0007.
- [6] Scott NA, Chedin A. A fast line-by-line method for atmospheric absorption computations: the automatized atmospheric absorption atlas. *J Appl Meteorol* 1981;20(7):802–12 ISSN 0021-8952.
- [7] Dudhia A. Reference forward model (RFM), URL (<http://www.atm.ox.ac.uk/RFM>); 2008 [accessed June 2015].
- [8] Curtis AR. Discussion of “A statistical model for water-vapour absorption” by R.M. Goody. *Q J R Meteorol Soc* 1952;78(338):638–40.
- [9] Godson WL. The evaluation of infrared radiative fluxes due to atmospheric water vapour. *Q J R Meteorol Soc* 1953;79(341):367–79 ISSN 1477-870X.
- [10] Goody RM, Yung YL. *Atmospheric radiation: theoretical basis*. 2nd ed. New York: Oxford University Press; 1989.
- [11] Tobin DC, Strow LL, Lafferty WJ, Olson WB. Experimental investigation of the self- and N₂-broadened continuum within the ν_2 band of water vapor. *Appl Opt* 1996;35(24):4724–34.
- [12] Tipping R, Ma Q. Theory of the water vapor continuum and validations. *Atmos Res* 1995;36(1–2):69–94 ISSN 0169-8095.
- [13] Koshchev M, Tretyakov M, Golubiatnikov G, Parshin V, Markov V, Koval I. Broadening and shifting of the 321-, 325- and 380-GHz lines of water vapor by pressure of atmospheric gases. *J Mol Spectrosc* 2007;241(1):101–8. <http://dx.doi.org/10.1016/j.jms.2006.11.005> ISSN 0022-2852, URL (<http://www.sciencedirect.com/science/article/pii/S0022285206003201>).
- [14] Rothman L, Gordon I, Babikov Y, Barbe A, Benner DC, Bernath P, et al. The HITRAN2012 molecular spectroscopic database. *J Quant Spectrosc Radiat Transf* 2013;130:4–50.
- [15] Galatry L. Simultaneous effect of Doppler and foreign gas broadening on spectral lines. *Phys Rev* 1961;122:1218–23 URL (<http://link.aps.org/doi/10.1103/PhysRev.122.1218>).
- [16] Ventress L. Atmospheric sounding using IASI [Ph.D. thesis]. Oxford University, URL (<http://ora.ox.ac.uk/objects/uuid:7ad570a3-35ad-4d98-93bb-7e1549afcdcd>); 2014.
- [17] Amato U, Masiello G, Serio C, Viggiano M. The σ -IASI code for the calculation of infrared atmospheric radiance and its derivatives. *Environ Model Softw* 2002;17(7):651–67 ISSN 1364-8152.
- [18] Clough S, Shephard M, Mlawer E, Delamere J, Iacono M, Cady-Pereira K, et al. Atmospheric radiative transfer modeling: a summary of the AER codes. *J Quant Spectrosc Radiat Transf* 2005;91(2):233–44 ISSN 0022-4073.
- [19] Hurtmans D, Coheur P-F, Wespes C, Clarisse L, Scharf O, Clerbaux C, et al. FORL radiative transfer and retrieval code for IASI. *J Quant*

- Spectrosc Radiat Transf 2012;113(11):1391–408 ISSN 0022-4073, three Leaders in Spectroscopy.
- [20] Rodgers CD. Inverse methods for atmospheric sounding: theory and practice, 2nd ed., vol. 2. Singapore: World Scientific; 2000.
- [21] Strow L, Motteler HE, Benson RG, Hannon SE, Souza-Machado SD. Fast computation of monochromatic infrared atmospheric transmittances using compressed look-up tables. J Quant Spectrosc Radiat Transf 1998;59(3–5):481–93. ISSN 0022-4073, Atmospheric Spectroscopy Applications 96.
- [22] Dudhia A, Morris PE, Wells RJ. Fast monochromatic radiative transfer calculations for limb sounding. J Quant Spectrosc Radiat Transf 2002;74(6):745–56. [http://dx.doi.org/10.1016/S0022-4073\(01\)00285-0](http://dx.doi.org/10.1016/S0022-4073(01)00285-0) ISSN 0022-4073, URL <http://www.sciencedirect.com/science/article/pii/S0022407301002850>.
- [23] Edwards DP. GENLN2: a general line-by-line atmospheric transmittance and radiance model, version 3.0 description and users guide. Technical Report NCAR/TN-367+STR, National Center for Atmospheric Research, Boulder, CO, USA; 1992.
- [24] Matricardi M. The generation of RTTOV regression coefficients for IASI and AIRS using a new profile training set and a new line-by-line database. Technical Report 564, ECMWF Research Dept, Shinfield Park, Reading, Berkshire RG2 9AX, England, URL <http://www.ecmwf.int/publications>; 2008.
- [25] Hollingsworth A, Engelen RJ, Benedetti A, Dethof A, Flemming J, Kaiser JW, et al. Toward a monitoring and forecasting system for atmospheric composition: the GEMS project. Bull Am Meteorol Soc 2008;89(8):1147–64 ISSN 0003-0007, URL <http://dx.doi.org/10.1175/2008BAMS2355.1>.
- [26] Sparks L. Efficient line-by-line calculation of absorption coefficients to high numerical accuracy. J Quant Spectrosc Radiat Transf 1997;57(5):631–50 ISSN 0022-4073.
- [27] Kuntz M, Höpfner M. Efficient line-by-line calculation of absorption coefficients. J Quant Spectrosc Radiat Transf 1999;63(1):97–114 ISSN 0022-4073, <http://www.sciencedirect.com/science/article/pii/S002240739800140X>.
- [28] Amato U, De Canditiis D, Serio C. Effect of apodization on the retrieval of geophysical parameters from Fourier-transform spectrometers. Appl Opt 1998;37(27):6537–43.