The NEMESIS planetary atmosphere radiative transfer and retrieval tool

P.G.J. Irwin[a], N.A. Teanby[a], R. de Kok[a], L. N. Fletcher[a], C.J.A. Howett[a], C.C.C. Tsang[a], C.F. Wilson[a], S.B. Calcutt[a], C.A. Nixon[b] and P.D. Parrish[c]

[a]Atmospheric, Oceanic, and Planetary Physics, University of Oxford, Clarendon Laboratory, Parks Rd, Oxford, OX1 3PU, United Kingdom

Tel: (+44) 1865 272933, *Fax:* (+44) 1865 272923.

E-mail: patrick.irwin@physics.ox.ac.uk

[b]University of Maryland /NASA GSFC, Solar System Exploration Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

[c]Institute of Atmospheric & Environmental Science, School of Geosciences, University of Edinburgh, Crew Building, King's Buildings, Edinburgh EH9 3JN

original manuscript submitted 21/9/07

revised manuscript submitted 21/11/07

Number of manuscript pages: 28 + acknowledgements, references and figures (40 in total)

Number of Tables: 1

Number of Figures: 4

Proposed running head:

The NEMESIS planetary atmosphere radiative transfer and retrieval tool

Editorial correspondence should be directed to:

Prof. Patrick G. J. Irwin,

Atmospheric, Oceanic and Planetary Physics, Clarendon Laboratory, Parks Rd,

Oxford OX1 3PU, United Kingdom.

Telephone: (+44) 1865 272933 (direct line: 272083)

Fax: (+44) 1865 272923

Email: patrick.irwin@physics.ox.ac.uk

Abstract

With the exception of *in situ* atmospheric probes, the most useful way to study the atmospheres of other planets is to observe their electromagnetic spectra through remote observations, either from ground-based telescopes or from spacecraft. Atmospheric properties most consistent with these observed spectra are then derived with retrieval models. All retrieval models attempt to extract the maximum amount of atmospheric information from finite sets of data, but while the problem to be solved is fundamentally the same for any planetary atmosphere, until now all such models have been assembled *ad hoc* to address data from individual missions.

In this paper we describe a new general purpose retrieval model, NEMESIS, which was originally developed to interpret observations of Saturn and Titan from the Composite Infrared Spectrometer on board the NASA Cassini spacecraft. NEMESIS has been constructed to be generally applicable to any planetary atmosphere and can be applied from the visible/near-infrared right out to microwave wavelengths, modelling both reflected sunlight and thermal emission in either scattering and non-scattering conditions. NEMESIS has now been successfully applied to the analysis of data from many planetary missions and also ground-based observations.

Keywords: Retrievals, Radiative Transfer, Correlated-*k*

1. Introduction

Apart from occasional entry probes, the best way to study the atmospheres of the other planets, and thus compare them with our own, is to observe their absorption, reflection and emission spectra remotely from spacecraft or from ground- or space-based telescopes. Particularly useful spectral regions are the visible and near-infrared, where sunlight reflected off clouds and hazes can be observed, and the thermal-infrared where thermal emission spectra can be used to constrain temperatures, gas and aerosol abundances. These processes are also important to correct for when studying the surfaces of planets such as Mars, Venus and Titan.

Interpretation of the measured near-infrared and thermal infrared spectra can be done in a number of ways. The simplest approach is to compare the observations to a representative range of synthetic spectra calculated from a plausible set of atmospheres and choose the one that fits best. This is fine if only a few atmospheric parameters are expected to vary and only a few spectra are to be analysed, but is not suitable for cases where the spectrum is dependent on multiple related factors, such as temperature and composition, and where there are thousands of spectra. In such cases an automated *retrieval* algorithm is desirable, which will process large quantities of data and return the fitted atmospheric states. Such retrieval models require two main components: 1) a radiative transfer, or *forward* model, which, given an assumed atmospheric structure, calculates a synthetic spectrum for comparison with that measured and 2) an inversion, or retrieval model, which compares the measured and modelled spectra and then adjusts the atmospheric parameters in such a way as to minimize any discrepancy.

A number of radiative transfer, or forward models, are already generally available, such as MODTRAN[1], FASCODE[2] and DISORT[3] to mention just a few. However, most of these models have been developed for conducting radiative transfer calculations in the Earth's atmosphere, containing hard-wired assumptions and parameters relevant for terrestrial conditions and so often require substantial modification before they can be applied generally to any other atmosphere. In addition, most models are set up for accuracy rather than speed and so are unsuitable for use in a retrieval model, which may require several iterations. Those radiative transfer models that are sufficiently fast, such as MODTRAN, usually use the band modelling technique and thus are not applicable to radiative transfer in scattering atmospheres.

A number of *retrieval* models have been developed over the years to analyse spacecraft and telescope data, all based on similar principles, but in general these have been generated with a particular planet and instrument in mind and thus have not been generally applicable, or indeed generally available, due to the wide range of atmospheres in the solar system and the different approximations that may be used to model their spectra.

All retrieval models face the same basic challenges: 1) the forward model must be very fast in order that test spectra can be calculated quickly enough to allow the retrieval to be done in a reasonably short space of time; 2) the forward model must, in addition, return the functional derivatives, which are the rates of change of the calculated spectrum with respect to each atmospheric parameter; 3) the retrieval model must then be able to modify the guessed atmospheric state in such a way as to improve the fit between the measured and calculated spectra, but must avoid noise on the spectra leading to physically unrealistic retrieved properties.

In 2001, we began setting up a new retrieval model for use in analysing the data expected from the NASA Cassini Composite Infrared Spectrometer (CIRS). From the start we decided to generate a general-purpose planetary atmospheric retrieval model that could be applied to any planet, any observing geometry and which could model transmission, thermal emission and multiple scattering. This paper describes the formulation and properties of that model: NEMESIS – **N**on-Linear Optimal **E**stimator for **M**ultivariat**E** Spectral Analy**SIS**, which is as far as we know the first generally available combined radiative transfer and retrieval tool that can be applied to any planet.

2. Forward Model

A retrieval model typically has to be iterated a number of times before the calculated spectra match those measured to within error. For each iteration, the forward model must take the latest estimate of the atmospheric profile concerned and compute not only a synthetic spectrum, but also the rate of change of that spectrum with respect to all the atmospheric parameters in order that the correction that needs to be applied to the trial atmospheric profile may be computed. This procedure is the most computationally expensive part of the retrieval process and thus a very fast forward model is required.

The most accurate way of calculating a synthetic atmospheric spectrum is to use a line-by-line model, which computes the absorption of each individual spectral line in the region of interest and then convolves the resulting spectrum with the instrument function of the instrument concerned. Lines are read from standard line databases such as HITRAN [4], or GEISA [5] and such a computation requires the addition of contributions from literally thousands of lines, whose individual absorptions must be calculated for the

different pressures and temperatures in the atmosphere. As a result, line-by-line models are extremely slow and not suitable for retrieval algorithms unless there are only a few lines in the spectral region concerned and/or if the spectral range is very small. Most infrared instruments have resolving powers of less than 5000 and so for our general-purpose retrieval model we required a forward model that was able to compute individual medium-resolution spectra much more rapidly than line-by-line models, but not at the expense of a substantial loss in accuracy.

One option was to use band models, but these have the disadvantage of not being suitable for radiative transfer calculations in multiple-scattering atmospheres. Instead, we chose to use the method of correlated-*k* [6]. The mean transmission $\bar{\tau}$ of an atmospheric path over a particular frequency interval ν_0 to $\nu_0 + \Delta\nu$ may be calculated from the equation

$$\bar{\tau}(m) = \frac{1}{\Delta\nu} \int_{\nu_0}^{\nu_0+\Delta\nu} \exp\left(-m \sum_j k_j(\nu) \right) d\nu \qquad (1)$$

where m is the total absorber amount (molecule cm^{-2}) and the sum is over all the lines of all the gases contributing to the absorption, where the absorption spectrum of the j^{th} line in the sum (taking into account the gas volume mixing ratios) is $k_j(\nu)$ (cm^2 molecule^{-1}). Eq. 1 must be calculated with sufficient spectral resolution, $d\nu$, to resolve the individual line shapes, which may be very much smaller than the interval width, $\Delta\nu$, thus requiring thousands of calculations. However, looking at Eq. 1 we can see that to calculate the mean transmission it is of no importance where in the interval a particular absorption coefficient occurs. Instead, it is sufficient to know what fraction of the frequency domain, $f(k)dk$, is

occupied by absorption coefficients between k and $k+dk$. It is thus possible to write the mean transmission function in an alternative form

$$\bar{\tau}(m) = \int_0^\infty f(k) \exp(-km) dk \,. \tag{2}$$

If we define a cumulative function $g(k)$ as

$$g(k) = \int_0^k f(k) dk \tag{3}$$

and note that since $g(k)$ is a smooth, single-valued, monotonically increasing function, it has a unique, smoothly varying inverse, $k(g)$, Eq.2 may then be further simplified to

$$\bar{\tau}(m) = \int_0^1 \exp(-k(g)m) dg \tag{4}$$

Herein lies the advantage of this approach: while $k(\nu)$ is a rapidly changing function of frequency, requiring $d\nu$ to be very small in a numerical integration, $k(g)$, known as the k-distribution function, is a smoothly varying function of g, as can be seen in Fig.1, requiring far fewer steps in a numerical integration. The integral in Eq. 4 may then be approximated by:

$$\bar{\tau}(m) = \sum_{i=1}^{N} \exp(-k_i m) \Delta g_i \tag{5}$$

where k_i is the value of the k-distribution of the path at the i^{th} quadrature point and Δg_i is the corresponding weight. NEMESIS uses a Gaussian quadrature scheme (e.g. [7]) with N = 10 to 20 ordinates, the number N chosen to achieve the best trade-off between accurate sampling of the k-distribution and computational speed.

Looking at the form of the equations, it becomes apparent that the change of variable and integration space need not be restricted to evaluation of the transmission, but in general may be applied to any analytic function of *k, m* and frequency. For example, suppose we want the mean value of a general function *X*:

$$\bar{X}(m) = \frac{1}{\Delta\nu}\int_{\nu_0}^{\nu_0+\Delta\nu} X\left(k(\nu),m\right)d\nu \tag{6}$$

then a similar transformation applies. i.e.:

$$\bar{X}(m) = \int_{0}^{1} X\left(k(g),m\right)dg \tag{7}$$

Hence, once the absorption coefficient frequency distribution of the gas absorption has been determined, any other spectral properties of a layer, such as emission and scattering, may also be determined in *g*-space leading to massive reductions in computation time.

To use NEMESIS to simulate the observations from a particular instrument observing a particular planet, a set of *k*-distribution functions $k(g)$ must first be calculated for each spectrally active gas likely to contribute to a planetary spectrum over a representative grid of pressures and temperatures for the planetary atmosphere concerned, and at an appropriate resolution to model the instrument properties, to form a set of *k*-tables. These *k*-tables are generated from line data, where possible, but may also be derived from band model parameters for cases where no line data exist (e.g. [8, 9]). For the case of line data, for each pressure and temperature in the grid, a high-resolution spectrum is computed for each spectral interval of width $\Delta\nu$ and then in the simplest scheme the

fraction of the interval with absorption coefficients less than a limiting value k_L is calculated for k_L varying between k_{min} and k_{max} as:

$$g(k_L) = \frac{\sum_{k_i \le k_L} \delta\nu_i}{\sum_i \delta\nu_i} = \frac{\sum_{k_i \le k_L} \delta\nu_i}{\Delta\nu} \tag{8}$$

where the summations are over the whole interval of width $\Delta\nu$ in very small steps $\delta\nu$, and the summation in the top line is for those parts of the total interval where k is less than or equal to k_L (Fig. 1). The resulting cumulative distribution function $g(k)$ is then inverted to give $k(g)$, and then sampled at the Gaussian quadrature ordinates to give k-tables suitable for modelling an instrument with a square instrument function of width, $\Delta\nu$, or a triangular instrument function of the same FWHM if the resulting spectrum is subsequently convolved with another square of width $\Delta\nu$. In some cases, other instrument functions are required, such as for example the Hamming function used in apodised CIRS spectra [10, 11]. The k-tables for such instrument functions can be calculated by:

$$g(k_L) = \frac{\sum_{k_i \le k_L} f_i \delta\nu_i}{\sum_i f_i \delta\nu_i} \tag{9}$$

where an extra parameter f_i has been introduced, which is the instrument function $f(\nu)$ at each point in the spectral interval $f_i = f(\nu_i - \nu_0)$ and ν_0 is the central frequency of the interval, i.e., of the apodisation function.

To calculate the spectrum of a path in a particular spectral interval, NEMESIS interpolates the k-tables for each gas to the path temperature and pressure. If more than one gas contributes to the absorption, the k-tables of the different gases have to be combined. Following [6], the k-tables are combined assuming the absorption lines of different gases

are not correlated with each other, using the gas volume mixing ratios as weights. Finally, NEMESIS then adds continuum absorptions to the path, such as cloud/aerosol opacity and collision-induced absorptions, such as H_2-H_2 and H_2-He (e.g. [12,13]) which depend on temperature and the ortho/para-H_2 ratio.

A real atmosphere is, of course, not homogenous as the temperature and pressure vary greatly with height. To use *k*-distribution tables for atmospheric calculations in an inhomogeneous atmosphere requires the use of the correlated-*k* approximation. Here, an atmosphere is split up into a suitable number of equivalent homogeneous layers and the *k*-distributions calculated for each layer. It is then found (e.g. [6]) that regions of high absorption in one layer are usually spectrally well correlated with regions of high absorption in all the other layers, and similarly for regions of middle and weak absorptions. Hence, once the *k*-distributions for all *M* layers have been computed, the mean properties of the total atmosphere, e.g., the transmission through all *M* layers can be approximated as:

$$\bar{\tau} = \sum_{i=1}^{N} \exp\left(-\sum_{j=1}^{M} k_{ij} m_j \right) \Delta g_i \tag{10}$$

where k_{ij} is now the value of the *k*-distribution of the j^{th} layer at the i^{th} quadrature point.

The accuracy of the correlated-*k* method depends on the number of quadrature points and on the particular atmosphere under investigation. For a given application test calculations with the correlated-k model are compared with those of a full line-by-line model. For the cases we have investigated so far, we have found that with 10-20 quadrature points, the correlated-k approximation is accurate to better than 5%, which is sufficient for a retrieval model considering there also exist similar magnitude uncertainties in the gas absorption data and the measured spectra.

To see how these principles are applied to a real radiative transfer calculation, suppose we wanted to compute the mean thermal emission spectrum to space from a planetary atmosphere. For non-scattering conditions, the mean radiance R (W m^{-2} sr^{-1} Hz^{-1}) emitted to space between frequencies ν_0 and $\nu_0 + \Delta\nu$ is calculated from the Schwarzchild equation (e.g. [14]), as:

$$R(\nu_0) = \frac{1}{\Delta\nu} \int_{\nu_0}^{\nu_0+\Delta\nu} \left\{ B_\nu(T_g)\tau_\nu(0) + \int_0^\infty B_\nu(T(z)) \frac{\mathrm{d}\tau_\nu(z)}{\mathrm{d}z} \mathrm{d}z \right\} \mathrm{d}\nu \tag{11}$$

where $\tau_\nu(z)$ is the transmission to space from an altitude z at frequency ν and $B_\nu(T)$ is the Planck function at that frequency for temperature, T. T_g is the temperature of the ground. Integrating Eq. 11 over frequency we derive:

$$R(\nu_0) = \overline{B}_{\nu_0}(T_g)\overline{\tau}_{\nu_0}(0) + \int_0^\infty \overline{B}_{\nu_0}(T(z)) \frac{\mathrm{d}\overline{\tau}_{\nu_0}(z)}{\mathrm{d}z} \mathrm{d}z \tag{12}$$

where $\overline{\tau}_{\nu_0}(z)$ is the mean transmission to space from an altitude z and $\overline{B}_{\nu_0}(T)$ is the mean Planck function, both averaged over the spectral bin; $\frac{\mathrm{d}\overline{\tau}_{\nu_0}(z)}{\mathrm{d}z}$ is known as the transmission weighting function $w(z)$. Splitting the atmosphere into M homogeneous layers and integrating over the spectral bin with our correlated-k scheme instead, Eq. 11 becomes:

$$R(\nu_0) = \sum_{i=1}^{N} \left\{ \overline{B}_{\nu_0}(T_g)\overline{\tau}_{i1} + \sum_{j=1}^{M} \overline{B}_{\nu_0}(T_j)\left(\overline{\tau}_{i,j+1} - \overline{\tau}_{i,j}\right) \right\} \Delta g_i \tag{13}$$

where $\overline{\tau}_{ij} = \exp\left(-\sum_{k=j}^{M} k_{ik} m_k \right)$ is the transmission to space from the bottom of the j^{th} level for the i^{th} g-space ordinate and $\overline{B}_{\nu_0}(T_j)$ is the mean Planck function of the bin at the temperature of the j^{th} level.

A very similar approach may be taken to any radiative transfer calculation and NEMESIS can calculate transmission, thermal emission in non-scattering cases (as outlined above) and can also calculate spectra in scattering conditions (of thermally emitted radiation, reflected sunlight, or a combination of both), where either a multiple-scattering Matrix Operator [15] or single-scattering model may be applied. For example, when using correlated-*k* with our multiple scattering model, the scattering model is run *N* times, with the gas optical depth of the *j* layers set to $k_{ij}m_j$, where *i* is the *g*-space ordinate. The temperature, aerosol optical depths and single scattering albedos are fixed to the same values for all *N* calculations and the resulting *N* radiances R_i are then summed as $R(\nu_0) = \sum_{i=1}^{N} R_i \Delta g_i$ where the individual multiple-scattered radiances R_i are computed with sufficient zenith angle quadrature points and azimuthal Fourier components to simulate the scattering conditions.

Functional derivatives

Retrieval models require not only the synthetic radiances calculated for a trial atmospheric state, but also the rate of change of radiance with respect to all atmospheric variables, known as the functional derivatives. In many retrieval models this is calculated by varying each atmospheric property to be fitted in turn, recalculating the entire spectrum and then dividing the change in spectrum by the change in atmospheric property. While this approach is straightforward to implement, it is very slow and inefficient, requiring the computation of many spectra for each iteration of the retrieval model. While for multiple-scattering conditions NEMESIS still currently calculates the functional derivatives using this inefficient numerical differencing technique, for non-scattering conditions, we have developed the code to compute the rates of change of radiance with respect to atmospheric

property at the deepest subroutine level analytically (i.e. through their functional dependence) and then track and combine these derivatives (through the chain rule) through to the final spectrum. For example, suppose the optical depth χ of a layer depends on temperature T and pressure p as $\chi = \alpha T^n p^m$, then as well as returning this optical depth, the subroutine concerned would also return $\frac{d\chi}{dT} = n\alpha T^{n-1} p^m$ and $\frac{d\chi}{dp} = m\alpha T^n p^{m-1}$. Similarly, all subroutines calculating properties dependent on temperature, pressure, gas abundance, cloud abundance, etc., have been coded to return not only those properties, but also the rates of change with respect to all dependent variables. In this way, the modelled spectrum and functional derivatives are computed in parallel by NEMESIS leading to an order of magnitude increase in the retrieval code's computational speed, which is main advantage of the model. NEMESIS is currently being developed to incorporate this implicit differentiation scheme into the multiple-scattering code.

To model the observed spectrum of a planet, we must first have an atmospheric model, which contains as many vertical profiles as are necessary to generate an accurate spectrum. For example, at thermal wavelengths, we require a temperature profile together with vertical profiles of the spectrally active gases, clouds and hazes, making up NPARAM profiles in total (Fig. 2). The number of vertical levels in these profiles, NPRO, is somewhat arbitrary, but must be sufficiently large to resolve any expected vertical variation.

For each iteration of the retrieval model (see next section) the NPARAM model atmospheric profiles are initially set to 'reference' profiles, which contain our best-guess estimate of all parameters. The NVAR profiles actually being retrieved (where NVAR ≤ NPARAM) are then updated using a state vector, **x**, of length NX, which contains a

parameterised representation of each of the profiles concerned. NEMESIS can parameterise a profile in many different ways. In its most variable form, the state vector **x** contains the NVAR profiles at all NPRO vertical levels in the atmosphere, giving NX = NVAR×NPRO. However, NEMESIS can also parameterise a profile as a mean value, or as mean value below a specified pressure level and a variable drop-off rate above, etc. If a single continuous vertical profile (NVAR=1) is to be retrieved such as temperature, then NX is equal to NPRO and the model temperature profile is set to the elements of the vector **x**. However, if instead we are just retrieving a mean atmospheric temperature then NX would be equal to 1, and all levels of the model temperature profile would be set to x_1. Depending on how the atmospheric profiles being retrieved are parameterised by the state vector, NEMESIS adjusts the model atmospheric profiles and also calculates a *Profile Property Gradient* matrix, **M**, of dimension (NX, NPARAM, NPRO), which contains the rate of change of the model atmospheric profiles with respect to each element of the state vector, **x**.

NEMESIS then splits up the atmosphere into an appropriate number of equivalent homogenous layers. Depending on the viewing geometry, the number of equivalent homogeneous layers required to model the radiative transfer with sufficient accuracy, NLAYER, is often slightly smaller than NPRO, but must be sufficient to resolve the transmission weighting functions, $w(z) = \frac{\mathrm{d}\,\bar{\tau}_{\nu_0}(z)}{\mathrm{d}z}$ (Eq. 12), over the total spectral range considered. In NEMESIS, layers may be created by a number of methods, e.g., splitting equally with height, or equally with logarithmic pressure and the mass-weighted mean atmospheric properties of each layer are calculated. During this process of calculating layer properties from profile properties, a *Layer Property Gradient* matrix **L** is computed (Table

1), containing the rate of change of the homogeneous layer properties with respect to atmospheric profile properties. For example, if the temperature of the n^{th} homogenous layer is a weighted average of the model temperature profile, $\overline{T}_n = \sum_{i=1}^{NPRO} \gamma_i T_i$, where γ_i are the relevant weights, then the rate of change of temperature of this layer with respect to the i^{th} level of the model temperature profile will be γ_i. The rows of the matrix **L** are made up of these coefficients γ_i and are calculated automatically for all NPARAM atmospheric profiles required to model the entire spectrum, regardless of the number actually being retrieved, NVAR. Since there are NPRO levels in the model atmosphere and NLAYER equivalent homogenous layers in the radiative transfer model, the *Layer Property Gradient* matrix, **L**, has the dimensions (NPARAM, NLAYER, NPRO).

Once the layers have been specified, together with their functional dependence on the model atmospheric profiles, NEMESIS then calculates the optical depth of each layer at the required wavelength or wavenumber, together with the rate of change of this optical depth with respect to each variable property. These derivatives are again calculated analytically at the deepest subroutine level and then passed up to the routines calling them, where they are combined as necessary, using the chain rule. The transmission to space from the bottom of each layer is then calculated together with the rate of change of that transmission with respect to the variable properties of <u>all</u> of the layers above that point. Finally, for the example of a thermal emission calculation, NEMESIS uses these optical depths to calculate the modelled radiance, together with the rate of change of that radiance with respect to each of the NPARAM model layer properties for each of the NLAYER layers, placing these derivatives in a *Layer Radiance Gradient* matrix, **J** (Table 1), of

dimension (NPARAM, NLAYER). For example, to calculate the rate of change of radiance with respect to the temperature of the l^{th} layer, referring to Eq.13, we write:

$$\frac{dR(\nu_0)}{dT_l} = \sum_{i=1}^{N} \left\{ \bar{B}_{\nu_0}(T_g) \frac{d\bar{\tau}_{i1}}{dT_l} + (\bar{\tau}_{i,l+1} - \bar{\tau}_{i,l}) \frac{d\bar{B}_{\nu_0}(T_l)}{dT_l} + \sum_{j=1}^{M} \bar{B}_{\nu_0}(T_j) \left(\frac{d\bar{\tau}_{i,j+1}}{dT_l} - \frac{d\bar{\tau}_{i,j}}{dT_l} \right) \right\} \Delta g_i \tag{14}$$

where

$$\frac{d\bar{\tau}_{ij}}{dT_l} = -\left(m_l \frac{dk_{il}}{dT_l} + k_{ik} \frac{dm_k}{dT_l} \right) \exp\left(-\sum_{k=j}^{M} k_{ik} m_k \right). \tag{15}$$

The *Profile Radiance Gradient* matrix, **P** (Table 1), containing the rate of change of radiance at each frequency, ν, with respect to the original NPARAM×NPRO profile parameters is then calculated as:

$$P_{ij}(\nu) = \sum_{k=1}^{\text{NLAYER}} L_{ikj} J_{ik}(\nu) \tag{16}$$

where **P** has dimension (NPARAM, NPRO).

The final stage is to calculate the *Radiance Derivative* vector **N** (Table 1), of length NX, containing the rate of change of the modelled radiance with respect to the elements of the state vector **x**, by multiplying **P** by the gradient matrix **M**, introduced earlier:

$$N_i(\nu) = \sum_{j=1}^{\text{NPARAM}} \sum_{k=1}^{\text{NPRO}} M_{ijk} P_{jk}(\nu). \tag{17}$$

Thus, our original NPARAM × NLAYER derivatives are reduced to NX derivatives.

If the synthetic spectrum is calculated at NY wavelengths, the gradient vectors **N** then make up the columns of the final *Functional Derivative* or *Jacobian* matrix **K** (Table 1), of dimension (NY, NX), whose use is described in more detail in the next section. The final functional derivatives are thus analytically correct and computed much more

efficiently than a numerical differencing scheme. NEMESIS is currently able to make use of this scheme for all radiative transfer calculations except multiple-scattering calculations, where the difficulty in coding such a scheme for the matrix operator model has delayed implementation. Hence, for multiple-scattering conditions, the trial spectrum has to be calculated from the state vector **x** and then NX additional spectra calculated by adjusting each element of **x**, with the difference between the perturbed and trial spectra used to calculate the columns of the Jacobian matrix **K**.

3. Retrieval Model

All retrieval algorithms must tackle the basic problem of trying to infer as much as possible about continuous atmospheric properties, such as temperature and gas or cloud abundance, from a finite set of noisy radiances. Mathematically, the problem is said to be *ill-posed*; thus, the number of degrees of freedom in the atmospheric representation must first be reduced to less than the number of individual measurements and there are a number of methods for doing this. A second problem is that measurements themselves have some noise, or measurement error and thus there are literally an infinite number of possible continuous atmospheric profiles that could fit the same set of noisy radiances equally well.

For a typical retrieval problem, an atmosphere is observed from space over a spectral range where the atmospheric opacity varies from very low to very high allowing radiation from a range of levels in the atmosphere to be observed. The goal is to extract an atmospheric profile of, say, temperature with as high a degree of vertical resolution as can be achieved. However, it is found that if it is attempted to extract too much vertical information from a spectral measurement, noise on the measured radiances can quickly

build up in the retrieved profile, leading to highly unrealistic vertical oscillations in the retrieved profiles, even though such a profile would generate the same spectrum as that measured. This phenomenon is known as *ill-conditioning*. All retrieval models, then, attempt to extract as much information as possible about an atmosphere from a finite set of measurements, without becoming ill-conditioned and it is found that there is basically a trade-off to be made between vertical resolution and retrieval error [16]. For example, if we assume that the temperature profile does not vary with height, then we can use the measurements to achieve a very precise least-squares estimate of the mean temperature. However, we find that as we try to extract more vertical resolution from the measurements, the errors in the retrieved profile increase until at some point, ill-conditioning is reached.

The NEMESIS retrieval code was initially developed from algorithms used for inversion of measurements of the Earth's atmosphere. The state of the Earth's atmosphere is reasonably well known and there are good statistics for its likely variability. Hence, we can start with a reasonably well known first guess, or *a priori*, profile and covariance matrix and use the satellite measurements to improve our estimate through the technique of optimal estimation [17]. In essence, this scheme attempts to minimise the difference between the measured and modelled spectra, subject to minimum departure from the *a priori* state vector, by minimising the 'cost' function:

$$\phi = \left(\mathbf{y}_m - \mathbf{y}_n\right)^{\mathrm{T}} \mathbf{S}_\varepsilon^{-1} \left(\mathbf{y}_m - \mathbf{y}_n\right) + \left(\mathbf{x}_n - \mathbf{x}_0\right)^{\mathrm{T}} \mathbf{S}_x^{-1} \left(\mathbf{x}_n - \mathbf{x}_0\right) \tag{18}$$

where $\mathbf{y}_m$ is the measured spectrum, $\mathbf{y}_n$, is the spectrum calculated for the trial atmosphere, represented by a model state vector $\mathbf{x}_n$, $\mathbf{S}_\varepsilon$ is the measurement covariance matrix (containing both measurement errors and estimated forward modelling errors), $\mathbf{x}_0$ is the *a priori* state vector and $\mathbf{S}_x$ is the *a priori* covariance matrix. Adding forward modelling

errors to $\mathbf{S}_\varepsilon$ is somewhat non-standard in that the optimal estimation formalism is developed assuming a Gaussian distributed random noise, whereas the forward modelling errors are systematic errors. However, adding the forward modelling errors like this allows us to incorporate their gross effect into the retrieval results.

For the retrieval of well measured properties such as temperature in Earth's atmosphere, there already exists a statistical record of the mean expected temperature profile and covariance and so $\mathbf{x}_0$ and $\mathbf{S}_x$ are determined from climatology. However, for constituents in planetary atmospheres, and indeed for some trace gases in the Earth's own atmosphere, no statistical information on the *a priori* profile and covariance matrix exists. Hence, the diagonal components of the *a priori* covariance matrix are set to the square of estimated *a priori* errors and, for the case of retrieving continuous vertical profiles, off-diagonal elements set to (after [17]):

$$S_{ij} = \left(S_{ii}S_{jj}\right)^{1/2} \exp\left(-\frac{\left|\ln p_i / p_j\right|}{c}\right) \tag{19}$$

where p_i and p_j are the i^{th} and j^{th} pressure levels and c is a 'correlation length', here equivalent to the number of scale heights over which we can assume the profile to be reasonably correlated. A value of $c = 1.5$ is commonly used. Since the covariance matrices are positive-definite, by definition, the inversion is performed by NEMESIS with a Cholesky decomposition routine (e.g. [7]), in double precision. However, in order to make the inversion of this matrix numerically stable, off-diagonal elements smaller than a certain prescribed factor are set to zero.

The difference between the spectrum computed from the trial measurement vector $\mathbf{x}_n$ and that measured is used to calculate a new estimate of the trial vector $\mathbf{x}_{n+1}$ through the equation [17]:

$$\mathbf{x}_{n+1} = \mathbf{x}_0 + \mathbf{S}_x \mathbf{K}_n^{\mathrm{T}} \left(\mathbf{K}_n \mathbf{S}_x \mathbf{K}_n^{\mathrm{T}} + \mathbf{S}_\varepsilon \right)^{-1} \left(\mathbf{y}_m - \mathbf{y}_n - \mathbf{K}_n \left(\mathbf{x}_0 - \mathbf{x}_n \right) \right) \tag{20}$$

where $\mathbf{K}_n$ is the matrix of functional derivatives (for the n^{th} iteration), or Jacobian, i.e. the rate of change of radiance with state vector elements for all the wavelengths in the spectrum, whose calculation was described in the previous section and whose columns contain the radiance derivative vectors, **N** (Eq. 17). It should be noted that the *a priori* vector, $\mathbf{x}_0$, is used at each iteration to ensure that the *a priori* constraints remain applied throughout the retrieval. However, in practice, $\mathbf{K}_n$ can vary greatly between iterations and the simple iteration scheme of Eq. 20 can become unstable. Instead, NEMESIS uses a modified iteration scheme, based on the Marquardt-Levenberg principal (e.g. [7]) where the actual modified state vector used in the next iteration, $\mathbf{x}'_{n+1}$, is calculated from $\mathbf{x}_{n+1}$ and $\mathbf{x}_n$ (Eq. 20) as:

$$\mathbf{x}'_{n+1} = \mathbf{x}_n + \frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{1+\lambda} . \tag{21}$$

The parameter λ is initially set to 1.0. If the spectrum calculated from $\mathbf{x}'_{n+1}$ is found to reduce the cost function ϕ (Eq. 18), then $\mathbf{x}_n$ is set to $\mathbf{x}'_{n+1}$, λ is multiplied by a factor of 0.3, and the next iteration started. If, however, the spectrum calculated from $\mathbf{x}'_{n+1}$ is found to increase the cost function ϕ, then $\mathbf{x}_n$ is left unchanged, the parameter λ increased by a factor of 10, and a new vector $\mathbf{x}'_{n+1}$ calculated. The choice of the multiplication parameters (0.3, 10) is somewhat arbitrary, although it is important to ensure that they are not reciprocals of one another, as this can potentially lead to endless loops. To ensure smooth

convergence, the lower number (0.3) was chosen to be greater than the reciprocal of the larger number so that λ did not decrease too quickly. As the retrieval approaches its final solution, $\lambda \to 0$, and the model tends to the optimal estimate, whose error is estimated as [17]:

$$\hat{\mathbf{S}} = \left(\mathbf{S}_x^{-1} + \mathbf{K}_n^{\mathrm{T}}\mathbf{S}_\varepsilon^{-1}\mathbf{K}_n\right)^{-1} \tag{22}$$

For cases where the $\mathbf{K}_n$ matrix does not change very rapidly with the state vector, the inversion is approximately linear and convergence is achieved in 2-3 steps. However, for volume mixing ratio retrievals, where $\mathbf{K}_n$ can vary greatly between iterations, convergence can be slower, requiring perhaps 10-20 steps. A further refinement in calculating these matrix operations with maximum speed is discussed in the appendix. In addition, although for retrieving temperatures, the elements of $\mathbf{x}_n$ are set to the actual temperatures, for all other variables, such as gas and cloud abundance, the elements of $\mathbf{x}_n$ contain logarithmic values. This ensures that the modelled atmospheric profiles can never become negative.

Since the optimal estimation technique was developed for inverting Earth observation data where the *a priori* knowledge of the expected atmospheric profiles is good, it is worth considering how applicable it is to inverting spectra from other planetary atmospheres (and some trace gases in Earth's atmosphere), where we have very little *a priori* knowledge. While we can apply some *a priori* information, such as the known saturated vapour pressure profiles of certain gaseous constituents, the *a priori* profiles and covariance matrices for planetary retrievals are really little more than first guesses. Hence, while NEMESIS uses the formalism of optimal estimation, it uses it in a modified manner. The measurement errors in $\mathbf{S}_\varepsilon$ are set by the measured instrument noise equivalent radiances and/or variance of the data measured and off-diagonal elements are assumed to

be zero. Since the forward model contains numerous assumptions, both in terms of the accuracy of the line data bases and also in the correlated-*k* model itself, an appropriate amount of forward modelling error is then added to $\mathbf{S}_\varepsilon$, as described earlier, to set how well we think we can really fit the spectra. The *a priori* profile $\mathbf{x}_0$ is then set to a reasonable first guess, based on previous measurements and/or modelling studies and the level of vertical correlation in the *a priori* covariance matrix set. Tests are then performed with a range of *a priori* errors. If the errors are too large, the solution is unconstrained and NEMESIS does everything it can to minimise the difference between the measured and modelled spectra leading to apparently small retrieval errors (Eq.22), but also ill-conditioning and large vertical oscillations in the retrieved vertical profile. If instead, the *a priori* errors are set to be very small, then the solution is over-constrained and the retrieved profile differs little from the *a priori* profile (i.e. by much less than the retrieval error) and the retrieval errors are close to the *a priori* errors. By tuning the *a priori* errors we search for intermediate conditions where the solution is constrained quasi-equally by the data and by the *a priori* profile leading to solutions that match the measured spectra well, but which still have sufficient smoothing, supplied by the *a priori* covariance matrix $\mathbf{S}_x$ to be well-conditioned. This balance can be found either by making a range of test retrievals with different degrees of constraint and selecting the constraint that best balances the retrieval or alternatively by considering the elements of the matrix inverted in Eq. 20, $\left(\mathbf{K}_n\mathbf{S}_x\mathbf{K}_n^{\mathrm{T}} + \mathbf{S}_\varepsilon\right)$, and noting that roughly equal weight between the measurements and the constraints can be achieved by ensuring that the diagonal elements of $\mathbf{K}_n\mathbf{S}_x\mathbf{K}_n^{\mathrm{T}}$ and $\mathbf{S}_\varepsilon$ are of the same order of magnitude.

It is interesting to compare and contrast the approach used by NEMESIS with the constrained linear inversion technique [18, 19] used by many research groups. In this

approach, the non-linear iterative solution may be written, using our optimal estimation formalism, as:

$$\mathbf{x}_{n+1} = \mathbf{x}_0 + \hat{\mathbf{S}}_x \mathbf{K}_n^{\mathrm{T}} \left(\mathbf{K}_n \hat{\mathbf{S}}_x \mathbf{K}_n^{\mathrm{T}} + \gamma \mathbf{S}_\varepsilon \right)^{-1} \left(\mathbf{y}_m - \mathbf{y}_n - \mathbf{K}_n \left(\mathbf{x}_0 - \mathbf{x}_n \right) \right) \tag{23}$$

where $\hat{\mathbf{S}}_x$ is now the *a priori* correlation matrix (which provides vertical smoothing), $\mathbf{S}_\varepsilon$ and $\mathbf{K}_n$ are as before, and γ is an adjustable parameter used to fine-tune the balance between measurement and *a priori* constraint. When retrieving a single profile, the *a priori* correlation matrix $\hat{\mathbf{S}}_x$ is equal to our *a priori* covariance matrix $\mathbf{S}_x$, but with each row divided by the value of the diagonal component. Thus, comparing Eq. 23 with Eq. 20 we can see that to all intents and purposes the two methods are identical when it comes to retrieving continuous vertical profiles.

However, as outlined earlier, NEMESIS is very flexible in its parameterisation of vertical profiles. In most cases, we wish to extract the maximum amount of vertical information from the measured data and thus the state vector $\mathbf{x}_n$ contains the complete profile of each of the NVAR variable quantities to be retrieved giving a length of NX = NVAR × NPRO in total. However, a variable quantity can be modelled by any function, and in cases where the data do not provide a great deal of height information, each variable quantity can instead by parameterised as a mean value, or by a scaling factor of an assumed vertical profile or as a variable value below a certain pressure level falling with a variable fractional scale height above. In such cases, the vertical smoothing is encompassed in the parameterisation itself and thus there is no need to constrain the solution to lie close to the *a priori* value to prevent ill-conditioning. Hence, the errors in the *a priori* covariance matrix can be set to be large and the cost function (Eq. 18) is dominated by the closeness

of fit to the measured spectra, resulting in a solution that is a simple least-squares fit to the data.

Another feature of NEMESIS is that it can be used for two stage retrievals, either of the same quantities from different spectral ranges, or by using one spectral range to retrieve, say, temperature, and another to retrieve, say, composition. In the former case, the *a priori* state vector and covariance matrix for the second stage are set to the last iterated solution $\mathbf{x}_n$ and retrieved error $\hat{\mathbf{S}}$ of the first stage respectively. In the latter case, the effect of retrieval errors of the first quantity on the retrieval of the second quantity are calculated as an extra effective forward modelling error, which is added to $\mathbf{S}_\varepsilon$.

4. Conclusions

The NEMESIS retrieval model, originally developed for Cassini/CIRS Saturn and Titan studies, has been highly successful in analysing the data for that mission (e.g. [20–24]). An example of its application to retrieve a vertical profile of temperature in Titan's atmosphere from a Cassini/CIRS observations [21,22] (in thermal emission conditions) is shown in Fig. 3. However, since NEMESIS was designed from its conception as being of general applicability it is now also being applied to observations from a wide range of planetary instruments including Mars Climate Sounder on Mars Reconnaissance Orbiter, the VIRTIS instrument on Venus Express (multiple-scattering), the OIR instrument on Pioneer Venus (thermal emission) [25], the NIMS instrument on Galileo (multiple scattering) [26] and also ground-based telescope observations of Uranus (multiple scattering) [27]. An example of NEMESIS' application to retrieve a vertical profile of cloud specific density

(particles/gram) in Uranus' atmosphere from a UIST/UKIRT observation [27] of reflected sunlight (multiple-scattering conditions) is shown in Fig. 4.

Key features of the NEMESIS code are:

- For transmission and thermal emission calculations, NEMESIS makes full use of an implicit differentiation correlated-*k* code forward model making it accurate and extremely fast.
- NEMESIS is very adaptable and may retrieve atmospheric properties from visible-infrared-microwave observations of any planet.
- NEMESIS allows the simultaneous retrieval of any combination of up to four different continuous variables: Temperature, gaseous volume mixing ratio, cloud opacity, and para-H_2 fraction (for giant planet atmospheres).
- NEMESIS allows multi-stage retrievals.
- NEMESIS can model nadir, off-nadir and limb viewing geometries.
- In addition to the atmospheric profiles discussed, NEMESIS can also retrieve surface temperature, a surface albedo correction and also a tangent height or pressure level correction for limb observing geometries.
- NEMESIS can model the measurements returned by spectrometers and also channel radiometers by convolving calculated spectra with the filter profiles. This can either be done by calculating high resolution spectra and then convolving, or alternatively by using *k*-tables which have been generated using the filter profiles through Eq. 9. The latter approach is quicker, but should, for each application, be validated for accuracy by comparison with the former method.

It should be noted that while NEMESIS can model and retrieve any physical characteristic of any atmosphere in question, it does not fit spectral properties. The gas line absorption spectra must be supplied by the user and used to pre-calculate the required *k*-tables. Programs for generating these *k*-tables are supplied with NEMESIS. All other spectral properties, such as the cloud extinction and scattering cross-sections, surface emissivity etc., must also be supplied at the outset and are not modified during a retrieval. Retrieving spectral properties of, for example, dust, must thus be done by conducting retrievals with a range of guessed dust spectral properties and selecting those which minimise the solution cost function (Eq. 18). NEMESIS has not yet been applied to UV wavelengths, but this could be attempted, either by including the UV absorptions of different gases in the calculation of the *k*-tables or alternatively by modifying that part of the code which computes continuum gas absorptions.

By being general purpose, improvements in parameterisation and forward modelling schemes developed for one particular application are then automatically available for all other observations. Hence, although the initial development of the model including its implicit differentiation scheme was slow, the model has been steadily refined and improved through the analysis of observations of several planets in recent years. Thus, new observations of any planetary atmosphere can now be rapidly reduced using the same validated and familiar retrieval scheme. In short, NEMESIS has proven itself to be a highly adaptable and useful tool and is now used by a number of teams across the world.

We welcome further collaborations with new research teams and NEMESIS is freely available from the principal author.

Acknowledgements

The development of NEMESIS was made possible by funding from the United Kingdom Science and Technology Facilities Council.

References

[1] Berk A, Bernstein LS, Anderson GP, Acharya PK, Robertson DC, Chetwynd, JH, Adler-Golden SM. MODTRAN cloud and multiple scattering upgrades with application to AVIRIS. Remote Sensing of Envionment 1998;65:367–375.

[2] Snell HE, Moncet J-L, Anderson GP,Chetwynd JH, Miller S, Wang JG. FASCODE for the environment (FASE). Proc. SPIE 1995:2471;88–95.

[3] Stamnes K, Tsay S, Wiscombe W, Jayaweera K. Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting layered media. Appl. Opt. 1998;27:2502–2509.

[4] Rothman LS, Jacquemart D, Barbe A, Benner DC, Birk DM, Brown LR, Carleer MR, Chackerian C, Chance K, Coudert LH, Dana V, Devi VM, Flaud JM, Gamache RR, Goldman A, Hartmann JM, Jucks KW, Maki AG, Mandin JY, Massie ST, Orphal J, Perrin A, Rinsland CP, Smith MAH, Tennyson J, Tolchenov RN, Toth RA, Vander Auwera J, Varanasi P, Wagner G. The HITRAN 2004 molecular spectroscopic database. JQSRT 2005;96:139–204.

[5] Jaquinet-Husson N, Scott NA, Chedin A, Garceran K, Armante R, Chrsin AA, Barbe A, Birk M, Brown LR, Camy-Peyret C, Claveau C, Clerbaux C, Coheur PF, Dana V, Daumont L, Debacker-Barilly MR, Flaud JM, Goldman A, Hamdouni A, Hess M, Jaquemart D, Kopke P, Mandin JY, Massi S, Mikhailenko S, Nemtchinov V, Nikitin A, Newnham D, Perrin A, Perevalov VI, Regalia-Jarlot L, Rublev A, Schreier F, Schult I, Smith KM, Tashkun SA, Teffo JL, Toth RA, Tyuterov VG, Vander Auwera J, Varanasi P, Wagner G. The 2003 edition of the GEISA/IASI spectroscopic database. JQSRT 2005;95:429–467.

[6] Lacis AA, Oinas V. A description of the correlated-*k* distribution method for modelling nongray gaseous absorption, thermal emission, and multiple scattering in vertically inhomogeneous atmospheres. J Geophys Res 1991;96:9027–9063.

[7] Press WH, Teukolsky SA, Vetterling WT, Flannery BP. Numerical Recipes in Fortran. 2nd Edition. Cambridge University Press;1992.

[8] Irwin PGJ, Calcutt SB, Sihra K, Taylor FW, Weir AL, Ballard J, Johnston WB. Band parameters and *k* coefficients for self-broadened ammonia in the range 4000-11000 cm^{-1}. JQSRT 1999;62:193–204.

[9] Irwin PGJ, Sromovsky LA, Strong EK, Sihra K, Teanby NA, Bowles N, Calcutt SB, Remedios JJ. Improved near-infrared methane band models and *k*-distribution parameters from 2000 to 9500 cm^{-1} and implications for interpretation of outer planet spectra. Icarus 2006;181:309–319.

[10] Nixon CA. Remote Sounding of the Atmosphere of Titan. D.Phil. Thesis. University of Oxford;1998.

[11] Flasar FM, Kunde VG, Abbas MM, Achterberg RK, Ade P, Barucci A, Bézard B, Bjoraker GL, Brasunas JC, Calcutt S, Carlson R, Césarsky CJ, Conrath BJ, Coradini A, Courtin R, Coustenis A, Edberg S, Edgington S, Ferrari C, Fouchet T, Gautier D, Gierasch PJ, Grossman K, Irwin P, Jennings DE, Lellouch E, Mamoutkine AA, Marten A, Meyer JP, Nixon CA, Orton GS, Owen TC, Pearl JC, Prange R, Raulin F, Read PL, Romani PN, Samuelson RE, Segura ME, Showalter MR, Simon-Miller AA, Smith MD, Spencer JR, Spilker LJ, Taylor

FW. Exploring the Saturn system in the thermal infrared: The Composite Infrared Spectrometer. Space Sci Rev. 2004;115:169–297.

[12] Borysow JL, Trafton L, Frommhold L, Birnbaum G. Modelling of pressure-induced far-infrared absorption spectra: Molecular hydrogen pairs. ApJ. 1985; 296:644-654.

[13] Borysow J, Frommhold L, Birnbaum G. 1988. Collision-induced roto-translational absorption spectra of H_2-He pairs at temperatures from 40 to 3000 K. ApJ. 1988;326:509-515.

[14] Andrews DR. An Introduction to Atmospheric Physics. Cambridge University Press;2000.

[15] Plass GN, Kattawar GW, Catchings FE. Matrix operator theory of radiative transfer. 1: Rayleigh Scattering. Appl. Opt. 1973;12:314–329.

[16] Backus GE, Gilbert JF. Uniqueness in the inversion of inaccurate gross earth data. Phil Trans Roy Soc Lond 1970;A266:123–192.

[17] Rodgers CD. Inverse methods for atmospheric sounding: Theory and practice. World Scientific;2000.

[18] Conrath BJ, Gierasch PJ, Ustinov EA. Thermal structure and para hydrogen fraction in the outer planets from Voyager IRIS measurements. Icarus 1998;135:501–517.

[19] Hanel RA, Conrath BJ, Jennings DE, Samuelson RE. Exploration of the Solar System by Infrared Remote Sensing: Second Edition. Cambridge University Press;2003.

[20] Irwin PGJ, Parrish P, Fouchet T, Calcutt SB, Taylor FW, Simon-Miller AA, Nixon CA. Retrievals of Jovian tropospheric phosphine from Cassini/CIRS. Icarus 2005;172:37–49.

[21] Teanby NA, Irwin PGJ, de Kok R, Nixon CA, Coustenis A, Bézard B, Calcutt SB, Bowles NE, Flasar FM, Fletcher L, Howett C, Taylor FW. Latitudinal variations of HCN, HC_3N, and C_2H_2 in Titan's stratosphere derived from Cassini CIRS data. Icarus 2006;181:243–255.

[22] Teanby NA, Irwin PGJ, de Kok R, Vinatier S, Bézard B, Nixon CA, Flasar FM, Calcutt SB, Bowles NE, Fletcher L, Howett C, Taylor FW. Vertical profiles of HCN, HC_3N, and C_2H_2 in Titan's atmosphere derived from Cassini/CIRS data. Icarus 2007;186:364–384.

[23] Fletcher LN, Irwin PGJ, Teanby NA, Orton GS, Parrish PD, Calcutt SB, Bowles N, de Kok R, Howett C, Taylor FW. The meridional phosphine distribution in Saturn's upper troposphere from Cassini/CIRS observations. Icarus 2007;188:72–88.

[24] de Kok R, Irwin PGJ, Teanby NA, Lellouch E, Bezard B, Vinatier S, Nixon CA, Fletcher L, Howett C, Calcutt SB, Bowles NE, Flasar FM, Taylor FW. Oxygen compounds in Titan's stratosphere as observed by Cassini CIRS, Icarus 2007;186:354–363.

[25] Koukouli ME, Irwin PGJ, Taylor FW. Water vapour abundance in Venus' middle atmosphere from Pioneer Venus OIR and Venera 15 FTS measurements. Icarus 2005;173:84-49.

[26] Howett CJA, Carlson RW, Irwin PGJ, Calcutt SB. Optical constants of ammonium hydrosulfide ice and ammonia ice. J Opt Soc America 2007;24:126–136.

[27] Irwin PGJ, Teanby NA, Davis GR. Latitudinal variations in Uranus' vertical cloud structure from UKIRT/UIST observations. ApJL 2007;665:L71-L74.

Appendix

The main equation of non-linear optimal estimation can be written as

$$\mathbf{x}_{n+1} = \mathbf{x}_0 + \mathbf{S}_x\mathbf{K}_n^{\mathrm{T}}\left(\mathbf{K}_n\mathbf{S}_x\mathbf{K}_n^{\mathrm{T}} + \mathbf{S}_\varepsilon\right)^{-1}\left(\mathbf{y}_m - \mathbf{y}_n - \mathbf{K}_n\left(\mathbf{x}_0 - \mathbf{x}_n\right)\right) \tag{A1}$$

or equivalently (Rodgers, 2000)

$$\mathbf{x}_{n+1} = \mathbf{x}_0 + \mathbf{G}_n\left(\mathbf{y}_m - \mathbf{y}_n\right) - \mathbf{A}_n\left(\mathbf{x}_0 - \mathbf{x}_n\right) \tag{A2}$$

where $\mathbf{G}_n$ is the gain matrix and $\mathbf{A}_n$ (= $\mathbf{G}_n\mathbf{K}_n$) is the averaging kernel matrix. The method of calculating the gain matrix $\mathbf{G}_n$ shown here, i.e.:

$$\mathbf{G}_n = \mathbf{S}_x\mathbf{K}_n^{\mathrm{T}}\left(\mathbf{K}_n\mathbf{S}_x\mathbf{K}_n^{\mathrm{T}} + \mathbf{S}_\varepsilon\right)^{-1} \tag{A3}$$

can be very slow for cases where the length m of the measurement vector is large. In such cases, NEMESIS makes use of an equivalent formulation of the gain matrix:

$$\mathbf{G}_n = \left(\mathbf{K}_n^{\mathrm{T}}\mathbf{S}_\varepsilon^{-1}\mathbf{K}_n + \mathbf{S}_x^{-1}\right)^{-1}\mathbf{K}_n^{\mathrm{T}}\mathbf{S}_\varepsilon^{-1} \tag{A4}$$

which, since $\mathbf{S}_\varepsilon$ is assumed to be diagonal, and $\mathbf{S}_x^{-1}$ is already calculated in order to work out the cost function, is much faster to calculate for $n < m$.

Figures

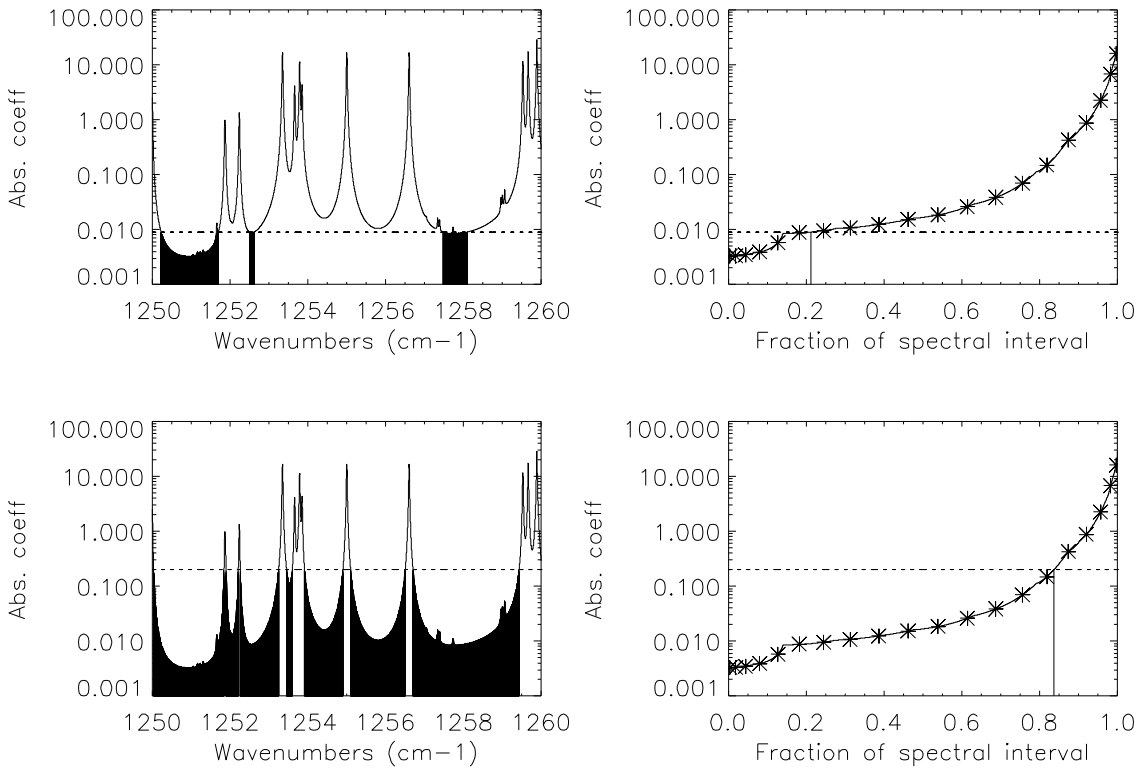

Figure 1. The *k*-distribution technique applied to the absorption of methane in Saturn's atmosphere. A high-resolution absorption coefficient spectrum of methane for the 1250-1260 cm^{-1} region is shown for typical conditions in Saturn's atmosphere of 0.1 bar, 85K and a volume mixing ratio of 4.5×10^{-3}. The left-hand plots show the highly detailed absorption spectra, together with the portions of the spectrum with absorption coefficients less than 0.009 and 0.2 respectively. Parts of the spectrum with absorption coefficients less than this are marked in black. The right-hand plots show the *k*-distributions, derived simply by ranking the absorption coefficients in increasing value, together with the sampled quadrature points (asterisks) for a 20-point numerical Gaussian integration. In this example we can see that just over 20% of the interval has absorption less than 0.009, while 84% has absorption less than 0.02.

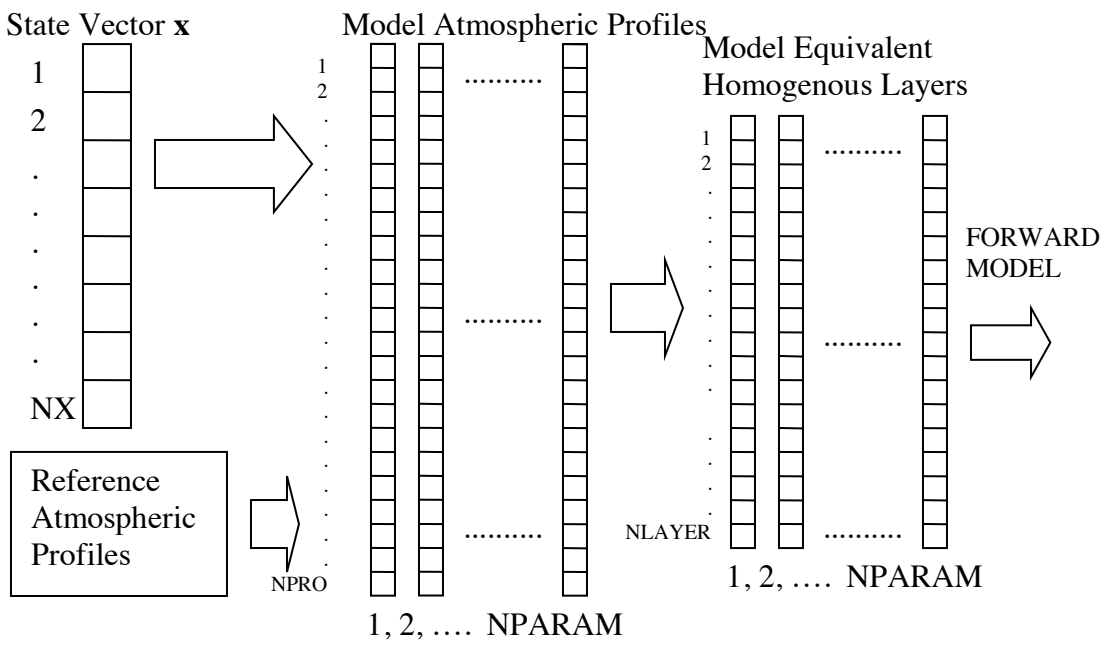

Figure 2. Schematic diagram showing how the reference atmospheric profiles and the model state vector are used at each stage of the retrieval to generate a set of NPARAM profiles with NPRO vertical points. These profiles are then split into NLAYER equivalent homogenous layers to allow the radiative transfer calculation to be conducted.

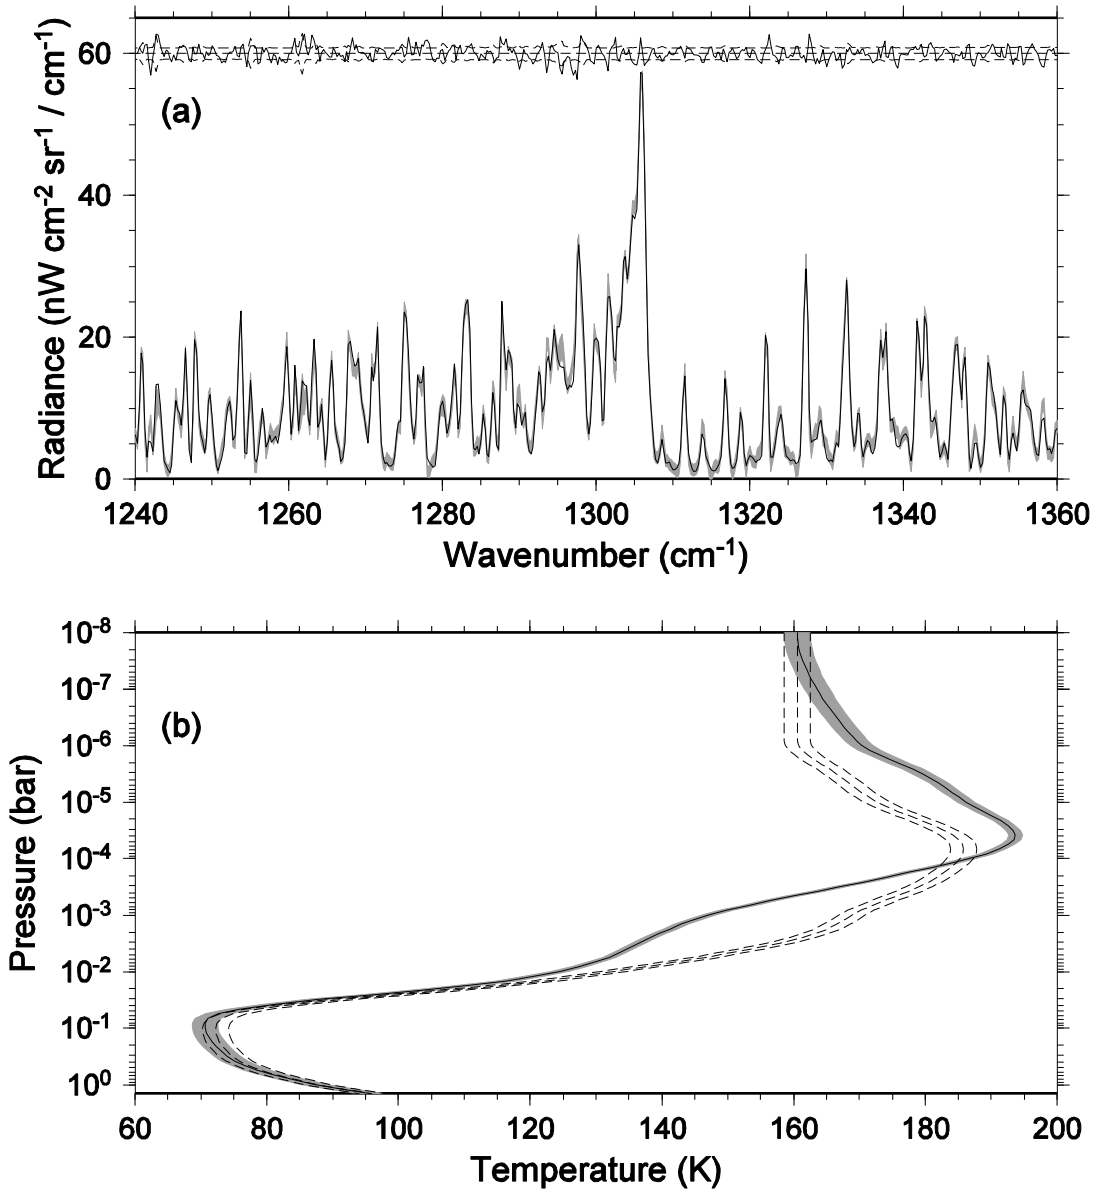

Figure 3. Example of a NEMESIS thermal emission temperature retrieval for a Cassini/CIRS near-nadir observation of Titan. (a) shows the measured spectrum and error limits (in grey) in the ν_4 methane absorption band between 1240 and 1360 cm^{-1}. The spectrum fitted by NEMESIS is shown as the solid line. (b) shows the retrieved vertical temperature profile, where the retrieved profile is the solid line and the error limits are

indicated by the grey shaded region. The *a priori* profile and error limits are indicated by the dashed lines. In this example the retrieved profile varies most greatly from the a priori profile between 10^{-2} and 10^{-6} bars, where the weighting functions for this observation peak and tend back towards the *a priori* profile above and below this region.

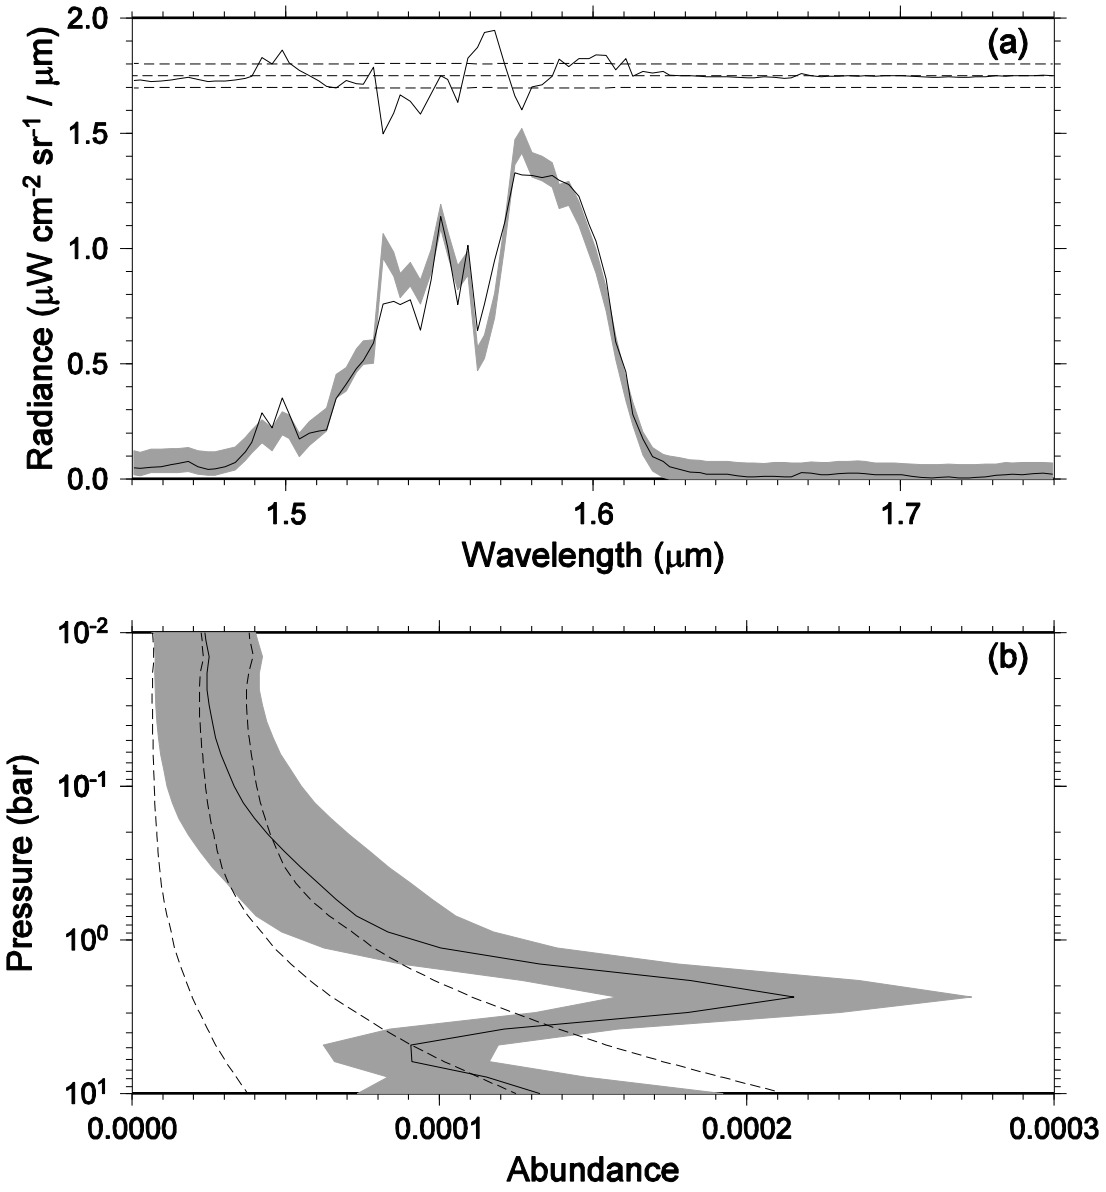

Figure 4. Example of a NEMESIS multiple-scattering cloud retrieval for a UKIRT/UIST near-nadir observation of reflected sunlight from Uranus' clouds. (a) shows the measured spectrum and (b) shows the retrieved cloud density profile (particles/gram). The identity of the lines and shaded regions are as described in Fig. 3.

Tables

Symbol	Name	Dimension	Purpose
Atmospheric profile and layer derivatives			
M	Profile Property Gradient Matrix	NX×NPARAM×NPRO	Contains rate of change of NPARAM model atmospheric profiles (represented with NPRO vertical levels) with respect to each of the NX elements of the state vector.
L	Layer Property Gradient Matrix	NPARAM×NLAYER×NPRO	Contains rate of change of the NPARAM mean properties of the NLAYER homogeneous layers with respect to the NPARAM model atmospheric profiles, represented with NPRO vertical levels.
Radiance Derivatives			
J	Layer Radiance Gradient Matrix	NPARAM×NLAYER	Contains rate of change of calculated radiance with respect to the NPARAM mean properties of the NLAYER homogeneous layers.
P	Profile Radiance Gradient Matrix	NPARAM×NPRO	Contains rate of change of calculated radiance with respect to the NPARAM properties of the NPRO model atmosphere vertical levels.
N	Radiance Derivative Vector	NX	Contains the rate of change of radiance for a with respect to each of the NX state vector elements.
K	Functional Derivative or Jacobian Matrix	NY×NX	Contains the rate of change of radiance at each of the NY wavelengths/wavenumbers with respect to each of the NX state vector elements.

Table 1. ‘gradient’ and ‘functional derivatives’ vectors and matrices defined at different stages within the NEMESIS retrieval model.