Carbon Monoxide Retrievals from MIPAS

Ben Fletcher
St. John's College
Oxford University, UK

November 13, 2002

Abstract

Potential microwindows are determined to investigate the feasibility of retrieving information on carbon monoxide concentration from the Michelson Interferometer for Passive Atmospheric Sounding, MIPAS. These microwindows are tested using simulated data. The significance of not having local thermal equilibrium (the NLTE effects) in determination of carbon monoxide levels is investigated. A comparison is made between the efficiency of using *a priori* vibrational temperature data with the idea of retrieving vibrational temperature from the spectrum as a means of taking into account the fact that carbon monoxide in the atmosphere is not in local thermal equilibrium. An attempt is made to effect a retrieval from live MIPAS data as a test of concept. Future avenues for research in this field are outlined.

1 Introduction

Largely due to human influences, levels of carbon monoxide have been increasing in the atmosphere since the industrial revolution. The gas is produced as a consequence of incomplete combustion of fossil fuels and is a major contributor to smog.

In order to determine how much of an influence carbon monoxide is having on atmospheric chemistry, it is important to be able to map how the gas is distributed, particularly in the troposphere. The goal of this project is to investigate the feasibility of extracting this information from the data recorded by MIPAS, the Michelson Interferometer for Passive Atmospheric Sounding.

1.1 Carbon Monoxide in the Atmosphere

Carbon monoxide is one of the many pollutants which are naturally regulated by reactions with hydroxyl ions in the troposphere. Unfortunately, the abundance of carbon monoxide produced by human industries means that more of the hydroxyl ions react with the gas, leaving fewer available to deal with other pollutants. Thus carbon monoxide not only serves as a pollutant in its own right, but hinders the natural mechanism for cleaning up other pollutants in the atmosphere.

Carbon monoxide also plays a key role in a reaction path that can produce ozone in the troposphere. Ironically, while the lack of ozone higher up in the atmosphere can cause significant problems with radiation, in the troposphere ozone is a pollutant and increased levels of it can lead to smog.

Production from carbon monoxide is one of several possible origins for tropospheric ozone: by measuring the concentration of ozone and carbon monoxide simultaneously, it becomes possible to determine how much of the ozone present in the troposphere was produced this way. These data provide a valuable perspective on the significance of carbon monoxide in ozone pollution of the troposphere.

1.2 About MIPAS

The Michelson Interferometer for Passive Atmospheric Sounding was launched in March, 2002, on board the European Space Agency satellite ENVISAT. It is a Fourier-transform spectrometer designed for studying trace gases in the atmosphere by observing their emissions in the mid infra-red. It has a high resolution (0.025 cm^{-1}) and completes a spectral scan in 4 seconds.

MIPAS was launched with three objectives in mind: [5]

- Simultaneous and global measurements of geophysical parameters in the middle atmosphere;
  Stratospheric chemistry: O_3, H_2O, CH_4, N_2O, NO_2 and HNO_3; and
  Climatology: Temperature, CH_4, N_2O, O_3;
- Study of chemical composition, dynamics, and radiation budget of the middle atmosphere;
- Monitoring of stratospheric O_3 and CFC's.

For known gases in local thermal equilibrium, an emission spectrum provides the information needed to determine the local temperature and the composition of the atmosphere.. In fact, it provides almost too much information - certainly too much to be analyzed quickly. To deal with this, the portions of the spectrum containing the most information about the factor to be studied are isolated, and only these regions are analyzed. The areas of the spectrum thus selected are known as microwindows: they are explained in more detail in section 3.1 The information from these microwindows can be used in combination with known variations in pressure to determine the concentrations of different gases at a range of heights through the troposphere.

Note that carbon monoxide is not one of the gases which MIPAS was intended to study. However, the instrument monitors a wide range of radiation, in bands with wavenumbers from 680 to 2410 cm^{-1}. Carbon monoxide's primary emission line lies at 2140 cm^{-1}, in the center of MIPAS' D-band. This suggests that there ought to be enough information contained within a spectrum from MIPAS to determine the concentration of carbon monoxide in that portion of the atmosphere at that time. The difficulty arises in compensating for the fact that carbon monoxide is not in local thermal equilibrium in the atmosphere (see section 2.3). -

Clearly if carbon monoxide can be retrieved from MIPAS data, it will advance both the study of the chemical composition of the middle atmosphere and the monitoring of ozone in the atmosphere.

2 Data Analysis and Retrievals

When working with spectral data from satellites, there are generally a very large number of measurements compared to the number of parameters to be retrieved. Data and error analysis in these circumstances requires some relatively complicated statistics, as described by Rodgers. This report will be limited to a general introduction, leaving the details to his work. [1, 3, 7]

2.1 Theory

Consider a set of measurements y_1, y_2, ..., y_n being used to obtain a set of parameters x_1, x_2, ..., x_m. In this case, for example, the $\mathbf{y}$ would represent a series of spectra obtained from MIPAS while the set $\mathbf{x}$ would represent a vertical profile of the carbon monoxide concentration in the area being studied.

Assume that the parameters are linearly related to the measurements:

$$\mathbf{y} = \mathbf{K}\mathbf{x} + \delta\mathbf{y}^{\text{random}} + \delta\mathbf{y}^{\text{systematic}}$$

$\mathbf{K}$, in this context, is known as a Jacobian matrix - it represents the relationship between $\mathbf{x}$ and $\mathbf{y}$. The two terms in $\delta(\mathrm{y})$ represent the random and systematic error in the measurements. Random error is due to instrument noise, while systematic error will stem from mis-calibration or an error in the forward model (discussed below).

By analogy to a Taylor series, it is clear that this assumption is valid even for non-linear relationships when considering small variations around a known value. This becomes important when considering radiance and temperature, as they are not linearly related except at very long wavelengths.

For use in analysis the inverse of this relationship is desired - that is, the operation that maps a set of measurements to the parameters behind the system. This can be done by a statistical technique of least squares solution. This technique finds the best guess at $\mathbf{x}$ by minimizing the difference between data predicted from a guess at $\mathbf{x}$ and the observed data $\mathbf{y}$. The software then adjusts its guess for $\mathbf{x}$ and repeats the least-squares analysis. After sufficient iterations, this will produce a profile which cannot be brought into better agreement with the observed data. The parameters used to calculate this profile may be combined in a weighted average with the *a priori* data, which is our best guess at the parameters before any observations were taken. This combination produces the optimal estimation, which is the best guess available for the actual values of these parameters. This process is known as *retrieval.*

Retrieval seems a somewhat awkward way of performing the calculation, but works well because it is vastly simpler to calculate a set of predicted measurements based on a given set of parameters than vice-versa. The set of software used to perform this calculation is known as a *forward model.*

2.2 The Reference Forward Model

The MIPAS Reference Forward Model is a radiative transfer model which uses a given atmospheric profile and viewing geometry to generate simulated MIPAS spectra. Technical details on its operation can be found in [4].

In this project, the RFM was used to provide sample data relating to carbon monoxide in the atmosphere. This data can then be analyzed using the process outlined in Dudhia [1] to determine the most suitable microwindows. The RFM will then be further used in the retrieval process to test the utility of the chosen microwindows.

2.3 Local Thermal Equilibrium

In analyzing data from MIPAS and other infrared remote sensing instruments, generally the gases are assumed to be in local thermal equilibrium. This means that the temperature used to describe the population of the vibrational and rotational states is the same as that used to describe the kinetic temperature. This is equivalent to saying that the population of the gas molecules in different energy levels will follow the Maxwell-Boltzmann distribution, at least approximately. It implies that the emissions from gases in the area will be described by the Planck function for the ambient temperature [6]. For carbon monoxide, the vibrational excitation energies are typical of the photon energies of radiation from the sun. This means that when sunlight is incident on carbon monoxide, a substantial number of molecules are pushed into vibrational excited states so that carbon monoxide cannot be considered to be in local thermal equilibrium, and the radiance will no longer be modelled by the Planck function.

The effects of this, known as NLTE, non-(local thermal equilibrium), are shown in Figure 1. This shows the difference between the spectra taken at night and the spectra taken in daylight. This difference arises largely because of the fact that when carbon monoxide is being bombarded with photons, it cannot be considered to be in local thermal equilibrium.

For the purposes of the RFM, carbon monoxide is modelled as any other gas except that it uses a higher vibrational temperature in place of the kinetic temperature. Emission is described by the Planck function for this new, higher temperature. The modelling is thus straightforward; the problem lies in determining the correct vibrational temperature profile since this depends on a variety of external factors such as solar zenith angle, cloud cover, temperature profile etc.

Under this assumption, the default way of handling NLTE effects in the RFM is to supply a vibrational temperature for each altitude. This vibrational temperature is then used to calculate the radiance. A more sophisticated method will also be tested, in which the vibrational temperature is determined along with the carbon monoxide concentration in order to get the best match to the observed spectrum. This is known as retrieving vibrational temperature - see section 3.2 for details on the retrieval process.

3 Microwindows

3.1 What are Microwindows?

For an illustration of the problems facing researchers, consider Figure 2. The first graph is an illustration of the combined emission spectrum in the MIPAS D-band, which is dominated by the emissions from carbon dioxide, ozone and water vapor. The second graph shows only the spectrum of carbon monoxide. Clearly there are great difficulties in distinguishing carbon monoxide from all of the other gases on the basis of the full spectrum.

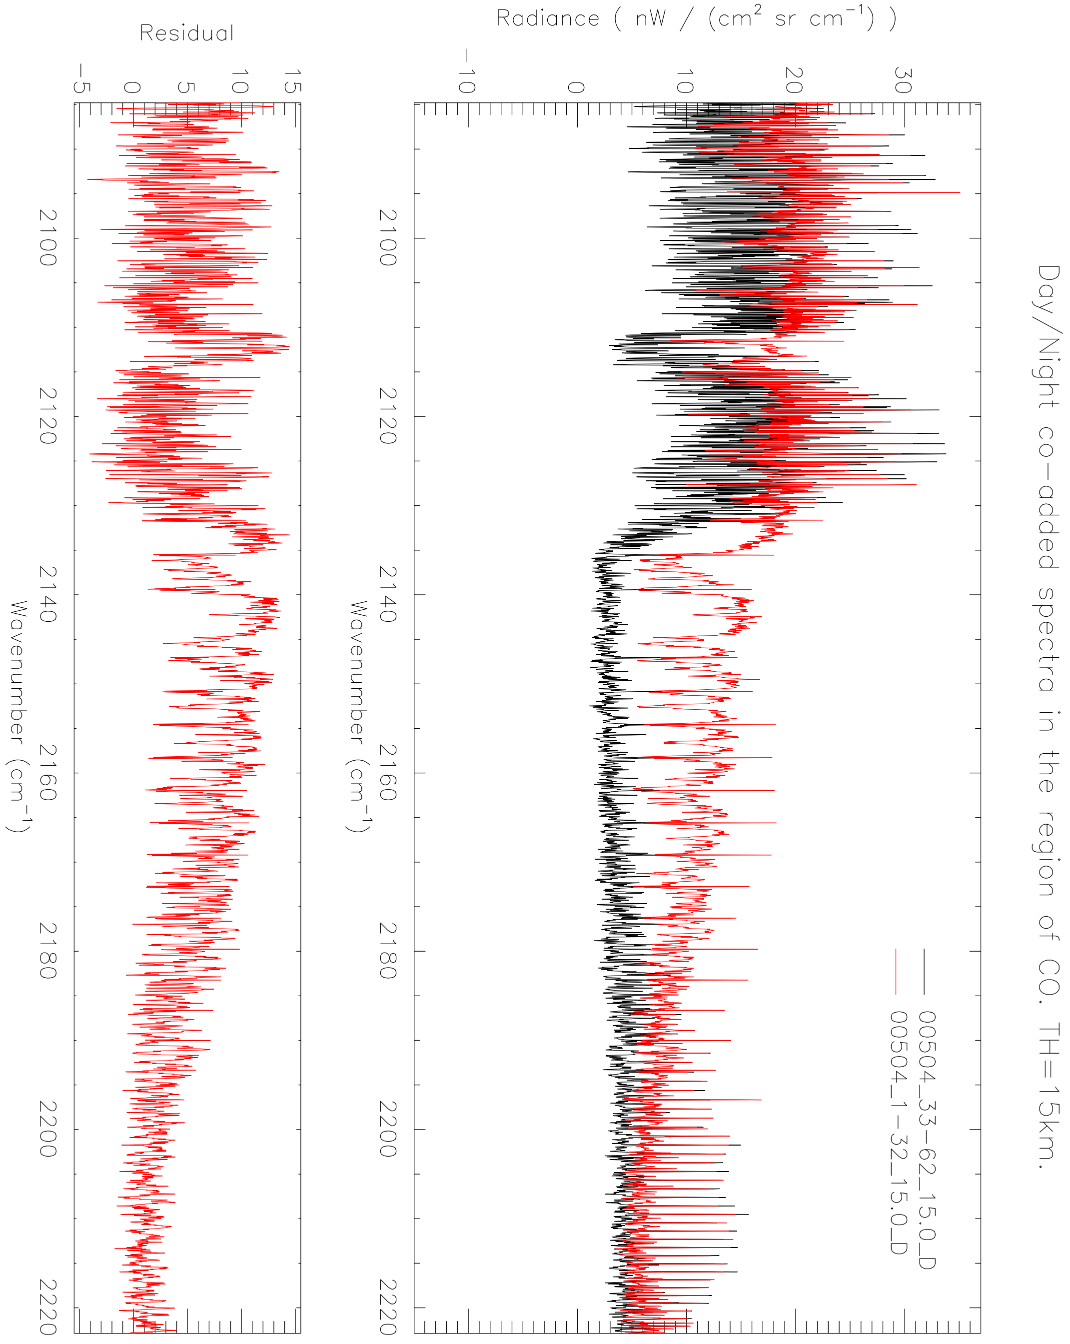

Figure 1: This figure illustrates the actual MIPAS spectra for 5th April 2002 by day and by night. This is due to the effects of gases - primarily carbon monoxide - not being in local thermal equilibrium (NLTE effects). These spectra were obtained at a tangent height of 15 km. In this as in all other graphs of radiance in this report, the radiance unit (RU) should be read as nW cm^{-2} sr^{-1} (cm^{-1})$^{-1}$. Source: Lopez-Puertes (pers. comm)

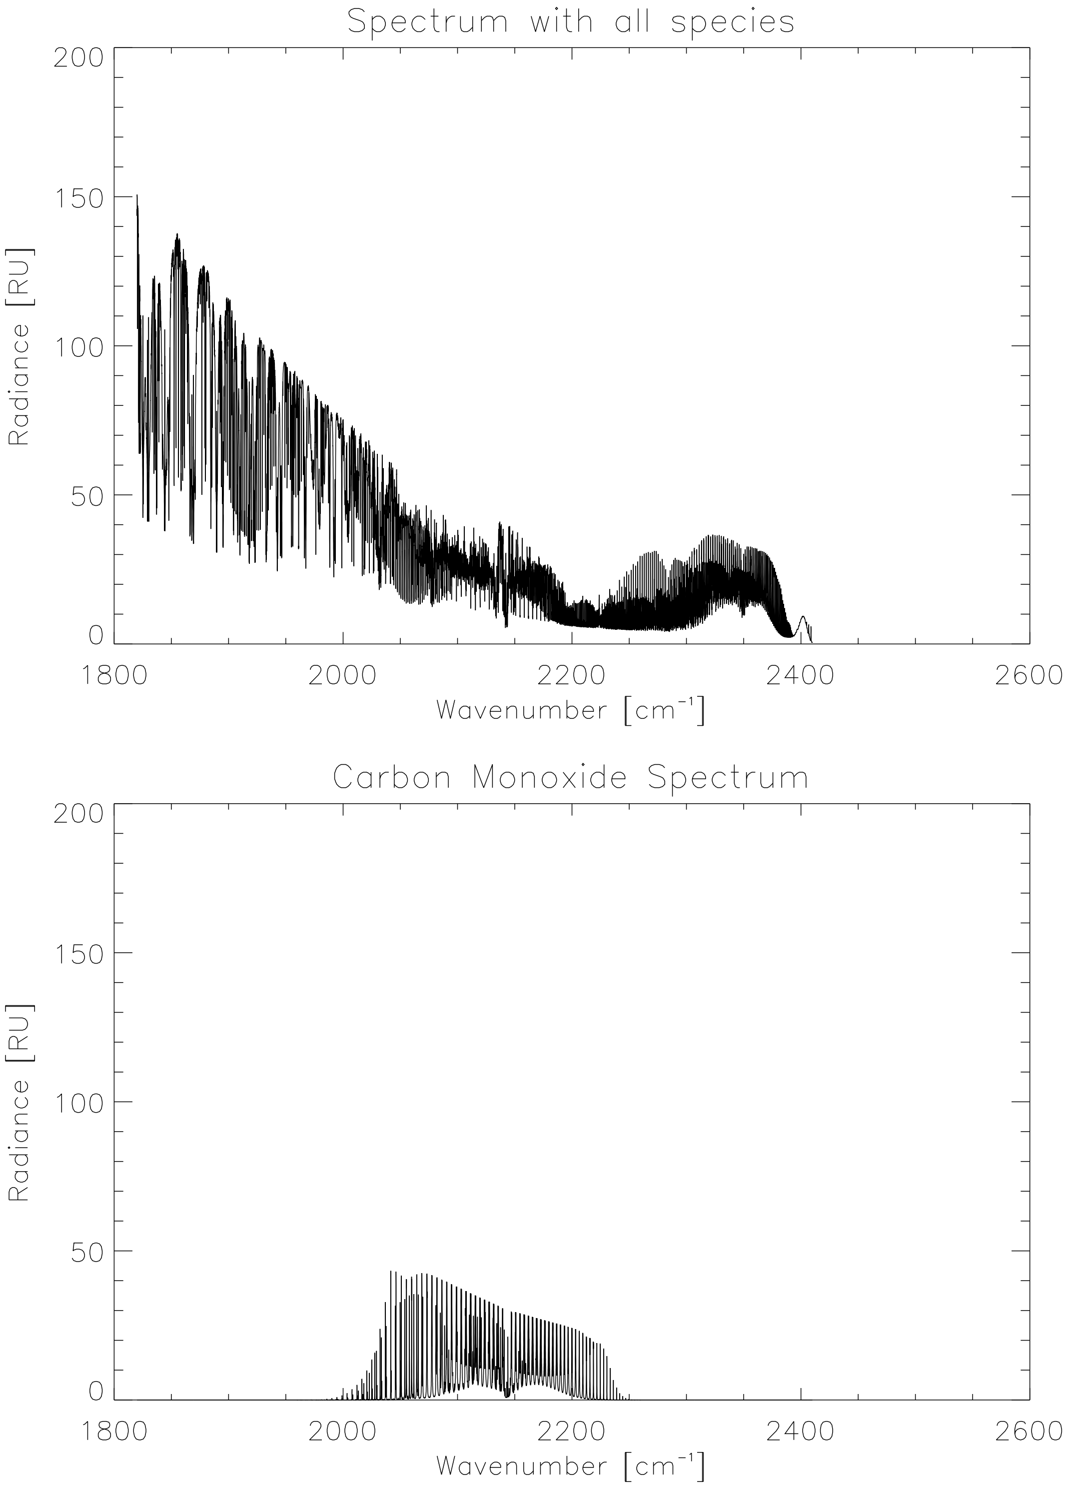

Figure 2: This figure compares a full spectrum typical of MIPAS to the carbon monoxide emission spectrum, both at a tangent height of 18 km

Indeed, this is a wider problem - instruments like MIPAS deliver a large quantity of information to researchers every day. It is simply not possible to process all of the data received in real time - the necessary computing power and time is prohibitive.

Fortunately, in examining the carbon monoxide spectrum in Figure 2 it becomes clear that it isn't necessary to consider all of the data that MIPAS receives. The information to be retrieved lies only in the portion of the spectrum due to carbon monoxide, which lies between 2050 and 2200 cm^{-1}. Even within this region, most of the significant data will clearly be found around the peaks of the spectrum.

Selecting microwindows formalizes this concept and carries it one step further forward - the selection chooses small areas of the spectrum which, when analysed, yield the majority of the total information available about carbon monoxide concentrations. This is, in fact, the definition of a microwindow: a contiguous subset of the spectrum that contains the most information about the quantity being examined (in this case, carbon monoxide concentration).

3.2 Selecting Microwindows

The process used for selecting microwindows is described in reference [1]. It simulates a full retrieval of data, including the propagation of both systematic and random errors; microwindows are then selected to maximize the information content with respect to carbon monoxide concentration.

To summarize the method, consider the data obtained by MIPAS in a single profile as being an overlay of a spectrum from each tangent height. This provides us with a grid of altitude-wavenumber points, each of which has a radiance value. A microwindow is then a box containing one or more points on this grid. The program used to perform the calculations - 'mwmake' - starts with microwindows consisting of single points in this grid. This tiny microwindow is then evaluated for how much random error contaminates the signal in it, relative to the amount of computing time required to retrieve it. This ratio is expressed as a scalar quantity known as the *figure of merit.*

The program next attempts to extend these small microwindows to cover neighbouring points in an attempt to improve the figure of merit. Individual points that improve the figure of merit will be added to the microwindow and the new perimeter of the microwindow will be considered for more points that will improve the figure of merit. On the altitude-wavenumber grid, then, a general microwindow will have an irregular shapes which can enclose excluded points. Expansion stops when the microwindow cannot be extended any further without reducing the figure of merit.

This process generates a list of potential microwindows for whatever parameter is to be retrieved. This list can then be sorted for information alone or to maximize the amount of information provided with limited computing time. Once the list of windows is selected and optimized, they can be tested by using them to retrieve a profile from data generated by a known atmosphere. Ideally, a retrieval using the microwindows will produce a profile that is identical with the known profile to within the limits of the noise on the data.

Microwindow Number	Lower Wavenumber [cm^{-1}]	Upper Wavenumber [cm^{-1}]
1	2138.525	2141.525
2	2163.725	2166.725
3	2133.100	2136.100
4	2153.200	2156.275
5	2149.125	2152.125
6	2160.700	2163.700
7	2110.300	2113.300
8	2064.250	2067.250

Table 1: The microwindows selected without considering NLTE effects. They are arranged in order of the information they convey - note that the first window in the list contains carbon monoxide's primary line. All microwindows apply over the full range of tangent altitudes from 6 to 68 km.

3.3 First Approximation - Local Thermal Equilibrium

Initially, carbon monoxide was considered without the complicating effects of lacking local thermal equilibrium. Using this assumption, the RFM will try to match the Planck function of the kinetic temperature to the observed spectrum. In other words, the vibrational temperature is assumed to be equal to the kinetic temperature. While this doesn't completely describe the behavior of carbon monoxide in the atmosphere, it does describe carbon monoxide at night and at low altitude (where pressure is higher, and thus there are more collisions to randomize the energy distribution in line with the kinetic temperature.)

In order to reduce error and ensure that the selected microwindows sampled the whole spectrum effectively, several factors were adjusted before we settled on the ideal set of microwindows. Among other things, the number of absorbers which the selection routine considered, the method of growing microwindows (as discussed above) and area searched were varied to get the best results.

From these trials, it was initially found that the microwindow selection program, MWMAKE, avoided the region between 2050 and 2150 cm^{-1}, probably due to the high activity of other species (particularly ozone) in that region. Since carbon monoxide's primary line is at 2140 cm^{-1}, this clearly discards a large amount of information on the concentration of the gas.

In order to achieve good coverage of the full spectrum, MWMAKE was configured to ignore the contributions from other species when choosing microwindows. This allowed the selection routine to choose the strongest CO lines without considering the impact of other features of a normal atmospheric spectrum. This obtained the set of microwindows illustrated here, and laid out in Table 1. The significance of the interference from other quantities can be evaluated when the microwindows are tested through a full retrieval.

3.4 Without LTE

Two typical profiles for vibrational temperature can be seen in Figure 5. These data came from the *a priori* data used to calculate MIPAS retrievals. Clearly,

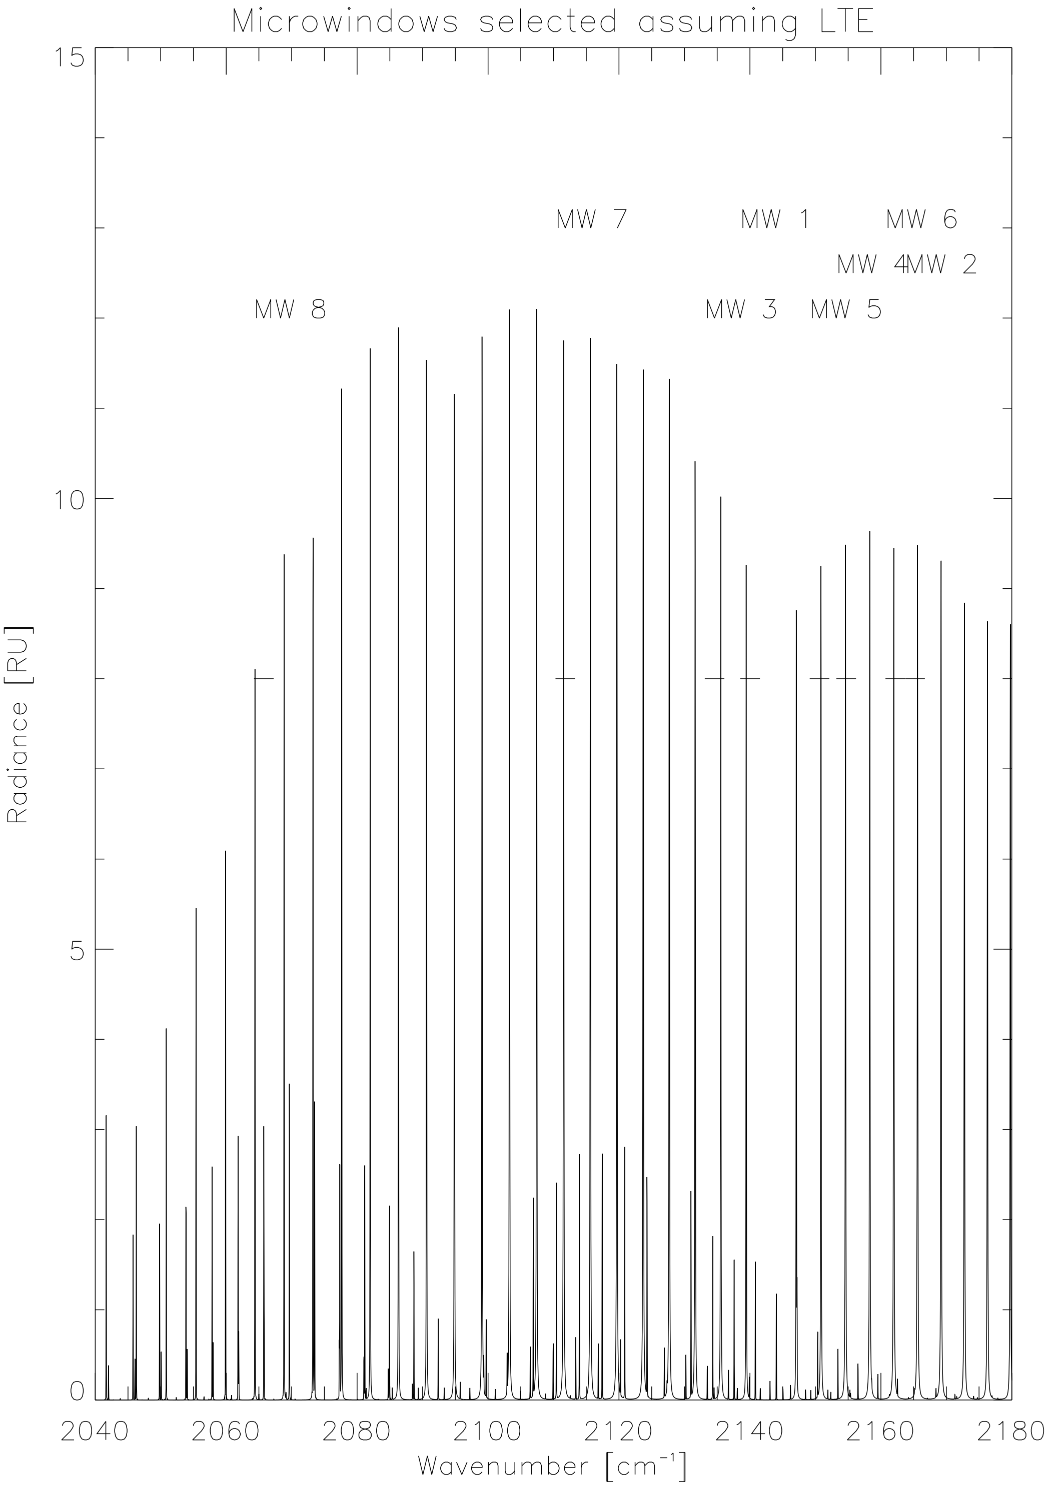

Figure 3: This figure illustrates the microwindows selected while ignoring NLTE effect and the contributions of other species

Window Number	Lower Wavenumber [cm^{-1}]	Upper Wavenumber [cm^{-1}]	Maximum Tangent Height [km]
1	2144.375	2147.375	68
2	2134.475	2135.950	68
3	2182.875	2183.500	68
4	2167.650	2169.825	60
5	2157.750	2159.650	60
6	2159.725	2162.625	68
7	2137.225	2140.225	68
8	2140.250	2142.275	68
9	2187.525	2190.525	68
10	2111.775	2114.775	68

Table 2: The microwindows selected with consideration of NLTE effects. All microwindows begin at a tangent-height of 6 km.

by the time the retrievals are complete at 68 km, there is already a substantial variation between ambient temperature and the vibrational temperature of the carbon monoxide. This would be expected this to produce a significant error in the retrieval if not accounted for.

To correct this error, a set of microwindows were designed to retrieve not only the concentration of carbon monoxide but also the vibrational temperature. This is achieved by calculating Jacobian spectra for both carbon monoxide concentration and vibrational temperature profiles. MWMAKE then uses these Jacobians to select a set of microwindows that optimise the information gained from a joint retrieval of the two parameters. This set of 57 microwindows was then sorted according to their figures of merit, and the top ten were chosen to be used. The microwindows thus selected are illustrated in Table 2 and in Figure 4.

It is interesting to note that the microwindow including the primary line at 2140 cm^{-1} is no longer the most useful in terms of information contained. This may be because that area of the spectrum is saturated, with the instrument unable to detect all of the information present. Instead, the most valuable microwindows are just to either side of the primary line, where a slightly lower intensity may allow more discrimination.

4 Microwindow Testing - Retrievals

4.1 Strategy

After microwindow selection there are two sets of microwindows ready for testing. The first was selected to minimize random error on a carbon monoxide retrieval that assumes no difference between kinetic and vibrational temperature, while the second was selected to minimize total error, taking into account the non-LTE effects. The task is now to examine both of these sets of microwindows to see how well they perform in an actual retrieval.

The retrievals themselves were performed using MORSE - MIPAS Oxford

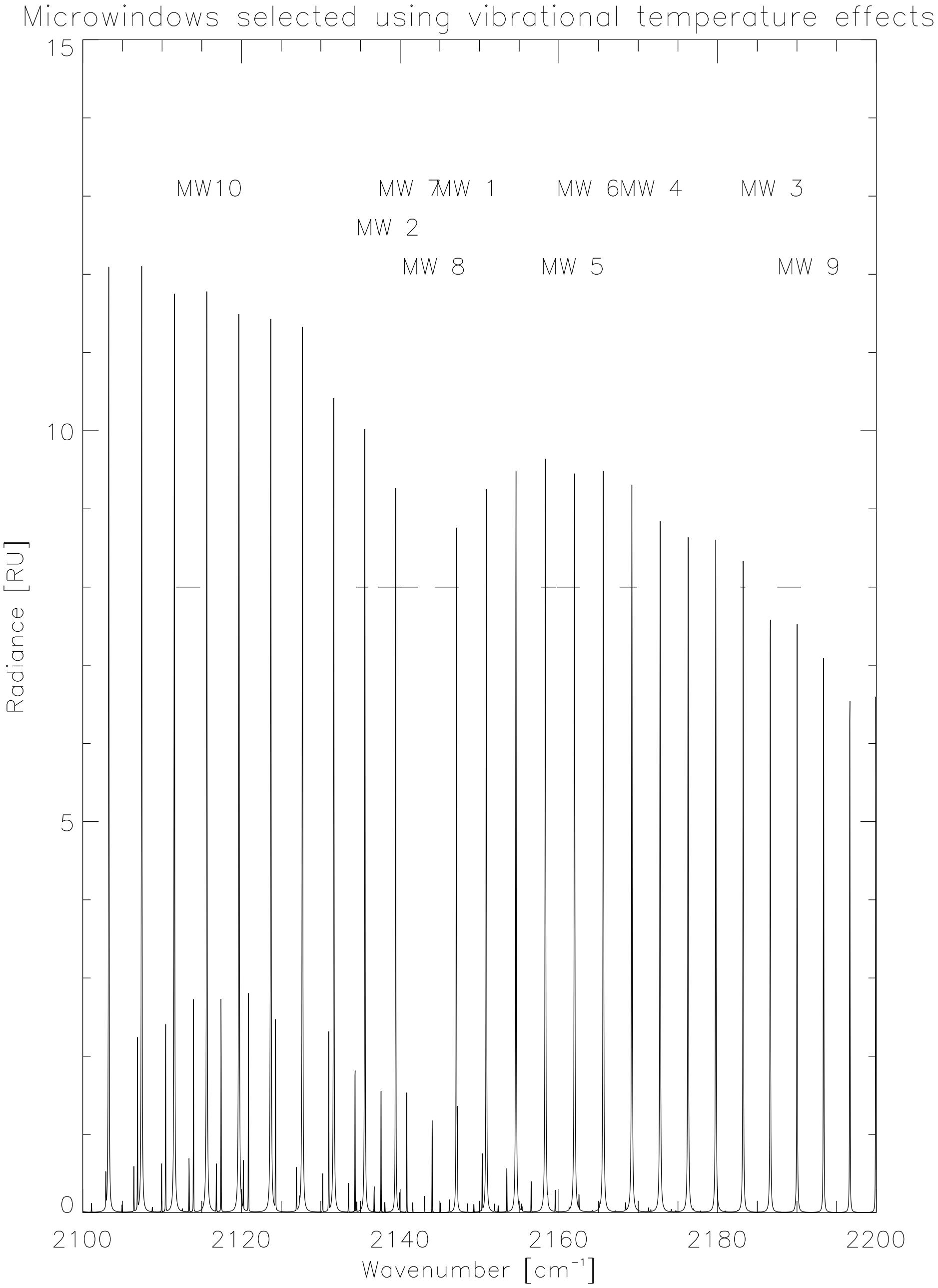

Figure 4: The same microwindows as in table 2 - the labels begin at the same wavenumber as the microwindow they refer to, but are placed at arbitrary heights in hopes of improving readability.

Retrievals using Sequential Estimation - a program developed at the Department of Atmospheric, Oceanic and Planetary Physics in Oxford. It utilizes the Reference Forward Model discussed earlier to calculate the parameters which give the simulated spectra that matches best to the observed spectra.

In order to test the microwindows thoroughly and to provide indications of where they might be improved, each set of microwindows was examined under several different conditions:

Local Thermal Equilibrium Retrievals were performed for each set of microwindows both with and without the simplifying assumption of local thermal equilibrium.

Number of absorbers For each set of microwindows, MORSE was configured to consider only CO once, and once to consider the distorting effects of other absorbers on the spectrum.

Vibrational Temperature Profiles In NLTE retrievals, MORSE uses a predetermined vibrational temperature profile to account for the NLTE effects. Two profiles were used for each test; one typical of mid-latitude daytime, the other typical of polar summer.

Vibrational Temperature Retrieval Each set of microwindows was tested through retrievals where the vibrational temperature is retrieved from the data instead of being read from an existing profile or assumed to be the same as the kinetic temperature.

Each of these conditions was varied separately - so, for example, both sets of microwindows were run with the assumption of local thermal equilibrium for carbon monoxide only and for the full set of absorbers.

In order to judge the accuracy of the retrieval process, it was run against a set of 'blind test' data created for the MIPAS project. This data was by taking a known temperature and carbon monoxide profile and simulating the observations that MIPAS would make of such an atmosphere by adding noise. The test data does take account of NLTE effects, but does not supply a vibrational temperature profile. All windows were tested by comparing the CO profile they retrieved to the known one. To give some indication of a real state of the atmosphere, typical kinetic and vibrational temperature profiles are shown in figure 5.

The investigation considered the significance of each of the factors discussed above and their effect on the quality of the retrieval. Finally, an attempt was made at a retrieval from real MIPAS data. Two profiles were selected at the same latitude, taken from MIPAS orbit 504 - one in daylight, the other at night. Generally, stratospheric gas concentrations are simple functions of latitude, so the two locations should have similar 'true' profiles. However, the one taken in daylight will be contaminated with significant NLTE effects - if the retrieval process produces consistent profiles for these two locations despite the very different NTLE conditions, it will be confirmation that carbon monoxide retrieval is viable.

4.2 LTE-selected microwindows

Figure 6 shows the results obtained using the microwindows selected under the assumption that the atmosphere was throughout in local thermal equilibrium

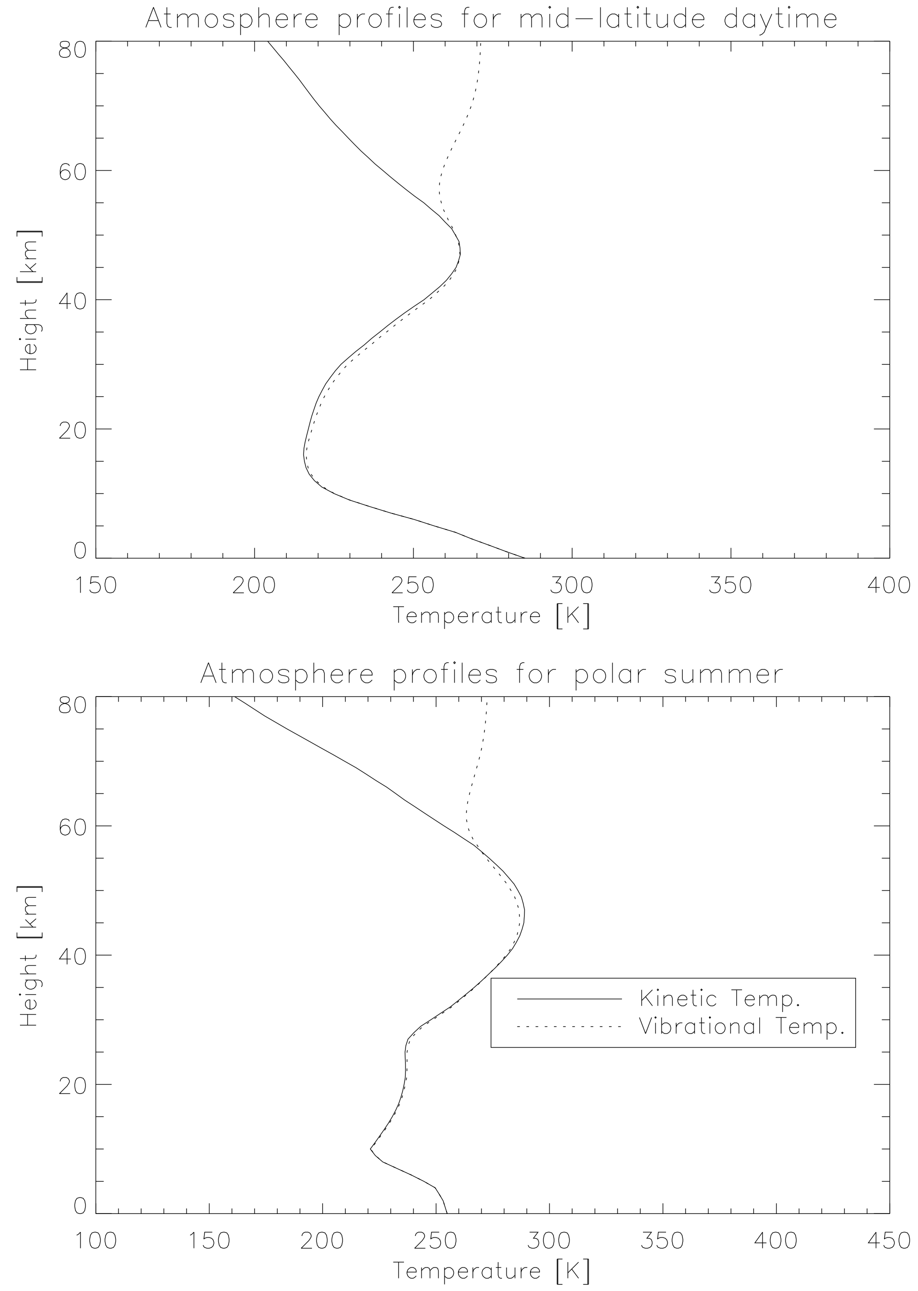

Figure 5: This illustrates two typical profiles of kinetic and CO vibrational temperature against altitude.

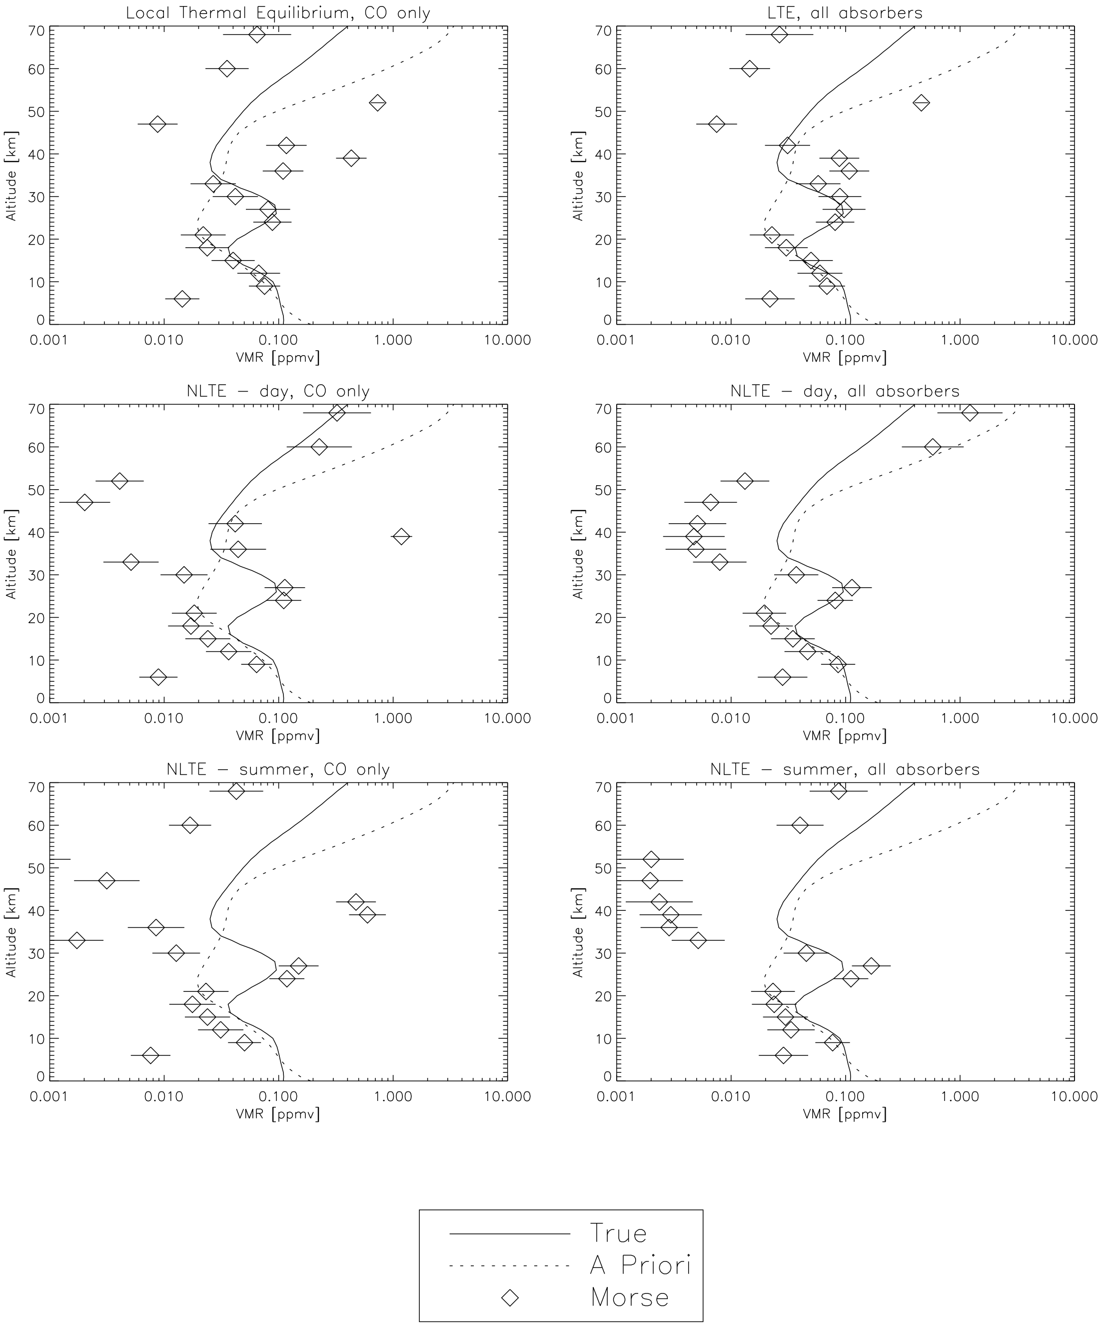
Local Thermal Equilibrium, CO only

LTE, all absorbers

NLTE – day, CO only

NLTE – day, all absorbers

NLTE – summer, CO only

NLTE – summer, all absorbers

Figure 6: This figure shows the *a priori*, true and retrieved profiles of carbon monoxide concentration against altitude for the microwindows selected on the assumption of local thermal equilibrium in the atmosphere.

to make a retrieval from data which includes the NLTE effects. For an ideal retrieval in which the error is primarily due to random noise, the error bars on each retrieved point (which represent the random error only) would overlap with the true profile.

In a less than perfect world, MIPAS has to retrieve its spectra from orbit, well above the atmosphere. This means that there will be some contribution to the spectra from the atmosphere above 68 km, particularly as NLTE effects reach their peak significance in the region of 90 km altitude. This effect will be adjusted into the 68 km retrieval, causing some error at that height but ideally removing the interference from all lower tangent heights.

It may also be difficult to get a strong enough signal from the high altitudes to detect carbon monoxide: obviously, as altitude increases, pressure (and thus gas concentrations) diminish accordingly. In the interests of speed 8-10 microwindows were used in this project - the sharply limited number might mean there is not enough information available to get a meaningful profile above 50 km. This is supported by the scattering - at high altitudes, the MORSE-retrieved points are not all grouped to one side of the true profile, as they are lower down, but rather fall to all sides.

Examining the profiles as a group, it's clear that there is some problem with the retrievals at 6 km. The MORSE retrieval is consistently lower than the actual value. This is probably a combination of several factors in the model, primarily the fact that while the 'true' atmosphere is modelled in 1-km bands down to the surface MORSE only retrieves to 6 km, matching this result to the *a priori* result at 3 km. This means that any divergence of the true atmosphere from the *a priori* in the 6 km down to the surface has to compensated for entirely in the 6 km spectrum. This will clearly introduce an overcompensation error - for this reason, the retrieval at 6 km should be ignored until the underlying causes of the error can be corrected in the model. The necessary corrections were underway at the time of writing, but not yet complete.

The top two graphs in figure 6 assume local thermal equilibrium. The first neglects all absorbers but carbon monoxide, whereas the second includes the (known) effects of ozone, water vapor and carbon dioxide, the other significant contributors to the spectrum in this area. Both show relatively poor agreement with the true atmosphere above 35 km tangent height - the measurements are not consistent with the known profile, nor are they even tracking the shape of the profile. This is a demonstration of how ignoring NLTE effects can be a poor assumption, as the significance these effects increases with increasing altitude. The retrieval accounting for the three additional absorbers shows the better agreement of the two, which is reasonable as removing the influence of the three principal contaminants should leave an almost pure carbon monoxide spectrum for analysis. The carbon monoxide only retrievals were primarily run for the advantage of speed in ironing out bugs in the code, and were not carried forward to the analysis of the NLTE microwindows except in the case of the vibrational temperature retrieval.

The remaining two pairs of graphs in Figure 6 do not assume local thermal equilibrium, and use respectively the mid-latitude daytime and the polar summer vibrational temperature profiles to model the NLTE effects. Up to a tangent height of approximately 35 km, these do not agree with the true profile even so well as the profiles taken assuming LTE. The profiles taken with MORSE only retrieving carbon monoxide are almost random above 35 km, while the profiles taken retrieving CO and the major contaminants track the shape of the true profile reasonably well, even if they are not consistent with the true values. This suggests that a systematic error may have some part to play in the difference between the retrieved profiles and the true one: all four graphs show a tendency to substantially underestimate the carbon monoxide concentration. This is particularly noticeable from 30-60 km in altitude, but is present at nearly all altitudes. This could be explained if the model used in the retrieval over-estimates the magnitude of NLTE effects.

Surveying all of the graphs, it must be concluded that this set of microwindows does not provide enough information to allow us to retrieve carbon monoxide through the entire MIPAS tangent-height range. For tangent heights up to approximately 35 km, the most accurate retrieval of carbon monoxide concentration appears to be that which considers all the major absorbers but does not consider NLTE effects.

This suggests that assuming a fixed vibrational temperature profile for every spectrum is inadequate to model the atmosphere as it is, so much so that ignoring the effects altogether results in a closer fit. Thus the vibrational temperature spectra should either be retrieved (see the next section) or modified for each retrieval.

4.3 NLTE Microwindows

Figure 7 shows the profiles retrieved using the NLTE microwindows, which were selected to be the best for a retrieval of both carbon monoxide concentration and of vibrational temperature. For comparison with the LTE microwindows further retrievals were also carried out; with the assumption of local thermal equilibrium, and using the *a priori* vibrational temperature profiles to model the NLTE effects . Except where specifically noted in the discussion, all retrievals were set to consider the major contaminants - ozone, water vapor and carbon dioxide - in addition to carbon monoxide.

The first chart shows a profile retrieved on the assumption of local thermal equilibrium. This shows some agreement between 20 and 50 km, which takes the agreement higher than previously possible. Below 20 km, however, there is little or no correlation between the retrieved profile and the true one; even above 20 the retrieved profile does not agree with the true profile to one standard deviation. If LTE is to be assumed, one would seem to be better off (unsurprisingly) with using the microwindows selected for that purpose, as in the previous section.

The following two charts repeat the tendency found in the previous section to underestimate levels of carbon monoxide in the atmosphere. This tendency is particularly strong in the profile calculated from the polar summer *a priori* files. There are two possible explanations: first, that the actual vibrational temperature profile is radically different from the assumed profile, or second, that the retrieval model is indeed overestimating NLTE effect. Figure 8 provides some perspective on the vibrational temperature profiles of the *a priori* files as compared to the (hopefully more accurate) retrieved profiles. The retrieved profiles are clearly slightly warmer than the reference files in the region above 30 km, where the greatest underestimates take place. Equally, however, this difference is a matter of a few degrees, and is not sufficient to explain the ten to hundred-fold underestimate. This argues strongly for a problem with the

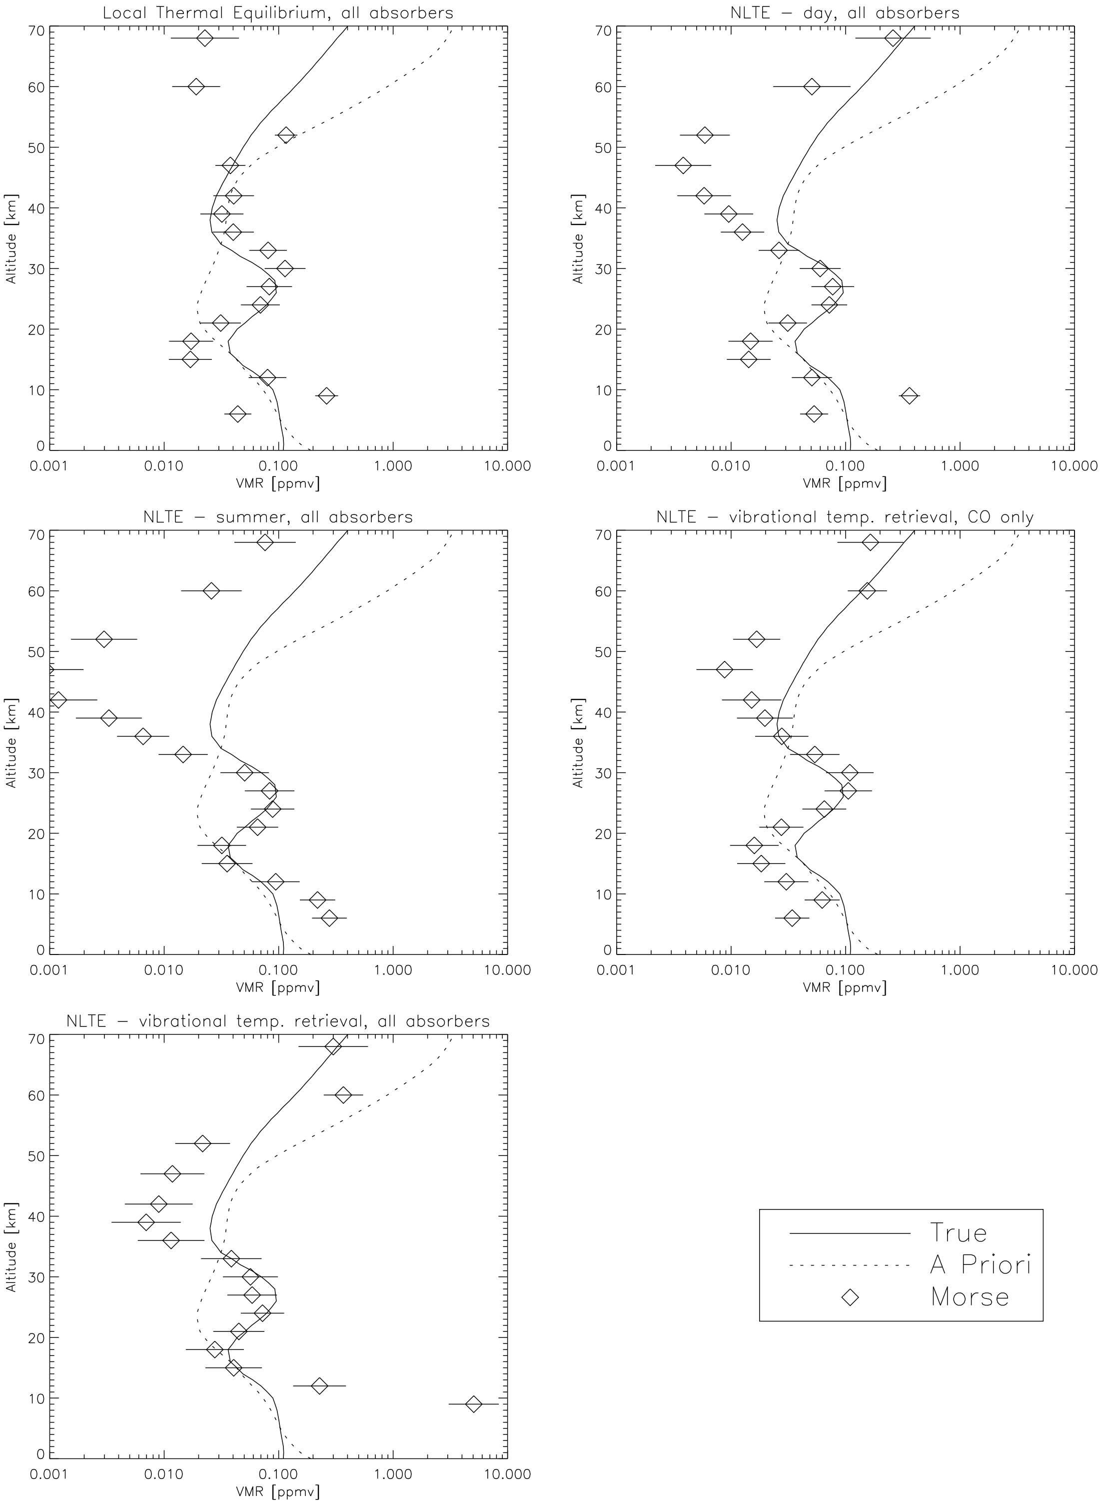

Figure 7: This graph shows the profiles retrieved by MORSE using the NLTE microwindows as compared to the *a priori* data and the true profile used to produce the test data.

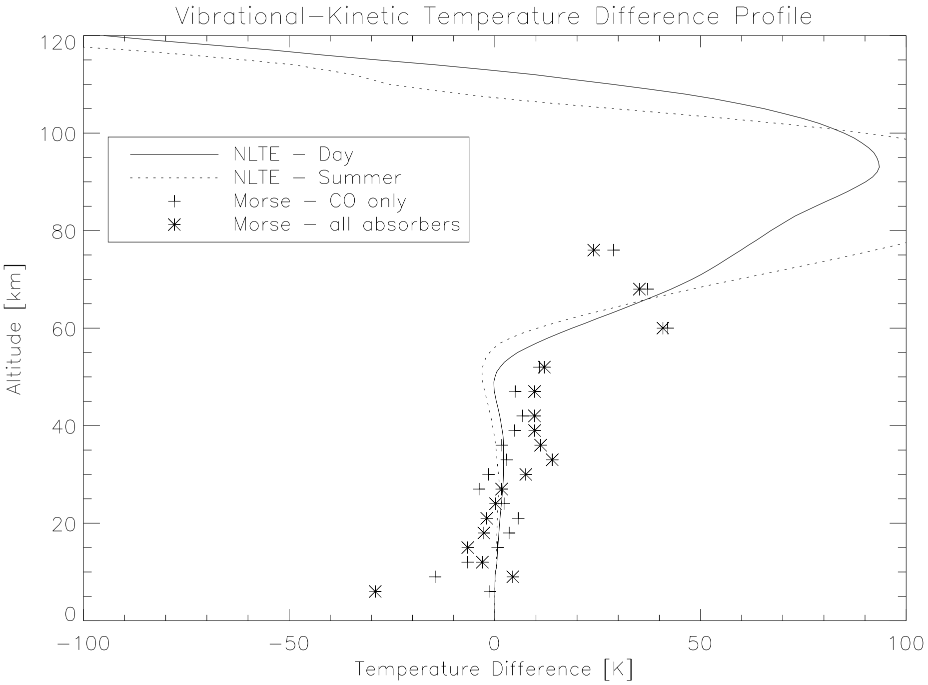

Figure 8: This figure shows the difference between vibrational temperature and kinetic temperature against altitude for both *a priori* data files and for the two retrieved profiles.

retrieval or modelling assumptions.

Pay particular note to the extreme underestimate in the chart for the polar summer profile. Figure 8 shows that the temperature difference between vibrational and kinetic temperature will be most significant in the polar summer profile; the strong underestimate can thus be seen as further evidence that there is some problem with the way the retrieval software handles NLTE effects. Above 60 km the retrieved vibrational temperature is less consistent, due to the boundary compensation errors as the retrieval tries to match up to the *a priori* data again.

The final pair of charts shows the profiles calculated with a simultaneous carbon monoxide and vibrational temperature retrieval. This still shows an underestimate of carbon monoxide at high altitudes, although the retrieved profile tracks the true profile up to about 45 km. Clearly both represent about an order of magnitude improvement over the pre-computed vibrational temperature method. Both still illustrate an underestimate of the carbon monoxide concentration between 35 and 55 km, providing more evidence for an overestimate of the NLTE effects.

There seems to be little to choose between the two simultaneous retrievals. The retrieval with the contaminants included seems to have tracked better with the true profile from 15 to 25 km. The carbon monoxide only retrieval tracked better from 25 to 45 km. Neither is consistent with the true profile at altitudes above 45 km, possibly due to the rarefied atmosphere at that altitude.

There is a clear fluke result in the retrieval which took account of the contaminants - the 6 km result is almost a hundred times greater than it ought to be. Further attention to the factors considered in the retrieval is likely to resolve this anomaly - for example, by modelling the effects of the instrument's field of view. The 9 km retrieval is also high, though more reasonable - this is likely to have been biased high to maintain some continuity with the 6 km result.

The results do not give us enough information to determine which of the two simultaneous retrievals is most suitable to determining a carbon monoxide profile. However, we can conclude that adding the simultaneous vibrational temperature retrieval does result in a significant advantage in the accuracy of the retrieval in comparison to the a priori method. As it is unreasonable not to include the effect of other absorbers in the atmosphere, and their inclusion does not reduce accuracy (which would suggest a problem in the modelling) there is no reason to use the carbon monoxide only retrieval in future.

4.4 Analysing Orbit 504

Figure 9 shows the carbon monoxide profiles retrieved from scans taken as part of MIPAS orbit 504. The first and most important thing to note is that the two profiles are consistent - this is good evidence that our retrieval process works. The scans that were used were taken from two points at the same latitude, with one taken in daylight and the other at night. One of the two will thus be significantly affected by NLTE effects, and the other will not. The fact that the two retrievals produced consistent scans means that the process is seeing 'through' NLTE effects to measure something at least proportional to carbon monoxide concentration.

This is not an unqualified success - both the low and high ends of the profile are two orders of magnitude higher than would be expected, which suggests

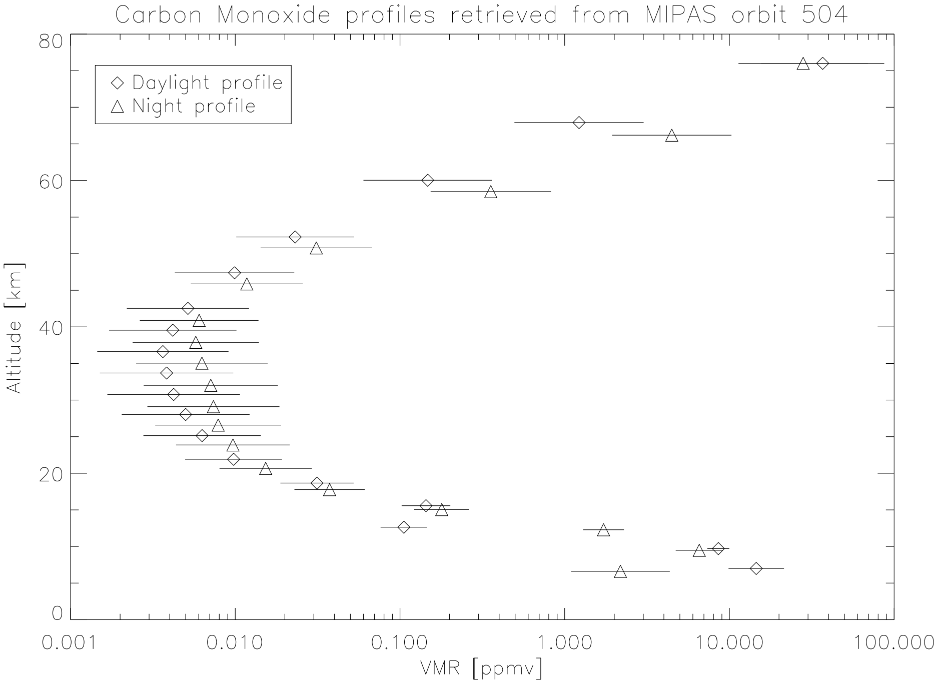

Figure 9: This figure shows the day time and night time carbon monoxide profiles retrieved from day-time and night-time scans of two points on orbit 504. As the points were on the same latitude, they should have approximately the same profile.

that they are flukes or artifacts of the retrieval process. It's also important to note that there is no way to determine how accurate these retrievals are in an absolute sense - systematic error would be the same for both, and thus they would remain consistent. Nevertheless, having self-consistent results is a good first step towards retrieving carbon monoxide using MIPAS. It should also be noted that these are the first attempts to retrieve carbon monoxide by any group working with MIPAS data and no other instruments routinely retrieve it, so intercomparisons are hard to come by.

5 Conclusions

It may be possible to retrieve accurate information on carbon monoxide concentration from the data collected by MIPAS. From a small set of microwindows we have demonstrated that it is possible to obtain profiles consistent up to 45 km and to retrieve consistent profiles from scans with and without NTE effects.

This is not to say that it is possible yet - more study is still necessary before this process can be made accurate enough for research. There are persistent problems in retrieving data below 10 km and above 60 km which need to be addressed, as well as a consistent tendency to underestimate the concentration of carbon monoxide present. Additionally, the error bars are broader than we could wish. In order to refine this technique, we would suggest several improvements.

Firstly, the selection of microwindows should be revised and extended. An increase in the number of microwindows would result in an increase in calculation time, but would also lead to a reduction in random error and an increase in sensitivity to faint signals. This would, it is hoped, improve the accuracy of the technique at high altitudes, where the low pressure leads to a weak signal.

Secondly, more work is necessary on the way the NLTE interactions are modelled in the MORSE/RFM system. Throughout our results there is a clear tendency to underestimate the concentration of carbon monoxide above 30 km. The consistency of this phenomena through profiles taken with different techniques and assumptions strongly suggests that the system itself overestimates the significance of NLTE effects.

Thirdly, better modelling at low altitudes is required - for all sets of microwindows and conditions the retrieved result at 6 km tangent height was a considerably poorer fit than other results. A good first step would be an evaluation of the accuracy of the *a priori* data used by MORSE at low altitudes, as this might be introducing a bias into the retrievals.

Finally, consideration should be given to the possibility of extending the microwindows themselves. As can be seen in Figures 3 and 4 the carbon monoxide lines are very closely spaced. A microwindow which covered two lines would have increased sensitivity to changes in relative intensity of the lines, which could allow a more precise determination of quantities like the vibrational temperature.

References

[1] Dudhia, Jay and Rodgers, *Microwindow Selection for High Spectral Resolution Sounders*, Applied Optics 41 3665

[2] Dudhia, A. *MWMAKE User's Guide* (Webpage 2002) http://www.atm.ox.ac.uk/RFM/mwmake.html

[3] Dudhia, A. *Retrievals and Error Analysis* Unpublished internal document, AOPP 2002

[4] Dudhia, A. *RFM Users' Guide* (Webpage 2002) http://www.atm.ox.ac.uk/RFM/sum.html

[5] ESA, Envisat, *MIPAS Introduction* (Webpage 2002) http://envisat.esa.int/instruments/mipas/index.html

[6] Houghton, John T. *The Physics of Atmospheres* Second Edition, CUP 1986

[7] Houghton, Taylor and Rodgers, *Remote Sounding of Atmospheres*, CUP 1984

[8] O'Brien, O. *Feasibility of Retrieval of Carbon Monoxide (CO) from the Michelson Interferometer for Passive Atmospheric Sounding (MIPAS)* M.Phys. Project, University of Oxford 2002