Calculation and Application of MOPITT Averaging Kernels

Merritt N. Deeter
Atmospheric Chemistry Division
National Center for Atmospheric Research
Boulder, Colorado 80307

July, 2002

I. Introduction

Retrieval results from remote sensing instruments are fundamentally different than measurements from more traditional instruments, including in−situ instrumentation. For example, most passive remote sensing instruments (including MOPITT) are typically incapable of making measurements with the same vertical resolution of most common in−situ instruments. In fact, although both techniques provide some information about the same quantity (trace gas concentration, for example), their sampling characteristics are generally very different. Such issues complicate comparisons of remote sensing results with other types of measurements as well as with model predictions. To make proper comparisons, these differences must be understood and explicitly treated using the averaging kernel formalism.

Broadly speaking, measurement characteristics which must be considered before comparing MOPITT data with other datasets fall into two categories. First, *retrieval resolution* effects define the relative vertical weighting of a particular measurement. For many in−situ techniques, the vertical resolution can usually be well approximated by a delta function. For most passive remote sensing techniques, however, the retrieval is sensitive to some non−uniformly weighted layer of the atmosphere. Second, the maximum likelihood retrieval algorithm used in the MOPITT data processing necessarily incorporates some amount of *a priori information* into the retrieval.[1] The incorporation of a priori information into the maximum likelihood retrieval is necessary because a unique solution (i.e. retrieval) based solely on the satellite−measured radiances generally does not exist. A priori information serves as a necessary constraint to make the retrieval problem soluble.

II. Averaging Kernel Formalism

Both retrieval resolution effects and a priori effects may be expressed mathematically in terms of the retrieval averaging kernel matrix. Neglecting error terms, the MOPITT retrieved CO profile x_{rtv} (specified on the standard MOPITT retrieval grid) can be expressed as a linear combination of the true profile x_{true} and an a priori profile x_a through the matrix relation

$$x_{rtv} \approx x_a + A(x_{true} - x_a)$$
$$\approx Ax_{true} + (I-A)x_a$$

where A is the averaging kernel matrix and I is the identity matrix. For an ideal retrieval, A tends toward I and thus x_{rtv} tends toward x_{true}. Each row of A defines the averaging kernel for one particular level in the retrieved profile. Each element in that row defines the relative weighting of the 'true' mixing ratio value at each level to the retrieved mixing ratio. For MOPITT retrievals, the averaging kernels exhibit significant variability. Generally, all variables which are required as inputs to the forward radiative transfer model have some bearing on A. In particular, though, variability of the atmospheric temperature profile, surface pressure and temperature, and actual CO profile are responsible for most of the variability in the averaging kernels. The observed variability in A indicates that it would not be reasonable to use 'fixed' averaging kernels when analyzing MOPITT data.

III. General Procedure for Calculating MOPITT Averaging Kernels

For each retrieved CO profile, the standard MOPITT Level 2 product includes both the retrieved profile x_{rtv} (specified on the standard MOPITT retrieval grid) and an associated retrieved error covariance matrix C_x. (Data formatting issues relevant to constructing C_x from the Level 2 Product are discussed in Appendix A.) The diagonal elements of this matrix indicate the variance (square of the uncertainty) of the retrieved error for each of the seven levels in the retrieved profile. The off–diagonal elements of this matrix indicate the correlation (or more specifically, the covariance) of the retrieved errors at the two levels corresponding to the row and column of each element. C_x is generally symmetric. The matrix C_x thus provides a complete description of the statistical uncertainties in the retrieved mixing ratio values at each level in the retrieved profile, as well as the expected correlations in the retrieved errors.

The MOPITT maximum likelihood retrieval algorithm requires an a priori covariance matrix C_a which describes the 'background' statistical variability of the CO profile. This matrix is identical in form to C_x. But whereas C_x describes the estimated uncertainties in a particular retrieved profile, C_a generally describes the statistical variability of tropospheric CO based on a large diverse set of observed CO profiles. Alternatively, C_a and C_x may be interpreted as respectively describing the covariance (or uncertainty) of the CO profile at a particular location before and after a remote sensing observation. Typically, C_a is derived from a large set of quality–controlled in–situ profiles representing broadly–varying atmospheric conditions. Currently, all MOPITT retrievals are based on the same C_a matrix. The decision to use the same C_a matrix for all retrievals was made to ensure that temporal and spatial features observed in the retrievals could not be traced back to features in the a priori. In the future, as confidence in the quality of the MOPITT products increases, this issue may be revisited. Both x_a and C_a are included in a downloadable ASCII file, which is documented in Appendix B. For a particular retrieval, the averaging kernel matrix A is readily calculable from the retrieved error covariance matrix C_x and the a priori covariance matrix C_a . The relation connecting these three quantities is simply

$$A = I - C_x C_a^{-1} .$$

IV. The CO Total Column Averaging Kernel

For every MOPITT retrieved CO profile, the Level 2 product also includes the corresponding CO total column. This quantity represents the integrated retrieved CO profile from the surface to the top of the atmosphere, and is expressed in units of molecules/cm^2. As described below, the averaging kernel associated with the CO total column or 'a' may be computed from the profile averaging kernel matrix A.

Mathematically, the CO total column c (a scalar) and retrieved profile x (a vector) are related by the total column operator t through the linear relation

$$c = t^T x$$

where the superscript 'T' denotes the transpose operation and both t and x are column vectors. For each element of x there exists an element of t which defines the 'layer width' associated with that particular level in the profile. Thus, each element of t multiplies a corresponding element of x to produce the CO 'partial column' associated with that level in the discrete profile. Elements of t are expressed in units of molecules/cm^2 or (molecules/cm^2)/ppbv. Once t is determined, the CO total column averaging kernel a is determined by the relation

$$a = t^T A \; .$$

In fact, because of the way c is actually calculated within the MOPITT operational retrieval algorithm, the linear relation shown above between the discretized CO profile x and the corresponding total column c is only an approximation (albeit a good one). In part, this is because the total column includes a small stratospheric contribution (i.e., where p < 150 mb) which is not directly calculable from the retrieved profile x.

Nevertheless, we have found that the relation for c shown above produces total column values which agree with values in the Level 2 product to within several percent. For the purposes of calculating the CO total column averaging kernel, this level of accuracy should be quite adequate.

Calculation of the total column operator t is straightforward. First, layers must be defined for each level in the retrieved profile. See Fig. 1. For simplicity, we define the pressure layer boundaries at the surface, the top of the atmosphere, and at the midpoints between the seven retrieved levels. For example, for a surface pressure of 1010 mb and fixed retrieval levels of 850, 700, 500, 350, 250, and 150 mb, the corresponding pressure layer widths Δp_i would be 80, 155, 175, 175, 125, 100, and 200 mb. Then, assuming the molecular weight of air as 28.97 g/mol (the standard value for 'dry air'), it is shown in Appendix C that

$$t = 2.120 \cdot 10^{13} \; \Delta p$$

where Δp is expressed in units of mb (or hPa) and t in units of (molecules/cm^2)/ppbv.

When used to calculate the total column, we find that this technique yields values which agree with values in the Level 2 product to within three to four percent. There is, however, a persistent bias using the technique just described because of the implicit assumption that the CO volume mixing ratio at the highest retrieved level (150 mb) extends up to the top of the atmosphere. Within the actual retrieval algorithm, the shape of the assumed CO profile at these levels is based on output from a chemistry model and decreases monotonically with increasing altitude. Thus, the 'effective' Δp value for this layer is substantially less than 200 mb. Empirically, we find that the total column bias for V3 retrievals can be virtually eliminated by setting Δp_{150} to 159 mb (and leaving the other Δp values unchanged). (It should be noted that this empirical correction is specific to V3 retrievals.) The resulting RMS total column error in this case (relative to the Level 2 product total column values) is approximately 0.5 percent.
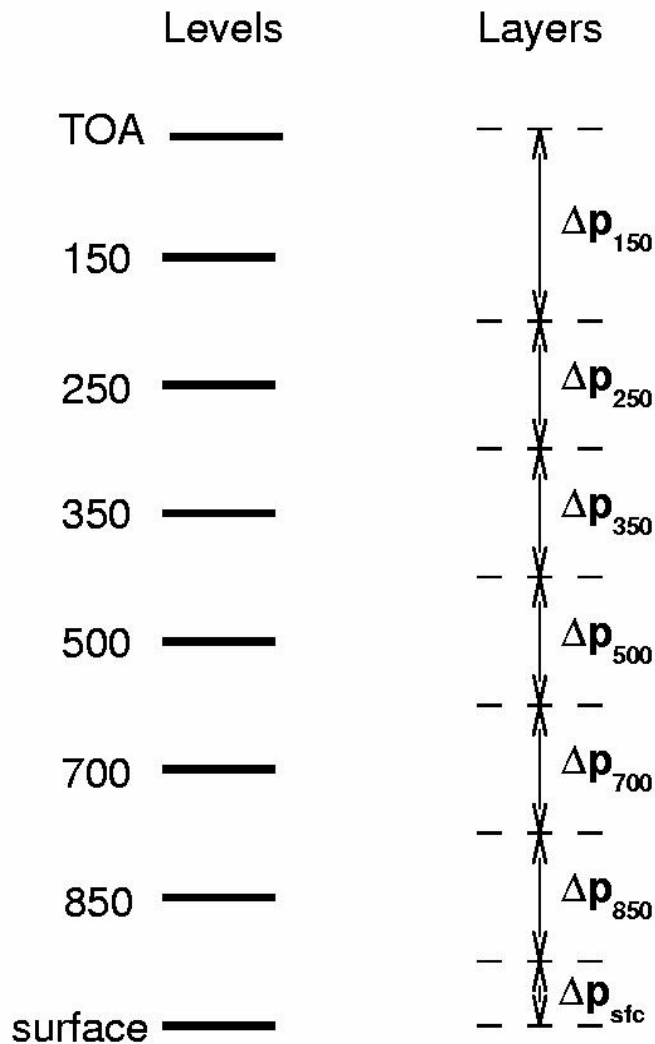

Figure 1. Layer boundaries used to calculate CO total column averaging kernel. Except for surface and TOA, all boundaries are set at midpoints of fixed retrieval grid levels.

As defined above, the CO total column averaging kernel element a_i specifies the sensitivity of the retrieved total column to a change in the true mixing ratio at level i. Alternatively, this quantity can easily be 'normalized' to represent the sensitivity of the retrieved total column to a change of the 'partial column' applied at level i. Elements of the resulting layer–normalized total column averaging kernel a′ are determined from the relation

$$a_i' = a_i / t$$

and are dimensionless.

V. General Procedure for Applying MOPITT Averaging Kernels to Other Datasets

Explicit consideration of the MOPITT averaging kernels is required to properly compare MOPITT data with trace gas profiles obtained from in–situ instrumentation or from chemistry model output. In the following, it is implicitly assumed that the comparative data source does not itself suffer from effects of finite vertical resolution or a priori incorporation. Thus, the following approach should be useful for comparing MOPITT data with in–situ data, or perhaps with chemistry model output, but not with other types of remote sensing data. As described in Section II, the MOPITT retrievals are, to first order, a weighted average of the true profile and the a priori profile. The averaging kernel matrix describes both the vertical resolution of each level in the retrieval and the relative weighting between the true profile and a priori profile. Thus, the averaging kernel matrix provides the information which, if properly applied to a particular trace–gas profile (such as in–situ profile data) transforms that measurement to have the same vertical resolution and a priori dependence as the MOPITT retrievals. Starting with a retrieved profile x_{rtv} and associated retrieved error covariance matrix C_x, the following procedure describes the proper technique for transforming a profile from some comparative data source such that the result may be properly compared against the MOPITT retrieval results.

Resampling. First, the comparative data must be resampled to the MOPITT standard seven–level pressure grid (p_{sfc}, 850, 700, 500, 350, 250, and 150 mb), presumably by some suitable interpolation technique.

Kernel Calculation. Second, the averaging kernel matrix A must be calculated from the retrieved error covariance matrix C_x as described in Section III.

Transformation. Finally, a 'pseudo–retrieval' x'_{comp} corresponding to the resampled comparative profile x_{comp} is calculated according to the linear transformation

$$x'_{comp} \equiv x_a + A(x_{comp} - x_a)$$

The quantity x'_{comp} represents the result of applying a linear transformation to the comparative profile in the same way that the remote sensing retrieval process is believed to transform the true profile. Thus, x'_{comp} can be directly compared against x_{rtv} in a manner that is not affected by varying vertical resolution or a priori dependence.

Similarly, retrievals of CO total column c_{rtv} may be compared against total column values derived from in–situ profiles. For this problem, the 'pseudo–retrieval' of the total column c'_{comp} is given by

$$c'_{comp} \equiv c_a + a(x_{comp} - x_a)$$

where

$$c_a \equiv t^T x_a.$$

V. Special Cases

Missing Levels in Retrieval Grid. The previous section describes the appropriate method for applying MOPITT averaging kernels to other sources of data whenever the surface pressure $p_{sfc} > 850$ mb, which is inaccurate only in regions of elevated topography. For these cases, i.e. where $p_{sfc} < 850$ mb, the MOPITT retrievals will be missing one or more standard levels. (The standard 'fixed' MOPITT pressure levels are 150, 250, 350, 500, 700, and 850 mb.) The technique for calculating and applying the MOPITT averaging kernels for these retrievals is exactly as described for the standard case, with the exception that all rows and columns of all relevant array quantities (i.e. C_a, C_x, A, x_a, and x_{comp}) associated with any missing levels must be eliminated. Thus, for a hypothetical retrieval where $p_{sfc} = 690$ mb, both the 700 and 850 mb levels would be missing in the retrieval grid and A would be calculated and applied as a 5 by 5 matrix (as opposed to a 7 by 7 matrix in the standard case).

Non–invertable C_a. In rare circumstances, the matrix inverse of C_a may be incalculable, preventing calculation of the MOPITT averaging kernel matrix A. The most likely cause of this problem would be duplicated (or linearly dependent) rows and/or columns in C_a. This could occur, for example, if the surface pressure p_{sfc} was very close or equal to some fixed level in the standard pressure grid (e.g. 700 or 850 mb) since, in that situation, the rows and columns of C_a identified with the surface level and the lowest fixed level would be nearly identical. When solving for a matrix inverse, commercial software packages typically provide some diagnostic that indicates the quality of the calculated matrix inverse. It is recommended that this diagnostic be checked after inverting C_a and used to determine whether or not to calculate A.

References

1. Clive D. Rodgers, *Inverse Methods for Atmospheric Sounding, Theory and Practice*, World Scientific, Singapore (2000).

Appendices

A. Formatting of C_x in Level 2 Product

Because C_x is symmetric, inclusion of all 49 elements in the Level 2 Product for each retrieval would require 21 duplicated elements, and would therefore unnecessarily inflate

the output file size. In addition, the diagonal elements of C_x can be calculated directly from the 'retrieval uncertainty' profile, and are therefore not listed explicitly in the Level 2 Product. Thus, C_x is constructed from the Level 2 Product in the following manner.

Diagonal Elements ($C_{x,ii}$). Associated with each retrieved profile x_{rtv} is an uncertainty profile σ_{rtv}. The diagonal elements of C_x are simply calculated by squaring the corresponding elements in σ_{rtv}.

Off–diagonal Elements ($C_{x,ij}$). The off–diagonal elements of the upper–right triangular section of C_x ($C_{x,[i,j]}$, $j > i$) are listed as a 21–element vector in the Level 2 Product in the following sequence: $Cx_{,[1,2]}$, $C_{x,[1,3]}$, ..., $C_{x,[1,7]}$, $C_{x,[2,3]}$, $C_{x,[2,4]}$, ..., $C_{x,[2,7]}$, $C_{x,[3,4]}$, ...,$C_{x,[3,7]}$, ...,$C_{x,[5,7]}$,$C_{x,[6,7]}$. Elements in the lower–left subdiagonal of C_x ($C_{x,[ij]}$, $j < i$) are then directly obtained by the symmetry relation $Cx_{,[ij]} = C_{x,[ji]}$.

B. Description of A Priori Data File

All a priori information used in the MOPITT retrievals is contained in a downloadable ASCII file. MOPITT data users should check the MOPITT website (www.eos.ucar.edu/mopitt) 'Data' page to locate and download the a priori file which corresponds to the particular data release (e.g. V2 or V3) they are using. For maximum flexibility within the retrieval algorithm, profiles and covariance matrix values are specified on a high–resolution 35–level grid, rather than the standard seven–level grid corresponding to the actual MOPITT retrieved CO profiles. With the exception of the lowest level in the retrieval grid (p_{sfc}), the MOPITT retrieval grid is a subset of the high–resolution pressure grid. Construction of C_a on the standard seven–level retrieval grid from the high–resolution version of C_a is a simple matter of identifying and selecting only the matrix elements corresponding to the standard MOPITT seven–level pressure grid. (Note, however, that pressure levels in the high–resolution pressure grid begin at the top of the atmosphere whereas pressure levels in the MOPITT retrieval grid begin at the surface.) Construction of x_a on the standard seven–level retrieval grid from the high–resolution version is performed in the same manner. For both C_a and x_a, array elements associated with the surface level (p_{sfc}) in the retrieval grid should be based on an interpolation of elements in the high–resolution versions of C_a and x_a. For example, for a true surface pressure of 990 mb, one should interpolate using array elements of the high–resolution versions of C_a and x_a at 975 mb and 1000 mb.

A Priori Data Format

(All floating–point values are listed in scientific notation with a maximum of five elements per line.)
Line 3. This line contains the number of levels in the high–resolution pressure grid on which the a priori profiles and CO covariance matrix are specified.
Lines 5–11. These lines list the pressure levels in the high–resolution pressure grid, starting at the top of the atmosphere.
Lines 13 –19. These lines list the CO a priori profile on the pressure grid listed in lines 5–11.
Lines 21–27. These lines list the CH_4 a priori profile on the pressure grid listed in lines

5–11.
Lines 29–203. These lines contain the CO a priori covariance matrix. Both the diagonal and lower–left off–diagonal elements are listed. For each level i in the high–resolution pressure grid, a header row–identifier line is followed by the a priori covariance matrix elements $C_{a,[ij]}$ (j = 1,..,i). Off–diagonal elements in the upper–right triangular section of C_a are found from the symmetry relation $C_{a,[ij]} = C_{a,[ji]}$.

C. Layer Thickness Calculation for Total Column Averaging Kernel

In order to calculate the total column averaging kernel, layers must be defined corresponding to each of the retrieved levels. See Fig. 1. Starting from the layer widths Δp_i defined in the figure, it is possible to determine t, the total column operator. Each element of the column vector t simply converts the dimensionless volume mixing ratio at a given level to a partial column amount (in molecules/cm^2).

The hydrostatic relation describes the derivative of atmospheric pressure p with respect to altitude z. For sufficiently small Δp and Δz, it can be written

$$| \Delta p / \Delta z | = \rho\, g$$

where ρ is density and g is the acceleration of gravity (9.80616×10^2 cm s^{-2}). (The absolute value has been applied to the standard hydrostatic relation to simplify the discussion.) The density ρ can be written as the product of the molar density m (in moles per unit volume) and the molecular weight M (28.97 grams/mole for dry air) such that

$$\rho = m\, M \; .$$

We wish to solve for the total column operator t (in molecules per unit area). For a uniform layer of thickness Δz, t can be written

$$t = m\, N_0\, \Delta z$$

where N_0 is Avogadro's number (6.02297×10^{23} molecs/mole). Eliminating m and Δz from this equation leads to the result that

$$t = \Delta p\, N_0 / (g\, M) = K\, \Delta p \; .$$

As derived, t is the total column operator whereas K is the proportionality constant between t and Δp, and has dimensions of molecules per unit area per unit pressure. In CGS units, $K = 2.120 \times 10^{19}$ (molecules/cm^2)/(dynes/cm^2) $= 2.120 \times 10^{22}$ (molecules/cm^2)/mb. However, as described in Section IV, the total column c is formed by the product of the total column operator t and the CO volume mixing ratio x. Thus, if x is expressed in terms of dimensionless volume mixing ratio, and Δp in units of mb,

$$t_i = 2.120 \times 10^{22} \, \Delta p_i \, .$$

If, however, the CO volume mixing ratio is expressed as parts per billion by volume (ppbv), the relevant formula is

$$t_i = 2.120 \times 10^{13} \, \Delta p_i \, .$$