AMIL2DA Error Analysis

Anu Dudhia & Vivienne Payne
Atmospheric, Oceanic & Planetary Physics
Oxford University, UK

April 6, 2002

1 Introduction

A retrieved profile $\hat{\mathbf{x}}$ can generally be represented as some smoothing of the 'true' profile $\mathbf{x}$ plus error terms:

$$\hat{\mathbf{x}} = \mathbf{A}_x\mathbf{x} + (\mathbf{I} - \mathbf{A}_x)\mathbf{x}_a + \mathbf{G}(\delta\mathbf{y} - \delta\mathbf{f}) \tag{1}$$

where $\mathbf{A}_x = \partial\hat{\mathbf{x}}/\partial\mathbf{x}$ is the *Averaging Kernel Matrix* for the retrieval grid, $\mathbf{x}_a$ is any *a priori* constraint (which may be zero), $\mathbf{G} = \partial\hat{\mathbf{x}}/\partial\mathbf{y}$ is the *Gain Matrix* representing the mapping of the measurements $\mathbf{y}$ into the retrieval, $\delta\mathbf{y}$ is the vector representing the error in each measurement, $\delta\mathbf{f}$ is the vector representing the error in each corresponding forward model calculation.

Rearranging this gives the difference between the retrieval and the 'true' profile $\mathbf{x}$:

$$\hat{\mathbf{x}} - \mathbf{x} = (\mathbf{A}_x - \mathbf{I})(\mathbf{x} - \mathbf{x}_a) + \mathbf{G}(\delta\mathbf{y} - \delta\mathbf{f}) \tag{2}$$

It is convenient to define $\delta\mathbf{y}$ to contain just the random noise component of measurement error, while including correlated or persistent effects such as calibration errors in $\delta\mathbf{f}$ along with the forward model errors. The random noise is routinely propagated through each retrieval and the covariance of $\mathbf{G}\delta\mathbf{y}$ gives the random errors reported with each profile (see under the individual retrieval descriptions).

In this analysis we examine the other terms in the above equations required for a full characterisation of a retrieval: the averaging kernel ($\mathbf{A}_x$) and the systematic error contribution ($\mathbf{G}\delta\mathbf{f}$). Finally, 'REC' analysis is used to attempt to identify the spectral signatures of various components of $\delta\mathbf{f}$ in the retrieval residuals.

2 Averaging Kernels

The averaging kernel for the retrieval grid can be calculated from the gain matrix (dependent on the retrieval scheme) and the Jacobian matrix $\mathbf{K}_x = \partial\mathbf{y}/\partial\mathbf{x}$,

$$\mathbf{A}_x = \mathbf{G}\mathbf{K}_x \tag{3}$$

For an unconstrained least squares fit,

$$\mathbf{G} = \left(\mathbf{K}^T\mathbf{S}_y^{-1}\mathbf{K}\right)^{-1}\mathbf{K}^T\mathbf{S}_y^{-1} \tag{4}$$

where $\mathbf{S}_y$ is the covariance of the measurement noise $\delta\mathbf{y}$. Therefore, in this idealised case, $\mathbf{A}_x = \mathbf{I}$, the identity matrix. In practice such a retrieval usually results in profiles with unacceptably large oscillations so it is necessary to introduce some additional smoothness constraint, either directly by 'regularisation' or indirectly through the introduction of *a priori* information. This has the effect of trading vertical resolution for improved precision, effectively averaging the vertical information in the profile.

However, even if $\mathbf{A}_x = \mathbf{I}$, the real atmospheric profile always contains structure at finer vertical resolution than the discrete levels represented by $\mathbf{x}$. We can represent this by defining a *high resolution* averaging kernel $\mathbf{A}_z$ (no longer a square matrix) representing the mapping of some detailed representation of the 'true' profile $\mathbf{z}$ into the retrieval on a coarser grid

$$\begin{aligned}\hat{\mathbf{x}} &= \mathbf{A}_z\mathbf{z} \\ &= \mathbf{A}_x\mathbf{L}\mathbf{z}\end{aligned} \tag{5}$$

where $\mathbf{L} = \partial\mathbf{x}/\partial\mathbf{z}$ is some type of averaging operator converting the high resolution profile $\mathbf{z}$ to the retrieval grid $\mathbf{x}$.

For these purposes, the high resolution profile will be defined as 1 km resolution, matching the 'true' input profiles for the blind test calculations. The high resolution averaging kernel $\mathbf{A}_z$ then represents the mapping of the true profile into the profile that should be retrieved given a perfect forward model and perfect (noiseless) measurements.

There are three methods for calculating these high resolution averaging kernels:

1. Numerical simulations: perturbing a 1 km level of an input profile, calculating the resulting measurements, passing these through the retrieval code and differencing the results compared to the results of an unperturbed profile. This yields a single column of $\mathbf{A}_z$.

2. Calculating high-resolution Jacobians $\mathbf{K}_z = \partial\mathbf{y}/\partial\mathbf{z}$

$$\mathbf{A}_z = \mathbf{G}\mathbf{K}_z \tag{6}$$

3. Using an estimate of $\mathbf{L}$ (combined with the readily available $\mathbf{A}_x$), which is not as simple an averaging process as it might seem. The retrieval forward model often uses an internal high resolution profile constructed by interpolating the retrieved profile,

$$\hat{\mathbf{z}} = \mathbf{H}\hat{\mathbf{x}} \tag{7}$$

where $\mathbf{H} = \partial\hat{\mathbf{z}}/\partial\hat{\mathbf{x}}$ is a model-specific interpolation (e.g., linear in VMR with altitude). The averaging of the true profile $\mathbf{z}$ to the retrieval grid can be thought of as the inverse of this process. Since $\mathbf{H}$ is not square there is not a unique solution, however if we assume that $\mathbf{A}_x \simeq \mathbf{I}$, the retrieval will converge to a profile $\hat{\mathbf{x}}$ which, when interpolated by $\mathbf{H}$, most closely matches the true profile $\mathbf{z}$. Minimising $(\mathbf{z} - \mathbf{H}\hat{\mathbf{x}})^2$ gives the standard least squares solution for $\mathbf{L}$

$$\mathbf{L} = \left(\mathbf{H}^T\mathbf{H}\right)^{-1}\mathbf{H}^T \tag{8}$$

Using one of the three methods outlined, each group has provided an estimate of the high resolution (1 km) averaging kernel matrix $\mathbf{A}_z$ for each retrieval.

The plots show curves representing both rows ('averaging kernels') and columns ('δ-function response') of $\mathbf{A}_z$. Each curve on the averaging kernel plot, $\partial\hat{x}_i/\partial\mathbf{z}$ shows where the information at each retrieved profile level i originates. Each curve on the δ-function response plot, $\partial\hat{\mathbf{x}}/\partial z_j$, shows how the retrieval responds to a perturbtion of the true atmosphere at level j (solid curves are where the perturbation level corresponds to a retrieval level).

3 Systematic Errors

Systematic errors are those errors which are correlated between measurements or forward model calculations (excluding the local correlations introduced by apodising the random noise). In general, these errors will also be correlated between one profile and the next according to various time/spatial scales.

Such errors are usually difficult to characterise, but some estimates have been made as input to various microwindow selection schemes. We can also use such a scheme to analyse the errors expected from the set of microwindows used in each blind test retrieval.

The following errors have been considered, calculated for mid-latitude day-time conditions:

- Calibration errors

  `GAIN` uncertainty in radiometric gain calibration (± 1 %, assumed uncorrelated between microwindows)

  `ILS` uncertainty in instrument line shape calibration (assumed uncorrelated between microwindows)

  `SHIFT` uncertainty in spectral calibration (± 1 cm^{-1})

- Forward model parameter errors

  `[gas]` for 28 different absorbers, this is either

1. the retrieval 1-σ variability (for previously retrieved gases where a sequential retrieval has been specified); or,
2. the climatological 1-σ variability (for gases listed as microwindow absorbers); or,
3. the error due to omitting the gas (all other cases)

The third case is actually a forward model error rather than a parameter error.

`CTMERR` uncertainty in modelling gaseous (mostly H_2O) continua ($\pm$25 %)

`HITRAN` estimated uncertainties in HITRAN line parameters (assumed uncorrelated between microwindows)

`TEM` effect of propagation of temperature retrieval errors ($\pm$ 1 K, assumed uncorrelated between tangent altitudes)

`PRE` effect of propagation of pressure retrieval errors ($\pm$ 2 %, assumed uncorrelated between tangent altitudes)

- Forward model errors

`NONLTE` effect of assuming local thermodynamic equilibrium applies at all altitudes

`CO2MIX` effect of ignoring CO_2 line-mixing effects

`GRA` effect of assuming horizontally homogeneous atmosphere in line-of-sight direction ($\pm$ 1 K/100 km temperature gradient)

Note that for the blind test simulations, there were no calibration errors and no horizontal gradients. Discrepancies in the HITRAN and gaseous continuum parameters between the simulations and the various retrieval forward models are also probably much smaller than represented by the above. For the BT1 retrievals, there were no temperature or pressure errors.

Various error components of $\delta\mathbf{f}$ are extracted from the appropriate subsets of the error spectra associated with the microwindows used for each retrieval. However, rather than attempt to simulate the contribution function matrices $\mathbf{G}$ characteristic of each group's retrieval, we have assumed a single retrieval model (optimal estimation, with loose *a priori* constraints) which differs only in the number of retrieval levels. Effectively this assumes that all retrieval schemes weight the measurements according to measurement noise in a similar fashion (i.e., inversely proportional to the noise covariance), so that the propagation of systematic errors does not depend too much on the details of the retrieval scheme.

The error analysis plots show the contribution to the retrieval error $\delta\mathbf{x}^j = \mathbf{G}\delta\mathbf{f}^j$ of each systematic error component j. Also plotted, as curves, are the sum of the systematic components, the random error and the total error = random + systematic errors.

Note that SO_2 often features as a major source of uncertainty. This is because the 'climatological' variability includes the significant enhancements that follow major volcanic eruptions. Under normal circumstances, the contribution of SO_2 and its variability is much smaller.

4 Residual and Error Correlation (REC) Analysis

The REC analysis involves fitting the pre-calculated spectral error signatures $\delta\mathbf{f}^j$ to the retrieval residuals, in order to assess whether any systematic error sources contribute more than expected (here, residuals are defined as measured−fitted spectra).

If the residuals $\mathbf{r}$ are a combination of the different systematic error components $\delta\mathbf{f}^j$, each multiplied by a coefficient c_i:

$$\mathbf{r} = \sum_j \delta\mathbf{f}^j c_j = \mathbf{K}_c\mathbf{c}$$

The error coefficients $\mathbf{c}$ can be obtained using:

$$\mathbf{c} = \mathbf{c}_a + \mathbf{G}_c(\mathbf{r} - \mathbf{K}_c\mathbf{c}_a)$$

where $\mathbf{G}_c$ is the optimal estimation gain matrix and $\mathbf{c}_a$ is a vector containing the *a priori* values of the error coefficients for each source.

Each error spectrum $\delta\mathbf{f}^j$ is the effect of a $\pm 1\sigma$ uncertainty, so each error source is given an *a priori* value of 0 and an *a priori* uncertainty of ± 1. Exceptions are the forward model errors `NONLTE` and `CO2MIX` where the *a priori* value is set to 1.

The following parameters are calculated for each error source:

c the error coefficient, expected to be be $\sim \pm 1$ if error assumptions are correct. Values far outside this range indicate a problem. Error sources representing known deficiencies in the forward model (i.e., `NONLTE`, `CO2MIX`) are expected to show an error coefficient of $\sim$+1.

σ the fit uncertainty. Since the *a priori* uncertainty is 1, a value smaller than 1 indicates that there is information about this error source in the microwindows used.

In the REC analyses presented here, systematic error sources for the calibration errors, spectroscopic errors and horizontal gradients are all omitted since it is known that these terms are negligible in the BT simulated spectra and this improves the reliability of the fit of the remaining error signatures. The temperature error signature is also omitted (in practice this is almost identical to gain errors) but pressure is retained, even for the residuals fitted to the BT1 spectra where no pressure error is expected. The water vapour climatological uncertainty is split into a stratospheric component `h2o_st` (altitudes $\geq$ 15 km), and separate tropospheric components for 6, 9 and 12 km altitudes.

For each group, the first table represents the results for the residuals from all the microwindows combined, while alongside the each error analysis plot is a separate table representing the fit to just the points used in the retrieval of that species. Only the best fit components are included, in decreasing order of confidence. The first table gives general information about how the BT test spectra differ from the assumed values of each parameter — values of $|c| > 1$ for any contaminant suggest a concentration outside the climatological 1σ variability, although a large positive value might also imply that the species has been omitted from the absorbers considered in the retrieval. In the analyses for each retrieval, one would normally expect the target species to be accurately fit ($\sigma \to 0$) but with small residual ($c \sim 0$), while other terms are generally also present in the adjacent plot of *expected* sources in the retrieval.