Atmos. Meas. Tech. Discuss.,
doi:10.5194/amt-2016-141-AC2, 2016

Atmospheric Measurement Techniques
Discussions

Interactive comment on "Remote sensing CO_2, CH_4 and CO emissions in a polluted urban environment" *by* Denis M. O'Brien et al.

Denis M. O'Brien et al.

denis.obrien@ggma.com.au

Received and published: 4 July 2016

We thank the reviewer for his constructive and helpful comments. The reviewer's comments are shown below in italics, while our responses are in normal font. Revised text is quoted (indented).

Printer-friendly version

Discussion paper

General Comments

General comment 1

The factor of 8 discrepancy between the posterior (i.e., noise-driven) CO uncertainties in comparison with the actual errors is curious. This is not explained by this paper and is potentially worrisome, especially because it is quite at odds with the results of the precursor geoCARB retrieval paper (Polonsky et al., AMT, 2014), which was closer to a factor of unity (typical errors 3 ppb, vs. 10-15 ppb here). The paper should both state this discrepancy, and attempt to explain the source of the discrepancy between this work and the previous work.

The actual errors are larger than those predicted by the linear error analysis because the former include biases while the latter describe only the random component of the error. The biases for X_{CO} are large because the L2 retrieval algorithm aborted unless an unrealistically tight prior error was provided for X_{CO}. The tight prior error pulled the CO profile towards the prior profile, which generally was too high, thereby producing the bias. Figure 1 attached to this reply shows a box-quantile plot of the true and retrieved values of X_{CO}, both weighted by the averaging kernel. The cyan dots are the individual points of the scatter plot, the red curve shows the median, the dark grey bars indicate the 25% and 75% quantiles, while the light grey bars show the 5% and 95% quantiles. The prior X_{CO} was approximately 175 ppb for all soundings. The plot shows that, during the retrieval, low values of X_{CO} were pulled up towards the prior, while high values were pulled down. As there were many more low values than high, the overall effect is a high bias in retrieved X_{CO}.

We offer the following conjecture as to the cause of the failure of the L2 algorithm. First, the CO spectral lines overlap those of CH_4 and H_2O; indeed detecting the CO requires high spectral resolution to identify CO lines between the CH_4 and H_2O lines.

Printer-friendly version

Discussion paper

Second, in response to an amplitude shift that varies slowly with frequency, such as that produced by aerosol, the retrieval algorithm mistakenly attempts to compensate by adjusting the gas concentrations. However, because the effect of CO on the spectrum is much smaller in this spectral region than that of CH_4 and H_2O, the adjustment made by the algorithm to X_{CO} often is too large, and results in negative X_{CO}. With a tight prior uncertainty, the CO profile cannot deviate far from the prior profile, resulting in fewer unphysical values. If this conjecture is correct the problem can be ameliorated by preconditioning the optimisation so that steps in X_{CO} are smaller. We note also that the masking of CO lines by CH_4 and H_2O lines will be reflected in the linear error analysis. It is reasonable to hope, then, that once numerical problems with the optimisation algorithm are overcome, the error statistics for X_{CO} will improve, perhaps considerably.

In the results section 7.1 titled “Trace gas concentrations”, we have added the following paragraphs after line 17 on page 12.

> The reason why a tight prior uncertainty for the CO profile assists convergence is not clear, but we offer the following conjecture. In response to an amplitude shift that varies slowly with frequency, such as that produced by aerosol, the retrieval algorithm mistakenly attempts to compensate by adjusting the gas concentrations. However, because the effect of CO on the spectrum is much smaller than that of CH_4 and H_2O in geoCARB’s 2323 nm band, the adjustment to CO concentrations made by the algorithm often is too large, and results in negative X_{CO}, which in turn causes the radiative transfer code to fail. With a tight prior uncertainty, the CO profile cannot deviate far from the prior profile, resulting in fewer unphysical values.
>
> If this conjecture is correct the problem can be ameliorated by preconditioning the optimisation so that steps in X_{CO} are smaller. We note also that the masking of CO lines by CH_4 and H_2O lines will be reflected in the linear error analysis. It is reasonable to hope, then, that once numerical problems

with the optimisation algorithm are overcome, the error statistics for X_{CO} will improve, perhaps considerably.

Our earlier paper (Polonsky et al., 2014) showed smaller X_{CO} errors than the present paper.

1. In the present paper both the aerosol loading and the air mass factor are higher. The tight prior uncertainty we placed on X_{CO} introduced a bias. Thus, the errors in the present paper inlude both bias and random components. This we have explained in the text.
2. In the earlier paper the aerosol loading was lower. The simulations assumed a geoCARB-like instrument in polar orbit, so the viewing geometry was more favourable, and the air mass factors were lower. Probably because the aerosol loading was lower, we were able to assign a more realistic prior uncertainty for the CO profile; it was ten times larger than in the present paper. Thus, the errors in the earlier paper were largely free of bias.
3. While trying to diagnose the origin of the apparent discrepancy, we found a configuration file for the earlier paper in which the resolving power in the CO band was set too high. Whether that file was used in the simulations is unclear; the simulations were performed some years ago. Therefore, we are in the process of recovering all the files and re-running the simulations and retrievals for the 2014 paper. If the earlier result proves to be incorrect, we will publish an erratum.
4. The same rigorous scrutiny has been applied to the present paper, and the procedures for providing configuration data have been improved. We are confident that the present calculations are correct.

General comment 2

Because a real flux inversion is not performed, only the statistics of the L2 retrieval errors are used, rather than the actual errors themselves. Thus, the outcome is still likely an optimistic assessment of the flux accuracy from using geoCARB, at least with the present-day trace gas retrieval system. Things could of course improve as improvements to the retrieval algorithm reduce systematic errors in the gas column measurements. This caveat is not mentioned in the discussion, and likely should be. (To be fair, it is mentioned in one sentence in the introduction).

The reviewer's comment is absolutely correct but needs to be understood carefully. We did, indeed, use the error statistics of the retrieval. This means we turned a retrieval bias into a random error, which is certainly a simplification. However, we did *not* ignore the bias as would happen with a simple linear error analysis. Secondly, a real inversion would include a term for a regional bias, which may absorb some of the deficiencies in the retrieval algorithm. Overall we agree that the results are likely to be optimistic, because the calculation of the posterior error assumes perfect meteorology along with other idealisations. To the opening paragraph of the discussion in Section 8, we have added:

> We caution that the estimates of posterior flux errors are likely to be optimistic, because the calculation assumes perfect meteorology along with other idealisations.

In addition, we have replaced the paragraph starting at line 32 on page 14 with the following.

> As with the study of Rayner et al. (2014), X_{CO} measurements provide much information on fluxes. It is important, therefore, to understand their

error characteristics. We noted already that the uncertainties generated by the L2 algorithm were a poor guide to the errors (retrieval minus truth). The differences between experiments with 15 ppb and 2.5 ppb for $\sigma(X_{CO})$ in Table 9 underline the importance of the question; if we could generate measurements with the characteristics implied by the L2 algorithm, the flux constraint would be far stronger. However, the mechanism we believe is causing the larger X_{CO} errors shows that eliminating bias will require substantial improvements to the L2 algorithm. Because there is interference between the CO lines and both CH_4 and H_2O lines in geoCARB's 2323 nm band, at the spectral resolution of geoCARB the sensitivity of the spectrum to changes in CO concentration is relatively small. In response to slowly varying amplitude changes in the spectrum (such as those caused by aerosol or cloud), the L2 algorithm frequently makes large adjustments to the CO concentration, sending it negative and causing the radiative transfer algorithm to fail. Our *ad hoc* remedy for this paper, tightening the prior error on the CO profile, improved the success rate for retrievals, but also added bias. Thus, the posterior uncertainty from the L2 algorithm understated the actual error, because the latter is the sum of both the bias and the random error. Future work should ensure that variables such as optical thickness always assume feasible values.

General comment 3

Maps of the true and retrieved fields of the different gases (in addition to aerosol) would be highly useful to give a sense of the variability of these fields.

We did not provide spatial distributions of the retrieved fields because so few soundings were analysed successfully that it was difficult to see the spatial structures. Figure 2 attached to this reply illustrates the point. It shows true and retrieved X$_{CO_2}$, X$_{CH_4}$ and X$_{CO}$ for observations at 04:15 UTC on 2010-08-05. Although this day was the clearest during the simulation, too few pixels were retrieved successfully to show the structures of the true fields. Nevertheless, the flux inversion, with its knowledge of the meterology, can take advantage of spatially dispersed pixels, as the results of the paper show. Consequently, we have omitted plots of the spatial distributions of retrieved column-averaged gas concentrations.

Specific Comments

Specific comment 1

Section 5.2 — Prior Aerosol: The logic of the adjusted aerosol scheme is not entirely clear. Is the idea that places where the AOD is very low is dominated by background aerosol, and where it is high it is dominated by local-source aerosol? Why were the phase functions of both aerosol types set to that of Kahn 2B?

Over dark targets, like Lake Taihu, the radiance scattered by aerosol and cloud is the dominant component at the top of the atmosphere. The hypotheses underlying the adjusted aerosol scheme are firstly that observations over dark targets can characterise aerosol optical properties, and secondly that those properties are representative of the whole region. The phase functions and single scattering albedos were not modified because they have a secondary effect upon the radiance; more important are the variations of the extinction coefficient with frequency estimated from the dark target observations.

Line 28 on page 9 has been replaced by the following:

> Neither the single scattering albedo nor the phase matrix was changed for two reasons. Firstly, in practice it is unlikely that information to guide such changes would be available. Secondly, the adjustment to the extinction coefficient ensures that aerosol contributions to the radiance are correctly proportioned across the geoCARB bands.

Specific comment 2

Line 16, page 9: Consider replacing the word "inversion" with "retrieval". Conventionally within this field, "inversion" implies a flux inversion.

We have replaced the sentence beginning on line 16 of page 9 with the following:

> After retrieving column-averaged gas concentrations from the simulated spectra with the standard algorithm, the retrievals were repeated with the wavelength dependence of the aerosol extinction coefficients adjusted to represent the regional aerosols more accurately.

However, please let's not get into a linguistic straight-jacket!

Specific comment 3

Section 7.1: The 3 retrieval tests are never formally described. Please do this. A table may help. For instance, it was not clear to this reviewer what "cloud disabled" meant. Does that mean disabled in the simulation, in the retrieval, or both? And what is the motivation for disabling the cloud?

Three experiments were conducted, as indicated in the first three columns of Table 7.

In experiment 1, the smoggy aerosol over Shanghai was simulated, but the optical thicknesses of water and ice clouds were set to zero. The retrieval algorithm used the ACOS aerosol scheme with water cloud, ice cloud and two types of smoggy aerosol. The retrieval algorithm did not assume that the atmosphere was cloud free; it retrieved optical thickness of smoggy aerosol, cloud water and cloud ice. Experiment 2 was similar, except that the radiance simulations of the forward model included the contributions of cloud water and cloud ice. Lastly, the forward simulations for experiment 3 included both aerosol and cloud, but the retrieval algorithm adjusted the optical properties for aerosol as described in the text.

No preliminary screening for cloud was performed. Instead, pixels with poor spectral fits or unreasonable returned values for the state vector were discarded, and pixels with cloud were assumed to be in this category. Running the retrieval algorithm on every pixel was expedient, though costly in computer time.

Clouds were disabled simply to assess the performance of the retrieval algorithm when faced with the high aerosol loading of Shanghai.

Specific comment 4

Section 7.2: Error vs. SNR section: Should make it clear that you could drive the OSSE with a fit to the posterior uncertainties, or the actual "rtrieval – truth" uncertainty. Which is done, and why? Please make it clear in this section what the goal is, and which approach (or both?) is taken.

Because we recognised that the "retrieved minus truth" errors involved bias, which for X_{CO} was large, the subsequent calculation of the amounts by which observations reduced the prior uncertainty of the fluxes did not use the errors shown in Section 7.2. Instead, nominal values for $\sigma(X_{CO})$ were chosen at 2.5 ppb, 10 ppb and 15 ppb, while the posterior errors for X$_{CO_2}$ and X$_{CH_4}$ were taken from the L2 retrieval algorithm. This was explained in Section 6. Table 9 lists results for the three values of σ(X$_{CO}$).

What then is the purpose of Section 7.2? OSSEs require estimates of the bias and random errors as functions of the signal-to-noise ratio, and various estimates have circulated within the geoCARB team. We thought it might be useful to record the actual errors that emerged from the retrievals. We have added the following to the last paragraph of Section 7.2.

> . . . factor of 8, although most of the X$_{CO}$ difference is due to bias rather than a serious under-estimation of the random error by the L2 algorithm.

Specific comment 5

Section 8: Regarding CO problems and the water vapor prior constraint. If you are using an ACOS-style retrieval, this means that you retrieve an essentially unconstrained scale factor to the prior h2o profile. Your hypothesis implies that the vertical profile of h2o really matters, not just the scale factor (which should be almost completely unconstrained by the prior). This seems strange since most of the water vapor lives in the bottom few km of the atmosphere, where moving h2o around shouldn't have much of an effect. Questions/comments about the CO error discrepancy:

- *What were the values of the correlation coefficients in the posterior covariance matrix between CO and H2O, and CO and CH4? Ie, were they highly correlated, such that the algorithm knows that uncertainty in H2O (for example) is truly leading to high uncertainty in CO?*
- *It would be useful to constrain CO completely in the retrieval, and retrieve everything else, since it likely has only a modest effect on the other retrieved parameters of the algorithm (aerosol, surface pressure, co2 profile, etc). Then, a*

> *secondary 2.3- micron only retrieval could be run, fixing all other variables to the retrieved value in the first step. If this is done, the errors against the truth may get better. This would help identify the source of the error discrepancy. If this is simple enough to do, the authors may consider it in the revision phase of this work. An even simpler approach would be to remove all the clouds, aerosols, h2o, and CH4 from the true profiles, and re-run the simulations and retrievals (the latter not attempting to retrieve aerosols, clouds, h2o, or CH4). In this highly simplified case, the theoretical and true errors should match. If they don't, it implies some kind of nonlinearity or even bug in the retrieval.*

The reasons underlying the "CO problem" are slightly different from those assumed by the reviewer, as explained in our response to general comment #1.

Addressing these issues is not easy. In this paper we deliberately chose to retrieve gas concentrations using all four geoCARB bands simultaneously. This we believed was the path least susceptible to error; in principle we only had to add CH_4 and CO to the gas list, calculate spectra in geoCARB's 2323 nm band, and run the retrieval algorithm with four channels. Splitting the retrieval into two stages, as suggested by the reviewer, might work well, but it requires more complex changes to the retrieval code with the attendant risk of error.

While the reviewer's suggestion is certainly a good one, we believe that acting on it for this manuscript will lengthen and delay an already long and late paper. This work is but a first step towards flux retrievals from simulated spectra over polluted environments, a necessary step in testing the accuracy of flux retrievals, and we believe that the results as listed in the conclusions are valuable as they stand.

Interactive comment on Atmos. Meas. Tech. Discuss., doi:10.5194/amt-2016-141, 2016.

Printer-friendly version

Discussion paper

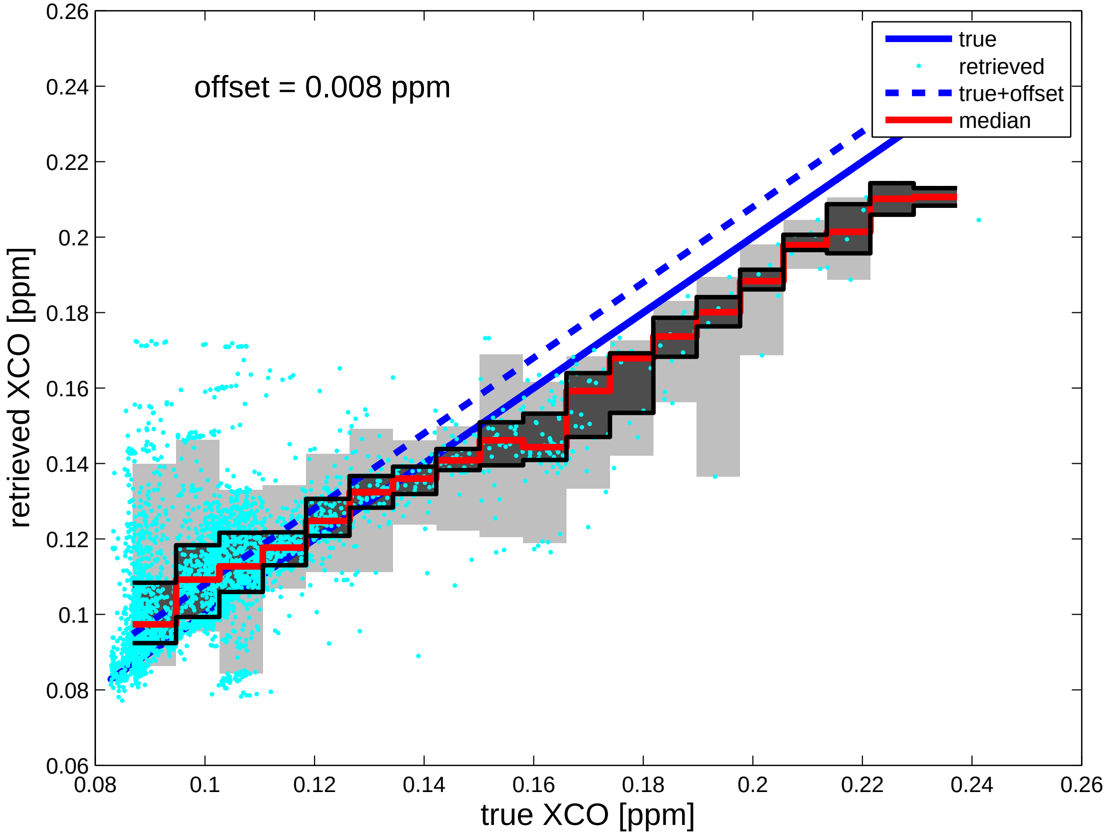

Fig. 1. Scatter plot of retrieved XCO, weighted by the averaging kernel, against the true value.

Printer-friendly version

Discussion paper

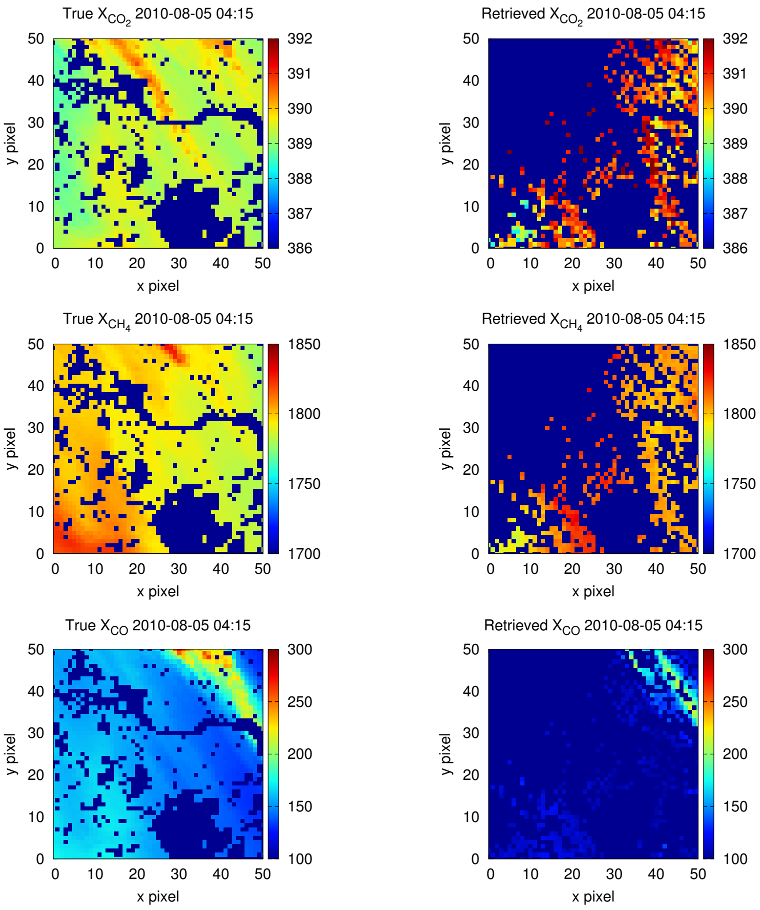

Fig. 2. Comparisons of XCO2, XCH4 and XCO weighted by the averaging kernel (left) with the retrieved fields (right) at 04:15∼UTC on 2010-08-05.

Printer-friendly version

Discussion paper