acp-2013-585 – author response to reviewer 2

This article describes a bromoform emission inversion study carried out using observations from two measurement sites in North Eastern Borneo. The material is definitely within the scope of ACP and the SHIVA special issue. The article is of scientific interest, they present several new findings, the conclusions are well supported, and the methodology also represents an advancement within the field of VSLS emissions. To my knowledge, they perform the first reported emission inversion for VSLS emissions that uses a mathematical technique to find the optimal emission distribution. The authors comprehensively discuss the limitations and demonstrate their effect where possible. They conclude that single measurement sites can only constrain emissions within a limited geographical region dependent on the prevailing meteorology, which has significant implications for other emission inversion studies and for studies that evaluate the emissions. They also clearly demonstrate the how the limited domain affects their analysis. I recommend that the article be published in ACP after having dealt with some minor comments.

We thank the reviewer for the positive review, and for the helpful suggestions below.

General comments

I think the authors have done a good job of discussing the different methodological decisions, and have made a good attempt to demonstrate the abilities of their inversion method in section 4. However, there are some further possible methodological choices that should at the very least be mentioned as possibilities. In addition to the use of an a priori it is possible to use an a priori constraint matrix in the cost function to weight the perceived quality of the a priori. The authors could have for instance used an a priori but given it a very low weighting based on the perceived poor quality of the current emission inventories and bottom-up knowledge. It would be good to see this possibility discussed. The authors make a good attempt to explore the effects of observational uncertainties. However, another methodological approach is to include an observational constraint matrix in the cost function that acts to weight the observations in the cost function. Combined, the a priori and observational constraint matrices act to smooth the topography of the cost function thus acting to limit the number of local minima and therefore to reduce the number of local minima. This has the effect of moderating the inversion solution in order to remove its sensitivity to observational noise and to force a solution at the global minimum. This should be discussed a little bit. In fact, the authors, performed a series of different inversions with different realisations of noise on the observations, so in some way they can demonstrate the solution insensitivity to observational noise. However, I could not find a mention of the resulting standard deviation across all 25 solutions that would go some way to demonstrating that the solution is stable and insensitive to noise. If this standard deviation was added with the corresponding discussion it would strengthen the conclusions of the paper.

The reviewer is correct, it is possible to include an a priori constraint as part of the cost function, and we have in fact tried tests of this type. The resulting solution is ultimately a blend of the a priori, and the solution obtained without the a priori. The mix depends on the weight(s) given to the a priori, which (as the reviewer notes) is/are based on a perception. For a range of experiments in which the weight of the a priori in the cost function is increased, it is not clear which experiment (if any) is best representation of reality. We concluded that this was an extra variable, in an already complicated system containing many sensitivities, that we could reasonably exclude.

The reviewer also mentions the possibility of varying the weight in the cost function associated with each individual observation. This is a useful approach as it allows the modelling and observational uncertainty to be assessed point by point. This functionality is now available in the latest version of the inversion code, InTEM, and so future work will allow this kind of flexibility. However, to subjectively assess the modelling uncertainty (how well does the model capture the current situation) for each time is challenging and subjective and would require further significant and careful thought.

In relation to the response above, the reviewer notes that we repeated each inversion experiment 25 times, with different realisations of noise. Their question relates to the variation within those 25 solutions (i.e. the stability to noise). We addressed a very similar point made by reviewer 1 at length, and hope that referring the reader to that response, where we suggest new text and a new figure, will be sufficient here.

Specific comments

Page 4, lines 5-10. It might help to describe the range of global emission estimates for $CHBr_3$. It might also help to show the range of estimates for the tropical region. The tropics are defined differently in many of the studies, but it would give the reader an idea of the uncertainties.

This would indeed be helpful, so we intend to insert some text summarising recent global (and tropical) emission estimates at line 4 on page 20466.

Page 5, lines 15-17. It might be an idea to mention the SHIVA cruise and data. I realise you have been working on this prior to the release of the SHIVA data and it is not practical to re-do the analysis with the SHIVA data, but it should be mentioned that it exists. You can then mention that those observations were not obtained during your analysis period.

We suggest rewording the final parts of this paragraph: “Our $CHBr_3$ observations are a major enhancement to the quantity of information available in the Maritime Continent, where until recently measurements have been collected only during occasional cruises through the region (e.g. Quack and Suess, 1999; Yokouchi et al., 1999). Following the period we consider in this study, further $CHBr_3$ data have been collected during a cruise through the South China Sea (Nadzir et al., in prep.) and near Borneo during the approximately month-long SHIVA campaign (e.g. the aircraft observations presented by Hossaini et al., 2013)."

Page 10 first paragraph. It might be clearer to readers if you refer to the modelled concentration as a simulated concentration at the measurement site. Also, which measurement site is this test for?

We are happy to use the term 'simulated concentration at the measurement sites'. The calculation actually uses a combined dilution matrix for both measurement sites, which we will make clear in this paragraph and in the caption for figure 4.

Page 10, line 13. Perhaps change "we will focus on" to "we will only solve for". I assume that this is what you are trying to say. Otherwise it is not clear if you are only showing results from the finest grids, or if you are only solving for the two finest grids.

We in fact meant 'focus on' to be read as 'show results for'. We intend to rephrase this sentence: "While the inversion process always returns estimated emissions for the entire solution domain, throughout the rest of the paper we will present only emissions in the smaller grid cells, of size 1 by 1 and 2 by 2 (coloured in red and orange in Fig. 3)."

Page 10, 2nd paragraph. In the inversions where you do not solve for the coarsest grids, i.e. >4 0 4, what do you use to simulate the contribution to the observed concentrations at the measurement sites from those coarser grids? How does this interact with the background that you choose for air masses older than 12 days described in the next section? These points could perhaps be made a bit clearer.

We hope that our previous suggestion will clear this issue up. The coarser grids are always solved for, but not presented, because as we note on line 18, they have much greater uncertainty due to their small impact on the observations.

Section 4. I really like the inclusion of this analysis, and I think it strengthens the conclusions in the paper. I might also be interesting to see what happens to the outcome of this analysis if the observations are degraded by noise. I realise that you perform some sensitivity tests on this in the real inversion, but it would be useful to see the point at which the inversion breaks down. The authors could even just discuss these tests in the text rather than adding another figure.

While we agree this is an interesting scenario, we do not feel that including a digression here would help the flow of the paper. We have repeated this pseudo observation experiment, but in each of the 25 individual inversions added noise of the same magnitude used in the real observation experiments (so a perfect solution no longer exists as the target 'observed' timeseries is now 'noisy', and not the same as the timeseries derived from the 'known' emissions). As we might expect the variability within the 25 solutions is increased, and the mean cost function score is higher (though still excellent, at ~0.01). It is clear we would need to add a significantly greater level of noise to find solutions that bear little resemblance to the ‘known’ p-TOMCAT emissions. An additional complication is that, for a range of experiments in which the magnitude of the noise added to the pseudo observations (derived from the p-TOMCAT emissions) is progressively increased, the point at which the solution is deemed to have ‘broken down’ would need to be chosen arbitrarily.

Section 5. In addition to table 2, it might be useful to add another table which summarises the setup used in each experiment (A through F). Further, the idea that there are specific inversion scenarios is first mentioned in this section. It would be better to have a summary at the end of section 3.2 that describes each of the scenarios relative to the uncertainties discussed. This would also remove the need to have as many methodological details introduced for the first time in section 5, which seems to be more about the results.

We are happy to follow the reviewers’ suggestion. We will add three columns to table 2, which reflect the differences between the six experimental set-ups: observations used; baseline mixing ratio; and use of ‘no land emissions’ a priori constraint. We will also move, and slightly re-word, the bulk of the summary of these experimental details from the first paragraph in section 5 to a new section 3.2.5.

Section 5.2, 2nd para. Please can you add a more precise explanation of how the emissions were scaled? For instance, do you only scale according to the difference in ocean area in the tropical band 20S-20N? Or do you scale for the total area, land and sea?

We suggest re-writing as follows: “In our first extrapolation, we simply assume that emissions are uniform across the entire tropics, and extrapolate the “fine” emissions from the mean solution in experiment D using the ratio of surface areas. We include land in this calculation, but would obtain a similar answer using only the oceanic areas because, according to our land/sea mask, land comprises a similar percentage of these two domains: 27% of the entire surface area between 20S-20N, and 32% of our “fine” grid. This results in a best estimate of a source strength for the global tropics of 225 Gg $CHBr_3$ yr^{-1}.”

Section 5.2. I think the attempt to up-scale the emissions to the entirety of the tropical band is worthwhile. However, I think that there should be some further mention of the limitations of this approach. Assuming that the truly oceanic emissions looked at here are associated with biological productivity, it is possible that the oceanic emissions in the region to the NE and SE of the measurement sites display an unrepresentatively large productivity compared to the open waters of the Pacific, for instance. The underlying driver of these differences are linked to the availability of nutrients in the surface waters (assuming there is a link to plankton). Since the oceanic areas close to Borneo are relatively productive, the up-scaled estimates may represent more of an upper bound. I would urge the authors to examine this issue and try to discuss it. This will therefore have some influence on the final sentence in the abstract.

This is an important issue and we suggest adding a new figure (9) and two accompanying paragraphs towards the end of section 5.2. The figure will present comparisons of ocean depth and chlorophyll a probability density functions for the tropics as a whole and for the 'fine' grid cells. The figure shows that the region we have focussed on contains more shallow ocean, and (as the reviewer suggests) is somewhat more biologically productive (as measured by chlorophyll a) than the tropics as a whole. In the text we will mention that if ocean depth or chlorophyll a were useful proxies for $CHBr_3$ emissions (previous studies have considered these relationships, but we suggest they are probably not – see your next question) this might indicate our estimates contain a slight high bias, and our extrapolation may be more consistent with an upper limit for total tropical emissions.

We will also make minor edits to the final paragraph of section 5.2, the third paragraph of section 6 and the abstract to reflect the above changes.

Section 6. Apart from mentioning sea weeds, is it possible for the authors to include any further discussion of other causes of the derived emission distribution, i.e., are there any correlations to the presence of shallow ocean shelves, areas of upwelling, or areas close to outflowing rivers. One might expect such environments to have higher biological productivity compared to the open ocean.

We have addressed a very similar comment from reviewer 1, and have suggested adding to the paper some text stating that the spatial distribution of our estimated emissions does not compare well with the distribution of two possible proxies for $CHBr_3$ emissions (ocean depth and chlorophyll a).

Technical comments

Page 15, line 11. Two instances of the.

Thanks.