

Interactive comment on “FLEXINVERT: an atmospheric Bayesian inversion framework for determining surface fluxes of trace species using an optimized grid” by R. L. Thompson and A. Stohl

R. L. Thompson and A. Stohl

rona.thompson@nilu.no

Received and published: 21 August 2014

General Comment

I think that the work and method described are good. It could be a very useful tool for people to use in understanding emissions. The key point is that many assumptions are required to be made and the user needs to understand what they are before this tool is used.

We thank Referee 2 for his/her constructive comments and reply to them below.

Specific Comments

C1457

p.3756 l.4: This is only true if the original LPDM run is on a fine enough resolution and large enough domain to meet all needs and also not if applied to gases with different loss processes - unless all particle info retained.

We have now modified this sentence to: “ the LPDM needs only be run once for each species and receptor to find the SRRs, as the output can be applied to optimize the fluxes for any domain and resolution (as long as the resolution is no finer than that of the LPDM run).”

p.3757 Eq2: I do not like the notation Hbg. It has different units to Hnest and Hout.

While it is true the Hbg has different units to Hnest and Hout, we prefer this notation as “H” indicates in general that this is a transport operator.

p.3758 Eq3: time/density has units time x volume/mass not time/volume as stated. What am I missing? Please clarify.

This was an error, it should read: “SRR is in units of residence time times volume per mass”. We have corrected this now.

p.3759 l.7: I assume that ni trajectories terminate in the grid cell? Please re-word to make clearer.

ni is the number of particles that terminate in a given grid cell (or in other words the number of particle trajectories that terminate in a given grid cell). We have changed this sentence to: “The sensitivity to mixing ratio in a grid cell at a given time (i) is calculated as the number of particle trajectories that terminate in the grid cell (ni) divided by the total number of particle trajectories released (J)”.

p.3759 l.22: Only works if site ‘sees’ background air 25 percent of the time. This part ignores the influence of latitude/altitude that is actually known. The next part also assumes that the prior is correct and the site not too polluted. It was unclear to me how the timeseries from lower quartile obs - (prior x SRR) was combined with just the time-series of (prior x SRR)? Are these steps sequential or exclusive?

C1458

Like any method that tries to extract the background from the observations (for a discussion, see Giostra et al., 2011), this assumes that the background signal is actually observed on a semi regular basis at the measurement site. This is, however, not always the case as the reviewer points out. We chose to select the lower quartile of values, which assumes that background is observed at least 25 percent of the time. The percentile of data selected for the background calculation can be varied to best represent the data. This method does not ignore the variation in the background with changing latitude/altitude of origin of the air masses as the change in mixing ratio with latitude/altitude will be reflected in the variability of the time series. A greater problem is that of resolving the variability in the background, e.g. on synoptic scales, which requires in situ measurements, rather than discrete ones, and selecting an appropriate window length over which to select the lower percentile and for the smoothing.

Once observations have been selected as being “background” (in our study we used the lower quartile), then the prior modelled values are calculated for these times. If the observation is truly background, there should be little to no influence from fluxes within the domain, thus the prior modelled value should be small. We subtracted these prior values from the “background” observations (thus the steps are sequential). In some cases, the prior value was higher than what would be expected for a background observation, in these cases, we subtracted a smoothed prior value instead. Although, this does imply that the prior is correct, the error in doing this is likely only small since the prior values themselves are small.

Determining the background reliably is very challenging, and requires assumptions somewhere about the validity of the prior model and what the observations represent.

p.3761 l.2: Helpful to add a comment that =0 as each fine grid can only be in one coarse grid.

We have added this now.

p.3762 l.2: Change 'row' to 'rows'.

C1459

We have corrected this.

p.3762 l.11: I assume that a measurement can have contributions from several latitudinal bands rather than just one as stated here, it is contradicted in the next paragraph, please clarify and re-word.

In the current set-up for the case using the observation-based method for determining the background, we do not disaggregate the background contribution to separate latitudinal bands as there is no information to constrain what the contribution from each latitudinal band is. For the case using the model-based method for determining the background, the contribution is disaggregated over latitudinal band (and optionally also by longitude) so that different latitudes (and longitudes) contribute to the background for each measurement. We have now changed the text to make this clearer.

P.3761-2, Eq 5 and 9: I would suggest the use of the word 'and' rather than commas

We have replaced the commas with “and”.

p.3763 l.8: Please either use 'modelled' or 'modeled' throughout or whatever the journal requires, both are used at the moment.

We have changed this now to “modelling” and “modelled” in accordance with British spelling (except in the bibliography where we use the original spelling).

p.3763 l.8: It does not need to stay within the bounds does it? Are they not expressed as standard deviation uncertainties and therefore have tails? What bounds are being described here?

The solution does not have to stay within any bounds, what we meant was only that the most probable solution depends also on the probability of the prior, which is described by its mean value and standard deviation uncertainties. We have now changed this sentence to: “Based on Bayes’ theorem, the most probable solution for x is the one that minimises the difference between the observed and modelled mixing ratios while also depending on the prior state variables, xb and their uncertainties”

C1460

p.3763 I.21: Out of interest, why if $M > N$ is it easier to invert $M \times M$ matrix? I am obviously missing something but it struck me as odd.

We apologize, the sentence should actually read: "if the number of observations is smaller than the number of unknowns". We have corrected this now.

p.3763 I.24: 'The' rather than 'This'

We have changed this.

p.3764 I.20: What happens if the uncertainty of the prior flux is provided? Is this method still imposed?

We have described the code as it currently is. The prior flux uncertainty may be provided instead, which would involve only a minor change to the code, basically to read this information in and replace the default uncertainties with those provided. However, unfortunately, most emission inventories nowadays still do not provide proper uncertainty values that could be used.

p.3766 I.8: I suspect that the errors in the meteorology e.g. Boundary Layer etc, will be much larger than stochastic uncertainty. So uncertainty will be under-estimated. Obviously not much can be done but a sentence describing this would be helpful.

Yes, we certainly agree. The stochastic uncertainty is likely to be small relative to the transport errors (from PBL height uncertainties etc). We have included the following sentence pointing this out: "Therefore, we do not quantify the full transport error, but only the part of it that can be estimated from the model FLEXPART, i.e. the stochastic uncertainty, which arises by the representation of transport with a limited number of particles (see Stohl et al. (2005)). The stochastic error, however, is likely to be much smaller than the full transport error."

p.3766 I.20: Is the 'minimum error' user defined, please state.

Yes, the minimum error is user-defined. We have now made this clearer in the text.

C1461

p.3767 Sect 2.8: Not sure this is a good idea to retain within the paper. I would suggest just removing this section, it seems rather unfinished work.

We possibly over-stated the uncertainty about this part of the code. We have tested this functionality in a limited number of cases and found that it works. However, we added the caution message, namely that the output should be evaluated carefully, because we cannot be sure of how stable this is in all cases. We have removed "still experimental".

p.3768 Sect 2.9: 'error-free observations' - this needs more discussion. What does it mean for the results? What assumptions are being made? Are the uncertainties being affected?

We have added the following sentences to section 2.9: "The inequality constraint does not only affect the grid cells with negative values but there is also some adjustment to other cells according to the correlations described by the posterior error covariance matrix, Aflux. The posterior error covariance matrix, however, is unchanged since the observation error covariance matrix in this case is zero."

p.3769 I.5: Why is the reaction with OH a good reason to choose CH4? Was it because it has a linear loss process?

We meant rather that because it is principally lost by OH reaction, which is approximately a linear process, that it is a suitable species to use for the test. We have reworded this as follows: "Methane was chosen, as it is an important greenhouse gas with an atmospheric lifetime of approximately 10 years (Denman et al. 2007) and since its loss in the troposphere is principally by reaction with the OH radical, which can be approximated as a linear process."

p.3769 I.12: Suggest replacing 'and largely wetlands and' with 'principally wetlands'

Done.

p.3769 I.25: Suggest replacing 'quite' with 'very'

C1462

Done.

p.3768 Sect 3.1.1: No mention of release height for mountain stations, agl or asl and a sentence to describe the issue here.

We released the particles from the approximate height of the sampling inlet i.e. the station height above sea level plus the height of the inlet above ground level. For mountain stations, we used the given altitudes. We have added the following sentence: "Particles were released from the sampling inlet height at each observation site (see Table 3)" Also, we specify in Table 3 that this is the height in meters above sea level.

p.3769: It would appear that LMP and CIB have the same impact as in-situ sites, were the different number of observations per site not used in calculating Fig 2a? I suggest this is important.

Fig. 2a shows the footprint calculated using FLEXPART assuming the same number of observations at each site. This is the information that is used to determine the variable grid, i.e. the different number of observations at each site is not taken into account in this step. We have now added this information to the caption of Fig. 2a.

p.3770 I.28: Add except high altitude sites to be consistent with p.3774 I.3.

We have added this.

p.3770 I.28: Boundary layer averages?

Yes.

p.3770 I.20: 'closest available' - I assume you mean 3-hr map that encompasses the observation?

We mean the SRR (or "footprint") corresponding to the closest 3-hourly particle ensemble release (we use the terminology "retro-plume" to distinguish from single trajectories). We have now tried to make this clearer in the text.

C1463

p.3770 I.22: It would appear that the in-situ obs are given greater uncertainty as they also have a representation error (variability error maybe better term). In fact the opposite is true, the flask data is only one point so nothing is known about the variability. This therefore should be higher than the in-situ data.

Our reasoning behind using the SD of the in-situ observations over the averaging interval for the representation error is that these are not repeat measurements of the same thing, therefore, the uncertainty does not decrease with averaging the observations. Instead, we are losing some information in averaging, so in this respect, it may also be thought of as a temporal aggregation error. On the other hand, we do not average the discrete (flask) measurements so there is no associated aggregation error.

p.3771 I.17: Therefore some observations are used twice - please state this and comment on the possible effect.

It is true that some observations may be used twice, i.e. first in the optimization of TM5, which is used to determine the background, and second, in the optimization of the fluxes on the nested domain. However, new observations were used in the Lagrangian inversion that were not used in the optimization of TM5 and the averaging interval at which they were used differs. We have added a discussion of this problem to section 3.4 (also in response to a comment by Reviewer 1).

p.3771 I.25: It was interesting to note that uncertainty in background was though more certain than measurement error - I would think the opposite?

The 0.2 percent refers to the uncertainty in the scalars of the background, which, if the background is optimized in the inversion, are included in the state vector. For background values of circa 2000 ppb this uncertainty is close to the measurement uncertainty. The concern was that increasing the uncertainty in the background scalars, would allow too many degrees of freedom for adjusting the background and, as our tests showed, can lead to erroneous background values a posteriori.

C1464

p.3772 I.25: Highlights that considerable care is required. Is this an issue with the prior or poor modelling of MHD? Prior has a dominate effect and the uncertainty in the prior is not propagated.

The determination of the background at MHD from the observations is complicated by the fact that the signal from within the domain is very small, i.e. the observations are mostly background. The overestimate of influence of emissions in the domain can be due to both errors in the transport and errors in the prior emissions. It is not possible to disentangle these two effects. The uncertainty in the prior may be propagated into the uncertainty in the background.

p.3773 I.11: Suggest removing 'slightly'

We have removed "slightly".

p.3774 I.25: Link this to discussion on other mountain sites used in the inversion JFJ and CMN etc as mentioned earlier.

JFJ and CMN are, compared to PUY, quite well modelled by FLEXPART and the normalized standard deviation at both these sites is very close to one, while it is close to 2 at PUY. We think, therefore, that PUY is not well modelled owing to both the topography (the station is located on a volcanic cone, which represents a very abrupt change in topography) as well as the fact that there are significant emissions in the prior around the station. A likely explanation is that FLEXPART overestimates the BL height at PUY and thus overestimates the influence of local emissions. We have amended the text accordingly.

p.3775 I.15: Relevant observations e.g. Cabauw used in Bergamaschi work which would have a dominate effect in Benelux.

Yes we agree, but unfortunately we could not get access to the observations from Cabauw. We have now included this fact in the discussion as well.

I could not find some of the references that are listed in the actual paper, probably
C1465

missed some e.g. Etiope, Houweling, Lambert and Sanderson.

Etiope et al. (2008); Lambert and Schmidt (1993) and Houweling et al. (1999) are referenced in Table 4.

Interactive comment on Geosci. Model Dev. Discuss., 7, 3751, 2014.