

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

We thank Referee 3 for his/her constructive comments and reply to them below.

The manuscript describes an inverse emission estimation framework “FLEXINVERT” for an analytical Bayesian inversion tailored for backward Lagrangian transport models like FLEXPART. Although most of the elements are well known and have been described elsewhere, there are two main reasons why I recommend publication of the manuscript: First of all, FLEXINVERT is presented as a comprehensive framework that combines essentially all of the elements required for such a system including a proper consideration of background concentrations (necessary because the backward

C1448

transport simulation accounts only for fluxes during the recent history of an air parcel), the computation of a variable resolution grid reflecting the true sensitivities of the observation network to upstream fluxes, the definition of observation errors and a priori uncertainties (including their spatial and temporal correlations), and the mathematical and numerical framework to solve for the optimized fluxes and their posterior uncertainties. I particularly appreciate the mathematical rigor the individual components are described with. This will greatly help any user of the framework to understand the individual steps.

The second reason why the manuscript deserves publication is that there are innovative elements that (to my knowledge) have not been presented before, at least not in the same way or to the same level of detail: These include the “Aggregation of background mixing ratios” in section 2.3 and the “Optimization of the fluxes to fine resolution” in section 2.8.

The paper is very well written and structured and well to the point. I thus have no specific recommendations regarding structure or content.

Nevertheless, I have one main concern which deserves more attention by the authors, in particular regarding future users that should also know the potential limitations of the framework: What are the computational costs and what are the corresponding limitations? 4DVAR systems have been developed because inverse problems can easily become too large to be solved analytically. FLEXINVERT is based on an analytical Bayesian inversion which involves operations with large matrices (including matrix inverse) that may become computationally expensive in particular in terms of required memory space. This will necessarily limit the applicability to small- to medium-size problems (limited number of observations/measurements sites, limited spatial and/or temporal resolution of the fluxes). Many operations involve sparse matrices but only standard linear algebra methods appear to be applied which will necessarily lead to many unnecessary computations. The observation error covariance matrix was chosen to be a diagonal matrix. Was this choice driven also by computational constraints?

C1449

Although the treatment of background concentration looks appealing at first sight, it may involve considerable computational cost: The matrix Hbg may be very (excessively?) large since it is dimensioned M (number of observations) \times P (number of grid cells of a global model). It would be useful if the authors could provide some information on the memory and computation time requirements e.g. for the CH₄ test case presented in Section 3. I found it difficult to judge which of the individual steps is particularly expensive and may therefore require special attention by a user.

Computational cost: we have now added the computational costs of running FLEXINVERT (memory usage and total computation time) for the test S1 (which is representative of all tests) to section 3.3:

Max. memory usage: 18 GB

Mean memory usage: 6.4 GB

Computation time (min): 1.8 days

FLEXINVERT is still being developed and we will develop a conjugate gradient version as well in order to solve larger problems (larger number of unknown variables and observations). For this, it will also be useful to have the analytical version to compare with as an accurate reference solution, which should be closely reproduced by numerical methods.

Standard linear algebra: in fact the matrices are not that sparse. The sparseness of the transport operator, H , depends on how long backwards in time the trajectories are run and the sparseness of the error covariance matrices depends on the correlation scale length used. Subsequent to submitting this paper, we have implemented the use of eigen-decomposition of the error covariance matrices, Bfluxnw, Bflux, and the spatial error covariance of Bfluxvg to avoid storing these in memory. Also, the full spatio-temporal error covariance matrix, Bfluxvg is not formed directly but rather calculations using this matrix are made using the eigen-values and vectors of the spatial

C1450

error covariance and the temporal error covariance matrix.

Observation error covariance matrix: this matrix, R , is in fact not diagonal since the aggregation error covariance matrix (in the observation space) is not diagonal. However, we did not account for error correlations between observations. As the reviewer states, this is not a problem for the infrequent discrete observations (flasks) but could be a problem for the in-situ observations. In the case study presented, we averaged the in-situ observations (daily day time average) and thus reduce the correlation between assimilated observations. If the user specifies the correlation of observations, however, this could be easily added to FLEXINVERT.

Background concentration matrix Hbg: this matrix is only stored in memory for one month at a time. We did not find the size of Hbg to be limiting in our case study.

FLEXINVERT makes a number of assumptions (which are often made), and a user will have to be aware that these may not necessarily be met:

Errors are assumed to be Gaussian. As shown e.g. by Stohl et al. (2010), measurement – model residuals may be highly skewed. FLEXINVERT may therefore be sensitive to extremes.

In the case study presented, we did not find the measurement-model residuals to be highly skewed. However, we agree that in such cases, i.e., with highly skewed distributions, that the inversion will be sensitive to extremes (as with any inversion framework). We have added the following text to section 2.7 pointing out the need to be careful about how to treat outlying observations and refer to the work of Stohl et al. (2009) (please note that the correct reference is 2009 and not 2010).

“Another assumption that is made is that the observed – modelled mixing ratio residuals have a Gaussian distribution (Eq. 10 is based on this assumption). Therefore, in cases where the distribution is highly skewed, observations corresponding to the tail of the distribution will have a strong influence on the result of the inversion. FLEXIN-

C1451

VERT does not include any component to deal with skewed distributions; however, the influence of observations in the tail of the distribution may be reduced by increasing their uncertainty. For more details about dealing with skewed distributions we refer the reader to Stohl et al. (2009).”

Observation errors are assumed to be uncorrelated in time: This is likely a good assumption for weekly flask samples, but it would certainly not be a good assumption for e.g. hourly data, mainly due to correlated transport errors (not only correlated PBL errors but also errors in the wind field). This assumption can be easily verified by analyzing the autocorrelation structure of the residuals. In this way the correlation length can be determined.

We agree that for e.g. hourly observations, it would be necessary to account for the error correlation between observations and we now point this out in section 2.7. In the case study we present, we use daily daytime averages of the in-situ observations, therefore, the error correlation between assimilated observations is less than for the hourly data.

The observation-based estimation of background concentrations as presented in Section 2.1.2 will not work for CO₂ which has strong negative fluxes and therefore has no clearly defined baseline. It would be good if the authors could add some word of caution on these points.

We certainly agree. The observation-based method for determining the background was included rather for anthropogenic species (e.g. halocarbons) for which there is a definable baseline. We have now added this information to section 3.4.

Finally, results of an inversion critically depend on the specification of a priori and model-data mis-match errors (and their correlation structure). FLEXINVERT provides a nice framework for solving the problem, but it provides little guidance with respect to the specification of these errors. The authors are obviously aware of the necessity to provide realistic error estimates as they have checked their inversion in the case study

C1452

(Sect. 3) for the chi-square statistics but it should probably be stated more clearly that it is the task of the user to define these errors in a realistic way. A recent publication addressing this issue is Berchet et al. (ACP, 13, 7115–7132, 2013, doi:10.5194/acp-13-7115-2013).

We agree that this is a very important consideration. FLEXINVERT includes a simple scheme for estimating the prior error covariance matrix (described in section 2.5). The prior flux errors are calculated as proportional to the magnitude of the flux with some spatial smoothing, i.e. grid cells with small fluxes but which are adjacent to grid cells with large fluxes will also have a large error. The user can then define the spatial and temporal correlation scale lengths for the describing the error correlations. The total error covariance matrix is then scaled to be consistent with some user-defined estimate of the total error for the domain. We considered this simple scheme to be the best default as it can be applied to many different species. However, the user may decide to use another, maybe more sophisticated scheme.

Minor points:

P3753, lines 14-16: Shouldn't it be “the partial derivative of the change in mixing ratio to the change in fluxes” rather than the reverse?

We refer to the adjoint, thus it is the change in flux that is calculated.

P3754, line 5: Other regional scale inverse modeling studies involving Lagrangian models would be: Keller, C. A., M. Hill, M. K. Vollmer, S. Henne, D. Brunner, S. Reimann, S. O'Doherty, J. Arduini, M. Maione, Z. Ferenczi, L. Haszpra, A. J. Manning and T. Peter, European Emissions of Halogenated Greenhouse Gases Inferred from Atmospheric Measurements, Environ. Sci. Technol., 46, 217-225, doi:10.1021/es202453j; Brunner, D., S. Henne, C. A. Keller, S. Reimann, M. K. Vollmer, S. O'Doherty, and M. Maione: An extended Kalman-filter for regional scale inverse emission estimation, Atmos. Chem. Phys., 12, 3455-3478, doi:10.5194/acp-12-3455-2012, 2012.

C1453

We have added these references to the list.

Page 3756, line 15: Variable resolution grids adapted to the average residence times probably have been introduced for the first time in Manning et al. (JGR, doi:10.1029/2002JD002312, 2003) and have also been applied in other studies such as Vollmer et al. (GRL, doi:10.1029/2009GL038659, 2009), Manning et al. (JGR, doi:10.1029/2010JD014763, 2011), etc. This should not be called the “method of Stohl et al.”.

We have changed “method” to “studies” and include the reference to Manning et al. 2003.

Page 3757, lines 15ff: It should be described more clearly that a typical setup of FLEXPART involves an outer (potentially global) domain and a finer, nested domain. The transport of particles is continued in the outer domain once they leave the nested domain. This setup is not always applicable. Consider e.g. a regional scale model such as FLEXPART-WRF where particles may terminate at the borders rather than being transported further. In this case, termination of particles may occur at any time before the end of the simulation. FLEXINVERT does not seem to be prepared for such a case.

The version of FLEXINVERT presented does require that a global domain is used for the FLEXPART runs. To use FLEXPART-WRF, i.e. with a regional domain would require some modifications to account for the fact that the particles terminate at the domain boundary. We have now added this information to section 1:

“FLEXINVERT, as it is presented here, requires that the LPDM is run on a global domain, or at least that the domain is large enough so that trajectories do not exit the domain.”

Page 3760, line 8: Estimating background concentrations from observations is a long-standing problem that has been addressed in numerous studies prior to Stohl et al. (2010) and more sophisticated methods have been developed than presented

C1454

here, e.g. Thoning et al., (JGR 94, 8549–8565, 1989), Ruckstuhl et al. (AMT, doi:10.5194/amt-5-2613-2012, 2012, 2012).

We agree that this is a long-standing problem, and it is not one that we have aimed to address in our paper. Rather, we have adopted a simple routine to calculate the baseline that we found to be fairly robust to the number of observations (i.e. can be used for in-situ as well as discrete measurements). We consider the method of Thoning et al. on the other hand not to be a suitable alternative. Thoning et al. use a Fourier transform to filter high frequency (synoptic and shorter timescales) components of the signal, however, their technique requires gap-filling of the data, which may introduce errors. Furthermore, it may in some cases overestimate the baseline if the baseline is at times lower than the low frequency signal (e.g. seasonal cycle) as it could be filtered out. On the other hand, the method of Ruckstuhl et al., could be an interesting alternative, as their method does not require any gap-filling and could be applied to data with differing sample frequency. We have included the following statement in any case to point out that there are a number of alternative methods that could be used:

“This method was chosen as it is robust to the number of observations (i.e. it can be used for in situ as well as discrete measurements) although other more sophisticated background selection algorithms exist (e.g. Ruckstuhl et al. 2001, Giostra et al., 2011).”

Page 3763, lines 19-21: I don't understand why this is the most efficient method when the number of observations M is large compared to the dimension of the state vector. The matrix to be inverted has dimension $M \times M$, and if M is large this is a large matrix.

This was a mistake, we meant that the number of observation is smaller than the number of unknowns.

Page 3764, line 21: How are the “8 surrounding grid cells” defined in a variable resolution grid?

This is done before the conversion to the variable grid. We have now clarified this in

C1455

the text.

Page 3765, line 2: Why is the dimension of B $P \times P$? P was introduced on page 3759 as the dimension of the global model grid.

This should be $K \times K$. We have now corrected this.

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

C1456