Replies to Anonymous Referee #1 for Manuscript acp-2013-297

Comparison of ensemble Kalman Filter and variational approaches for CO_2 data assimilation

A. Chatterjee and A. M. Michalak

Major Comment

For CO_2 atmospheric inversion, the main advantage of 4D-VAR is the capability to address very large inversion problems, and in particular to process large inversion windows seamlessly, respecting both the physics (mass conservation) and the statistics (the prior/posterior uncertainty about each state variable is rigorously propagated forward) throughout the period of interest.

In this context, it is very surprising that the authors make use of overlapping windows (p. 12850, l. 6, actually without any information given about the window length). The reason given in p. 12848, l. 20 (computational cost associated with calculating the inverse of the prior error covariance matrix), does not hold: the computational cost of the transport model in the 4D-VAR inversion makes the cost of the matrix operations negligible (see, e.g., Le Quéré et al. 2007 or Chevallier et al. 2010 for the application and Yadav and Michalak 2012 for the principles). Note that the existence of temporally- correlated prior errors stated in l. 22 (same page) does not favour splitting the inversion problem into small windows.

For this reason (and, to a smaller extent, for those given in my first two minor comments), the particular 4D-VAR algorithm used here should not be called "state of the art" (p. 12828, l.24), but in that case the paper loses much of its appeal. The algorithm should be upgraded for the final version.

> **Response 1.1:** We thank the reviewer for their feedback and suggestions. A primary target of this study is to guide selection of the DA methods for real applications in the future. Hence, in designing the 1D framework, we intended to mirror a realistic inverse problem where we are able to use a suite of atmospheric CO_2 observations (i.e., satellite and in situ network) to estimate global surface fluxes at a high spatial and temporal resolution. The rationale for estimating at fine-scale fluxes is already discussed in p. 12829, lines 10-18 (also see *Chatterjee et al*. [2012]).
>
> If we are to estimate fluxes globally over even just one year at a spatial resolution of 1° × 1° and at a temporal resolution of daily, for example, then the dimensions of the state vector would be ~2e7. Most existing studies in the literature estimate fluxes at a coarser spatio-temporal resolution. For example, *Chevallier et al*. [2010] solved for weekly fluxes on a 3.75° × 2.5° grid over a period of 21 years, which indicates that the state vector was of dimensions ~7.5e6 over the entire 21 year period. If we expand the daily, 1° × 1° state vector to cover a 21-year estimation period, the state vector will be ~67 times larger than *Chevallier et al*. [2010]. It is reasonable to expect that one shall certainly run into a computational bottleneck when it comes to calculating the inverse or square root of the prior covariance matrix using a single inversion window over the entire period.
>
> In order to tackle this computational problem, the reviewer is correct that matrix

manipulation techniques outlined in *Yadav and Michalak* [2013] may be used. An alternate approach, as outlined in this study, is to break the problem into smaller overlapping (moving) windows. In order to determine the length of the moving window, however, one needs to keep in mind two primary factors: (a) observations typically inform fluxes over a finite period of time (for e.g. 4 to 6 month timeframe based on the analysis reported in *Bruhwiler et al*. [2005] and *Michalak et al*. [2008]), and (b) temporal correlation in the prior flux errors (for e.g., for the prior land flux errors, *Chevallier et al*. [2012] has shown the temporal correlation to be strongly positive for lags <85 days and mildly positive for lags >275 days). If one does not assume a priori temporal correlations in the fluxes that go beyond seasonal scales, then the length of the moving window can be specified to be 4-6 months to account for both the factors above. Given a sufficiently long window, therefore, a 4D-VAR approach using a moving window will emulate a 4D-VAR approach using a single window, while at the same time providing substantial computational savings. We foresee this as a realistic scenario for future global inversions, and this is the setup that the simplified case study presented in this work aims to represent.

For the toy problem used in this study, observations inform fluxes for 5 [T] (p. 12833, lines 18-20) and consequently, the length of each moving window is specified to be 5 [T]. Note that we did carry out the optimization using a single long window as a sensitivity test, and the conclusions regarding the performance of the variational approach were the same as those reported in the original manuscript. We decided to report the moving window approach to make the reader aware that this technique is applicable for a large-scale flux estimation problem as well. We agree with the reviewer, however, that a more thorough discussion of these points is warranted, and will include a description of the definition of a 'large-scale' problem (e.g., daily, 1° × 1°) along with the issues dictating the length of the window that is used.

Finally, we acknowledge the reviewer's concern that the 4D-VAR scheme used in the study has a few simplifications compared to applications outlined in *Chevallier et al*. [2007] or *Baker et al*. [2006]. For the revised manuscript, we are updating the minimization routine to a Lanczos algorithm as suggested by the reviewer in their Minor Comment #2. Preconditioning will be applied for all the 4D-VAR runs. We will also recover posterior uncertainties based on the Monte Carlo approach outlined in *Chevallier et al*. [2007].

In addition, in the revisions we shall revise the discussion on the 4D-VAR setup, and the availability of posterior covariances from 4D-VAR (p. 12839, Section 4, Section 5, Appendix) to allow an unbiased comparison with the EnSRF scheme.

Minor Comment

1. In p. 12840, the text discusses preconditioning for 4D-VAR and restricts it to "very specific assumptions of the correlation structure and sparsity". I was actually not aware of 4D-VAR applications without preconditioning. Further, I would argue that such "specific" assumptions are more relaxed for 4D-VAR, that can access both small and high spatial resolutions in the state vector and both short and long inversion windows (with long temporal correlations), than for EnSRF, that has to be run at coarse resolution and with short windows.

Response 2.1: In p. 12840, lines 17-20, the text discusses preconditioning for a large-scale CO_2 inversion problem and not for general applications of 4D-VAR.

We agree that the ensemble Kalman filters also require specific assumptions similar to 4D-VAR, and this has been acknowledged throughout the manuscript. We do not agree, however, that it is necessary to run EnSRF "*at coarse resolution and with short windows*". Some studies within the CO_2 community did choose to solve at coarse resolutions and with short windows (e.g. *Peters et al*. [2005; 2007]; *Feng et al*. [2009]) to balance the number of ensemble members and the resultant sampling error. But more recent applications have demonstrated that with judicious use of localization and inflation algorithms one may use a small number of ensemble members to solve at scales equivalent to 4D-VAR applications (e.g. *Miyazaki et al*. [2011]; *Chatterjee et al*. [2012]; *Kang et al*. [2012]) as well as use lag windows that are physically valid for the CO_2 flux estimation problem (e.g. *Chatterjee et al*. [2012]).

2. Little is written about the two algorithms themselves. For instance, from p. 12848, l. 10, the reader may guess that the minimiser of the 4D-VAR is L-BFGS. Actually, if this guess is correct, it is not the best choice and would explain why the authors need so many iterations: the Lanczos algorithm (e.g., Desroziers and Berre 2012, and references therein) converges much faster for linear problems.

Response 2.2: The reviewer is correct in their assessment that in the current study the L-BFGS algorithm is used as the optimization method. The Lanczos algorithm (or similar algorithms that work with the orthornormal basis spanning the Krylov space of the Hessian) may certainly converge faster. In p. 12841, lines 26-30, we do acknowledge that the performance of the ensemble filter and the 4D-VAR can be improved upon by further tuning of the algorithms. We thank the reviewer for their suggestion, and in the revisions we shall update the optimization method to use a Lanczos algoritm.

We are aware that we have glossed over several algorithmic details in Appendix A. But the main motivation for this study was to assess the broad science question of whether numerically approximate methods, such as EnSRF and 4D-VAR, can serve as a suitable long-term replacement for batch technique under different inversion conditions. This study is definitely not intended to be a primer on data assimilation methods and hence, the reader is referred to review papers (p.12850, lines 13-14) for a more detailed discussion on these issues. We have only tried to present generic flavors for both the ensemble Kalman filter (skipping over details on the type of filter, localization and inflation algorithms etc.) and the 4D-VAR (skipping over the types of minimization algorithm, preconditioning, weak vs. strong-constraint etc.), so that the reader is not caught up in the algorithmic details and specific choices made for a particular version of the ensemble or the variational approach.

As outlined earlier, in the revisions we shall make clearer that a variety of choices can be made regarding how the DA methods are set up. And although each different filter or optimization setting may impact the results marginally, we do not expect this to substantially change the overall conclusions in Section 5.

3. The text repeatedly considers the absence of second order statistical moments in output of 4D-

VAR as a weakness compared to EnSRF (p. 12838, p. 12839, p. 12845, p. 12849). Indeed the raw algorithm does not provide an error covariance matrix directly. In p. 12849, l. 22, the text mentions that solutions exists but are 'extremely expensive'. However, such solutions are routinely used and have formed the basis of many papers (e.g., Chevallier et al. 2007, 2013). The existence of these papers demonstrates that such methods are not 'extremely expensive'. It can also be shown that these approaches can fasten convergence of the reference inversion (Desroziers and Berre 2012). In other words, the second order statistical moments can be obtained from 4D-VAR if they are needed, even though it is more (only more) expensive. On the EnSRF side, the paper emphasises its capability to produce these quantities, but without studying the quality of these products. These two points bias the discussion and invalidate the conclusion given in p. 12845 ("the ENSRF is more desirable for attribution purposes, wherein source/sink estimates with confidence bounds can be used to gain better. . .").

Response 2.3: We do not agree with the reviewer that the paper does not study the quality of the posterior covariances obtained from EnSRF. For example, as discussed in p. 12837, lines 17-25 (also see p. 12839, lines 14-17, p. 12841, lines 19-25), we have made a concerted effort to compare the posterior covariances from EnSRF to the posterior covariances from the batch scheme by looking at the ratio of the two quantities. If the EnSRF posterior covariances over or under-estimate the batch posterior covariances (p. 12841, lines 19-25), we have discussed the possible reasons for this behavior.

As stated in Response 1.1, in the revisions we will include the posterior covariances from the 4D-VAR scheme, and compare that with the posterior covariance obtained from the batch perspective. We agree that this will balance out the discussion and it will provide a good comparison between the performance of the EnSRF and the 4D-VAR relative to the batch scheme.

4. *In p. 12841, operational constraints are simulated by reducing the ensemble size to 100, given a domain of 300 pixels. In proportion, this is still much larger than the ensemble systems quoted in the introduction and therefore dramatically damps the performance loss. Tests with a smaller number (e.g., 10) should be shown.*

Response 2.4: Even though the domain size is 300 pixels, we estimate for 35 time periods, which makes the total state vector size 300*35=10500 (p. 12834, lines 14-16). The lag window size is 5 [T] (p. 12833, lines 18-20), which implies that each observation that is assimilated impacts a state vector of size 300*5 = 1500. Thus, the relative size of the state space to the ensemble size is 1500/100 =15.

Proportionately this is low compared to "real" applications. For example, our recent experiments for solving daily global surface fluxes at 1° × 1° using the GOSAT-ACOS data use a state space/ensemble size ratio close to 10,000. Conversely, *Chatterjee et al.* [2012] for a synthetic data application over North America employed a state space/ensemble size ratio of 500. There is no mathematical agreement over what constitutes a good or a bad ratio and is clearly based on the computational size of the problem being studied, and the number of ensemble members that can be efficiently run.

In this study, when operational constraints are not introduced, the ensemble size (for the

EnSRF) and the number of descent iterations (for the 4D-VAR) are set to 1000 and 250, respectively. The number of descent iterations is prescribed to be lower than the number of ensemble members, keeping in mind that 4D-VAR typically requires more model integrations (i.e., both forward and adjoint model run) than EnSRF (p. 12836, Lines 3-5). Thus, when operational constraints are introduced, both the ensemble size and the descent iterations are reduced by a factor of 10. If we reduce the ensemble size further by another factor of 10 (i.e. down to 10 ensemble members), then to be consistent the number of 4D-VAR iterations should be equivalently reduced to 2. While the EnSRF would require vigorous localization and inflation to ameliorate the sampling error, 4D-VAR would of course not converge with only 2 iterations. Clearly, both the DA methods would be at a severe disadvantage with such *stringent* operational constraints and have huge performance losses. In response to the reviewer's concern, we plan to conduct experiments where we impose operational constraints by reducing the ensemble size and the number of descent iterations by a factor of 20, and re-assessing whether the conclusions presented in Section 3.3 remain consistent. Accordingly the discussion will be modified in the revised version of the manuscript.

5. In p. 12842, l. 11, the authors suggest using a dynamical flux model. This is a fait point, but the discussion should state that this is already achieved in CCDASs (e.g., Rayner 2010) at the cost of adding significant model errors (e.g., Kuppel et al. 2013), therefore reducing the performance of the inversion in terms of inverted fluxes.

Response 2.5: We thank the reviewer for this suggestion. In the revised manuscript, this discussion and suggested references will be included.

6. p. 12827, l. 13 and elsewhere: the exclusion of batch schemes from the DA group is not appropriate since batch schemes should yield the same solution than the "DA" schemes if perfectly implemented.

Response 2.6: We agree with the reviewer's comment that batch schemes should yield the same solution as the DA schemes if perfectly implemented, and this is exactly why we use the batch inversion as a benchmark in the comparison. In terms of nomenclature, however, we attempted to make a subtle distinction between the *traditional batch* (i.e., either regular Bayesian Inverse Modeling or the Geostatistical Inverse Modeling) scheme and the *DA* (4D-VAR, ensemble Kalman filter, hybrid) schemes. The traditional batch scheme uses the analytical expression arising from the maximum likelihood estimation for Gaussian pdfs and a linear model to calculate the parameters of the posterior pdf. For linear systems and Gaussian pdfs, the DA methods are numerical approximations of this. By using two different terms (*traditional batch* vs. *DA*) we wanted to highlight the aspect of the numerical approximation.

We acknowledge the reviewer's concern that inadvertently we have made it sound that the *traditional batch* and the *DA* schemes are more different than they really are. In the revised manuscript, we will add the above discussion to make the reader aware that the batch scheme is an example of DA too.

7. P. 12849, l.7: It is not clear why the study by Gejadze et al. (2012) is referenced in this context, since it applies to very non-linear systems.

Response 2.7: Our primary intention here was to make the reader aware that computationally efficient techniques for obtaining the posterior covariance from a 4D-VAR setup indeed exist, and are the subject of ongoing research. Even though the test case in the *Gejadze et al*. [2012] paper is non-linear, for a linear problem (such as the CO_2 flux estimation problem), the technique will still deliver a posterior covariance that is *approximately* equivalent to the posterior covariance from the batch perspective. In the revisions we will modify the line to read: "*Recent applications of 4D-VAR for NWP problems have shown that computationally efficient alternatives do exist (e.g. Cheng et al., 2010; Gejadze et al., 2012). Although some of these techniques are more suited for non-linear problems (e.g. Gejadze et al., 2012), these retain potential applicability to the linear CO_2 flux estimation problem as well*".

References:

Baker, D. F., S. C. Doney, and D. S. Schimel (2006), Variational data assimilation for atmospheric CO(2), *Tellus Series B-Chemical and Physical Meteorology*, 58(5), 359-365, doi: 10.1111/j.1600-0889.2006.00218.x

Bruhwiler, L. M. P., A. M. Michalak, W. Peters, D. F. Baker, and P. Tans (2005), An improved Kalman Smoother for atmospheric inversions, *Atmospheric Chemistry and Physics*, 5, 2691-2702, doi:10.5194/acp-5-2691-2005

Chatterjee, A., A. M. Michalak, J. L. Anderson, K. L. Mueller, V. Yadav (2012), "Towards reliable ensemble Kalman filter estimates of CO_2 fluxes", *Journal of Geophysical Research – Atmospheres*, 117, D22306, doi:10.1029/2012JD018176

Cheng, H. Y., M. Jardak, M. Alexe and A. Sandu (2010), A hybrid approach to estimating error covariances in variational data assimilation, *Tellus Series a-Dynamic Meteorology and Oceanography*, 62(3), 288-297, doi: 10.1111/j.1600-0870.2010.00442.x

Chevallier, F., et al. (2012), What eddy-covariance measurements tell us about prior land flux errors in CO2-flux inversion schemes, *Global Biogeochemical Cycles*, 26, GB1021, doi:10.1029/2010GB003974

Chevallier, F., et al. (2010), CO_2 surface fluxes at grid point scale estimated from a global 21 year reanalysis of atmospheric measurements, *Journal of Geophysical Research-Atmospheres*, 115, D21307, doi:10.1029/2010jd013887.

Chevallier, F., F.-M. Breon, and P. J. Rayner (2007), Contribution of the Orbiting Carbon Observatory to the estimation of CO_2 sources and sinks: Theoretical study in a variational data assimilation framework, *Journal of Geophysical Research-Atmospheres*, 112, D09307, doi:10.1029/2006jd007375.

Feng, L., P. I. Palmer, H. Boesch, and S. Dance (2009), Estimating surface CO(2) fluxes from space-borne CO(2) dry air mole fraction observations using an ensemble Kalman Filter, *Atmospheric Chemistry and Physics*, 9(8), 2619-2633, doi: 10.5194/acp-9-2619-2009

Gejadze, I. Y., V. P. Shutyaev and F. X. Le Dimet (2012), Analysis error covariance *versus* posterior covariance in variational data assimilation, *Quarterly Journal of the Royal Meteorological Society*, doi: 10.1002/qj.2070

Kang, J.-S., E. Kalnay, T. Miyoshi, J. Liu, and I. Fung (2012), Estimation of surface carbon fluxes with an advanced data assimilation methodology, *Journal of Geophysical Research-Atmospheres*, 117, D24101, doi:10.1029/2012JD018259.

Michalak, A. M. (2008), Technical Note: Adapting a fixed-lag Kalman smoother to a geostatistical atmospheric inversion framework, *Atmospheric Chemistry and Physics*, 8(22), 6789-6799, doi:10.5194/acp-8-6789-2008

Miyazaki, K., T. Maki, P. Patra, and T. Nakazawa (2011), Assessing the impact of satellite, aircraft, and surface observations on CO_2 flux estimation using an ensemble-based 4-D data assimilation system, *Journal of Geophysical Research-Atmospheres*, 116, D16306, doi:10.1029/2010JD015366

Peters, W., J. B. Miller, J. Whitaker, A. S. Denning, A. Hirsch, M. C. Krol, D. Zupanski, L. Bruhwiler, and P. P. Tans (2005), An ensemble data assimilation system to estimate CO_2 surface fluxes from atmospheric trace gas observations, *Journal of Geophysical Research-Atmospheres*, 110, D24304, doi: 10.1029/2005JD006157

Peters, W., et al. (2007), An atmospheric perspective on North American carbon dioxide exchange: CarbonTracker, *Proceedings of the National Academy of Sciences of the United States of America*, 104(48), 18925-18930, doi:10.1073/pnas.0708986104.

Yadav, V. and Michalak, A. M. (2013), Improving computational efficiency in large linear inverse problems: an example from carbon dioxide flux estimation, *Geosciences Model Development*, 6, 583-590, doi:10.5194/gmd-6-583-2013.