

Interactive comment on “Balancing aggregation and smoothing errors in inverse models” by A. J. Turner and D. J. Jacob

Anonymous Referee #2

Received and published: 17 March 2015

Choosing an optimal dimension for the state vector to be optimized in an atmospheric inversion has been a long standing issue. As the authors of this manuscript point out, solving for the state vector at the native resolution of the CTM can introduce smoothing errors, while pre-aggregating and solving for only certain patterns is sure to introduce aggregation errors. Therefore, a systematic study of how to balance the two, which is what the authors have presented, is a welcome addition to the field, and should be published.

My biggest complaint, however, is the choice of the journal. When I read or review a paper in Atmospheric Chemistry and Physics, my first question is "What have I learned about the physics or chemistry of the atmosphere from this paper?" Unfortunately for this manuscript, the answer to that question is "Nothing!". This is not to say that the

C858

work is not good or not important; it is both, and should be published. However, it is a technical study that will be of relevance only to a class of modelers during their model development, and therefore I think Geoscientific Model Development (from the same publishers) is a much better journal for publishing this work. I would strongly urge the authors to consider submitting this specific work to that journal instead. I do not think this suggestion should come as a surprise to the authors. Previous work on the same problem (which they cite) was published in the Quarterly Journal of the Royal Meteorological Society, and similar technical developments are routinely published in Geoscientific Model Development.

My second biggest complaint is the applicability of the technique detailed here. As someone who does atmospheric inversions off and on, my first impulse upon coming across a manuscript of this sort is to wonder "This looks great! Can I apply this technique to my inversions?" From the manuscript, it is not clear that I or any other atmospheric inverse modeler will be able to use the results presented here in real-world inversions. The authors choose the optimal number of state vector elements as the number which minimises the total error in Figure 3. If I understood correctly, generation of Figure 3 required performing the same inversion over and over again with different restriction operators Γ , to get the posterior covariance matrices. This was possible for the authors because their native resolution state vector was small, owing to their choice of focussing on the annual average emission over N America. In most real world inversions spanning multiple years with daily/weekly variability in the fluxes, performing the inversion is the most time consuming part, and so performing many inversions just to figure out the optimal size of the state vector seems like a waste of resources. After all, since the authors show that even at the native resolution the smoothing error does not become significant compared to the observational error, what's wrong with just solving at the native (CTM) resolution? I would be happy to be proved wrong on this point, and to be shown that one doesn't need to execute a bunch of inversions to estimate the optimal size. From the current manuscript, however, I do not see how one could use this technique in a real-world inversion, for example any of

C859

the CO₂ inversions in Peylin et al (Biogeosciences, 2013), or any of the CH₄ inversions in Kirschke et (Nature Geoscience, 2013). This is one more reason why I would prefer to have this manuscript published in a journal dedicated to technical developments (such as GMD) instead of ACP.

Apart from the above (major) points, I have a few comments I would like the authors to address:

(1) In the abstract and in section 5 (bottom of p1017), the authors make the point that the GMM method retains resolution of major local features in the state vector. This is true, but only if the prior already has that particular feature. Further, this is not always an advantage, since those major features can sometimes be wrongly located in the prior emission estimate (less of an issue with coal mines and power plants, big issue for wetlands and bovine methane). I would like the authors to mention this.

(2) On page 1003, near line 25, the authors say that an additional cost of using a large state vector is the increased computational cost of the inversion. This is not correct. In fact, in most inversions beyond TRANSCOM-style basis region inversions, the costliest part of the inversion is the evaluation of the forward model F (and its adjoint, if needed), be it a CTM in variational/EnKF systems, or an LPDM for "batch" inversions. Irrespective of the aggregation chosen for the state vector, the atmospheric transport still needs to be run at the native resolution, which is the time limiting step.

(3) On page 1008, near line 15, the authors mention the assumption that the prior is unbiased. While this is an assumption widely adopted theoretically, in practice it is rarely true. A biased prior leads to a biased posterior, a fact inverse modellers grudgingly live with, as long as they think that the posterior bias is lower than their posterior uncertainty estimate. I would like to know what the consequence of a biased prior is for determining the optimal length of the state vector. Is that estimate expected to change?

(4) On page 1011, near line 15, the authors have a caveat, which, if I understand

C860

correctly, says that one of the assumptions is that the error covariance matrix of the true state is the same as the error covariance matrix for the prior state. Did I understand correctly? If so, then that's a big assumption; knowing the error covariance of the true state before doing an inversion seems like a big ask! If I misunderstood, I will be happy to be corrected.

(5) One aggregation technique the authors do not discuss is K-means clustering. If we choose the number of clusters to be equal to the optimal number of state vector elements, and use the same 14 variables as the GMM model to determine the clusters, how would the smoothing and aggregation errors compare to the GMM+RBF case? Did the authors already look into that? If so, I would love to see the results.

(6) On page 1016, line 4, the authors say that equations (32)-(35) are iterated until convergence. What counts as convergence, i.e., what is the convergence criterion?

(7) On page 1017, line 26, the authors say that RBF weighting performs slightly better than GMM clustering. Is this a general statement about RBF vs clustering, or is it because the 14 variables used to construct the similarity matrix (table 1) are strongly correlated with CH₄ fluxes?

Interactive comment on Atmos. Chem. Phys. Discuss., 15, 1001, 2015.

C861