Response to Reviewer Comments:

We thank Dr. Marc Bocquet and an Anonymous Reviewer for their thorough comments, we think they have greatly improved the content of the manuscript.

Reviewer #1 (Dr. Marc Bocquet) Comments:

The authors overlooked the findings that have been reported by Bocquet et al. (2011); Wu et al. (2011). From the results of Bocquet et al. (2011); Wu et al. (2011), I believe that the optimal resolution as seen by the authors is the results of suboptimal choices. What is reported in the present manuscript is nevertheless interesting since those suboptimal choices could be made for the sake of numerical efficiency. It is problematic that the authors are (unintentionally) hiding what actually leads to the appearance of a minimum in the total error curve as a function of resolution.

We have expanded the discussion of Bocquet et al. (2011), Bocquet and Wu (2011), and Wu et al. (2011). (see response to Dr. Bocquet's second comment below).

1.) Frankly, the notations are unfriendly. I understand the authors follow those of Rodgers (2000). Yet, they diverge a lot from standard data assimilation or inverse modelling notations that have been widely adopted in atmospheric chemistry data assimilation. For instance "a" usually refers to the analysis while the authors use it to refer to the prior, when "b" ("f" in a sequential context) is very often chosen. The gain is usually designated as **K**, not **G**; "**H**" is much preferred to "**K**" for the observation/Jacobian/source-receptor operator. That said, the choice of notations belongs to the authors. But, I guess that the present notations would significantly distract potential readers.

We draw on notation from the inverse modelling and trace gas retrieval communities. This notation is standard in the retrieval communities (e.g., Rodgers 2000) where most of the smoothing error discussion has previously taken place (e.g., von Clarmann 2014). This is also the notation that has been used by the Jacob group for the better part of a decade (e.g., Jacob et al. 2002; Jones et al. 2003; Heald et al. 2004; Palmer et al. 2006; Kopacz et al. 2009; Drury et al. 2011; Wecht et al. 2012; Zoogman et al. 2013; Wecht et al. 2014). As such, we prefer to keep the present notation given the history of use in both inverse modelling and trace gas retrieval communities.

2.) One of the results of Bocquet et al. (2011) is that with a proper choice of prolongation operator, one can reduce the smoothing error as much as possible, so that the total error (smoothing+aggregation) is actually a monotonically decreasing function of the resolution. If this is correct, there is no optimal resolution but the finest one (CTM's for instance), except from a numerical efficiency standpoint or if one introduces other sources of scale-dependant errors (such as model errors). The authors presumably obtain such (discrete) optimum because they make an arbitrary choice in the prolongation operator which restricts the transfer of information through scales. Mathematically speaking, this can be seen as an artifact. Had the authors made another implicit choice for the prolongation operator, they

would have found a different result, possibly leading to the finest grid being optimal. If this correct, the authors should clearly acknowledge this and give a fair account of the findings of (Bocquet et al., 2011).

We appreciate Dr. Bocquet's insightful discussion of the error and expanding on the findings of Bocquet et al (2011) and Wu et al. (2011). However, we believe there may have been a slight misunderstanding in the interpretation of the local minimum. Dr. Bocquet assumes the local minimum was due to the use of suboptimal restriction/prolongation operators. As Dr. Bocquet pointed out, this issue was discussed in Bocquet et al. (2011) and Wu et al. (2011). In practice, we are generally unable to specify the true off-diagonal terms in the covariance matrices and typically resort to using a single correlation length scale for the entire domain. In reality, these length scales are not constant. For example, regions dominated by wetland sources (e.g., the Hudson Bay Lowlands) will have large error correlation length scales because the underlying emissions for a large region are driven by a parameterized wetland model while a region like Los Angeles will be largely independent of the surrounding region and should not have a large error correlation length scale. So, while we may not be able to specify the true off-diagonal terms in the covariance matrices we may be able to approximate them. Thus, it was our goal to design our state vector such that it accounts for these off-diagonal terms that are generally missing from native-resolution inversions.

Therefore, the appearance of the local minimum is due to the coarser state vector accounting for off-diagonal terms while the native-resolution inversion does not, thus reducing the error. That said, this is not inconsistent with Dr. Bocquet's reasoning that our local minimum is an "artifact". If the covariance matrices at the native resolution included realistic off-diagonal terms then this approach would indeed be "suboptimal" because our state vector design is not dependent on the Jacobian at the native resolution.

As for Dr. Bocquet's discussion of the total error, in many cases the state vector is not designed using any sort of "intelligent" method like tilings, qtrees, ftrees, PCA, GMM, *k*-means, etc. but is instead designed based on predetermined regions like the TRANSCOM regions (as Review #2 touches on) or a simple coarse-graining scheme. In this case, the simple derivations presented here are useful because they do not require any assumptions about the prolongation operator that Dr. Bocquet claims is hidden in $\mathbf{K}_\omega$ (minor comment #11). We construct the $\mathbf{K}_\omega$'s by perturbing the elements of our reduced state vector, thus explicitly constructing $\mathbf{K}_\omega$ for each different case at many resolutions. We make a choice in designing the restriction operator ($\mathbf{\Gamma}_\omega$) that gives us $\boldsymbol{x}_\omega$ and $\mathbf{K}_\omega$. Correct me if I'm wrong, but I fail to see how we could simply "choose another prolongation operator". The only choice we make is the restriction operator. Thus, the derivations presented here are valuable in that they do not require us to make a choice about a prolongation operator.

We have updated the text to explain this:

Lines 383-394: "Previous work by Bocquet (2009), Bocquet et al. (2011), Bocquet and Wu (2011), Wu et al. (2011), and Koohkan et al. (2012) analyzed the scale-dependence of different grids using the degrees of freedom for signal: $\mathrm{DFS} = \mathrm{Tr}(\mathbf{I} - \mathbf{S}_{a,\omega}^{-1}\hat{\mathbf{S}}_\omega)$. These past works found this error metric to be monotonically increasing. This implies that the native resolution grid will have the least total error and there is no optimal resolution, except from a numerical efficiency standpoint. Here we find a local minimum that is, seemingly, at odds with this previous work. However, the reasoning for this local minimum is that we have

allowed the aggregation to account for spatial error correlations that we are unable to specify at the native resolution. As such, we are taking more information into account and obtaining a minimum total error at a state vector size that is smaller than the native resolution. If the native resolution error covariance matrices were correct then, as previous work showed, the only reason to perform aggregation would be to reduce the computational expense and the grid used here would be suboptimal because it does not depend on the native-resolution grid."

In addition to updating the text to clarify this local minimum, we have added a paragraph discussing the findings of Bocquet (2009), Bocquet et al. (2011), Bocquet and Wu (2011), Wu et al. (2011), and Koohkan et al. (2012) and a distinction of our work:

Lines 57-63: "Previous work by Bocquet (2009), Bocquet et al. (2011), Bocquet and Wu (2011), Wu et al. (2011), and Koohkan et al. (2012) developed optimal grids that allow the transfer of information across multiple scales. These computationally efficient methods (Bocquet and Wu, 2011) generally require the use of the native-resolution grid to derive the optimal representation. They also assume that the native-resolution prior error covariance matrices can be accurately constructed. However, in practice we are generally unable to specify realistic prior error correlations and must resort to simple assumptions."

Minor Comments:

1.) Title: We all know there is no such thing as an "inverse model". This is an abuse of language that I would personally avoid in a title. "Inverse modelling" is almost always preferred.

"Inverse modelling" would awkwardly add another gerund in the title.

2.) p. 1002, l. 4-6: "When the observation vector is large, such as with satellite data, selecting a suitable dimension for the state vector is a challenge". Selecting a suitable dimension for the state vector space is always a challenge, even, and perhaps even more so when the observation vector is small. Let me just mention one paper directly related to what you are discussing and where the observations are in situ and far less abundant than in a satellite retrieval context: Koohkan et al. (2012).

Indeed but that is not really the problem we are addressing. We agree that choosing a suitable state vector is always challenging. However, this work uses satellite data for the example problem and the companion paper performs a "real-inversion" using satellite data. The information content from *in situ* data is, generally, far more intuitive. The observations provide a lot of information near the site and upwind. One could conceivably design a decent state vector by placing many grid cells near the site with fewer upwind. With satellite data (and total column observations) the information content is not immediately clear.

3.) p. 1003, l. 6-7: Same remark as above.

See above.

4.) p. 1003 l. 18-19: "and may not be able to depart from that knowledge". It all depends on the balance between the observation and background statistics. If the background is not informative enough, the solution may be highly oscillating. In a flux inversion context, the retrieved fluxes would increase around the observations sites, which is all but smoothing. In my humble opinion, the appellation is partially misleading. But I might not have understood its interpretation very clearly (from your manuscript or even Rogders' book), in spite of some experience with inverse modelling.

We have rephrased this and added additional citations:

Lines 34-40: "The inverse solution must then rely on some prior estimate for the state vector and may not be able to depart sufficiently from that knowledge. The associated error is known as the smoothing error (Rodgers, 2000; von Clarmann, 2014) and increases with size of the state vector (Bousquet et al., 2000; Kaminski and Heimann, 2001; Kaminski et al., 2001; von Clarmann, 2014)."

5.) p. 1003 l. 18-19: "smoothing error" lacks a proper definition (although it is given later) and interpretation.

See above.

6.) p. 28-29: "Numerical solutions using variational methods circumvent this problem but not inherently provide error characterisation as part of the solution": we know that this is not true. If that kind of statement was fine a few years ago, I believe it should be nowadays mitigated. Several researchers are using conjugate-gradient and quasi-Newton methods such as BFGS that inherently provide estimation of the posterior errors (for instance Bousserez et al., 2015).

We have updated the text to mention these approximate methods in Sections 1 and 2:

Lines 47-48: "Approximate error statistics can be obtained (e.g., Bousserez et al., 2015) but at the cost of additional computation."

Lines 115-118: "Several approaches have been presented to obtain approximate error characterization (e.g., Courtier et al., 1994; Desroziers et al., 2005; Chevallier et al., 2007; Bousserez et al., 2015) but they can be computationally expensive."

Furthermore, we compared the exact posterior covariance matrix to some of these approximate posterior covariance matrices and found the discrepancies were large to be useful for our analysis, thus we decided not to pursue that approach further.

7.) p. 1004: the literature is incomplete. I believe you have to mention Wu et al. (2011), given it is very close to your objective and analysis and also related to greenhouse gas flux inversions.

Done. See, for example, our response to major comment #2.

8.) p. 1006, l. 5-12: This is incomplete or partially incorrect. The Jacobian can also be computed using the model adjoint, requiring m runs. By the Sherman-Morisson-Woobury lemma, the matrix algebra will scale like m^3. Also, sequential updating by serial processing of observations usually (unless the scheme is sub-optimal) leads to the same numerical cost.

Both points are true. However, typically $m >> n$ so using the adjoint to construct the Jacobian is usually not the most efficient method. As for the latter point, the benefit comes in the reduced memory cost for the matrix operations (i.e. break a large matrix up into smaller matrices so you can perform the necessary matrix operations).

9.) p. 1007, l. 11: "Probabilistic" is one word too many. Bocquet et al. (2011) additionally provide a probabilistic interpretation. But it can be seen as an entirely deterministic process just as the Best Linear Unbiased Estimator (BLUE) formalism. Please remove the word "probabilistic" which conveys the wrong idea in the context of this sentence.

Done. See minor comment #10 (below) for updated text.

10.) p. 1007, l. 12: "However, construction of this prolongation operator is not a well-posed problem because the operator is not unique". Please rephrase the sentence. The construction as defined by Bocquet et al. (2011) is well-defined and well-posed. But in general the choice of the prolongation operator is not unique. Incidentally, you do make a choice for the operator without acknowledging it! That is why I disagree and think that your method might be less robust. But maybe you meant "more practical" rather than "more robust", did you? If you did intend "less robust", please justify your statement with precision.

Our apologies. Yes, "more practical" was the intended meaning. We have corrected this phrasing.

We have updated the text as:

Lines 132-134: "Their analysis relies heavily on the construction of a prolongation operator ($\mathbf{\Gamma}_\omega^\star$) mapping $\boldsymbol{x}_\omega$ back to $\boldsymbol{x}$: $\boldsymbol{x} = \mathbf{\Gamma}_\omega^\star \boldsymbol{x}_\omega$. However, construction of this prolongation operator is not unique. We present here a simpler and more practical method."

11.) p. 1007, Eq.(12): Please define $\mathbf{K}_\omega$ (the source-receptor matrix). That is where you put the definition of the prolongation operator under the carpet... This must be discussed.

We explicitly constructed the $\mathbf{K}_\omega$'s through perturbations to the reduced state vectors for all the different cases. Thus, $\mathbf{K}_\omega$ was constructed in the same manner as $\mathbf{K}$. So all we needed was a reduced state vector ($\boldsymbol{x}_\omega$) which we defined in Eq. 10. There was no use of a prolongation operator in the construction of $\mathbf{K}_\omega$.

We have updated the text as:

Lines 145-147: “Here $\boldsymbol{y}$ is the observation vector (common in both cases), $\boldsymbol{x}$ and $\boldsymbol{x}_\omega$ are the true values of the native-resolution and aggregated state vectors, and $\mathbf{K}$ and $\mathbf{K}_\omega$ are the native resolution and the reduced-dimension Jacobians.”

12.) p. 1008, l. 7-8: The introduction of the concept of ensemble is cumbersome (just as it is in Rodgers (2000) to be fair). It requires more justification. It appears as a *deus ex machina*.

This concept is not critical to our derivation. Furthermore, this concept has been extensively discussed in the retrieval community. We have included additional references directing the reader to more complete discussions of this concept.

Lines 152-153: “Obtaining the error statistics for $\boldsymbol{\epsilon}_\mathrm{A}$ requires knowledge of the pdf of $\boldsymbol{x}$ for the ensemble of possible true states (cf. Rodgers, 2000; von Clarmann, 2014).”

13.) p. 1008; l. 5: "A" for aggregation, and "a" for background. Really? Why not use "b" for background instead of “a”?

“a” for the a priori is consistent with the notation of Rodgers (2000).

14.) p. 1010, l. 10: Please define the gain $\mathbf{G}_\omega$ explicitly. As I explained, several choices can be made, one being more consistent. Without an explicit definition, you hide what is at the origin of the appearance of the fittest resolution.

Done.

Line 197: “ $\mathbf{G}_\omega = \left(\mathbf{K}_\omega^T \mathbf{S}_\mathrm{O}^{-1} \mathbf{K}_\omega + \mathbf{S}_{\mathrm{a},\omega}^{-1}\right)^{-1} \mathbf{K}_\omega^T \mathbf{S}_\mathrm{O}^{-1}$ ”

15.) p. 1011, l. 6-8: The sum of an increasing and decreasing function does not always possess a minimum.

We have rephrased this line:

Lines 210-213: “Because the smoothing error increases with state vector dimension while the aggregation error decreases, analysis of the error budget can point to the optimal dimension where the total error is minimum.”

16.) Koohkan et al. (2012) discuss how to choose the optimal resolution and how it is impacted by the error balance (observation versus background, section 2.2). Since, ultimately, you end up making the same choice as all the papers I am referring to, that is to say choosing the resolution on a numerical cost basis, Koohkan et al. (2012)’ discussion is relevant and perhaps a bit more precise than only adjustment with respect to the

observation error only.

Thank you for pointing us to Section 2.2 of Koohkan et al. (2012) for the discussion of the error balance, however Section 3 of our manuscript is mostly concerned with simply deriving the different error components. We have added a brief discussion of Koohkan et al. (2012) (See response to minor comments #2 and #19).

17.) p. 1011, l. 12-18: Again, this discussion appears like a *deus ex machina*.

See response to minor comment #4 from Reviewer #2.

18.) p. 1012, l. 5: Actually the adaptive grid method based on tiling was introduced in Bocquet (2009). Moreover, it's worth mentioning that these grid are built to be optimal for the purpose of the inversion.

We have added Bocquet (2009) to the discussion and amended the description:

Lines 225-227: "Analyzing the off-diagonal structure of a precisely constructed prior error correlation matrix would provide the best objective way to carry out the aggregation, as described by Bocquet (2009), Bocquet et al. (2011), and Wu et al. (2011)."

19.) p. 1012, l. 8: Bocquet and Wu (2011) also use PCA coupled to the hierarchical grid to compute an optimal grid in a numerically efficient way yet capturing the variability of the prior. This should be acknowledged.

We have added Bocquet et al. (2011), Bocquet and Wu (2011), Wu et al. (2011), and Koohkan et al. (2012) to the discussion:

Lines 231-234: "Previous work by Bocquet et al. (2011), Wu et al. (2011), and Koohkan et al. (2012) used tiling and tree-based aggregation methods, while Wecht et al. (2014) used a hierarchal clustering method based on prior error patterns. Bocquet and Wu (2011) also used principal component analysis (PCA) coupled to the hierarchal grid to compute an optimal grid."

20.) p. 1013, l. 15-21: Rodgers (2000) also suggests projection over a specific function basis albeit in a different context.

21.) p. 1016, l. 5: Could you please briefly discuss the numerical cost of the approach?

See response to minor comment #6 from Reviewer #2.

22.) p. 1016: The application of the GMM methods is very interesting. From the methodological standpoint, I believe the fact that the control space is defined with a *probabilistic* mixture is quite novel in this context.

23.) p. 1017: What about the time dimension? Do you apply aggregation in time? I assume you didn't, but you could have.

This same approach should be applicable to the time dimension:

Line 240: "However, the same methods can be used for temporal aggregation."

24.) p. 1017: What if the background error covariance matrices were not diagonal? Could you discuss the issue a little? Apart from the numerical problem, we can see from Eq. (2) that properly transferring information through the scales is more tricky. If one chooses a pragmatical $\mathbf{\Gamma}_{\omega}^{\star}$ as you do (or as I could as well for a very high-dimensional application), it is possible that the resulting "optimal" resolution would be more pronounced.

This is an excellent question and goes back to the main issue discussed in your second major comment. The two main benefits of these methods would be: (1) computational cost and (2) accounting for spatial correlations that are difficult to specify. If one could specify realistic off-diagonal error correlations then the only benefit to a multi-scale approach would be the computational benefits. However, in practice we are unable to specify those realistic off-diagonal error correlations and, thus, have neglected valuable information. See our response to major comment #2 for the added text on this.

25.) p. 1017, l. 18: Please spell out SD (standard deviation?).

Done.

26.) p. 1018, l. 8-13: Your result is not surprising. Because of the baseline results of Bocquet et al. (2011), I was expected that kind of results with a non-pronounced minimum (unless your implicit prolongation operator is badly chosen). Above all, you end up choosing the optimal resolution on a numerical efficiency criterion, just as we did (for not only practical but also theoretical reasons). This should be acknowledged.

See response to major comment #2.

27.) p. 1018-1019: The conclusion should be amended.

Reviewer #2 Comments:

1.) 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 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.

We considered Geoscientific Model Development (GMD) but ultimately chose to Atmospheric Chemistry and Physics (ACP) for three reasons:

1.) This is a companion paper to Turner et al. ACPD (2015) that performs a "real-world" inversion with 2.5 years of GOSAT methane data. It seems fitting to keep the companion papers in the same journal.
2.) The focal point of this paper is about a methodology, not code development. Furthermore, based on the journal scopes, this seems to be a more natural fit for ACP than GMD.
    ACP scope: "The journal scope is focused on studies with general implications for atmospheric science" (see: "http://www.atmospheric-chemistry-and-physics.net/about/aims_and_scope.html").
    GMD scope: "dedicated to the publication and public discussion of the description, development, and evaluation of numerical models of the Earth system and its components." (see: "http://www.geoscientific-model-development.net/about/aims_and_scope.html").
    There have been previous "methods" papers that were published in ACP.
3.) This is a methodology that is widely applicable to both atmospheric chemistry and physics. We also directly apply this methodology to an atmospheric chemistry problem as an example.

2.) 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 \Gamma, 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 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.

We suspect the reviewer may not have noticed that the time period for this manuscript was greatly reduced from our "real-world" inversion presented in the companion paper (Turner et al. ACPD 2015). So the reviewer is correct that we do use Figure 3 to guide the choice of state vector size but we do so using a shorter time period. We have rephrased the text to mention this sooner:

Lines 165-167: "Application of Eq. 17 requires computation of the native-resolution Jacobian **K** but this can be done for a limited test period only. We will give an example below."

> 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?

We agree that the native CTM resolution would be ideal if we could specify the off-diagonal error correlations in a realistic manner. However, in practice the off-diagonal error correlations are neglected or simply treated with a single correlation length scale. In this framework we are able to, in essence, prescribe different correlation lengths for different regions (e.g., LA is distinct from the surrounding region whereas the HBL wetlands have a longer correlation length). See our response to Dr. Bocquet's major comment #2.

> From the current manuscript, however, I do not see how one could use this technique in a real-world inversion, for example any of the CO2 inversions in Peylin et al (Biogeosciences, 2013), or any of the CH4 inversions in Kirschke et (Nature Geoscience, 2013)

As for the applicability to "real-world" inversions, I believe this is a perfect example of the applicability. In this manuscript we sampled the full range of possible state vector sizes and determined a reasonable state vector size. We then used that state vector in a "real-world" inversion (Turner et al. ACPD 2015).

Minor Comments:

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.

Thank you for bringing this up. The GMM method can retain major local features if they are

in the prior *or* the adjoint-constraint ("Similarity Vector" number 3; Table 1). Our motivation for including the adjoint-constraint dataset was to allow the state vector to include potential "missing sources" seen by the observations. Further, one could add other datasets (or replace sectors in the prior) if they have reason to believe the other dataset is representative of prior error correlations. We have expanded on this in the manuscript:

Lines 255-256: "the third [similarity vector] represents the scaling factors from the first iteration of an adjoint-based inversion at native resolution"

Lines 259-263: "We choose here to include initial scaling factors from the adjoint-based inversion because we have them available and they can serve to correct any prior patterns that are grossly inconsistent with observations, or to identify local emission hotspots missing from the prior. One iteration of the adjoint-based inversion is computationally inexpensive and is sufficient to pick up major departures from the prior."

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.

We were specifically referring to analytical inversions like the one performed here. In an analytical inversion a larger state vector could increase the computational cost of the inversion in two ways: (1) it will increase the size of the matrices that need to be multiplied and inverted and (2) it may increase the number of forward run "batches". For example, simulating atmospheric transport at high resolution generally requires a lot of memory. A modeler can quickly reach the allowable memory limits if they are running thousands of forward runs to construct the Jacobian. This was the main limiting factor in our simulations; ultimately, we could only run a few hundred forward simulations at a time due to memory constraints on the cluster. Thus, we had to run multiple batches of forward runs. The second point does not apply to LPDM inversions. We have rephrased the text to clarify this.

Lines 43-44: "An additional drawback of using a large state vector is that analytical solution to the inverse problem may not be computationally tractable. Analytical solution requires…"

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?

Absolutely, a biased prior will bias the posterior. The only reason that an additional error term would impact the optimal state vector length is if it exhibited scale-dependence. Intuitively, a bias term would not exhibit scale-dependence and would, presumably, behave

like the observation error term. As such, a bias in the prior will bias the posterior but not affect the choice of the optimal state vector size. We have added text to reflect this.

Lines 208-209: "A bias term should exhibit similar scale-dependence to the observation error term and could be included by following the derivation from Rodgers (2000)."

Lines 101-103: "We have assumed here that errors are unbiased, as is standard practice in the inverse modeling literature. An observational error bias $\boldsymbol{b}_O$ would propagate as a bias $\mathbf{G}\boldsymbol{b}_O$ in the solution $\hat{x}$ in Eq. 8."

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

This is, indeed, correct. In Section 5 we present a simple experiment where we pretend to know the true emissions, so for the purposes of this experiment our assumptions are valid. These expressions are useful for similar experiments where one wants to diagnose the different error components. However, these expressions should not be used to diagnose errors in a "real-world" inversion because that assumption will not hold. Rodgers (2000, p. 49) and von Clarmann (2014) present a detailed discussion of this exact issue. We have rephrased this paragraph:

Lines 215-221: "A caveat in the above expressions for the aggregation and smoothing error covariance matrices is that they are valid only if the prior $\boldsymbol{x}_a$ is the mean value $\bar{x}$ for the pdf of true states and if the error covariance matrix Sa is the covariance matrix for that pdf ($\mathbf{S}_e = \mathbf{S}_a$). Rodgers (2000, p. 49) and von Clarmann (2014) provide a detailed discussion of the errors induced by failing to meet this assumption. As such, these conditions define the assumption for the prior, so the expressions can be taken as valid for the purpose of selecting an appropriate state vector dimension in an inverse problem. However, they should not be used to diagnose errors on the inversion results."

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.

Excellent question. As the reviewer may have surmised, there was a very large computational expense associated with explicitly constructing the Jacobian multiple state vector sizes with multiple methods. As such, we considered including *k*-means clustering and performed some preliminary analysis with *k*-means. The figure below shows example clusters created using the same similarity matrix and criteria as in Figure 1. The *k*-means clustering used 100 replicates (different initializations) with 1000 iterations per replicate. We ultimately decided to include the course-graining method instead of *k*-means. However, the different methods perform comparably so, presumably, *k*-means would also give similar results.

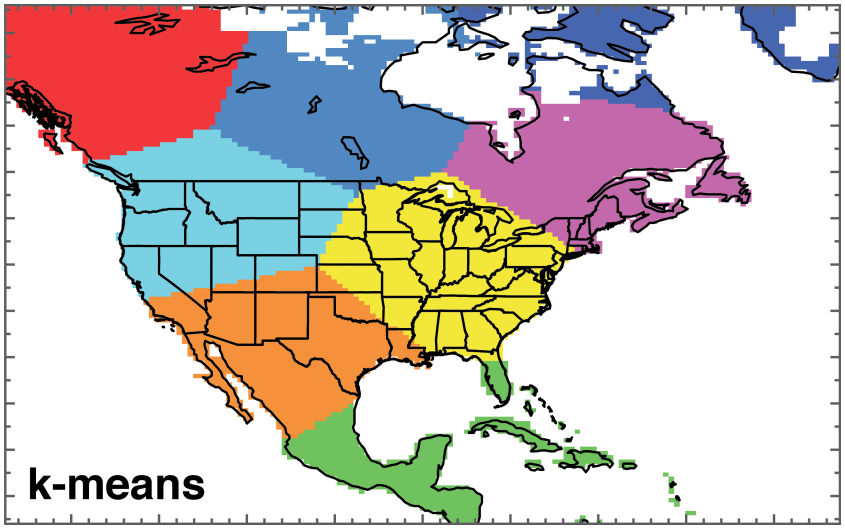

Figure 1: Same as Figure 1 from the manuscript but for *k*-means clustering.

We have added the following text:

Lines 365-368: "The different aggregation methods of Sect. 4 yield very similar smoothing errors, suggesting that any reasonable aggregation scheme (such as *k*-means clustering (cf. Bishop 2007)) would perform comparably."

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?

We used an absolute tolerance of $\tau < 10^{-10}$ where:

$$\tau = \sum_i \sum_j \left| \mathcal{M}_{i,j} - \mathcal{M}^{\star}_{i,j} \right| + \sum_i \sum_j \sum_k \left| \mathcal{L}_{i,j,k} - \mathcal{L}^{\star}_{i,j,k} \right| + \sum_i \left| \pi_i - \pi^{\star}_i \right|$$

and the superscript star indicates the value from the previous iteration. We didn't use a relative tolerance because the true value of one of the parameters could, potentially, be zero. In any case, preliminary tests were insensitive to using a tolerance of 10^{-4}. In response, we have added the following text to the manuscript:

Lines 319-326: "The computational complexity for the expectation-maximization algorithm is $\mathcal{O}(nK + pn^2)$ (Chen et al., 2007), however the actual runtime will be largely dictated by the convergence criteria. Here we use an absolute tolerance of $\tau < 10^{-10}$ where

$$\tau = \sum_i \sum_j |\mathcal{M}_{i,j} - \mathcal{M}^{\star}_{i,j}| + \sum_i \sum_j \sum_k |\mathcal{L}_{i,j,k} - \mathcal{L}^{\star}_{i,j,k}| + \sum_i |\pi_i - \pi^{\star}_i|$$

and the superscript star indicates the value from the previous iteration."

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?

This is a general statement about RBF weighting vs. GMM clustering (as well as coarse-graining and PCA clustering). However, it's not necessarily a general statement about RBF weighting vs. other clustering methods. That said, we suspect that RBF weighting would perform favorably against more clustering methods but we have only tested a small subset of clustering methods here.

Manuscript prepared for Atmos. Chem. Phys.
with version 2014/09/16 7.15 Copernicus papers of the LaTeX class copernicus.cls.
Date: 27 March 2015

Balancing aggregation and smoothing errors in inverse models

A. J. Turner[1] and D. J. Jacob[1,2]

[1]School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA
[2]Department of Earth and Planetary Sciences, Harvard University, Cambridge, Massachusetts, USA
Correspondence to: A. J. Turner (aturner@fas.harvard.edu)

Abstract. Inverse models use observations of a system (observation vector) to quantify the variables driving that system (state vector) by statistical optimization. When the observation vector is large, such as with satellite data, selecting a suitable dimension for the state vector is a challenge. A state vector that is too large cannot be effectively constrained by the observations, leading to smoothing error. However, reducing the dimension of the state vector leads to aggregation error as prior relationships between state vector elements are imposed rather than optimized. Here we present a method for quantifying aggregation and smoothing errors as a function of state vector dimension, so that a suitable dimension can be selected by minimizing the combined error. Reducing the state vector within the aggregation error constraints can have the added advantage of enabling analytical solution to the inverse problem with full error characterization. We compare three methods for reducing the dimension of the state vector from its native resolution: (1) merging adjacent elements (grid coarsening), (2) clustering with principal component analysis (PCA), and (3) applying a Gaussian mixture model (GMM) with Gaussian pdfs as state vector elements on which the native-resolution state vector elements are projected using radial basis functions (RBFs). The GMM method leads to somewhat lower aggregation error than the other methods, but more importantly it retains resolution of major local features in the state vector while smoothing weak and broad features.

1 Introduction

Inverse models quantify the state variables driving the evolution of a physical system by using observations of that system. This requires a physical model $\mathbf{F}$, known as the forward model, that relates

a set of input variables $\boldsymbol{x}$ (state vector) to a set of output variables $\boldsymbol{y}$ (observation vector),

$$\boldsymbol{y} = \mathbf{F}(\boldsymbol{x}) + \boldsymbol{\epsilon} \tag{1}$$

The observational error $\boldsymbol{\epsilon}$ includes contributions from both the forward model and the measurements. Solution to the inverse problem involves statistical optimization to achieve a best error-weighted estimate of $\boldsymbol{x}$ given $\boldsymbol{y}$.

A critical step in solving the inverse problem is determining the amount of information contained in the observations and choosing the state vector accordingly. This is a non-trivial problem when using large observational datasets with large errors. An example that will guide our discussion is the inversion of methane emissions on the basis of satellite observations of atmospheric methane concentrations (Turner et al., 2015). Methane concentrations can be predicted on the basis of emissions by using a chemical transport model (CTM) that solves the 3-D continuity equation for methane concentrations. Here the CTM is the forward model $\mathbf{F}$, the satellite provides a large observation vector $\boldsymbol{y}$, and we need to choose the resolution at which to optimize the methane emission vector $\boldsymbol{x}$.

The simplest approach would be to use the native resolution of the CTM in order to extract the maximum information from the observations. However, the observations may not be sufficiently dense or precise to optimize emissions at that level of detail, resulting in an underdetermined problem. Bocquet et al. (2011) refer to this as the "resolution problem". The inverse solution must then rely on some prior estimate for the state vector and may not be able to depart sufficiently from that knowledge. ~~This~~ The associated error is known as the smoothing error ~~(Rodgers, 2000) and increases~~ (Rodgers, 2000; von Clarmann, 2014) and increase with size of the state vector ~~(Bousquet et al., 2000; Kaminski and Heimann, 2001~~ Wecht et al. (2014) illustrate the severity of this problem in their inversion of methane emissions using satellite data.

An additional drawback of using a large state vector is ~~the computational cost of the inversion. Analytical~~ that analytical solution to the inverse problem may not be computationally tractable. Analytical solution requires calculation of the Jacobian matrix, $\nabla_{\boldsymbol{x}}\mathbf{F}$, and inversion and multiplication of the error covariance matrices (Rodgers, 2000). It has the major advantage of providing complete error statistics ~~on~~ as part of the solution but it becomes impractical as the state vector becomes too large. Numerical solutions using variational methods circumvent this problem but do not ~~inherently~~ provide error characterization as part of the solution. Approximate error statistics can be obtained (e.g., Bousserez et al., 2015) but at the cost of additional computation.

Reducing the dimensionality of the state vector in the inverse problem thus has two advantages. It improves the observational constraints on individual state vector elements and it facilitates analytical solution. Reduction can be achieved by aggregating state vector elements. For a state vector of gridded time-dependent emissions, the state vector can be reduced by aggregating grid cells and time periods. However, this introduces error in the inversion as the underlying spatial and temporal patterns of the aggregated emissions are now imposed from prior knowledge and not allowed to be

optimized as part of the inversion. The resulting error is called the aggregation error (Kaminski and Heimann, 2001; Kaminski et al., 2001; Schuh et al., 2009).

Previous work by Bocquet (2009) , Bocquet et al. (2011) , Bocquet and Wu (2011) , Wu et al. (2011) , and Koohkan et al. (2012) developed optimal grids that allow the transfer of information across multiple scales. These computationally efficient methods (Bocquet and Wu, 2011) generally require the use of the native-resolution grid to derive the optimal representation. They also assume that the native-resolution prior error covariance matrices can be accurately constructed. However, in practice we are generally unable to specify realistic prior error correlations and must resort to simple assumptions.

Here we present a method for optimizing the selection of the state vector in the solution of the inverse problem for a given ensemble of observations ~~.~~ without requiring an accurate specification of the native-resolution prior error covariance matrix. Instead, we use the expected error correlations between native-resolution state vector elements as criteria in the aggregation process. Relative to Bocquet et al. (2011) , our method is suboptimal but is more practical to implement. As the dimension of the state vector decreases, the smoothing error decreases while the aggregation error increases. ~~We show how to derive an optimum~~ There is therefore an optimum dimension where the overall error is minimized. We derive an analytical expression for the aggregation error covariance matrix and show how this can guide selection of a reduced-dimension state vector where the aggregation error remains below an acceptable threshold. We also show how intelligent selection of the state vector can extract more information from the observations for a given state vector dimension.

2 Formulating the inverse problem

Inverse problems are commonly solved using Bayes' theorem,

$$P(\boldsymbol{x}|\boldsymbol{y}) \propto P(\boldsymbol{y}|\boldsymbol{x})P(\boldsymbol{x}) \tag{2}$$

where $P(\boldsymbol{x}|\boldsymbol{y})$ is the posterior probability density function (pdf) of the state vector $\boldsymbol{x}$ ($n \times 1$) given a vector of observations $\boldsymbol{y}$ ($m \times 1$), $P(\boldsymbol{x})$ is the prior pdf of $\boldsymbol{x}$, and $P(\boldsymbol{y}|\boldsymbol{x})$ is the conditional pdf of $\boldsymbol{y}$ given the true value of $\boldsymbol{x}$. Assuming Gaussian distributions for $P(\boldsymbol{y}|\boldsymbol{x})$ and $P(\boldsymbol{x})$ allows us to write the posterior pdf as

$$P(\boldsymbol{x}|\boldsymbol{y}) \propto \exp\left\{-\frac{1}{2}\left(\boldsymbol{y}-\mathbf{F}(\boldsymbol{x})\right)^T \mathbf{S}_{\mathrm{O}}^{-1}\left(\boldsymbol{y}-\mathbf{F}(\boldsymbol{x})\right) - \frac{1}{2}\left(\boldsymbol{x}_{\mathrm{a}}-\boldsymbol{x}\right)^T \mathbf{S}_{\mathrm{a}}^{-1}\left(\boldsymbol{x}_{\mathrm{a}}-\boldsymbol{x}\right)\right\} \tag{3}$$

where $\boldsymbol{x}_{\mathrm{a}}$ is the $n \times 1$ prior state vector, $\mathbf{S}_{\mathrm{O}}$ is the $m \times m$ observational error covariance matrix, and $\mathbf{S}_{\mathrm{a}}$ is the $n \times n$ prior error covariance matrix. Here and elsewhere, our notation and terminology follow that of Rodgers (2000) . The most probable solution $\hat{\boldsymbol{x}}$ (called the maximum a posteriori or MAP) is defined by the maximum of $P(\boldsymbol{x}|\boldsymbol{y})$, i.e., the minimum of the cost function $\mathcal{J}(\boldsymbol{x})$:

$$\mathcal{J}(\boldsymbol{x}) = \frac{1}{2}\left(\boldsymbol{y}-\mathbf{F}(\boldsymbol{x})\right)^T \mathbf{S}_{\mathrm{O}}^{-1}\left(\boldsymbol{y}-\mathbf{F}(\boldsymbol{x})\right) + \frac{1}{2}\left(\boldsymbol{x}_{\mathrm{a}}-\boldsymbol{x}\right)^T \mathbf{S}_{\mathrm{a}}^{-1}\left(\boldsymbol{x}_{\mathrm{a}}-\boldsymbol{x}\right) \tag{4}$$

This involves solving

$$\nabla_{\boldsymbol{x}}\mathcal{J} = \nabla_{\boldsymbol{x}}\mathbf{F}(\boldsymbol{x})^T\mathbf{S}_{\mathrm{O}}^{-1}\left(\mathbf{F}(\boldsymbol{x}) - \boldsymbol{y}\right) + \mathbf{S}_{\mathrm{a}}^{-1}\left(\boldsymbol{x}_{\mathrm{a}} - \boldsymbol{x}\right) = \mathbf{0} \tag{5}$$

Solution to Eq. (5) can be done analytically if $\mathbf{F}$ is linear, i.e., $\mathbf{F}(\boldsymbol{x}) = \mathbf{K}\boldsymbol{x} + \boldsymbol{c}$ where $\mathbf{K} \equiv \nabla_{\boldsymbol{x}}\mathbf{F} = \partial\boldsymbol{y}/\partial\boldsymbol{x}$ is the Jacobian of $\mathbf{F}$ and $\boldsymbol{c}$ is a constant that can be set to zero in the general case by subtracting $\boldsymbol{c}$ from the observations. This yields

$$\hat{\boldsymbol{x}} = \boldsymbol{x}_{\mathrm{a}} + \mathbf{G}\left(\boldsymbol{y} - \mathbf{K}\boldsymbol{x}_{\mathrm{a}}\right) \tag{6}$$

where $\mathbf{G} = \hat{\mathbf{S}}\mathbf{K}^T\mathbf{S}_{\mathrm{O}}^{-1}$ is the gain matrix and $\hat{\mathbf{S}}$ is the posterior error covariance matrix,

$$\hat{\mathbf{S}} = \left(\mathbf{K}^T\mathbf{S}_{\mathrm{O}}^{-1}\mathbf{K} + \mathbf{S}_{\mathrm{a}}^{-1}\right)^{-1} \tag{7}$$

The MAP solution can also be expressed in terms of the true value $\boldsymbol{x}$ as

$$\hat{\boldsymbol{x}} = \boldsymbol{x}_{\mathrm{a}} + \mathbf{A}\left(\boldsymbol{x} - \boldsymbol{x}_{\mathrm{a}}\right) + \mathbf{G}\boldsymbol{\epsilon} \tag{8}$$

where $\mathbf{A}$ is the averaging kernel matrix that measures the error reduction resulting from the observations:

$$\mathbf{A} = \mathbf{G}\mathbf{K} = \mathbf{I} - \hat{\mathbf{S}}\mathbf{S}_{\mathrm{a}}^{-1} \tag{9}$$

and $\mathbf{G}\boldsymbol{\epsilon}$ is the observation error in state space with error covariance matrix $\mathbf{G}\mathbf{S}_{\mathrm{O}}\mathbf{G}^T$. We have assumed here that errors are unbiased, as is standard practice in the inverse modeling literature. An observational error bias $\boldsymbol{b}_{\mathrm{O}}$ would propagate as a bias $\mathbf{G}\boldsymbol{b}_{\mathrm{O}}$ in the solution $\hat{\boldsymbol{x}}$ in Eq. 8.

The analytical solution to the inverse problem thus provides full error characterization as part of the solution. It does require that the forward model be linear. The Jacobian matrix must ~~be generally~~ generally be constructed numerically, requiring n sensitivity simulations with the forward model~~, and subsequent~~. Subsequent matrix operations are also of dimension n. This limits the practical size of the state vector. The matrix operations also depend on the dimension m of the observation vector but this can be easily addressed by splitting that vector into uncorrelated packets, a method known as sequential updating (Rodgers, 2000).

The limitation on the state vector size can be lifted by finding the solution to $\nabla_{\boldsymbol{x}}\mathcal{J} = \mathbf{0}$ numerically, rather than analytically, for example by using the adjoint of the forward model to calculate $\nabla_{\boldsymbol{x}}\mathcal{J}$ iteratively at successive approaches to the solution (e.g., Henze et al., 2007). This variational method allows for optimization of state vectors of any size because the Jacobian is not explicitly constructed. But it only yields the MAP solution, $\hat{\boldsymbol{x}}$, with no error statistics. Several approaches have been presented to obtain approximate error characterization ~~(e.g., Desroziers et al., 2005; Chevallier et al., 2007) but they are~~ (e.g., Courtier et al., 1994; Desroziers et al., 2005; Chevallier et al., 2007; Bousserez et al., 2015) but they can be computationally expensive. An excessively large state vector relative to the strength of the observational constraints also incurs smoothing error, as discussed above.

3 Quantifying aggregation and smoothing errors

The resolution of the forward model (e.g., grid resolution of the CTM) places an upper limit on the dimension for the state vector, which we call the native dimension. As we reduce the dimension of the state vector from this native resolution, the smoothing error decreases while the aggregation error increases. Here we present analytical expressions for the aggregation and smoothing error covariance matrices and show how they can be used to select an optimal state vector dimension.

3.1 Aggregation error

As in Bocquet et al. (2011), we define a restriction (aggregation) operator that maps the native-resolution state vector $\boldsymbol{x}$ of dimension n to a reduced-resolution vector $\boldsymbol{x}_\omega$ of dimension p. We assume a linear restriction operator $\boldsymbol{\Gamma}_\omega$ as a $p \times n$ matrix relating $\boldsymbol{x}_\omega$ to $\boldsymbol{x}$:

$$\boldsymbol{x}_\omega = \boldsymbol{\Gamma}_\omega \boldsymbol{x} \tag{10}$$

Bocquet et al. (2011) provide a detailed analysis of aggregation error for reduced-resolution state vectors. Their analysis relies heavily on the ~~probabilistic~~ construction of a prolongation operator ($\boldsymbol{\Gamma}^\star$) mapping $\boldsymbol{x}_\omega$ back to $\boldsymbol{x}$: $\boldsymbol{x} = \boldsymbol{\Gamma}^\star \boldsymbol{x}_\omega$. However, construction of this prolongation operator is not ~~a well-posed problem because the operator is not~~ unique. We present here a simpler and more ~~robust~~ practical method.

Aggregation error is the error introduced by aggregating state vector elements in the inversion. The relationship between the aggregated elements is not optimized as part of the inversion anymore and instead becomes an unoptimized parameter in the forward model, effectively increasing the forward model error and inhibiting the ability of the model to fit the observations. The aggregation error is thus a component of the observational error.

The aggregation error can be quantified by comparing the observational error incurred by using the native-resolution state vector,

$$\boldsymbol{\epsilon} = \boldsymbol{y} - \mathbf{K}\boldsymbol{x} \tag{11}$$

to that using the aggregated state vector,

$$\boldsymbol{\epsilon}_\omega = \boldsymbol{y} - \mathbf{K}_\omega \boldsymbol{x}_\omega \tag{12}$$

Here $\boldsymbol{y}$ is the observation vector (common in both cases), ~~and~~ $\boldsymbol{x}$ and $\boldsymbol{x}_\omega$ are the true values of the native-resolution and aggregated state vectors, and $\mathbf{K}$ and $\mathbf{K}_\omega$ are the native resolution and the reduced-dimension Jacobians. The only difference between $\boldsymbol{\epsilon}$ and $\boldsymbol{\epsilon}_\omega$ is the aggregation of state vector elements. As such,

$$\boldsymbol{\epsilon}_\omega = \boldsymbol{\epsilon} + \boldsymbol{\epsilon}_\mathrm{A} \tag{13}$$

where $\boldsymbol{\epsilon}_\mathrm{A}$ is the aggregation error. Rearranging,

$$\boldsymbol{\epsilon}_\mathrm{A} = (\mathbf{K} - \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega)\, \boldsymbol{x} \tag{14}$$

Obtaining the error statistics for $\boldsymbol{\epsilon}_\mathrm{A}$ requires knowledge of the pdf of $\boldsymbol{x}$ for the ensemble of possible true states (cf. Rodgers, 2000; von Clarmann, 2014). Let $\overline{\boldsymbol{x}}$ represent the mean value of this ensemble and $\mathbf{S}_\mathrm{e}$ the corresponding covariance matrix. The aggregation error covariance matrix is:

$$\mathbf{S}_\mathrm{A} = E\left[(\boldsymbol{\epsilon}_\mathrm{A} - E[\boldsymbol{\epsilon}_\mathrm{A}])\,(\boldsymbol{\epsilon}_\mathrm{A} - E[\boldsymbol{\epsilon}_\mathrm{A}])^T\right] \tag{15}$$

where $E[\,]$ is the expected value operator. $E[\boldsymbol{\epsilon}_\mathrm{A}] = (\mathbf{K} - \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega)\,\overline{\boldsymbol{x}}$ is the bias introduced by the aggregation. Replacing into Eq. (15):

$$\begin{aligned}\mathbf{S}_\mathrm{A} &= (\mathbf{K} - \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega)\, E\left[(\boldsymbol{x} - \overline{\boldsymbol{x}})\,(\boldsymbol{x} - \overline{\boldsymbol{x}})^T\right] (\mathbf{K} - \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega)^T \\ &= (\mathbf{K} - \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega)\, \mathbf{S}_\mathrm{e}\, (\mathbf{K} - \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega)^T \end{aligned} \tag{16}$$

In designing our inversion system we use $\boldsymbol{x}_\mathrm{a}$ as our best estimate of $\overline{\boldsymbol{x}}$ and $\mathbf{S}_\mathrm{a}$ as our best estimate of $\mathbf{S}_\mathrm{e}$. ~~If~~ Indeed, if $\boldsymbol{x}_\mathrm{a} = \boldsymbol{x}$ there ~~is~~ would be no aggregation error since the prior relationship assumed between state vector elements ~~is~~ would be correct, thus $\mathbf{K} = \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega$ and the aggregation bias ~~is zero. Furthermore, assuming~~ would be zero. Assuming $\mathbf{S}_\mathrm{a} = \mathbf{S}_\mathrm{e}$ allows us to calculate the aggregation error covariance matrix as

$$\mathbf{S}_\mathrm{A} = (\mathbf{K} - \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega)\, \mathbf{S}_\mathrm{a}\, (\mathbf{K} - \mathbf{K}_\omega \boldsymbol{\Gamma}_\omega)^T \tag{17}$$

and we will use this expression in the analysis that follows. Application of Eq. 17 requires computation of the native-resolution Jacobian $\mathbf{K}$ but this can be done for a limited test period only. We will give an example below.

3.2 Smoothing error

Following Rodgers (2000), we can express the smoothing error on $\hat{\boldsymbol{x}}$ by rearranging Eqs. (6) and (1):

$$\hat{\boldsymbol{x}} - \boldsymbol{x} = (\mathbf{I} - \mathbf{A})\,(\boldsymbol{x}_\mathrm{a} - \boldsymbol{x}) + \mathbf{G}\boldsymbol{\epsilon} \tag{18}$$

where $\boldsymbol{\epsilon}_\mathrm{S} = (\mathbf{I} - \mathbf{A})\,(\boldsymbol{x}_\mathrm{a} - \boldsymbol{x})$ is the smoothing error. As pointed out by Rodgers (2000), the smoothing error statistics must be derived from the pdf of possible true states, in the same way as for the aggregation error and characterized by the error covariance matrix $\mathbf{S}_\mathrm{e}$. For purposes of designing the inverse system we assume that $\mathbf{S}_\mathrm{e} = \mathbf{S}_\mathrm{a}$. Thus we have

$$\mathbf{S}_\mathrm{S} = (\mathbf{I} - \mathbf{A})\, \mathbf{S}_\mathrm{a}\, (\mathbf{I} - \mathbf{A})^T \tag{19}$$

We can also express the smoothing error in observation space, $\boldsymbol{\epsilon}^*_\mathrm{S}$, (i.e., as a difference between $\boldsymbol{y}$ and $\mathbf{K}\hat{\boldsymbol{x}}$) by multiplying both sides of Eq. (18) by the Jacobian matrix:

$$\mathbf{K}\,(\hat{\boldsymbol{x}} - \boldsymbol{x}) = \mathbf{K}\,(\mathbf{I} - \mathbf{A})\,(\boldsymbol{x}_\mathrm{a} - \boldsymbol{x}) + \mathbf{K}\mathbf{G}\boldsymbol{\epsilon} \tag{20}$$

so that

$$\boldsymbol{\epsilon}_{\mathrm{S}}^{*} = \mathbf{K}\left(\mathbf{I}-\mathbf{A}\right)\left(\boldsymbol{x}_{\mathrm{a}}-\boldsymbol{x}\right) \tag{21}$$

The corresponding smoothing error covariance matrix in observation space is

$$\mathbf{S}_{\mathrm{S}}^{*} = \mathbf{K}\left(\mathbf{I}-\mathbf{A}\right)\mathbf{S}_{\mathrm{a}}\left(\mathbf{I}-\mathbf{A}\right)^{T}\mathbf{K}^{T} \tag{22}$$

This expression can be generalized to compute the smoothing error covariance matrix in observation space for any reduced-dimension state vector $\boldsymbol{x}_{\omega}$ with Jacobian $\mathbf{K}_{\omega}$, prior error covariance matrix $\mathbf{S}_{\mathrm{a},\omega}$, and averaging kernel matrix $\mathbf{A}_{\omega}$:

$$\mathbf{S}_{\mathrm{S}}^{*} = \mathbf{K}_{\omega}\left(\mathbf{I}-\mathbf{A}_{\omega}\right)\mathbf{S}_{\mathrm{a},\omega}\left(\mathbf{I}-\mathbf{A}_{\omega}\right)^{T}\mathbf{K}_{\omega}^{T} \tag{23}$$

3.3 Total error budget

From Eq. (18) we can see that the total error on $\hat{\boldsymbol{x}}$ without aggregation is $\boldsymbol{\epsilon}_{\mathrm{T}} = \boldsymbol{\epsilon}_{\mathrm{S}} + \mathbf{G}\boldsymbol{\epsilon}$ in the state space, or $\boldsymbol{\epsilon}_{\mathrm{T}}^{*} = \boldsymbol{\epsilon}_{\mathrm{S}}^{*} + \mathbf{KG}\boldsymbol{\epsilon}$ in the observation space. The $\mathbf{KG}$ term in the observation space appears because we are interested in the error on $\hat{\boldsymbol{x}}$. If $\hat{\boldsymbol{x}} = \boldsymbol{x}$ then $\mathbf{KG} = \mathbf{I}$ and $\mathbf{A} = \mathbf{I}$, thus $\boldsymbol{\epsilon}_{\mathrm{S}} = \mathbf{0}$ and our total error reverts to $\boldsymbol{\epsilon}$,

$$\boldsymbol{\epsilon}_{\mathrm{T}}^{*}|_{\hat{\boldsymbol{x}}=\boldsymbol{x}} = \mathbf{K}\left(\mathbf{I}-\mathbf{A}\right)\left(\boldsymbol{x}_{\mathrm{a}}-\boldsymbol{x}\right) + \mathbf{KG}\boldsymbol{\epsilon} = \boldsymbol{\epsilon} \tag{24}$$

Additional consideration of aggregation error for a reduced-dimension state vector $\boldsymbol{x}_{\omega}$ yields a total error in the state space

$$\boldsymbol{\epsilon}_{\mathrm{T}} = \boldsymbol{\epsilon}_{\mathrm{S}} + \mathbf{G}_{\omega}\boldsymbol{\epsilon} + \mathbf{G}_{\omega}\boldsymbol{\epsilon}_{\mathrm{A}} \tag{25}$$

where ~~$\mathbf{G}_{\omega}$~~

$$\mathbf{G}_{\omega} = \left(\mathbf{K}_{\omega}^{T}\mathbf{S}_{\mathrm{O}}^{-1}\mathbf{K}_{\omega} + \mathbf{S}_{\mathrm{a},\omega}^{-1}\right)^{-1}\mathbf{K}_{\omega}^{T}\mathbf{S}_{\mathrm{O}}^{-1} \tag{26}$$

is the gain matrix for the reduced-dimension state vector. In the observation space we get

$$\boldsymbol{\epsilon}_{\mathrm{T}}^{*} = \boldsymbol{\epsilon}_{\mathrm{S}}^{*} + \mathbf{K}_{\omega}\mathbf{G}_{\omega}\boldsymbol{\epsilon} + \mathbf{K}_{\omega}\mathbf{G}_{\omega}\boldsymbol{\epsilon}_{\mathrm{A}} \tag{27}$$

From these relationships we derive the total error covariance matrix as

$$\begin{aligned}\mathbf{S}_{\mathrm{T},\omega} &= \underbrace{\left(\mathbf{I}-\mathbf{A}_{\omega}\right)\mathbf{S}_{\mathrm{a},\omega}\left(\mathbf{I}-\mathbf{A}_{\omega}\right)^{T}}_{\text{Smoothing Error}} \\ &+ \underbrace{\mathbf{G}_{\omega}\left(\mathbf{K}-\mathbf{K}_{\omega}\boldsymbol{\Gamma}_{\omega}\right)\mathbf{S}_{\mathrm{a}}\left(\mathbf{K}-\mathbf{K}_{\omega}\boldsymbol{\Gamma}_{\omega}\right)^{T}\mathbf{G}_{\omega}^{T}}_{\text{Aggregation Error}} \\ &+ \underbrace{\mathbf{G}_{\omega}\mathbf{S}_{\mathrm{O}}\mathbf{G}_{\omega}^{T}}_{\text{Observation Error}}\end{aligned} \tag{28}$$

in the state space and

$$\mathbf{S}^{*}_{\mathrm{T},\omega} = \underbrace{\mathbf{K}_\omega \left(\mathbf{I}-\mathbf{A}_\omega\right)\mathbf{S}_{\mathrm{a},\omega}\left(\mathbf{I}-\mathbf{A}_\omega\right)^T\mathbf{K}_\omega^T}_{\text{Smoothing Error}} + \underbrace{\mathbf{K}_\omega\mathbf{G}_\omega\left(\mathbf{K}-\mathbf{K}_\omega\mathbf{\Gamma}_\omega\right)\mathbf{S}_\mathrm{a}\left(\mathbf{K}-\mathbf{K}_\omega\mathbf{\Gamma}_\omega\right)^T\mathbf{G}_\omega^T\mathbf{K}_\omega^T}_{\text{Aggregation Error}} + \underbrace{\mathbf{K}_\omega\mathbf{G}_\omega\mathbf{S}_\mathrm{O}\mathbf{G}_\omega^T\mathbf{K}_\omega^T}_{\text{Observation Error}} \tag{29}$$

in the observation space. A bias term should exhibit similar scale-dependence to the observation error term and could be included by following the derivation from Rodgers (2000) .

Each of ~~these~~ the three error terms ~~depends~~ above depend on state vector dimension. Because the smoothing error increases with state vector dimension while the aggregation error decreases, ~~we expect to find an~~ analysis of the error budget can point to the optimal dimension where the total error is minimum. ~~To enable an analytical inversion we may wish to decrease the~~ It can also point to the minimum state vector dimension ~~further within a tolerance on aggregation error, such as requiring that~~ needed for the aggregation error ~~remain~~ to be below a certain tolerance, e.g., smaller than the observation error. We give an example in Sect. 5.

A caveat in the above expressions for the aggregation and smoothing error covariance matrices is that they are valid only if the prior $\boldsymbol{x}_\mathrm{a}$ is the mean value $\overline{\boldsymbol{x}}$ for the pdf of true states and if the error covariance matrix $\mathbf{S}_\mathrm{a}$ is the covariance matrix for that pdf ~~. These conditions define the assumption for the prior, so the expressions~~ ($\mathbf{S}_\epsilon = \mathbf{S}_\mathrm{a}$). Rodgers (2000, p. 49) and von Clarmann (2014) provide a detailed discussion of the errors induced by failing to meet this assumption. Since these assumptions define our prior, they can be taken as valid for the purpose of selecting an appropriate state vector dimension in an inverse problem. However, they should not be used to diagnose errors on the inversion results.

4 Aggregation methods

Aggregation of state vector elements to reduce the state vector dimension introduces aggregation ~~errors~~error, as described in Sect. 3.1. The aggregation error can be reduced by grouping elements with correlated errors. Analyzing the off-diagonal structure of a precisely constructed prior error correlation matrix would provide the best objective way to carry out the aggregation, as described by Bocquet (2009) , Bocquet et al. (2011) , and Wu et al. (2011) . We generally lack such information but do have some qualitative knowledge of prior error correlation that can be used to optimize the aggregation. ~~Bocquet et al. (2011)~~ By aggregating regions that have correlated errors we can exploit additional information that would otherwise be neglected in a native-resolution inversion assuming (by default) uncorrelated errors.

Previous work by Bocquet et al. (2011) , Wu et al. (2011) , and Koohkan et al. (2012) used tiling and tree-based aggregation methods, while Wecht et al. (2014) used a ~~-~~hierarchal clustering method based on prior error patterns. Bocquet and Wu (2011) also used principal component analysis (PCA) coupled to the hierarchal grid to compute an optimal grid. Here we compare three aggregation methods: (1) simple grid coarsening, (2) ~~principal component analysis (PCA)~~ PCA clustering, and (3) a Gaussian mixture model (GMM) with radial basis functions (RBFs) to project native-resolution state vector elements to Gaussian pdfs. A qualitative illustration of these methods is shown in Fig. 1 for the aggregation of a native-resolution state vector of methane emissions with $\frac{1}{2}^\circ \times \frac{2}{3}^\circ$ native grid resolution over North America (Turner et al., 2015). We focus here on spatial aggregation and assume that the state vector has no temporal dimension. However, the same methods can be used for temporal aggregation.

The simplest method for reducing the dimension of the state vector is to merge adjacent elements, i.e., neighboring grid cells. This method considers only spatial proximity as source of error correlation. It may induce large aggregation errors if proximal but otherwise dissimilar regions are aggregated together. In the case of methane emissions, aggregating neighboring wetlands and farmland would induce large errors because different processes drive methane emissions from these two source types.

The other two methods enable consideration of additional similarity factors besides spatial proximity when aggregating state vector elements. These similarity factors are expressed by vectors of dimension n describing correlative properties of the original native-resolution state vector elements. In the case of a methane source inversion, for example, we can choose as similarity vectors latitude and longitude to account for spatial proximity, but also wetland fraction to account for error correlations in the bottom-up wetland emission estimate used as prior.

4.1 Similarity matrix for aggregation

Table 1 lists the similarity vectors chosen for our example problem of estimating methane emissions (Turner et al., 2015). The first two vectors account for spatial proximity, the third represents the scaling factors from the first iteration of an adjoint-based inversion at native resolution (Wecht et al., 2014), and the others are the source type patterns from the bottom-up inventories used as prior. All similarity vectors are normalized and then weighted by judgment of their importance. We choose here to include initial scaling factors from the adjoint-based inversion because we have them available and they can serve to correct any prior patterns that are grossly inconsistent with the observations, or to identify local emission hotspots missing from the prior. One iteration of the adjoint-based inversion is computationally inexpensive and is sufficient to pick up major departures from the prior.

Let $\{\boldsymbol{c}_1, \ldots, \boldsymbol{c}_K\}$ represent the K similarity vectors chosen for the problem ($K = 14$ in our example of Table 1). We assemble them into a $n \times K$ similarity matrix $\mathbf{C}$. We will also make use of the

ensemble of similarity vector values for individual state vector elements, which we assemble into vectors $\{\boldsymbol{c}'_1, \ldots, \boldsymbol{c}'_n\}$ representing the rows of $\mathbf{C}$. Thus:

$$\mathbf{C} = \left[\begin{pmatrix} \vdots \\ \boldsymbol{c}_1 \\ \vdots \end{pmatrix} \quad \begin{pmatrix} \vdots \\ \boldsymbol{c}_2 \\ \vdots \end{pmatrix} \quad \cdots \quad \begin{pmatrix} \vdots \\ \boldsymbol{c}_{\mathrm{K}} \\ \vdots \end{pmatrix} \right] = \begin{bmatrix} \begin{pmatrix} \cdots & \boldsymbol{c}'_1 & \cdots \end{pmatrix} \\ \begin{pmatrix} \cdots & \boldsymbol{c}'_2 & \cdots \end{pmatrix} \\ \vdots \\ \begin{pmatrix} \cdots & \boldsymbol{c}'_n & \cdots \end{pmatrix} \end{bmatrix} \tag{30}$$

In this work all of the aggregation methods except for grid coarsening will use the same similarity matrix to construct the restriction operator.

This approach of using a similarity matrix $\mathbf{C}$ to account for prior error covariances bears some resemblance to the geostatistical approach for inverse modeling (e.g., Michalak et al., 2004, 2005; Gourdji et al., 2008; Miller et al., 2012). The geostatistical approach specifies the prior estimate as $\boldsymbol{x}_{\mathrm{a}} = \mathbf{C}\boldsymbol{\beta}$ where $\boldsymbol{\beta}$ is a vector of unknown drift coefficients to be optimized as part of the inversion. Here we use the similarity matrix to reduce the dimension of the state vector, rather than just as a choice of prior constraints.

4.2 Clustering with principal component analysis

In this method we cluster state vector elements following the principal components of the similarity matrix. It is generally not practical to derive the principal components in state vector space because the n-dimension is large. Instead we derive them in in similarity space (dimension K) as the eigenvectors of $\mathbf{C}^T\mathbf{C}$ sorted in order of importance by their eigenvalues. The leading j principal components are kept for clustering. The reduced state vector is then constructed by grouping state vector elements that have the same sign patterns for all j principal components. Each unique j-dimensional sign pattern constitutes a cluster. The number of clusters defined in that way ranges between j and 2^j. Figure 1b shows an example of applying this method to methane emissions in North America with reduction of the state vector to $n = 8$. The separation into four quadrants reflects the importance of latitude and longitude as error correlation factors. The additional separation within each quadrant isolates large from weak sources as defined by the prior.

4.3 Gaussian mixture model (GMM)

Here we use a Gaussian mixture model (GMM; Bishop, 2007) to project the native-resolution state vector onto p Gaussian pdfs using radial basis functions (RBFs). Mixture models are probabilistic models for representing a population comprised of p subpopulations. Each subpopulation is assumed to follow a pdf, in this case Gaussian. The Gaussians are K-dimensional where K is the number of similarity criteria. Each native-resolution state vector element is fit to this ensemble of Gaussians using RBFs as weighting factors.

The first step in constructing the GMM is to define a $p \times n$ weighting matrix $\mathbf{W} = [\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_p]^T$. Each element $w_{i,j}$ of this weighting matrix is the relative probability for native-resolution state vector element j to be described by Gaussian subpopulation i, i.e., "how much does element j look like Gaussian i?". It is given by

$$w_{i,j} = \frac{\pi_i \mathcal{N}(\boldsymbol{c}'_j | \boldsymbol{\mu}_i, \boldsymbol{\Lambda}_i)}{\sum_{k=1}^{p} \pi_j \mathcal{N}(\boldsymbol{c}'_j | \boldsymbol{\mu}_k, \boldsymbol{\Lambda}_k)} \tag{31}$$

Here $\boldsymbol{c}'_j$ is the jth row of the similarity matrix $\mathbf{C}$, $\boldsymbol{\mu}_i$ is a $1 \times K$ row vector of means for the ith Gaussian, $\boldsymbol{\Lambda}_i$ is a $K \times K$ covariance matrix for the ith Gaussian, and $\boldsymbol{\pi} = [\pi_1, \ldots, \pi_p]^T$ is the relative weight of the p Gaussians in the mixture. $\mathcal{N}\left(\boldsymbol{c}'_j | \boldsymbol{\mu}_i, \boldsymbol{\Lambda}_i\right)$ denotes the probability density of vector $\boldsymbol{c}'_j$ on the normal distribution of Gaussian i. We define a $p \times K$ matrix $\boldsymbol{\mathcal{M}}$ with rows $\boldsymbol{\mu}_i$ and a $K \times K \times p$ third-order tensor $\boldsymbol{\mathcal{L}} = [\boldsymbol{\Lambda}_1, \ldots, \boldsymbol{\Lambda}_p]$ as the set of covariance matrices.

Projection of the native-resolution state vector onto the GMM involves four unknowns: $\mathbf{W}$, $\boldsymbol{\pi}$, $\boldsymbol{\mathcal{M}}$, and $\boldsymbol{\mathcal{L}}$. This is solved by constructing a cost function to estimate the parameters of the Gaussians in the mixture model using maximum likelihood:

$$\mathcal{J}_{\text{GMM}}(\mathbf{C} | \boldsymbol{\pi}, \boldsymbol{\mathcal{M}}, \boldsymbol{\mathcal{L}}) = \sum_{j=1}^{n} \ln \left\{ \sum_{i=1}^{p} \pi_i \mathcal{N}(\boldsymbol{c}'_j | \boldsymbol{\mu}_i, \boldsymbol{\Lambda}_i) \right\} \tag{32}$$

Starting from an initial guess for $\boldsymbol{\pi}$, $\boldsymbol{\mathcal{M}}$, and $\boldsymbol{\mathcal{L}}$ we compute the weight matrix $\mathbf{W}$ using Eq. (31). We then differentiate the cost function with respect to $\boldsymbol{\pi}$, $\boldsymbol{\mathcal{M}}$, and $\boldsymbol{\mathcal{L}}$, and set the derivative to zero to obtain (see Bishop, 2007):

$$\boldsymbol{\mu}_i = \Psi_i \sum_{j=1}^{n} w_{i,j} \boldsymbol{c}'_j \tag{33}$$

$$\boldsymbol{\Lambda}_i = \Psi_i \sum_{j=1}^{n} w_{i,j} \left(\boldsymbol{c}'_j - \boldsymbol{\mu}_i\right)^T \left(\boldsymbol{c}'_j - \boldsymbol{\mu}_i\right) \tag{34}$$

$$\pi_i = \frac{1}{n \Psi_i} \tag{35}$$

where:

$$\Psi_i = \sum_{j=1}^{n} \frac{1}{w_{i,j}} \tag{36}$$

The weights are re-calculated from the updated guesses of $\mathbf{W}$, $\boldsymbol{\pi}$, $\boldsymbol{\mathcal{M}}$, and $\boldsymbol{\mathcal{L}}$ from Eqs. (33)–(36), and so on until convergence. The final weights define the restriction operator as $\boldsymbol{\Gamma}_\omega = \mathbf{W}$. The computational complexity for the expectation-maximization algorithm is $\mathcal{O}(nK + pn^2)$ (Chen et al., 2007), however the actual runtime will be largely dictated by the convergence criteria. Here we use an

absolute tolerance of $\tau < 10^{-10}$ where

$$\tau = \sum_i \sum_j |\mathcal{M}_{i,j} - \mathcal{M}^\star_{i,j}| \tag{37}$$

$$+ \sum_i \sum_j \sum_k |\mathcal{L}_{i,j,k} - \mathcal{L}^\star_{i,j,k}| \tag{38}$$

$$+ \sum_i |\pi_i - \pi^\star_i| \tag{39}$$

and the superscript star indicates the value from the previous iteration.

The GMM allows each native-resolution state vector element to be represented by a unique linear combination of the Gaussians through the RBFs. For a state vector of a given dimension, defined by the number of Gaussian pdfs, we can achieve high resolution for large localized sources by sacrificing resolution for weak or uniform source regions where resolution is not needed. This is illustrated in Fig. 2 with the resolution of ~~southern~~ Southern California in an inversion of methane sources for North America. The figure shows the three dominant Gaussians describing emissions in Southern California and the corresponding RBF weights for each native-resolution grid square. Gaussian 1 is centered over Los Angeles and is highly localized, Gaussian 2 covers the Los Angeles Basin, and Gaussian 3 is a Southern California background. The sum of these three Gaussians accounts for most of the emissions in ~~southern~~ Southern California and Nevada (which is mostly background). Additional Gaussians (not shown) resolve the southern San Joaquin Valley (large livestock and oil/gas emissions) and Las Vegas (large emissions from waste).

5 Application

We apply the aggregation methods described above to our example problem of estimating methane emissions from satellite observations of methane concentrations, focusing on selecting a reduced-dimension state vector that minimizes aggregation and smoothing errors. The inversion is described in detail in Turner et al. (2015) and uses GOSAT satellite observations for 2009–2011 over North America. The forward model for the inversion is the GEOS-Chem CTM with $\frac{1}{2}^\circ \times \frac{2}{3}^\circ$ grid resolution. The native-resolution state vector of methane emissions as defined on that grid includes 7366 elements.

For purpose of selecting an aggregated state vector for the inversion we consider a subset of observations for May 2010 ($m = 6070$) so that we can afford to construct the corresponding Jacobian matrix $\mathbf{K}$ at the native resolution; this is necessary to derive the aggregation error covariance matrix following Eq. (17). The prior error covariance matrix is specified as diagonal with 100 % uncertainty at the native resolution, decreasing with aggregation following the central limit theorem (Turner et al., 2015). The observational error covariance matrix is also diagonal and specified as the scene-specific retrieval error from Parker et al. (2011), which dominates the total observational

error as shown by Turner et al. (2015). We compare the three methods presented in Sect. 4 for aggregating the state vector in terms of the implications for aggregation and smoothing errors for different state vector dimensions. In addition to the GMM with RBFs, we also consider a "GMM clustering" method where each native resolution state vector element is assigned exclusively to its dominant Gaussian pdf. This yields sharp boundaries between clusters (Fig. 1) as in the grid coarsening and PCA methods.

Figure 3 shows the mean error ~~SD~~ standard deviation in the aggregation and smoothing error covariance matrices, computed as the square root of the mean of the diagonal terms, as a function of state vector dimension. The aggregation error is zero by definition at the native resolution (7366 state vector elements), and increases as the number n of state vector elements decreases, following a roughly $n^{-0.7}$ dependence. Conversely, the smoothing error increases as the number of state vector elements increases, following roughly a $\log(n)$ dependence. The different aggregation methods of Sect. 4 yield very similar smoothing errors~~but the aggregation~~, suggesting that any reasonable aggregation scheme (such as k-means clustering (cf. Bishop, 2007)) would perform comparably. The aggregation error is somewhat improved using the GMM method. RBF weighting performs slightly better than GMM clustering (sharp boundaries). As discussed above, a major advantage of the GMM method is its ability to retain resolution of large localized sources after aggregation.

Figure 4 shows the sum of contributions from aggregation, smoothing, and observational error ~~SD~~ standard deviations as a function of state vector aggregation using the GMM with RBF weighting. In this application, aggregation error dominates for small state vectors ($n < 100$), but drops below the observation error for $n > 100$ and below the smoothing error for $n > 1000$. The smoothing error remains smaller than the observational error even at the native resolution ($n = 7366$). The observational error is not independent of aggregation, as shown in Eq. (29), but we find here that the dependence is small.

From Fig. 4 we can identify a state vector dimension for which the total error is minimum ($n = 2208$; circle in Fig. 4). However, error growth is small until $n \approx 200$, below which the aggregation error grows rapidly. A state vector of 369 elements, as adopted by Turner et al. (2015), does not incur significant errors associated with aggregation or smoothing, and enables computation of an analytical solution to the inverse problem with full error characterization.

Previous work by Bocquet (2009) , Bocquet et al. (2011) , Bocquet and Wu (2011) , Wu et al. (2011) , and Koohkan et al. (2012) analyzed the scale-dependence of different grids using the degrees of freedom for signal: $\mathrm{DFS} = \mathrm{Tr}(\mathbf{I} - \mathbf{S}_{a,w}^{-1}\hat{\mathbf{S}}_w)$. These past works found this error metric to be monotonically increasing. This implies that the native resolution grid will have the least total error and there is no optimal resolution, except from a numerical efficiency standpoint. Here we find a local minimum that is, seemingly, at odds with this previous work. However, the reasoning for this local minimum is that we have allowed the aggregation to account for spatial error correlations that we are unable to specify at the native resolution. As such, we are taking more information into account and obtaining

a minimum total error at a state vector size that is smaller than the native resolution. If the native resolution error covariance matrices were correct then, as previous work showed, the only reason to perform aggregation would be to reduce the computational expense and the grid used here would be suboptimal because it does not depend on the native-resolution grid.

6 Conclusions

We presented a method for optimizing the selection of the state vector in the solution of the inverse problem for a given ensemble of observations. The optimization involves minimizing the total error in the inversion by balancing the aggregation error (which increases as the state vector dimension decreases), the smoothing error (which increases as the state vector dimension increases), and the observational error. We further showed how one can reduce the state vector dimension within the constraints from the aggregation error in order to facilitate an analytical solution to the inverse problem with full error characterization.

We explored different methods for aggregating state vector elements as a means of reducing the dimension of the state vector. Aggregation error can be minimized by grouping state vector elements with the strongest correlated prior errors. We showed that a Gaussian mixture model (GMM), where the state vector elements are multi-dimensional Gaussian pdfs constructed from prior error correlation patterns, is a powerful aggregation tool. Reduction of the state vector dimension using the GMM retains fine-scale resolution of important features in the native-resolution state vector while merging weak or uniform features.

Acknowledgements. For advice and discussions, we thank Kevin Wecht (Harvard University). Special thanks to Robert Parker and Hartmut Boesch (University of Leicester) for providing the GOSAT observations. This work was supported by the NASA Carbon Monitoring System and by a Department of Energy (DOE) Computational Science Graduate Fellowship (CSGF) to AJT. We ~~also~~ thank the Harvard SEAS Academic Computing center for access to computing resources. We also thank Dr. Marc Bocquet and an anonymous reviewer for their thorough comments.

References

Bishop, C. M.: Pattern Recognition and Machine Learning, Springer, 1st Edn., New York, 2007.

Bocquet, M.: Towards optimal choices of control space representation for geophysical data assimilation, Mon. Wea. Rev., 137, 2331–2348, doi:10.1175/2009MWR2789.1, 2009.

Bocquet, M., Wu, L., and Chevallier, F.: Bayesian design of control space for optimal assimilation of observations. Part I: Consistent multiscale formalism, Q. J. Roy. Meteor. Soc., 137, 1340–1356, doi:10.1002/qj.837, 2011.

Bocquet, M. and Wu, L.: Bayesian design of control space for optimal assimilation of observations. II: Asymptotics solution, Q. J. Roy. Meteor. Soc., 137, 1357–1368, doi:10.1002/qj.841, 2011.

Bousserez, N., Henze, D. K., Perkins, A., Bowman, K. W., Lee, M., Liu, J., Deng, F., and Jones, D. B. A.: Improved analysis-error covariance matrix for high-dimensional variational inversions: application to source estimation using a 3D atmospheric transport model, Q. J. Roy. Meteor. Soc., n/a–n/a, doi:10.1002/qj.2495, 2015.

Bousquet, P., Peylin, P., Ciais, P., Le Quere, C., Friedlingstein, P., and Tans, P. P.: Regional changes in carbon dioxide fluxes of land and oceans since 1980, Science, 290, 1342–1346, doi:10.1126/Science.290.5495.1342, 2000.

Chen, Z., Haykin, S., Eggermont, J. J., and Becker, S.: Correlative Learning: A Basis for Brain and Adaptive Systems, John Wiley & Sons, 1st Edn., New York, 2007.

Chevallier, F., Breon, F. M., and Rayner, P. J.: Contribution of the Orbiting Carbon Observatory to the estimation of CO_2 sources and sinks: theoretical study in a variational data assimilation framework, J. Geophys. Res.-Atmos., 112, D09307, doi:10.1029/2006jd007375, 2007.

Courtier, P., Thepaut, J., and Hollingsworth, A.: A strategy for operational implementation of 4D-Var, using an incremental approach, Q. J. Roy. Meteor. Soc., 120, 1367–1387, doi:10.1002/qj.49712051912, 1994.

Desroziers, G., Berre, L., Chapnik, B., and Poli, P.: Diagnosis of observation, background and analysis-error statistics in observation space, Q. J. Roy. Meteor. Soc., 131, 3385–3396, doi:10.1256/qj.05.108, 2005.

Gourdji, S. M., Mueller, K. L., Schaefer, K., and Michalak, A. M.: Global monthly averaged CO_2 fluxes recovered using a geostatistical inverse modeling approach: 2. Results including auxiliary environmental data, J. Geophys. Res., 113, D21115, doi:10.1029/2007jd009733, 2008.

Henze, D. K., Hakami, A., and Seinfeld, J. H.: Development of the adjoint of GEOS-Chem, Atmos. Chem. Phys., 7, 2413–2433, doi:10.5194/acp-7-2413-2007, 2007.

Kaminski, T. and Heimann, M.: Inverse modeling of atmospheric carbon dioxide fluxes, Science, 294, 259, doi:10.1126/science.294.5541.259a, 2001.

Kaminski, T., Rayner, P. J., Heimann, M., and Enting, I. G.: On aggregation errors in atmospheric transport inversions, J. Geophys. Res., 106, 4703, doi:10.1029/2000jd900581, 2001.

Koohkan, M. R., Bocquet, M., Wu, L., and Krysta, M.: Potential of the International Monitoring System radionuclide network for inverse modelling, Atmos. Env., 54, 557–567, doi:10.1016/j.atmosenv.2012.02.044, 2012.

Michalak, A. M., Bruhwiler, L., and Tans, P. P.: A geostatistical approach to surface flux estimation of atmospheric trace gases, J. Geophys. Res., 109, D14109, doi:10.1029/2003jd004422, 2004.

Michalak, A. M., Hirsch, A., Bruhwiler, L., Gurney, K. R., Peters, W., and Tans, P. P.: Maximum likelihood estimation of covariance parameters for Bayesian atmospheric trace gas surface flux inversions, J. Geophys. Res., 110, D24107, doi:10.1029/2005jd005970, 2005.

Miller, S. M., Kort, E. A., Hirsch, A. I., Dlugokencky, E. J., Andrews, A. E., Xu, X., Tian, H., Nehrkorn, T., Eluszkiewicz, J., Michalak, A. M., and Wofsy, S. C.: Regional sources of nitrous oxide over the United States: seasonal variation and spatial distribution, J. Geophys. Res., 117, D06310, doi:10.1029/2011jd016951, 2012.

Parker, R., Boesch, H., Cogan, A., Fraser, A., Feng, L., Palmer, P. I., Messerschmidt, J., Deutscher, N., Griffith, D. W. T., Notholt, J., Wennberg, P. O., and Wunch, D.: Methane observations from the Greenhouse Gases Observing SATellite: comparison to ground-based TCCON data and model calculations, Geophys. Res. Lett., 38, L15807, doi:10.1029/2011gl047871, 2011.

Rodgers, C. D.: Inverse Methods for Atmospheric Sounding, World Scientific, Singapore, 2000.

Schuh, A. E., Denning, A. S., Uliasz, M., and Corbin, K. D.: Seeing the forest through the trees: recovering large-scale carbon flux biases in the midst of small-scale variability, J. Geophys. Res., 114, G03007, doi:10.1029/2008jg000842, 2009.

Turner, A. J., Jacob, D. J., Wecht, K., Maasakkers, J. D., Biraud, S. C., Boesch, H., Bowman, K. W., Deutscher, N. M., Dubey, M. K., Griffith, D. W. T., Hase, F., Kuze, A., Notholt, J., Ohyama, H., Parker, R., Payne, V. H., Sussmann, R., Velazco, V. A., Warneke, T., Wennberg, P. O., and Wunch, D.: Estimating global and North American methane emissions with high spatial resolution using GOSAT satellite data, Atmos. Chem. Phys. Discuss., 15, 4495–4536, doi:10.5194/acpd-15-4495-2015, ~~available at: , submitted, 2014.~~ 2015.

von Clarmann, T.: Smoothing error pitfalls, Atmos. Meas. Tech., 7, 3023–3034, doi:10.5194/amt-7-3023-2014, 2014.

Wecht, K. J., Jacob, D. J., Frankenberg, C., Jiang, Z., and Blake, D. R.: Mapping of North American methane emissions with high spatial resolution by inversion of SCIAMACHY satellite data, J. Geophys. Res.-Atmos., 119, 7741–7756, doi:10.1002/2014jd021551, 2014.

Wu, L., Bocquet, M., Lauvaux, T., Chevallier, F., Rayner, P., and Davis, K.: Optimal representation of source-sink fluxes for mesoscale carbon dioxide inversion with synthetic data, J. Geophys. Res., 116, D21304, doi:10.1029/2011jd016198, 2011.

Table 1. Similarity vectors for inverting methane emissions in North America[a].

Similarity Vector		Weighting factor[b]
1.	Latitude[c]	1.00
2.	Longitude[d]	1.00
3.	Initial scaling factors[e]	0.15
4.	Wetland [f]	0.31
5.	Livestock [f]	0.22
6.	Oil/gas [f]	0.16
7.	Waste [f]	0.15
8.	Coal [f]	0.06
9.	Soil absorption [f]	0.05
10.	Termites [f]	0.02
11.	Biomass burning [f]	0.02
12.	Biofuel [f]	0.01
13.	Rice [f]	0.01
14.	Other [f]	0.01

[a] The $K = 14$ similarity vectors describe prior error correlation criteria for the native-resolution state vector, representing here the methane emission in North America at the $\frac{1}{2}^\circ \times \frac{2}{3}^\circ$ resolution of the GEOS-Chem chemical transport model. The criteria are normalized and then weighted (weighting factor). Criteria 4-14 are prior emission patterns used in the GEOS-Chem model (Wecht et al., 2014; Turner et al., 2015) .

[b] The weighting factors (dimensionless) measure the estimated relative importance of the different similarity criteria in determining prior error correlations in the state vector. For the prior emission patterns these weighting factors are the fractional contributions to total prior emissions in North America.

[c] Distance in kilometers from the equator.

[d] Distance in kilometers from the prime meridian.

[e] Initial scaling factors from one iteration of an adjoint inversion at the native resolution.

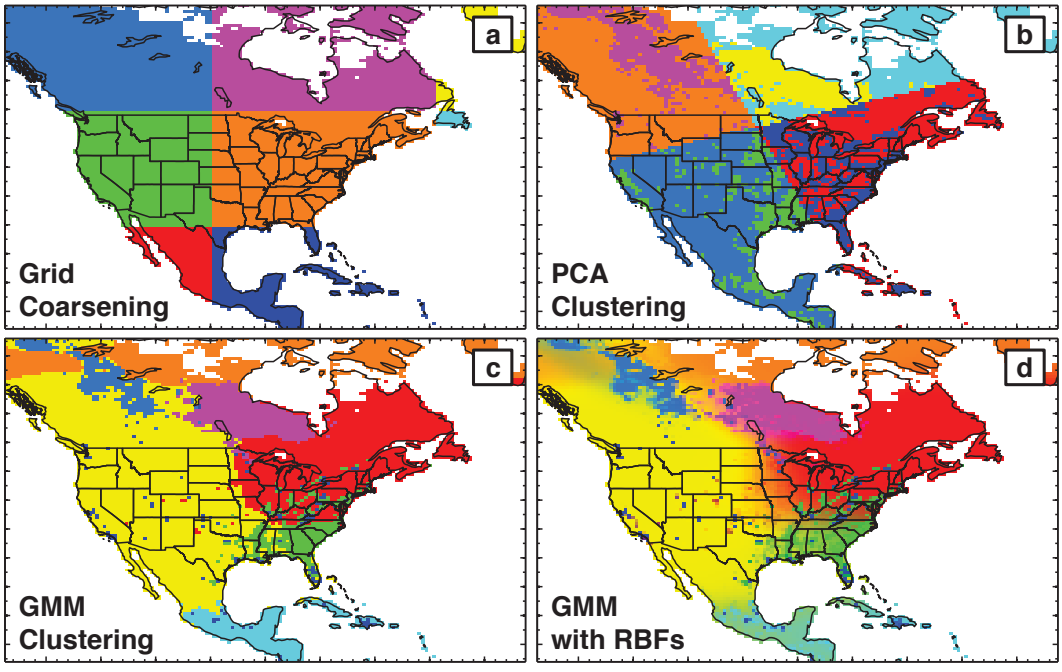

Figure 1. Illustration of different approaches for aggregating a state vector. Here the native resolution state vector is a field of gridded methane emissions at $\frac{1}{2}^\circ \times \frac{2}{3}^\circ$ resolution over North America. Extreme reduction to 8 state vector elements is shown with individual elements distinguished by color.

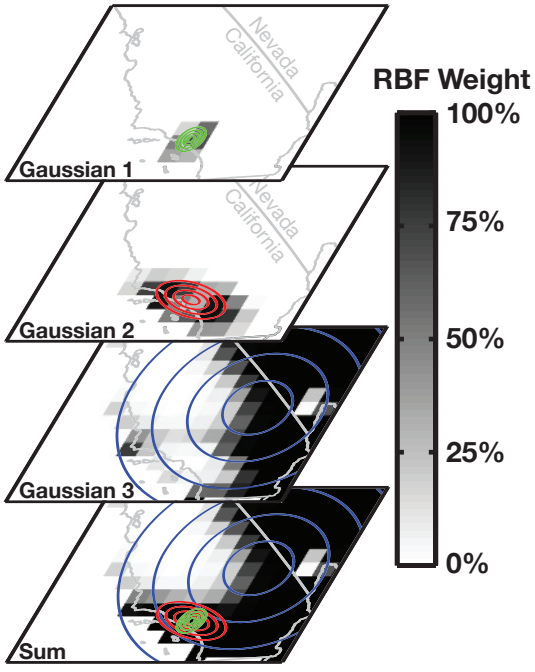

Figure 2. Gaussian Mixture Model (GMM) representation of methane emissions in ~~southern~~ Southern California with Gaussian pdfs as state vector elements. The Gaussians are constructed from a similarity matrix for methane emissions on the $\frac{1}{2}^\circ \times \frac{2}{3}^\circ$ horizontal resolution of the GEOS-Chem CTM used as forward model for the inversion. The figure shows the dominant three Gaussians for ~~southern~~ Southern California with contours delineating the 0.5, 1.0, 1.5, and 2.0 σ spreads for the latitude–longitude dimensions. The RBF weights $\boldsymbol{w}_1$, $\boldsymbol{w}_2$, and $\boldsymbol{w}_3$ of the three Gaussians for each $\frac{1}{2}^\circ \times \frac{2}{3}^\circ$ grid square are also shown along with their sum.

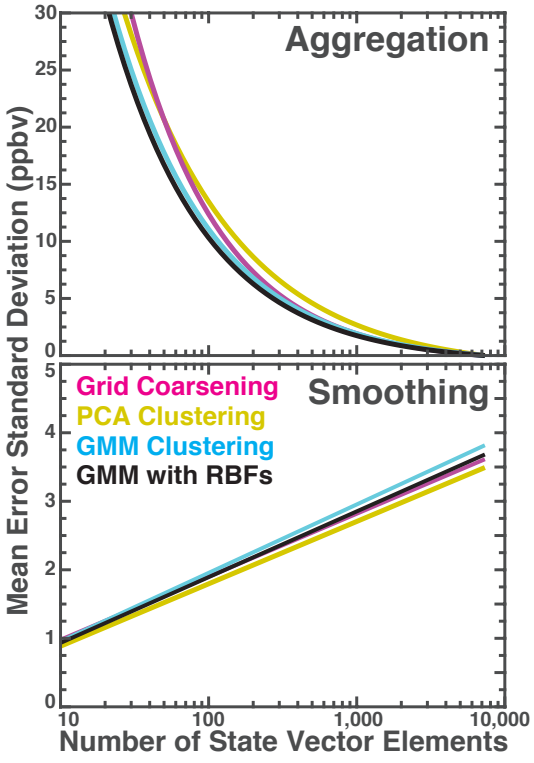

Figure 3. Aggregation and smoothing error dependences on the aggregation of state vector elements in an inverse model. The application here is to an inversion of methane emissions over North America using satellite methane data with 7366 native-resolution state vector elements (Sect. 5 and Turner et al., 2015). Results are shown as the square roots of the means of the diagonal terms (mean error ~~SD~~standard deviation) in the aggregation and smoothing error covariance matrices. Different methods for aggregating the state vector (Sect. 4) are shown as separate lines. Note the log-scale on the x axis.

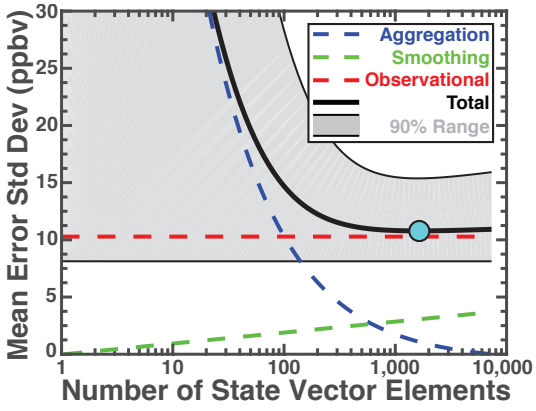

Figure 4. Total error budget from the aggregation of state vector elements in an inverse model. The application here is to an inversion of methane emissions over North America using satellite methane data with 7366 native-resolution state vector elements (Sect. 5 and Turner et al., 2015). Results are shown as the square roots of the means of the diagonal terms (mean error ~~SD~~standard deviation) in the aggregation, smoothing, and observational error covariance matrices, and for the sum of these matrices. Aggregation uses the GMM with RBF weighting (Sect. 4). There is an optimum state vector size for which the total error is minimum and this is shown as the circle. Gray shading indicates the 90 % ~~confidence intervals~~ range for the total error on individual elements as diagnosed from the 5th and 95th quantiles of diagonal elements in the total error covariance matrix. Note the log-scale on the x axis.