Response to reviewers

A sparse reconstruction method for the estimation of multiresolution emission fields via atmospheric inversion

Ray, Lee, Yadav, Lefantzi, Michalak and van Bloemen Waanders, GMDD 7:5623-5659, 2014.

Response to Reviewer No. 1

We thank the reviewer for his/her suggestions and comments. Our responses to his/her suggestions are below.

Overall response: At the very outset we would like to clarify that our study is focused on investigating the algorithmic aspects of a sparse reconstruction method (based on Stagewise Orthogonal Matching Pursuit, StOMP) for estimating rough emission fields, such as that of fossil-fuel CO2 (ffCO2). A sparse reconstruction method is necessary since the spatial parameterization for rough fields tends to be high dimensional (many parameters). The parameters that can be estimated depend on the information content of the observation data, which can change with time/season. The method is not customized to a particular tracer, measurement network or a transport model. Customization to a tracer occurs when we choose a spatial parameterization (a wavelet-based random field model in this study) for use with our sparse reconstruction method. It also occurs when we choose an observational dataset. The method can accommodate prior information on the field being estimated, but only uses its spatial pattern; thus, by design, it is insensitive to under/overestimation of the emissions in the prior information.

The paper investigates which formulations of the inverse problem do and do not work, and explains why. It develops a metric (mutual coherence) to quantify the information content in the observations collected by our measurement network. It finds the information content lacking, which motivates the need to introduce prior information into the inverse problem. We then identify a way to do so; the obvious/intuitive ways do not work. We also show how a wavelet-based field model, designed for modeling fields in rectangular geometries, can be used to estimate emission fields in an irregular region $\mathcal{R}$ (the Lower 48 states of the US). Finally, we show how StOMP can be extended to enforce non-negativity on the estimates. Sparse reconstruction methods are typically not used in atmospheric inversions.

Our motivation to develop this method arose from a need to construct and/or validate gridded inventories of ffCO2. Fortunately, many gridded ffCO2 inventories are available and a wavelet-based spatial parameterization also exists. We demonstrated the method in an idealized, synthetic-data inversion. The idealizations include: (1) assuming ffCO2 to be a radiocarbon-like tracer and ignoring interference by biospheric CO2 which can make ffCO2 estimation impossible except in winter (see Shiga et al., [2014]); (2) using a model-data mismatch ε that is smaller than the one used in real-data inversions and (3) assuming the same distribution for ε for all towers (i.e., ignoring transport model errors). These idealizations allowed us to explore issues related to the algorithm and formulations in a relatively "clean" setting. We also use an observational dataset collected from a measurement network that was sited with biospheric CO2 fluxes, not ffCO2, in mind (the towers are usually far from locations with high ffCO2 emissions); a network for ffCO2 does not currently exist.

Due to these idealizations adopted in our test, we do not claim that the method can be used to estimate ffCO2 emissions fields in a realistic setting using measurement techniques and infrastructure that are currently available (or could be in the near future). At the very least, our method has to be extended to include the estimation of biospheric fluxes as well as larger and tower-dependent model – data mismatches. This is a substantial body of work and outside the scope of this study. In order to check how accurate the estimates would be, we would have to conduct an OSSE (Observational System Simulation Experiment) or design an ideal network. Our tests also provide no information on the best method to collect information on estimation of ffCO2 emissions over regional scales (tower, airplane transects etc.).

We check our inversion method using the following metrics:

1. As part of our algorithmic development, we modify StOMP to incorporate prior information to improve estimates. We check whether it indeed does so, since the information content of the observations are found to be poor.
2. The aim of sparse reconstruction is to estimate parameters that are supported by data (usually large spatial patterns in the emission field) and remove the details that are not. We check whether this "sparsification" characteristic of the algorithm is still present after including prior information.
3. Our method restricts emission fields in an irregular region $\mathcal{R}$ (while using a wavelet-based model); this incurs a computational cost that can be limited by a user-defined setting. We check if the behavior of the algorithm provides a principled way of computing that setting (e.g., if improvement of results shows a "diminishing returns" behavior with the computational cost).

Note that in this study we do not use the accuracy of the estimated emission fields as a metric for evaluating our inversion method. This is because accuracy of estimation is determined primarily by two factors (once we have specified a model – data mismatch): (1) the suitability of the spatial parameterization for the rough fields being estimated and (2) the information content of the observational dataset. In our previous paper Ray et al, [2014] we fixed the observational data and used the accuracy of the emission estimates to gauge the quality of the spatial parameterization. The converse procedure – fixing the spatial parameterization and varying the quantity of observational data – is not a very useful direction for investigation, for our StOMP-based method, because of the following reasons:

1. The estimation accuracy of StOMP, as the quantity of observational data is varied, has been investigated in Donoho et al, [2012]
2. If the aim is to obtain a very accurate reconstruction of the ffCO2 field (when we have full discretion to design an ideal observation network/technique), then we are limited only by what the spatial parameterization can capture. As reported in our previous paper [Ray et al, 2014], the spatial parameterization with 1023 wavelet coefficients (parameters) has a relative error of 10% at the 1-degree resolution; this would be recovered (modulo the small model – data mismatch) in case of an ideal network. If we retain all wavelet coefficients that can be described on a 1-degree mesh in the spatial parameterization (4096 coefficients), the reconstruction will be perfect (modulo the model – data mismatch).

Atmospheric inversion could be a way of estimating/verifying self-reported ffCO2 emissions in countries where the uncertainty is high. The uncertainty in emission reports from China is estimated to be 15-20% [Andres et al, 2012], though studies based on the TRACE-P campaign proposed a 54% revision of inventory estimates for 2000 [Suntharalingam et al, 2004] (it was officially revised upwards by 23% between 2006-2007). Other countries have larger variations. These uncertainties affect inventories, but do not affect our inversion method (see paragraph 1). Even if our variable of interest were to be total emissions over a region (nation or province), estimating a spatially variable emission field before spatially aggregating it reduces the aggregation error. However in order to do this, a measurement infrastructure designed with ffCO2 in mind is a requirement. Its size will be determined by whether we are interested estimating total national emissions or we seek fine scale details.

In addition, as mentioned above, our method can be used with other tracers provided we have a spatial parameterization for them.

The introduction section in our paper does not describe the idealized nature of our tests or the limits/caveats on the conclusions that can be drawn from them. It also does not describe the reasoning behind the metrics that were adopted for evaluating our algorithm. We will add them in the revised paper.

Overarching comments

The reviewer states his/her first general issue –"Accessibility: Much of the paper uses very technical wording, and I worry that much of this phrasing may not be very accessible to atmospheric scientists. Most existing atmospheric inversion studies use a single framework for inverse modeling – Eq. 1 listed in this article. Additional frameworks, like the one presented in this article, could be incredibly useful. However, I suspect that most atmospheric scientists will be unfamiliar with this type of sparse reconstruction in the same way that most are familiar with Eq. 1. I might focus on making this article more accessible to that audience. The authors could do this in a number of ways: (1) by removing technical phrases or terms of art that are not strictly necessary, (2) by providing more descriptive explanation of some of the methods, or (3) by more explicitly guiding the reader through some of the equations.

In particular, I might focus on re-wording the abstract and introduction in a less technical way – in a manner that provides more physical intuition for a reader who may not be familiar with this type of sparse reconstruction method. This re-wording would help broaden the article's appeal to a wider audience and will clearly motivate the subsequent sections that, by necessity, are more technical. To this end, I might focus on giving the reader a holistic, descriptive overview of why under-constrained problems can be challenging, what sparse reconstruction methods are, and how those methods can provide an attractive solution."

We will re-write the abstract to be simpler and jargon-free (as far as possible). We will include a paragraph on how sparse reconstruction methods help in the solution of inverse problems.

The reviewer states his/her second general issue – "I am somewhat concerned about the choice of synthetic data study. A recent paper by Shiga et al. (2014, doi:10.1002/2014GL059684) indicates that existing atmospheric measurements have difficulty identifying ffCO2 fluxes above biospheric fluxes. As a result, I wonder if ffCO2 is necessarily the best species for a synthetic case study. In the real world, these emissions are *often obscured by fluxes from the biosphere. The authors might instead want to consider a gas with both natural and anthropogenic emissions that are largely non-negative. For example, methane, nitrous oxide, or one of several fluorinated greenhouse gases could make for a good synthetic case study.*

I think this issue may actually be cursory to the central objective of the paper – to present a new inversion method. However, it may nonetheless distract the reader or detract from the perceived applicability of the method."

We agree that idealizations we adopted for in our test case (no biospheric CO2 fluxes, very small model-data mismatch etc.) could distract readers (especially in view of Shiga et al.'s recent paper), even though they may be tangential to the usefulness and novelty of a sparse reconstruction method in atmospheric inversion. See "Overall response" for the reasons behind our use of the idealizations.

However, we did investigate other pollutants and ultimately decided on ffCO2. Our choice of ffCO2 (in our idealized inversion) was based on the fact that we had the spatial parameterization for ffCO2 ready (unlike the other gases). Further, anthropogenic CH4 and N2O emissions also suffer interference from their natural counterparts (about 40% and 70% of the total respectively). Lastly, the issues related to ffCO2 (scientific, economic and political) are well known, which allowed us to describe them in a concise manner, and devote most of the paper to mathematical and algorithmic issues.

Detailed comments

1. The reviewer states "The beginning of the abstract is somewhat technical and may be challenging for the reader to follow. For example, the terms "wavelet-based random field models" and "non-rectangular geometries" may not be familiar to the reader. The authors could instead open with a non-technical sentence that communicates how this sparse reconstruction scheme is advantageous or how it represents advancement."

We will re-write the abstract

2. Pg 5625, Line 6-8: The reviewer asks "Would it be possible to cite a reference that illustrates an example of this?" i.e. whether we could cite a reference that shows that non-stationary fields cannot be modeled using multivariate Gaussian fields

We will re-write this sentence. Non-stationary fields can be modeled using multivariate Gaussian fields but require complex covariance models. In practice, these are difficult to construct and we adopt a model (typically, a variogram) that has one length-scale. Such a model suffices if the emission fields being modeled are smooth. In some cases, a non-stationary field can be represented as a non-stationary mean field with an additive correction modeled as a stationary multivariate Gaussian field. In such a case, modeling effort shifts to constructing the mean field representation.

We will add this clarification in the text.

3. Pg. 5625, Line 16-17: The reviewer asks "Could one theoretically use a model selection method (like AIC, BIC, or DIC) to decide whether a parameterization is too simple or too complex?"

Yes, one could start with an over-parameterized model, create multiple variations of it by retaining a subset of the parameters and use AIC (or BIC or DIC) to choose the best model after fitting to data. It is an alternative to sparse reconstruction to simplify the spatial parameterization, conditioned on data.

We have added this explanation in the paper.

4. Pg. 5625, Line 25-28: The reviewer points out "These examples add a lot of extra technical detail to the description of sparse reconstruction methods. I wonder if this level of detail is necessary when giving the reader a broad, holistic definition of sparse reconstruction."

Our method for estimating ffCO2 involves formulating a linear inverse problem in a way that could be solved using StOMP. Since StOMP is so central to our method we felt that it would be natural to introduce it early when discussing what sparse reconstruction methods are and what they do. The other methods cited (Matching Pursuit and its variations) are all related to StOMP and identify the class of sparse reconstruction methods we will use.

Note that there are other l_1 minimization methods such as LASSO (Least Absolute Shrinkage and Selection Operator) and LARS (Least Angle RegreSsion) that can also solve the sparse reconstruction problem. Since we will not be using this "family" of algorithms in our paper, we have omitted them from our brief introduction to sparse reconstruction methods. Sparse reconstruction (or shrinkage regression) is too large a topic to be adequately summarized in our paper.

5. Pg 5626, Lines 2-6: The reviewer asks: "What is an l1 and l2 norm? Furthermore, what is the "offline construction of a spatial parameterization" and why would we want to dispense of it?"

l_1 norm of $\mathbf{x} = \|\vec{x}\|_1 = \sum |x_i|$. l_2 norm of $\mathbf{x} = \|\vec{x}\|_2 = \sum_i x_i^2$. We will add this definition in the text.

The conventional way of estimating a field is as follows:

1. *Determine a model for the field ("spatial parameterization")*: It could be based on Gaussian Process (or Gaussian random fields) in which case one develops a model for the field's mean as well as a covariance model for deviations from the mean field model. Alternatively, one could opt for a field model constructed as a linear sum of orthogonal bases (wavelets, eigenvectors of the covariance matrix etc.), in which case one decides on which bases to retain in the field model. There are many methods to do so. This is the "offline construction of the spatial parameterization".
2. *Solving the inverse problem:* Here one fits the spatial parameterization to the observed data using whatever method is convenient.
3. *Refining the model:* The spatial model constructed could be too complex or too simple for the observations used to solve the inverse problem. To address this, one proposes a set of competing spatial parameterizations and solves the inverse problem again. Thereafter, using a metric such as AIC, one picks the best spatial parameterization, which provides accuracy without overfitting the data. This is an iterative process, since the proposed set of competing spatial parameterizations could very well be incomplete.
4. *Repeat if the quality of the data is time-dependent:* We estimate weekly emission fields for a year. Since wind directions change with time, the information content of

the observations (the "quality" of the observations) is time-dependent. The process described above has to be repeated for every week. This is tedious.

Consequently, we were motivated to construct a method where we could dispense with the offline construction of the spatial parameterization. Rather, we desired a method where we could start with a spatial parameterization that had too many parameters and have the sparse reconstruction algorithm identify, given the observations available, the parameters that could actually be estimated (and set the rest to zero). Our sparse reconstruction method does this efficiently. The set of "estimate-able" parameters could vary week to week. One could call this process an "online construction of a sparse spatial parameterization" or "online sparsification of the spatial parameterization".

6. Page 5626, Line 19-20: The reviewer asks, "What do you mean by the choice of the proxy used for spatial disaggregation"?

Most fossil-fuel fuel CO2 emission inventories are constructed by obtaining fossil-fuel emissions declared by a country (only national and/or regional aggregate emissions are provided) and disaggregating them down to much finer resolutions e.g. 10km. This spatial disaggregation/downscaling is performed using a proxy of human activity that is available at such a fine spatial scale. To date, population density, measures of economic activity (e.g., per capita GDP) and images of lights at night have been used as the disaggregation / downscaling proxies. These proxies are quite different at fine spatial scales and these differences are naturally manifest in the downscaled emissions. Consequently, the ffCO2 inventories do not agree at fine spatial scales.

7. Page 5627, Line 12-13: The reviewer asks "What is a wavelet-based random field model?"

Wavelets are an orthogonal basis set. Any function (e.g. emission field) defined on a grid with N gridcells can be decomposed (by projection) on the N wavelets that can be described on the grid. Not all projections are large/numerically significant, and consequently the function could be approximated by a weighted linear combination of k wavelets, $k << N$. This creates a wavelet-based field model.

The k weights constituting the wavelet-based field model can be considered to be random numbers. Depending on the values assumed by the weights, we can create an infinite number of fields ("random fields"). However, since these fields are constructed only using k wavelets on an N grid-cell grid, they fall in a certain class of fields i.e., we cannot create all fields that could be described on the grid. This is called a wavelet-based *random* field model.

We will add this description in the paper.

8. Pg 5627, Line 22-23: The reviewer asks "What kind of a spatial parameterization does this paper develop? i.e. it is not clear what the term "spatial parameterization refers to here".

The parameterization refers to a spatial parameterization that was developed in our previous work (we begin the sentence with a citation of our previous work). It is a wavelet-based random field model, constructed out of Haar wavelets, for representing ffCO2 emissions. The wavelets retained in the model were selected by projecting radiances of lights at night on to a set of Haar wavelets described on a 1 degree X 1 degree grid; only about 25% of the wavelets were retained. The weights of the wavelets are deemed to be random numbers (uniformly distributed between plus and minus infinity) and are estimated by the inverse problem that we describe in this paper.

We will add this description of what we mean by spatial parameterization in the manuscript.

9. Pg 5629, Line 4-11: The reviewer asks "How different are EDGAR and Vulcan? This difference would help the reader understand whether the estimated emissions shown in Fig. 1 more closely resemble the prior (EDGAR) or the true fluxes (Vulcan). The authors may want to consider adding a plot of EDGAR fluxes to Fig. 1."

EDGAR emissions are 7.1% higher than Vulcan emissions when aggregated over the US. The RMSE between the two is 0.035 μmoles $m^{-2}s^{-1}$ of C and the Pearson correlation between the two is 0.726.

We will add a figure of EDGAR emissions to the text. Remember, though, that Vulcan emissions change with time while EDGAR is an annual average.

10. Pg 5630, Eq. 1 The reviewer ask "It may be useful to the reader to define the dimensions of each matrix or vector."

Updated in the new manuscript.

11. Pg. 5630, Line 26: The reviewer asks "What are orthogonal bases with compact support? Some readers may understand this term, but I worry that many atmospheric scientists may not fully understand this technical term."

Orthogonal bases are a set of functions, which are mutually orthogonal i.e. they have zero projection on each other, in the domain where they are defined. Sine waves are a common example. Because of orthogonality, they can be linearly combined to model arbitrary functions.

Compact support refers to the fact that the bases are non-zero only in a part of the domain where they are described. This is unlike sine functions that are defined between plus/minus infinity. Wavelets are orthogonal and have limited support. The limited spatial extent of the basis allows us to use them to define localized structures and discontinuities with ease, unlike sine waves.

We will add this description of orthogonal bases with compact support in the manuscript.

12. Pg 5631, Eq 3. The reviewer asks "It might be useful to the reader to explain in words what the components of this equation mean. I.e., it may be helpful to guide the reader through this equation. The current text does not provide much explanation of what this equation means. In addition, does μ refer to the mean here?"

We will do so in the updated manuscript.

13. Pg 5631: The reviewer asks "Are "incoherence" and "mutual coherence" the same thing? It may be useful to explain the relationship of these terms to the reader."

No, quite the opposite. Zero mutual coherence is (perfect) incoherence. Loosely, we refer to a small mutual coherence (small compared to $N^{1/2}$) as "incoherence".

We've clarified this in the new manuscript.

14. *Pg. 5632, Line 11: The reviewer asks "What is the l0 norm of $\vec{w}$ and what is the l2 norm of the measurement-model discrepancy? Do they refer to the "1" and "2" subscripts in Eq. 4? If so, it may be useful to clarify here."*

l_0 norm of w = $\|w\|_0 = \sum_i w_i^0$ = the number of non-zero elements in vector **w**. The l_2 norm has been described in Response no. 5. It is indeed the "2" subscript in Eq. 4. The "1" subscript refers to l_1 norm (see Response 5 above).

15. Pg 5635, Eq. 7: The reviewer states "Is there any way to guide the reader through this equation? I know that this equation is, to some degree, an extension of Eq. 6, but I worry that the authors may lose the reader here."

We will do so. It was described much better in our previous paper (which had the full context within which it was developed), but we realize that the current paper has to stand on its own.

16. Pg 5637, Line 3-6: The reviewer asks "What is the ultimate effect of fine versus coarse wavelets? i.e., what effect would these differences have on the estimated fluxes or what desirable quality would these properties confer?"

The coarse wavelets capture the large-scale patterns in the emission field e.g., the urban sprawl on the East Coast whereas the fine wavelets capture local structures e.g. intense regions of ffCO2 emissions like cities in the High Desert e.g., Salt Lake City. As mentioned in the paper, it is easy to estimate the emission in the sprawl as a number of measurement towers lie inside it. On the other hand, an isolated hot spot like Salt Lake City can be difficult to estimate unless it lies close to a tower.

17. Pg 5637: The reviewer asks "Approach C: Could you potentially describe in more qualitative terms how this approach differs from B?"

Approach B models the true emission as essentially our prior belief (based on EDGAR) with an additive correction field ΔF. This is an intuitive way of modeling the emission field, but leads to problems when using sparse reconstruction.

Approach C normalizes the wavelet coefficients of the emission field being estimated using the wavelet coefficients from the prior (EDGAR). This works well.

We will add this explanation in the updated manuscript.

The reason for the failure of Approach B is explained on Pg 5641, Line 22:28 and Response 16. It arises from the fact that wavelets with smaller support (fine wavelets) have wavelet weights (or wavelet coefficients) that are large – they tend to model isolated emission sources or abrupt changes in the emission field. When we model the emissions with a correction over EDGAR, the wavelet coefficients of correction, corresponding to the fine wavelets, also tend to be large. Since the sparse reconstruction method seeks to reduce

$l_1(\Delta w)$, where Δw are the wavelet coefficients of the correction field, the large coefficients of the fine wavelets quickly come under scrutiny. Unless they can induce a large change in predicted concentrations at the measurement towers (only possible if there's one nearby), they are quickly removed. In fact, Δw becomes very sparse indeed, and the estimated emission field looks a lot like EDGAR.

The reason for this discrepancy is the l_1 ($\mathbf{w}$) minimization implicitly assumes that all elements of **w** are of the same order of magnitude. In our problem, they are not.

Consequently, in Approach C we normalize the wavelet coefficients of the modeled field with that of the prior emission field. It renders the wavelet coefficients non-dimensional and the same order of magnitude. The sparse reconstruction proceeds as expected. However, in normalizing, we have assumed that the prior's wavelet coefficients' proportions to each other are similar to those of the true emission field's wavelet coefficients. This implies that normalized wavelet coefficients are non-dimensional and of the same order of magnitude, though not necessarily O(1). If the prior is not a gross under/over-estimate of the true emissions, the normalized wavelet coefficients are of O(1), in addition. Thus Approach C provides a way of incorporating prior information. This is explained on Pg. 5643, Line 17:24.

18. Pg 5638 Eq. 11 The reviewer asks "Have these variables already been defined elsewhere in the manuscript? I may have missed this definition. If not, it would be helpful to explicitly define these variables for the reader (or explicitly state how they relate to the equations in Approaches A-C)."

This is our mistake – they have not been defined. We will fix this in our updated manuscript.

19. The reviewer states "The authors do a great job of leading the reader through this section in a structured and informative way"

20. The reviewer states "Conclusions: The conclusion section is well written. The authors are adept at summarizing their method and its advantages in a way that is likely to be accessive to many readers

Thank you!

21. Pg. 5648 Line 3-8: The reviewer asks "The authors may want to remind the reader which sections discuss each step. The reader may not remember each step exactly, and a reference to individual sections would help the reader jump to this information quickly."

The paragraph describes a test that we did (it involved omitting certain steps in the inversion method) and whose results are not presented in the paper. However, we will refer back to the sections as we describe the steps that were omitted, leading to the failure of the tests.

We will add this clarification in the updated manuscript.

22. Pg 5648, Line 11: The reviewer asks "How would a Kalman filter rectify this problem? I would either clarify this logic or omit the statement."

Using a Kalman filter (or probably an ensemble Kalman filter, given the large dimensionality) would provide us with a mean emission field and a covariance matrix. The covariance matrix would allow us to compute standard deviations on the estimated fields.

23. Pg. 5648, Line 17:20: The reviewer states "I think this Matlab code will make the inversion method much more accessible to most readers."

We agree, and we are pleased to announce that we have finally received permission to distribute the code. The software has been freely downloadable on the website since October 10th.

24. The reviewer states "Figure 1: I might set zero values to white or use a color scale that uses light colors or shades for low emissions values. This could make the figure more visually appealing and easy to see."

We will do so.

25. The reviewer states "Figure 5: It could be more informative to list actual dates on the x-axis in panel a. The current label ("observation number") is not very informative to the reader."

We will do so.

References

[Andres et al, 2012] R. J. Andres et al, "A synthesis of carbon dioxide emissions from fossil-fuel combustion", Biogeosciences, 9, 1845-1871, 2012. doi: 10.5194/bg-9-1845-2012.

[Donoho et al, 2013] D. Donoho et al, "Sparse solutions of underdetermined linear equations by Stagewise Orthogonal Matching Pursuit", IEEE Transactions on Information Theory, 58, 1094-1121, 2012.

[Ray et al, 2014] J. Ray et al, "A multiresolution spatial parameterization for the estimation of fossil-fuel carbon dioxide emissions via atmospheric inversions", Geoscientific Model Development, Geosci. Model Dev., 7, 1901-1918, 2014. doi:10.5194/gmd-7-1901-2014.

[Shiga et al, 2014] Y. P. Shiga et al, "Detecting fossil fuel emissions patterns from sub-continental regions using North American in-situ CO2 measurements", Geophysical Research Letters, 41(12):4381-4388, 2014.

[Suntharalingam et al, 2004] P. Suntharalingam et al, "Improved quantification of Chinese carbon fluxes using CO2/CO correlations in Asian outflows", Journal of Geophysical Research, 109, D18S18, 2004. doi:10.1029/2003JD004362.