



A sparse reconstruction method for the estimation of multi-resolution emission fields via atmospheric inversion

J. Ray¹, J. Lee¹, V. Yadav², S. Lefantzi¹, A. M. Michalak², and B. van Bloemen Waanders³

¹Sandia National Laboratories, P.O. Box 969, Livermore, CA 94551, USA

²Carnegie Institution for Science, Stanford, CA 94305, USA

³Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM 87185-0751, USA

Correspondence to: J. Ray (jairay@sandia.gov)

Received: 23 July 2014 – Published in Geosci. Model Dev. Discuss.: 20 August 2014

Revised: 7 March 2015 – Accepted: 7 April 2015 – Published: 29 April 2015

Abstract. Atmospheric inversions are frequently used to estimate fluxes of atmospheric greenhouse gases (e.g., biospheric CO₂ flux fields) at Earth's surface. These inversions typically assume that flux departures from a prior model are spatially smoothly varying, which are then modeled using a multi-variate Gaussian. When the field being estimated is spatially rough, multi-variate Gaussian models are difficult to construct and a wavelet-based field model may be more suitable. Unfortunately, such models are very high dimensional and are most conveniently used when the estimation method can simultaneously perform data-driven model simplification (removal of model parameters that cannot be reliably estimated) and fitting. Such sparse reconstruction methods are typically not used in atmospheric inversions. In this work, we devise a sparse reconstruction method, and illustrate it in an idealized atmospheric inversion problem for the estimation of fossil fuel CO₂ (ffCO₂) emissions in the lower 48 states of the USA.

Our new method is based on stagewise orthogonal matching pursuit (StOMP), a method used to reconstruct compressively sensed images. Our adaptations bestow three properties to the sparse reconstruction procedure which are useful in atmospheric inversions. We have modified StOMP to incorporate prior information on the emission field being estimated and to enforce non-negativity on the estimated field. Finally, though based on wavelets, our method allows for the estimation of fields in non-rectangular geometries, e.g., emission fields inside geographical and political boundaries.

Our idealized inversions use a recently developed multi-resolution (i.e., wavelet-based) random field model developed for ffCO₂ emissions and synthetic observations of

ffCO₂ concentrations from a limited set of measurement sites. We find that our method for limiting the estimated field within an irregularly shaped region is about a factor of 10 faster than conventional approaches. It also reduces the overall computational cost by a factor of 2. Further, the sparse reconstruction scheme imposes non-negativity without introducing strong nonlinearities, such as those introduced by employing log-transformed fields, and thus reaps the benefits of simplicity and computational speed that are characteristic of linear inverse problems.

1 Introduction

The estimation of spatially resolved fields, e.g., permeability fields in aquifers or CO₂ fluxes in the biosphere, from limited observations, are required for many scientific or engineering analyses. These fields are generally represented on a grid whose spatial resolution is dictated by the analyses. The observations are usually too scarce to allow the estimation of the field's values in each grid cell independently. If the field is known to be smooth, one can impose a spatial correlation between the grid cells (e.g., model the field as a realization from a stationary multi-variate Gaussian distribution) and reduce the effective dimensionality of the estimation problem so that the limited observations suffice. In contrast, if the field is complex, i.e., non-smooth or non-stationary (in the statistical sense, implying different characteristic length scales at different locations), multi-variate Gaussian models are difficult to construct and an alternative parameterization may be preferable. The parameterization has to be low dimensional,

i.e., have few independent parameters, so that they can be estimated from limited observations.

The construction of the spatial parameterization for complex fields poses a stiff challenge. The parameterization is usually problem dependent and sometimes based on heuristics. One may use an easily observed covariate (or predictor) of the field being estimated to construct such a model; for example, see Ray et al. (2014) for a description on how images of lights at night were used to create a spatial parameterization for fossil fuel CO₂ (ffCO₂) emissions. However, one is never quite sure if the resultant parameterization is too simple or too complex – in the former case, the estimates will be needlessly inaccurate, while in the latter case, one may not obtain a unique solution or the estimates may reproduce the noise in the observations (overfitting). Consequently, one uses some method, for example, Akaike information criterion, to devise models of suitable complexity. However, if the quality of the observations changes with time, then, ideally, a different parameterization is constructed for each time instant. In practice, often the simplest model that can be used with all the observations is employed. This degrades estimation accuracy.

Sparse reconstruction methods can allow one to circumvent these problem which arise from the dimensionality of spatial parameterization (also called the random field model). Sparse reconstruction methods such as matching pursuit (MP; Mallat and Zhang, 1993), orthogonal matching pursuit (OMP; Tropp and Gilbert, 2007) and Stagewise OMP (StOMP; Donoho et al., 2012) are optimization methods that are used to fit high-dimensional models to limited observations. Unlike other optimization methods, these methods enforce sparsity, i.e., they identify the model parameters that are not informed by the observations and set them to zero. This is accomplished by augmenting the objective function (usually a ℓ_2 norm of the data – model discrepancy or residuals) with a penalty formulated as a ℓ_1 norm over the parameters being estimated. (The ℓ_2 norm of a vector \mathbf{x} is defined as $\|\mathbf{x}\|_2 = \sqrt{\sum_i x_i^2}$, while the ℓ_1 norm is $\|\mathbf{x}\|_1 = \sum_i |x_i|$.) An optimizer is used to manipulate model parameters to minimize the objective function. The parameters that do not impact the residual appreciably are quickly driven to zero, as it minimizes the ℓ_1 penalty, i.e., the optimizer performs dimensionality reduction while it fits the model to data. This model simplification characteristic of sparse reconstruction methods allows one to dispense with the offline construction of a spatial parameterization and postulate a general, high-dimensional random field model instead; thereafter, the optimization method simplifies (reduce the dimensionality of) the random field model in a data-driven manner. In the case of observations with time-variant quality, sparse reconstruction methods have the potential to be particularly useful.

Our interest in sparse reconstruction methods arises from a need to develop accurate spatially resolved estimates of

emissions that are not smoothly distributed in space; ffCO₂ emissions are one such example. Estimates of ffCO₂ emissions are used to assess regional contributions to greenhouse gas emissions and to drive climate change simulations (Andres et al., 2012). Currently, spatially resolved estimates of ffCO₂ emissions are typically derived from national-level emissions inventories, and are mapped spatially using population density or some other proxy of human activity; examples of such spatially resolved inventories are described in Gurney et al. (2009), Olivier et al. (2005), Rayner et al. (2010) and Oda and Maksyutov (2011). Their shortcomings arise from errors in national/provincial reporting and the choice of the proxy used in spatial disaggregation (Andres et al., 2012). Recently, the possibility of using atmospheric observations to constrain fossil fuel emissions, and thereby improve inventories, has been explored (Pacala et al., 2010). Such applications involve the solution of an inverse problem driven by ffCO₂ concentration measurements (Rayner et al., 2010). Note that such improvements would be contingent on a good representation of an estimation problem within the context of an inversion, including the use of a suitable parameterization for the emissions, the characterization of transport model errors, and the availability of an observational network that is sufficient to provide an adequate constraint on ffCO₂ emissions; ffCO₂ emissions for individual urban domes have been estimated using atmospheric measurements (Turnbull et al., 2011; McKain et al., 2012; Kort et al., 2012), i.e., without solving an inverse problem, but existing methods do not offer a scalable approach to updating entire inventories in this manner.

As a step towards enabling such applications, we constructed a wavelet-based spatial parameterization, called the multiscale random field (MsRF; Ray et al., 2014), to represent ffCO₂ emission fields. The MsRF was used to model ffCO₂ emissions in the lower 48 states of the USA at $1^\circ \times 1^\circ$ spatial resolution. The MsRF covers a rectangular region described by the corners 24.5° N, 63.5° W and 87.5° N, 126.5° W. The emissions are modeled using Haar wavelets, which provide the sparsest representation of ffCO₂ emissions in the relevant region. The model has $O(10^3)$ independent model parameters which were selected using images of lights at night. Due to its high dimensionality, the MsRF model cannot be used directly given realistic in situ observational limitations. However, a data-driven dimensionality reduction of the MsRF model, using a sparse reconstruction method, could help constrain the inverse problem and make it possible to capture coarse spatial patterns of ffCO₂ emissions (and, perhaps, finer details in the vicinity of the sensors), conditioned on atmospheric measurements.

The use of sparse reconstruction methods poses certain methodological challenges. First, these reconstruction methods do not provide a mechanism for imposing non-negativity, which is a requirement when estimating emission fields. Second, sparse reconstruction methods have, to date, been used with wavelet-based random field models which can

only model rectangular domains; in contrast, the geometry of emission fields could be decided by geographical or political boundaries. (A random field model is a spatial parameterization for a field defined on a grid. It can be constructed using orthogonal bases such as wavelets; the wavelets' weights are the model parameters and are treated as random variables. Realizations of these random variables produce a realization of the field. Depending upon the choice of the basis set, e.g., if it contains only a subset of wavelets that can be supported by the grid, the random field model may be able to produce only a subset of the infinite number of fields that the grid can support). Finally, sparse reconstruction methods do not provide a simple mechanism to incorporate prior information or guesses of the field being estimated, a common technique to ensure a unique solution to an inverse problem. This is because methods such as OMP and StOMP were largely developed for the reconstruction of compressively sensed images (Candes and Wakin, 2008) where prior information is weak. In contrast, many emission fields of an anthropogenic nature have inventories that can serve as very informative priors and reconstruction methods could profitably use them.

Our previous work (Ray et al., 2014) focused on the spatial parameterization (the MsRF described above) for estimating ffCO₂ emission fields via atmospheric inversion. In this paper, we describe the methodological innovations in sparse reconstruction techniques that allowed us to perform the inversion, despite the high dimensionality of the parameterization. These innovations result in an extension of StOMP which can address the peculiarities of reconstructing an emission field. The StOMP extension will be demonstrated in a top-down inversion, using synthetic observations generated from a known, ground-truth emission field so that we may examine certain algorithmic and numerical aspects of the estimation technique, as described below. The novel algorithmic developments addressed in this paper are

1. *Incorporation of a prior model of spatially rough emissions*: we demonstrate a novel and simple method to introduce prior information on spatially rough emission fields (in the form of an approximate field f_{pr}) into StOMP. Currently, sparse reconstruction methods employ no other prior information beyond the phenomenological observation that most fields can be represented quite accurately with a sparse set of judiciously chosen wavelet bases (Candés and Romberg, 2007).

Note, that the term *prior model* or *prior information* is used somewhat loosely here since our method is not strictly Bayesian. However, f_{pr} serves a similar function by providing regularization in the inverse problem.

2. *Estimating fields in irregularly shaped regions*: the MsRF model, being based on wavelets, can only model fields in rectangular domains, whereas our emission field is distributed over an irregular region \mathcal{R} , the lower 48 states of the USA. We demonstrate how this geomet-

rical constraint can be imposed efficiently using random projections, a technique that underlies much of compressive sensing. The reconstruction of fields in non-rectangular geometries has no parallel in the compressive sensing of images and the method discussed in this paper is the first of its kind.

3. *Imposition of non-negativity*: the estimation of the emission field is posed as a linear inverse problem (see Sect. 2). Non-negativity of emissions can be enforced by log transforming the field, but converts the problem into a nonlinear one, requiring computationally expensive, iterative sparse reconstruction methods, like the one developed in Li and Jafarpour (2010). We develop a simple, iterative post-processing method to enforce non-negativity on the estimated ffCO₂ emissions. The non-negativity enforcement mechanism uses StOMP but does not use the MsRF model. The imposition of non-negativity in the sparse reconstruction of an emission field has never been explored before; for example, in Hirst et al. (2013), the non-negativity constraint was not applied to CH₄ emissions from landfills.

In this study, we demonstrate our method on an idealized atmospheric inversion of a spatially rough emission field in \mathcal{R} . The method is general, but we use ffCO₂ as the test case. The idealizations are enumerated below.

1. We assume that ffCO₂ can be measured independently without interference from biospheric CO₂ fluxes. As described in Ray et al. (2014), this could be performed using $\Delta^{14}\text{CO}_2$ (radiocarbon) or other non-CO₂ tracers, but the measurement technology is expensive and far from being widely deployable.
2. Inversions require us to adopt a statistical error model for the mismatch between observations and model predictions using the estimated emission field. This error quantifies the aggregate of measurement uncertainties and errors introduced by the approximations in the transport model, among others. It varies between measurement locations. In this study, we model this mismatch as i.i.d. (independent and identically distributed) Gaussian random variables. We assume a value for the standard deviation of the distribution that is too small compared to what is possible using existing transport models and measurement technologies; further, we use the same error model for all the measurement locations (details in Sect. 4). The small error allowed us to investigate the numerical aspects of our formulation and solution algorithms without being substantially affected by observational noise. The small error was also required due to the nature of the measurement network employed in the synthetic data test (see below).

The synthetic measurements are obtained at a set of 35 towers, a network that existed in 2008 (see Ray et al., 2013,

for their locations). This network, sited with biospheric CO₂ measurements in mind, has towers which tend to be far from urban areas and thus sources of ffCO₂ emissions. Consequently, the modeled ffCO₂ concentrations at these towers tend to be low, forcing us to employ an error model that is unrealistically small. These idealizations lead to limits on the inferences that can be drawn regarding the use of our method in a real-data inversion for ffCO₂ emission fields; they are discussed in Sect. 5.

We will estimate the emission field at 1° × 1° resolution. Emission fields are averaged over 8 days, and estimated over 360 days, i.e., we estimate 360/8 = 45 fields. ffCO₂ emission from the Vulcan inventory (version 1) (<http://vulcan.project.asu.edu/index.php>; Gurney et al., 2009) serve as the ground truth, to generate the synthetic or pseudo-observations y^{obs} of time-variant ffCO₂ concentrations. (The Vulcan inventory provides hourly ffCO₂ emissions at 10 km resolution for the lower 48 states of the USA for 2002; it can also be downloaded at 0.1° resolution.) The prior model f_{pr} will be constructed using the Emission Database for Global Atmospheric Research (source: European Commission, Joint Research Centre/Netherlands Environmental Assessment Agency; Emission Database for Global Atmospheric Research (EDGAR), release version 4.0, <http://edgar.jrc.ec.europa.eu>, 2009; Olivier et al., 2005), which provides a single emission field at 1° resolution for 2005. These choices were driven solely by the easy availability of data.

We evaluate our inversion method using the following metrics. First, we check whether the incorporation of prior information into our modification of the StOMP algorithm improves estimates. Second, we investigate the *sparsifying* nature of our algorithm. The aim of sparse reconstruction is to estimate parameters supported by data (usually large-scale spatial patterns in the rough emission field) and remove details that are not. We check whether this property of StOMP is retained after our modifications. Finally, we check the efficiency with which our method reconstructs emission fields inside an irregular \mathcal{R} . Our use of a wavelet-based spatial parameterization incurs a computational cost which can be limited by a user-defined setting. We check if there is a principled way of computing this setting, e.g., if improvements in results follow a diminishing returns behavior with the computational cost.

Note that in this study, we do not use the accuracy of the estimated field as a metric for evaluating our method; we only use estimation accuracy to select between competing formulations of the inverse problem. The estimation accuracy depends on (1) the spatial parameterization (the MsRF) and (2) the information content of the data set, and was explored in detail in our previous paper (Ray et al., 2014). There, we fixed the observational data and used the accuracy of the estimated emission field to gauge the quality of the MsRF. The converse problem – fixing the MsRF and varying the quantity of data – is not very useful for our StOMP-based algorithm,

since StOMP's sensitivity to data was addressed in Donoho et al. (2012).

The paper is structured as follows. In Sect. 2, we review sparse reconstruction techniques, their use with wavelet models of fields and the tenets of compressive sensing that establish the necessary conditions for successful sparse reconstructions. In Sect. 3, we pose the inverse problem and describe the numerical method used to solve it. Three formulations, differing in the manner in which they incorporate f_{pr} are examined. In Sect. 4 we perform inversion tests with synthetic data to select the best formulation. We also explain, using the properties required for sparse reconstruction, why the selected formulation performed better than the others. The efficacy of limiting the estimated field within \mathcal{R} using random projections is also investigated. Conclusions are in Sect. 5.

2 Background

In this section, we review techniques used to estimate CO₂ fluxes, compressive sensing and the use of sparse reconstruction in inverse problems.

Estimation of CO₂ fluxes: let the vector f be the CO₂ flux defined on a grid with $N_{\mathcal{R}}$ grid cells. Let f be of size $KN_{\mathcal{R}}$, representing a flux field defined over K time periods. The flux is assumed to be time invariant during a given time period. The transport of CO₂ is modeled as that of a passive scalar, i.e., the concentration of CO₂ due to f at an arbitrary set of sites, is given by $y = \mathbf{H}f$. Here, y is a vector $K_s N_s$ long, N_s being the number of locations where measurements are collected K_s times over the K time periods. The matrix \mathbf{H} ($K_s N_s \times KN_{\mathcal{R}}$) contains the sensitivity of measurements to a CO₂ source in each grid cell and is computed using an atmospheric transport model such as the Stochastic Time-Inverted Lagrangian Transport Model (STILT; Lin et al., 2003). In an atmospheric inversion, CO₂ concentration y^{obs} are measured at a limited set of locations, usually a set of measurement towers (as in our case) or as column-averaged satellite soundings. The measurements are too few or too uninformative to estimate f , with each grid cell treated independently. In case of biospheric fluxes, a prior flux f_{pr} (with the same dimensions as f) can be obtained from a biogeochemical process-based model such as CASA (Carnegie–Ames–Stanford Approach; Potter et al., 1993). The discrepancy ($f - f_{\text{pr}}$) is usually modeled as a multi-variate Gaussian field with covariance \mathbf{Q} (a $KN_{\mathcal{R}} \times KN_{\mathcal{R}}$ matrix), and the estimation of f is typically performed by minimizing the objective function

$$J = \underbrace{(y^{\text{obs}} - \mathbf{H}f)^T \mathbf{R}_e^{-1} (y^{\text{obs}} - \mathbf{H}f)}_{\text{Observation term}} + \underbrace{(f - f_{\text{pr}})^T \mathbf{Q}^{-1} (f - f_{\text{pr}})}_{\text{Prior term}}, \quad (1)$$

where \mathbf{R}_e is a diagonal matrix with the data – model variances and includes many sources of errors including measurement errors, aggregation errors and transport model inaccuracies. Methods to solve this linear inverse problem are reviewed in Ciais et al. (2010). A comparison of biogenic CO₂ fluxes and ffCO₂ emissions (Fig. 1 in Ray et al., 2014) shows that ffCO₂ are multiscale in nature and a multi-variate Gaussian field approximation of $(f - f_{pr})$ is unlikely to be accurate. This motivated us to construct the MsRF model for ffCO₂ emission fields (Ray et al., 2014). The solution of an inverse problem using MsRF requires the use of a sparse reconstruction method that, to date, has been used in the reconstruction of compressively sensed images.

Compressive sensing of images: compressive sensing (Romberg, 2008; Candes and Wakin, 2008) is a very efficient means of representing images using wavelets. Wavelets are a family of orthogonal bases with compact support that are routinely used to model complex fields, including ffCO₂ emissions (Ray et al., 2014). Compact support refers to the fact that a wavelet is defined over a finite region (compact support). This is in contrast to other commonly used bases, e.g., $\sin(k\pi x)$, $-\infty \leq x \leq \infty, k \in \mathbb{Z}$, which have infinite support. However, like Fourier bases, wavelets are orthogonal; the scalar product of two different wavelets of the same type and order, e.g., Daubechies wavelets of order 4, is 0. Compressive sensing (CS) is based on two key tenets: *compressible representation* and *encoding via random projections*. CS assumes that an image, projected onto a suitable wavelet basis set, will yield wavelet weights (represented by a vector \mathbf{w}) that are mostly very small (i.e., a *compressible representation*) and can be set to zero. Removing the small wavelets results in a *sparse* approximation of the image. Encoding via random projections is more involved and determines the necessary conditions for successful sampling. Random encoding is central to our method for applying boundary conditions, viz., limiting ffCO₂ emissions within complex, non-rectangular boundaries.

Consider an image \mathbf{g} of size N , that can be represented sparsely using $L \ll N$ wavelets. Random encoding, as used in CS, asserts that the image may be sampled by projecting it onto a set of random vectors ψ_j , to obtain compressive measurements \mathbf{g}' , of size $N_m, L < N_m \ll N$:

$$\mathbf{g}' = \Psi \mathbf{g} = \Psi \Phi \mathbf{w} = \mathbf{A} \mathbf{w}, \quad (2)$$

where the rows of the sampling matrix Ψ consist of the random vectors ψ_j , the columns of Φ consist of the orthonormal basis vectors (the wavelets) ϕ_i and \mathbf{w} are the weights (or coefficients) of the wavelets. Φ is a $N \times N$ matrix while Ψ is $N_m \times N$. The bulk of the theory was established in Candes and Tao (2006), Donoho (2006), Candes et al. (2006) and Baraniuk et al. (2008).

In order that one may recover the original image \mathbf{g} from \mathbf{g}' using sparse reconstruction, Ψ and Φ must satisfy *incoherence* and a *restricted isometry property* (Candes and Wakin, 2008). Incoherence implies that no row ψ_k in Ψ is co-aligned

with column ϕ_l in Φ and thus collects information on all bases. It is ensured by choosing some well-known wavelets bases (e.g., Haars or Daubechies 4 and 8) for Φ and random vectors for Ψ (Tsaig and Donoho, 2006; Coifman et al., 2001). This is formally quantified by the *mutual coherence* $\mu(\Psi, \Phi)$ of Ψ and Φ :

$$\begin{aligned} \mu(\Psi, \Phi) &= \sqrt{N} \max_{1 \leq (k,l) \leq N} | \langle \psi_k, \phi_l \rangle | \\ &= \sqrt{N} \max(|A_{kl}|), \end{aligned} \quad (3)$$

where A_{kl} are elements of \mathbf{A} . Each row of Ψ is normalized to a unit vector. The term $| \langle \psi_k, \phi_l \rangle |$ is the projection of row ψ_k on a wavelet basis ϕ_l . Co-alignment of a row $\psi_{k'}$ with a wavelet $\phi_{l'}$ would lead to $A_{k'l'} = 1$ and $A_{k'l} = 0$ for all $l \neq l'$, indicating that a random vector $\psi_{k'}$ collects information on only one wavelet. The scaling by \sqrt{N} is conventional. A small mutual coherence ensures that all projections of Φ on the rows of Ψ are of moderate magnitude ($O(10^{-1})$ – $O(10^{-3})$). A small mutual coherence aids accurate reconstruction. When $\mu(\Psi, \Phi) \ll \sqrt{N}$, we loosely refer to Ψ and Φ as being incoherent. The restricted isometry property (RIP) is a condition imposed on \mathbf{A} which ensures that \mathbf{w} can be recovered from \mathbf{g}' uniquely without the use of priors (except sparsity). We did not pursue this thread since the use of a prior – making the inventory that supplies f_{pr} consistent with observations – is the motivation behind this investigation.

Sparse reconstruction of images from compressive measurements: the aims of reconstruction in CS are to (1) recover the sparsity pattern (alternatively, identify the components of \mathbf{w} that can be estimated from \mathbf{g}') and (2) estimate those elements of \mathbf{w} that are informed by \mathbf{g}' while setting the rest to zero. The former can be realized by minimizing the ℓ_0 norm of \mathbf{w} while the latter is typically achieved by minimizing the ℓ_2 norm of the measurement – model discrepancy. However, an objective function that contains a ℓ_0 norm is discontinuous, and consequently ℓ_0 is replaced by an ℓ_1 norm, which is more tractable (Donoho et al., 2012). Thus, the optimization problem is posed as

$$\underset{\mathbf{w} \in \mathbb{R}^N}{\text{minimize}} \|\mathbf{w}\|_1, \quad \text{subject to } \|\mathbf{g}' - \mathbf{A}\mathbf{w}\|_2 < \epsilon_2. \quad (4)$$

This optimization problem can be solved using methods like MP, OMP and StOMP. Bayesian equivalents also exist (Ji et al., 2008; Babacan et al., 2010), where Laplace priors are used to enforce sparseness in the inferred \mathbf{w} . Algorithms based on convex optimization that serve the same purpose are reviewed in Jafarpour (2013). All these algorithms are general and do not exploit any particular structure in \mathbf{g} except sparsity. However, one may also create a prior model for wavelet distributions, e.g., by using a database of similar images, for higher quality reconstructions (Duarte et al., 2005; La and Do, 2005; Baraniuk et al., 2010; He and Carin, 2009). In order to do so, sparse reconstruction methods have to be modified to incorporate prior information.

Sparsity is sometimes used to solve inverse problems in physics, with the Ψ operator representing the physical process. Most of these inverse problems have been in the estimation of log-transformed permeability fields (Li and Jafarpour, 2010; Jafarpour, 2013), seismic tomography (Loris et al., 2007; Simons et al., 2011; Gholami and Siahkoobi, 2010) and estimation of point and distributed emissions (Hirst et al., 2013; Martinez-Camara et al., 2013). A more detailed review of the sparse reconstruction methods can be found in our previous paper (Ray et al., 2014). Most of these inverse problems involved nonlinear models, i.e., $\mathbf{y} = \mathbf{a}(\mathbf{w})$, rather than $\mathbf{y} = \mathbf{A}\mathbf{w}$, for which incoherence (and RIP) are not well defined and consequently were not investigated.

To summarize, sparse reconstruction techniques and wavelet-based random field models have been used in nonlinear inverse problems. In contrast, the problem of estimation of spatially rough emission fields is linear, raising the possibilities that (1) the same approach may offer a solution to the emission estimation problem and (2) mutual incoherence may provide analytical metrics for the quality of observations and, consequently, solutions. We build on the principles of compressive sensing and sparse reconstruction methods to design an inversion scheme for rough emission fields. In particular, we show (using coherence metrics) why the use of f_{pr} was necessary. We also show the degree of computational saving achieved when we use random projections to limit ffCO₂ emissions within \mathcal{R} .

3 Formulation of the estimation problem

Ray et al. (2014) developed a MsRF model for ffCO₂ emissions in the USA. The MsRF model allows ffCO₂ emissions to be represented as $\mathbf{f} = \Phi\mathbf{w}$, where Φ is a collection of Haar wavelets. Consequently, the observational term in Eq. (1) can be written as $\|\mathbf{y}^{\text{obs}} - \mathbf{H}\Phi\mathbf{w}\|_2^2$. Compared with Eq. (2), we see that the transport model \mathbf{H} serves as the sampling matrix Ψ . Since we seek to estimate the wavelet weights \mathbf{w} from \mathbf{y}^{obs} , an optimization problem like Eq. (4) could be posed with the constraint $\|\mathbf{y}^{\text{obs}} - \mathbf{A}\mathbf{w}\|_2 < \epsilon_2$, $\mathbf{A} = \mathbf{H}\Phi$. In order to solve this problem via sparse reconstruction, one requires that \mathbf{H} and Φ be incoherent. As we will show in Sect. 4.2, the incoherence requirement is not met, and sparsity (solely) is not sufficient to solve the problem accurately (as tested in Sect. 4.1). Consequently, we modify StOMP to incorporate a prior emission field f_{pr} . We also adapt it to accommodate fields defined over irregularly shaped domains as well as to ensure non-negativity of the estimated field.

Let \mathbf{f} be a time-variant, non-negative field defined in an irregular region \mathcal{R} , gridded with $N_{\mathcal{R}}$ grid cells. In our case \mathbf{f} models ffCO₂ emission fields. The field is averaged over a time period T and covers K time periods, i.e., it is a vector $N_{\mathcal{R}}K$ long. \mathbf{f} drives a linear model of observations of ffCO₂ concentrations:

$$\mathbf{y}^{\text{obs}} = \mathbf{y} + \epsilon = \mathbf{H}\mathbf{f} + \epsilon, \quad (5)$$

where \mathbf{H} is the sensitivity matrix obtained from an atmospheric transport model (see Sect. 2), ϵ is the model–data mismatch due to measurement and transport model errors and \mathbf{y}^{obs} is a vector of time-variant measurements collected at N_s measurement towers. Each tower collects K_s measurements over the K time periods, i.e., \mathbf{y}^{obs} is a vector $K_s N_s$ long. The \mathbf{H} matrix is $(K_s N_s) \times (N_{\mathcal{R}} K)$.

3.1 Prior models

We employ two prior models in our work – the MsRF model for ffCO₂ emissions and a time-invariant approximation of ffCO₂ emissions f_{pr} . The MsRF is a collection of wavelets and model emissions in the logically rectangular domain given by the corners 24.5° N, 63.5° W and 87.5° N, 126.5° W. The MsRF discretizes the domain using a dyadic $2^M \times 2^M$ mesh. Haar wavelets are defined on all M levels of this dyadic grid, but not all of them are retained in the MsRF. Wavelets constituting the MsRF model are chosen using radiance-calibrated images of lights at night (http://ngdc.noaa.gov/eog/data/web_data/v4composites/F152002.v4.tar; Cinzano et al., 2000), which serve as a proxy for human activity and thus capture the spatial patterns of ffCO₂ emissions. The emission field is allowed to assume non-zero values only within \mathcal{R} , the lower 48 states of the USA. We denote the field during the k th time period as f_k and model it as

$$\mathbf{f}_k = w'_k \phi' + \sum_{s=1}^M \sum_{i,j} w_{s,i,j,k} \phi_{s,i,j}, \quad \{s, i, j\} \in W^{(s)} = \Phi \mathbf{w}_k, \quad (6)$$

where $W^{(s)}$ contains the L wavelets that constitute the MsRF model, and L is a fraction of the 4^M wavelets that can be supported by a $2^M \times 2^M$ mesh.

The MsRF is also the starting point for developing the second prior model f_{pr} . The MsRF provides a sparse representation of the radiances $\mathbf{X}^{(s)}$:

$$\mathbf{X}^{(s)} = w'_{(X)} \phi' + \sum_{l,i,j} w_{(X),l,i,j} \phi_{l,i,j}, \quad \{l, i, j\} \in W^{(s)}. \quad (7)$$

$\mathbf{X}^{(s)}$ is used to calculate a time-invariant prior model for ffCO₂ emissions as $f_{\text{pr}} = c\mathbf{X}^{(s)}$. c is computed such that

$$\int_{\mathcal{R}} \overline{f_V} dA = \int_{\mathcal{R}} f_{\text{pr}} dA = c \int_{\mathcal{R}} \mathbf{X}^{(s)} dA = c \int_{\mathcal{R}} \left(w'_{(X)} \phi' + \sum_{l,i,j} w_{(X),s,i,j} \phi_{l,i,j} \right) dA, \quad \{l, i, j\} \in W^{(s)}. \quad (8)$$

Equation (8) implies that c is calculated such that both $\overline{f_V}$ and f_{pr} provide the same value for the total emissions in \mathcal{R} .

$\overline{f_V}$ in our case is the annually averaged 2005 emission field obtained from EDGAR. The leftmost term $\int_{\mathcal{R}} \overline{f_V} dA$ quantifies the total emissions as predicted by EDGAR over \mathcal{R} . The term $c \int_{\mathcal{R}} \mathbf{X}^{(s)} dA$ is an estimate of emissions over the same region but modeled using f_{pr} . The rightmost term simply replaces $\mathbf{X}^{(s)}$ with its wavelet model, in accordance with Eq. (7). The details of how the MsRF and f_{pr} were constructed are in Ray et al. (2014). f_{pr} differs from the ground truth (Vulcan emissions aggregated over the lower 48 states) by 5–25 % (see Fig. 9 in Ray et al., 2014).

3.2 Posing and solving the inverse problem

We seek emissions over an entire year (360 days), i.e., we seek $\mathbf{F} = \{f_1, f_2, \dots, f_K\} = \{\Phi w_1, \Phi w_2, \dots, \Phi w_K\} = \tilde{\Phi} w$. \mathbf{F} models the field in $\mathcal{R} \cup \mathcal{R}'$, where \mathcal{R}' models the region outside \mathcal{R} (but inside the rectangular domain modeled by the MsRF) with zero ffCO₂ emissions. We separate out the fluxes in \mathcal{R} and \mathcal{R}' by permuting the rows of $\tilde{\Phi}$

$$\mathbf{F} = \begin{pmatrix} \mathbf{F}_{\mathcal{R}} \\ \mathbf{F}_{\mathcal{R}'} \end{pmatrix} = \begin{pmatrix} \tilde{\Phi}_{\mathcal{R}} \\ \tilde{\Phi}_{\mathcal{R}'} \end{pmatrix} w,$$

where $\tilde{\Phi}_{\mathcal{R}}$ and $\tilde{\Phi}_{\mathcal{R}'}$ are $(N_{\mathcal{R}}K) \times (LK)$ and $(N_{\mathcal{R}'}K) \times (LK)$ matrices, respectively. Here $N_{\mathcal{R}'}$ is the number of grid cells in \mathcal{R}' . The modeled concentrations at the measurement towers, caused by $\mathbf{F}_{\mathcal{R}}$, can be written as $\mathbf{y} = \mathbf{H}\mathbf{F}_{\mathcal{R}}$. For arbitrary w , $\mathbf{F}_{\mathcal{R}'}$ (the emissions in \mathcal{R}') are not zero and $\mathbf{F}_{\mathcal{R}'} = 0$ will have to be imposed as a constraint in the inverse problem.

Specifying the constraint in individual grid cells is not very efficient since it leads to $N_{\mathcal{R}'}K$ constraints. This can get very large in a global inversion at high spatial resolutions. Instead, we adapt an approach from compressive sensing to enforce this constraint approximately. Consider a $M_{cs} \times (N_{\mathcal{R}'}K)$ matrix \mathbf{R} , whose rows are direction cosines of random points on the surface of $N_{\mathcal{R}'}K$ -dimensional unit sphere. This matrix is called a uniform spherical ensemble (Tsaig and Donoho, 2006). The M_{cs} projections of the emission field $\mathbf{F}_{\mathcal{R}'}$ on \mathbf{R} , i.e., $\mathbf{R}\mathbf{F}_{\mathcal{R}'}$, compressively samples $\mathbf{F}_{\mathcal{R}'}$ and setting them to zero during inversion allows us to enforce zero emissions outside \mathcal{R} . In Sect. 4.3, we will investigate the degree of computational saving afforded by imposing the $\mathbf{F}_{\mathcal{R}'} = 0$ constraint in this manner. The problem is now modeled as

$$\mathbf{Y} = \begin{pmatrix} \mathbf{y}^{\text{obs}} \\ 0 \end{pmatrix} \approx \begin{pmatrix} \mathbf{H} & \tilde{\Phi}_{\mathcal{R}} \\ \mathbf{R} & \tilde{\Phi}_{\mathcal{R}'} \end{pmatrix} w = \mathbf{G}w. \quad (9)$$

In this equation, \mathbf{G} is akin to \mathbf{A} in Eq. (2). The left hand side \mathbf{Y} is approximately equal to $\mathbf{G}w$ since the observations \mathbf{y}^{obs} contain measurement errors that cannot be modeled with \mathbf{H} .

The case where \mathcal{R}' contains non-zero emissions requires the use of boundary fluxes and is discussed in Ray et al. (2014).

The wavelet coefficients w in Eq. (9) are not normalized and usually display a large range of magnitudes. The wavelets in $W^{(s)}$ at finer scales, i.e., those with a small support, tend to have coefficients with a large magnitude. Their

small support cause the fine-scale wavelets to impact only neighboring measurement towers. In contrast, wavelets at the coarser scales have large *footprints* that span multiple measurement locations. Total emissions in \mathcal{R} , as well as \mathbf{y}^{obs} , are very sensitive to their coefficients. Solving Eq. (9) as it is incorporates no information from f_{pr} beyond the selection of wavelets to be included in $\tilde{\Phi}$. We explore the incorporation of f_{pr} in the estimation of w using three different approaches:

Approach A: this is the baseline approach and solves Eq. (9) as it is. The lack of normalization of w , in conjunction with the sparse reconstruction procedure described below, leads to artifacts that will be described in Sect. 4.1.

Approach B: in this formulation, we include f_{pr} as a prior. We write the emissions as $\mathbf{F} = f_{pr} + \Delta\mathbf{F}$. Substituting into Eq. (9), we get $\mathbf{Y} \approx \mathbf{H}f_{pr} + \mathbf{G}\Delta w$, where $\Delta w = w - w_{(X)}$. Here, $w_{(X)} = c\{w'_{(X)}, w_{(X),s,i,j}\}$, $\{s, i, j\} \in W^{(s)}$, where c is obtained from Eq. (8). Simplifying, we get

$$\Delta\mathbf{Y} = \mathbf{Y} - \mathbf{H}f_{pr} \approx \mathbf{G}\Delta w. \quad (10)$$

Approach C: in approach B we expressed the true flux \mathbf{F} as an additive correction over f_{pr} , thus incorporating the prior information in f_{pr} . In approach C, we use the spatial pattern of f_{pr} , as captured by its wavelet coefficients $w_{(X)}$, to normalize w . We rewrite Eq. (9) as

$$\mathbf{Y} \approx \mathbf{G}\mathbf{B}\mathbf{B}^{-1}w = \mathbf{G}'w' = \begin{pmatrix} \mathbf{H}\tilde{\Phi}'_{\mathcal{R}} \\ \mathbf{R}\tilde{\Phi}'_{\mathcal{R}'} \end{pmatrix} w', \quad (11)$$

$$\mathbf{B} = c \text{diag}(w_{(X)}), \quad (12)$$

where $w' = \{w_{s,i,j}/(c w_{(X),s,i,j})\}$, $\{s, i, j\} \in W^{(s)}$ are the wavelet coefficients normalized by those of f_{pr} , $\tilde{\Phi}'_{\mathcal{R}} = \tilde{\Phi}_{\mathcal{R}} \text{diag}(w_{(X)})$ and $\tilde{\Phi}'_{\mathcal{R}'} = \tilde{\Phi}_{\mathcal{R}'} \text{diag}(w_{(X)})$. If f_{pr} is close to \mathbf{F} , the elements of w' will be $O(1)$. If f_{pr} is a gross underestimate, the elements of w' will still be of the same order of magnitude, but not $O(1)$. Thus, normalization with $w_{(X)}$ removes the large differences that exist between the wavelet coefficients at different scales.

In all the three cases, we obtain an underdetermined set of linear equations of the form

$$\mathbf{Y} \approx \mathbf{\Gamma}\zeta. \quad (13)$$

Here \mathbf{Y} represents \mathbf{Y} in approaches A and C, and $\Delta\mathbf{Y}$ in approach B. $\mathbf{\Gamma}$ represents \mathbf{G} in approaches A and B and \mathbf{G}' in approach C. ζ represents w in approach A, Δw in approach B and w' in approach C.

Since \mathbf{y}^{obs} is obtained from a set of locations sited with an eye towards biospheric CO₂ fluxes (see Ray et al., 2013), it is unlikely that it will allow the estimation of all the elements of ζ . Further, a priori, we do not know the identity of these unestimateable elements and so we use sparse reconstruction to find and compute them. Equation (13) is recast similar to Eq. (4):

$$\underset{\zeta \in \mathbb{R}^N}{\text{minimize}} \quad \|\zeta\|_1, \quad \text{subject to} \quad \|\mathbf{Y} - \mathbf{\Gamma}\zeta\|_2^2 < \epsilon_2. \quad (14)$$

We solve Eq. (14) using StOMP. $\|\zeta\|_1$ is minimized by setting to zero as many elements of ζ as possible, thus enforcing sparsity. Meanwhile, the constraint $\|\Upsilon - \Gamma\zeta\|_2$ ensures that the solutions being proposed by the optimization procedure provide a good reproduction of the observations. Note ζ contains only the wavelets in $W^{(s)}$. The StOMP algorithm is detailed in Donoho et al. (2012). We will refer to this step in the estimation procedure as step I.

3.3 Enforcing non-negativity on $F_{\mathcal{R}}$

Estimates of \mathbf{w} calculated by StOMP do not necessarily provide $F_{\mathcal{R}} = \tilde{\Phi}_{\mathcal{R}}\mathbf{w}$ that are non-negative. In practice, negative values of $F_{\mathcal{R}}$ occur in only a few grid cells and are usually small in magnitude. A large fraction of elements of \mathbf{w} are set to zero by StOMP. Having identified the sparsity pattern, i.e., the spatial scales that can be estimated from \mathbf{y}^{obs} , we devise an iterative procedure for enforcing non-negativity on $F_{\mathcal{R}}$. We discard $F_{\mathcal{R}'}$ and manipulate the field (the emissions) in \mathcal{R} directly, rather than via the wavelet coefficients.

We seek the non-negative vector $\mathbf{E} = \{E_i\}$, $i = 1 \dots Q$, $Q = (N_{\mathcal{R}}K)$ such that

$$\frac{\|\mathbf{y}^{\text{obs}} - \mathbf{H}\mathbf{E}\|_2}{\|\mathbf{y}^{\text{obs}}\|_2} \leq \epsilon_3. \quad (15)$$

\mathbf{E} is constructed iteratively through a sequence $\mathbf{E}_1, \mathbf{E}_2, \dots, \mathbf{E}_N$. \mathbf{E}_0 is initialized by using the absolute values of $F_{\mathcal{R}}$ calculated by solving Eq. (14). At each iteration m , we seek a correction $\xi = \{\xi_i\}$, $i = 1 \dots Q$, where $|\xi_i| \leq 1$, such that

$$\begin{aligned} \mathbf{E}^{(m)} &= \text{diag}(\exp(\xi_1), \exp(\xi_2), \dots, \exp(\xi_Q))\mathbf{E}^{(m-1)} \\ &\approx \text{diag}(1 + \xi_1, 1 + \xi_2, \dots, 1 + \xi_Q)\mathbf{E}^{(m-1)} \\ &= \mathbf{E}^{(m-1)} + \Delta\mathbf{E}^{(m-1)}, \text{ where } \Delta\mathbf{E}^{(m-1)} = \xi^T \mathbf{E}^{(m-1)}. \end{aligned}$$

Since the field must satisfy $\mathbf{y}^{\text{obs}} \approx \mathbf{H}\mathbf{E}^{(m)}$, we get

$$\mathbf{y}^{\text{obs}} - \mathbf{H}\mathbf{E}^{(m-1)} = \Delta\mathbf{y} \approx \mathbf{H}\Delta\mathbf{E}^{(m-1)}. \quad (16)$$

This is an underconstrained problem, and we seek the sparsest set of updates $\Delta\mathbf{E}^{(m-1)}$ using StOMP. The corrections are calculated, and the field updated as

$$\begin{aligned} \xi_i &= \text{sgn}\left(\frac{\Delta E_i^{(m-1)}}{E_i^{(m-1)}}\right) \max\left(1, \left|\frac{\Delta E_i^{(m-1)}}{E_i^{(m-1)}}\right|\right), \\ E_i^{(m)} &= E_i^{(m-1)} \exp(\xi_i) \end{aligned} \quad (17)$$

to obtain $\mathbf{E}^{(m)}$. The convergence requirement in Eq. (15) is checked with $\mathbf{E}^{(m)}$, and if not met, the iteration count is updated $m := m + 1$ and Eq. (16) is solved again.

We will refer to this step in the estimation procedure as step II.

4 Numerical results

In this section, we test the sparse estimation technique in Sect. 3, using synthetic observations. The time period T over

which the ffCO₂ emissions are averaged is 8 days. $K = 45$, i.e., we estimate emissions over $8 \times 45 = 360$ days. $N_s = 35$ towers, which are a subset of NOAA's Earth System Research Laboratory (ESRL) Global Monitoring Division's cooperative air sampling network (Tans and Conway, 2005); their locations are in Ray et al. (2013). These towers provide continuous observations of CO₂ concentrations (in parts per million by volume, ppmv), and 3-hourly averaged synthetic observations are used here (i.e., $K_s = 24/3 \times 8 \times 45 = 2880$). We discretize the domain covered by the MsRF using $1^\circ \times 1^\circ$ grid cells i.e., $M = 6$. The number of grid cells in the entire domain (the rectangle with the corners 24.5° N, 63.5° W and 87.5° N, 126.5° W), N , is $4^M = 4096$, which is also equal to the number of wavelets that can be defined on the mesh. The number of wavelets retained in the MsRF, L , is 1031. \mathcal{R} denotes the lower 48 states of the USA. They are covered with $N_{\mathcal{R}} = 816$ grid cells. The number of grid cells outside \mathcal{R} , $N_{\mathcal{R}'} = N - N_{\mathcal{R}} = 3280$.

The \mathbf{H} matrix in Eq. (5) is calculated per the description in Gourdji et al. (2012). We use the Stochastic Time-Inverted Lagrangian Transport Model (Lin et al., 2003), with wind fields from the Weather Research & Forecasting model (Skamarock and Klemp, 2008), version 2.2, driven by 2008 meteorology to compute \mathbf{H} . Concentration sensitivities are calculated at 3 h intervals over a North American grid, at a resolution of $1^\circ \times 1^\circ$. The sensitivity of the CO₂ concentration at each observation location due to the flux at each grid cell is calculated in units of ppmv $\mu\text{mol}^{-1} \text{m}^2 \text{s}^{-1}$. The sensitivity of \mathbf{y} to the 8-day-averaged emissions were obtained from the 3 h sensitivities by simply adding the $8 \times 24/3 = 64$ sensitivities that span the 8-day period.

The true ffCO₂ emissions in \mathcal{R} are obtained, for 2002, from the Vulcan inventory. Hourly Vulcan fluxes are coarsened from 0.1° resolution to 1° , and averaged to 8-day periods. These 8-day-averaged fluxes at 1° resolution are multiplied by \mathbf{H} to obtain ffCO₂ concentrations at the measurement towers. Note that averaging over 8 days removes the diurnal variations of ffCO₂ emissions in Vulcan. Observations are generated every 3 h and span a full year. A measurement error $\epsilon \sim N(0, \sigma^2)$ is added to the concentrations to obtain \mathbf{y}^{obs} (see Eq. 5), as used in Eq. (9). The same σ is used for all towers. We use $\sigma = 0.1$ ppmv, which is too small and represents an idealized inversion scenario that is used here to test the quality of the proposed numerical method. Realistic values of transport model errors for some of the towers used in this study are in Gourdji et al. (2012). Radiocarbon measurement errors can be found in Turnbull et al. (2011).




Figure 1. Plot of ffCO₂ emissions (micromoles of C m⁻² s⁻¹) as reported by EDGAR for 2005. Emissions below 0.02 micromoles of C m⁻² s⁻¹ are grayed out.

4.1 Comparison of optimization formulations

We choose between approaches A, B and C by solving the inverse problem for the ffCO₂ emission field. The inversion is performed for the emissions $\mathbf{F} = \{f_k\}, k = 1 \dots K$, for the entire year. The following parameters are used in the inversion process: $\epsilon_2 = 10^{-5}$, $\epsilon_3 = 5.0 \times 10^{-4}$, $M_{cs} = 13\,500$, i.e., 300 random projections for each 8-day period. The rationale for these values can be found in our previous paper (Ray et al., 2014).

In Fig. 2 we plot the estimated emissions during the 31st 8-day period, as calculated using approaches A, B and C. The true emissions are also plotted for reference. Four quadrants are also plotted for easier comparison and reference. The distribution of measurement towers is very uneven, with most of the towers being concentrated in the northeast (NE) quadrant, where we expect the reconstruction to be most accurate. We see that approach A (Fig. 2, top right) provides estimates that have large areas in the northwest (NW) and southwest (SW) quadrants with moderate levels of ffCO₂ emissions. In contrast, the true emissions (Fig. 2, top left) are mostly empty. Thus, we see that the minimization of $\|\zeta\|_1$ (alternatively $\|\mathbf{w}\|_1$) drives the wavelet coefficients to small values, but not identically to zero. In Fig. 2 (bottom left), approach B provides estimates that show much structure in the eastern quadrants, and the patterns seen in f_{pr} (see Ray et al., 2014) are reproduced. The reason is as follows. While f_{pr} captures the broad, coarse-scale patterns of ffCO₂ emissions, it incurs significant errors at the finer scales. Equation (10) seeks to rectify the discrepancy between f_{pr} and true emissions using observations. However, as mentioned in Sect. 3.2, fine-scale wavelets tend to have large wavelet coefficients and the minimization of $\|\zeta\|_1$ (alternatively $\|\Delta\mathbf{w}\|_1$) removes them since the constraint $\|\Upsilon - \Gamma\zeta\|_2^2 < \epsilon_2$ is not very sensitive to individual wavelets at the fine scale. (See Gerbig et al. (2009) for a discussion on the largely local impact of a CO₂

flux source.) The inability to rectify the fine-scale discrepancies led to a final ffCO₂ estimate that resembles f_{pr} in the finer details. Figure 2 (bottom right) plots the estimates obtained using approach C, which uses normalized wavelet coefficients \mathbf{w}' . The estimates from approach C show large areas of little or no emissions in the western quadrants, similar to the true emissions in the top-left figure. In the eastern quadrants, the emissions show less spatial structure than the true emissions as well as those obtained using approach A.

The quality of the estimate is due to both the MsRF model and the new sparse reconstruction scheme. The limited observations are sufficient to allow the estimation of the coarse MsRF wavelets, and in certain areas, e.g., the NE quadrant, finer details. The MsRF model is sufficiently flexible to accommodate the spatial heterogeneity in detail, but requires a sparse reconstruction method to address the high dimensionality that such flexibility entails. Further, the multi-resolution nature of MsRF model allows for the accurate estimation of coarse-scale patterns of ffCO₂ emissions, i.e., we expect that aggregate measures of emission quality, such as integrated emissions in \mathcal{R} , will be accurate. It will incur larger errors as the domain of integration has shrunk.

In Fig. 3 (top) we evaluate the accuracy of the reconstruction quantitatively. We integrate the emissions in \mathcal{R} to obtain the country-level ffCO₂ emissions and compare that with the emissions from Vulcan. We plot a time series of errors defined as a percentage of total, country-level Vulcan emissions

$$\text{Error}_k (\%) = \frac{100}{K} \sum_{k=1}^K \frac{E_k - E_{V,k}}{E_{V,k}},$$

$$\text{where } E_k = \int_{\mathcal{R}} \mathbf{E}_k \, dA \text{ and } E_{V,k} = \int_{\mathcal{R}} \mathbf{f}_{V,k} \, dA. \quad (18)$$

Here, $\mathbf{f}_{V,k}$ are Vulcan emissions averaged over the k th 8-day period and \mathbf{E}_k are the non-negativity enforced emission estimates in the same time period. A positive error denotes an overestimation by the inverse problem. In Fig. 3 (bottom) we plot the Pearson correlation coefficient between the true and reconstructed emissions in \mathcal{R} over the same duration. We define the Pearson correlation coefficient between \mathbf{E}_k and $\mathbf{f}_{V,k}$ as

$$C(\mathbf{E}_k, \mathbf{f}_{V,k}) = \frac{\text{cov}(\mathbf{E}_k, \mathbf{f}_{V,k})}{\sigma_{\mathbf{E}_k} \sigma_{\mathbf{f}_{V,k}}},$$

where $\sigma_{\mathbf{E}_k}^2$ and $\sigma_{\mathbf{f}_{V,k}}^2$ are the variances of the true and reconstructed fluxes and $\text{cov}(Z_1, Z_2)$ is the covariance between two random variables Z_1 and Z_2 . It is clear that approach B provides the worst reconstructions, with the largest errors and smallest correlations. Approach C tends to over-predict emissions a little more than approach A, but has better spatial correlation with the Vulcan emissions.

In Fig. 4 we see the essential difference between approach A and C. We plot the reconstruction error (top) and

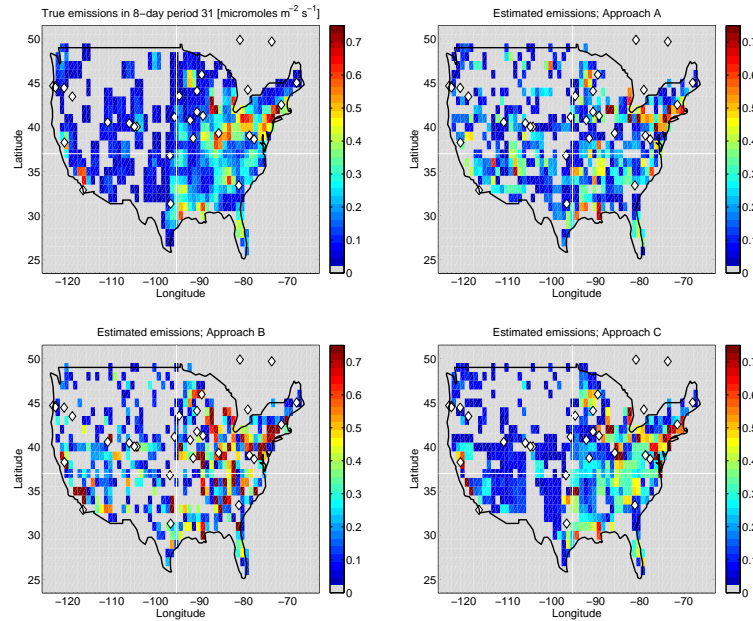


Figure 2. Plots of ffCO_2 emissions during the 31st 8-day period. The units are micromoles of $\text{C m}^{-2} \text{s}^{-1}$. Emissions below $0.02 \text{ micromoles of C m}^{-2} \text{s}^{-1}$ are grayed out. Top left, we plot true emissions from the Vulcan inventory. Top right, the estimates from approach A. Bottom left and right figures contain the estimates obtained from approaches B and C, respectively. Each figure contains the measurement towers as white diamonds. Each figure is also divided into quadrants. We see that approach A, unconstrained by f_{pr} provides low levels of (erroneous) emissions in large swathes of the western quadrants. Approach B reflects f_{pr} very strongly. Approach C provides a balance between the influence of f_{pr} and the information in y^{obs} .

correlation between true and reconstructed emissions (bottom) in the northeast (NE) and northwest (NW) quadrants. Errors in the emissions are represented as a percentage of the total (true) emissions in that quadrant. We see the approach C has smaller errors in both the quadrants. It also provides higher correlation in the NW quadrant, which does not have many measurement towers (white diamonds in Fig. 2). Both the approaches have errors of opposite signs in the quadrants which largely cancel out when errors are assessed over \mathcal{R} as a whole, leading to approximately similar estimation accuracies by both the approaches in Fig. 3. However, the estimates produced by approach A (without the use of f_{pr}) show larger spatial variability and error than approach C. This is because normalization using $w_{(X)}$ and minimization of $\|\xi\|_1$ (alternatively $\|w'\|_1$) prevents large departures from f_{pr} and also rectifies the tendency to remove large wavelet coefficients belonging to the finer wavelets. Approach C therefore provides a formulation that is more accurate and robust at the quadrant scale, even though both have similar fidelity at the scale of \mathcal{R} .

4.2 Evaluating formulation using compressive sensing metrics

Having established empirically that approach A is less accurate than approach C, we can explain why this is the case. We employ coherence metrics for this purpose.

In compressive sensing, random matrices such as Gaussians, Hadamard, Circulant/Toeplitz or functions such as noiselets (Tsaig and Donoho, 2006; Gan et al., 2008; Yin et al., 2010; Tuma and Hurley, 2009) serve as Ψ . In Fig. 5, we plot the distribution of $\log_{10}(|A_{i,j}|)$, the elements of $\mathbf{A}_\Psi = \Psi\Phi$ for these standard sampling matrices. Φ contains only the wavelets in $W^{(s)}$. Note that $\max(|A_{i,j}|)$ specifies the mutual coherence, and small values of $\max(|A_{i,j}|)$ indicate informative measurements. We see that $\log_{10}(|A_{i,j}|)$ may assume continuous (Gaussian and circulant sampling matrices) or discrete (Hadamard, scrambled-block Hadamard and noiselets) distributions, and generally lie between -3 and -1 . This provides a range for the level of coherence observed in theoretical CS analyses.

In Eq. (9), \mathbf{H} serves a similar sampling purpose, and the efficiency of sampling depends on the incoherence between \mathbf{H} and Φ . We construct a new \mathbf{H}' by picking the rows of \mathbf{H} corresponding to two towers and for the 21st and 22nd 8-day periods. We compute $\mathbf{A}_{\mathbf{H}'} = \mathbf{H}'\Phi$, and in Fig. 5, plot the log-transformed magnitudes of the elements of $\mathbf{A}_{\mathbf{H}'}$. The distributions for the two towers are almost identical. We clearly see that, unlike \mathbf{A}_Ψ , $\mathbf{A}_{\mathbf{H}'}$ contains a significant number of elements that are close to 1, and a large number of elements that are close to 0 (e.g., near 10^{-6}). This is a consequence of the rows of \mathbf{H}' being approximately aligned to some of the columns of Φ and, consequently, nearly orthogonal to oth-

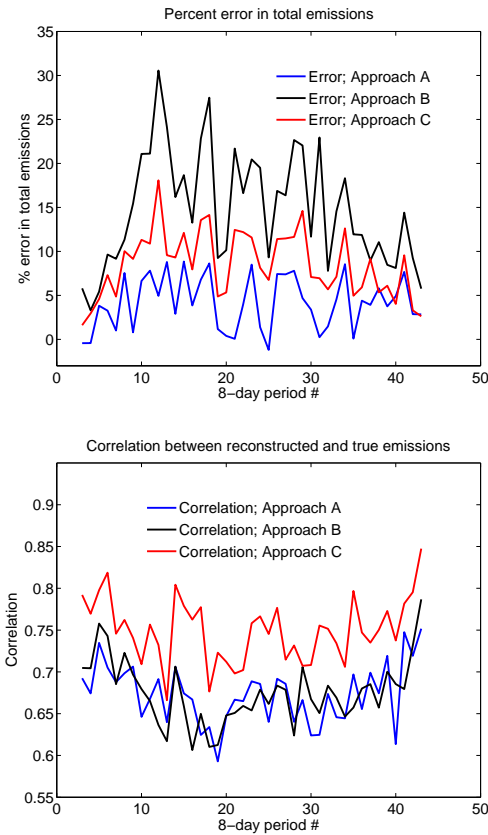


Figure 3. Comparison of estimation error (top) and the correlation between true and estimated emissions (bottom) using approaches A, B and C. It is clear that approach B is inferior to the others.

ers. The small values in \mathbf{A}_H indicate that the CO_2 concentration prediction \mathbf{y} at the two selected towers are insensitive to many of the wavelets, i.e., to many scales and locations, as observed in Sect. 4.1. Further, the coherence $\mu(H', \Phi)$ is larger than $\mu(\Psi, \Phi)$, indicating a sampling efficiency a few orders of magnitude inferior to those achieved in the CS of images. Consequently, approach A, based solely on sparsity, and identical to the method adopted in CS, would not work well. Thus, approach C, which employed both sparsity and f_{pr} , proved superior to approach A.

4.3 Numerical consistency and computational efficiency

We now address some of the numerical aspects of the solution. The results presented here are *not* tests of accuracy of the estimated emission field; estimation accuracy also depends on the MsRF and was investigated in Ray et al. (2014). Here we empirically verify that certain necessary conditions of our sparse reconstruction are satisfied.

In Fig. 6 (top), we plot \mathbf{y} predicted by the reconstructed emissions at two towers, BAO (Boulder Atmospheric Observatory, Colorado) and MAP (Mary’s Peak, Oregon). These

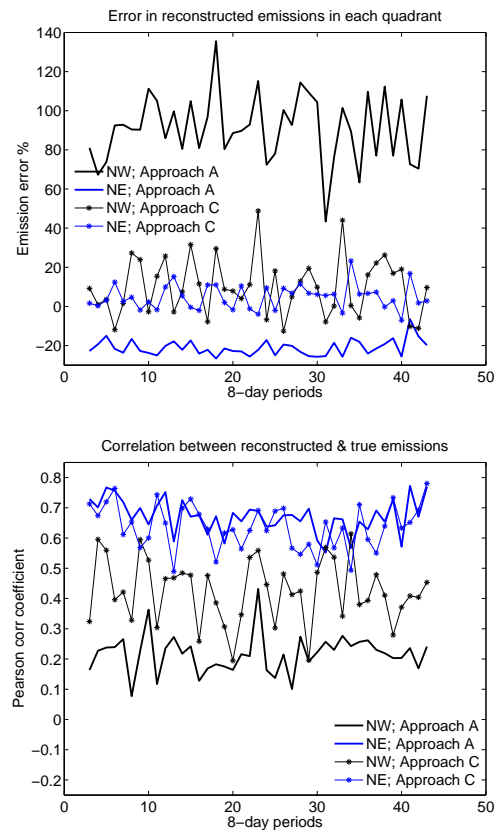


Figure 4. Reconstruction error (top) and correlation between the true and estimated emissions, using approaches A and C, for the northeast (NE) and northwest (NW) quadrants. We see that approach C, which includes information from f_{pr} , leads to lower errors in both the quadrants and better correlations in the less instrumented NW quadrant.

towers were included in the inversion and are *not* being used as an out-of-sample test of the accuracy of the estimated emission field. Rather, the MsRF for rough fields allows the estimation of local sources which can help reproduce a tower’s measurements very closely, unless neighboring towers provide a constraint; in a sparse network, this is not always possible. Thus, an accurate reproduction of a tower’s observations is not necessarily a sign of an accurately estimated emission field, but a bad reproduction can be a sign of a malfunctioning sparse reconstruction method. We see that the ffCO_2 concentrations are well reproduced by the estimated emissions. In Fig. 6 (bottom) we plot the wavelet coefficients obtained by projecting the emissions (both the true and reconstructed) on the wavelet bases. The wavelet coefficient values have been subjected to a hyperbolic tangent transformation for ease of plotting. The true wavelet coefficients with a magnitude above 0.01 are plotted with red symbols. The true (Vulcan) emissions have a large number of coefficients with small magnitude; these are usually for small-scale features, i.e., have coefficient indices in the right half of




Figure 5. Comparison of the distribution of the elements of \mathbf{A}_Ψ and \mathbf{A}_Φ . We see that Gaussian and circulant random matrices lead to continuous distributions, whereas Hadamard, scrambled-block Hadamard (sbHadamard) and noiselets serving as sampling matrices lead to \mathbf{A}_Ψ where the elements assume discrete values. In contrast, the elements of \mathbf{A}_Φ assume values which are spread over a far larger range, some of which are quite close to 1 while others are very close to 0.

the range (red symbols in Fig. 6, bottom). During sparse reconstruction, these coefficients are set to zero (blue symbols in Fig. 6, bottom). The low-index coefficients, which represent large structures, are estimated accurately. The explicit separation of scales is thus leveraged into omitting fine-scale details which are difficult to inform with data and focusing model-fitting effort on the large scales instead. Sparse reconstruction achieves this in an automatic, purely data-driven manner, rather than via a pre-processing, scale-selection step.

Finally, we address the issue of enforcing the $\mathbf{F}_{\mathcal{R}'} = 0$ constraint via random M_{cs} projections. Naively, the constraint can be enforced for every individual grid cell, resulting in $N_{\mathcal{R}'} = 3280$ linear equations per 8-day period in Eqs. (9) and (13). Considering that $\mathbf{y}^{obs} = \mathbf{H}\Phi_{\mathcal{R}}$ results in $64 \times 35 = 3240$ linear equations per 8-day period, we see that enforcing the constraint is as expensive as computing $\mathbf{F}_{\mathcal{R}}$. Instead, we set $M_{cs} \ll N_{\mathcal{R}'}$ random projections of $\mathbf{F}_{\mathcal{R}'}$ to zero in Eqs. (9) and (13), exploiting the basic efficiency-via-random-sampling tenet of CS. Since Eq. (13) is solved approximately, and due to the small number of wavelets in $W^{(s)}$ that span \mathcal{R}' , the constraint $\mathbf{F}_{\mathcal{R}'} = 0$ is not satisfied exactly. This error varies with M_{cs} ; a larger M_{cs} results in a closer realization of the constraint. Errors in the enforcement of the $\mathbf{F}_{\mathcal{R}'} = 0$ constraint lead to commensurate errors in $\mathbf{F}_{\mathcal{R}}$. Here we check the trade-off between M_{cs} (computational efficiency) and accuracy of the estimated emissions ($\mathbf{F}_{\mathcal{R}}$ and $\mathbf{F}_{\mathcal{R}'}$). In practice, this affects only step I of the procedure, where an approximation of ffCO₂ emissions is calculated, and thereafter it is used as a guess in step II. However, a good estimate of the emission field accelerates the second step. The quality of the solution from step I, quantified as the

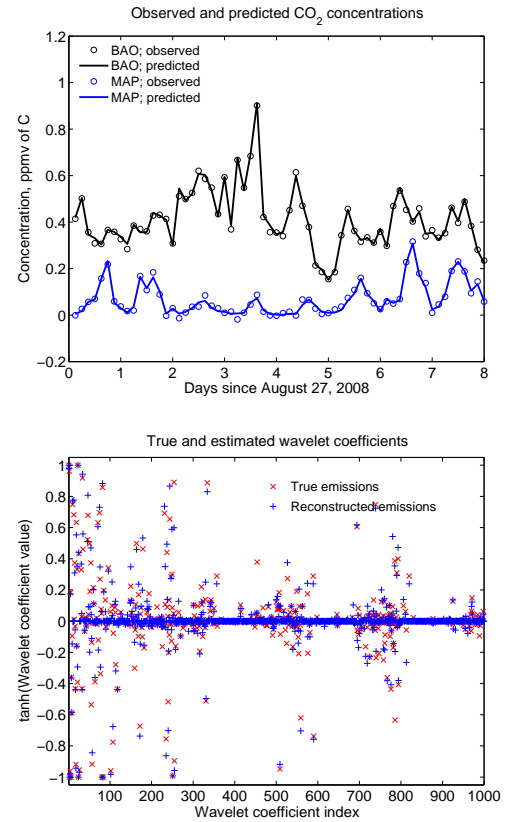


Figure 6. (Top) predictions of ffCO₂ concentrations at two measurement locations, using the true (Vulcan) and reconstructed emissions (blue lines) over an 8-day period (period no. 31). Observations occur every 3 h. We see that the concentrations are accurately reproduced by the estimated emissions. (Bottom) Projection of the true and estimated emissions on the wavelet bases for the same period. Coarse wavelets have lower indices, and they progressively get finer with the index number. We see that the true emissions have a large number of wavelets with small, but not zero, coefficients. In the reconstruction (plotted in blue), a number of wavelet coefficients are set to very small values (almost zero) by the sparse reconstruction. The larger scales are estimated accurately.

cumulative distribution function of the fluxes can be found in Ray et al. (2013, 2014). There are only a few grid cells with negative emissions and their magnitudes are small.

In Fig. 7, we plot the impact of M_{cs} on the reconstruction. We perform sparse reconstruction of the emission field, for the 31st 8-day period and compute the ratios

$$\eta_{\mathcal{R}} = \frac{\|\mathbf{f}_{k,\mathcal{R}}\|_2}{\|\mathbf{f}_{V,k}\|_2} \text{ and } \eta_{\mathcal{R}'} = \frac{\|\mathbf{f}_{k,\mathcal{R}'}\|_2}{\|\mathbf{f}_{V,k}\|_2} \text{ for } k = 31, \quad (19)$$

where $\mathbf{f}_{k,\mathcal{R}}$ and $\mathbf{f}_{k,\mathcal{R}'}$ are the emissions over \mathcal{R} and \mathcal{R}' from step I. $\mathbf{f}_{V,k}$ is the true (Vulcan) emission field during the same period. These ratios are plotted as a function of $\log_{10}(M_{cs})$ per 8-day period. We see that 10 projections per 8-day period is too few, leading to around 20 % errors in $\mathbf{f}_{k,\mathcal{R}'}$ ($\eta_{\mathcal{R}'} \approx 0.2$). Beyond about 100 projections per 8-day

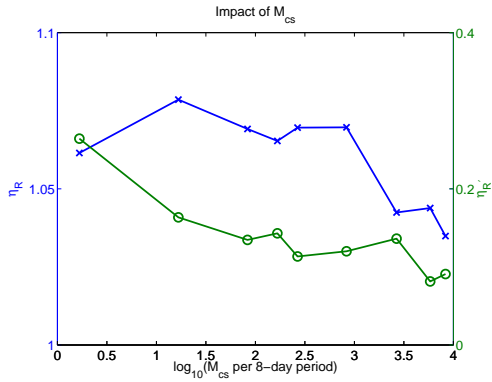


Figure 7. The impact of the number of compressive samples M_{cs} on the reconstruction of $F_{\mathcal{R}}$ ($\eta_{\mathcal{R}}$) and $F_{\mathcal{R}'}$ ($\eta_{\mathcal{R}'}$). $\eta_{\mathcal{R}}$ and $\eta_{\mathcal{R}'}$ are plotted on the Y1 and Y2 axes, respectively. Results are plotted for the 31st 8-day period. We see that $M_{cs} > 10^3$ does not result in an appreciable increase in reconstruction quality. Also, $M_{cs} < 10^2$ shows a marked degradation in $\eta_{\mathcal{R}'}$.

period, $\eta_{\mathcal{R}'}$ oscillates around 0.1. The corresponding errors in $f_{k,\mathcal{R}}$ are about 5 % ($\eta_{\mathcal{R}} \approx 1.05$). In our study we used 300 random projections for each 8-day period. This is about 10 % of the 3280 linear constraints that we would have enforced under a naive implementation of the $F_{\mathcal{R}'} = 0$ constraint. It also halves the computational cost of step I.

5 Conclusions

In this study, we have developed a sparse reconstruction scheme that could be used for solving physics-based linear inverse problems. Our method is an extension of stage-wise orthogonal matching pursuit (StOMP) (Donoho et al., 2012) and borrows many concepts from the compressive sensing (CS) and sparse reconstruction of images (Candes and Wakin, 2008). This scheme is useful for estimating non-stationary fields, e.g., permeability or flux fields, provided their random field model consists of independent parameters. This is typically achieved by representing the fields in terms of orthogonal bases, e.g., wavelets or Karhunen–Loève modes, if a prior covariance is available. The dimensionality of the resultant representation is not an issue; the sparse reconstruction method estimates only those parameters that are informed by the observations while setting the rest to zero.

Our new method has three novel characteristics. First, it can impose non-negativity on the estimated field, without resorting to log transformations. This retains the linear nature of the inverse problem and consequently, its computational efficiency. Second, it allows one to estimate geometrically irregular fields while using a random field model designed for rectangular domains. Third, it allows us to incorporate a prior model of the field being estimated into the sparse reconstruction procedure. While other model-based sparse reconstruction methods exist (Baraniuk et al., 2010; He and

Carin, 2009; La and Do, 2005), our method is simple and is seen empirically to recover the correct solution.

We have demonstrated our method in an atmospheric inverse problem for the estimation of a spatially rough emission field. It is an idealization of the estimation of ffCO₂ emissions in \mathcal{R} , the lower 48 states of the USA. The emissions were modeled in a square domain, with a 64×64 grid, using a recently developed multiscale random field (MsRF) model (Ray et al., 2014). It uses Haar wavelets and images of lights at night to capture the spatial patterns of ffCO₂ emission fields. The observational data consists of ffCO₂ measurements at a limited set of towers, which are linked to the emission field via a CO₂ transport model (the forward model). We draw parallels between our physics-based inverse problem and the sparse reconstruction of images in CS, and show that a fundamental CS tenet – incoherence – holds only approximately. Consequently, such inverse problems may not bear an accurate solution if they are regularized solely using sparsity. We demonstrate this in our study and show how incorporation of prior information, in the form of spatial patterns in images of lights at night, and a prior model of ffCO₂ emissions can enable a solution. We also demonstrate how CS concepts can be used to restrict the estimated field to an irregular region (in our case, \mathcal{R}) with a factor-of-ten less computational effort than a naive approach. Finally, we show how non-negativity of ffCO₂ emissions can be imposed using a simple post-processing step.

We also tested whether step I (Sect. 3.2) was necessary by bypassing it completely, and starting step II (Sect. 3.3) with E_0 initialized using an inventory. We do not present results of these tests in this paper, but find that the iterative scheme converges only when E_0 is very close to the true results. For example, initializing using perturbed Vulcan emissions led a converged solution, whereas f_{pr} did not. Thus, step I is required for robustness and generality. This is particularly relevant for developing countries where inventories contain larger errors.

Our sparse reconstruction scheme suffers from one serious drawback – it does not provide uncertainty bounds on the estimated field due to the paucity of data, and/or the shortcomings of the models. While this can be rectified using a Kalman filter, it does not provide any mechanism for reducing the dimensionality of the random field model, should the observational data prove inadequate. This is currently being investigated. Also, we assumed that there were no emissions outside \mathcal{R} , but in reality, there are. See our previous paper (Ray et al., 2014) on how they could be accommodated as boundary fluxes. Our use of the MsRF in the inversion is a second source of error; in the limit of a very informative measurement network, the accuracy of the inversion is limited by the ability of the MsRF to represent ffCO₂ fields accurately.

Due to the lack of a good tracer for ffCO₂ emissions, we demonstrated our method in an idealized inversion problem. The idealizations include a very small model – data mismatch ϵ (much smaller than what can be supported by

contemporary transport models and radiocarbon measurement technology) and an ability to measure ffCO₂ accurately, without interference from biospheric CO₂ fluxes (i.e., we treated ffCO₂ like a radiocarbon tracer). In order that our method could be used in a real-data inversion for ffCO₂ emissions, our method would need to be extended in a number of ways. First, we would require a measurement network better suited for ffCO₂ measurements, with sensors near large sources; one could be designed by conducting an observation system simulation experiment, perhaps using the method described here. Second, we would have to extend the method to perform a joint ffCO₂–biospheric CO₂ inversion, by including a spatial parameterization and priors for biospheric CO₂. Finally, we would have to devise a separate ϵ for each tower to reflect transport model errors.

In conjunction with this paper, we are also providing, at our website (Ray, 2013), the MATLAB[®] code required to construct the MsRF model for ffCO₂ emissions and perform the inversion using synthetic observations. The website also contains links to the (free) MATLAB[®] toolkits that our code depends on, along with a user's manual.

Acknowledgements. This work was supported by Sandia National Laboratories' LDRD (Laboratory Directed Research and Development) funds, sponsored by the Geosciences Investment Area. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the US Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Edited by: A. Stenke

References

- Andres, R. J., Boden, T. A., Bréon, F.-M., Ciais, P., Davis, S., Erickson, D., Gregg, J. S., Jacobson, A., Marland, G., Miller, J., Oda, T., Olivier, J. G. J., Raupach, M. R., Rayner, P., and Treanton, K.: A synthesis of carbon dioxide emissions from fossil-fuel combustion, *Biogeosciences*, 9, 1845–1871, doi:10.5194/bg-9-1845-2012, 2012.
- Babacan, S. D., Molina, R., and Katsaggelos, A. K.: Bayesian compressive sensing using Laplace priors, *IEEE T. Signal Proces.*, 19, 55–63, doi:10.1109/TIP.2009.2032894, 2010.
- Baraniuk, R., Davenport, M., DeVore, R., and Wakin, M.: A simple proof of the restricted isometry property for random matrices, *Constr. Approx.*, 28, 253–263, 2008.
- Baraniuk, R., Cevher, V., Duarte, M., and Hegde, C.: Model-based compressive sensing, *IEEE T. Inform. Theory*, 56, 1982–2001, 2010.
- Candes, E. and Tao, T.: Near optimal signal recovery from random projections: universal encoding strategies?, *IEEE T. Inform. Theory*, 52, 5406–5425, 2006.
- Candes, E. and Wakin, M.: An introduction to compressive sampling, *IEEE Signal Proc. Mag.*, 25, 21–30, 2008.
- Candes, E., Romberg, J., and Tao, T.: Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, *IEEE T. Inform. Theory*, 52, 489–509, 2006.
- Candés, E. and Romberg, J.: Sparsity and incoherence in compressive sampling, *Inverse Probl.*, 23, 969, doi:10.1088/0266-5611/23/3/008, 2007.
- Ciais, P., Rayner, P., Chevallier, F., Bousquet, P., Logan, M., Peylin, P., and Ramonet, M.: Atmospheric inversions for estimating CO₂ fluxes: methods and perspectives, *Climate Change*, 103, 69–92, 2010.
- Cinzano, P., Falchi, F., Elvidge, C. D., and Baugh, K. E.: The artificial night sky brightness mapped from DMSP satellite Operational Linescan System measurements, *Mon. Not. R. Astron. Soc.*, 318, 641–657, 2000.
- Coifman, R., Geshwind, F., and Meyer, Y.: Noiselets, *Appl. Comput. Harmon. A.*, 10, 27–44, doi:10.1006/acha.2000.0313, 2001.
- Donoho, D.: Compressed sensing, *IEEE T. Inform. Theory*, 52, 1289–1306, 2006.
- Donoho, D. L., Tsaig, Y., Drori, I., and Starck, J.-L.: Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit, *IEEE T. Inform. Theory*, 58, 1094–1121, 2012.
- Duarte, M. F., Wakin, M. B., and Baraniuk, R. G.: Fast reconstruction of piecewise-smooth signals from random projections, in: *Proceedings of Signal Processing with Adaptive Sparse Structured Representations*, Rennes, France, 2005.
- Gan, L., Do, T., and Tran, T.: Fast compressive imaging using scrambled block Hadamard ensemble, 16th European Signal Processing Conference, 2008.
- Gerbig, C., Dolman, A. J., and Heimann, M.: On observational and modelling strategies targeted at regional carbon exchange over continents, *Biogeosciences*, 6, 1949–1959, doi:10.5194/bg-6-1949-2009, 2009.
- Gholami, A. and Siahkoobi, H. R.: Regularization of linear and non-linear geophysical ill-posed problems with joint sparsity priors, *Geophys. J. Int.*, 180, 871–882, 2010.
- Gourdji, S. M., Mueller, K. L., Yadav, V., Huntzinger, D. N., Andrews, A. E., Trudeau, M., Petron, G., Nehr Korn, T., Eluszkiewicz, J., Henderson, J., Wen, D., Lin, J., Fischer, M., Sweeney, C., and Michalak, A. M.: North American CO₂ exchange: inter-comparison of modeled estimates with results from a fine-scale atmospheric inversion, *Biogeosciences*, 9, 457–475, doi:10.5194/bg-9-457-2012, 2012.
- Gurney, K. R., Mendoza, D. L., Zhou, Y., Fischer, M. L., Miller, C. C., Geethakumar, S., and de la Rue de Can, S.: High resolution fossil fuel combustion CO₂ emission fluxes in the United States, *Environ. Sci. Technol.*, 43, 5535–5541, 2009.
- He, L. and Carin, L.: Exploiting structure in wavelet-based bayesian compressive sensing, *IEEE T. Signal Proces.*, 57, 3488–3497, 2009.
- Hirst, B., Jonathan, P., del Cueto, F. G., Randell, D., and Kosut, O.: Locating and quantifying gas emission sources using remotely obtained concentration data, *Atmos. Environ.*, 74, 141–158, 2013.
- Jafarpour, B.: Sparsity-promoting solution of subsurface flow model calibration inverse problems, in: *Advances in Hydrogeology*, edited by: Mishra, P. K. and Kuhlman, K. L., Springer, 2013.

- Ji, S., Xue, Y., and Carin, L.: Bayesian compressive sensing, *IEEE T. Signal Proces.*, 56, 2346–2356, doi:10.1109/TSP.2007.914345, 2008.
- Kort, E. A., Frankenberg, C., Miller, C. E., and Oda, T.: Space-based observations of megacity carbon dioxide, *Geophys. Res. Lett.*, 39, 4493, doi:10.1029/2002JD003161, 2012.
- La, C. and Do, M. N.: Signal reconstruction using sparse tree representation, in: *Proc. Wavelets XI at SPIE Optics and Photonics*, San Diego, CA, USA, 2005.
- Li, L. and Jafarpour, B.: A sparse Bayesian framework for conditioning uncertain geologic models to nonlinear flow measurements, *Adv. Water Resour.*, 33, 1024–1042, 2010.
- Lin, J. C., Gerbig, C., Wofsy, S. C., Andrews, A. E., Daube, B. C., Davis, K. J., and Grainger, C. A.: A near-field tool for simulating the upstream influence of atmospheric observations: the Stochastic Time-Inverted Lagrangian Transport (STILT) model, *J. Geophys. Res.*, 108, 4493, doi:10.1029/2002JD003161, 2003.
- Loris, I., Nolet, G., Daubechies, I., and Dahlen, F. A.: Tomographic inversion using ℓ_1 -norm regularization of wavelet coefficients, *Geophys. J. Int.*, 170, 359–370, 2007.
- Mallat, S. and Zhang, Z.: Matching pursuit with time-frequency dictionaries, *IEEE T. Signal Proces.*, 41, 3397–3415, 1993.
- Martinez-Camara, M., Dokmanic, I., Ranieri, J., Scheibler, R., M. Vetterli, M., and Stohl, A.: The Fukushima inverse problem, in: *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4330–4334, 2013.
- McKain, K., Wofsy, S., Nehrkorn, T., Eluszkiewicz, J., Ehrlinger, J. R., and Stephens, B. B.: Assessment of ground-based atmospheric observations for verification of greenhouse gas emissions from an urban region, *P. Natl. Acad. Sci. USA*, 109, 8423–8428, 2012.
- Oda, T. and Maksyutov, S.: A very high-resolution (1 km \times 1 km) global fossil fuel CO₂ emission inventory derived using a point source database and satellite observations of nighttime lights, *Atmos. Chem. Phys.*, 11, 543–556, doi:10.5194/acp-11-543-2011, 2011.
- Olivier, J. G. J., Aardenne, J. A. V., Dentener, F. J., Pagliari, V., Ganzeveld, L. N., and Peters, J. A. H. W.: Recent trends in global greenhouse gas emissions: regional trends 1970–2000 and spatial distribution of key sources in 2000, *Journal of Integrative Environmental Science*, 2, 81–99, 2005.
- Pacala, S. W., Breidenich, C., Brewer, P. G., Fung, I. Y., Gunson, M. R., and co authors: *Verifying Greenhouse Gas Emissions: Methods to Support International Climate Agreements*, Committee on Methods for Estimating Greenhouse Gas Emissions, National Research Council, The National Academies Press, available at: http://www.nap.edu/openbook.php?record_id=12883 (last access: 13 August 2014), 2010.
- Potter, C. S., Randerson, J. T., Field, C. B., Matson, P. A., Viousek, P. M., Mooney, H. A., and Klooster, S. A.: Terrestrial ecosystem production: a process model based on global satellite and surface data, *Global Biogeochem. Cy.*, 7, 811–841, 1993.
- Ray, J.: Estimating ffCO₂ using a MsRF and sparse reconstruction, available at: <http://www.sandia.gov/~jairay/software.html> (last access: 13 August 2014), 2013.
- Ray, J., Lee, J., Lefantzi, S., Yadav, V., Michalak, A. M., Bloemen-Waanders, B., and McKenna, S. A.: A multiresolution spatial parametrization for the estimation of fossil-fuel carbon dioxide emissions via atmospheric inversions, SAND Report SAND2013-2919, Sandia National Laboratories, Livermore, CA 94551-0969, available at: <http://www.sandia.gov/~jairay/Presentations/sand2013-2919.pdf> (last access: 13 August 2014), unclassified and unlimited release, 2013.
- Ray, J., Yadav, V., Michalak, A. M., van Bloemen Waanders, B., and McKenna, S. A.: A multiresolution spatial parameterization for the estimation of fossil-fuel carbon dioxide emissions via atmospheric inversions, *Geosci. Model Dev.*, 7, 1901–1918, doi:10.5194/gmd-7-1901-2014, 2014.
- Rayner, P. J., Raupach, M. R., Paget, M., Peylin, P., and Koffi, E.: A new global gridded data set of CO₂ emissions from fossil fuel combustion: methodology and evaluation, *J. Geophys. Res.*, 115, D19306, doi:10.1029/2009JD013439, 2010.
- Romberg, J.: Imaging via compressive sampling, *IEEE Signal Proc. Mag.*, 25, 14–20, 2008.
- Simons, F. J., Loris, I., Nolet, G., Daubechies, I. C., Voronin, S., Judd, J. S., Vetter, P. A., Vetter, P. A., Charlety, J., and Vonesch, C.: Solving or resolving global tomographic models with spherical wavelets and the scale and sparsity of seismic heterogeneity, *Geophys. J. Int.*, 187, 969–988, 2011.
- Skamarock, W. C. and Klemp, J. B.: A time-split nonhydrostatic atmospheric model for weather research and forecasting applications, *J. Comput. Phys.*, 227, 3465–3485, 2008.
- Tans, P. and Conway, T. J.: Monthly atmospheric CO₂ mixing ratios from the NOAA CMDL Carbon Cycle Cooperative Global Air Sampling Network, 1968–2002, in: *Trends: A Compendium of Data on Global Change*, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, Oak Ridge, TN, 2005.
- Tropp, J. and Gilbert, A. C.: Signal recovery from partial information via orthogonal matching pursuit, *IEEE T. Inform. Theory*, 53, 4655–4666, 2007.
- Tsaig, Y. and Donoho, D.: Extensions of compressed sensing, *Signal Process.*, 86, 533–548, 2006.
- Tuma, T. and Hurley, P.: On the incoherence of noiselet and Haar bases, in: *International Conference on Sampling Theory and Applications*, 2009.
- Turnbull, J. C., Karion, A., Fischer, M. L., Faloona, I., Guilderson, T., Lehman, S. J., Miller, B. R., Miller, J. B., Montzka, S., Sherwood, T., Saripalli, S., Sweeney, C., and Tans, P. P.: Assessment of fossil fuel carbon dioxide and other anthropogenic trace gas emissions from airborne measurements over Sacramento, California in spring 2009, *Atmos. Chem. Phys.*, 11, 705–721, doi:10.5194/acp-11-705-2011, 2011.
- Yin, W., Morgan, S., Yang, J., and Zhang, Y.: Practical compressive sensing with Toeplitz and circulant matrices, in: *Proc. SPIE*, 7744, 77440K-77440K-10, doi:10.1117/12.863527, 2010.