SANDIA REPORT

SAND2013-7817

Unlimited Release

Printed September 2013

Kalman-Filtered Compressive Sensing for High Resolution Estimation of Anthropogenic Greenhouse Gas Emissions from Sparse Measurements

J. Ray, J. Lee, S. Lefantzi, V. Yadav, A. M. Michalak, B. van Bloemen Waanders and S. A. McKenna

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>

Kalman-Filtered Compressive Sensing for High Resolution Estimation of Anthropogenic Greenhouse Gas Emissions from Sparse Measurements

J. Ray, J. Lee and S. Lefantzi

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

V. Yadav and A. M. Michalak,

Carnegie Institution for Science, Stanford, CA 94305

B. van Bloemen Waanders,

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

S. A. McKenna,

IBM Research, Smarter Cities Technology Centre, Bldg 3

Damastown Industrial Estate, Mulhuddart, Dublin 15, Ireland.

{jairay,jlee3,slefant,bartv@sandia.gov}, {vineety,michalak@stanford.edu},
{seanmcke@ie.ibm.com}

Abstract

The estimation of fossil-fuel CO₂ emissions (ffCO₂) from limited ground-based and satellite measurements of CO₂ concentrations will form a key component of the monitoring of treaties aimed at the abatement of greenhouse gas emissions. The limited nature of the measured data leads to a severely-underdetermined estimation problem. If the estimation is performed at fine spatial resolutions, it can also be computationally expensive. In order to enable such estimations, advances are needed in the spatial representation of ffCO₂ emissions, scalable inversion algorithms and the identification of observables to measure.

To that end, we investigate parsimonious spatial parameterizations of ffCO₂ emissions which can be used in atmospheric inversions. We devise and test three random field models, based on wavelets, Gaussian kernels and covariance structures derived from easily-observed proxies of human activity. In doing so, we constructed a novel inversion algorithm, based on compressive sensing and sparse reconstruction, to perform the estimation. We also address scalable ensemble Kalman filters as an inversion mechanism and quantify the impact of Gaussian assumptions inherent in them. We find that the assumption does not impact the estimates of mean ffCO₂ source strengths appreciably, but a comparison with Markov chain Monte Carlo

estimates show significant differences in the variance of the source strengths. Finally, we study if the very different spatial natures of biogenic and ffCO₂ emissions can be used to estimate them, in a disaggregated fashion, solely from CO₂ concentration measurements, without extra information from products of incomplete combustion e.g., CO. We find that this is possible during the winter months, though the errors can be as large as 50%.

Acknowledgment

Thanks are due to Sharon Gourdji of Stanford University, and Kim Mueller and Abhishek Chatterjee of University of Michigan, Ann Arbor for fruitful discussions regarding algorithms and the challenges of working with field observations as well as some of the data used in this work.

This work was funded under LDRD (Laboratory Directed Research and Development) Project Number 151293 and Title “Kalman-Filtered Compressive Sensing for High Resolution Estimation of Anthropogenic Greenhouse Gas Emissions from Sparse Measurements at Global Scale”. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000.

This page intentionally left blank

Contents

1	Introduction	17
2	A sparse reconstruction technique to estimate ffCO₂ emissions	19
2.1	Introduction	19
2.2	Background	21
2.2.1	Estimation of CO ₂ fluxes	21
2.2.2	Wavelet modeling	23
2.2.3	Sparse reconstruction	24
2.3	Constructing a multiscale random field model	26
2.4	Formulation of the estimation problem	29
2.4.1	Transport model	29
2.4.2	The inverse problem	30
2.4.3	Solving the inverse problem	32
2.4.4	Enforcing non-negativity of $\mathbf{F}_{\mathcal{R}}$	33
2.5	Tests	35
2.5.1	Comparison of optimization formulations	35
2.5.2	Evaluating formulation using compressive sensing metrics	37
2.5.3	Inversions with nightlights as proxy	39
2.5.4	Inversions with built-up area maps as proxy	42
2.5.5	Impact of M_{cs}	44
2.5.6	Impact of ϵ	47
2.5.7	Impact of the number of observation towers	48
2.6	Conclusions	48

3	Estimation of ffCO₂ emissions using ensemble Kalman filters	53
3.1	Introduction	53
3.2	Mathematical Formulation	55
3.2.1	Bayesian Inference Review	55
3.2.2	Comparison Formulation	56
3.2.3	Sequential Processing	57
3.2.4	Kalman Filter	58
3.2.5	Ensemble Kalman Filter	60
3.3	Overview of the Physics and Numerical Setup	61
3.3.1	Problem setup	61
3.3.2	Inverse modeling and dimensionality reduction	62
3.3.3	Enforcing non-negativity in EnKF	63
3.4	Numerical Results	64
3.4.1	Inversions with EnKF	64
3.4.2	EnKF vs MCMC	64
3.5	Conclusions	67
4	Joint estimation of biogenic and fossil-fuel CO₂ fluxes	71
4.1	Introduction	71
4.2	Objectives and case studies	72
4.3	Numerical details	73
4.3.1	Flux domain and resolution	73
4.3.2	Methods	73
4.3.3	Biospheric and ffCO ₂ covariance structure	74
4.3.4	Model of the trend	75
4.3.5	Covariance, covariate selection and parameter optimization	75
4.3.6	Non-negativity constraints on ffCO ₂ fluxes	76
4.4	Results and discussion	76

4.4.1	Model selection	76
4.4.2	Disaggregated ffCO ₂ and biospheric fluxes	77
4.5	Conclusions	81
5	Conclusions	83
	References	85
 Appendix		
A	Glossary	93
B	Measurement locations	95
C	Second-best covariates	97

This page intentionally left blank

List of Figures

2.1	Differences in the nature of spatial distribution of biospheric (left) and fossil-fuel (right) CO ₂ fluxes. The biospheric fluxes are stationary, whereas ffCO ₂ emissions are multiresolution (non-stationary) and correlated with human habitation. The biospheric fluxes are for June 1 - June 8, 2004, obtained from CASA-GFED [1]. The post-processing steps to obtain the fluxes as plotted are described in [2]. The units of fluxes/emissions are $\mu\text{moles s}^{-1} \text{m}^{-2}$ of C. The ffCO ₂ emissions are obtained from the Vulcan database [3, 4]. Note the different colormaps; ffCO ₂ emissions can assume only non-negative values.	20
2.2	Left: The fraction of wavelet coefficients which are non-zero on each scale l , when subjected to a wavelet transform using Haars, Daubechies 4 and 6 and Symlet 4 and 6 wavelets. We see that Haar wavelets provide the sparsest representation. Right: We plot the average value of the non-zero coefficients (solid lines) and their standard deviation (dashed line). We find that while Haars may provide the sparsest representation, the non-zero values tend to be large and distinct.	26
2.3	Top row: Maps of nightlight radiances (left) and built-up area percentage (right), for the US. Middle row: The sparsity of representation, the correlation between $\mathbf{X}$ and $\mathbf{f}_V$ and the normalized error ϵ_f between the Vulcan emissions $\mathbf{f}_V$ and the sparsified form obtained by projecting it on $\mathbf{X}$. These values are plotted for nightlights (left) and the built-up area maps (right). Bottom row: Plots of $(\mathbf{f}_{pr} - \mathbf{f}_V)$ obtained from nightlights (left) and built-up area maps (right).	28
2.4	Flowchart of the procedure for performing the inversion of ffCO ₂ emissions. The top half of the figure was addressed in Sec. 2.3. We start with a collection of wavelet families that <i>could</i> be used for modeling the emissions, and find that Haars provide the sparsest representation. Thereafter, we use the Haars to create a “prior” model $f_{pr} = c\tilde{\Phi}'_{\mathcal{R}}w$ for the emissions. In the lower half of the figure, we illustrate the posing and solution of the inverse problem. Step A, the posing of the inverse problem is described in Sec. 2.4.2. Three slightly different formulations (Step B), based on how the “prior” is incorporated into the inverse problem, are also described in the same section. Step C, the StOMP algorithm, is described in Sec. 2.4.3. Sec. 2.4.4, the description of non-negativity on the ffCO ₂ emissions, corresponds to Step D. .	30
2.5	Plots of ffCO ₂ emissions during the 31 st 8-day period. 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 $\mathbf{f}_{pr}$ provides low levels of (erroneous) emissions in large swathes of the Western quadrants. Approach B reflects $\mathbf{f}_{pr}$ very strongly. Approach C provides a balance between the influence of $\mathbf{f}_{pr}$ and the information in $\mathbf{y}^{obs}$	36

2.6	Comparison of estimation error (left) and the correlation between true and estimated emissions (right) using Approaches A, B and C. It is clear that Approach B is inferior to the others.	37
2.7	Reconstruction error (left) 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 $\mathbf{f}_{pr}$, leads to lower errors in both the quadrants and better correlations in the less instrumented NW quadrant.	38
2.8	Comparison of the distribution of the elements of $\mathbf{A}_{\Psi}$ and $\mathbf{A}_{\Phi}$. 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 $\mathbf{A}_{\Psi}$ where the elements assume discrete values. In contrast, the elements of $\mathbf{A}_{\mathbf{H}}$ assume values which are spread over a far larger range, some of which are quite close to 1 while others are very close to zero.	39
2.9	Reconstruction of the ffCO ₂ emissions from the 35 towers (plotted as diamonds) over one year. The true emissions are on the left and the reconstructions on the right. We see that the large scale structure of the emissions have been captured, as seen by the comparisons for 8-day periods number 9 (mid-March), 33 (end of August) and 42 (early November). The west coast of the US has few towers near heavily populated regions and thus is not very well estimated. In the middle and bottom rows, we see that the emissions in the Los Angeles-San Diego region (32.87N, -117.26W) is underestimated. On the other hand, due to the higher density of towers in the Northeast, the true and estimated emissions are qualitatively similar. Emissions have units of $\mu\text{mol m}^{-2} \text{s}^{-1}$ of C (not CO ₂).	40
2.10	Comparison of reconstruction error and correlations. Left: We plot the error between the reconstructed and true (Vulcan) emissions in blue. In red, we plot the error between $\mathbf{f}_{pr}$ using nightlights and Vulcan emissions. We see that assimilation of $\mathbf{y}^{obs}$ leads to significantly improved accuracy. Right: We plot the accuracy of the spatial distribution of the reconstructed emissions. The Pearson correlations $C(\mathbf{E}_k, \mathbf{f}_{V,k})$ and $C(\mathbf{f}_{pr}, \mathbf{f}_{V,k})$ show that incorporating $\mathbf{y}^{obs}$ improves the spatial agreement of estimated emissions versus the true one from 0.7 to around 0.9. If the emissions are averaged over 32-day periods, rather than 8-day periods, the correlation with true (Vulcan) emissions rises even higher.	41
2.11	Plot of the error in, and the correlation between reconstructed emissions and true (Vulcan) ones. Left: We see that the reconstruction error in the NE quadrant is small compared to the others. However, the error in each of the quadrants is much larger than the error at the country level (see Fig. 2.10). Right: Correlation in each of the quadrants. The NE quadrant is substantially better than the others due to the higher density of towers.	42

2.12	Left: Prediction of ffCO ₂ concentrations at 3 measurement locations, using the true (Vulcan) and reconstructed emissions (blue lines) over an 8-day period (Period no. 31). Observations occur every 3 hours. We see that the concentrations are accurately reproduced by the estimated emissions. Right: 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.	43
2.13	CDF of emissions in $\mathcal{R}$, before and after the imposition of non-negativity, as described in Sec. 2.4.4. We see that the CDF of the emissions without non-negativity imposed contains a few grid-cells with negative fluxes; further, the magnitude of the negative emissions is small. Thus the spatial parameterization, with sparse reconstruction provides a good approximation of the final, non-negative emissions.	43
2.14	Left: Plot of the error in the reconstruction, as performed with built-up area maps (blue line) and nightlights (black line) as the proxies. Right: Correlation between the true and reconstructed emissions, as performed with built-up area maps and nightlights as the proxies. The nightlights-based reconstruction has slightly less reconstruction error and better correlation with the true emissions, when compared at 8-day temporal resolution. When averaged to a 32-day temporal resolution, emissions obtained via the two methods are very similar.	45
2.15	Left: Emission reconstruction error in the NE (blue) and NW (black) quadrants, when performed with BUA (line) and nightlights (symbols) as proxies. We see that the NW quadrant is very badly constrained and a change in proxies changes the sign of the error. This is also seen in the NE quadrant; however, the error magnitudes are far smaller. Right: The comparison of correlations between true and reconstructed emissions shows similar trends; nightlights-based estimation produces better reconstructions.	45
2.16	Comparison of emission estimates developed using $\mathbf{f}_{pr}$ constructed from nightlight radiances and built-up area maps. Top: Estimated ffCO ₂ emissions for the 34 th 8-day period developed using nightlight prior (left) and the prior from built-up areas maps (right). Bottom left: we plot the difference between the two estimates. Bottom right: We plot a scatter plot between the estimated and prior emissions, for the two prior models.	46
2.17	The impact of the number of compressive samples M_{cs} on the reconstruction of $\mathbf{F}_{\mathcal{R}}$ ($\eta_{\mathcal{R}}$) and $\mathbf{F}_{\mathcal{R}'}$ ($\eta_{\mathcal{R}'}$). $\eta_{\mathcal{R}}$ and $\eta_{\mathcal{R}'}$ are plotted on the Y1 and Y2 axes respectively. Results are plotted for the 31 st 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 $\eta_{\mathcal{R}'}$	47
2.18	Impact of measurement error ϵ . Top left: We plot the true ffCO ₂ emissions from the Vulcan inventory. Top right: We plot the estimates calculated using $\epsilon = 2.5 \times 10^{-2}$. Bottom left, we plot the reconstructed emissions using $\epsilon = 10^{-1}$; we see a clear degradation of the reconstruction. Bottom right: We plot the reconstruction error (%) and the correlation between the reconstructed and true emissions for various values of ϵ ; a clear degradation is seen. Reconstruction errors and correlations are plotted on opposing Y-axes. All results are for the 31 st 8-day period.	49

2.19	Left column: Estimated ffCO ₂ emissions for the 33 rd 8-day period. Inversions were done with 35 (top), 25 (middle) and 15 (bottom) towers. As the number of towers decreases, we see that the resolution of estimated emissions decreases. Right column: At the top, we plot the error in the reconstructed emissions for the three tower sets. There is not much difference between reconstructions with 25 and 35 towers, but the 15-tower reconstruction is poor. Right (middle): The correlation between reconstructed and true emissions shows much the same trend as the reconstruction, with the 15-tower reconstruction having a far lower correlation over the entire 360-day duration. Right column, bottom: We plot the sparsity of the estimated wavelet coefficients. We see that as the number of towers decrease, the wavelet coefficient estimates deviate further from the true values (red crosses).	50
3.1	Comparison of MCMC and EnKF approach to estimating the posterior density in Eq. 3.10. Note the potential for parallelism in the EnKF.	58
3.2	Plots of inferred (left) and true (right) ffCO ₂ concentrations (top row) and ffCO ₂ sources (bottom row) using no observations. The true velocity field was used. We see that the prior is not very informative.	65
3.3	Plots of inferred (left) and true (right) ffCO ₂ concentrations (top row) and ffCO ₂ sources (bottom row) using ffCO ₂ concentration observations collected on a 15 × 15 sensor grid. The true velocity field was used. We see an immense improvement in the inferred ffCO ₂ concentration field, and a more informative source inversion.	66
3.4	Plots of inferred and true concentration fields (top row) and ffCO ₂ sources (bottom row) when assimilating velocity and concentration measurements.	67
3.5	Final concentration fields, reconstructed using limited CO ₂ measurements. The true velocity field is used. The MCMC reconstruction is on the left and EnKF on the right. The true solution is in the middle.	68
3.6	Reconstruction of sources using the truth velocity. The MCMC reconstruction is on the left and the EnKF on the right. The true solution is in the middle.	68
3.7	Plot of the mean of posterior distribution of the amplitudes of 10 GK, inferred using the true velocity. There is not much difference between the means calculated with MCMC versus EnKF.	69
3.8	Plot of the variance of posterior distribution of GK amplitudes. There is significant difference in the results of MCMC versus EnKF.	69
3.9	Covariance matrix of posterior distribution of GK amplitudes, calculated using MCMC (left) and EnKF (middle). On the right is the difference between the two.	70
4.1	Estimated January biospheric (top) and ffCO ₂ fluxes (bottom) for June for the real-data study. Units: moles m ² s ⁻¹	79
4.2	Estimated June biospheric (top) and ffCO ₂ fluxes (bottom) for June for the real-data study. Units: moles m ² s ⁻¹	80

List of Tables

4.1	Best covariates and covariance model in the real-data case study. Selections for January–December 2008 are shown. Note: An asterisk indicates a covariate.	77
4.2	Best covariates and covariance models for the pseudo-data case study. Selections for January–December 2008 are shown. Note: An asterisk indicates a covariate.	78
4.3	Root mean squared error (RMSE) and correlation coefficient of the estimated biospheric and ffCO ₂ fluxes in the pseudo-data case for January and June 2008. CDIAC and CASA-GFED are considered to provide the true ffCO ₂ and biospheric CO ₂ fluxes.	78
A.1	Definitions of Greek symbols	93
A.2	Definitions of symbols	94
B.1	List of observation towers and their locations.	96
C.1	Second-best covariates and covariance model in the real-data case from January to December 2008. Note: The asterisk denotes a covariate.	97

Chapter 1

Introduction

CO₂, produced from the combustion of fossil fuels (ffCO₂) form the largest component of *net* land-atmosphere exchange in the carbon cycle. The accumulation of ffCO₂ in the atmosphere is also held responsible for climate change. The US National Academy of Science has issued a report [5] that proposes monitoring of ffCO₂ emissions as part of a CO₂ abatement treaty.

Currently, ffCO₂ emission estimates are obtained from inventories which provide ffCO₂ fluxes at fine spatiotemporal scales e.g., Vulcan [4], EDGAR [6], ODIAC [6] etc. These inventories begin national or provincial estimates of fossil-fuel production and consumption and convert them to ffCO₂ emissions based on estimates of combustion efficiency. The downscaling to fine spatial and temporal scales is performed using proxies of human activity e.g., images of light at night [7], population density [8] etc. However, the bottom-up estimates, when used in atmospheric transport do not reproduce observations of CO₂ concentrations (after accounting for biogenic fluxes). This could be due to errors in the atmospheric transport models or errors in the inventory itself. Being able to infer ffCO₂ emissions from CO₂ concentrations would provide independent verification of the inventory. Such an inversion capability (called Greenhouse Gas Information System in [5]) could also be used to monitor a CO₂ abatement treaty.

The inference of spatially-resolved CO₂ fluxes from CO₂ concentrations (called atmospheric inversions) are routine for biogenic fluxes and methods to do so are reviewed in [9]. However, the estimation of ffCO₂ emissions is rare [10, 11] and limited to estimates for individual cities. The estimation of spatially-resolved ffCO₂ emissions, at the regional scale (e.g., with 100 km resolution) encounters the following challenges:

1. *Random field models:* The use of limited CO₂ concentration measurements in inversions require a low dimensional model of ffCO₂ emissions i.e., a model with few parameters that can nevertheless reproduce the strongly non-stationary nature of ffCO₂ emissions as revealed by inventories. Such random field (RF) models do not exist.
2. *Scalable inversion techniques:* ffCO₂ emissions are best estimated at global scale. This allows one to use all measurement modalities. It also does not require one to impose any boundary conditions. Regional estimates of biogenic CO₂ fluxes surround the region of interest with a “control volume” on which boundary velocities and CO₂ fluxes, as obtained from a meteorological model, are imposed as boundary conditions. This can incur a significant error and contributes to uncertainty in CO₂ flux estimates. However, a global inversion is computationally expensive, and Bayesian methods, which are used to estimate biogenic CO₂ fluxes (and their uncertainties) are reaching their scalability limits. Thus a need exists to identify an inversion mechanism that is scalable to large-scale global inversions.
3. *Measurement variables:* ffCO₂ signatures may be measured via trace species of incomplete combustion e.g., CO, SO₂ etc or via radiocarbon ¹⁴C. The introduction of reactive species, e.g., CO, in the inversion increases the computational cost and limits the inversion to low spatial resolutions [12].

In contrast, the inclusion of radiocarbon into inversions poses no technical challenges, except that radiocarbon measurements are very expensive. Thus a requirement exists to establish what needs to be measured to infer ffCO₂ emissions.

4. *Accurate atmospheric transport models:* These models are perhaps the largest sources of error and uncertainty in CO₂ flux inversions.

In this investigation, we will address aspects related to the first three challenges since they impact the inverse problem that underlies the ffCO₂ estimation problem. The fourth challenge, improving transport models, is a field in itself.

In Chapter 2 we address the construction of a RF model for ffCO₂ emissions. We hypothesize that wavelets are well suited to represent the strongly non-stationary nature of ffCO₂ emissions. We identify the type of wavelet that provides the most parsimonious model of ffCO₂ emissions, and devise a sparse reconstruction scheme that performs data-driven simplification of the model, and its fitting of CO₂ observations simultaneously.

In Chapter 3 we explore ensemble Kalman filters (EnKF) as a means of estimating ffCO₂ emission fields. EnKF is a scalable inference capability, but assumes a Gaussian distribution for the variables being inferred. In this chapter, we explore the ramifications of that assumption. We also develop a spatial parameterization for ffCO₂ emissions, based on Gaussian kernels, which provides a very parsimonious representation when the sources are localized e.g., in the Great Plains.

In Chapter 4 we address the question of whether ffCO₂ and biogenic CO₂ fluxes can be estimated jointly from CO₂ observations without tracer/radiocarbon information. The spatial distributions of biogenic CO₂ and ffCO₂ fluxes are very different and we investigate the hypothesis that the two types of fluxes may be inferred, in a disaggregated fashion, based solely on their spatial natures. In all the chapters mentioned above, we will use easily-observed proxies of ffCO₂ emissions to construct parsimonious RF models.

In Chapter 5 we summarize our findings.

Chapter 2

A sparse reconstruction technique to estimate ffCO₂ emissions

2.1 Introduction

The potential role of anthropogenic greenhouse gas emission in climate change has led to a host of policies that seek to reduce them e.g., the UN-REDD program [13] and the UNFCCC [14]. Emissions of CO₂ from the burning of fossil fuels (for power generation, transportation, industrial and residential use) form the largest component of these greenhouse gas emissions and any policies aimed at the abatement of fossil-fuel (ffCO₂) emissions are expected to have significant socioeconomic impacts. Monitoring the ability of such policies to reduce greenhouse gas emissions is a complex endeavor and is expected to receive increasing attention in the future. To this end, The United States National Academy of Sciences considered the design of a Greenhouse Gas Information System that would combine measurements of CO₂ concentrations and fluxes, with reports of national inventories of fossil fuels, to estimate ffCO₂ emissions [5].

In this report, we present a method for estimating ffCO₂ emissions from CO₂ concentration measurements at observation towers. We develop a spatial parameterization for the emissions and an inversion scheme, based on sparse reconstruction, to calibrate the parametrization. The method is demonstrated on ffCO₂ in the United States using synthetic data. Its extension to global scale inversion is conceptually simple. It will, naturally, require more data. It will also be significantly more computationally demanding, thus necessitating the development of scalable counterparts of the sparse reconstruction methods described in this report.

Currently, the best data on ffCO₂ emissions are obtained from inventories/databases that are constructed from national reporting of the consumption of fossil fuels. Some inventories [15, 8, 16] start with national or provincial estimates of fossil-fuel consumption and disaggregate them using proxies of human habitation (e.g., population density), whereas others aggregate them from local information of fossil-fuel use patterns [4]. Their accuracy depends on the data used for constructing the inventories, as well as the method; see [17, 18] for a discussion on their uncertainties and revisions when more data becomes available. In some cases, these revisions can be large [19].

An alternate way of estimating CO₂ emissions is via inversion of atmospheric measurements. In this process, time-varying CO₂ fluxes $\mathbf{f}$ are estimated from measurements of time-varying CO₂ concentrations at certain ground-based locations (towers) or satellite soundings of column-integrated CO₂ concentrations ($\mathbf{y}^{obs}$). The CO₂ fluxes $\mathbf{f}$ are defined on a gridded domain and related to $\mathbf{y}^{obs}$ via an atmospheric transport model. Since the magnitude of $\mathbf{f}$ does not impact the wind velocity field and dispersion, CO₂ is modeled as a passive scalar

Figure 2.1: Differences in the nature of spatial distribution of biospheric (left) and fossil-fuel (right) CO₂ fluxes. The biospheric fluxes are stationary, whereas ffCO₂ emissions are multiresolution (non-stationary) and correlated with human habitation. The biospheric fluxes are for June 1 - June 8, 2004, obtained from CASA-GFED [1]. The post-processing steps to obtain the fluxes as plotted are described in [2]. The units of fluxes/emissions are $\mu\text{moles s}^{-1} \text{m}^{-2}$ of C. The ffCO₂ emissions are obtained from the Vulcan database [3, 4]. Note the different colormaps; ffCO₂ emissions can assume only non-negative values.

and a linear relationship holds

$$\mathbf{y}^{obs} = \mathbf{y} + \boldsymbol{\varepsilon} = \mathbf{H}\mathbf{f} + \boldsymbol{\varepsilon}, \quad (2.1)$$

where $\mathbf{H}$ is the transport or sensitivity matrix, obtained from a transport model like Weather Research and Forecasting Model (WRF, [20]). $\mathbf{y}$ is the CO₂ concentration predicted by the atmospheric model which differs from its measured counterpart by an error $\boldsymbol{\varepsilon}$. Atmospheric inversion has become a routine tool for estimating *biospheric* CO₂ fluxes, conditioned on both satellite and ground-based observations [21, 22, 23, 24, 25]. The inverse problem is set up as an optimization between the need to reproduce observations while adhering to a prior belief regarding the spatio-temporal distribution of fluxes. Since biospheric CO₂ fluxes vary smoothly in space (see Fig. 2.1), the prior modeling of CO₂ fluxes is performed with a stationary, multivariate Gaussian field. The prior covariance provides the regularization in the inverse problem and allows high-dimensional gridded emissions to be estimated from relatively few measurements, at the cost of smoothing out fine-scale spatial structures in the estimated fluxes.

There are two methods for calculating these emission estimates : (1) deterministic methods, based on a variational formulation (which provide a “mean” or a “best-fit” estimation) and (2) Bayesian methods that provide a measure of the uncertainty in the estimate in addition to the “best-fit”. These methods have been reviewed in [9]. Kalman filters, too, have been used in these inversions [26]. All these methods are strongly influenced by the choice of the prior model, and the *geostatistical inversion* method [27] was developed to lessen the dependence.

Fig. 2.1 contrasts the spatial distribution of biospheric CO₂ fluxes against $\overline{\mathbf{f}_V}$, the ffCO₂ emissions from the Vulcan database [3, 4]. Vulcan provides emissions for the lower 48 states of the US, on a 0.1° grid and at an hourly resolution. The emissions are coarsened to a 1° resolution and averaged over the entire year, to obtain $\overline{\mathbf{f}_V}$. It is clear that the spatial distribution of ffCO₂ is correlated with human habitation, is strongly multiscale,

and may be difficult to represent with a variogram. Instead, we seek a different spatial parameterization that can represent its strongly multiresolution character. Further, the spatial parametrization must necessarily be low-dimensional (i.e., have few free/independent parameters) so that it can be used for estimation within the context of an inverse problem with sparse observations.

In this paper, we construct a spatial parametrization for ffCO₂ based on wavelets. We will refer to it as the Multiscale Random Field (MsRF) model. Wavelets have compact support, form an orthogonal basis set and are widely used to model non-stationary fields e.g. images [28, 29]. We will reduce the dimensionality of the MsRF using an easily-observed proxy of human habitation e.g., images of lights at night (henceforth called nightlights), maps of built-up areas etc. The MsRF so formed will be used in a synthetic atmospheric inversion test using a novel, sparsity-enforcing optimization method. “Synthetic” here refers to the fact that CO₂ concentration observations will be generated using a model (*not* the one used in the estimation procedure) rather obtained from field measurements. The inversion assumes a pure ffCO₂ signal i.e., we can measure the fossil-fuel contribution to the CO₂ concentration in a sample, as done in [10, 11]. This procedure will identify the subset of wavelets in the MsRF that can be actually estimated from the synthetic observations, while “turning off” the rest. In doing so, it will ensure that the MsRF, as designed, has sufficient flexibility to extract the information on ffCO₂ in the observations. We conclude with a discussion on the efficiency of the inversion/ffCO₂ estimation. The Vulcan inventory will serve as the ‘ground-truth’ in our synthetic data inversions.

The paper is structured as follows. In Sec. 2.2, we review existing literature on the construction of ffCO₂ inventories, with emphasis on how proxies of human habitation (population density, nightlights etc) are used to disaggregate national and provincial-level fossil-fuel emissions to finer resolution. We will also review existing literature on atmospheric inversions (both biospheric and fossil-fuel CO₂ fluxes) and *compressive sensing*, a wavelet-based image-processing technique that we adapt to our inversion problem. In Sec. 2.3, we construct two MsRF models, based on nightlights and maps of built-up areas. Sec. 2.4 contains the formulation of the inverse problem and the algorithm to solve it. In Sec. 2.5, we perform tests with synthetic data, discuss the quality of the competing MsRFs and examine the impact of various features of the inverse problem (e.g. number of measurements, errors in measurements etc.). Conclusions are in Sec. 2.6.

2.2 Background

In this section, we review how ffCO₂ emission inventories are currently calculated and the shortcomings in them. This is followed by a discussion of CO₂ flux (both fossil-fuel and biospheric) estimation via atmospheric inversion. We identify the essential difference between fossil-fuel and biospheric CO₂ fluxes and motivate the need for a spatial parameterization for ffCO₂ emissions. We then review existing literature on compressive sensing, which provides the kernel of the spatial parametrization, as well as the optimization techniques which exploit the regularization such random field models can provide.

2.2.1 Estimation of CO₂ fluxes

ffCO₂ emissions, calculated from the consumption of fossil fuels, are reported by most countries and published by the International Energy Agency (<http://www.iea.org>) as well the United Nations Statistics Division [30]. They are usually published after aggregation to the national (and sometimes provincial/state) levels. ffCO₂ emissions can be thought of as a combination of intense point-sources (e.g., electricity gen-

eration and cement production) and diffuse spatially distributed emissions associated with transportation, residential and commercial activities. Emissions from individual, geotagged power plants can be obtained from CARMA [31] whereas [32] provides cement production data; both can be subtracted from national inventories. The remainder, constituting the diffuse sources, can be disaggregated onto a grid based on a number of easily-observed proxies of human activity. Emission inventories are generally employed in their gridded form.

Images of lights at night [33] have been correlated to many socioeconomic parameters [7] and are a common tool for disaggregating national or regional ffCO₂ emission onto a grid. This is discussed in [15], where gridded inventory with a kilometer resolution was constructed. In contrast, EDGAR (Emission Database for Global Atmospheric Research, [34]) and CDIAC (Carbon Dioxide Information Analysis Center, [35]) use population density to perform the disaggregation [16, 36]. In [8] both population density and nightlights are jointly assimilated to perform the disaggregation. The Vulcan inventory [3, 4] and its follow-on, Hestia [37, 38], follow a more complex method, consisting of both aggregation from census-tract data, as well as disaggregation to obtain their gridded inventory. They are not directly dependent on nightlights or population density. Nightlights and population density differ in their spatial distribution at the small (10 kilometer) scale and therefore the two disaggregation methods produce different results (see [8, 15] for discussion); further, the inventories are ultimately dependent on the accuracy of national reporting of emissions, which raises its own set of uncertainties [39, 18, 17].

ffCO₂ emissions or source strengths can also be back-calculated from measurements of CO₂ concentrations using a CO₂ transport model. Such an approach requires one to measure the concentration of ffCO₂, rather than CO₂, in a given sample. This can be performed by either measuring $\Delta^{14}\text{CO}_2$ or measuring CO and estimating ffCO₂ from CO/CO₂ ratios observed in incomplete combustion; see [10] for an example of the estimation of emissions from Sacramento, California, using airborne measurements. Alternatively, in conjunction with an atmospheric transport model, one can scale the emissions from an inventory, e.g., Vulcan, to reproduce observations in the vicinity of an urban area to estimate its emissions, as performed for Salt Lake City in [11]. In both cases a single variable was evaluated from the data - the source strength in case of Sacramento and the scaling factor for Salt Lake City. There was no attempt to back-calculate a spatially variable quantity.

Spatially variable *biospheric* CO₂ fluxes are routinely obtained via atmospheric inversions, using both ground-based and satellite measurements. The spatial domain (a region or the entire Earth) is discretized with a grid; each grid-cell hosts a CO₂ flux source. A Bayesian inverse problem is posed using a transport model that linearly relates CO₂ concentration measurements at a set of locations to the strength of each of the sources. The resolution of the grid and the temporal discretization employed for temporally-varying – diurnal and seasonal – biospheric CO₂ fluxes result in more unknowns that can be constrained by the observations and regularization is used to reduce the effective dimensionality of the problem. The “true” fluxes are assumed to be distributed around a “prior”/guess (often obtained from process-based models of biospheric CO₂ fluxes e.g. CASA [40]); the discrepancy between the “true” and prior fluxes is modeled as a multivariate Gaussian field, whose covariance is calculated beforehand/offline. The inverse problem thereafter reduces to finding a CO₂ flux distribution that is a compromise between reproducing CO₂ observations and deviation from the prior fluxes. The multivariate Gaussian field, along with the prior fluxes, provide the regularization. Since the inverse problem is linear, the posterior distribution of the fluxes is also a multivariate Gaussian field. A review of inversion methods can be found in [9]. This approach has been successfully used with satellite retrievals [21] and ground-based measurements [41]. A scalable variational approach has successfully obtained global estimates of CO₂ fluxes at high resolution using satellite measurements [24, 42], and jointly with ground data [25]. Inversion methods based on ensemble Kalman filters

and its derivatives have also been successfully used [23, 43, 26].

Given the paucity of observations, the prior fluxes influence the inferred CO₂ fluxes strongly [21]. This is especially true when inversions are required at a fine spatial resolution. The *geostatistical inverse* method was constructed to lessen its effect [27]. In [44, 45, 22], the prior fluxes were constructed as a linear model which included some easily observed proxies of biospheric and fossil-fuel fluxes (e.g., leaf-area index, population density, per capita GDP etc) with the sensitivities treated as unknowns to be estimated from data. The inversions inferred spatial patterns in CO₂ emissions which reflected both biospheric and anthropogenic contributions.

To summarize, biospheric CO₂ fluxes are routinely estimated from ground and satellite measurements using a variety of methods. Spatial resolutions finer than $1^\circ \times 1^\circ$ and temporal ones that resolve the diurnal cycle, can be achieved. This capability is due, in large part, to the smooth spatial variation of biospheric fluxes, which allows them to be modeled as multivariate Gaussian fields, and serves as a regularization in the inverse problem. In contrast, no such spatial parameterization exists for ffCO₂ emissions, limiting current work to estimating single parameters. In this paper, we will construct a MsRF model for ffCO₂ emissions, as a first step towards enabling their inference via atmospheric inversion. The spatial patterns in nightlights (and other proxies of human activity) will serve as the source of regularization. The approach will be tested in a synthetic data problem, where Vulcan [3, 4], which does not use these proxies, will supply the “ground truth” emissions.

2.2.2 Wavelet modeling

Wavelets are a family of orthogonal bases with compact support i.e., unlike Fourier bases (sines and cosines) which are defined over $[-\infty, \infty]$, individual wavelets are defined only over a bounded region (the “compact support”). Compact support enables wavelets to serve as a flexible parameterization for complicated functions i.e., one can approximately model a complex shape or distribution with a few wavelets. Thus they form an attractive modeling framework for strongly multiscale ffCO₂ emission fields.

Wavelets are generated using a scaling function ϕ' which obeys the recursive relationship

$$\phi'(x) = \sum_i c_i \phi'(2x - i).$$

A wavelet ϕ is generated from the scaling function by taking differences in the following manner:

$$\phi(x) = \sum_i (-1)^i c_{1-i} \phi'(2x - i).$$

Note that the wavelet has compact support. The choice of the filter coefficients c_i and ϕ' determine the type of the resulting wavelets. Wavelets with a large number of non-zero filter coefficients have larger support and are smoother. The simplest are the Haars, which are symmetrical in shape, but not smooth and not differentiable. They have only their first 2 moments equal to zero. Daubechies order 4 and order 6 (Daubechies 4 and 6) are smoother, with up to 4th – (respectively 6th)-order moments that are zero. They have wider support and are asymmetrical in shape. Symlet wavelets are a modified form of Daubechies wavelets, which are more symmetrical. Wavelets can be shrunk and translated to model functions

$$\phi_{s,i} = 2^{\frac{s}{2}} \phi(2^s x - i)$$

where s is the dilation scale and i refers to translation (location). For each increment in scale, the support of the wavelet halves.

Wavelets are defined on dyadic (power-of-two) hierarchical or multi-resolution grids. Consider a domain of size D , discretized by a hierarchy of meshes with resolutions $\Delta D/D = \{1, 1/2, 1/2^2, \dots, 1/2^M\}$. The coarsest mesh, with $\Delta D = D$ is called the scale (or level) $s = 0$, whereas the finest mesh, with $\Delta D = D/2^M$ is the scale M mesh. Wavelets are defined on each of the levels of the hierarchical mesh. Their support spans the same number of grid-points on each scale i.e., wavelets' supports halve every level as we traverse down the hierarchical mesh. Wavelets can be translated or positioned at any grid-cell $i, 0 \leq i \leq 2^s - 1$, on any scale s of the hierarchical mesh. Thus a wavelet $\phi_{s,i}(x)$ requires the specification of s and i to denote its position in the hierarchical grid.

Consider a 1D function $g(x)$ defined on the hierarchical grid. It can be represented as

$$g(x) = w' \phi'(x) + \sum_{s=1}^M \sum_{i=0}^{2^s-1} w_{s,i} \phi_{s,i}(x).$$

The coefficients (or weights) $w_{s,i}$ and w' are obtained by taking projections of $f(x)$,

$$w_{s,i} = \int_{-\infty}^{\infty} g(x) \phi\left(\frac{x-b}{a}\right) dx$$

where $b = iL/2^s$, $a = 2^s$ and $\phi(\xi)$ is the wavelet on the coarsest level. The compact support of $\phi_{s,i}(x)$ and the large number of coefficients $w_{s,i}$ allow one to represent arbitrary $f(x)$ accurately, though not necessarily efficiently (i.e., with few non-zero $w_{s,i}$). Fast methods to compute the projection, called wavelet transforms, exist.

Wavelet transforms of a function are performed using Mallat's pyramid [46]. Wavelet coefficients c_i are determined by calculating weighted averages and differences, as we proceed recursively from the finest scale (grid resolution) to the coarsest (where the grid is reduced to 1 cell). The transform of a function $g(x)$, discretized on the mesh, can be represented by $\mathbf{w} = \mathcal{W}\mathbf{g}$, though $\mathcal{W}$ is never actually formed. $\mathbf{w}$ are the wavelet coefficients or weights. $\mathbf{w}$ contains 2^M wavelet coefficients. Details on wavelet decompositions and transforms are in [47, 28].

An arbitrary 2D field e.g., an image, can, in general, be expressed as a linear superposition of wavelets [47, 28]. A function $f(x, y)$, defined on a $D \times D$ domain and discretized on a hierarchical $2^M \times 2^M$ mesh, can be wavelet transformed by applying 1D wavelet transforms repeatedly, e.g., first by rows and then by columns. In 2D, wavelets of scale s have a support $2^{M-s} \times 2^{M-s}$, $0 \leq s \leq M$. A wavelet of scale s can also be "translated" in space and positioned (in 2D space) at location (i, j) , $0 \leq (i, j) < 2^s$. Compact support and the ability to translate wavelets allow them to model complex, non-stationary fields. A 2D wavelet transform results in $2^M \times 2^M$ wavelet coefficients. If the type of wavelets is chosen judiciously, many $w_{s,i,j}$ may be small and can be approximated as zero. In such a case we obtain an approximate, but *sparse* representation of $f(x, y)$ on the mesh.

2.2.3 Sparse reconstruction

Compressive Sensing (CS) [48, 49] is an efficient means of encoding sparse images. Consider an image $\mathbf{g}$ of size N . Assume, too, that it can be represented sparsely using $L \ll N$ wavelets. CS asserts that a sparse image may be sampled compactly by projecting on to a set of random vectors Ψ_i , to obtain a measurement $\mathbf{g}'$, of size N_m , $L < N_m \ll N$,

$$\mathbf{g}' = \Psi \mathbf{g} = \Psi \Phi \mathbf{w}, \quad (2.2)$$

where the rows of Ψ consist of the random vectors ψ_j , and columns of Φ consist of the orthonormal basis vectors ϕ_i . Φ is a $N \times N$ matrix while Ψ is $N_m \times N$. The bulk of the theory was established in [50, 51, 52]. Sparsity expresses the idea that the “information rate” of a signal may be smaller than its bandwidth and this “information sparsity” may find a reflection in its concise representation in a suitable set of bases Φ (see Eq. 2.2). Ψ and Φ should be “incoherent”. Incoherence relates to the fact that a concise signal in Φ must be measured in a space Ψ where it is spread out, so that a few measurements may suffice to capture its behavior. Further, while sparsity is assumed, the support – the identity of basis vectors with (appreciably) non-zero weights – is not known *a priori* and is inferred along with the values of $\mathbf{w}$. The incoherence between Ψ and Φ is ensured by choosing some well known wavelets bases (Haars, Daubechies 4 and 8) for Φ and random vectors like noiselets [53], random sign ensembles, uniform spherical ensembles, partial Fourier and Hadamard ensembles [54] for Ψ . In [55] the authors establish the degree to which compression might be performed

$$M \geq CL \log(N/L). \quad (2.3)$$

C is generally 3-5.

In CS, the reconstruction of $\mathbf{g}$ (alternatively, $\mathbf{w}$) can be performed using a number of methods. It can be posed as an optimization of the cost function

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \left[\frac{1}{2} \|\mathbf{g}' - \mathbf{A}\mathbf{w}\|_2 + \lambda \|\mathbf{w}\|_1 \right]$$

which is a trade-off between the ℓ_1 norm of $\mathbf{w}$ (to enforce sparsity) and the ℓ_2 norm of the misfit between $\mathbf{g}'$ and the signal reconstructed from $\mathbf{w}$. $\mathbf{A} = \Psi\Phi$. This can be solved using Basis Pursuit (BP) [56] or LASSO (Least Absolute-Shrinkage and Selection Operator) [57]. A slightly different formulation,

$$\begin{aligned} & \underset{\mathbf{w} \in \mathbb{R}^N}{\operatorname{minimize}} && \|\mathbf{w}\|_1 \\ & \text{subject to} && \|\mathbf{g}' - \mathbf{A}\mathbf{w}\|_2 < \epsilon_2, \end{aligned} \quad (2.4)$$

which achieves the same effect, can be addressed using Matching Pursuit (MP) [58], Orthogonal Matching Pursuit (OMP, [59]) and Stagewise OMP [60] (StOMP). StOMP is much faster and offers many of the theoretical guarantees of BP. Bayesian equivalents also exist [61, 62], where Laplace priors are used to enforce sparseness in the inferred $\mathbf{w}$. All these algorithms are general and do not exploit any particular structure in the signal (except sparsity).

One may also create a model of the wavelet distribution and use it along with sparsity-enforcement when constraining $\mathbf{w}$. In a typical wavelet transform, the coefficients of the wavelets at different scales and translations can be stored in a binary tree. The tree can be pruned, *a priori*, based on a learning set of images and thereafter employed with MP or OMP [63, 64, 65]. A Bayesian modification removed the need for a learning set [66]. Comparison with “regular” methods that did not exploit the tree-structure of wavelet coefficients showed a marked improvement in the quality and speed of reconstruction.

To summarize, CS has some striking parallels with the problem of atmospheric inversion. Fig. 2.1 (right) reveals that ffCO₂ emissions are clustered near the coasts and there are vast areas with little emissions. This indicates that a sparse wavelet representation of ffCO₂ emissions may be possible. Further, in Eq. 2.1 the transport matrix $\mathbf{H}$ serves much the same function as the measurement matrix Ψ in Eq. 2.2 – it samples the $\mathbf{f}$ and aggregates the contribution of each grid cell to the ffCO₂ concentration $\mathbf{y}^{obs}$ measured at the sensing locations. While $\mathbf{H}$ is not random, wind flow patterns over the United States are unlikely to align along wavelets and the incoherence requirement might be met. Given the limited number of measurement sites, we may only be able to reconstruct a very sparse subset of the wavelet coefficients required to represent ffCO₂ emissions accurately. Sparse reconstruction methods that allow us to do so were reviewed above.

2.3 Constructing a multiscale random field model

Here, we examine wavelet models of ffCO₂ emissions. We seek an approximate representation, which is low dimensional or sparse i.e., many of the $w_{s,i,k}$ may be set to zero. We subject $\bar{\mathbf{f}}_V$ to wavelet analysis. The emissions are described on a $2^M \times 2^M$ grid, $M = 6$ i.e., our hierarchical mesh has 6 levels. The spatial resolution is $1^\circ \times 1^\circ$ and the rectangular domain extents are given by the corners (24.5N, -63.5W) and (87.5N, -126.5W). ffCO₂ emissions are restricted to $\mathcal{R}$ the lower 48 states of the US.

We select a wavelet type e.g., Daubechies 4, and perform a wavelet transform of $\bar{\mathbf{f}}_V$. At each scale s , we identify the wavelet coefficient $w_{max,s}$ with the largest magnitude and set all wavelets $|w_{s,i,j}| < w_{max,s}/1000$ equal to zero. In Fig. 2.2 (left) we plot the fraction of non-zero wavelet coefficients, on each scale s , for each of the wavelet types. We see a significant decrease in the number of wavelet coefficients when using Haar wavelets, whereas the others show a decrease only at scale 4 and higher. Clearly Haars provide the sparsest representation of ffCO₂ emissions due to their non-smooth distribution in space. In Fig 2.2 (right), we plot the average and standard deviation of the non-zero wavelet coefficients. Most of the wavelet coefficients at the finer scales are set to zero, and the means are small, regardless of the wavelet type. We see that the means and standard deviations shrink, especially after scale $s = 3$; further, the distributions of wavelet coefficients arising from the different wavelet types begin to resemble each other. This arises from the fact that there are sharp boundaries around the areas where ffCO₂ emissions occur; when subjected to a wavelet transform, the region in the vicinity of a sharp boundary gives rise to large wavelet coefficients down to the finest scale. Thus the few non-zero wavelet coefficients at the finer scales assume similar values, irrespective of the wavelet type.

Figure 2.2: Left: The fraction of wavelet coefficients which are non-zero on each scale l , when subjected to a wavelet transform using Haars, Daubechies 4 and 6 and Symlet 4 and 6 wavelets. We see that Haar wavelets provide the sparsest representation. Right: We plot the average value of the non-zero coefficients (solid lines) and their standard deviation (dashed line). We find that while Haars may provide the sparsest representation, the non-zero values tend to be large and distinct.

Henceforth, we will proceed with Haar wavelets as the basis set of choice for representing ffCO₂ emissions,

since they provide the sparsest representation. We seek a spatial parametrization for ffCO₂ emissions, of the form

$$\mathbf{f} = w' \phi' + \sum_{s=1}^M \sum_{i,j} w_{s,i,j} \phi_{s,i,j}, \quad \{s, i, j\} \in W^{(s)} \quad (2.5)$$

where $W^{(s)}$ is a set containing a small number of Haar bases. We will select the components of $W^{(s)}$ using an easily observed proxy $\mathbf{X}$ of human activity (which correlates with ffCO₂ emissions). We will use radiance calibrated nightlights [67, 33, 68, 69] for the proxy. However, nightlight radiances are also affected by economic factors [70], and we will explore maps of built-up area [71, 72] as an alternative. As mentioned in [71], the map of built-up areas uses nightlight radiances in its computations, and so these are *not* independent proxies; however the built-up area map also includes information from IGBP (International Geosphere-Biosphere Programme [73]) land-cover map. The two choices for $\mathbf{X}$ will be compared with respect to (1) sparsity, i.e., the size of $W^{(s)}$ compared to 4^M , size of the full complement of wavelet bases on a $2^M \times 2^M$ grid, (2) the correlation between $\mathbf{X}$ and $\bar{\mathbf{f}}_V$ and (3) the ability of $W^{(s)}$ to capture $\bar{\mathbf{f}}_V$.

In Fig. 2.3 (top row), we plot maps of the two proxies, coarsened to 1° resolution. Comparing with Fig. 2.1 (right), we see that they bear a strong resemblance to $\bar{\mathbf{f}}_V$. We then subject $\mathbf{X}$ to a wavelet transform and set all wavelet coefficients $|w_{s,i,j}| < \delta$ to zero, where δ is a user-defined threshold. The bases with non-zero coefficients are selected to constitute $W^{(s)}$. We reconstitute a “sparsified” proxy, $\mathbf{X}^{(s)}$, using just the bases in $W^{(s)}$, and compute the correlation between $\mathbf{X}^{(s)}$ and $\bar{\mathbf{f}}_V$. Finally, we project $\bar{\mathbf{f}}_V$ onto $W^{(s)}$, obtain its “sparsified” form $\bar{\mathbf{f}}_V^{(s)}$, and compute the error $\epsilon_f = \|\bar{\mathbf{f}}_V^{(s)} - \bar{\mathbf{f}}_V\|_2 / \|\bar{\mathbf{f}}_V\|_2$, where $\|\cdot\|_2$ denotes the ℓ_2 norm. In Fig. 2.3 (middle row), we plot the sparsity, correlation and ϵ_f for various values of δ , for both nightlights and built-up areas. For nightlights, we achieve a sparsity of around 0.25 for $\delta < 10^{-2}$ i.e., we need retain only a quarter of the Haar bases to represent nightlights. The nightlights so represented bear a correlation of around 0.7 with $\bar{\mathbf{f}}_V$, and achieve an error ϵ_f of around 0.1. Note that this measure of error reflects the inability of the MsRF to represent fine-scale details i.e., if we were interested only in spatially-aggregated quantities, the error using the sparsified representation could be far less. In contrast, using built-up area as a proxy, we see that while the sparsity achieved is similar, the correlation between $\mathbf{X}^{(s)}$ and $\bar{\mathbf{f}}_V$ is slightly higher. The behavior of ϵ_f is similar, except the error increases faster with δ , as compared to nightlights. However both nightlights and built-up area maps show significant correlation with $\bar{\mathbf{f}}_V$ and the sparsified set of Haar bases that they (i.e., the proxies) provide (using $\delta = 10^{-2}$ in both the cases) allow us to construct a low dimensional parametrization of ffCO₂ emissions.

Finally, we use $\mathbf{X}^{(s)}$ to create a “prior model” $\mathbf{f}_{pr} = c\mathbf{X}^{(s)}$ for ffCO₂ emissions, $\mathbf{f}$. c is computed such that

$$\int_{\mathcal{R}} \bar{\mathbf{f}}_V dA = \int_{\mathcal{R}} \mathbf{f}_{pr} dA = c \int_{\mathcal{R}} \left(w'^{(X)} \phi' + \sum_{l,i,j} w_{s,i,j}^{(X)} \phi_{l,i,j} \right) dA, \quad \{l, i, j\} \in W^{(s)} \quad (2.6)$$

where $\mathcal{R}$ denotes the Lower 48 states of USA and $w_{l,i,j}^{(X)}$ are coefficients from a wavelet transform of $\mathbf{X}$. This implies that c is calculated such that both $\bar{\mathbf{f}}_V$ and $\mathbf{f}_{pr}$ provide the same value for the total emissions for the US. In Fig. 2.3 (bottom row), we plot the error $(\mathbf{f}_{pr} - \bar{\mathbf{f}}_V)$. We see that neither nightlights nor the built-up area map provide a $\mathbf{f}_{pr}$ that is an accurate representation of $\bar{\mathbf{f}}_V$, though they share similar spatial patterns i.e., $\mathbf{f}_{pr}$ may be used to provide regularization for $\mathbf{f}$ in an inverse problem, but, by itself, is a poor predictor, regardless of the proxy $\mathbf{X}$ used to create it.

Figure 2.3: Top row: Maps of nightlight radiances (left) and built-up area percentage (right), for the US. Middle row: The sparsity of representation, the correlation between $\mathbf{X}$ and $\mathbf{f}_V$ and the normalized error ϵ_f between the Vulcan emissions $\mathbf{f}_V$ and the sparsified form obtained by projecting it on $\mathbf{X}$. These values are plotted for nightlights (left) and the built-up area maps (right). Bottom row: Plots of $(\mathbf{f}_{pr} - \mathbf{f}_V)$ obtained from nightlights (left) and built-up area maps (right).

2.4 Formulation of the estimation problem

In this section, we pose the inverse problem to estimate ffCO₂ emissions from limited measurements and describe a method to solve it. Fig. 2.4 provides a flowchart of the estimation procedure. Sec. 2.3 addressed the top half of the figure where we (1) selected Haar wavelets for modeling ffCO₂ emissions and (2) derived “prior” models for emissions based on nightlights and built-up area maps. The lower half of Fig. 2.4 illustrates the issues addressed in this section. We will use MsRF constructed from Haars to create a model for the (unknown) ffCO₂ fluxes and concentrations (which will also involve an atmospheric transport model). The fluxes will be constrained by observations of ffCO₂ concentrations. We then use sparse reconstruction, along with the “prior” model to estimate the ffCO₂ emissions.

Below, we briefly describe the transport model used to link observations at a set of measurement towers to emissions on a numerical grid. Thereafter, we pose the estimation problem, followed by a review of the method used to solve it. Finally, we describe how we enforce non-negativity in the estimated ffCO₂ emission estimates.

2.4.1 Transport model

The transport of CO₂ in the atmosphere is modeled using Eq. 2.1. The calculation of the sensitivities $\mathbf{H}$ is described in detail in the supplement to [2], and we provide a summary below.

The elements of the $\mathbf{H}$ matrix are calculated using the Stochastic Time-Inverted Lagrangian Transport Model (STILT) model [74]. STILT has already been used in a number of inversion studies for estimating biospheric CO₂ fluxes [75, 2]. It represents air arriving at the observation locations as an ensemble of particles, and transports them in reverse using wind fields. These wind fields are generated by numerical weather prediction models, in this particular case, the Weather Research & Forecasting (WRF) model [76] version 2.2. WRF has been customized for STILT and other transport models [77]. The details of the submodels used in the WRF-STILT calculations are in the supplementary material of [2]. The grid used for generating $\mathbf{H}$ was three-level, with a 2 kilometer resolution level around the 3 tallest towers (LEF, AMT and WKT; see Sec. B for details). The 2 km grid was embedded in a 10 km resolution grid over northern the Midwest and Gulf Coast region and New England, extending to approximately 105° W. An outermost, 40 km resolution, grid covered the rest of the domain (see Fig. 1 in [75] for a diagram). At each measurement location, 500 particles were released every hour and their 10-day back-trajectories calculated using meteorology from 2008.

Concentration footprints (or sensitivities) were calculated at 3-hour intervals by integrating the trajectories over the North American $1^\circ \times 1^\circ$ grid as described in [74]. The sensitivity of the CO₂ concentration at each observation location due to the flux at each grid-cell (the “footprint”) is calculated in units of ppmv / $\mu\text{mol m}^{-2} \text{ s}^{-1}$ (ppmv: parts per million by volume). The footprint, which is the adjoint of the transport field, is calculated by counting the number of particles in a surface-influenced region (defined as a fraction of the estimated planetary boundary layer height at that grid-cell) for a given time-period spent in the region (for details, see [74]). When multiplied by a flux field, the sensitivities provide the contribution of fluxes to the mixing ratios (ppmv) at the measurement location/receptor. These sensitivities, calculated for 2008, were used in the estimation of biospheric fluxes [78, 26]. The sensitivities for CO₂ fluxes with an 8-day resolution (i.e., the ones used in this work) were obtained, from the 3-hour sensitivities described above by simply adding the $8 \times 24/3 = 64$ sensitivities that spanned the 8-day period. Thereafter, the grid-cells outside $\mathcal{R}$ were removed to obtain the $\mathbf{H}$ matrix used in this study. The size of the $\mathbf{H}$ matrix is $(K_s N_s) \times (N_{\mathcal{R}} K)$, where

Figure 2.4: Flowchart of the procedure for performing the inversion of ffCO₂ emissions. The top half of the figure was addressed in Sec. 2.3. We start with a collection of wavelet families that *could* be used for modeling the emissions, and find that Haars provide the sparsest representation. Thereafter, we use the Haars to create a “prior” model $f_{pr} = c \tilde{\Phi}'_R w$ for the emissions. In the lower half of the figure, we illustrate the posing and solution of the inverse problem. Step A, the posing of the inverse problem is described in Sec. 2.4.2. Three slightly different formulations (Step B), based on how the “prior” is incorporated into the inverse problem, are also described in the same section. Step C, the StOMP algorithm, is described in Sec. 2.4.3. Sec. 2.4.4, the description of non-negativity on the ffCO₂ emissions, corresponds to Step D.

K_s is the number of tower measurements every year ($24/3 \times 45 = 2880$), N_s is the number of sensors/towers, $N_R = 816$ is the number of grid-cells in $\mathcal{R}$, the part of the domain covered by the lower 48 states of the US and $K = 45$ is the number of 8-day periods in 360 days (approximately a year).

2.4.2 The inverse problem

We seek to estimate ffCO₂ emissions over the region constituting the lower 48 states of the US, averaged over 8-day periods. The spatial distribution of emissions during an arbitrary 8-day period k is denoted by $\mathbf{f}_k$. We are interested in estimating emissions over an entire year i.e., we seek $\mathbf{F} = \{\mathbf{f}_k\}, k = 1 \dots K$.

We will model emissions on the $2^M \times 2^M, M = 6$ mesh with wavelets:

$$\begin{aligned} \mathbf{f}_k &= w'_k \phi' + \sum_{s=1}^M \sum_{i,j} w_{s,i,j,k} \phi_{s,i,j}, \quad \{s,i,j\} \in W^{(s)} \\ &= \Phi \mathbf{w}_k. \end{aligned} \quad (2.7)$$

Note that Φ comprises of only those wavelets selected using $\mathbf{X}$ and contained in $W^{(s)}$, and not the 4^M wavelets that the grid can support. For the entire year, the expression for emissions becomes

$$\mathbf{F} = \begin{pmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_K \end{pmatrix} = \text{diag}(\Phi, \Phi, \dots, \Phi) \begin{pmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \\ \vdots \\ \mathbf{w}_K \end{pmatrix} = \tilde{\Phi} \mathbf{w}.$$

Since $\Phi \mathbf{w}_k$ models the emissions over all grid-cells, and not just $\mathcal{R}$, $\mathbf{F}$ contains emissions over the lower 48 states, as well as the region outside it (where we have assumed that the emissions are non-existent). We separate out the two fluxes by permuting the rows of $\tilde{\Phi}$

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

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 L is the number of wavelets in $W^{(s)}$ and $N_{\mathcal{R}'} = 4096 - 816 = 3280$ 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 $\mathbf{w}$, $\mathbf{F}_{\mathcal{R}'}$, the emissions in the region outside $\mathcal{R}$, are not zero. Consequently, it will be necessary to specify $\mathbf{F}_{\mathcal{R}'} = 0$ as a constraint in the inverse problem.

Specifying the constraint $\mathbf{F}_{\mathcal{R}'} = 0$ directly is not very efficient since it leads to $N_{\mathcal{R}'}K$ constraints. In a global inversion, or at resolutions higher than $1^\circ \times 1^\circ$, this could get very large. Consequently, 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 is called a uniform spherical ensemble and is used in compressive sampling [54]. The projection of the emission field $\mathbf{F}_{\mathcal{R}'}$ on $\mathbf{R}$ i.e., $\mathbf{R} \mathbf{F}_{\mathcal{R}'}$ compressively samples $\mathbf{F}_{\mathcal{R}'}$. Setting this projection to zero during inversion allows us to enforce zero emissions outside $\mathcal{R}$. However, to do so, we add only M_{cs} constraint equations rather than $N_{\mathcal{R}'}K$, which would be the case if we set the emission in each grid-cell in $\mathcal{R}'$ to zero. In Sec. 2.5.5, we will investigate the degree of computational saving afforded by imposing the $\mathbf{F}_{\mathcal{R}'} = 0$ constraint in this manner.

The optimization problem can be written as

$$\mathbf{Y} = \begin{pmatrix} \mathbf{y}^{obs} \\ 0 \end{pmatrix} \approx \begin{pmatrix} \mathbf{H} \tilde{\Phi}_{\mathcal{R}} \\ \mathbf{R} \tilde{\Phi}_{\mathcal{R}'} \end{pmatrix} \mathbf{w} = \mathbf{G} \mathbf{w}. \quad (2.8)$$

In this equation, $\mathbf{G}$ is akin to $\mathbf{A}$ discussed in Sec. 2.2.3. It is a “sampling” of a basis set. The $\mathbf{R} \tilde{\Phi}_{\mathcal{R}'}$ component of $\mathbf{G}$ is formally a random sampling. The $\mathbf{H} \tilde{\Phi}_{\mathcal{R}}$ component of $\mathbf{G}$, on the other hand, “samples” the emission field as guided by the transport processes encoded in $\mathbf{H}$. The left hand side $\mathbf{Y}$ is approximately equal to $\mathbf{G} \mathbf{w}$ since the observations $\mathbf{y}^{obs}$ contain measurement errors that cannot be modeled with $\mathbf{H}$. This completes Step A in Fig. 2.4.

The wavelet coefficients $\mathbf{w}$ in Eq. 2.8 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 / “footprint” tend to have coefficients with a large magnitude. Their small footprint 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. 2.8, as-is, incorporates no information from $\mathbf{X}$ beyond the selection of wavelets to be included in $\tilde{\Phi}$. We explore the incorporation of $\mathbf{X}$ in the estimation of $\mathbf{w}$ using three different approaches:

Approach A : This is the baseline approach and solves Eq. 2.8 as-is. The lack of normalization of $\mathbf{w}$, in conjunction with the optimization procedure described below in Sec. 2.4.3, leads to artifacts which will be described in Sec. 2.5.1.

Approach B : In this formulation, $\mathbf{X}$ is included in the optimization as a “prior”. We write the emissions as $\mathbf{F} = \mathbf{f}_{pr} + \Delta\mathbf{F}$. Substituting into Eq. 2.8, we get $\mathbf{Y} \approx \mathbf{H}\mathbf{f}_{pr} + \mathbf{G}\Delta\mathbf{w}$, where $\Delta\mathbf{w} = \mathbf{w} - \mathbf{w}_{(X)}$. Here, $\mathbf{w}_{(X)} = c\{w'^{(X)}, w_{s,i,j}^{(X)}\}$, $\{s, i, j\} \in W^{(s)}$, where c is obtained from Eq. 2.6. It provides a wavelet decomposition of $\mathbf{f}_{pr}$ using the bases in $W^{(s)}$. Simplifying, we get

$$\Delta\mathbf{Y} = \mathbf{Y} - \mathbf{H}\mathbf{f}_{pr} \approx \mathbf{G}\Delta\mathbf{w}, \quad (2.9)$$

The term “prior” model is used somewhat loosely since this is not a Bayesian method. However, $\mathbf{f}_{pr}$ serves a similar function by providing regularization in the inverse problem.

Approach C : The incorporation of the spatial patterns in $\mathbf{X}$ into the estimation procedure can be performed in an alternative manner. We note that $\mathbf{w}_{(X)}$ can be used to normalize $\mathbf{w}$. We rewrite Eq. 2.8 as

$$\mathbf{Y} \approx \mathbf{G} \text{diag}(\mathbf{w}_{(X)}) \text{diag}(\mathbf{w}_{(X)}^{-1}) \mathbf{w} = \mathbf{G}' \mathbf{w}' = \begin{pmatrix} \mathbf{H} \tilde{\Phi}'_{\mathcal{R}} \\ \mathbf{R} \tilde{\Phi}'_{\mathcal{R}'} \end{pmatrix} \mathbf{w}', \quad (2.10)$$

where $\mathbf{w}' = \{w_{s,i,j}/w_{(X),s,i,j}, \{s, i, j\} \in W^{(s)}\}$, is the normalized set of wavelet coefficients, $\tilde{\Phi}'_{\mathcal{R}} = \tilde{\Phi}_{\mathcal{R}} \text{diag}(\mathbf{w}_{(X)})$ and $\tilde{\Phi}'_{\mathcal{R}'} = \tilde{\Phi}_{\mathcal{R}'} \text{diag}(\mathbf{w}_{(X)})$.

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

$$\mathbf{Y} \approx \mathbf{\Gamma}\boldsymbol{\zeta}. \quad (2.11)$$

This completes Step B in Fig. 2.4. We outline the solution procedure in Sec. 2.4.3 and compare the performance of the three formulations in Sec. 2.5.1.

2.4.3 Solving the inverse problem

Eq. 2.11 is solved using an optimization procedure. Since we obtain our observations from a set of measurement locations that were sited with an eye towards biospheric CO₂ fluxes (see Appendix B), it is unlikely that they will constrain all the elements of $\boldsymbol{\zeta}$. Here we fall back on our multiscale model of $\mathbf{F}_{\mathcal{R}}$, which explicitly parametrizes orthogonal spatial patterns (the wavelets) of different scales. Further, only those patterns that were observed in the $\mathbf{X}$ field were retained in $W^{(s)}$. $\mathbf{y}^{obs}$ is probably sufficient to estimate the larger patterns, whereas it may not be possible to constrain the finer details. However, *a priori*, we do not know the identity of these “un-constrainable” details.

In our solution of Eq. 2.11, we will attempt to identify the wavelets coefficients that cannot be estimated from $\mathbf{Y}$ and drive them to zero i.e, the solution method will impose sparsity on $\boldsymbol{\zeta}$. This will provide us with a reconstruction of $\mathbf{F}_{\mathcal{R}}$, containing only those features/patterns that are supported by data. The sparse solution of Eq. 2.11 is performed using Stagewise Orthogonal Matching Pursuit (StOMP, [60]). Eq. 2.11 is recast similar to Eq. 2.4

$$\begin{aligned} & \underset{\boldsymbol{\zeta} \in \mathbb{R}^N}{\text{minimize}} && \|\boldsymbol{\zeta}\|_1 \\ & \text{subject to} && \|\mathbf{Y} - \boldsymbol{\Gamma}\boldsymbol{\zeta}\|_2^2 < \varepsilon_2. \end{aligned} \quad (2.12)$$

$\|\boldsymbol{\zeta}\|_1$ is minimized by setting as many elements of $\boldsymbol{\zeta}$ to zero as possible, thus enforcing sparsity. Meanwhile, the constraint $\|\mathbf{Y} - \boldsymbol{\Gamma}\boldsymbol{\zeta}\|_2$ ensures that the solutions being proposed by the optimization procedure provide a good reproduction of the observations. While the details of StOMP are in [60], we summarize the algorithm below.

StOMP operates iteratively via S stages, building up a sequence of approximations $\boldsymbol{\zeta}_0, \boldsymbol{\zeta}_1, \dots$ by removing detected structure from a sequence of residual vectors $r_1, r_2, \dots$. We start with an initial guess $\boldsymbol{\zeta}_0 = 0$ and initial residual $r_0 = \mathbf{Y}$. The stage counter is initialized $s = 1$. The algorithm maintains a sequence of estimates $I_1, I_2, \dots$ of the locations of non-zeros in $\boldsymbol{\zeta}$.

In the l^{th} stage, we obtain a vector of residual correlation

$$c_l = \boldsymbol{\Gamma}^T r_{l-1}$$

which we consider to be a vector with a few significant non-zeros and the rest containing Gaussian noise of small magnitude. Based on a threshold calculated from the assumption that c_l is mostly i.i.d. Gaussian noise, we identify a small set J_l of “large” coordinates:

$$J_l = \{j : |c_l(j)| > t_l \sigma_l\}$$

where $\sigma_l = \|r_l\|_2 / \sqrt{|r_l|}$, $2 \leq t_l \leq 3$ and $|r_l|$ is the length of the vector r_l . The newly detected locations of non-zeros are added to the running list of the support of $\boldsymbol{\zeta}$

$$I_l = I_{l-1} \cup J_l.$$

We then project $\mathbf{Y}$ on the columns of $\boldsymbol{\Gamma}$ contained in the enlarged support. Let $\boldsymbol{\Gamma}_l$ denote the $(K_s N_s + M_{cs}) \times |I_l|$ matrix constructed with the columns chosen using the index set I_l . The new approximation $\boldsymbol{\zeta}_l$, with support I_l is given by

$$(\boldsymbol{\zeta}_l)_{I_l} = (\boldsymbol{\Gamma}_l^T \boldsymbol{\Gamma}_l)^{-1} \boldsymbol{\Gamma}_l^T \mathbf{Y}$$

and the updated residual is

$$r_l = \mathbf{Y} - \boldsymbol{\Gamma}\boldsymbol{\zeta}_l.$$

We check a stopping criterion, and proceed to the next iteration after setting $l := l + 1$.

This completes Step C in Fig. 2.4.

2.4.4 Enforcing non-negativity of $\mathbf{F}_{\mathcal{R}}$

Estimates of $\mathbf{w}$ calculated by StOMP do not necessarily provide $\mathbf{F}_{\mathcal{R}} = \mathbf{H} \tilde{\boldsymbol{\Phi}}_{\mathcal{R}} \mathbf{w}$ that are non-negative. In practice the negative ffCO₂ emissions occur in only a few grid-cells and are usually small in magnitude. A

large fraction of elements of $\mathbf{F}_{\mathcal{R}}$ are set to zero by StOMP. Having identified the sparsity pattern, i.e., the spatial scales that can and cannot be estimated from $\mathbf{Y}$, we devise an iterative procedure for enforcing non-negativity on $\mathbf{F}_{\mathcal{R}}$. We discard $\mathbf{F}_{\mathcal{R}'}$ and manipulate the emissions in $\mathcal{R}$ directly, rather than via the wavelet coefficients.

We seek non-negative ffCO₂ emissions $\mathbf{E} = \{E_i\}, i = 1 \dots Q, Q = (N_{\mathcal{R}}K)$ such that

$$\frac{\|\mathbf{y}^{obs} - \mathbf{H}\mathbf{E}\|_2}{\|\mathbf{y}^{obs}\|_2} \leq \epsilon_3. \quad (2.13)$$

$\mathbf{E}$ is constructed iteratively through a sequence $\mathbf{E}_1, \mathbf{E}_2, \dots$. $\mathbf{E}_0$ is initialized by using the absolute values of $\mathbf{F}_{\mathcal{R}}$ calculated by solving Eq. 2.12.

At each iteration m , we seek a correction $\xi = \{\xi_i\}, i = 1 \dots Q$, where $|\xi_i| \leq 1$, such that

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

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

$$\mathbf{y}^{obs} - \mathbf{H}\mathbf{E}^{(m-1)} = \Delta \mathbf{y} \approx \mathbf{H}\Delta \mathbf{E}^{(m-1)} \quad (2.14)$$

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 emissions updated as

$$\begin{aligned} \xi_i &= \text{sgn}\left(\frac{\Delta E_i^{(m-1)}}{E_i^{(m-1)}}\right) \max\left(1, \left|\frac{\Delta E_i^{(m-1)}}{E_i^{(m-1)}}\right|\right), \\ E_i^{(m)} &= E_i^{(m-1)} \exp(\xi_i), \end{aligned} \quad (2.15)$$

to obtain $\mathbf{E}^{(m)}$.

The convergence requirement Eq. 2.13 is checked with $\mathbf{E}^{(m)}$, and if not met, the iteration count is updated $m := m + 1$ and Eq. 2.14 is solved again. This completes Step D in Fig. 2.4.

To summarize, the solution of the inverse problem proceeds in two steps:

1. *Step I – Sparse estimation of $\mathbf{F}_{\mathcal{R}}$ using the wavelet model:* This is performed by StOMP, which explicitly identifies small scales that cannot be constrained by $\mathbf{y}^{obs}$ and drives them to zero. The estimation of $\mathbf{F}_{\mathcal{R}}$ is performed by calculating the wavelet coefficients $\mathbf{w}$. We obtain an approximation to the emissions, $\mathbf{F}_{\mathcal{R}}$, which may be negative in certain grid-cells.
2. *Step II – Enforcement of non-negativity:* Non-negative emissions $\mathbf{E}$ are calculated by iteratively updating $|\mathbf{F}_{\mathcal{R}}|$ with a multiplicative correction. Updates are performed on fluxes in grid-cells, not the wavelet coefficients. Each iteration includes an invocation of StOMP.

2.5 Tests

In this section, we test the multiscale parameterization for ffCO₂ emissions, developed in Sec. 2.3, using the sparse estimation technique in Sec. 2.4.

Generating synthetic observations: The true ffCO₂ emissions are obtained, for 2002, from the Vulcan inventory [3, 4]. Vulcan provides emissions in $\mathcal{R}$ and we assume that there are no emissions elsewhere. Hourly Vulcan fluxes are coarsened from 0.1° resolution to 1°, and averaged to 8-day periods. These fluxes are multiplied by $\mathbf{H}$ to obtain ffCO₂ concentrations at $K_s = 35$ measurement towers (see Appendix B). Observations are available every 3 hours and span a full year. A measurement error $\varepsilon \sim N(0, \sigma^2)$ is added to the concentrations to obtain $\mathbf{y}^{obs}$, as used in Eq. 2.8. The same σ is used for all towers and is set to 0.01 ppm (approximately, 1% of the average concentration observed at the 35 towers).

2.5.1 Comparison of optimization formulations

We solve Eq. 2.8 using the StOMP method (Sec. 2.4.3) and enforce non-negativity on $\mathbf{F}_{\mathcal{R}}$ to obtain $\mathbf{E}$. The coefficients $\mathbf{w}_{(\mathbf{x})}$ used in Eq. 2.9 and Eq. 2.10 are obtained from a wavelet decomposition of $\mathbf{f}_{pr}$ based on nightlights (Sec. 2.3). The constant c in Eq. 2.6 is obtained by using fluxes from the EDGAR inventory [34] for 2005 i.e., instead of using emissions from Vulcan to calculate $\bar{\mathbf{f}}_V$, we use EDGAR. We thus ensure that Vulcan is not used in any capacity during the inverse modeling (except to generate the synthetic observations). The inversion is performed for $k = 1 \dots 45$, for the entire year. The following parameters are used in the inversion process (Sec. 2.4.3 & 2.4.4): $\varepsilon_2 = 10^{-5}$, $\varepsilon_3 = 5.0 \times 10^{-4}$, $M_{cs} = 13,500$ i.e., 300 random projections for each 8-day period.

In Fig. 2.5 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 towers is very uneven, with most of the towers being concentrated in the Northeast quadrant. We see that Approach A (Fig. 2.5, top right) provides estimates that have large areas in the Northwest (NW) and Southwest (SW) quadrants with low levels of ffCO₂ emissions. In contrast, the true emissions (Fig. 2.5, top left) are mostly empty, which is also borne out by nightlight and built-up area maps in Fig. 2.3. Thus we see that the minimization of $\|\boldsymbol{\zeta}\|_1$ (alternatively $\|\mathbf{w}\|_1$) drives the wavelet coefficients to small values, but not identically to zero. In Fig. 2.5 (bottom left), Approach B provides estimates that show much structure in the Eastern quadrants, and the patterns seen in nightlights (Fig. 2.3) are clearly reproduced. The reason is as follows. While $\mathbf{f}_{pr}$ captures the broad, coarse scale patterns of ffCO₂ emissions, it incurs significant errors at the finer scales. Eq. 2.9 essentially seeks to rectify these errors. However, as mentioned in Sec. 2.4.2, fine-scale wavelets tend to have large wavelet coefficients and the minimization of $\|\boldsymbol{\zeta}\|_1$ (alternatively $\|\Delta\mathbf{w}\|_1$) removes them since the constraint $\|\mathbf{Y} - \mathbf{\Gamma}\boldsymbol{\zeta}\|_2^2 < \varepsilon_2$ is not very sensitive to individual wavelets at the fine scale (they only affect neighboring towers significantly). The inability to rectify the fine-scale discrepancies lead to a final ffCO₂ estimate that resembles $\mathbf{f}_{pr}$ in the finer details. Fig. 2.5 (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.

In Fig. 2.6 (left) we evaluate the accuracy of the reconstruction quantitatively. We total the emissions in $\mathcal{R}$ to obtain the country-level ffCO₂ emissions and compare that with the emissions from Vulcan. We plot a

Figure 2.5: Plots of ffCO_2 emissions during the 31st 8-day period. 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 $\mathbf{f}_{pr}$ provides low levels of (erroneous) emissions in large swathes of the Western quadrants. Approach B reflects $\mathbf{f}_{pr}$ very strongly. Approach C provides a balance between the influence of $\mathbf{f}_{pr}$ and the information in $\mathbf{y}^{obs}$.

time-series of errors defined as a percentage of total, country-level Vulcan emissions

$$\text{Error}_k (\%) = \frac{100}{K} \sum_{k=1}^K \frac{E_k - E_{V,k}}{E_{V,k}}, \quad \text{where} \quad E_k = \int_{\mathcal{R}} \mathbf{E}_k dA \quad \text{and} \quad E_{V,k} = \int_{\mathcal{R}} \mathbf{f}_{V,k} dA. \quad (2.16)$$

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. 2.6 (right) we plot the correlation between the true and reconstructed emissions over the same duration. This was done by calculating the Pearson correlation coefficient between the vector of grid-cells covering the United States. It is clear that Approach B provides the worst reconstructions, with the largest errors and smallest correlations. Approach C tends to over-predict emissions, whereas Approach B tends to under-predict them, often by approximately the same (small) amount. The correlations of the reconstructed emissions from these two emissions are about the same.

In Fig 2.7 we see the essential difference between Approach A and C. We plot the reconstruction error

Figure 2.6: Comparison of estimation error (left) and the correlation between true and estimated emissions (right) using Approaches A, B and C. It is clear that Approach B is inferior to the others.

(left figure) and correlation between true and reconstructed emissions (right figure) 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. 2.5). Thus normalization using $\mathbf{w}_{(X)}$ and minimization of $\|\boldsymbol{\zeta}\|_1$ (alternatively $\|\mathbf{w}'\|_1$) prevents large departures from $\mathbf{f}_{pr}$ 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. Note, however, that both Approach A and C have similar fidelity at the scale of $\mathcal{R}$.

In the tests below, we will restrict ourselves to Approach C.

2.5.2 Evaluating formulation using compressive sensing metrics

In this section we compare Approach A versus C in terms of certain compressive sensing metrics to explain the lower accuracy of the estimates developed using Approach A. Note that in Approach A, sparsity is the sole source of regularization.

As specified in Eq. 2.2, the process of acquiring compressive measurements $\mathbf{g}'$ involves project wavelet bases Φ on the rows of Ψ . Since the rows $\psi_{i,\cdot}$ of Ψ are random unit vectors, they are neither aligned with nor orthogonal to the bases $\phi_{\cdot,j}$ and consequently the elements of $\mathbf{A}_\Psi = \Psi\Phi$ tend to assume a distribution of magnitudes which are generally not very close to 0 or 1. This *coherence* $\mu(\Psi, \Phi)$ is defined as [79]

$$\mu(\Psi, \Phi) = \sqrt{N} \max_{1 \leq i, j \leq N} |\langle \psi_{i,\cdot}, \phi_{\cdot,j} \rangle|, \quad 1 \leq \mu(\Psi, \Phi) \leq \sqrt{N} \quad (2.17)$$

Values of μ near 1 (alternative $\max |\langle \psi_{i,\cdot}, \phi_{\cdot,j} \rangle|$ near $1/\sqrt{N}$) indicate incoherence between Ψ and Φ and more efficient sampling. Note that the lower bound assumes that Ψ is a random matrix.

In compressive sensing, random matrices such as Gaussians, Hadamard, Circulant/Toeplitz or functions such as noiselets [54, 80, 81, 82] serve as Ψ . In Fig. 2.8, we plot the distribution of $\log_{10}(|A_{i,j}|)$, the elements

Figure 2.7: Reconstruction error (left) 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 $\mathbf{f}_{pr}$, leads to lower errors in both the quadrants and better correlations in the less instrumented NW quadrant.

of $\mathbf{A}\Psi$ for some “standard” sampling matrices. Φ contains the wavelets used in Sec. 2.5.1 i.e., those selected using nightlights. 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. The samples collected by these projection schemes are sufficiently informative, due to the incoherence between Ψ and Φ , to allow reconstruction of the original image with sparsity as the sole regularization.

In Eq. 2.8, $\mathbf{H}$ serves a similar sampling purpose. $\mathbf{H}$ is determined by atmospheric transport processes and is not a random matrix. However, 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 2 towers and for the 21st and 22nd 8-day periods. We compute $\mathbf{A}_{\mathbf{H}'} = \mathbf{H}'\Phi$, and in Fig. 2.8, 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 $\mathbf{A}\Psi$, $\mathbf{A}_{\mathbf{H}'}$ contains a significant number of elements that are close to 1, and a large number of elements which 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 others. The small values in $\mathbf{A}_{\mathbf{H}'}$ 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. In fact, the dominance of near-field CO_2 fluxes on tower measurements is well known [83] and is responsible for the particular structure of $\mathbf{A}_{\mathbf{H}'}$.

Further, the coherence $\mu(\mathbf{H}', \Phi)$ is larger than $\mu(\Psi, \Phi)$, indicating that the sampling efficiency of $\mathbf{H}'$ is inferior to what is generally achieved in compressive sensing. Consequently a combination of sparsity and a “prior” emission model $\mathbf{f}_{pr}$ were required to regularize the problem and enhance the accuracy of the emission estimates in Approach C.

Figure 2.8: Comparison of the distribution of the elements of $\mathbf{A}_{\Psi}$ and $\mathbf{A}_{\Phi}$. 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 $\mathbf{A}_{\Psi}$ where the elements assume discrete values. In contrast, the elements of $\mathbf{A}_H$ assume values which are spread over a far larger range, some of which are quite close to 1 while others are very close to zero.

2.5.3 Inversions with nightlights as proxy

In Fig. 2.9, we plot the true and reconstructed emissions for 3 8-day periods ($k = 9, 33, 42$). We see that due to the abundance of measurement towers in the NE quadrant, the reconstruction there is accurate. On the other hand, the tower density on the West coast is sparse and gives rise to significant inaccuracies. For example, in 8-day periods 33 and 42, we see that the Los Angeles–San Diego region (Southwest quadrant) is underestimated, since the only tower is in La Jolla (32.87N, -117.26W). The estimated emissions in the center of the country (Continental Divide and Great Plains, in the Western quadrants) show some underestimation, and far less structure than the true ffCO₂ emissions. This is due to the presence of just 3 towers, which are also clustered in a relatively small region. The region around the Gulf of Mexico is not well estimated since there are few towers there.

In Fig. 2.10 (left) we plot a time-series of errors defined as a percentage of total, country-level Vulcan emissions. Percent errors in reconstructed emissions are calculated using Eq. 2.16. The “prior” errors are computed as

$$\text{Error}_{pr,k} (\%) = \frac{100}{K} \sum_{k=1}^K \frac{E_{pr} - E_{V,k}}{E_{V,k}} \quad \text{where} \quad E_{pr} = \int_{\mathcal{R}} \mathbf{f}_{pr} dA.$$

We see 25% errors in $\mathbf{f}_{pr}$. This is a consequence not only of the disagreement between EDGAR (in 2005) and Vulcan (in 2002), but also the manner in which they account for emissions. Since we are only interested in obtaining a rough guess of US emissions with $\mathbf{f}_{pr}$, we did not perform a careful analysis and comparison. As

Figure 2.9: Reconstruction of the ffCO₂ emissions from the 35 towers (plotted as diamonds) over one year. The true emissions are on the left and the reconstructions on the right. We see that the large scale structure of the emissions have been captured, as seen by the comparisons for 8-day periods number 9 (mid-March), 33 (end of August) and 42 (early November). The west coast of the US has few towers near heavily populated regions and thus is not very well estimated. In the middle and bottom rows, we see that the emissions in the Los Angeles-San Diego region (32.87N, -117.26W) is underestimated. On the other hand, due to the higher density of towers in the Northeast, the true and estimated emissions are qualitatively similar. Emissions have units of $\mu\text{mol m}^{-2} \text{s}^{-1}$ of C (not CO₂).

Figure 2.10: Comparison of reconstruction error and correlations. Left: We plot the error between the reconstructed and true (Vulcan) emissions in blue. In red, we plot the error between f_{pr} using nightlights and Vulcan emissions. We see that assimilation of y^{obs} leads to significantly improved accuracy. Right: We plot the accuracy of the spatial distribution of the reconstructed emissions. The Pearson correlations $C(E_k, f_{V,k})$ and $C(f_{pr}, f_{V,k})$ show that incorporating y^{obs} improves the spatial agreement of estimated emissions versus the true one from 0.7 to around 0.9. If the emissions are averaged over 32-day periods, rather than 8-day periods, the correlation with true (Vulcan) emissions rises even higher.

can be seen, assimilation of y^{obs} reduces the error to a maximum of 4%. The least accurate reconstructions are during spring. In order to check the accuracy of the spatial distribution of E_k , we calculate the Pearson correlations $C(E_k, f_{V,k})$ and $C(f_{pr}, f_{V,k})$. We see that data assimilation increases the correlation from around 0.7 to 0.85-0.95, with a mean around 0.9. When the emissions are aggregated/averaged over 32-day periods, the correlation increases to 0.9-0.95, with a mean around 0.93. Thus, the ffCO₂ emissions obtained using a nightlight proxy are substantially improved by the incorporation of y^{obs} .

Next, we address the impact of tower density. As was clear in Fig. 2.9, the reconstruction in the Northeast (NE) quadrant is more accurate compared to the reconstruction elsewhere. We compute the percent error between reconstructed and true (Vulcan) emissions on a quadrant basis, and plot them in Fig. 2.11. The correlation between the reconstructed and Vulcan emissions are plotted too. In Fig. 2.11 (left), we see that the reconstruction error in the NE quadrant is far smaller than elsewhere. Further, the reconstruction error in individual quadrants is far higher than in the country as a whole (Fig. 2.10), where the total reconstruction error never exceeded 4%. Thus, while large scale structures (e.g., at the scale of the US) are being estimated quite accurately, significant errors are seen when we address regions the size of the individual quadrants. The consequences of not being able to constrain quadrant-scale structures is seen in Fig. 2.11, where the NE quadrant displays better correlation with Vulcan than the rest. The most inaccurate spatial patterns are seen in the Northwest (NW) quadrant that includes the Great Plains and the Continental Divide. Here, the lack of strong emissions, nightlights and towers make the reconstruction particularly poor.

We now address some of the numerical aspects of the solution. In Fig. 2.12 (left) we plot y predicted by the reconstructed emissions at 3 towers. We see that the ffCO₂ concentrations are well reproduced by the estimated emissions. Thus, the lack of fidelity at the smaller scales (seen in Fig. 2.11) do not substantially impact the measurements. In Fig. 2.12 (right) 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

Figure 2.11: Plot of the error in, and the correlation between reconstructed emissions and true (Vulcan) ones. Left: We see that the reconstruction error in the NE quadrant is small compared to the others. However, the error in each of the quadrants is much larger than the error at the country level (see Fig. 2.10). Right: Correlation in each of the quadrants. The NE quadrant is substantially better than the others due to the higher density of towers.

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 the range (Fig. 2.12, right; red symbols). During the sparse reconstruction from sparse data, these coefficients are set to zero (blue line in Fig. 2.12, right). The coefficients corresponding to the low-index coefficients, which represent large structures are estimated accurately. This explains the good correlation between true and estimated emissions at the country scale and a far poorer one at the quadrant scale (Fig. 2.11). The explicit separation of scales is thus leveraged into ignoring unimportant, fine-scale details (which are difficult to constrain with data) and focusing model-fitting effort on the scales that can be resolved by the data. Sparse reconstruction achieves this in an automatic, purely data-driven manner, rather than via a pre-processing, scale-selection step.

Finally we see the impact of the enforcement of non-negativity. In Fig. 2.13 we plot the cumulative distribution function (CDF) of the ffCO_2 emissions in $\mathcal{R}$ before and after the imposition of non-negativity, as described in Sec. 2.4.4. The emissions are from the 31st 8-day period (depicted in Fig. 2.12). We see from the CDF that before the imposition of non-negativity, the number of grid-cells with negative emissions is small; further, the negative emissions are small in magnitude. Thus the sparse reconstruction (Step I in Sec. 2.4.4) provides a very good approximation to the final estimated fluxes, by detecting the spatial patterns/wavelets that could be constrained by observations. Given a good approximation, the non-negativity enforcement converges quickly.

2.5.4 Inversions with built-up area maps as proxy

We investigate the effect of built-up area (BUA) maps, instead of images of nightlights, as the proxy. Changing the proxy results in a different set of wavelets being chosen (nightlights resulted in a $W^{(s)}$ of 1031

Figure 2.12: Left: Prediction of ffCO₂ concentrations at 3 measurement locations, using the true (Vulcan) and reconstructed emissions (blue lines) over an 8-day period (Period no. 31). Observations occur every 3 hours. We see that the concentrations are accurately reproduced by the estimated emissions. Right: 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.

Figure 2.13: CDF of emissions in $\mathcal{R}$, before and after the imposition of non-negativity, as described in Sec. 2.4.4. We see that the CDF of the emissions without non-negativity imposed contains a few grid-cells with negative fluxes; further, the magnitude of the negative emissions is small. Thus the spatial parameterization, with sparse reconstruction provides a good approximation of the final, non-negative emissions.

wavelets; the corresponding number for BUA was 1049); further one was *not* a strict subset of the other. It also results in a different normalization in Eq. 2.10. The inversion was performed in a manner identical to that in Sec. 2.5.3. In Fig. 2.14 (left) we see that emissions reconstructed from BUA as the proxy are more variable; however, averaged over the entire 45 8-day periods, the reconstruction errors are not very different. In Fig. 2.14 (right), we see that the correlation between estimated and true emissions is better in case of nightlight-based reconstruction (when performed at a 8-day temporal resolution). However, when averaged to a 32-day temporal resolution, the difference between the two reconstructions (nightlights- versus BUA-based) is minimal. In Fig. 2.15 we investigate the differences between the nightlight- and BUA-based reconstructions at the quadrant level. We see in Fig. 2.15 (left) that the difference between nightlight- and BUA-based reconstruction errors in the NE quadrant are smaller than those for the NW quadrant. Thus, while the prior emissions from nightlights and BUA are quite different (see the last row of Fig. 2.3), the estimated emissions are well constrained by $\mathbf{y}^{obs}$ in the NE quadrant and the impact of the priors/proxies is small. This is not the case for the NW quadrant, where not only are the errors high, the two $\mathbf{f}_{pr}$ lead to reconstruction errors of different signs. Thus the region is badly constrained, which is not surprising given the paucity of towers (see Fig. 2.9). However, on the whole, BUA-based reconstruction is a less accurate one. In Fig. 2.15 (right) we plot the correlation of the reconstructed and true emissions in the NE and NW quadrants. We see that the correlations for the nightlight- and BUA-based estimated emissions are somewhat similar, though BUA-based estimation is more variable, and has lower correlation as a whole. Thus, while Fig 2.3 (middle row) showed that BUA-had a slightly better correlation with true (Vulcan) emissions, its larger errors, as seen in Fig. 2.3 (bottom row) lead to a less accurate reconstruction. This result is also a testament to the inadequacy of $\mathbf{y}^{obs}$ over the whole country; had there been sufficient data to constrain $\mathbf{E}$, the impact of $\mathbf{f}_{pr}$ would have been minimal.

In Fig. 2.16 we compare the estimated emissions developed from the two competing prior models. In the top row we plot the estimated emissions for the 34th 8-day period using the nightlight prior (left) and built-up area maps (right). The difference between the two estimates, plotted in the bottom left subfigure, shows differences spread over a large area, though their magnitudes are not very big. Thus the “prior” model has a measurable impact on the spatial distribution of the emissions. Bottom right, we plot the degree to which the measurements update/change the prior model. On the horizontal axis, we plot the emissions predicted by the “prior” model, while the vertical axis represents the estimated emission in the corresponding grid-box. We see some correlation between the “prior” and estimated emission when the emissions are large (more than $0.25 \mu\text{mol m}^{-2} \text{s}^{-1}$ of C).

2.5.5 Impact of M_{cs}

In Sec. 2.4.2 we had used M_{cs} random projections of $\mathbf{F}_{\mathcal{R}'}$ to implement a $\mathbf{F}_{\mathcal{R}'} = 0$ constraint efficiently. Since Eq. 2.11 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 and $\mathbf{F}_{\mathcal{R}'}$ is usually small. This error varies with M_{cs} ; a larger number of random projections result in a closer realization of the constraint. However, they are never driven to zero, primarily because the wavelets used to model $\mathbf{F}_{\mathcal{R}'}$ and $\mathbf{F}_{\mathcal{R}}$ were chosen using $\mathbf{X}$ (and thus may not form a complete basis set for ffCO₂ emissions). Errors in the enforcement of the $\mathbf{F}_{\mathcal{R}'}$ constraint lead to commensurate errors in $\mathbf{F}_{\mathcal{R}}$. In practice, this affects only Step I of the emission estimation procedure, where a sparse approximation of $\mathbf{F}_{\mathcal{R}}$ is calculated; thereafter it is used as a guess in Step II, the enforcement of non-negativity of emissions. However, a good estimate of $\mathbf{F}_{\mathcal{R}}$ accelerates Step II.

In Fig. 2.17, we plot the impact of M_{cs} on the reconstruction. We perform Step I of the emission estimation

Figure 2.14: Left: Plot of the error in the reconstruction, as performed with built-up area maps (blue line) and nightlights (black line) as the proxies. Right: Correlation between the true and reconstructed emissions, as performed with built-up area maps and nightlights as the proxies. The nightlights-based reconstruction has slightly less reconstruction error and better correlation with the true emissions, when compared at 8-day temporal resolution. When averaged to a 32-day temporal resolution, emissions obtained via the two methods are very similar.

Figure 2.15: Left: Emission reconstruction error in the NE (blue) and NW (black) quadrants, when performed with BUA (line) and nightlights (symbols) as proxies. We see that the NW quadrant is very badly constrained and a change in proxies changes the sign of the error. This is also seen in the NE quadrant; however, the error magnitudes are far smaller. Right: The comparison of correlations between true and reconstructed emissions shows similar trends; nightlights-based estimation produces better reconstructions.

Figure 2.16: Comparison of emission estimates developed using $\mathbf{f}_{pr}$ constructed from nightlight radiances and built-up area maps. Top: Estimated ffCO_2 emissions for the 34th 8-day period developed using nightlight prior (left) and the prior from built-up areas maps (right). Bottom left: we plot the difference between the two estimates. Bottom right: We plot a scatter plot between the estimated and prior emissions, for the two prior models.

Figure 2.17: The impact of the number of compressive samples M_{cs} on the reconstruction of $\mathbf{F}_{\mathcal{R}}$ ($\eta_{\mathcal{R}}$) and $\mathbf{F}_{\mathcal{R}'}$ ($\eta_{\mathcal{R}'}$). $\eta_{\mathcal{R}}$ 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 $\eta_{\mathcal{R}'}$.

procedure (see Sec. 2.4), for the 31st and 32nd 8-day periods and compute the ratios

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

Here $\mathbf{f}_{k,\mathcal{R}}$ and $\mathbf{f}_{k,\mathcal{R}'}$ are the emissions (from Step I) over $\mathcal{R}$ and $\mathcal{R}'$. $\mathbf{f}_{V,k}$ is the true (Vulcan) emission field during the same period. These ratios are plotted as a function of the log-transformed number of compressive samples 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 period, $\eta_{\mathcal{R}'}$ oscillates around 0.1. The corresponding errors in $\mathbf{f}_{k,\mathcal{R}}$ are about 5% ($\eta_{\mathcal{R}} \approx 1.05$). In our study we used 300 random projections for each 8-day period. In contrast, had we imposed $\mathbf{F}_{\mathcal{R}'} = 0$ in all grid cells in $\mathcal{R}'$, we would have generated over 3000 constraints per 8-day period (our 64×64 mesh has 4096 grid cells, of which 816 are in $\mathcal{R}$ and the rest, 3184 are in $\mathcal{R}'$). This economy of computational effort in the imposition of the constraint is only partially due to the efficiency of random projections; a major simplification is achieved by the lower-dimensional model of $\mathbf{F}_{\mathcal{R}}$ using the wavelets identified by $\mathbf{X}$.

2.5.6 Impact of ϵ

In this section we study the impact of the measurement error ϵ . The nominal value used in this study is $\epsilon = 10^{-2}$; we explore the impact of $\epsilon = \{2.5, 5.0, 7.5, 10\} \times 10^{-2}$. The results, for the 31st 8-day period are shown in Fig. 2.18. Top left, we plot the true emissions from the Vulcan inventory for reference. Top right, we plot results using $\epsilon = 2.5 \times 10^{-2}$. We see enhanced “blockiness” as some of the finer wavelets cannot be estimated; whereas the true emissions are represented using 363 wavelets coefficients with a magnitude above 0.01, the reconstruction with $\epsilon = 2.5 \times 10^{-2}$ recovers only 325 such coefficients. Bottom left, at a higher level of measurement errors ($\epsilon = 10^{-1}$), the solution seems qualitatively different; low emission regions in the Western quadrants (Continental Divide and the deserts of the Southwest) show erroneous (and higher) levels of emissions. This leads to the recovery for 377 wavelet coefficients with a magnitude above

0.01. Bottom right: we summarize the quality of the reconstruction via the reconstruction error (%) and the correlation between the true and estimated emissions. We see that the degradation of reconstruction quality is roughly linear in the measurement error ϵ .

2.5.7 Impact of the number of observation towers

In this section, we investigate the impact of reducing the number of observation towers. In Fig. 2.19 we plot the results from inversions performed with 35-, 25- and 15-tower configurations. In the left column, we plot reconstructed emissions during the 33rd 8-day period, with the different tower configurations. The resolution at which the emissions can be estimated decreases with the number of towers; the differences between the 35- and 15-tower reconstructions are easy to detect. In the right column, we quantify the differences. In Fig. 2.19, top right, we plot the reconstruction error over 45 8-day periods. We see that while the difference in the reconstruction error between 35- and 25-tower inversions is not much, the 15-tower inversion is clearly inferior. This is also borne out in the plot of the correlation of the true and estimated emissions (Fig. 2.19, right middle), where the estimation performed with 15 towers provides the lower bound. Finally, in Fig. 2.19, bottom right, we plot the wavelet coefficients. The coefficients have been subjected to a hyperbolic tangent transformation for plotting clarity. The true wavelet coefficients with a magnitude above 0.01 are plotted with red crosses. The 35-tower reconstruction comes close to estimating the true coefficients; this is especially true for the wavelets with larger spatial support (i.e., low index). As the number of towers drop, the errors in the estimated wavelet values grow. Further, the number of wavelet coefficients with large magnitudes (above 0.01) steadily decreases. Thus while the true emissions are represented by 363 wavelet coefficients with magnitudes above 0.01, the estimated emissions, using 35, 25 and 15 towers have, respectively, 315, 292 and 287 coefficients. Therefore, in the absence of observations, the solution to Eq. 2.11 sets the wavelet coefficients to values near zero.

2.6 Conclusions

We have devised a multiresolution parametrization (also known as a multiscale random field or MsRF model) for modeling ffCO₂ emissions at 1° resolution. It is based on Haar wavelets, and is designed for use in atmospheric inversions. It uses easily observable proxies of human activity e.g., images of lights at night and maps of built-up areas to reduce its dimensionality as well as to build “prior” models of ffCO₂ emissions.

The MsRF model was tested in a set of synthetic-data inversions. Time-dependent ffCO₂ emission fields were estimated over the lower 48 states of the US, conditioned on 360 days of observations of time-varying ffCO₂ concentrations at a set of observation towers. In conjunction with an atmospheric transport model, a sparsity-enforcing optimization method, Stagewise Orthogonal Matching Pursuit (StOMP), was used to fit MsRF models and reconstruct a time-series of emission fields. It was found that less than half the parameters of the MsRF model could be estimated from the sparse data; the rest were identified by StOMP and set to zero. We also identified an efficient way of incorporating a prior model of emissions into the inversion.

The MsRF, being based on wavelets, models emissions on a dyadic square grid, whereas ffCO₂ emissions were restricted to an irregular region $\mathcal{R}$ (the lower 48 states of the US). Instead of specifying zero emissions in the grid-cells outside $\mathcal{R}$, we used concepts from compressive sensing (random projections) to achieve the same effect, but at a tenth of the computational effort and memory requirements (Sec. 2.5.5). The

Figure 2.18: Impact of measurement error ϵ . Top left: We plot the true ffCO₂ emissions from the Vulcan inventory. Top right: We plot the estimates calculated using $\epsilon = 2.5 \times 10^{-2}$. Bottom left, we plot the reconstructed emissions using $\epsilon = 10^{-1}$; we see a clear degradation of the reconstruction. Bottom right: We plot the reconstruction error (%) and the correlation between the reconstructed and true emissions for various values of ϵ ; a clear degradation is seen. Reconstruction errors and correlations are plotted on opposing Y-axes. All results are for the 31st 8-day period.

Figure 2.19: Left column: Estimated ffCO₂ emissions for the 33rd 8-day period. Inversions were done with 35 (top), 25 (middle) and 15 (bottom) towers. As the number of towers decreases, we see that the resolution of estimated emissions decreases. Right column: At the top, we plot the error in the reconstructed emissions for the three tower sets. There is not much difference between reconstructions with 25 and 35 towers, but the 15-tower reconstruction is poor. Right (middle): The correlation between reconstructed and true emissions shows much the same trend as the reconstruction, with the 15-tower reconstruction having a far lower correlation over the entire 360-day duration. Right column, bottom: We plot the sparsity of the estimated wavelet coefficients. We see that as the number of towers decrease, the wavelet coefficient estimates deviate further from the true values (red crosses).

emission fields reconstructed using the MsRF were found to contain some regions (grid-cells) with negative emissions. Using these as a starting guess, we devised a simple method to iteratively impose non-negativity on the ffCO_2 emissions.

Primarily, our work demonstrated that observations of ffCO_2 concentrations could be used to update models of ffCO_2 emissions and improve their accuracy (see Fig. 2.10). Further, we assumed that these measurements could be obtained from existing towers, which were sited with a view of reconstructing biospheric, rather than ffCO_2 , fluxes. This accounts for some of our inability to estimate fine-scale spatial structure in the emissions. Note that currently the network has expanded beyond the 35 towers included in this work, and could potentially furnish better estimates of ffCO_2 than reported here. Future work will examine how the inversion method described here could be used to design a monitoring network for ffCO_2 , rather than biospheric CO_2 fluxes.

The MsRF models constructed using images of nightlights and built-up area maps provided reconstructions that differed in their fine-scale details (at the quadrant scale and below). This is a reflection of the paucity of observations and their inability to constrain the fine-scale features of the emission fields. The formulation also correctly showed that as the number of observation towers were decreased, larger numbers of MsRF parameters were set to zero and the reconstructed emission fields progressively lost their spatial fidelity, starting from the fine-scale details.

Our inversion formulation suffers from two drawbacks. It requires measurements of ffCO_2 concentrations at the measurement towers. While these concentrations can be back-calculated using $\Delta^{14}\text{CO}_2$ or CO measurements, they are less common and more expensive than CO_2 concentration measurements. The second drawback is the deterministic nature of the reconstruction - we do not provide error bounds on the estimates of the MsRF parameters (the wavelet coefficients). This can be rectified by adopting a Bayesian approach e.g., Kalman filters, but it is unclear how one would preserve the non-negative property of ffCO_2 emissions. This investigation is currently underway.

Chapter 3

Estimation of ffCO₂ emissions using ensemble Kalman filters

3.1 Introduction

Many dynamical systems in engineering and science are being instrumented with sensors to help calibrate and understand uncertainties associated with numerical predictions. For instance, weather and climate prediction rely on a large range of detailed temporal and spatial observations, from stationary pressure and velocity readings to satellite images. All these data measurements and observations are collected for the specific purpose of calibrating numerical models so that more accurate forecasts can be issued. At the heart of that problem lies an inverse problem where the goal is to reconcile the differences between data and numerical predictions subject to complex dynamics.

Our application focus is on climate science where measurements are used to help characterize global warming. In particular, atmospheric scientists are measuring the concentration of many trace gases, with CO₂ being the most important one, to evaluate the affect of rising temperatures. Emissions from fossil fuel based sources have caused rapid increases in CO₂ concentration in our atmosphere and have been linked to significant changes to global climate patterns. In an attempt to mitigate the increase in CO₂ levels, a treaty verification process may eventually be deployed to help manage global inventories and among a range of technical mechanisms, a measurement-based inversion process could prove to be instrumental in supporting such an effort [5]. In contrast to a bottom-up approach [4, 84, 6], where sources are estimated using socio-economic data, a measurement process minimizes the difference between measurements and numerical predictions by solving an optimization problem in an attempt to reconstruct the source signatures at some point in time.

Inverting for fossil fuel sources at different spatial scales is a complex endeavor plagued by modeling approximations, data errors, complex source signatures, large computational domains, multi-physics, and limited measurements. The character of fossil fuel emissions sources, in particular, has complicated features that pose many inversion challenges. In addition to temporal variations from diurnal and seasonal cycles, source signatures have a variety of character and spatial distribution at different geographical scales. A view at the country (e.g., US) spatial scale is considerably different from regional scales in which cities, powerplants, and roadways can be distinguished. Finally, the biospheric signal can overwhelm the fossil fuel signal depending on location and time; the estimated fluxes reflect the smooth spatial distributions characteristic of biospheric land-atmosphere exchanges. Previous studies show that the effect of ffCO₂ emissions can only be detected during the dormant (winter) season [2].

In this chapter, we investigate ensemble Kalman Filter (EnKF) algorithms for multi-physics dynamics in

which different streams of data are available to help calibrate the numerical models. Our goal is to invert for source terms using sparse data from different dynamics to eventually provide accurate forecasting capabilities. In our target inversion problem, the source terms are spread throughout the computational domain and as such the solution approach must manage many inversion parameters. In addition, considerable uncertainty is associated with the transport model and the measurements, which an inversion process must be able to account for. EnKF is selected as our inversion algorithm to address the uncertainty issues, ensure scalability to high-dimensional inversions and to allow for non-linearities in the model linking sources and observations. Furthermore, our intent is to evaluate the statistical quality of the EnKF algorithm for this problem by comparing with Markov Chain Monte Carlo methods.

The ensemble Kalman filter, first proposed by Evensen in [85] for meteorological applications, uses an ensemble of states to represent the probability distribution of the model state; when observations are introduced, each member is updated using only the mean and covariance of the predicted distribution. The method has become popular in the meteorological and hydrological communities, as well as for oil reservoir simulation, as an alternative to adjoint or four-dimensional variational (4DVAR) methods for handling large state spaces and error evolution through nonlinear models. EnKF requires no derivation of adjoint equations or tangent linear operators, and is relatively simple to implement [86].

EnKF have primarily been employed in the estimation of biospheric CO₂ fluxes. Since CO₂ concentration observations are a consequence of emissions in the recent past, the data assimilation procedure is a smoothing operation. In [87] Bruhwiler *et al.* developed a Kalman smoother to reconstruct CO₂ fluxes over a 5 year duration at monthly resolution. The formulation was a conventional Bayesian one with informative priors. A version with less informative priors was developed by Michalak [88]. Scalable approaches based on ensemble methods have been constructed. In [43] Peters *et al.* developed an Ensemble Square Root Filter (EnSRF) to estimate CO₂ fluxes at a weekly resolution over a year. They introduced an exponentially decaying correlation function model for the forecast covariance. Tests with synthetic data simulating observations at the NOAA/CMDL sites and a few flights were performed with 200–1500 ensemble members; runs generally took 50–100 processors. In [23], Feng *et al.* constructed an ensemble transform Kalman filter for the assimilation of (synthetic) satellite retrievals. The method was used to infer regional fluxes (at 1000km × 1000km resolution and 8-day periods. The state vector spanned 96 days, and observations over a 7-month period was assimilated. A related construction, based on a Maximum Likelihood Ensemble Filter (MLEF) is described in [89], is more general than the ensemble techniques described, being able to address non-Gaussian prior distribution.

Current methods for estimating atmosphere-land exchanges e.g., estimating biospheric CO₂ fluxes [9] use an atmospheric transport model which uses wind fields as an input. These wind fields are generally re-analysis products, constructed by assimilating measurements into meteorological models. The uncertainties in the re-analysis products, though calculated during the data assimilation process, are not used in the atmospheric transport models; when used in estimating atmosphere-land exchanges, this leads, at the very least, to an underestimation of the uncertainty in flux estimates. Further, EnKF assume that all distributions are Gaussians, which within the context of nonlinear models, is not correct. In this chapter, we explore the two methodological issues associated with the current approach for estimating fluxes: (1) the underestimation of flux uncertainties due to the convention of using a deterministic wind field and ignoring the uncertainties in it and (2) the impact of a Gaussian assumption on the second-order moments of the flux distributions calculated by an EnKF-based inversion. We will address the first by performing a joint inversion for the velocity field and ffCO₂ sources and comparing with the conventional approach. The second issue will be studied by performing the inversion using a MCMC method.

The remainder of this paper will describe the mathematical formulation of our inversion techniques, present

a simple parameterization, and demonstrate the effectiveness of our simultaneous inversion approach using the EnKF algorithm, in addition to comparing the statistics to MCMC generated values.

3.2 Mathematical Formulation

In this section, we derive the EnKF algorithm from Bayes Theorem. It will provide a mathematical exposition from which we can identify key aspects of the algorithm and draw connections to other methods. In particular, we evaluate the inversion results from EnKF by sampling the posterior density with a Markov Chain Monte Carlo (MCMC) method. Furthermore, the quality of simultaneous inversions is compared to explicit inversion streams where a mean solution from one inversion is used in the forward simulation of the second inversion process.

3.2.1 Bayesian Inference Review

Bayesian inference is a flexible and statistically rigorous way to model parameters and states. The inference problem is usually formulated to statistically estimate parameters or inputs to a model, based on noisy observations of the model output. The process of propagating error through a model is referred to as the forward problem and can be defined as

$$d(x, t) \approx G(x, t; m) \quad (3.1)$$

where m is a set of model parameters, d is a vector of observable output quantities, x is position, t is time, and G is potentially a nonlinear function of x and t , parameterized by m . The model parameters m are unknown and need to be estimated from noisy measurements of d . To handle the observation uncertainty, both d and m will be modeled as random variables in a proper probability space.

To estimate the parameter, we invoke Bayes' rule to condition the model parameters on d :

$$f(m|d) = \frac{f(d|m)f(m)}{\int f(d|m)f(m)dm} \quad (3.2)$$

where $f(m)$ is the prior density of m , which represents previously held degrees of belief about m . This information could come from physical meanings of the model parameters or simply expert opinion. In a Bayesian setting, as opposed to the frequentist paradigm, probabilities represent degrees of belief that the model parameters will take particular values. An informative prior acts as a regularization term as well, allowing for the solution of ill posed and underdetermined problems. While the prior distribution represents previously held beliefs about model parameters, the likelihood function, $f(d|m)$ measures how well a particular set of model parameters match the available measurements. By using both the prior distribution and likelihood function, Bayes rule updates previously held beliefs with the current measurements in d . In a static problem, only one update will be used because all information is presented at once. However, Bayes rule can also be used iteratively to update the densities as each new measurement arrives. Thus, the posterior at time t becomes the prior for an update at time $t + \delta t$.

Simplifying Bayes rule by disregarding the normalization constant gives:

$$f(m|d) \propto f(d|m)f(m) \quad (3.3)$$

A major challenge of using Bayes' rule for inferences is sampling from the posterior density $p(m|d)$. The posterior evaluations require a computationally expensive forward simulation to compute the likelihood function, making methods such as Markov chain Monte Carlo (MCMC, [90]) methods difficult to apply. MCMC methods use a random walk to construct a Markov chain with a stationary distribution equal to $f(m|d)$. These methods are guaranteed to converge to $f(m|d)$ in the limit of infinite samples and can therefore capture highly non-Gaussian and multimodal distributions. However, constraints on computation time often do not allow users to run MCMC methods for many samples. Significant research has been devoted to variants of the MCMC algorithms in an attempt to accelerate the convergence [91, 92, 93]. However, the sequential nature of this random walk algorithm poses computational challenges that currently prevents consideration of inverting large numbers of parameters constrained by complex dynamics.

In Bayes rule, the likelihood function is equivalent to a probabilistic description of the measurement errors. If Gaussian noise is assumed as in (also assumes i.i.d.):

$$d = G + \eta \quad (3.4)$$

where $\eta \sim N(0, \sigma^2)$, the likelihood function can be represented as a Gaussian function:

$$f(d|m) \sim \prod_i \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(G_i - d_i)^2}{2\sigma^2}\right) \quad (3.5)$$

and provides one of the components of the mathematical foundation to consider less statistically robust but computationally more efficient methods such as the Kalman filter [94, 95]. This algorithm recursively averages, in a weighted fashion, a prediction of a system's state with new measurements and represents the uncertainty associated with the predictions of the systems state. Gaussianity is assumed and the dynamics must be linear. To address the nonlinearities in G , the extended Kalman Filter (EKF, [96]) makes use of Taylor's series to linearize the system. Both algorithms however require an adjoint of the system's Jacobian to calculate the covariance. This can be difficult to implement and simply not tractable in complex simulators. The Ensemble Kalman Filter, on the other hand, does not depend on any linearizations or adjoint calculations. In addition it can capture some nonlinear behavior in the model G because the covariance is based on post-processing the ensemble members which depend entirely on the dynamics. The following sections provides a mathematical formulation starting from Bayes Theory to help compare solution strategies using MCMC and EnKF solvers.

3.2.2 Comparison Formulation

MCMC and EnKF handle data in fundamentally different ways, necessitating an appropriate inference problem that can be formulated as a dynamic data assimilation problem for the EnKF and a batch assimilation (static) inference problem for the MCMC. Consider the general dynamical system

$$\frac{\partial s(x, t)}{\partial t} = G(s(x, t), m(x)) \quad (3.6)$$

where $s(x, t)$ is the system state, G is a potentially nonlinear model operator parameterized by the model parameters m . Although appropriate boundary and initial conditions also exist, these are not explicitly stated for conciseness. In addition to the dynamical system, measurements of the system are available through some linear functional M :

$$M[s, m] = d + \epsilon \quad (3.7)$$

where ε represents additive measurement noise. The goal here is to obtain probabilistic estimates of the model parameters $m(x)$ using the observations of the measurements $d + \varepsilon$. Formulating this inference problem in a Bayesian framework gives the posterior density:

$$f(m|d) \propto f(m)f(d|m) \quad (3.8)$$

(Note that spatial and temporal dependencies have been dropped here to simplify notation.) The EnKF is generally used to estimate the system state in addition to the model parameters, according to

$$f(s, m|d) \propto f(s, m)f(d|s, m) \quad (3.9)$$

However, attempting to estimate the state and model parameters simultaneously is intractable for MCMC on most applications due to large state dimensions on fine resolution grids. Thus, the formulation in Eq. 3.8 will be used for comparison.

The measurements in d will generally become available at several discrete times $\{t_0, t_1, \dots, t_k\}$, $k > 0$. Denote these measurements as $d = \{d_0, d_1, \dots, d_k\}$. MCMC will handle all of these observations at once with batch processing, and will use Eq. 3.8 directly. The EnKF on the other hand, will assimilate the data sequentially, updating the estimate of m after each observation. This requires the likelihoods and posterior density in Eq. 3.8 to be expanded in time.

For both MCMC and the EnKF, measurements are assumed to be uncorrelated in time, and model integration is assumed Markov, implying that given the state and model parameters at time t_i , the model state at $t_i + \delta t$ is completely determined. No information from $t < t_i$ is necessary. With this assumption, applying Bayes' rule to all available data gives the joint posterior density for the model parameters:

$$f(m|d) \propto f(m) \prod_{j=1}^k f(d_j|m) \quad (3.10)$$

3.2.3 Sequential Processing

The static description in Eq. 3.10 gives the joint posterior density after all data have been accumulated. This expression can be used in a MCMC sampling of the posterior; however, the additional sequential structure of the posterior can also be exploited to sequentially update the posterior as in EnKF. At the first time step, the intermediate density will be

$$f(m|d_1) \propto f(m)f(d_1|m) \quad (3.11)$$

and then at the second time step,

$$f(m|d_1, d_2) \propto f(m|d_1)f(d_2|m) \quad (3.12)$$

In general, the intermediate density for the first j measurements is given by

$$f(m|d_1, d_2, \dots, d_j) \propto f(m|d_{j-1})f(d_j|m) \quad (3.13)$$

This recursive update of the posterior can continue until all observations have been processed, resulting in the posterior described by Eq. 3.10. The EnKF uses this formulation with a variance minimizing update of ensemble locations to efficiently approximate the full posterior.

Figure 3.1: Comparison of MCMC and EnKF approach to estimating the posterior density in Eq. 3.10. Note the potential for parallelism in the EnKF.

3.2.4 Kalman Filter

The EnKF approach depends on fundamental components from the original Kalman Filter algorithm which is described next. Consider a simple case where the operator G is linear and both the error model and prior distribution are Gaussian. The standard Kalman filter fully captures the posterior statistics, which consequently will also be Gaussian. After temporal discretization, this linear system becomes:

$$s_k^t = Gs_{k-1}^t + q_{k-1} \quad (3.14)$$

where k is the index in time, the superscript t denotes that this is the true system state, and q is the error in the model dynamics. First, consider a time-step when no observations are present. In this case, the model dynamics provide the best estimate of the state at the next time step. The prediction, $s_k^f(t_k)$, is then given by:

$$s_k^f = Gs_{k-1}^a \quad (3.15)$$

where s^a is the best estimate from the previous time-step. For time steps without observations, $s^f = s^a = E[s]$. Thus, the Kalman filter simply integrates the mean in time to obtain the predictions. The covariance can be propagated forward by first subtracting the best numerical estimate from the unknown true field:

$$\begin{aligned} s_k^t - s_k^f &= Gs_{k-1}^t + q_{k-1} - Gs_{k-1}^a \\ &= G(s_{k-1}^t - s_{k-1}^a) + q_{k-1} \end{aligned} \quad (3.16)$$

This is the model forecast error at time t_k when no new observations have occurred. Now, the model error covariance becomes

$$C_{ss}^f = \overline{(s_k^t - s_k^f)(s_k^t - s_k^f)^T} = GC_{ss}^a(t_{k-1})G^T + C_{qq}(t_{k-1}) \quad (3.17)$$

where the model state has error covariance

$$C_{ss}^a(t_{k-1}) = \overline{(s_{k-1}^t - s_{k-1}^a)(s_{k-1}^t - s_{k-1}^a)^T}$$

and model error covariance is

$$C_{qq}(t_{k-1}) = \overline{q_{k-1} q_{k-1}^T}$$

In most cases, the model G is not linear and variants of the standard Kalman filter need to be used. A popular alternative is the extended Kalman filter (EKF, [96]) which simply linearizes the nonlinear operator with a Taylor expansion around the mean, s_{k-1}^a .

The best estimate from the EKF is also computed by integrating the mean forward in time,

$$s_k^f = G(s_{k-1}^a) \quad (3.18)$$

Not surprisingly, nonlinearities in G can cause this estimate to be biased, as thoroughly discussed in [97]. The EKF covariance update step follows the usual approach when extending linear methods to nonlinear cases and uses a truncated Taylor expansion of G , to arrive at the approximate covariance update:

$$C_{ss}^f \approx G'_{k-1} C_{ss}^a(t_{k-1}) G_{k-1}'^T + C_{qq}(t_{k-1}) \quad (3.19)$$

where G'_{k-1} is the Jacobian matrix at the previous time step.

The Kalman filter and EKF updates shown above are used when no data are available. When data are available, the sequential Bayes' formulation in Eq. 3.12 is used in addition to the model dynamics to update the mean and covariance. The “best” estimate in this case is the *maximum a posteriori* (MAP) estimate, obtained by maximizing the log posterior distribution in Eq. 3.12.

With a Gaussian error model and a continuous state, Bayes' rule gives the posterior distribution of the model parameters and system state when observations are present:

$$\begin{aligned} f(s, m|d) &\propto f(s|m) f(d|s, m) \\ &= \exp \left[-\frac{1}{2} \iint (s^f(x_1) - s(x_1)) W_{ss}^f(x_1, x_2) (s^f(x_2) - s(x_2)) dx_1 dx_2 \right. \\ &\quad \left. - \frac{1}{2} (d - M[s]) W_{\epsilon\epsilon} (d - M[s]) \right] \end{aligned} \quad (3.20)$$

where M is a linear measurement functional relating the field s to d , W_{ss} is the functional inverse of the covariance function, and $W_{\epsilon\epsilon} = C_{\epsilon\epsilon}^{-1}$ is the inverse of the error covariance. Here, s^f is the system state integrated forward from the previous time-step as if no observations were present. The first term in the exponential is the Gaussian prior and the second term is the likelihood. In a discretized system where the likelihood is a multivariate Gaussian distribution and not a Gaussian process, the integrals would be unnecessary and $W_{ss}(x_1, x_2)$ would be replaced by the matrix $W_{ss} = C_{ss}^{-1}$. The integration simply replaces the summation that occurs when pre and post multiplying W_{ss} by vectors of discretized states. From this posterior distribution, it is clear that when data are available, the best estimate of s occurs at the maximum of this density. This is the MAP estimate for s .

To find the MAP estimate, it is generally easier to minimize the negative log posterior

$$\begin{aligned} -2 \log[f(\psi, m|d)] &= \iint (\psi^f(x_1) - \psi(x_1)) W_{\psi\psi}^f(x_1, x_2) (\psi^f(x_2) - \psi(x_2)) dx_1 dx_2 \\ &\quad + (d - M[\psi]) W_{\epsilon\epsilon} (d - M[\psi]) \end{aligned} \quad (3.21)$$

than to operate on the posterior directly. In [98] a variational derivative and the Euler-Lagrange equation approach was used to minimize this functional and find the data assimilation update step. The result is:

$$s^a = s^f + K (d - Ms^f) \quad (3.22)$$

where the Kalman gain matrix K is defined as:

$$K = C_{ss}^f M^T (M C_{ss}^f M^T + C_{\epsilon\epsilon})^{-1} \quad (3.23)$$

where $C_{\epsilon\epsilon}$ is the covariance of the error term ϵ and C_{ss}^f is:

$$C_{ss}^f = G C_{ss}^a(t_{k-1}) G^T + C_{qq}(t_{k-1}). \quad (3.24)$$

Note that s^f and s^a are two estimates for the same time step, the former is the estimate before data assimilation and the latter is the analyzed estimate that includes observations. It is the covariance of the model error that differentiates the EnKF from the standard Kalman Filter, which requires the transpose (or the adjoint operator) of G and is not always easily implemented, especially in parallel code. The EnKF approach calculates this covariance by post-processing a collection (the ensemble) of forward runs per time-step.

3.2.5 Ensemble Kalman Filter

In addition to the adjoint implementation and the errors that arise when using a truncated Taylor expansion as an approximation to a nonlinear function, the standard Kalman filter algorithms also have to store a potentially large covariance matrix. Many realistic problems, especially in geophysics, have millions of degrees of freedom, making it impossible to store the covariance matrix required by the KF and EKF. The ensemble Kalman filter (EnKF) resolves two of the major issues with the KF and EKF – it does not need to store the covariance matrix and no linearization of the model operator is required. Instead, the EnKF uses ensemble averaging to compute the mean and covariance after each update. Randomly instantiated ensemble members are propagated between observation times with the deterministic model G . Using the ensemble members as Monte Carlo samples of the model parameters, m , and system states, s , the ensemble mean and covariance are used instead of propagating the mean and covariance directly. Data assimilation is similar to the Kalman filter, except each of the ensemble members is updated individual by replacing the mean and covariance in Eq. 3.22 with ensemble equivalents. The ensemble mean is chosen as the best estimate although it may be biased. Thus,

$$s_k^f = \bar{s}_k = \frac{1}{N_e} \sum_{j=1}^{N_e} s_{j,k} = \frac{1}{N_e} \sum_{j=1}^{N_e} G(s_j, k-1) \quad (3.25)$$

where the subscript s_j denotes the j^{th} ensemble member, and N_e is the ensemble size. Now, defining the ensemble error covariances as:

$$\begin{aligned} (C_{ss}^e)^f &= \overline{(s^f - \bar{s}^f)(s^f - \bar{s}^f)^T} \\ (C_{ss}^e)^a &= \overline{(s^a - \bar{s}^a)(s^a - \bar{s}^a)^T} \end{aligned} \quad (3.26)$$

where the ensemble average defined by $\bar{\cdot}$ is defined in Eq. 3.25. To incorporate observations, ensemble members are individually updated according to the KF assimilation step:

$$s_j^a = s_j^f + (C_{ss}^e)^f M^T \left(M (C_{ss}^e)^f M^T + C_{\epsilon\epsilon} \right)^{-1} (d - Ms_j^f). \quad (3.27)$$

By using the mean and covariances in this way, the ensemble Kalman filter assumes the system state and parameters have a Gaussian PDF; however, ensemble members are individually propagated through the model G and can therefore capture some of the model nonlinearity. By using the KF update during assimilation, the EnKF maximizes a Gaussian posterior for the update step. For a highly nonlinear G , this may not be a sufficient approximation. In problems with bimodal or highly skewed distribution, EnKF may not be able to capture all nonlinearities and non-Gaussian statistics.

3.3 Overview of the Physics and Numerical Setup

In this section, we describe the physical problem, and models which will be used in our exploration of multiphysics systems, the handling of different data streams and the comparison of EnKF and MCMC. We also discuss practical issues of estimating ffCO₂ sources e.g., imposing non-negativity in concentration and source-magnitude estimates and developing a low-dimensional representation of the target inversion parameters.

3.3.1 Problem setup

We consider a square domain with dimensions $L \times L$ containing the US, Canada and Mexico (Fig. 3.2). Our numerical experiments use synthetic data which are generated by driving a convection-diffusion system (Eq. 3.29) with source terms f that in turn are generated from the Vulcan inventory [3, 4]).

$$\frac{\partial c}{\partial t} - D\Delta c + v\nabla c = f \quad \in \Omega \quad (3.28)$$

$$c = 0 \quad \text{on } \Gamma \quad (3.29)$$

where c is concentration, D represents the diffusion coefficient, and v is velocity. The computational and boundaries are represented by Ω and Γ , respectively. The Vulcan map represents anthropogenic CO₂ sources calculated from socio-economic data and we assume this to be the “truth” model. We introduce a simple Gaussian kernel (GK) parameterization for these source terms and magnitudes of each kernel are our inversion targets. This parameterization reduces the variables being inferred far below the grid resolution. The Vulcan-based image is thresholded to provide a binary image discriminating regions with significant night lights versus background. For each contiguous region with significant night lights, a single ellipse is fit to the pixels exceeding the threshold using the Matlab `regionprops` function (in the Image Processing Toolbox). The orientation (θ), major and minor radii (r_1 and r_2) and the centroid coordinates of the ellipse are recorded. The centroid coordinates define the mean of the Gaussian function, $\mu = (x_0, y_0)$

The geometric description of the ellipse is used to define the covariance matrix, Σ , of the GK where the major and minor radii define the ± 2 standard deviations of the kernel. This conversion is done through

calculation of geometric relationships with the inverse of the covariance matrix:

$$\Sigma_{xx}^{-1} = \frac{\cos(\theta)^2}{r_1^2} + \frac{\sin(\theta)^2}{r_2^2}, \quad (3.30)$$

$$\Sigma_{yy}^{-1} = \frac{\sin(\theta)^2}{r_1^2} + \frac{\cos(\theta)^2}{r_2^2}, \quad (3.31)$$

$$\Sigma_{xy}^{-1} = \sin(\theta)\cos(\theta) \left(\frac{1}{r_1^2} - \frac{1}{r_2^2} \right). \quad (3.32)$$

Each GK is calculated as a bivariate Gaussian function:

$$f(x, \mu, \Sigma) = \frac{1}{\sqrt{(2\pi)^2 |\Sigma|}} \exp \left(-\frac{1}{2} (x - \mu) \Sigma^{-1} (x - \mu)' \right) \quad (3.33)$$

and ranked using the integrated light intensity weighted by the local values of the Gaussian function. The GK parameterization is also applied to a satellite image of North America at night [33, 69], which is used as the initial guess in our inversion process.

A velocity field is set up by imposing a horizontal velocity $u_l(x)$ on the top boundary (henceforth the lid velocity), whereas the rest of the walls of $\mathcal{R}^*$ act as no-slip boundaries. This induces a vortical flow inside $\mathcal{R}^*$, which is modeled using the unsteady, 2D Navier-Stokes equations (Eq. 3.34):

$$\begin{aligned} \frac{\partial v}{\partial t} + v \cdot \nabla v &= -\nabla p + \mu \nabla^2 v + f \in \Omega \\ \nabla \cdot v &= 0 \in \Omega \\ v &= 0 \text{ on } \Gamma_1 \\ v &= \zeta \text{ on } \Gamma_2 \end{aligned} \quad (3.34)$$

where p is pressure, μ is viscosity, f the body force (set to zero for our tests), and ζ is a prescribed velocity on a boundary. The equation is solved on a square grid with a resolution (per side) ranging between 60 and 400. For the inversion tests, we generated synthetic observations by simulating at a finer resolution (600 grid cells per side) and collect sparse set of measurements on a square grid of sensors. I. I. D Gaussian noise is then added to the data. The magnitude of the noise depends on the variable concerned (concentration versus velocity), but the standard deviation of the noise distribution varies between 10^{-3} to 10^{-2} .

3.3.2 Inverse modeling and dimensionality reduction

In the EnKF inversion, the state vector includes both dynamic states and static parameters as specified in Eq. 3.35.

$$\Psi = \begin{pmatrix} u, & \text{horizontal velocity at vertical cell interfaces} \\ v, & \text{vertical velocity at horizontal cell interfaces} \\ c, & \text{concentrations cell centers} \\ b, & \text{horizontal velocities at lid} \\ f, & \text{sources (amplitudes)} \end{pmatrix} \quad (3.35)$$

For the initial ensemble, velocity fields were generated using a point vortex placed in a random location the center of the domain. Initial concentration fields were constant, with their level drawn from a uniform distribution not quite centered around the true initial concentration. Initial boundary velocities at the lid were described as parabolas centered at the center of the edge, with coefficients drawn from uniform distributions. Initial source amplitudes were each drawn independently from a half-normal distribution.

The velocities are first advanced forward to the next time when observations are made, whether they be of concentrations or velocity, or both. The velocity fields at increments in between are stored and passed to the convection-diffusion equation solver to advance the concentrations forward as well. The measurement matrix is recalculated each time data is available, depending on its nature, and the entire state vector Ψ is updated. The model used to propagate the ensemble solved for the lid-cavity flow at a lower spatial resolution using the predicted lid velocities at the lid. The peak of the GK i.e., its amplitude, was left as an unknown to be inferred and included in the EnKF state vector; see Eq. 3.35. This reduced the number of unknown sources significantly; the number of GKs used in the inversion are mentioned in the individual test cases. Images of nightlights provide the prior information on spatial patterns of ffCO₂ emissions.

In the EnKF inversions studies included in this chapter, we will assimilate both velocities and CO₂ observations. While the assimilation of velocity provides a multivariate Gaussian distribution (realized as an ensemble of velocity fields) for the reconstructed (a.k.a. re-analyzed) velocity field, only its mean will be used in the convection-diffusion system for the transport of CO₂ emissions. This is no different from the manner in which inversions are performed today, by neglecting the uncertainty in the re-analyzed velocity fields. However, the impact of the Gaussian assumption in EnKF will be quantified by comparing against an inversion using a MCMC method.

3.3.3 Enforcing non-negativity in EnKF

For nonlinear models, the EnKF can create nonphysical ensemble members during the update step. The simplest way to address these nonphysical values is to simply truncate them to the nearest acceptable value. Phale et al. [99] discusses three alternative methods in the context of estimating the properties of a hydrocarbon reservoir, the simplest of which involves assimilating violated constraints as perfect observations. [100] also encountered nonphysical values in using the EnKF with a nonlinear convection-diffusion-reaction fire propagation model, and proposed including a regularization term penalizing large gradients by incorporating the regularization term as an artificial observation in the EnKF update step. [101] addresses non-physical values near the displacement front of an immiscible flow through a change in parameterization. Phale’s method of constraint assimilation was implemented, but found to take up to four times as long to run as ignoring constraints, and did not necessarily produce better RMSE than the simple truncation method. Thus, negative values of concentration and amplitude parameters were truncated after each round of data assimilation.

Since the update step of EnKF uses only the mean and variance of the distribution, it may not correctly update the distribution conditional on the new data. To obtain a “correct” distribution, we used MCMC; for the sake of time, we examined a lower-resolution case with fewer GKs.

3.4 Numerical Results

In this section, we investigate the ramifications of the modeling choices made when using EnKF. First, we investigate the impact of the prior, which in our case is an informative one. We solve the inversion problem i.e., we reconstruct the ffCO_2 concentration field using the prior sources without assimilating in ffCO_2 concentrations, while using the true velocity field. Secondly, we investigate the impact of introducing nonlinearities by inferring velocities and ffCO_2 concentrations. Whereas in the first case (when we use the true velocity field) one has a linear problem, the second test involves solving the nonlinear Navier-Stokes model. Finally we investigate the impact of Gaussian assumptions by comparing EnKF inferences with those obtained using MCMC.

3.4.1 Inversions with EnKF

We perform an EnKF inversion on a 200×200 grid with ffCO_2 sources modeled with 1000 GKs. A 15×15 grid of concentration sensors are used (see Fig. 3.3). We first perform a run without data assimilation. Fig. 3.2 (top row) plots the inferred and true ffCO_2 concentration field, while the bottom row contains results for the sources. The true velocity field was used. We see that the prior is not very informative; the reconstructed ffCO_2 concentration and source fields bear little resemblance to the true solution if no observations are used to update them. Fig. 3.3 shows the impact of assimilating data (the 15×15 grid of concentration sensors is also shown). Data assimilation reduced the RMSE of the concentration field by over 55%. However, the RMSE of the *source* field actually increased after data assimilation. This is due to the Vulcan-derived source map having very small, concentrated sources, while the Gaussian kernel approximations are larger and vary smoothly within each ellipse. Qualitatively, in the reconstruction the cluster of sources along the northeast coast and the amplitudes of the larger Gaussian kernels located in Canada, which does not appear in the Vulcan-derived (truth-data) source map, are diminished.

Next, we introduce nonlinearities into the inversion. We assimilate both velocity and concentration measurements, using the Navier-Stokes equation to reconstruct the velocity field. These are plotted in Fig. 3.4. The EnKF performed similarly to the case when the velocity field was known, reducing the RMSE of the predicted final concentration field by about 55% from that predicted by the initial ensemble. In addition, the RMSE of the predicted final velocity fields and boundary velocities were reduced by about 75% and 65%, respectively.

Consistent with [102], it was found that assimilating data more frequently did not necessarily improve the final prediction. It was also noted that having a smaller measurement error sometimes worsened the final prediction, perhaps because this reduced the spread of the ensemble excessively after a few assimilation's while also causing the algorithm to try to match the observed (and not quite true) data too closely. This can be seen in the example in [94] where a very small measurement error led to a very erroneous prediction of a constant truth value.

3.4.2 EnKF vs MCMC

EnKF uses a Gaussian update step and although the post-processed of the ensembles introduce nonlinear effects from the dynamics, it is not clear how accurate the final statistics are. We therefore compare EnKF results to those from a delayed-rejection adaptive Metropolis (DRAM, [91]) MCMC run on a 60×60 grid

Figure 3.2: Plots of inferred (left) and true (right) ffCO₂ concentrations (top row) and ffCO₂ sources (bottom row) using no observations. The true velocity field was used. We see that the prior is not very informative.

with 10 GK (largest 10 of the 60 originally generated), in the case where the true velocity field is known. We used a single chain of 401400 forward simulations, with a burn-in time of 20,000 samples, and adapting every 1000 steps. The same measurement noise, sensors (8×8), initial concentration, velocity fields, and prior parameter distributions were used for both the EnKF and MCMC runs. It should be noted that MCMC is computationally intensive, forcing the use of small datasets and a small number of GKs. The coarseness of this dataset severely limits our ability to resolve velocity and concentration dynamics and this also places some uncertainty on the overall conclusions of the MCMC comparison. Nevertheless, the results of this comparison seem in line with our intuition and can therefore be used to guide the robustness of the EnKF results.

Fig. 3.5 presents the EnKF versus MCMC comparison for concentration. Qualitatively the mean solutions are very similar, although they both are not accurate in comparison to the truth model, which can be explained by the limited grid resolution and the small number of GKs. Similar conclusions can be drawn from the comparison of the GK magnitudes (Fig. 3.6).

Our goal is to assess the statistical quality of the EnKF results and potentially determine the appropriate adjustment factor. We are ultimately interested in an accurate mean solution and a measure of the certainty

Figure 3.3: Plots of inferred (left) and true (right) ffCO₂ concentrations (top row) and ffCO₂ sources (bottom row) using ffCO₂ concentration observations collected on a 15×15 sensor grid. The true velocity field was used. We see an immense improvement in the inferred ffCO₂ concentration field, and a more informative source inversion.

of the solutions, which can be extracted from the trace of the covariance matrix. Fig 3.7 shows the mean solution for different GK amplitude values, for EnKF and MCMC. As expected the mean solution is accurate and well within an acceptable quality range.

However, Fig 3.8 shows considerable differences between the variance calculations. The EnKF overestimates the standard deviation which is a result of the limited interrogation of the overall stochastic space combined with the effects of the Gaussian update step. Given the range of the difference, it is not clear if an adjustment factor is possible.

Fig 3.9 shows the difference between EnKF and MCMC covariances. As discussed the diagonal of the covariance shows a difference which is a measure of the certainty of the solutions. In this case, EnKF suggests a higher level of uncertainty of the solution.

Figure 3.4: Plots of inferred and true concentration fields (top row) and ffCO₂ sources (bottom row) when assimilating velocity and concentration measurements.

3.5 Conclusions

In this chapter, we have developed a random field model, based on Gaussian kernels, for the representation of ffCO₂ emissions. This representation is better suited for representing very localized emission fields e.g., the ones observed in the Midwest and the Great Plains. In such a case, a wavelet-based representation may provide a very compact representation only if the emissions are aligned along wavelet boundaries; otherwise, it is likely that GK provide a lower-dimensional approximation for the emissions. Further, the elliptical cross-section of the kernels allow some adaptation to the amorphous shape of most population centers, while the Gaussian shape allows approximation to the urban dome of ffCO₂ emissions.

We demonstrated the GK random field in a synthetic data inversion using EnKF. Comparing Figs. 3.2 and Fig. 3.3, it is clear that the limited observations of CO₂ concentrations significantly assist in the reconstruction of the concentration field, and to a large extent, also constrain the source/emission field. The joint inference of velocity and concentration fields (using velocity and concentration observations) do not yield much of an improvement over the true-velocity field inversion (compare Fig. 3.3 versus Fig. 3.4 with respect to the reconstructed emission and concentration fields), indicating, as expected, that the errors in the

Figure 3.5: Final concentration fields, reconstructed using limited CO₂ measurements. The true velocity field is used. The MCMC reconstruction is on the left and EnKF on the right. The true solution is in the middle.

Figure 3.6: Reconstruction of sources using the truth velocity. The MCMC reconstruction is on the left and the EnKF on the right. The true solution is in the middle.

reconstructed velocity are overwhelmed by the reconstruction error in CO₂ concentrations.

The comparison of EnKF versus MCMC do not provide a clear measure of the drawbacks of using EnKF in the estimation of ffCO₂ emissions. In the GK parameterization, the estimates of the GK amplitudes (means of the distributions) are very close, regardless of whether MCMC or EnKF was used. However, the estimates of the variances are very different, with EnKF over-estimating the variance in the inferred values of the GK amplitudes. The lack of variance in the MCMC solution could be because of the lack of samples (less than 500,000 samples drawn from a 10-dimensional parameter space) or to numerical difficulties in the EnKF (over-inflating the covariance after the data-assimilation step). The posterior, as reconstructed by the MCMC, was Gaussian, and in the limit of a correct reconstruction, the MCMC and EnKF results should have agreed. The good agreement in means, and a worse one in variances, indicate that the true distribution is symmetric. The root cause behind the differences in variance is thus numerical rather than in the model or the choice of algorithm.

Figure 3.7: Plot of the mean of posterior distribution of the amplitudes of 10 GK, inferred using the true velocity. There is not much difference between the means calculated with MCMC versus EnKF.

Figure 3.8: Plot of the variance of posterior distribution of GK amplitudes. There is significant difference in the results of MCMC versus EnKF.

Figure 3.9: Covariance matrix of posterior distribution of GK amplitudes, calculated using MCMC (left) and EnKF (middle). On the right is the difference between the two.

Chapter 4

Joint estimation of biogenic and fossil-fuel CO₂ fluxes

4.1 Introduction

Rising atmospheric concentrations of CO₂ are the main driver of anthropogenic climate change. The spatial and temporal variations in the global atmospheric CO₂ emissions can be inferred from bottom-up forward models or inverse models that have the capability to infer a combined global distribution of ffCO₂ and biospheric emissions. In recent past, inverse estimates of CO₂ emissions were inferred from in-situ measurements. These measurements were combined with atmospheric transport models (based on observed meteorology) in an inverse modeling framework that generates estimates of net exchange of land and oceanic CO₂ fluxes [103, 27, 104, 105]. However, due to sparse spatial coverage of in-situ measurements, the problem of atmospheric inversions is mathematically underdetermined and hence global CO₂ fluxes can be reasonably inferred only for large regions spread over multiple latitudes and longitudes. Satellite measurements of atmospheric CO₂ are a promising source for overcoming in-situ data limitations [106]. These data, combined with in-situ measurements, provide an opportunity to infer sub-grid-scale CO₂ fluxes, thereby leading to improved modeling, monitoring, and understanding of sinks and sources of CO₂ [107]. In recent years, these data have been utilized to perform global inversions for improved understanding of CO₂ exchange. However significant effort is required to reduce the uncertainty and refine the estimates of CO₂ exchange even at regional scale. This improvement in inverse estimates can aid in verification and monitoring of ffCO₂ emissions and formulation of policies for managing risks associated with climate change.

An important facet of the inversions for inferring global CO₂ exchange is that they cannot separate ffCO₂ and biospheric fluxes without the use of tracers like the radiocarbon isotope ¹⁴C. Recently, atmospheric inverse models have been suggested as a potential tool for independent verification of the bottom up estimates of ffCO₂ emissions; however due to large uncertainties, constraining continental or national level ffCO₂ emissions through atmospheric inverse models remains infeasible [5]. Improvements in terms of increasing in-situ and satellite measurements of CO₂ concentration and radiocarbon isotope ¹⁴C have been suggested to reduce this uncertainty. Targeted efforts (e.g. Megacities Carbon Project [108, 109, 110]; Hestia Project [37]; Indianapolis Flux Experiment (INFLUX), <http://influx.psu.edu/>) are ongoing with their focus on constraining the fossil fuel CO₂ (ffCO₂) fluxes in urban areas. However, these efforts cannot be replicated at large regional or national scales, especially in developing countries. Thus, they can only act as verification base cases for refining inverse modeling methodology predicated on a large number of CO₂ concentration measurements. In the next decade, remote sensing of CO₂ will provide a large number of observations which would reduce the ratio of observations to estimated fluxes, thereby reducing the uncertainty of ffCO₂ fluxes estimated within an inverse modeling framework. Other than reducing the ratio of observations to estimated fluxes, improvements are also required in transport and inverse models to realize

the full potential of the extra information provided by satellite observations.

In case of inverse models, these methodological improvements include, but are not limited to, estimation of fluxes over variable grids (see [111]) and designing improved random field models to account for the non-stationarity in the fluxes. A random field model proposed in Chapter 2 is one such approach. Another approach to modeling a non-stationary and non-smooth random field (e.g. distribution of ffCO_2) is suggested in this chapter. This method utilizes the information of the variability of the covariates of the ffCO_2 fluxes to develop covariance structures that define the underlying distribution. It is a novel approach as covariates of the dependent variables are used to define covariance and a prior for ffCO_2 fluxes. In this study, we assess the ability of these covariate-based covariance structures to disaggregate ffCO_2 and biospheric fluxes without using any information from radioisotope tracers.

4.2 Objectives and case studies

The primary objective of this study is to disaggregate biospheric and fossil-fuel fluxes at grid-scale ($1^\circ \times 1^\circ$) on the basis of error covariance models rather than through tracer-based CO_2 inversions. This study also makes several novel methodological contributions. It presents methods for (1) defining covariance structures for discontinuous non-smooth fields based on auxiliary variables, (2) selecting the *best* covariance structure from several competing covariance structures describing the same field, and (3) demonstrates the applicability of a simultaneous optimization scheme for obtaining covariance parameters and a set of *best* auxiliary variables within a geostatistical inverse modeling framework. The scientific goal of this research is to more accurately define the underlying structure of the processes governing the spatio-temporal distribution of the biospheric and ffCO_2 fluxes.

To this end, a separable stationary spatio-temporal exponential covariance model is prescribed for biospheric fluxes. A semi-stationary *diagonal* covariance model, derived by comparing auxiliary variables correlated with ffCO_2 fluxes and inventory estimates of ffCO_2 emissions, is used to represent the underlying distribution of ffCO_2 fluxes. The covariance matrices for ffCO_2 are derived from the spatial distribution of auxiliary variables at sub-grid-scale spatial resolution. After selecting covariance models associated with biospheric and ffCO_2 fluxes, we perform a joint-inversion to disaggregate these fluxes. As both negative and positive fluxes are obtained for ffCO_2 , a non-negativity constraint is imposed by Lagrange multipliers. The quality of disaggregation is examined by comparing the inferred spatial distribution of biospheric fluxes and fossil-fuel emissions in synthetic-data and real-data inversions. In addition to disaggregation of fluxes, the ability of the covariance models to explain the fossil-fuel emissions over North America is also examined.

Two case studies are performed to assess the methodology employed for disaggregating biospheric and ffCO_2 fluxes. In the first case study, we assess the quality of disaggregation within a pseudo-data setup whereby synthetic CO_2 observations are generated by transporting *combined* biospheric and ffCO_2 fluxes forward. This is done by multiplying them by the transport sensitivity matrix ($\mathbf{H}$) obtained from the Stochastic Time Inverted Lagrangian Model, driven by meteorological fields from the Weather Research and Forecasting model [76]. We refer to this as WRF-STILT. For this purpose, the biospheric fluxes from Carnegie Ames Stanford Approach terrestrial carbon cycle model (CASA), as configured for the Global Fire Emissions Database v2 project [112, 113] and ffCO_2 fluxes from the Carbon Dioxide Information Analysis Center (CDIAC) global monthly fossil fuel inventory at $1^\circ \times 1^\circ$ were used (see [114]). Diurnal variability in the CDIAC ffCO_2 fluxes were introduced by using the scaling factors given in [115]. CASA-GFED v2 fluxes can only be obtained at monthly scale, hence for inversion they were downscaled using the method in [116] which is based on net shortwave radiation and near-surface temperature data from the NASA Global Land

Data Assimilation System (GLDAS, [117]). The model-data mismatch variances for measurements were included in this case study and they varied by observation locations and month.

In the second case study, the same transport sensitivity matrix ($\mathbf{H}$) was used to disaggregate biospheric and ffCO₂ fluxes. Only 3 afternoon observations with data gaps were included while performing inversions (in both the case studies). The covariance and model-data mismatch parameters in this case study were obtained within a dual optimization framework as outlined in the next section. The competing covariance structures and auxiliary variables for ffCO₂ estimation in both case studies were selected by comparing Bayesian Information Criterion (BIC) and the results for these are presented in Sec. 4.4.

The inferred grid-scale fluxes in the pseudo-data study are compared with the true monthly biospheric and ffCO₂ fluxes (CASA-GFED and CDIAC) by computing root mean squared error (RMSE) and correlation coefficient at grid-scale. Inversion results are also compared qualitatively by examining the spatial patterns of inferred fluxes to those of the true fluxes.

4.3 Numerical details

4.3.1 Flux domain and resolution

In the joint inversion, biospheric fluxes are estimated at a 3-hourly temporal resolution and ffCO₂ fluxes are estimated at 8-day interval for each month in 2008. They are then aggregated up to monthly resolution. These fluxes are obtained at grid-scale for a land area between 10° N to 70° N and 50° W to 170° W. The sensitivity of fluxes to measurements was obtained from WRF-STILT (for details see [2]). As suggested in [78], only afternoon measurements from 35 towers in North America in 2008 were used in this research.

4.3.2 Methods

The geostatistical inverse modeling (GIM, [27, 2]) is a Bayesian method in which a model of the trend of the CO₂ flux distribution replaces the prior flux estimates used in traditional Bayesian inverse modeling. The criterion of remaining close to a prior flux estimate is replaced by a criterion of preserving a spatial and/or temporal correlation in the distribution of the CO₂ flux residuals. The objective function for the GIM can be written as is:

$$2L(\mathbf{s}, \boldsymbol{\beta}) = (\mathbf{z} - \mathbf{H}\mathbf{s})^T \mathbf{R}^{-1} (\mathbf{z} - \mathbf{H}\mathbf{s}) + (\mathbf{s} - \mathbf{X})^T \mathbf{Q}^{-1} (\mathbf{s} - \mathbf{X}) \quad (4.1)$$

where $\mathbf{z}$ are CO₂ concentrations, $\mathbf{H}$ is the sensitivity matrix, $\mathbf{s}$ are the estimated CO₂ fluxes, $\mathbf{R}$ is the model-data mismatch, $\mathbf{X}$ are auxiliary variables, $\boldsymbol{\beta}$ are estimated drift coefficients and $\mathbf{Q}$ is the error covariance matrix. In this study, we modify this objective function to solve for two different flux fields, i.e., ffCO₂ and biospheric fluxes. To estimate these flux fields we modify the objective function given in Eq. 4.1. The modified objective function can be written as:

$$\begin{aligned} 2L(\mathbf{s}_{bio}, \mathbf{s}_{ff}, \boldsymbol{\beta}_{bio}, \boldsymbol{\beta}_{ff}) &= (\mathbf{z} - [\mathbf{H}_{bio}\mathbf{s}_{bio} + \mathbf{H}_{ff}\mathbf{s}_{ff}])^T \mathbf{R}^{-1} (\mathbf{z} - [\mathbf{H}_{bio}\mathbf{s}_{bio} + \mathbf{H}_{ff}\mathbf{s}_{ff}]) \\ &+ (\mathbf{s}_{bio} - \mathbf{X}_{bio}\boldsymbol{\beta}_{bio})^T \mathbf{Q}_{bio}^{-1} (\mathbf{s}_{bio} - \mathbf{X}_{bio}\boldsymbol{\beta}_{bio}) \\ &+ (\mathbf{s}_{ff} - \mathbf{X}_{ff}\boldsymbol{\beta}_{ff})^T \mathbf{Q}_{ff}^{-1} (\mathbf{s}_{ff} - \mathbf{X}_{ff}\boldsymbol{\beta}_{ff}) \end{aligned} \quad (4.2)$$

where subscripts *bio* and *ff* represent the biospheric and fossil fuel components of the symbols defined in Eq. 4.1. The sensitivities of CO₂ concentrations (**z**) to fluxes obtained from WRF-STILT were aggregated at 3-hourly and 8-day temporal resolution to get **H_{bio}** and **H_{ff}**. The model-data mismatch (**R**) and error covariance parameters (**Q_{ff}** and **Q_{bio}**) were obtained within a Restricted Maximum Likelihood (REML) framework as described in [75].

4.3.3 Biospheric and ffCO₂ covariance structure

A stationary, separable spatio-temporal covariance function was used to model the underlying correlation of biospheric flux residuals. This function which can be written as (for details see [2]):

$$\mathbf{Q}_{bio} = \sigma^2 \left[\exp \left(-\frac{h_{temp,bio}}{l_{temp,bio}} \right) \otimes \exp \left(-\frac{h_{spat,bio}}{l_{spat,bio}} \right) \right] \quad (4.3)$$

where σ^2 is the variance in space and time, $h_{temp,bio}$ represents the separation distances/lags between estimation locations in space and time for biospheric fluxes and $l_{temp,bio}$ and $l_{spat,bio}$ are the spatial and temporal correlation range parameters for biospheric fluxes. Temporal covariance for biospheric fluxes was only defined across days and no covariance was prescribed for fluxes estimated within a day.

In the case of ffCO₂ we assumed no spatio-temporal covariance between flux residuals and this covariance structure was defined as:

$$\mathbf{Q}_{ff} = \left(a \begin{bmatrix} k_1 & 0 & 0 \\ 0 & k_2 & 0 \\ \dots & \dots & \dots \\ 0 & 0 & k_r \end{bmatrix} + b \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \dots & \dots & \dots \\ 0 & 0 & 1 \end{bmatrix} \right) \quad (4.4)$$

where a and b are tuning parameters to be estimated and $k_1 \dots k_r$, describes the independent covariance of flux residuals at each grid-cell. The parameters associated with **Q_{ff}** also govern the intensity of ffCO₂ estimates at grid-scale.

A novel approach was adopted in this study to obtain $k_i, i = 1 \dots r$ shown in Eq. 4.4. We assumed that k_i in Eq. 4.4 can be prescribed by the covariates of ffCO₂ fluxes. Three of these covariates i.e., nightlight intensity, population density and percentage of built up area were available at a spatial resolution higher than $1^\circ \times 1^\circ$. Hence we computed the mean, variance and maximum value at grid-scale of (1) nightlight intensity (year: 2008; spatial resolution 1 km), (2) population density (year: 2008; spatial resolution 2.5 km), and (3) percent built up area (year: 2008; spatial resolution 10 km) from available higher resolution estimates of these quantities. This resulted in a set of 9 k_i (i.e., $i = 1 \dots 9$) that can be used to define **Q_{ff}**. In addition to these 9 k_i we also added another k_i representing the mean ffCO₂ fluxes obtained by forming a temporally scaled combined fossil fuel inventory of Vulcan [4] and ODIAC [118, 15] for our estimation domain (see [2]). The use of the inventory to prescribe k_i was done to assess the ability of the corresponding covariance in disaggregating ffCO₂ fluxes from biospheric fluxes. To select an appropriate set of k_i for performing joint inversion, we employed an iterative dual-optimization scheme through which we obtained an “optimal k_i ” based on BIC.

Model-data mismatch errors were assumed to be uncorrelated in space and time, yielding a diagonal matrix **R**. A different variance was estimated for each tower as part of the covariance parameter optimization.

4.3.4 Model of the trend

For explaining the variability of ffCO₂ fluxes as seen through atmospheric concentration measurements we included mean nightlight intensity, population density, percent built up area, percent urban area and combined Vulcan and ODIAC fossil fuel inventory in $\mathbf{X}_{ff}$. These covariates were included to aid in disaggregation of ffCO₂ fluxes from biospheric fluxes. $\mathbf{X}_{bio}$ in this study was represented as a vector of ones associated with 3-hourly estimation time periods. The estimated β_{bio} associated with them represents the constant mean value in space and time of fluxes across grid cells.

4.3.5 Covariance, covariate selection and parameter optimization

A novel dual simultaneous optimization scheme was employed to (1) obtain parameters of $\mathbf{Q}_{ff}$, $\mathbf{Q}_{bio}$ and model data mismatch and (2) for selecting the best covariates of ffCO₂ fluxes for inclusion in $\mathbf{X}_{ff}$. In this optimization scheme, the REML (see [27]) objective function was used for optimizing covariance parameters. This objective function can be written as:

$$L = \log |\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R}| + \log |(\mathbf{H}\mathbf{X})^T (\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1} (\mathbf{H}\mathbf{X})| + [\mathbf{z}^T ((\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1} - (\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1} \mathbf{H}\mathbf{X} - (\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1} \mathbf{H}\mathbf{X}) \mathbf{z}] \quad (4.5)$$

where all symbols are as defined earlier and $\mathbf{Q}$ and $\mathbf{X}$ can be given as:

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}_{bio} & 0 \\ 0 & \mathbf{Q}_{ff} \end{bmatrix} \quad (4.6)$$

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_{bio} \\ \mathbf{X}_{ff} \end{bmatrix} \quad (4.7)$$

A large-scale trust region, reflective Newton's algorithm, was used to optimize the function given in Eq. 4.5. In addition, a branch-and-bound discrete optimization (for details see [119]) algorithm was employed to select covariates in $\mathbf{X}_{ff}$ to avoid overfitting of the model of the trend. In this study, BIC allowed comparison of all possible combinations of a superset of auxiliary variables and covariance models. Thus, as part of each iteration of the optimization of the likelihood function (Eq. 4.5), the branch-and-bound algorithm was used to select the best model associated with a covariance structure in a backward elimination framework. A likelihood function was used to obtain an optimal set of auxiliary variables associated with each covariance model. The likelihood function used as part of discrete optimization framework can be given as:

$$BIC = \log |\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R}| + [\mathbf{z}^T ((\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1} - (\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1} \mathbf{H}\mathbf{X} - (\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1} \mathbf{H}\mathbf{X}) \mathbf{z}] + p \log(n) \quad (4.8)$$

where n represents the number of observations and p is the number of parameters included in the inverse model. This process of optimizing covariance parameters and selecting a model for $\mathbf{X}_{ff}$ was repeated for 10 different covariance structure of $\mathbf{Q}_{ff}$ described in Sec. 4.3.3 for all 12 months of 2008. Out of these 10 different covariance structures, the set of k_i and $\mathbf{X}_{ff}$ that resulted in minimum BIC were selected for performing the joint inversion.

4.3.6 Non-negativity constraints on ffCO₂ fluxes

The joint inversion performed after estimating covariance parameters and obtaining a set of “best” covariates for $\mathbf{X}_{ff}$ resulted in some grid cells with negative ffCO₂ fluxes. To avoid this non-physical result, we imposed non-negativity on the ffCO₂ fluxes during the joint inversion. No constraints were imposed on the biospheric fluxes as they can be negative or positive depending on the seasonal variation in respiration and gross primary productivity.

Though there are many methods like data transformation and Markov Chain Monte Carlo for inferring $\mathbf{s}$ (see [120]), none of them are suitable for imposing non-negativity constraints on a large state space. Hence, in this work, we used Lagrange multipliers to obtain non-negative ffCO₂ fluxes. This method consists of rewriting the original function into a Lagrange formulation as given in Eq. 4.9.

$$h[L(\mathbf{s}_{ff}, \mathbf{\beta}_{ff}), \boldsymbol{\lambda}] = f[L(\mathbf{s}_{ff}, \mathbf{\beta}_{ff})] - \sum_{i=1}^p \lambda_i [\delta_i (L(\mathbf{s}_{ff}, \mathbf{\beta}_{ff})) - b_i] \quad (4.9)$$

where $L(\mathbf{s}_{ff}, \mathbf{\beta}_{ff})$ must satisfy the constraints such that $\delta_i [L(\mathbf{s}_{ff}, \mathbf{\beta}_{ff})] \geq b_i$, where m is the total number of active constraints, and $\boldsymbol{\lambda} = \{\lambda_1, \lambda_2, \dots, \lambda_p\}$ are the Lagrange multipliers. This involves setting the derivative of the Lagrange function equal to zero by satisfying the first-order Kuhn-Tucker conditions (for additional details see [121]). Note that since non-negativity constraints were imposed on the ffCO₂ fluxes we have only included those components of Eq. 4.2 on which non-negativity was imposed in Eq. 4.9. Non-negativity constraints were imposed for an inversion performed using the best covariance structure and auxiliary variables obtained in the previous step. Since only fluxes are updated while imposing non-negativity, the uncertainty associated with the fossil fuel and biospheric fluxes were not changed.

4.4 Results and discussion

In this section, we first describe results of model selection of auxiliary variables and covariance structure. This is followed by the performance of inversions in estimating ffCO₂ fluxes for January and June which are representative of winter and summer. We perform two case studies: (1) using pseudo-data (synthetic CO₂ observations) and (2) using real CO₂ observations.

4.4.1 Model selection

In the real-data case study we find a consistent pattern for all months. From the 10 covariance structure described in Sec. 4.3.5, the Vulcan and ODIAC-based covariance structure is selected by BIC for all months. However, the results vis-à-vis covariates across months are not consistent. Thus mean population density is selected for six months, followed by mean nightlight intensity and fossil fuel inventory which are selected for four months. Percent built up area and percent urban area are least important in explaining ffCO₂ fluxes and are only selected during two and one months respectively (for details see Table 4.1). Given these results, we can conclude that mean population density and nightlight intensity are the most important factors associated with ffCO₂ fluxes. In this study, we used Vulcan and ODIAC in $\mathbf{X}_{ff}$ and in defining $\mathbf{Q}_{ff}$. However, in cases where independent verification of ffCO₂ fluxes are required, these data cannot be used and it can be concluded from the results in Table 4.1 that population density and nightlight intensity are most reliable indicators of ffCO₂ emissions. This is not surprising as other studies have found similar results. It is

Table 4.1: Best covariates and covariance model in the real-data case study. Selections for January–December 2008 are shown. Note: An asterisk indicates a covariate.

Month	Mean nightlight*	Mean population density*	Percent built-up area*	Percent urban area*	Vulcan & ODIAC*	Best covariance model
1					Yes	Vulcan & ODIAC
2		Yes		Yes		Vulcan & ODIAC
3	Yes	Yes		Yes		Vulcan & ODIAC
4	Yes		Yes			Vulcan & ODIAC
5					Yes	Vulcan & ODIAC
6			Yes			Vulcan & ODIAC
7		Yes				Vulcan & ODIAC
8		Yes				Vulcan & ODIAC
9					Yes	Vulcan & ODIAC
10	Yes	Yes				Vulcan & ODIAC
11	Yes	Yes				Vulcan & ODIAC
12					Yes	Vulcan & ODIAC

also important to analyze second-best covariance models, since in verification studies fossil fuel inventories cannot be used to prescribe covariance in inversions. There is no single covariance model that performs consistently across months, but, overall, a covariance based on maximum nightlight intensity is selected in most months (see Appendix C).

In the pseudo-data case study we found that Vulcan and ODIAC-based covariance structure is not selected by BIC (Table 4.2). Additionally it also not selected as a significant covariate in any month. This clearly indicates that there are significant differences between these two inventories. However, as “Vulcan and ODIAC” inventory is selected in the real-data case study, it correlates more with the underlying spatio-temporal fossil fuel distribution in North America. The spatial distribution of CDIAC was found to be more correlated with variance in population density.

4.4.2 Disaggregated ffCO₂ and biospheric fluxes

In both real and pseudo-data cases, disaggregating biospheric and ffCO₂ fluxes on the basis of covariance structure proved to be more difficult in summer versus winter. In both case studies the magnitude of fluxes was not accurate across months, though this problem was more acute in summer. ffCO₂ fluxes were only detected in major cities. This primarily happened due to the difference in the signal of biospheric and ffCO₂ fluxes which was higher in summer. Other than seasonal variation of gross primary productivity, the magnitude of the ffCO₂ signal (in case of real data) was extremely sensitive to a and b parameters in the $\mathbf{Q}_{ff}$. A small a resulted in smaller magnitude of fossil fuel fluxes across all months in the year whereas a larger a increased the magnitude of the ffCO₂ fluxes and also resulted in increased number of negative values. Even though a was determined through REML, the uncertainty on a , estimated by computing the Hessian of the objective function, was extremely high. Thus, the accuracy of a in $\mathbf{Q}_{ff}$ in real-data inversions is doubtful. This problem was not observed in the pseudo-data study as the parameter a and b were known and pre-computed by the likelihood method described in [27, 122]

Table 4.2: Best covariates and covariance models for the pseudo-data case study. Selections for January–December 2008 are shown. Note: An asterisk indicates a covariate.

Month	Mean nightlight*	Mean population density*	Percent built-up area*	Percent urban area*	Vulcan & ODIAC	Best covariance model
1	Yes	Yes				Variance, pop. density
2	Yes	Yes		Yes		Variance, pop. density
3	Yes			Yes		Variance, pop. density
4	Yes					Variance, pop. density
5	Yes	Yes				Variance, pop. density
6		Yes				Variance, pop. density
7		Yes				Variance, pop. density
8				Yes		Variance, pop. density
9		Yes				Variance, pop. density
10	Yes			Yes		Variance, pop. density
11			Yes			Variance, pop. density
12	Yes			Yes		Variance, pop. density

Table 4.3: Root mean squared error (RMSE) and correlation coefficient of the estimated biospheric and ffCO₂ fluxes in the pseudo-data case for January and June 2008. CDIAC and CASA-GFED are considered to provide the true ffCO₂ and biospheric CO₂ fluxes.

	Correlation	RMSE
Estimated ffCO ₂ fluxes with CDIAC (January)	0.62	0.52
Estimated biospheric fluxes with CASA-GFED (January)	0.76	0.58
Estimated ffCO ₂ fluxes with CDIAC (June)	0.54	0.68
Estimated biospheric fluxes with CASA-GFED (June)	0.79	0.55

In the real-data case study, the correlation coefficient of fossil fuel fluxes (to the real fluxes) was higher in winter. A similar trend was also observed with respect to RMSE which was lower in winter. Inferred fluxes are shown in Fig. 4.1 and 4.2 for January and June.

In the pseudo-data case study, in comparison to summer, the correlation coefficient of ffCO₂ fluxes was higher and a similar trend was also observed with respect to RMSE which was lower in winters (see Table 4.3). However, for the estimated biospheric fluxes, the correlation coefficient and RMSE did not vary considerably between January and June. Total disaggregation between ffCO₂ and biospheric fluxes was not achieved in June and some residual ffCO₂ fluxes were observed in the biospheric component of the disaggregated fluxes (see Fig. 4.2). This was not observed in January when the gross primary productivity was lower. Overall, even with the aid of covariates, we were not able to detect the summer signal of ffCO₂ fluxes which clearly points to the need of having a radioisotope tracer and more measurements in inversions.

Figure 4.1: Estimated January biospheric (top) and ffCO₂ fluxes (bottom) for June for the real-data study. Units: moles m²s⁻¹.

Figure 4.2: Estimated June biospheric (top) and ffCO₂ fluxes (bottom) for June for the real-data study. Units: moles m²s⁻¹.

4.5 Conclusions

The patterns of the disaggregated ffCO_2 fluxes in January are similar to the inventory data. The Vulcan-ODIAC, nightlight and population density-based covariance structures are able to capture the underlying distribution of the ffCO_2 fluxes in disaggregated inversions. However, the magnitude of the fluxes is not adequately captured by these diagonal covariance structures. These result indicates that the use of tracers is essential for isolating fossil fuel CO_2 emissions from biospheric fluxes (e.g., radiocarbon, carbon monoxide [123, 110]). Other than tracer information, ongoing efforts of estimating ffCO_2 emissions at fine scale in urban areas (e.g., Megacities Project) can provide information for constructing better ffCO_2 covariance models. The covariance structures proposed in this research, on their own, cannot disaggregate ffCO_2 and biospheric fluxes from CO_2 concentrations; however, they can help define the covariance in tracer-based inversions that attempt to disaggregate biospheric and ffCO_2 fluxes. Additionally, they can also be used to construct covariance matrices in inversions that only solve for biospheric fluxes.

Chapter 5

Conclusions

In this study, we addressed the question of estimating ffCO_2 emissions from limited CO_2 concentration measurements. We investigated the construction of spatial parameterizations (random field models) for ffCO_2 emissions, scalable inversion techniques for inferring them and whether they could be estimated from CO_2 concentration measurements alone, without recourse to species like CO (which are markers of incomplete combustion) or radiocarbon (^{14}C) which are not present in ffCO_2 , but are available, in a known proportion, in the biosphere.

Random field models: ffCO_2 emissions are characterized by a strongly non-stationary or multiscale spatial distribution. We constructed and tested three separate random field models for such emissions. The first, based on wavelets, is general i.e., it can be used at regional scale while representing ffCO_2 emissions at a 1° resolution. The model is not particularly low-dimensional, and when used in a limited-measurement scenario, requires a novel sparse reconstruction algorithm to guard against over-fitting. The sparse reconstruction algorithm employs elements of compressive sensing, maintains the non-negativity of ffCO_2 emissions and uses spatial patterns encoded in images of lights at night (a proxy for ffCO_2 emissions) to perform the atmospheric inversion. The second random field model is based on Gaussian kernels and provides an efficient parameterization when ffCO_2 sources are very localized, e.g., on the Great Plains. The third parameterization uses a multivariate Gaussian field model for ffCO_2 emissions, with the covariance developed from an amalgamation of several proxies of ffCO_2 emissions e.g., lights at night, population density etc. The method for selecting these proxies was also developed in our investigations.

The three models, the estimation problem formulated with them, and in two cases, the inversion algorithms, are unique since they are used for estimating spatially-resolved ffCO_2 emissions at the *regional* scale. They were demonstrated using synthetic and real data. There are no competing models and methods; current ffCO_2 emission estimates have been drawn only for individual urban regions.

Scalable inversion techniques: We investigated ensemble Kalman filters (EnKF) as a potential mechanism for solving the ffCO_2 estimation problem, while quantifying the uncertainty in the estimates. A particular aspect of EnKF, the Gaussian assumptions built into them, was investigated with respect to their impact on ffCO_2 estimates. We perform this check by comparing with solutions of the ffCO_2 estimation problem using a Markov chain Monte Carlo (MCMC) method. We find that EnKF and MCMC yield similar mean values of ffCO_2 emissions but the variance of the estimates are quite different. We believe that the root cause of the difference is numerical (too few MCMC samples or over-inflation in EnKF using data assimilation), rather than in modeling assumptions.

Measurement variables: We find that in order to estimate spatially-resolved ffCO_2 emissions at a regional scale, it will be necessary to measure radiocarbon or some other marker of fossil-fuel combustion (e.g., CO). ffCO_2 and biogenic CO_2 fluxes display very different spatial distributions. We hypothesized that this difference, as captured by their respective covariance models, would be sufficient to estimate them jointly

from CO₂ concentration measurements, without radiocarbon or other tracers. This hypothesis held only during the winter months, when gross primary productivity is comparable to ffCO₂ emissions. In pseudo-data inversions, we obtained correlations of around 0.7 with true emissions. In summer, the correlation degraded to around 0.5, indicating that even the spatial distribution was wrong. Apart from spatial distribution, the actual error in the inferred ffCO₂ emissions was around 50%, indicating that a large fraction of ffCO₂ emissions were being estimated as biogenic fluxes. This indicates that future estimates of ffCO₂ emissions may only be possible with a far higher density of measurement locations (we used 35 towers) *as well as* the use of tracer (e.g., CO, radiocarbon etc.) information, so that the ffCO₂ signal may be uniquely captured.

References

- [1] Global Fire Emissions Database. <http://www.globalfiredata.org/index.html>.
- [2] S. M. Gourdji, K. L. Mueller, V. Yadav, D. N. Huntzinger, A. E. Andrews, M. Trudeau, G. Petron, T. Nehrkorn, J. Eluszkiewicz, J. Henderson, D. Wen, J. Lin, M. Fischer, C. Sweeney, and A.M. Michalak. North American CO₂ Exchange: Inter-comparison of modeled estimates with results from a fine-scale atmospheric inversion. *Biogeosciences Discussions*, 9:457–475, 2012. doi:10.5194/bg-9-457-2012.
- [3] Project Vulcan webpage. <http://www.purdue.edu/eas/carbon/vulcan/index.php>.
- [4] K. R. Gurney, D. L. Mendoza, Y. Zhou, M. L. Fischer, C. C. Miller, S. Geethakumar, and S. de la Rue de Can. High resolution fossil fuel combustion CO₂ emission fluxes in the United States. *Environmental Science and Technology*, 43:5535–5541, 2009.
- [5] Committee on Methods for Estimating Greenhouse Gas Emissions; National Research Council. *Verifying Greenhouse Gas Emissions: Methods to Support International Climate Agreements*. The National Academies Press, 2010.
- [6] J. G. J. Olivier, J. Peters, J. Bakker, J. J. M. Berdowski, A. J. H. Visschedijk, and J. J. Bloos. *EDGAR 3.2: Reference database with trend data of global greenhouse gas emissions for 1970–1995*. Milpress Science Publishers, Rotterdam, 2002.
- [7] C. N. H. Doll, J-P. Muller, and C. D. Elvidge. Night-time imagery as a tool for global mapping of socioeconomic parameters and greenhouse gas emissions. *Ambio*, 29:157–162, 2000.
- [8] P. J. Rayner, M. R. Raupach, M. Paget, P. Peylin, and E. Koffi. A new global gridded data set of CO₂ emissions from fossil fuel combustion: Methodology and evaluation. *Journal of Geophysical Research*, 115, 2010. D19306.
- [9] P. Ciais, P. Rayner, F. Chevallier, P. Bousquet, M. Logan, P. Peylin, and M. Ramonet. Atmospheric inversions for estimating CO₂ fluxes: methods and perspectives. *Climate Change*, 103:69–92, 2010.
- [10] J. C. Turnbull, A. Karion, M. L. Fisher, I. Faloona, T. Guilderson, S. J. Lehman, B. R. Miller, J. B. Miller, S. Montzka, T. Sherwood, S. Saripalli, C. Sweeny, and P. P. Tans. Assessment of fossil fuel carbon dioxide and other anthropogenic trace gas emissions from airborne measurements over Sacramento, California in spring 2009. *Atmospheric Chemistry and Physics*, 11:705–721, 2011.
- [11] K. McKain, S. Wofsy, T. Nehrkorn, J. Eluszkiewicz, J. R. Ehrlinger, and B. B. Stephens. Assessment of ground-based atmospheric observations for verification of greenhouse gas emissions from an urban region. *Proceedings of the National Academy of Science*, 109(22):8423–8428, 2012.
- [12] M. Kopacz, D. J. Jacob, J. A. Fisher, J. A. Logan, L. Zhang, I. A. Megretskaya, R. M. Yantosca, K. Singh, D. K. Henze, J. P. Burrows, M. Buchwitz, I. Khlystova, W. W. McMillan, J. C. Gille, D. P. Edwards, A. Eldering, V. Thouret, and P. Nédélec. Global estimates of CO₂ sources with high resolution by adjoint inversion of multiple satellite datasets (MOPITT, AIRS, SCIAMACHY, TES). *Atmospheric Chemistry and Physics*, 10(3):855–876, 2010.

- [13] The United Nations Collaborative Program on Reducing Emissions from Deforestation and Forest Degradation in Developing Countries. <http://www.un-redd.org/>.
- [14] United Nations Framework Convention on Climate Change. <http://unfccc.int/>.
- [15] T. Oda and S. Maksyutov. A very high-resolution ($1\text{km} \times 1\text{km}$) global fossil fuel CO_2 emission inventory derived using a point source database and satellite observations of nighttime lights. *Atmospheric Chemistry and Physics*, 11:543–556, 2011.
- [16] J. G. J. Olivier, J. A. V. Aardenne, F. J. Dentener, V. Pagliari, L. N. Ganzeveld, and J. A. H. W. Peters. 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.
- [17] G. Marland. Uncertainties in accounting for CO_2 from fossil fuels. *Journal of Industrial Ecology*, 12(2):136–139, 2008.
- [18] G. Marland, K. Hamal, and M. Jonas. How uncertain are estimates of CO_2 emissions? *Journal of Industrial Ecology*, 13(1):4–7, 2009.
- [19] J. E. Sinton. Accuracy and reliability of China’s energy statistics. *China Economic Review*, 12:373–383, 2001.
- [20] Weather research and forecasting model. <http://www.wrf-model.org/index.php>.
- [21] S. Houwelling, F.-M. Breon, I. Aben, C. Rodenbeck, M. Gloor, M. Heimann, and P. Ciais. Inverse modeling of CO_2 sources and sinks using satellite data: A synthetic inter-comparison of measurement techniques and their performance as a function of space and time. *Atmospheric Chemistry and Physics Discussions*, 3:5237–5274, 2003.
- [22] K. L. Mueller, S. M. Gourdj, and A. M. Michalak. Global monthly averaged CO_2 fluxes recovered using a geostatistical inverse modeling approach: 1. Results using atmospheric measurements. *Journal of Geophysical Research*, 2008. D21114.
- [23] L. Feng, P. I. Palmer, H. Bosch, and S. Dance. Estimating surface CO_2 fluxes from space-borne CO_2 dry air mole fraction observations using ensemble Kalman Filter. *Atmospheric Chemistry and Physics*, pages 2619–2633, 2009.
- [24] F. Chevallier, R. J. Engelen, C. Carouge, T. J. Conway, P. Peylin, C. Pickett-Heaps, M. Ramonet, P. J. Rayner, and I. Xueref-Remy. AIRS-based versus flask-based estimation of carbon surface fluxes. *Journal of Geophysical Research*, 114, 2009. D20303.
- [25] R. Nassar, D. B. A. Jones, S. S. Kulawik, J. R. Worden, K. W. Bowman, R. J. Andres, P. Suntharalingam, J. M. Chen, C. A. M. Brenninkmeijer, T. J. Schuck, T. J. Conway, and D. E. Worthy. Inverse modeling of CO_2 sources and sinks using satellite observations of CO_2 from TES and surface flask measurements. *Atmospheric Chemistry and Physics*, 11:6029–6047, 2011.
- [26] A. Chatterjee, A. M. Michalak, K. L. Mueller, J. L. Anderson, and V. Yadav. Towards reliable ensemble Kalman filter estimates. *Journal of Geophysical Research - Atmospheres*, 117, 2012. D22306.
- [27] A. M. Michalak, L. Bruhwiler, and P. Tans. A geostatistical approach to surface flux estimation of atmospheric trace gases. *Journal of Geophysical Research*, 2004. D14109.
- [28] G. Strang and T. Nguyen. *Wavelets and filter banks*. Wellesley-Cambridge Press, Wellesley, MA, 1997.

- [29] Tony F. Chan and Jackie (Jianhong) Shen. *Image Processing and Analysis – Variational, PDE, Wavelet, and Stochastic Methods*. Society of Applied Mathematics, 2005.
- [30] Energy Statistics Database. <http://unstats.un.org/unsd/energy/edbase.htm>.
- [31] Carbon Monitoring and Action website. <http://carma.org>.
- [32] T. D. Kelly and G. R. Matos. Historical statistics for mineral and material commodities in the United States. *U.S. Geological Survey Data*, 140, 2007.
- [33] C. D. Elvidge, K. E. Baugh, E. A. Kihn, H. W. Kroehl, and E. R. Davis. Mapping city lights with nighttime data from DMSP operational linescan system. *Photogrammetric Engineering and Remote Sensing*, 63:727–734, 1997.
- [34] European Commission, Joint Research Centre (JRC)/Netherlands Environmental Assessment Agency (PBL). Emission Database for Global Atmospheric Research (EDGAR), release version 4.1. <http://edgar.jrc.ec.europa.eu/index.php>. EDGAR/EUROPA Web-site.
- [35] Carbon dioxide information analysis center webpage. <http://cdiac.ornl.gov>.
- [36] R. J. Andres, J. S. Gregg, L. Losey, G. Marland, and T. A. Boden. Monthly, global emissions of carbon dioxide from fossil fuel consumption. *Tellus B*, 63(3):309–327, 2011.
- [37] K. R. Gurney, I. Razlivanov, Y. Song, Y. Zhou, B. Benes, and M. Abdul-Massih. Quantification of fossil fuel CO₂ emissions on the building/street scale for a large city. *Environmental Science and Technology*, 46:12194–12202, 2012.
- [38] Hestia webpage. <http://www.purdue.edu/climate/hestia>.
- [39] J. S. Gregg, R. J. Andres, and G. Marland. China: Emission pattern of the world leader in CO₂ emissions from fossil fuel consumption and cement production. *Geophysical Research Letters*, 35, 2008.
- [40] C. S. Potter, J. T. Randerson, C. B. Field, P. A. Matson, P. M. Vitousek, H. A. Mooney, and S. A. Klooster. Terrestrial ecosystem production: A process model based on global satellite and surface data. *Global Biogeochemical Cycles*, 7(4):811–841, 1993.
- [41] C. Rodenbeck, S. Houwelling, M. Gloor, and M. Helmann. CO₂ flux history 1982–2001 inferred from atmospheric transport data using a global inversion of atmospheric transport. *Atmospheric Chemistry and Physics*, 3:1919–1964, 2003.
- [42] F. Chevalier, M. Fisher, P. Peylin, S. Serrar, P. Bousquet, F.-M. Breon, A. Chédin, and P. Ciais. Inferring CO₂ sources and sinks from satellite observations: Method and application to TOVS data. *Journal of Geophysical Research*, 110, 2005. D24309.
- [43] W. Peters, J. B. Miller, J. Whitaker, A. S. Demming, A. Hirsch, M. C. Krol, D. Zupanski, L. Bruhwiler, and P. P. Tans. An ensemble data assimilation system to estimate CO₂ surface fluxes from atmospheric gas observations. *Journal of Geophysical Research*, 110, 2005. D24304.
- [44] S. M. Gourdji, K. L. Mueller, K. Schaefer, and A. M. Michalak. Global monthly averaged CO₂ fluxes recovered using a geostatistical inverse modeling approach: 2. Results including auxiliary environmental data. *Journal of Geophysical Research*, 113, 2008. D21115.

- [45] S. M. Gourdji, A. I. Hirsch, K. L. Mueller, A. E. Andrews, and A. M. Michalak. Regional-scale geostatistical inverse modeling of North American CO₂ fluxes: a synthetic data study. *Atmospheric Chemistry and Physics Discussions*, 9:22407–22458, 2009.
- [46] S. Mallat. *A wavelet tour of signal processing*. Academic Press, San Diego, CA, USA, 1999.
- [47] J. R. Williams and K. Amaratunga. Introduction to wavelet engineering. *International Journal for Numerical Methods in Engineering*, 37:2365–2388, 1994.
- [48] J. Romberg. Imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2):14–20, 2008. <http://www-dsp.rice.edu/cs>.
- [49] Emmanuel Candes and Michael Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, 2008.
- [50] E. Candes and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions Information Theory*, 52(12):5406–5425, 2006.
- [51] D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [52] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.
- [53] R. Coifman, F. Geshwind, and Y. Meyer. Noiselets. *Applied and Computational Harmonic Analysis*, 10(1), 2001.
- [54] Y. Tsaig and D. Donoho. Extensions of compressed sensing. *Signal Processing*, 86(3):533–548, 2006.
- [55] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [56] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.
- [57] R. Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society*, 58(1):267–288, 1996.
- [58] S. Mallat and Z. Zhang. Matching pursuit with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, 1993.
- [59] J. Tropp and A. C. Gilbert. Signal recovery from partial information via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, 2007.
- [60] D. L. Donoho, Y. Tsaig, I. Drori, and J-L Starck. Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit. March 2006.
- [61] Shihao Ji, Ya Xue, and Lawrence Carin. Bayesian compressive sensing. *IEEE Transactions on Signal Processing*, 56(6), 2008.
- [62] S. D. Babacan, R. Molina, and A. K. Katsaggelos. Bayesian compressive sensing using Laplace priors. *IEEE Transactions on Signal Processing*, 19(1), 2010.

- [63] M. F. Duarte, M. B. Wakin, and R. G. Baraniuk. Fast reconstruction of piecewise-smooth signals from random projections. In *Proceedings of Signal Processing with Adaptive Sparse Structured Representations*. Rennes, France, 2005.
- [64] C. La and M. N. Do. Signal reconstruction using sparse tree representation. In *Proc. Wavelets XI at SPIE Optics and Photonics*. San Diego, CA, USA, 2005.
- [65] Richard Baraniuk, Volkan Cevher, Marco Duarte, and Chinmay Hegde. Model-based compressive sensing. *IEEE Transactions on Information Theory*, 56(4):1982 – 2001, 2010.
- [66] L. He and L. Carin. Exploiting structure in wavelet-based bayesian compressive sensing. *IEEE Transactions on Signal Processing*, 57(9):3488–3497, 2009.
- [67] Global radiance calibrated nighttime lights. http://www.ngdc.noaa.gov/dmsp/download_radcal.html.
- [68] P. Cinzano, F. Falchi, C. D. Elvidge, and K. E. Baugh. The artificial night sky brightness mapped from DMSP satellite Operational Linescan System measurements. *Monthly Notices of the Royal Astronomical Society*, 318:641–657, 2000.
- [69] C. D. Elvidge, K. E. Baugh, J. B. Dietz, T. Bland, P. C. Sutton, and H. W. Kroehl. Radiance calibration of DMSP-OLS low-light imaging data of human settlements - a new device for portraying the Earth’s surface entire. *Remote Sensing of Environment*, 68:77–88, 1999.
- [70] M. R. Raupach, P. J. Rayner, and M. Page. Regional variations in spatial structure of nightlights, population density and fossil-fuel CO₂ emissions. *Energy Policy*, 38:4756–4764, 2009.
- [71] Atlas of the biosphere. <http://www.sage.wisc.edu/atlas/maps.php?datasetid=18&includerelatedlinks=1&dataset=18>.
- [72] Billie Miteva. billie_m@rocketmail.com. Personal communication.
- [73] International Geosphere-Biosphere Programme. <http://www.igbp.net/>.
- [74] J. C. Lin, C. Gerbig, S. C. Wofsy, A. E. Andrews, B. C. Daube, K. J. Davis, and C. A. Grainger. A near-field tool for simulating the upstream influence of atmospheric observations: the Stochastic Time-Inverted Lagrangian Transport (STILT) model. *Journal of Geophysical Research*, 108(D16), 2003. 4493, doi:10.1029/2002JD003161.
- [75] S. M. Gourdj, A. I. Hirsch, K. L. Mueller, A. E. Andrews, and A. M. Michalak. Regional-scale geostatistical inverse modeling of North American CO₂ fluxes: a synthetic data study. *Atmospheric Chemistry and Physics*, 10:6156–6167, 2010. doi:10.5194/acp-10-6151-2010.
- [76] W. C. Skamarock and J. B. Klemp. A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. *Journal of Computational. Physics*, 227:3465–3485, 2008.
- [77] T. Nehrkorn, J. Eluszkiewicz, S. C. Wofsy, J. C. Lin, C. Gerbig, M. Longo, and S. Freitas. Coupled Weather Research and Forecasting/Stochastic Time-Inverted Lagrangian Transport (WRF-STILT) model. *Meteorology and Atmospheric Physics*, 107:51–64, 2010. doi:10.1007/s00703-010-0068-x.
- [78] Kimberly L. Mueller. *A data-driven multi-scale statistical investigation of regional sources and sinks to improve knowledge of terrestrial carbon cycling*. PhD thesis, University of Michigan, Ann Arbor, MI, 2011. Environmental Engineering.

- [79] Emmanuel Cands and Justin Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23(3):969, 2007.
- [80] L. Gan, T.T. Do, and T.D. Tran. Fast compressive imaging using scrambled block Hadamard ensemble. *16th European Signal Processing Conference*, 2008.
- [81] Wotao Yin, Simon Morgan, Junfeng Yang, and Yin Zhang. Practical compressive sensing with Toeplitz and circulant matrices. *Signals*, 7744:77440K–77440K–10, 2010.
- [82] Tomas Tuma and Paul Hurley. On the incoherence of noiselet and Haar bases. In *International Conference on Sampling Theory and Applications*, 2009.
- [83] C. Gerbig, A. J. Dolman, and M. Heimann. On observational and modeling strategies targeted at regional carbon exchange over continents. *Biogeosciences*, 6:1949–1059, 2009.
- [84] R. J. Andres, T. A. Boden, F.-M. Bréon, et al. A synthesis of carbon dioxide emissions from fossil-fuel combustion. *Biogeosciences*, 9:1845–1871, 2012.
- [85] Geir Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99:10143–10162, 1994.
- [86] Geir Evensen. The Ensemble Kalman Filter for Combined State and Parameter Estimation. *IEEE Control Systems Magazine*, pages 83–104, June 2009.
- [87] L. M. P. Bruhwiler, A. M. Michalak, W. Peters, D. F. Baker, and P. Tans. An improved Kalman Smoother for atmospheric inversions. *Atmospheric Chemistry and Physics*, pages 2691–2702, 2005.
- [88] A. M. Michalak. Technical note: Adapting a fixed lag Kalman smoother to a geostatistical atmospheric inversion framework. *Atmospheric Chemistry and Physics*, pages 6789–6799, 2008.
- [89] R. S. Lokupitiya, D. Zupanski, A. S. Denning, S. R. Kawa, K. R. Gurney, and M. Zupanski. Estimation of global CO₂ fluxes at regional scale using the maximum likelihood ensemble filter. *Journal of Geophysical Research*, 113, 2008. D20110.
- [90] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, 1996.
- [91] Heikki Haario, Marko Laine, Antoinietta Mira, and Eero Saksman. DRAM-Efficient adaptive MCMC. *Statistics and Computing*, 16(4):339–354, 2006.
- [92] J. A. Vrugt, C. J. F. ter Braak, C. G. H. Diks, B. A. Robinson, J. M. Hyman, and D. Higdon. Accelerating markov chain monte carlo simulation by self-adaptive differential evolution with randomized subspace sampling. *International Journal of Nonlinear Scientific Numerical Simulation*, 10(3), 2009.
- [93] E. Laloy and J. A. Vrugt. High-dimensional posterior exploration of hydrologic models using multiple-try dream-zs and high-performance computing. *Water Resources Research*, 48, 2012. W01526.
- [94] Greg Welch and Gary Bishop. An introduction to the kalman filter. www.cs.unc.edu/~welch/media/pdf/kalman_intro.pdf, July 2006. Accessed: 2013-06-24.
- [95] B. Anderson and J. Moore. *Optimal Filtering*. Prentice Hall, 1979.

- [96] G. Einecke. Smoothing, filtering and prediction: Estimating the past, present and future. 2012.
- [97] W. Nowak. Best unbiased ensemble linearization and the quasi-linear kalman ensemble generator. *Water Resources Research*, 45, 2009.
- [98] Geir Evensen. *Data Assimilation: The Ensemble Kalman Filter, Second edition*. Springer-Verlag, New York, 2009.
- [99] Hemant A. Phale. *Data assimilation using the ensemble kalman filter with emphasis on the inequatlity constraints*. PhD thesis, University of Oklahoma, 2010.
- [100] Craig J. Johns and Jan Mandel. A two-stage ensemble Kalman filter for smooth data assimilation. *Environ Ecol Stat*, 15:101–110, 2008.
- [101] Yan Chen, Dean S. Oliver, and Dongxiao Zhang. Data assimilation for nonlinear problems by ensemble kalman filter with reparameterization. *Journal of Petroleum Science and Engineering*, 66:1–14, 2009.
- [102] Jia Li and Dongbin Xiu. On numerical properties of the ensemble Kalman filter for data assimilation. *Comput. Methods Appl. Mech. Engrg.*, 197:3574–3583, 2008.
- [103] K. R. Gurney, R. M. Law, A. S. Denning, P. J. Rayner, D. Baker, P. Bousquet, L. Bruhwiler, Y. H. Chen, P. Ciais, S. Fan, I. Y. Fung, M. Gloor, M. Heimann, K. Higuchi, J. John, T. Maki, S. Maksyutov, K. Masarie, P. Peylin, M. Prather, B. C. Pak, J. Randerson, J. Sarmiento, S. Taguchi, T. Takahashi, and C. W. Yuen. Towards robust regional estimates of CO₂ sources and sinks using atmospheric transport models. *Nature*, 415(6872):626–630, 2002.
- [104] P. P. Tans, I. Y. Fung, and T. Takahashi. Observational constraints on the global atmospheric CO₂ budget. *Science*, 247(4949):431–438, 1990.
- [105] P. J. Rayner, I. G. Enting, R. J. Francey, and R. Langenfelds. Reconstructing the recent carbon cycle from atmospheric CO₂, delta c-13 and o2/n2 observations. *Tellus Series B-Chemical and Physical Meteorology*, 51(2):213–232, 1999.
- [106] E. T. Olsen, M. T. Chahine, L. L. Chen, and T. S. Pagano. Retrieval of mid-tropospheric CO₂ directly from airs measurements. In S. S. Shen and P. E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XIV*. Bellingham: SPIE-International Society of Optical Engineering, 2008.
- [107] P. J. Rayner and D. M. O’Brien. The utility of remotely sensed CO₂ concentration data in surface source inversions. *Geophysical Research Letters*, 28(12):175–, 2001.
- [108] R. M. Duren and C. E. Miller. Measuring the carbon emissions of megacities. *Nature Climate Change*, 2(8):560–562, 2012.
- [109] E. A. Kort, C. Frankenberg, C. E. Miller, and T. Oda. Space-based observations of megacity carbon dioxide. *Geophysical Research Letters*, 39(17), 2012.
- [110] J. B. Miller et al. Linking emissions of fossil fuel CO₂ and other anthropogenic trace gases using atmospheric 14CO₂. *Journal of Geophysical Research*, 117(23), 2012.
- [111] L. Wu, M. Bocquet, T. Lauvaux, F. Chevallier, P. Rayner, and K. Davis. Optimal representation of source/sink fluxes for mesoscale carbon dioxide inversion with synthetic data. *Journal of Geophysical Research*, 116:D21304, 2011.

- [112] J. T. Randerson, M. V. Thompson, T. J. Conway, I. Y. Fung, and C. B. Field. The contribution of terrestrial sources and sinks to trends in the seasonal cycle of atmospheric carbon dioxide. *Global Biogeochemical Cycles*, 11(4):535–560, 1997.
- [113] G. R. van der Werf, J. T. Randerson, L. Giglio, G. J. Collatz, P. S. Kasibhatla, and A. F. Arellano Jr. Interannual variability in global biomass burning emissions from 1997 to 2004. *Atmospheric Chemistry and Physics*, 6:3423–3441, 2006.
- [114] T. A. Boden, G. Marland, and R. J. Andres. Global, regional, and national fossil-fuel co₂ emissions. In *Trends: A Compendium of Data on Global Change*. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tenn., U.S.A, 2012.
- [115] R. Nassar, L. Napier-Linton, K. R. R. Gurney, R. Andres, T. Oda, F. R. Vogel, and F. Deng. Improving the temporal and spatial distribution of co₂ emissions from global fossil fuel emission data sets. *Journal of Geophysical Research: Atmospheres*, 118(2):917–933, 2013.
- [116] J. C. Olsen and J. T. Randerson. Differences between surface and column atmospheric CO₂ and implications for carbon cycle research. *Journal of Geophysical Research – Atmosphere*, 109:D02301, 2004.
- [117] M. Rodell, P. R. Houser, U. Jambor, et al. The global land data assimilation system. *Bulletin of the American Meteorological Society*, 85:381–394, 2004.
- [118] T. Oda, S. Maksyutov, and Elvidge C. D. Disaggregation of national fossil fuel CO₂ emissions using a global power plant database and dmsp nightlight data. *Proceedings of the Asia Pacific Advanced Network*, 30:220–229, 2010.
- [119] V. Yadav, K. M. Mueller, and A. M. Michalak. A backward elimination discrete optimization algorithm for model selection in spatio-temporal regression models. *Environmental Modeling and Software*, 42:88–98, 2013.
- [120] A. M. Michalak and P. K. Kitanidis. A method for enforcing parameter nonnegativity in Bayesian inverse problems with an application to contaminant source identification. *Water Resources Research*, 39(2), 2003.
- [121] P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization*. Academic, San Diego, California, 1986.
- [122] P. K. Kitanidis. Quasi-linear geostatistical theory for inversing. *Water Resources Research*, 31:2411–2419, 1995.
- [123] J. Turnbull, J. Miller, S. Lehman, P. Tans, R. Sparks, and J. Southon. Comparison of (CO₂)-C-14, CO, and SF₆ as tracers for recently added fossil fuel CO₂ in the atmosphere and implications for biological co₂ exchange. *Geophysical Research Letters*, 33(1), 2006.

Appendix A

Glossary

Table A.1: Definitions of Greek symbols

ε	Measurement error
ε_2	Convergence tolerance for StOMP
ε_3	Convergence tolerance for non-negativity imposition
ϕ	Basis vector
Φ	Matrix of basis vectors, $N \times L$
$\tilde{\Phi}$	Matrix of basis vectors for K periods, $\text{diag}(\Phi)$
$\tilde{\Phi}_{\mathcal{R}}$	Basis matrix for modeling $\mathbf{F}_{\mathcal{R}}$, $(N_{\mathcal{R}}K) \times (LK)$
$\tilde{\Phi}_{\mathcal{R}'}$	Basis matrix for modeling $\mathbf{F}_{\mathcal{R}'}$, $(N_{\mathcal{R}'}K) \times (LK)$
$\tilde{\Phi}'_{\mathcal{R}}$	$\tilde{\Phi}_{\mathcal{R}}$ with columns multiplied by weights from $\mathbf{f}_{pr}$
$\tilde{\Phi}'_{\mathcal{R}'}$	$\tilde{\Phi}_{\mathcal{R}'}$ with columns multiplied by weights from $\mathbf{f}_{pr}$
$\tilde{\Phi}'_{\mathcal{R}'}$	Basis matrix for modeling $\mathbf{F}_{\mathcal{R}'}$, $(N_{\mathcal{R}'}K) \times (LK)$
Ψ	Random measurement vector for CS
$\mathbf{\Psi}$	Random projection matrix / Measurement matrix for CS
Υ	“Observations”; $\mathbf{Y}$ or $\Delta\mathbf{Y}$
Γ	$\mathbf{G}$ or $\mathbf{G}'$

Table A.2: Definitions of symbols

CS	Compressive Sensing
D	Domain size, along one axis
$\mathbf{E}$	Non-negative ffCO ₂ emissions, over K 8-day periods, in $\mathcal{R}$
$\mathbf{f}$	CO ₂ emission, over a spatial domain. This is not a time-series.
$\mathbf{f}_k$	CO ₂ emission averaged over a 8-day period k
$\mathbf{f}_{k,\mathcal{R}}$	ffCO ₂ emissions averaged over a 8-day period k inside $\mathcal{R}$
$\mathbf{F}_{\mathcal{R}}$	CO ₂ emissions time-series $\mathbf{F}_{\mathcal{R}} = \{\mathbf{f}_k\}, k = 1 \dots K$, in the grid-cells covering the Lower 48 states of US, $\mathcal{R}$; $(N_{\mathcal{R}}K) \times 1$
$\mathbf{F}_{\mathcal{R}'}$	CO ₂ emissions time-series in the grid-cells other than those covering the Lower 48 states of US; the complement of $\mathbf{F}_{\mathcal{R}}$; $(N_{\mathcal{R}'}K) \times 1$
$\mathbf{f}_{pr}$	Prior flux or a guess of what the $\mathbf{f}$ might be
$\mathbf{f}_V$	Vulcan emissions at 1 degree resolution, averaged over 2002
$\overline{\mathbf{f}_V}^{(s)}$	Vulcan emissions at 1 degree resolution, averaged over 2002, as represented by a sparse basis set
$\mathbf{f}_{V,k}$	Vulcan emissions at 1 degree resolution, but averaged over the k^{th} 8-day period
$\mathbf{g}$	A 1D signal
$\mathbf{g}'$	A measured 1D signal
$\mathbf{G}$	The gain matrix; $(K_s N_s + M_{cs}) \times (LK)$
$\mathbf{G}'$	The “normalized” gain’ matrix; $(K_s N_s + M_{cs}) \times (nzwK)$
$\mathbf{H}$	Transport model
K	Number of 8-day time-periods in a year, $K = 45$
K_s	The number of times a tower sensor senses in a year. Towers measure once every 3 hours
L	Number of non-zero weights in a wavelet decomposition; sparsity
M	The number of wavelet levels the grid can be decomposed to
M_{cs}	The number of random projections used to enforce zero flux outside $\mathcal{R}$
N	Number of grid cells; equal to $2^M \times 2^M$
$N_{\mathcal{R}}$	The number of grid-cells in $\mathcal{R}$
$N_{\mathcal{R}'}$	The number of grid-cells in $\mathcal{R}'$
N_s	Number of sensors
$\mathbf{R}$	Random projection matrix, used to enforce zero flux outside $\mathcal{R}$; $M_{cs} \times (N_{\mathcal{R}'}K)$
$\mathcal{R}$	Region of interest, lower 48 states of US
$\mathcal{R}'$	Region that is in the grid, but outside the lower 48 states of US
StOMP	Stagewise Orthogonal Matching Pursuit, [60]
US	United States
$\mathbf{w}$	Weights of a wavelet basis set
$\mathbf{w}'$	Weights of a wavelet basis set, normalized by weights of a guessed flux from a proxy
$\mathbf{W}^{(s)}$	Set of weights corresponding to the sparsified set of wavelet bases
$\mathbf{X}$	A proxy for CO ₂ emissions, e.g., nightlights
$\mathbf{X}^{(s)}$	Proxy, sparsified
$\mathbf{y}$	Modeled CO ₂ concentrations
$\mathbf{y}^{obs}$	Measured CO ₂ concentrations
$\mathbf{Y}$	LHS of optimization problem; $\mathbf{Y} = \{\mathbf{y}^{obs}, 0\}$; $(K_s N_s + M_{cs}) \times 1$

Appendix B

Measurement locations

This is a list of CO₂ measurement towers that were used in this study. They are targeted at biospheric CO₂ fluxes and their locations are not optimized for ffCO₂ emissions.

Table B.1: List of observation towers and their locations.

Tower symbol	Full name	Latitude & Longitude	Height above ground (meters)
AMT	Argyle	45.03 N, -68.68 W	107
ARM	Southern Great Plains	36.80 N, -97.50 W	60
BAO	Boulder Atmospheric Observatory	40.05 N, -105.01 W	300
BRW	Barrow	71.32 N, -156.61 W	17
CDL	Candle Lake	53.99 N, -105.12 W	30
CEN	Centerville	40.79 N, -92.88 W	110
CHI	Chibougamau	49.69 N, -74.34 W	30
CVA	Canaan Valley	39.06 N, -79.42 W	7
EGB	Egbert	44.23 N, -79.78 W	3
ETL	East Trout Lake	54.35 N, -104.99 W	105
FIR	Fir	44.65 N, -123.55 W	38
FRD	Fraserdale	49.88 N, -81.57 W	40
GAL	Galesville	44.09 N, -91.34 W	122
HDP	Hidden Peak, Snowbird	40.56 N, -111.65 W	18
HFM	Harvard Forest	42.54 N, -72.17 W	30
KEW	Kewanee	41.28 N, -89.97 W	140
LEF	Park Falls	45.95 N, -90.27 W	396
LJA	La Jolla	32.87 N, -117.26 W	5
LLB	Lac LaBiche	54.95 N, -112.45 W	10
MAP	Mary's Peak	44.50 N, -123.55 W	8
MEA	Mead	41.14 N, -96.46 W	122
MET	Metolius	44.45 N, -121.56 W	34
MOM	Morgan Monroe	39.32 N, -86.41 W	48
NGB	NGBER	43.47 N, -119.69 W	7
NWR	Niwot Ridge	40.05 N, -105.58 W	5
OZA	Ozark	38.74 N, -92.20 W	30
ROL	Round Lake	43.53 N, -95.41 W	110
SBL	Sable Island	43.93 N, -60.02 W	25
SCT	South Carolina Tower	33.41 N, -81.83 W	305
SNP	Shenandoah National Park	38.62 N, -78.35 W	17
SPL	Storm Peak Lab	40.45 N, -106.73 W	9
WBI	West Branch	41.73 N, -91.35 W	379
WGC	Walnut Grove	38.27 N, -121.49 W	483
WKT	Moody	31.32 N, -97.33 W	457
YAH	Yaquina Head	44.67 N, -124.07 W	13

Appendix C

Second-best covariates

Table C.1: Second-best covariates and covariance model in the real-data case from January to December 2008. Note: The asterisk denotes a covariate.

Month	Mean nightlight*	Mean population density*	Percent built-up area*	Percent urban area*	Vulcan & ODIAC*	Best covariance model
1					Yes	Variance, pop. density
2					Yes	Variance, pop. density
3	Yes			Yes	Yes	Variance, pop. density
4	Yes		Yes			Variance, nightlights
5		Yes				Variance, built up area
6		Yes				Mean, built up area
7		Yes				Mean, built up area
8			Yes			Max. nightlight
9					Yes	Max. pop. density
10	Yes				Yes	Variance, built up area
11	Yes		Yes			Max. nightlights
12	Yes				Yes	Max. nightlights

DISTRIBUTION:

1	Jaideep Ray, 08954	MS 9159
1	B. van Bloemen Waanders, 01442	MS 1318
1	J. Lee, 08112	MS 9406
1	S. Lefantzi, 08954	MS 9152
1	Technical Library, 08944 (electronic)	MS 0899