MethaneSat: Detecting Methane Emissions From the Barnett Shale Region

The Harvard community has made this article openly available. **Please share** how this access benefits you. Your story matters

Acknowledgements

This senior thesis project is most definitely my proudest accomplishment from my Harvard career, and one of the most enjoyable experiences I have had thus far. For this, I am deeply grateful.

This project would not have been possible without the contributions of several individuals. I thank Arlyn Andrews[1], Kirk Thoning[1], and Joshua Benmergui[2] for their development of the inversion framework that I used in this experiment. I thank Marikate Mountain[3] and Thomas Nehrkorn[3] for developing the STILT footprints I used in my OSSEs. I thank Steven Wofsy[2,4], Kelly Chance[5], and Xiong Xiu[5], for the development of MethaneSat. I thank Joannes Maasakkers[2] for creating the emission files containing the Barnett Shale region inventory. I thank HUCE and HCRP for providing me with the financial means to conduct my research this past summer.

Additionally, I thank Cody Floerchinger, Yanina Berrera, and the rest of the Wofsy group for welcoming me into their community, teaching me about their research, and providing excellent conversation over 3 pm cookies. I thank the EPS senior thesis writers for welcoming me into the EPS senior thesis tutorial, and the EPS preceptors for assisting me throughout this process. I thank my friends for bringing me food when I was working for long hours, and of course my parents for their endless support.

I cannot express enough thanks to my mentor, Joshua Benmergui. Not only did he teach me so much about statistics, computer programing, atmospheric science, and more, but he helped me develop confidence in my research abilities. His endless knowledge and excitement about his research are inspiring, and I am grateful to have had the opportunity to work with him.

Finally, I would like to thank my advisor, Professor Steven C. Wofsy, for welcoming me into his lab, directing me towards a project that so closely fit my interests, and making this entire experience possible. His guidance over the past year and a half has been indispensable. More, he has created a research environment that is stimulating, supportive, and inspiring - one that I will certainly miss when I leave the Harvard community.

---

[1]National Oceanic and Atmospheric Administration
[2]School of Engineering and Applied Sciences, Harvard University
[3]Atmospheric and Environmental Research
[4]Department of Earth and Planetary Sciences, Harvard University
[5]Smithsonian Astrophysical Observatory. Harvard-Smithsonian Center for Astrophysics.

Contents

List of Figures

List of Tables

Introduction

As concern surrounding climate change grows, there will be increased efforts to curb environmentally harmful activity worldwide. In the United States, there has already been a serious push in this direction. In his 2013 Climate Action Plan, President Obama set forth goals to cut carbon pollution, lead international efforts to address climate change, and prepare the country for its impact (White House, Executive Office of the President [2013]). As part of this plan, Obama called for reductions in methane (CH_4) emissions from the oil and gas industry by 40-45% from their 2012 levels by the year 2025. In order to track progress towards this goal and similar plans globally, it will be necessary to advance methods of measurement, monitoring, and reporting of greenhouse gases.

In this project, we explore MethaneSat as a possible solution to the gap in currently available methods of monitoring methane emissions. This chapter discusses the motivation and provides context for the project. First, we consider the benefits of methane as a target for climate change mitigation. This discussion will highlight the need for deep understanding and extensive monitoring of global methane emissions. We then review currently available methods of detecting methane emissions, identifying their strengths and short-

comings. Finally, we introduce MethaneSat, our proposed solution to the identified gap in understanding methane emissions.

1.1 Motivation

Methane is an ideal target for climate change mitigation efforts for several reasons.

Methane is an extremely important greenhouse gas, second only to carbon dioxide. The importance of a greenhouse gas is dependent on three factors: abundance, radiative efficiency, and lifetime. Methane has an abundance of about 1850 ppb, making it the most abundant long-lived greenhouse gas in the atmosphere after carbon dioxide (CO_2) (Montzka et al. [2011]). The radiative efficiency of methane is $3.7\,e{-}4$ W m^{-2} ppb^{-1}, even higher than that of carbon dioxide ($1.5\,e{-}5$ W m^{-2} ppb^{-1}) (Intergovernmental Panel on Climate Change [2007]). The lifetime of methane is just under 10 years (Prinn et al. [1995]). This is short, relative to carbon dioxide; half of carbon dioxide is short lived, while the other half has a lifetime of about 10,000 years (Howarth et al. [2011]). This gives rise to two arguments about methane's relative importance. Some argue that its short lifetime makes methane less important in the long run, the timescale on which climate change will occur (Shoemaker and Schrag [2013]). However, methane already accounts for 25% of anthropogenic radiative forcing; its short lifetime therefore means that efforts to decrease emissions in the coming years could not only slow increases in radiative forcing, but reverse them. Thus, due to its abundance, radiative efficiency, and lifetime, CH_4 is an important greenhouse gas worthy of great attention.

There is strong evidence that CH_4 is becoming more abundant in the atmosphere (Figure 1). Ice core observations indicate that methane has nearly tripled in atmospheric concentration since pre-industrial times, implying a large anthropogenic component (Hausmann et al. [2016]; Schaefer et al. [2016]). Although the global average methane concentration did plateau after the year

GLOBAL MONTHLY MEAN CH₄
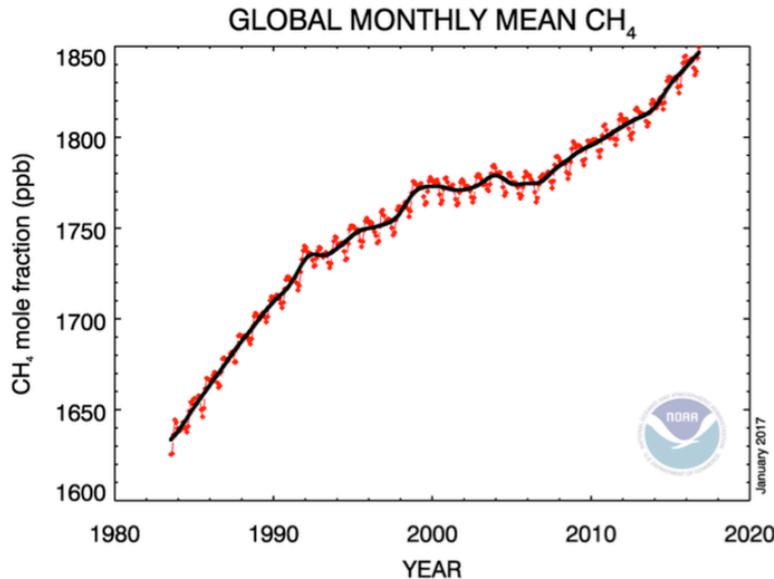
Figure 1: Global monthly CH_4 mole fraction (ppb). Atmospheric concentrations of CH_4 rose steadily until about 2000, then plateaued, then began to increase again around 2006. Source: NOAA Trends in Atmospheric Methane, 2017.

2000, there has been a renewed increase in CH_4 concentrations since 2006 (Dlugokencky et al. [2009]; Rigby et al. [2008]). There exist several theories about why and how this increase is occurring, made more complicated by the fact that methane is produced from both natural and anthropogenic sources. Many predict that methane emissions will continue to increase from both types of sources due to climate change and population growth (Montzka et al. [2011]). There is thus a clear need for mitigation of methane emissions, but also a need for better understanding of why methane concentrations are rising. Locating

and characterizing the sources of atmospheric methane will provide important insight into the causes of its recent increase in concentration, which, in turn, will allow for targeted mitigation strategies to be formulated.

In particular, the recent boom and projected growth of the shale gas industry has important implications for global methane emissions. Since hydraulic fracturing and horizontal drilling were proven successful in the Barnett Shale region in the 1990s, shale gas has grown dramatically as a proportion of natural gas production in the United States (Figure 2)(Britannica Academic [2016]; U.S. Energy Information Administration [2017]). For various economic,
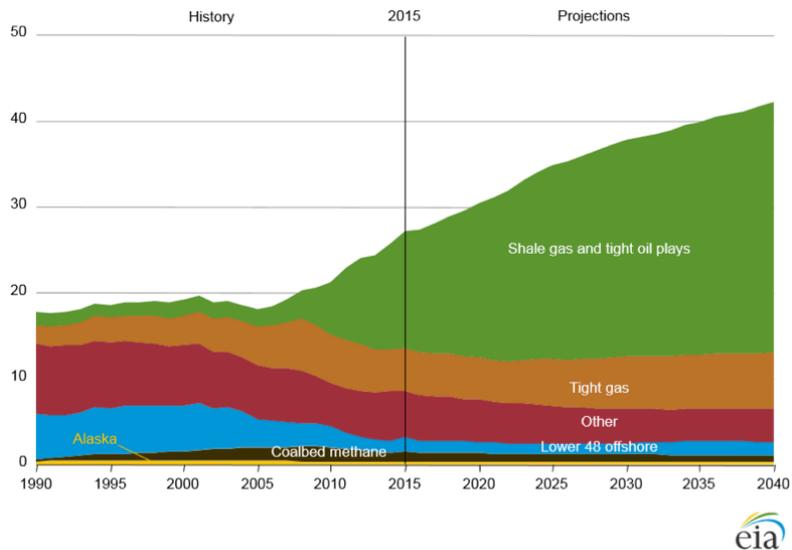
Figure 2: U.S. dry natural gas production by source, 1990-2040 (trillion cubic feet). Shale gas has grown dramatically as a proportion of natural gas production in the United States over the last decade, and this trend is expected to continue. Source: U.S. Energy Information Administration [2017].

political, and logistical reasons, shale gas has not yet seen such a boom elsewhere in the world. However, with predicted shale gas resources in 14 regions outside the US, comprising 32 countries, the industry is likely to grow on a global scale (Britannica Academic [2016]).

While natural gas is often promoted as a bridge fuel or green energy source

due to its low CO_2 emissions (U.S. Energy Information Administration [1999]), this ignores the role of CH_4, the major component of natural gas (Britannica Academic [2016]). According to Howarth et al. [2011], significantly more methane escapes as fugitive emissions from shale gas production than from conventional gas production. Due to the high global warming potential of CH_4, these emissions could significantly contribute to climate change. For this reason, it is essential that methane emissions from shale gas plays be closely monitored in order to track their impact on the global methane budget, and also to locate leaks, if possible.

In the United States, this is possible through various observing systems, including in situ networks. However, as the gas industry is expanding on a global scale, including in regions with less developed scientific programs, these observing systems are no longer sufficient. A tool capable of monitoring methane on a global scale and differentiating between sources will be necessary in order to monitor the impact of the shale gas industry on the global methane budget.

Critically, CH_4 is a realistically feasible target for climate change mitigation efforts due to support from both the private and public sectors.

The oil and gas industry is financially motivated to reduce its contribution to the global methane budget. The EDF estimated that \$2 billion worth of natural gas is lost each year from leaks and intentional releases (Environmental Defense Fund [2016]). Methane is the largest component of natural gas, so these leaks are likely a large source of anthropogenic methane. Locating and

characterizing these leaks would benefit private oil and gas companies as well as the environment.

There is also growing support in the public sector for methane mitigation initiatives. Several states have established policies to monitor or regulate methane emissions (National Conference of State Legislatures [2014]). Furthermore, the short lifespan of CH_4 could encourage political action, as emission reductions today could impact greenhouse gas concentrations in the near future. An improved ability to quantify the impact of natural gas on the global methane budget will provide further support for policy seeking to reduce methane emissions.

1.2   Existing Observing Systems

Greenhouse gas emissions are estimated through bottom-up or top-down studies. Bottom-up estimates are based on site-level emissions measurements, while top-down estimates are based on atmospheric concentration measurements and chemical transport models (Harriss et al. [2015]). Bottom-up estimates are capable of providing source attribution, and are less likely to be biased than top-down estimates. However, top-down estimates can be less labor-intensive and provide greater coverage, and are therefore often used to acquire greenhouse gas emission estimates on a global scale.

Atmospheric concentration data for greenhouse gases can be obtained through several methods. Methane is most often observed through long-term surface observation networks, surface campaigns, flight campaigns, or with satellites, each method having unique benefits and drawbacks. Factors such as spatial

resolution, measurement density, and measurement precision affect the resulting emissions estimates, and can vary greatly between observing systems.

At continental scales, flask and in-situ networks, as well as the satellites GOSAT and soon-to-be-launched TROPOMI, can estimate emissions. These provide high precision and long term observations. While this is valuable, these systems lack the spatial resolution to discriminate between emission sources (Brasseur and Jacob [2015]). Because methane has both natural and anthropogenic sources, this is problematic when attempting to quantify emissions from oil and gas production sites. Furthermore, flask and in-situ observations are sparse outside of North America and Europe.

At the regional scale (less than 1000 km), measurements from field campaigns can be used to estimate methane emissions. Several studies (e.g., Goetz et al. [2015]; Peischl et al. [2016]; Karion et al. [2015]) have recently been completed using this strategy. These campaigns constrain emissions from particular sources at a particular point in time, allowing for some discrimination of source types. This is extremely valuable for policymakers, and for oil and gas producers looking to reduce their carbon footprint. However, these campaigns are extremely labor intensive and require access to sites, which is impossible in many important regions.

At local scales (less than 200 km), intensive ground and aircraft campaigns can be used to estimate methane emissions. These measurements are generally highly accurate and can be taken over long periods of time (Brasseur and Jacob [2015]). While these campaigns are common in North America and Europe, they are rare in other regions for which we are interested in quantifying methane emissions from oil and gas production, such as Russia, Africa, and

the Middle East.

GHGSat is a private company that recently launched the CLAIRE satellite in the summer of 2016. It focuses on measuring greenhouse gas emissions from individual industrial facilities for customers who order measurements. CLAIRE has extremely high resolution of 100 m, but a low precision of 1-5% (Fairley [2016]). This precision on a total column measurement translates into 18-90 ppb, which is equivalent to a large methane plume signal. It will thus be difficult to rigorously constrain emissions with this instrument for all but the largest emitters. There is great value in measuring emissions from individual facilities, but these results cannot be extrapolated to be broadly applicable. Moreover, the data are not publicly available, and thus not applicable to policy decisions.

Therefore, there is a major gap in available observing systems' abilities to quantify methane emissions from oil and gas production sites around the globe, and to discriminate them from other sources. We hope to address this gap using a proposed satellite, MethaneSat.

1.3   Our Proposed Solution: MethaneSat

MethaneSat, a proposed satellite for measuring methane emissions, is a large-scale collaboration by the Wofsy group, the Smithsonian Astrophysical Observatory, Atmospheric and Environmental Research (AER), and Environmental Defense Fund (EDF). In order to obtain observations with global reach that are still capable of distinguishing between source types, a satellite is the most promising tool. MethaneSat will measure the total column dry-air mole

fraction of methane, exclusively, and will not be tasked with measuring the concentrations of other gases. It will target regions with sides of 200-300 km in length, with 1x1 km or 2x2 km spatial resolution – 10-100 times higher than TROPOMI – with high precision (0.1-0.2% random error).

MethaneSat's mission will most likely take place over a span of 12-18 months, plus six months for emission analysis. In that time, MethaneSat will sample about 100 target sites around the globe, each of which will be sampled up to three times. The satellite sensor will initially likely be capable of imaging up to 10 or more sites per day. This provides the potential for expansion of the project, increasing the number of target sites to up to 4000 per year.

Part of the strategy of MethaneSat is to minimize costs by using existing technology such as the Bella Terra, a Google SkyBox satellite. This will minimize the amount of custom design required. In the Phase 0 feasibility study, completed by Steven C. Wofsy, Kelly Chance, and Xiong Xiu in November 2015, it was estimated that the end-to-end cost of the satellite would be about $30 million on orbit, including a year of data collection.

Phase A, the design stage, is the next step. This stage involves making important design decisions, predicting cost and scheduling, establishing mission requirements, and predicting its final performance. The project discussed in this paper is meant to complement the planned Phase A study for MethaneSat. By performing observation system simulation experiments (OSSEs), we intend to quantify the limitations associated with methane emission estimates using existing observing systems, and the new information that will be provided by MethaneSat.

Methods

We performed observation system simulation experiments (OSSEs) to demonstrate the new information that MethaneSat will provide. OSSEs test the ability of an observing system, consisting of an instrument and inverse model, to constrain emissions, given a reasonable model of error. We performed these OSSEs for the Barnett Shale region of Texas, and tested the National Oceanic and Atmospheric Administration (NOAA) surface and aircraft network, Greenhouse Gases Observing Satellite (GOSAT), and MethaneSat. We hypothesized that the OSSEs would illustrate the advantages of using MethaneSat, as compared to available observations, to constrain emissions from oil and gas production regions.

This section will provide a brief overview of the OSSEs and the experimental setup. For a more detailed explanation of OSSEs, the reader may consult the Cooperative Institute for Meteorological Satellite Studies (CIMSS) and the National Aeronautics and Space Administration (NASA) (The Cooperative Institute for Meteorological Satellite Studies [1999]; National Aeronautics and Space Administration [2012]).
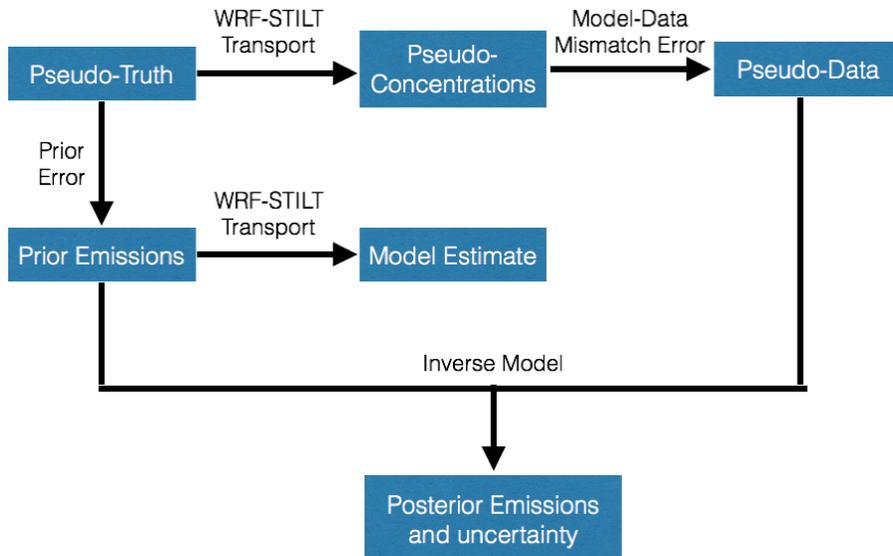
Figure 3: OSSE flow diagram. An OSSE derives a posterior emission estimate using a prior emission estimate and simulated observations in a Bayesian inversion model framework.

2.1 Observation System Simulation Experiments (OSSEs)

OSSEs provide information about how well an observing system can constrain emissions. They attempt to retrieve "target" emissions, or "pseudo-truth," through a Bayesian inversion framework, using a prior emission estimate and simulated observations of atmospheric mole fractions (in this case, of CH_4) (Figure 3). OSSEs are frequently conducted to evaluate proposed observing systems and determine their ability to deliver new information (Hungershoefer et al. [2010], Houweling et al. [2016], Brasseur and Jacob [2015]).

The inversion improves upon our prior knowledge of the unknown fluxes by reconciling the prior estimate with the simulated observations, using an atmospheric transport model to relate fluxes to atmospheric mole fractions.

The prior estimate regularizes the inverse model, preventing over-fitting of the data. The OSSE's best attempt to recover the target is called the "posterior" emission estimate. We can compare the posterior estimate with the target to predict the ability of the observing system to retrieve emissions.

Both the prior estimate and simulated observations are generated using the target emissions and realistic models of error. However, OSSEs are optimistic in assuming unbiased error, random sampling of errors, and full knowledge of error correlations, and will therefore produce better posterior estimates than can be expected in real life. However, OSSEs are still extremely useful as a tool for comparing observing systems' abilities to constrain emissions (Brasseur and Jacob [2015]).

2.1.1 Our OSSE Setup

We used the Stochastic Time-Integrated Lagrangian Transport (STILT) model (Lin et al. [2003]) to infer the sensitivities of receptors to upstream fluxes. This information allows us to relate fluxes and concentrations in a linear expression,

$$\mathbf{z} = \mathbf{Hs} + \mathbf{b} + \boldsymbol{\epsilon}, \tag{1}$$

where $\mathbf{s}$ are the unknown fluxes, $\mathbf{z}$ are the known concentrations, $\mathbf{H}$ is a Jacobian matrix containing the receptor sensitivity information from STILT "footprints", $\mathbf{b}$ are the background concentrations, and $\boldsymbol{\epsilon}$ is model-data mismatch error. While we do not explicitly simulate the estimation of background concentrations in this study, the expected error is included in the model-data mismatch error.

12
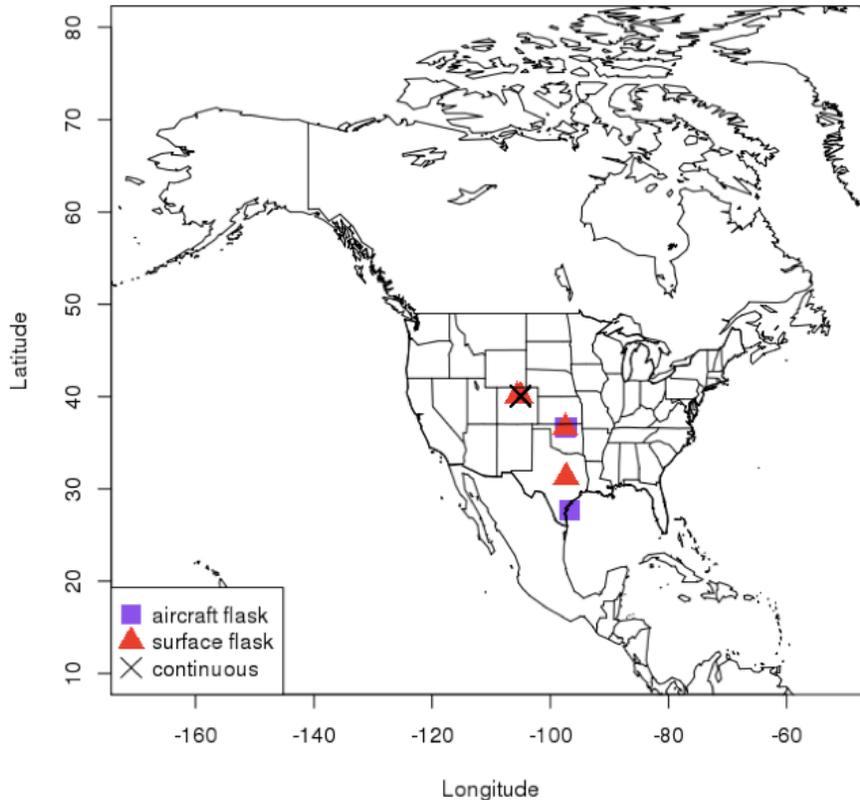
Figure 4: NOAA continuous flask and in-situ receptor sites used in this experiment. We chose these receptors due to their proximity to the Barnett Shale region.

The STILT model simulates sending particles backwards in time and space from a receptor location, guided by a meteorological model. In this experiment, we used two meteorological models to guide a 500-particle ensemble: WRF (Weather Research Forecast) for the NOAA network and GOSAT OSSEs, and GDAS (Global Data Assimilation System) for the MethaneSat OSSEs. The WRF fields were run as part of the NASA Carbon Monitoring System project, "CarbonTracker Lagrange." Joshua Benmergui[1] derived and ran the algorithms to integrate GDAS fields, available online, into footprints prior to

―――――――――――――――――――――
[1]School of Engineering and Applied Sciences, Harvard University

this experiment. While it would have been ideal to use the same transport fields for all three observing systems, both are representative of the information contained in each system, and the difference should have a minimal impact on the experiment. The resulting particle densities provided "footprints," or sensitivities of observations at the receptors to upstream fluxes, given with units of $\frac{ppm}{\mu mol \cdot m^{-2} \cdot s^{-1}}$. We used footprints for the NOAA sites surrounding the Barnett Shale region (Figure 4), and for all available GOSAT and MethaneSat retrievals in our time period of interest.

The "target" emissions map, also called the "pseudo-truth" is the emissions map that the OSSE seeks to recover through assimilation of a prior estimate and simulated observations. In this experiment, the target is a modified version of a CarbonTracker-Lagrange methane posterior for October 2013, developed by Joshua Benmergui[1] and Arlyn Andrews[2] (Figure 5). We replaced the emissions in the Barnett Shale region and its immediate surroundings, bounded by -102° to -95° longitude and 30° to 35° latitude, with an inventory for October 2013 based on the results of Lyon et al. [2015], updated by Zavala-Araiza et al. [2015]. The Lyon et al. [2015] study used a combination of available data, including bottom-up measurements from the Barnett Coordinated Campaign. This inventory implies higher emissions from the Barnett Shale region than other currently available inventories, such as those produced by the EPA, but has been directly validated by an intensive aircraft campaign (Karion et al. [2015]). The inventory is resolved with 4km x 4km grid cells, which we aggregated into 1° x 1° grid cells. The updated Zavala-Araiza emission files report total methane emissions in the region to be 0.7483 Tg year^{-1}.

---

[2]National Oceanic and Atmospheric Administration
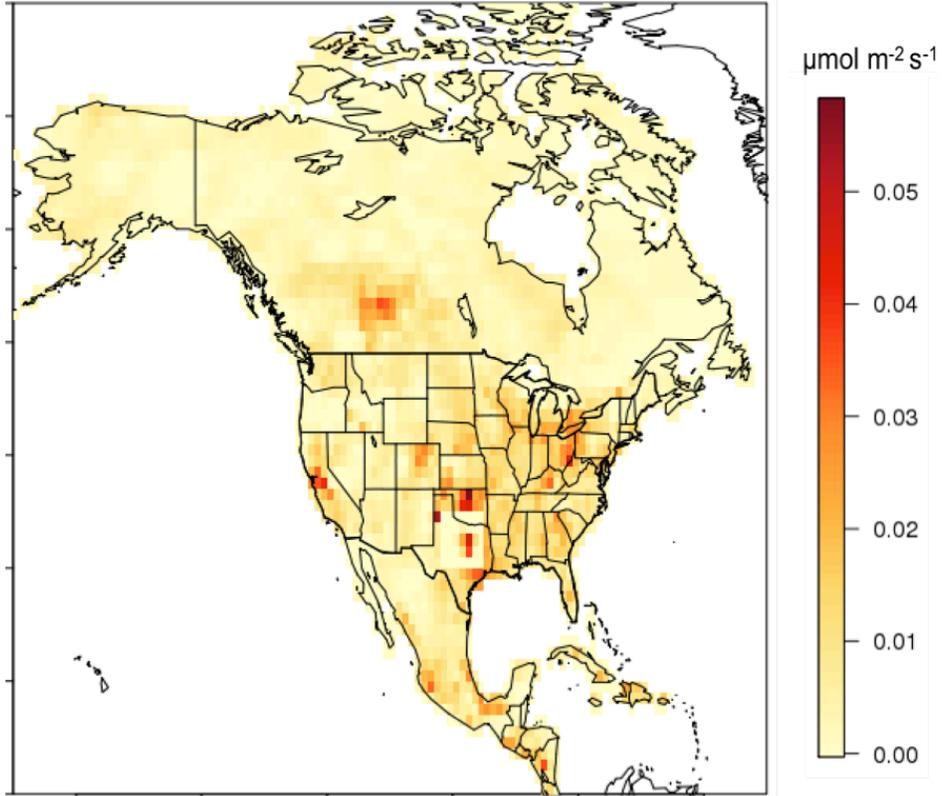
Figure 5: Target emissions map used in this experiment, combining a CarbonTracker-Lagrange methane posterior with the Barnett Shale inventory developed by Zavala-Araiza et al. [2015].

We created our prior emissions estimate by perturbing these target emissions. The perturbation was a Gaussian random field with mean zero, accounting for variance and covariance in time and space. The covariance matrix can be written as

$$\mathbf{Q} = \sigma_Q^2 \exp(\frac{\mathbf{X}_\tau}{l_\tau}) \otimes \exp(\frac{\mathbf{X}_s}{l_s}) \qquad (2)$$

where σ_Q^2 is the variance in time and space, $\mathbf{X}_\tau$ and $\mathbf{X}_s$ represent the separation distances/lags between estimation locations in time and space, respectively, and l_τ and l_s are the temporal and spatial correlation range parameters, respectively (Yadav and Michalak [2013]). We used 36 days as the temporal correlation parameter, and 100 km as the spatial correlation parameter. We
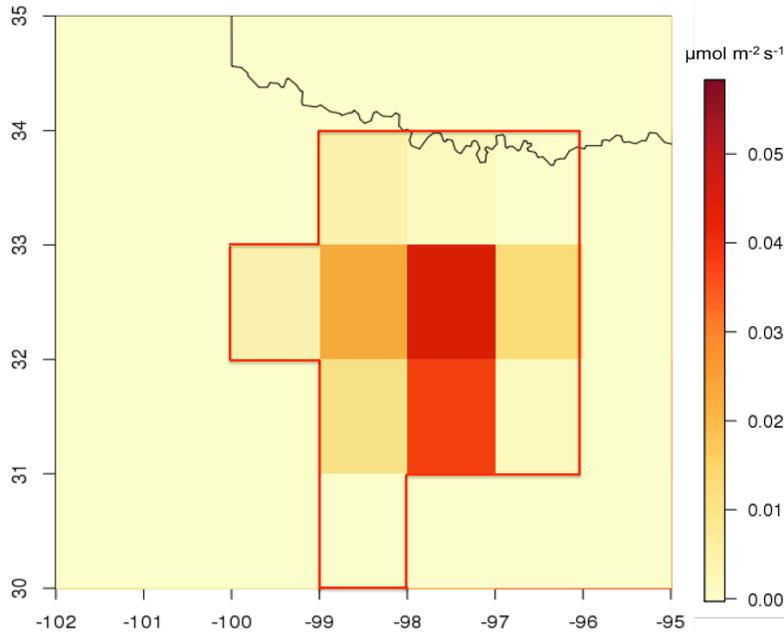
Figure 6: A closeup of the target emissions for the Barnett Shale region, taken from the Zavala-Araiza inventory. The 11 grid cells comprising the Barnett Shale region are outlined in red. Total emissions for these cells are 0.7483 Tg year^{-1}.

performed OSSEs for five distinct values of σ_Q^2, from 0.005 to 0.00005 μmol m^{-2}s^{-1}, with step sizes of half an order of magnitude (Table 1).

Table 1: Covariance matrix parameters.

Parameter	Value
σ_Q^2 (μmol m^{-2}s^{-1})^2	0.005 - 0.00005
l_s (km)	100
l_τ (days)	36
σ_R^2 (ppm^2)	0.01 - 0.0001

We chose these parameters to reflect a reasonable model of error in a prior estimate of emissions, ranging from an excellent prior that provides a great deal of information (error of about 0.007 μmol m^{-2} s^{-1}), to a useless prior that provides no information (error of about 0.07 μmol m^{-2} s^{-1}). In order to keep the prior estimate physically meaningful, we truncated the prior at zero, converting all

16

negative values to zero. Although it is unlikely that the prior error would actually be distributed as a Gaussian, it is reasonable to approximate it as such; the prior estimate for each grid cell is as likely to have positive error or negative error, thus rendering the prior emissions map unbiased. However, a limitation is that truncation of the prior estimate at zero means that total error in the negative direction may be less than total error in the positive direction, increasing the likelihood that the prior estimate is positively biased. The risk of positive bias grows with prior error variance.

For this experiment, we created two types of prior estimates. In Method A, we perturbed just the Barnett Shale region, and left the rest of the domain perfectly specified by the target. In Method B, we perturbed the entire domain. In both cases, if the perturbation exceeded the target emission rate in a grid cell, the value of the prior at that location was truncated at zero. This strategy leaves room for improvement, as negative values do not necessarily need to be removed in the context of satellite data. Other potential methods for generating a prior estimate are considered in the Conclusion.

Method A is a contrived numerical experiment, as it is unrealistic to assume perfect prior information for the domain outside the Barnett Shale region. However, we hypothesized that it would provide interesting insights to complement Method B, isolating the effect of error on a hotspot of emissions. Method B is a realistic numerical experiment, in which we seek to detect a hotspot of emissions in a noisy domain. We expected that this would reveal the difficulties of constraining a small region of interest with sparse data.

We generated the pseudo-data for each observing system by applying the corresponding STILT footprints to the target emissions map, and adding error

17

sampled from the model-data mismatch covariance matrix, $\mathbf{R}$, with diagonal elements σ_R^2. In this experiment, we used five distinct values for the error variance, σ_R^2, from 0.01 ppm^2 to 0.0001 ppm^2, corresponding to error between 10 ppb and 100 ppb (Table 1). This interval of error is similar to, though slightly wider than, that used by Miller et al. [2014a].

The posterior estimate, $\hat{s}$, is found by minimizing a cost function that penalizes deviation from both the prior estimate, $\mathbf{s}_p$, and pseudo-data, $\mathbf{z}$, according to the error on each:

$$L_s = \frac{1}{2}(\mathbf{z} - \mathbf{Hs})^T \mathbf{R}^{-1}(\mathbf{z} - \mathbf{Hs}) + \frac{1}{2}(\mathbf{s} - \mathbf{s}_p)^T \mathbf{Q}^{-1}(\mathbf{s} - \mathbf{s}_p) \qquad (3)$$

Thus, when the error on the prior estimate, $\mathbf{Q}$, is large and the error on the pseudo-data, $\mathbf{R}$, is small, the inversion favors a posterior estimate that would agree well with the pseudo-data. When the error on the prior estimate is small but the error on the pseudo-data is large, the inversion favors a posterior estimate close to the prior estimate. The posterior estimate, $\hat{s}$, is the solution to the inverse problem, and is the emissions estimate that minimizes the cost function:

$$\hat{s} = \mathbf{s}_p + (\mathbf{HQ})^T(\mathbf{HQH}^T + \mathbf{R})^{-1}(\mathbf{z} - \mathbf{Hs}_p) \qquad (4)$$

with error covariance matrix:

$$\mathbf{V}_{\hat{s}} = \mathbf{Q} - (\mathbf{HQ})^T(\mathbf{HQH}^T + \mathbf{R})^{-1}(\mathbf{HQ}) \qquad (5)$$

Because the second quantity in $\mathbf{V}_{\hat{s}}$ is positive, the error on $\hat{s}$ ($\mathbf{V}_{\hat{s}}$) is strictly less than the error on $\mathbf{s}_p$ ($\mathbf{Q}$).

Michalak [2008], as cited in Miller et al. [2014a], discusses the possibility that allowing for unrealistic negative emissions in the posterior estimate may distort the surrounding emissions due to conservation of mass. In order to investigate this, we also calculated a posterior estimate using a non-negativity algorithm, developed by Scot Miller[3] and Joshua Benmergui. This algorithm uses an iterative solver (L-BFGS-B) to find the minimum of the cost function under the constraint of positive emissions. A limitation of this method is that it may lead to overestimation of emissions due to its strategy of correcting negative fluxes. For this reason, we performed the OSSEs using both the regular algorithm and the non-negativity algorithm throughout the study.

2.2 Performing the OSSEs

In order to perform the OSSEs, we used a Python inversion framework developed by Arlyn Andrews[4], Kirk Thoning[5], and Joshua Benmergui. We made several modifications, customizing it for our experimental setup.

The large computational cost of the OSSEs required great consideration. Each inversion involves computation with an emissions map for the entire North American region with grid size of $1°$ x $1°$, for up to 91 time steps and tens of thousands of measurements. We ran the OSSEs on Odyssey, Harvard's supercomputer cluster in Holyoke, Massachusetts, to gain access to sufficient storage capacity and computational power. Even so, efficiency had to be maximized within the inversion framework. For example, explicit calculation of **HQ** was impossible, even on Odyssey. Thus, the framework

---

[3]Department of Global Ecology, Carnegie Institution for Science, Stanford University
[4]National Oceanic and Atmospheric Administration
[5]National Oceanic and Atmospheric Administration

created Jacobian "slices," breaking H into pieces for each time step, and calculating **HQ** in blocks, according to the method proposed by Yadav and Michalak [2013]. Similarly, explicit calculation of a Gaussian random field to perturb the entire North American region was not possible. In this case, we created realizations of the perturbation matrix in R, using the package 'RandomFields' (Schlather et al. [2015]).

For the NOAA network and GOSAT, we performed inversions for 91 time steps, between September 1 2012 and November 30 2012. We did this partly because we wanted a buffer on either side of the Lyon et al. [2015] inventory time period (October 16 - 30 2013). Unfortunately, however, we had to use footprints for 2012, as none were available for 2013. For MethaneSat, we performed inversions for the one time step available, September 19 2013. However, this single time step provided more pseudo-data than all 91 time steps of the NOAA network footprints or GOSAT footprints. This reflects a key aspect of the strategy of MethaneSat relative to satellites like GOSAT. MethaneSat collects many high precision observations in a single shot, while GOSAT collects fewer observations over a longer period of time (Table 2). Ideally, we would have used footprints from the same time period for all three observing systems, but this was not possible at the time of the study. It is unlikely, however, that the differences between 2012 and 2013 or September and October would make a large impact on the OSSEs' ability to demonstrate the new information gained from MethaneSat.

Table 2: Simulated observations for use in the OSSEs.

Observing System	Time Period of Footprints	Number of Observations
NOAA Network	9/1/2012 - 11/30/2012	552
GOSAT	9/1/2012 - 11/30/2012	2116
MethaneSat	10/19/2013	37975

2.3   Experimental Setup

First, in order to verify that the inversion framework operated as expected, we performed an OSSE without error on the prior estimate or error on the pseudo-data. This OSSE was an attempt to recover an emissions map using that very emissions map and the corresponding modeled concentrations. Mathematically, it can be determined that the posterior estimate will perfectly match the target under these circumstances, as $\mathbf{z} = \mathbf{H}\mathbf{s}_p$.

$$\hat{s} = \mathbf{s}_p + (\mathbf{H}\mathbf{Q})^T(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}(\mathbf{z} - \mathbf{H}\mathbf{s}_p)$$

$$\hat{s} = \mathbf{s}_p + \mathbf{0}$$

$$\hat{s} = \mathbf{s}_p$$

which, in this case, means the posterior emissions equal the target emissions, $\mathbf{s}_t$, as there was no prior error. Thus,

$$\mathbf{s}_p = \mathbf{s}_t + \boldsymbol{\epsilon} = \mathbf{s}_t$$

$$\hat{s} = \mathbf{s}_p = \mathbf{s}_t$$

This OSSE served as a test for the inversion framework. This type of analysis would likely reveal a mathematical error or bug in the code.

Next, we considered an OSSE with error on the pseudo-data but no error on the prior estimate. This translates to perfect prior information, and flawed

observations. Mathematically, it can again be determined that the posterior estimate will perfectly match the target under these circumstances, as the prior error covariance matrix, $\mathbf{Q}$, will be the zero matrix.

$$\hat{\boldsymbol{s}} = \mathbf{s}_p + (\mathbf{HQ})^T(\mathbf{HQH}^T + \mathbf{R})^{-1}(\mathbf{z} - \mathbf{Hs}_p)$$

$$\hat{\boldsymbol{s}} = \mathbf{s}_p$$

Besides serving as another check on the inversion system, this type of OSSE does not offer any new information.

We also considered an OSSE with error on the prior estimate but no error on the pseudo-data. This would be equivalent to having perfect observations, but a flawed prior estimate. Unlike the previous two types of OSSEs, this type does not mathematically reduce to $\hat{\boldsymbol{s}} = \mathbf{s}_p$. We performed these OSSEs with a near-zero value of σ_R^2 (using zero would create a singularity in calculation of the cost function), and prior error as described in Table 1, with all 5 values of σ_Q^2. These OSSEs provide insight about the effect of prior error on the posterior estimate, and how density of observations can affect an observing system's ability to separate signal from noise. We performed 30 trials of this type of OSSE for each observing system, for each value of σ_Q^2, and for both Method A and Method B.

Finally, we performed full OSSEs with error on both the prior estimate and the pseudo-data. We performed the OSSEs for each of the 25 categories of error (every combination of the 5 values of σ_Q^2 and 5 values of σ_R^2). These are the most realistic simulations for the observing systems. We performed 30 trials of this type of OSSE for each observing system, for each combination of σ_Q^2 and σ_R^2, and for both Method A and Method B.

2.4 Analysis

The goal of this study is to demonstrate the new information that could be gained from MethaneSat. Specifically, we hoped to show the relative strength of MethaneSat over other observing systems in constraining methane emissions from oil and gas production regions, such as the Barnett Shale.

One important metric is the ability of the observing systems to constrain total emissions for the Barnett Shale region, relative to the target total emissions of 0.74 Tg year^{-1}. We are also interested in the total error in our posterior estimate for the Barnett Shale region. To calculate this, we take the posterior error covariance matrix, $\mathbf{V}_{\hat{s}}$ (dimension: grid cells x grid cells), extract the grid cells for the Barnett Shale region, sum the covariance terms, and take the square root.

It is also useful to investigate the ability of the observing systems to constrain the spatial distribution of CH_4 emissions, and the sharp gradients seen in the Barnett Shale region. We can learn about this by plotting the differences between the average posterior estimates and target emissions. To evaluate how closely each observing system retrieves the sharp peak in emissions, we measure the gradient between the highest-emitting grid cell and lowest-emitting grid cell in the target emissions, and compare to the gradient given in the posterior estimates.

Results

We performed OSSEs to demonstrate the advantages of MethaneSat above existing observing systems for constraining methane emissions from oil and gas production regions. Results from the OSSEs support our hypothesis that MethaneSat is able to retrieve emissions for the Barnett Shale region more successfully than both the NOAA network and GOSAT. We found that MethaneSat was able to successfully constrain total emissions in the Barnett Shale region, their spatial distribution, and their steep gradient.

3.1 Framework test - Perfect prior, perfect pseudo-data

The results of the framework test were as expected for all observing systems, given that we introduced no error into the system. The posterior estimate produced by the inversion framework was identical to the target. In some cases there was minor truncation error (on the order of e−13). The result of this test case increased our confidence in the reliability and accuracy of the inversion framework.

3.2 Method A: Barnett Shale region perturbed

In Method A, we applied prior error only to a small portion of the domain: the region bounded by -102° to -95° longitude and 30° to 35° latitude, including the Barnett Shale region. These OSSEs were designed to asses the effect of optimizing a mis-specified hotspot in emissions on well-specified emissions outside the hotspot. We were also interested in assessing the observing systems' abilites to constrain the hotspot without external interfering noise.

3.2.1 Error on prior, perfect pseudo-data

We first conducted these OSSEs without error on the pseudo-data in order to isolate the effect of error on the prior estimate.

When we plot the pseudo-data against the prior modeled enhancement and posterior modeled enhancement, we see that the posterior modeled enhancement perfectly matches the pseudo-data (Figure 7). This was true for all observing systems, all levels of prior error, and both algorithms. Thus, with no error on the pseudo-data, but error on the prior estimate, the inversion completely favors agreement with the pseudo-data.

Plotting the total average posterior emissions for the 11-grid cell Barnett Shale region (refer to Figure 6) reveals differences between the observing systems' abilities to constrain emissions (Figure 8). As expected, MethaneSat constrained emissions best, while the NOAA continuous and in situ network performed the worst. The difference was most noticeable when the error on the prior estimate was high. Furthermore, the difference was most drastic for the estimate produced using the non-negativity algorithm (Figure 8a).

Importantly, MethaneSat produced average total emissions that were very close to the target emissions for all values of prior error variance used in these OSSEs, and using both algorithms. However, using the regular algorithm, GOSAT constrained total emissions almost as well as MethaneSat.

The error bars in Figures 8a and 8b denote the interquartile range (IQR) for the results of the 30 iterations. The IQR is a robust statistic for the variability of the result, which does not make distributional assumptions. In every case, the NOAA network has the widest IQR, followed by GOSAT, and finally MethaneSat. The difference is most drastic when the error on the prior estimate was high. However, at every level of prior error variance, the IQR for MethaneSat is so narrow that the error bars are barely visible, indicating low variability across the 30 iterations. In contrast, even when GOSAT constrained total emissions well on average, its IQR was quite wide (Figure 8b).

Plotting the difference between the average posterior emissions and target
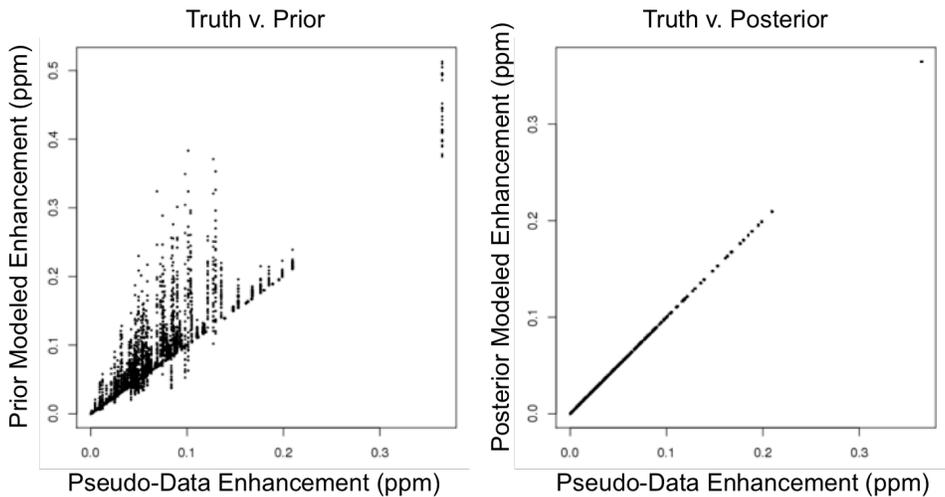
Figure 7: Pseudo-Data v. Prior Modeled Enhancement, Posterior Modeled Enhancement for the NOAA Network, with perfect pseudo-data and high error on the prior ($\sigma_Q^2 = 0.005$). The inversion has corrected all error in the prior modeled enhancement, and the posterior modeled enhancement perfectly match the pseudo-data.

emissions demonstrates the observing systems' abilities to constrain the spatial

distribution of emissions (Figures 9, 10). With only the Barnett Shale region

perturbed and no error on the pseudo-data, MethaneSat was able to almost
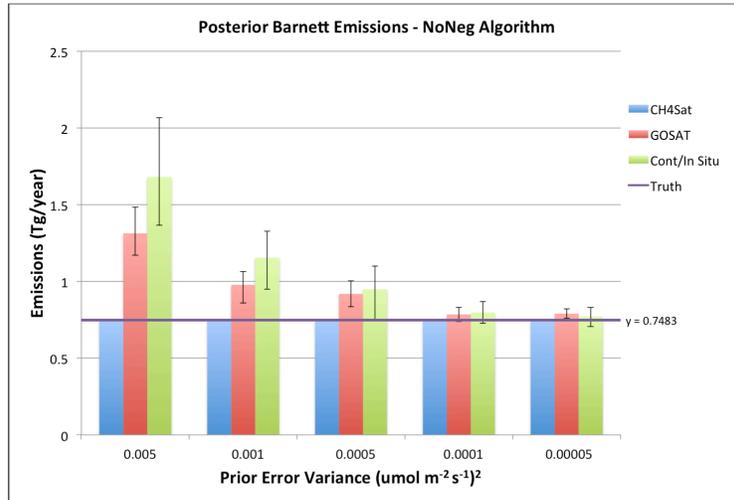
(a) Using the non-negativity algorithm, there is a marked difference in MethaneSat's ability to constrain emissions relative to the other observing systems.
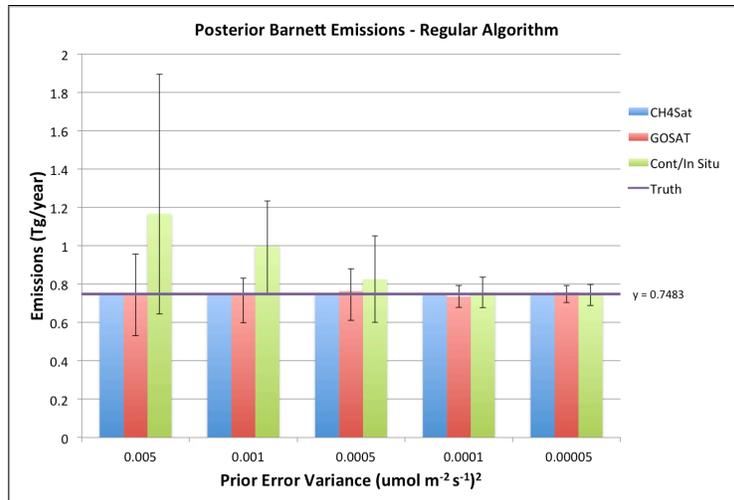
(b) Using the regular algorithm, MethaneSat constrains emissions much better than the NOAA network. GOSAT is able to constrain emissions almost as well as MethaneSat, but not as consistently (note the wide IQR).

Figure 8: The averaged total emissions for the 11-grid cell Barnett Shale region, with error on the prior but no error on the pseudo-data. The purple lines show the target flux of 0.7483 Tg year^{-1}. MethaneSat was able to closely constrain emissions for all levels of prior error. The error bars denote the IQR.

perfectly constrain emissions in the Barnett Shale region. The NOAA network and GOSAT were not able to do so, underestimating the peak in emissions and overestimating emissions in the surrounding gridcells. This was the case using both algorithms, though the differences were most drastic using the non-negativity algorithm.

The results of these OSSEs also show differences in how the observing systems distinguish between signal and noise. Although we only added error to the Barnett Shale region, the posterior estimates from the NOAA network and GOSAT included corrections to the prior estimate outside of the Barnett Shale region. MethaneSat showed almost no correction to the prior estimate outside of the Barnett Shale region, except for a small region to the south. This was the case for both the regular algorithm and the non-negativity algorithm, and can be seen when we plot mean prior error (average prior emissions minus true emissions) and mean correction to the prior (average posterior emissions minus average prior emissions) (Figure 11). In the plots of mean correction to the prior estimate, the NOAA network and GOSAT show color variation throughout the domain, while MethaneSat's domain beyond the Barnett Shale region is completely flat (except for the small region to the south).

3.2.2   Error on prior, error on pseudo-data

The OSSEs with prior error in the Barnett Shale region and error on the pseudo-data are a more realistic test for the observing systems. This isolates the effect of a mis-specified hostpot with realistic error. When we plot the difference between average posterior emissions and target emissions, we see an interesting difference in the observing systems' abilities to constrain the spatial

distribution of emissions (Figures 12-14). While all observing systems show some error with a poor prior estimate and poor pseudo-data, MethaneSat is the only observing system that shows improvement with decreasing error on the pseudo-data. With a poor prior but excellent pseudo-data, MethaneSat is able to constrain emissions in the grid cells corresponding to the Barnett Shale region, almost perfectly capturing the steep gradient, despite some error in the surrounding cells. The other observing systems smoothed the peak into the surrounding low-emission cells, even with excellent pseudo-data.
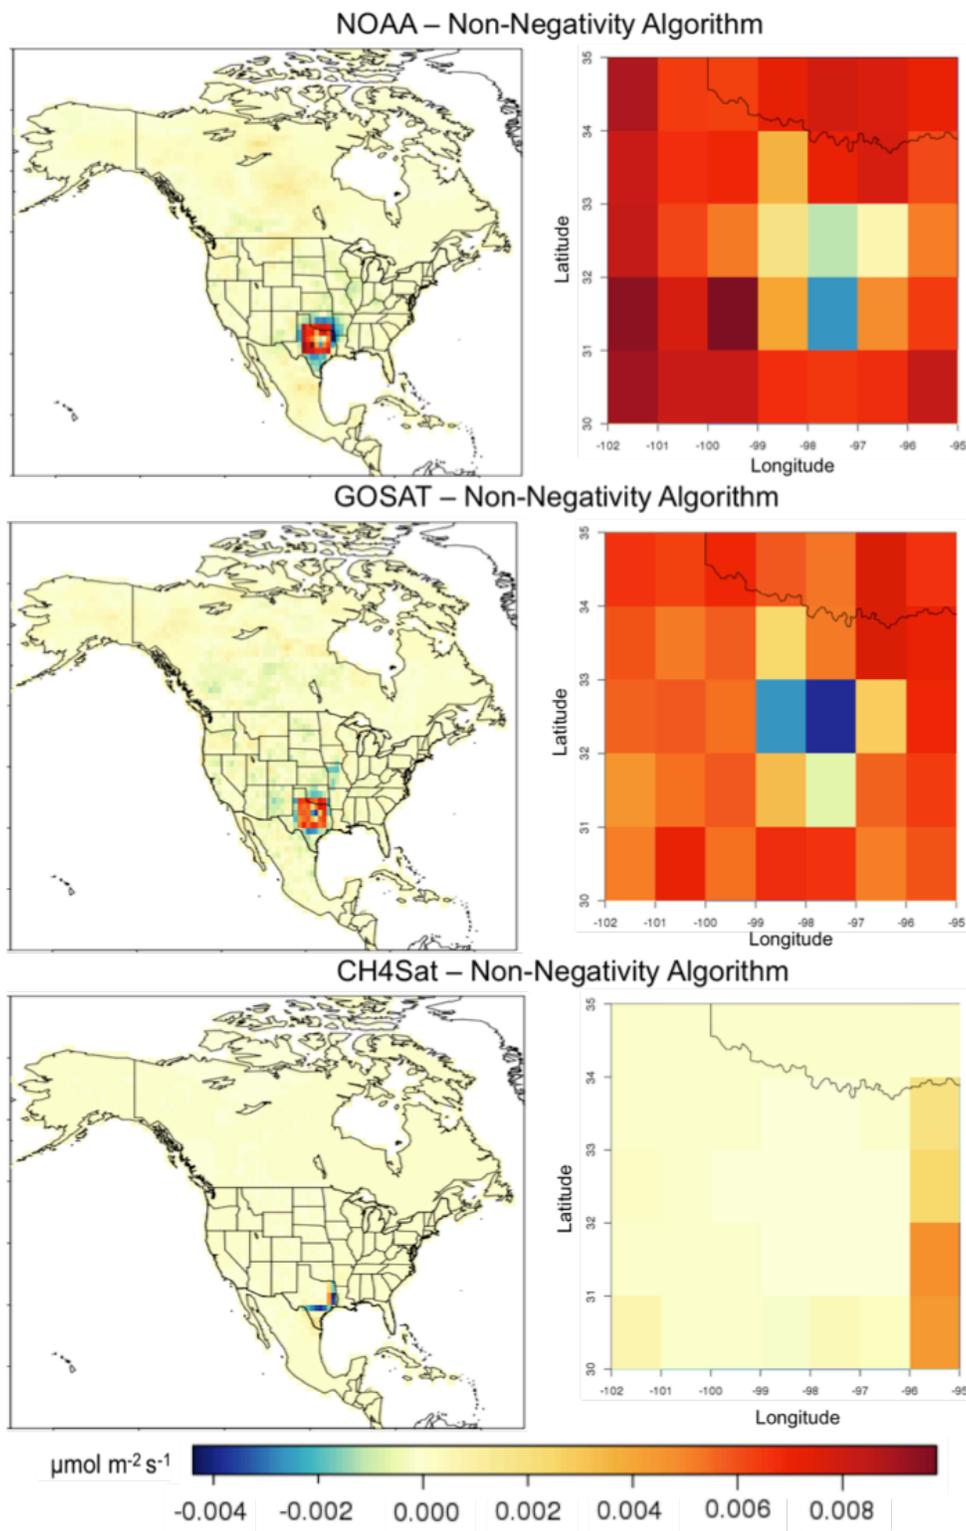
Figure 9: The average posterior error (calculated as average posterior emissions - true emissions), with the middle level of error on the prior but no error on the pseudo-data, using the non-negativity algorithm. MethaneSat was able to closely constrain emissions, capturing the high peak in the middle of the region.
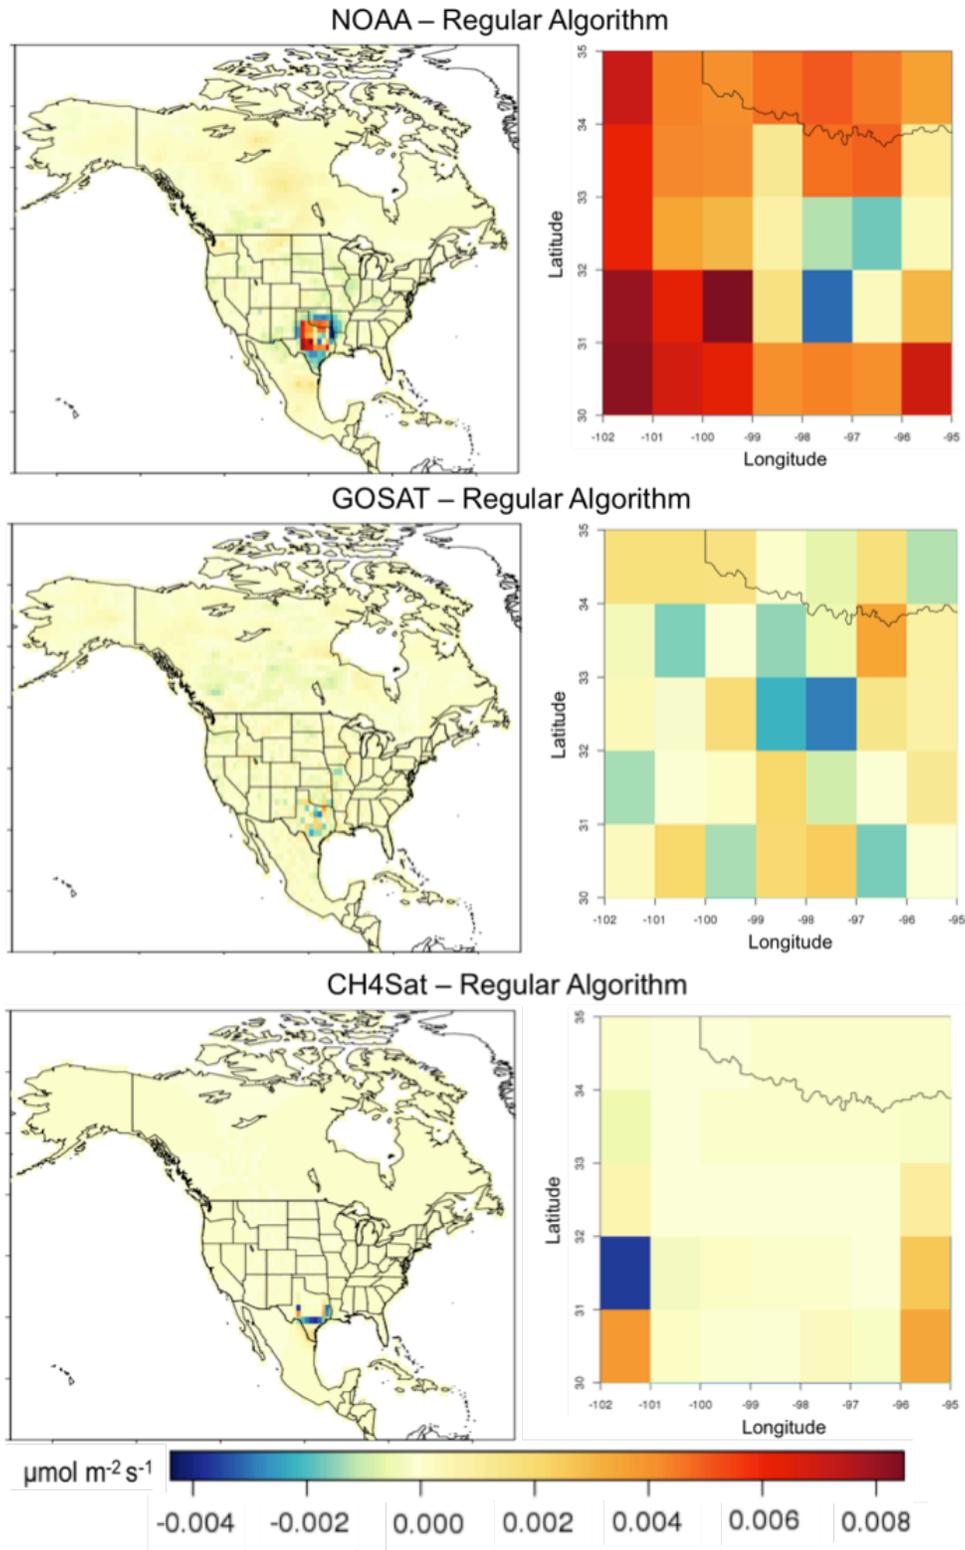
Figure 10: The average posterior error (calculated as average posterior emissions - true emissions), with the middle level of error on the prior but no error on the pseudo-data, using the regular algorithm. MethaneSat was able to closely constrain emissions, capturing the high peak in the middle of the region.

31
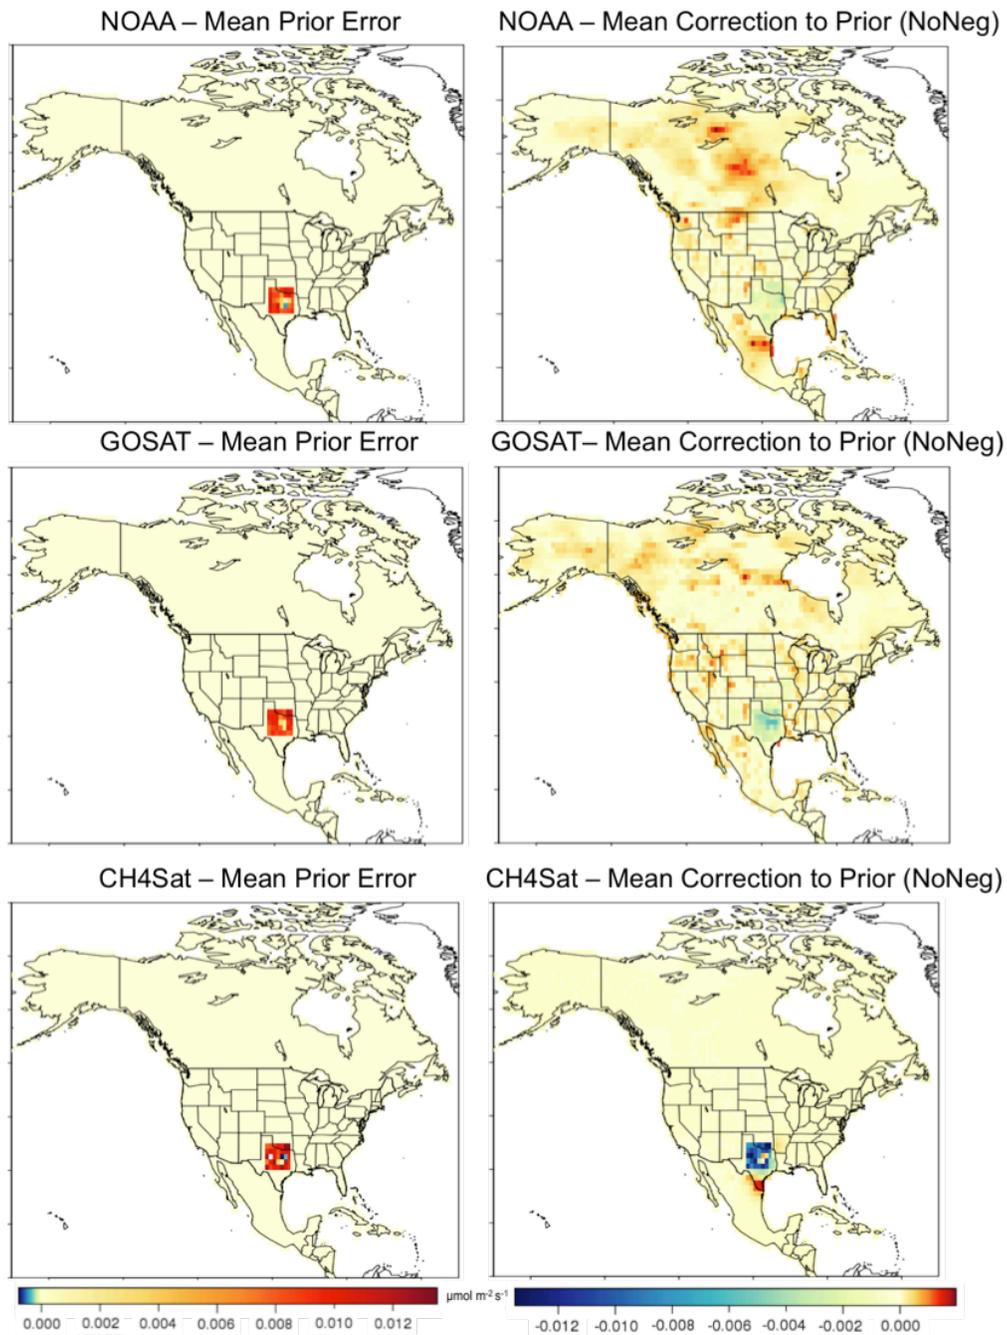
Figure 11: Mean prior error and mean correction to the prior with perfect pseudo-data, using the non-negativity algorithm. The NOAA network and GOSAT show extensive correction to the prior outside of the Barnett Shale region, even though the rest of the domain was left perfectly specified. MethaneSat only corrects the prior in the Barnett Shale region, and a small region to the south.
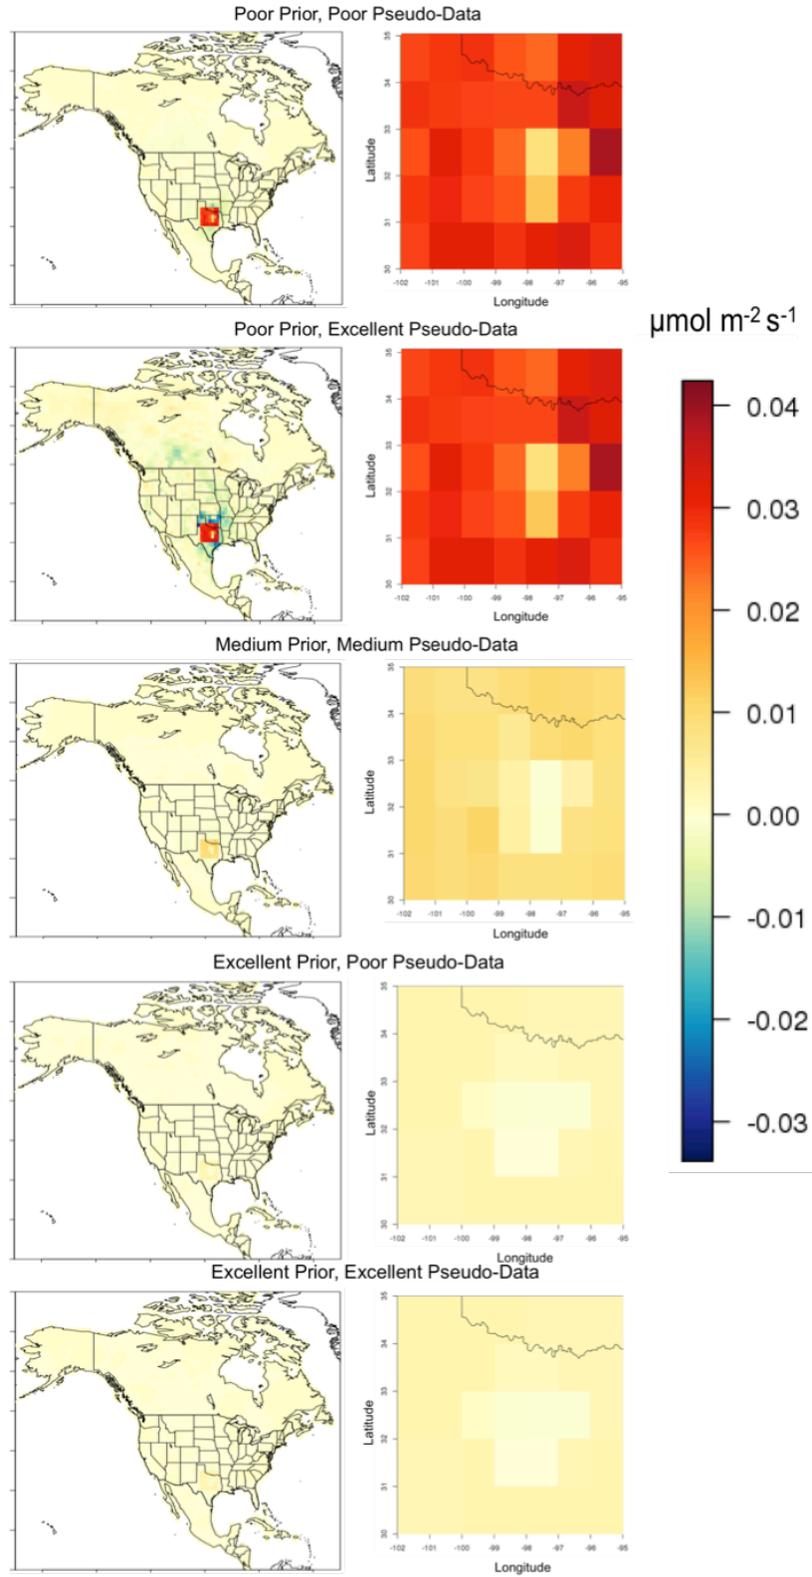
Figure 12: NOAA: difference between average posterior emissions and target emissions, using the regular algorithm. The NOAA network fails to retrieve the peak in emissions with high error on the prior estimate.

33
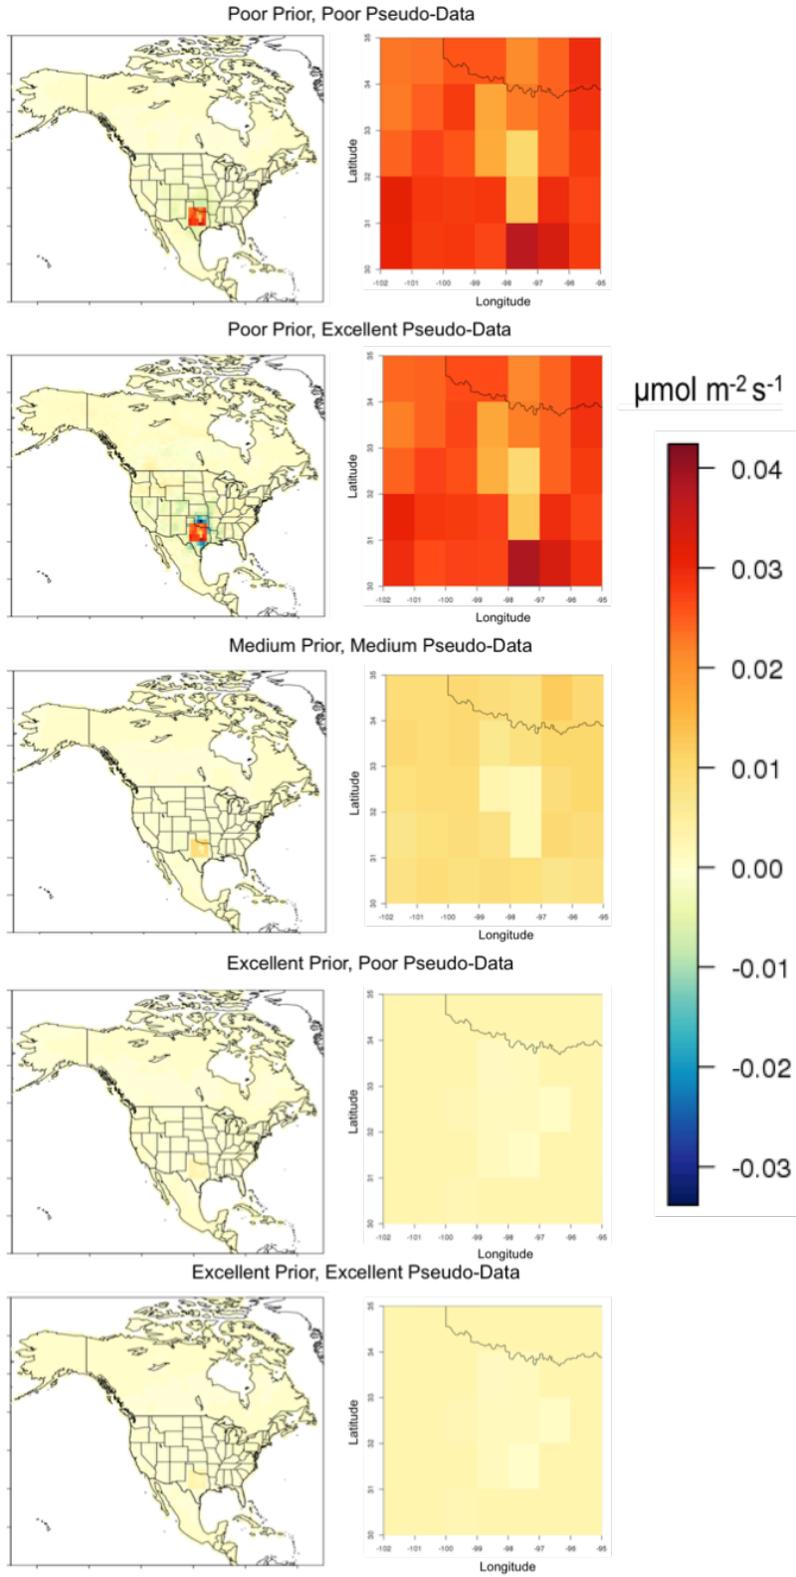
Figure 13: GOSAT: difference between average posterior emissions and target emissions, using the regular algorithm. GOSAT network fails to retrieve the peak in emissions with high error on the prior estimate.
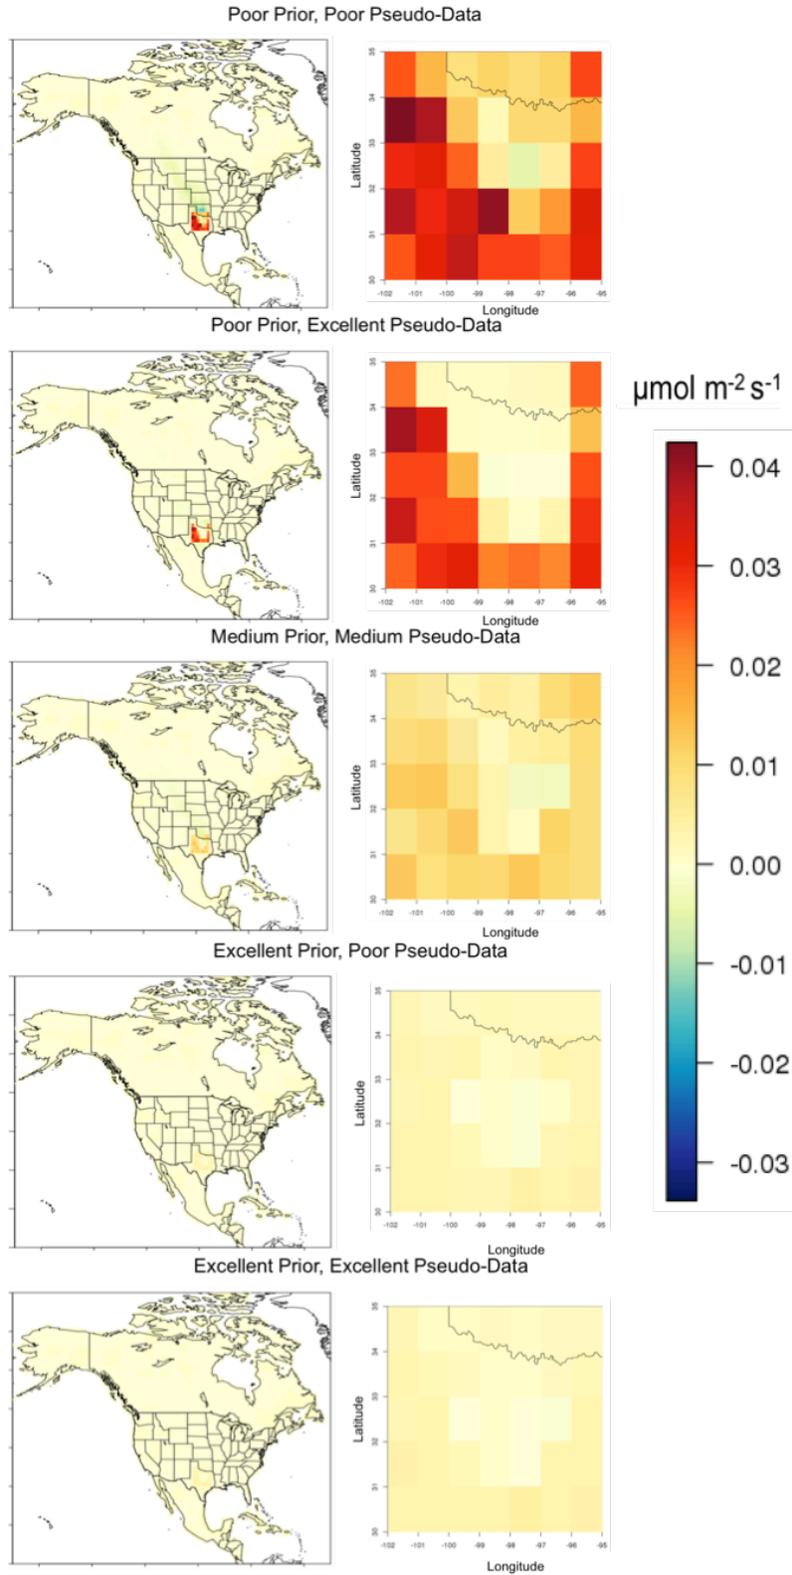
Figure 14: MethaneSat: difference between average posterior emissions and target emissions, using the regular algorithm. CH4Sat successfully retrieves the peak in emissions with high error on the prior estimate, particularly when error on the pseudo-data is low.

3.3 Method B: Entire domain perturbed

Perturbing the entire domain is our most realistic test of the observing systems' ability to retrieve emissions. We hypothesized that the NOAA network and GOSAT would not be capable of constraining emissions in the Barnett Shale region, but that MethaneSat would be able to do so. We found that, under conditions of perfect pseudo-data, MethaneSat was, far more capable of constraining total emissions and their spatial distribution than the other observing systems. When error was added to the pseudo-data, MethaneSat demonstrated a superior ability to capture the high peak in emissions for large values of prior error variance. MethaneSat was also able to retrieve the steep gradient in the Barnett Shale region at nearly every level of error, while the NOAA network and GOSAT could only do so in cases of low prior error.

3.3.1 Error on prior, perfect pseudo-data

Without error on the pseudo-data, these OSSEs isolate the effect of error on the prior estimate. We saw similar results here as with Method A: MethaneSat was able to constrain total emissions for the Barnett Shale region much more effectively than the other observing systems (Figure 15). MethaneSat performed the best, constraining average emissions almost perfectly in every case, while the NOAA continuous and in situ network performed the worst. The difference was most noticeable when the error on the prior estimate was high. Furthermore, the difference was most drastic for the posterior estimate produced using the non-negativity algorithm.

Importantly, MethaneSat produced total emissions that were very close to
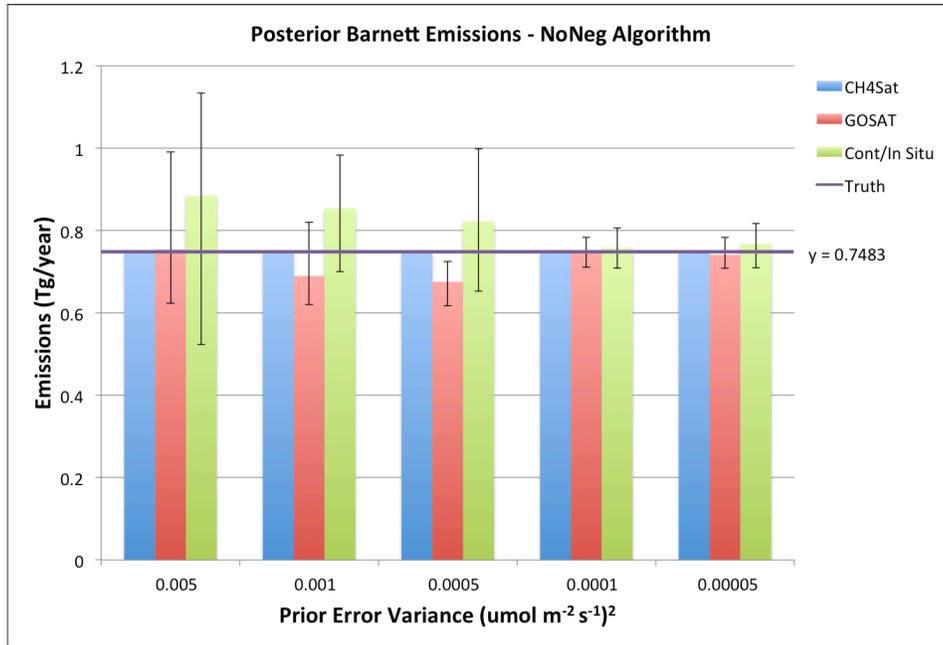
(a)
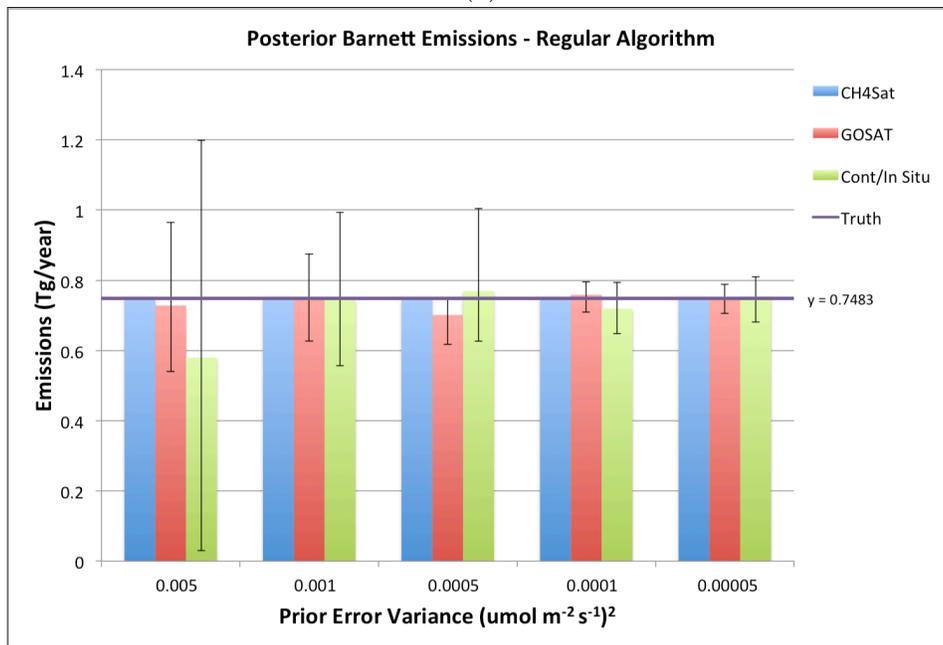
(b)

Figure 15: Total posterior emissions in the 11-grid cell Barnett Shale region, averaged over the 30 trials. All observing systems are able to predict posterior emissions that closely match the true emissions, using both algorithms. However, the error bars, indicating the IQR for the 30 trials, are much larger for the NOAA network and GOSAT than for MethaneSat.

the target emisions for all values of prior error variance used in these OSSEs, and using both algorithms. However, using the regular algorithm, GOSAT constrained total emissions almost as well as MethaneSat (Figure 15b). The results were comparable for the Barnett Shale region plus the cells immediately surrounding it (Figure 16).

The error bars in Figures 15 and 16 denote the interquartile range (IQR) for the results of the 30 iterations. In every case, the NOAA network has the widest IQR, followed by GOSAT, and finally MethaneSat — the same trend as with Method A. Though GOSAT is sometimes able to closely constrain emissions on average, its IQR is much wider than that of MethaneSat for every value of prior error variance, indicating greater variability.

These differences between observing systems can likely be attributed to the low sampling density of the NOAA network and GOSAT, diminishing their ability to constrain emissions. In other words, the observations do not correct the prior estimate as well as with MethaneSat. The low sampling density results in more reliance on the prior estimate, which translates to greater variability in the posterior estimate and wide IQRs for the NOAA network and GOSAT. This high variability is particularly concerning considering that the perturbations to the truth that produced the prior estimate had a mean of zero — therefore, even totally unconstrained priors would have correct total emissions over a large number of trials.

Interestingly, the NOAA network and GOSAT were unsuccessful at constraining the low-emitting grid cells surrounding the Barnett Shale region when using the non-negativity algorithm (note the differences between Figure 15a and Figure 16a). This is likely due to the non-negativity algorithm's

strategy for dealing with negative values in the posterior estimate, which may cause overestimation of total emissions. This topic is discussed further in the Discussion.

The results also indicate a clear advantage in MethaneSat's ability to constrain the spatial distribution of emissions in the Barnett Shale region. Figures 17 and 18 show the results of the OSSEs performed using $\sigma_Q^2 = $ 5e-4, the median level of prior error used in this study. We plotted the difference between average posterior emissions and target emissions. Using both algorithms, MethaneSat was able to constrain the spatial distribution of emissions in the Barnett Shale region almost perfectly, as well as the large region to the north-west of the Barnett Shale region. Neither the NOAA network nor GOSAT was able to retrieve the peak in emissions in the center of the region, smoothing it into the surrounding grid cells instead. The results were comparable for higher and lower levels of prior error.

Finally, MethaneSat showed extremely low posterior error (calculated using $\mathbf{V}_{\hat{s}}$, from Equation 5) for all values of prior error variance (Figure 19). It was far below that of the NOAA network or GOSAT, with the NOAA network having the highest posterior error.

3.3.2 Error on prior, error on pseudo-data

The OSSEs with error on the prior and the pseudo-data are the most realistic tests for the observing systems. As expected, total posterior error for the Barnett Shale region (calculated using $\mathbf{V}_{\hat{s}}$) decreased as error on the prior estimate and error on the pseudo-data decreased. When we plot this in a heat map, we see a smooth gradient across both error axes for all observing

Table 3: Method B, Full OSSE: Posterior error for the Barnett Shale region, given in Tg year^{-1}.

Observing System	Max Posterior Error	Min Posterior Error
NOAA Network	3.235	0.323
GOSAT	3.414	0.341
CH$_4$Sat	3.127	0.313

systems (Figure 20). However, posterior error decreases most dramatically with decreasing error on the prior. Thus, at least for the values of error tested in these OSSEs, error on the prior has a larger impact on posterior error than does error on the pseudo-data.

Interestingly, there was not a large difference in the calculated posterior error between observing systems. For a given level of error on the prior and error on the pseudo-data, the errors for the three observing systems were comparable, although MethaneSat generally had slightly lower error than both the NOAA network and GOSAT, as well as lower maximum and minimum errors.

The trend in average total emissions for the 11-grid cell Barnett Shale region was as expected - as error decreased in either direction, the posterior estimate typically grew closer to the target total emissions (Figure 21). Overall, MethaneSat's posterior estimates were closer to the target emissions than those of the other observing systems. In particular, using the regular algorithm, MethaneSat retrieved total emissions almost perfectly at all levels of prior error variance and pseudo-data error variance. The improvement over the other observing systems is most drastic for the two highest levels of prior error, for which the NOAA network severely underestimated emissions. GOSAT was closer to the target total, but often overestimated emissions.

Again, considering the IQRs of the 30 iterations makes this comparison even more meaningful. MethaneSat had much narrower IQRs than both the NOAA network and GOSAT at every level of error variance, indicating more consistent results (Figure 22). While the other observing systems constrained emissions at some levels of error, they demonstrated much more variability, likely indicating more reliance on the prior estimate. MethaneSat was therefore able to constrain total emissions in the region most successfully and most consistently, regardless of the prior estimate.

We plotted the difference between average posterior emissions and target emissions to investigate the ability of the observing systems to constrain the spatial distribution of emissions in the region (Figures 23-25). We found that with high error on the prior and low error on the pseudo-data, MethaneSat was able to retrieve the large peak in the region far better than both the NOAA network and GOSAT (as demonstrated by the dark blue squares in the NOAA network and GOSAT plots), although there was slightly more error in the surrounding regions for MethaneSat. With low prior error, all three observing systems were able to constrain the spatial distribution well. Interestingly, upon investigation of the results of individual iterations, it was found that the high average error in the surrounding regions for MethaneSat was driven by a few iterations with extremely high prior emission estimates in that region, but that individual inversions were generally able to retrieve the entire Barnett Shale region and its surroundings almost perfectly. These results were consistent between the regular algorithm and non-negativity algorithm.

It is also useful to consider how well the observing systems can capture the emission gradients from the Barnett Shale region. The steepest gradient

41

in the region is 0.30765 μmol m^{-2} s^{-1}/m. We hypothesized that, due to its high sampling density, MethaneSat would best retrieve this gradient. This was indeed the case — the NOAA network and GOSAT both smoothed the gradient, except at very low levels of prior error variance (Figure 26). MethaneSat captured the gradient successfully for nearly every combination of prior error variance and pseudo-data error variance, and captured it nearly perfectly when pseudo-data error variance was low. In contrast, the NOAA network and GOSAT gradients did not improve much with decreasing pseudo-data error variance, and actually worsened with decreasing pseudo-data error variance when using the non-negativity algorithm. Interestingly, and unsurprisingly, the non-negativity algorithm caused more smoothing of the sharp gradient than the regular algorithm, for all observing systems and all levels of error.

Focusing on the gradient retrieved by MethaneSat, we see in Figure 26 that it was able to closely retrieve the steep gradient for all levels of error — decreasing the error on the pseudo-data did not drastically affect the result. However, considering the width of the IQR across all 30 iterations, we see that there was much less variability when error on the pseudo-data was low. This was particularly true for high values of prior error variance. At lower levels of prior error variance, decreasing the error on the pseudo-data had less of an impact (although at all levels of prior error variance, the lowest level of pseudo-data error variance gives extremely narrow IQRs).
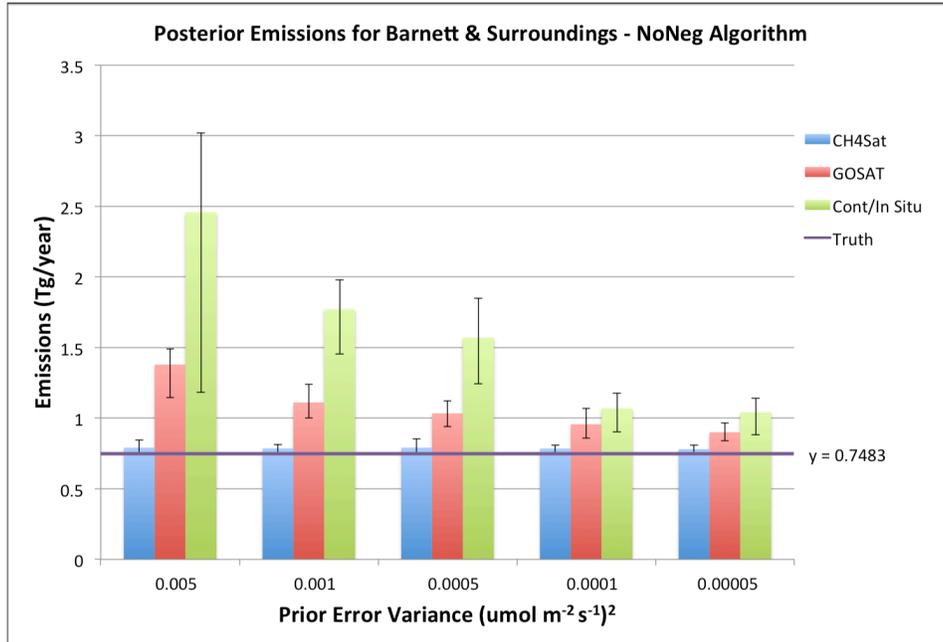
(a) Using the non-negativity algorithm, there is a marked difference in MethaneSat's ability to constrain emissions relative to the other observing systems.
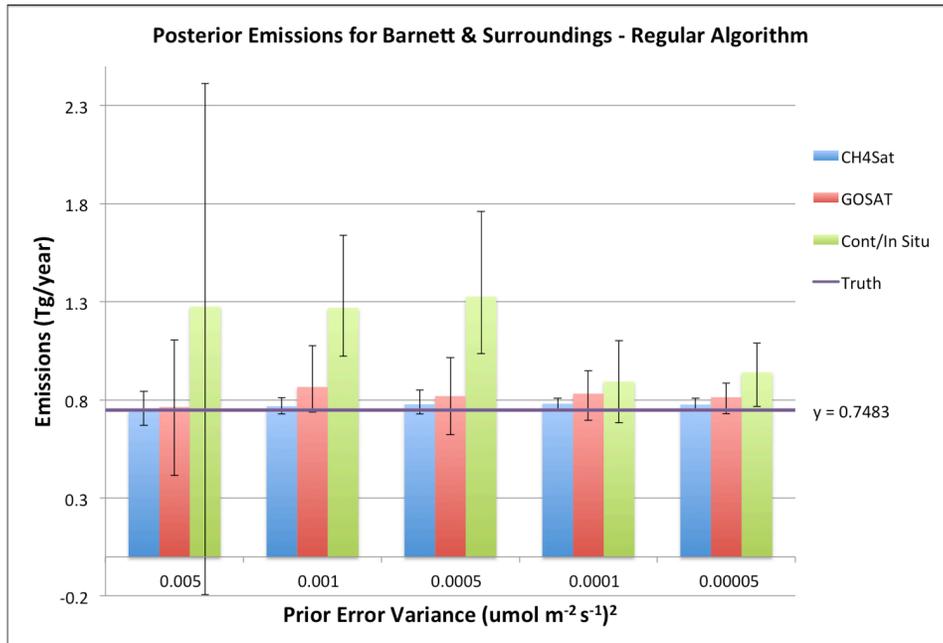
(b) Using the regular algorithm, MethaneSat constrains emissions much better than the NOAA network. GOSAT is able to constrain emissions almost as well as MethaneSat, but not as consistently (note the wide IQR).

Figure 16: Total posterior emissions in the Barnett Shale region and surrounding gridcells, averaged over the 30 trials. MethaneSat predicts posterior emissions that closely match the true emissions for all levels of prior error variance using both algorithms.

43
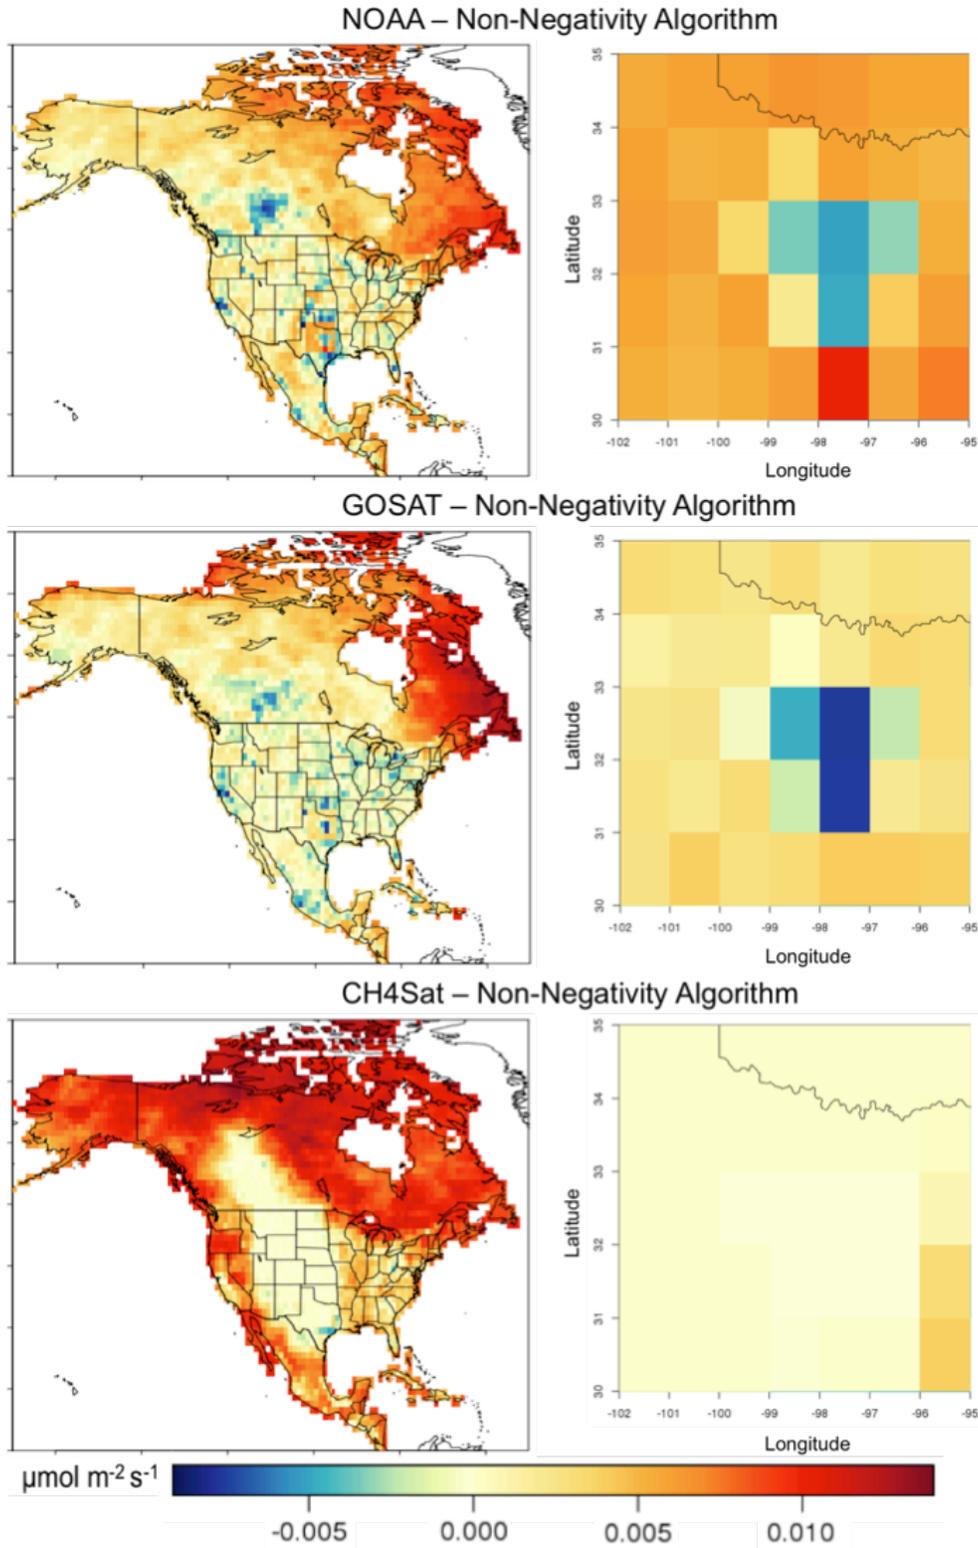
Figure 17: The difference between average posterior emissions and true emissions, with $\sigma_Q^2 = 5e{-}4$ but no error on the pseudo-data, using the non-negativity algorithm. MethaneSat was able to closely constrain emissions, capturing the high peak in the middle of the region.

44
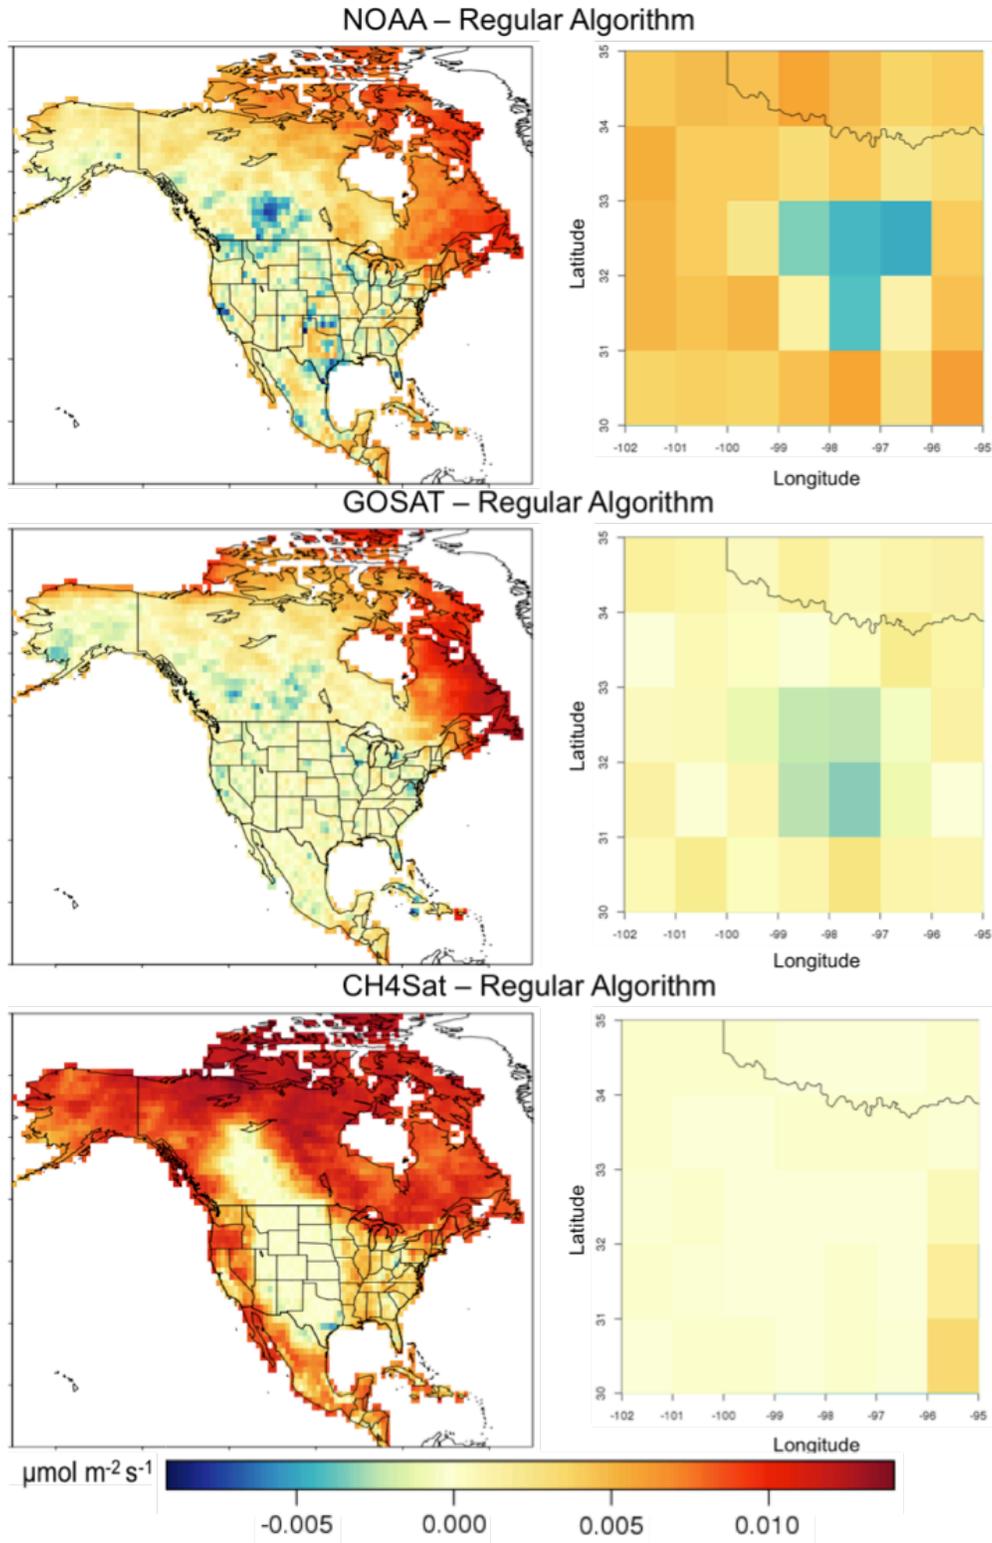
Figure 18: The difference between average posterior emissions and true emissions, with $\sigma_Q^2 = 5e{-}4$ but no error on the pseudo-data, using the regular algorithm. MethaneSat was able to closely constrain emissions, capturing the high peak in the middle of the region.
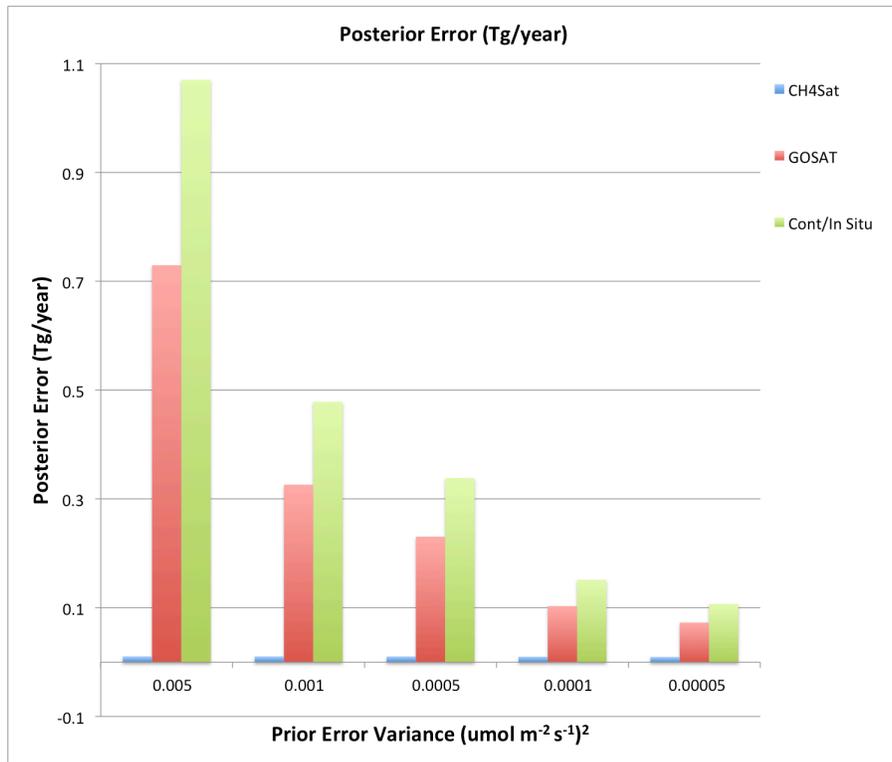
Figure 19: Posterior error in the 11-grid cell Barnett Shale region (calculated using $\mathbf{V}_{\hat{s}}$) for Method B with no error on the pseudo-data. The calculated posterior error for MethaneSat is minuscule relative to that of the other systems, for all values of prior error variance.
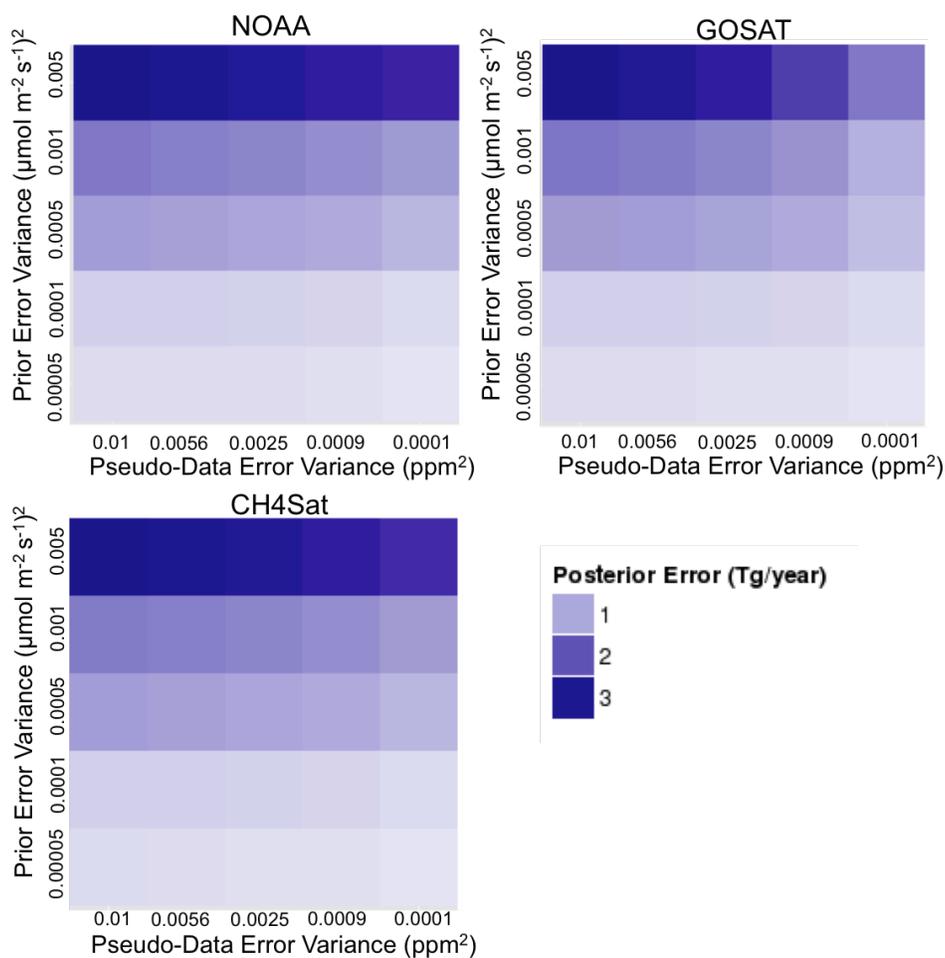
Figure 20: Posterior error, calculated using $\mathbf{V}_{\hat{s}}$. As expected, posterior error decreases along both error axes, but more drastically along the axis of prior error variance. Posterior error is comparable for all observing systems. Darker color indicates higher posterior error.
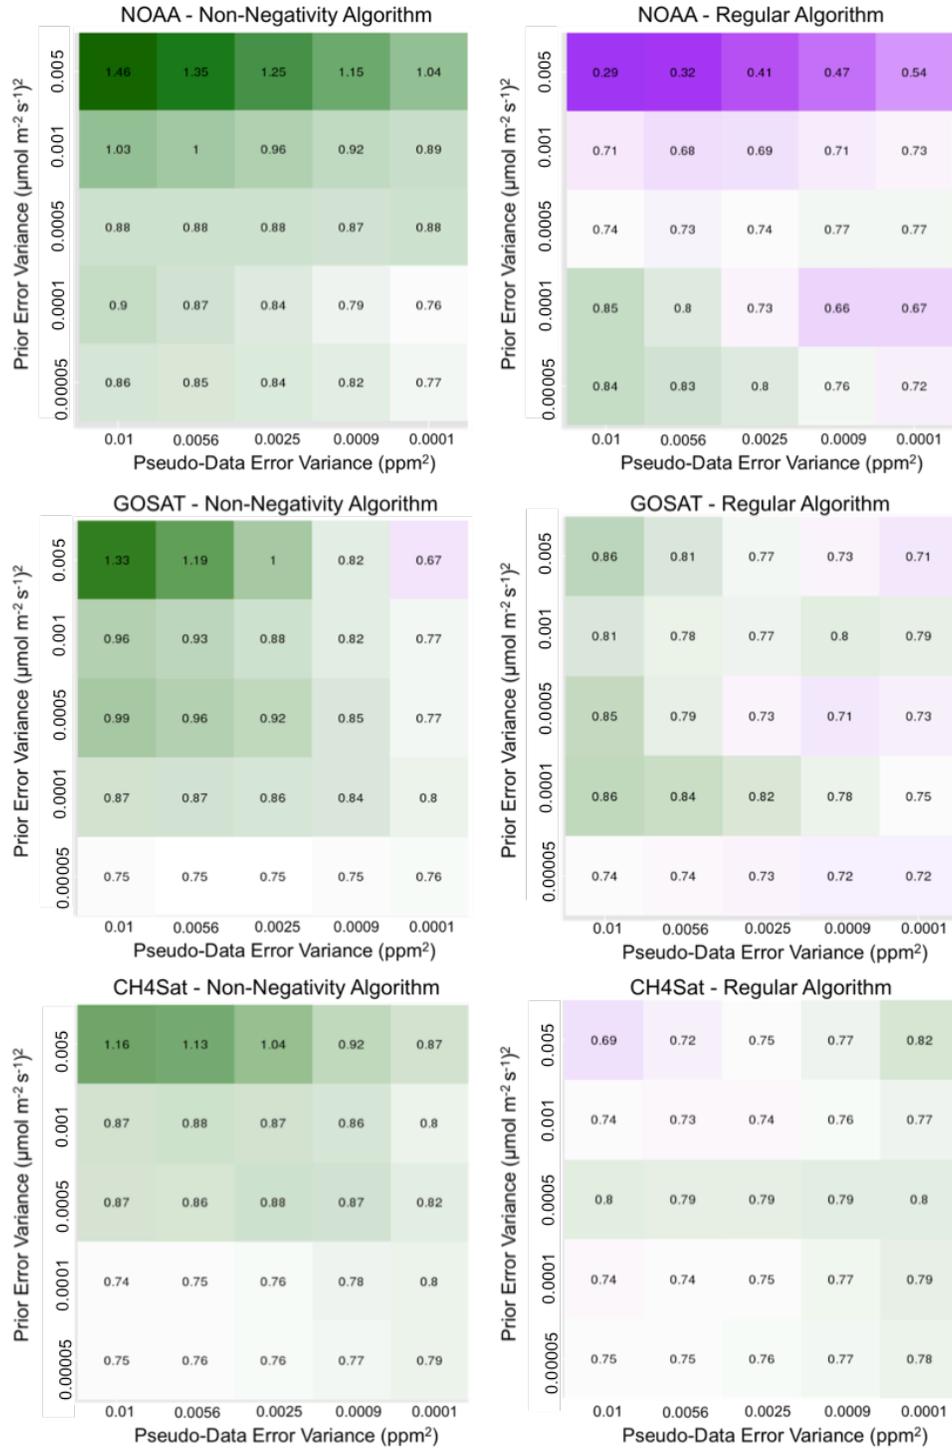
Figure 21: Average total posterior emissions for the 11-grid cell Barnett Shale region. The white-colored squares denote estimates close to the target. Green-colored squares are overestimated, while purple are underestimated. MethaneSat shows the greatest number of white squares, and overall less extreme color than the other two observing systems, particularly for high prior error variance. Overall, the estimates using the regular algorithm are closer to the target than those using the non-negativity algorithm.

48
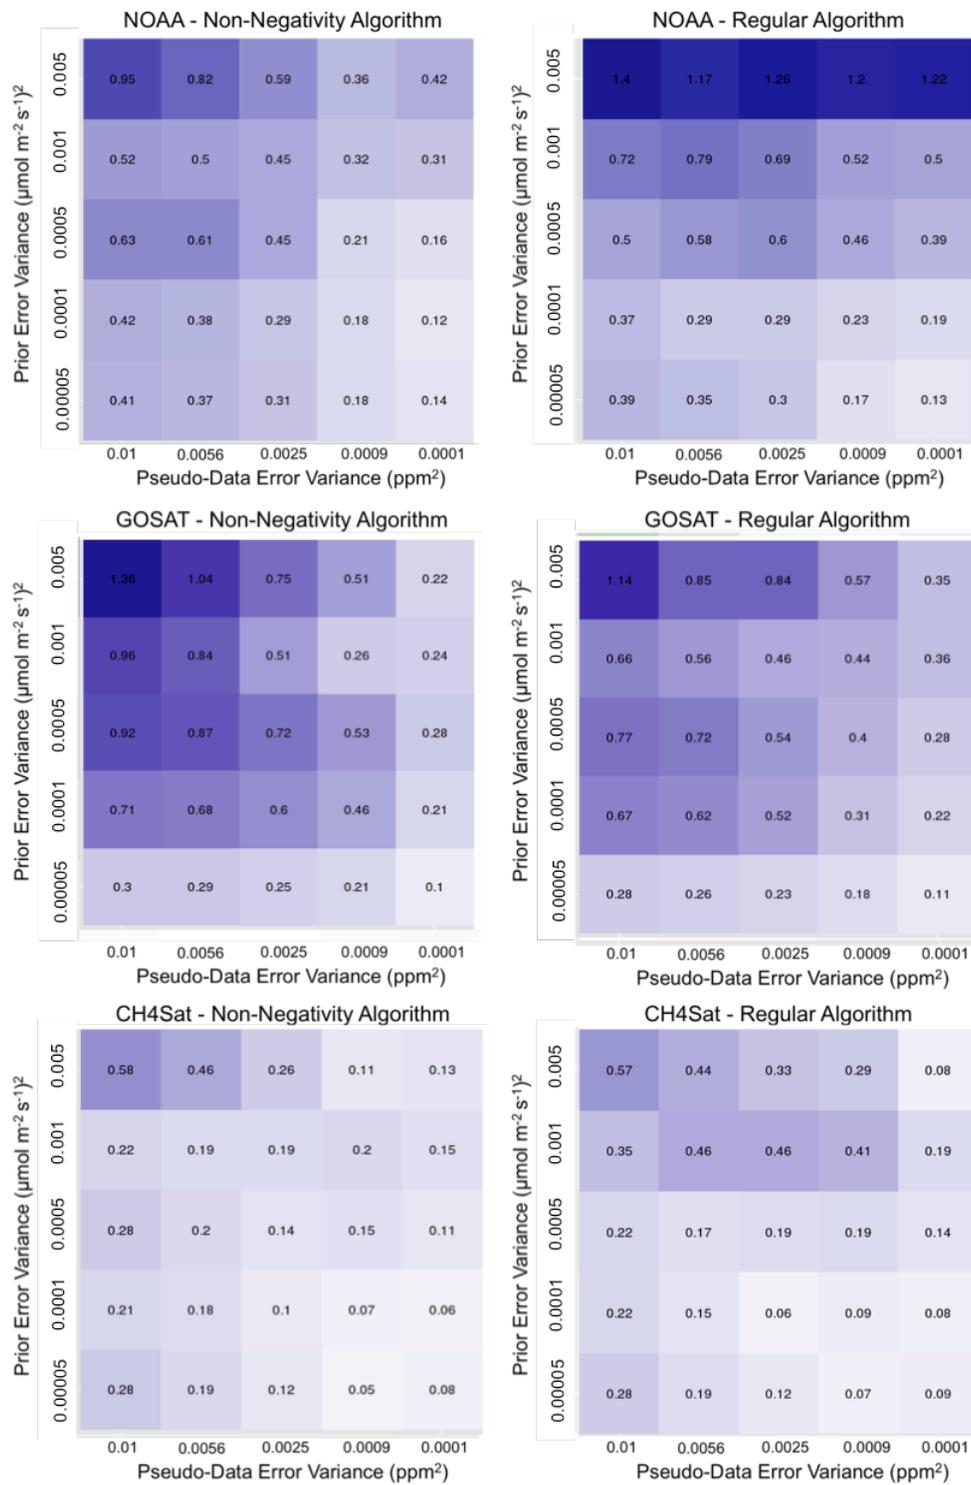
Figure 22: The widths of the IQRs for the total posterior emission estimates for the 11-grid cell Barnett Shale region, where darker color indicates a wider IQR. MethaneSat has much narrower IQRs than the other observing systems for every level of error.
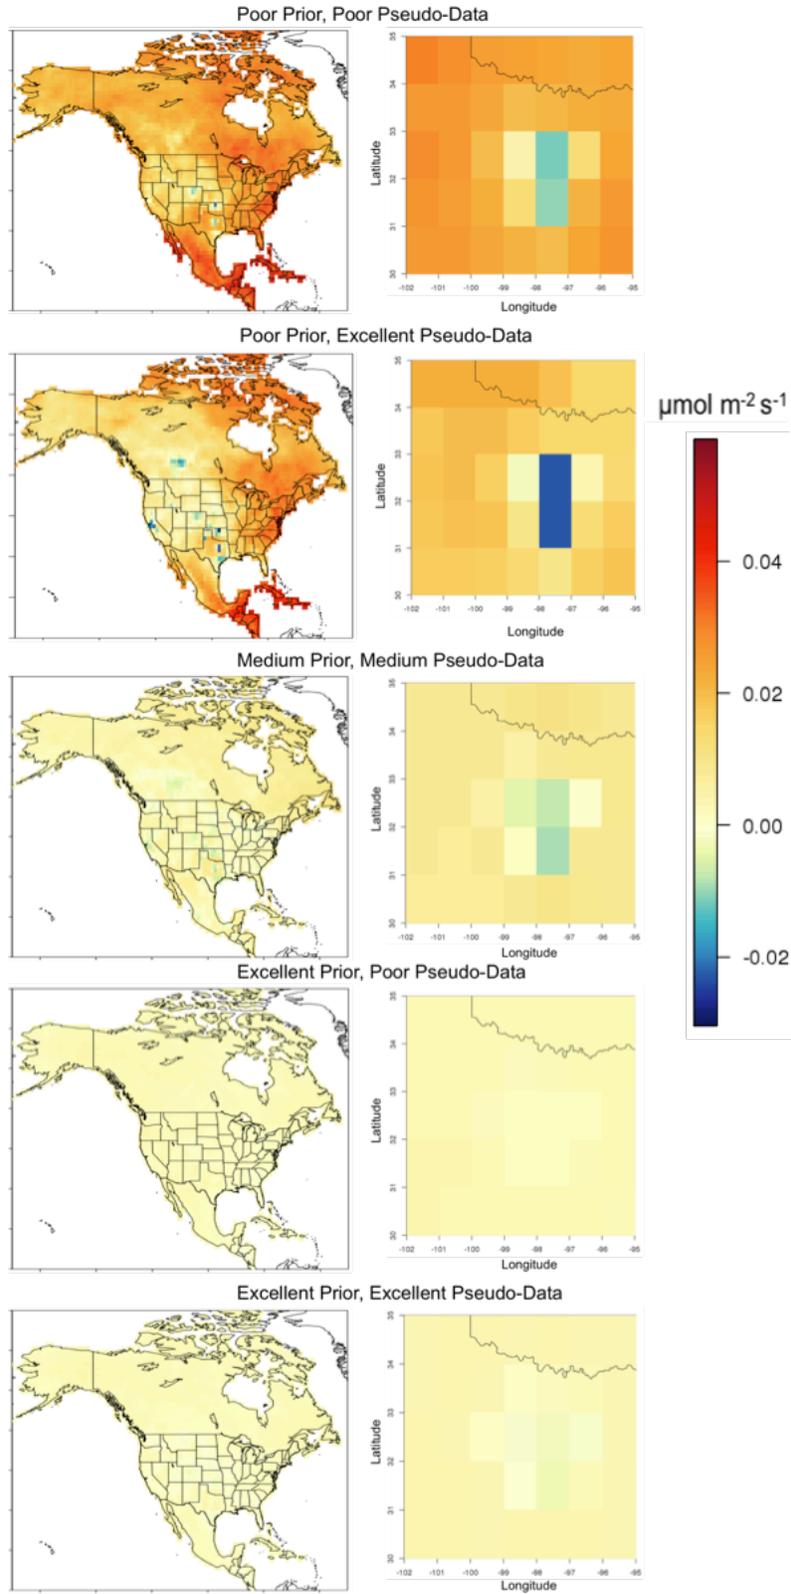
Figure 23: NOAA network: Average posterior emissions minus target emissions, using the non-negativity algorithm. The NOAA network fails to retrieve the peak in emissions in the middle of the Barnett Shale region.
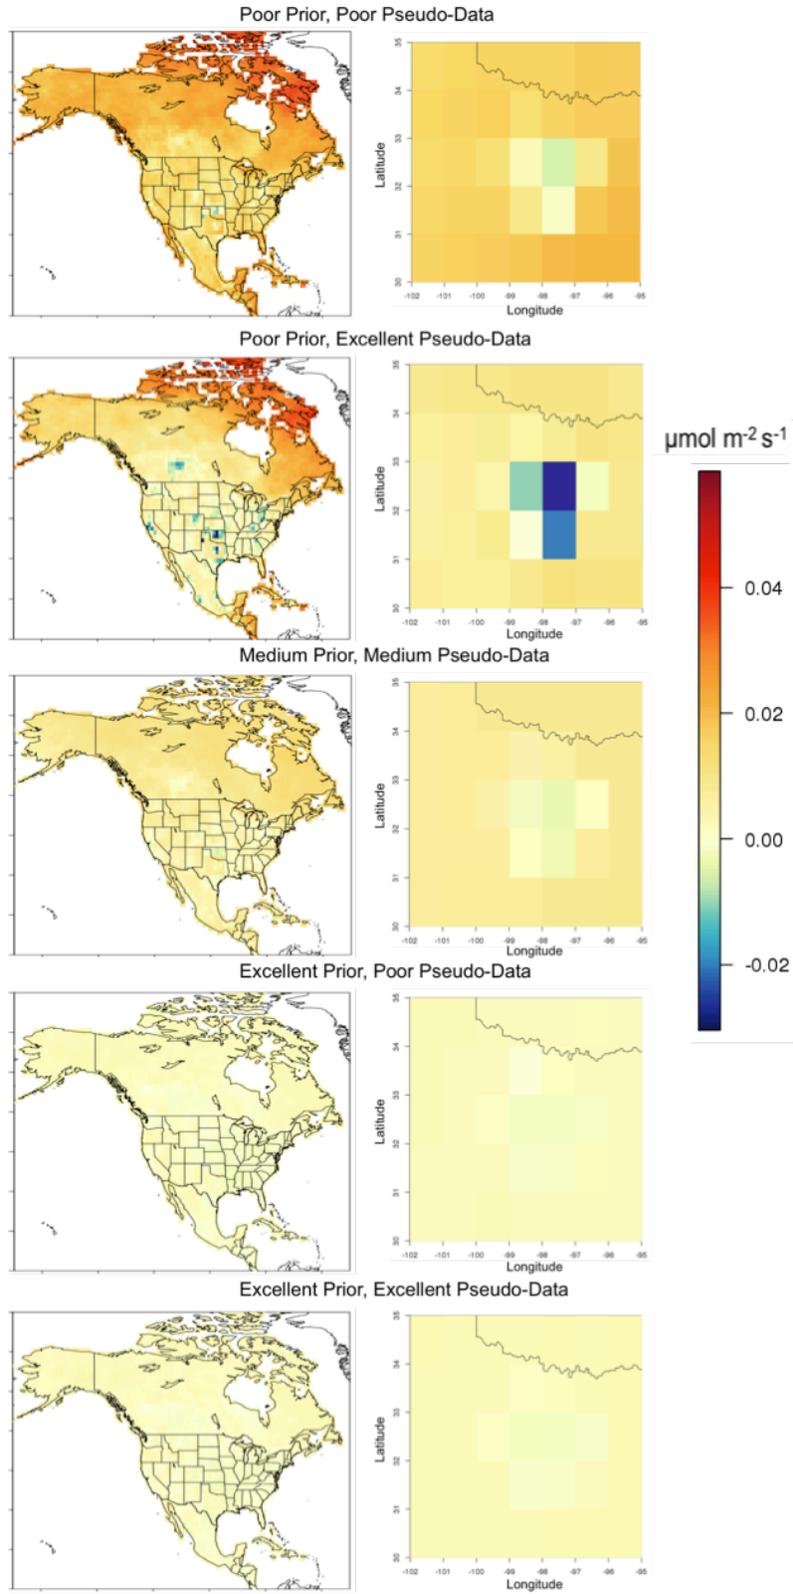
Figure 24: GOSAT: Average posterior emissions minus target emissions, using the non-negativity algorithm. GOSAT fails to retrieve the peak in emissions in the middle of the Barnett Shale region.
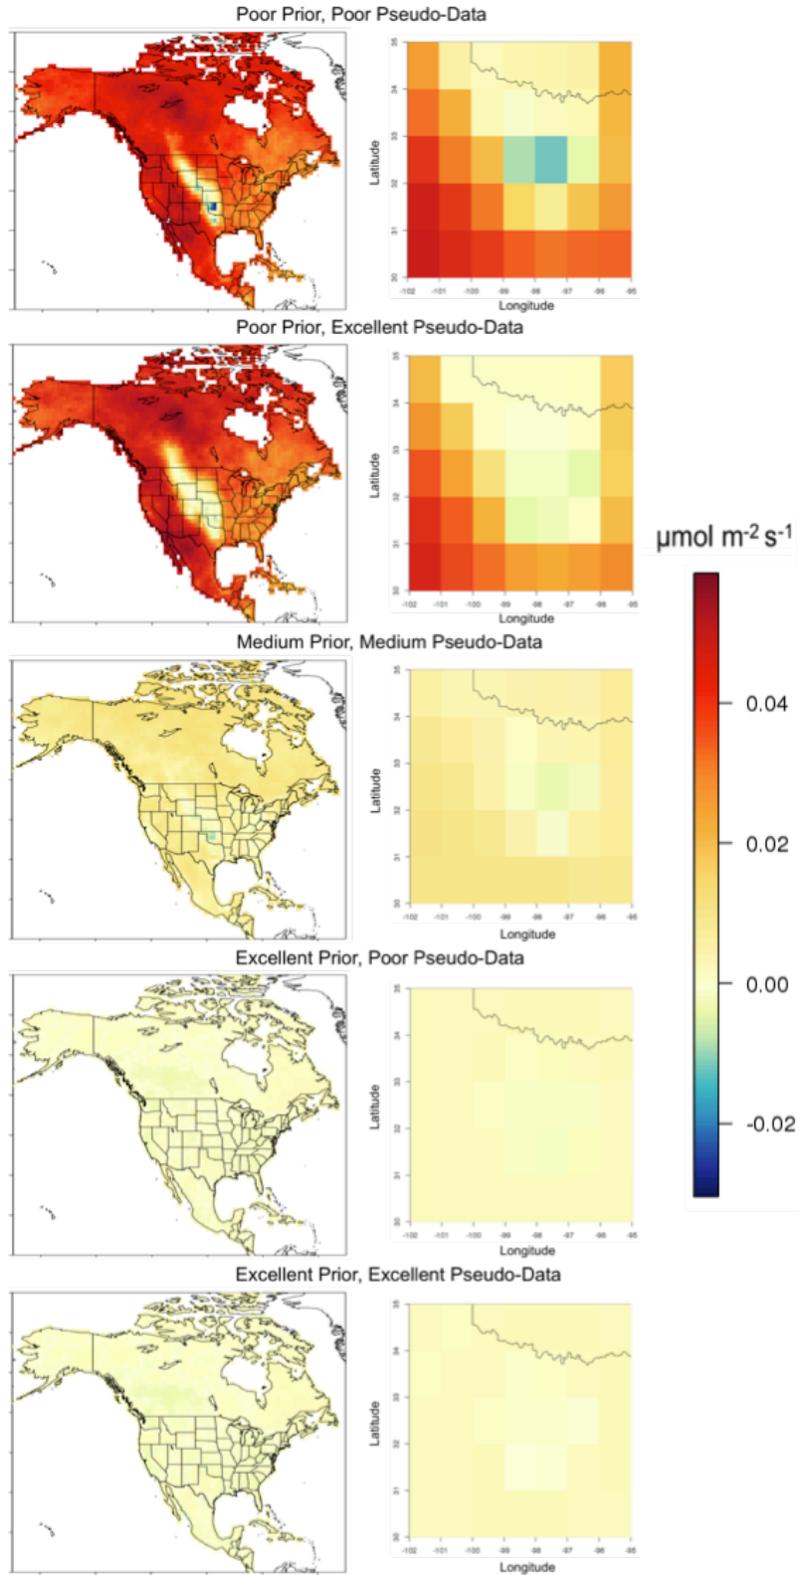
Figure 25: MethaneSat: Average posterior emissions minus target emissions, using the non-negativity algorithm. MethaneSat retrieves the peak in emissions in the middle of the Barnett Shale region, particularly with excellent pseudo-data, despite some error in the surrounding regions.
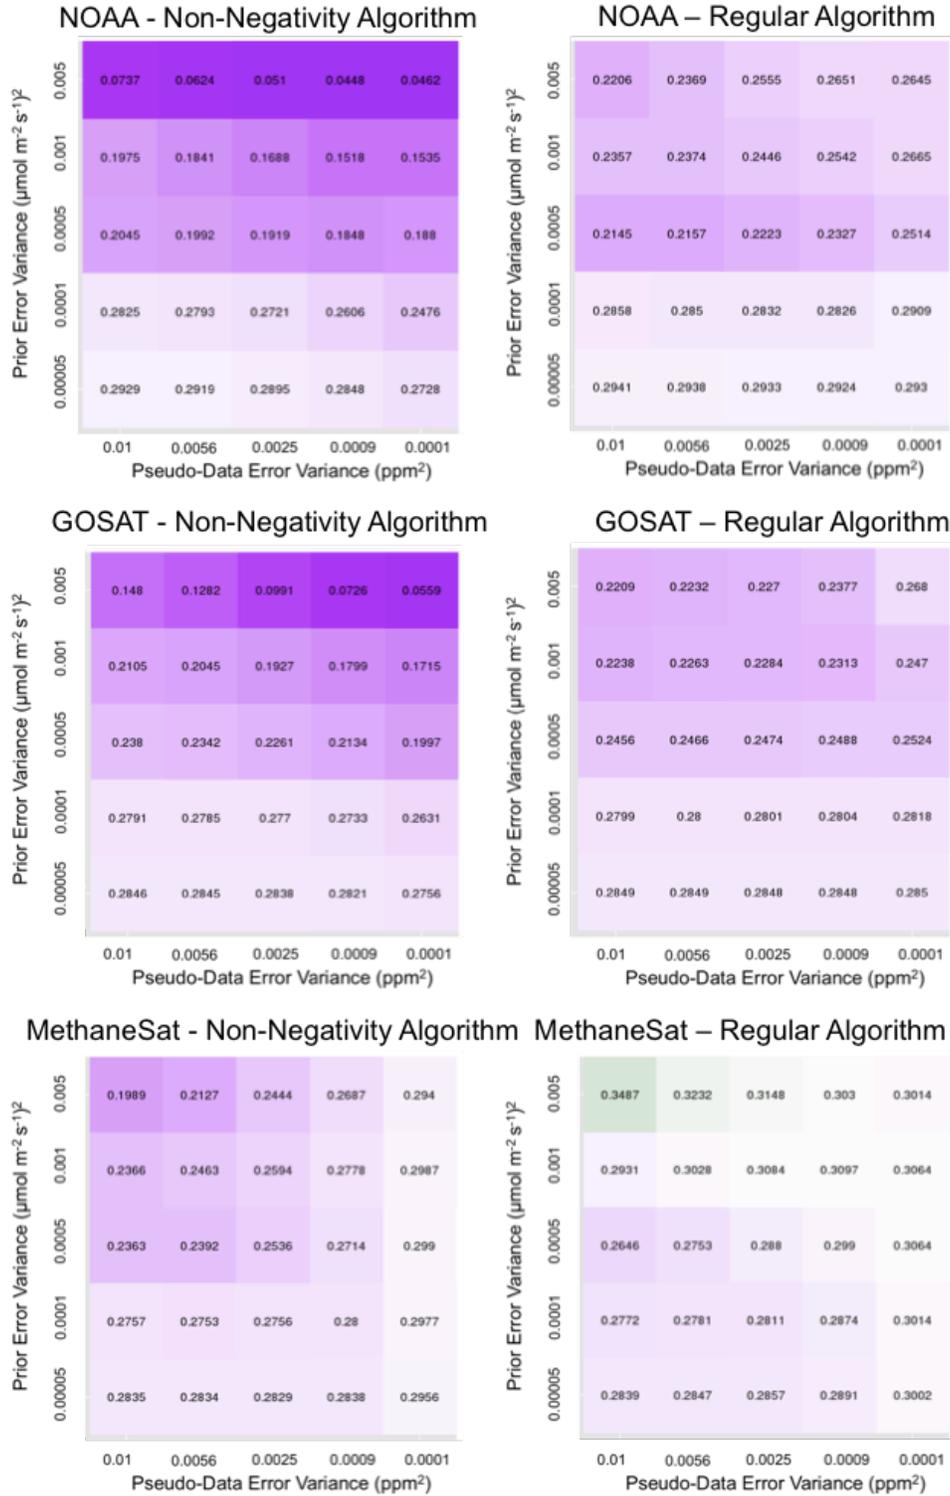
Figure 26: Average posterior emission gradients. White indicates closeness to the true emission gradient of 0.3077 μ mol m^{-2} s^{-1}/m, purple squares are underestimated, and green are overestimated. MethaneSat captured the steep gradient in the Barnett Shale region, while the other systems failed to do so, except for with low prior error variance. The non-negativity algorithm caused more smoothing of the gradient than did the regular algorithm.

53
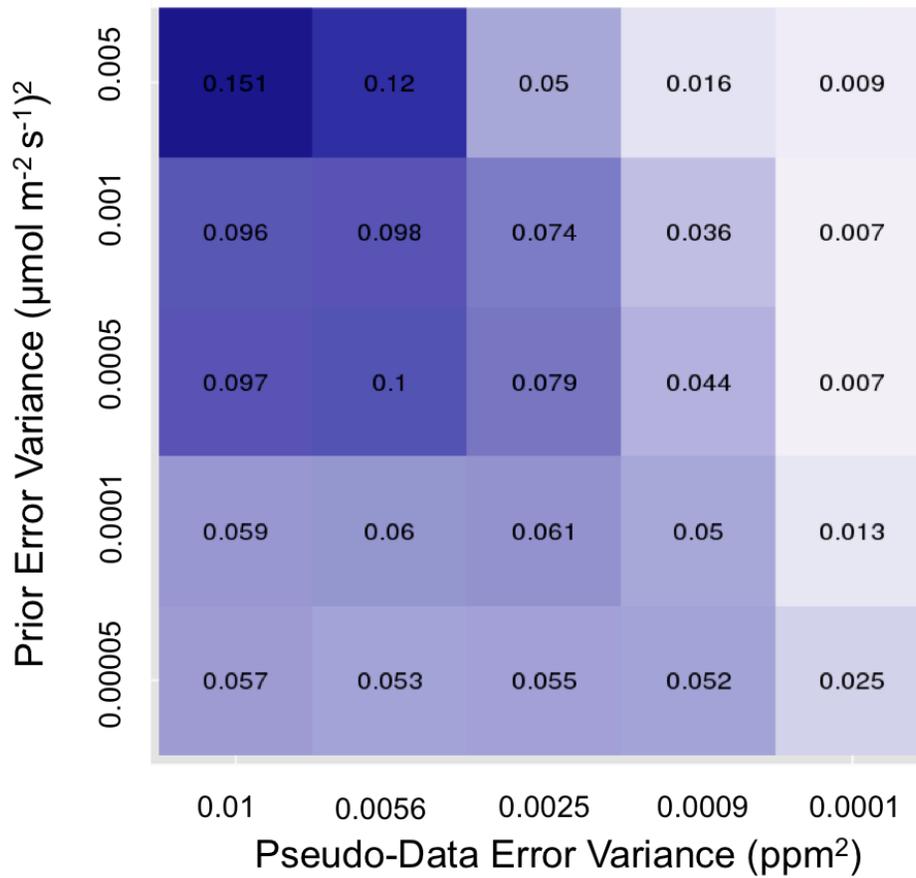
Figure 27: IQR widths for the average posterior emission gradients from MethaneSat. For high values of prior error variance, decreasing pseudo-data error variance dramatically narrows the IQR. Values given in μ mol m^{-2} s^{-1}/m.

Discussion

The results of the OSSEs provide insights about the information that can be gained from the proposed MethaneSat. Comparison against the NOAA surface and aircraft network and GOSAT allowed us to evaluate the advantage of MethansSat over existing observing systems for constraining methane emissions. We performed OSSEs for the Barnett Shale region of Texas, which is an example of the type of oil and gas production region that MethaneSat will target.

The results of the study can be broadly divided into two main categories: support for MethaneSat as a tool for constraining total emissions, their spatial distribution, and sharp gradients; and insights about the role of the prior estimate in constraining emissions.

In every type of OSSE performed in this study, MethaneSat was able to provide information not captured by the NOAA network and GOSAT.

In the idealized scenarios with no error on the pseudo-data, MethaneSat was able to constrain total emissions and their spatial distributions, as well as distinguish signal from noise. MethaneSat captured the total emissions nearly perfectly at every level of prior error, and had extremely narrow interquartile

ranges (IQRs) in each case. This was true for both Method A and Method B. Even when the other observing systems captured the total emissions on average (as GOSAT often was able to), the IQR was quite wide, indicating great variability between iterations. MethaneSat was also able to retrieve the spatial distribution in the Barnett Shale region almost perfectly, while both the NOAA network and GOSAT failed to do so, instead smoothing out the emissions profile. Finally, MethaneSat was better able to distinguish between signal and noise; when prior error was added only to the Barnett Shale region in Method A, the NOAA network and GOSAT posterior estimates included extensive corrections to the prior estimate in regions without prior error. MethaneSat's high density of observations allowed it to distinguish between the perfectly specified emissions outside the Barnett Shale region and the perturbed emissions within the Barnett Shale region.

In the OSSEs with error applied to the pseudo-data, MethaneSat again provided information not provided by the NOAA network and GOSAT. For low levels of prior error, all three observing systems were able to constrain total emissions and their distributions quite well. However, for high levels of prior error, MethaneSat was the only observing system capable of constraining total emissions and capturing the peak in the center of the region. The NOAA network and GOSAT greatly underestimated the high-emitting grid cells, although they were able to constrain the low-emitting grid cells. Furthermore, MethaneSat was the only observing system capable of retrieving the steep emission gradient of 0.3076 μmol m^{-2} s^{-1}/m at every level of error. For all but the lowest two values of prior error variance, the NOAA network and GOSAT smoothed the steep gradient. MethaneSat was able to retrieve the

gradient within a reasonable margin of error even with the two highest levels of prior error variance, particularly when error on the pseudo-data was low.

This contrast between the abilities of the observing systems to constrain emissions is likely a function of the interior sampling density. The high sampling density of MethaneSat allows it to capture more information about the emissions profile in the region, translating to tighter constraint of total emissions, their distribution, and the size of their gradients. The NOAA network and GOSAT rely on much sparser pseudo-data, and are therefore less capable of correcting a poor prior estimate.

The results of this study also provide interesting insights into the role of the prior emission estimate in this type of inversion.

In these OSSEs, changing the level of uncertainty on the prior generally impacted the posterior estimate more than did changing the level of uncertainty on the pseudo-data. In both Method A and Method B, decreasing the prior error variance allowed all observing systems to successfully retrieve total emissions and their spatial distribution. However, decreasing the pseudo-data error variance did not have the same effect, barely improving the posterior estimate for the NOAA network and GOSAT. Furthermore, posterior error (calculated using $\mathbf{V}_{\hat{s}}$ from Equation 5) fell faster with decreasing error on the prior than with decreasing error on the pseudo-data, for all three observing systems.

These results seem to indicate that the uncertainty on the prior estimate exerts more influence on the posterior estimate than does uncertainty on the pseudo-data. Indeed, the levels of both types of error variance chosen for this

study reflect realistic levels of uncertainty. However, the prior uncertainty levels that we tested range from providing a nearly perfect prior estimate to an almost completely uninformative prior estimate. This range is important in these OSSEs in order to acquire a full picture of the information provided by MethaneSat. In the real world, however, the error variance on a prior estimate for a given region would likely be a smaller subset of the range tested here. Thus, the posterior estimate is not necessarily more sensitive to the prior uncertainty – rather, the range we tested is effectively wider than the range tested for pseudo-data uncertainty.

Another interesting result is the wide variability in the posterior total emission estimates produced with the NOAA network and GOSAT. Even when there was no error applied to the pseudo-data, meaning that the only difference between iterations was the prior estimate, both systems gave a wide IQR – much wider than that of MethaneSat. This indicates that the NOAA network and GOSAT rely a great deal on the prior emission estimate. In contrast, MethaneSat's high sampling density was able to constrain emissions regardless of the prior estimate. This could present problems in using the NOAA network or GOSAT to constrain emissions in regions where we have little prior knowledge. Although MethaneSat did not always provide perfect results at high levels of prior error variance, it performed well in constraining total emissions, their distribution, and steep gradients in the region of interest, at all levels of error tested.

Other Findings

Another interesting result of this study is the performance of the non-negativity algorithm. Overall, the non-negativity algorithm produced worse results than the regular algorithm. Emissions were overestimated, and the sharp gradient was smoothed. The effect was more severe with high levels of prior error variance. These results are not unsurprising, considering the strategy of the non-negativity algorithm. When there is a sharp peak in emissions, the optimal estimate ($\hat{\boldsymbol{s}}$, from Equation 4) will sometimes include negative values in some grid cells. The non-negativity algorithm recalculates these values, and therefore does not necessarily conserve the total emissions or the gradients in the same way that the regular algorithm does. Overall, the non-negativity algorithm resulted in more error when constraining total emissions, their distribution, and sharp gradients. These results should be considered when using this algorithm in future studies.

A final interesting result was that, while MethaneSat performed better than the NOAA network and GOSAT in constraining the total mean, spatial distribution, and gradient of emissions, posterior error (calculated from $\mathbf{V}_{\hat{\boldsymbol{s}}}$) was comparable for all three systems for the full OSSEs. This is possibly a result of error in the transport model. However, despite the similar pixel-to-pixel posterior error across all three observing systems, MethaneSat's high spatial resolution was able to overcome this limitation for aggregated properties, such as total mean emissions and their spatial gradient. This is a promising result, given the importance of these aggregated properties in real-world applications.

Conclusion

There is currently a gap in available observing systems' abilities to quantify methane emissions from oil and gas production regions around the world, and discriminate them from other sources. We performed observation system simulation experiments (OSSEs) to demonstrate the new information that would be provided by MethaneSat, a satellite that has been proposed to fill this gap.

Our results were informative and encouraging. The OSSEs demonstrated the ability of MethaneSat to constrain total emissions in the Barnett Shale region, as well as their spatial distribution and sharp gradients. The high variability of the results for the NOAA network and GOSAT indicate heavy reliance on the prior estimate, and the insufficiency of those observations for correcting a poor prior estimate. Even with high uncertainty on the prior estimate, MethaneSat was able to closely constrain methane emissions in the Barnett Shale region, likely due to its high sampling density. This is support for the great value that MethaneSat would add in constraining emissions from regions for which we have little information, or from regions in which we suspect the presence of steep emission gradients.

The results from this experiment invite further investigation. Another

valuable topic to address would be how well the various observing systems distinguish between source types, such as oil and gas production sites, cattle farms, and landfills. The Barnett Shale region itself includes more methane sources than just oil and gas, so this could be a potential extension of the present study. Another extension would be to potentially increase the resolution of the inversion, given the high proposed resolution of MethaneSat. This would allow us to investigate the question of at how high a resolution MethaneSat can retrieve emissions.

Another topic to consider is the generation of the prior emission estimate. We generated the prior estimate by perturbing the true emission field with a Gaussian random field. Another option could be to generate the perturbations from a Gamma distribution, which is positive but can have zero bias around a positive mean. A third possibility would be to generate a prior estimate based on another existing emissions inventory, such as the one produced by the Emissions Database for Global Atmospheric Research (EDGAR), scaled to be unbiased with respect to our target emissions inventory. In either case, it might also be useful to use the same prior estimates across observing systems, in order to eliminate the effect of these differences on the results.

Appendix

The total posterior emissions for the Barnett Shale region, not including any surrounding cells, can be found in Tables 4-9. The units for pseudo-data error variance are ppm^2, and the units for prior error variance are $(\mu\text{mol m}^{-2}\text{ s}^{-1})^2$.

Table 4: Total posterior emissions (Tg/year) in the Barnett Shale region from the NOAA network, using the non-negativity algorithm.

NOAA Network (No-Neg Algorithm) - Barnett Exclusively					
Prior Error	Pseudo-Data Error Variance				
Variance	0.01	0.0056	0.0025	0.0009	0.0001
0.005	1.458	1.355	1.249	1.146	1.039
0.001	1.033	1.001	0.959	0.915	0.889
0.0005	0.880	0.882	0.882	0.871	0.877
0.0001	0.899	0.873	0.836	0.788	0.756
0.00005	0.858	0.851	0.839	0.819	0.771

Table 5: Total posterior emissions (Tg/year) in the Barnett Shale region from GOSAT, using the non-negativity algorithm.

GOSAT (No-Neg Algorithm) - Barnett Exclusively					
Prior Error	Pseudo-Data Error Variance				
Variance	0.01	0.0056	0.0025	0.0009	0.0001
0.005	1.328	1.191	0.997	0.821	0.672
0.001	0.963	0.928	0.879	0.822	0.773
0.0005	0.992	0.963	0.915	0.847	0.773
0.0001	0.874	0.869	0.860	0.845	0.800
0.00005	0.746	0.747	0.748	0.752	0.759

Table 6: Total posterior emissions (Tg/year) in the Barnett Shale region from MethaneSat, using the non-negativity algorithm.

CH4Sat (No-Neg Algorithm) - Barnett Exclusively					
Prior Error	Pseudo-Data Error Variance				
Variance	0.01	0.0056	0.0025	0.0009	0.0001
0.005	1.156	1.129	1.043	0.923	0.874
0.001	0.874	0.885	0.873	0.859	0.798
0.0005	0.865	0.863	0.880	0.869	0.819
0.0001	0.744	0.754	0.765	0.785	0.800
0.00005	0.754	0.757	0.763	0.774	0.794

Table 7: Total posterior emissions (Tg/year) in the Barnett Shale region from the NOAA network, using the regular algorithm.

NOAA Network - Barnett Exclusively					
Prior Error	Pseudo-Data Error Variance				
Variance	0.01	0.0056	0.0025	0.0009	0.0001
0.005	0.291	0.325	0.407	0.678	0.543
0.001	0.710	0.684	0.690	0.714	0.734
0.0005	0.745	0.727	0.740	0.766	0.769
0.0001	0.847	0.802	0.730	0.663	0.670
0.00005	0.845	0.829	0.803	0.762	0.724

Table 8: Total posterior emissions (Tg/year) in the Barnett Shale region from GOSAT, using the regular algorithm.

	GOSAT - Barnett Exclusively				
Prior Error Variance	Pseudo-Data Error Variance				
	0.01	0.0056	0.0025	0.0009	0.0001
0.005	0.864	0.811	0.768	0.729	0.709
0.001	0.810	0.784	0.775	0.797	0.786
0.0005	0.849	0.789	0.731	0.710	0.725
0.0001	0.859	0.845	0.816	0.779	0.753
0.00005	0.741	0.738	0.730	0.719	0.722

Table 9: Total posterior emissions (Tg/year) in the Barnett Shale region from MethaneSat, using the regular algorithm.

	CH4Sat - Barnett Exclusively				
Prior Error Variance	Pseudo-Data Error Variance				
	0.01	0.0056	0.0025	0.0009	0.0001
0.005	0.691	0.724	0.751	0.774	0.821
0.001	0.742	0.735	0.736	0.756	0.772
0.0005	0.797	0.792	0.790	0.795	0.796
0.0001	0.738	0.744	0.753	0.770	0.790
0.00005	0.752	0.753	0.757	0.765	0.782

Bibliography

M. Alexe, P. Bergamaschi, A. Segers, R. Detmers, A. Butz, O. Hasekamp, S. Guerlet, R. Parker, H. Boesch, C. Frankenberg, R. A. Scheepmaker, E. Dlugokencky, C. Sweeney, S. C. Wofsy, and E. A. Kort. Inverse modelling of ch4 emissions for 2010 2011 using different satellite retrieval products from gosat and sciamachy. *Atmospheric Chemistry and Physics*, 15(1):113–133, 2015. doi: 10.5194/acp-15-113-2015.

G. P. Brasseur and D. J. Jacob. *Mathematical Modeling of Atmospheric Chemistry*. 2015.

Britannica Academic. Shale gas, 2016. URL `https://www.britannica.com/science/shale-gas`. Accessed: 2017-01-11.

L. M. Bruhwiler, E. Dlugokencky, K. Masarie, M. Ishizawa, A. Andrews, J. Miller, C. Sweeney, P. Tans, and D. Worthy. Carbontracker- ch4: an assimilation system for estimating emissions of atmospheric methane. *Atmospheric Chemistry and Physics Discussions*, 14(2):2175–2233, 2014. doi: 10.5194/acpd-14-2175-2014.

E. J. Dlugokencky, L. Bruhwiler, J. W. C. White, L. K. Emmons, P. C. Novelli, S. A. Montzka, K. A. Masarie, P. M. Lang, A. M. Crotwell, J. B. Miller, and

L. V. Gatti. Observational constraints on recent increases in the atmospheric ch4 burden. *Geophysical Research Letters*, 36(18):n/a–n/a, 2009. doi: 10.1029/2009GL039780.

Environmental Defense Fund. Epa's new source rules for oil and gas operations, 2016. URL `https://www.edf.org/sites/default/files/content/methane_rule_fact_sheet_in_brief_spring_2016_final.pdf`. Accessed: 2017-01-03.

Environmental Protection Agency. Greenhouse gas emissions: Overview of greenhouse gases, 2017. URL `https://www.epa.gov/ghgemissions/overview-greenhouse-gases`. Accessed: 2017-01-12.

Peter Fairley. Private satellite to track carbon polluters. IEEE Spectrum, 2016. URL `http://spectrum.ieee.org/energy/environment/private-satellite-to-track-carbon-polluters`. Accessed: 2017-01-12.

J. D. Goetz, C. Floerchinger, E. C. Fortner, J. Wormhoudt, P. Massoli, W. B. Knighton, S. C. Herndon, C. E. Kolb, E. Knipping, S. L. Shaw, and P. F. DeCarlo. Atmospheric emission characterization of marcellus shale natural gas development sites. *Environmental Science & Technology*, 49(11):7012–7020, 2015. doi: 10.1021/acs.est.5b00452. URL `http://dx.doi.org/10.1021/acs.est.5b00452`. PMID: 25897974.

R. Harriss, R. A. Alvarez, D. Lyon, D. Zavala-Araiza, D. Nelson, and S. P. Hamburg. Using multi- scale measurements to improve methane emission estimates from oil and gas operations in the barnett shale region, texas.

Environmental Science and Technology, 49(13):7524, 2015. doi: 10.1021/acs.est.5b02305.

P. Hausmann, R. Sussmann, and D. Smale. Contribution of oil and natural gas production to renewed increase in atmospheric methane (2007 2014): top down estimate from ethane and methane column observations. *Atmospheric Chemistry and Physics*, 16(5):3227–3244, 2016. doi: 10.5194/acp-16-3227-2016.

J. Hegarty, R. R. Draxler, A. F. Stein, J. Brioude, M. Mountain, J. Eluszkiewicz, T. Nehrkorn, F. Ngan, and A. Andrews. Evaluation of lagrangian particle dispersion models with measurements from controlled tracer releases. *Journal of Applied Meteorology and Climatology*, 52(12): 2623–2637, 2013. doi: 10.1175/JAMC-D-13-0125.1.

S. Houweling, P. Bergamaschi, F. Chevallier, M. Heimann, T. Kaminski, M. Krol, A. M. Michalak, and P. Patra. Global inverse modeling of ch4 sources and sinks: An overview of methods. *Atmospheric Chemistry and Physics Discussions*, pages 1–30, 2016. doi: 10.5194/acp-2016-572.

R. Howarth, R. Santoro, and A. Ingraffea. Methane and the greenhouse-gas footprint of natural gas from shale formations. *Climatic Change; An Interdisciplinary, International Journal Devoted to the Description, Causes and Implications of Climatic Change*, 106(4):679–690, 2011. doi: 10.1007/s10584-011-0061-5.

K. Hungershoefer, F. M. Breon, P. Peylin, F. Chevallier, P. Rayner, A. Klo-

necki, S. Houweling, and J. Marshall. Evaluation of various observing systems for the global monitoring of co2 surface fluxes. 2010.

Intergovernmental Panel on Climate Change. Ipcc fourth assessment report. 2:580–595, 2007. URL `https://www.ipcc.ch/publications_and_data/publications_ipcc_fourth_assessment_report_synthesis_report.htm`.

A. Karion, C. Sweeney, E. A. Kort, P. B. Shepson, A. Brewer, M. Cambaliza, S. A. Conley, K. Davis, A. Deng, M. Hardesty, et al. Aircraft-based estimate of total methane emissions from the barnett shale region. *Environmental science & technology*, 49(13):8124–8131, 2015.

E. Kintisch and K. Buckheit. Along the road from kyoto: global greenhouse gas emissions keep rising. *Science*, 311(5768):1702, 2006.

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:ACH 2–ACH 2–17; 1, 2003. doi: 10.1029/2002JD003161.

D. R. Lyon, D. Zavala-Araiza, R. A. Alvarez, R. Harriss, V. Palacios, X. Lan, R. Talbot, T. Lavoie, P. Shepson, T. I. Yacovitch, S. C. Herndon, A. J. Marchese, D. Zimmerle, A. L. Robinson, and S. P. Hamburg. Constructing a spatially resolved methane emission inventory for the barnett shale region. *Environmental Science and Technology*, 49(13):8147, 2015.

A. M. Michalak. A gibbs sampler for inequality–constrained geostatistical interpolation and inverse modeling. *Water Resources Research*, 44(9), 2008.

A. M. Michalak, L. Bruhwiler, and P. P. Tans. A geostatistical approach to surface flux estimation of atmospheric trace gases. *Journal of Geophysical Research*, 109, 2004. doi: 10.1029/2003JD004422.

S. M. Miller, A. M. Michalak, and P. J. Levi. Atmospheric inverse modeling with known physical bounds: an example from trace gas emissions. *Geoscientific Model Development*, 7(1):303–315, 2014a. doi: 10.5194/gmd-7-303-2014.

S. M. Miller, D. E. J. Worthy, A. M. Michalak, S. C. Wofsy, E. A. Kort, T. C. Havice, A. E. Andrews, E. J. Dlugokencky, J. O. Kaplan, P. J. Levi, H. Tian, and B. Zhang. Observational constraints on the distribution, seasonality, and environmental predictors of north american boreal methane emissions. *Global Biogeochemical Cycles*, 28(2):146–160, 2014b. doi: 10.1002/2013GB004580.

S. A. Montzka, E. J. Dlugokencky, and J. H. Butler. Non-co2 greenhouse gases and climate change, 2011.

National Aeronautics and Space Administration. Gmao observing system simulation experiments (osse), 2012. URL `https://gmao.gsfc.nasa.gov/projects/osse/`. Accessed: 2017-01-09.

National Conference of State Legislatures. State methane policies, 2014. URL `http://www.ncsl.org/research/environment-and-natural-resources/state-methane-policies.aspx`.

J. Peischl, A. Karion, C. Sweeney, E. A. Kort, M. L. Smith, A. R. Brandt, T. Yeskoo, K. C. Aikin, S. A. Conley, A. Gvakharia, et al. Quantifying atmospheric methane emissions from oil and natural gas production in the bakken shale region of north dakota. *Journal of Geophysical Research: Atmospheres*, 121(10):6101–6111, 2016.

R. G. Prinn, R. F. Weiss, B. R. Miller, J. Huang, F. N. Alyea, D. M. Cunnold, P. J. Fraser, D. E. Hartley, and P. G. Simmonds. Atmospheric trends and lifetime of ch₃ccl₃ and global oh concentrations. *Science*, 269(5221):187–192, 1995. URL `http://www.jstor.org.ezp-prod1.hul.harvard.edu/stable/2888434`.

M. Rigby, R. G. Prinn, P. J. Fraser, P. G. Simmonds, R. L. Langenfelds, J. Huang, D. M. Cunnold, L. P. Steele, P. B. Krummel, R. F. Weiss, S. O'Doherty, P. K. Salameh, H. J. Wang, C. M. Harth, J. Mhle, and L. W. Porter. Renewed growth of atmospheric methane. *Geophysical Research Letters*, 35(22):n/a–n/a, 2008. doi: 10.1029/2008GL036037.

H. Schaefer, S. E. M. Fletcher, C. Veidt, K. R. Lassey, G. W. Brailsford, T. M. Bromley, E. J. Dlugokencky, S. E. Michel, J. B. Miller, I. Levin, D. C. Lowe, R. J. Martin, B. H. Vaughn, and J. W. C. White. A 21st-century shift from fossil- fuel to biogenic methane emissions indicated by sup.13]ch.sub.4].(atmospheric methane)(author abstract). *Science*, 352 (6281):80, 2016. doi: 10.1126/science.aad2705.

Martin Schlather, Alexander Malinowski, Peter J. Menck, Marco Oesting, and Kirstin Strokorb. Analysis, simulation and prediction of multivariate

random fields with package RandomFields. *Journal of Statistical Software*, 63(8):1–25, 2015. URL `http://www.jstatsoft.org/v63/i08/`.

J. K. Shoemaker and D. P. Schrag. The danger of overvaluing methanes influence on future climate change. *Climate Change*, 2013.

Thomas Sterner. Higher costs of climate change. *Nature*, 527(7577):177–178, 2015.

The Cooperative Institute for Meteorological Satellite Studies. Observing system simulation experiments (osse) for geostationary satellite retrievals, 1999. URL `https://cimss.ssec.wisc.edu/model/osse/osse.html`.

A. J. Turner, D. J. Jacob, J. Benmergui, S. C. Wofsy, J. D. Maasakkers, A. Butz, O. Hasekamp, and S. C. Biraud. A large increase in u.s. methane emissions over the past decade inferred from satellite data and surface observations. *Geophysical Research Letters*, 43(5):2218–2224, 2016. ISSN 1944-8007. doi: 10.1002/2016GL067987. URL `http://dx.doi.org/10.1002/2016GL067987`. 2016GL067987.

U.S. Energy Information Administration. Natural gas 1998: Issues and trends, 1999. URL `https://www.eia.gov/pub/oil_gas/natural_gas/analysis_publications/natural_gas_1998_issues_trends/pdf/it98.pdf`. Accessed: 2017-01-03.

U.S. Energy Information Administration. Natural gas explained: Where our natural gas comes from, 2017. URL `https://www.eia.gov/energyexplained/index.cfm?page=natural_gas_where`. Accessed: 2016-12-22.

A.L. Waddams, L. H. Solomon, J. E. Carruthers, J. P. Riva, and G. I. Atwater. Natural gas, 2012. URL `https://www.britannica.com/science/natural-gas`.

White House, Executive Office of the President. The presidents climate action plan. *Washington, DC*, 2013. URL `https://obamawhitehouse.archives.gov/sites/default/files/image/president27sclimateactionplan.pdf`.

Lin Wu, Grgoire Broquet, Philippe Ciais, Valentin Bellassen, Felix Vogel, Frdric Chevallier, Ir Xueref-Remy, and Yilong Wang. What would dense atmospheric observation networks bring to the quantification of city co2 emissions? *Atmospheric Chemistry and Physics*, 16(12):7743–7771, 2016. doi: 10.5194/acp-16-7743-2016.

V. Yadav and A. M. Michalak. Improving computational efficiency in large linear inverse problems: an example from carbon dioxide flux estimation. *Geoscientific Model Development*, 6(3):583–590, 2013. doi: 10.5194/gmd-6-583-2013.

D. Zavala-Araiza, D. R. Lyon, R. A. Alvarez, K. J. Davis, R. Harriss, S. C. Herndon, A. Karion, E. A. Kort, B. K. Lamb, X. Lan, A. J. Marchese, S. W. Pacala, A. L. Robinson, P. B. Shepson, C. Sweeney, R. Talbot, A. Townsend-Small, T. I. Yacovitch, D. J. Zimmerle, and S. P. Hamburg. Reconciling divergent estimates of oil and gas methane emissions. *Proceedings of the National Academy of Sciences of the United States of America*, 112 (51):15597–15602, Dec 22 2015. doi: 10.1073/pnas.1522126112[doi]. LR: