

Lawrence Berkeley National Laboratory

LBL Publications

Title

Reduced-cost construction of Jacobian matrices for high-resolution inversions of satellite observations of atmospheric composition

Permalink

<https://escholarship.org/uc/item/2663x93q>

Journal

Atmospheric Measurement Techniques, 14(8)

ISSN

1867-1381

Authors

Nesser, Hannah

Jacob, Daniel J

Maasackers, Joannes D

et al.

Publication Date

2021

DOI

10.5194/amt-14-5521-2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Reduced-Cost Construction of Jacobian Matrices for High-Resolution Inversions of Satellite Observations of Atmospheric Composition

Hannah Nesser¹, Daniel J. Jacob¹, Joannes D. Maasakkers², Tia R. Scarpelli³, Melissa P. Sulprizio¹,
5 Yuzhong Zhang^{4,5}, Chris H. Rycroft¹

¹School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA

²SRON Netherlands Institute for Space Research, Utrecht, the Netherlands

³Department of Earth and Planetary Sciences, Harvard University, Cambridge, MA, USA

⁴Key Laboratory of Coastal Environment and Resources of Zhejiang Province, School of Engineering, Westlake University,
10 Hangzhou, Zhejiang, China

⁵Institute of Advanced Technology, Westlake Institute for Advanced Study, Hangzhou, Zhejiang, China

Correspondence to: Hannah Nesser (hnesser@g.harvard.edu)

Abstract. Global high-resolution observations of atmospheric composition from satellites can greatly improve our understanding of surface emissions through inverse analyses. Variational inverse methods can optimize surface emissions at
15 any resolution but do not readily quantify the error and information content of the posterior solution. The information content of satellite data may be much lower than its coverage would suggest because of failed retrievals, instrument noise, and error correlations that propagate through the inversion. Analytical solution of the inverse problem provides closed-form characterization of posterior error statistics and information content but requires the construction of the Jacobian matrix that relates emissions to atmospheric concentrations. Building the Jacobian matrix is computationally expensive at high
20 resolution because it involves perturbing each emission element, typically individual grid cells, in the atmospheric transport model used as forward model for the inversion. We propose and analyze two methods, reduced-dimension and reduced-rank, to construct the Jacobian matrix at greatly decreased computational cost while retaining information content. Both methods are two-step iterative procedures that begin from an initial native-resolution estimate of the Jacobian matrix constructed at no computational cost by assuming that atmospheric concentrations are most sensitive to local emissions. The reduced-
25 dimension method uses this estimate to construct a Jacobian matrix on a multiscale grid that maintains high resolution in areas with high information content and aggregates grid cells elsewhere. The reduced-rank method constructs the Jacobian matrix at native resolution by perturbing the leading patterns of information content given by the initial estimate. We demonstrate both methods in an analytical Bayesian inversion of GOSAT methane satellite data with augmented information content over North America in July 2009. We show that both methods reproduce the results of the native-resolution
30 inversion while achieving a factor of four improvement in computational performance. The reduced-dimension method produces an exact solution at lower spatial resolution while the reduced-rank method solves the inversion at native resolution in areas of high information content and defaults to the prior estimate elsewhere.

1 Introduction

Satellite observations of atmospheric composition provide a powerful resource to improve our knowledge of emissions (Streets et al., 2013). However, the inverse analyses used to infer emissions from observed atmospheric concentrations are subject to large errors from the measurements and from the inversion procedure. Conducting inverse analyses of satellite data to quantify emissions at high resolution is of considerable interest but may be limited by data quality in ways that are difficult to quantify and that may compromise the results. Here we present two methods to conduct high-resolution inversions of satellite observations that optimize the information content of the observations while providing full error statistics and minimizing computational cost.

Inverse analyses infer emissions by fitting the observed atmospheric concentrations to a chemical transport model (CTM) that simulates atmospheric concentrations as a function of emissions (Brasseur and Jacob, 2017). The CTM represents the forward model for the inverse problem. The solution is generally obtained in a Bayesian framework by minimizing a cost function regularized by a prior emissions estimate. The optimal (posterior) emissions estimate corresponds to the minimum of the cost function. This minimum is typically found using a numerical (variational) method, often employing the adjoint of the CTM to compute the cost function gradient (e.g., Daescu et al., 2000; Elbern and Schmidt, 2001; Quélo et al., 2005; Henze et al., 2007). However, the numerical solution provides no explicit characterization of the solution’s error or information content. Methods of estimating the error exist (e.g., Chevallier et al., 2007; Meirink et al., 2008; Koohkan et al., 2013), but these approaches are computationally expensive, incomplete, and rarely applied in practice.

In the common case where the observed atmospheric concentrations depend linearly on emissions and the error statistics can be assumed to be normally or log-normally distributed, the Bayesian optimization problem has an analytical solution including closed-form expressions for the posterior emissions estimate, its error statistics, and its information content (Rodgers, 2000; Maasakkers et al., 2019). The analytical solution requires explicit construction of the Jacobian matrix of the forward model, $\mathbf{K} = \partial\mathbf{y}/\partial\mathbf{x} \in \mathbb{R}^{m \times n}$, which represents the sensitivity of the simulated concentrations $\mathbf{y} \in \mathbb{R}^m$ to the emission state vector $\mathbf{x} \in \mathbb{R}^n$ (Brasseur and Jacob, 2017). The elements of \mathbf{y} are individual observations and the elements of \mathbf{x} are the emissions optimized by the inversion, often grid cells in a two-dimensional emissions field. When $m \gg n$, as for inversions of satellite observations, the Jacobian can be constructed column-wise by conducting $n + 1$ CTM simulations to perturb each of the state vector elements x_i and obtain the corresponding column $\partial\mathbf{y}/\partial x_i$. Even on massively parallel computing clusters, the computational cost of conducting these simulations can limit the size of the state vector \mathbf{x} and therefore the resolution at which inversions are conducted (Turner and Jacob, 2015). However, once the Jacobian matrix is constructed, inversions can be conducted at essentially no additional computational cost, allowing study of the solution’s sensitivity to changes in the specification of inversion parameters, error statistics, prior assumptions, and the number and type of observations.

An illustrative example is the inversion of satellite observations to infer methane emissions. Methane is an important greenhouse gas but the spatial and temporal distribution of emissions is highly uncertain (Saunois et al., 2019). Satellite observations of atmospheric methane columns can inform emission estimates (Jacob et al., 2016). This was first shown with data from the SCIAMACHY satellite instrument (2003 – 2012) with nadir pixel resolution of 30 x 60 km² (Bergamaschi et al., 2009, 2013; Houweling et al., 2014; Wecht et al., 2014). More recent inversions used observations from the TANSO-FTS instrument aboard the GOSAT satellite (2009 – present) with 10-km diameter pixels approximately 250 km apart along- and cross-track (Monteil et al., 2013; Alexe et al., 2015; Turner et al., 2015; Maasackers et al., 2019). The Tropospheric Monitoring Instrument (TROPOMI) aboard the Sentinel-5 precursor satellite, launched in October 2017, now provides daily, global retrievals of atmospheric methane columns at 5.5 x 7 km² nadir pixel resolution, increasing coverage by orders of magnitude relative to GOSAT (Veefkind et al., 2012). However, TROPOMI’s methane retrieval has only a ~3% success rate for daytime scenes limited by dark surfaces (water), clouds, high aerosol loadings, and variable surface albedo and topography, resulting in heterogeneously distributed observations (Hu et al., 2018; Hasekamp et al., 2019). Inversions of TROPOMI data must attempt to capture the high resolution and density of observations where appropriate while recognizing the limitations in information content resulting from data sparsity and observational errors.

Several methods have been proposed to decrease the computational cost of high-resolution analytical inversions by optimally reducing the dimension or rank of the observations or state vector. Approaches that decrease the dimension of the observation vector (e.g., Xu, 2007) reduce the computational cost of solving the inversion but not of constructing the Jacobian matrix. Approaches that decrease the dimension of the state vector lower the cost of both computations. Reduced-dimension methods solve inversions on a multiscale emission grid of dimension $k < n$ for which the construction of the Jacobian matrix $\mathbf{K} \in \mathbb{R}^{m \times k}$ is computationally tractable. Bocquet et al. (2011) and Bocquet and Wu (2011) defined a method to select a multiscale grid from a limited array of allowable grids that preserve resolution where the observations have the highest information content. Turner and Jacob (2015) used prior emissions information to group together similar grid cells using a Gaussian mixture model, but the criteria used to define similarity were subjective and did not consider the information content of the forward model or the observations. Other approaches that decreased the dimension of the state vector assumed knowledge of the Jacobian matrix (e.g., Rigby et al., 2011; Thompson and Stohl, 2014; Ray et al., 2015; Lunt et al., 2016; Liu et al., 2017). Reduced-rank methods generate an approximation of the posterior solution at the original dimension n by solving the inversion in the directions of highest information content. The reduced-rank method proposed by Spantini et al. (2015) assumed knowledge of the Jacobian matrix. Bousserez and Henze (2018) and Miller et al. (2020) avoided explicit construction of the Jacobian matrix by estimating the directions of highest information content, but their approach is effective only if a small number of directions explain most of the information content.

Here we present two methods to construct the Jacobian matrix for a native n -dimensional state vector and maximize the information content of the inverse analysis using $k < n$ forward model simulations. The reduced-dimension method generates a multiscale grid that preserves native resolution where information content is highest and aggregates grid cells elsewhere. The resulting reduced-dimension Jacobian matrix $\mathbf{K}_{\text{RD}} \in \mathbb{R}^{m \times k}$ solves the inversion exactly on the multiscale grid. The reduced-rank method constructs a Jacobian matrix $\mathbf{K}_{\text{rr}} \in \mathbb{R}^{m \times n}$ along the dominant patterns of information content in the system, allowing the approximation of the inverse solution at native resolution. In both cases, a low-cost initial estimate of the Jacobian matrix is updated using k forward model simulations where k is selected by the user based on the information content of the observing system and the available computational resources. We demonstrate both methods in a 1-month inversion of satellite data.

2 Methods

This section presents the reduced-dimension and reduced-rank methods of constructing the Jacobian matrix. Following a review of the standard analytical inverse framework (Section 2.1), we describe optimal reductions in both dimension and rank for an inverse system with a known native-resolution Jacobian matrix $\mathbf{K} \in \mathbb{R}^{m \times n}$ (Section 2.2). We then present a two-step approach to approximate an initially unknown Jacobian matrix using reductions in dimension and rank (Sections 2.3 through 2.5). For the purposes of illustration, we take the state vector to be a gridded field of static emissions, but the methods apply to temporally variable emissions and more generally to any state vector.

2.1 Analytical solution to the inverse problem

The optimal estimate $\hat{\mathbf{x}}$ of a state vector \mathbf{x} given a prior estimate \mathbf{x}_A , observation vector \mathbf{y} , and normal error statistics given by prior and observational error covariance matrices \mathbf{S}_A and \mathbf{S}_O , respectively, is obtained by the minimization of the Bayesian scalar cost function $J(\mathbf{x})$ (Brasseur and Jacob, 2017):

$$J(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_A)^T \mathbf{S}_A^{-1} (\mathbf{x} - \mathbf{x}_A) + (\mathbf{y} - \mathbf{F}(\mathbf{x}))^T \mathbf{S}_O^{-1} (\mathbf{y} - \mathbf{F}(\mathbf{x})). \quad (1)$$

Here $\mathbf{F}(\mathbf{x})$ represents the forward model that simulates the observations \mathbf{y} given \mathbf{x} . In our application, the forward model is a CTM. The observational error covariance matrix \mathbf{S}_O includes errors from both the measurement and the forward model, which collectively form the observing system. If the forward model is linear so that $\mathbf{F}(\mathbf{x}) = \mathbf{K}\mathbf{x} + \mathbf{c}$, where $\mathbf{K} = \partial\mathbf{y}/\partial\mathbf{x}$ is the Jacobian matrix calculated by finite difference (see Introduction) and \mathbf{c} is a constant, then an analytical solution to the cost function minimum exists that yields both the posterior estimate $\hat{\mathbf{x}}$ and its error covariance matrix $\hat{\mathbf{S}}$:

$$\hat{\mathbf{x}} = \mathbf{x}_A + \mathbf{S}_A \mathbf{K}^T (\mathbf{K} \mathbf{S}_A \mathbf{K}^T + \mathbf{S}_O)^{-1} (\mathbf{y} - \mathbf{K} \mathbf{x}_A) = \mathbf{x}_A + \hat{\mathbf{S}} \mathbf{K}^T \mathbf{S}_O^{-1} (\mathbf{y} - \mathbf{K} \mathbf{x}_A), \quad (2)$$

$$\hat{\mathbf{S}} = (\mathbf{K}^T \mathbf{S}_O^{-1} \mathbf{K} + \mathbf{S}_A^{-1})^{-1}. \quad (3)$$

130 Comparison of $\hat{\mathbf{S}}$ and \mathbf{S}_A defines the information content of the observing system, quantified by the averaging kernel matrix $\mathbf{A} = \partial\hat{\mathbf{x}}/\partial\mathbf{x}$ that represents the sensitivity of the posterior emissions estimate $\hat{\mathbf{x}}$ to the true state \mathbf{x} . \mathbf{A} can be calculated as $\mathbf{A} = \mathbf{I} - \hat{\mathbf{S}}\mathbf{S}_A^{-1}$ or equivalently as

$$\mathbf{A} = \mathbf{S}_A\mathbf{K}^T(\mathbf{K}\mathbf{S}_A\mathbf{K}^T + \mathbf{S}_0)^{-1}\mathbf{K}. \quad (4)$$

135

Equation (4) expresses the dependence of the averaging kernel matrix on the forward model and both error covariance matrices. The diagonal elements of \mathbf{A} are commonly referred to as the averaging kernel sensitivities. They are highest in highly observed locations with uncertain, high emissions and lowest in poorly observed areas or in regions known to have low emissions. The sum of the sensitivities, or the trace of \mathbf{A} , measures the number of pieces of information that can be independently quantified by the observing system, known as the degrees of freedom for signal or DOFS (Rodgers, 2000).

140

2.2 Optimal reductions in dimension and rank of inverse systems

We first consider the problem of optimally reducing the dimension and rank of an inverse system as described in Section 2.1 with a known Jacobian matrix $\mathbf{K} \in \mathbb{R}^{m \times n}$. Figure 1 illustrates dimension and rank reductions for an emission grid over North America. The top left panel represents the original n -dimensional state space, i.e., the native-resolution grid. A linear transformation $\mathbf{\Gamma} \in \mathbb{R}^{k \times n}$ reduces the dimension of the state space from n to k . This transformation may reduce dimension discretely, as in the case of grid cell aggregation (top right panel), or non-discretely, in which case the k state vector components are themselves n -dimensional vectors (bottom right panel). A second linear transformation $\mathbf{\Gamma}^* \in \mathbb{R}^{n \times k}$ restores the dimension of the state space from k back to the original n . The resulting space, depicted in the bottom left, is a low-rank approximation of the original state space. The matrix $\mathbf{\Pi} = \mathbf{\Gamma}^*\mathbf{\Gamma}$ transforms the original state space to the low-rank subspace.

145 The inverse problem can be solved in any of these four spaces, although the eigenvector corrections generated in the non-discrete reduced-dimension space (bottom right panel) would be difficult to interpret.


150

We wish to define matrices $\mathbf{\Gamma}$ and $\mathbf{\Gamma}^*$ that minimize the information loss associated with reducing the dimension or rank of the state vector. Bousserez and Henze (2018) show that the projection $\mathbf{\Pi}$ that maximizes the probability of restoring the original full dimension state vector \mathbf{x} given the reduced dimension state vector $\mathbf{\Gamma}\mathbf{x}$ is given by $\mathbf{\Pi} = \mathbf{S}_A^{1/2}\mathbf{U}\mathbf{U}^T\mathbf{S}_A^{-1/2}$ where $\mathbf{U} = \mathbf{S}_A^{1/2}\mathbf{\Gamma}(\mathbf{\Gamma}\mathbf{S}_A\mathbf{\Gamma}^T)^{-1/2}$. For a projection of this form, they show that information loss is minimized by maximizing $\text{DOFS}_{\mathbf{\Pi}} = \text{Tr}(\mathbf{A}_{\mathbf{\Pi}}) = \text{Tr}(\mathbf{U}^T\mathbf{S}_A^{-1/2}\mathbf{A}\mathbf{S}_A^{1/2}\mathbf{U})$ where $\mathbf{A}_{\mathbf{\Pi}}$ and \mathbf{A} are the reduced-rank and native-resolution averaging kernel matrices, respectively. Defining

155

160

$$\mathbf{Q} = \mathbf{S}_A^{-1/2}\mathbf{A}\mathbf{S}_A^{1/2} = \mathbf{W}\mathbf{\Sigma}\mathbf{W}^T, \quad (5)$$



165 **Figure 1. Dimension and rank reductions of a gridded emissions field. The linear transformation matrix Γ reduces the dimension of the original state space (upper left) either discretely by aggregating grid cells to generate a multiscale grid (upper right) or non-discretely by projecting along the patterns given by the rows of Γ (lower right, with positive values in red and negative in blue). The reverse transformation Γ^* restores the dimension but not the rank, producing a low-rank subspace of the original state space (lower right). The projection $\Pi = \Gamma^* \Gamma$ reduces rank but not dimension.**

where the columns of \mathbf{W} are the eigenvectors of \mathbf{Q} and $\mathbf{\Sigma}$ is a diagonal matrix of the corresponding eigenvalues ranked in descending order, Bousserez and Henze (2018) show that $\text{Tr}(\mathbf{A}_\Pi)$ is maximized for a rank- k subspace when $\mathbf{U} = \mathbf{W}_k$ where \mathbf{W}_k is the matrix of the first k columns of \mathbf{W} . The corresponding optimal projection is then

170

$$\mathbf{\Pi} = \mathbf{S}_A^{1/2} \mathbf{W}_k \mathbf{W}_k^T \mathbf{S}_A^{-1/2}. \quad (6)$$

This projection applies a dimension-reducing transformation Γ followed by a dimension-restoring transformation Γ^* :

175

$$\mathbf{\Gamma} = \mathbf{W}_k^T \mathbf{S}_A^{-1/2}, \quad (7)$$

$$\mathbf{\Gamma}^* = \mathbf{S}_A^{1/2} \mathbf{W}_k. \quad (8)$$

The columns of $\mathbf{\Gamma}^*$ give an eigenvector basis for the averaging kernel matrix while the diagonal of $\mathbf{\Sigma}$ give its eigenvalues, together defining the dominant patterns of information content. The fraction of information content explained by the first i columns of $\mathbf{\Gamma}^*$ is the sum of the i largest eigenvalues divided by the total DOFS (Bousserez and Henze, 2018). The eigenvalues can also be related to other measures of information content, including the Shannon and relative entropy differences (Rodgers, 2000; Xu, 2007). We will refer to the ordered list of the eigenvalues as the information content spectrum. On the basis of this spectrum, we can select k so that most of the information content is explained by the first k eigenvectors. Alternatively, we can select k so that all eigenvectors have a sufficiently large signal-to-noise ratio. The signal-to-noise ratio SNR of the i th eigenvector is given by the i th singular value of the pre-whitened Jacobian matrix $\tilde{\mathbf{K}} = \mathbf{S}_0^{-1/2} \mathbf{K} \mathbf{S}_A^{1/2}$ and is calculated as

$$\text{SNR}_i = \sqrt{\frac{\sigma_i}{1 - \sigma_i}}, \quad (9)$$

where σ_i is the i th ordered eigenvalue of \mathbf{Q} (Rodgers, 2000).

190 2.3 Approximating the Jacobian matrix

Section 2.2 described optimal reductions in dimension and rank of a state vector assuming knowledge of the native-resolution Jacobian matrix \mathbf{K} . However, the $n + 1$ forward model simulations needed to construct \mathbf{K} may be prohibitively expensive. Here we present a two-step approach to construct a reduced-dimension or reduced-rank Jacobian matrix at much lower computational cost. We start from a low-cost, native-resolution estimate $\mathbf{K}^{(0)}$ (see below) and calculate the corresponding averaging kernel matrix $\mathbf{A}^{(0)}$. In the reduced-dimension method, we use $\mathbf{A}^{(0)}$ to construct a multiscale grid that maintains resolution in the areas of highest information content (top right panel of Figure 1). We generate the updated, reduced-dimension Jacobian matrix $\mathbf{K}_{\text{RD}}^{(1)} \in \mathbb{R}^{m \times k}$ on the resulting grid using the forward model. In the reduced-rank method, we construct $\mathbf{K}_{\Pi}^{(1)} \in \mathbb{R}^{m \times n}$ on the basis of the k dominant eigenvectors of $\mathbf{A}^{(0)}$ by perturbing those patterns in the forward model, generating an approximation of the Jacobian matrix in a reduced-rank state space (bottom left panel of Figure 1). In both methods, the updated Jacobian matrix improves the estimate of the averaging kernel matrix and its eigenvectors by incorporating information content from the forward model. We use either $\mathbf{K}_{\text{RD}}^{(1)}$ or $\mathbf{K}_{\Pi}^{(1)}$ to conduct a second update and construct the final Jacobian matrix.

In our demonstration case, we generate an initial estimate of the native-resolution Jacobian matrix $\mathbf{K}^{(0)}$ at no cost by assuming that a local perturbation of methane emissions Δx [$\text{kg m}^{-2} \text{s}^{-1}$] produces local dry column mixing ratio

enhancements Δy [mol mol⁻¹] as determined by a simple column mass balance dependent on local wind speed and parameterized turbulent diffusion. We construct $\mathbf{K}^{(0)}$ column-wise by assuming that observation i responds to emissions in grid cell j as

$$210 \quad \Delta y_i = \alpha_{ij} \frac{M_{\text{air}} Lg}{M_{\text{CH}_4} Up} \Delta x_j \quad (10)$$

so that the elements $k_{ij}^{(0)}$ of $\mathbf{K}^{(0)}$ are given by

$$k_{ij}^{(0)} = \frac{\partial y_i}{\partial x_j} = \alpha_{ij} \frac{M_{\text{air}} Lg}{M_{\text{CH}_4} Up}, \quad (11)$$

215

where $\alpha_{ij} \in [0, 1]$ is a dimensionless coefficient providing a crude parameterization of turbulent diffusion, M_{air} and M_{CH_4} are the molecular weights of dry air and methane, respectively, L is a ventilation length scale taken as the square root of the grid cell area, g is gravitational acceleration, U is the local wind speed taken here as 5 km h⁻¹, and p is the surface pressure. We assume $\alpha_{ij} = 0.4$ for observations in grid cell j and distribute the remaining mass over the three concentric rings surrounding that cell with $\alpha_{ij} = 0.3/8, 0.2/16, \text{ and } 0.1/24$ from the inner to outer ring. Including a representation of turbulent diffusion increases the spatial coverage of the dominant patterns of information content; the exact form of the parameterization (e.g., the number of rings used or the values of α_{ij}) is unimportant.

220

The reduced-dimension and reduced-rank methods rely on characterizing the dominant patterns of information content of the observing system using the initial estimate of the averaging kernel matrix $\mathbf{A}^{(0)}$ corresponding to $\mathbf{K}^{(0)}$. $\mathbf{A}^{(0)}$ can provide a good approximation of \mathbf{A} even if the initial estimate of the Jacobian matrix $\mathbf{K}^{(0)}$ is crude because the averaging kernel matrix depends strongly on the specified prior and observational error covariance matrices \mathbf{S}_A and \mathbf{S}_O (Eq. (4)) and because, by assuming that observed concentrations are most sensitive to local emissions, $\mathbf{K}^{(0)}$ generates the highest information content where the observations are densest. This information content structure can then be refined by a two-step update.

225

230 **2.4 Constructing the reduced-dimension Jacobian matrix**

In an inverse system with a known native-resolution Jacobian matrix \mathbf{K} , a reduced-dimension Jacobian matrix \mathbf{K}_{RD} can be constructed on a multiscale grid that maintains native resolution where information content is highest and aggregates grid cells elsewhere (top right panel of Figure 1). We refer to the state vector elements of this multiscale grid as clusters. An optimal multiscale grid maximizes the total DOFS and the averaging kernel sensitivities of each state vector element, referred to here as the DOFS per cluster. To construct this grid, we first define the state vector as a single element

235

encompassing the inversion domain. We then add the native-resolution grid cells with the highest averaging kernel sensitivities to the state vector one-by-one, removing them from the original state vector element. For each new element x_i , we calculate the corresponding Jacobian matrix column $\partial \mathbf{y} / \partial x_i$ and the resulting increase in DOFS. When the DOFS stabilize, we add instead clusters of two or more native-resolution grid cells and repeat this procedure. Clusters can be generated by, for example, K-means clustering, which aggregates spatially proximate grid cells. We repeat this process, increasing cluster size, until all native-resolution grid cells are allocated to the multiscale grid and the corresponding reduced-dimension Jacobian matrix \mathbf{K}_{RD} is constructed. The DOFS convergence criteria and the sequence of cluster sizes can be selected to achieve the desired state vector dimension.

We apply this approach beginning with our initial estimate $\mathbf{K}^{(0)}$ (Section 2.3) in a two-step update that iteratively improves the multiscale grid. Algorithm 1 describes this process in detail. Briefly, the information content for the initial multiscale grid is given by $\mathbf{A}^{(0)}$, which we use to identify the grid cells with the highest information content and construct a multiscale grid as described above. We compute the corresponding reduced-dimension Jacobian matrix $\mathbf{K}_{\text{RD}}^{(1)}$, introducing information content from the forward model. We identify the state vector elements where the forward model contributes the most information content by comparing the sensitivities given by the updated reduced-dimension averaging kernel matrix $\mathbf{A}_{\text{RD}}^{(1)}$ to the sensitivities given by $\mathbf{A}^{(0)}$. We disaggregate the clusters with the largest differences and update the reduced-dimension Jacobian, generating $\mathbf{K}_{\text{RD}}^{(2)}$. Convergence is rapid and we find no need for further iteration. The analytical inversion can then be solved exactly on the multiscale grid using $\mathbf{K}_{\text{RD}}^{(2)}$.

Algorithm 1: Reduced-dimension Jacobian matrix construction

Given a native-resolution state vector with dimension n , a state vector encompassing the entire domain with dimension $n_{\text{RD}} = 1$, and $\mathbf{A}^{(0)}$ and $\mathbf{A}_{\text{RD}}^{(0)}$ the $n \times n$ and 1×1 initial estimates of the averaging kernel matrix, respectively, and an initial cluster size of one native-resolution grid cell:

- 1: Add the J clusters with the highest diagonal values of $\mathbf{A}^{(0)}$ to the state vector and update its dimension $n_{\text{RD}} = n_{\text{RD}} + J$;
- 2: Perturb in the forward model those J clusters and the background cluster to generate the $m \times n_{\text{RD}}$ reduced-dimension Jacobian matrix $\mathbf{K}_{\text{RD}}^{(1)}$ and the corresponding $n_{\text{RD}} \times n_{\text{RD}}$ averaging kernel matrix $\mathbf{A}_{\text{RD}}^{(1)}$;
- 3: If the difference in DOFS per cluster $\Delta \text{DPC} = \frac{\text{Tr}(\mathbf{A}_{\text{RD}}^{(1)})}{n_{\text{RD}}} - \frac{\text{Tr}(\mathbf{A}_{\text{RD}}^{(0)})}{n_{\text{RD}} - J} < \varepsilon$, where ε is a set threshold, increase the cluster size (e.g., by aggregating together non-allocated, native-resolution grid cells using K-means clustering);
- 4: Let $\mathbf{A}_{\text{RD}}^{(0)} = \mathbf{A}_{\text{RD}}^{(1)}$ and repeat steps 1 to 4 until all native-resolution grid cells are allocated to the state vector;
- 5: Disaggregate the clusters with the largest increase in the diagonal values from $\mathbf{A}^{(0)}$ to $\mathbf{A}_{\text{RD}}^{(1)}$ as measured on the multiscale grid and update the state vector dimension n_{RD} ;
- 6: Perturb in the forward model the disaggregated grid cells to generate the final $m \times n_{\text{RD}}$ reduced-dimension Jacobian matrix $\mathbf{K}_{\text{RD}}^{(2)}$.

2.5 Constructing the reduced-rank Jacobian matrix

In an inverse system with a known native-resolution Jacobian matrix \mathbf{K} , a reduced-rank approximation of the Jacobian matrix \mathbf{K}_Π can be constructed by calculating the linear relationship between emissions and observations for the most important patterns of information content rather than for individual or aggregate grid cells. A low-rank Jacobian corresponds to the state space shown in the bottom left panel of Figure 1. We showed in Section 2.2 that the leading patterns of information content are given by the columns of the dimension-restoring transformation $\mathbf{\Gamma}^*$ (Eq. (8)). For any selected value of k , the k leading patterns span a rank- k , dimension- n subspace of the original information content space. A Jacobian matrix can be constructed within this space by calculating the model response to perturbations of these patterns. The response of the forward model \mathbf{F} to the j th normalized eigenvector $\mathbf{y}_j^* \in \mathbb{R}^n$, given by the j th column of $\mathbf{\Gamma}^*$, is

$$\mathbf{y}_j = \frac{\mathbf{F}(\mathbf{x}_A + \beta \mathbf{y}_j^*) - \mathbf{F}(\mathbf{x}_A)}{\beta}, \quad (12)$$

where β is any scalar sufficiently large to ensure numerical stability. The model responses $\mathbf{y}_j, j \in \{1, \dots, k\}$ form the columns of the matrix $\mathbf{K}_\omega \in \mathbb{R}^{m \times k}$, which is the Jacobian matrix for an inverse system with a reduced-dimension state space spanned by the first k eigenvectors of the information content, illustrated by the bottom right panel of Figure 1. This reduced-dimension Jacobian must be transformed to the original state dimension to enable physical interpretation of the posterior results. Bousserrez and Henze (2018), following Bocquet et al. (2011), show that the reduced-dimension Jacobian matrix \mathbf{K}_ω is given by $\mathbf{K}_\omega = \mathbf{K}\mathbf{\Gamma}^*$ and the reduced-rank Jacobian matrix \mathbf{K}_Π by $\mathbf{K}_\Pi = \mathbf{K}\mathbf{\Pi} = \mathbf{K}\mathbf{\Gamma}^*\mathbf{\Gamma}$. Thus, the reduced-rank Jacobian can be calculated from the reduced-dimension Jacobian by $\mathbf{K}_\Pi = \mathbf{K}_\omega\mathbf{\Gamma}$. The resulting Jacobian has dimension $m \times n$ and rank k .

In an inverse system without a known Jacobian matrix, the reduced-rank Jacobian matrix approximation can be constructed in a two-step update that iteratively improves the patterns of information content used as perturbations. Algorithm 2 describes this process in detail. Briefly, we use the initial estimate of the Jacobian matrix $\mathbf{K}^{(0)}$ (Section 2.3) to calculate the corresponding averaging kernel matrix $\mathbf{A}^{(0)}$ and the matrix of its eigenvectors $\mathbf{\Gamma}^{*(0)}$. When calculating $\mathbf{\Gamma}^{*(0)}$, we select the $k^{(0)}$ eigenvectors that have a signal-to-noise ratio greater than some threshold. We use the signal-to-noise criterion, which is stricter than the information content criterion, to account for the errors in the initial estimate of the information content. We compute the forward model response to each of the eigenvectors using Eq. (12) and transform the resulting reduced-dimension Jacobian $\mathbf{K}_\omega^{(1)}$ to the full-dimension state space with $\mathbf{K}_\Pi^{(1)} = \mathbf{K}_\omega^{(1)}\mathbf{\Gamma}^{(0)}$. We calculate the associated averaging kernel matrix $\mathbf{A}_\Pi^{(1)}$ and the matrix of its eigenvectors $\mathbf{\Gamma}_\Pi^{*(1)}$. Because $\mathbf{K}_\Pi^{(1)}$ is a reduced-rank approximation, its spectrum of information content is discontinuous at $k^{(0)}$. We therefore use the spectrum of information content associated with the initial, full-rank estimate $\mathbf{A}^{(0)}$ to select the rank $k^{(1)}$ of the second update and calculate $\mathbf{\Gamma}_\Pi^{*(1)}$. We use the $k^{(1)}$ eigenvectors that

span most of the information content from the initial estimate and construct an updated reduced-rank Jacobian matrix approximation $\mathbf{K}_{\Pi}^{(2)}$ as above. The resulting Jacobian matrix $\mathbf{K}_{\Pi}^{(2)}$ is a rank $\approx k^{(1)}$ approximation that accurately quantifies the forward model where the observing system has high information content and loses accuracy in areas with lower information content where the observations are least able to constrain emissions. The resulting posterior solution is accurate in areas with high information content and defaults to the prior estimate elsewhere.

305 **Algorithm 2: Reduced-rank Jacobian matrix construction**

Given a native-resolution state vector with dimension n , $\mathbf{A}^{(0)}$ an $n \times n$ initial estimate of the averaging kernel matrix, and \mathbf{S}_A the $n \times n$ prior error covariance matrix:

1: Complete the eigendecomposition of the $n \times n$ matrix $\mathbf{Q}^{(0)} = \mathbf{S}_A^{-1/2} \mathbf{A}^{(0)} \mathbf{S}_A^{1/2} = \mathbf{W}^{(0)} \mathbf{\Sigma}^{(0)} \mathbf{W}^{(0)T}$;

2: Select k so that $1 - \frac{\text{Tr}(\mathbf{\Sigma}_k^{(0)})}{\text{Tr}(\mathbf{\Sigma}^{(0)})} < \epsilon$ where $\mathbf{\Sigma}_k^{(0)}$ is the $k \times k$ subset of $\mathbf{\Sigma}^{(0)}$ containing the largest diagonal values and ϵ is a set threshold; or,

310 select k so that the k th eigenvector has a signal-to-noise ratio $\text{SNR}_k = \sqrt{\frac{\sigma_k^{(0)}}{1 - \sigma_k^{(0)}}} > \sim 1$, where $\sigma_k^{(0)}$ is the k th largest diagonal value of $\mathbf{\Sigma}^{(0)}$;

3: Form the $n \times k$ matrix $\mathbf{\Gamma}^{*(0)} = \mathbf{S}_A^{1/2} \mathbf{W}_k^{(0)}$ where $\mathbf{W}_k^{(0)}$ is a matrix of the first k columns of $\mathbf{W}^{(0)}$ as ranked by the diagonal values of $\mathbf{\Sigma}^{(0)}$;

4: Perturb in the forward model the columns of $\mathbf{\Gamma}^{*(0)}$ to generate the $m \times k$ reduced-dimension Jacobian matrix $\mathbf{K}_{\omega}^{(1)}$;

5: Form the $m \times n$ reduced-rank Jacobian matrix $\mathbf{K}_{\Pi}^{(1)} = \mathbf{K}_{\omega}^{(1)} \mathbf{\Gamma}^{(0)}$, where $\mathbf{\Gamma}^{(0)} = \mathbf{W}_k^{(0)T} \mathbf{S}_A^{-1/2}$, and the corresponding $n \times n$ reduced-rank

315 averaging kernel matrix $\mathbf{A}_{\Pi}^{(1)}$;

6: Let $\mathbf{A}^{(0)} = \mathbf{A}_{\Pi}^{(1)}$ and repeat steps 1 to 5 to generate the final $m \times n$ reduced-rank Jacobian matrix $\mathbf{K}_{\Pi}^{(2)}$.

3 Results and discussion

We demonstrate the reduced-dimension and reduced-rank Jacobian matrix construction methods in an analytical Bayesian inversion of atmospheric methane columns observed by the GOSAT satellite over North America in July 2009. Although TROPOMI now provides higher density observations, using GOSAT allows us to use the inversion framework of Maasackers (2019). We construct a ‘‘native-resolution’’ inverse system at $1^\circ \times 1.25^\circ$ grid cell resolution ($n = 2098$, top left panel of Figure 1) against which we compare our reduced-dimension and reduced-rank methods. To demonstrate the applicability of the methods to higher-information observing systems such as TROPOMI, we artificially increase the information content of the GOSAT data by introducing an amplification factor $\lambda > 1$ to the cost function that increases the weight of the observational terms:

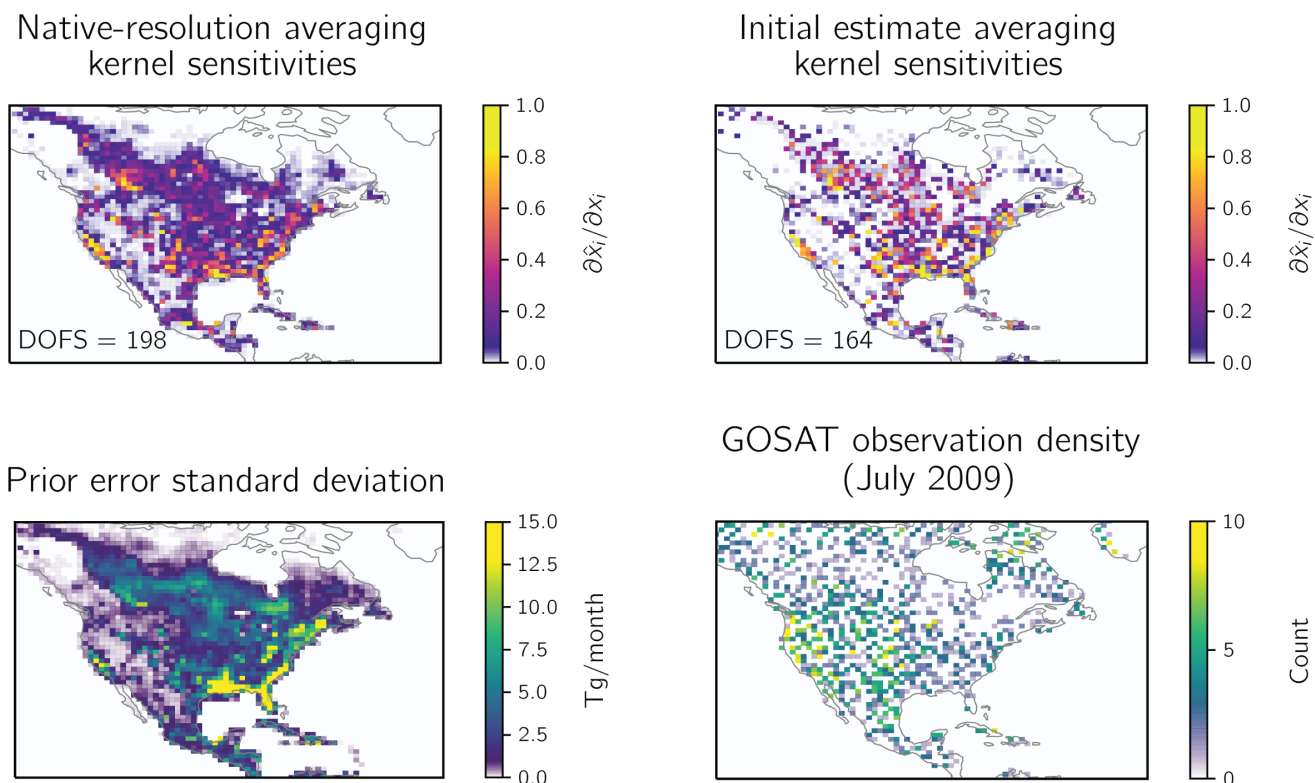
$$\mathcal{J}(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_A)^T \mathbf{S}_A^{-1} (\mathbf{x} - \mathbf{x}_A) + \lambda (\mathbf{y} - \mathbf{F}(\mathbf{x}))^T \mathbf{S}_0^{-1} (\mathbf{y} - \mathbf{F}(\mathbf{x})). \quad (13)$$

The amplification factor functionally decreases the observational error covariance, increasing the DOFS. We set $\lambda = 5$,
330 which increases the native-resolution DOFS from 82 to 198. Because of noise in the GOSAT data, this results in an overfit
solution that is irrelevant in our demonstration.

We use the nested North American GEOS-Chem CTM version 12.4.0 as forward model to simulate atmospheric methane
column concentrations at $1^\circ \times 1.25^\circ$ resolution for July 2009. The 2098 $1^\circ \times 1.25^\circ$ grid cells constitute our native-resolution
335 state vector. The model is driven with MERRA-2 meteorological fields (Bosilovich et al., 2016) from the NASA Global
Modeling and Assimilation Office. We use boundary conditions and initial conditions from a global posterior GEOS-Chem
 $4^\circ \times 5^\circ$ simulation from Maasakkers et al. (2019). The GOSAT observations are from the University of Leicester version 7
CH₄ proxy retrieval over land (Parker et al., 2011; Parker et al., 2015; ESA CCI GHG project team, 2018) for July 2009.
Prior emissions and observational error covariances are as described in Maasakkers et al. (2019). We assume uniform
340 relative prior errors of 50%. The demonstration is sufficiently coarse-resolution and short that the native-resolution Jacobian
matrix \mathbf{K} can be explicitly computed with 2099 model runs. After constructing \mathbf{K} , we use it as the forward model in lieu of
additional GEOS-Chem simulations.

Figure 2 (top left panel) shows the native-resolution averaging kernel sensitivities, i.e., the diagonal elements of the native-
345 resolution averaging kernel matrix \mathbf{A} . As discussed in Section 2.3, the structure of the averaging kernel matrix is largely
determined by the prior error covariance matrix \mathbf{S}_A and by the observation density as reflected in both the observational
covariance matrix \mathbf{S}_O and the Jacobian matrix \mathbf{K} . This is apparent in the bottom panels of Figure 2, which show the
distribution of the prior error standard deviations (left) and observation density (right). The absolute errors on the prior
emissions estimate are largest for wetlands along the southeastern coast of the U.S. (Bloom et al., 2017). The variability in
350 the observation density is driven by sampling frequency and retrieval success (Parker et al., 2015).

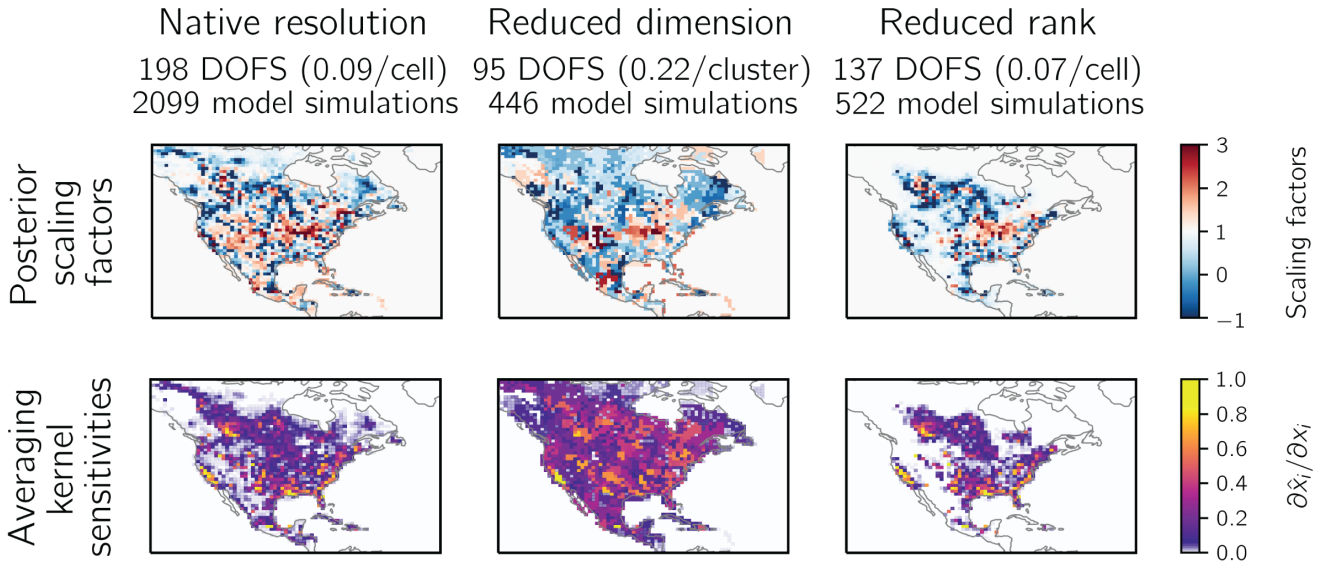
Figure 2 (top right panel) also shows the initial estimate of averaging kernel sensitivities of $\mathbf{A}^{(0)}$ derived from the initial
estimate of the Jacobian matrix $\mathbf{K}^{(0)}$ constructed as described in Section 2.3. No forward model simulations were conducted
to construct this initial estimate, yet the patterns of information content reproduce those given by the native-resolution
355 averaging kernel matrix \mathbf{A} because of the strong dependence on the prior error standard deviations and the observation
density. \mathbf{A} has a smoother structure than $\mathbf{A}^{(0)}$ because of the effect of long-range transport in the CTM, but this has little
effect on the leading patterns of information content.



360 **Figure 2. Averaging kernel sensitivities for the demonstration inversion of GOSAT observations with enhanced information**
content for July 2009. The top panels show the sensitivities given by the diagonal elements of the averaging kernel matrix \mathbf{A} of the
native-resolution inversion (left) and of the initial estimate of the averaging kernel matrix $\mathbf{A}^{(0)}$ (right). The DOFS for each
averaging kernel matrix are inset in the corresponding panel. The lower left panel shows the error standard deviations on the
prior emissions estimate given by the square roots of the diagonal elements of \mathbf{S}_Λ . The lower right panel shows the GOSAT
365 **observation density.**

For this demonstration, we aim to reduce the number of forward model runs needed to construct the Jacobian matrix by at least a factor of four relative to the native-resolution inversion, from 2099 to ≈ 525 simulations. We first apply the reduced-dimension method to construct a Jacobian matrix on a multiscale grid generated with K-means clustering following Section 2.4. The resulting initial multiscale grid and reduced-dimension Jacobian matrix $\mathbf{K}_{\text{RD}}^{(1)}$ constrain 380 clusters and required 381
370 model simulations. We disaggregate 11 clusters with a sensitivity increase greater than 0.6, adding 65 native-resolution grid cells and model simulations. The resulting multiscale grid is shown in the top right panel of Figure 1. It has dimension 434 and the corresponding reduced-dimension Jacobian matrix $\mathbf{K}_{\text{RD}}^{(2)}$ required 446 forward model simulations across 17 parallelized batches, representing a factor of five decrease in computational cost relative to the native-resolution solution. The grid has 137 native-resolution grid cells and clusters of between 2 and 55 grid cells.

375



380 **Figure 3. Results from the demonstration inversion at native-resolution compared to the reduced-dimension and reduced-rank methods. The figure shows the posterior scaling factors and averaging kernel sensitivities with respect to the prior emissions estimate for each inversion. The degrees of freedom for signal (DOFS) give the number of pieces of information that the inversion can independently constrain.**

Figure 3 shows the posterior emission scaling factors relative to the prior estimate (top) and averaging kernel sensitivities (bottom) for the reduced-dimension solution (center column) compared to the native-resolution solution (left column). Both solutions are exact on the grids used. The reduced-dimension solution generates fewer DOFS (95) than the native-resolution solution (198) because the DOFS depend on the dimension of the state vector. When comparing the DOFS per cluster, a dimension-independent measure, the reduced-dimension solution produces more than twice the value of the native-resolution solution (0.22 compared to 0.09), reflecting the consolidation of information content. This is reflected in the reduced-dimension averaging kernel sensitivities, which are more uniform than the native-resolution values. The reduced-dimension posterior scaling factors exhibit less variability than the native-resolution solution, which exhibits checkerboard patterns resulting in part from the overfit of the posterior solution to observational noise. The posterior scaling factors agree on regional scales.

390

We next apply the reduced-rank method (Section 2.5) to construct a reduced-rank approximation of the Jacobian matrix. We calculate the dominant eigenvectors of the initial averaging kernel matrix estimate $\mathbf{A}^{(0)}$, requiring that the signal-to-noise ratio of all eigenvectors be greater than 1.25. This yields $k^{(0)} = 90$ eigenvectors, which account for 43% of the initial-estimate DOFS. We perturb these eigenvectors in the forward model and construct the reduced-rank Jacobian matrix $\mathbf{K}_{\Pi}^{(1)}$. We then calculate the averaging kernel matrix $\mathbf{A}_{\Pi}^{(1)}$ and its dominant eigenvectors, defining $k^{(1)} = 431$ by requiring that the improved eigenvectors capture 97% of the information content defined by $\mathbf{A}_{\Pi}^{(0)}$. The resulting Jacobian matrix $\mathbf{K}_{\Pi}^{(2)}$ has rank

395

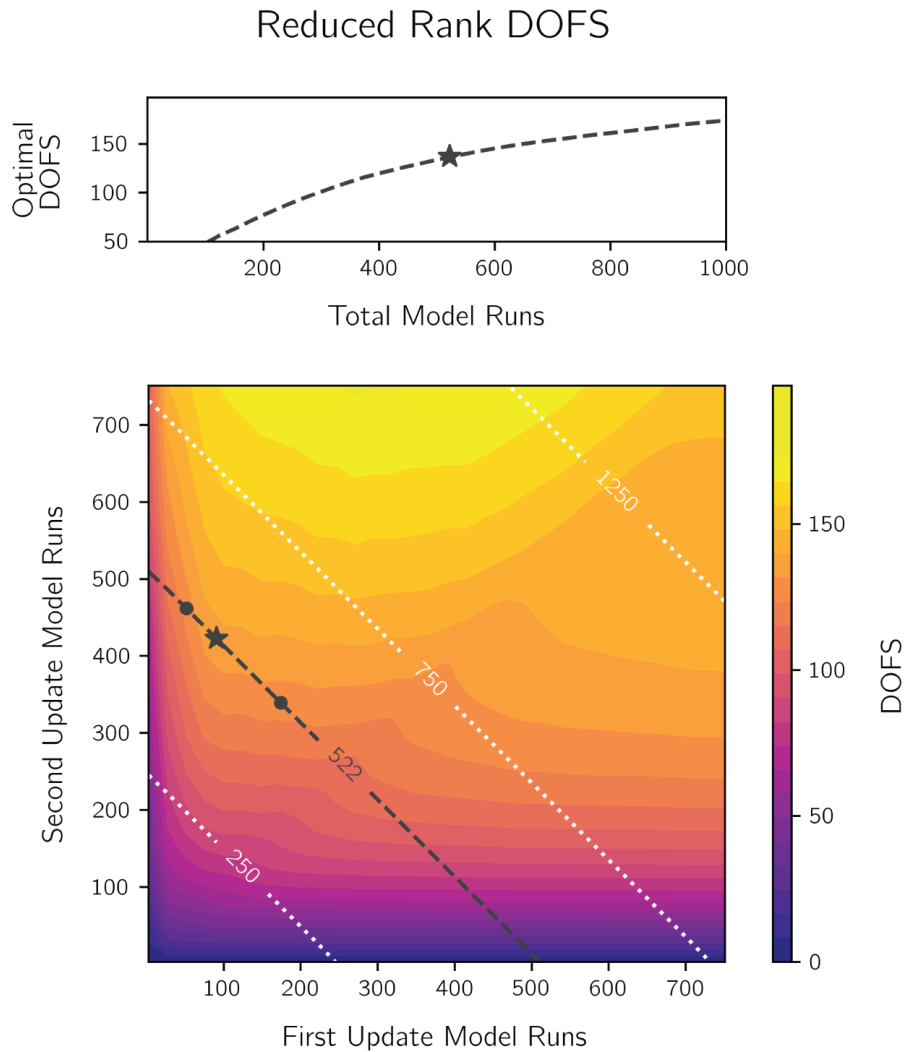
400 ≈ 431 and required 522 forward model simulations across two parallelized batches. We solve the inversion with $\mathbf{K}_{\Pi}^{(2)}$ and find 137 DOFS compared to the 198 DOFS generated in the native-resolution inversion, achieving 69% of the DOFS at a quarter of the computational cost.

The DOFS of the reduced-rank inversion are only moderately sensitive to the first- and second-update thresholds, with a stronger dependence on the number of model runs conducted in the second update. Figure 4 shows the reduced-rank DOFS as a function of the number of first- and second-update forward model runs. Among all possible partitions of 522 total model runs (dashed line), our update scheme (starred) nearly maximizes the DOFS, but the DOFS has only moderate sensitivity to the choice of partition. Using a signal-to-noise ratio threshold of 0.75 or 1.75 instead of 1.25 (dots), decreases the reduced-rank DOFS by 7%. Lowering the signal-to-noise ratio threshold increases the number of eigenvectors drawn from $\mathbf{A}^{(0)}$, which increases the effect of errors in the initial Jacobian matrix estimate $\mathbf{K}^{(0)}$. Increasing the threshold fails to exploit the information content of $\mathbf{A}^{(0)}$. More generally, applying a signal-to-noise ratio threshold of 1.25 in the first update maximizes the DOFS regardless of the number of model runs conducted in the second update. We show the DOFS generated by these optimal configurations as a function of the total number of forward model runs in the top panel of Figure 4. After only 300 simulations, the optimal reduced-rank inversion generates 101 DOFS, achieving half of the native-resolution DOFS at 14% of the computational cost.

415 We solve the inversion (Eqs. (2) – (4)) using the reduced-rank Jacobian matrix $\mathbf{K}_{\Pi}^{(2)}$ and compare the posterior to the native-resolution solution. Figure 3 (right column) shows the distribution of the reduced-rank posterior scaling factors (top) and averaging kernel sensitivities (bottom) compared to the native-resolution inversion (left column). Because $\mathbf{K}_{\Pi}^{(2)}$ was constructed on the basis of the dominant patterns of information content, it solves for the posterior scaling factors accurately in the areas of highest information content and defaults to the prior value (a scaling factor of one) elsewhere. As a result of the exclusion of grid cells with low native-resolution information content, the reduced-rank DOFS (137) are lower than native-resolution DOFS (198). However, in grid cells with large averaging kernel sensitivities, the reduced-rank inversion preserves most information content. 755 grid cells have reduced-rank averaging kernel sensitivities greater than 0.01 and generate 136 DOFS, amounting to 83% of the 163 DOFS generated by the same grid cells in the native-resolution inversion.

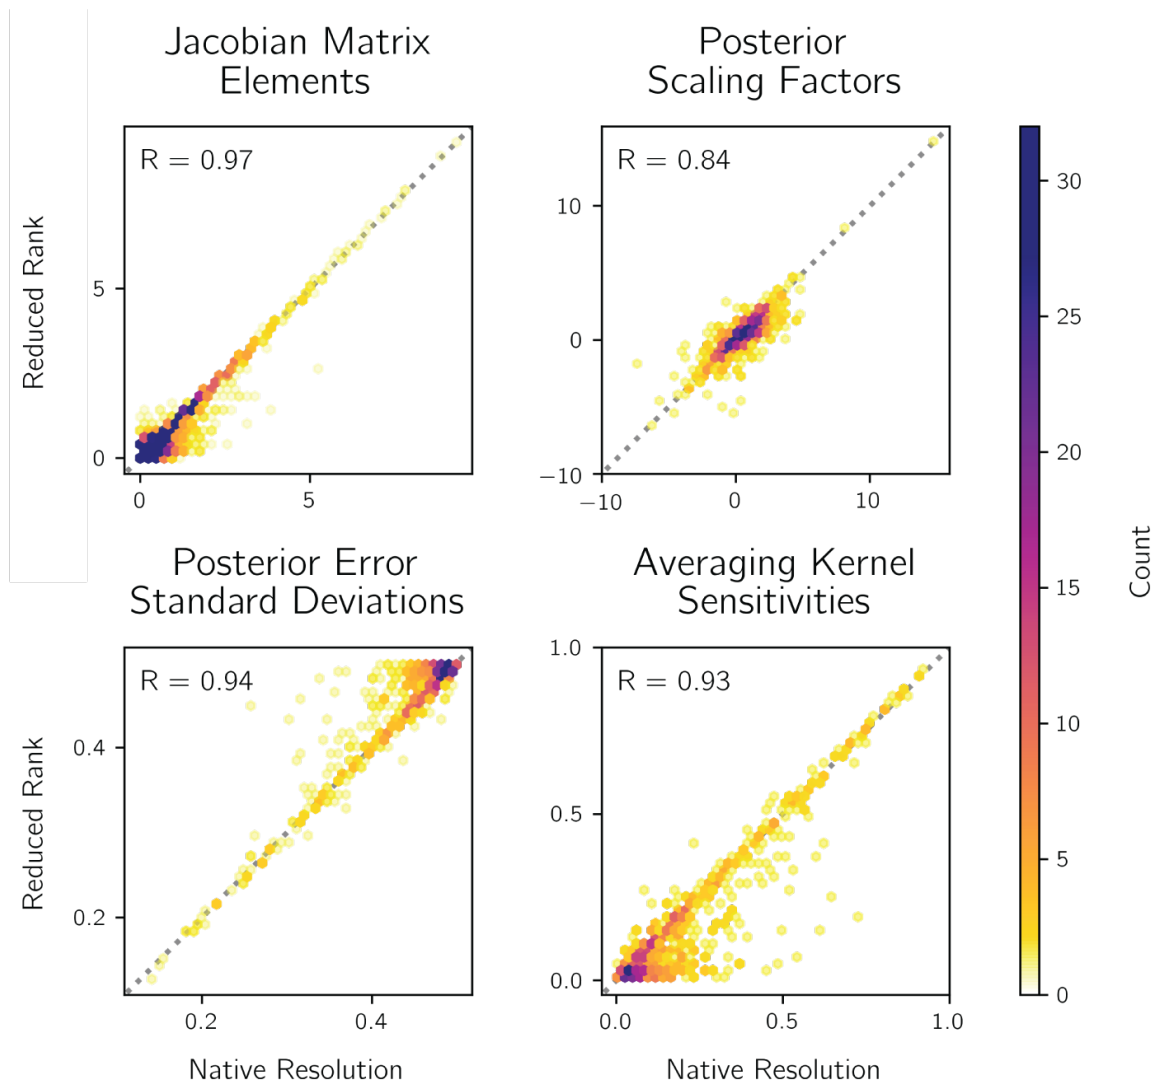
425 Figure 5 shows a statistical comparison of the reduced-rank and native-resolution inversion results for grid cells with a reduced-rank averaging kernel sensitivity above 0.01. None of the reduced-rank quantities exhibit significant bias, as shown by comparison to the 1:1 line. The elements of the reduced-rank Jacobian matrix $\mathbf{K}_{\Pi}^{(2)}$ correspond closely with those of the native-resolution Jacobian matrix \mathbf{K} (correlation coefficient $R = 0.97$). The strong correlation of the averaging kernel sensitivities ($R = 0.93$) confirms that the reduced-rank inversion accurately identifies the native-resolution grid cells with the

430



435 **Figure 4.** The sensitivity of the reduced-rank inversion DOFS to the number of forward model runs. The bottom panel shows the sensitivity of the DOFS to the partitioning of model runs between the first (x-axis) and second (y-axis) update. The lines represent the total number of simulations. Our inversion uses a signal-to-noise ratio of 1.25 for the first update and an information content threshold of 97% for the second update (star), requiring 522 forward model runs and generating 137 DOFS, accounting for 69% of the native-resolution DOFS at a quarter of the computational expense. Using a signal-to-noise ratio of 0.75 or 1.75 with the same total number of model simulations (dots) does not substantially decrease the DOFS. The top panel shows the DOFS as a function of the total number of model runs for all optimal first- and second-update partitions.

440 highest information content. The posterior errors and scaling factors agree well in these grid cells. The posterior error standard deviations correlate strongly ($R = 0.94$) due in part to the common contribution of the prior and observational error covariance matrices (Eq. (3)). The outlier reduced-rank standard deviations tend to be larger than the native-resolution values, reflecting the error introduced by discarding information content. The posterior scaling factors also agree well but the



445 **Figure 5. Comparison statistics between the reduced-rank and native-resolution inversions for the grid cells with reduced-rank averaging kernel sensitivity greater than 0.01. Individual panels compare binned counts for Jacobian matrix elements [ppb], posterior scaling factors [dimensionless], posterior error standard deviations [dimensionless], and averaging kernel sensitivities [dimensionless]. Correlation coefficients are inset and the 1:1 line (dotted) is shown.**

correlation coefficient is smaller ($R = 0.84$) because of the smaller dynamical range and the propagation of errors from the
 450 posterior error covariance and Jacobian matrices (Eq. (2)). Negative scaling factors reflect the overfit from artificially increasing the information content but are of no consequence for our demonstration.

The reduced-dimension and reduced-rank methods reproduce the native-resolution inversion with a factor of at least four reduction in total computational cost. The reduced-dimension method generates lower DOFS but higher DOFS per state
 455 vector element due to the clustering of grid cells. The resulting posterior solution is exact on the multiscale grid and provides

460 better spatial coverage than the reduced-rank method at lower resolution. The reduced-rank method generates a higher-DOFS, higher-resolution approximation where the averaging kernel sensitivities are large. While the calculation of large Jacobian matrices can take advantage of parallel computing environments (Maasackers et al., 2019), the iterative nature of both methods proposed here puts some limit on parallelization. The limit is greater for the reduced-dimension method, which requires an iteration for each cluster size added to the state vector. The reduced-rank method requires only two iterations. In both cases, these limitations may not be meaningful because the native-resolution Jacobian matrix is rarely generated in a fully parallel environment in practice.

4 Conclusions

465 We proposed two methods to conduct analytical high-resolution inversions of satellite observations of atmospheric composition to infer emissions while maximizing information content and minimizing computational cost. The computational cost of analytical inversions is driven by the construction of the Jacobian matrix, which expresses the sensitivity of the observed concentrations to emissions. The Jacobian matrix is constructed numerically by conducting perturbation simulations with a chemical transport model (CTM) that serves as forward model for the inversion. Our methods exploit the dominant patterns of information content in the observing system to build the Jacobian matrix. The 470 reduced-dimension method constructs the Jacobian matrix on a multiscale grid that aggregates grid cells where information content is lowest. The reduced-rank method approximates the Jacobian matrix using the dominant patterns of information content, discarding the weaker patterns. Beyond the atmospheric application presented here, both methods can be applied more generally to the problem of efficient numerical approximation of high-dimension Jacobian matrices.

475 Both methods use a two-step update to improve an initial, no-cost estimate of the Jacobian matrix and the corresponding averaging kernel matrix. The initial estimate of the Jacobian matrix is constructed here by assuming that observed atmospheric concentrations are most sensitive to local emissions. Because the averaging kernel matrix has a strong dependence on the prior error covariance matrix and observation density, this initial estimate can accurately quantify the fine structure of information content. The reduced-dimension method uses the initial estimate of the averaging kernel matrix to 480 generate a multiscale grid that maintains native resolution where information content is highest and consolidates grid cells elsewhere. The forward model is used to build the Jacobian matrix on this grid, and the resulting reduced-dimension averaging kernel matrix is compared to the initial estimate to identify the state vector elements where the forward model contributed the most information content. These elements are disaggregated to generate a second and final update. The reduced-rank method uses the initial estimate of the averaging kernel matrix to identify the dominant patterns of information 485 content. These patterns are perturbed in the forward model, generating a first update of the Jacobian matrix. This update serves as the basis for a second and final update. In both methods, rapid convergence occurs after two updates.

We applied both methods in a demonstration inversion of GOSAT column methane observations over North America for July 2009 with artificially enhanced information content. We compared the results to a native-resolution inversion
490 optimizing emissions on a $1^\circ \times 1.25^\circ$ grid. Both methods successfully approximated the native-resolution results and decreased total computational cost by a factor of at least four. The reduced-dimension method produced only 50% of the native-resolution information content as measured by the degrees of freedom for signal (DOFS) due to spatial averaging but generated twice the DOFS per state vector element and avoided the correlated errors found in the native-resolution inversion. The reduced-rank method retained 70% of the native-resolution DOFS by solving the inversion accurately in the grid cells
495 with the highest information content and defaulting to the prior emissions estimate elsewhere. In sensitivity tests, the reduced-rank method retained 50% of the native-resolution DOFS while decreasing computational cost by a factor of seven.

Code and data availability

All code and data are available at https://github.com/hannahnesser/reduced_rank_jacobian.

Author contribution

500 HN and DJJ designed the study. HN performed the analysis. JDM provided all demonstration inversion data and supporting guidance. HN and MPS performed the simulations. HN, DJJ, JDM, TRS, MPS, YZ, and CHR discussed the methods and results. HN and DJJ prepared the manuscript with contributions from all co-authors.

Competing interests

The authors declare that they have no conflict of interest.

505 Acknowledgements

This work was funded by the NASA Carbon Monitoring System and by an NSF Graduate Fellowship to H. Nesser. Y. Zhang is supported by NSFC Project 42007198 and Westlake University. We thank Daven Henze, Kevin Bowman, Michael Brenner, Cynthia Randles, Jeremy Brandman, and Laurent White for helpful discussions and Marc Bocquet for an insightful review.

510 References

Alexe, M., Bergamaschi, P., Segers, A., Detmers, R., Butz, A., Hasekamp, O., Guerlet, S., Parker, R., Boesch, H., Frankenberg, C., Scheepmaker, R. A., Dlugokencky, E., Sweeney, C., Wofsy, S. C. and Kort, E. A.: Inverse modelling of

- CH4 emissions for 2010-2011 using different satellite retrieval products from GOSAT and SCIAMACHY, *Atmos. Chem. Phys.*, 15(1), 113–133, doi:10.5194/acp-15-113-2015, 2015.
- 515 Bergamaschi, P., Frankenberg, C., Meirink, J. F., Krol, M., Villani, M. G., Houweling, S., Dentener, F., Dlugokencky, E. J., Miller, J. B., Gatti, L. V., Engel, A. and Levin, I.: Inverse modeling of global and regional CH4 emissions using SCIAMACHY satellite retrievals, *J. Geophys. Res. Atmos.*, 114(D22), doi:10.1029/2009JD012287, 2009.
- Bergamaschi, P., Houweling, S., Segers, A., Krol, M., Frankenberg, C., Scheepmaker, R. A., Dlugokencky, E., Wofsy, S. C., Kort, E. A., Sweeney, C., Schuck, T., Brenninkmeijer, C., Chen, H., Beck, V. and Gerbig, C.: Atmospheric CH4 in the first
520 decade of the 21st century: Inverse modeling analysis using SCIAMACHY satellite retrievals and NOAA surface measurements, *J. Geophys. Res. Atmos.*, 118(13), doi:10.1002/jgrd.50480, 2013.
- Bloom, A. A., Bowman, K. W., Lee, M., Turner, A. J., Schroeder, R., Worden, J. R., Weidner, R., McDonald, K. C. and Jacob, D. J.: A global wetland methane emissions and uncertainty dataset for atmospheric chemical transport models (WetCHARTs version 1.0), *Geosci. Model Dev.*, 10(6), 2141–2156, doi:10.5194/gmd-10-2141-2017, 2017.
- 525 Bocquet, M. and Wu, L.: Bayesian design of control space for optimal assimilation of observations. Part II: Asymptotic solutions, *Q. J. R. Meteorol. Soc.*, 137(658), 1357–1368, doi:10.1002/qj.841, 2011.
- Bocquet, M., Wu, L. and Chevallier, F.: Bayesian design of control space for optimal assimilation of observations. Part I: Consistent multiscale formalism, *Q. J. R. Meteorol. Soc.*, 137(658), 1340–1356, doi:10.1002/qj.837, 2011.
- Bousserez, N. and Henze, D. K.: Optimal and scalable methods to approximate the solutions of large-scale Bayesian
530 problems: theory and application to atmospheric inversion and data assimilation, *Q. J. R. Meteorol. Soc.*, 144(711), 365–390, doi:10.1002/qj.3209, 2018.
- Brasseur, G. P. and Jacob, D. J.: Inverse Modeling for Atmospheric Chemistry, in *Modeling of Atmospheric Chemistry*, pp. 487–537, Cambridge University Press, Cambridge., 2017.
- Chevallier, F., Bréon, F. M. and Rayner, P. J.: Contribution of the Orbiting Carbon Observatory to the estimation of CO2
535 sources and sinks: Theoretical study in a variational data assimilation framework, *J. Geophys. Res. Atmos.*, doi:10.1029/2006JD007375, 2007.
- Daescu, D., Carmichael, G. R. and Sandu, A.: Adjoint Implementation of Rosenbrock Methods Applied to Variational Data Assimilation Problems, *J. Comput. Phys.*, 165(2), doi:10.1006/jcph.2000.6622, 2000.
- Elbern, H. and Schmidt, H.: Ozone episode analysis by four-dimensional variational chemistry data assimilation, *J. Geophys. Res. Atmos.*, 106(D4), doi:10.1029/2000JD900448, 2001.
540
- ESA CCI GHG project team: ESA Greenhouse Gases Climate Change Initiative (GHG_cci): Column-averaged CH4 from GOSAT generated with the OCPR (UoL-PR) Proxy algorithm (CH4_GOS_OCPR), v7.0, 2018.
- Hasekamp, O., Lorente, A., Hu, H., Butz, A., Aan de Brugh, J. and Landgraf, J.: Algorithm Theoretical Baseline Document for Sentinel-5 Precursor Methane Retrieval, , 1(10), 1–67, 2019.
- 545 Henze, D. K., Hakami, A. and Seinfeld, J. H.: Development of the adjoint of GEOS-Chem, *Atmos. Chem. Phys.*, 7, 2413–2433, doi:10.5194/acp-7-2413-2007, 2007.

- Houweling, S., Krol, M., Bergamaschi, P., Frankenberg, C., Dlugokencky, E. J., Morino, I., Notholt, J., Sherlock, V., Wunch, D., Beck, V., Gerbig, C., Chen, H., Kort, E. A., Röckmann, T. and Aben, I.: A multi-year methane inversion using SCIAMACHY, accounting for systematic errors using TCCON measurements, *Atmos. Chem. Phys.*, 14(8), 3991–4012, doi:10.5194/acp-14-3991-2014, 2014.
- 550
- Hu, H., Landgraf, J., Detmers, R., Borsdorff, T., Aan de Brugh, J., Aben, I., Butz, A. and Hasekamp, O.: Toward Global Mapping of Methane With TROPOMI: First Results and Intersatellite Comparison to GOSAT, *Geophys. Res. Lett.*, 45(8), 3682–3689, doi:10.1002/2018gl077259, 2018.
- Jacob, D. J., Turner, A. J., Maasackers, J. D., Sheng, J., Sun, K., Liu, X., Chance, K., Aben, I., McKeever, J. and Frankenberg, C.: Satellite observations of atmospheric methane and their value for quantifying methane emissions, *Atmos. Chem. Phys.*, 16(22), 14371–14396, doi:10.5194/acp-16-14371-2016, 2016.
- 555
- Koohkan, M. R., Bocquet, M., Roustan, Y., Kim, Y. and Seigneur, C.: Estimation of volatile organic compound emissions for Europe using data assimilation, *Atmos. Chem. Phys.*, doi:10.5194/acp-13-5887-2013, 2013.
- Liu, Y., Haussaire, J. M., Bocquet, M., Roustan, Y., Saunier, O. and Mathieu, A.: Uncertainty quantification of pollutant source retrieval: comparison of Bayesian methods with application to the Chernobyl and Fukushima Daiichi accidental releases of radionuclides, *Q. J. R. Meteorol. Soc.*, 143(708), doi:10.1002/qj.3138, 2017.
- 560
- Lunt, M. F., Rigby, M., Ganesan, A. L. and Manning, A. J.: Estimation of trace gas fluxes with objectively determined basis functions using reversible-jump Markov chain Monte Carlo, *Geosci. Model Dev.*, 9(9), doi:10.5194/gmd-9-3213-2016, 2016.
- 565
- Maasackers, J. D., Jacob, D. J., Sulprizio, M. P., Scarpelli, T. R., Nesser, H., Sheng, J. X., Zhang, Y., Hersher, M., Anthony Bloom, A., Bowman, K. W., Worden, J. R., Janssens-Maenhout, G. and Parker, R. J.: Global distribution of methane emissions, emission trends, and OH concentrations and trends inferred from an inversion of GOSAT satellite data for 2010–2015, *Atmos. Chem. Phys.*, 19(11), 7859–7881, doi:10.5194/acp-19-7859-2019, 2019.
- Meirink, J. F., Bergamaschi, P. and Krol, M. C.: Four-dimensional variational data assimilation for inverse modelling of atmospheric methane emissions: Method and comparison with synthesis inversion, *Atmos. Chem. Phys.*, doi:10.5194/acp-8-6341-2008, 2008.
- 570
- Monteil, G., Houweling, S., Butz, A., Guerlet, S., Schepers, D., Hasekamp, O., Frankenberg, C., Scheepmaker, R., Aben, I. and Röckmann, T.: Comparison of CH₄ inversions based on 15 months of GOSAT and SCIAMACHY observations, *J. Geophys. Res. Atmos.*, 118(29), doi:10.1002/2013JD019760, 2013.
- 575
- Parker, R., Boesch, H., Cogan, A., Fraser, A., Feng, L., Palmer, P. I., Messerschmidt, J., Deutscher, N., Griffith, D. W. T., Notholt, J., Wennberg, P. O. and Wunch, D.: Methane observations from the Greenhouse Gases Observing SATellite: Comparison to ground-based TCCON data and model calculations, *Geophys. Res. Lett.*, 38(15), doi:10.1029/2011GL047871, 2011.
- Parker, R. J., Boesch, H., Byckling, K., Webb, A. J., Palmer, P. I., Feng, L., Bergamaschi, P., Chevallier, F., Notholt, J., Deutscher, N., Warneke, T., Hase, F., Sussmann, R., Kawakami, S., Kivi, R., Griffith, D. W. T. and Velasco, V.: Assessing 5
- 580

- years of GOSAT Proxy XCH₄ data and associated uncertainties, *Atmos. Meas. Tech.*, 8(11), 4785–4801, doi:10.5194/amt-8-4785-2015, 2015.
- Quélo, D., Mallet, V. and Sportisse, B.: Inverse modeling of NO_x emissions at regional scale over northern France: Preliminary investigation of the second-order sensitivity, *J. Geophys. Res. Atmos.*, 110(24), doi:10.1029/2005JD006151, 585 2005.
- Ray, J., Lee, J., Yadav, V., Lefantzi, S., Michalak, A. M. and Van Bloemen Waanders, B.: A sparse reconstruction method for the estimation of multi-resolution emission fields via atmospheric inversion, *Geosci. Model Dev.*, 8(4), doi:10.5194/gmd-8-1259-2015, 2015.
- Rigby, M., Manning, A. J. and Prinn, R. G.: Inversion of long-lived trace gas emissions using combined Eulerian and 590 Lagrangian chemical transport models, *Atmos. Chem. Phys.*, 11(18), doi:10.5194/acp-11-9887-2011, 2011.
- Rodgers, C. D.: *Inverse Methods for Atmospheric Sounding: Theory and Practice.*, 2000.
- Saunois, M., Stavert, A. R., Poulter, B., Bousquet, P., Canadell, J. G., Jackson, R. B., Raymond, P. A., Dlugokencky, E. J., Houweling, S., Patra, P. K., Ciais, P., Arora, V. K., Bastviken, D., Bergamaschi, P., Blake, D. R., Brailsford, G., Bruhwiler, L., Carlson, K. M., Carrol, M., Castaldi, S., Chandra, N., Crevoisier, C., Crill, P. M., Covey, K., Curry, C. L., Etiope, G., 595 Frankenberg, C., Gedney, N., Hegglin, M. I., Höglund-Isakson, L., Hugelius, G., Ishizawa, M., Ito, A., Janssens-Maenhout, G., Jensen, K. M., Joos, F., Kleinen, T., Krummel, P. B., Langenfelds, R. L., Laruelle, G. G., Liu, L., Machida, T., Maksyutov, S., McDonald, K. C., McNorton, J., Miller, P. A., Melton, J. R., Morino, I., Müller, J., Murgia-Flores, F., Naik, V., Niwa, Y., Noce, S., O’Doherty, S., Parker, R. J., Peng, C., Peng, S., Peters, G. P., Prigent, C., Prinn, R., Ramonet, M., Regnier, P., Riley, W. J., Rosentreter, J. A., Segers, A., Simpson, I. J., Shi, H., Smith, S. J., Steele, P. L., 600 Thornton, B. F., Tian, H., Tohjima, Y., Tubiello, F. N., Tsuruta, A., Viovy, N., Voulgarakis, A., Weber, T. S., van Weele, M., van der Werf, G. R., Weiss, R. F., Worthy, D., Wunch, D., Yin, Y., Yoshida, Y., Zhang, W., Zhang, Z., Zhao, Y., Zheng, B., Zhu, Q., Zhu, Q. and Zhuang, Q.: The Global Methane Budget 2000-2017, *Earth Syst. Sci. Data*, 12(3), 1561–1623, doi:10.5194/essd-2019-128, 2019.
- Streets, D. G., Canty, T., Carmichael, G. R., De Foy, B., Dickerson, R. R., Duncan, B. N., Edwards, D. P., Haynes, J. A., 605 Henze, D. K., Houyoux, M. R., Jacob, D. J., Krotkov, N. A., Lamsal, L. N., Liu, Y., Lu, Z., Martin, R. V., Pfister, G. G., Pinder, R. W., Salawitch, R. J. and Wecht, K. J.: Emissions estimation from satellite retrievals: A review of current capability, *Atmos. Environ.*, 77, 1011–1042, doi:10.1016/j.atmosenv.2013.05.051, 2013.
- Thompson, R. L. and Stohl, A.: FLEXINVERT: An atmospheric Bayesian inversion framework for determining surface fluxes of trace species using an optimized grid, *Geosci. Model Dev.*, 7(5), doi:10.5194/gmd-7-2223-2014, 2014.
- 610 Turner, A. J. and Jacob, D. J.: Balancing aggregation and smoothing errors in inverse models, *Atmos. Chem. Phys.*, 15(12), 7039–7048, doi:10.5194/acp-15-7039-2015, 2015.
- Turner, A. J., Jacob, D. J., Wecht, K. J., Maasackers, J. D., Lundgren, E., Andrews, A. E., Biraud, S. C., Boesch, H., Bowman, K. W., Deutscher, N. M., Dubey, M. K., Griffith, D. W. T., Hase, F., Kuze, A., Notholt, J., Ohyama, H., Parker, R., Payne, V. H., Sussmann, R., Sweeney, C., Velazco, V. A., Warneke, T., Wennberg, P. O. and Wunch, D.: Estimating

- 615 global and North American methane emissions with high spatial resolution using GOSAT satellite data, *Atmos. Chem. Phys.*, 15(12), 7049–7069, doi:10.5194/acp-15-7049-2015, 2015.
- Veefkind, J. P., Aben, I., McMullan, K., Förster, H., de Vries, J., Otter, G., Claas, J., Eskes, H. J., de Haan, J. F., Kleipool, Q., van Weele, M., Hasekamp, O., Hoogeveen, R., Landgraf, J., Snel, R., Tol, P., Ingmann, P., Voors, R., Kruizinga, B., Vink, R., Visser, H. and Levelt, P. F.: TROPOMI on the ESA Sentinel-5 Precursor: A GMES mission for global
620 observations of the atmospheric composition for climate, air quality and ozone layer applications, *Remote Sens. Environ.*, 120, 70–83, doi:10.1016/j.rse.2011.09.027, 2012.
- Wecht, K. J., Jacob, D. J., Frankenberg, C., Jiang, Z. and Blake, D. R.: Mapping of North American methane emissions with high spatial resolution by inversion of SCIAMACHY satellite data, *J. Geophys. Res. Atmos. Res.*, 119(12), 7741–7756, doi:10.1002/2014JD021551, 2014.
- 625 Xu, Q.: Measuring information content from observations for data assimilation: Relative entropy versus shannon entropy difference, *Tellus, Ser. A Dyn. Meteorol. Oceanogr.*, 59(2), doi:10.1111/j.1600-0870.2006.00222.x, 2007.