Supplement of Atmos. Meas. Tech., 17, 5429–5454, 2024
https://doi.org/10.5194/amt-17-5429-2024-supplement

Atmospheric Measurement Techniques
Open Access
EGU

Supplement of

Methane retrieval from MethaneAIR using the CO_2 proxy approach: a demonstration for the upcoming MethaneSAT mission

Christopher Chan Miller et al.

Correspondence to: Christopher Chan Miller (cmiller@g.harvard.edu)

S1 Forward Model Details

Let $F(\lambda; \mathbf{x})$ denote the forward model used to simulate the MethaneAIR radiance at wavelength λ. $\mathbf{x}$ is the state vector to be optimized. It consists of a set of sub-vectors;

$$\mathbf{x} = \begin{pmatrix} \mathbf{x_{co2}} \\ \mathbf{x_{ch4}} \\ \mathbf{x_{h2o}} \\ \mathbf{x_{ps}} \\ \mathbf{x_T} \\ \mathbf{x_{alb}} \\ \mathbf{x_{sqz}} \\ \mathbf{x}_\lambda \\ \mathbf{x_b} \end{pmatrix} \tag{S1}$$

Here $\mathbf{x_{co2}}$ and $\mathbf{x_{ch4}}$ are the CO_2 and CH_4 dry-air mole fraction profile scaling factors, $\mathbf{x_{h2o}}$ is a H_2O column scaling factor, $\mathbf{x_{ps}}$ is the surface pressure, and $\mathbf{x_T}$ is an offset from the a priori temperature profile. $\mathbf{x_{alb}}$ consists of two sets of polynomial coefficients corresponding to each fit window $\mathbf{x_{alb}} = (\mathbf{x_{alb,co2}}, \mathbf{x_{alb,ch4}})$. These act to scale the a priori albedo calculated as the Lambertian equivalent reflectance from the MethaneAIR observation at 1622.5nm.

The forward model consists of the convolution of radiative transfer model (RTM) simulated radiance ($I(\lambda)$) with the instrument spectral response function ($\Gamma(\lambda)$), plus a baseline offset intended to account for an additive bias present in the measured radiance, such as broad-band stray light.

$$F(\lambda; \mathbf{x}) = I(\lambda; \mathbf{x_{co2}}, \mathbf{x_{ch4}}, \mathbf{x_{h2o}}, \mathbf{x_{ps}}, \mathbf{x_T}, \mathbf{x_{alb}}) \circledast \Gamma(\lambda; \mathbf{x_{sqz}}) + P_b(\lambda; \mathbf{x_b}) \tag{S2}$$

S1.1 Radiative Transfer Model

Currently there are four radiative transfer options for the solar back-scatter problem in SPLAT, namely (in order of additional complexity):

1. Absorption Only
2. Stand-alone VLIDORT First Order Model - Adds direct contribution from first order scattering
3. Stand-alone LIDORT Two stream model - Adds two stream diffuse radiation field (Intensity only)
4. VLIDORT v2.8.3 - Full n-stream quadrature and polarization

The stand-alone VLIDORT First-Order and two-stream produce identical results to VLIDORT, but have been computationally optimized for the simpler assumptions. The choice of RTM can easily be changed and SPLAT can be configured to run in both forward and inverse modes, enabling it to be used as a single environment for performing realistic observing system simulation experiments (e.g. Chan Miller et al. (2013)).

The CO_2-proxy retrieval described here uses the simplest absorption-only model - this is due to the fact that the retrieval does not explicitly simulate aerosols, relying on CO_2 light-path normalization to account for aerosol scattering. Since Rayleigh scattering in the CH_4 band is small, all scattering can be neglected. Thus $I(\lambda)$ can be written as:

$$I(\lambda) = \frac{\cos(SZA)A(\lambda)}{\pi} \exp\left(-\sum_l S_l \tau_l(\lambda)\right) \tag{S3}$$

In the above $A(\lambda)$ is the Lambertian surface albedo, τ_l absorption optical depth, SZA the solar zenith angle, and S_l is the geometry air mass factor. Let p_{l-1} and p_l be the pressure levels at the bottom and top of layer l, and the p_{obs} the pressure at the point of observation. Then the geometric air mass is

$$S_l = \begin{cases} \frac{1}{\cos(SZA)} & \dots p_{l-1} < p_{obs} \text{ (layer above aircraft)} \\ \frac{1}{\cos(SZA)} + \frac{1}{\cos(VZA)} & \dots p_l > p_{obs} \text{ (layer below aircraft)} \\ \frac{1}{\cos(SZA)} + \frac{1}{\cos(VZA)} \frac{p_{l-1} - p_{obs}}{p_l - p_{obs}} & \text{else (aircraft in layer)} \end{cases} \tag{S4}$$

where VZA is the viewing zenith angle.

S1.2 Profile Definitions and Optical Properties

In this section we describe the preparation of optical properties in SPLAT. The apriori profile pressure and temperature levels are a function of the a priori surface and tropopause pressure ($p_{S,0}$ and $p_{T,0}$ respectively). The pressure at level l is defined as

$$p_{0,l} = a_l(p_{S,0} - p_{T,0}) + b_l p_{T,0} + c_l \tag{S5}$$

The coefficients a_l,b_l, and c_l are those used by the University of Leicester GOSAT CO_2-Proxy retrieval (Parker et al., 2020), and are provided below:

Hybrid Pressure Grid Coefficients			
Level	a_l	b_l	c_l
0	1.0	1.0	0.0
1	0.9230769230769231	1.0	0.0
2	0.8461538461538463	1.0	0.0
3	0.7692307692307693	1.0	0.0
4	0.6923076923076923	1.0	0.0
5	0.6153846153846154	1.0	0.0
6	0.5384615384615385	1.0	0.0
7	0.4615384615384616	1.0	0.0
8	0.3846153846153846	1.0	0.0
9	0.3076923076923077	1.0	0.0
10	0.23076923076923078	1.0	0.0
11	0.15384615384615385	1.0	0.0
12	0.07692307692307693	1.0	0.0
13	0.0	1.0	0.0
14	0.0	0.5	40.0
15	0.0	0.0	80.0
16	0.0	0.0	50.0
17	0.0	0.0	10.0
18	0.0	0.0	1.0
19	0.0	0.0	0.1

The absorption optical depth at layer l is computed as

$$\tau_l(\lambda) = \sum_i v_l^{[AIR]} X_{i,l} \sigma_i(\lambda, p_{l-1}, p_l, T_{l-1}, T_l) \tag{S6}$$

In the above equation X_i is the mole fraction of gas i and $\sigma_i(\lambda, p_{l-1}, p_l, T_{l-1}, T_l)$ is the corresponding absorption cross section. Here they are computed from a GGG2020-based pressure/temperature lookup table on a 4-layer even-pressure sub-grid, and the result is averaged to better account for pressure broadening within the layer. $v_l^{[AIR]}$ is the column density within layer l. It is given by

$$v_l^{[AIR]} = \frac{N_A}{M_l g_l}(p_{l-1} - p_l) \tag{S7}$$

where N_A is Avogadro's number, M_l is the mean molecular weight of air in layer l, and g_l is the gravitational acceleration, computed using the latitude-dependent parameterization found in GGG2020, based on equation I.2.4-(17) in US Standard Atmosphere 1962. This parameterization accounts for centripetal acceleration due to Earth's rotation and the variation in gravity due to Earth's deviation from an ideal sphere.

S1.3 Radiative Transfer Model State Vector Elements

Whilst the initial pressure grid uses the hybrid parameterization in the previous section, surface pressures are optimized using sigma coordinates. The fitted surface pressure at layer l is

$$p_l = \frac{x_{ps}}{p_{S,0}} p_{0,l} \tag{S8}$$

Thus the surface pressure changes the spectrum in two ways - by adding more mass through its effect on $v_l^{[AIR]}$, and changing the gas absorption cross sections. A temperature offset is also included in the fit. In terms of the a priori temperature in level l, $T_{0,l}$

$$T_l = T_{0,l} + x_T \tag{S9}$$

The main effect of temperature is changing the absorption cross sections. In the retrieval the absorption cross sections are re-interpolated from their lookup tables each iteration to account for changes from the pressure/temperature state vector elements.

CO_2 and CH_4 are optimized using scaling factors that operate on the a priori profile of dry-air mole fractions $X_{0,i,l}$. For i =CO_2, CH_4

$$X_{i,l} = x_{i,l} X_{0,i,l} \tag{S10}$$

A single column scaling factor is applied to the apriori H_2O mole fraction profile $X_{0,H_2O,l}$

$$X_{H_2O,l} = x_{H_2O} X_{0,H_2O,l} \tag{S11}$$

The albedo for fit window w is parameterized using Chebyshev Polynomials of the first kind (V_i).

$$A_w(\lambda) = A_0 \sum_{i=0}^{n} x_{alb,w,i} V_i \left(2 \frac{\lambda - \lambda_{w,min}}{\lambda_{w,max} - \lambda_{w,min}} - 1 \right) \tag{S12}$$

A_0 is the albedo derived from the radiance at 1622.5 nm. $\lambda_{w,min}$ and $\lambda_{w,max}$ are the fit window lower and upper limits, with an additional 2 nm buffer added to allow for ISRF convolution.

S1.4 ISRF and ISRF Squeeze State Variables

The ISRF has been calibrated in-lab with a tuneable laser, stepped at 5 nm increments (Staebell et al., 2021). To evaluate the convolution at pixel wavelength λ_p

$$I \circledast \Gamma(\lambda_p) = \int_{\lambda_p - \delta_\lambda}^{\lambda_p + \delta_\lambda} I(\lambda_p) \Gamma(\lambda_p - \lambda) d\lambda \tag{S13}$$

The above integral is evaluated numerically on a 0.01 nm grid. Let λ_j be the closest wavelength in the lookup table that is less than λ_p. The ISRF is evaluated as follows:

$$\Gamma(\lambda_p - \lambda) = w_j \Gamma_{\lambda_j}^{LUT}(x_{sqz}(\lambda_p - \lambda)) + w_{j+1} \Gamma_{\lambda_{j+1}}^{LUT}(x_{sqz}(\lambda_p - \lambda)) \tag{S14}$$

In the above expression $\Gamma_{\lambda_j}^{LUT}$ represents the function that linearly interpolates over the ISRF at center wavelength λ_j. In the above equation, x_{sqz} represents the fitted ISRF squeeze factor ($x_{sqz,co2}$ or $x_{sqz,ch4}$ depending on the fit window). w_j and w_{j+1} are the linear weightings between the lookup table center wavelength nodes.

$$w_j = \frac{\lambda_p - \lambda_j}{\lambda_{j+1} - \lambda_j} \tag{S15}$$

$$w_{j+1} = 1 - w_j \tag{S16}$$

S1.5 Wavelength Shift State Variable

Upon simulation, the forward model is linearly interpolated from the 0.01 convolution grid to the wavelength grid of MethaneAIR. A wavelength shift is included in the spectral fit. Letting $\lambda_{0,i}$ represent the wavelength of the i^{th} spectral pixel, then the forward model simulation corresponding to the observation element y_i is:

$$F_i = F(\lambda_{0,i} + x_\lambda) \tag{S17}$$

S1.6 Solving the Inverse Problem

The inverse problem is formulated using Bayesian inference following Rodgers (2000), and we use the same notation as the reference. Here we fit the forward model, $\mathbf{F}(\mathbf{x})$, to an observed m-pixel radiance spectrum $\mathbf{y} \in \mathbb{R}^m$

$$\mathbf{y} = \mathbf{F}(\mathbf{x}) + \boldsymbol{\epsilon} \tag{S18}$$

If the observed radiance pixel errors $\boldsymbol{\epsilon}$ are distributed by a multivariate normal ($\mathcal{N}(\mathbf{0}, \mathbf{S_o})$), and the a priori uncertainty of the state variables is distributed $\mathcal{N}(\mathbf{x_a}, \mathbf{S_a})$, the *maximum a posteriori* solution $\hat{\mathbf{x}}$ (that is the solution maximizing $P(\mathbf{x}|\mathbf{y})$) is found by minimizing:

$$J(\mathbf{x}) = (\mathbf{y} - \mathbf{F}(\mathbf{x}))^T {\mathbf{S_o}}^{-1} (\mathbf{y} - \mathbf{F}(\mathbf{x})) + (\mathbf{x} - \mathbf{x_a})^T {\mathbf{S_a}}^{-1} (\mathbf{x} - \mathbf{x_a}) \tag{S19}$$

$J(\mathbf{x})$ is optimized using the Levenberg–Marquardt algorithm, following the approach used by the OCO2 full-physics algorithm (Crisp et al., 2021). Here the change in state vector on the $i + 1^{th}$ iteration, $\mathbf{dx_{i+1}}$, is given by:

$$\mathbf{dx_{i+1}} = \left((1+\alpha) {\mathbf{S_a}}^{-1} + {\mathbf{K_i}}^T {\mathbf{S_o}}^{-1} \mathbf{K_i} \right)^{-1} \left({\mathbf{K_i}}^T {\mathbf{S_o}}^{-1} (\mathbf{y} - \mathbf{F}(\mathbf{x_i})) + {\mathbf{S_a}}^{-1} (\mathbf{x_a} - \mathbf{x_i}) \right) \tag{S20}$$

$\mathbf{K_i}$ is the Jacobian of the forward model at the i^{th} iteration ($\mathbf{K_i} = \frac{\partial \mathbf{F}(\mathbf{x_i})}{\partial \mathbf{x}}$). α is the Levenberg-Marquardt damping parameter, here initialized at 100. For large values, the algorithm approximates the gradient descent method. As the solution approaches the optimum, the parameter can be relaxed, approaching a Gauss-Newton step, accelerating convergence. The parameter is based on the following metric measuring forward model nonlinearity, computed as the ratio of actual cost function reduction relative to the situation where the forward model is assumed linear:

$$R = \frac{J(\mathbf{x_i}) - J(\mathbf{x_{i+1}})}{J(\mathbf{x_i}) - J_{lin}(\mathbf{x_i}, \mathbf{dx_{i+1}})} \tag{S21}$$

$$J_{lin}(\mathbf{x}, \mathbf{dx}) = (\mathbf{y} - (\mathbf{F}(\mathbf{x}) + \mathbf{K dx}))^T {\mathbf{S_o}}^{-1} (\mathbf{y} - (\mathbf{F}(\mathbf{x}) + \mathbf{K dx})) + (\mathbf{x} - \mathbf{x_a})^T {\mathbf{S_a}}^{-1} (\mathbf{x} - \mathbf{x_a}) \tag{S22}$$

Values of $R > 0.75 (< 0.25)$ are considered convergent(divergent), and α is scaled by a factor of 0.5(10). The convergence criteria are based on the error variance derivative. For the i^{th} iteration this is:

$$d\sigma_i^2 = {\mathbf{dx_{i+1}}}^T \hat{\mathbf{S}}^{-1} \mathbf{dx_{i+1}} \tag{S23}$$

$$\approx {\mathbf{dx_{i+1}}}^T \left({\mathbf{K_i}}^T {\mathbf{S_o}}^{-1} (\mathbf{y} - \mathbf{F}(\mathbf{x_i})) + {\mathbf{S_a}}^{-1} (\mathbf{x_i} - \mathbf{x_a}) \right) \tag{S24}$$

This essentially measures the ratio of the state vector update to the a posteriori uncertainty. These will be approximately equal when $d\sigma_i^2 = n$, where n is the dimension of the state vector. Here pixels are considered converged when $d\sigma_i^2 < 0.02n$.

S1.7 Tuning of the A Priori Covariance

Here we tune the regularization of the MethaneAIR retrieval by scaling the a priori covariance matrix. In terms of the cost function $J(\mathbf{x})$ this is achieved by introducing the regularization parameter γ:

$$J(\mathbf{x}) = (\mathbf{y} - \mathbf{F}(\mathbf{x}))^T {\mathbf{S_o}}^{-1} (\mathbf{y} - \mathbf{F}(\mathbf{x})) + \gamma^{-2} (\mathbf{x} - \mathbf{x_a})^T {\mathbf{S_a}}^{-1} (\mathbf{x} - \mathbf{x_a}) \tag{S25}$$

For the observations to be useful for estimating methane emissions, to first order the appropriate choice of γ is guided by the need to generate a 1:1 response to boundary layer CH_4 concentrations. The vertical sensitivity of the retrieval is quantitatively expressed through the column averaging kernel $\mathbf{a_{CH4}}$, which represents the fractional change in retrieved methane vertical column $\hat{N}_{CH4}$ relative to a perfect observing system where departures from the prior are perfectly fitted. Following Connor et al. (2008) the column averaging kernel element for the l^{th} layer is related to the CH_4 profile component of the state vector $\mathbf{x}_{CH4}$ by

$$a_{CH4,l} = \frac{1}{h_l} \frac{\partial \hat{N}_{CH4}}{\partial x_{CH4,l}} = \frac{1}{h_l} \mathbf{h}^T \mathbf{A}_{CH4} \tag{S26}$$

Here $\mathbf{A_{CH4}}$ corresponds to the block of the retrievals averaging kernel covered by the profile scale factors. $\mathbf{h}$ is the column operator that maps the retrieved scale factors to the methane column (i.e. $N_{CH4} = \mathbf{h}^T \mathbf{x}_{CH4}$).

Similar expressions exist for the CO_2 averaging kernel. The a priori covariance matrices for CH_4 and CO_2 largely control the averaging kernel responses. Here we use those from the UoL GOSAT retrieval algorithm (Figure S1, Parker et al. (2020)). Early testing on 15×3 (across×along track) pixel aggregates suggested that CH_4 was under-regularized relative to CO_2 (not shown here), likely due to the lower spectral information content for MethaneAIR relative to GOSAT. As a result, the original CH_4 matrix was reduced by a factor of 9 to produce unit boundary layer column averaging kernel values. After the scaling, the uncertainty is still loose relative to that expected from nature, with standard deviations >300 ppbv in the boundary layer (Figure S1).

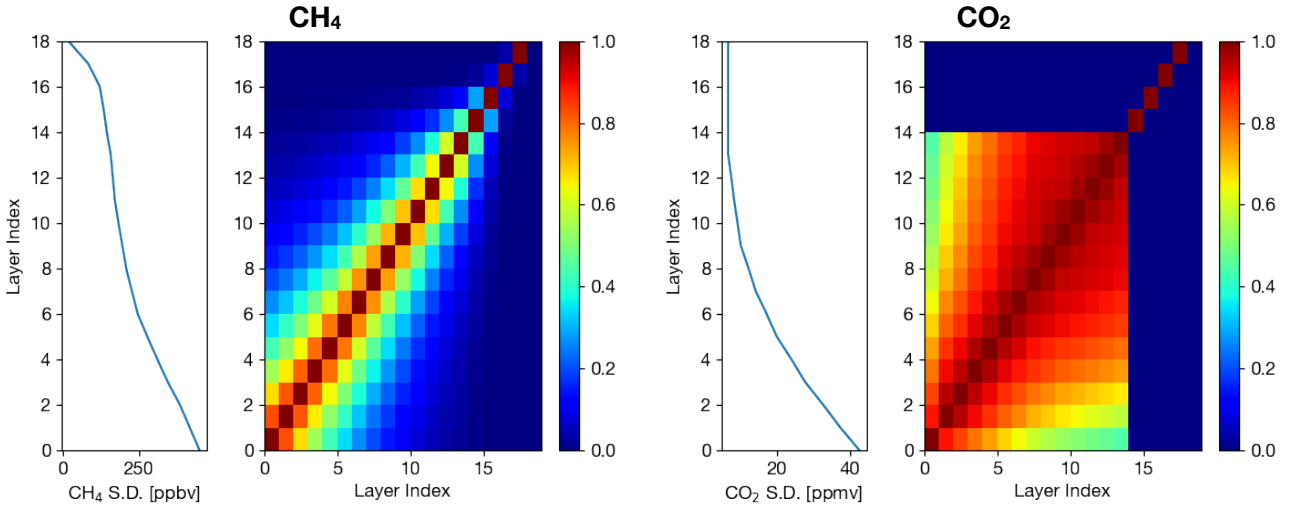

Figure S1: A priori error covariance matrices for CH_4 and CO_2 used by the MethaneAIR retrieval. The covariance matrices have been decomposed into layer standard deviation profiles (left) and correlation matrices (right).

In this paper, retrievals were performed at native and 5×1 pixel aggregations. To maintain a uniform surface response γ must be re-tuned due to the noise increase from the lower pixel aggregation. If γ_n corresponds to the tuned value for the n-pixel aggregate, it can easily be shown that to obtain the same averaging kernel for an m pixel aggregate, γ_m should be scaled as follows;

$$\gamma_m^2 = \frac{m}{n}\gamma_n^2 \tag{S27}$$

To show this, the averaging kernel $\mathbf{A}$ is related to the retrieval gain matrix $\mathbf{G}$ by

$$\mathbf{G} = \left(\mathbf{K}^T {\mathbf{S_o}}^{-1}\mathbf{K} + \frac{1}{\gamma^2}{\mathbf{S_a}}^{-1}\right)^{-1}\mathbf{K}^T {\mathbf{S_o}}^{-1} \tag{S28}$$

$$\mathbf{A} = \mathbf{GK} \tag{S29}$$

Where $\mathbf{K}$ is the jacobian of the forward model, let $\mathbf{S_{o},}_n$ be the observation error covariance for an n-pixel aggregation. By the central limit theorem $\mathbf{S_{o},}_n = \frac{m}{n}\mathbf{S_{o},}_m$. Matching the averaging kernel requires matching the gain.

$$\mathbf{G}_n = \left(\mathbf{K}^T\mathbf{S}_{\mathbf{o},n}^{-1}\mathbf{K} + \frac{1}{\gamma_n^2}{\mathbf{S_a}}^{-1}\right)^{-1}\mathbf{K}^T\mathbf{S}_{\mathbf{o},n}^{-1} \tag{S30}$$

$$= \frac{n}{m}\left(\frac{n}{m}\mathbf{K}^T\mathbf{S}_{\mathbf{o},m}^{-1}\mathbf{K} + \frac{1}{\gamma_n^2}{\mathbf{S_a}}^{-1}\right)^{-1}\mathbf{K}^T\mathbf{S}_{\mathbf{o},m}^{-1} \tag{S31}$$

substituting equation S27

$$\begin{aligned}\mathbf{G}_n &= \frac{n}{m}\left(\frac{n}{m}\left(\mathbf{K}^T\mathbf{S}_{\mathbf{o},m}^{-1}\mathbf{K} + \frac{1}{\gamma_m^2}{\mathbf{S_a}}^{-1}\right)\right)^{-1}\mathbf{K}^T\mathbf{S}_{\mathbf{o},m}^{-1} \\ &= \left(\mathbf{K}^T\mathbf{S}_{\mathbf{o},m}^{-1}\mathbf{K} + \frac{1}{\gamma_m^2}{\mathbf{S_a}}^{-1}\right)^{-1}\mathbf{K}^T\mathbf{S}_{\mathbf{o},m}^{-1} \\ &= \mathbf{G}_m\end{aligned} \tag{S32}$$

To tune the regularization we perform a set of retrievals, where γ^2 is systematically varied over a range of 10^{-4} to 10^4 using a 5×1 aggregated MethaneAIR observation from the return leg of RF06. We select an observation with relatively low signal to ensure that signal-dependent artifacts caused by over-regularization are minimized. In this case, the solar zenith angle was 46°, and albedo 0.17, making it within the lowest 5 percentile of albedos observed during the return leg. The resulting dependence of the CO_2 and CH_4 averaging kernels is shown in Figure S2 (middle, right). As γ^2 increases, the column averaging kernels tend towards 1, with both species having 1:1 responses for $\gamma^2 \geq 10$. This confirms the original tuning performed on the 15×3 pixel aggregates, as based on Equation S27 we would expect to loosen the regularization by a factor of 9. In all cases, there remains a steep drop off in the averaging kernel sensitivity above 200 hPa. This corresponds to the airmass above the aircraft, and occurs because light is only traversing the layers above once.

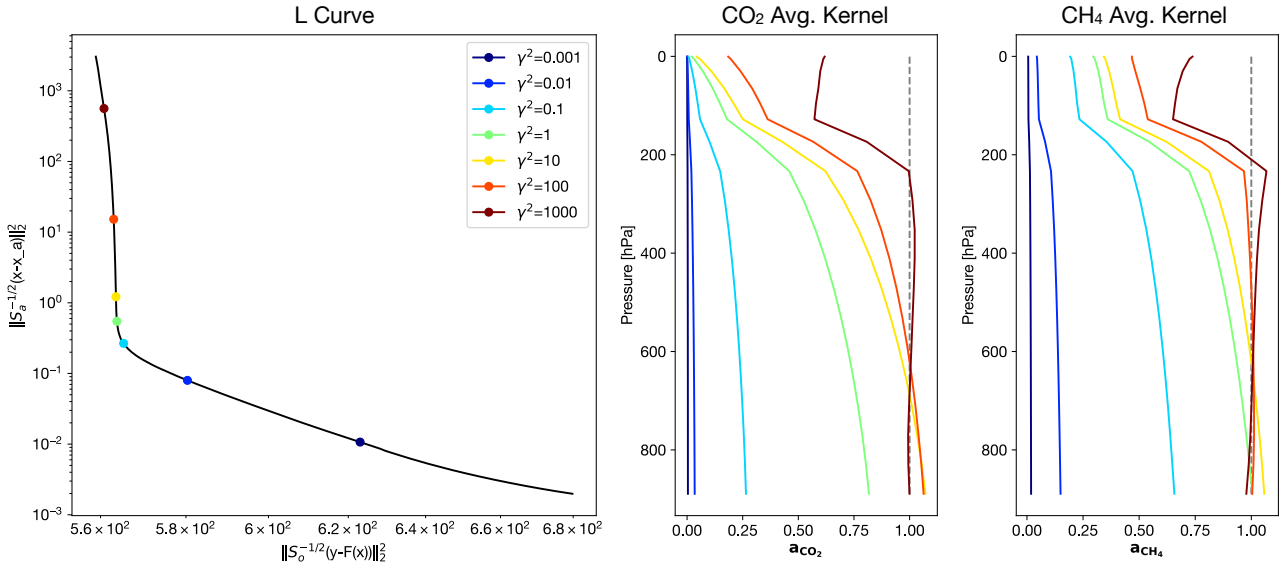

Figure S2: L-Curve and CO_2/CH_4 column averaging kernels for the test pixel used to tune the MethaneAIR regularization factor for the 5×1 pixel aggregation (see text).

Selecting values of γ larger than 10 marginally improves the vertical sensitivity at the cost of worsening the measurement precision. The appropriate trade-off point can be determined heuristically using the L-curve method (Hansen, 1993), which compares the two components of $J(\mathbf{x})$ (Equation S25) (i.e. the norm of the spectral residuals versus the norm of the difference between the retrieved and a priori state) at different γ. Figure S2 (left) shows the L-curve for the test pixel. At high regularization (low γ), increasing γ reduces the spectral fit residuals without significant increases in the regularization norm, indicating over-regularization. At the other extreme (high γ), increasing γ mostly increases the regularization norm, reflecting the situation where measurement noise starts dominating the retrieved state. The trade-off point where the retrieval does not under/overfit the spectrum can be estimated as the point of maximum curvature between these two extremes. From Figure S2 this occurs within $\gamma^2 = 0.1 - 10$. In consideration of this and the averaging kernel boundary layer response, we select $\gamma^2 = 10$ for the 5×1 observation and scale to $\gamma^2 = 50$ for the native resolution (Equation S27).

Figure S3 shows the CO_2 and CH_4 degrees of freedom for signal (DoFS) for all pixels for the RF06 return leg as a function of albedo. The DoFS are the trace of the averaging kernel CO_2/CH_4 sub-matrices, and are a measure of the individual pieces of profile information that can be constrained by the retrieval. Figure S3 shows that there is a steep drop off in the DoFS curve for CO_2 starting at ~ 0.10 albedo. Below this value the retrieved CO_2 column will increasingly be dominated by the prior. The drop-off in CH_4 DoFS occurs at smaller signal levels, indicating that the retrieval is limited by the CO_2 light path constraint. The majority of scenes corresponding to the prior-dominated signal levels correspond to those over water.

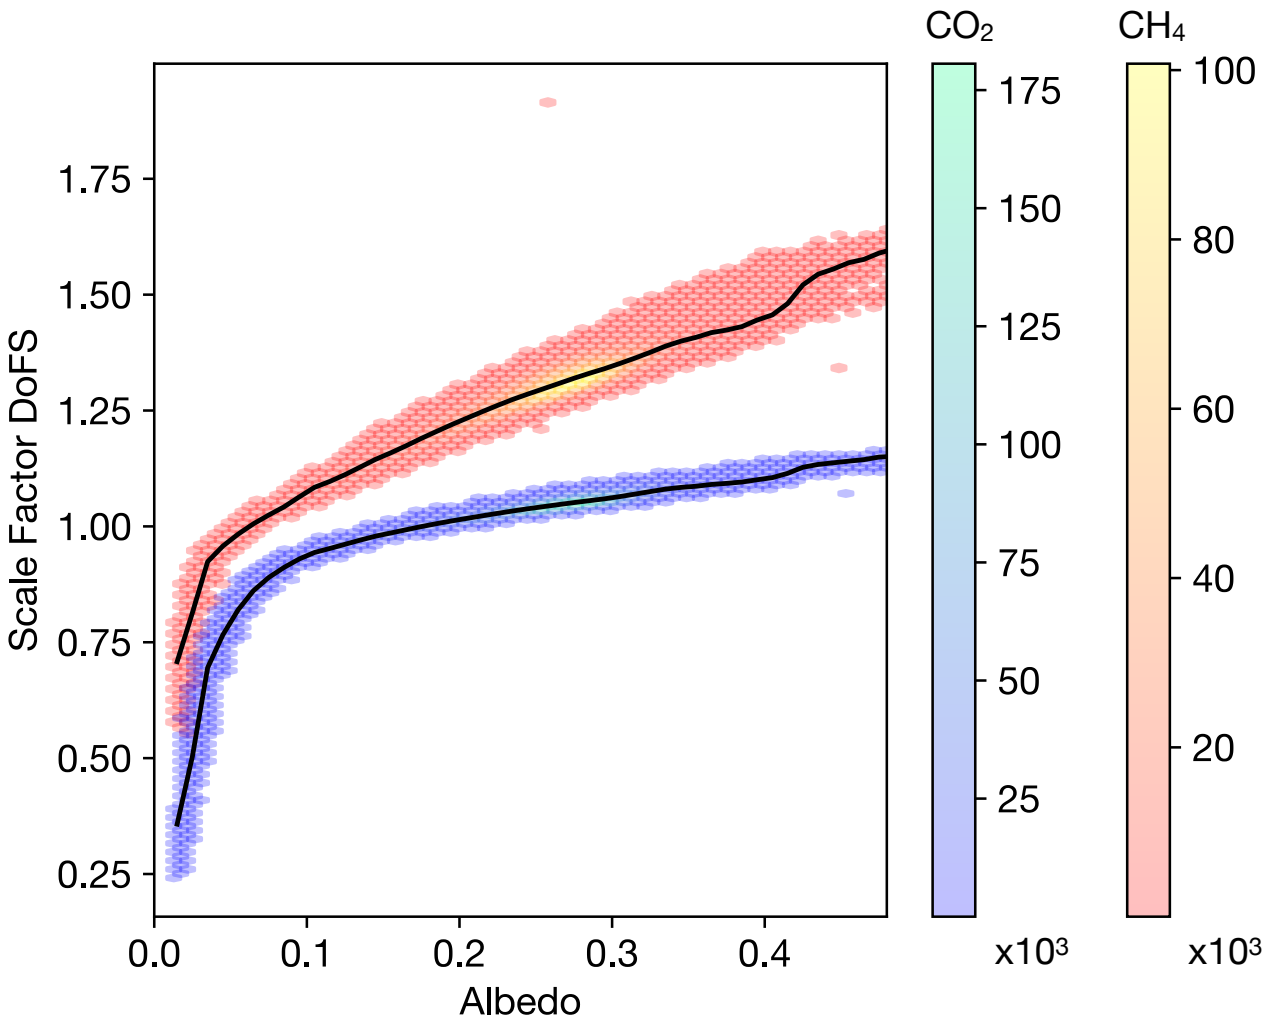

Figure S3: CO_2 and CH_4 degrees of freedom for signal as a function of scene albedo for the return segment used to determine albedo XCH_4 (see main text). Color maps represent observation density.

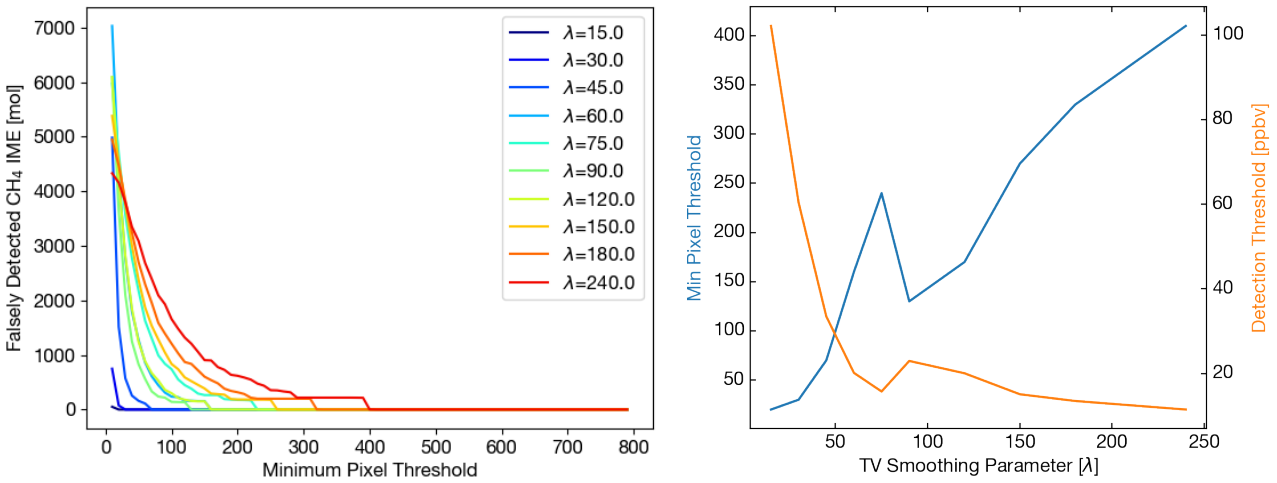

Figure S4: Determination of plume mask parameters from plume-free synthetic MethaneAIR retrievals at native resolution. The left panel shows the falsely detected plume mass for different levels of TV smoothing as a function of the minimum pixel cluster threshold. The right panel shows the minimum pixel threshold required to exclude all falsely detected plume mass as a function of smoothing weight (blue line). The plume XCH_4 threshold versus smoothing weight is also shown (orange line).

S2 Implications for MethaneSAT

S2.1 Estimation of MethaneSAT Precision

Here we provide an estimate of the MethaneSAT XCH_4 precision via a linear sensitivity analysis. The precision of the state vector $\mathbf{x}$, is characterized by the measurement error covariance $\hat{\mathbf{S}}_\mathbf{m}$. $\hat{\mathbf{S}}_\mathbf{m}$ is determined via the retrieval model gain ($\mathbf{G}$, Equation S28), and radiance error described by the observation error covariance matrix $\mathbf{S_o}$:

$$\hat{\mathbf{S}}_\mathbf{m} = \mathbf{G}\mathbf{S_o}\mathbf{G}^T \tag{S33}$$

Here $\mathbf{S_o}$ is computed using the SNR model provided by Ball Aerospace to define the MethaneSAT instrument requirements in the mission design phase. We scale these values to match an SNR of 190 for an input radiance of 1.4×10^{13} photons cm^{-2} s^{-1} sr^{-1} nm^{-1}, based on preliminary MethaneSAT calibration measurements which show the instrument outperforming the initial noise requirements. The linear sensitivity analysis uses the Jacobian computed from $\mathbf{x_a}$, thus making the assumption that this is a good approximation for the Jacobian at the final fit iteration. To estimate XCH_4 precision (σ_{XCH_4}), we must do one final round of error propagation via the column operator $\mathbf{k} = \frac{\partial XCH_4}{\partial \mathbf{x}}$:

$$\sigma^2_{XCH_4} = \mathbf{k}^T \hat{\mathbf{S}}_\mathbf{m} \mathbf{k} \tag{S34}$$

$\mathbf{k}$ accounts for the errors in both the retrieved CH_4 and CO_2 vertical columns:

$$\mathbf{k} = XCO_{2,0} \frac{\partial}{\partial \mathbf{x}} \left(\frac{N_{CH_4}(\mathbf{x})}{N_{CO_2}(\mathbf{x})} \right) \tag{S35}$$

$$= XCO_{2,0} \left(\frac{1}{N_{CO_2}} \frac{\partial N_{CH_4}}{\partial \mathbf{x}} - \frac{N_{CH_4}}{N^2_{CO_2}} \frac{\partial N_{CO_2}}{\partial \mathbf{x}} \right) \tag{S36}$$

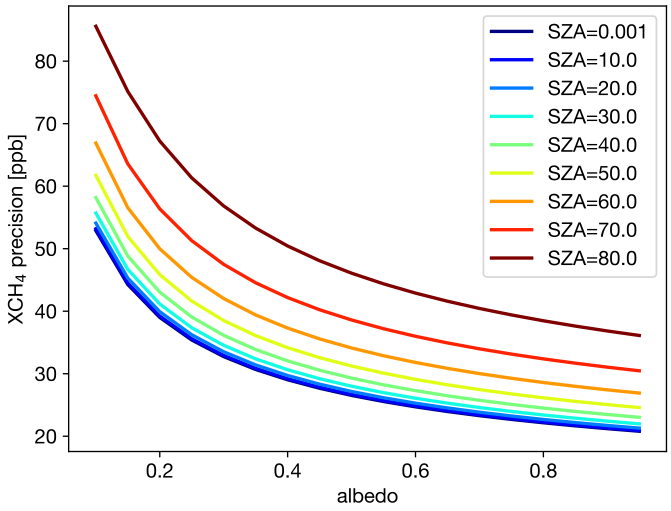

Figure S5: Estimated XCH_4 precision via the linear sensitivity analysis (described in text) as a function of Lambertian albedo. Lines correspond to different solar zenith angles.

Figure S5 shows the XCH_4 precision estimated via the linear sensitivity analysis, as a function of albedo and solar zenith angle. The calculation used a sampled GGG2020 a priori profile from RF06 over the Permian, and with a viewing zenith angle of 0 degrees. The precision curves look similar at other viewing zenith angles as the precision is primarily determined by the radiance, which itself is primarily a function of solar zenith angle and surface reflectance. For a typical surface albedo of 0.3, the random measurement error is less than 35 ppb for solar zenith angles below 30 degrees, and 40 ppb below 50 degrees.

S2.2 Impact of Colocated CO_2 Emissions

The CO_2 proxy approach relies on knowledge of the true CO_2 dry air total column mole fraction (XCO_2). Here we focus on errors induced by the a priori ($XCO_{2,0}$). To isolate this error source, we assume that the retrieval perfectly retrieves the CH_4 and CO_2 vertical densities (N_{CH_4} and N_{CO_2} respectively), then the true (XCH_4) and retrieved ($X\hat{C}H_4$) CO2 dry air total column mole fractions can be written

$$XCH_4 = \frac{N_{CH_4}}{N_{CO_2}} XCO_2 \tag{S37}$$

$$X\hat{C}H_4 = \frac{N_{CH_4}}{N_{CO_2}} XCO_{2,0} \tag{S38}$$

$$= XCH_4 \frac{XCO_{2,0}}{XCO_2} \tag{S39}$$

Thus the error scales with the ratio of the a priori and true XCO_2. If there is little XCO_2 variability within a scene, this will not cause a significant error because the a priori error in the XCO_2 background is usually within 1% of the truth (Laughner et al., 2023), leading to a small scaling bias. However, if there is CO2 variation within the scene, it may induce errors of similar magnitude to the CH_4 enhancements above the background due to surface emissions. To illustrate this, we use the extreme example of a gas flare, a point source of both species. Let E represent the rate of gas flowing through the flare in mol C h^{-1}. The emissions of CH_4 and CO_2 (E_{CH_4} and E_{CO_2}) may be written

$$E_{CO_2} = E \times DRE \tag{S40}$$

$$E_{CH_4} = E(1 - DRE)F_C \tag{S41}$$

Here F_C represents the mole fraction of C that is methane, and DRE is the destruction removal efficiency of the flare. In order to simulate realistic values, we use DRE and F_C values obtained from a recent aircraft survey of gas flares of major US oil and gas basins (Plant et al., 2022). We simulate XCH_4 and XCO_2 enhancements using the same WRF LES simulations described in the main text, and assume background XCH_4 and XCO_2 of 1900 ppb and 400 ppm, respectively. We use Equation S39 combined with the WRF-LES XCH_4 and XCO_2 of to estimate the retrieved XCH_4, under the assumption that $XCO_{2,0}$ is equal to the 400 ppm background. Figure S6 shows an example realization plume from a vent releasing 1000 kg/h of methane unlit, with Permian gas composition, at 3 DREs ranging from 0-100%.

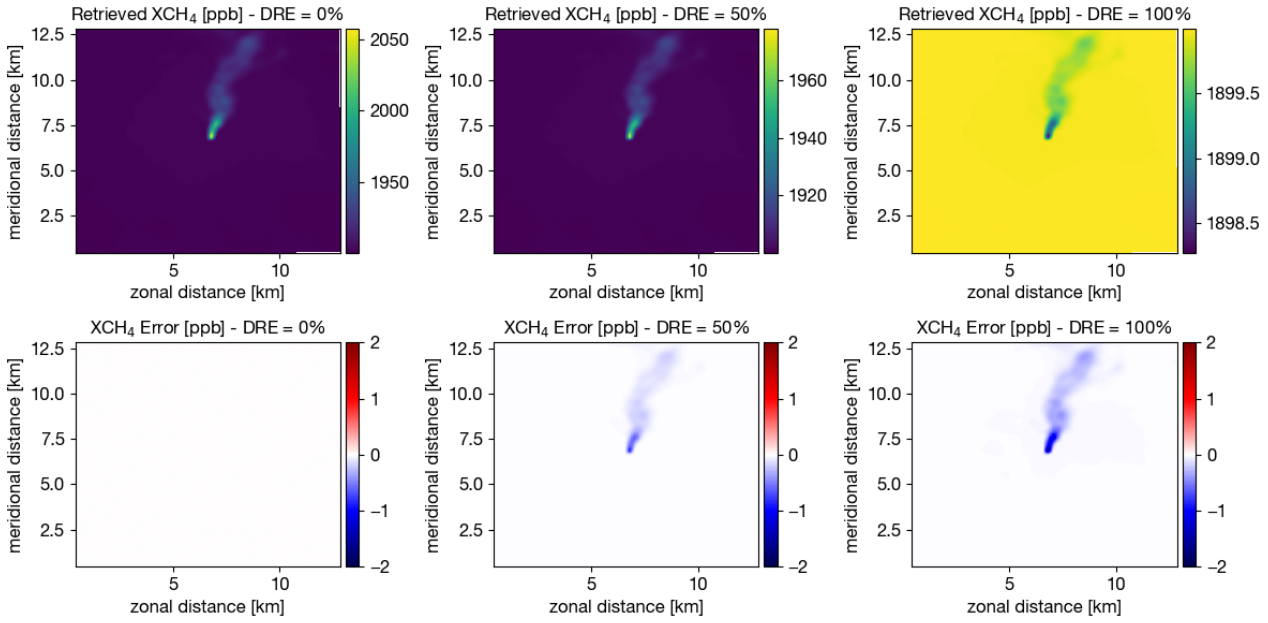

Figure S6: Retrieved XCH_4 from a gas flare simulated with WRF-LES for 3 different destruction removal efficiencies (top row). When unlit, the flare emits 1000 kg/h methane (see text). The bottom row shows errors induced by co-emitted CO_2.

For the 0 and 50% DRE cases, the XCH_4 error is small relative to the enhancement. When combustion is 100% efficient, the CO_2 enhancement reduces the retrieved XCH_4 below the background. To quantify how

this would impact emissions estimates, we can compare the integrated mass enhancements for the true and retrieved XCH_4 . Figure S7 shows the ratio of the retrieved IMEs, computed from the ensemble of WRF LES realizations. The individual curves correspond to the gas compositions of the mean Permian, Eagle Ford, and Bakken formation, which represent 80% of US oil and gas flaring activity (Plant et al., 2022).

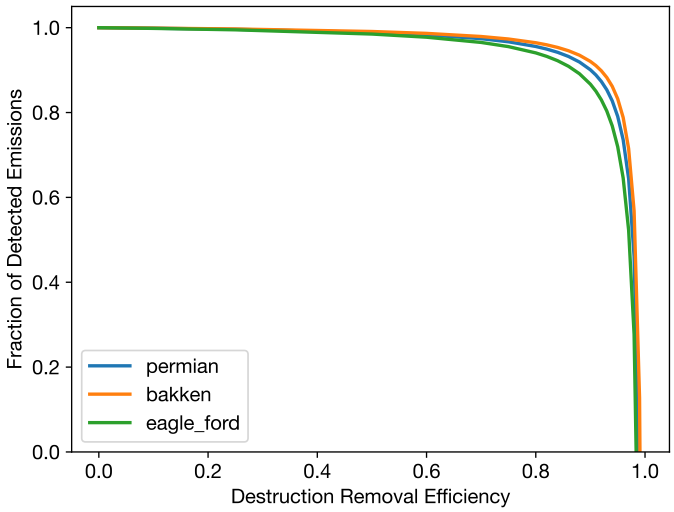

Figure S7: Fraction of integrated CH_4 mass detected (after considering error due to co-emitted CO_2) as a function of destruction removal efficiency for 3 different basin compositions.

Figure S7 shows that the fraction of detected emissions precipitously declines at high DREs. Whilst well over 80% of total emissions are detected for a 95% DRE, this drops to 28-58% by a DRE of 98%, the nominal value used by the EPA GHGRP (GHGRP). However, the current assumed rate likely underestimates current flare efficiencies, with Plant et al. (2022) estimating a US-average flaring efficiency $\sim 91.1\%$ based on their basin surveys. We can use their reported flaring distributions to estimate the fraction of emissions missed due to the CO2 proxy normalization error. We calculate the total true (E_{true}) and observed (E_{obs}) flaring emissions for each basin as follows:

$$E_{true} = f_{unlit} + (1 - f_{unlit}) \int_0^1 p(DRE)(1 - DRE)dDRE \tag{S42}$$

$$E_{obs} = f_{unlit} + (1 - f_{unlit}) \int_0^1 p(DRE)(1 - DRE)f_{obs}(DRE)dDRE \tag{S43}$$

f_{unlit} is the fraction of unlit flares within the basin (4.9%, 4.1%, and 3.2% for the Permian, Eagle Ford, and Bakken respectively), f_{obs} is the fraction of detected emissions (Figure S7), and $p(DRE)$ is the probability distribution of the flare DREs. Here we use the log-normal distribution parameters provided by Plant et al. (2022). From this we estimate that the error caused by co-emitted CO_2 causes gas-flaring emissions in the Eagle Ford, Bakken, and Permian basins by 16%, 13%, and 8% respectively.

Whilst under current conditions, it appears the majority of flaring emissions can be detected with the CO_2 proxy method, the potential of improved flaring efficiencies spurred by recent global agreements may require this to be revisited. We have also not studied co-emitted area sources, which may be important for certain targets (e.g. cities). For such cases, we are currently exploring the use of un-normalized CH_4 columns or proxy-normalization from O_2 retrieved from the nearby singlet-delta band. This may work because of the spectral proximity to the CH_4 band; however, requires an evaluation through observing system simulation experiments. For co-emitted area sources, we are also considering performing a full-physics CO_2 retrieval, at a coarser spatial resolution due to computational constraints. The retrieved CO_2 could be used in place of the prior in the higher spatial resolution CO_2 proxy product.

S2.3 Impact of MethaneSAT vs. MethaneAIR spatial resolution differences

S2.3.1 Cloud contamination

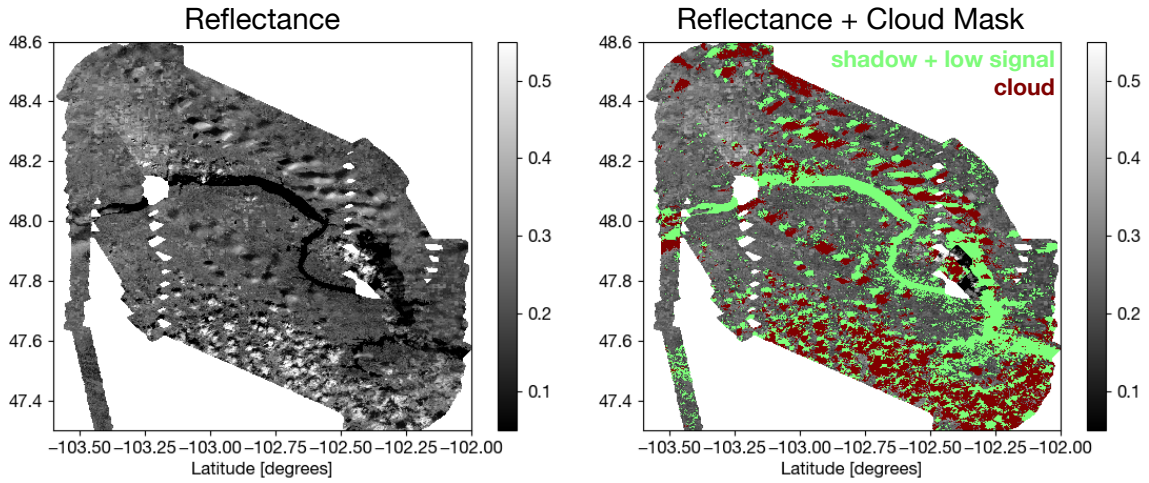

Figure S8: Example scene (Bakken Formation, RF09) used to assess the dependence of cloud contamination on instrument spatial resolution. The left panel shows the greyscale reflectance at 1622.5 nm. The right panel overlays the cloud (red) and cloud shadow/low signal (green) masks from the MethaneAIR cloud detection algorithm.

Instruments with finer spatial resolution can better resolve areas between clouds, yielding a greater fraction of useable data. To assess the reduction in coverage due to the coarser resolution of the satellite we use observations from RF09 (Figure S8). The flight was conducted over the Bakken formation, at a time when there was significant low-altitude cumulus clouds present, clearly seen from the reflectance map in Figure S8. This provides an ideal test case to assess resolution differences. Figure S8 (right panel) shows the corresponding cloud and low-signal masks for the same scene. From these, approximately 67% of the observed area is cloud-free.

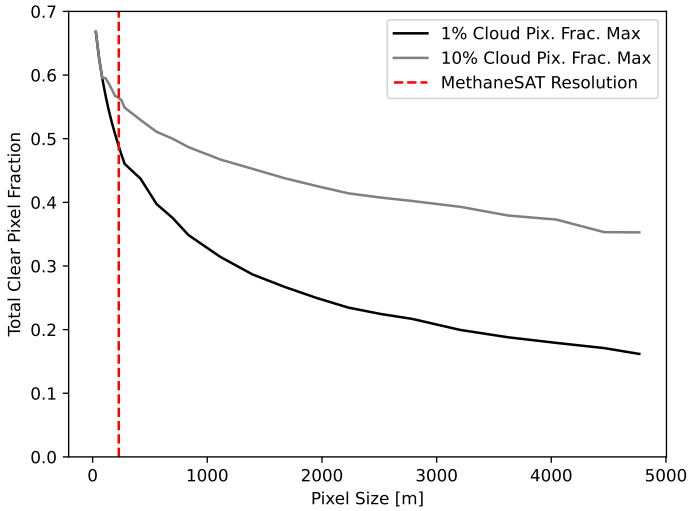

Figure S9: Fraction of clear pixels as a function of spatial pixel size, using the cloud-flagged data from RF09. See text for method details. The pixel size used in the x-axis is the length of the equivalent square pixel corresponding to the aggregate pixel used in the calculation.

To assess the effect of spatial resolution, we aggregate pixels in square $n \times n$ blocks, and for these aggregates, compute the fraction of the aggregate that is covered by clouds. We define a new cloud mask for the aggregate pixels based on a threshold cloud fraction. Figure S9 shows these at two cloud fraction maximums (1% and 10%) as a function of pixel size. At MethaneAIR spatial resolution, an additional 18% and 10% of pixels are lost to cloud contamination for the 1% and 10% thresholds, respectively. At the largest pixel size ($\sim 5 \times 5$ km^2,

approaching the spatial resolution of TROPOMI), significantly more pixels are lost (51% and 32% respectively). For MethaneSAT, the amount of pixels lost to cloud contamination will be further mitigated by its flexible target selection strategy, which factors forecasted cloud cover into its decision tree (Benmergui, 2019).

S2.3.2 Inhomogeneous Illumination

The ISRF used in the MethaneAIR retrievals is derived from calibration measurements whereby the instrument slit is homogeneously illuminated. Large sub-pixel differences in radiance can induce changes to the ISRF and shifts in the wavelength-to-pixel mapping (Noël et al., 2012; Landgraf et al., 2016). Not accounting for these ISRF changes typically yields changes in retrieved columns of up to a few percent (e.g. TROPOMI CO, Landgraf et al. (2016); GeoCarb CO_2, CH_4, and CO, Nivitanont et al. (2019)).

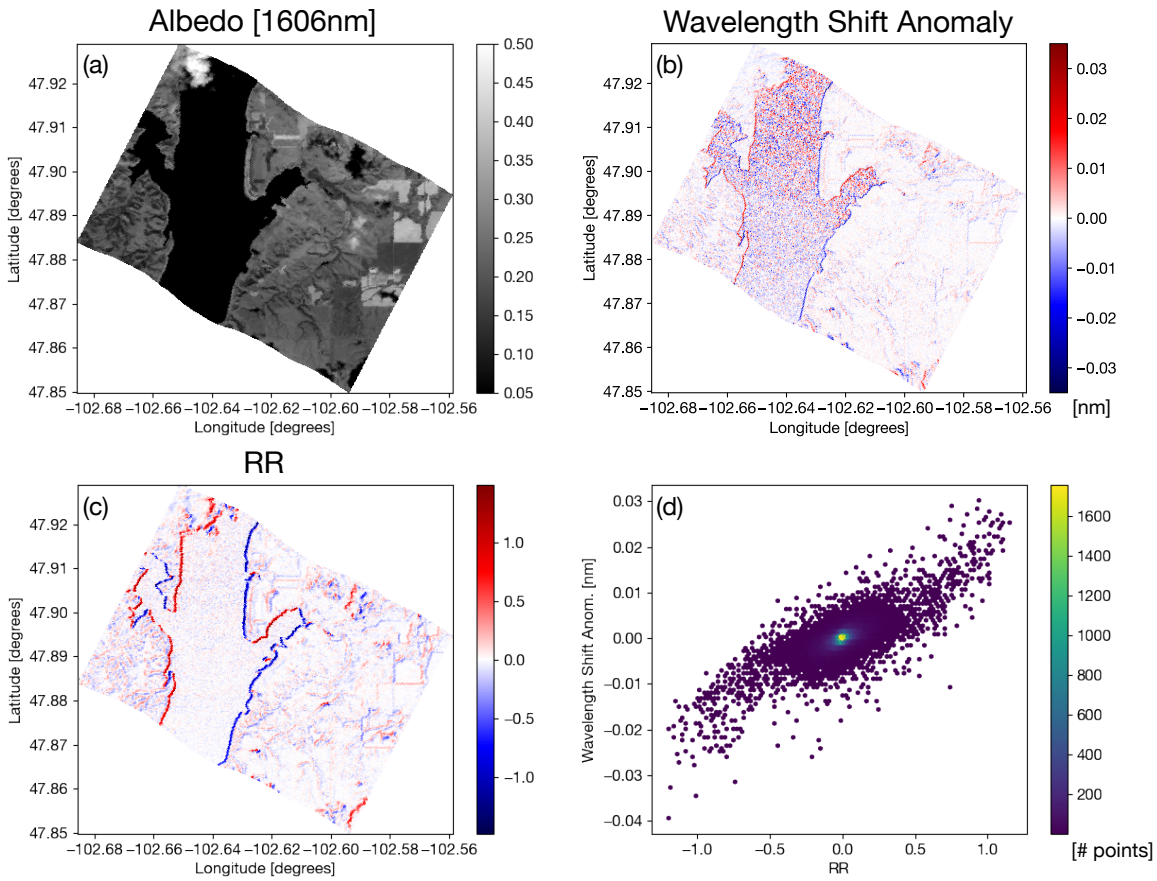

Figure S10: Example scene from RF09 demonstrating the effect of inhomogeneous illumination on the retrieved wavelength grid. (a) MethaneAIR retrieved albedo at 1606nm. (b) Retrieved wavelength shift (with median shift subtracted) (c) Reflectance ratio (RR) (see text) (d) relationship between RR and retrieved wavelength shift, observations filtered for CH_4 DoFS > 1.

In the 1.6 micron band, the two main situations yielding significant inhomogeneous illumination are the edges of clouds and water bodies, again making RF09 a useful dataset for this purpose. Figure S10 shows an example of this from a scene in RF09. The river in the image (Figure S10(a)) induces a retrievable wavelength shift at its edges (Figure S10(b)), clearly demonstrating that inhomogeneous illumination is occuring in the data.

The inhomogeneity can be quantified by dividing each pixel in the along-track direction into a lower and upper half (indexed l and u respectively). The pixel is divided in the along-track direction because this is the direction that strongly impacts the ISRF. A reflectance ratio metric (RR) can be used to predict these ISRF changes.

$$RR = \frac{A_l - A_u}{0.5(A_l + A_u)} \tag{S44}$$

A_l and A_u are the albedos for the lower and upper halves of the pixel respectively. Figure S10(c) shows RR estimated from the example scene. Since the scene is at the retrieved resolution, we have estimated RR via

linearly interpolating the albedos between pixels in the along track direction. Figure S10(d) shows that there is a strong relationship between the retrieved wavelength shift and RR, indicating that it is a reasonable metric for indicating ISRF changes from inhomogeneous illumination.

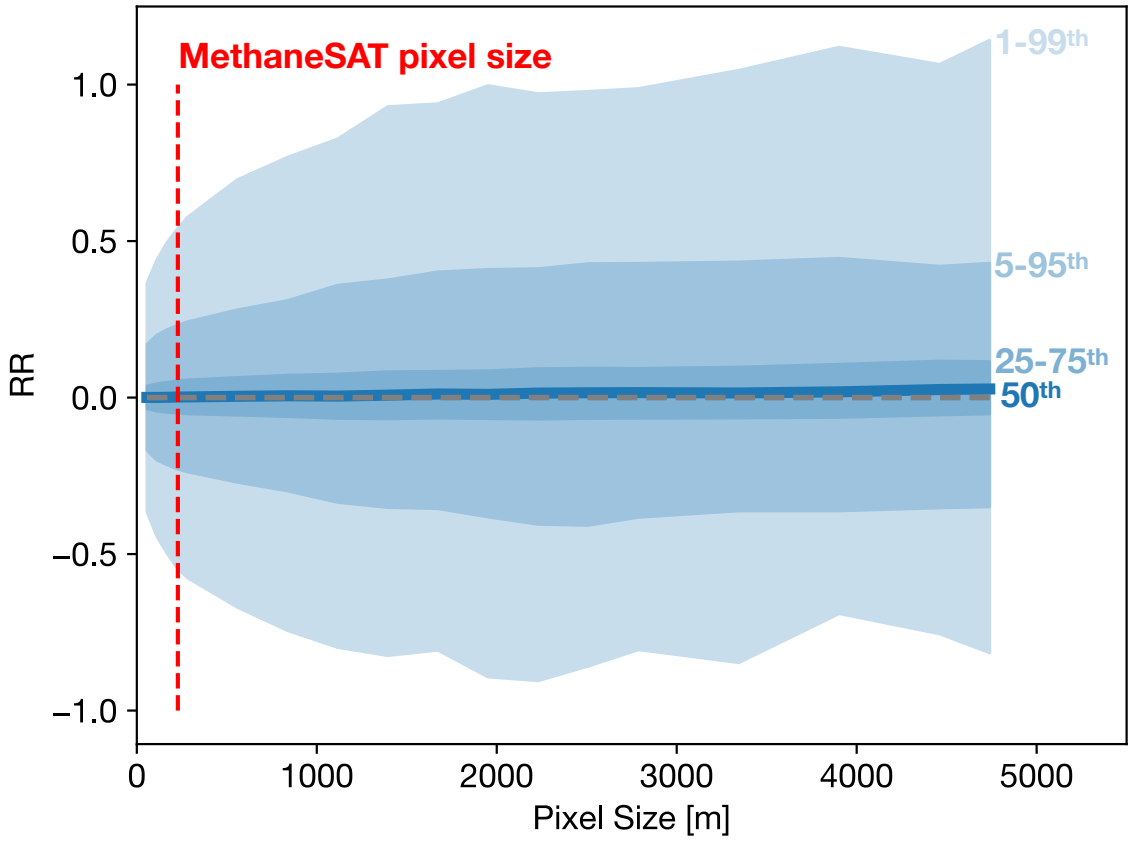

Figure S11: Distributions of RR as a function of spatial resolution computed from reflectance data from RF09.

We now use RR to assess how inhomogeneous illumination depends on pixel size. Figure S12 shows the distribution of RR for different pixel resolutions using the same RF09 dataset from Section S2.3.1. These are computed from aggregate pixels in increasing $n \times n$ blocks in the same manner as Section S2.3.1. At the size of a MethaneSAT pixel, distributions of RR are similar to MethaneAIR, indicating that has a similar susceptibility to its effects. At larger pixel sizes the spread of the RR increases in a quasi-symmetrical manner, whilst the median value is close to 0. This suggests that overall the errors associated with inhomogeneous illumination will be quasi-random. However spatially systematic biases will be correlated, as shown in Figure S10(b) where wavelength shifts on the east/west banks of the river are positive/negative respectively. Similar systematic effects on the retrieved CH_4 and CO_2 columns are also likely, as demonstrated by previous work (e.g. TROPOMI CO, Landgraf et al. (2016)), caused ultimately by changes in the ISRF. In MethaneAIR these are partially mitigating by the fitting of the ISRF squeeze factors. For MethaneSAT, these squeeze parameters may not be required due to improved thermal control. In this case deviations in retrieved wavelength shift can be used as an quality filter for pixels impacted by significant inhomogeneous illumination.

S2.3.3 Sub-pixel CH_4 gradients

Satellite retrievals assume that the concentration of the retrieved species is homogeneously distributed horizontally across the pixel being retrieved. As spatial resolution coarsens, heterogeneity in methane within a pixel becomes more pronounced. This will be especially pronounced in observations of plumes over point sources - In these cases, not accounting for the sub-pixel distribution of methane will yield a systematic underestimate in the retrieval, as "smearing" the actual enhancement over the entire pixel area effectively makes absorption more efficient. This can be argued simply from Beers law, where confining the CH_4 to a fraction of the pixel yields smaller overall radiance changes due to greater saturation in the absorption line depths.

To quantify this effect, we performed a simple OSSE where we confine a boundary layer methane enhancement to a fraction of the total pixel. We use a GGG2020 a priori profile sampled from RF04 over the Permian (20210730, 15:40:35) for the background conditions. We simulate two radiances, corresponding to the background, and one with the enhancement, confined in the lowest four model layers (681-918 hPa). A retrieval is then performed on the the combined radiance, which is weighted based on the plume area fraction.

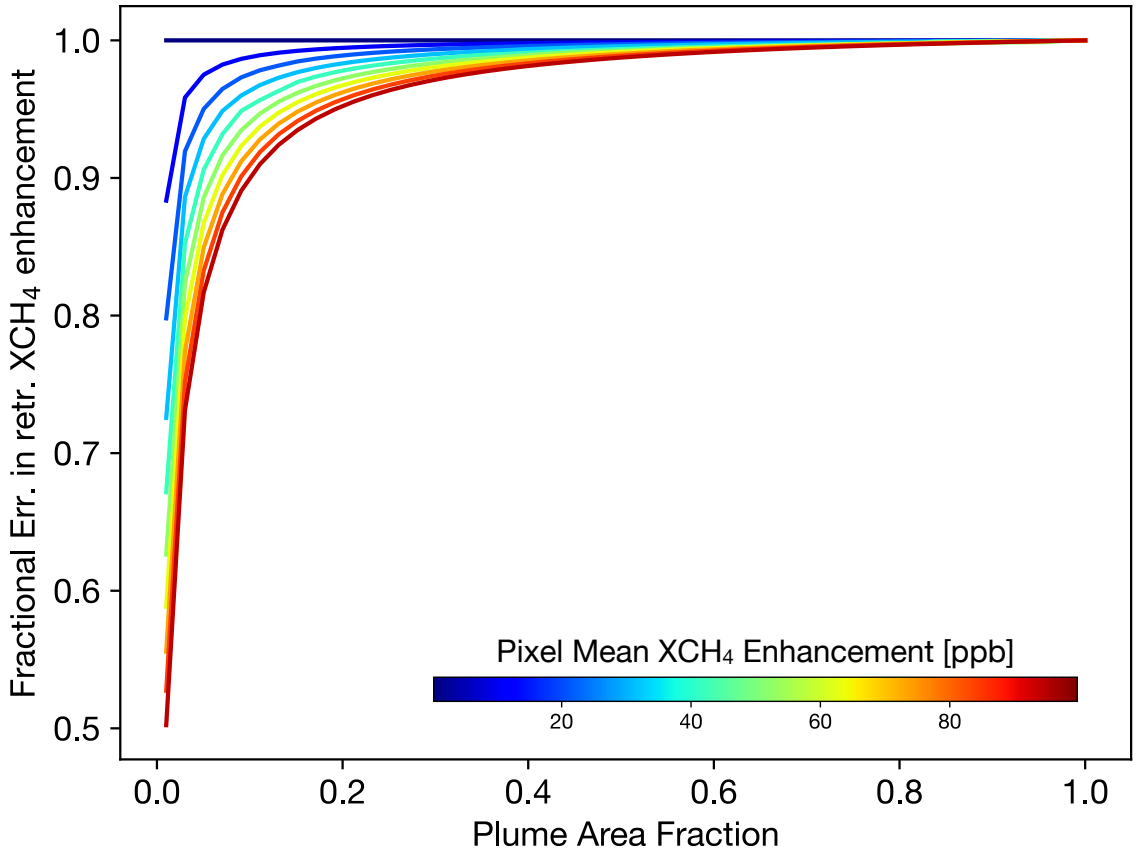

Figure S12: Fractional error in the retrieved XCH_4 enhancement, after errors in sub-pixel horizontal heterogeneity in CH_4 are accounted for (see text).

Figure S12 shows the fractional error in the retrieved XCH_4 enhancement as a function of the fractional area of the pixel covered by the plume. The color of the curves corresponds to the pixel-mean enhancement. The figure shows that the error increases sizeably when the plume only fills $\sim 0-15\%$ of the pixel. The error is more pronounced the larger the plume enhancement, with errors as large as 50% at plume area fractions of 1% for a 100 ppb pixel-average enhancement.

To create a simple estimate of the magnitude of these errors induced by sub-pixel heterogeneity on MethaneSAT, we use the values in Figure S12 with an example WRF-LES plume (Figure S13). We aggregate the data to the approximate distribution of MethaneSAT (140×400 m^2. Let $\delta_1, \ldots, \delta_n$ represent the XCH_4 values of the n WRF-LES pixels in an aggregate MethaneSAT pixel. We estimate the XCH_4 concentration within the plume as:

$$\delta_{plm} = \max(\delta_1, \ldots, \delta_n) \tag{S45}$$

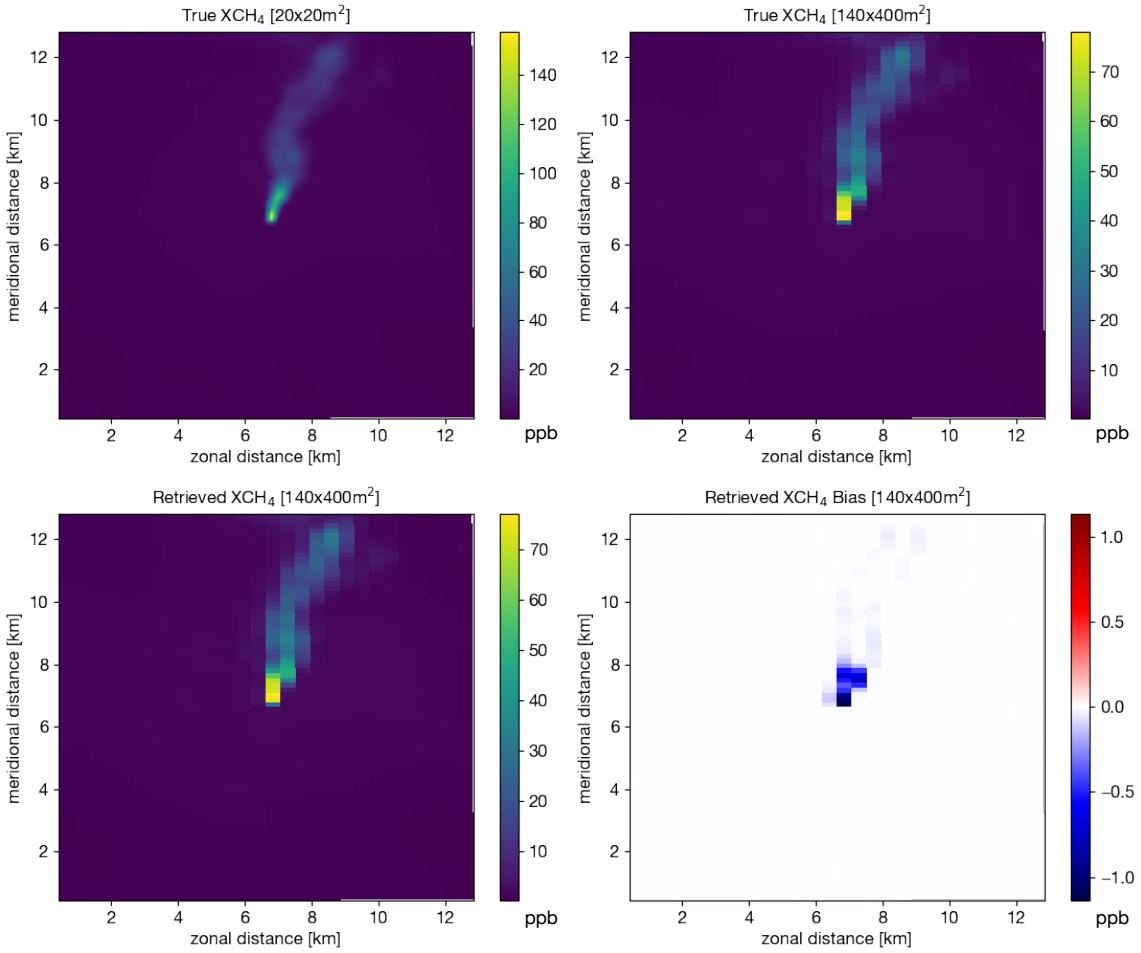

Figure S13: Example induced by sub-pixel XCH_4 heterogeneity estimated using WRF LES. (Top Right) The XCH_4 enhancement due to a 1000 kg/h plume simulated with WRF LES. (Top Left): The same plume aggregated to approximate MethaneSAT resolution. (Bottom Left): Retrieved plume factoring in error from subgrid XCH_4 variability (Bottom Left): Retrieval bias due to subgrid XCH_4 variability

The plume fraction (f_{plm}) is then

$$f_{plm} = \frac{n\delta_{plm}}{\sum_{i=1}^{n} \delta_i} \tag{S46}$$

In this manner, the plume is maximally concentrated in the smallest possible area of the pixel. This creates an upper bound for the estimate based on the curves in Figure S12. The resulting error is shown in Figure S13 (bottom right). The error peaks close to the source, but is only around 1 ppb, much smaller than the peak enhancement ($> 70ppb$). This suggests sub-pixel heterogeneity is not a huge issue at MethaneSAT native resolution.

References

Benmergui, J.: Innovations in Modeling and Statistical Analysis for Greenhouse Gas Flux Monitoring, undefined, URL https://www.semanticscholar.org/paper/Innovations-in-Modeling-and-Statistical-Analysis-Benmergui/ 2019.

Chan Miller, C., Liu, X., Abad, G. G., Zoogman, P., and Spurr, R.: Synthetic spectra generation for TEMPO, in: 5th TEMPO SCience Team Meeting, URL https://tempo.si.edu/presentations/May-June2017/10-cmiller$_t empo_s tm_2$0130531.*pdf*, 2013.

Connor, B. J., Boesch, H., Toon, G., Sen, B., Miller, C., and Crisp, D.: Orbiting Carbon Observatory: Inverse method and prospective error analysis, Journal of Geophysical Research: Atmospheres, 113, https://doi.org/10.1029/2006JD008336, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2006JD008336, 2008.

Crisp, D., O'Dell, C., Eldering, A., Fisher, B., Oyafuso, F., Payne, V., Drouin, B., Toon, G., Laughner, J., Somkuti, P., McGarragh, G., Merrelli, A., Nelson, R., Gunson, M., Frankenberg, C., Osterman, G., Boesch, H., Brown, L., Castano, R., Christi, M., Connor, B., McDuffie, J., Miller, C., Natraj, V., O'Brien, D., Polonski, I., Smyth, M., Thompson, D., and Granat, R.: Orbiting Carbon Observatory-2 3 (OCO-2 OCO-3) LEVEL2 FULL PHYSICS ALGORITHM Theoretical Basis Document Version 3.0, Rev 1, 2021.

GHGRP: U.S. EPA Greenhouse Gas Reporting Program (GHGRP), URL https://www.epa.gov/ghgreporting.

Hansen, P. C.: The L-curve and its use in the numerical treatment of inverse problems, 1993.

Landgraf, J., aan de Brugh, J., Scheepmaker, R., Borsdorff, T., Hu, H., Houweling, S., Butz, A., Aben, I., and Hasekamp, O.: Carbon monoxide total column retrievals from TROPOMI shortwave infrared measurements, Atmospheric Measurement Techniques, 9, 4955–4975, https://doi.org/10.5194/amt-9-4955-2016, 2016.

Laughner, J. L., Roche, S., Kiel, M., Toon, G. C., Wunch, D., Baier, B. C., Biraud, S., Chen, H., Kivi, R., Laemmel, T., McKain, K., Quéhé, P.-Y., Rousogenous, C., Stephens, B. B., Walker, K., and Wennberg, P. O.: A new algorithm to generate a priori trace gas profiles for the GGG2020 retrieval algorithm, Atmospheric Measurement Techniques, 16, 1121–1146, https://doi.org/10.5194/amt-16-1121-2023, publisher: Copernicus GmbH, 2023.

Nivitanont, J., Crowell, S., O'Dell, C., Burgh, E., McGarragh, G., O'Brien, D., and Moore, B.: Characterizing the Effects of Inhomogeneous Scene Illumination on the Retrieval of Greenhouse Gases from a Geostationary Platform, ESS Open Archive, https://doi.org/10.1002/essoar.10501288.1, 2019.

Noël, S., Bramstedt, K., Bovensmann, H., Gerilowski, K., Burrows, J. P., Standfuss, C., Dufour, E., and Veihelmann, B.: Quantification and mitigation of the impact of scene inhomogeneity on Sentinel-4 UVN UV-VIS retrievals, Atmospheric Measurement Techniques, 5, 1319–1331, https://doi.org/10.5194/amt-5-1319-2012, 2012.

Parker, R. J., Webb, A., Boesch, H., Somkuti, P., Barrio Guillo, R., Di Noia, A., Kalaitzi, N., Anand, J. S., Bergamaschi, P., Chevallier, F., Palmer, P. I., Feng, L., Deutscher, N. M., Feist, D. G., Griffith, D. W. T., Hase, F., Kivi, R., Morino, I., Notholt, J., Oh, Y.-S., Ohyama, H., Petri, C., Pollard, D. F., Roehl, C., Sha, M. K., Shiomi, K., Strong, K., Sussmann, R., Té, Y., Velazco, V. A., Warneke, T., Wennberg, P. O., and Wunch, D.: A decade of GOSAT Proxy satellite CH_4 observations, Earth System Science Data, 12, 3383–3412, https://doi.org/10.5194/essd-12-3383-2020, publisher: Copernicus GmbH, 2020.

Plant, G., Kort, E. A., Brandt, A. R., Chen, Y., Fordice, G., Negron, A. M. G., Schwietzke, S., Smith, M., and Zavala-Araiza, D.: Inefficient and unlit natural gas flares both emit large quantities of methane, Science, 377, 1566–1571, https://doi.org/10.1126/science.abq0385, 2022.

Rodgers, C. D.: Inverse Methods for Atmospheric Sounding, WORLD SCIENTIFIC, https://doi.org/10.1142/3171, 2000.

Staebell, C., Sun, K., Samra, J., Franklin, J., Chan Miller, C., Liu, X., Conway, E., Chance, K., Milligan, S., and Wofsy, S.: Spectral calibration of the MethaneAIR instrument, Atmospheric Measurement Techniques, 14, 3737–3753, https://doi.org/10.5194/amt-14-3737-2021, publisher: Copernicus GmbH, 2021.