Supplementary Material for: *First Methane Retrievals from MethaneSAT: Demonstrating Sensor Performance for Constraining Regional Methane Emissions*

Chan Miller et al.

June 9, 2026

S1 Column Averaging Kernels

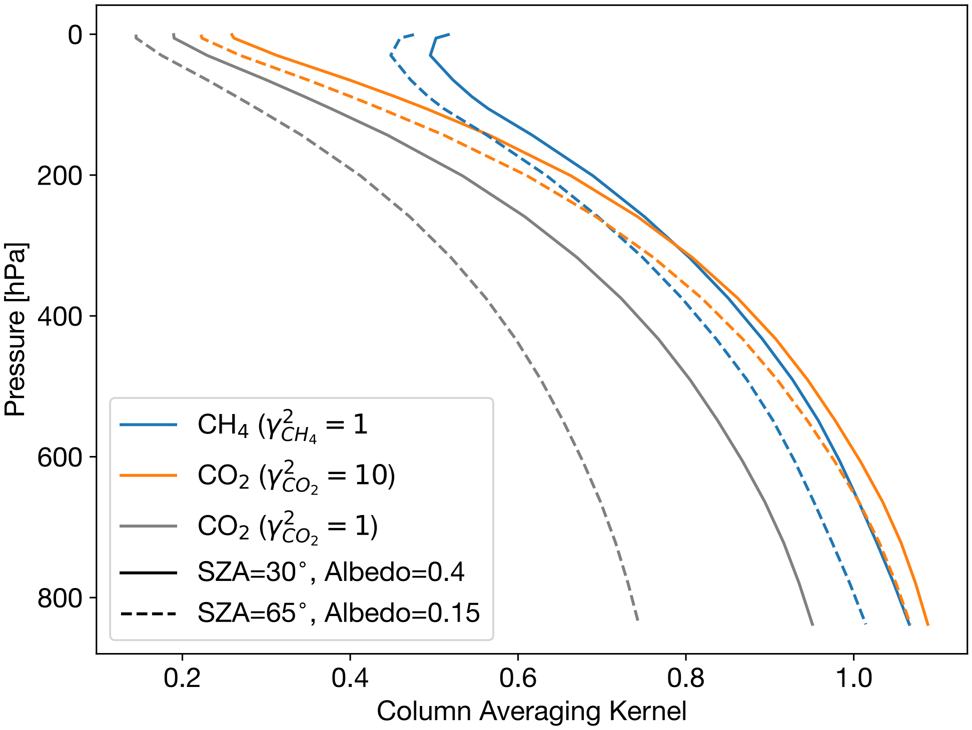

Figure S1: Comparison of MethaneSAT CH_4 and CO_2 column averaging kernels from the linear sensitivity analysis (Section 3.1, main text) before ($\gamma^2_{\mathrm{CO_2}} = 1$) and after ($\gamma^2_{\mathrm{CO_2}} = 10$) scaling of the CO_2 prior covariance matrix. Solid lines correspond to a typical bare-ground scene (SZA= 30°, albedo= 0.4) and dashed lines represent the lowest-signal scene targeted by the satellite (SZA= 65°, albedo= 0.15).

S2 Stripe Correction

S2.1 PLS Model Tuning

We used PLS regression to predict X_{CH_4} cross-track biases from retrieved ISRF squeeze factors. PLS projects predictors and responses onto latent components that maximize covariance, which is advantageous for handling collinearity among predictors. The number of components was initially assessed using 5-fold cross-validation to minimize prediction error. As shown in Figure S2, the mean squared error (MSE) falls from approximately 57 ppb^2 with one component to below 40 ppb^2 by 15 components, but the improvement levels off after about nine components. This behavior is common in small datasets ($n = 54$) because each validation fold contains few samples, making error estimates noisy and less reliable. With many components relative to sample size, the model can increasingly fit random variation in the training data, which appears as improved performance in cross-validation but does not generalize, indicating overfitting.

To mitigate this, we implemented a permutation test (200 iterations) that randomly permuted X_{CH_4} cross-track bias response values to generate a chance distribution of model performance. This test provides a statistical baseline, ensuring that the observed improvement is not due to random association. Importantly, we based component selection on incremental improvement rather than absolute error. This approach is favorable when cross-validation error decreases monotonically, as it identifies whether each additional component provides a meaningful gain beyond what is already achieved. By comparing incremental improvements against the chance model, we avoid adding components that capture random variation rather than true signal. Figure S3 illustrates the resulting p-values: the blue curve corresponds to p-values computed from the RMSE between true and permuted datasets, while the orange curve shows p-values based on incremental improvements from adding successive components. Comparing the absolute improvements (blue curve), p-values are always below 0.01, suggesting the overall model is significant. However, the incremental improvement p-values (orange curve) increase sharply after nine components, indicating diminishing returns. We therefore selected nine components as the optimal balance between performance and complexity ($p < 0.01$).

S2.2 Wavelet Model Tuning

Here we seek to tune the parameters used in the wavelet step of the destriping algorithm detailed in the main paper. Let $\mathcal{W}$ represent the operator for the wavelet destriping step, and $f(x,t)$ represent the X_{CH_4} field with stripe errors, where x and t indicate the along- and across-track pixel indices. The destriped field $\hat{f}(x,t)$ is thus:

$$\hat{f}(x,t) = \mathcal{W}(f; L, \sigma_{f_n}) \tag{1}$$

We seek to tune the parameters L, the number of wavelet decomposition levels, and σ_n, the width of the gaussian mask used by the low-pass filter applied to the wavelet horizontal-detail coefficients. The first thing to note is that $\mathcal{W}$ is linear. If we divide the overall image

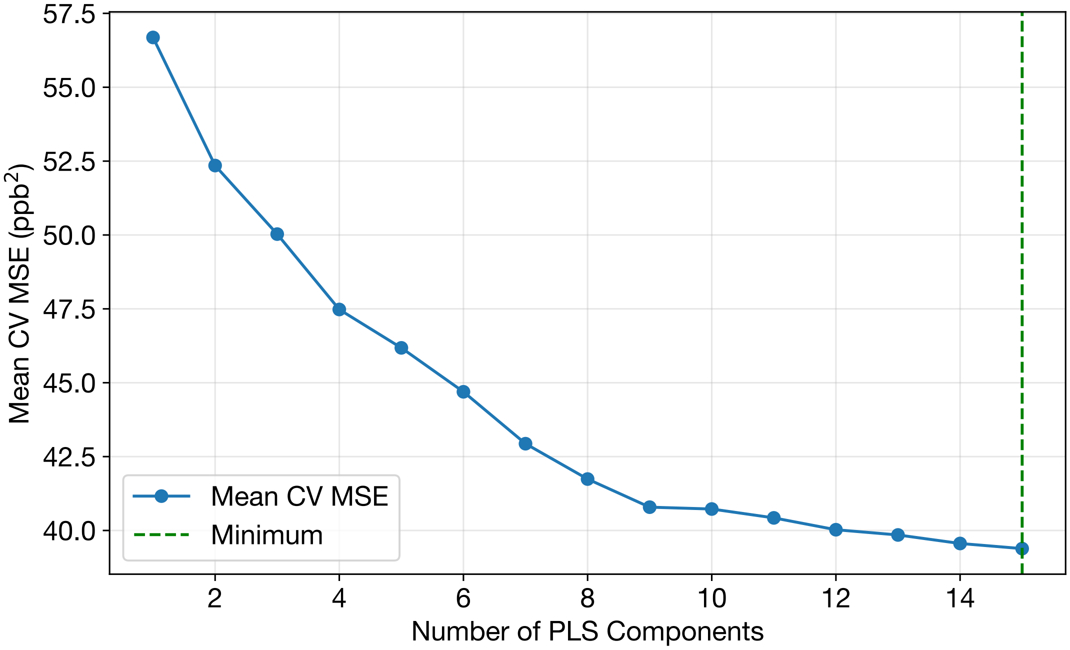

Figure S2: Cross-validation MSE averaged over all folds versus number of PLS components. Error decreases from ∼57 ppb^2 to <40 ppb^2 by 15 components, with little improvement beyond 9 components.

into a stripe ($b(x,t)$), signal ($s(x,t)$), and random noise ($r(x,t)$) then we can analyze the effects of the filter on these components separately.

$$\hat{f}(x,t) = \hat{b}(x,t) + \hat{s}(x,t) + \hat{r}(x,t) = \mathcal{W}(s; L, \sigma_{f_n}) + \mathcal{W}(b; L, \sigma_{f_n}) + \mathcal{W}(r; L, \sigma_{f_n}) \tag{2}$$

First let us consider the most optimistic stripe case, where the stripes are constant in the along track direction. Thus for some stripe pattern $\bar{b}(x)$

$$b(x,t) = \bar{b}(x) \tag{3}$$

Since the stripes do not vary in the along-track direction, $\hat{b}(x,t)$ does not depend on σ_{f_n}. Therefore L controls the degree to which a constant stripe pattern will be suppressed. To tune L we run 200 Monte Carlo simulations where $\bar{b}(x)$ is drawn from a normal distribution with a standard deviation of 6.5 ppb; This value corresponds to the 95$^{\text{th}}$ percentile value of the remaining residual stripe noise after the PLS correction (Figure 8, main text).

Figure S4 shows the standard deviation of the corrected stripe pattern $\hat{b}(x,t)$ as a function of L. We select $L = 7$; This corresponds to the level where residual stripe variation reduces below 1 ppb for 95% of cases, and beyond this there is not much further improvement. This value was also chosen in the TROPOMI/WFMD XCO/XCH$_4$ retrieval (Schneising et al., 2023), from which the method is based.

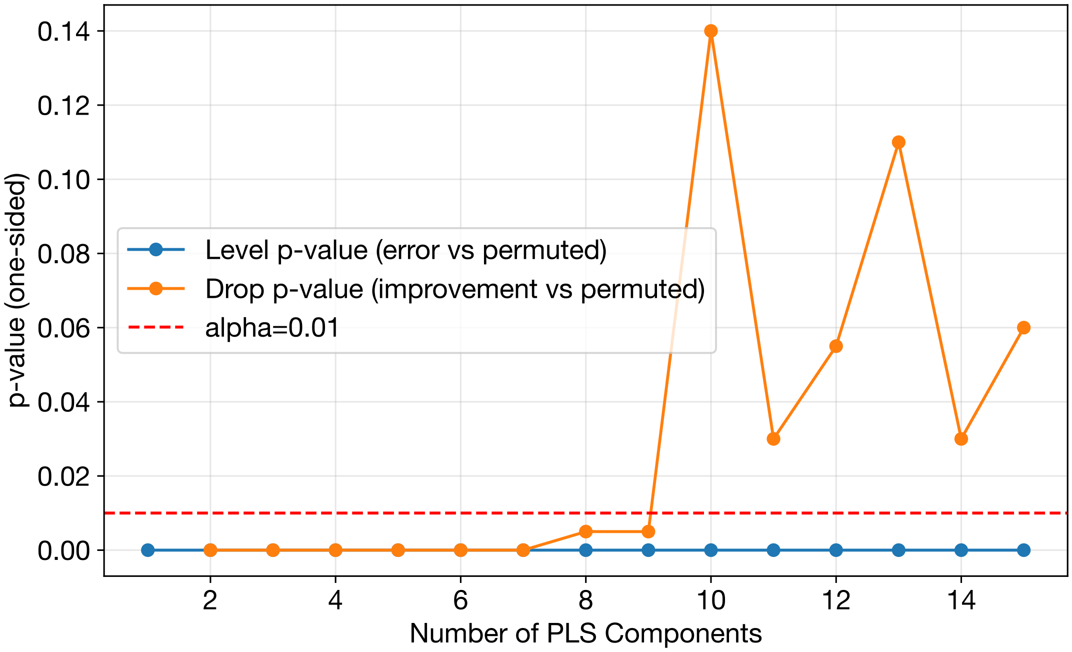

Figure S3: Permutation test p-values for model significance. Absolute improvements (blue) remain <0.01 for all components, while incremental improvements (orange) rise sharply after 9 components

In reality the stripe pattern can vary along-track. The degree to which the filtered stripes can vary is controlled by σ_{f_n}, with larger values allowing more variation. This comes at a cost of dampening real signal, especially if that signal aligns along-track.

To illustrate this we show the impact of the wavelet filter on a gaussian plume (Figure S5 for three cases: 90°, 135°, and 180°, angles ranging from where the plume is fully aligned to perpendicular to the along-track direction. The plume is simulated in the center of the scene following Nassar et al. (2021), with a 300 kg h^{-1} emission rate and 5 m s^{-1} wind speed. This represents a pernicious edge case, as the gaussian plume assumption combined with high wind leads to a geometry very similar to a real stripe. As expected, error is largest when the wind is perfectly aligned with the along-track direction (90°), creating a wave-pattern in the across-track direction, rippling from a large negative anomaly aligned in the along-track direction. The errors are much smaller as the plume moves out of along-track alignment.

We use the rotating gaussian plume to tune σ_{f_n}. Figure S6 shows the root mean square error (RMSE) of the destriped field for various σ_{f_n} as a function of wind rotation angle. The inset panel shows RMSE as a function of σ_{f_n} for the fully along-track aligned plume. The RMSE tends from a lower limit of 1 ppb RMSE to an upper limit of 3.5 ppb RMSE, corresponding to σ_{f_n} values that preserve all along track frequencies ($f_n \neq 0$) to removing all. The region between σ_{f_n} of 10^{-2} to 10 cycles/pixel corresponds to cases between these extremes. For $\sigma_{f_n} = 10^{-1}$, the error for the perfectly along-track aligned wind doubles, but produces similar errors to the most conservative case after a 10° rotation.

The optimal value of σ_{f_n} must also consider the true along-track variation in stripes, from

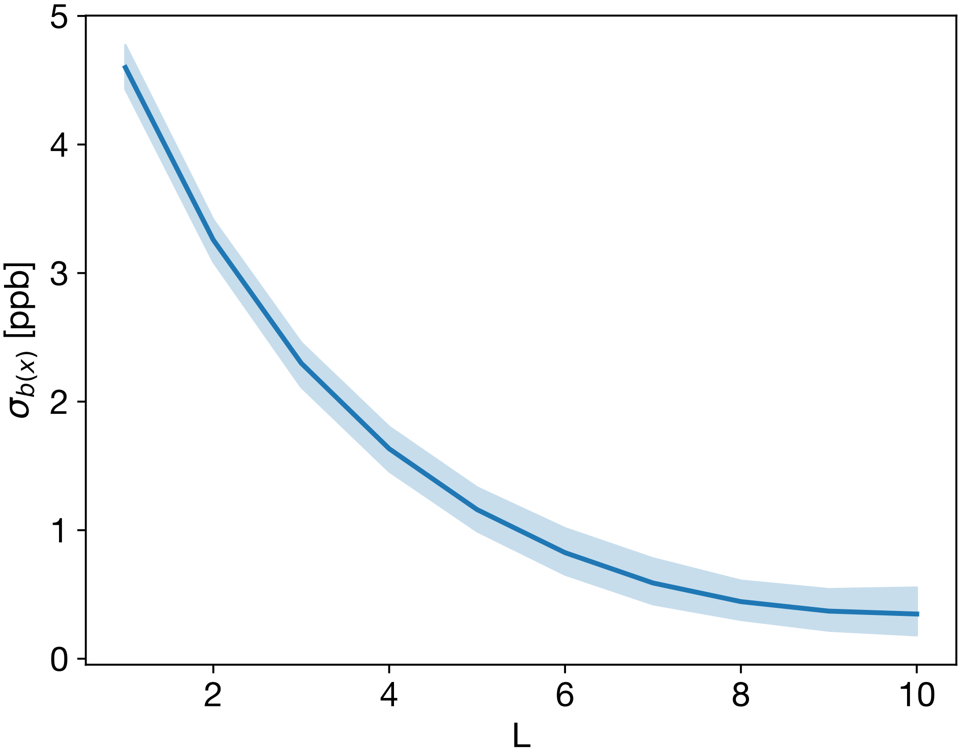

Figure S4: Dependence of remaining stripe noise as a function of L for constant stripes computed from 200 Monte Carlo simulations for a normally distributed stripe pattern with standard deviation of 6.5 ppb. The solid blue line corresponds to the mean of the trials, and shaded light blue region the 5^{th} to 95^{th} percentile values.

which we can only assess using the actual data. Figure S7 shows results from the wavelet filter at four levels of smoothing, spanning the filter limits discussed above. The Permian scene displayed here (Collection ID:c01D50380 , 29 September 2024) is chosen because it contains point sources in the lower-middle portion of the image aligned along-track, testing the worst-case of the filter. Since the native pixel data is dominated by noise, we also show results aggregated to $\sim 2 \times 2\text{km}^2$, the aggregation level used by the emissions inversion algorithm. At low levels of attenuation ($\sigma_{f_n} = 10^{-3}$ cycles/pixel) there is a clear stripe present, as indicated in Figure S7. This feature is removed by $\sigma_{f_n} = 5 \times 10^{-2}$ cycles/pixel. At this level the plumes at the bottom of the image are still preserved, but are removed at high attenuation levels ($\sigma_{f_n} = 10$ cycles/pixel). We therefore choose $\sigma_{f_n} = 5 \times 10^{-2}$, as there appears to be a good balance between stripe removal and signal preservation. We note that for the gaussian plume case this would yield an increase in error of 50% for the perfectly aligned plume. In practice, this is unlikely to ever be the case, as real plume structures are unlikely to be perfectly aligned, and structures formed by turbulent eddies likely lessen the error.

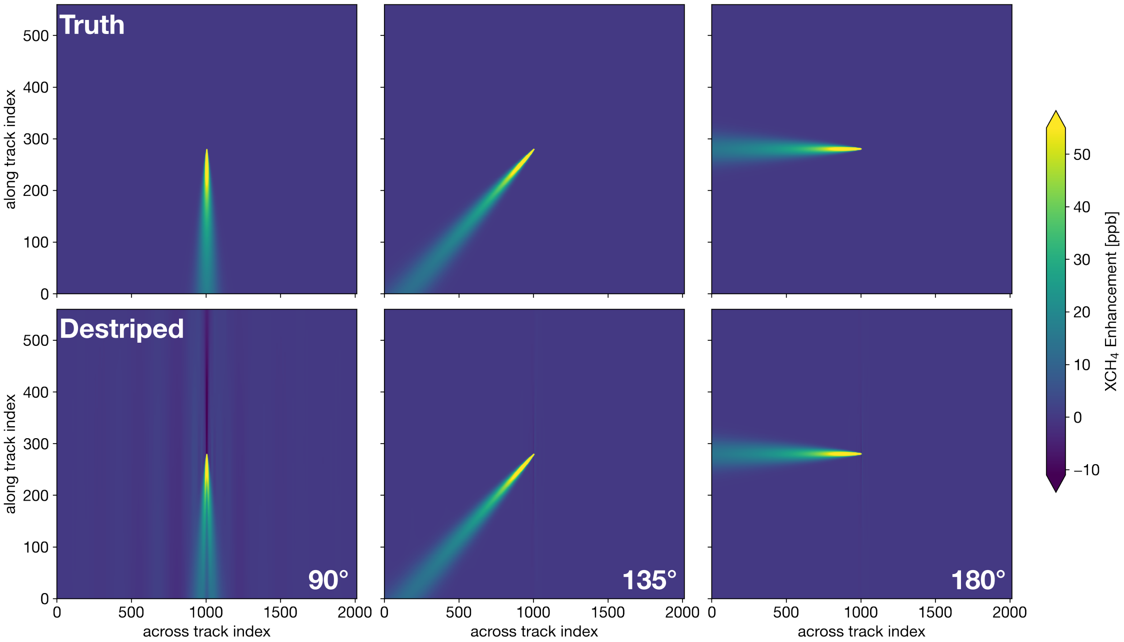

Figure S5: Distortion induced by a wavelet filter on a Gaussian plume. The top row shows the gaussian plume for a 300 kg s^{-1} point source transported by a 5 ms^{-1} wind. The wind direction angle is defined clockwise relative to the across-track direction. The bottom row shows the resulting plume after application of the wavelet filter.

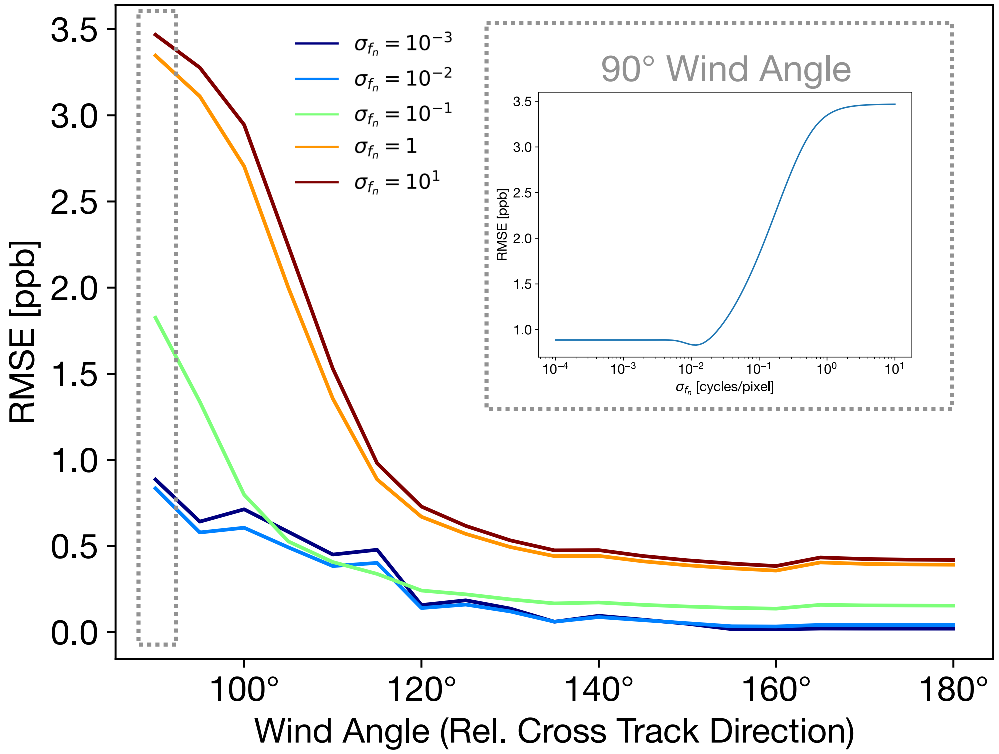

Figure S6: Root mean square error (RMSE) induced by the wavelet filter on a gaussian plume (300 kg s^{-1} emission, 5 ms^{-1} wind) for different Fourier attenuation coefficient (σ_{f_n}) shown as a function of wind direction. The inset panel shows the RMSE as a function of σ_{f_n} for the worst case scenario, in which a wind aligned in the along-track direction.

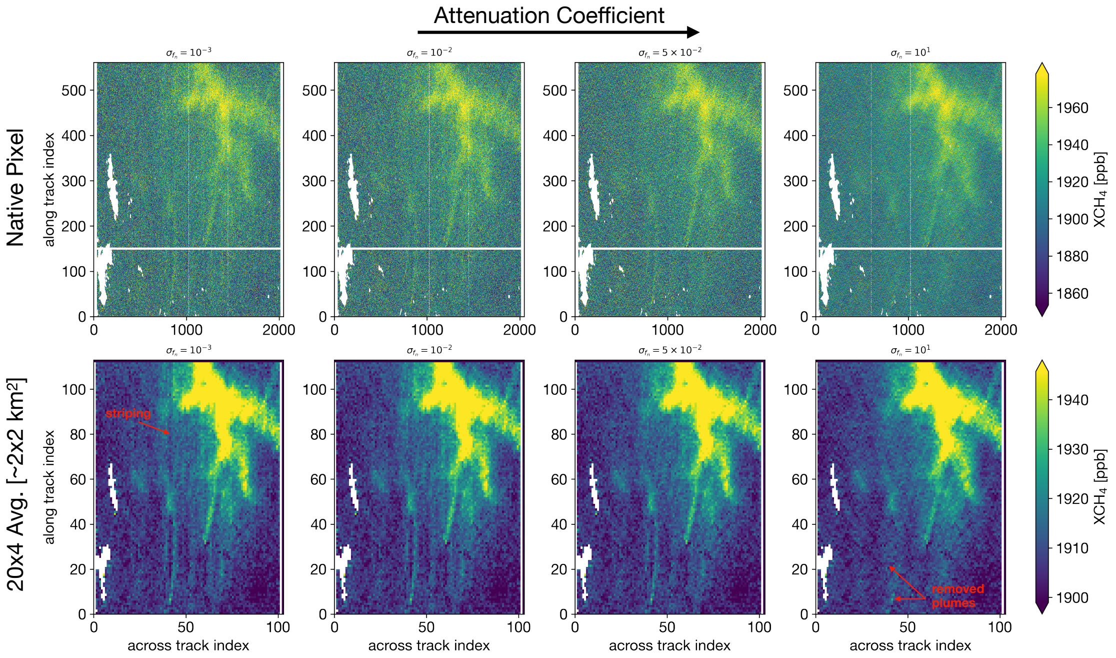

Figure S7: Sensitivity of wavelet-destriped X_{CH_4} to attenuation coefficient (σ_{f_n} for a real scene over the Permian (Collection ID:c01D50380 , 29 September 2024). The top and bottom rows show X_{CH_4} with pixels at native and 20×4 across×along track averages (yielding an aggregate pixel size of $\sim 2 \times 2\mathrm{km}^2$), respectively. σ_{f_n} increases in each panel moving from left to right, corresponding to increased stripe attenuation. The location of a prominant stripe that remains unfiltered at lower attenuation levels is marked int he bottom left panel. The location of two plumes that are inadvertently removed by the stripe filter at higher attenuation levels is marked in the bottom right panel.

References

Nassar, R., Mastrogiacomo, J.-P., Bateman-Hemphill, W., McCracken, C., MacDonald, C. G., Hill, T., O'Dell, C. W., Kiel, M., and Crisp, D.: Advances in quantifying power plant CO2 emissions with OCO-2, Remote Sensing of Environment, 264, 112 579, https://doi.org/10.1016/j.rse.2021.112579, 2021.

Schneising, O., Buchwitz, M., Hachmeister, J., Vanselow, S., Reuter, M., Buschmann, M., Bovensmann, H., and Burrows, J. P.: Advances in retrieving XCH_4 and XCO from Sentinel-5 Precursor: improvements in the scientific TROPOMI/WFMD algorithm, Atmospheric Measurement Techniques, 16, 669–694, https://doi.org/10.5194/amt-16-669-2023, 2023.