Algorithm Theoretical Basis Document

L3/L4: Methane and Carbon Dioxide Emission Quantification for Satellites

Version 3.0.1
Release date February 17th, 2026

Carbon Mapper, Inc.
info@carbonmapper.org

TABLE OF CONTENTS

1 Background

Methane (CH4) and Carbon Dioxide (CO2) are greenhouse gases that have primarily driven anthropogenic warming since the pre-industrial era. High emission CH4 and CO2 point sources make up a disproportionate amount of the anthropogenic budget. Carbon Mapper's mission is to detect, quantify, and publish these sources using airborne and satellite remote sensing platforms. Carbon Mapper supports policymakers and stakeholders by providing decision support tools and analyses that synthesize satellite and airborne remote sensing data into actionable insights.

The Carbon Mapper data platform is a full-scale operational implementation of a science data system that builds on 10+ years of research and development projects led by Carbon Mapper team members, initially at NASA's Jet Propulsion Laboratory supported by funding from NASA, the California Air Resources Board, and the University of Arizona. Those research projects included multiple airborne field campaigns, satellite and surface observations, and development of CH4 retrieval algorithms, machine learning tools, multi-scale analytic frameworks, data pipelines, open data portals and synthesis analysis.

The Carbon Mapper data platform is designed to rapidly process and publish point-source CH4 and CO2 data from multiple satellite and airborne imaging spectrometers. The platform has been routinely processing data from airborne surveys using NASA JPL's AVIRIS-NG and the Arizona State University Global Airborne Observatory since 2022 and expanded in early 2023 to include observations from NASA's EMIT mission on the International Space Station. In 2024, the platform began operational processing of Planet's first Tanager satellite as well as NASA JPL's AVIRIS-3 airborne imaging spectrometer. The platform is anticipated to continue to process data from additional Tanager launches and potentially data from other imaging spectrometer remote sensing observations.

Carbon Mapper is dedicated to providing CH4 and CO2 data that is transparent, trusted and actionable. Here we provide an overview of our methods and procedures to quantify CH4 and CO2 plumes, along with relevant quality control (QC) review. Other Carbon Mapper Coalition documentation will describe the theoretical basis for other key retrieval and detection processes.

2 Overview of data products and data processing

Figure 1 below lists Carbon Mapper Coalition data products and processing levels. In brief, the L2b data product is an estimate of column CH4 concentrations that are derived from L1b top of the atmospheric calibrated radiance using CH4 and CO2 absorption features at shortwave infrared (SWIR) wavelengths. L2b data then undergo plume detection and attribution procedure (L2c). Each plume is individually

segmented (L3) and emission rates are quantified (L4). In this document we describe the quantification processes (L3 and L4 products).

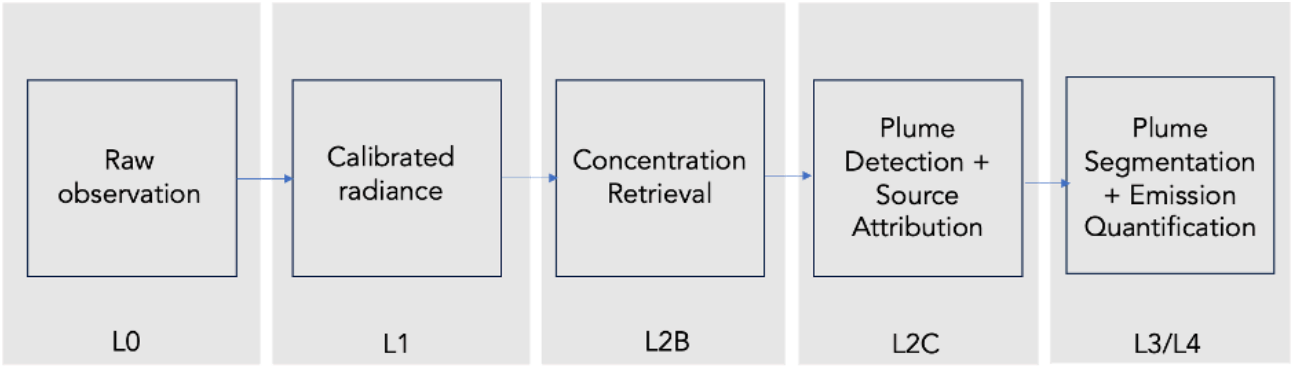

Figure 1: Simplified data flow indicating the Carbon Mapper data processing pipeline and product levels.

Instrument specifications for satellites that Carbon Mapper routinely processes for CH4 and CO2, Carbon Mapper Coalition's Tanager and NASA Earth Mineral Dust Investigation (EMIT), are listed in Table 1.

Table 1: Instrument specifications for satellites that Carbon Mapper routinely processes for CH4 and CO2.

Instrument Name	**Carbon Mapper Coalition Tanager**	**NASA EMIT** †
Swath width	18.6-24.2 km (varies with look angle)	75 km
Off-nadir pointing ability ("look angle")	30 degrees	None
Ground Sample Distance (GSD)	30-43 meters (varies with look angle)	60 m
Spectral response (FWHM)	5.5 nm	8.5 nm
Spectral sampling	5 nm	7.5 nm
Spectral range	400 - 2500 nm	381-2493 nm
Signal-to-noise @ 2200 nm	310 – 655 (varies with imaging mode)*	450

*35 deg Solar Zenith Angle, 25% albedo

† Values taken or extrapolated from Thompson et al., 2024

3 Plume segmentation

Carbon Mapper implements an automated plume segmentation and delineation process on plumes identified and geolocated during the L2c data processing step. The L2c processing step output is an origin location (latitude, longitude) for a plume given a unique L2b concentration map (units in column enhancement CH4 or CO2 (units ppm-m)). The L3 process then segments a plume around this origin point to create a masked plume boundary that is used for mass and emission quantification. The segmentation algorithm proceeds as follows (visual example in Figure 2):

1. The L2b concentration map is cropped radially around the origin of a plume using a 2500 m radius
2. A background concentration is determined for the plume to separate enhanced pixels from background pixels. This is done by subdividing the crop around the origin of the plume into 72 wedges (5° wedge). For each wedge (i), mean ($\boldsymbol{\mu}_i$) and standard deviation ($\boldsymbol{\sigma}_i$) are calculated from the the retrieved concentrations that pertain to that wedge. A global sigma value ($\boldsymbol{\sigma}_{all}$) is derived as the median of all 72 $\boldsymbol{\sigma}_i$. The candidate thresholds are parameterized as $\boldsymbol{\mu}_i$+k*$\boldsymbol{\sigma}_{all}$. For CH4 we let k=2 and for CO2 we let k=1, which we find produces the closest correspondence to truth datasets (controlled releases for CH4 and power plant continuous emissions monitoring (CEMS) data for CO2). Using these wedge-derived thresholds results in 72 candidate thresholds from which 72 candidate masks are eventually produced.
3. For each candidate mask, connected pixels of enhanced concentration above the estimated threshold (Step 2) are grouped together. A cluster must contain 3 pixels to be considered part of the plume. Lastly, a cluster closest to the origin, within 8-pixels, is selected as a candidate mask.
4. The probability of plume presence (p) at each pixel is estimated using 72 dynamic mask candidates, where p = N / 72, where N = number of candidate masks where that pixel was included in the plume mask.
5. We generate an initial mask by applying a relaxed threshold (0.3) and retaining the connected plume component nearest to the estimated origin. Any mask candidate not overlapping with the initial mask is excluded and probability is newly estimated using filtered mask candidates. We then refine the initial mask using a stricter threshold (0.7).

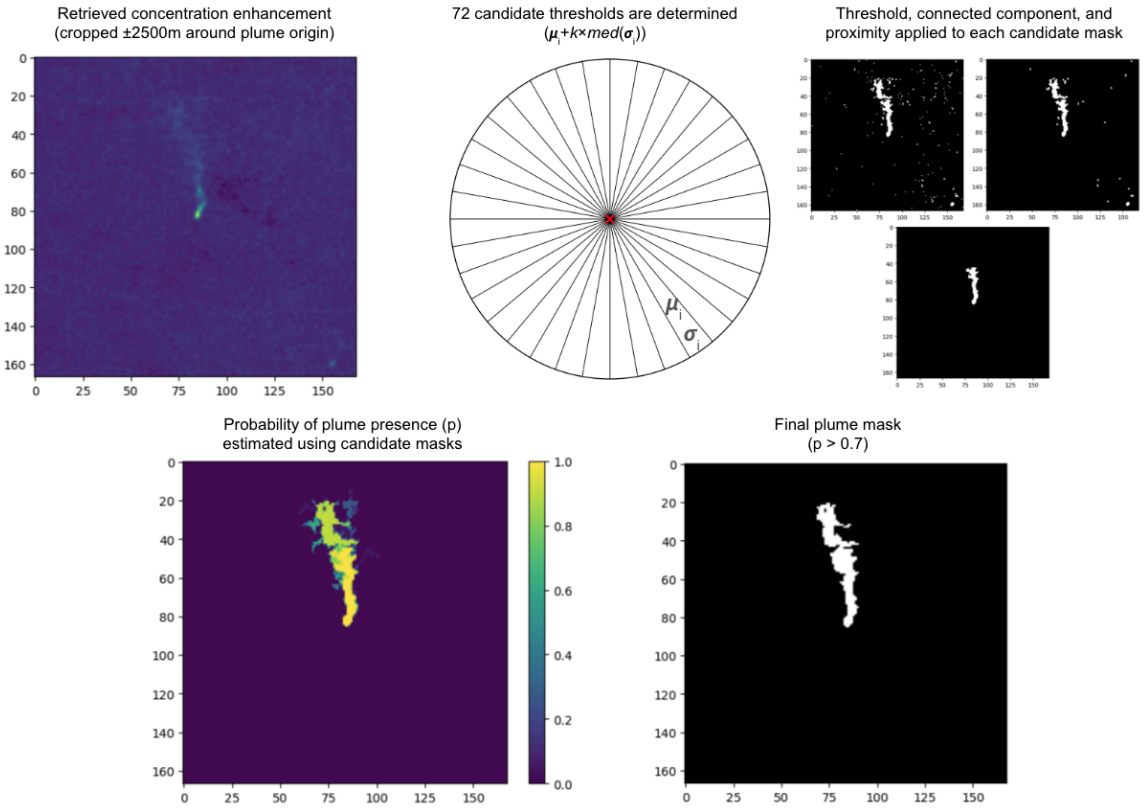

Figure 2. Visual example of plume segmentation processed applied to a CH4 plume detected from a Tanager CH4 retrieval.

During Step 2 we calculate candidate thresholds using mean concentration values calculated per wedge ("wedge-wise"), but opt to further refine these thresholds using globally calculated concentration standard deviations. The justification is as follows: in the presence of localized plumes, wedge-wise mean values vary less than wedge-wise standard deviation values that are prone to high fluctuations due to both retrieval artifacts and inconsistent plume extents per wedge. To prevent wedge-dependent variability in $\boldsymbol{\sigma}_i$ from propagating into the probabilistic plume mask, we use a global sigma value ($\boldsymbol{\sigma}_{all}$) for stable and consistent segmentation.

For Step 5, we adopt a two-step thresholding approach based on the probability of plume presence. Plume presence is estimated through application of the 72 candidate plume masks: the number of candidate masks (divided by 72) that encompass each pixel. We start by applying a low probability threshold (0.3), then identifying an appropriate plume cluster that is nearest the origin of the identified plume, and then further refine the plume mask boundary by applying a more strict probability threshold (0.7). We must start from a more expansive probability threshold to ensure sufficient pixels are captured. This is demonstrated visually in Figure 3: The top-left panel shows the probability of plume preference; no pixel exceeds the 0.7 probability threshold. This is due to the fact that several of the candidate masks only capture the plume cluster to the west of the origin of the plume (bottom-left panel), with the effect of diluting per-pixel mask probabilities generally. This is mitigated by applying a relaxed threshold of 0.3 and retaining only the connected region closest to the plume origin (top-right panel) as this step enables the exclusion of non-overlapping mask candidates. A refined

probability map is then recalculated using only the retained candidates, resulting in a larger probability (bottom-right panel). Finally, a stricter threshold of 0.7 is applied to the refined map to yield the final plume mask.

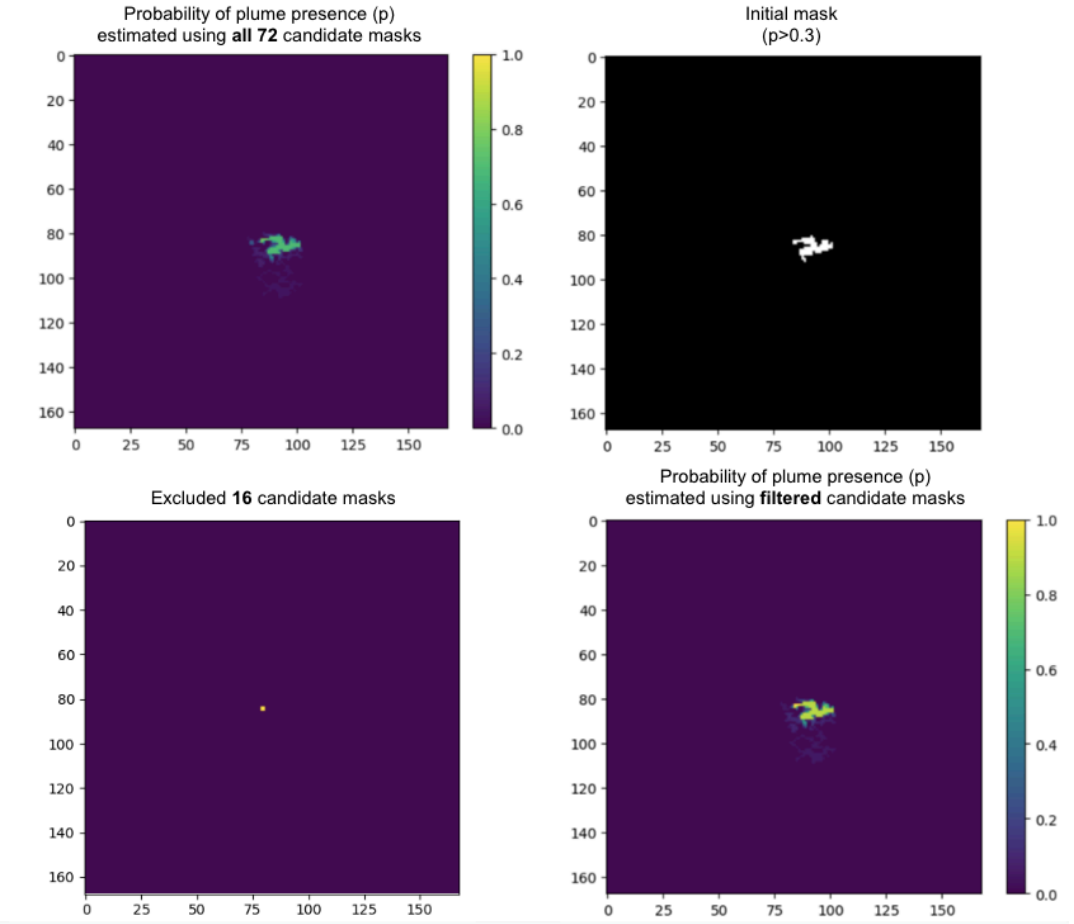

Figure 3. Example of the necessity of performing a two-step probability thresholding procedure during the segmentation process (step 5).

The methods above describe the Tanager and EMIT plume segmentation process. A nearly identical process is applied to AVIRIS-3 concentration data, with the exception of step 3. For AVIRIS-3 identified plumes, a cluster must contain 5 pixels to be considered part of the plume and is then dilated by 1 pixel. Lastly, a cluster closest to the origin, within 16-pixels, is selected as a candidate mask. These parameters were selected through visual inspection of plume masks and comparison to validation datasets.

4 Emission Quantification

We estimate plume emission rates (units kg h^{-1}) using the Integrated Mass Enhancement (IME; units kg; Thompson et al., 2016) approach, which calculates the excess mass emitted to the atmosphere from a source:

$$IME = \alpha \sum_{i=1} \Omega_i A_i \quad (1)$$

Where i refers to a single plume pixel, Ω is the concentration enhancement of that pixel, α is a unit conversion scalar (from ppm-m to kg m^{-2}), and A is the area of that pixel (m^2). Scene specific α is defined using the pressure and temperature at the ground elevation queried from MODTRAN standard atmosphere, the same conditions used for L2B processing. We calculate an emission rate Q using the following relationship (Ayasse et al., 2023):

$$Q = \frac{IME}{L} U \quad (2)$$

Where U is the 10-m wind speed (m s^{-1}) and L is the plume length (m). Here U is taken from the High Resolution Rapid Refresh (HRRR) 3km, 60 minute reanalysis product in the U.S. and the ECMWF IFS 9km outside the U.S. Comparisons (Figure 4) of these gridded reanalysis products to 10-m weather station data in the U.S. (via the Synoptic Weather Data API: synopticdata.com) show high scatter for any data point, and roughly 1 m/s mean absolute bias.

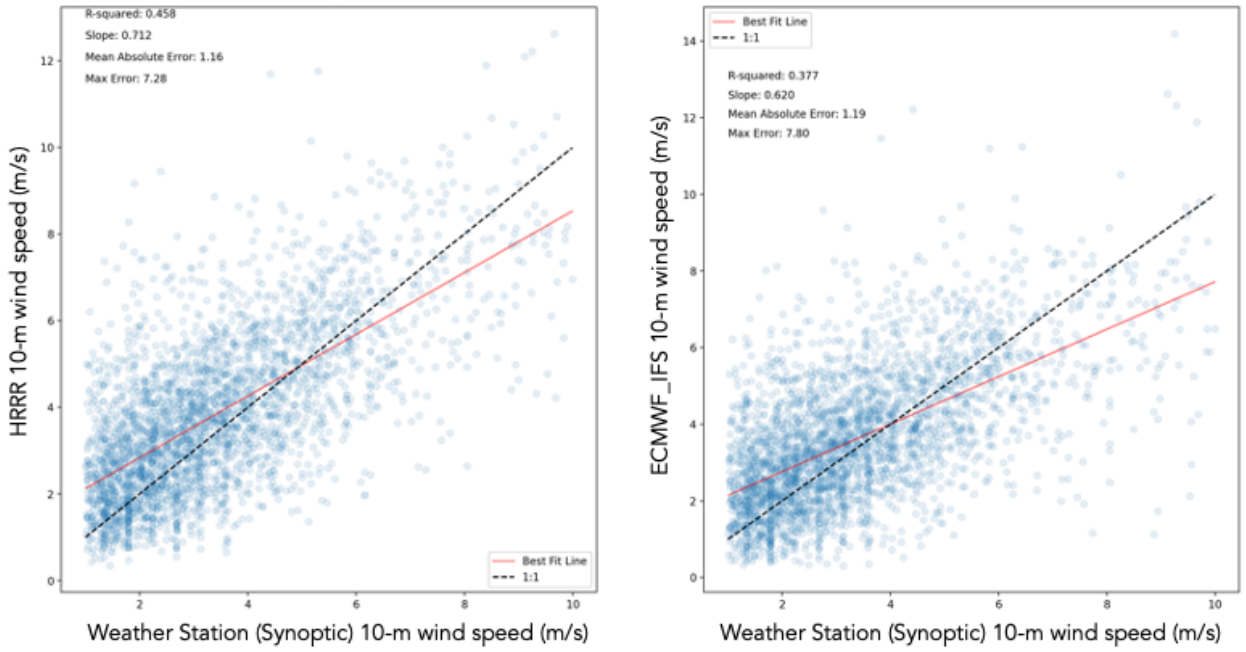

Figure 4. Comparison of U10 wind speeds from the HRRR and ECMWF IFS global data products to 10-m weather station data from the Synoptic Weather Data API (synopticdata.com).

In Equation 2, L is estimated as the max distance from one point of the plume to another point along the segmented plume's convex hull (i.e., maximum Feret diameter). For plumes covering large spatial distances, we impose a distance constraint such that the segmented plume mask is clipped to not exceed a 2500 m radial extent from the origin of the plume. By cropping the mask as described in step 1 of plume segmentation, operational plume lengths cannot exceed 5000m. The IME (Equation 1) is also only calculated within this clipped plume mask. This clipping procedure is employed to reduce bias that may enter into IME

quantification due to differing surface and meteorological conditions across large plumes, dispersed methane at the edge of the plume, intermittency of the emission rate of the source, and to limit potential merging of multiple plumes downwind of their sources.

5 Uncertainty Quantification

Uncertainties in emission estimates are calculated by summing in quadrature elements that contribute to variability in emissions:

$$\sigma_Q = \sqrt{\left(\frac{\partial Q}{\partial U}\sigma_U\right)^2 + \left(\frac{\partial Q}{\partial IME}\sigma_{IME}\right)^2 + \left(\frac{\partial Q}{\partial L}\sigma_L\right)^2} \quad (4)$$

Where

$$\sigma_{IME} = \frac{\partial Q}{\partial IME}\sigma_N + \frac{\partial Q}{\partial \Omega}\sigma_\Omega \quad (5)$$

In Equation 4, the $(\frac{\partial Q}{\partial U}\sigma_U)$ term represents the uncertainty due to wind speed, which we estimate by computing the standard deviation of 10-m wind estimates across the hour before, during, and after the plume detection. For global ECMWF IFS products this is simply the standard deviation of three 9km tiles. For US based HRRR products we take an average of nine adjacent 3km tiles per hour (creating an equal area to the 9km tiles used globally) and take the standard deviation between these three averages. This keeps our wind uncertainties comparable across US and global targets. The $(\frac{\partial Q}{\partial IME}\sigma_{IME})$ term is decomposed into two components, uncertainty due to masking and due to retrieval. The first uncertainty due to masking $(\frac{\partial Q}{\partial IME}\sigma_N)$ is parameterized as the standard deviation of IME estimates across dynamic plume masks which have lengths within median ± 1×std derived from the segmentation approach described in Figure 2. The second uncertainty due to the retrieval $(\frac{\partial Q}{\partial \Omega}\sigma_\Omega)$ is estimated as the quadrature sum of pixel-wise L2b concentration uncertainty within the plume mask calculated following Fahlen et al. (2024) for CH4. For CO2 the standard deviation of concentration enhancements outside of the segmented plume mask, but within a 2500 m crop of the plume define retrieval uncertainty $(\frac{\partial Q}{\partial \Omega}\sigma_\Omega)$. Finally, the $(\frac{\partial Q}{\partial L}\sigma_L)$ represents an irreducible uncertainty term due to the pixel resolution of the satellite instrument and how it affects the estimate of plume length L. Figure 5 shows the effect of this term: the edge of a plume may manifest as concentration enhancement in a single or multiple pixels depending on the true geolocation of the plume and the spatial resolution of the instrument.

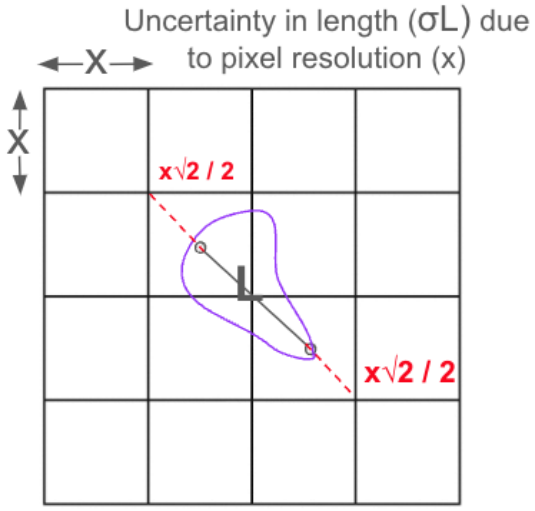

Figure 5. Irreducible uncertainty due to pixel resolution for IME quantification approaches.

6 Emission Quantification Issues and Quality Checks

Not all instances of plume detections are quantified for emissions. A quality control procedure rules out publishing emission rates if there are issues severely affecting quantification, including (1) the overwhelming presence of artifacts in the retrieval; (2) non-standard plume shapes that may violate mass balance assumptions of IME (high wind shear, concentration pooling, etc) or make plume segmentation difficult (e.g. large gaps in plume); (3) plumes that appear at the boundaries of images; (4) overlapping plumes where emission rates from visibly distinct sources can not be partitioned. In these cases, a detection with no emission rate may be published.

7 Validation of Emission Rate Quantification

Carbon Mapper's emission rate validation procedure follows several independent protocols (Figure 6). This includes controlled release testing, aircraft underflight, distributional assessment, and application to synthetic data. All of these experiments were performed for Tanager. For EMIT and AVIRIS-3, we complete a subset of these experiments and results for EMIT and AVIRIS-3 are included for the experiments we completed. We use results from these protocols to build confidence and refine algorithm development.

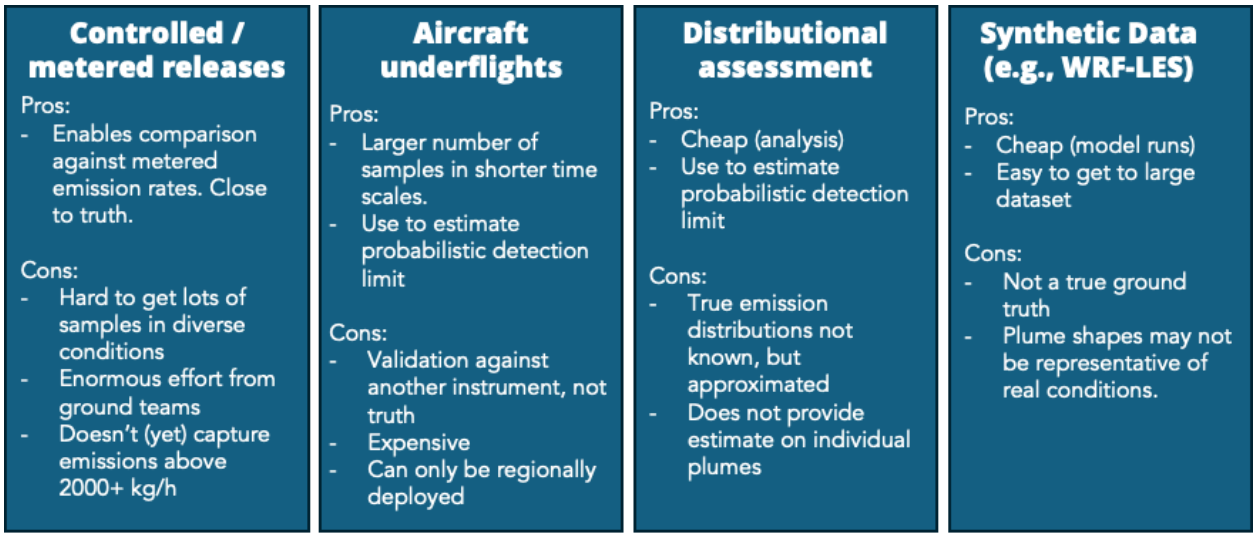

Figure 6: Experiments used to validate point source methane emissions with relative pros and cons of each experiment.

7.1 Controlled Releases

Carbon Mapper participated in a controlled release experiment performed by Stanford University that included an unblinded phase, and subsequently 4 blinded phases between 2024-2025. Tanager quantified emission rates compared to metered releases show close agreement (y=1.04x; R^2=0.89) under both blinded and unblinded conditions at sites in Wyoming and Arizona between 2024-2025, providing high confidence in quantification methods described for those environments (Figure 7). For Tanager, the plume detection with the lowest metered emission rate was 49 kg/h. There were no false positive detections or false negative detections during the experiment for Tanager. AVIRIS-3 similarly performs well (y=1.09x; R^2=0.92), which not only provides high confidence in the emission rate quantification from AVIRIS-3 but also helps provide additional confidence in the results from the aircraft underflights shown in section 7.2. The EMIT controlled release results show a slight high bias but with all cases except two the truth values are within uncertainties, and the correlation is robust (y=1.38x; R^2=0.95).

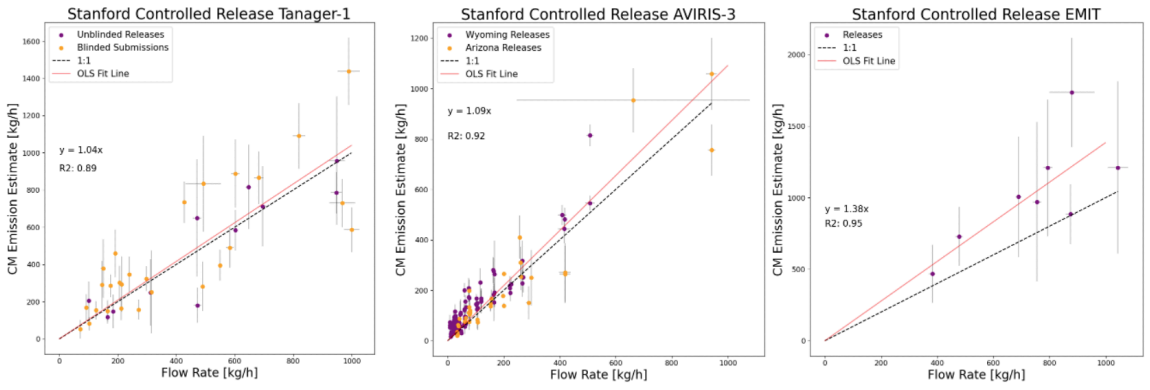

Figure 7: Controlled release testing for Tanager, AVIRIS-3, and EMIT, representing both unblinded and blinded conditions for Tanager. AVIRIS-3 and EMIT were both performed under unblinded conditions. The x-axis (Flow Rate) represents metered values (averaged 90 seconds prior to observation) and the y-axis represents the Carbon Mapper emission estimates. Carbon Mapper uncertainties were calculated as described in Section 5, displayed as 1-sigma values.

7.2 Aircraft Under Flights

In October and November of 2024, Carbon Mapper coordinated AVIRIS-3 flights to coincide with Tanager overpasses in the Permian Basin and the controlled release site. Results from the plumes observed by AVIRIS-3 and Tanager are shown in Figure 8. Generally the agreement is good, but the agreement degrades the longer the spacing between aircraft and satellite observing. In addition to a direct comparison of emission rates, the coincident AVIRIS-3 and Tanager observations allow for a qualitative assessment of the plume morphology from a plume with similar emission rates (Figure 9).

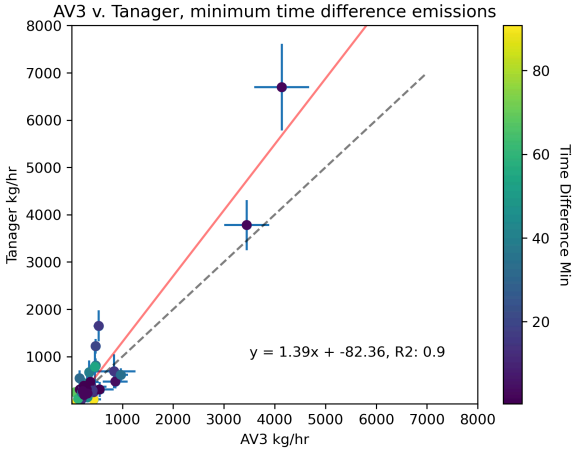

Figure 8. Comparison of simultaneously observed plumes by Tanager and AVIRIS-3 plumes during coordinated flights. The plumes are colored by the time difference (minutes) between the Tanager and AVIRIS-3 acquisitions. Uncertainties are calculated following Section 5 and represent 1-sigma.

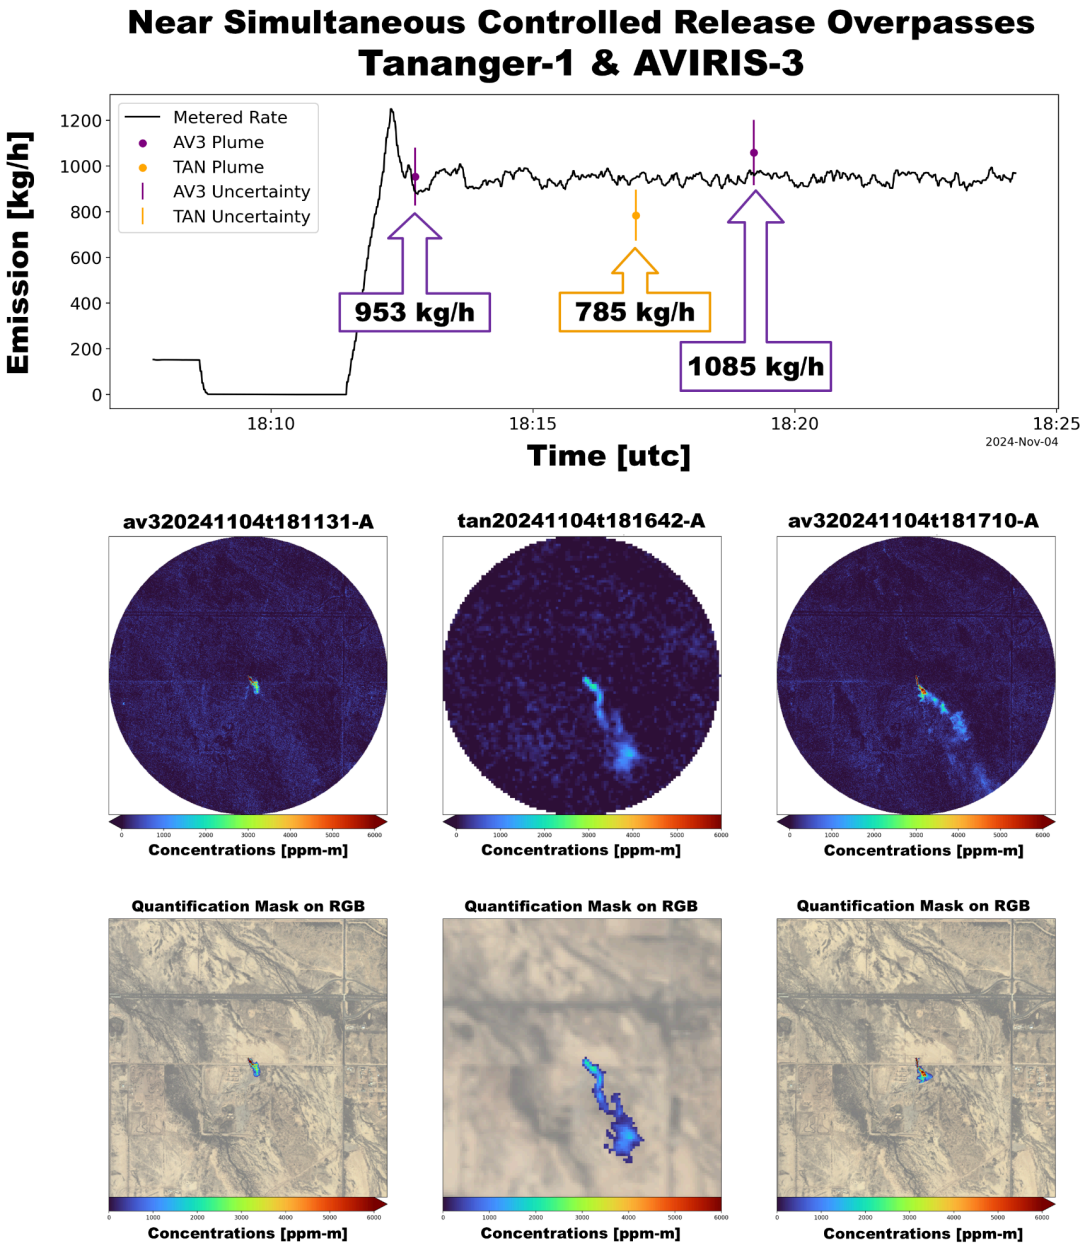

Figure 9: AVIRIS-3 and Tanager observation of a controlled release plume. The top panel shows the continuous flow rate (the truth emission; y-axis) with the AVIRIS-3 and Tanager emission rates at the time of their respective acquisitions (x-axis; time). The middle panel shows the matched filter retrievals for each plume and the bottom panel shows the segmented masks overlayed on the RGB.

7.3 Distributional Analysis

Distributional assessments of quantified emissions can help to assess systematic biases in aggregate plume quantification, and can help verify probabilistic detection limit estimates. Multiple previous studies have shown that oil and gas CH4 emissions above super-emitter thresholds (100 kg/h) generally follow a heavy-tailed distribution (Sherwin et al., 2024). Under this assumption, the peak of an emission rate cumulative distribution function (CDF) represents an estimate of detection rate (e.g., 90% probability of detection (POD)) for an

instrument, and all emissions left of the peak are not representative of the true distribution of emissions but also the partial detection limit of the instrument. Theoretical and empirical assessments of Tanager's reliable detection range result in estimates of the 90% POD between 200-300 kg/h. Figure 10 shows the distribution of oil and gas plumes quantified by Tanager using the algorithms described above and shows a peak near the theoretical and empirical derived limits.

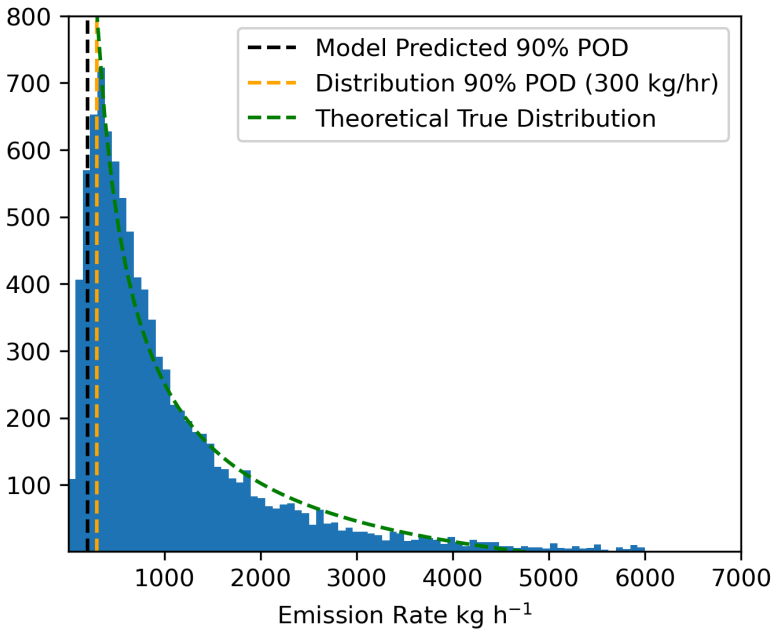

7.4 Synthetic Images

Given relatively sparse validation data, we tested optimized algorithms against simulated data from the Weather and Research Forecasting Model - Large Eddy Simulation (as in Varon et al., 2018). These simulations were performed under five unique wind speed regimes at a 25m spatial resolution with scalable concentrations within noiseless backgrounds.

To better simulate Tanager and EMIT data, simulations were downsampled to 30m and 60m respectively to match each sensor's average ground sample distance. Emission distributions of published plumes within the wind bounds of the initial simulations were randomly sampled 1,000 times for each sensor. To each emission rate representative gaussian noise was inserted. Noise was generated by determining a maximum noise value for the given emission which would still allow detection, based on theoretical minimum detection limit (MDL) calculations, and randomly sampling the true sensor noise distribution below this value. A random LES time stamp was then selected for each of the 1,000 emission and noise combinations which were applied to all five wind regimes. This resulted in 5,000 representative plumes per instrument for assessment as outlined in Figure 11.

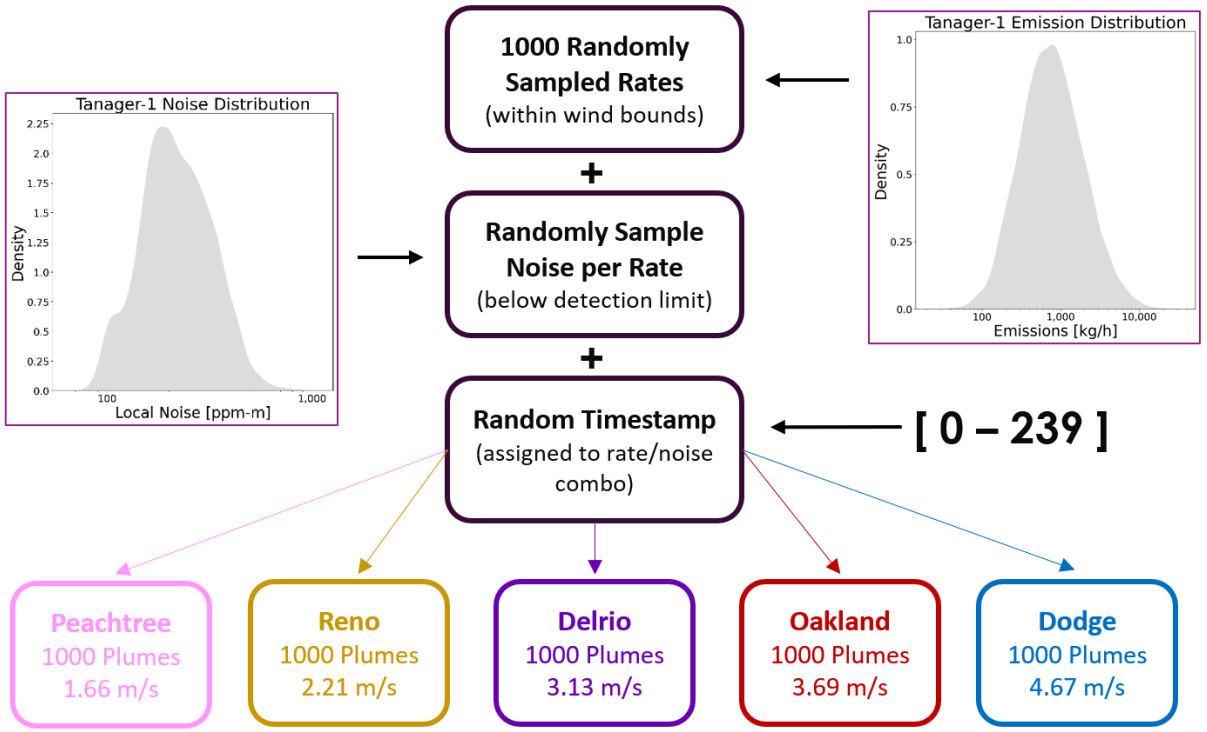

Figure 11. Processing steps to produce 5,000 representative LES outputs.

Quantification results from applying standard Carbon Mapper algorithms to these datasets are shown in Figure 12. For each sensor segmentation and quantification produces some scatter, especially at large (>5,000 kg/h) emission rates. These results are unbiased in aggregate with an ordinary least squares (OLS) fit to the data without an intercept (Sherwin et al., 2023) resulting in an $R^2 = 0.91$ and $y = 0.91x$ for Tanager based data and an $R^2 = 0.84$ and $y = 0.87x$ for EMIT based data. All plumes intersect the 1:1 line within uncertainties. There is a minor low bias, particularly apparent at rates above 5,000 kg/h, possibly incurred by the downsampling process. Due to potential biases inherent in any simulated dataset we do not use these results to derive an additional bias-correction factor (often called the “effective wind”).

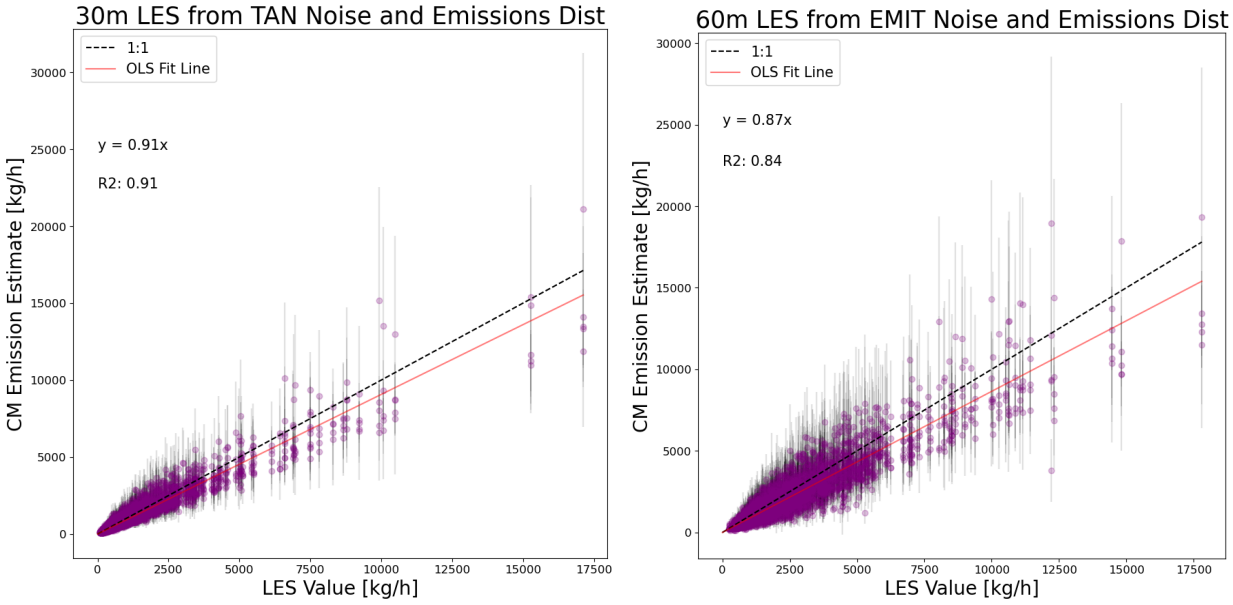

Figure 12. LES testing for Tanager and EMIT. The x-axis (LES Value) are the simulated emission rates and the y-axis are Carbon Mapper emission estimates for each output.

8 Outputs

8.1 Plume image:

Segmented plume image (GeoTIFF) that provides the boundaries and concentrations (units ppm-m) of pixels that were used for quantification

8.2 Plume information:

Each plume that has passed quality control procedures may include the following information: geographic coordinate of plume origin (process described in Carbon Mapper Plume Detection Quality Control Protocol), emission rate, emission rate uncertainty, IPCC sector attribution, IME, wind speed, quality flags.

9 References

Ayasse, A., Cusworth, D., O'Neill, K., Thorpe, A. and Duren, R., 2023. Performance and sensitivity of column-wise and pixel-wise methane retrievals for imaging spectrometers.

Ayasse, A., Cusworth, D.H., Howell, K., O'Neill, K., Conrad, B., Johnson, M.R., Asner, G. and Duren, R., 2024. Probability of Detection and Multi-Sensor Persistence of Methane Emissions from Coincident Airborne and Satellite Observations.

Loizzo, R., Guarini, R., Longo, F., Scopa, T., Formaro, R., Facchinetti, C. and Varacalli, G., 2018,

July. PRISMA: The Italian hyperspectral mission. In IGARSS 2018-2018 IEEE international geoscience and remote sensing symposium (pp. 175-178). IEEE.

Sherwin, E.D., Rutherford, J.S., Zhang, Z., Chen, Y., Wetherley, E.B., Yakovlev, P.V., Berman, E.S., Jones, B.B., Cusworth, D.H., Thorpe, A.K. and Ayasse, A.K., 2024. US oil and gas system emissions from nearly one million aerial site measurements. Nature, 627(8003), pp.328-334.

Sherwin, E.D., Rutherford, J.S., Chen, Y., Aminfard, S., Kort, E.A., Jackson, R.B. and Brandt, A.R., 2023. Single-blind validation of space-based point-source detection and quantification of onshore methane emissions. Scientific Reports, 13(1), p.3836.

Sherwin, E.D., El Abbadi, S.H., Burdeau, P.M., Zhang, Z., Chen, Z., Rutherford, J.S., Chen, Y. and Brandt, A.R., 2024. Single-blind test of nine methane-sensing satellite systems from three continents. Atmospheric Measurement Techniques, 17(2), pp.765-782.

Thompson, D.R., Thorpe, A.K., Frankenberg, C., Green, R.O., Duren, R., Guanter, L., Hollstein, A., Middleton, E., Ong, L. and Ungar, S., 2016. Space-based remote imaging spectroscopy of the Aliso Canyon CH4 superemitter. Geophysical Research Letters, 43(12), pp.6571-6578.