Methane emission detection, localization, and quantification using continuous point-sensors on oil and gas facilities

William S Daniels,* Meng Jia, and Dorit M. Hammerling

Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, Colorado, United States

E-mail: wdaniels@mines.edu

Abstract

We propose a generic, modular framework for emission event detection, localization, and quantification on oil and gas facilities that uses concentration data collected by point-in-space continuous emissions monitoring systems (CEMS). The framework uses a gradient-based spike detection algorithm to estimate emission start and end times (event detection) and pattern matches simulated and observed concentrations to estimate emission source location (localization) and rate (quantification). We test the framework on a month of single-source controlled releases ranging from 0.50 to 8.25 hours in duration and 0.18 to 6.39 kg/hr in size conducted at the Methane Emissions Technology Evaluation Center in Fort Collins, Colorado. All controlled releases are identified and 82% are localized correctly. For emissions $\leq$ 1 kg/hr, the framework underestimates by 37.2% on average, with 90% of rate estimates within a factor of [-4.6, 2.8] or a percent difference of [-78.1%, 178.6%]; for emissions $>$ 1 kg/hr, the framework overestimates by 1.5% on average, with 90% of rate estimates within a factor of [-2.0, 1.8] or a percent difference of [-49.6%, 77.4%] from the true rates. Potential uses for

the proposed framework include near real-time alerting for rapid emissions mitigation and emission quantification for data-driven inventory estimation on production-like facilities.

1 Introduction

Reducing methane emissions is a key component of short-term climate action and would greatly increase the feasibility of the 1.5 degree temperature goal from the 2015 Paris Climate Agreement.[1,2] The oil and gas sector accounts for 22% of global anthropogenic methane emissions[3,4] and 32% within the U.S.,[5] and hence it provides a promising avenue for emissions reduction.

Recent attention and advances in measurement technology have revealed a number of important characteristics of methane emissions from the oil and gas sector. First, emissions exhibit high temporal variability, with intermittent events such as liquids unloadings and blowdowns representing a significant portion of total emissions.[2,6–8] Second, emissions can vary by orders of magnitude across basins and facilities.[9–11] Third, infrequent super-emitter events can represent a large portion of total emissions, and hence measuring these events is critical for accurate emissions accounting.[9,12,13] Bottom-up inventories have been found to underestimate total emissions and are poorly suited to accommodate these complicated emission characteristics.[14–19] Therefore, direct measurements are needed to better understand and mitigate emissions. In fact, the Inflation Reduction Act will soon require that all reported methane emissions from the United States oil and gas sector are derived from empirical data rather than conventional bottom-up inventories.[20]

A range of methane measurement technologies exist, including satellites, aircraft, and ground-based continuous emissions monitoring systems (CEMS). Current public domain satellite data is too coarse to retrieve all but the largest emission sources, with a detection threshold of around 63 kg/s (226,800 kg/hr).[21,22] Commercial satellites designed to target specific locations have lower detection thresholds, around 4 kg/s (14,400 kg/hr), but the

data is (for the most part) not publicly available.[21,22] Aircraft overflights have been shown to reliably estimate emission rates,[7,16,22,23] but occur too infrequently to consistently capture the intermittent, short-lived emission events that play an important role in the overall emissions profile. This leaves CEMS as the main avenue for providing measurement-based information on methane emissions in near-real time.

There are three broad classes of CEMS: solutions that create actionable information using (1) a sensor network that provides point-in-space concentration observations at fixed time intervals, (2) scanning lasers to measure long, integrated open-path concentrations, and (3) cameras to create 2D images of concentration enhancements.[24,25] In this work, we propose a new analytical framework for methane emission detection, localization, and quantification that uses data from the first class of CEMS listed above. Our framework is intended to provide rapid feedback to oil and gas operators so that emissions can be quickly mitigated, and as such is designed to operate in near real-time. Finally, we focus on production-like facilities that are around 150m in diameter and contain a number of distinct equipment groups (e.g., wellheads, separators, and tanks).

A number of groups have proposed inversion techniques for emission localization and quantification in this setting. Some of these solutions rely on a large sensor network (20 plus), which would be cost prohibitive in practice,[26,27] or on long term (multiple hours to days) aggregates, which would be unable to identify short-lived events, even if they are large.[28,29] Additionally, many solutions rely on the Gaussian plume atmospheric dispersion model for modeling atmospheric transport.[27,30] This model was optimized for far-field applications with distance scales much larger than the typical production-like facility and assumes steady state wind conditions over short ($\sim$10 minute) time intervals, an assumption which breaks down in many practical scenarios where wind conditions are variable. On the other end of the complexity spectrum, some solutions rely on large eddy simulations for modeling atmospheric transport,[26,31] which are far more accurate than the Gaussian plume model. They do, however, require special expertise to operate and implement for specific sites, are

computationally expensive, and are not generally publicly available. Finally, none of the solutions discussed here attempt to perform emission event detection, and hence they are all tested in rounds with known start and stop times.[26–31] Note that proprietary solutions targeting emission event detection, localization, and quantification on oil and gas facilities are developing rapidly in the private sector, which we are unable to assess.

We propose two methods related to emission event detection, localization, and quantification that use concentration data from a point-in-space sensor network. First, a spike detection algorithm that filters identified spikes by their background corrected amplitude. Second, an emission detection, localization, and quantification framework that contributes the following advancements. First, no information on emission start or end times is needed to use the framework. The framework can be run continuously and will only estimate emission location and rate if an emission event is detected (via the spike detection algorithm). Second, the framework uses a Gaussian puff atmospheric transport model that provides a balance between simpler models whose assumptions are rarely met in practice (e.g., the Gaussian plume) and more complex simulations that are computationally expensive and may require customization for different locations. Third, with the goal of providing rapid alerts, the framework can provide a location and rate estimate using as little as 15 minutes of data. Finally, the framework does not require an expansive sensor network, but rather can operate using only as many sensors as are required to surround all potential emission sources on a typical production-like facility (approximately 4-8 sensors). Note that the proposed framework is generic and works with any sensor network that provides point-in-space methane concentration, wind speed, and wind direction measurements.

We test the proposed methods on approximately one month of single-source controlled releases at the Methane Emissions Technology Evaluation Center (METEC) during a range of meteorological conditions. Note that the controlled release data used to evaluate the proposed framework were not blinded, meaning that the leak locations and rates were provided to us along with the raw CEMS data. We only use this information for evaluation; the pro-

posed framework does not rely on any site-specific parameters that are fit to training data. While there are a number of inputs to the framework, they control site-agnostic settings related to, e.g., how often a rate estimate is produced. See Section S1 in the supporting information file for a list of framework inputs.

2 Methods

Here we describe the two main contributions of this work: (1) a gradient-based spike detection algorithm and (2) a modular framework for methane emission event detection, localization, and quantification.

2.1 Spike detection algorithm

We present an algorithm for flagging sharply elevated values ("spikes") in a univariate time series, which we assume to be methane concentrations. This algorithm is used in the event detection, localization, and quantification framework to estimate background methane concentrations, determine when an emission is occurring, and isolate time steps in which a given sensor is recording a relevant signal.

The algorithm proceeds as follows. Iterate once through the univariate time series. When an observation is **"going up threshold"** parts per million (ppm) greater than the previous observation, start a spike. Remain in the spike until the concentration values return to **"return threshold"** percent of the maximum concentration encountered during the spike. After exiting the spike, average the concentration values immediately preceding and following the spike and use this value as the background estimate for the spike, which imposes no assumptions on the spatial or temporal homogeneity of the methane background. Subtract the background estimate from all observations within the spike, and if the maximum background-corrected concentration in the spike is less than **"amplitude threshold"** ppm, discard the spike.

This loop results in a mask that records if each entry in the concentration time series is a spike or background. The algorithm relies on three parameters: (1) the **going up threshold** in ppm, which is used to identify the start of an event, (2) the **return threshold** as a percent, which is used to identify the end of an event, and (3) the **amplitude threshold** in ppm, which is used to filter spikes by their background-corrected amplitude. These values are set to 0.25ppm, 5%, and 1ppm by default and can be adjusted for different use-cases. The algorithm requires only a single loop through the time series, and hence it can be run on historical or real-time data. See Section S2 in the supporting information file for a detailed description of the spike detection algorithm.

2.2 Event detection, localization, and quantification framework

We now propose a framework for emission detection, localization, and quantification. We design the framework to be highly modular, meaning that it is broken up into distinct steps that can be executed using any number of metrics or methods. Again, the framework is generic and works with any sensor network that provides point-in-space methane concentration, wind speed, and wind direction measurements. Note that the quality of the sensor data (i.e., accuracy and precision) will impact the quality of the framework output, but for brevity, we do not analyze or discuss this point further in this work.

At a high level, the framework pattern matches observations from the CEMS to simulated observations assuming different potential sources. For each identified emission event, the potential source whose simulated observations most closely match the CEMS observations is taken to be the estimated source for that event. We then estimate emission rate by minimizing mean squared error (MSE) between the simulated observations from the most likely source and the CEMS observations. We only produce a rate estimate when there is reasonably good alignment between the simulation and CEMS observations, a measure we have found useful in practice given that the framework relies on an atmospheric transport model that does not directly model turbulence. The framework is broken up into four steps

that are described in Sections 2.2.1 through 2.2.4.

Note that while event detection, localization, and quantification can be framed as an inverse problem, we do not perform a full inversion to retrieve source location and emission rate. Instead, we greatly reduce the dimension of the problem by specifying potential emission sources. Doing so is a useful choice in many practical applications on oil and gas facilities, as potential emission sources are often well-known (e.g., tank thief hatches, separators, and wellheads).

2.2.1 Remove background and perform event detection

The simulated methane concentrations described in the following section do not include background concentrations that are present in the atmosphere. Therefore, we must first remove background concentrations from the CEMS observations before comparing them to simulation predictions. We do so for each sensor separately by using the spike detection algorithm proposed in Section 2.1. Summarizing this procedure at a high level, we group identified spikes that are close together in time, background correct each group using the concentration observations immediately preceding and following the group, and then set all observations that are not in a group to zero.

Next we define time intervals over which to perform localization and quantification. We do so by taking the minute-by-minute maximum across the background-corrected concentration time series from each sensor, which collapses the signal from each sensor into one time series that captures the concentration spikes across the entire facility. If the sensors sufficiently surround the facility, then any on-site emission will cause an enhancement in this maximum value time series, regardless of wind direction. Since all non-spike observations have already been set to zero (during background removal), we define groups of non-zero values in this maximum value time series as emission events. As a last step, we combine events that are separated by less than 30 minutes, as there are often small gaps between concentration enhancements that do not correspond to gaps in emissions, but rather to periods in which the

methane plume is being blown between sensors. We further discard events that are less than 15 minutes long, as these events typically correspond to noise rather than actual emissions.

See Section S3 in the supporting information file for details and an alternative method for defining time intervals that is better suited for more complex facilities.

2.2.2 Simulate methane concentrations at sensor locations

We now simulate methane concentrations at each sensor location during the events identified in the previous module. We run a separate simulation for all potential sources on the oil and gas facility, which allows us to pick the most likely source for each event in the following module. The framework is not intrinsically tied to any atmospheric transport model, and for this paper we select the Gaussian puff model as a balance between simulation accuracy and computational expense, availability, and ease of use. This model approximates a continuous cloud of methane released from a point source as a sequence of many small "puffs," each of which is modeled using a three-dimensional Gaussian-like function. This provides a reasonable approximation of atmospheric transport within a short time period after release and over short distances barring any major obstructions that would block the transport of methane (e.g., a large building), and hence we think it is suitable for use in our framework when applied to oil and gas production facilities.

Predicted concentrations from the Gaussian puff model are a linear function of the selected emission rate, meaning that the pattern of the predictions does not depend on the choice of emission rate. Therefore, we simulate with a unitary emission rate (1 g/s), as true emission rates are unknown in practice. With this in mind, we design the pattern matching module described in Section 2.2.3 to be scale-independent by evaluating patterns between observations and predictions rather than overall amplitudes. However, to avoid issues with code-internal quality check thresholds, the simulated concentrations should be on approximately the same order of magnitude as the sensor observations. We have found that simulating with an emission rate of 1 g/s works well on typical production facilities, but a

larger value may be better if, e.g., the sensors are located extremely close to the emission source. An error handling check has been built into the code that alerts the user when an adjustment to the simulation scale is recommended.

Section S4 in the supporting information file and Jia et al.[32] contain a detailed description of the Gaussian puff atmospheric dispersion model and our implementation for use in the framework proposed here.

2.2.3 Source localization

We now estimate the source for each emission event identified in Section 2.2.1. For each event, the potential source whose simulated observations most closely match the actual CEMS observations is taken to be the estimated source. As mentioned in Section 2.2.2, the true emission rate is unknown in practice, and hence we simulate with a unitary emission rate and pattern match using a scale-independent metric.

For each potential source, we compute the correlation between a stacked vector of simulation predictions at all sensor locations and a stacked vector of the corresponding CEMS observations. This results in a correlation value for each potential source for all identified emission events. For each emission event, the source with the highest positive correlation is taken to be the localization estimate for that event. Note that we tested a wide range of pattern matching metrics, including a number of custom metrics that rewarded spike alignment and penalized spike misalignment, and found that correlation had the best performance.

2.2.4 Emission rate quantification

We now estimate emission rates by comparing the amplitude of the CEMS observations to the amplitude of the simulation predictions, assuming the localization estimate from Section 2.2.3. We only perform this comparison using time steps in which both the observations and predictions are in a spike, which we identify using the spike detection algorithm described in Section 2.1. This drastically reduces the impact of transport model inadequacies on the

emission rate estimate and in part justifies the use of a relatively simple model. If there are at least four time steps during which both the observations and predictions are in a spike, then a rate estimate is computed for that emission event. Otherwise, we deem the event unsuitable for quantification. We require at least four time steps because the impact of outliers on the resulting rate estimate is amplified when using fewer points. Note that the framework is not highly sensitive to this threshold, meaning that slightly different values do not result in drastically different rate estimates. However, if more rate estimates are desired (potentially at the expense of accuracy), the threshold can be lowered, and vice versa.

For each identified emission event, we sample with replacement 1000 times from the time steps in which both predictions and observations are in a spike. For each sample, we minimize a loss function over a range of emission rates between a stacked vector of simulation predictions at all sensor locations and a stacked vector of the corresponding CEMS observations. This optimization can be done using the simulation output from Section 2.2.2, as the Gaussian puff dispersion model is a linear function of emission rate. Hence, to compute simulation predictions at an arbitrary emission rate, q, we simply need to multiply the predictions simulated with a unitary rate by q. A number of different loss functions were tested, with mean squared error (MSE) having the best performance. The overall rate estimate for the event is then taken to be the mean of the sample-specific rate estimates.

There are two reasons why we perform this sampling procedure rather than simply minimizing MSE between predictions and observations. First, it reduces the impact of outliers on the resulting rate estimate, as many samples will not include the outliers. Second, the spread of the sample-specific rate estimates provides uncertainty quantification on the overall rate estimate. Specifically, we set the error bound on the overall rate estimate as the 5^{th} and 95^{th} percentiles of the sample-specific rate estimates. Note that this approach does not impose any assumptions on the symmetry of the error bound.

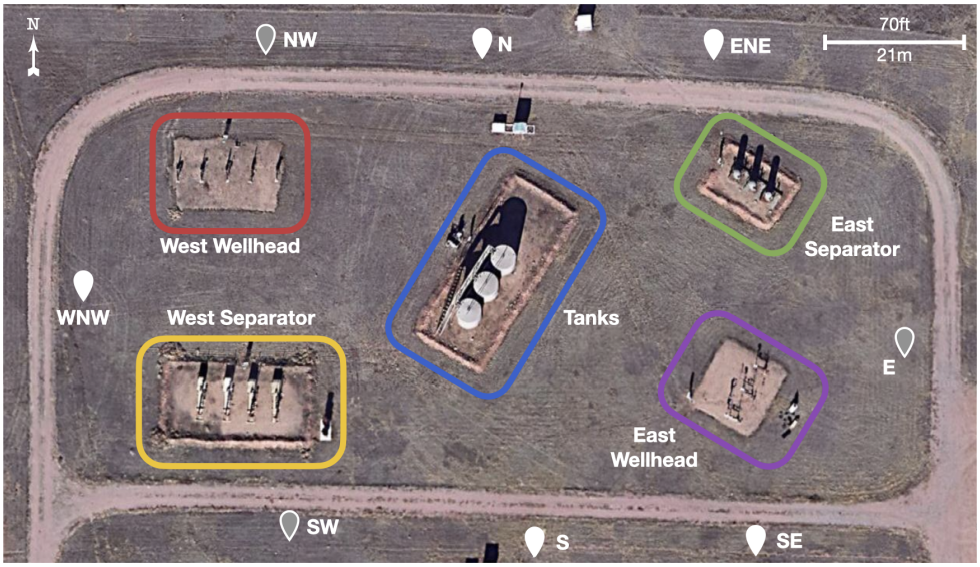

Figure 1: Satellite imagery of the Methane Emissions Technology Evaluation Center (METEC) facility. Source locations are marked with colored boxes. CEMS sensor locations are marked with pins and are named based on their geographical position. Pins with gray interiors indicate that the corresponding sensor measures wind speed and direction in addition to methane concentrations.

3 Results

We test the methods proposed in Section 2 using data collected by Project Canary[33] CEMS sensors at the Methane Emissions Technology Evaluation Center (METEC) in Fort Collins, Colorado during the Advancing Development of Emissions Detection (ADED) research program,[25] which Project Canary participated in. METEC is a testing center that resembles an oil and gas production facility and performs controlled methane releases from multiple pieces of equipment. We focus on roughly one month of ADED controlled releases during which there are primarily single-source emissions (April 17 to May 16, 2022). This study period contains 85 single-source emission events ranging from 0.50 to 8.25 hours in duration and 0.18 to 6.39 kg/hr in size. Emission rates are constant during each event. Controlled releases have variable start and end times and are separated by periods of no emissions with different durations. Emissions can come from one of five potential emission sources, indicated with colored boxes in Figure 1. Note that there are seven multi-source emission events during this study period, which we do not address in this work.

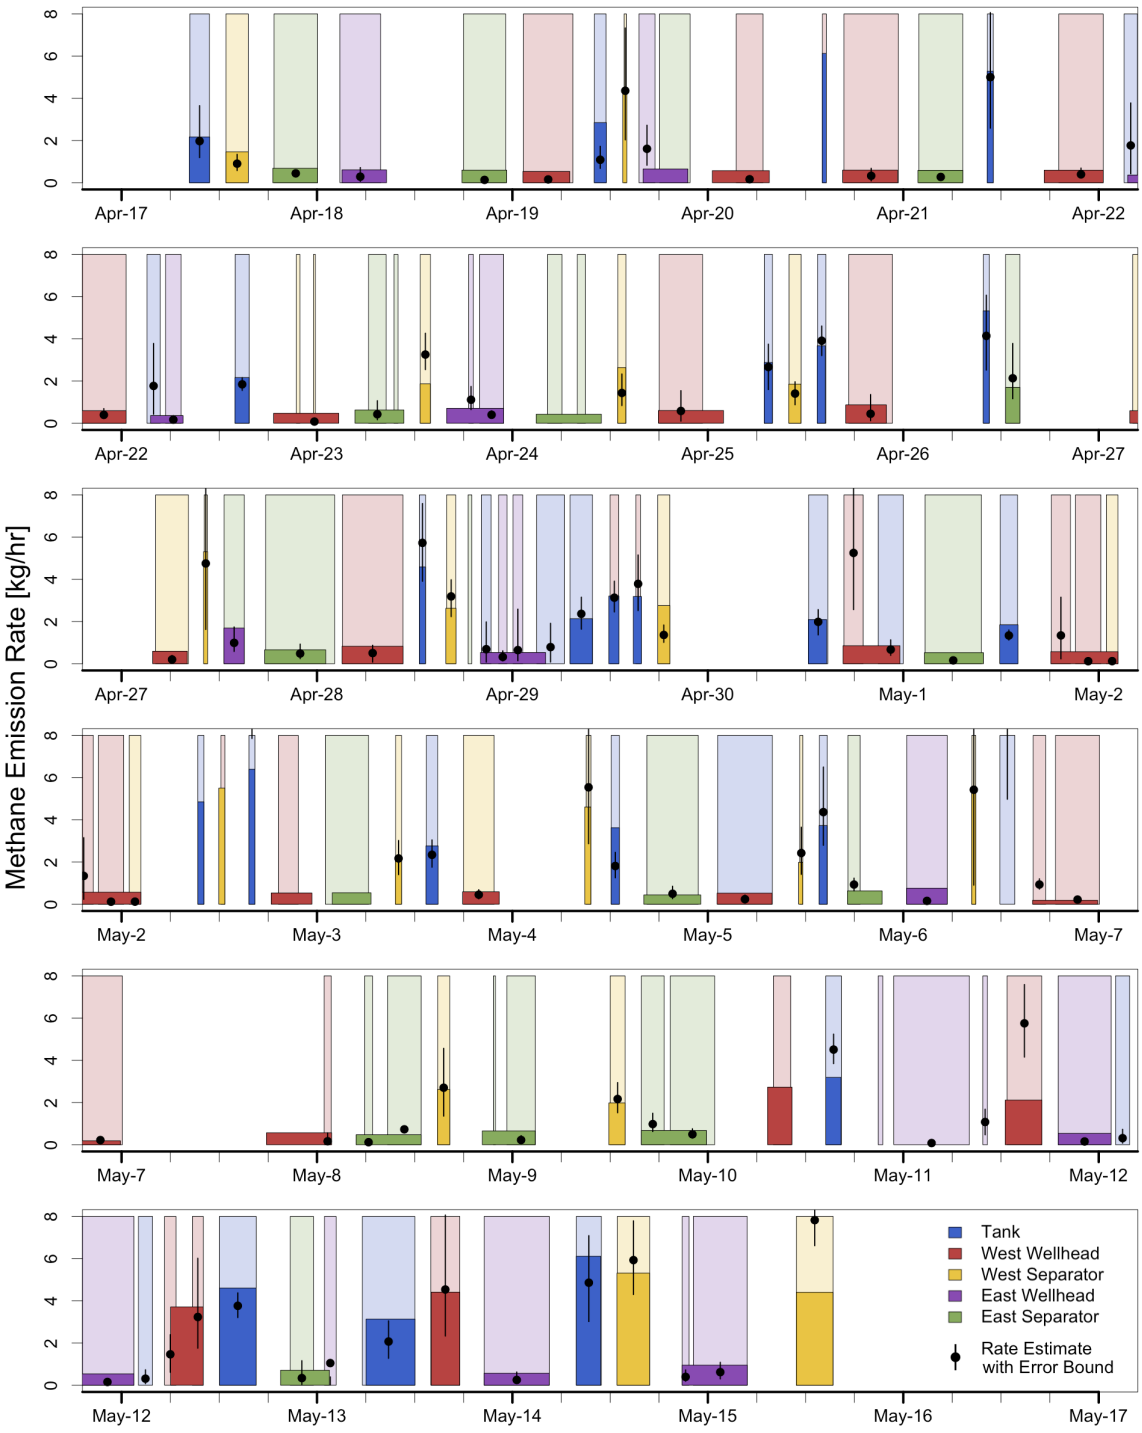

Figure 2: Detection, localization, and quantification results over the one month study period from April 17 to May 16, 2022. Controlled release events are indicated with solid rectangles, with the height of the rectangle corresponding to the true emission rate and the color corresponding to the true source. Timing and location of estimated events are indicated with transparent rectangles, with color corresponding to the estimated source. The heights of the transparent rectangles are fixed at an arbitrary value for visual clarity. Estimated emission rates and error bounds are indicated with black circles and lines, respectively.

Methane concentrations are measured by eight CEMS sensors placed around the perimeter of the METEC facility, three of which also measure wind speed and direction. The locations of the sensors that measure methane only and the sensors that measure both methane and wind are shown by white and gray pins, respectively, in Figure 1. Methane concentration, wind speed, and wind direction are measured every second and are averaged every minute by Project Canary, resulting in one data point per minute. We assume a homogeneous wind field across the site, and hence we take the median of the wind speed and direction measurements across the three sensors that measure these fields (taking into account the fact that wind direction is a circular variable). We found that using the median reduces the impact of sensor noise, which can otherwise result in highly variable wind speed and direction data.

Figure 2 provides an overview of the timing, source, and emission rate for each of the 85 emission events along with the event detection, localization, and quantification results from the framework proposed in Section 2. The remainder of this section will discuss specific aspects of the performance of the framework: emission event start and end time error, event detection and localization error, and quantification error.

Most start times are predicted accurately, with a median start time error of 4 minutes (meaning the predicted start time came 4 minutes late) and a distribution tightly centered around the median (25^{th} percentile = 1 minute and 75^{th} percentile = 11 minutes). Similarly, most event end times are predicted accurately, with a median end time error of zero minutes and a distribution tightly centered around the median (25^{th} percentile = -1 minute and 75^{th} percentile = 2 minutes). Overall, the framework correctly predicts the presence or absence of an emission event for 86.2% of time steps. See Sections S5 and S6 in the supporting information file for details on start and end time errors and minute-by-minute event detection errors, respectively.

We now consider a number of event-level performance metrics, which are summarized in Figure 3. The framework correctly identifies 100% of emission events, with 93% deemed

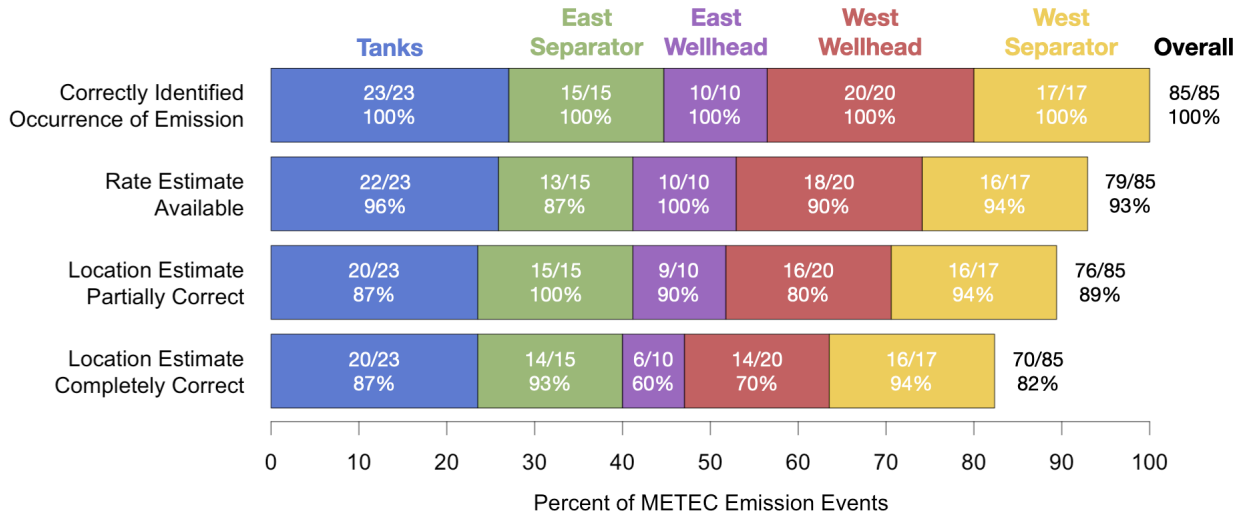

Figure 3: Summary of event-level performance metrics. The bar marked "correctly identified occurrence of emission" shows the percent of true emission events that overlap with a predicted emission event. The bar marked "rate estimate available" shows the percent of true emission events that are deemed suitable for quantification. Recall that the proposed framework can predict multiple emission events that all occur during one true event. Therefore, the bar marked "location estimate partially correct" shows the percent of true emission events where at least one of the overlapping predicted events has correctly predicted the source location, while the bar marked "location estimate completely correct" shows the percent of true emission events where all overlapping predicted events have correctly predicted the source location. Within each bar, the events are grouped by their true emission source.

suitable for quantification. Furthermore, 89% of emission events overlap with at least one predicted emission event that was correctly localized and 82% have a completely correct localization estimate. The performance in all four categories is largely consistent across the source locations. This level of event detection and localization performance suggests that the framework could be used to provide informative alerts to operators when emissions are occurring on production facilities similar in complexity to METEC.

We now consider quantification performance. Figure 4 shows a parity plot of the true and estimated emission rates, with the true rate plotted on the horizontal axis and the corresponding estimate on the vertical axis. The best fit linear model is shown in magenta. A slope of 1 and an intercept of zero indicates zero bias in the estimates (while saying nothing about the spread of the emission estimates around the best fit line). The proposed framework results in a best fit line with slope of 1.20 and intercept of -0.30. Estimates for the larger

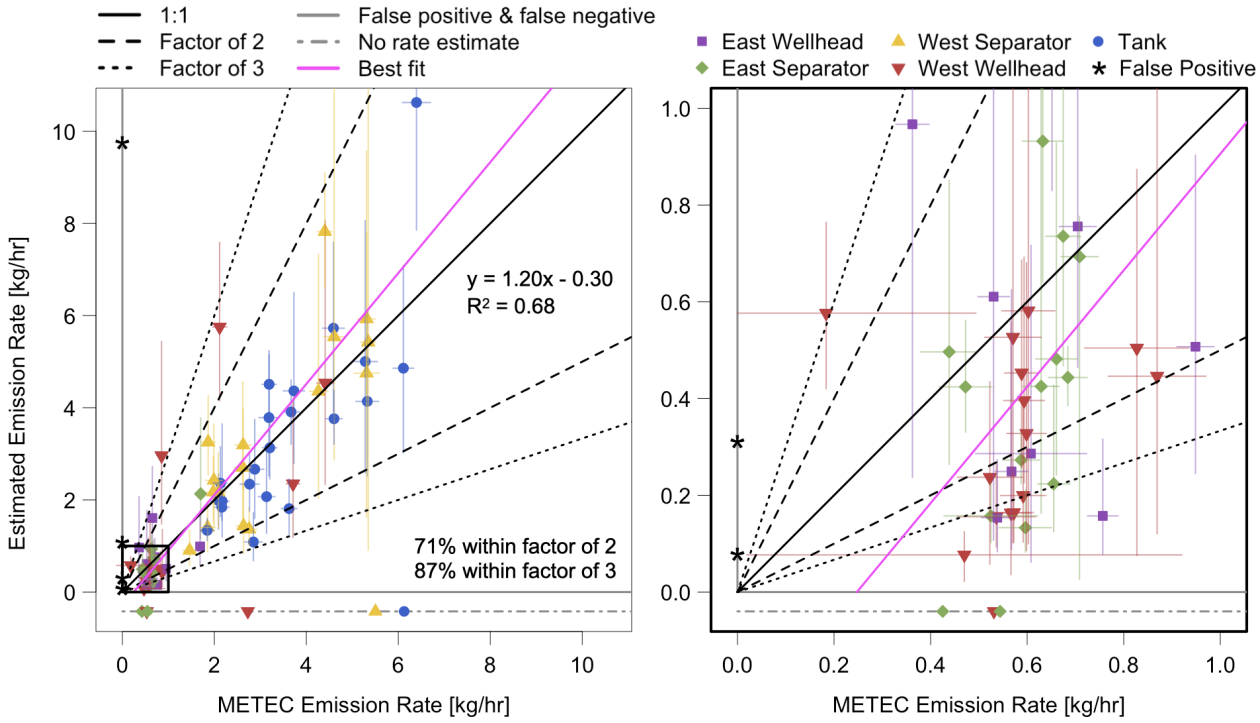

Figure 4: Parity plot of true and estimated emission rates. True emission rates are plotted on the horizontal axis and estimated rates are plotted on the vertical axis. Point color and symbol correspond to the true emission source. Points falling on the 1:1 line are perfect estimates. Points between the black dashed lines are within a factor of two of the true rate. Points between the black doted lines are within a factor of three of the true rate. Points along the vertical solid gray line are false positives, and points along the horizontal solid gray line are false negatives. Points along the horizontal dot-dashed gray line are emission events that were identified but were not deemed suitable for quantification. The magenta line shows the linear model fit to all identified events. One estimate of 14.4 kg/hr corresponding to a METEC emission of 4.9 kg/hr is excluded for visual clarity.

emissions ($>$ 1 kg/hr) appear much less biased than estimates for the smaller emissions ($\leq$ 1 kg/hr), where we see systematic underestimation. This suggests that the framework is better able to quantify larger emissions, which is expected as these emissions have a stronger signal (more separation from baseline and noise) and hence should be easier to quantify. Note that true emission rate is the primary driver of quantification accuracy. There is very little relationship between quantification error and emission duration, wind speed, or wind direction variability. See Section S7 in the supporting information file for details.

Next consider the R^2 value of the best fit linear model. A high R^2 value indicates that any single emission rate estimate is likely to be very close to the average relationship between the

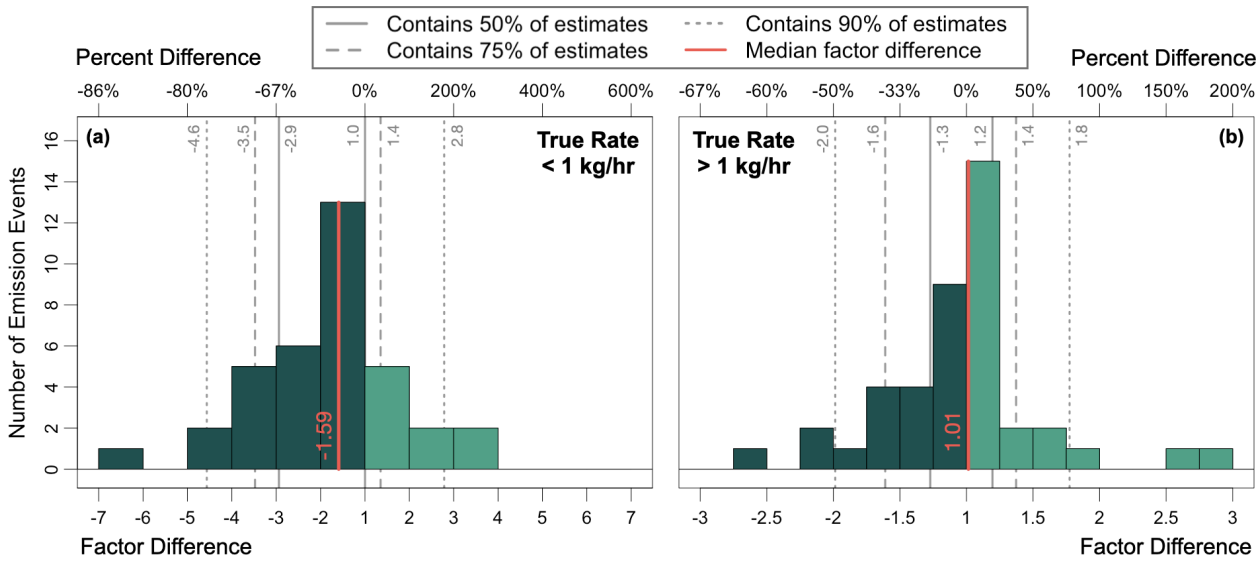

Figure 5: Relative error between estimated and true emission rate, separated by small ($\leq$ 1 kg/hr) and large ($>$ 1 kg/hr) true emission rates. Factor difference is shown on the bottom axis, with negative values corresponding to underestimation, and percent difference is shown on the top axis. Note the different scales between the two subfigures. Vertical gray lines show the error bounds of the middle 50, 75, and 90% of all individual estimates, and vertical red line shows the median error. Darker shading indicates underestimation and lighter shading indicates overestimation.

true and estimated rates represented by the best fit line. The proposed framework results in a best fit line with $R^2 = 0.68$. Note that false positives and false negatives are excluded when computing the best fit coefficients, as we want to use these values to diagnose error in the estimated emission events that properly align with a controlled release.

Figure 5 further explores the difference in quantification accuracy between small ($\leq$ 1 kg/hr) and large ($>$ 1 kg/hr) emissions. Subfigures (a) and (b) show histograms of the relative error between estimated and true rates for small and large emissions, respectively. The factor difference is shown along the bottom axis, with negative values meaning underestimation, and percent difference is shown along the top axis. Both factor and percent difference are measures of relative error, but unlike percent difference, factor difference is not bounded by -100% and hence we think better captures underestimation error. Note that a factor difference of 1 indicates perfect alignment, while a percent difference of 0% indicates perfect alignment.

Estimates for small emissions tend to underestimate, with a median factor difference of -1.59 (-37.2%). Estimates for large emissions are nearly unbiased with a very slight tendency to overestimate, having a median factor difference of only 1.01 (1.5%). For small emissions ($\leq$ 1 kg/hr), 90% of the estimates have error within a factor difference of [-4.6, 2.8] or a percent difference of [-78.1%, 178.6%] from the true rate. The performance is markedly better for the large emissions ($>$ 1 kg/hr), where 90% of the estimates have error within a factor difference of [-2.0, 1.8] or a percent difference of [-49.6%, 77.4%] from the true rate.

Finally, we consider cumulative quantification error over the one-month study period. When considering only small emissions ($\leq$ 1 kg/hr), the framework underestimates cumulative emissions by 41.9%. When considering only large emissions ($>$ 1 kg/hr), the framework overestimates cumulative emissions by 0.8%. These percent differences are close to the median bias values discussed earlier (-37.2% for small emissions and 1.5% for large emissions), which is expected, as there were no event-level false negatives during the one-month study period. The presence of false negatives would cause the gap between estimated and true cumulative emissions to be larger than the median bias values, as the bias values are based on correctly identified events only. Overall, when including both false positives and false negatives and both small and large emissions, the framework underestimates cumulative emissions by 6.7%. See Section S8 in the supporting information file for details.

4 Discussion

We have proposed a framework for emission detection, localization, and quantification on production-like oil and gas facilities that uses concentration data from a point-in-space sensor network. Performance of the framework (especially for emissions $>$ 1 kg/hr) suggests that it could be used to monitor methane emissions and provide informative alerts to oil and gas operators in near real-time. Note, however, that larger, more complex facilities (e.g., midstream compressor stations) will likely require a more nuanced monitoring and modeling

approach. Large buildings and equipment groups block the transport of methane and introduce downwash effects, both of which are not captured by the Gaussian puff model, and hence a more sophisticated atmospheric transport model may be necessary for these facilities. Additionally, a single ring of sensors around the perimeter of larger, complex facilities likely provides insufficient signal separation, as it becomes harder to distinguish between two nearby sources the further away the sensors are placed. Hence, for monitoring purposes, these sites may need to be divided into smaller sectors that are each surrounded by sensors.

Future work will introduce methods for identifying multi-source and off-site emissions. Multi-source emissions can be modeled as the sum of two or more single source emissions (as methane concentrations from the Gaussian puff model are linearly additive), and hence would require no additional simulations to accommodate. However, this will notably increase the search space during localization and quantification and requires further study to assess practical feasibility. Off-site emissions are an important consideration for dense production settings (e.g., the Permian basin) and could be identified by simulating from a ring of potential sources surrounding the facility. Events localized to any of these off-site sources could be excluded from the cumulative emissions tally. The fidelity of this strategy would need to be balanced against the extra computational cost of simulating from many additional sources.

Finally, we note that this framework could be used to update or even replace traditional bottom-up inventories for production-like facilities by computing cumulative emissions on a, e.g., monthly cadence. This use-case will be demonstrated in future papers.

Acknowledgement

The authors thank Project Canary for sharing their raw sensor data.

Supporting Information Available

The following files are available free of charge.

- Supporting information: additional details on the spike detection algorithm and framework performance.

References

(1) Collins, W. J.; Webber, C. P.; Cox, P. M.; Huntingford, C.; Lowe, J.; Sitch, S.; Chadburn, S. E.; Comyn-Platt, E.; Harper, A. B.; Hayman, G.; Powell, T. Increased importance of methane reduction for a 1.5 degree target. *Environmental Research Letters* **2018**, *13*, 054003.

(2) Schleussner, C. F.; Rogelj, J.; Schaeffer, M.; Lissner, T.; Licker, R.; Fischer, E. M.; Knutti, R.; Levermann, A.; Frieler, K.; Hare, W. Science and policy characteristics of the Paris Agreement temperature goal. *Nature Climate Change 2016 6:9* **2016**, *6*, 827–835.

(3) O'Rourke, P. R.; Smith, S. J.; Mott, A.; Ahsan, H.; McDuffie, E. E.; Crippa, M.; Klimont, Z.; McDonald, B.; Wang, S.; Nicholson, M. B.; Feng, L.; Hoesly, R. M. CEDS v-2021-02-05 Emission Data 1975-2019. `http://doi.org/10.5281/zenodo.4509372`, 2021; Zenodo.

(4) Crippa, M.; Guizzardi, D.; Muntean, M.; Schaaf, E.; Lo Vullo, E.; Solazzo, E.; Monforti-Ferrario, F.; Olivier, J.; Vignati, E. EDGAR v6.0 Greenhouse Gas Emissions. `http://data.europa.eu/89h/97a67d67-c62e-4826-b873-9d972c4f670b`, 2021; European Commission, Joint Research Centre (JRC).

(5) EPA, *Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2020*; 2022.

(6) Allen, D. T.; Cardoso-Saldaña, F. J.; Kimura, Y. Variability in Spatially and Temporally Resolved Emissions and Hydrocarbon Source Fingerprints for Oil and Gas Sources in Shale Gas Production Regions. *Environmental Science and Technology* **2017**, *51*, 12016–12026.

(7) Lavoie, T. N.; Shepson, P. B.; Cambaliza, M. O.; Stirm, B. H.; Conley, S.; Mehrotra, S.; Faloona, I. C.; Lyon, D. Spatiotemporal Variability of Methane Emissions at Oil and Natural Gas Operations in the Eagle Ford Basin. *Environmental Science and Technology* **2017**, *51*, 8001–8009.

(8) Vaughn, T. L.; Bell, C. S.; Pickering, C. K.; Schwietzke, S.; Heath, G. A.; Pétron, G.; Zimmerle, D. J.; Schnell, R. C.; Nummedal, D. Temporal variability largely explains top-down/bottom-up difference in methane emission estimates from a natural gas production region. *Proceedings of the National Academy of Sciences of the United States of America* **2018**, *115*, 11712–11717.

(9) Brandt, A. R.; Heath, G. A.; Cooley, D. Methane Leaks from Natural Gas Systems Follow Extreme Distributions. *Environmental Science and Technology* **2016**, *50*, 12512–12520.

(10) O'Connell, E.; Risk, D.; Atherton, E.; Bourlon, E.; Fougère, C.; Baillie, J.; Lowry, D.; Johnson, J. Methane emissions from contrasting production regions within Alberta, Canada: Implications under incoming federal methane regulations. *Elementa* **2019**, *7*.

(11) Robertson, A. M.; Edie, R.; Snare, D.; Soltis, J.; Field, R. A.; Burkhart, M. D.; Bell, C. S.; Zimmerle, D.; Murphy, S. M. Variation in Methane Emission Rates from Well Pads in Four Oil and Gas Basins with Contrasting Production Volumes and Compositions. *Environmental Science and Technology* **2017**, *51*, 8832–8840.

(12) Caulton, D. R.; Lu, J. M.; Lane, H. M.; Buchholz, B.; Fitts, J. P.; Golston, L. M.; Guo, X.; Li, Q.; McSpiritt, J.; Pan, D.; Wendt, L.; Bou-Zeid, E.; Zondlo, M. A. Importance of Superemitter Natural Gas Well Pads in the Marcellus Shale. *Environmental Science and Technology* **2019**, *53*, 4747–4754.

(13) Cusworth, D. H.; Duren, R. M.; Thorpe, A. K.; Olson-Duvall, W.; Heckler, J.; Chapman, J. W.; Eastwood, M. L.; Helmlinger, M. C.; Green, R. O.; Asner, G. P.; Den-

nison, P. E.; Miller, C. E. Intermittency of Large Methane Emitters in the Permian Basin. *Environmental Science and Technology Letters* **2021**, *8*, 567–573.

(14) Alvarez, R. A. et al. Assessment of methane emissions from the U.S. oil and gas supply chain. *Science* **2018**, *361*, 186–188.

(15) Brandt, A. R. et al. Methane leaks from North American natural gas systems. *Science* **2014**, *343*, 733–735.

(16) Karion, A. et al. Methane emissions estimate from airborne measurements over a western United States natural gas field. *Geophysical Research Letters* **2013**, *40*, 4393–4397.

(17) Pétron, G. et al. A new look at methane and nonmethane hydrocarbon emissions from oil and natural gas operations in the Colorado Denver-Julesburg Basin. *Journal of Geophysical Research: Atmospheres* **2014**, *119*, 6836–6852.

(18) Wang, J. L.; Daniels, W. S.; Hammerling, D. M.; Harrison, M.; Burmaster, K.; George, F. C.; Ravikumar, A. P. Multiscale Methane Measurements at Oil and Gas Facilities Reveal Necessary Frameworks for Improved Emissions Accounting. *Environmental Science & Technology* **2022**, *56*, 14743–14752.

(19) Zavala-Araiza, D. et al. Reconciling divergent estimates of oil and gas methane emissions. *Proceedings of the National Academy of Sciences of the United States of America* **2015**, *112*, 15597–15602.

(20) Yarmuth, J. A. H.R.5376 - 117th Congress (2021-2022): Inflation Reduction Act of 2022. `http://www.congress.gov/`, 2022.

(21) Jacob, D. J.; Turner, A. J.; Maasakkers, J. D.; Sheng, J.; Sun, K.; Liu, X.; Chance, K.; Aben, I.; McKeever, J.; Frankenberg, C. Satellite observations of atmospheric methane and their value for quantifying methane emissions. *Atmospheric Chemistry and Physics* **2016**, *16*, 14371–14396.

(22) Fox, T. A.; Barchyn, T. E.; Risk, D.; Ravikumar, A. P.; Hugenholtz, C. H. A review of close-range and screening technologies for mitigating fugitive methane emissions in upstream oil and gas. *Environmental Research Letters* **2019**, *14*, 053002.

(23) Schwietzke, S. et al. Improved Mechanistic Understanding of Natural Gas Methane Emissions from Spatially Resolved Aircraft Measurements. *Environmental Science and Technology* **2017**, *51*, 7286–7294.

(24) Alden, C. B.; Coburn, S. C.; Wright, R. J.; Baumann, E.; Cossel, K.; Perez, E.; Hoenig, E.; Prasad, K.; Coddington, I.; Rieker, G. B. Single-Blind Quantification of Natural Gas Leaks from 1 km Distance Using Frequency Combs. *Environmental Science and Technology* **2019**, *53*, 2908–2917.

(25) Bell, C.; Ilonze, C.; Duggan, A.; Zimmerle, D. Performance of continuous emission monitoring solutions under single-blind controlled testing protocol. *ChemRxiv* **2022**,

(26) Keats, A.; Yee, E.; Lien, F. S. Bayesian inference for source determination with applications to a complex urban environment. *Atmospheric Environment* **2007**, *41*, 465–479.

(27) Sharan, M.; Issartel, J. P.; Singh, S. K.; Kumar, P. An inversion technique for the retrieval of single-point emissions from atmospheric concentration measurements. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **2009**, *465*, 2069–2088.

(28) Luhar, A. K.; Etheridge, D. M.; Leuning, R.; Loh, Z. M.; Jenkins, C. R.; Yee, E. Locating and quantifying greenhouse gas emissions at a geological CO2 storage site using atmospheric modeling and measurements. *Journal of Geophysical Research: Atmospheres* **2014**, *119*, 10,959–10,979.

(29) Feitz, A. et al. The Ginninderra CH4 and CO2 release experiment: An evaluation of gas detection and quantification techniques. *International Journal of Greenhouse Gas Control* **2018**, *70*, 202–224.

(30) Kumar, P. et al. Near-field atmospheric inversions for the localization and quantification of controlled methane releases using stationary and mobile measurements. *Quarterly Journal of the Royal Meteorological Society* **2022**, *148*, 1886–1912.

(31) Travis, B.; Dubey, M.; Sauer, J. Neural networks to locate and quantify fugitive natural gas leaks for a MIR detection system. *Atmospheric Environment: X* **2020**, *8*, 100092.

(32) Jia, M.; Daniels, W. S.; Hammerling, D. M. Efficient implementations of the Gaussian puff atmospheric dispersion model applied to methane emissions on oil and gas facilities. *ChemArxiv* **2022**,

(33) Project Canary. `https://www.projectcanary.com/`, 2022.

TOC Graphic

Some journals require a graphical entry for the Table of Contents. This should be laid out "print ready" so that the sizing of the text is correct.
Inside the `tocentry` environment, the font used is Helvetica 8 pt, as required by *Journal of the American Chemical Society*.
The surrounding frame is 9 cm by 3.5 cm, which is the maximum permitted for *Journal of the American Chemical Society* graphical table of content entries. The box will not resize if the content is too big: instead it will overflow the edge of the box.
This box and the associated title will always be printed on a separate page at the end of the document.