Point Sensor Network Detects Short Releases Under Favorable Wind Conditions

David Ball, Ph.D[1], Ali Lashgari, Ph.D[1], Umair Ismail, Ph.D[1], Noah Metzger[1], and Nathan Eichenlaub[1]

[1]Project Canary

July 2024

Abstract

In this study, we apply a recursive Bayesian state updater algorithm to assimilate meteorological data with point-sensor measurements of methane concentrations to infer timeseries of methane emission rates at three operating oil and gas facilities. These calculations are performed over a timeframe with known numerous short ($\sim$ 30 minute) controlled releases, allowing for "ground truth" data to compare our emission estimates against. The highly-varying and unknown operational emissions pose challenges in analyzing quantification results when trying to determine whether there is evidence in the emission estimates of a given controlled release. Ultimately, we find that despite the non-ideal conditions at these sites (poor sensor placement and the presence of large obstructions that the quantification model does not account for) that the site-level emission estimates show evidence for 31 out of the 60 controlled releases and that the majority of nondetections were due to the wind simply not pointing from the source to any sensor in the network during a short release event.

1 Introduction

Methane has a shorter atmospheric lifetime, but significantly higher heat-trapping potential compared to carbon dioxide (CO2). As a result, methane reduction actions taken today will have a more significant short-term impact than CO2 reduction efforts. Human activities are responsible for approximately 60% of global methane, with the oil and gas, agriculture (including fermentation, manure management, and rice cultivation), landfills, wastewater treatment, and coal mining industries as the largest contributors to anthropogenic methane emissions [1].

In 2022, natural gas and petroleum systems were responsible for 28% of US methane emissions [1]. Tackling oil and gas methane emissions is considered as one of the most economically and technically effective ways of lowering greenhouse gas (GHG) emissions in the near term [2]. In addition to its climate impact, leaked and vented methane from the oil and gas sector results in substantial economic losses in terms of lost sales and social costs [3].

Several studies using various emissions measurement techniques report that the existing emission inventories underestimate actual methane emissions from the oil and gas sector [4–17] with significant regional variability in the scale of discrepancies between inventories and measurement-based estimates. A recent study reported that methane emission loss rates range from approximately 0.75% to 9.63% of natural gas production in different geographical areas in the US [18]. The wide range of methane emissions intensity in various geographical regions and among different facility categories and operators indicates that the effectiveness of methane mitigation strategies relies on a comprehensive understanding of emission events characteristics, including source, size (emission rates), duration, and frequency of emission events, which can only be understood via direct measurement.

Direct methane emissions measurement can be done from space (using satellites), in the air (by drones or piloted aircraft), or on the ground (by point sensor networks, laser-based scanning, or imaging systems). Due to the limited revisit rate of space and aerial-based measurement technologies, these platforms can identify emissions above the minimum detection limit for the specific method but

offer little to no information regarding the frequency or duration of events. In principle, continuous monitors can address the shortcomings associated with the limited observation frequency of space and aerial-based measurements, however the detection, localization, and quantification algorithms for these systems are highly dependent on gas dispersion models, which introduces significant uncertainty in their emissions estimates. Additionally, these systems rely on the wind to "cooporate", i.e., point from emission sources towards sensors in the network with sufficient frequency to be able to make reliable estimates.

Recent studies suggest that 72% of the continental US midstream and upstream oil and gas methane emissions originate from facilities that emit less than 100 kg/hr, with 30% of total emissions coming from facilities emitting less than 10 kg/hr [15]. With current satellite technology, the lower detection limit is at best around 100 kg/hr, but in many cases can be significantly higher [19]. These measurements are impacted by atmospheric conditions (including cloud cover), solar geometry, and surface characteristics [20]. As a result, relying solely on point-in-time satellite methane observations results in an incomplete estimation of the emissions spectrum for both regional and site-level emissions inventories.

Under ideal measurement conditions, aircraft methane measurement methods can detect emissions below 100 kg/h, and in some cases, below 10 kg/h [21]. These methods are highly weather-dependent and due to the snapshot nature of these methods, there is no practical way to determine emission event start and end times. When attempting to calculate regional emissions inventories, these methods rely on statistical methods to estimate event duration and frequency using large sample sizes, which introduces significant uncertainty that depends highly on the specific revisit frequency and detection limit of the deployed technology. In addition, down-scaling basin-level estimates that use large-scale sample sizes to determine site-level emissions measurements is challenging [22, 23]. Recent studies recommend using operational data and continuous monitoring as two useful methods to complement snapshot methods in estimating emission event duration and frequency [23, 24].

Continuous monitoring systems (CMS) have been rapidly evolving since the inception of the U.S. Department of Energy (DOE) ARPA-E MONITOR program. Since these systems provide real-time or near-real-time monitoring of facility emissions at a far higher frequency compared to intermittent methods, they have the potential to address temporal variability and offer better insight into emissions event frequency and duration. It should be noted that other variables such as wind variability, stability class of the atmosphere, and facility complexity may impact the performance of CMS in detecting, localizing, and quantifying emission events. With sufficient sensor density and placement, proper measurement sensitivity, and advanced data assimilation algorithms, CMS have proven to be an effective technology in detecting, localizing, and quantifying emission events.

Recent studies suggest that employing multiple types of measurements is likely to improve the accuracy of emissions estimates. This approach is known as multi-scale measurement, where in the case of developing measurement-based site-level inventories, a CMS is used for validation of specific snapshot measurements and determining how snapshot measurements relate to the temporal emission profile of a given site [25, 26]. This approach requires independent testing of various technologies to comprehensively understand the applications and limitations of each emission measurement method and validate the quality of event detections and emissions estimates.

In 2016, the DOE sponsored the creation of a facility at Colorado State University (CSU), operated by the Energy Institute's Methane Emissions Technology Evaluation Center (METEC), to simulate the equipment configurations on oil and natural gas production sites with controlled release capabilities. The METEC facility is frequently used to evaluate the performance of methane measurement technologies for emissions event detection, localization, and quantification. While testing at METEC is valuable in assessing the performance of various technologies, it represents something of a best-case-scenario: the facility is small relative to operational sites, there is very little terrain variation or obstructive features (buildings, large tank batteries, numerous compressors), and solutions often deploy a higher number of sensors than they would in the field. As such, when considering published results from studies such as [27], one should consider them to be the most optimistic assessment of each solution's capabilities, and take care in extrapolating this performance to the field.

A team affiliated with METEC performed an extensive study, [28], in which they took a mobile controlled-release rig to 11 operational facilities equipped with methane point-sensor networks from several companies, performed controlled release tests with varying locations and rates, analyzed the data from the associated dashboards, and attempted to characterize the probability of detection (PoD)

and quantification accuracy of the solutions. They dubbed these "challenge releases" as they are not performed in a controlled environment, and the presence of significant background makes them challenging to detect. We will henceforth adopt this nomenclature for the releases from this study. The primary focus of their analysis was on solutions that provided site-level emissions estimates. At the time of this study, Project Canary (henceforth, PC) did not have quantification enabled in their dashboard at these sites, and as such, were excluded from the bulk of the analysis and discussion from this study, which focused on assessing each solutions' PoD and quantification accuracy.

Controlled release experiments at operational facilities are incredibly important to fully understand the performance of point-sensor networks and associated algorithms in a field setting: the majority of controlled-release studies are performed in relatively sterile environments (e.g., METEC), and do not fully replicate the complexity of emissions and the logistics of deploying point sensors at operational sites. Testing the accuracy of these solutions in conditions that more closely mimic an operational environment is paramount to understanding their performance when deployed. Such testing, however, poses additional challenges: because solutions are distributed across many sites (rather than all deployed at one central testing location like METEC), it takes a significant amount of time and person-power to conduct an appreciable number of these studies across different sites and associated solutions. As such, it is impractical to perform the same number of releases to build up robust statistics (such as the ADED campaigns, where they typically employ ~ 500 unique releases, which are seen by all of the participating solutions). Another major challenge is the presence of unknown operational emissions that may fluctuate on timescales comparable to the duration of controlled release experiments. These unknown emissions make it difficult to disentangle a signal associated with a controlled release from the background, and extra care must be taken in exactly how this background subtraction is performed. In contrast, at a facility like METEC, there are no interfering nearby sources with significant impact on the background methane levels, so any methane concentration enhancement or associated emissions estimate can be directly attributed to the controlled releases, naturally giving rise to an evaluation framework in which one can distinguish true event detections from false alarms and missed event detections.

In [28], they define two distinct frameworks for defining whether a given solution detected a controlled release. The first step in both processes is to estimate a static background level (BL) of the site-level emissions. This is accomplished by taking a period of time (excluding times with controlled releases), that varies from 1-6 weeks depending on the availability of data from a given solution, and computing the mean of the solution's estimate of site-level emissions over this timeframe. In their PoD analysis, if a solution's estimate exceeds the baseline (which turns out to be 0 for many of the solutions deployed during this study) during the time that a controlled release occurs, then this is considered a detection of the release. In another framework they use for generating a classification matrix and calculating the associated χ^2 to determine whether the data rules out a relationship between emissions estimates and controlled releases, they define the detection threshold to be one standard deviation above the baseline value (where the standard deviation is computed over the same timeframe as the mean baseline).

As noted before, PC was not evaluated as one of the solutions in [28] that provided quantification estimates. We aim to run a similar post-facto analysis by applying PC's current quantification algorithms to historical data over the time period that these studies were performed and evaluate whether PC's emission estimates show evidence for upticks in site-level emissions that are consistent with the controlled releases.

2 Historical Quantification Estimates at Controlled Release Sites

PC had sensors deployed at 3 of the 11 sites studied in [28] numbered 9, 10, and 11. Figures 1-3 show the relative positioning of site equipment, color coded by equipment type, with circles, PC's sensors with black x's, and the challenge release locations with red triangles. The associated wind roses computed from the months of January and February (the months when these challenge releases took place) are shown in the right panels. A visual inspection of the sensor location with respect to the dominant wind directions at these sites suggests non-ideal placement. For example, considering Site 10 (Figure 2), the dominant wind directions are from the southwest/west, yet no sensors are deployed

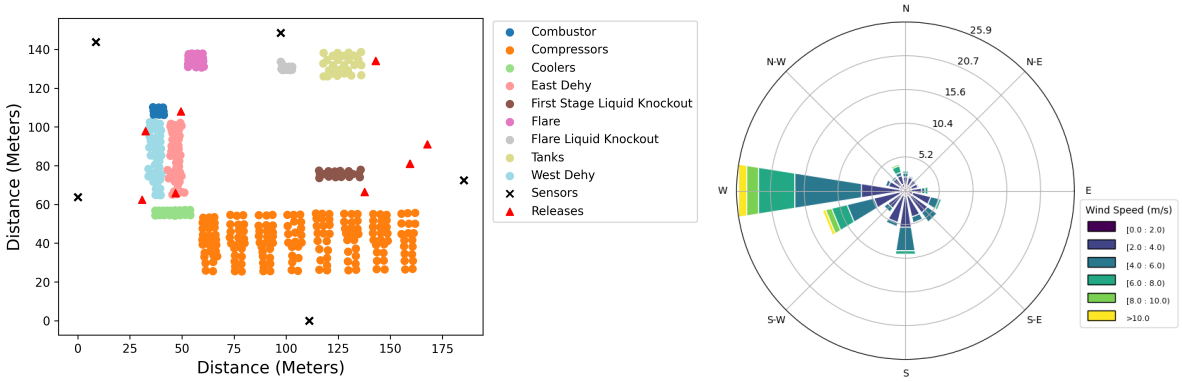

(a) (b)

Figure 1: (left) Schematic of site equipment (colored dots), sensor positions (black x's) and challenge release locations (red triangles) and associated wind rose (right) of Site 9.

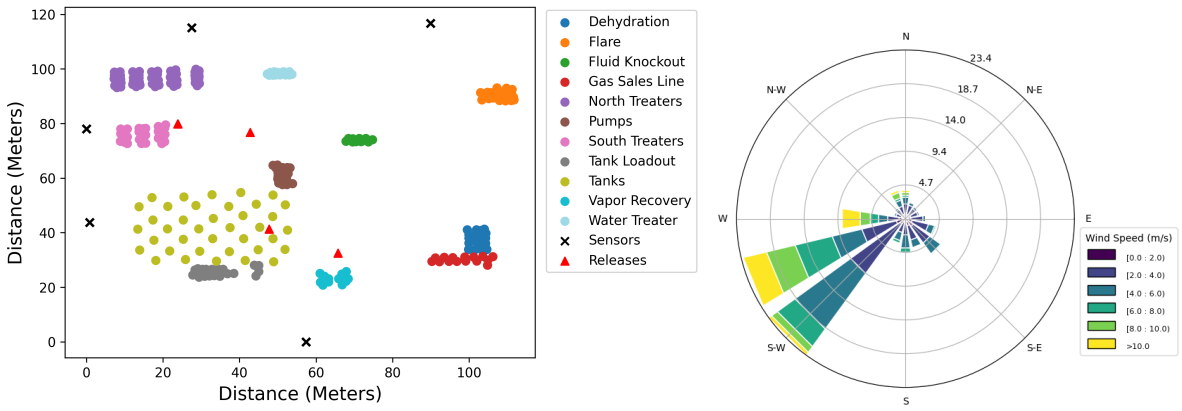

(a) (b)

Figure 2: (left) Schematic of site equipment (colored dots), sensor positions (black x's) and challenge release locations (red triangles) and associated wind rose (right) of Site 10.

on the east side of the facility, while two are placed along the western fenceline.

One anticipated effect of poor sensor placement is that the system will be less likely to detect a given emission event. This impact is pronounced if release events are short, as the wind does not have as much time to vary and point from the release source to any sensor in the network. As such, for a set of short releases at these sites (such as those in this study), many may be missed simply due to the wind never blowing towards one of the sensors. It should be noted, however, that if releases are longer (such as a persistent leak at a site), then the system is much more likely to detect the event. To illustrate this, imagine a system that, given its sensor placement and wind direction statistics, has a probability p every minute of detecting an enhancement from a given source on the site (for argumentative purposes, let's assume that this probability encapsulates both the wind blowing from source to sensor and the sensor's response informed by its detection limit). The probability that this system detects evidence of a m-minute long release is simply $1-(1-p)^m$. For a system with poor sensor positioning, p will be quite small, as the probability that the wind blows from source to sensor is low. Let us consider a system with $p = 0.01$. The probability of such a system detecting a 30 minute long release is 26 percent (a fairly low detection probability). If the release represents a persistent problem that is not mitigated until it is physically addressed, then the probability of detecting it within a day is almost guaranteed: plugging in 60×24 for m (the number of minutes in a day), the probability of detecting the emissions event is 99.9999 percent. While this is a simplistic argument and does not fully represent the response of a quantification system to input concentration and wind data, it highlights the effect of release length on detection probability.

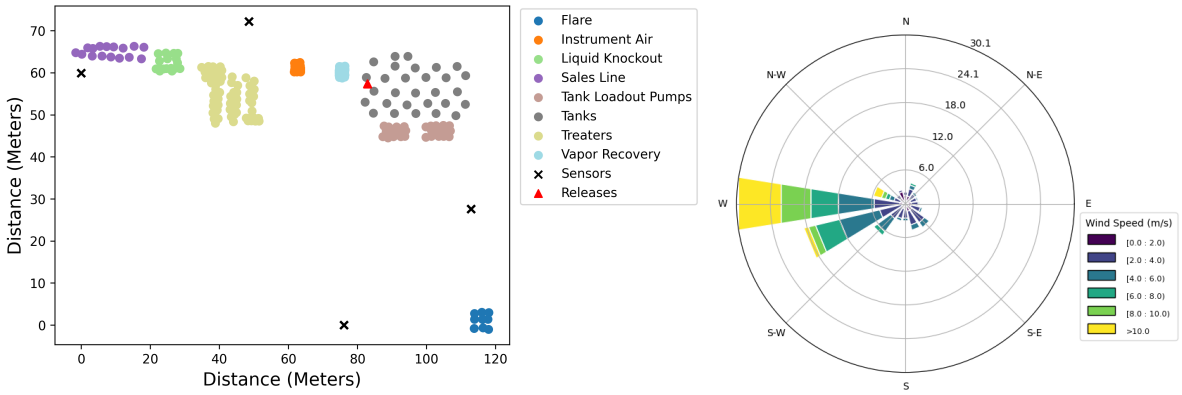

(a) (b)

Figure 3: (left) Schematic of site equipment (colored dots), sensor positions (black x's) and challenge release locations (red triangles) and associated wind rose (right) of Site 11.

In general, PC does not recommend performing quantification on sites that contain a significant amount of complex infrastructure (such as the several large compressor stations at Site 9 or the large tank battery at Site 10), due to the limited capability of computationally-tractable (i.e., ones that can run in real-time) physical dispersion models to account for these obstructions. Furthermore, the sensor placement at these sites was dubious, with many sensors being deployed upwind of sources when considering the statistically dominant wind directions. For these reasons, a sub-optimal quantification performance was expected at these sites. Nevertheless, we thought it would be illuminating to analyze the results of running PC's quantification algorithm against these challenge releases. Figure 4 shows the output of PC's quantification algorithm applied to each site over a 3-week time-span centered on the challenge release dates for that specific site. This timeframe was selected to be representative of what [28] used to estimate the background level and standard deviation in site level emissions that are employed to calculate the detection threshold.

Figure 4 depicts the mean of the site-level quantification estimate over each timeframe (excluding controlled release times) with an orange dashed line, and the mean plus one standard deviation with a dashed red line. These correspond to the detection criteria defined in [28] using these 3 weeks of data to estimate the background. A visual inspection of these plots leads to some immediate and important observations: (i) **these sites have substantial variation in their operational emissions and the characteristic scale of these variations can be larger than controlled release rates**, (ii) **if the controlled releases happen during a period of relatively low operational emissions when compared to the time period that the baseline was computed over, the previously-described detection algorithm will always yield non-detections because the baseline is artificially high.** For example, in sub panels (a) and (c) of Figure 4 the mean baseline (orange line) over the 3-week timeframe is substantially higher than the site-level emissions right before the controlled releases start. This is entirely due to the fact that the timespan over which the baselines were computed included some high-rate emissions events that are not present during the time of the controlled releases. As such, the baseline for these two sites is overestimated by a factor of about 5 when considering the timeframe associated with the controlled release experiments, which will invariably result in the detection methodology described in [28] erroneously determining that the solution did not detect controlled releases, even if there was clear evidence of increased rates during challenge release times.

3 Alternative Baseline Calculation and Definition of Detection

We propose an alternative to computing the baseline over several weeks: use the emissions estimate associated with the timestep immediately prior to the beginning of a series of challenge releases. The challenge releases performed in this study tend to be clustered together, often occurring back-to-back, with little to no time gap in between consecutive short releases. For this reason, when estimating the

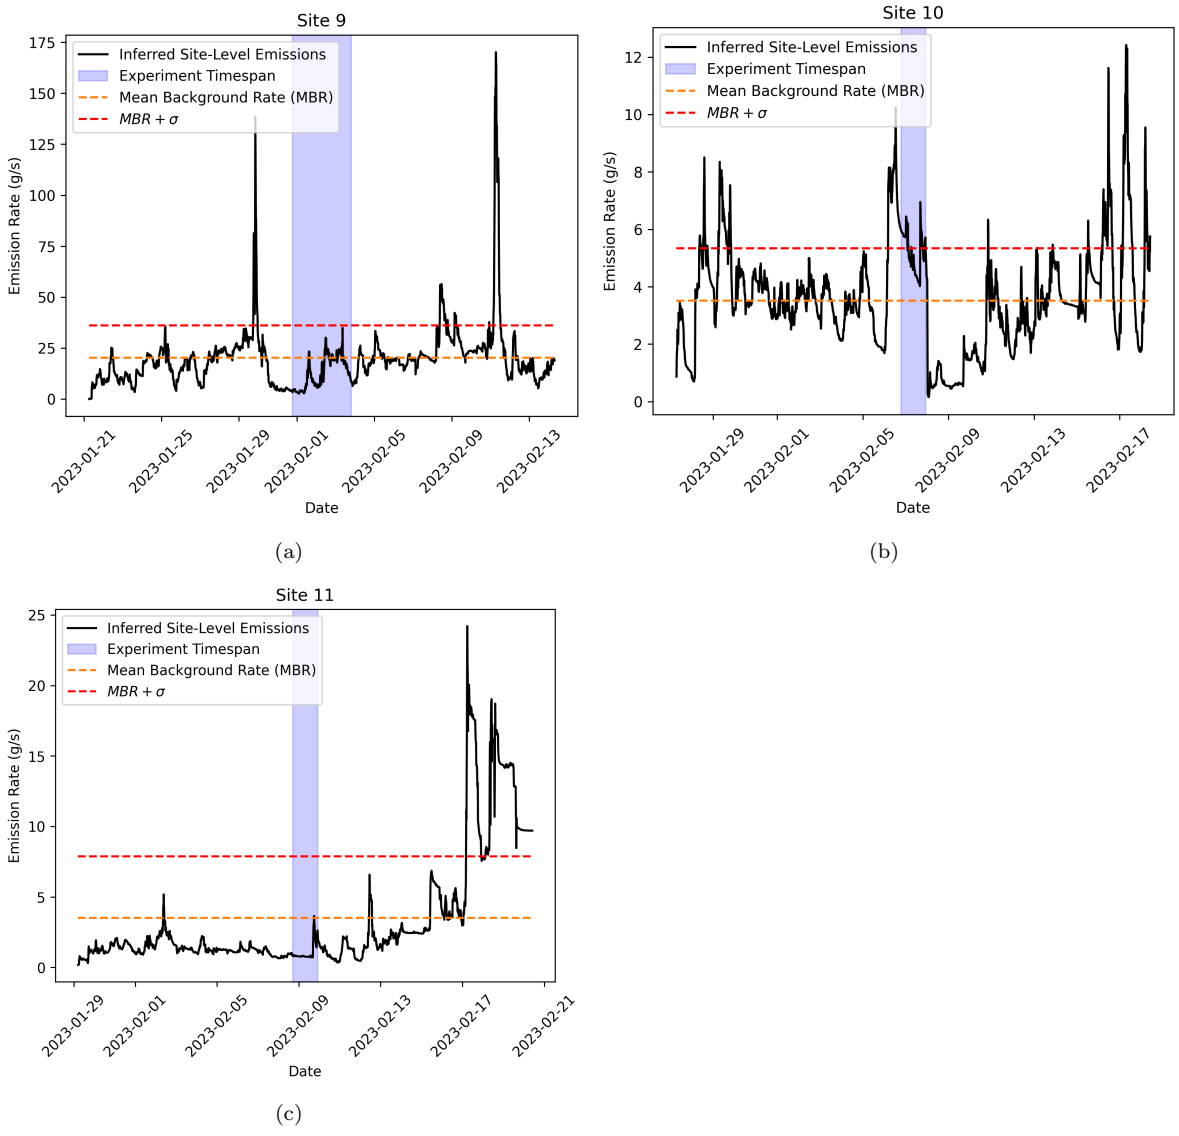

Figure 4: Site-level emissions estimate (black line) over 3-week timespan centered on controlled release experiments for Sites 9 (a) 10 (b) and 11 (c). The mean emission rate (excluding the controlled release time window) over this timespan is depicted with a dashed orange line and the mean plus one standard deviation is shown with a dashed red line. Significant variation in operational emissions result in computed baselines that are highly dependent on the choice of time window and do not necessarily reflect the actual operational background during the time of controlled release experiments.

background level, we first apply a temporal clustering algorithm to identify groups of releases that are closely spaced in time. We then identify the timestep immediately prior to this grouping and use the emissions estimate at that time as the estimated background value for following group of challenge releases. In this way, the baseline represents the site-level emissions at a time relevant to the controlled release experiment, rather than aggregated over a large timeframe with varying emissions and potential events (e.g., blowdowns, leaks, and fugitives) that may drive up the baseline values in a way that is inconsistent with the actual baseline during the time period of interest.

Rather than calculating the standard deviation of the baseline emissions and defining a threshold based on the mean baseline and standard deviation (which may be highly affected by a few large emissions events outside of the timeframe of interest), we propose tying the detection threshold directly to the controlled release rate. A release is considered "detected" if the quantification estimate exceeds the baseline by 25% of the controlled release rate at any point during the release. In this way, the detection threshold is defined by a more temporally-relevant baseline value (taken right before the controlled release start) and tied to the physical release rate.

4 Release Detection Results

Figures 5-7 show the results of applying the alternative detection scheme described in Section 3 to the quantification estimates from PC. In these figures, the black line depicts PC's site-level emissions estimate and the shaded regions highlight controlled releases, with orange corresponding to controlled releases that were not detected (i.e., the black line showed no evidence of an uptick beyond the background values, which are depicted with orange dots), and green corresponding to controlled releases that were detected (i.e., the black line showed evidence of an uptick over the background value by at least 25% of the controlled release rate). The controlled release rate is depicted with blue dashes.

There is substantial variation in the background at Site 9 (Figure 5), making it particularly challenging to identify evidence of short releases. Nevertheless, there appears to be evidence of increased emissions during half of the releases. We note that the significant variation in the unknown background on short timescales may trigger positive detections that are not actually associated with a challenge release, or result in missed detections if the background decreases during a challenge release. This is an unavoidable challenge of these setups, but may be addressed via a statistical analysis of a much larger sample of releases, performing studies with longer releases, or ideally knowing the background perfectly in a more sophisticated controlled release experiment that includes controlled time-varying emissions to mimic an operational site.

At Site 10 (Figure 6), PC's quantification estimates show evidence for 8 out of the 12 unique releases. The first set of contiguous releases are missed entirely, due to poor sensor placement and the wind never pointing to any of the sensors during the $\sim$ 4 hour period (see Section 5 for an analysis of wind direction with respect to nondetections). The background variation at this site is much less significant than at Site 9 (the standard deviation in the emissions estimate is less than the controlled release rates), making it fairly obvious to see why the algorithm determined that 8 out of the 12 releases were detected: during the first set of releases, the background is estimated to be at about 5.5 grams per second (g/s), and the site-level emission estimate stays within 0.1 g/s of this value during the entirety of the challenge release experiments on this day. In the second set of challenge releases, however, the background is estimated to be around 4.5 g/s. As soon as the first release begins, there is an obvious uptick in the emissions estimate by about 4 g/s. The emissions estimate remains elevated above the background during the entirety of the challenge releases and declines shortly thereafter. Note that the quantification estimates for these challenge releases are significantly lower than the release rates: the uptick above the background for these releases ranges from 1-4 g/s, while the controlled release rates are between 2 and 7.5 g/s. As we will elaborate on in Section 6, this is likely due the presence of the large tank battery, a significant obstruction, the effects of which are not currently accounted for in this version of PC's quantification algorithm.

Considering Site 11 (Figure 6), there are three windows of contiguous releases during which the emission rate varies a total of 20 times. Of these 20 releases, PC's quantification estimates show evidence for 9 of them. During the first grouping of contiguous releases (from roughly hours 15 to 22 on February 8th), the wind never points towards any of PC's sensors, and as such, there is no evidence of any emissions on site. During the following two release groupings that occur on the same day, and hence are assigned the same background level, PC's quantification estimate clearly show evidence for

upticks in the emission rate. On this day, the first release is missed. This is likely due to a combination of two factors: the first has to do with the statistical variance of the wind direction and the likelihood that it is pointing from the source to a sensor at any given time: it is not surprising that there may be a delay in the uptick of site-level emissions with respect to the start of the challenge release window, as the wind is not guaranteed to immediately blow the enhanced emissions towards one of the sensors. The second potential reason for a slight delay in the detection is that PC employs a recursive Bayesian estimator that effectively performs a weighted average of the last known state of the system (here, the state of the system is a vector of source rates), and the best state estimate from the new information. As such, it can take a couple recursions for rate changes to fully manifest in the state estimate. For the detected releases, the rate estimates at this site are quite comparable to the calculated upticks; the blue dashes (controlled release rates) look fairly well correlated to the black line (emissions estimate) during the green shaded regions (positive detections).

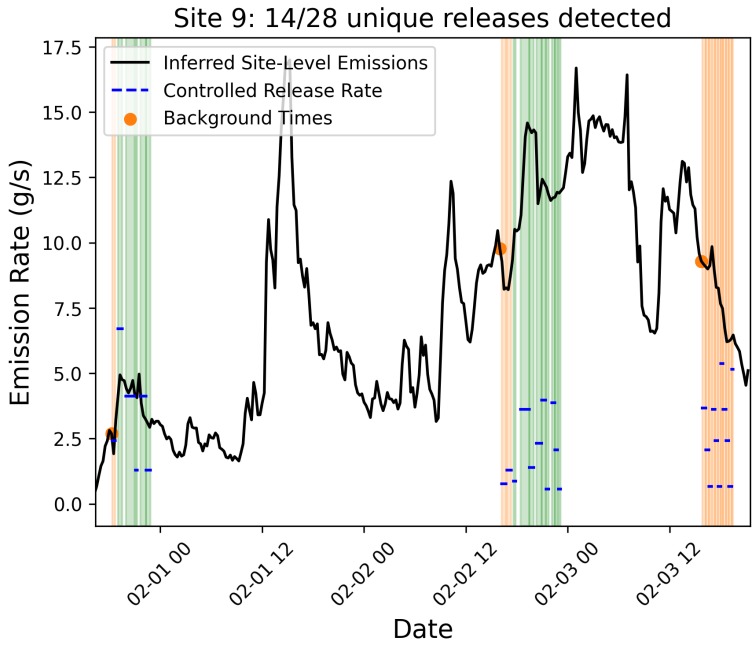

Figure 5: Missed (orange) and detected (green) controlled releases at Site 9.

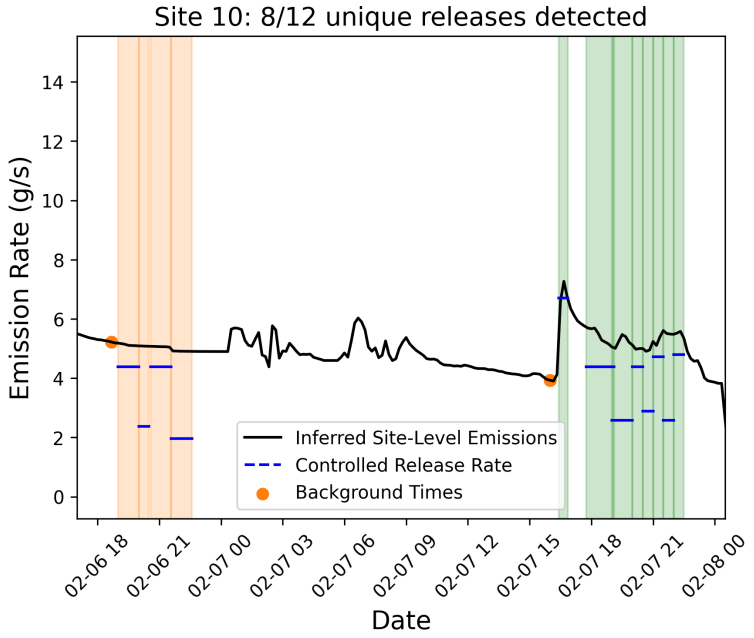

Figure 6: Missed (orange) and detected (green) controlled releases at Site 10.

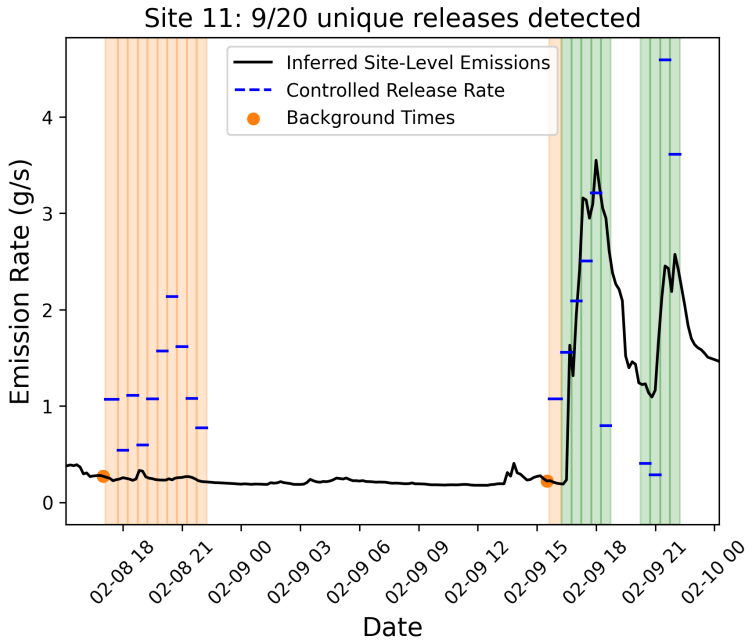

Figure 7: Missed (orange) and detected (green) controlled releases at Site 11.

5 Nondetections and Wind Direction Analysis

In order to determine the precise reason for the nondetection of some of the controlled releases, we examine the statistics of the wind direction with respect to controlled release location and sensor locations. At each minute during every challenge release, we compute the minimum angular separation between the wind direction and the source-sensor angle to every sensor on the site. By performing this calculation on every minute of data during a challenge release and then taking an average, we obtain $\bar{\theta}_{min}$, a statistic that describes the mean angular offset between the downwind direction from the source and the nearest (in angle) sensor. Smaller values of $\bar{\theta}_{min}$ during a challenge release indicate

that the wind was frequently pointing towards the vicinity of a sensor, while a large value indicates that the wind tended not to point towards any sensors. For the purposes of this analysis, we define any challenge release experiment with $\bar{\theta}_{min} < 30$ degrees as a "detectable" experiment, meaning that the wind direction was close enough to being along a source-sensor axis that our system should, in theory, be able to detect the source. If this statistic is greater than 30 degrees, then we deem the release "undetectable" due to the large angular offset between the wind direction and nearest source-sensor axis.

We show a plot of $\bar{\theta}_{min}$ versus the detection binary (0 for nondetection, 1 for detection) for all of the challenge releases, colored by site number in Figure 8. For Sites 10 and 11, PC's system detects every "detectable" release (releases with $\bar{\theta}_{min} < 30$). For Site 9, however, there are several nondetections below this threshold. This could be due to a couple factors: first, the rapidly fluctuating background at this site may cause issues with the background estimations and associated detection framework. For example, if there is an active source with a decreasing emission rate during a challenge release, then the site-integrated emissions estimate may not show evidence for a challenge release because it is counteracted by the (unknown) operational emissions decreasing. Another potential reason why Site 9 may show more nondetections at low $\bar{\theta}_{min}$ is the presence of several large obstructive compressor stations at the facility. As previously mentioned, due to the large obstructions at this facility, applying PC's existing quantification models to the data collected from this site is not generally recommended. Setting aside the detection statistics from Site 9 for the reasons mentioned above, we turn our focus to analyzing the detections from Sites 10 and 11. It is encouraging to see that all of the nondetections from these sites can be attributed to the wind simply not pointing towards sensors. This has two key implications: first, if the sensor placement were better at these sites (such that the characteristic $\bar{\theta}_{min}$ values were smaller), PC's system may have detected substantially more of the challenge releases. Second, this system is clearly capable of detecting the challenge releases at Sites 10 and 11, provided that the wind points towards sensors in the network at some point during the release. This suggests that if the challenge releases were longer, that PC's system may detect a larger fraction of them.

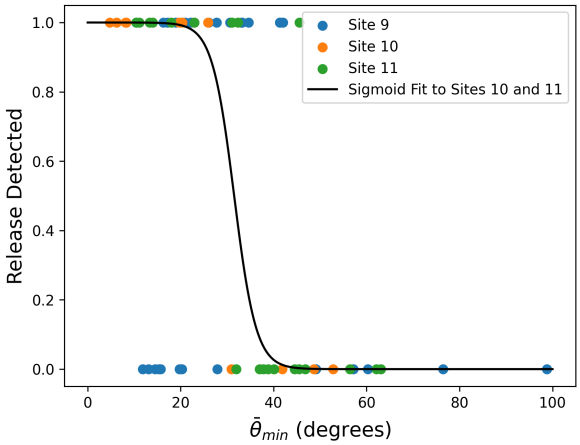

Figure 8: Detection binary (vertical axis) as a function of $\bar{\theta}_{min}$ and best-fit logistic regression applied to detection data from Sites 10 and 11.

In order to test the hypothesis that improved sensor location may yield better detection probability, PC employed their "siting tool", an algorithm that simulates methane emission from every potential source at a site and uses historical atmospheric time series data to produce time-varying concentration field predictions. The algorithm selects sensor locations that maximize the amount of information the sensor network receives from the simulated plumes while minimizing the fractional "blind time" (the fraction of 12 hour periods that every source on a site is not observed by the sensor network a sufficient number of times to generate a reliable quantification estimate). We note that the deployment of sensors at the sites in this study predated PC's development of this tool. Figure 9 depicts the locations of the deployed sensor network at the time of the controlled releases with blue dots and the output of the siting tool with red dots. The potential emission source locations that are used as an input to the siting tool is shown with black dots. Site 9 (left panel) shows a relatively small difference between deployed and optimal sensor locations: this site had reasonably good sensor coverage to begin with, so only

minor changes are evident when comparing the deployed versus optimal positions. At Sites 10 and 11 (middle and right panels), we see a starker contrast between the deployed and optimal locations, and indeed these were the sites that we suspected sensor placement may be more dubious when comparing the wind roses to the deployed sensor locations in Figures 2 and 3. At Site 10, one of the sensors along the western fenceline is effectively moved over to the eastern side to capture emissions when the wind is coming from the west, which is one of the statistically dominant directions, as shown in Figure 2. At Site 11, the most salient difference in sensor location is that the eastmost sensor is moved to the north so that it is positioned directly east from the large groups of equipment present at the site. This also makes logical sense when considering the wind rose in Figure 3, which shows that the statistically dominant wind direction at this site comes from the west, and no sensors were deployed directly east of any of the equipment.

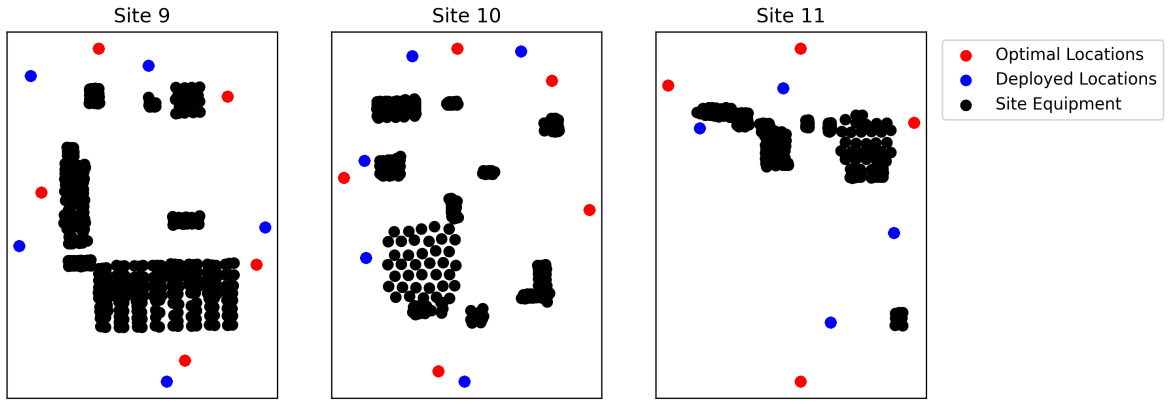

Figure 9: Deployed (blue) and optimal (red) sensor locations with respect to site equipment (black).

To analyze what effect these improved sensor locations may have had on the controlled release study, we compute the $\bar{\theta}_{min}$ parameter for each of the controlled releases for both the deployed and optimal sensor locations, and compare what fraction of the releases fell into the "detectable" regime ($\bar{\theta}_{min} < 30$ degrees) when the sensor placement is optimized using PC's new siting practices. Figure 10 shows histograms of $\bar{\theta}_{min}$ for all of the controlled release experiments for the sensor positions deployed in this study (top row) and for the optimized sensor positions using PC's siting tool (bottom row). There are a few key conclusions to draw from this figure: the detectable fraction of controlled releases for the deployed sensor locations closely mimics the actual fraction of releases detected in this study. At Site 9, the detectable fraction of releases based on the $\bar{\theta}_{min}$ analysis is 0.57, and 15/30 (fraction of 0.5) of the releases were actually detected. Similarly, the detectable fraction of releases at Site 10 was 0.67, and PC's network detected 8/12 (fraction of 0.66) of the releases. Finally, at Site 11, the detectable fraction of releases with the deployed sensor locations was 0.3, and PC's system detected 8/20 (0.4) of the releases. The obvious correspondence between the "detectable" metric of $\bar{\theta}_{min} < 30$ degrees and the fraction of releases PC actually detected at these sites suggests that this system is capable of detecting releases of these magnitudes, and whether a given release is actually detected simply depends on the wind direction.

We now turn to analyzing the improvement that may be recognized if sensor placement were better optimized in future studies. The bottom panel of Figure 10 shows analogous histograms of each controlled release's $\bar{\theta}_{min}$ if the sensors had been positioned according to PC's most up-to-date siting practices (the locations of which are shown with red dots in Figure 9). We clearly see an improvement in the fraction of detectable releases across all sites, but especially at Sites 10 and 11, where we had originally pointed out that the deployed sensor positioning was poor. Site 9 shows a marginal improvement in the detectable fraction, going from 0.57 for the deployed system to 0.71 for the optimal system. Sites 10 and 11 show more significant improvements: with the detectable fraction of releases going from 0.67 to 1 for Site 10, and from 0.3 to 1 for Site 11.

These results highlight the critical importance of sensor positioning to improving probability of detection metrics for these systems and also suggest that, had PC employed better sensor positioning at the time of this study, that a higher fraction of the controlled releases would have been detected. Future blind studies that employ PC's new siting tool to optimize sensor locations at operational

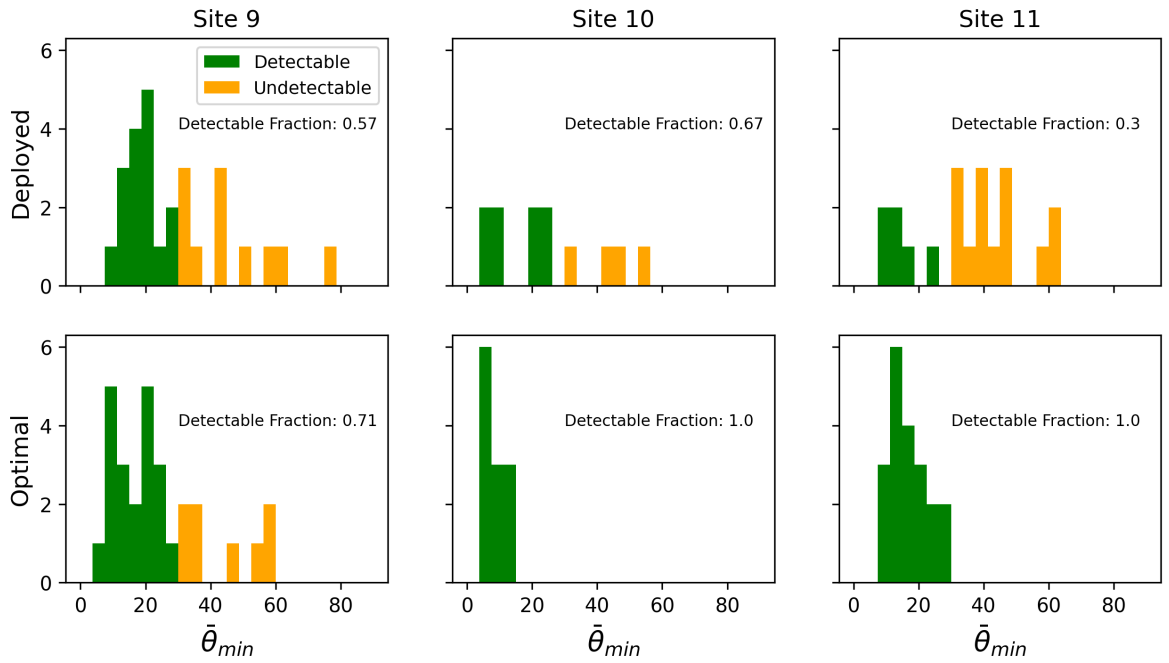

Figure 10: Histograms of $\bar{\theta}_{min}$ for the controlled releases at each site for the deployed sensor locations (top row) versus optimal sensor locations (bottom row).

facilities will confirm whether this is indeed the case, however we consider this analysis to be strong evidence for the argument that PC's hardware is sensitive enough to detect releases of the magnitudes used in this controlled release study, and that the data assimilation and quantification algorithms are sophisticated enough to process those raw measurements into quantification estimates that show evidence for releases, provided that the wind direction is close to a source-sensor axis, on average.

6 Quantification Accuracy

Due to the unknown background, making robust claims about the quantification accuracy of PC's system from the present study is challenging. Nevertheless, we present a brief analysis of PC's quantified site-level emissions for the controlled releases that were positively detected. We show in Figure 11 a scatter plot of the controlled release rate versus the detected uptick (computed as the maximum rate estimate during a controlled release minus the estimated background). We see that there is significant scatter about the parity relation (depicted with a dashed black line). This is likely due to a combination of many factors: i) the quantification estimate is inaccurate due to uncertainties inherent in the state estimation calculation, ii) the background is not perfectly estimated, iii) there were fluctuations in the operational emissions during the challenge release that result in the site-level emissions estimate changing due to both the challenge release and the non static background, thereby breaking the expected 1-to-1 relationship between uptick and release rate, and finally iv) PC's quantification algorithm requires the potential emission source locations as an input. If any of the challenge release locations were not exactly collocated with existing site infrastructure, then the release location may not be reflected in the list of sources used for the quantification calculation, leading to erroneous attribution of source fluxes and error in the site-integrated emissions rate.

Without a better understanding of the background emissions, it is impossible to completely disentangle the factors that may cause the computed uptick to differ from the actual release rate and directly assess the error of PC's quantification system. We can, however, examine specific trends in the quantified values and attempt to ascertain why they may be present. For instance, although there is significant scatter both above and below the parity line, the mean "error" (i.e., uptick minus release rate) is -0.51 kg/hr. The fact that these upticks are systematically lower than the actual release rates may be explainable by the fact that PC's dispersion model does not account for the obstructions

present at some of these sites. In general, obstructions will serve to increase the amount of dispersion by inducing more turbulence and associated mixing. Because PC's gas dispersion models that are employed at scale do not currently account for this effect, they will be **underdispersive** at the sites in this study which results in **higher** concentration predictions, which in turn result in **lower** emissions estimates than if they were appropriately capturing the effects of obstructions. As such, improving the underlying physical models to account for these obstructions should result in improved quantification accuracy and detection performance of challenge releases at more complex sites. Note that the site that shows the most systematic underprediction of source rates is Site 10 (orange dots), which also happens to be the site with the largest obstructions (battery of 15 tanks in the middle of a relatively small site).

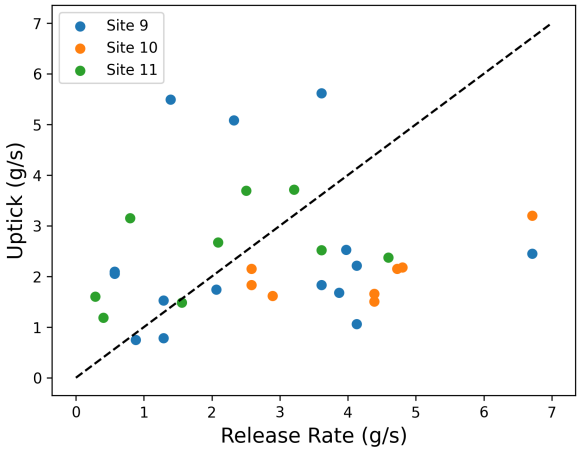

Figure 11: Controlled release rate versus detected uptick (quantified value minus background).

6.1 Cumulative Emissions Comparison

In addition to analyzing the error associated with individual releases, we compute the cumulative emissions over each day of challenge releases and compare the total mass emitted during that time period to the time-integrated uptick. For the purposes of this analysis, we focus on Sites 10 and 11, where the background is more stable than the highly-fluctuating unknown background present at Site 9. We show in Figures 12 and 13 the cumulative emission curves from the 2nd day of challenge releases at Site 10 and 11, respectively. In these figures, blue shows the actual cumulative release of gas in total kilograms over the course of the experiments on that day while orange depicts the integrated uptick (IU) inferred by PC, computed simply as:

$$IU(t) = \sum_{t'=0}^{t'=t} (R_{t'} - BL)\Delta t. \tag{1}$$

Here, t represents time, R_t represents the quantified rate at time t, Δt is the recursion interval that rates are computed over, and BL represents the background rate for that given day, taken as the the quantified rate, R, immediately proceeding the set of releases (the orange dots in Figures 6 and 7). In this way, the integrated uptick represents the portion of PC's emission estimate from the facility that are attributable to the controlled releases. Considering Figures 12 and 13, we see that, as expected, the cumulative emissions at Site 10 are underestimated with respect to the actual release mass: the releases over that 5-hour period totaled to about 77 kg, while the integrated uptick only estimated 30 kg of methane released over that time period above the baseline. In contrast, at Site 11, we see very good agreement between the actual and estimated cumulative emissions over the $\sim$ 6 hour period over which releases were occurring: the actual total emitted mass during this time was 37 kg while the system estimated 34 kg, corresponding to an error of 8 percent. We note that Site 11 is the most representative of the three sites in this study of a facility that would be considered eligible to deploy PC's quantification system on. Sites with too much obstructive infrastructure (such as Sites 9 and 11)

are explicitly **not** recommended to perform quantification on due to exactly the reason we demonstrate here: the presence of large obstructions result in inaccurate quantified rates. PC is actively working on computationally tractable parametrized dispersion models that will account for these effects in the future, but for now, the quantification method that is deployed at scale does not account for these obstacles.

To investigate whether the systematic under prediction of source rates is indeed due to the lack of accounting for the large obstructions present at some of these sites, we employ three dimensional computational fluid dynamics (CFD) simulations in the following section to understand the dispersion of methane on Site 10.

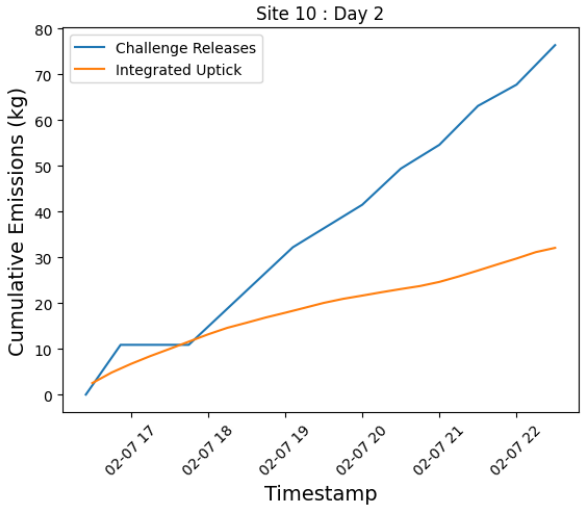

Figure 12: Cumulative emissions from the 2nd day of challenge releases at Site 10.

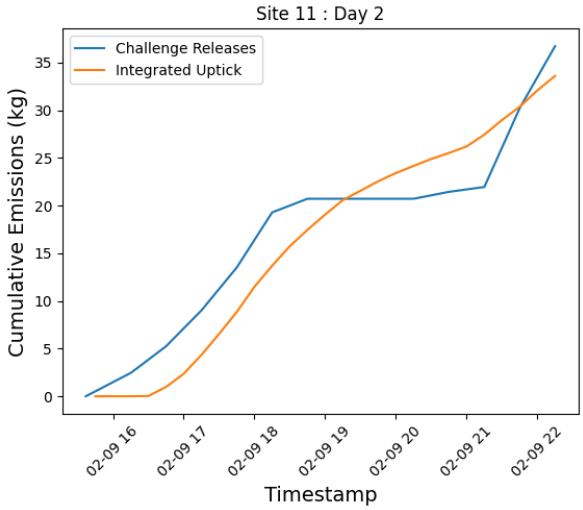

Figure 13: Cumulative emissions from the 2nd day of challenge releases at Site 11.

6.2 Influence of Obstructions on Methane Transport

In this section, we present results from Large Eddy Simulations (LES) [29, 30] of pollutant dispersion at Site 10. This facility has the most significant obstructions relative to its area and also shows the most systematic signs of under predicted source rates: all of the detected upticks were lower than the controlled release rates, whereas at the other two sites, there is some scatter both above and below the known rates. At this site, there were two distinct sets of challenge releases that occurred on different

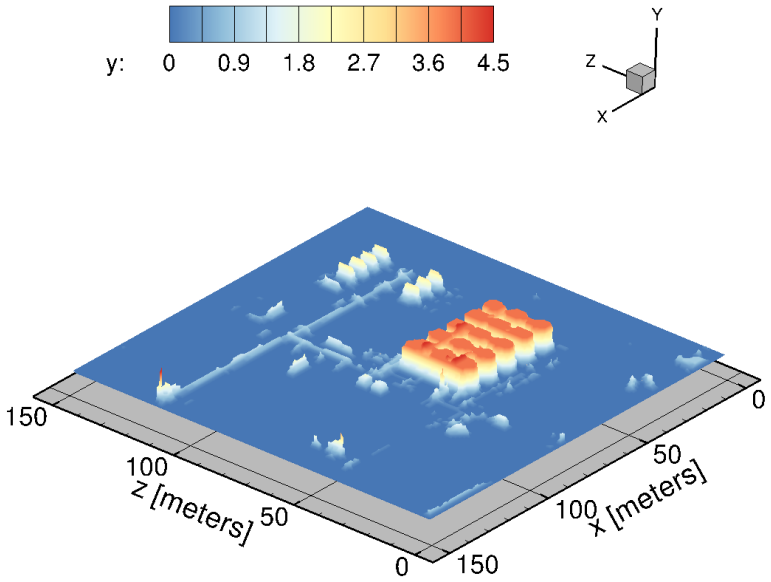

Figure 14: Visualizing the digital elevation map of Site 10 after it is cast onto a 3D Cartesian grid via the iso-surface of the local terrain height (y) in meters. The vertical axis is stretched to highlight terrain undulations.

days (Febuary 6th and 7th). As we have demonstrated, the nondetections from the challenge releases on Febuary 6th were due to the wind not pointing from the controlled release source towards a sensor. On the 7th, however, the controlled release location was directly north of the tanks (the leftmost triangle in Figure 2), and the wind blew predominantly south. As such, PC's southernmost sensor detected significant concentration enhancements during this time, and the magnitude of these enhancements is likely affected by the large tank battery between the source and the sensor, making this time period a good case study for understanding the effects of obstructions on measured concentrations and, in turn, quantified rates. For these reasons, we simulate the period of time corresponding to the challenge releases on February 7th.

In order to investigate the effects of surface obstructions on tracer dispersion at Site 10, we employ two simulations: one that includes the obstructive infrastructure at this site, and another where the surface is completely flat. These simulations are otherwise identical. The obstructions are incorporated into our CFD framework using a digital surface model of the site, which was generated via drone-based photogrammetry. Figure 14 depicts the digital surface model that is used in the LES simulations. The effect of surface roughness on the overlying atmospheric boundary layer (ABL) is included in both simulations by introducing a roughness lengthscale, z_0. Visual inspection suggests that Site 10 and its surroundings resemble an open terrain, which corresponds to $z_0 \simeq 0.03$ meters. We mimic the realistic flow conditions at this site by forcing the time-varying instantaneous surface layer towards a wind velocity profile estimated using in-situ measurements by anemometers. A detailed description of the numerical approach employed here to perform LES simulations accounting for measured wind conditions and facility obstructions is deferred to a future publication, although we provide a brief explanation of the LES framework in Appendix A.

Figure 15 shows the minute-averaged concentrations at the location of the southern sensor from the simulations with and without obstructions in red and green, respectively. The actual measurements from the southern sensor during this time period are shown in blue. Examining this figure, it is clear that the concentration predictions from the simulation that includes the obstructions (green) are closer to the field measurements (blue) than the concentration predictions from the simulation that does not

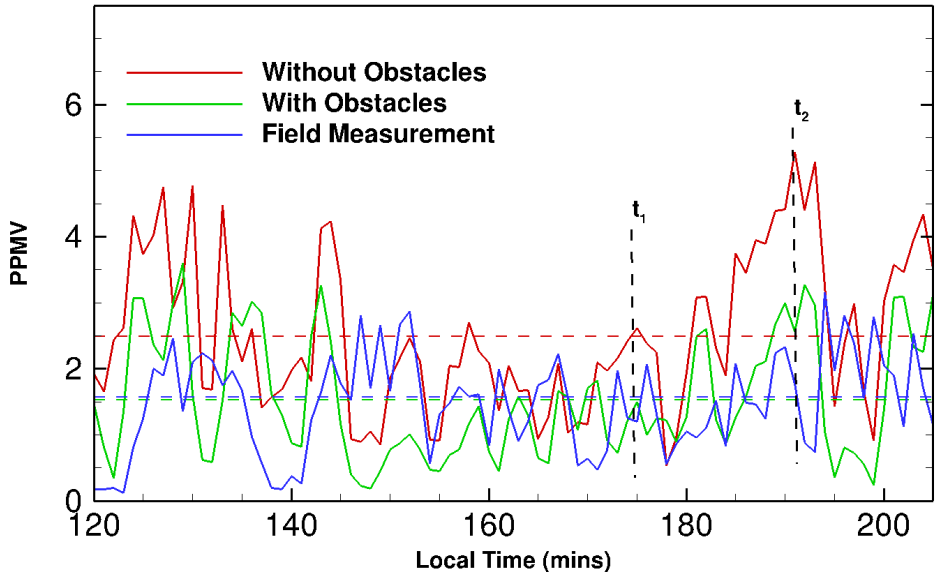

Figure 15: Time traces of minute-averaged concentration levels C at the southern sensor for the field measurement (blue) and the LES simulations with and without obstructions (green and red, respectively). The dashed horizontal lines at 2.49 parts per million (ppm), 1.53 ppm and 1.57 ppm indicate the average concentration over the time period shown ($\overline{C}$) for the LES without obstructions, LES with obstructions and field measurements, respectively. The local time label of 120 minutes on the horizontal axis corresponds to 3:00pm on Feb 7^{th}.

account for the obstructions (red). The mean predicted concentration,

$$\overline{C} = \frac{1}{T}\int_0^T C(x_i, t)\ \mathrm{d}t,$$

at the position of the southern sensor during the continuous releases are shown in Figure 15 with dashed horizontal lines. We see very good agreement between the $\overline{C}$ of the simulation with obstacles and the field measurements (1.53 and 1.57 ppm, respectively), while $\overline{C}$ of the simulation without obstructions is significantly higher, at 2.49 ppm.

These results are consistent with what one would expect: bluff bodies induce stronger and sharper horizontal and vertical gradients in the wind velocity field, which in turn enhances the three dimensional transport of methane away from regions of high concentration. In order to further demonstrate the effect of obstructions at this site, we show instantaneous snapshots of the predicted concentration at the receptor height (2 meters) from the two simulations in Figure 16. In this figure, the color scale depicts the predicted methane enhancement due to the controlled release, with blue corresponding to high concentrations and yellow corresponding to low concentrations. We show the positions of the sensors in the network with black circles, and the position of the controlled release with an encircled magenta marker. The red regions indicate computational cells associated with obstructions. Each row in this figure corresponds to a different instance in time (shown in Figure 15 with black dashed vertical lines and denoted t_1 and t_2). The leftmost column corresponds to the simulation with no obstructions while the rightmost column corresponds to the simulation with obstructions. By comparing the predicted concentration fields between the unobstructed (left) and obstructed (right) plumes, we can immediately see that concentrations downwind of the tank battery are significantly lower for the simulation that includes the obstructions due to the increased dispersion in the wake of the tank battery.

PC's quantification algorithm employed for this study did not account for the additional dispersion that occurs when large obstructions are present at a facility. As demonstrated in this section, properly accounting for the obstructions decreases the predicted concentration at sensors that are downwind of an obstruction. Because the dispersion model used in the quantification algorithm did not account for these effects, it tended to **overpredict** concentrations, which naturally leads to an **underprediction**

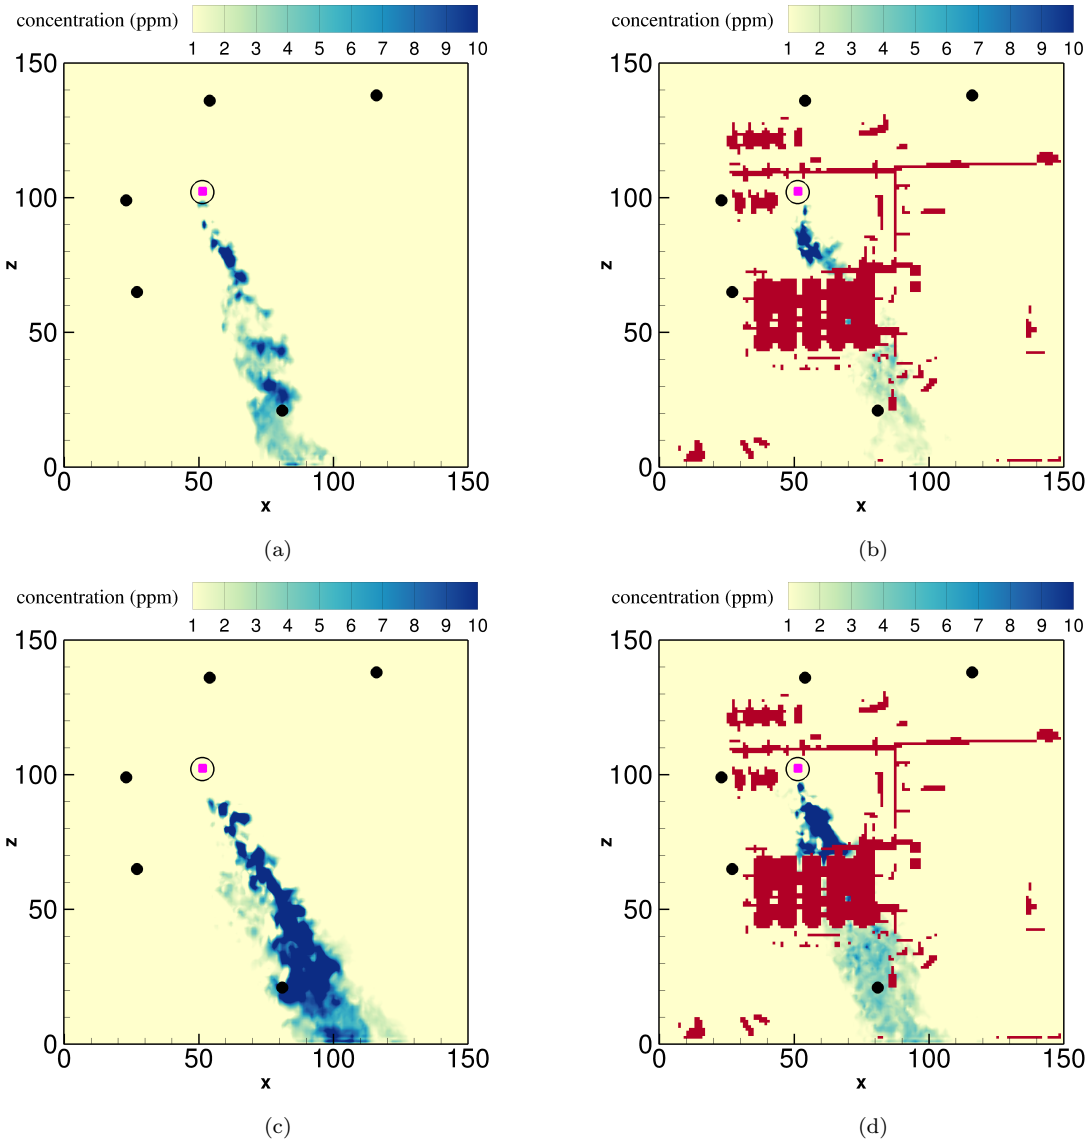

Figure 16: Horizontal instantaneous snapshots of the concentration levels in ppm at the receptor height. (a, c) Without obstructions; (d, e) with obstructions. (a, b) At time instant t_1 in Figure 15; (c, d) at time instant t_2 in Figure 15. The positions of the five point sensors are highlighted using black filled circles, the source location is identified by an encircled magenta marker and the obstructions are represented via red regions. Methane release height during this challenge release was 4m. All dimensions are in meters.

in source rates, a trend that is clearly evident when comparing the quantified upticks to known release rates at Site 10 (see Figure 11).

7 Discussion

In [28], the main conclusion is that point sensor networks and associated quantification algorithms failed to detect and quantify the vast majority of the controlled releases at operational facilities. The reasons for this are twofold: first, the quantification-enabled solutions that were a part of this study were not giving reliable estimates, as the majority of rate estimates were 0, even at active facilities, where there should obviously be a nonzero baseline emission rate. Second, even if the tested systems had produced more nonzero rate estimates, analyses presented in Section 3 suggests that using multi-week baselines to define detection thresholds for challenge releases lasting at most a couple hours results in baseline values that are not representative of the actual operational emissions over the several hour timespan that the challenge releases occur. Despite the poor performance of quantification systems in the previous study, [28] points out that there seems to be evidence in the raw mixing ratio signals for some of the controlled releases, and that improved quantification algorithms may be able to process these into emissions estimates that show evidence of the releases. By applying PC's most up-to-date quantification algorithm to the mixing ratio data, applying a new background estimation technique, and analyzing the results, we see that this is indeed the case: while detecting short releases is challenging and PC's system does not perfectly extract information about every individual release, we show that the site-level emissions estimates display enhancements coincident with 31 out of short releases, with the majority of nondetections being explainable by the wind direction simply not pointing towards a sensor during these short release periods.

When interpreting these results in the context of what these systems are deployed to accomplish (detecting persistent leaks and accurately quantifying facility emissions over long averaging times), we should consider the timescales of these individual releases and ask whether they accurately represent emission events of import. We are curious from an operators' perspective: how important is it to detect a 30 minute emission event? It seems probable that the emissions they most care about would persist for an extended period of time (e.g., an unmitigated equipment failure or leak), and as such, the analysis of these systems' capabilities to detect individual 30-minute releases, although interesting, may not be precisely relevant to the most important tasks that these systems are deployed to address. Furthermore, such short duration emission events are not likely to significantly impact the cumulative emissions computed over the 7 or 90 day timeframes associated with recent regulatory standards.

As previously mentioned, PC does not suggest using their quantification calculations on sites that are as complex as many of those that were a part of the [28] study: the presence of large obstructions (such as tank batteries or compressor stations) are not accounted for in the dispersion model PC currently uses for quantification. This is reflected in the fact the detected upticks in emissions from PC's estimate tended to be less than the challenge release rates, indicating that the quantification algorithm was systematically underestimating source rates, consistent with the expected effect of neglecting the physics of obstructions. We corroborate these conclusions by performing CFD simulations of Site 10 and demonstrate the effect that including the obstructions has on the predicted concentrations. In a future study, PC will describe in more detail the computational fluid dynamics tooling coupled to a quantification inverse solver that is capable of more accurately inferring source rates at complex facilities.

In addition to the challenges posed by these obstructions, the sensor placement at these sites was non-ideal, with too many of the sensors being upwind from sources when considering the dominant wind directions. In the past year, PC has improved their practices regarding sensor placement by developing an automated "siting tool" that finds optimal sensor locations given the wind statistics and potential source locations. We ran this tool against the three sites and found that it recommended sensor locations that deviated substantially from what PC had deployed at the time. As such, we suspect that if this study were to be repeated with sensor locations informed by PC's most up-to-date practices, that more of the releases would be detected.

Despite the challenges mentioned above, PC's quantification estimates showed evidence for just over half of the challenge releases from [28]. In order to ensure that future studies are as informative and useful as possible, we suggest using a background estimation technique that does not rely on uncontrolled multi-week baselines, and also employing longer releases to more closely mimic the

emissions events that these systems are deployed to detect.

Acknowledgements

The authors thank Ethan Emerson for his thoughtful review, comments, and insights during the preparation of this manuscript.

A Appendix: Background on the LES Setup

We employ a CFD utility that is specifically tailored for performing unsteady 3D simulations of pollutant dispersion around complex industrial facilities. The inputs to this tool are a digital surface model of the facility, a timeseries of wind speeds and directions, and a set of release locations and rates (which may vary in time). The model is able to effectively reconstruct the full chaotic and rapidly-changing 3D flow field and associated tracer dispersion at a facility accounting for the obstructive effects and undulating terrain informed by the digital surface model.

The flow solver makes use of the finite-difference framework to solve the incompressible Navier-Stokes equation on a staggered Cartesian mesh using a predictor-corrector based approach. This includes solving a Poisson equation to enforce the divergence-free condition on the velocity vector field. The effect of thermal buoyancy is modeled via the Boussinesq Approximation [31] and physical obstructions are mapped on the Cartesian grid by an immersed boundary method (IBM) [32]. The code is parallelized using the Message-Passing Interface (MPI), and has been employed successfully in the past to perform direct numerical simulations of both equilibrium and non-equilibrium turbulent channels flows and boundary layers over explicitly resolved rough surfaces [33]. The pollutant (methane) is treated as a passive tracer, which requires the solution of an advection-diffusion transport equation in which the Schmidt number (ratio of momentum diffusivity to mass diffusivity) is fixed at 0.7 [34]. The dynamic Smagorinsky model is selected for representing the subgrid-scale stresses and closing the filtered momentum and scalar transport equations [35], whereas an equilibrium surface-shear-stress based model is added to represent the wall effect [36]. Spatial derivatives in the momentum and passive scalar equations are discretized via the Central Difference and the Quadratic Upstream Interpolation for Convective Kinematics (QUICK) schemes, respectively, while time advancement is accomplished via a Newton-Raphson based iterative method [37]. The overall numerical approach is globally second-order accurate in both space and time.

The minute-averaged boundary layer is reconstructed by including a measurement-informed time-varying forcing term on the right-hand side of the momentum equation. The measurements by the on-site anemometer coupled with the surface heat flux estimates are utilized to reconstruct the horizontally-averaged boundary layer profile, which in turn is employed to interpolate the horizontal forcing term needed at each time step for the U and W components of the momentum equation. Time traces of instantaneous and minute-averaged horizontal wind speed ($U_{mag} = \sqrt{U^2 + W^2}$) and wind direction ($WD = 270^\circ - tan^{-1}(W/U)$) at 2m height (i.e., the approximate height at which the anemometer is installed) from the LES of Site 10 are compared with field measurements in Figure 17.

Periodic and Neumann boundary conditions are used on the vertical boundary planes for the wind velocity and the tracer fields, respectively. As our focus is on near-surface pollutant dispersion, we do not attempt to construct the entire boundary layer. The ceiling of the domain is capped at 100m and a no-flux condition is imposed at this height. While a uniform horizontal mesh was selected for the present study, the grid was stretched in the vertical direction to accommodate the ground-imposed anisotropy. The grid spacing at the ground was 0.2m, which was gradually relaxed to 3m towards the top boundary. A grid sensitivity analysis was performed by comparing the LES predictions at different virtual sensor locations using three different horizontal grid resolutions: (*i*) 5.0m $\times$ 5.0m, (*i*) 3.0m $\times$ 3.0m and (*iii*) 1.5m $\times$ 1.5m. While there is an expected dependence of the minute-averaged concentration levels on the horizontal spatial resolution, it is safe to suggest that this dependence is relatively small. Notwithstanding these differences, the results discussed herein are from the LES performed at the finest horizontal grid spacing, i.e. 1.5m $\times$ 1.5m.

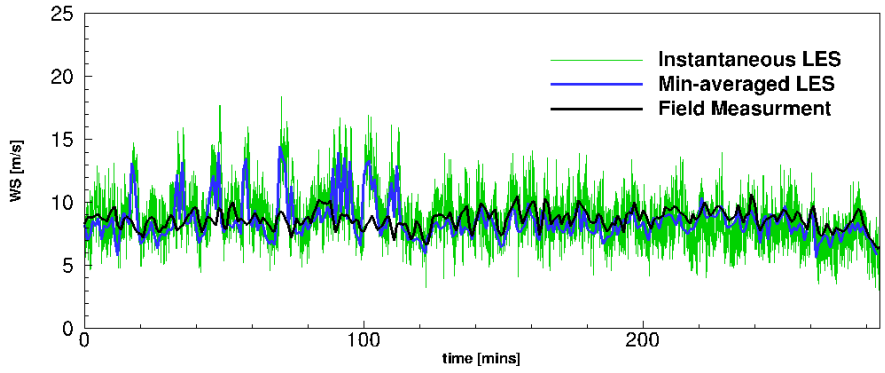

(a)

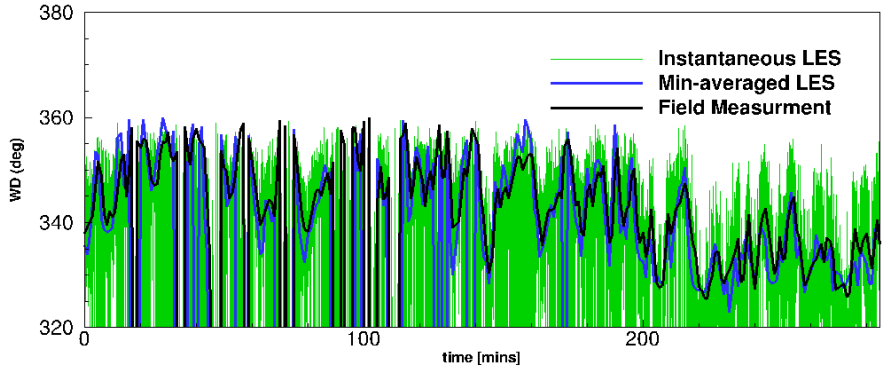

(b)

Figure 17: Time traces of instantaneous LES, minute-averaged LES and field measurements at North East sensor location for (a) wind speed U_{mag} and (b) wind direction WD. The LES results plotted in this figure include the effect of obstructions for Site 10 during the 285-minute long continuous release interval of the challenge release on Feb 7^{th}.

References

[1] *IEA Global Methane Tracker 2022: Methane and Climate Change*. https://www.iea.org/reports/global-methane-tracker-2022/methane-and-climate-change. Accessed: 2024-05-23.

[2] *IEA Global Methane Tracker 2023: Strategies to reduce emissions from oil and gas operations*. https://www.iea.org/reports/global-methane-tracker-2023/strategies-to-reduce-emissions-from-oil-and-gas-operations. Accessed: 2024-05-23.

[3] Evan D Sherwin et al. "US oil and gas system emissions from nearly one million aerial site measurements". In: *Nature* 627.8003 (2024). ISSN: 1476-4687. DOI: 10.1038/s41586-024-07117-5. URL: https://www.nature.com/articles/s41586-024-07117-5.

[4] Matthew R. Johnson et al. "Comparisons of Airborne Measurements and Inventory Estimates of Methane Emissions in the Alberta Upstream Oil and Gas Sector". In: *Environmental Science & Technology* 51.21 (2017). PMID: 29039181, pp. 13008–13017. DOI: 10.1021/acs.est.7b03525. eprint: https://doi.org/10.1021/acs.est.7b03525. URL: https://doi.org/10.1021/acs.est.7b03525.

[5] Ramón A. Alvarez et al. "Assessment of methane emissions from the U.S. oil and gas supply chain". In: *Science* 361.6398 (2018), pp. 186–188. DOI: 10.1126/science.aar7204. eprint: https://www.science.org/doi/pdf/10.1126/science.aar7204. URL: https://www.science.org/doi/abs/10.1126/science.aar7204.

[6] Anna M. Robertson et al. "New Mexico Permian Basin Measured Well Pad Methane Emissions Are a Factor of 5–9 Times Higher Than U.S. EPA Estimates". In: *Environmental Science & Technology* 54.21 (2020). PMID: 33058723, pp. 13926–13934. DOI: 10.1021/acs.est.0c02927. eprint: https://doi.org/10.1021/acs.est.0c02927. URL: https://doi.org/10.1021/acs.est.0c02927.

[7] J. D. Maasakkers et al. "2010–2015 North American methane emissions, sectoral contributions, and trends: a high-resolution inversion of GOSAT observations of atmospheric methane". In: *Atmospheric Chemistry and Physics* 21.6 (2021), pp. 4339–4356. DOI: 10.5194/acp-21-4339-2021. URL: https://acp.copernicus.org/articles/21/4339/2021/.

[8] Genevieve Plant et al. "Inefficient and unlit natural gas flares both emit large quantities of methane". In: *Science* 377.6614 (2022), pp. 1566–1571. DOI: 10.1126/science.abq0385. eprint: https://www.science.org/doi/pdf/10.1126/science.abq0385. URL: https://www.science.org/doi/abs/10.1126/science.abq0385.

[9] L. Shen et al. "Satellite quantification of oil and natural gas methane emissions in the US and Canada including contributions from individual basins". In: *Atmospheric Chemistry and Physics* 22.17 (2022), pp. 11203–11215. DOI: 10.5194/acp-22-11203-2022. URL: https://acp.copernicus.org/articles/22/11203/2022/.

[10] Xiao Lu et al. "Observation-derived 2010-2019 trends in methane emissions and intensities from US oil and gas fields tied to activity metrics". In: *Proceedings of the National Academy of Sciences* 120.17 (2023), e2217900120. DOI: 10.1073/pnas.2217900120. eprint: https://www.pnas.org/doi/pdf/10.1073/pnas.2217900120. URL: https://www.pnas.org/doi/abs/10.1073/pnas.2217900120.

[11] Stuart N. Riddick and Denise L. Mauzerall. "Likely substantial underestimation of reported methane emissions from United Kingdom upstream oil and gas activities". In: *Energy Environ. Sci.* 16 (1 2023), pp. 295–304. DOI: 10.1039/D2EE03072A. URL: http://dx.doi.org/10.1039/D2EE03072A.

[12] Stuart N. Riddick et al. "Potential Underestimate in Reported Bottom-up Methane Emissions from Oil and Gas Operations in the Delaware Basin". In: *Atmosphere* 15.2 (2024). ISSN: 2073-4433. DOI: 10.3390/atmos15020202. URL: https://www.mdpi.com/2073-4433/15/2/202.

[13] Kushal Tibrewal et al. "Assessment of methane emissions from oil, gas and coal sectors across inventories and atmospheric inversions". In: *Communications Earth & Environment* 5.1 (Jan. 2024), p. 26. DOI: 10.1038/s43247-023-01190-w. URL: https://hal.science/hal-04387996.

[14] M. Omara et al. "Constructing a measurement-based spatially explicit inventory of US oil and gas methane emissions". In: *Earth System Science Data Discussions* 2024 (2024), pp. 1–25. DOI: 10.5194/essd-2024-72. URL: https://essd.copernicus.org/preprints/essd-2024-72/.

[15] J. P. Williams et al. "Small emission sources disproportionately account for a large majority of total methane emissions from the US oil and gas sector". In: *EGUsphere* 2024 (2024), pp. 1–31. DOI: 10.5194/egusphere-2024-1402. URL: https://egusphere.copernicus.org/preprints/2024/egusphere-2024-1402/.

[16] X. Lu et al. "Methane emissions in the United States, Canada, and Mexico: evaluation of national methane emission inventories and 2010–2017 sectoral trends by inverse analysis of in situ (GLOBALVIEWplus CH4 ObsPack) and satellite (GOSAT) atmospheric observations". In: *Atmospheric Chemistry and Physics* 22.1 (2022), pp. 395–418. DOI: 10.5194/acp-22-395-2022. URL: https://acp.copernicus.org/articles/22/395/2022/.

[17] J. R. Worden et al. "The 2019 methane budget and uncertainties at 1 degree resolution and each country through Bayesian integration Of GOSAT total column methane data and a priori inventory estimates". In: *Atmospheric Chemistry and Physics* 22.10 (2022), pp. 6811–6841. DOI: 10.5194/acp-22-6811-2022. URL: https://acp.copernicus.org/articles/22/6811/2022/.

[18] Evan Sherwin et al. "Quantifying oil and natural gas system emissions using one million aerial site measurements". In: (Dec. 2022). DOI: 10.21203/rs.3.rs-2406848/v1.

[19] D. J. Jacob et al. "Quantifying methane emissions from the global scale down to point sources using satellite observations of atmospheric methane". In: *Atmospheric Chemistry and Physics* 22.14 (2022), pp. 9617–9646. DOI: 10.5194/acp-22-9617-2022. URL: https://acp.copernicus.org/articles/22/9617/2022/.

[20] Mozhou Gao et al. "Global observational coverage of onshore oil and gas methane sources with TROPOMI". In: *Scientific Reports* 13.16759 (Oct. 2023). ISSN: 2045-2322. DOI: 10.1038/s41598-023-41914-8. eprint: https://www.nature.com/articles/s41598-023-41914-8. URL: https://doi.org/10.1038/s41598-023-41914-8.

[21] Sahar H. El Abbadi et al. "Technological Maturity of Aircraft-Based Methane Sensing for Greenhouse Gas Mitigation". In: (Jan. 1753). DOI: 10.1021/acs.est.4c02439.s001. URL: https://acs.figshare.com/articles/journal_contribution/Technological_Maturity_of_Aircraft-Based_Methane_Sensing_for_Greenhouse_Gas_Mitigation/25852731.

[22] Benjamin Hmiel et al. "Empirical quantification of methane emission intensity from oil and gas producers in the Permian basin". In: *Environmental Research Letters* 18.2 (Jan. 2023), p. 024029. DOI: 10.1088/1748-9326/acb27e. URL: https://dx.doi.org/10.1088/1748-9326/acb27e.

[23] Shaun Higgins et al. "A Practical Framework for Oil and Gas Operators to Estimate Methane Emission Duration Using Operational Data". In: *SPE Journal* 29.05 (May 2024), pp. 2763–2771. ISSN: 1086-055X. DOI: 10.2118/219445-PA. eprint: https://onepetro.org/SJ/article-pdf/29/05/2763/3408678/spe-219445-pa.pdf. URL: https://doi.org/10.2118/219445-PA.

[24] William Daniels, Meng Jia, and Dorit Hammerling. "Estimating methane emission durations using continuous monitoring systems". In: *ChemRxiv* (May 2024). DOI: 10.26434/chemrxiv-2024-f921j. eprint: https://chemrxiv.org/engage/chemrxiv/article-details/6633e52891aefa6ce1ffaf3b. URL: https://doi.org/10.26434/chemrxiv-2024-f921j.

[25] Jiayang Lyra Wang et al. "Multiscale Methane Measurements at Oil and Gas Facilities Reveal Necessary Frameworks for Improved Emissions Accounting". In: *Environmental Science & Technology* 56.20 (2022). PMID: 36201663, pp. 14743–14752. DOI: 10.1021/acs.est.2c06211. eprint: https://doi.org/10.1021/acs.est.2c06211. URL: https://doi.org/10.1021/acs.est.2c06211.

[26] William S. Daniels et al. "Toward Multiscale Measurement-Informed Methane Inventories: Reconciling Bottom-Up Site-Level Inventories with Top-Down Measurements Using Continuous Monitoring Systems". In: *Environmental Science & Technology* 57.32 (2023). PMID: 37506319, pp. 11823–11833. DOI: 10.1021/acs.est.3c01121. eprint: https://doi.org/10.1021/acs.est.3c01121. URL: https://doi.org/10.1021/acs.est.3c01121.

[27] Clay Bell et al. "Performance of Continuous Emission Monitoring Solutions under a Single-Blind Controlled Testing Protocol". In: *Environmental Science & Technology* 57.14 (2023). PMID: 36977200, pp. 5794–5805. DOI: 10.1021/acs.est.2c09235. eprint: https://doi.org/10.1021/acs.est.2c09235. URL: https://doi.org/10.1021/acs.est.2c09235.

[28] Rachel Elizabeth Day et al. "Point Sensor Networks Struggle to Detect and Quantify Short Controlled Releases at Oil and Gas Sites". In: *Sensors* 24.8 (2024). ISSN: 1424-8220. DOI: 10.3390/s24082419. URL: https://www.mdpi.com/1424-8220/24/8/2419.

[29] Stephen B. Pope. *Turbulent Flows*. Cambridge University Press, 2000.

[30] PA Durbin and BA Pettersson Reif. *Statistical Theory and Modeling for Turbulent Flows*. Wiley Online Library, 2011.

[31] Roland B Stull. *An introduction to boundary layer meteorology*. Vol. 13. Springer Science & Business Media, 2012.

[32] Jungwoo Kim, Dongjoo Kim, and Haecheon Choi. "An Immersed-Boundary Finite-Volume Method for Simulations of Flow in Complex Geometries". In: *Journal of Computational Physics* 171.1 (2001), pp. 132–150. ISSN: 00219991. DOI: 10.1006/jcph.2001.6778.

[33] Umair Ismail. "Direct Numerical Simulation of a Turbulent Boundary Layer Encountering a Smooth-to-Rough Step Change". In: *Energies* 16.4 (2023). ISSN: 1996-1073. DOI: 10.3390/en16041709. URL: https://www.mdpi.com/1996-1073/16/4/1709.

[34] Wenjun Qian and Akula Venkatram. "Performance of Steady-State Dispersion Models Under Low Wind-Speed Conditions". In: *Boundary-Layer Meteorology* 138.3 (Mar. 2011), pp. 475–491. DOI: 10.1007/s10546-010-9565-1.

[35] D. K. Lilly. "A proposed modification of the Germano subgrid-scale closure method". In: *Physics of Fluids A: Fluid Dynamics* 4.3 (Mar. 1992), pp. 633–635. ISSN: 0899-8213. DOI: 10.1063/1.858280. eprint: https://pubs.aip.org/aip/pof/article-pdf/4/3/633/12443162/633_1_online.pdf. URL: https://doi.org/10.1063/1.858280.

[36] Teresa Salomone, Ugo Piomelli, and Giuliano De Stefano. "Wall-Modeled and Hybrid Large-Eddy Simulations of the Flow over Roughness Strips". In: *Fluids* 8.1 (2023), pp. 1–16. ISSN: 23115521. DOI: 10.3390/fluids8010010.

[37] Charles David Pierce. "Progress-variable approach for large-eddy simulation of turbulent combustion". PhD thesis. Stanford University, California, Jan. 2001.