APPLICATION OF CONTINUOUS MONITORING SYSTEMS IN METHANE EMISSIONS MEASUREMENT AND QUANTIFICATION

David Ball, Diego Leon-Quintero, Ali Lashgari
Project Canary

ABSTRACT

Methane emissions measurement technologies have been evolving rapidly and becoming increasingly efficient over the last few years. The purpose of this paper is to introduce recent technological advancements that have helped operators in the US obtain more in-depth methane leak insights, improving the performance of emissions mitigation programs, ensuring proper management of associated risks, and delivering measurement-based methane emissions inventories. Technological advancements include both measurement hardware and emissions data processing algorithms and software tools. However, emission source detection, localization, and quantification are still areas of ongoing research and need further improvement.

Technology companies, such as Project Canary, have developed novel models for detecting, localizing, and quantifying the total site emissions from oil and gas production facilities using continuous monitoring data. This model uses real-time and historical data to quantify emissions from various intermittent and continuous sources while differentiating any offsite emissions. A machine learning model is employed to build a unique model for each methane monitoring device to determine how the wind direction affects the concentration readings, simulating plumes from all potential emission sources and matching the plumes to the device model with a mixture model. This model is currently used to quantify emissions on hundreds of operating well pads across the United States. These models are complemented with operator notification and alerting systems to ensure timely actions by operators that result in reducing their environmental footprint and help keep the gas in pipelines. The most recent updates to the operator notification systems, called Smart Alerts, employ machine learning algorithms to eliminate unnecessary notifications and avoid alert fatigue.

INTRODUCTION

Managing methane emissions has become a serious challenge for oil and gas professionals in the U.S. as the industry landscape evolves quickly. Precise measurement of methane emissions is not only important; it is essential, with far-reaching implications for the sustainability of the environment, regulatory compliance, and operational efficiency of oil and gas production facilities. We have seen a profound change in the technology available to measure methane emissions over the last few years, driven by major advances in equipment and software. With these technological breakthroughs, operators have been given powerful tools to learn more about methane leaks and improve the effectiveness of emissions reduction programs.

This paper aims to provide insight into recent technological breakthroughs and their significant impact on the efficiency of emissions mitigation strategies. It gives operators a high-level understanding of emissions management by analyzing existing technologies to measure methane emissions continuously. The advances discussed include the measurement equipment and software tools developed to process emissions data. These improvements will allow a more accurate and effective approach to managing methane emissions, enabling operators to manage related risks better to comply with regulatory requirements.

Nevertheless, there is still work to be done to improve methane emissions management. Further research and development are still needed to detect, localize, and quantify methane emissions. Addressing these challenges is essential to refining emission management strategies and improving the accuracy of methane inventories.

A groundbreaking model for the continuous monitoring and management of methane emissions from oil and gas production facilities is an important advance in this area. This model is based on real-time and historical data to accurately quantify emissions from different sources, including machine learning techniques to improve the detection, localization, and quantification of methane to provide unprecedented insight into managing methane emissions. It separates on-site and off-site emissions and simulates emission streams. The potential of these technological innovations is underlined by the model's integration with advanced operator notification and alerting systems, including innovative Smart Alerts that use machine learning algorithms to reduce unnecessary notifications. These systems reduce environmental impact and enhance operational efficiency by minimizing product loss and ensuring pipeline gas containment.

The paper will discuss in detail these technical developments, their implementation difficulties, and possible new strategies to tackle methane emissions as they evolve. The aim is to give oil and gas professionals a comprehensive overview of how these innovations can be used for more sustainable and efficient production practices, which will, in turn, lead to greater awareness about the potential future benefits of such progress.

METHODOLOGY

Hardware

Emissions monitoring companies have a suite of methane sensing technology that can be deployed, ranging from metal oxide and spectroscopic-based point sensors to IR imaging cameras. For this discussion, we will focus on laser-based spectroscopy, as it offers the highest-fidelity continuous monitoring solutions that enable reliable leak alerting and provide relevant data to feed into the localization and quantification algorithms. As an example of the current capabilities of the monitors, Project Canary currently develops and deploys the “Canary X” device, which uses a near-infrared tunable diode laser to measure the absorption spectrum of the ambient gas and estimate the concentration of different molecules. This device can measure down to 0.4 parts per million, with a precision of around 0.125 parts per million. In contrast, Project Canary’s “Sentinel” device is based on mid-infrared spectroscopy and has a precision of less than 0.001 parts per million and a lower detection limit of 0.01 parts per million (the tradeoff is that the technology required to get down to this level of accuracy is more expensive). Both devices can be deployed for leak detection, localization, and quantification, and the choice between the two depends on the required level of precision and accuracy for the given application.

Data Acquisition

To monitor the emissions from a given site, at least three in-situ methane sensors (either Canary X or Sentinels) must be deployed along the facility's fence line. One of these sensing stations is also equipped with an anemometer that measures the wind speed and direction. All these data (methane concentrations from each sensor and wind measurements) are reported via cellular to cloud-based databases every minute.

Real-Time Alerting

The high temporal sampling of continuous monitoring coupled with highly scalable cloud computing infrastructure allows the performance of real-time analytics and alerts operators to anomalous methane signals shortly after they arise. These anomalies may indicate a large leak that not only costs money in lost products (natural gas molecules) but also poses significant risks of incurring regulatory fines if the leak persists. Innovative alerting system, such as Project Canary’s novel “Smart Alerts,” combines observed methane concentrations with on-site wind data to identify anomalous events. Since higher wind speeds dilute concentrations, the algorithm uses wind-normalized methane concentrations, which are the product of wind speed and a baseline-adjusted methane concentration, to compute alerting thresholds. Alerting thresholds are computed based on ten to thirty days of recent history at the site and updated nightly. For the operator to receive an alert, there must be sufficient observations above the threshold coming from the same sensor and a similar wind direction. This is indicative of sustained emissions from a consistent source on site. For the alert to end, there must be several consecutive observations with wind-normalized concentrations close to baseline. This helps ensure that just one alert is generated for a consistent event, even if the sensors pick up evidence inconsistently as the wind changes direction, reducing the risk of alert fatigue for the operator.

Localization and Quantification

In addition to providing real-time alerts to operators, there has been significant progress regarding novel localization and quantification services that generate emissions insights at the equipment level. These services have many applications. To name a few: these insights aid in localizing potential leaks so that they can be mitigated as rapidly as possible, and also provide emissions statistics for each piece of equipment so that targeted improvements can be made when trying to prioritize which pieces of equipment to upgrade, allowing the operators to prioritize the most problematic pieces of equipment in a data-driven manner rather than making ad-hoc assumptions about which pieces of equipment may be responsible for most of the site's emissions. Finally, these insights can be aggregated to a site-level emissions profile that can be analyzed on various timeframes to ensure compliance with regulatory limits.

Project Canary's localization and quantification algorithm is composed of two distinct components: the "forward model" that describes the transport of gas from sources to sensors, and the "inverter" that combines the forward model with sensor methane measurements to estimate the best-fit state vector (i.e., source rates).

Forward Model

As previously mentioned, the so-called "forward model" is the physical description of gas transport from sources to sensors under specified atmospheric conditions. It must account for wind speed and direction, source release rates, turbulent dispersion, and the relative positioning and heights of sources and sensors. In principle, there are copious forward models of various levels of complexity that one could employ for this purpose. At the extreme level of complexity is a fluid simulation using meteorological data as inputs that co-evolves a passive scalar (i.e., methane) as it is advected along with the simulated fluid flow. Such an approach would represent a high-fidelity forward model but is not practical for implementation in real-time systems due to the high computational expense.

A variety of simplifications exist that parametrize gas transport in a form that is more tractable for rapid computation. For instance, the well-known Gaussian Plume model represents a closed-form solution to the steady-state advection-diffusion equation, using an empirically derived estimate for the dispersion of gas as a function of downwind distance that depends on the stability class of the atmosphere. In short, the vast amount of scientific research on this model and the various extensions makes it flexible and trustworthy, while the ease of computation makes it highly efficient to drive the calculations necessary for the localization and quantification algorithms (for reference, running a high-fidelity fluid simulation as the forward model would take on the order of hours, while the Gaussian Plume model is computed on the order of milliseconds). We have found that for most cases, the benefits of the Gaussian Plume outweigh the slight accuracy improvements one can realize by employing a more sophisticated forward model, such as the time-dependent Gaussian Puff or numerically solving the advection-diffusion equation. For a facility with dramatic terrain variation or significant obstructions between sources and sensors, however, these other methods will likely be important for capturing the more complex dynamics of gas transport that the Gaussian Plume cannot account for as easily. As such, we are working to include them in our system as optional forward models we can enable if deemed necessary for a given facility.

The Gaussian Plume equation is given by:

Equation 1

$$C(x,y,z) = \frac{Q}{2\pi\sigma_y\sigma_z u}\exp\left(\frac{-y^2}{2\sigma_y{}^2}\right)\left[\exp\left(\frac{-(z-h)^2}{2\sigma_z{}^2}\right) + \exp\left(\frac{-(z+h)^2}{2\sigma_z{}^2}\right)\right].$$

On the left-hand side of this equation, *C* represents the predicted concentration at a point in space, *x,* and *y* are the downwind and crosswind distances, respectively, and *z* is the height. On the right-hand side, *Q* represents the source rate, and sigma sub y and sigma sub z are the horizontal and vertical components of the turbulent dispersion, which are functions of the downwind distance. To estimate these dispersion parameters, we turn towards empirically derived functional forms and associated coefficients from well-known controlled release studies that characterized the dispersion of passive tracers as a function of downwind distance and atmospheric conditions (e.g., Martin, 1976). Finally, *u* represents the average wind speed.

Applying the forward model to every minute of data at every sensor for every potential source with a unit rate results in the so-called "source-sensor sensitivity matrix," **A**. The rows (indexed by *i*) of this matrix correspond to each sensor measurement and associated atmospheric conditions at that time, and the columns (indexed by *j*) correspond to each source. The value at each element (*i,j*) in the matrix represents the predicted concentration at sensor/time *i* from source *j* via the forward model.

We emphasize here that any forward model can be used to compute the source-sensitivity matrix. The following steps in the localization and quantification calculation (the inversion to source rates) are independent of the particular choice of the forward model. As such, the forward model can be viewed as a modular component of the system that can be adjusted depending on the details of the facility (i.e., we can swap out the Gaussian Plume for a higher-fidelity model at complex facilities, compute the **A** matrix, and move forward in the calculation in the same way).

Inverter

The next step in the localization and quantification algorithm is the so-called "inversion" of methane measurements to a state vector (or source rates). We employ a recursive Bayesian estimator in the form of a standard Kalman Filter (see, e.g., Welch & Bishop 2006 for a modern overview of Kalman Filters), where the previously described "source-sensitivity matrix," **A**, is the measurement function (typically denoted as **H** in control theory literature) of the Kalman filter that is responsible for transforming a state-space estimate to measurement-space. This framework is run every 15 minutes and computes the state of the system at the time being considered, q_t, via a recursive Bayesian update. In other words, the existing state of the facility, q_{t-1}, is used as the "prior" and is combined with the new data to give the optimal state estimate at the current time, q_t. This current estimate is then saved and used as q_{t-1} in the next iteration, and the calculation continues in perpetuity.

This process generates an updated estimate of the emissions rate for each potential source at a site every 15 minutes. The historical state estimates are all saved so that the source rates as a function of time can be used for further analysis (leak detection and alerting, facility-level quantification of emissions over time, analyzing individual pieces of equipment's emissions distribution, etc.).

RESULTS

Controlled Release Testing

In order to evaluate the accuracy of the previously described algorithms, we turn towards controlled-release testing to quantify a variety of performance metrics, including localization accuracy, quantification accuracy, and the leak detection rate and detection limit. We have participated in several controlled release studies, but for the current purposes, we will focus on the “Advancing Development of Emissions Detection” (ADED) testing from 2023 run by Colorado State University’s Methane Emission Technology Evaluation Center (METEC). At this facility, they have 18 different pieces of equipment, including tanks, separators, and wellheads, from which they release methane in a highly controlled and metered manner. We have sensors placed around the perimeter of the facility and attempt to identify, localize, and quantify their controlled leaks. By comparing our estimates to their known leak rates, we can evaluate our current algorithms across a variety of performance metrics. We will now summarize some of the key results from applying our current algorithms to the known controlled releases from the ADED 2023 testing.

We will begin by examining the event-level statistics from ADED 2023, including an analysis of the event detection performance, the detection limit, and the cumulative amount of methane released. Over the course of ADED, there were 254 experiments, each of which had between 1 and 5 release sources. Of these 254 experiments, our leak detection algorithm correctly identified 240, corresponding to a detection rate of 94%. The missed detections (or false negatives) were all low-release rate experiments or high wind speed (around 10 meters/second), corresponding to an event-level detection limit of around 0.3 kg/hr. In Figure 1, we show the cumulative emissions curves from the controlled releases (blue) compared to our quantified emissions estimates (orange). These results demonstrate that our current quantification calculations are highly accurate when integrated over long timescales; over the course of the three-month ADED campaign, our total emissions estimate was off by only 1.34 percent.

We now turn toward the source-level metrics. During these 254 experiments, 536 individual releases were recorded. Of these, we detected 472 leak sources (corresponding to 64 false negatives), corresponding to a source-level detection rate of 88%. Of these correctly identified sources, our localization algorithms correctly identified the equipment group of the individual releases 95% of the time and correctly identified the piece of equipment responsible for the emissions about 55% of the time.

We note that the results presented here are derived from our own analyses after applying an updated version of the detection, localization, and quantification algorithms to the ADED 2023 data. In other words, these are not the same algorithms that were evaluated by the test center during the testing campaign. As such, the true test of these algorithms will be from the ADED 2024 campaign that is currently ongoing.

ANALYSIS

By applying our detection, localization, and quantification algorithms to the ADED 2023 campaign and performing a posthoc analysis of the results, we demonstrated that current technology and associated algorithms can produce reliable leak alerts down to a site-emission rate of 0.3 kg/hr, localize leaks to the equipment-group level with an accuracy of 95%, and estimate cumulative emissions over long (multi-week) timescales with extremely small (sub 5%) error. These results are based on analysis and development that occurred after the results from the blind testing were released, so these statistics need to be validated at the conclusion of the ADED 2024 campaign. Nevertheless, these promising results demonstrate that this technology is effective for practical applications, having a high detection rate, low detection limit, and highly accurate estimate of cumulative emissions). We note that one significant area for improvement is in the equipment-level localization, which is significantly worse than the equipment-group localization (55% accurate compared to 95%). This stark difference is likely due to the inherent error associated with our current dispersion modeling (the Gaussian Plume). We are working on incorporating higher-frequency measurements coupled with more advanced time-dependent modeling techniques that should improve the accuracy of the forward model and, hence, improve the localization accuracy significantly.

CONCLUSIONS

In this paper, we provided an overview of continuous methane emissions technology, from hardware to algorithms. We described how the data collected by high-fidelity methane sensors can be employed to generate emissions insights, including leak detection, localization, and quantification estimates. We also described how the most up-to-date algorithms against the controlled release experiments from ADED 2023 correctly detected 95% of all emission events with a 90% detection limit of around 0.3 kg/hr. The estimated cumulative emissions over the three-month-long testing campaign were higher than the amount released by a small margin (1.34%). At the source level, the technology correctly identified 88% of all equipment leaks and correctly localized them to their respective equipment group at a rate of 95%. Finally, we note that an area of future work and improvement is in incorporating more advanced modeling techniques in calculating the dispersion of gas from source to sensors. We postulate that by leveraging higher temporal frequency data (i.e., second-level measurements instead of minute-level) and combining these with dispersion models that explicitly account for time dependence (e.g., Gaussian Puff, advection-diffusion solver, or another reduced-order fluid simulation), we should get significant improvements in the accuracy of the forward model, which will likely propagate to a significant improvement in the localization accuracy at the equipment-level.

These promising results from state-of-the-art continuous monitoring hardware and associated quantification algorithms demonstrate the utility of these technologies for aiding operators in a variety of applications, from reducing the time-to-detection of leaks and minimizing the manual effort in leak detection and repair to optimizing equipment upgrades in a data-driven way that takes into account the emissions profile from each piece of equipment at a site, to generating site-level emissions insights on a variety of timescales to ensure compliance with regulatory requirements, thereby reducing the risk of government-imposed fines.

REFERENCES

Martin, D. O. (1976, February). Comment On"The Change of Concentration Standard Deviations with Distance." *Journal of the Air Pollution Control Association, 26*, 145–147. doi:10.1080/00022470.1976.10470238

Welch, G. and Bishop, G. (2006) An Introduction to the Kalman Filter. Department of Computer Science, University of North Carolina, Chapel Hill, TR 95-041

TABLES, GRAPHS, AND FIGURES

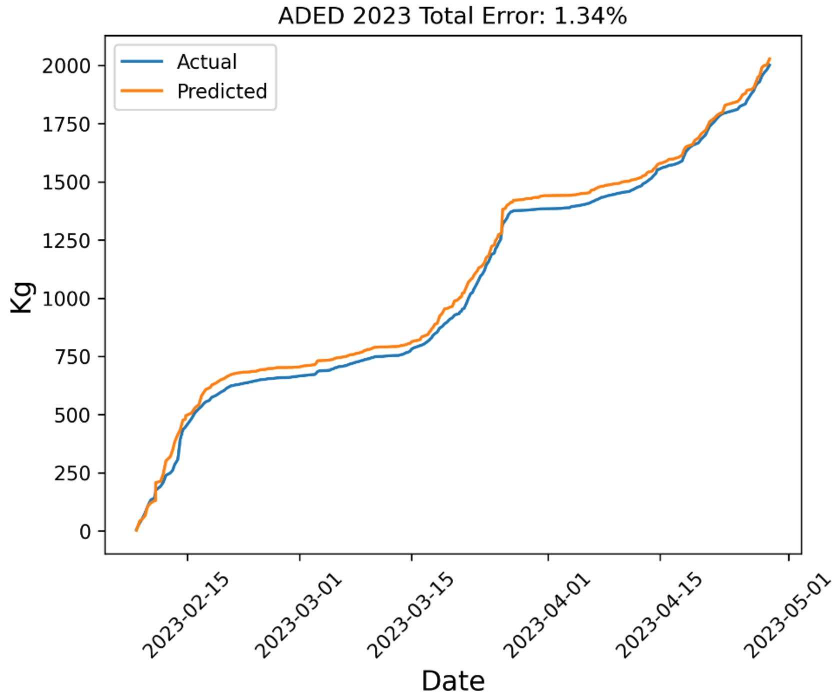

Figure 1 - Cumulative emissions over time from ADED 2023. The actual releases are shown with blue while Project Canary's quantified emission estimates are shown with the orange line.