

Data Informed Multipoint Ground Flare Evaluation Using Measurement and Theory

J. Thornock¹, M. Hradisky¹, M. Zhou², M. Cremer²

Abstract

This paper will explore the comparison of experimental (PFTIR) and simulation data obtained from assisted multipoint ground flares operating at high turndown. The flare assist is intended to encourage air entrainment for best combustion efficiency as well as reduce visible smoke. Operational envelopes of the assist to vent gas ratios are regulated on flare specific parameters. Experimental measurements have, however, measured low combustion efficiency even when operating within flare regulatory guidelines. This result is most likely due to the nonlinear nature of the flare combustion process and its response to regulatory constraints while being exposed to chaotic environmental conditions. Reliable and accurate methods for measuring, predicting and controlling flare performance across a wide range of conditions could offer significant environmental and economical advantages through better informed regulation, operation and ultimately high combustion efficiency.

PFTIR measurements represent a promising technology for remote flare sensing, as they may be placed safely away from the combustion zone and report critical species concentrations along a line of sight from which a combustion efficiency may be computed. However, environmental conditions, experimental error and bias significantly affect the measured combustion efficiency. Advanced simulation techniques have also offered a promising avenue for evaluating flare performance. However, these simulation tools may require large computational resources or require high level expertise to run and evaluate. While both the experimental and simulation approach have their limitations, the combination of the two data sources can offer significant insight into the flare's performance and even offer real-time feedback for operation and control.

Here we will present simulation predictions of assisted multipoint ground flares combined with PFTIR measurements. We will also explore the possibility of using the model in concert with the PFTIR measurements to provide real-time, data informed predictions for flare operations.

We would like to acknowledge the participation and help of Scott Evans and Dan Pearson from Clean Air and Sean Smith of the University of Utah

¹University of Utah, Institute for Clean and Secure Energy

²Reaction Engineering International (REI)

Introduction

This work focuses on combining experimentally PFTIR measured data on multi-point ground flares (MPGF) with data from theory/modeling. While comparisons of measured data and simulated versions of the same physical system are the norm, the current approach reframes this process in an artificial intelligence, machine learning (ML) based approach. The intent of the reframing is to build data upon data, rather than the traditional approach of contrasting data with data. The current Bayesian Machine Learning (BML) approach views all data as potentially contributing to better informed decision making as long as the uncertainty of the data source can be quantified either rigorously or adopted from prior information of the system.

By combining the data with help from the computer (machine), it is possible to have the data be cross-informed to allow inferences (learning) to occur. The approach also potentially reduces uncertainty in both the inputs and outputs values, benefitting both experimental and modeled data sources all within a sound mathematical theory. This approach differs from the more common Neural Network (NN) ML approaches because it requires a physical model of the system to generate data. Typical NNML approaches rely on collections of large amount of data with data fitting techniques for inference, while lacking a physical description of the system. The Bayesian based approach allows us to avoid some bias errors that naturally arise in NNML approaches because of the specific lack of a grounding physical model. Our approach also allows inference of quantities not easily measured (e.g., integrated combustion efficiency) which would be difficult or impossible to measure in the NNML approach because of the inclusion of the physical model.

The specific case study of this work focuses on the combustion efficiency (CE) of a John Zink SKEC steam-assisted flares at high turndown. Data used for this work were collected as part of Marathon Petroleum Company's Flare Consent Decree (as recorded August 30, 2012) at the John Zink Flare Facility on SKEC and LRGO flares. The high turndown scenario is particularly compelling because the ground flares spend most of the time in this state, only being fully utilized in process upset scenarios. This data identifies regions of low combustion efficiency in over-assisted scenarios and has characterized the DRE as a function of the net heating value in the combustion zone, which includes the steam flow rates. The data illustrate that the assist rate need be considered in determining operating parameters for these steam-assisted flares. For this work, the application of the BML method will be on test suite SN1 as termed in the reported data.

Data Sources

"A theory is something nobody believes, except the person who made it. An experiment is something everybody believes, except the person who made it."

-Albert Einstein

Experimental Data

The data of interest were provided by Clean Air in reported form entitled "Performance Test of Steam-Assisted and Pressure Assisted Ground Flare Burners with Passive FTIR – Garyville", dated March 21, 2013. In the report, five groups of tests were described exploring the operating envelope of two assisted ground flares; the steam-assisted John Zink SKEC flare and the pressure-assisted John Zink LRGO flare. These five groups of tests were carried out at the John Zink flare test facility. In all tests, Passive Fourier transform Infrared (PFTIR) measurements were made to determine the combustion efficiency. Scenario measurements, such as wind speed, wind direction, fuel compositions, vent gas feed rates, etc. were reported.

The PFTIR measurement theory is based on the fact that hot gasses emit and absorb a specific radiative signal, which the device captures within a narrow view angle of the subject. Different molecular species emit at specific wavelengths that act as a fingerprint of that specific molecule. The intensity of the signal allows one to determine the relative concentration of that species along the line of sight of the instrument. This allows the instrument to be placed at a safe distance from the combustion zone with an operator aiming the instrument behind the combustion zone to enable the measurements of the species of interest. The flare combustion efficiency is the quantity of interest and is computed as,

$$\eta = \frac{\phi_{CO_2}}{\phi_{CO_2} + \phi_{CO} + \phi_{HC}} \quad (1)$$

where η is the flare combustion efficiency and ϕ_i is the concentration of each species along the line ($ppm \cdot m$). The PFTIR instrument captures the radiative signal of the downstream conditions of the flare combustion zone at a point, identifies the relative species concentrations of CO₂, CO, and total HC (hydrocarbons), and computes the predicted value of η .

Errors or uncertainty in the PFTIR measurement of combustion efficiency, as in all measured data, can arise from a variety of sources. Here, we distinguish between instrument and measurement error and will use the terms “error” and “uncertainty” interchangeably. Instrument uncertainty is that error that is typically reported by the manufacturer. It is usually quantified in a well controlled experiment or calibration procedure in which the true value of the measured quantity is known to a good degree of accuracy. The instrument uncertainty is ascertained by repeated measurements and calibrations in the well controlled setting. As a result, instrument uncertainty is typically small. Measurement error, on the hand, is often the largest source of uncertainty. This error source results from variables beyond the control of the experimentalist or those that may have been overlooked. Overlooked variables are sometimes observed during the experimentation procedure at best, but often remain hidden from view, resulting in unknown and significant sources of variability in the observed result.

The PFTIR predictions of flare combustion efficiency has both instrument and measurement uncertainty. In this work, we have not quantified the instrument uncertainty. However, we recognize that the PFTIR instrument itself doesn’t measure concentrations directly (as described above). A physics based model (or instrument model, IM) is used to convert the signal to the final concentration measurements. In this process, the plume temperature must be inferred. Constants also must be calibrated regularly, which are used in the conversion of the signal to concentrations. The data are also filtered in the case of weak CO₂ signals, with that data being rejected. In the case of low combustion efficiency (high HC and low CO₂ signals), data were rejected because the IM would report zero combustion efficiency. With enough information and details of the PFTIR IM, the instrument uncertainty could be quantified. However, in this presentation we will assume that the mean of that uncertainty distribution and the standard deviation is small relative to the measurement uncertainty.

Environmental challenges pose the biggest challenge for the measurement uncertainty. These uncontrolled variables require the PFTIR operator adjust the aim of the instrument so that it ideally captures the downstream spectral signal one flame length downstream from the combustion zone. This is not always possible given that wind shifts continuously throughout the tests. Additionally, the point-wise positioning in the downstream combustion zone might not characterize the entire combustion zone given the narrow angle of the PFTIR without making a well-mixed assumption of the combustion products. Here, we will attempt to characterize the measurement uncertainty given the data provided.

Modeled Data

The raw modeled data is obtained from our in-house Arches computational fluid dynamics code. Arches is a large eddy simulations (LES) combustion simulator that runs on distributed parallel computers. Arches solves equations of mass, momentum, and energy, including radiative transport, for single and multiphase combustion applications. Arches resolves a range of timescales capturing time dependent information as well as spatial information of the system of interest. A variety of combustion models are available in Arches. Here, we use the Rate Controlled Constrained Equilibrium (RCCE) model. RCCE resolves a global combustion rate that constrains a chemical equilibrium assumption allowing for combustion quenching to be represented, which is crucial for computing combustion efficiency. A dynamic LES turbulence closure is used that uses local turbulent information to close the equation set. Various

quantities of interest from Arches simulations are easily extracted, time averaged and reported for visualization or data reduction.

There are three major sources of uncertainty in the Arches simulations. These are uncertainties in the model or model parameters, numerical uncertainty arising from discrete representation of continuous equations, and uncertainty in the scenario parameters. Scenario parameters are those environmental related inputs required to construct a simulation. For example, variations in the wind speed, flare geometry, vent gas flow rates, steam flow rates, and so-on are typical examples that result in boundary condition uncertainty for the simulation. For the work presented here, we will assume that the simulations are of sufficiently fine resolution ($O(1\text{mm})$) that the numerical error is small relative to the model and scenario uncertainty. The characterization of the model and scenario uncertainty is done in context of the BML method presented below.

The modeled combustion efficiency from Arches is obtained by shooting a line of sight through the time-dependent LES data. The target of the line of sight is held constant. Values of CO_2 , CO , and hydrocarbon are time-averaged along the line. This process is done to replicate the PFTIR instrument. For this virtual instrument, the position of the PFTIR camera and view angle is set independently as the operator of the real instrument would do. Because the virtual instrumentation has no feedback on the simulation results, several view angles can be explored as an independent, after-the-fact parameter without increasing the cost of the simulation. This allows us to explore the uncertainty of the PFTIR target in the simulation. For any of the virtual lines-of-sight, various windows of time-averaging are explored to ensure statistical steady-state.

Geometric details of the SKEC flare head were scarce. The simulated version was constructed using photos, online promotional information, and flow area information reported with the PFTIR data. Wind data including speed and direction was collected during the test but it was uncertain where the information was collected relative to the flare stacks. The average windspeed and direction were used for the simulated boundary conditions. These were held constant though the simulation. Other important boundary conditions, such as the vent gas conditions, flow rates and steam flow rates were provided in the report.

Given that both the simulated and measured data sources have sources of uncertainty, it is our hope that the combination of the data together as valid data sources results in a net information gain more so than any data source on its own. This process is done using Bayes Theorem and is discussed next.

Bayesian Machine Learning

In our analysis of the data, we use the following basis:

- Our LES model is a function which maps a set of inputs, x , to an output, y_m . So, $y_m = f(x)$.
- The x vector is potentially large, containing model parameters, environmental scenario parameters, and numerical parameters, all of which have uncertainty.
- The observed data also maps a set of conditions to an output, y_e . The input parameter space, x_e , for the measured output is also large, with many parameters being unmeasured or unknown. These parameters also have uncertainty.
- For both y_e and y_m , the uncertainty in the input parameters maps to uncertainty in the output.
- The difference between the two, $\varepsilon = y_m - y_e$, is uncertainty and is characterized by a distribution.

It is upon this framework which we build the Bayesian Machine Learning (BML) approach.

BML operates on the principle of using Bayes Theorem which, mathematically, is stated as,

$$f(x|y) \approx f(x)f(y|x) \quad (2)$$

where f represents a distribution, x are the input parameters, and y are the observations. In words, the left hand side of the equation is distribution of the input parameters conditioned on the data and is called the posterior. On the right is the prior uncertainty distribution of the inputs ($f(x)$) multiplied with the data conditioned on those inputs ($f(y|x)$). This last function is typically called the likelihood function. In both the construction of the likelihood function and so-called predictive posterior, a relationship is needed to map inputs (x) to the output observables (y). The likelihood construction also requires observed data with a prior uncertainty distribution. The relationship that maps inputs to outputs is the physical model of the system, which can come in any form with as much or as little complexity as the system demands. Once the posterior is obtained, one can use it to predict new values of y . This prediction is called the predictive posterior and, once obtained, contains the updated state of knowledge using all available data - both modeled and observed. It is worth pointing out that often times the model may be able to map other information of interest that aren't easily measured or observed. Thus, the predictive posterior can potentially provide more accurate information on y , but it may also provide additional information to further evaluate the state of the system and draw inference from the physical model.

In the case of the steamed flare, there is one quantity of interest examined here; the combustion efficiency. The inputs, x_e , are many. Here, we employ the sparsity-of-effects principle, which states that the system is typically dominated by a small subset of lower-order interactions of the input space. This eases the curse of dimensionality, which can rapidly increase the cost of application of BML due to the cost of the LES model. Thus, to keep the scope of this project down, we examine three important input parameters that we have deemed significant; The windspeed and direction, the steam flow rate (or net heating value of the combustion zone), and the position of the PFTIR target. Each input is characterized by a distribution, which is an assumed prior. We also assume that the experimental PFTIR data is also normally distributed.

As mentioned, the typical cost of one LES simulation is not insignificant. Given one set of inputs, a single simulation runs for about three days on roughly 400 cores. To sample the input space with enough points to construct a good representation of the posterior would be prohibitively expensive. As a result, we carefully select a range of input parameters, evaluate the full LES model, and finally fit a surface to the simulation output. This response surface then serves as a surrogate model for the full LES simulation and can be sampled at a high rate with very low cost (e.g., fractions of seconds for one sample).

Once the posterior predictive is obtained, additional information may be extracted from the physical model. For example, an integrated combustion efficiency across a volume is impossible to obtain in the physical system, but easily computed from the model form. Given that the measured data has informed model through the input distributions, the *inferred* integrated combustion efficiency can be obtained with an uncertainty distribution. This could extend to other quantities that the LES model could capture. With additional collection of measured data, the posterior and predictive posterior would be updated. The process just described can be extended to one that provides real-time data for evaluation and control of an operating flare system. Such a system would be continuously updated with current, measured data, which could improve the accuracy of the controller over time.

Results

Sample Simulation

Here we discuss an example of a typical LES simulation of the SKEC flare.

The geometry of the flare head was constructed using information from the Marathon report and with whatever little information could be gleaned from online and openly available sources. Figure 1 shows the comparison of a photo of the actual flare head with the reconstructed flare head used in the LES.

The Arches simulations are based on the SN1 (as titled in the report) test series. The primary objective of the experiments were to determine the performance envelope of the burner at the minimum flow rate (very high turndown). This scenario tests the performance of the flare at a point of near flame extinction. The minimum flow rating for this particular tip is approximately 4 ft/sec. 100 % Tulsa natural gas was used as the vent gas. The steam flow rate was adjusted to vary the net heating value of the combustion zone (NHV_{CZ}), computed as,

$$NHV_{CZ} = \frac{Q_{vg} \cdot NHV_{vg}}{Q_{vg} + Q_s} \quad (3)$$

where Q is the volumetric flow rate of the vent gas (vg) and steam (s). In this example, the NHV_{CZ} was around 100 BTU/scf given the flow rate conditions.

Experimental measurements were obtained using PFTIR. The PFTIR camera was aimed at a location about one flame length downstream from the combustion zone. The line-of-sight measurement returns signal from which the relative amounts of various species can be deduced. Specifically, using the CO_2 , CO , and unburnt hydrocarbon signal a combustion efficiency is computed as the ratio of the CO_2 to the sum CO_2+CO +unburnt hydrocarbon. As the PFTIR is a line-of-sight measurement technique, naturally the value is dependent on several variables, such as the location of the camera, the positioning of the target, and environmental conditions affecting the plume of the flare to name a few. Alignment of the camera and positioning of the target becomes critical in the final combustion efficiency signal. For example, for test SN1, several measurements were obtained across a range of cross winds and cross wind directions for constant vent gas to assist ratios. These efficiencies are shown in Figure 2. Failed measurements are indicated with a zero combustion efficiency reading. Note that the

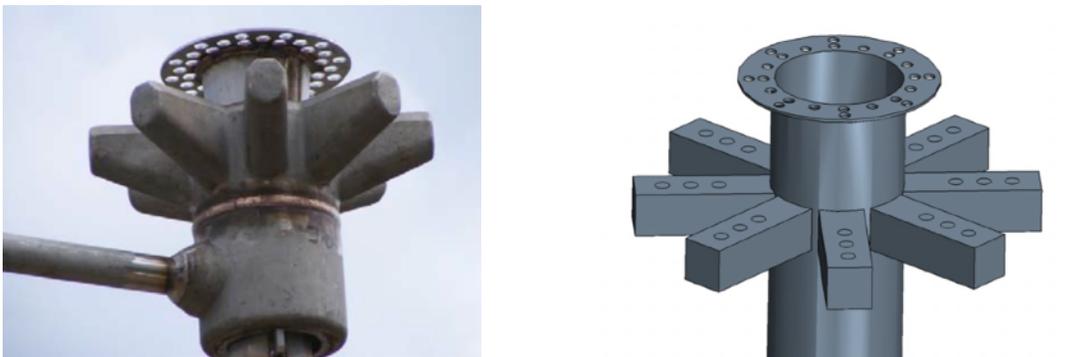


Figure 1: Image of the flare actual flare head (left) and the reconstructed geometry used in the LES simulation (right).

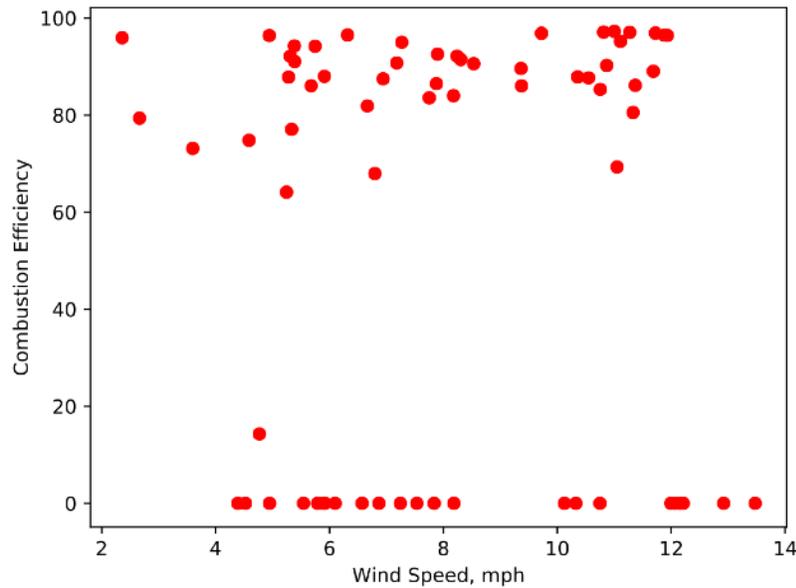


Figure 2: Measured combustion efficiency percentages as a function of the observed wind speed. Values at zero indicate failed readings.

range of observed efficiencies range from a low of 14% to as high as 97%. Excluding data below 50%, an average combustion efficiency is computed as 90%.

Using the CAD mockup as described above, Arches LES simulations were performed and combustion efficiencies were obtained for a wind speed of 10mph. Simulations were run on about 320 processors and consumed about 23,000 CPUhrs/case. Combustion efficiencies were obtained by extracting time-averaged data along lines of sight downstream from the flare. An example is shown in Figure 3, where the temperature field at a specific time has been volume rendered and a line of sight is shown by the red cylinder.

The simulated line-of-sight combustion efficiency varies as a function of elevation. This is demonstrated in Figures 4 and 5. Figure 4 shows the combustion efficiency as a function of line-of-sight elevation with the average line-of-sight temperature reading. For reference, Figure 5 is a two-dimensional contour plot of the combustion efficiency at the same downstream position. The simulation appears to be predicting values within the range of the experimentally measured data for the same wind conditions.

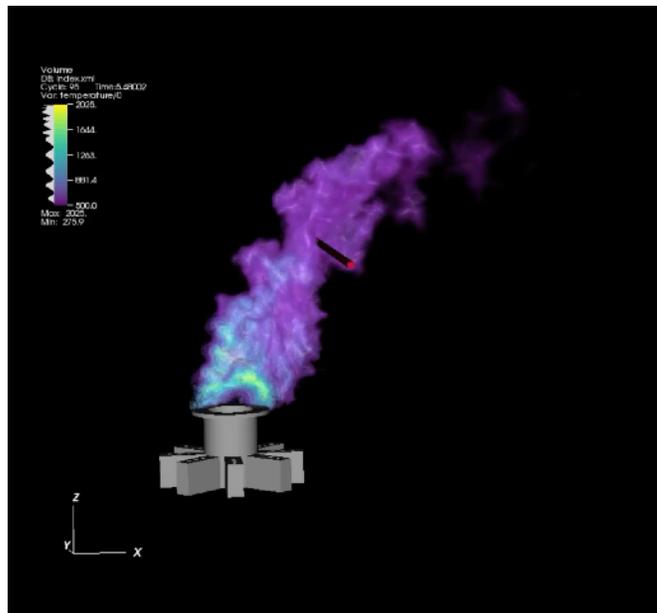


Figure 3: Volume rendered temperature with the red cylinder showing an example of the ID line-of-sight extraction from the LES data.

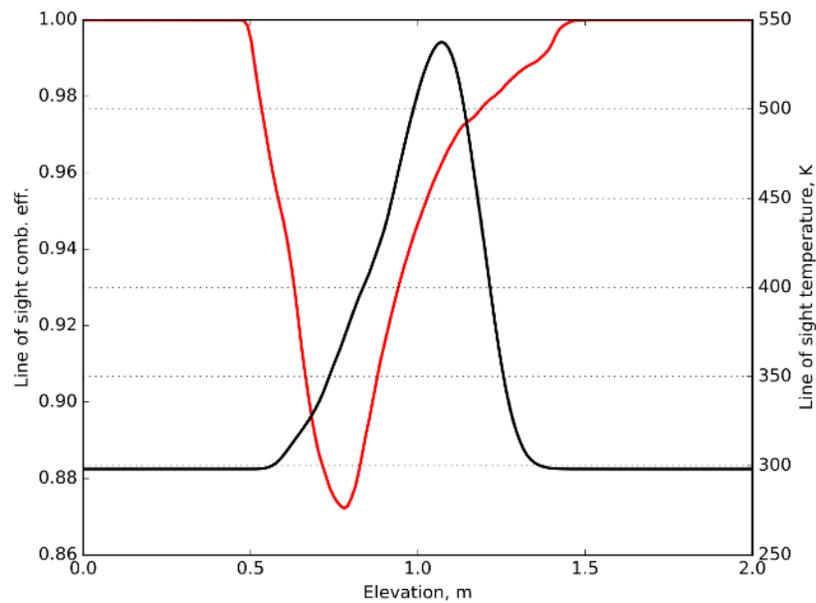


Figure 4: Combustion efficiency (red line) and temperature (black line) line of sight LES predictions at a location downstream from the flare.

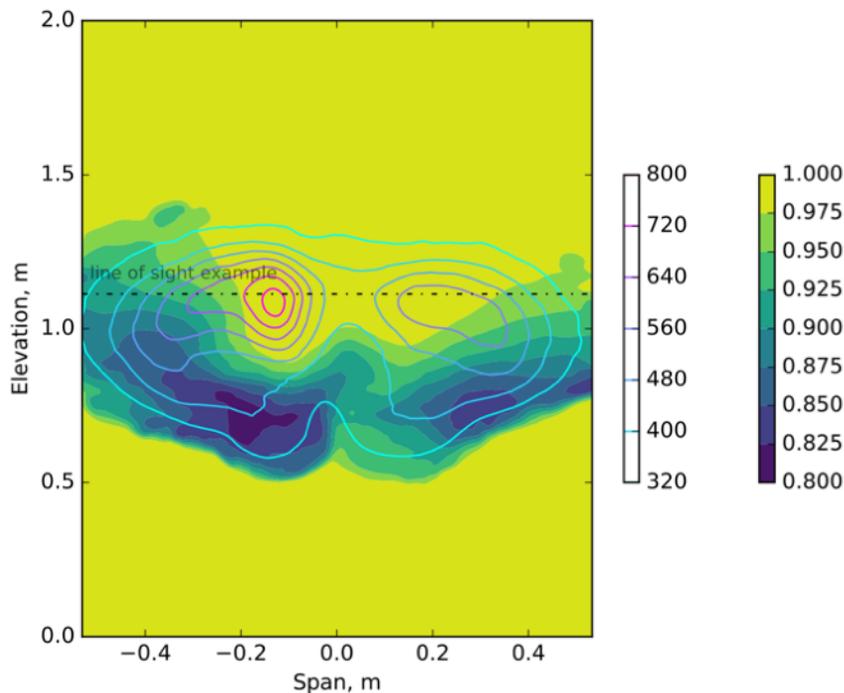


Figure 5: A two-dimensional filled contour plot of the combustion efficiency and line contour plot of temperature at the same downstream location.

Instrument Modeling Results - Validation with UQ

An instrument model (IM) of the PFTIR physical instrument is constructed with three inputs. These inputs are the physical location of the PFTIR instrument relative to the flare, the environmental wind direction, and the location of the target for the PFTIR instrument. Note that several other parameters could be included, including calibration parameters. However, these were omitted to keep the scope limited.

Each input to the IM representation of the PFTIR instrument has a measure of uncertainty. The uncertainty is characterized as a distribution and an assumed distribution on each parameter is assigned. The assumed distribution is the prior distribution and is made with a best engineering guess. Given a single LES simulation at a specific wind speed condition and steam flow rate, the averaged CO_2 , CO , and unburnt hydrocarbon data is sampled over a range of target locations and wind directions. From this data, combustion efficiencies are computed. This maps the input variables to a response in combustion efficiency. From this response, a function is fit using Gaussian Processes. The fit function will then serve as a surrogate to the data sampling. The surrogate is required to speed up the sampling of the distributions in the construction of the likelihood estimation.

With the assumption on the priors and a method for computing the likelihood, we can proceed with computing the posterior parameter distributions. Recall that these parameter distributions represent the informed parameters because they used the measured data to arrive at their updated state.

A total of 20 separate LES simulations were performed at different windspeed/steam flow rate conditions. Using one of the runs, we perform the BML method to arrive at the posterior distributions and the predictive posterior combustion efficiency. This can be done for all cases. Here, we show an example of two of the 20 cases.

Figure 6 shows the IM targeting superimposed over an instantaneous and time averaged LES temperature field. The procedure for picking the target was to time average the temperature field, and then place the target near the end of the combustion zone, not unlike the technique used in the actual experiment. Once the target was placed, it was then varied about that point to generate a sample set from which the surrogate model could be constructed.

After the targeting samples were chosen, the PFTIR IM was placed at an observation point. This positioning is estimate, but Google satellite imagery and distance estimates were used while cross referencing the qualitative reporting from the Marathon Data Report on the actual PFTIR placement for the experiments. After the placement of the virtual instrument, the instrument was then rotated about the flare stack to mimic the effect of the changing wind direction. This approach allows us to sample a wide range of wind direction angles without rerunning an LES simulation. In other words, the LES flare stays in a fixed reference frame. For a typical sampling, the angles ran from -150 to +150 degrees relative to the initial positioning of the PFTIR placement. Figure 7 shows an example of such sampling with the combination of the variation in the target position for a single case.

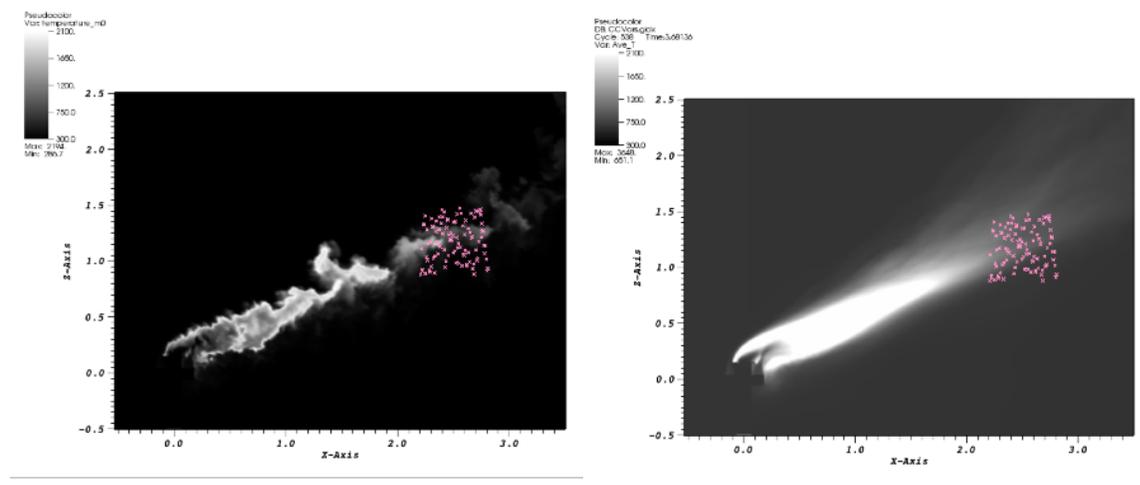


Figure 6: Instantaneous (left) and time-averaged (right) temperature plots with the PFTIR instrument model target samples shown with pink markers.

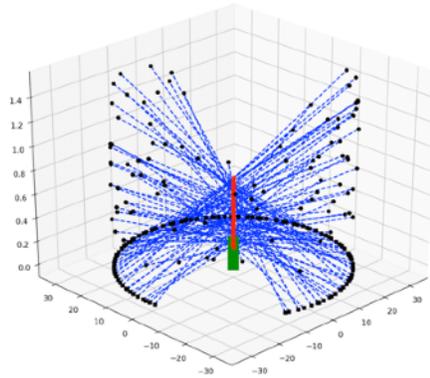


Figure 7: Example of the PFTIR instrument model sampling showing the variation in the targeting position and the wind direction effect. The flare stack is indicated with the green bar. The intersecting targets are shown with red dots. The starting and stopping locations of the PFTIR ray are indicated with black dots while the ray path is shown with the dashed blue line.

Although it wasn't explicitly reported with the data, an assumption made here is that the target positioning in the experiment is held constant throughout the data acquisition period and is likewise held constant in the IM. If our assumption is incorrect and if the targeting was adjusted during the experiment, we would expect this to generate a bias in the analysis.

The results of the IM are now discussed for two experimental conditions; a case with a moderate wind/low steam rate (Condition 3 from the report) and moderate wind/high steam rate (Condition 6 from the report). For each individual case, the steam flow rate and crosswind speed condition is held constant at the experimentally reported mean condition. Once the LES case is complete, the data is post-processed to include variation of the IM parameters of wind direction target placement. Note that the target placement is parameterized by a downstream location and an elevation above the flare head. A total of 100 samples are taken varying the IM parameters using a Latin Hypercube sampling technique. After the sampling, a response surface is generated from the data. This response surface is then passed into the BML algorithm to produce posteriors from the priors distributions. For the BML algorithm, we choose four parameters:

- The variation of the angle of the wind, σ_ϕ
- The downstream target location relative to a mean, \bar{d}
- The target elevation relative to a mean, \bar{h}
- All other sources of uncertainty, σ_u

Regarding σ_u , this parameter represents a measure of the variation on the experimental and modeled data and represents uncertainty from unaccounted sources. When σ_u is quantified in the posterior, it serves as a representation of the standard deviation of the uncertainty about the measured point when the modeled and measured data are combined through the BML algorithm. Also, regarding the target position, because the target is a spatial location we center and normalize this data so that it represents a relative distance from some given target. This is done by simply subtracting the mean of all values from any specific value and also scaling it by the mean value.

Each of the four IM parameters are characterized with a prior distribution. Ideally, the prior distribution should be realistic and use the best state of knowledge for each parameter. For the variance in the wind, we use a lognormal distribution. For the target location, we use a normal distribution, and for σ_u we use a Jeffrey's prior, which is an uninformed prior. These distributions are fed into the BML algorithm resulting in a set of posteriors on the parameter space. Since this is a four-dimensional space, we show the joint and marginal distributions in Figure 8. We note that the σ_u value for Condition 3 is more than 2.5 times less than that for Condition 6. It appears from the data that Condition 6 experienced the most variability of any of the six conditions. This is demonstrated in Figure 9 where we have plotted the 95% confidence interval on the raw experimental data after removing missing entries and reported values of zero combustion efficiency.

Upon collection of the marginals from the BML posteriors, we then use this information in a forward manner through the physics or surrogate model to predict values of combustion efficiency at various conditions. Here, the logical choice for the physics model is the surrogate model is used to predict combustion efficiency given a range of inputs. Because we are dealing with probability distributions, we can analyze the data in the form of distributions or simply moments of the distribution. Figure 10 shows a sampling of the distributions for a few points in the measurement space to illustrate how the distributions compare from the measured data distribution (assumed normal with mean as the measured point and $\sigma = \sigma_u$) and the predicted distributions. Note that efficiencies greater than one can be eliminated with a proper variable transformation which was not done here.

Conclusion

In this paper we have outlined an approach for pairing PFTIR data with advanced LES simulation to move beyond the usual model validation viewgraph norms that are typically presented. In this approach, the view of data sources is more encompassing, in that all data sources are viewed as potentially valuable as long as they can be characterized with uncertainty distributions. The aim is to combine these methods in a Bayesian Machine Learning framework to promote learning on all data sources. This is accomplished using Bayes law. With the expensive LES simulations, we also employ response surface modeling to serve as a surrogate to the expensive LES once the input parameter space has been sampled to a good degree. The end result of this analysis is (potentially) an increase in data learning of all sources. Additionally, because of the use of a physics-based model, other unmeasured (or not easily measured) quantities can be inferred. The entire methodology can be packaged into code that rapidly provides feedback in a matter of seconds when the response surface is used. Such a form would be useful for analysis and control of the flare in real-time.

Work is underway to expand the parameter space using the BML approach. This expanded parameter will result in an instrument model that could take online data for the steam flow rate, wind speed, and wind direction, and produce a line of sight combustion efficiency. It will also allow inference of other modeled parameters that the LES could provide, such as temperature predictions, heat flux measurements, and integrated combustion efficiency over a volume. This model, sometimes referred to as a digital twin, could be used for control, operational insight, or design of flares.

The instrument developed here executes in the matter of seconds and could provide real-time feedback to flare operations. Furthermore, additional gathered data from PFTIR

measurements could update the posterior distributions through the BML algorithm and propagate new measurements downstream to the IM through the posterior predictive. The propagation occurs in minutes. If flare operations go beyond the bounds explored by the LES

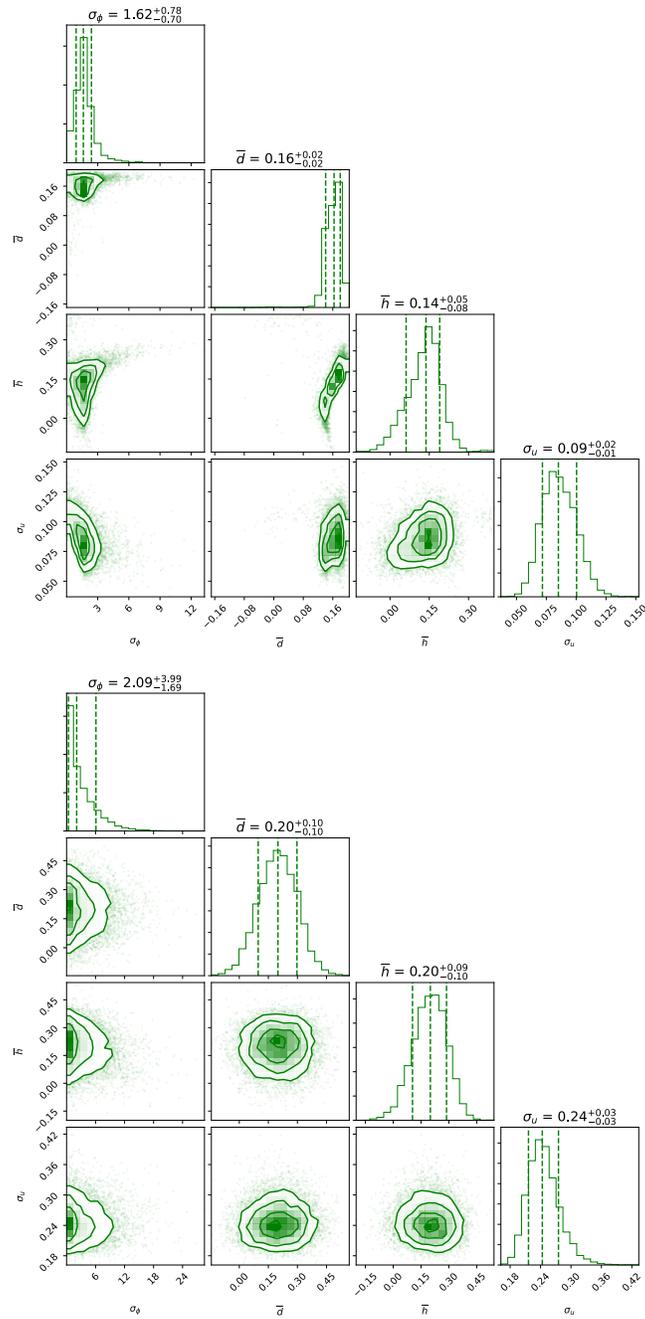


Figure 8: The posterior joint and marginal distributions for Conditions 3 (top, moderate wind/low steam) and Condition 6 (bottom, moderate wind/high steam). Shown on the plot are the numerical values of the 15th, 50th, and 85th quartiles.

simulations and included in the response surface model, then new LES simulations should be added to expand the parameter space. A typical LES simulation of the flare examined require about 2-3 days on roughly 1.5K computing cores.

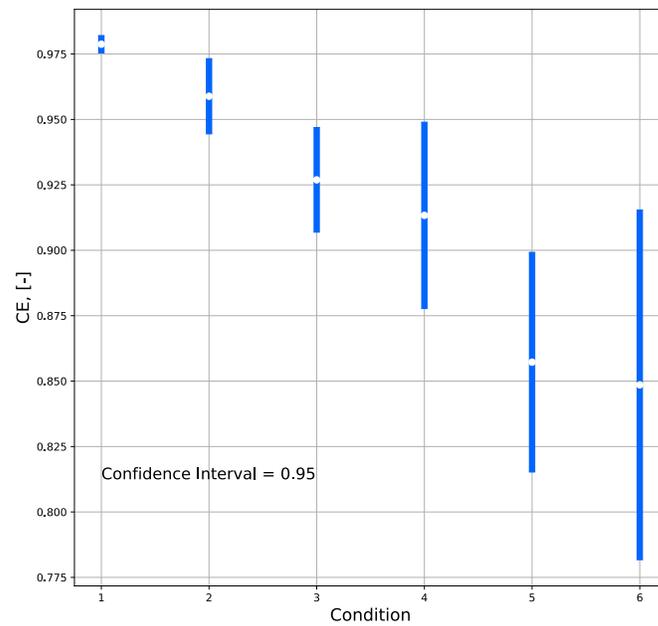


Figure 9: A plot of the 95% confidence interval from the raw measured data used to represent the relative variations in the combustion efficiency per case. The mean values are shown with the white dots.

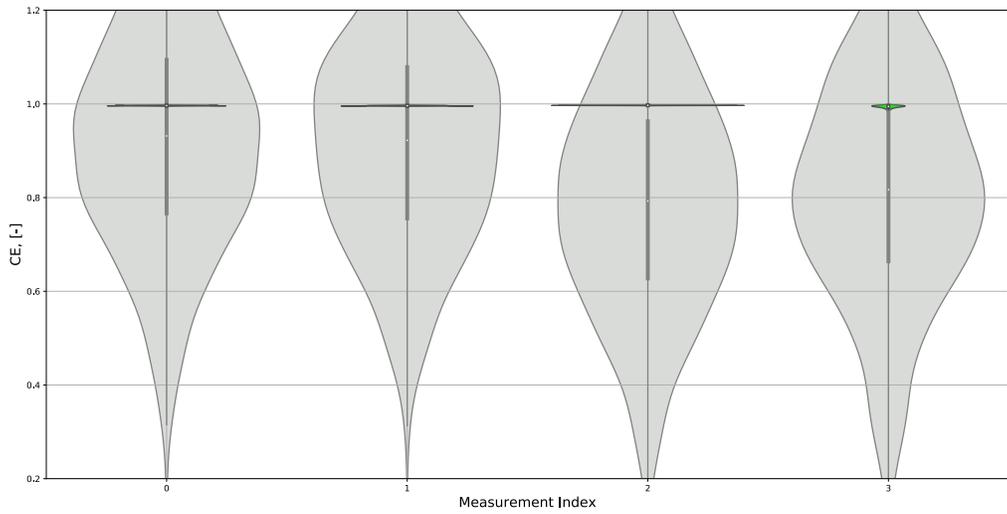
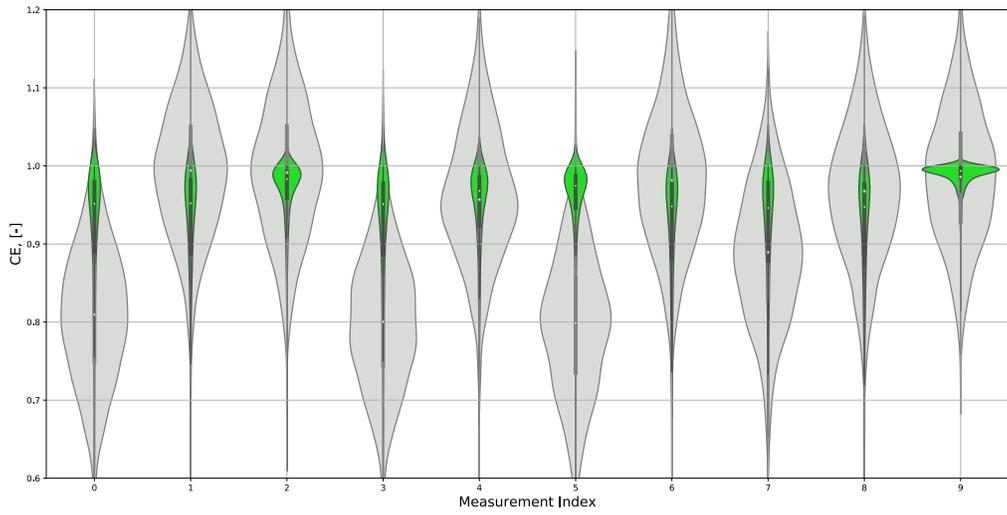


Figure 10: Distribution predictions of measured (gray) vs. predicted (green) values of combustion efficiency for Conditions 3 (top) and 6 (bottom). Note that Condition 6, Replicate 1 had only four valid readings of CE.

This material is based upon work supported by the U.S. Department of Energy, Office of Science, under Award Number DE-SC0017039. This paper was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.