Fire Technology, 56, 445–467, 2020

https://doi.org/10.1007/s10694-019-00886-w

Model Considerations for Fire Scene Reconstruction Using a Bayesian Framework

Andrew Kurzawski, Jan-Michael Cabrera and Ofodike A. Ezekoye, Department of Mechanical Engineering, The University of Texas at Austin, Austin, TX, USA*

Received: 17 January 2019/**Accepted:** 25 June 2019

Abstract. Towards the development of a more rigorous approach for coupling collected fire scene data to computational tools, a Bayesian computational strategy is presented in this work. The Bayesian inversion technique is exercised on synthetic, time-integrated data to invert for the location, size, and time-to-peak of an unknown fire using two well-known forward models; Consolidated Model of Fire and Smoke Transport (CFAST) and Fire Dynamics Simulator (FDS). A Gaussian process surrogate model was fit to coarse FDS simulations to facilitate Markov Chain Monte Carlo sampling. The inversion framework was able to predict the total energy release by all fire cases except for one CFAST forward model, a 1000 kW steady fire. It was found that insufficient information was available in the time-integrated data to distinguish the temporal variations in peak times. FDS performed better than CFAST in predicting the maximum energy release rate with the posterior mean of the best configurations being 0.05% and 2.77% of the true values respectively. Both models performed equally well on locating the fire in a compartment.

Keywords: Fire forensics, Bayesian statistics, Inverse problems, Surrogate modeling, Fire dynamics simulator

1. Introduction

The fire forensic analyst has a challenging job. The analyst is charged with identifying and collecting data/evidence in the severely damaged post-fire environment with a goal of using these data to determine the fire evolution. The data gathering process includes visually recording the fire scene and taking measurements and samples at the scene [1]. Despite the challenges, the technical foundations continue to improve for fire forensic analysis. For example, researchers are quantifying the statistical properties of damage patterns generated by fires in otherwise similar fire conditions [2–5]. These patterns are commonly used to determine the fire's origin and evolution. Hicks et al. conducted a controlled set of experiments with duplicate fuel packets to determine the repeatability of producing three commonly

* Correspondence should be addressed to: Ofodike A. Ezekoye, E-mail: dezekoye@mail.utexas.edu

noted fire damage patterns on the walls of compartments: triangle, column, and conical [2, 3]. It was found that the shapes are repeatable for these controlled tests. Madrzykowski and Fleischmann examined the repeatability of damage patterns on walls for three different fuels and quantified the uncertainty in the heat release rate, flame height, and dimensions of the fire pattern [4, 6]. In similar work, Mealy et al. evaluated floor damage patterns and found that they were not consistent enough to be useful in room fires with heavy soot deposition [7]. As improvements in fire characterization have taken place, there has also been increasing awareness of the general need to improve the scientific and engineering foundations in many other areas of forensic science.

One area in which improvements have been made in forensic science is the standard by which experts and evidence can be used in trials (e.g. Daubert vs. Merrell Dow Pharmaceuticals [8]). Despite these improvements, there are still holes in the level of reliability and validity of evidence that can be introduced in prosecutions. To identify these holes, various engineering and science organizations including the U.S. National Academies and the AAAS have published reports on deficiencies in the application of forensic science in the US [9, 10]. Among other issues, a theme evolves on the need to improve the statistical weighting of any given hypothesis and/or elements of evidence. Even for areas of forensic science that are more strongly supported by scientific research (e.g., use of DNA analysis), there continue to be issues related to the probability of match as noted by Nordgaard et al. [11]. In that paper, Nordgaard et al. discuss and show by example the value of a Bayesian approach (i.e. use of the Bayes factor) in developing a more formal statistical process for hypothesis testing. This methodology for hypothesis testing could aid hypothesis testing in many areas of forensic science, including fire forensics.

Bayesian methods have been used successfully in a variety of fire problems, for example, Wang and Zabaras use Bayesian inversion techniques to find the magnitude of a heat source from temperature data [12]. Biedermann et al. note the importance of uncertainty in drawing inferences and applied Bayesian networks to forensic fire investigation problems involving the detection of flammable liquid in debris [13, 14]. Guo et al. used Bayesian inversion with a zone model as the forward model along with measured gas temperatures to find the fire size and room of origin within a multi-compartment structure [15]. Overholt and Ezekoye [16], use Bayeseian techniques to invert for the heat release rate of the fire as a function of time using temperature data from room-scale experiments. Overholt and Ezekoye also investigated the use of heat flux data to determine either the size or the location of a fire in a compartment [17, 18].

In the types of studies cited above, inversion techniques for fire properties rely on the availability of temporally evolving measurements of temperature or heat flux. In reality, a fire investigator does not have access to these types of temporal variations of detectors within the compartment. More often, the investigators have an indication of the total amount of heating in the form of damage to many of the surfaces in the compartment. As an abstraction of the type of damage experienced by common items in a compartment fire, a surrogate sensor system can be used to translate thermal exposures to damage metrics. This research seeks to

address this by building an inversion framework that uses a relatively simple measure of thermal exposure, total energy per unit area (abbreviated TEA), as a data source to issue predictions on the location and size of a single unknown fire.

Fire models are increasingly used to aid in hypothesis testing in fire forensic cases. There is need for guidance on the overall predictive capability of these models and the quality of data required to support a hypothesis. Towards the goal of improvements in the use of models for hypothesis testing in fire scenarios, we discuss in this paper the feasibility of a Bayesian statistics framework to identify the location and heat release rate of a fire. Admittedly, this work has a theoretical focus and uses data and processes not currently used by practicing fire investigators. The intent of the paper is to clarify what aspects of models and data must be improved to issue predictions and lay a path forward for more rigorous and statistically credible use model based hypothesis testing. Better connection to the data used by practitioners will be added as the basic framework is better developed. In this paper, several types of models are used to process the synthetic fire damage data to understand the costs associated with this framework and evaluate model performance. Additionally, an overview of Bayesian inversion methods for forensic analysis is provided.

2. Bayesian Inversion Framework

Bayesian inversion can, in general, be used to estimate the probability of unknown fire parameters when conditioned on observed data. Here, the methods outlined in Overholt and Ezekoye [16–18] are expanded to simultaneously estimate the distributions of fire location, heat release rate (HRR), and time to peak for various scenarios. Additionally, these methods relied on time-resolved data to invert for properties of the fire. As fire investigators do not have access to this level of time-resolved data, the methods of Overholt and Ezekoye have been extended to only use time-integrated data in the form of a simple damage model.

For this work, it is assumed that some true fire with an unknown location and size produces a damage signature (in the form of total energy per unit area or TEA) on the walls and ceiling of a room. Various methods exist to assess the extent of damage to objects, and TEA is perhaps the simplest of these. Perhaps more accurate than TEA are methods that evaluate the damage to a material as in the use of kinetic processes to describe material degradation (i.e. core samples in degraded gypsum, wood char growth etc.). To exercise this methodology, we have chosen to first explore the use of TEA for a damage estimate in lieu of a material specific kinetic model as it represents a total, time-integrated impact at a set of locations in a room following a fire.

The Bayesian inversion framework was then used to estimate the probability distribution of the unknown fire parameters conditioned on observed TEA data via Markov chain Monte Carlo (MCMC) sampling with a Metropolis-Hastings stepping algorithm as implemented in the software package PyMC [19, 20]. We seek to estimate the posterior distribution of the unknown model parameters ($\boldsymbol{\theta}$) given some data ($\mathbf{d}$) by using Bayes' formula:

$$P(\boldsymbol{\theta}|\mathbf{d}) = \frac{P(\mathbf{d}|\boldsymbol{\theta})P(\boldsymbol{\theta})}{P(\mathbf{d})} \tag{1}$$

with a Gaussian likelihood given by

$$P(\mathbf{d}|\boldsymbol{\theta}) = \frac{1}{\sigma_e\sqrt{2\pi}} exp\left(-\frac{(\mathbf{d_{obs}} - \mathbf{d_{pre}})^2}{2\sigma_e^2}\right) \tag{2}$$

where $\mathbf{d_{obs}}$ is the observed data (from the true fire), $\mathbf{d_{pre}}$ is the predicted by a forward model, and σ_e is the standard deviation of the likelihood function. The forward model is a numerical representation of the fire behavior that takes the unknown model parameters as inputs and produces predictions of the fire behavior (TEA in this case).

The unknown model parameters ($\boldsymbol{\theta}$) are the fire location (x_{fire}, y_{fire}), peak heat release rate ($\dot{Q}$), and standard deviation of the likelihood (σ_e). Additionally, we introduce a time to peak parameter (t_{peak}) to represent a fire with a triangular HRR profile in some cases.

In the the simplest case, the prior distributions ($P(\boldsymbol{\theta})$) for each model parameter are assumed to be uniform (annotated $U(l, u)$) with lower bound l and upper bound u:

$$\dot{Q} \sim U(Q_{low}, Q_{high}), \quad x_{fire} \sim U(x_{low}, x_{high}), \quad y_{fire} \sim U(y_{low}, y_{high}) \tag{3}$$

with the additional t_{peak} parameter only in the triangular fire case

$$t_{peak} \sim U(t_{low}, t_{high}) \tag{4}$$

and a uniform prior on the hyper-parameter σ_e with a high upper bound to reflect little knowledge of the standard deviation of the error in the data. If experimental uncertainties were available, they could be used to estimate σ_e.

$$\sigma_e \sim U(0, \sigma_{high}) \tag{5}$$

These uniform priors encompass the belief that the fire could be within the specified HRR range Q_{low} to Q_{high} and located anywhere within the room. An expert investigator could examine a fire scene and incorporate their beliefs on the location and size of the fire in to the prior distributions. This would change the form of the priors to more heavily weight the expert's opinion.

Draws from the posterior distribution are generated using a Metropolis-Hastings sampler initialized at a maximum a-posteriori (MAP) estimate from Powell's method [21]. Tuning of the standard deviations of the proposal distribution for each parameter occurs every 1000 samples. The tuning process targets a ratio of 0.2 to 0.5 overall accepted proposals by reducing the standard deviation if the ratio is below the target window and increasing the standard deviation if the ratio is above the target window so that greater steps are taken between points [19].

We choose to discard (burn) samples until the diminishing tuning condition from Roberts is met such that tuning should be small and infrequent as the sampler progresses [22]. Burning of samples is stopped once the tuning of all parameters changes less than the threshold of 0.1, after which all samples are saved. In practice, this occurs after approximately ten to fifteen thousand samples for this specific model. A termination point for the sampler is chosen using the Geweke score for early chain means [23].

3. Models

The Bayesian framework is exercised for inversion in the single room compartment geometry depicted in Fig. 1. For this initial examination of the framework, synthetic data generated by a high fidelity model are used as the observed data while two lower fidelity models are used as the forward models in the inversion process.

3.1. Domain and Synthetic Data

The Bayesian room-scale localization framework was tested on synthetic data generated from a series of Fire Dynamics Simulator version 6 (FDS) simulations with moderately fine meshes (5 cm mesh). One case was run with a finer, 2.5 cm, mesh to check grid sensitivity and the TEA values were found to be within 6.6% of the values from the 5 cm mesh. If one wished to compare the FDS results to experiments the 2.5 cm mesh would be more desirable, however 5 cm was selected as a

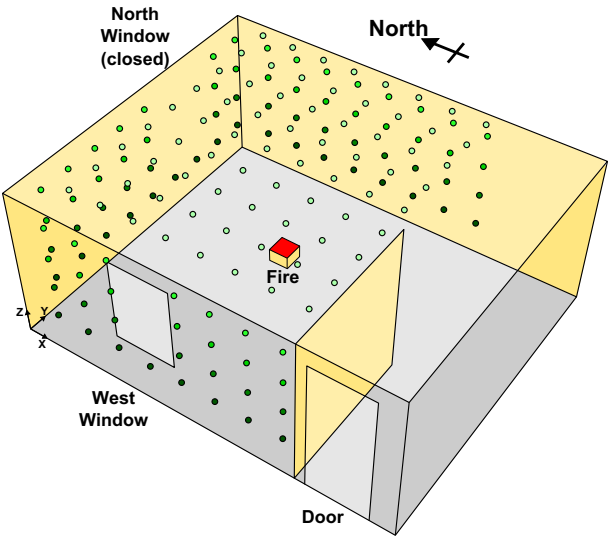

Figure 1. Layout of experimental structure used for model. Dots represent sensor locations where dots become darker with decreased elevation.

sufficiently "fine" mesh for this work to balance accuracy and computational cost as the 2.5 cm simulation requires approximately 20 times the computing time of the 5 cm mesh. FDS has an extensive set of verification and validation cases that aid the user in determining if FDS is acceptable for their application and is widely used for modeling compartment fires [24, 25]. The FDS simulations in this work each consisted of a single fire at varied locations in a room with fixed ventilation conditions. Simulations take upwards of 24 h of wall clock time on 4–6 cores running at 2.2 GHz for 5 min of simulated time. Synthetic data were extracted from the simulation in the form of time-integrated net heat flux at discrete points on the room's gypsum walls and ceiling with a specified level of noise added to the raw values of either 25% of the measured value or no noise. As previously noted, the time-integrated net heat flux is considered to be a surrogate for damage.

The room dimensions are 4.2 by 4.58 m with a ceiling height of 2.24 m as seen in Fig. 1. There are two windows on the north and west walls of the room with a hallway on the south side of the room leading to a doorway on the west wall. For this work, the window on the north wall was closed while the door and west window were left open. All interior surfaces of the structure are assumed to be lined with 1.6 cm thick gypsum wallboard. Virtual heat flux sensors were spaced as close as 0.5 m apart along the west, north, and east walls and ceiling of the main compartment resulting in a maximum total of 163 sensors.

To begin moving toward the ultimate goal of finding the origin of a fire in a room with many fuel sources, we simplify the problem to finding the location of a single fire in a compartment to develop the inversion framework and understand the scope of the framework's capabilities. Secondary to the location of the fire, we seek to determine how much temporal information on the fire can be extracted using only time-integrated data, i.e. with what level of detail can we resolve the heat release rate (HRR) as a function of time. As is well known in the fire research community, the HRR is the most important quantity characterizing a fire.

It is common to represent the HRR of a burning item as one of a set of simplified shapes. The simplest parameterization of a fire is a steady HRR defined by two parameters (HRR and duration), however this is generally only applicable to fires generated in a laboratory. We will use the case of the steady fire as it represents the simplest experiment that can be conducted in a lab and can function as a test case for exploring the behavior of the Bayesian inversion framework. The so called "triangle" fire is the next simplest parameterization and is defined by three parameters: peak HRR, time to peak HRR, and duration. This parameterization can be used to approximate the HRR of items such as trash bags and Christmas trees [26]. It is anticipated that this framework is extendable to other fire conditions and fuel types. For example, items such as upholstered furniture and office workstations generally require more complicated HRR parametrization and often use the t^2 fire model which includes growth, steady burning, and decay times. Further, post-flashover fires will not be covered in this work, but they will be important for future versions of the inversion framework. In a real compartment fire under forensic investigation, if only a single item were involved, finding its loca-

tion is trivial. It is important to first evaluate the feasibility of using a damage model to locate and size a fire before applying the model to more realistic cases where multiple items become involved and flash-over may have occurred.

For this study, four synthetic data sets were generated with FDS featuring two HRR profiles (steady and triangular fires) and two maximum HRRs (200 kW and 1000 kW). Each case was run for 300 s of simulation time, where the steady fire is at the maximum HRR for the entire duration of the simulation and the triangular fire begins at 0 kW, increases linearly to the maximum HRR at 90 s and then decreases to 0 kW at 300 s. In each case the fire was centered at 1.95 by 2.95 m from the northwest corner of the room.

An example of the TEA on the walls and ceiling for the steady 200 kW fire is shown in Fig. 2. The shaded contours show the area on each surface covered by the sensors and the boundaries of the walls and windows are marked with black lines. Windows marked with an "X" are closed and windows with a white fill are open. The fire has a diameter of 0.32 m and its outline is marked on the top-down view of the ceiling with double-width black lines.

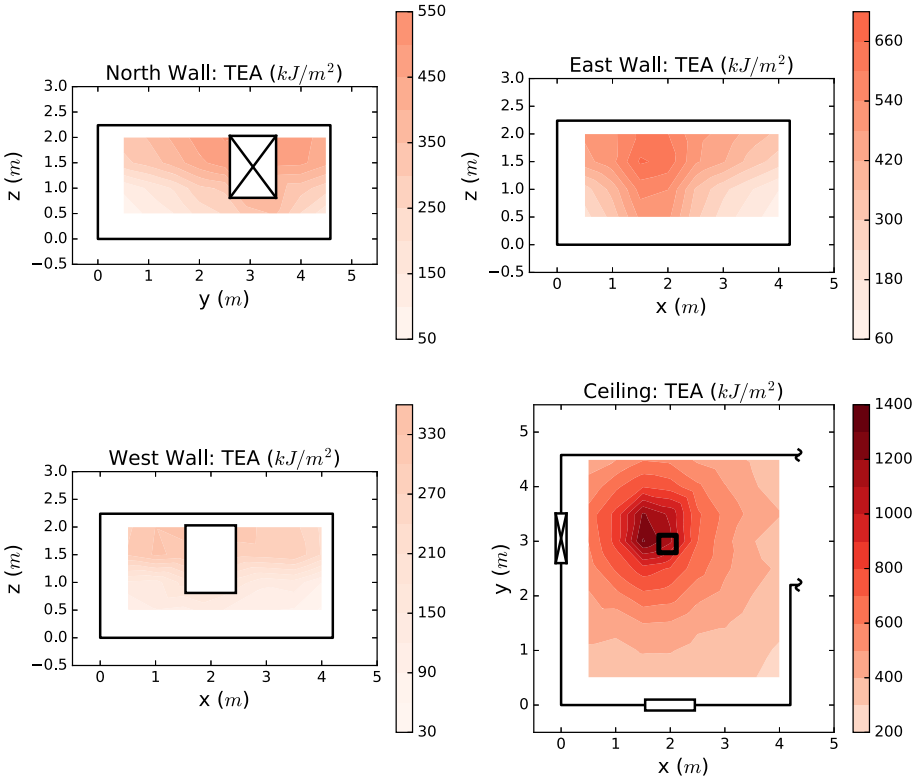

Figure 2. Synthetic TEA data (kJ/m²) for a 200 kW steady fire where the small square represents the fire location.

The heating patterns seen in each HRR size and shape combination are qualitatively similar, differing primarily in the magnitude of the TEA imparted to the surfaces which scales with the total energy released by the fire. There is a clear maximum in the TEA where the plume impinges on the ceiling, indicating that the sensors on the ceiling will be important for determining the location of the fire. The plume tilts in the negative x direction due to the flow pattern that develops within the compartment where most of the inflow comes from the hallway and moves through the compartment counter-clockwise until exiting through the open window. In the 200 kW fire case, there is a tilt in the positive y direction that is not present in the 1000 kW cases indicating greater influence from the inflow through the open window on the weaker fire plume.

We consider a control volume around the gas in the main compartment containing the fire with an energy balance of

$$\frac{d}{dt}E_{stored} = \sum \dot{m}_{in}h_{in} - \sum \dot{m}_{out}h_{out} - \dot{Q}_{surf} - \dot{Q}_{rad} + \dot{Q}_{fire} \tag{6}$$

where the rate of change in the energy stored within the control volume is equal to the rate of energy entering the control volume minus the rate of energy exiting the control volume. The energy flows are classified as flow enthalpies ($\dot{m}h$), heat loss to the solid surfaces ($\dot{Q}_{surf}$), heat loss via radiation through openings ($\dot{Q}_{rad}$), and the HRR of the fire ($\dot{Q}_{fire}$).

Two additional triangular fire simulations (one for each peak HRR) were run with a time to peak HRR of 210 s, and the energy budget was compared to the 90 second peak fire. Note that the change in time to peak does not change the total energy released by the fire as the compartment does not become under-ventilated in any of the 200 kW or 1000 kW HRR triangular fire cases. The energy flows are the temporally and spatially integrated components of the loss terms from Eq. 6 separated into the energy imparted to the ceiling, walls, and floor as well as the flow energy exiting through the open window and doorway. These quantities are integrated in time from the inception of the fire up to the extinction time to make general comparisons between bulk energy flows. Each term is calculated as a percentage of the total energy released by the fire (i.e. the time integral of $\dot{Q}_{fire}$).

Approximately 40% of the energy from the fire is absorbed by the solid surfaces, while over 50% of the energy exits the compartment in the form of gas flowing out of the open window and in to the hallway. The remainder of the energy either exits the open vents through radiation to the surroundings or remains inside the compartment. There is an approximately 1% difference in flow energy between different peak times of 90 and 210 s for both the 200 and 1000 kW cases and less than 1% difference in the heat imparted to the ceiling and walls. This indicates that even without noisy data, it may be difficult to determine the time to peak HRR from TEA data, however one can imagine that this breakdown of heat flows would be different for a more restrictive ventilation condition resulting in a greater amount of heating on the walls and ceiling.

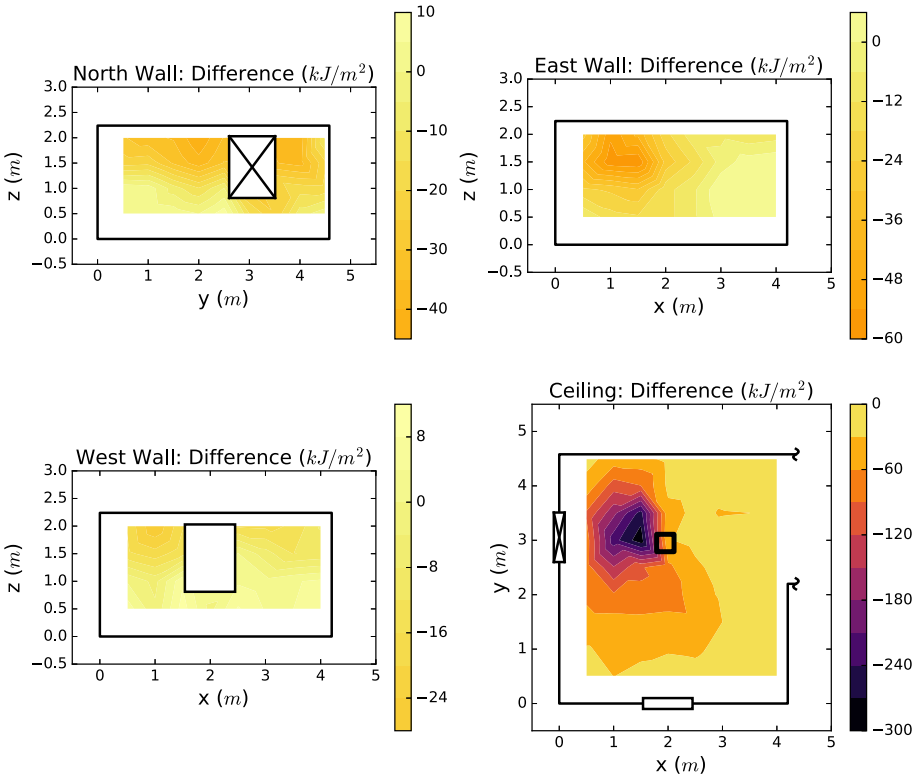

Figure 3. Difference in TEA over each surface between peaks at 90 s and 210 s for a 1000 kW triangle fire.

Figure 3 compares the difference in local heating for the 1000 kW HRR fire with different peak times (i.e. $TEA_{90s} - TEA_{210s}$). The maximum difference in TEA occurs at the sensor to the left of the fire and it is on the order of 300 kJ/m^2 or about 7% of the TEA at that point. This indicates that it will be challenging to invert for the time to peak given that the difference in total heat imparted to the solid surfaces between the two peak times is small (<5%).

3.2. Forward Models

As previously noted, the MCMC process requires a large number of forward model calculations to exercise the Bayesian sampling algorithm. As such, significant effort is required in generating computationally fast models. For this work, we describe two types of models used for the MCMC computations. The first is a physics based model, CFAST. The second is a data generated model constructed as a Gaussian process model.

3.2.1. CFAST The first forward model to be used for the inversion process is the Consolidated Model of Fire and Smoke Transport Version 7 (CFAST) which is a

two-zone fire model primarily used for simulating the behavior of compartment fires. The zone model is a simplified representation of a compartment fire that relies on a combination of physics and correlations to predict the bulk mass and energy transport processes for fires that occur in a compartment or series of connected compartments. Net heat flux data (convection from surrounding gas, radiative exchange with surfaces, gas, and fire) is collected at virtual sensors located on the walls of the compartment and integrated in time with the trapezoidal rule in a post-processing step after the simulation is complete to approximate the TEA.

To understand the expected performance of CFAST as a forward model in the Bayesian inversion framework, an error surface was explored for the triangular HRR cases with peaks at 200 kW and 1000 kW. As this requires searching over four input parameters, two input parameters were fixed at the true values while the other two parameters were divided into an evenly spaced grid. At each grid point, CFAST was run for the corresponding parameter set and the l^2 error (E_{grid}) between the CFAST data ($y_{i,grid}$) and the synthetic data generated by FDS ($y_{i,synth}$) was evaluated at $P = 163$ sensors.

$$E_{grid} = \sqrt{\sum_{i=1}^{P} (y_{i,synth} - y_{i,grid})^2} \tag{7}$$

Additionally, Nelder-Mead optimization was used to find the input parameter pair with the minimum error by starting at several randomly sampled points, running the optimizer until convergence, and choosing the point with the lowest error [27]. Local minima will be reported where applicable.

In the l^2 error contour plots of Fig. 4a, the fire location was varied with a fixed peak HRR and time to peak. The minimum error is located to the left of the true fire origin (as denoted by the bold square) and upon examination of the synthetic TEA data on the ceiling, it can be seen that this coincides with the location of the maximum TEA imparted to the ceiling due to the tilting of the fire plume. CFAST does not model the flow field within the compartment and assumes that the fire plume travels straight up from the fire's origin. Therefore, we expect this bias to be reflected in the inversion results. Note that there is a local minimum near location (0, 0) in both HRR cases that is rarely encountered by the optimizer.

The l^2 error contour plot for a fixed location with varied peak HRR and time to peak for both HRRs is shown in Fig. 4b. There is a clear valley near the true peak HRR (i.e. 200 kW) that extends in the direction of the "time to peak" for all potential peak times. The findings of the energy budget analysis from Sect. 3.1 corroborate the existence of this valley as there was little difference in the energy imparted to the walls for peak times differing by 110 s. The minimum error is located at a peak time of 1 second, however the optimizer frequently became stuck at other points along the time to peak valley.

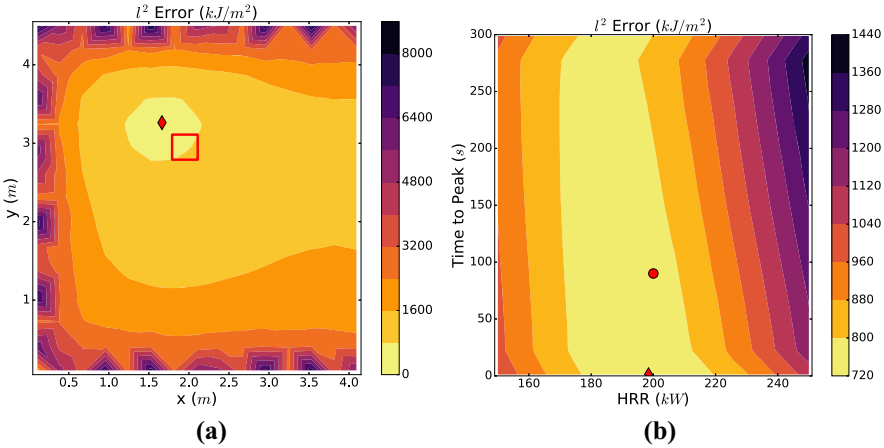

Figure 4. CFAST fire location error surface (a) and HRR and time to peak HRR error surface (b) for a 200 kW fire. The outlined box in (a) represents the boundary of the true fire from the fine mesh FDS simulation. The circle in (b) is the true HRR and time to peak. In both figures, the diamond is the minimum from the Nelder-Mead optimization with CFAST.

3.2.2. Surrogate FDS The importance of ventilation effects on the behavior of a compartment fire is a common theme found in the literature [2, 28]. The ability of CFAST to model flow is limited to openings in the compartment that can only add and remove mass from the two gas zones while the fire drives mass from the lower layer to the upper layer. A computational fluid dynamics model like Fire Dynamics Simulator (FDS) can be used with a coarse mesh to predict flow physics at the expense of increased computational time compared to CFAST. Surrogate models are often used to represent the output of an expensive model over a small range of input parameters. MCMC generally requires large numbers of samples and an FDS simulation with a coarse 20 cm mesh takes on the order of 30 min of computational time for 5 min of simulated time. Creating a Markov chain is a serial process so it may take a chain with FDS as the forward model on the order of years to converge. Generating a surrogate model can be done in parallel by running a set of simulations sampled from the input parameter space and fitting a model such as a Gaussian process regression (kriging), polynomial regression, or support vector machines [29, 30]. A naive approach would be to use linear interpolation to approximate values not directly simulated. If only a prediction of the point estimate of the most likely input parameters were desired, the grid search used to generate the surrogate model plus a local optimizer would be sufficient. However, construction of the surrogate model facilitates the generation of uncertainty information on that prediction.

Higdon et al. use Gaussian processes to represent temperature as a function of time for regulatory compliance experiments noting the computational time savings gained by fitting the surrogate model [31]. While Higdon et al. generate predic-

tions by sampling from a surrogate model constructed from Gaussian processes, we use a Gaussian process regression as a linear smoother by modeling the simulator (FDS) as a Gaussian process that is a function of an input vector $\boldsymbol{\theta}$ in $\mathbb{R}^4$, where

$$\boldsymbol{\theta} = (x_{fire}, y_{fire}, \dot{Q}, t_{peak}). \tag{8}$$

The standard linear regression model is

$$\boldsymbol{y} = \boldsymbol{f}(\mathbf{X}) + \epsilon_{\boldsymbol{n}}, \quad \boldsymbol{f}(\mathbf{X}) = \mathbf{X}\mathbf{w} \tag{9}$$

where the function $\boldsymbol{f}(\mathbf{X})$ is simply a product of some matrix $\mathbf{X}$ and weight vector ($\mathbf{w}$) that approximates observed data ($\boldsymbol{y}$). The noise, $\epsilon_{\boldsymbol{n}}$, is commonly assumed to be normally distributed with mean zero and covariance σ_n^2.

A Gaussian process is a non-parametric, multidimensional regression model whose target predictions are informed by input-output pairs of training data. We must construct a matrix of weights that depends on observed data and corresponding inputs. The model is trained on N realizations of the simulator with input vectors $\boldsymbol{\theta}_i$ and output vectors $\mathbf{y}_i$ for $i = 1, \ldots, N$. The output vectors are in $\mathbb{R}^P$ where P is the number of sensors at which the TEA are calculated. The input vectors are stored in the rows of matrix $\mathbf{X}$, and the output vectors are stored in the rows of matrix $\mathbf{Y}$. We define the model function, $\mathbf{f}(\mathbf{X})$, for a mean-zero Gaussian process prior as follows

$$\mathbf{f}(\mathbf{X}) \sim GP(\mathbf{0}, C(\mathbf{X}, \mathbf{X})) \tag{10}$$

that is parameterized by the covariance function C which depends on $\mathbf{X}$.

If we have acquired a set of data ($\mathbf{Y}$) at a set of input points ($\mathbf{X}$), we can make a prediction at a new point ($\boldsymbol{\theta}^*$) by sampling from the predictive distribution of $\mathbf{f}(\boldsymbol{\theta}^*)$ given $\mathbf{Y}$, $\mathbf{X}$, and $\boldsymbol{\theta}^*$ [32]. The predictive distribution is also a Gaussian process with a mean given by

$$\bar{\mathbf{f}}(\boldsymbol{\theta}^*) = C(\boldsymbol{\theta}^*, \mathbf{X})[C(\mathbf{X}, \mathbf{X}) + \sigma_{\mathbf{n}}^{\mathbf{2}}\mathbf{I}]^{-\mathbf{1}}\mathbf{Y} \tag{11}$$

and variance

$$v(\boldsymbol{\theta}^*) = C(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) - C(\boldsymbol{\theta}^*, \mathbf{X})[C(\mathbf{X}, \mathbf{X}) + \sigma_{\mathbf{n}}^{\mathbf{2}}\mathbf{I}]^{-\mathbf{1}}\mathbf{C}(\mathbf{X}, \boldsymbol{\theta}^*) \tag{12}$$

In this case we will simplify this process by treating the mean of the predictive distribution as the deterministic surrogate model prediction, effectively creating a linear smoother such that $\mathbf{f}(\boldsymbol{\theta}^*) = C(\boldsymbol{\theta}^*, \mathbf{X})\mathbf{W}$ where $\mathbf{W}$ is an N by P matrix of weights. Note that we could sample from the full predictive distribution in the Bayesian inversion scheme which would roll in the uncertainty associated with fitting the surrogate model, however this has been omitted to keep this process in line with using the deterministic model CFAST. Through preliminary investiga-

tions it was found that sampling from the full Gaussian process resulted in a change of less than one percent in the posterior means and 95% credible intervals when compared to deterministic surrogate model for all parameters except for the time to peak.

The underlying model that we are approximating is deterministic and there is no measurement error in the observations (ϵ_n) so the prediction simplifies to

$$\bar{\mathbf{f}}(\boldsymbol{\theta}^*) = C(\boldsymbol{\theta}^*, \mathbf{X})C(\mathbf{X}, \mathbf{X})^{-1}\mathbf{Y} \tag{13}$$

Therefore, the weight matrix is $C(\mathbf{X}, \mathbf{X})^{-1}\mathbf{Y}$ which can be calculated before MCMC sampling commences as it does not depend on the new point $\boldsymbol{\theta}^*$.

The covariance function (C) is chosen to be the Matern 5/2 covariance function with a regularization term given by

$$\begin{aligned} C(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) &= \tau^2\left(1 + \frac{\sqrt{5}d}{b} + \frac{5d^2}{3b^2}\right)exp\left(-\frac{\sqrt{5}d}{b}\right) + \delta(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2), \\ d &= ||\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2||_2 \end{aligned} \tag{14}$$

where $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are rows in X corresponding to input vectors with a distance (d) between them, and τ and b are hyper-parameters that must be fit to the data [32].

Cross-validation was used to fit hyper-parameters of the surrogate model wherein a training set of coarse-mesh FDS simulations was run at an evenly spaced grid of points in the input parameter space defined by the bounds of the uniform priors from Eqs. 3 and 4 to form matrix $\mathbf{X}$. The optimal hyper-parameters were determined for two HRR ranges, 50–800 kW and 850–1500 kW, due to memory limitations of having 30,000 sample points. The location of the fire was constrained to the main compartment and the time to peak was constrained between 0 s and 300 s. The grid spacing for the cross-validation training set was 50 kW for the HRR, 30 s for the time to peak, 0.43 m in the x direction, and 0.47 m in the y direction resulting in $N_{low} = 16{,}000$ and $N_{high} = 14{,}000$ for the low and high HRR ranges respectively. The TEA is collected at sensors on the walls and ceiling for a total of $P = 163$ sensors.

A smaller number of input parameter vectors ($N_{test} = 100$) was sampled randomly from within the input space for each HRR range and run in coarse-mesh FDS to serve as the test set for cross validation. The l^2 norm between the Gaussian process prediction from Eq. 11 evaluated at the test points ($\bar{\mathbf{f}}(\mathbf{X}_{test})$) and the coarse-mesh FDS output ($\mathbf{Y}_{test}$) was minimized with the Nelder-Mead method to determine the best set of hyper-parameters [27]. The optimal hyper-parameter values can be found in Table 1.

$$min\ \left|\left|\bar{\mathbf{f}}(\mathbf{X}_{test}) - \mathbf{Y}_{test}\right|\right|_2 \tag{15}$$

Error contours were generated in the same manner as with the CFAST model to gain insight in to the expected performance of the Bayesian inversion framework

Table 1
Optimal Hyper-Parameters for the Surrogate FDS Model

HRR range	b	τ
50–800 kW	0.16528	3.49135
850–1500 kW	0.3503	91.37655

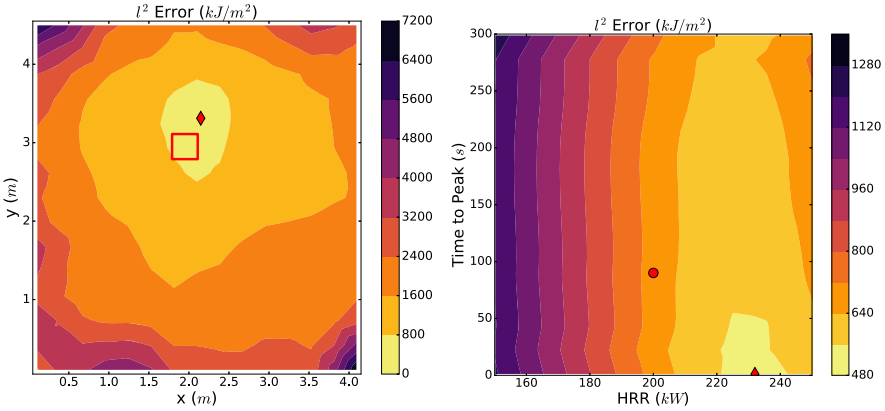

Figure 5. Surrogate FDS fire location error surface (a) and HRR and time to peak HRR error surface (b) for a 200 kW fire. The outlined box in (a) represents the boundary of the true fire from the fine mesh FDS simulation. The circle in (b) is the true HRR and time to peak. In both figures, the diamond is the minimum from the Nelder-Mead optimization with surrogate FDS.

(Fig. 5). In the case of the 200 kW fire, the minimum error is located to the top right of the true fire origin while the predicted location of the 1000 kW fire is close to the boundary of the true fire. This is a small improvement over CFAST, however the bias in the x direction is now to the right of the true origin indicating that the tilt of the plume is being over-predicted in this direction. We note that the bias errors are on the order of one to two grid cells (20 cm). We therefore expect the prediction accuracy of the Bayesian framework with the surrogate FDS to be on the same order as CFAST for the fire location.

Figure 5b shows the error contour for fire size and time to peak. In both peak HRR cases, the HRR with the lowest error (represented by a circle marker) is higher than the true value (diamond marker) due to numerical diffusion effects from the coarse mesh size, i.e. the hot gases from the fire incorrectly diffuse quicker than the fine-mesh FDS simulation. Therefore, a higher HRR is required to match the TEA imparted to the ceiling. For the 200 kW case, there is an indi-

cation that the time to peak occurs within the first half of the fire's burn time but the minimum is not at 90 s, and instead it is at 1 second. This is a small improvement over CFAST, but will likely not be sufficient for resolving the time dependent HRR profile of the fire. For the 1000 kW case, there appear to be several local minima which could be problematic for the inversion process resulting in chains that could become stuck in local minima.

Based on this exploration of the error surfaces, we expect that inverting for the location of the fire and peak HRR with the Bayesian framework will be feasible but biased with both CFAST and FDS. There will be high uncertainty in the time to peak as the total energy released by the fire is not a function of the time to peak and there is not a large enough difference in the TEA imparted to the surfaces between different peak times to overcome the model bias error.

4. Inversion Results

The Bayesian inversion process outlined in Sect. 2 was conducted on the following forward models and HRR profile combinations: a steady fire with CFAST as the forward model, a triangular fire with CFAST as the forward model, and a triangular fire with the surrogate FDS formulation as the forward model. In each combination, 200 kW and 1000 kW fires were examined. Each of these case studies has the same ventilation conditions (north window closed), fire location (1.95 m, 2.95 m from the northwest corner), time to peak HRR (90 s), and test duration (300 s).

The steady case is a four parameter inversion problem for the fire location (x_{fire},y_{fire}), constant HRR ($\dot{Q}$), and likelihood standard deviation (σ) while the triangular fire case has five parameters: substituting peak HRR ($\dot{Q}_{peak}$) and time to peak HRR (t_{peak}) for the constant HRR. In both cases, the output quantity used for evaluation of the likelihood function is the TEA at each sensor. The prior bounds for the location parameters are set such that the edge of the fire is constrained by the walls. For the room with dimensions of 4.2 m by 4.58 m and a fire diameter of 0.32 m, the location priors are

$$x_{fire} \sim U(0.16, 4.04), \quad y_{fire} \sim U(0.16, 4.42) \tag{16}$$

The prior for t_{peak} is a uniform distribution with bounds of 1 to 299 s indicating that there is no prior knowledge of when the peak HRR occurs. The upper bound of the likelihood standard deviation prior (σ_{high}) is set at 100, reflecting little knowledge in the misfit of the TEA data. Note that all TEA data were scaled by dividing by the average TEA from the respective synthetic simulation. Future work will consider different scaling scenarios for the different surfaces. In all cases where CFAST was the forward model, the HRR prior was uniform 50 to 1500 kW. When the surrogate FDS forward model was used, the prior bounds depended on the peak HRR being inverted for as two surrogate models were created. For the 200 kW case, the bounds were 50 to 800 kW while the bounds for the 1000 kW case were set at 850 to 1500 kW.

Three variables relating to the sensors were explored for each fire type and forward model combination: sensor density, noise level, and sensor configuration. The sensor density corresponds to the spacing of the sensors where the 0.5 m sensor grid uses all 163 sensors spaced on a regular grid and the 1.0 m sensor grid uses only 40 sensors. The noise level was informed by work done by Kokel et al. on directional flame thermometers (DFT) as these devices are used in the real experimental structure for measuring heat flux [33]. Kokel et al. found a maximum of 25% discrepancy between the DFT measurements and a Schmidt-Boelter heat flux gage with an average discrepancy of around 4% to 6% of the measured heat flux. Two noise levels were applied to the synthetic data that bound the extremes of the error measured by Kokel et al. where the low bound was set to no added noise and the high bound was set to normally distributed noise with a standard deviation of 25% of the measured value. Sensor configuration relates to the sensor locations used in the inversion problem. One can imagine a real fire scenario where sections of the walls or ceiling are completely destroyed by the fire and no data can be collected in those locations. Likewise, we experimented with using only the ceiling sensors, only the wall sensors, and all of the sensors and compared the results of the inversion process.

The results of inverting for the unknown parameters under each sensor configuration were compared to the true value by plotting the posterior mean as a dot and the 95% Bayesian credible interval as a set of whiskers on the dot representing the posterior mean. A 95% Bayesian credible interval is defined as the region that contains the true value of a random parameter with a 95% probability. This interval is calculated as the smallest interval that contains 95% of the posterior samples [34]. The sample chain is sorted from lowest to highest and the interval is taken as the two points in the sorted chain with the smallest distance between them that contain $0.95N_{samp}$ samples, where N_{samp} is the total number of samples in the posterior chain after burn-in. This interval is also referred to as the highest probability density (HPD) interval and is not necessarily equal-tailed.

4.1. CFAST Inversion Results

Figure 6 shows the mean and 95% Bayesian credible interval for the steady HRR, x location, and y location of the fire for each sensor configuration and noise level. The true value of each parameter used for generating the synthetic data is represented by the vertical dashed line. Predictably, the bounds of the 95% Bayesian credible interval widen with increasing noise and sensor sparsity, and we note that this occurs for all other cases considered in this work.

The fire location predictions correspond to the observations made in the exploration of the error space from Sect. 3.2.1, where we noted that the tilt of the fire plume is not modeled in CFAST whereas in reality the effects of ventilation cause the plume to lean and pulse. The MCMC sampler narrows in on the place where the plume impinges on the ceiling when the ceiling sensors are used, however using only the wall sensors performs better in this respect as the predicted fire locations are closer to the true location. There remains a bias toward the wall

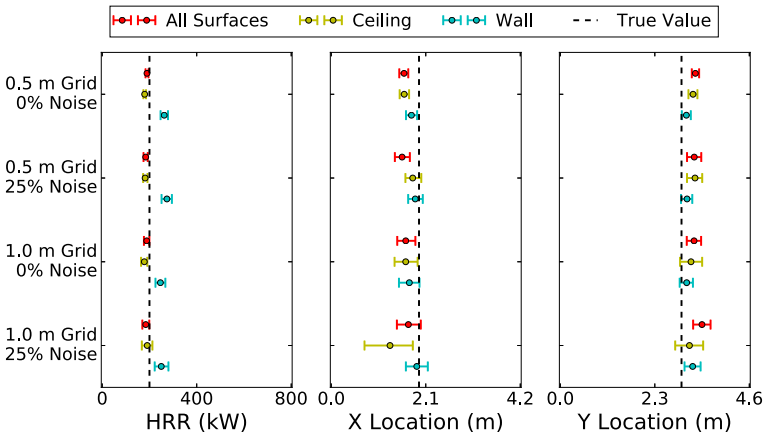

Figure 6. CFAST 200 kW steady fire summary of the posterior mean and 95% Bayesian credible intervals.

closest to the fire as CFAST uses a point source model to calculate radiation from the center of the fire while the true fire has an approximately conical shape.

The HRRs predicted by the cases where all of the sensors were used are a blend of the ceiling and wall cases. This trend becomes more apparent in some of the case studies with larger fires. There may be a weighting of the wall and ceiling TEA values that will more accurately predict the HRR by balancing the known biases of CFAST as discussed in Sect. 3.2.1. Inclusion of bias errors would necessarily expand the uncertainty around these inverted values. For some applications, work on model inadequacy has been investigated by Kennedy and O'Hagan as well as Hidgon et al. and this could be incorporated in to future analysis of the fire forensics problem [31, 35, 36].

A similar trend in the fire location bias toward the point where the most heating occurred was observed in the inversion results for a steady 1000 kW fire (Fig. 7). We note that the plume impingement point of the synthetic data in the 1000 kW fire is shifted lower in the y direction than in the 200 kW fire. There is overall greater uncertainty in the location parameters for the 1000 kW case than in steady 200 kW fire case. There is a hard upper limit on the plume temperature rise in CFAST and the 1000 kW fires have reached this limit. This temperature limit results in under-prediction of the heat flux to the ceiling, which negatively affects the ability of the inversion algorithm to find the correct maximum HRR using CFAST as the forward model as a higher HRR would be required to overcome this hard limit.

Figure 8 shows the posterior results of inverting for a 200 kW triangular fire with CFAST as the forward model. The same trends are observed as in the steady cases for the fire size and location. As expected based on the error exploration from Sect. 3.2.1, there seems to be little indication of where the peak in the HRR ramp lies as the chains either get stuck on one end of the "time to peak" prior or sample from almost the full extent of the prior bounds. However, we note that

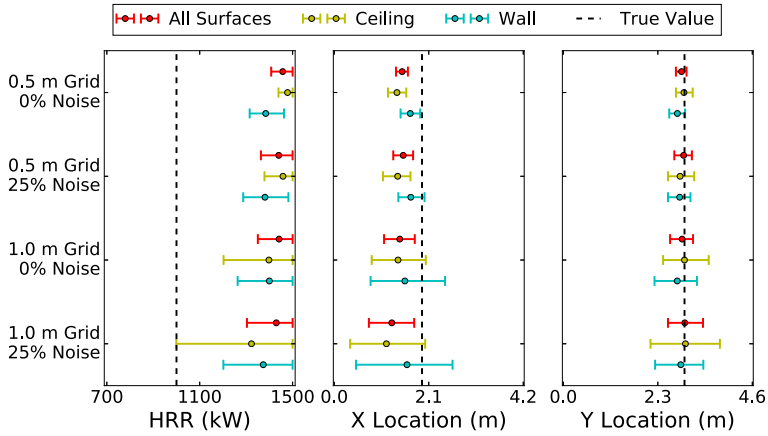

Figure 7. CFAST 1000 kW steady fire summary of the posterior mean and 95% Bayesian credible intervals.

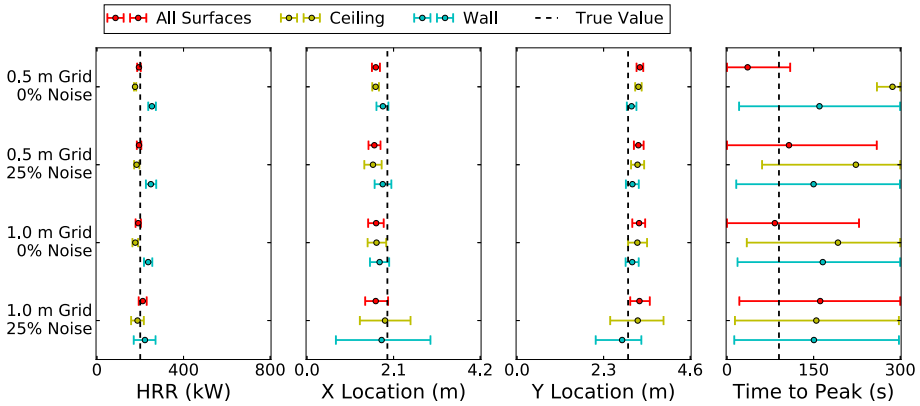

Figure 8. CFAST 200 kW triangular fire summary of the posterior mean and 95% Bayesian credible intervals.

even without predicting the time to peak, in most cases the peak HRR is still correctly predicted with at least one sensor configuration indicating that the total energy released by the fire is also correctly predicted. This finding is promising as we know that it is possible to invert for the HRR using one parameter that scales with the total energy released. In future versions of this inversion framework, HRR curves could be generated as single draws from a family of curves corresponding to the given burned item instead of drawing multiple parameters that describe the fire.

In the 1000 kW triangular HRR fire inversion, the posterior statistics for the location parameters show the same bias as the previous CFAST case studies (Fig. 9). Interestingly, using only the ceiling sensors improves the predictions of

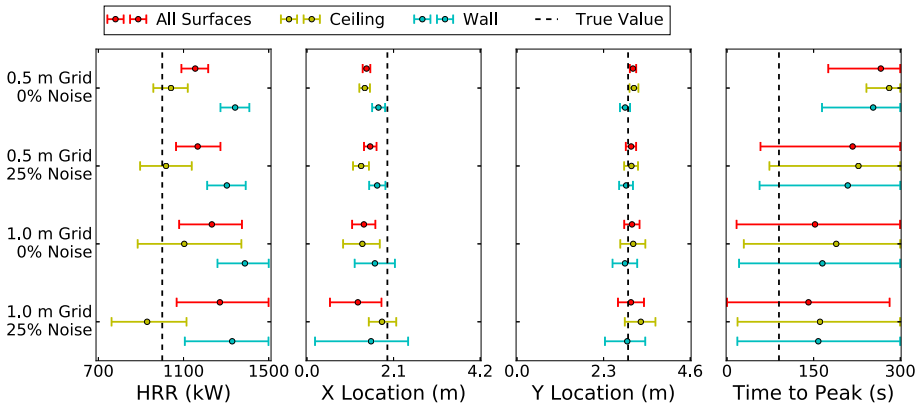

Figure 9. CFAST 1000 kW triangular fire summary of the posterior mean and 95% Bayesian credible intervals.

the peak HRR, whereas using only the wall sensors results in over-prediction of the peak HRR due to the consistent under-prediction of heat transfer to the walls in CFAST. Time to peak predictions are either skewed towards later times or completely unidentifiable. The inadequacy in the ability of using CFAST in the inversion framework to predict time to peak leads to the examination of the coarse FDS surrogate model with the hope that better prediction of the fluid flow physics will increase the identifiability of the time to peak HRR.

4.2. Surrogate FDS Inversion Results

Based on the error space explorations in Sect. 3.2.2 we expect there to be a small bias in the location predictions, over-prediction on the order of 10% in the peak HRR, and improved time to peak predictions compared to CFAST for the 200 kW peak HRR case. There is also the possibility of encountering local minima in the 1000 kW peak HRR case. Upon examining the Bayesian inversion results, the majority of these expectations held true with a few notable exceptions.

In Fig. 10, the uncertainty in the location predictions for the 25% noise and 1 m sensor spacing cases increased dramatically when only the wall or ceiling sensors were used. There is likely a critical sensor density needed to overcome high noise levels and recover the fire behavior. The time to peak predictions are greatly improved compared to CFAST for the 0.5 m sensor grid, no noise case but higher levels of noise and sensor sparsity show that the inversion process only has an indication that the peak is more likely to occur before the half time mark at 150 s.

In the 1000 kW peak HRR cases there are similar trends in the location parameter mean values compared to the 200 kW peak HRR case, however there is no longer large uncertainty in the 25% noise and lower sensor density predictions possibly due to a higher signal to noise ratio from the 1000 kW fire that overcomes the decreased information density (Fig. 11). The wall sensors perform the best at predicting the peak HRR which differs from the CFAST inversion results. This could be the result of the ceiling sensors seeing lower heat fluxes as the

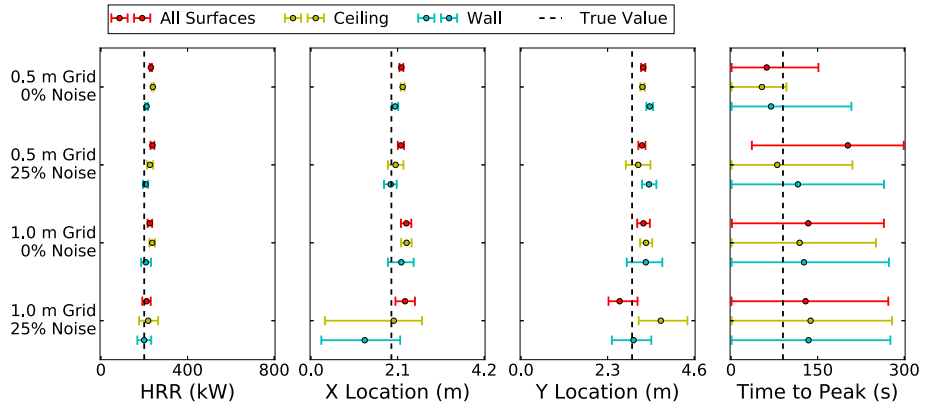

Figure 10. Surrogate FDS 200 kW triangular fire summary of the posterior mean and 95% Bayesian credible intervals.

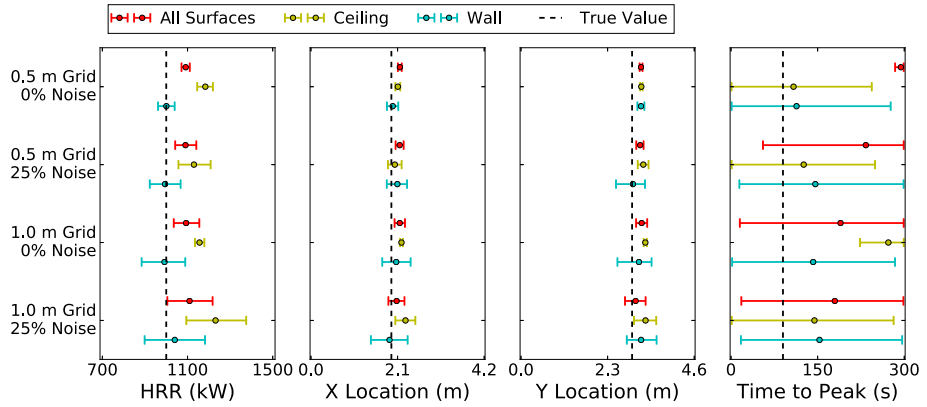

Figure 11. Surrogate FDS 1000 kW triangular fire summary of the posterior mean and 95% Bayesian credible intervals.

coarse grid smooths out the gas temperature thus lowering the temperature near the ceiling. Therefore, when the ceiling sensors are used in the likelihood calculation a HRR higher than the true value is needed to match the TEA seen by the ceiling sensors. As expected, the time to peak is unidentifiable as the chains either become trapped in local minima as seen in Fig. 5 or traverse the extent of the prior distribution.

5. Conclusions

This work has shown that by using Bayesian inversion with surrogate FDS and CFAST as forward models and TEA data we can consistently locate a fire in a compartment. When inverting for the peak HRR, CFAST breaks down for HRRs

approaching 1000 kW for this geometry whereas the surrogate FDS model performs well when using the wall sensors. With the exception of the 1000 kW CFAST cases, the inversion framework was able to locate the fire within 0.76 m of the true origin (i.e. the posterior mean was within 0.76 m of the true origin). This study pushed the limits of using time-integrated data to determine temporal parameters by seeking to invert for the time to peak HRR of a triangular fire. This is ultimately a difficult task and while there are differences in the data between different peak times, they are small compared to the magnitude of the measurements and any noise can easily obscure all differences between temporal variations in peak times. Despite this shortcoming, the inversion framework was able to reasonably predict the total energy released by the fire for all cases except for the 1000 kW fires with CFAST as the forward model. FDS performed better as the forward model when inverting for the maximum HRR where the posterior mean of the best configuration was within 0.05% of the true value whereas the best CFAST configuration of the inversion framework predicted the maximum HRR within 2.77% of the true value.

At this point, using FDS for new geometries and scenarios would not be practical as it requires a large amount of computational resources compared to CFAST. However, it has been shown that zone models are inadequate for identification of fire properties for moderately large fires. On the order of 10,000 forward solves are required to get to a converged posterior distribution for the quantities of interest using a conventional Markov chain Monte Carlo method, necessitating the use of a surrogate model for FDS. To get the MAP, tens to low hundreds of calls are required. While computational speed continues to increase, we learned that there is a need to develop fast surrogate models for the forward model to exercise the methodology.

In the future, the inversion framework could be tailored to switch between the FDS and CFAST forward models when the HRR passes a specified threshold where CFAST become unreliable. To fully realize this inversion framework and take the analysis from data collection to hypothesis testing, there is a need for validation experiments and a data source that can be measured at a fire scene. As it proved to be difficult to invert for the time variation of the HRR of an item given no knowledge of the item, future endeavors will seek to incorporate prior knowledge from fire investigators in the form of a classification system for items that draws from a family of likely HRR curves.

Acknowledgements

This work was supported by the U.S. National Science Foundation under grant number 1707090.

References

1. De Haan JD (2008) Kirk's fire investigation. Pearson/Brady

2. Hicks W, Gorbett GE, Kennedy PM, Hopkins RL, Abney WM (2006) Advanced fire pattern research project: single fuel package fire pattern study. In: 2006 International symposium on fire investigation proceedings, National Association of Fire Investigators, Sarasota, FL
3. Hicks W, Gorbett GE, Hopkins MC, Kennedy PM, Hopkins RL, Thurman JT (2008) Full-scale single fuel package fire pattern study. In: 2008 International symposium on fire investigation proceedings, National Association of Fire Investigators, Sarasota, FL
4. Madrzykowski D, Fleischmann C (2012) Fire pattern repeatability: a study in uncertainty. J Test Eval 40(1):1
5. Yeager RW (1986) Uncertainty analysis of energy release rate measurement for room fires. J Fire Sci 4(4):276–296
6. Madrzykowski D (2017) Repeatability of pre-flashover fire patterns on gypsum wallboard. Ph.D thesis, University of Canterbury
7. Mealy CL, Wolfe AJ, Gottuk DT (2013) Forensic analysis of ignitable liquid fuel fires in buildings. Hughes Associates, Incorporated
8. Farrell MG (2012) Daubert v. merrell dow pharmaceuticals Inc.: epistemilogy and legal process. Cardozo Law Rev 15(6):2183–2218
9. National Research Council (2009) Strengthening forensic science in the united states: a path forward. Technical report, National Research Council
10. Lentini J, Mowrer F, Pawliszyn J, Almirall J, Arkes H (2017) Forensic science assessments: a quality and gap analysis. Technical report
11. Nordgaard A, Hedell R, Ansell R (2011) Assessment of forensic findings when alternative explanations have different likelihoods—"blame-the-brother"-syndrome. Sci Justice 52:226–236
12. Wang J, Zabaras N (2004) A bayesian inference approach to the inverse heat conduction problem. Int J Heat Mass Transf 47:3927–3941
13. Biedermann A, Taroni F, Delemont O, Semadeni C, Davison AC (2005) The evaluation of evidence in the forensic investigation of fire incidents (part I): an approach using bayesian networks. Forensic Sci Int 147(1):49–57
14. Biedermann A, Taroni F, Delemont O, Semadeni C, Davison AC (2005) The evaluation of evidence in the forensic investigation of fire incidents. part II. Practical examples of the use of bayesian networks. Forensic Sci Int 147(1):59–69
15. Guo S, Yang R, Zhang H, Zhang X (2010) New inverse model for detecting fire-source location and intensity. J Thermophys Heat Transf 24(4):745
16. Overholt KJ, Ezekoye OA (2012) Characterizing heat release rates using an inverse fire modeling technique. Fire Technol 48:893–909
17. Overholt KJ (2013) Forward and inverse modeling of fire physics towards fire scene reconstructions. PhD Dissertation, University of Texas at Austin. http://hdl.handle.net/2152/21971
18. Overholt KJ, Ezekoye OA (2014) Quantitative testing of fire scenario hypotheses: a bayesian inference approach. Fire Technol 51(2):335–367
19. Patil A, Huard D, Fonnesbeck CJ (2010) PyMC: Bayesian stochastic modelling in Python. J Stat Softw 35(4):1
20. Bolstad WM (2010) Understanding computational Bayesian statistics, vol 644. John Wiley & Sons
21. Powell MJD (1964) An efficient method for finding the minimum of a function of several variables without calculating derivatives. Comput J 7(2):155–162
22. Roberts GO, Rosenthal JS (2009) Examples of adaptive MCMC. J Comput Graph Stat 18(2):349–367

23. Geweke J (1991) Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments, vol 196. Federal Reserve Bank of Minneapolis, Research Department, Minneapolis, MN
24. McGrattan K, Hostikka S, McDermott R, Floyd J, Weinschenk C, Overholt K (2013) Fire dynamics simulator, technical reference guide, vol 2: verification. NIST Special Publication, 1018-2
25. McGrattan K, Hostikka S, McDermott R, Floyd J, Weinschenk C, Overholt K (2013) Fire dynamics simulator, technical reference guide, vol 3: validation. NIST Special Publication, 1018-3
26. Babrauskas V (2002) SFPE handbook of fire protection engineering, chapter heat release rates, 3rd edn. National Fire Protection Association, Quincy
27. Nelder JA, Mead R (1965) A simplex method for function minimization. Comput J 7(4):308–313
28. Gorbett GE, Meacham BJ, Wood CB, Dembsey NA (2015) Use of damage in fire investigation: a review of fire patterns analysis, research and future direction. Fire Sci Rev 4(1):4
29. O'Hagan A (1978) Curve fitting and optimal design for prediction. J R Stat Soc: Ser B (Methodological) 40(1):1–24
30. Sacks J, Welch WJ, Mitchell TJ, Wynn HP (1989) Design and analysis of computer experiments. Stat Sci 4(4):409–423
31. Higdon D, Nakhleh C, Gattiker J, Williams B (2008) A bayesian calibration approach to the thermal problem. Comput Methods Appl Mech Eng 197(29):2431–2441
32. Rasmussen CE, Williams CKI (2006) Gaussian processes for machine learning. MIT Press, Cambridge
33. Kokel P, Weinschenk C, Ezekoye OA (2010) Evaluation of directional flame thermometer for real-time inversion of heat flux. In: 2010 14th International heat transfer conference. American Society of Mechanical Engineers, pp 83–92
34. Chen M-H, Shao Q-M (1999) Monte Carlo estimation of Bayesian credible and HPD intervals. J Comput Graph Stat 8(1):69–92
35. Kennedy MC, O'Hagan A (2001) Bayesian calibration of computer models. J R Stat Soc: Ser B (Statistical Methodology) 63(3):425–464
36. Higdon D, Gattiker J, Williams B, Rightley M (2008) Computer model calibration using high-dimensional output. J Am Stat Assoc 103(482):570–583