

COMBINED DATA AND DEEP LEARNING MODEL UNCERTAINTIES: AN APPLICATION TO THE MEASUREMENT OF SOLID FUEL REGRESSION RATE

Georgios Georgalis,^{1,*} Kolos Retfalvi,² Paul E. Desjardin,² & Abani Patra^{1,3}

¹Data Intensive Studies Center, Tufts University, Medford, Massachusetts 02155, USA

²Department of Mechanical and Aerospace Engineering, University at Buffalo, The State University of New York, Buffalo, New York 14260, USA

³Department of Mathematics, Tufts University, Medford, Massachusetts 02155, USA

*Address all correspondence to: Georgios Georgalis, Data Intensive Studies Center, Tufts University, Medford, Massachusetts, 02155, USA, E-mail: georgios.georgalis@tufts.edu

Original Manuscript Submitted: 10/20/2022; Final Draft Received: 03/12/2023

In complex physical process characterization, such as the measurement of the regression rate for solid hybrid rocket fuels, where both the observation data and the model used have uncertainties originating from multiple sources, combining these in a systematic way for quantities of interest (QoI) remains a challenge. In this paper, we present a forward propagation uncertainty quantification (UQ) process to produce a probabilistic distribution for the observed regression rate \dot{r} . We characterized two input data uncertainty sources from the experiment (the distortion from the camera U_c and the non-zero-angle fuel placement U_γ), the prediction and model form uncertainty from the deep neural network (U_m), as well as the variability from the manually segmented images used for training it (U_s). We conducted seven case studies on combinations of these uncertainty sources with the model form uncertainty. The main contribution of this paper is the investigation and inclusion of the experimental image data uncertainties involved, and how to include them in a workflow when the QoI is the result of multiple sequential processes.

KEY WORDS: *uncertainty characterization, deep learning model uncertainty, image data uncertainty, combustion experiments, regression rate density estimation*

1. INTRODUCTION

When analyzing a physical system, we use the “input data” (knowledge we have) to run the “model” and evaluate its outputs (knowledge we want to obtain). All of the components involved in the process carry their own assumptions, limitations, and uncertainties. When the outputs are part of a larger computational framework or are used in decision making, reporting the associated uncertainty is required, because the modeling process and the model evaluations are approximations of the underlying true physical process. Probability models are common representations of such uncertainty in both the inputs and outputs. Forward-propagation of the uncertainty of input data and estimation of the probability distribution of the simulation output via sampling-based Monte Carlo methods or functional approximations is abundant in the literature with recent applications in many fields (e.g., see [1–3]). Similarly, observation data-driven calibration procedures using Markov chain Monte Carlo or variants are common [4].

However, in cases where the “input data” are results of a complex physical process that is inherently variable (e.g., from a complex experiment), and the “modeling” includes multiple sequential models and dependencies, propagating and combining all uncertainty sources in a systematic way for a required quantity of interest (QoI) remains a challenge [5]. This is especially true when using operations like image segmentation models [6] for identification of surfaces and interfaces which implicitly and explicitly use modeling and calibration. In particular, we focus here on the characterization of the “input data” uncertainty when it is acquired in a complex experiment and used within a model-based interpretation that transforms the raw observation into a QoI.

In hybrid rocket motor combustion, one QoI is the rate at which the fuel surface recedes during the burn, defined as the regression rate \dot{r} . The regression rate has a direct impact on the geometrical design of the rocket motor and its performance ([7,8]). Due to the high cost of building a full-scale hybrid motor, it is common practice to estimate the regression rate *a priori* via a smaller-scale experiment of a slab burner [9]. Tracking the surface of the fuel from experimental images during the burn with time data is one way to estimate \dot{r} [10]. In this case, “inputs” are all of the assumptions and limitations of the experiment, the equipment, and any processing to get the fuel surface images. “Modeling” is the process of translating the images to the regression rate, which is practically a sequence of multiple models: a deep learning convolutional network model to segment the fuel masks of an entire experiment (e.g., Monte Carlo dropout (MCD) U-net from [11]), a process to detect the boundaries of the segmented fuel masks, and a process to estimate \dot{r} from the changing boundaries with time data (Fig. 1).

The regression rate and its measure of uncertainty are necessary to either validate multiscale computational models (e.g., simulations of turbulent combustion chemistry coupled with flow dynamics [12]), or for decision making on the rocket motor design and performance. In these coupled systems, there is the additional requirement that the uncertainty information of the QoI be expressed in a simple way (e.g., as a probability distribution) because simulations can be sequential: the distribution of the regression rate represented by an ensemble is an input to another simulation (e.g., see [13] for an application of sequential uncertainty quantification in combustion systems).

Reports in the literature has focused on separating the different types of uncertainty within a model as aleatoric (AU) or epistemic uncertainty (EU) and how to measure them in physical models such as hypersonic flows [14], in machine learning [15], in Bayesian convolutional networks [16], and in safety decision making [17,18] among others. In the context of a complex “input data” and “modeling” process as in the regression rate case (Fig. 1), these two distinct uncertainty definitions are often challenging to separate or measure. For example, the inherent camera structure and how it represents an image captured in the experimental environment and whether the researcher misplaced the fuel specimen by a few degrees include some inherent randomness, but we have the capability to estimate *some* parts of that experimental randomness by analyzing the setup and the optical errors. Therefore, in this work, we do not classify uncertainty as aleatoric or epistemic, but rather characterize and quantify the individual uncertainty measures based on their source (i.e., experiment or modeling).

In this paper, we present an uncertainty quantification (UQ) process to produce a probabilistic distribution for the regression rate \dot{r} from the input experimental images, by combining the input data uncertainty with inference model uncertainty. We characterize the input data uncertainties from the experiment $U_{data} = \{U_c, U_\gamma\}$, the model form uncertainty from the MCD U-net (U_m), and the variance in the model prediction from the variability of manually tracing the masks used to train the MCD U-net (U_s). The process is shown in Fig. 1(a) we introduce the input data uncertainty U_{data} to create sequenced ensembles of experimental images, Fig. 1(b) pass the ensembles through the U-net to get an estimate of the mean predicted masks μ_{mask} and corresponding uncertainty maps U_{mask} which carry both the data and model uncertainties, and lastly Fig. 1(c) process each sequenced ensemble of predicted fuel masks and uncertainty maps to get an estimate for the regression rate \dot{r} and its bounds.

The sources of input data uncertainty are most relevant for this problem because optical distortion from the camera (U_c) or non-zero-angle placement of the fuel (U_γ) may misrepresent the fuel boundaries in the images that are tracked to accurately measure the regression rate. The model form uncertainty (U_m) is computed as the entropy of the probability prediction vector of the MCD U-net. The possible variation in the manually segmented masks used for training the MCD U-net is added as a prediction variance U_s directly to the final resulting uncertainty map U_{mask} .




FIG. 1: Overview of the workflow and uncertainties at each step to estimate the fuel regression rate \dot{r} from a slab burner experiment. Hybrid rocket image (left) is NASA’s Peregrine rocket motor.

The paper is organized as follows. Section 2 includes details about the experimental setup and the characterization for each of the studied uncertainty measures. Section 3 presents a series of cases that investigate how the different types of uncertainties impact the overall flow of uncertainty measure in the forward-propagation process (i.e., how the input data uncertainties individually and together combine with the model form uncertainty and the manual segmentation variance to form the output uncertainty map used in the measuring \dot{r}). Section 4 is an application of the entire workflow shown in Fig. 1, and includes the resulting distribution for \dot{r} via forward-propagation after including all the uncertainty sources in sequence. Section 5 summarizes the conclusions of the paper with directions on future work.

2. FORMULATION OF UNCERTAINTY SOURCES AND REGRESSION RATE MEASUREMENT

In this section, we describe the experimental setup and equipment used, the characterization of the input data uncertainties $U_{data} = \{U_c, U_\gamma\}$, and how they are introduced to the experimental images. U_c corresponds to error from optical distortion that is inherent to the camera placement in respect to the fuel specimen. U_γ corresponds to the non-zero-angle fuel placement error, which may result from the researcher not perfectly placing the fuel orthogonal to the camera axis, thus having the camera misrepresent the distance between the fuel boundaries correctly. We also show the formulation for the model form uncertainty U_m and the manual segmentation variance U_s , as well as how the model outputs μ_{mask}, U_{mask} are used to measure the regression rate \dot{r} and its bounds.

2.1 Experimental Setup and Its Uncertainties

The experimental setup follows closely the one developed by Dunn et al. [19] and can be seen in Fig. 2. The fuel specimen is placed in a chamber consisting of two stainless steel plates on the top and bottom and high temperature borosilicate glasses on the side of the experiment for optical access. The chamber is 15.24 cm long, 2.54 cm tall, and 2.54 cm wide, with an oxidizer inlet pipe 1.83 cm in length and 2.54 cm in diameter. Based on the entrance length, the inlet flow is assumed to be fully developed as it enters the chamber. We used lab-grade paraffin wax from the Carolina Biological Company as the fuel during the experiments and it was cast in a stainless steel mold. The temperature of the mold was monitored during the solidification process to avoid impurities in the samples. The average dimensions for the fuel specimens used in the experiments were 9.4 mm in width, 80.7 mm in length, and 11.2 mm in height. Each sample had a 45° slant in the front to guide the flow. The oxidizer used for the experiment was 100% gaseous oxygen which was regulated with solenoids and measured with an Omega FMA 1744a mass flow meter. The measurement range of the flowmeter was 5–500 SLMs with $\pm 1.5\%$ accuracy.

To obtain the image dataset used in this paper, we conducted two experiments with measured oxidizer mass fluxes of $G_1 = 6.96 \text{ kg/m}^2\text{s}$ and $G_2 = 10.96 \text{ kg/m}^2\text{s}$, respectively. During the automated experimental sequence the oxidizer was first introduced and the slab ignited using a Bosch diesel glow plug which was lowered to make contact with the slab and later retracted using a stepper motor. The experiment continued until the flame was extinguished. The experiment was recorded with a Chronos 2.1 high speed camera during the experimental runs, with a frame rate of 1000 frames per second. The lens used was a NIKON AF Nikkor 50 mm lens with apertures ranging from 1.8 to 8 with an exposure of 1 μs . The full burn of the specimen took 8–12 s. Some example images from the experiment with oxidizer flux $G_2 = 10.96 \text{ kg/m}^2\text{s}$ at different times can be seen in Fig. 3.




FIG. 2: Slab burner test chamber with paraffin wax sample




FIG. 3: Example images from the experiment with oxidizer flux $G_2 = 10.96 \text{ kg/m}^2\text{s}$ at early (a), middle (b), and later (c) stages

2.1.1 Optical Distortion Error U_c

The image data from the experiment are not exactly accurate representations of the real phenomena during the burn. One reason is that despite best efforts to set the camera correctly for best image quality, the fuel surface moves with respect to the camera during the experiment, adding a possibility of optical distortions. To account for the possible camera distortions as part of the input data uncertainty, we followed a calibration process by using checkerboard patterns. Assuming the pinhole model for the camera, i.e., the real points of the captured space are imaged by rays that pass through a single origin of the camera lens [20], we estimated the intrinsic and extrinsic parameters of the camera and the distortion coefficients as outlined in [21].

Given the global coordinate system (x, y, z) and the camera coordinate system (i, j, k) , a point in the real world is represented by the position vector $\mathbf{p} = [p_x, p_y, p_z]^T$. The same point in the camera coordinate system is represented by the position vector $\mathbf{q} = [q_i, q_j, q_k]^T = R^T(\mathbf{p} - \boldsymbol{\tau})$, where R is the rotation matrix and $\boldsymbol{\tau}$ the translation vector (extrinsic parameters). The point represented in the camera coordinates as \mathbf{q} is then mapped into the image plane

using the intrinsic parameters matrix K and represented as $\mathbf{r} = K\mathbf{q} = \begin{bmatrix} f_x & 0 & 0 \\ s & f_y & 0 \\ c_x & c_y & 1 \end{bmatrix} [q_i, q_j, q_k]^T$, where f_x, f_y are

the focal lengths, s is the skew coefficient, and c_x, c_y are the coordinates of the optical center. The rotation matrix R , the translation vector $\boldsymbol{\tau}$, and the intrinsic matrix K are all estimated during the camera calibration process. In addition, we also estimate the camera distortion coefficients k_1, k_2 that characterize radial distortion in the image coordinates [22]: $r_{x,distorted} = (r_x/f_x)(1 + k_1d^2 + k_2d^4)$ and $r_{y,distorted} = (r_y/f_y)(1 + k_1d^2 + k_2d^4)$, where $d = (r_x/f_x)^2 + (r_y/f_y)^2$.

To calibrate the camera, we took ten images of a square checkerboard pattern with a side of 2 cm under different angles (Fig. 4), and processed them using MATLAB's camera calibrator app [23]. The resulting estimated parameters and camera resolution are shown in Table 1. The rotation matrices and translation vectors are not shown because they are different for each calibration image. During the calibration process, the true locations of the checkerboard pattern are detected and compared with their locations on the reconstructed images from the camera model. The reconstruction error is then $e_{i,j} = s_{i,j} - \hat{s}_{i,j}$, where $s_{i,j}$ are the true detected points on the calibration patterns and $\hat{s}_{i,j}$ the reconstructed points. To arrive at a measure of distortion uncertainty from the camera that can be used to introduce uncertainty to the slab burner images, we count how many of the calibration points have a total error that is greater than 1 pixel in distance, for each of the calibration images. The threshold is set at 1 pixel because when we represent the experimental images as tensors for the deep learning model, an error greater than 1 pixel means the image point is distorted enough to not populate that field in the tensor anymore. The total number of distorted pixels from each calibration image is expressed as a percentage of the total pixels in the image. The result of this process is shown in Fig. 5: the expected maximum number of distorted points with an error greater than 1 pixel corresponds to 0.7% of the image, and the minimum number to 0.133% of the image. Therefore, based on the calibration results, we expect that an experimental image taken with our camera has a number of distorted pixels $U_c \sim \mathcal{U}(0.133, 0.7)\%$, expressed as a percentage of the image. Since there are only two options for the pixels, distorted or not, for a given experimental image X_k of resolution 512×64 , there is a corresponding distortion map $D_{map,k}$. The distortion map follows the binomial distribution $D_{map,k} \sim B(n = 64 \times 512, p = U_c)$. The pixels in X_k that have a success trial on the distortion map $D_{map,k}$, are distorted, and therefore have their intensity reduced to zero when we introduce distortion uncertainty U_c to the original image (Algorithm 1).




FIG. 4: Relative positioning between the camera and ten 2-cm-side checkboard calibration images

TABLE 1: Camera resolution and parameter estimation from calibration

Camera Parameter	
Resolution	1920×1080 [pix]
f_x	4386.76 ± 27.41 [pix]
f_y	4374.34 ± 26.68 [pix]
s	0.0 (perpendicular axes)
c_x	971.65 ± 1.25 [pix]
c_y	503.64 ± 1.22 [pix]
k_1	4.30 ± 0.24 [—]
k_2	28.74 ± 13.40 [—]

Algorithm 1: Add distortion uncertainty U_c to experimental images

Given experimental image $X_k : \mathbb{P}_{i,j} \Rightarrow [0, 255]^3$ with $i = [0; 512]$ and $j = [0; 64]$

Step 1. Draw a sample for the percentage of distorted pixels $U_c \sim \mathcal{U}(0.133, 0.7)[\%]$

Step 2. Create the corresponding distortion map $D_{map,k} \sim B(n = 64 \times 512, p = U_c)$

Step 3. Create the uncertain image copy $\widetilde{X}_k = \begin{cases} 0 & \text{for } D_{map,k} = 1 \\ X_k & \text{otherwise} \end{cases}$

2.1.2 Non-Zero-Angle Fuel Placement Error U_γ

Under ideal circumstances, the slab burner is positioned perfectly orthogonal to the axis of the camera. However, the researcher who places the fuel specimen may not always do so with precision, and place the fuel at a small offset angle

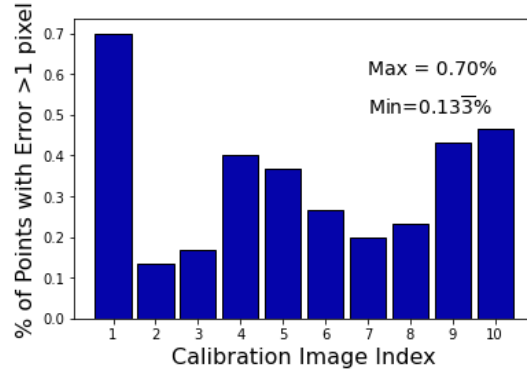


FIG. 5: For each of the calibration images, we expect between 0.133% and 0.7% of the image to be distorted for a distance greater than 1 pixel. From this result, we assume a random experimental image from our camera to have a number of distorted pixels $U_c \sim \mathcal{U}(0.133, 0.7)[\%]$.

$\gamma \neq 0$, which then adds uncertainty to the experimental images via over- or underestimating the distances between borders for the specimen [24]. The relative border location is important in accurate estimation of the regression rate, since we track the border changes to measure it. The schematic in Fig. 6 shows the geometry of this phenomenon. Given the true length of the fuel specimen $y = 8.069$ cm, the length of the specimen to the camera is $L = 37.46$ cm, the angle of the camera to the end of the specimen if placed correctly is $\beta = \tan^{-1} y/2L = 6.147^\circ$, and solving the geometry, we get an estimate for the perceived specimen length of $y_m = y(\cos \gamma + \tan \beta \sin \gamma)$. We express the error in the perceived fuel length, caused by the non-zero-angle placement of the specimen, as $U_\gamma = (y_m - y)/y[\%]$. The extreme values for U_γ are directly related to the extreme values for the angle γ . We assume that $-5^\circ \leq \gamma \leq 5^\circ$, because a placement by a larger angle would likely be visually noticed by the researcher and corrected before the experiment, which corresponds to $-1.319 \leq U_\gamma[\%] \leq 0.558$ as shown in Fig. 7. Given that we have no prior knowledge about the distribution of the error, we assume the non-zero-angle placement uncertainty to be uniform $U_\gamma \sim \mathcal{U}(-1.319, 0.558)[\%]$.

To introduce U_γ to a given experimental image X_k , we first estimate the pixel density for the fuel specimen ρ_k . Given the area of the fuel surface (which can be taken from the manually segmented or predicted mask) and the length




FIG. 6: If the fuel specimen is not placed orthogonally to the camera, but with an offset small angle γ , the distance between the fuel boundaries is misrepresented ($y_m \neq y$)




FIG. 7: Change of non-zero-angle placement error U_γ for different placement angles γ . For this analysis, we choose $-5^\circ \leq \gamma \leq 5^\circ$ to be a possible range of angle values, which corresponds to $-1.319 \leq U_\gamma[\%] \leq 0.558$.

of the slab, we estimate the density ρ_k = fuel length in pixels/ y [pix/cm] for that image. We then translate U_γ into [cm] units $U_{\gamma,cm} = U_\gamma/100 \times y$, expressing the uncertainty as the lengthwise error between the boundaries. Then the quantity $U_{\gamma,cm} \times \rho_k$ provides the number of pixels that the right fuel boundary has to shift to introduce the sampled value of U_γ . If $U_\gamma > 0$ then the right boundary is shifted to the right, and if $U_\gamma < 0$ then the right boundary is shifted to the left. The process of shifting includes either removing a vertical strip in the middle of the image and stitching the remaining image together adding background pixels to the right ($U_\gamma < 0$), or duplicating the vertical strip in the middle of the image and removing some of the preexisting background on the right ($U_\gamma > 0$). The manipulation happens in the horizontal middle (at 256 pixels) since the right and left boundaries are far away from that location. The reader is referred to Algorithm 2 for the process of introducing U_γ to the images and two visual examples of the extreme cases in Fig. 8.

2.2 Monte Carlo Dropout (MCD) U-Net for Segmentation

2.2.1 Model Form Uncertainty U_m

We created the U-net architecture using the Keras library [25]. Figure 9 shows a schematic of the U-net architecture with five dropout layers. The U-net has an encoding path that initially extracts the most important features from the image as resolution decreases and depth increases, followed by a decoding path to return to higher resolutions by

Algorithm 2: Add non-zero-angle placement uncertainty U_γ to experimental images

Given experimental image $X_k : \mathbb{P}_{i,j} \rightarrow [0, 255]^3$ with $i = [0; 512]$ and $j = [0; 64]$

Step 1. Draw a sample for the non-zero-angle uncertainty $U_\gamma \sim \mathcal{U}(-1.319, 0.558)[\%]$

Step 2. Translate to lengthwise units and find pixel density $U_{\gamma,cm} = U_\gamma y/100$ and ρ_k = fuel length in pixels/ y

Step 3. If $U_\gamma > 0$:

$$\begin{cases} \text{Loc1} = \text{floor}(256 - U_{\gamma,cm}\rho_k/2) \\ \text{Loc2} = \text{floor}(256 + U_{\gamma,cm}\rho_k/2) \\ \text{Added slice } S_k = X_k[\text{loc1} : \text{loc2}] \\ \widetilde{X}_k = \text{append}([X_k[: 256], S_k, X_k[256 : 512 - \text{size}(S_k)]) \end{cases}$$

If $U_\gamma < 0$:

$$\begin{cases} \text{Loc1} = \text{floor}(256 - U_{\gamma,cm}\rho_k/2) \\ \text{Loc2} = \text{floor}(256 + U_{\gamma,cm}\rho_k/2) \\ \widetilde{X}_k = \text{append}([X_k[: \text{loc1}], X_k[\text{loc2} : 512], [0]_{\text{loc2}-\text{loc1}}]) \end{cases}$$




FIG. 8: Different versions of the same experimental image at extreme values: (a) $U_\gamma = 0.558\%$, (b) $U_\gamma = 0$, and (c) $U_\gamma = -1.319\%$. The red (vertical) line is a reference point on the edge of the right boundary in the image with no uncertainty to aid the visualization of the difference in the cases with uncertainty present.

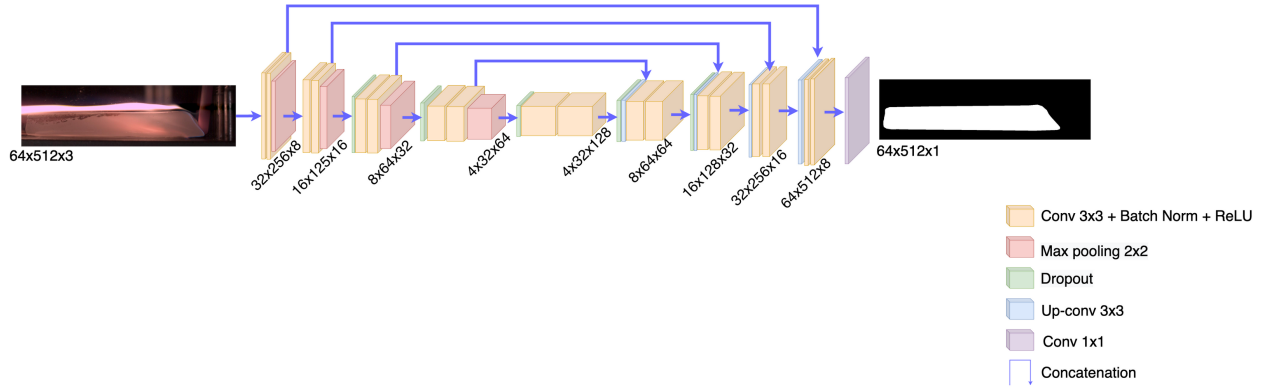


FIG. 9: The U-net neural network architecture with dropout in the intermediate layers. Given an input experimental image, the U-net outputs the predicted fuel segmentation mask.

reconstructing the image using up-sampling [26]. The encoding path consists of four sequences of the following: two 3×3 convolutions, batch normalization, ReLU activation, and max pooling. The decoding path consists of four sequences of 3×3 up-convolutions with a concatenation of the feature map from the corresponding level in the encoding path, two 3×3 convolutions, batch normalization, and lastly a ReLU activation. Finally, the U-net includes a 1×1 convolution operation to output the segmented fuel mask. The total number of trainable parameters in the network is 492,609. The parameters of a convolution layer are equal to $out \times [in \times (3 \times 3) + 1]$, where out is the number of features after the convolution, in is the number of features before the convolution, the nine parameters of the 3×3 filter, and one parameter at each node for the bias term. For example, the first convolution layer takes the original image ($64 \times 512 \times 3$), which after the convolution becomes ($64 \times 512 \times 8$); therefore that layer has $8 \times [3 \times (3 \times 3) + 1] = 224$ parameters. Each of the batch normalization layers adds four additional parameters per feature of the previous layer [27]. For example, since the output after the first convolution layer has eight features, the corresponding batch normalization layer will have 32 parameters (16 of which are nontrainable).

To estimate the prediction model form uncertainty U_m , we implemented Monte Carlo dropout (MCD), a Bayesian approximation in deep learning models [28]. The dropout process randomly silences neurons with a probability of p_D in the intermediate layers of the U-net. Then, at inference, the model is sampled with dropout which is an approximation to sampling from the posterior weight distribution of a fully Bayesian network [29]. Given a new test image X_k , and the training set (X_{train}, Y_{train}) , the resulting mean segmented mask can be approximated using Monte Carlo integration as

$$\hat{p}(y = c | X_{k,test}, X_{train}, Y_{train}) \approx \frac{1}{T} \sum_{t=1}^T \sigma(f^{\hat{W}_t}(X_k)), \quad (1)$$

where $\hat{p}(y)$ corresponds to the predicted probability that a pixel in the test image X_k is classified as $c = \{\text{Fuel}\}$, T is the number samples for the weights set to 20 [30], $\sigma(\cdot)$ is the softmax function, and f is the model output given the weights $\hat{W}_t \sim q_\theta(W_t)$, sampled from the dropout approximating variational distribution $q_\theta(W_t)$ [16], which is defined for every layer t :

$$W_t = w_t \cdot \text{diag}([z_{t,m}]_{m=1}^{K_t}), \quad (2)$$

$$z_{t,m} \begin{cases} \sim \text{Bernoulli}(p_D) & \text{for layers } t = 7, 10, 13, 15, 18 \text{ and } m = 1, \dots, K_{t-1}, \\ = 1 & \text{otherwise} \end{cases}, \quad (3)$$

where w_t are the weights of the trained network. The $\text{diag}(\cdot)$ operator maps vectors to diagonal matrices whose diagonals are the elements of the vectors [31]. The variable $z_{t,m} = 0$ corresponds to node m in layer $t - 1$ to be dropped out as an input to the next layer, whereas $z_{t,m} = 1$ corresponds to the trained weights w_t to be used as is, which is also the case for layers without dropout, $p_D = 0.5$ is the dropout probability, and K_t is the number of nodes/units within layer t . We note that the dropout probability is non-zero only in the intermediate blocks of the U-net, which corresponds to layers 7, 10, 13, 15, and 18 of the network.

The model form uncertainty for the predicted probability \hat{p} can be expressed as the corresponding map of the cross-entropy of the two class dimensions [16,32,33]:

$$U_m = - \sum_{c=1}^2 \hat{p}_c \log(\hat{p}_c) = -[\hat{p}_{fuel} \log(\hat{p}_{fuel}) + \hat{p}_{background} \log(\hat{p}_{background})], \quad (4)$$

where U_m is the model form uncertainty map, \hat{p}_{fuel} is the predicted probability that a given pixel corresponds to fuel, and $\hat{p}_{background}$ is the predicted probability that a given pixel corresponds to noise/background.

The U-net was trained for 79 epochs with 179 images from the early, mid, and late phases from the first experiment with oxidizer flux $G_1 = 6.96 \text{ kg/m}^2\text{s}$, accompanied with their manually segmented fuel masks that were traced using VGG [34]. Figure 10 shows the loss progression during training. Observing the results from the testing set (Fig. 11), the U-net predicts the segmentation masks correctly, even for highly saturated images and without including any other components of the chamber as part of the mask (e.g., the glow plug). The uncertainty maps show the largest level of uncertainty to be around the fuel boundary geometries. The model form uncertainty further inside and outside of the identified fuel boundary is very small.

2.2.2 Manual Segmentation Uncertainty U_s

In the process of training the U-net, we inherently make the assumption that the manually segmented fuel masks are ground truth. However, there is expected to be some variability between individuals that segment the fuel masks, and repeating the process multiple times to investigate the level of this variability is beyond the research scope of

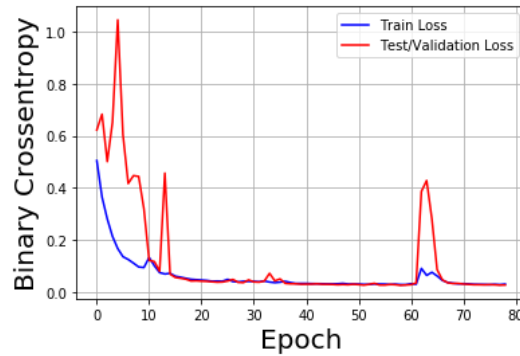


FIG. 10: Binary cross-entropy loss during the training of the U-net used for this work. Some abrupt changes in the loss are due to dropout effects. The training stops once the validation loss has not improved after 20 consecutive epochs.

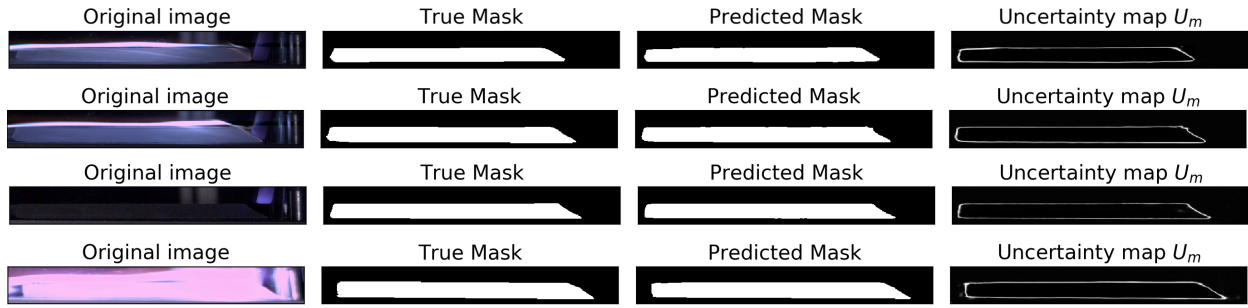


FIG. 11: Results from the baseline U-net model without any data uncertainty introduced to the images. The U-net performs well on all images, including low, medium, and high saturation levels during the phases of the experiment.

this paper. There have been studies to quantify the effect of improperly segmented fuel masks to the U-net outputs (e.g., see [35,36]), which found that the error in the segmentation accuracy for the MNIST [37] and MSLesion [38] dataset is 0.83% and 0.56%, respectively. The MNIST dataset consists of gray-scale images of digits from 0 to 9, and the MSLesion dataset consists of 21 medical scans. Given that the fuel segmentation masks are simple with only two labels, and without any prior knowledge, we assume the manual segmentation prediction uncertainty to be similar to those datasets and uniform $U_s \sim \mathcal{U}(0.56, 0.83)\%$, expressed as a percentage variance of the corresponding segmentation probability \hat{p}_{fuel} . Therefore, the complete resulting uncertainty map for the mean prediction μ_{mask} of the U-net after considering all the uncertainty sources is

$$U_{mask} = U_m | (U_c, U_\gamma) + \hat{p}_{fuel} U_s, \quad (5)$$

where U_{mask} is the complete resulting uncertainty map after all sources of uncertainty are included, U_m is the model form uncertainty, U_c is the distortion uncertainty, U_γ is the non-zero-angle placement uncertainty, \hat{p}_{fuel} is the corresponding segmentation probability for the fuel class, and U_s is the manual segmentation prediction uncertainty.

2.3 Fuel Surface Detection to Measure \dot{r}

To approximate the regression rate from the binary fuel masks, we follow the approach described in [10]. For a given sequence of fuel masks $\mu_{mask,k}$ representing the fuel profile between time intervals Δt , we track the height of the top surface of the fuel profile through calibration points. At each point, the localized regression rate is

$$r_{l,k} = \frac{h(t_{k+1}) - h(t_k)}{\Delta t}, \quad (6)$$

where $h(\cdot)$ is the fuel height for a given time measured from the bottom of the fuel mask and Δt is the time difference between two frames. The total regression rate for the experiment is then the average localized regression rate from all points:

$$\dot{r} = \frac{\sum_{k=1}^M r_{l,k}}{M}. \quad (7)$$

Given that each of the predicted fuel masks $\mu_{mask,k}$ also has an accompanying uncertainty map $U_{mask,k}$ after the inclusion of all uncertainties, we calculate two extreme values for the total regression rate estimate \dot{r} . The upper bound is found by adding the uncertainty maps $U_{mask,k}$ to the corresponding fuel masks $\mu_{mask,k}$, flipping the class from “background” to “fuel” for any pixels that now have $\hat{p}_{fuel} \geq 0.5$ with the added uncertainty, and then recalculating a total regression rate estimate \dot{r}^+ . The lower bound comes from the opposite process: subtracting the uncertainty maps $U_{mask,k}$ from the corresponding fuel masks $\mu_{mask,k}$, flipping the class from “fuel” to “background” for any pixels that now have $\hat{p}_{fuel} < 0.5$ after the subtracted uncertainty, and recalculating a total regression rate estimate \dot{r}^- . From this entire process, for a given sequence of experimental images (with or without added data uncertainties U_c, U_γ) and using the U-net (with or without the manual segmentation uncertainty U_s), we calculate a mean estimate for the total regression rate from the experiment \dot{r} with its bounds (\dot{r}^-, \dot{r}^+) .

3. COMBINATION AND FORWARD-PROPAGATION OF UNCERTAINTY SOURCES

In this section, we study all possible combinations of how each of the individual data uncertainty sources shown in Fig. 1 impacts the overall uncertainty propagation process. For each of the three uncertainty sources U_c , U_γ , U_s , we consider two scenarios: either they do not contribute (i.e., are zero) or do contribute (i.e., are uniformly distributed as described in the previous section). All possible cases are shown in the Table 2.

First, we observe how the error from the camera distortion U_c and the non-zero-angle placement error U_γ impact the U-net model form uncertainty map U_m separately and together (Cases 1, 2, and 3). Secondly, we observe how the manual segmentation error U_s changes the complete uncertainty map U_{mask} individually and together with U_c , U_γ (Cases 4, 5, and 6). Lastly, we showcase an example where all three uncertainties are included (Case 7), which is representative of the approach followed in the calculation of the full probability distribution of the regression rate in the next section, which is the end goal of the paper. For these studies, we used three representative experimental images from the testing set of the U-net that correspond to three different image saturation levels in the experiment: “low,” “mid,” and “high,” as shown in Fig. 12. The histograms are the model form uncertainty near the boundary for each image. Observing the model form uncertainty *near* the fuel boundary, we identify three different modes depending on the saturation level in the experimental images. Images with low saturation [Fig. 12(a)], show an approximately uniform profile across all uncertainty ranges. Images of medium saturation [Fig. 12(b)] show a bimodal behavior for U_m with a uniform profile in between the peaks. Lastly, images with high saturation [Fig. 12(c)] have again a bimodal profile similar to the medium saturation images, but the intermediate levels of uncertainty are higher, which is attributed to the higher saturation from the more violent burn.

We follow a standard forward-propagation uncertainty approach: if U_c and U_γ are present, we create an ensemble of perturbed/uncertain copies of a given image by following the process in Algorithm 1 or Algorithm 2, respectively, and pass these uncertain copies through the U-net to get the model form uncertainty map U_m . If U_s is present, that is added to the model form uncertainty map U_m proportionally to the predicted probability of each pixel [Eq. (5)]. For each of these cases, we sampled the corresponding uniform distributions $N = 10,000$ times and show the resulting uncertainty distribution compared to the baseline case where no uncertainties are present.

TABLE 2: Studied cases of uncertainty sources

Case #	Distortion U_c	Non-zero-angle U_γ	Manual segmentation U_s
Baseline	0	0	0
1	$\mathcal{U}(0.133, 0.7)\%$	0	0
2	0	$\mathcal{U}(-1.319, 0.558)\%$	0
3	$\mathcal{U}(0.133, 0.7)\%$	$\mathcal{U}(-1.319, 0.558)\%$	0
4	0	0	$\mathcal{U}(0.56, 0.83)\%$
5	$\mathcal{U}(0.133, 0.7)\%$	0	$\mathcal{U}(0.56, 0.83)\%$
6	0	$\mathcal{U}(-1.319, 0.558)\%$	$\mathcal{U}(0.56, 0.83)\%$
7	$\mathcal{U}(0.133, 0.7)\%$	$\mathcal{U}(-1.319, 0.558)\%$	$\mathcal{U}(0.56, 0.83)\%$




FIG. 12: Example images from the experiment with different saturation levels at low (a), medium (b), and high (c) with their corresponding model form uncertainty histograms. We used these representative images to account for different image qualities, since the uncertainty distribution has a different profile for each of them.

For Case 1, where only the uncertainty from the distorted pixels U_c is present, we find that distortion adds some variability to the U_m probability distribution, but still follows the overall pattern of the distribution without distortion in the low and medium saturated images. The small variability is expected since the distortion is added uniformly to the entire image. However, in the highly saturated case, distortion appears to be a moderating factor to the large peaks of small uncertainty in the original image. The explanation for this irregularity is justified in Fig. 13. When distortion is present in the highly saturated experimental image, the uncertainty around distorted points with high density is much larger compared to the image without distortion. The distorted points outside the fuel boundary do not contribute to the uncertainty since they were already part of the background. However, the distorted pixels that are present within the fuel boundary add additional uncertainty, which changes the bimodal profile for the probability density present in the baseline case to a more uniform profile for the lower uncertainty values.

For Case 2, where only the non-zero-angle fuel placement uncertainty U_γ is present, we find that some variability is added, but the overall pattern of the distribution of uncertainty compared to the baseline case is maintained for all saturation levels (Fig. 14). Case 3, which includes both data uncertainties U_c, U_γ , also appears to maintain the overall

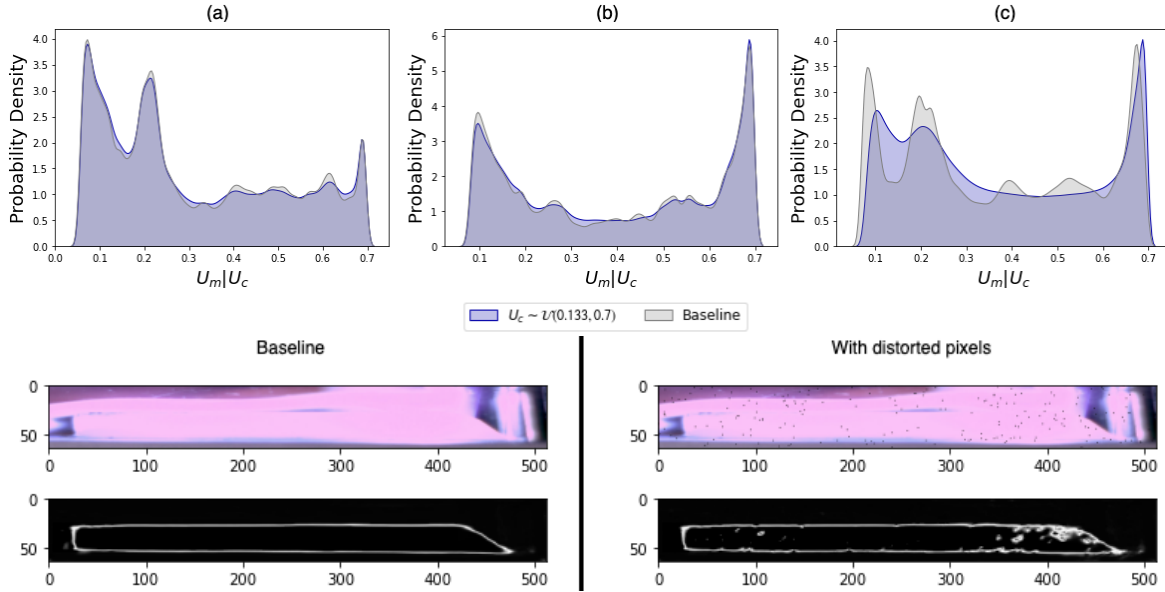


FIG. 13: Case 1: Impact of the distorted pixel uncertainty U_c on the model form uncertainty U_m for three levels of saturation for the experimental images: low (a), medium (b), and high (c). For the image with high saturation, distortion changes the uncertainty profile for the lower uncertainty values.

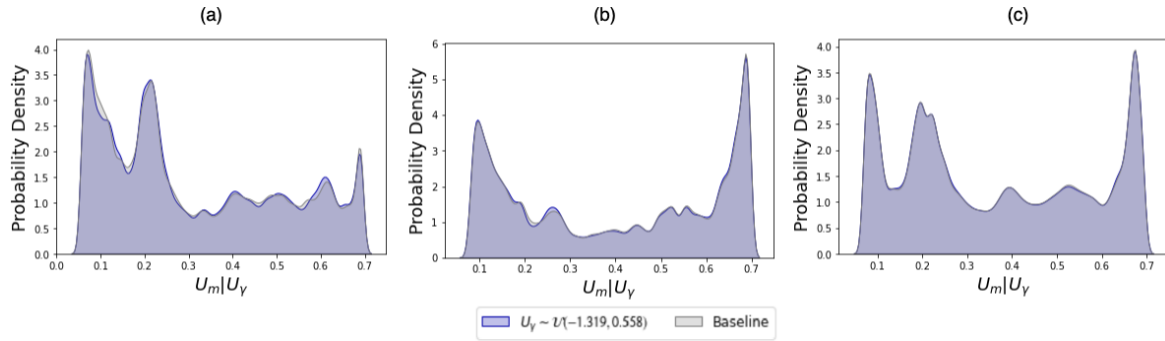


FIG. 14: Case 2: Impact of the non-zero-angle fuel placement uncertainty U_γ on the model form uncertainty U_m for three levels of saturation for the experimental images: low (a), medium (b), and high (c). The overall distribution of uncertainty is maintained.

profile, but the non-zero-angle effects appear to negate some of the variability caused by distortion, likely due to the removal or addition of parts of the same image (Fig. 15).

Cases 4–6 show the effects of the manual segmentation uncertainty U_s individually and together with the other uncertainty sources. U_s has an added effect based on the predicted probability of a given pixel; therefore the additive effects are present for those pixels with high fuel probability (i.e., within and on the fuel boundary). The result appears as a right-shift to the uncertainty distribution profile (Figs. 16–18).

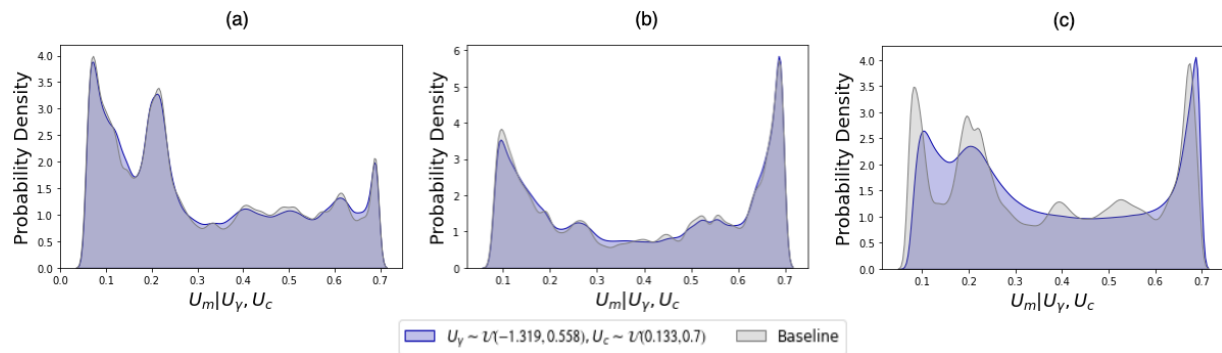


FIG. 15: Case 3: Impact of the non-zero-angle fuel placement uncertainty U_γ and distorted pixel uncertainty U_c on the model form uncertainty U_m for three levels of saturation for the experimental images: low (a), medium (b), and high (c). Compared to Case 1, the presence of non-zero-angle uncertainty U_γ negates some of the variability caused by the distortion uncertainty U_c .

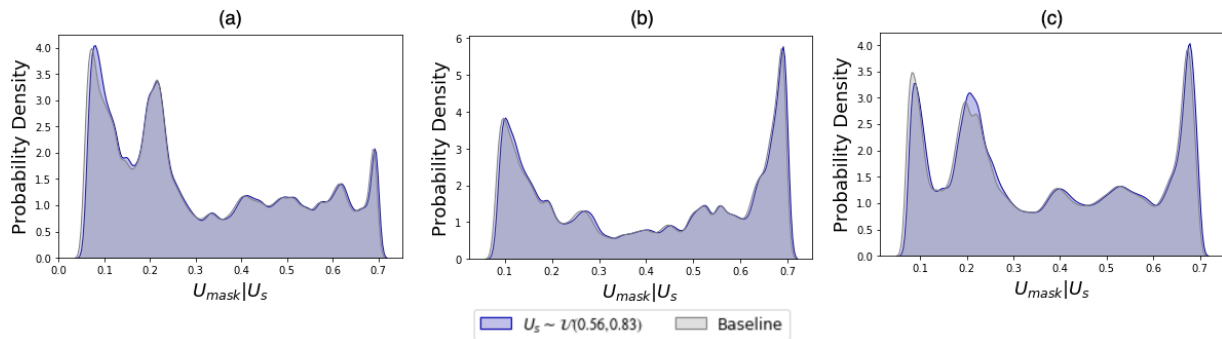


FIG. 16: Case 4: Impact of the manual segmentation uncertainty U_s on the overall uncertainty map U_{mask} for three levels of saturation for the experimental images: low (a), medium (b), and high (c). The added effect of this uncertainty appears as a right-shift to the uncertainty distribution profile.




FIG. 17: Case 5: Impact of the distorted pixel uncertainty U_c and the manual segmentation uncertainty U_s to the overall uncertainty map U_{mask} for three levels of saturation for the experimental images: low (a), medium (b), and high (c)




FIG. 18: Case 6: Impact of the non-zero-angle fuel placement uncertainty U_γ and the manual segmentation uncertainty U_s to the overall uncertainty map U_{mask} for three levels of saturation for the experimental images: low (a), medium (b), and high (c)

Case 7 shows the combined effects of the studied uncertainties, which appear to be mixed: a right-shift caused by the manual segmentation uncertainty U_s , a moderating effect caused by removing or replicating parts of these images from U_γ , and added variability (and profile change for the highly saturated images) due to more areas with uncertainty from U_c (Fig. 19).

4. ESTIMATION OF THE DISTRIBUTION OF \dot{r}

The goal of this section is to implement the framework shown in Fig. 1 for the experiment with oxidizer mass flux $G_2 = 10.96 \text{ kg/m}^2\text{s}$, and find the probability distribution of the QoI: the total regression rate \dot{r} . To begin, we first capture a sequence of images from the experiment at fixed frame intervals (Fig. 20).




FIG. 19: Case 7: Impact of the distorted pixel uncertainty U_c , the non-zero-angle fuel placement uncertainty U_γ , and the manual segmentation uncertainty U_s to the overall uncertainty map U_{mask} for three levels of saturation for the experimental images: low (a), medium (b), and high (c)




FIG. 20: Images from experiment 2 with oxidizer mass flux $G_2 = 10.96 \text{ kg/m}^2\text{s}$. The frames shown are taken at a fixed interval of 0.32 s. The sequence of the burn starts from the top left and ends at the bottom right. The goal is to measure the probability distribution of the average regression rate, the quantity that describes how fast the fuel recedes during the burn.

Then, we follow a forward-propagation Monte Carlo approach: the data uncertainties U_c, U_γ are uniformly sampled and added to the sequence of images from experiment 2 (Fig. 20), the images are passed through the U-net to get the segmented fuel masks $\mu_{mask,k}$, and the manual segmentation uncertainty U_s is uniformly sampled and added to the model form uncertainty U_m [Eq. (5)] to get the overall uncertainty maps $U_{mask,k}$ for all the experimental images (see Fig. 21 for one example of an ensemble of fuel masks with profile tracking). From one sequenced ensemble of $\mu_{mask,k}$ and $U_{mask,k}$, one point measurement of the total regression rate \dot{r} is computed together with its bounds (\dot{r}^+, \dot{r}^-). The process is repeated until the variance of the final probability distribution of the regression rate $\sigma_{\dot{r}}^2$ converges to a tolerance of e^{-6} (Algorithm 3).

The distribution of the regression rate is shown in Fig. 22. We compare two instances: a probability distribution if we consider the uncertainty maps U_{mask} that provide the bounds for the regression rate with each ensemble or if we do not. The probability distribution considering only the mean masks μ_{mask} from the ensembles resembles a Gaussian distribution with a mean regression rate $\mu_{\dot{r}} = 0.759954$ mm/s. Considering the uncertainty maps that provide the bounds for the regression rate from each ensemble adds variability to both tail ends of the distribution with a shift of the mean to $\mu_{\dot{r}} = 0.740032$ mm/s. In previous work where we only considered the model form uncertainty U_m , we found the regression rate for an experiment with $G = 9.58$ kg/m²s to be 0.74 ± 0.09 mm/s. We note that when the uncertainty maps are included in the calculation, a significant amount of density is shifted towards the smaller




FIG. 21: One ensemble of predicted fuel masks from experiment 2 with oxidizer mass flux $G_2 = 10.96$ kg/m²s, with visible fuel tracing. The process of calibration points ensures that in cases where the U-net misclassifies some part of the picture as fuel, which is likely in the later parts of the experiment, the algorithm approximates the boundary correctly for the regression rate calculation.

Algorithm 3: Estimation of the distribution of the regression rate \dot{r}

Given a sequence of experimental images $X_k : \mathbb{P}_{i,j} \Rightarrow [0, 255]^3$ with $i = [0; 512]$ and $j = [0; 64]$

While $\sigma_{\dot{r}}^{2,new} - \sigma_{\dot{r}}^{2,old} > e^{-6}$ **do:**

- Step 1.** Draw a sample for the non-zero-angle uncertainty $U_\gamma \sim \mathcal{U}(-1.319, 0.558)[\%]$
- Step 2.** Draw a sample for the distortion uncertainty $U_c \sim \mathcal{U}(0.133, 0.7)[\%]$
- Step 3.** Introduce U_c, U_γ to the ensemble X_k (Algorithms 1 and 2)
- Step 4.** Pass the ensemble through U-net to get the mean predicted fuel masks $\mu_{mask,k}$ and the mean model form uncertainty maps $U_{m,k}$
- Step 5.** Draw a sample for the manual segmentation uncertainty $U_s \sim \mathcal{U}(0.56, 0.83)[\%]$
- Step 6.** Compute overall uncertainty map for the ensemble $U_{mask,k} = U_{m,k}[(U_c, U_\gamma) + \hat{p}_{fuel,k}U_s]$
- Step 7.** Compute point estimate for regression rate \dot{r} and bounds \dot{r}^+, \dot{r}^- from $\mu_{mask,k}$ and $U_{mask,k}$
- Step 8.** Append probability distribution $p(\dot{r}) = + \{\dot{r}, \dot{r}^+, \dot{r}^-\}$ and update variance $\sigma_{\dot{r}}^{2,new}$

End While

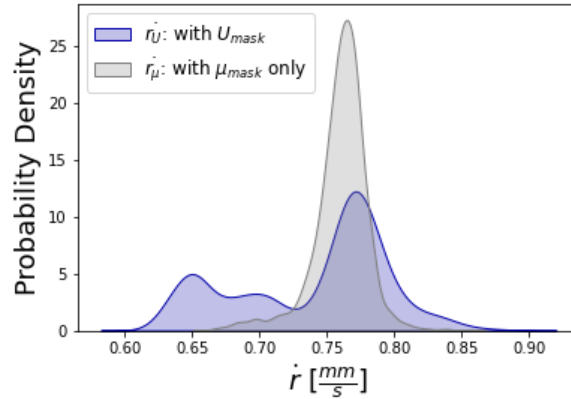


FIG. 22: The probability distribution of the regression rate for the experiment with oxidizer mass flux $G_2 = 10.96 \text{ kg/m}^2\text{s}$. The gray (higher) curve only considers the mean fuel masks μ_{mask} for the calculation (i.e., does not include the regression rate bounds using U_{mask}), whereas the blue (lower) curve includes the information from the uncertainty maps. The mean of the probability distribution from the mean masks only is $\mu_{r_\mu} = 0.759954 \text{ mm/s}$ (higher curve) and considering the uncertainty maps as well is $\mu_{r_U} = 0.740032 \text{ mm/s}$. The uncertainty maps translate to significant variability to both tail ends of the distribution.

regression rate values, because of how the surface tracking algorithm works. The surface tracking algorithm accounts for a physical constraint: it does not allow for the fuel surface to increase in size during the burn, which would be a violation of combustion laws. Some images in the later stages of the experiment ensembles can often result in larger fuel surface than previous time steps when we add the uncertainty masks, and that effect is neglected. Uncertainty masks tend to have higher values as the burn proceeds (images are more saturated). In the subtracting case, the equivalent scenario does not happen; the higher uncertainty maps in the later stages of the burn will be considered normally. Therefore, the process penalizes the higher regression rates when uncertainties are added, resulting in a more significant density shift to lower values. The contributions of this paper agree with trends from previous reports in the literature, but add the additional step of characterization and quantification of data uncertainty which translates to more accurate representation of the distribution of the regression rate when measured from images.

5. CONCLUSIONS AND FUTURE WORK

In this work, we presented a full forward-propagation framework for uncertainty quantification of the regression rate from experimental images. We characterized uncertainty sources that are pertinent to this case: distortion uncertainty U_c from the camera, non-zero-angle placement of the fuel U_γ , manual segmentation training uncertainty U_s , and model form uncertainty from the Monte Carlo dropout U-net U_m . The image data generated for this work came from two slab burner experiments with measured oxidizer mass fluxes of $G_1 = 6.96 \text{ kg/m}^2\text{s}$ and $G_2 = 10.96 \text{ kg/m}^2\text{s}$. We conducted seven case studies of all possible combinations between the three uncertainty sources and found that their effects combine: a right-shift in the model form uncertainty distribution caused by the additive effect of the manual segmentation uncertainty U_s , a moderating effect from U_γ caused by the effect of placing the fuel under a small angle, and the overall added variability throughout the model form uncertainty distribution (and profile change for the highly saturated images) due to more areas with uncertainty caused from the distortion uncertainty U_c . We completed our study by using the images from the experiment with oxidizer flux $G_2 = 10.96 \text{ kg/m}^2\text{s}$, generating a large ensemble of their uncertain copies by sampling the three uncertainty sources, and tracking the fuel profiles in each case to estimate the full regression rate probability distribution. The probability distribution considering only the mean masks μ_{mask} from the ensembles resembles a Gaussian distribution with a mean regression rate $\mu_{r_\mu} = 0.759954 \text{ mm/s}$. Considering the uncertainty maps that provide the bounds for the regression rate from each ensemble adds variability to both tail ends of the distribution with a shift of the mean to $\mu_{r_U} = 0.740032 \text{ mm/s}$.

In future work, we will use the resulting probability distribution of the regression rate from experimental data shown in this work, together with experiments on other oxidizer fluxes, to validate an equivalent estimate for the

regression rate from coupled flow and chemistry simulations for the slab burner combustion phenomena using our open-source software framework ABLATE (<https://ablate.dev/>).

ACKNOWLEDGMENTS

The authors would like to thank the reviewers for comments that greatly improved the manuscript.

This work was supported by the United States Department of Energy's (DoE) National Nuclear Security Administration (NNSA) under the Predictive Science Academic Alliance Program III (PSAAP III) at the University at Buffalo, under contract number DE-NA0003961.

Computational resources: the authors acknowledge the Tufts University High Performance Compute Cluster which was used for training the U-net models and running the simulations for the uncertainty analysis.

REFERENCES

1. Moreno-Rodenas, A.M., Tscheikner-Gratl, F., Langeveld, J.G., and Clemens, F.H., Uncertainty Analysis in a Large-Scale Water Quality Integrated Catchment Modelling Study, *Water Res.*, **158**:46–60, 2019.
2. Tan, J., Villa, U., Shamsaei, N., Shao, S., Zbib, H.M., and Faghihi, D., A Predictive Discrete-Continuum Multiscale Model of Plasticity with Quantified Uncertainty, *Int. J. Plasticity*, **138**:102935, 2021.
3. Jones, R.E., Redle, M.T., Kolla, H., and Plews, J.A., A Minimally Invasive, Efficient Method for Propagation of Full-Field Uncertainty in Solid Dynamics, *Int. J. Numer. Methods Eng.*, **122**(23):6955–6983, 2021.
4. Smith, R., *Uncertainty Quantification*, Philadelphia: SIAM, 2014.
5. Psaros, A.F., Meng, X., Zou, Z., Guo, L., and Karniadakis, G.E., Uncertainty Quantification in Scientific Machine Learning: Methods, Metrics, and Comparisons, *J. Comput. Phys.*, **477**:111902, 2023.
6. Krygier, M.C., LaBonte, T., Martinez, C., Norris, C., Sharma, K., Collins, L.N., Mukherjee, P.P., and Roberts, S.A., Quantifying the Unknown Impact of Segmentation Uncertainty on Image-Based Simulations, *Nat. Commun.*, **12**:5414, 2021.
7. Zilliac, G. and Karabeyoglu, M., Hybrid Rocket Fuel Regression Rate Data and Modeling, in *Proc. of the 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, 2006.
8. Zilliac, G., Story, G.T., Karp, A.C., Jens, E.T., and Whittinghill, G., Combustion Efficiency in Single Port Hybrid Rocket Engines, in *Proc. of the AIAA Propulsion and Energy 2020 Forum*, 2020.
9. Karabeyoglu, M., Cantwell, B., and Altman, D., Development and Testing of Paraffin-Based Hybrid Rocket Fuels, in *Proc. of the 37th Joint Propulsion Conf. and Exhibit*, 2001.
10. Budzinski, K., Aphale, S.S., Ismael, E.K., Surina, G., and DesJardin, P.E., Radiation Heat Transfer in Ablating Boundary Layer Combustion Theory Used for Hybrid Rocket Motor Analysis, *Combust. Flame*, **217**:248–261, 2020.
11. Surina, G., Georgalis, G., Aphale, S.S., Patra, A., and DesJardin, P.E., Measurement of Hybrid Rocket Solid Fuel Regression Rate for a Slab Burner Using Deep Learning, *Acta Astronaut.*, **190**:160–175, 2022.
12. Hawkes, E.R., Sankaran, R., Sutherland, J.C., and Chen, J.H., Direct Numerical Simulation of Turbulent Combustion: Fundamental Insights towards Predictive Models, *J. Phys.*, **16**:65–79, 2005.
13. Yousefian, S., Bourque, G., and Monaghan, R.F., Bayesian Inference and Uncertainty Quantification for Hydrogen-Enriched and Lean-Premixed Combustion Systems, *Int. J. Hydrogen Energy*, **46**(46):23927–23942, 2021.
14. Bettis, B. and Hosder, S., Uncertainty Quantification in Hypersonic Reentry Flows Due to Aleatory and Epistemic Uncertainties, in *Proc. of the 49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2011.
15. Hüllermeier, E. and Waegeman, W., Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods, *Mach. Learn.*, **110**(3):457–506, 2021.
16. Kendall, A. and Gal, Y., What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?, in *Advances in Neural Information Processing Systems*, Vol. 30, 2017.
17. Hora, S.C., Aleatory and Epistemic Uncertainty in Probability Elicitation with an Example from Hazardous Waste Management, *Reliab. Eng. Syst. Saf.*, **54**(2):217–223, 1996.
18. Kiureghian, A.D. and Ditlevsen, O., Aleatory or Epistemic? Does It Matter?, *Struct. Saf.*, **31**(2):105–112, 2009.

19. Dunn, C., Gustafson, G., Edwards, J., Dunbrack, T., and Johansen, C., Spatially and Temporally Resolved Regression Rate Measurements for the Combustion of Paraffin Wax for Hybrid Rocket Motor Applications, *Aerosp. Sci. Technol.*, **72**:371–379, 2018.
20. Bouguet, J. and Perona, P., Camera Calibration from Points and Lines in Dual-Space Geometry, in *Proc. of the 5th European Conf. on Computer Vision*, pp. 2–6, 1998.
21. Juarez-Salazar, R., Zheng, J., and Diaz-Ramirez, V.H., Distorted Pinhole Camera Modeling and Calibration, *Appl. Opt.*, **59**(36):11310–11318, 2020.
22. Heikkila, J. and Silven, O., A Four-Step Camera Calibration Procedure with Implicit Image Correction, in *Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 1106–1112, 1997.
23. Bouguet, J.Y., Camera Calibration Toolbox for Matlab, from http://www.vision.caltech.edu/bouguetj/calib_doc, 2008.
24. Robbe, C., Nsiampa, N., Oukara, A., and Papy, A., Quantification of the Uncertainties of High-Speed Camera Measurements, *Int. J. Metrol. Qual. Eng.*, **5**(2):201, 2014.
25. Chollet, F., Keras, from <https://keras.io>, 2015.
26. Ronneberger, O., Fischer, P., and Brox, T., U-Net: Convolutional Networks for Biomedical Image Segmentation, in *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2015*, N. Navab, J. Hornegger, W.M. Wells, and A.F. Frangi, Eds., Cham, Switzerland: Springer International Publishing, pp. 234–241, 2015.
27. Ioffe, S. and Szegedy, C., Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift, in *Proc. of the 32nd Int. Conf. on Machine Learning*, F. Bach and D. Blei, Eds., Lille, France, pp. 448–456, 2015.
28. Gal, Y. and Ghahramani, Z., Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning, in *Proc. of the 33rd Int. Conf. on Machine Learning*, Vol. 48, pp. 1050–1059, 2016.
29. Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., Fieguth, P., Cao, X., Khosravi, A., Acharya, U.R., Makaremkov, V., and Nahavandi, S., A Review of Uncertainty Quantification in Deep Learning: Techniques, Applications and Challenges, *Inf. Fusion*, **76**:243–297, 2021.
30. DeVries, T. and Taylor, G.W., Leveraging Uncertainty Estimates for Predicting Segmentation Quality, *Comput. Vis. Pattern Recognit.*, arXiv:1807.00502, 2018.
31. Gal, Y. and Ghahramani, Z., Bayesian Convolutional Neural Networks with Bernoulli Approximate Variational Inference, *Proc. of the 6th Int. Conf. on Learning Representations*, 2016.
32. Hastie, T., Tibshirani, R., and Friedman, J., *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, New York: Springer, 2009.
33. Shannon, C.E., A Mathematical Theory of Communication, *SIGMOBILE Mob. Comput. Commun. Rev.*, **5**(1):3–55, 2001.
34. Dutta, A. and Zisserman, A., The VIA Annotation Software for Images, Audio and Video, *Proc. of the 27th ACM Int. Conf. on Multimedia*, Vol. 72, pp. 2276–2279, 2019.
35. Zhang, L., Tanno, R., Xu, M.C., Jin, C., Jacob, J., Ciccarrelli, O., Barkhof, F., and Alexander, D., Disentangling Human Error from Ground Truth in Segmentation of Medical Images, in *Advances in Neural Information Processing Systems*, Red Hook, NY: Curran Associates, Inc., Vol. 33, pp. 15750–15762, 2020.
36. Kohl, S., Romera-Paredes, B., Meyer, C., De Fauw, J., Ledsam, J.R., Maier-Hein, K., Eslami, S.M.A., Jimenez Rezende, D., and Ronneberger, O., A Probabilistic U-Net for Segmentation of Ambiguous Images, in *Advances in Neural Information Processing Systems*, Red Hook, NY: Curran Associates, Inc., Vol. 31, 2018.
37. Deng, L., The MNIST Database of Handwritten Digit Images for Machine Learning Research, *IEEE Signal Process. Mag.*, **29**(6):141–142, 2012.
38. Commowick, O., Kain, M., Casey, R., Ameli, R., Ferré, J.C., Kerbrat, A., Tourdias, T., Cervenansky, F., Camarasu-Pop, S., Glatard, T., Vukusic, S., Edan, G., Barillot, C., Dojat, M., and Cotton, F., Multiple Sclerosis Lesions Segmentation from Multiple Experts: The MICCAI 2016 Challenge Dataset, *Neuroimage*, **244**:118589, 2021.