



Convolutional neural networks for image analysis of high-speed videos from two slab burners

Oliver Assenmacher^{a,*}, Riccardo Gelain^b, Alexander Rüttgers^a, Anna Petrarolo^{c,1}, Patrick Hendrick^b

^a German Aerospace Center (DLR), Institute for Software Technology, High-Performance Computing Department, Linder Hoehe, 51147, Cologne, Germany

^b Université Libre de Bruxelles, Aero-Thermo-Mechanics Department, Avenue F.D. Roosevelt, 50, 1050, Brussels, Belgium

^c German Aerospace Center (DLR), Institute of Space Propulsion, Satellite and Orbital Propulsion Department, Im Langen Grund, 74239, Hardthausen, Germany

ARTICLE INFO

Keywords:

Hybrid rockets
Combustion
Machine learning
Computer vision

ABSTRACT

High-speed video recordings of slab burner experiments were analyzed using a machine learning approach with convolutional neural networks in order to compute the regression rate of hybrid rocket fuels over time. Combustion tests of paraffin-based fuel grains performed in two different hybrid rocket slab burners were recorded with high-speed video cameras and the resulting image data are analyzed in order to determine the height of the fuel in each frame. To this end, a deep neural network with U-net architecture is trained in a supervised fashion to segment the shape of the fuel slab. It is demonstrated that this approach is more capable to segment combustion images in unsteady flow conditions than classical computer vision methods based on thresholding or edge detection. Furthermore, methods in the area of uncertainty quantification of neural networks are applied to estimate the errors in the neural network prediction to new previously unseen data. Finally, the regression rate of the fuel is computed as the rate of change of this height. This method enables automatic analysis of a large amount of video data, taking full advantage of the optical access capabilities of slab burners. Additionally, the method delivers not only the time and space average values of the fuel regression rate, but also quantifies its variation over time and over the length of the slab, providing deeper insights into the combustion mechanics of hybrid rockets.

1. Introduction

The development of hybrid rocket propulsion promises some advantages compared to the more mature technologies of solid and liquid propulsion systems. First, the less complex design allows for a cost reduction in comparison to liquid engines. Second, in contrast to solid motors, they allow for controllable thrust, including shut off and restart capability. Furthermore, hybrid engines are safer than solid ones because the propellants are stored in two different states of matter. However, this characteristic also leads to typically slower burning rates in hybrid propulsion systems. Indeed, the two propellants need to gasify and mix together in the boundary layer before reacting. For this reason, hybrid systems that use conventional polymeric fuels often achieve relatively low regression rate performance, which results in a low thrust level. One viable solution is using liquefying hybrid rocket fuels, for instance, on a paraffin basis, which are considered in the following. Paraffin-based fuels are able to liquefy at lower temperatures and to form a thin melt layer on the fuel surface. The high-speed oxidizer flow

in the fuel port can break the fuel liquid layer into unstable waves, which, in turn, break-up into droplets. The combustion of fuel droplets outside of the boundary layer, where the main combustion process takes place, increases the burning rate of hybrid engines [1,2].

For a better understanding of the combustion mechanics of liquefying hybrid rocket solid fuels and their relation to the regression rate, optical investigations on the combustion behavior of paraffin-based fuels burning with gaseous oxygen have been performed in hybrid rocket slab burners. The essential flow phenomena can be captured by using a high-speed camera to visualize the combustion process. The large image dataset can be used to derive relevant secondary quantities of interest, such as a time-dependent regression rate. The following study focuses on optical regression rate measurements computed from the combustion tests performed at the German Aerospace Center (DLR) and at the Aero-Thermo-Mechanics department of the Université libre de Bruxelles (ULB). The space and time averaged regression rate of a fuel grain of known geometry can be measured by comparing the

* Corresponding author.

E-mail address: oliver.assenmacher@dlr.de (O. Assenmacher).

¹ Current employment at: Gilmour Space Technologies, Gold Coast, Australia.

mass of the fuel before and after the combustion. There are also several options for obtaining time- and space-dependent quantities, enabling a better understanding of the combustion mechanics, e.g. methods based on X-ray radiography as discussed in [3]. The slab burners considered for this study both include an optical access, which can be used to accurately measure the geometry of the solid fuel as it changes during the combustion. Furthermore, by using a high-speed camera to capture the combustion process, a high temporal resolution can be achieved. The main challenge in using an optical approach is accurately separating the flow field and the image background from the solid fuel. This separation is achieved with a convolutional neural network, using an U-Net architecture, which produces binary image masks that indicate the shape of the solid fuel. The regression rate can be computed from these masks by tracking the height of the fuel over time. A similar approach is discussed in [4] for a slab burner experiment for a paraffin wax-gaseous hybrid combustion. However, the results presented in the following differ from the previous ones in several aspects: first, the high-speed videos have been acquired at a frame rate of 2000 and 10,000 fps, respectively. In total, 110,000 frames have to be analyzed and the high frame rate of the camera does not allow for using a high intensity flash for a correct image exposure. Therefore, the obtained images are, in general, underexposed, as the flame tends to saturate the camera and the flame profile hides the contours of the slab. This makes the task of separating the solid fuel from the image background a lot more challenging for both the domain expert and the neural network. Second, the images were captured at two different test facilities so that the results in this study give a more realistic picture of how accurate this method is under varying lab conditions. Third, time-dependent regression rates are derived over a sequence of positions on the surface of the solid fuel to allow for a better analysis of the geometric changes over time.

The remainder of this article is organized as follows: first, the experimental setup is described in Section 2.1. Then, in Section 2.2 the neural network model is introduced and Monte Carlo dropout, an approach to measure the uncertainty of the model prediction, is presented. The details for the training of this neural network are given in Section 2.3 and Section 2.4 describes how the binary masks, produced by the neural network are processed to compute the regression rate. Finally, Section 3.1 presents the results on the accuracy of the trained segmentation model and Section 3.3 contains the derived regression rate measurements.

2. Methods

2.1. Experimental setup

The results presented in this article are based on the data acquired from two different hybrid rocket slab burners. The first batch of combustion tests comes from the test complex M11 of the Institute of Space Propulsion, German Aerospace Center (DLR), in Lampoldshausen. Fig. 1 shows a side view of the hybrid slab burner, that was adapted from a combustion chamber set-up used in previous studies to investigate the combustion behavior of solid fuel ramjets [5]. The optically accessible combustion chamber has a rectangular cross-section that is 450 mm long, 150 mm wide and 90 mm high. The pre-chamber with the flow straighteners has a length of 450 mm, while the post chamber is 150 mm long. The selected oxidizer is gaseous oxygen. Before entering the combustion chamber, it passes through two flow straighteners that ensure homogeneous flow conditions in the burner. The mass flow rate is adjusted by a flow control valve and is measured with a Coriolis flow meter. Furthermore, an oxygen/hydrogen torch igniter is positioned at the left bottom of the chamber. At the end of the combustion, when the main oxidizer valve is closed, nitrogen is used for purging. No nozzle is present at the end of the post-chamber, therefore all the tests performed with this set-up are at atmospheric pressure. All tests are run automatically by the test bench control system, that also triggers




Fig. 1. Side view of the atmospheric combustion chamber set-up at DLR. Source: Adapted from [7].

the high-speed camera. A more detailed description of the test bench and test settings can be found in [6].

The second test bench exploited in this research is the MOUETTE (Moteur OptiqUe pour Étudier et Tester Ergols hybrides) slab burner, developed at the Aero-Thermo-Mechanics department of the Université libre de Bruxelles (ULB), in Brussels. The test chamber has cylindrical cross section with two parallel optical accesses. The test section has a maximum envelope of about 120 mm length, 75 mm width and 50 mm height. The oxidizer is again gaseous oxygen, and the mass flow rate can be adjusted using a choking orifice in the feed lines and by varying the pressure of the gas fed from the reservoir. A pyrotechnic squib is used for the ignition of the solid fuel slab, and the pressure in the combustion chamber is maintained using a nozzle with a graphite throat insert. These design solutions allow the test bench to operate from atmospheric conditions up to 15 bar pressure in the chamber, to investigate the effect of pressure on the fuel combustion. A photograph of the test bench during a test is shown in Fig. 2, while a detailed description of the design rationale can be found in [8].

The first set of combustion tests analyzed in this study was performed at the German Aerospace Center (DLR). Three combustion tests, selected from the tests performed in [10], have been analyzed to investigate liquid layer combustion instabilities. The test matrix is presented in Table 1, respecting the original enumeration. All the tests performed with this set-up were run at atmospheric combustion pressure and with an oxidizer mass flow ranging from 50 to 100 g/s. In this way, the influence of the oxidizer flow speed on the break-up process of the fuel melt layer and, therefore, on the regression rate could be investigated. Combustion tests were performed using a single-slab paraffin-based fuel with a 30° forward facing ramp angle (see Fig. 3), in combination with gaseous oxygen. For this analysis, only tests with plain paraffin were considered in order to maximize the regression rate and being able to clearly distinguish the regression of the fuel surface. Indeed, the regression rate of the higher viscosity fuel is so low that the height and, therefore, pixel variation of the fuel surface over the entire test duration is below two pixels and, consequently, measuring the variation of the regression rate over time is hardly possible with this optical approach. All fuel slabs, produced and machined according to the same procedure, were 200 mm long, 100 mm wide and 20 mm high. Burning time was 3 s for each test. For video data acquisition a Photron Fastcam SA 1.1 high speed video camera with a resolution of $N = 1024 \times 336$ pixels was used. Due to the high frame rate of 10,000 frames per second all relevant combustion and flow phenomena could be captured. The shutter speed and lens aperture of the camera was adjusted for each test, according to the test conditions (especially brightness) and to the position of the camera.

The second set of tests taken into account for this research comes from a test campaign conducted with the slab burner at ULB. In particular, two tests have been selected from the campaign described in [11], with different oxidizer mass flow and combustion chamber pressure. The conditions of the two tests are listed in Table 1. Both the tests have been conducted with gaseous oxygen as oxidizer and




Fig. 2. MOUETTE test.




Fig. 3. Fuel slab configuration with initial slope on the left-hand side before (top) and after (bottom) combustion test. Source: Adapted from [9].

Table 1
Test matrix of combustion tests run at DLR and ULB.

Test id.	Burn time [s]	Pressure [bar]	\dot{m}_{O_x} [g/s]	G_{O_x} [g/cm ² s]
d1r252	3	0.98	85	1.60
d1r253	3	0.98	50	0.91
d1r254	3	0.98	50	0.91
ulb04	5	1.07	26.75	1.30
ulb08	5	10.02	59.81	2.91

with a fuel slab made of pure paraffin (Tudamelt 52/54 supplied by H&R) with a forward facing ramp angle of 30°, 125 mm long, 75 mm wide and 30 mm high. The main difference between the two tests lies in the oxidizer mass flow rate (\dot{m}) and in the combustion chamber pressure. While test ulb04 is comparable to the DLR tests both in mass flux (G_{O_x}) and chamber pressure, test ulb08 has been performed at a much higher combustion pressure. This different conditions allows to assess the reliability and performance of the convolutional neural network on different test conditions. The high-speed videos have been acquired using a Photron Fastcam SA4 camera, with a resolution of 896 × 512 pixels and a frame rate of 2000 fps. The shutter speed and lens aperture have been adjusted between tests to account for the change of brightness due to the different conditions.

2.2. Convolutional neural networks for image segmentation

In order to compute the regression rate from the high-speed video the fuel slab first needs to be separated from the rest of the image so

that its height can be computed at each time-point. The task of separating a given image into distinct regions, e.g. to separate foreground from background or to identify the shape of objects in the image, is called *image segmentation* and a convolutional neural network based on the U-net architecture [12] is used to tackle this problem.

In general, artificial neural networks are sequences of basic transformations applied to some input data, e.g. an image, that can approximate complex, non-linear functions. Each basic transformation in this sequence is called a layer and the intermediate results are called neurons, features or activations depending on the context. Some of these layers may contain free parameters, called weights, that are randomly initialized and then optimized, using the backpropagation algorithm, to fit the neural network to a predefined task. This optimization procedure is called training and usually relies on a set of examples called training data. After training the neural network can be applied to new data to make predictions and this process is called inference.

The U-net architecture, as depicted in Fig. 4, is a fully convolutional neural network, i.e. its weights are almost entirely in convolutional layers. In these layers a linear sliding-window filter is applied to the input and the values of the kernel in this filter are free parameters that are automatically computed during training. Furthermore, the layers are separated into an encoder and a decoder path. In the encoder path the input image goes through a sequence of convolutional layers and downsampling layers. In the downsampling layers a 2 × 2 non-overlapping sliding window is moved over the input and for each position the maximum value inside the window is computed. This operation is called *max pooling* and it halves the resolution of the intermediate results. To make up for this, the number of convolutional

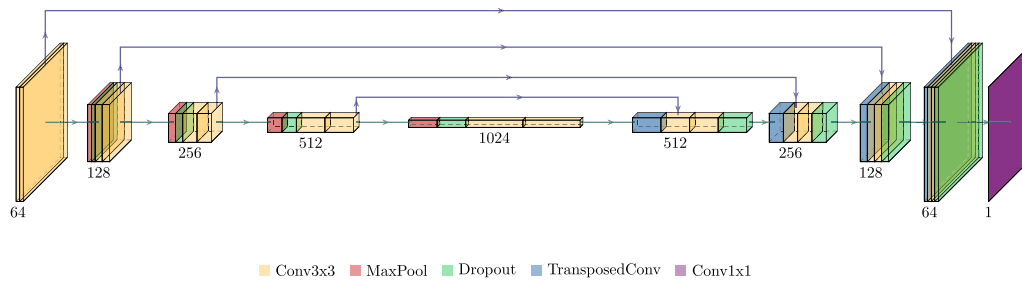


Fig. 4. U-net architecture in which the number of filters in each layer is indicated. The width and height of the input layer (left) and output layer (right) exactly corresponds with the original resolution of the combustion images. Each 3×3 convolution is followed by an instance normalization layer and a LeakyReLU activation function.

filters is doubled after each downsampling layer. The results of the encoder path, that is the final output and the intermediate results from just before each downsampling step, are passed to the decoder path. Starting from the results at the lowest resolution, the features from the encoder are increased in resolution and decreased in depth using transposed convolutions. The output is concatenated with the encoder features at this new resolution (“skip connection”) followed by a block of convolutional layers. This pattern is repeated until arriving at the original resolution of the input. Finally, the last layer of the model is a 1×1 convolution followed by a sigmoid function that maps the pixelwise features to a single probability per pixel. By choosing a threshold one can then generate a binary mask segmenting the slab. The convolutional layers except for the very last one use 3×3 filters and “same” padding, that is their input is padded with zeros so that the output does not lose resolution. Moreover, they are each followed by instance normalization, as proposed in [13], and a LeakyReLU non-linearity (with negative slope of 0.01).

As Fig. 4 indicates, dropout is included at several places in the architecture. Dropout is a common building block in neural networks in which some of the results of a layer of a neural network are dropped, i.e. their value is set to zero, with a certain probability. Initially proposed in [14], this operation is usually used to avoid weight configurations during training that rely on only a few neurons to make their prediction which is a sign of overfitting on the training data. During inference, i.e. when making predictions for new data, this random zeroing of neurons is usually turned off and the results are scaled appropriately to match their mean value during training. Monte Carlo dropout introduced in [15] is a modification of this technique, where this random aspect is kept active during inference. Consequently, repeated evaluations of the same model on the same data can lead to slightly different results. These differences can be used to assess the uncertainty of the prediction.

2.3. Training details

To train a neural network on the task of segmenting the shape of the fuel slab in the frames of the high-speed video, some training examples have to first be created manually. Fig. 6 depicts a typical mask. Fig. 6(a) also illustrates the three different positions along the slab that are analyzed in this study. Since the shape of the slab changes relatively slowly compared to the frame-rate of the high-speed cameras, each mask is used for a small range of frames from the video. Specifically, each mask is used for an interval of 11 frames, corresponding to a time interval of 1.1 ms and 5.5 ms, respectively. Table 2 shows the total number of analyzed frames and resulting number of labeled data points that are available for each test. The time-window available for the analysis is shorter than the total burn time (see Table 1). This is due to the fact that, during the first milliseconds of the ignition transient, the flame is not yet visible and fully developed along the fuel slab length, therefore the fuel slab profile is not yet detectable. The same applies to the shut-down transient.

This modified version of the U-net architecture [12] described in Section 2.2 is implemented in Pytorch for this study and trained to

Table 2
Number of analyzed frames and manually created masks for each test.

Test id.	Analyzed time window [s]	No. of frames	No. of masks	No. of labeled frames
d1r252	2.62	26 200	12	132
d1r253	2.78	27 800	16	176
d1r254	2.88	28 800	12	132
ulb04	3.61	7223	8	88
ulb08	3.40	6800	8	88

perform the task of segmenting the fuel slab using these manually created masks. For all experiments the AdamW optimizer, as first proposed in [16], is used with a starting learning rate of 3×10^{-4} and default values for all other hyperparameters. Furthermore, the learning rate is scheduled with cosine annealing, cf. [17], without restarts or warmup. The dropout rate of the dropout layers indicated in Fig. 4 is fixed at 0.5 for all experiments. Each model is trained for 10 000 steps with a batch size of 16. The training objective is defined as the sum of binary cross entropy and dice loss, which was first used as a loss function for training segmentation models in [18].

During training the input images are first cropped to a region around the starting shape of the slab so that the top of the slab is roughly in the center of the frame. Then, the images are randomly augmented to artificially increase the size of the training data. Two types of data augmentation are used. The first type are affine transformations, i.e. rotation, translation, scaling or shearing. The parameters of these operations, e.g. the degree of rotation, are sampled randomly and the manually created masks are transformed in the same way to preserve correctness of the segmentation. This type of augmentation is used to improve the model robustness with respect to small variations in the shape and positioning of the slab in the images. The second type of data augmentation is based on plasma fractals, which are fractal shapes that can be efficiently computed with the so called diamond-square algorithm. Using this algorithm for data augmentation was first proposed in [19]. For a given input image a random fractal shape is generated and added to the image as a bright area, cf. Fig. 5 for an example of the resulting effect. This emulates visual obstructions by flames forming at the side of the slab, which is a common point of failure for optical measurement of the height of the slab. The type of augmentation applied to an input image is sampled randomly and for both transformations the implementation in the software package Kornia [20] is used.

2.4. Computation of regression rate

After training the neural network, the resulting model is applied to all frames to predict a segmentation of the slab. For each column of the segmented frame the number of pixels classified as part of the slab is counted, i.e. the summation is in direction of the frame height. This leads to a height measurement for each column of each frame. A mean filter with a window size of 101 time points is applied in the time dimension to reduce variance. The applied filter has a negligible





Fig. 5. Example of data augmentation using plasma fractals.



(a) Original frame 21245.



(b) Manual segmentation of frame 21245.



(c) U-net segmentation of frame 21245.



(d) U-net uncertainty map of frame 21245.

Fig. 6. Example of a frame from dlr253 with manually created binary mask and output of U-net.

impact on the accuracy of the method due to the small time-steps between each frame. Furthermore, for each column, all time-steps in which the height of the slab increases are neglected. Indeed, since the liquid layer is not visible in these tests (due to its small thickness and high flame brightness), the height increase in the fuel slab could only be due to spurious events that are not important for the sake of this analysis. Finally, the mean along the column dimension is computed (possibly only some subset of columns are considered, cf. Section 3.3 for more details) producing the final height measurement at each time-point. To compute the regression rate spline smoothing is used, based on [21], that is a piece-wise polynomial function is fitted to the height measurements with an error tolerance of a tenth of a pixel times the number of time-points. The derivative of the resulting approximation can be computed analytically which yields the regression rate. The polynomials on each sub interval in the spline approximation are cubic, so that this approach is a direct extension of the one presented in [22] but allows for explicit control on the level of approximation.

3. Results

3.1. Image segmentation

The neural network described in Section 2.2 is trained to segment the shape of the fuel slab in the frames of the high-speed videos. To quantify the accuracy of the resulting segmentation model, the following metrics are used:

1. Mean Intersection over Union (IoU): Area of the intersection between model segmentation of the slab and ground truth segmentation divided by the area of their union.
2. Average Error (AE): Error of the height of the slab computed from the model output compared to the ground truth, averaged over the length of the slab.

To train the model 75% of the labeled data is randomly selected as training data, i.e. the weights of the neural network are optimized using these examples, and the remaining 25% of data is used to validate the

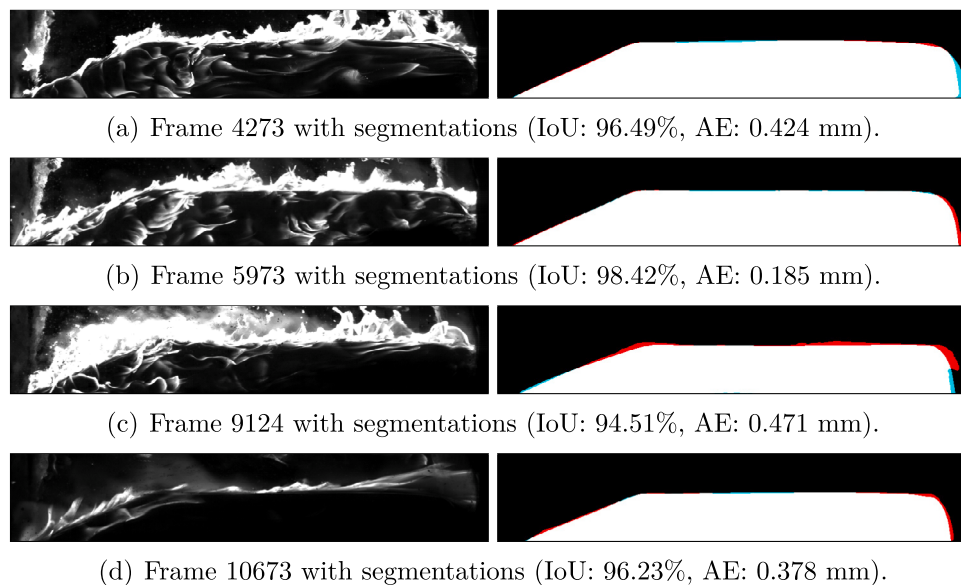


Fig. 7. Examples of frames from ulb08 and a comparison of its manual segmentation with U-Net output. White and black areas are where both U-Net and manual masks agree on fuel and background, respectively. Red areas are marked as fuel by U-Net but they are background in manual mask, and blue areas denote the reverse. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3

Results of four-fold cross-validation. Metrics are only computed on the validation data. Results are reported as mean and standard deviation across the four folds.

Test id.	IoU [%]	AE [mm]
dlr252	98.96 ± 0.18	0.137 ± 0.025
dlr253	99.12 ± 0.26	0.132 ± 0.040
dlr254	99.12 ± 0.11	0.130 ± 0.015
ulb04	98.82 ± 0.18	0.177 ± 0.025
ulb08	96.38 ± 0.19	0.405 ± 0.033

Table 4

Height error [mm] of respective method in computing the average fuel height compared to the manually traced masks.

Test id.	U-Net error	Otsu's error
dlr252	0.111 ± 0.059	0.271
dlr253	0.078 ± 0.021	0.295
dlr254	0.094 ± 0.061	0.268
ulb04	0.076 ± 0.040	0.340
ulb08	0.282 ± 0.161	0.349

model's accuracy. In order to account for variance due to the random initialization of the neural network and the random data split, four-fold cross validation is used to evaluate the accuracy of the model. That means that each dataset is randomly split into four folds and four versions of the neural network are trained, each using three of the four folds.

The results of this procedure are summarized in Table 3. The model achieves high IoU on all four folds at roughly 99% with a slight drop to 96.4% on the test ulb08. Furthermore, the average error corresponds to a discrepancy of less than one pixel compared to the manually created masks with ulb08 again being a slight outlier at an average error between two to three pixels. One reason for the larger error for ulb08 is that the end of the fuel grain is not traced correctly, as can be seen in Fig. 7, contributing to both lower IoU and higher AE. For most frames the top of the fuel grain is still traced reasonably well.

3.2. Height of the fuel slab

The height of the fuel slab can be computed from the binary masks produced by the models trained in Section 3.1. The accuracy of this method is compared with Otsu's thresholding [23], a widely used approach in computer vision. In this method, a moving average with window size 100 is first applied to every pixel along the time dimension. Otsu's thresholding is then used on a small crop of the image, centered on the top of the slab, to segment the shape of the flame on top of the slab. The position of this crop is chosen manually for each test to avoid obstruction by flames on the side of the slab as well as reflections or soot on the window. An opening filter is applied, which is used in Computer Vision to remove isolated segmented pixels, and the lowest point of the flame in the resulting segmentation is computed for each

column of the cropped image. Here it is assumed that the lower border of the flame coincides with the height of the slab though in practice the flame is not always fully attached to the slab. Finally, the height of the slab is averaged over all columns of the cropped region. Fig. 8 shows an example of this procedure being applied.

To compare the two methods, the average height in the cropped region is computed using both methods and compared to the value obtained from manual segmentations of the fuel slab. For all tests the proposed method using a U-Net to segment the shape of the fuel slab outperforms the one based on segmenting the shape of the flame using Otsu's thresholding as can be seen in Table 4, which reports the height error of the U-Net and Otsu's thresholding with respect to the manually traced masks. Results for U-Net are reported as mean and standard deviation across the 4 folds of cross-validation with 4 passes of Monte-Carlo dropout each. The reason for using 4 passes of MC dropout and 4 folds for cross-validation is that these experiments were run on a server with 4 GPUs so that 4 models can be trained or evaluated in parallel. The models are only evaluated on the data used for validation during its training. For the test ulb08 the difference in error is within one standard deviation across the folds of cross-validation from Section 3.1. Since ulb08 is the only pressurized test and the amount of labeled data for this test alone is very limited, cf. Table 2, this high variance across different data splits can be reduced by increasing the amount of labeled data for this or similar tests. Furthermore, the pixelwise moving average applied as a preprocessing step to Otsu's thresholding is well suited to smooth out the more transient flow conditions of test ulb08. Overall, the U-Net provides a more accurate measurement of the fuel height which can then be used to compute the regression rate.




Fig. 8. Close crop of the top of the fuel slab (left), result of moving average (center) and segmentation after Otsu's thresholding and opening filter (right).

Table 5

Comparison of average regression rate [mm s^{-1}] computed from U-Net segmentation over entire flat top, front (first fourth after ramp) and the ramp of the slab. Results are reported as mean and standard deviation across the 16 computed samples.

Test id.	flat top	front	ramp
d1r252	0.276 ± 0.007	0.283 ± 0.026	0.384 ± 0.028
d1r253	0.232 ± 0.025	0.257 ± 0.007	0.344 ± 0.017
d1r254	0.203 ± 0.029	0.237 ± 0.035	0.250 ± 0.003
ulb04	0.201 ± 0.006	0.234 ± 0.005	0.310 ± 0.010
ulb08	0.554 ± 0.021	0.543 ± 0.023	0.517 ± 0.007

3.3. Regression rate

The regression rate of each test is computed as described in Section 2.4. To this end, each of the four models trained via four-fold cross-validation, cf. Section 3.1, is evaluated on the datasets four times using Monte-Carlo dropout. The regression rate is then computed at several spatial positions along the length of the slab. The three positions considered are the flat top of the slab, the ramp of the slab and the front of the slab corresponding to the first fourth of the flat top right after the ramp. In total, this results in 16 possible trajectories for the regression rate for each of these three positions and Table 5 shows the mean and standard deviation across these 16 computed trajectories for the time-averaged regression rate.

For all tests, except for ulb08, the regression rate at the front is higher than the average value over the entire flat top of the slab. In some cases the differences between the mean values at the front and over the entire top may be within one standard deviation. However, these quantities are not uncorrelated as the front is a subset of the entire flat top. Indeed, a higher regression rate at the front of the slab clearly contributes to a higher regression rate when considering the entire top of the slab. Furthermore, the regression rate measured at the ramp of the slab is, again with the exception of ulb08, notably higher than that measured along the top of the slab.

From a combustion point of view, these numerical results match the trend of the experimental regression rates. From theoretical expectations, as well as from observations after the combustion, it is clear that fuels burning at atmospheric pressure regress slower than those burning at super-critical pressures (above ca. 6.5 bar [24–27]). Especially for low G_{O_x} values, the chamber pressure is strongly influencing the burning behavior and, consequently, the regression rate. This trend is clear from the numerical results: all tests performed at atmospheric pressure (d1r252, d1r253, d1r254, ulb04) show regression rate values ca. 50% lower than the test performed at super-critical pressure (ulb08) at all positions along the fuel length. Moreover, by comparing the results of the tests performed in DLR, it is observed that the tests d1r253 and d1r254 have a lower regression rate than test d1r252 at all the fuel positions. This is due to the higher oxidizer mass flow of d1r252, which, under the same operating conditions, directly leads to higher regression rate, in accordance to the hybrid combustion theory. Furthermore, for all the atmospheric tests, it is possible to notice that the ramp regresses faster than the rest of the fuel slab. The reason for this is that the flame is attached to the fuel ramp, thus creating a hotter zone at the slab front which helps the combustion and entrainment process. This trend is inverted in ulb08, where the rest of the fuel slab has a higher regression rate than the ramp. This can be explained by the fact that this pressurized test was performed

without any honeycomb grids to break the flow turbulent structures before entering the combustion chamber. This caused the appearance of longitudinal pressure waves traveling along the chamber and, thus, affecting the combustion behavior. In particular, the flame attachment point appears to periodically oscillate, thus causing the flame not to be steadily attached to the ramp, generating waves structure that detach from the burning surface, reducing the regression rate of the ramp. This can be observed close to the leading edge of the slab in Fig. 9. For what concerns the top surface of the fuel slab, from the segmentation results, it is possible to observe that, for all the atmospheric tests, the first fourth after the ramp (front) always regresses slightly faster than the whole fuel top surface. This is related to the typical boundary layer combustion process of hybrid rocket engines: by moving to the back of the fuel slab, the flame located in the boundary layer moves away from the fuel surface, thus slowing down the thermo-dynamics processes driving and controlling the combustion. The differences are mainly within one standard deviation, meaning that the observed effect is relatively small, though one has to keep in mind that the error for the two regression rate measurements will not be uncorrelated, i.e. a higher measurement for the regression rate at the front will contribute to a higher measurement for the entire top. Furthermore, this effect can be observed for all atmospheric tests, which reduces the chance of it being caused by numerical errors.

For what concerns the evolution in time, it is expected that the maximum regression rate is achieved towards the middle of the combustion process, when the oxidizer mass flow reaches its target value. This is true for tests d1r254, ulb04 (at least for the top surface) and ulb08 (with the exception of the ramp), cf. Fig. 10. The different trend in test d1r252 and d1r253 could be due to a slightly longer ignition transient (observed from test data and combustion videos, see Table 2), which delayed the development of the flame along the whole fuel surface, thus resulting in lower regression rates. Regarding test d1r252, the different numerical results (in particular for the slab front) are likely an artifact of the data processing. Since time-steps with an increase in the height of the slab are discarded in a preprocessing step, if the shape of the slab is underestimated in an early time-step this can lead to a relatively sharp drop in height at the beginning followed by a longer flat region in the resulting height curve. The corresponding regression rate will follow a trend like the one for the front portion of d1r252. Additionally, the regression rate at the front is more susceptible to this type of error because it considers a significantly smaller region of the image. As for the curves in ulb04 and ulb08 that reach the maximum regression rate at the end of combustion, a possible explanation could be found in the test sequence. In particular, the nitrogen purge at the end of combustion causes a sudden expulsion of melt paraffin from the fuel surface (in particular from the ramp and the front of the slab), thus resulting in a fake increase in regression rate. This is pretty clear in test ulb04. For test ulb08, it must be also considered that the visualization and contour detection of the fuel slab, in particular of the ramp, was not as good as for the other tests, cf. Table 3 and Fig. 7. Therefore, numerical error cannot entirely be ruled out as the cause for the observed qualitative differences in the results for ulb08.

4. Conclusion

This study shows that a convolutional neural network can be used to detect the geometry of a solid fuel grain in the frames of a high-speed




Fig. 9. Frame acquired from test ulb08.

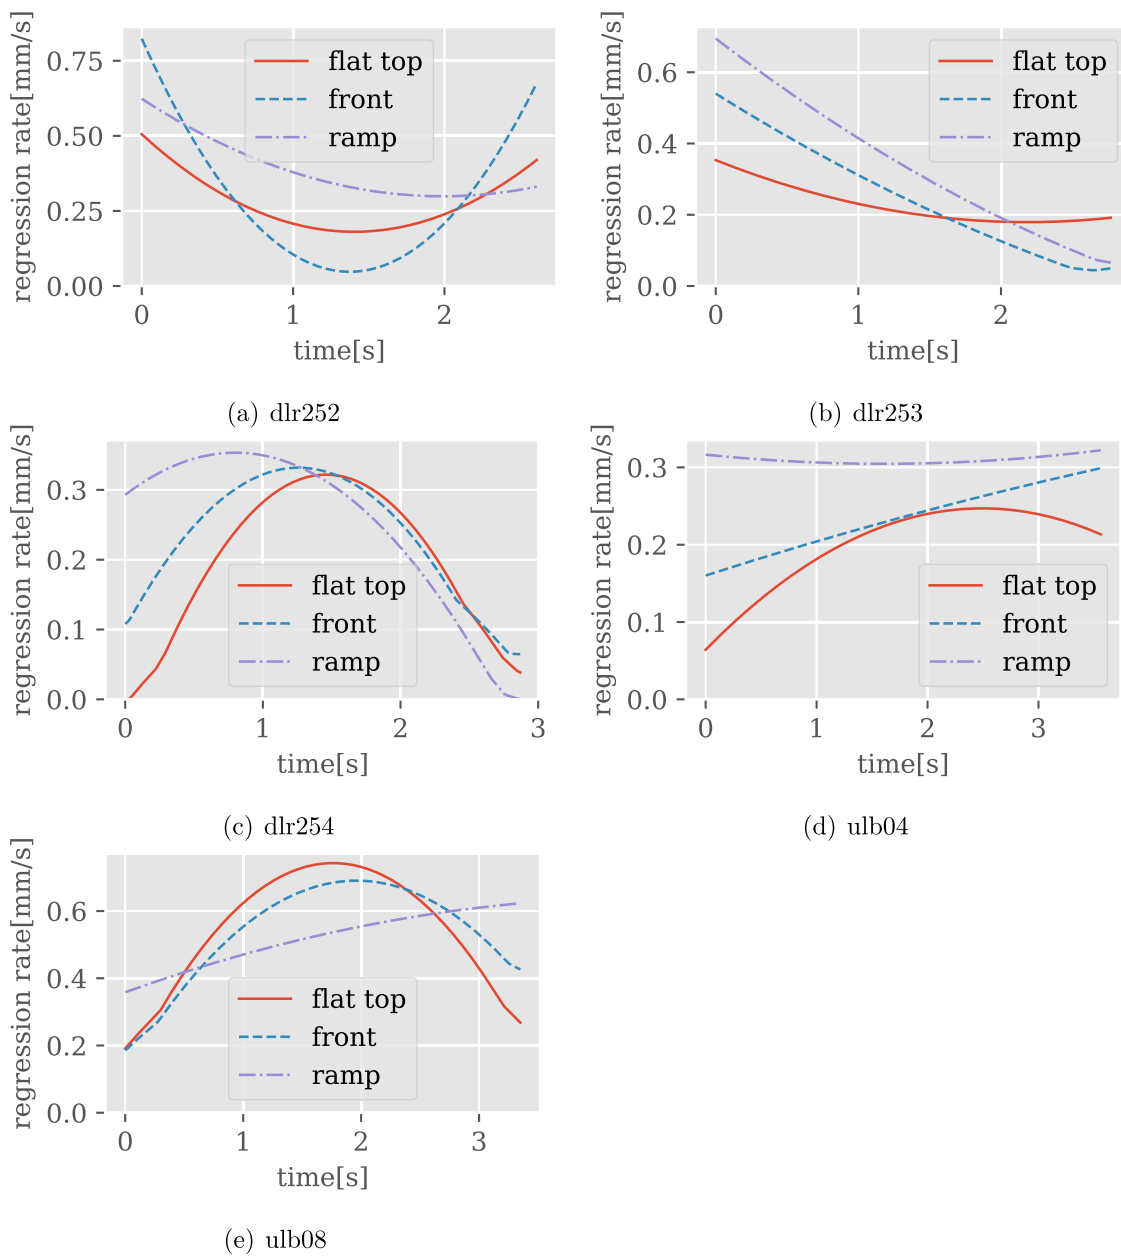


Fig. 10. Comparison of regression rate over time at the entire flat top, the front and the ramp of the slab.

video of the combustion in a hybrid slab burner. This is demonstrated on five tests that were performed at two different test facilities under varying test conditions, specifically with different oxidizer mass flow rates and chamber pressures. The regression rate can be computed from the detected geometry as a space- and time-dependent quantity, allowing for better understanding of the combustion mechanics of solid fuels for hybrid propulsion systems. The regression rate calculated with the neural network increases with increasing oxidizer mass flow rate and increasing pressure, all others factors being equal, as observed experimentally and as predicted by hybrid combustion theory. It is observed for all tests executed at atmospheric pressure that the regression rate varies along the length of the fuel grain and is on average higher at the first quarter of the flat top of the slab than at the rest of the flat top and it is again higher at the forward-facing ramp. This effect is not observed for the pressurized test and several explanations are presented, although, due to the difficult nature of analyzing the videos from the pressurized tests, more research is required to confidently explain this observation. Finally, the variation of the regression rate over time has been analyzed and the observed trends are explained qualitatively in connection with the different phases of the experiment.

As a final result, this research demonstrates how a neural network can be applied to the analysis of large datasets of combustion images acquired through high-speed video techniques in hybrid rocket slab burners, testing the capability of detecting the fuel contour not only at different oxidizer mass flow rates and combustion chamber pressure, but also with different test facilities. Moreover, the ability to automatically analyze thousands of images and evaluate the space and time estimation of the regression rate proves to be a valuable tool to support the experimental research on hybrid rocket combustion.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

R. Gelain received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 860956. The research was partly carried out under the project Antriebstechnologien und Komponenten für Trägersysteme (ATEK) by the German Aerospace Center (DLR).

References

- [1] M.A. Karabeyoglu, D. Altman, B.J. Cantwell, Combustion of liquefying hybrid propellants: Part 1, general theory, *J. Propuls. Power* 18 (3) (2002) <http://dx.doi.org/10.2514/2.5975>.
- [2] M.A. Karabeyoglu, B.J. Cantwell, Combustion of liquefying hybrid propellants: Part 2, stability of liquid films, *J. Propuls. Power* 18 (3) (2002) <http://dx.doi.org/10.2514/2.5976>.
- [3] M.J. Chiaverini, N. Serin, D.K. Johnson, Y.-C. Lu, K.K. Kuo, G.A. Risha, Regression rate behavior of hybrid rocket solid fuels, *J. Propuls. Power* 16 (1) (2000) 125–132, <http://dx.doi.org/10.2514/2.5541>.
- [4] G. Surina, G. Georgalis, S.S. Aphale, A. Patra, P.E. DesJardin, Measurement of hybrid rocket solid fuel regression rate for a slab burner using deep learning, *Acta Astronaut.* 190 (2022) 160–175, <http://dx.doi.org/10.1016/j.actaastro.2021.09.046>.
- [5] H.K. Ciezki, J. Sender, W. Clauß, A. Feinauer, A. Thumann, Combustion of solid-fuel slabs containing boron particles in step combustor, *J. Propuls. Power* 19 (6) (2003) 1180–1191, <http://dx.doi.org/10.2514/2.6938>.
- [6] M. Kobald, A. Petrarolo, S. Schleichtriem, Combustion visualization and characterization of liquefying hybrid rocket fuels, in: 51st AIAA/SAE/ASEE Jt. Propuls. Conf., American Institute of Aeronautics and Astronautics, 2015, pp. 1–24, <http://dx.doi.org/10.2514/6.2015-4137>.
- [7] A. Thumann, H.K. Ciezki, Comparison of PIV and Colour-Schlieren measurements of the combustion process of boron particle containing solid fuel slabs in a rearward facing step combustor, *Int. J. Energ. Mater. Chem. Prop.* 5 (2002) 742–752, <http://dx.doi.org/10.1615/IntJEnergeticMaterialsChemProp.v5.i1-6.770>.

- [8] R. Gelain, F. Angeloni, A.E. De Morais Bertoldi, P. Hendrick, Design and commissioning of the MOUETTE hybrid rocket slab burner, in: 9th Eur. Conf. for Aeronaut. and Space Sci., EUCASS, 2022, <http://dx.doi.org/10.13009/EUCASS2022-6055>.
- [9] A. Rüttgers, A. Petrarolo, M. Kobald, Clustering of paraffin-based hybrid rocket fuels combustion data, *Exp. Fluids* 61 (1) (2020) 1–17, <http://dx.doi.org/10.1007/s00348-019-2837-8>.
- [10] A. Petrarolo, *Liquid Layer Combustion Instabilities in Paraffin-Based Hybrid Rocket Fuels* (Ph.D. thesis), Universität Stuttgart, 2020.
- [11] R. Gelain, A.E. De Morais Bertoldi, P. Hendrick, M. Lefebvre, Optical investigation of paraffin-based fuel combustion in a hybrid rocket slab burner, in: *Aerosp. Eur. Conf.* 2023, Jt. 10th EUASS 9th CEAS Conf., 2023, <http://dx.doi.org/10.13009/EUCASS2023-041>.
- [12] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *Med. Image Comp. and Comp.-Assist. Interv.*, MICCAI 2015, 2015, pp. 234–241, http://dx.doi.org/10.1007/978-3-319-24574-4_28.
- [13] D. Ulyanov, A. Vedaldi, V. Lempitsky, Instance normalization: The missing ingredient for fast stylization, 2017, Preprint at <https://arxiv.org/abs/1607.08022>.
- [14] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* 15 (56) (2014) 1929–1958.
- [15] Y. Gal, Z. Ghahramani, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning, in: *Proc. of the 33rd Int. Conf. on Mach. Learn.*, in: *Proc. of Mach. Learn. Res.*, vol. 48, 2016, pp. 1050–1059.
- [16] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, in: *Int. Conf. on Learn. Represent.*, ICLR, 2019.
- [17] I. Loshchilov, F. Hutter, SGDR: Stochastic gradient descent with warm restarts, in: *Int. Conf. on Learn. Represent.*, ICLR, 2017.
- [18] F. Milletari, N. Navab, S.-A. Ahmadi, V-Net: Fully convolutional neural networks for volumetric medical image segmentation, in: 2016 Fourth. Int. Conf. on 3D Vision, 3DV, 2016, pp. 565–571, <http://dx.doi.org/10.1109/3DV.2016.79>.
- [19] A. Nicolau, V. Christlein, E. Riba, J. Shi, G. Vogeler, M. Seuret, TorMentor: Deterministic dynamic-path, data augmentations with fractals, in: *Proc. of the IEEE/CVF Conf. on Com. Vision and Pattern Recognit. (CVPR) Workshops*, 2022, pp. 2707–2711, <http://dx.doi.org/10.1109/CVPRW56347.2022.00305>.
- [20] E. Riba, D. Mishkin, D. Ponsa, E. Rublee, G. Bradski, Kornia: an open source differentiable computer vision library for PyTorch, in: *Winter Conf. on Appl. of Comp. Vision*, 2020, <http://dx.doi.org/10.1109/WACV45572.2020.9093363>.
- [21] P. Dierckx, *Curve and Surface Fitting with Splines*, Monographs on Numerical Analysis, Oxford University Press, 1993.
- [22] K. Budzinski, S.S. Aphale, E.K. Ismael, G. Surina, P.E. DesJardin, Radiation heat transfer in ablating boundary layer combustion theory used for hybrid rocket motor analysis, *Combust. Flame* 217 (2020) 248–261, <http://dx.doi.org/10.1016/j.combustflame.2020.04.011>.
- [23] R. Gelain, A. Petrarolo, O. Assenmacher, A. Bertoldi, A. Rüttgers, P. Hendrick, Estimation of regression rate from image analysis in hybrid rocket slab burners, in: 13th Int. Symposium on Special Top. in Chem. Propuls. and Energ. Mater., 2023.
- [24] J.J. Marano, G.D. Holder, General equation for correlating the thermophysical properties of n-paraffins, n-olefins, and other homologous series. 1. Formalism for developing asymptotic behavior correlations, *Ind. Eng. Chem. Res.* 36 (5) (1997) 1887–1894, <http://dx.doi.org/10.1021/ie960511n>.
- [25] J.J. Marano, G.D. Holder, General equation for correlating the thermophysical properties of n-paraffins, n-olefins, and other homologous series. 2. Asymptotic behavior correlations for PVT properties, *Ind. Eng. Chem. Res.* 36 (5) (1997) 1895–1907, <http://dx.doi.org/10.1021/ie960512f>.
- [26] J.J. Marano, G.D. Holder, A general equation for correlating the thermophysical properties of n-paraffins, n-olefins, and other homologous series. 3. Asymptotic behavior correlations for thermal and transport properties, *Ind. Eng. Chem. Res.* 36 (6) (1997) 2399–2408, <http://dx.doi.org/10.1021/ie9605138>.
- [27] A. Karabeyoglu, B. Cantwell, J. Stevens, Evaluation of the homologous series of normal alkanes as hybrid rocket fuels, in: 41st AIAA/ASME/SAE/ASEE Jt. Propuls. Conf. & Exhib., 2005, <http://dx.doi.org/10.2514/6.2005-3908>.

Glossary

Artificial Neural Network: Sequence of simple mathematical operations that can approximate complicated non-linear functions, commonly abbreviated as Neural Network

Convolutional Layer: Possible operation in a neural network applying a linear kernel as a sliding window

Dropout Layer: Possible operation in a neural network, where some of the values of the output of the previous layer are set to zero with a given probability

Training: Optimizing the free parameters of an artificial neural network for a given task

Inference: Using an artificial neural network that was optimized for a given task to make a prediction for new data

Layer: One step in an artificial neural network

Image Segmentation: Task of separating an image into discrete regions, e.g. to separate foreground from background

Data Augmentation: Processing step during training in which input data is randomly transformed to artificially increase the number of training examples

Training Data: Set of data points that are used as examples during the training of an artificial neural network

Loss Function: Objective function to be minimized during training with respect to some set of training data

Weights: Free parameters of an artificial neural network that are adjusted automatically during training



Oliver Assenmacher works as a researcher in the group “Scalable Machine Learning” of the High Performance Computing Department at the Institute for Software Technology of the German Aerospace Center (DLR). He studied mathematics at the University of Bonn in Germany, where he obtained his master’s degree in 2021. In his current work he develops algorithms for data analysis in space transport applications.



Riccardo Gelain is a Ph.D. student at the Aero-Thermo-Mechanics department of Université Libre de Bruxelles. He holds a master’s degree in aerospace engineering from University of Padua and a Second Level Master Degree in Space transportation Systems: Launchers and Re-Entry Vehicles from Sapienza University of Rome. His research focuses on experimental investigation of hybrid rocket engines combustion and performance.



Alexander Rüttgers is team lead of the research group “Scalable Machine Learning” and deputy head of the department “High-Performance Computing” at the Institute for Software Technology of the German Aerospace Center (DLR). His team conducts research on large-scale machine learning, with a focus on anomaly detection, analyzing data from Earth observation, space transport, and aero medicine. He has studied mathematics and obtained his Ph.D. in 2016 at the University of Bonn in Germany.



Anna Petrarolo studied Aerospace Engineering at Politecnico di Milano and obtained her Ph.D. from the University of Stuttgart in 2020, with a thesis about paraffin-based hybrid rocket fuels. From 2014 till 2023, she worked as a Research Engineer at the German Aerospace Center (DLR), Institute of Space Propulsion. From 2018, she also worked as a Propulsion Engineer at HyImpulse Technologies, developing a LOX/paraffin-based hybrid rocket motor. Since June 2023, she is a Senior Propulsion Engineer at Gilmour Space Technologies in Australia. She is responsible for the development of the hybrid rocket engine, which will propel the Gilmour’s orbital launcher.



Patrick Hendrick is the head of the Department Aero-Thermo-Mechanics at the Faculty of Engineering at Université libre de Bruxelles and Chairman of the Belgian Energy Research Alliance (BERA), which is part of the European Energy Research Alliance (EERA). He is heavily involved in research projects and activities related to Hybrid Rocket Engines and Hydrogen-Fueled Gas Turbine Engines mixing experimental activities and numerical activities on both topics.