Semantic Image Segmentation of Hybrid Rocket Fuel Combustion Data using Convolutional Neural Networks

Oliver Assenmacher* and Alexander Rüttgers†
German Aerospace Center (DLR), Institute for Software Technology, High-Performance Computing Department, Linder Hoehe, 51147 Cologne, Germany

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

Riccardo Gelain§
Université Libre de Bruxelles, Aero-Thermo-Mechanics Department, Avenue F.D. Roosevelt, 50, 1050 Brussels, Belgium

Semantic image segmentation using a convolutional neural network was applied to image data of hybrid rocket combustion tests to accurately compute the fuel regression rate over time. Combustion tests with different paraffin-based fuels have been performed at the German Aerospace Center (DLR) and have been captured with a high-speed video camera leading to large image datasets. The main task to allow for the further experimental evaluation with an optical approach is to create binary masks of the solid fuel. For this purpose, a neural network model to segment 120,000 images is presented and is justified by a thorough analysis. This analysis includes the generalization capabilities of the neural network to new image data and an analysis of the model uncertainty. As a result, time-dependent regression rates are computed for the combustion tests over a sequence of different spatial positions. This allows for a detailed time-dependent and spatial comparison of the different experimental configurations and gives valuable insights into phenomena that appear during combustion.

I. Introduction

Hybrid rockets are a promising but complex propulsion technology. First, they allow for a cost reduction like solid rocket engines due to the less complex and therefore cheaper design. Second, similar to liquid rocket engines they allow for controllable thrust, including shut off and restart capability. Furthermore, since the propellants are stored in two different states of matter, hybrid rockets are also safer than solid rockets.

However, hybrid systems that use conventional polymeric fuels often achieve relatively low regression rate performance which results in a low thrust level. One viable solution are liquefying hybrid rocket fuels, for instance, on a paraffin-basis that are considered in the following. For a better understanding of the combustion mechanics of 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 a hybrid rocket slab burner. Using a high-speed camera to visualize the combustion process, it is possible to capture the essential flow phenomena. Using this large image dataset, relevant secondary quantities of interest such as a time dependent regression rate can be derived. In the following, we focus on an optical regression rate measurement for the combustion tests performed at the German Aerospace Center (DLR).

The typical approach to obtain an averaged regression rate of a fuel grain of known geometry is to determine the mass of the grain before and after the combustion test and divide this by the combustion time. However, for a more detailed analysis of the regression rate, there are several viable options to obtain time- and space-dependent quantities, for instance, methods based on x-ray radiography as discussed in [1]. Thanks to the capabilities given by the optical access of the slab burner, we employ an optical approach as it delivers accurate geometric information of the solid fuel over the observation time, as long as the flow can be accurately separated from the fuel, and it delivers a high

*Researcher, oliver.assenmacher@dlr.de.
†Team Lead "Scalable Machine Learning", alexander.ruettgers@dlr.de.
‡Postdoctoral Researcher, anna.petrarolo@dlr.de, Member AIAA.
§Ph.D. Candidate, riccardo.gelain@ulb.be

temporal resolution (due to a high-speed camera). The main task using an optical approach is to separate the flow field and the image background from the solid fuel. For this purpose, a so-called U-net, basically a convolutional neural network architecture, is used for semantic image segmentation, i.e. applying a binary mask where "1" indicates the solid fuel and zero otherwise. Using the binary masks generated by the U-net, the regression rate can be accurately computed by comparing the number of segmented pixels over time. A similar approach is discussed in [2] for a slab burner experiment for a paraffin wax-gaseous hybrid combustion. However, the results presented in the following differ in several aspects: First, the high-speed videos have been acquired at a frame rate of 10,000 fps. For an average test of three seconds duration, this translates to 30,000 images to be analyzed per combustion test. In this research, four tests have been analyzed, for a total of 120,000 frames. The dataset is much larger in size so that the aspect of generalization (i.e. using the neural network model to new images it was not trained on before) becomes more relevant. Second, since the high frame rate of the camera does not allow for using a high intensity flash for a correct image exposure, 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 image segmentation more difficult for the domain expert and for the algorithm. Third, we derive time-dependent regression rates 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 Sec. II.A. Then, in Sec. II.B the neural network model is introduced and Monte Carlo dropout, an approach to measure the uncertainty of the model prediction, is presented. Finally, Sec. III presents the preliminary results and discusses the expected results that are intended for the final manuscript.

II. Methods

A. Experimental setup

All combustion tests were performed at the test complex M11 from the Institute of Space Propulsion, German Aerospace Center (DLR), in Lampoldshausen. Figure 1 shows a side view of the hybrid rocket slab burner, that was adapted from a combustion chamber set-up used in previous studies to investigate the combustion behavior of solid fuel ramjets [3]. The optically accessible combustion chamber has a rectangular cross-section and is 450 mm long, 150 mm wide and 90 mm high. The pre-chamber with the flow straighteners is 450 mm long and the post chamber is 150 mm long.

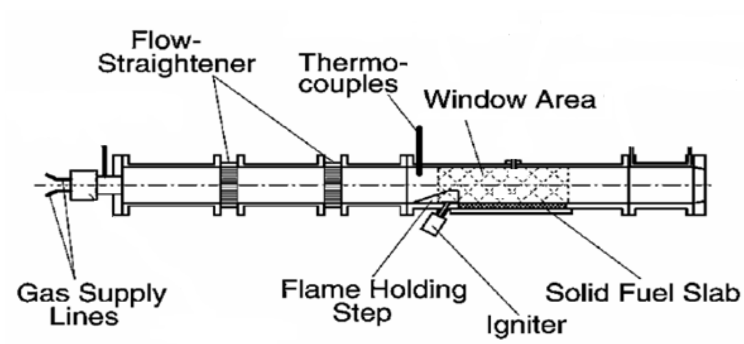

Fig. 1 Side view of the atmospheric combustion chamber set-up, adapted from Thumann and Ciezki [4].

The selected oxidizer is gaseous oxygen, which is fed in the combustion chamber from the left side of the drawing. In order to obtain homogeneous flow conditions in the burner, two flow straighteners are used. 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. Finally, at the end of the combustion, when the main oxidizer valve is closed, nitrogen is used for purging. All tests are run automatically by the test bench control system. A more detailed description of the test bench and test settings can be found in [5].

In this study, four combustion tests have been analyzed, selected from the tests performed in [7] to investigate liquid layer combustion instabilities. The test matrix is presented in Table 1, respecting the original enumeration. All the tests investigated in this study were run at an atmospheric combustion pressure and with an oxidizer mass flow ranging from 50 to 80 g/s, with a mass flux ranging from 9.1 to 14.6 $\mathrm{kg/(m^2 s)}$. Combustion tests were performed

Table 1 Test matrix of combustion tests

Test no.	Fuel		$\dot{m}_{Ox}$ [g/s]	
	6805	6805+5% polymer	50	80
203		✓	✓	
252	✓			✓
253	✓		✓	
254	✓		✓	

Fig. 2 Fuel slab configuration with initial slope on the left-hand side before (top) and after (bottom) combustion test, adapted from Rüttgers et al. [6].

using a single-slab paraffin-based fuel with a 30° forward facing ramp angle (see Fig. 2), in combination with gaseous oxygen. Two different fuel compositions were considered to allow for analyzing the effect of the fuel composition on the regression rate. First, pure paraffin 6805 from the manufacturer Sasol Wax was used. Second, 5% of a commonly available polymer was added to it in order to change the fuel viscosity. 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 seconds 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.

B. U-net neural network architecture with Monte Carlo dropout

An artificial neural network in general is a sequence of linear operations and nonlinear element-wise operations that are applied to some input data, e.g. images. Each operation in this sequence is called a layer and the intermediate results are called neurons, features or activations depending on the context. Furthermore each layer may contain some free parameters called weights. Using backpropagation these weights can be automatically adjusted to fit the neural network to some given task and this optimization is called training. 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 first proposed in [8] is a fully convolutional neural network designed for dense prediction tasks, that is for a given image it produces a per-pixel prediction. The architecture, which is depicted in Fig. 3 for this specific application, can be divided into an encoder and a decoder path. The encoder is similar to the first stages of most convolutional neural networks in that the input image goes through a sequence of convolutional layers and downsampling layers. In the convolutional layers a linear sliding-window filter is applied to the input and the values of the kernel used in this filter are free parameters automatically computed during training. In the downsampling layers a

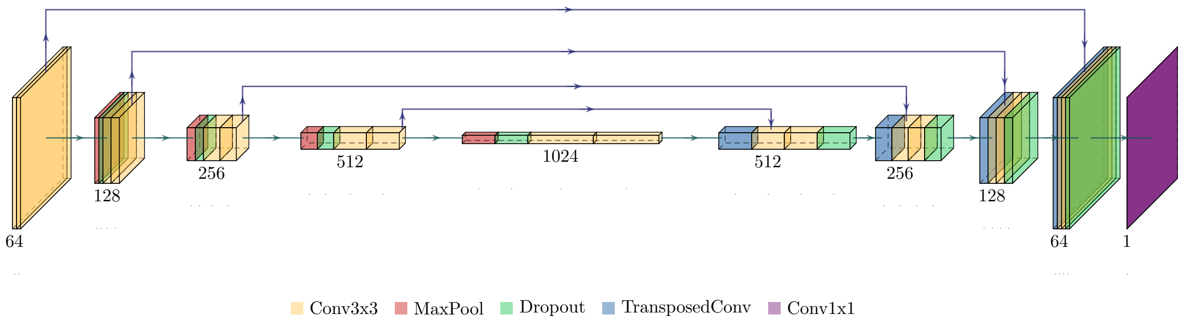

Fig. 3 U-net architecture in which the number of filters in each layer is indicated. The width and height of the input layer (left) and output layer (right) exactly corresponds with the original resolution of the combustion images. Each 3×3 convolution is followed by batch normalisation and a ReLU non-linearity.

2×2 non-overlapping sliding window is moved over the input and for each position the maximum value inside the window is computed. This downsampling halves the resolution of the intermediate results and to make up for this the number of convolutional filters is doubled after each downsampling layer. The output of the encoder path is passed to the decoder path together with intermediate results from just before each downsampling step. The decoder path uses transposed convolutions to increase the resolution of the features again while reducing their depth. The result of this upsampling and the intermediate encoder result at the same resolution are then passed to a block of convolutional layers. This pattern is repeated until arriving at the original resolution. Ultimately, 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. All convolutional layers except for the very last one use 3×3 filters and are each followed by batch normalisation, a detailed description is given in [9], and a ReLU non-linearity.

Dropout is a common building block in neural networks in which some of the neurons of in a hidden layer of a neural network are randomly dropped, i.e. their value is set to zero. This operation was initially proposed in [10] to avoid weight configurations during training that overly rely on a few neurons to make their prediction which is a sign of overfitting on the training data. After training the model this random zeroing of neurons is usually turned off and the activations are scaled appropriately to match their mean value during training. Monte Carlo dropout introduced in [11] refers to keeping this random aspect during inference. This means that repeated evaluations of the same model on the same data can lead to slightly different results. The differences between the different evaluations of the model can then be used to assess the uncertainty of the prediction.

III. Preliminary and expected results

Table 2 Average regression rate $[\mathrm{mm/s}]$ **measured by difference in mass/thickness before and after experiment**

test no.	mass	thickness
203	0.085	0.036
252	0.55	1.1
253	0.49	0.97
254	0.49	1.1

In a first step, a domain expert manually created a binary mask for a small subset of 16 images of the full dataset of test 253 containing 30,000 images (cf. Fig. 4 (b)). The 16 manual segmentations are distributed over the full range of test 253 to cover all relevant flow phases such as ignition, steady combustion and extinction phase. For the determination of the different flow phases in the dataset, we employ a clustering approach that is described in detail in [12].

Due to the usage of a high-speed video camera, the fuel slab in adjacent images in the dataset is almost identical in

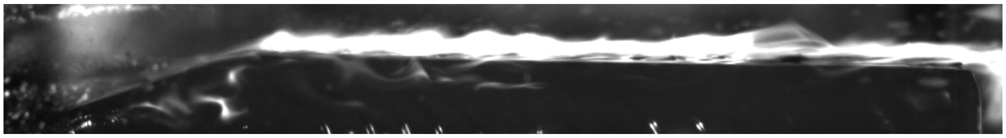

(a) image 29,365 from test 253

(b) manual segmentation by domain expert (ground truth)

(c) U-net prediction

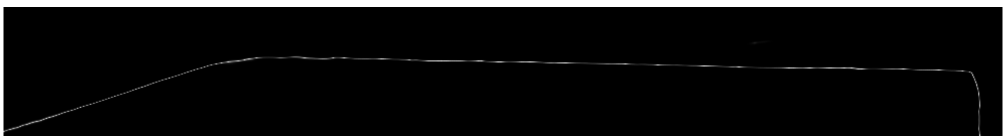

(d) uncertain pixels using Monte Carlo dropout

Fig. 4 Binary segmentation of the fuel slab in test 253 at time $t = 2.9365$**.**

size even if flow behavior varies strongly. Therefore, a manually created binary mask can be reused several times for the surrounding five neighbouring images without incurring significant additional error. As an example, the binary mask of image 29,365 for test 253 was reused for the images 29,360, . . . , 29,370. This approach has a further advantage since the neural network is forced to separate the transient flow field and the background from the fuel slab. Thus we end up with a total number of 176 datapoints, each consisting of an image and a binary mask, as reference data.

On this dataset we perform 4-fold crossvalidation, that is we split the 16 masks into 4 blocks of 4 masks each and train a U-net on 3 of these blocks corresponding to 132 images. We then evaluate the resulting model on the remaining images that were left out during training. This process is repeated for each block resulting in 4 models. To quantify the performance of the models we use the mean intersection over union (mIoU), also known as the Jaccard Index, and the mean average error of the height of the slab computed from the prediction relative to the manually created mask (MAE). There is some variation in the resulting metrics due to the relatively small size of the training data and the random initialization of the neural network but nonetheless the U-net achieves on average a mIoU of 0.99 and a MAE of 1.2% on the validation data, which confirms the potential of the U-net approach for high accuracy.

Table 3 Results for 4-fold crossvalidation on test 253

fold no.	mIoU	MAE[%]
1	0.993	1.131
2	0.987	1.600
3	0.986	1.252
4	0.995	0.737

Figure 4 (c) shows the binary mask prediction of the final neural network model of fold 1 after training for the unseen image 29,365 from the validation data set. It can be observed that the prediction closely matches the ground truth from the manual segmentation. In more detail, the U-net is able to distinguish the surface combustion flame from the fuel slab.

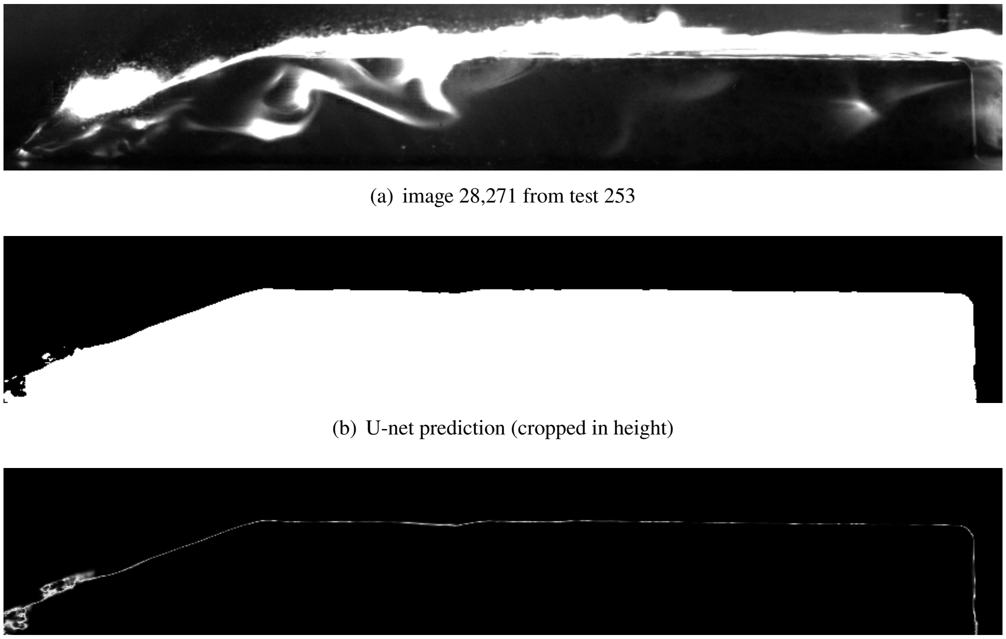

(a) image 28,271 from test 253

(b) U-net prediction (cropped in height)

(c) uncertain pixels using Monte Carlo dropout

Fig. 5 Binary segmentation of the fuel slab in test 203 at time $t = 2.8271$ **using the U-net prediction. The U-net trained with test 253 has not seen any image of test 203 in training.**

Moreover, the U-net is also able to recognize the part of the fuel slab that is behind the side flames, which are closer to the camera in the direction of view. This obstruction is a major challenge for classical filters from computer vision such as thresholding which was verified in a preliminary study that will be shown in the final paper. Therefore, this example demonstrates the advantage of machine learning image segmentation compared to classical computer vision filters.

Although the difference between the manual segmentation (Fig. 4 (b)) and the neural network prediction is small (Fig. 4 (c)) and only visible on the surface of the fuel slab, it is necessary to quantify the magnitude of the error. Since ground truth data does not exist for the majority of the dataset, a direct comparison in general is not possible. Therefore, we employ Monte Carlo dropout, explained in detail in Section II.B, as a way to determine the pixels with the highest uncertainty regarding a correct classification. Figure 4 (d) visualizes the pixel cells that were classified differently in different model predictions generated with Monte Carlo dropout. As expected, the uncertain pixels are located on the surface of the fuel slab. By identifying these uncertain pixels, one can then estimate the segmentation error.

Next, we investigate whether a U-net trained on a specific combustion such as test 253 can also be used to segment further combustion tests. The main motivation is to have a general software tool in the end that can be directly applied to new combustion tests. Table 1 lists the different tests that will be considered. Test 254 coincides with test 253 in all relevant physical parameters such as fuel type and oxidizer mass flow whereas test 252 uses a higher oxidizer mass flow and test 203 uses a different fuel composition. The average regression rate over the entire experiment was determined by measuring the difference in height/mass before and after the experiment for each test (cf. Table 2). Figure 5 shows the prediction of the binary mask for an example image in test 203. Although the U-net prediction for the binary mask is still plausible in general, a larger error is observed on the left-hand side (LHS) of the fuel slab (s. Fig. 5 (b)). This is due to the fact no image from this test was used for the training of the neural network. However, the larger error region is also visible on the LHS of the pixelwise uncertainty depicted in Fig. 5 (d). Using this uncertainty estimation as an error indicator for each image of the 30 000 images of test 203, we can determine those images with the highest expected error and segment them manually by a domain expert. Some preliminary tests indicate that adding 2-3 of manual segmentations for test 203 to the training data can already help to obtain a low overall error for this new combustion test. In the final publication, the generalization property of the U-net approach to new combustion tests such as test 203, 252 and 254 will be discussed in more detail.

As mentioned in the beginning, the main motivation for the binary segmentation of the images is to compute secondary quantities of interest such as the time-dependent regression rate. Finally, Fig. 6 (a) shows the height for each

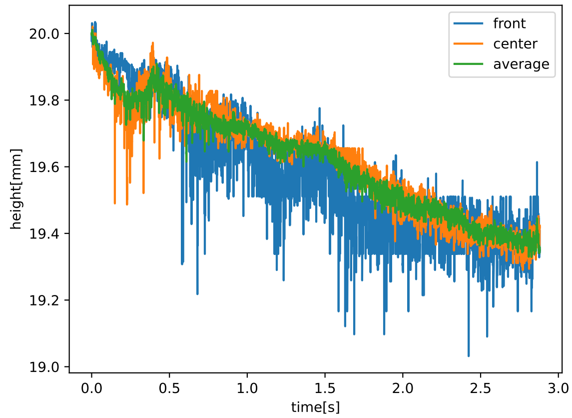

(a) Mean pixel height over time for each image

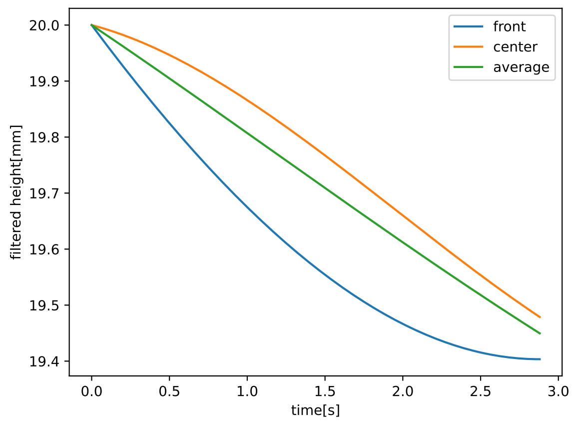

(b) Smoothed mean pixel height after filtering

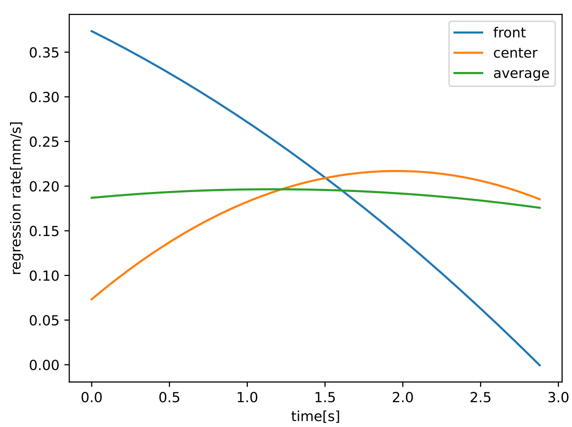

(c) Regression rate over time as primary quantity of interest

Fig. 6 Computation of the decreasing height and the regression rate in test 253 at different spatial positions.

image of test 253 at different spatial positions along the length of the block. The graphs "front" and "center" refer to the average height of a small section of the slab at the front, i.e. near the forward ramp, and in the center of the slab, respectively, whereas "average" refers to the average height along the entire length of the slab excluding the forward ramp. The plot contains several spikes or artifacts, for instance at about $t = 0.5$ s, which are caused by misclassifications of some pixels in these frames. Nonetheless a clear downward trend can be observed showing that the U-net is able to detect the regression of the fuel slab. Due to the usage of a high-speed video camera, we know that the shape of the slab cannot change dramatically between adjacent images so that some of these artifacts can also be removed in a post-processing step. Figure 6 (b) shows a smoothed plot of the pixel height using a polynomial fit. Finally, the regression rate is determined as derivative of the pixel height (Fig. 6 (c)). It can be observed that the regression near the forward ramp is largest at the beginning of the combustion and decreases over time while at the center the regression starts slower but increases until reaching a maximum at around $t = 2$ s. This coincides with our observation since the combustion flame starts at the LHS of the slab and moves further right over the course of the experiment. Indeed it covers the complete surface of the fuel slab in the steady combustion phase (between 0.8 s and 2.5 s) and only covers the right-hand side (RHS) of the fuel slab in the final extinction phase.

For the final publication, it is intended to further expand on this discussion of the time-dependent regression rates computed at different spatial positions. Using this approach, a better comparison between the different test configurations can be achieved and a more detailed analysis is possible. Furthermore, other postprocessing techniques will be investigated to avoid some of the problems that can be caused by a simple polynomial fit.

IV. Conclusion

In the presented study, a neural network was used for semantic image segmentation of hybrid rocket combustion data. Compared to classical approaches from computer vision such as thresholding, the neural network achieved better results for the specific application due its better handling of noise and complex flow phenomena. This shows the potential of an optical determination of the fuel regression rate for future combustion tests.

Acknowledgments

R. Gelain received funding from the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie grant agreement No 860956.

References

[1] Chiaverini, M. J., Serin, N., Johnson, D. K., Lu, Y.-C., Kuo, K. K., and Risha, G. A., "Regression Rate Behavior of Hybrid Rocket Solid Fuels," *Journal of Propulsion and Power*, Vol. 16, No. 1, 2000, pp. 125–132. https://doi.org/10.2514/2.5541, URL https://doi.org/10.2514/2.5541.

[2] 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 Astronautica*, Vol. 190, 2022, pp. 160–175. https://doi.org/https://doi.org/10.1016/j.actaastro.2021.09.046, URL https://www.sciencedirect.com/science/article/pii/S0094576521005324.

[3] Ciezki, H. K., Sender, J., Clauß, W., Feinauer, A., and Thumann, A., "Combustion of Solid-Fuel Slabs Containing Boron Particles in Step Combustor," *J Propuls Power*, Vol. 19, No. 6, 2003, pp. 1180–1191. https://doi.org/10.2514/2.6938.

[4] Thumann, A., and Ciezki, H. K., *Combustion of energetic materials, chap. Comparison of PIV and Colour-Schlieren measurements of the combusiton process of boron particle containing soild fuel slabs in a rearward facing step combustor*, Vol. 5, Begell House Inc., 2002. https://doi.org/10.1615/IntJEnergeticMaterialsChemProp.v5.i1-6.770.

[5] Kobald, M., Petrarolo, A., and Schlechtriem, S., "Combustion Visualization and Characterization of Liquefying Hybrid Rocket Fuels," *51st AIAA/SAE/ASEE Joint Propulsion Conference*, American Institute of Aeronautics and Astronautics, 2015, pp. 1–24. https://doi.org/10.2514/6.2015-4137.

[6] Rüttgers, A., Petrarolo, A., and Kobald, M., "Clustering of paraffin-based hybrid rocket fuels combustion data," *Exp Fluids*, Vol. 61, No. 1, 2020, pp. 1–17.

[7] Petrarolo, A., "Liquid Layer Combustion Instabilities in Paraffin-Based Hybrid Rocket Fuels," Ph.D. thesis, Universität Stuttgart, 2020.

[8] Ronneberger, O., Fischer, P., and Brox, T., "U-Net: Convolutional Networks for Biomedical Image Segmentation," *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, edited by N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Springer International Publishing, Cham, 2015, pp. 234–241.

[9] Ioffe, S., and Szegedy, C., "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," *International conference on machine learning*, 2015, pp. 448–456. URL http://jmlr.org/proceedings/papers/v37/ioffe15.pdf.

[10] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R., "Dropout: A Simple Way to Prevent Neural Networks from Overfitting," *Journal of Machine Learning Research*, Vol. 15, No. 56, 2014, pp. 1929–1958. URL http://jmlr.org/papers/v15/srivastava14a.html.

[11] Gal, Y., and Ghahramani, Z., "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," *Proceedings of The 33rd International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 48, edited by M. F. Balcan and K. Q. Weinberger, PMLR, New York, New York, USA, 2016, pp. 1050–1059. URL https://proceedings.mlr.press/v48/gal16.html.

[12] Debus, C., Ruettgers, A., Petrarolo, A., Kobald, M., and Siggel, M., "High-performance data analytics of hybrid rocket fuel combustion data using different machine learning approaches," *AIAA Scitech 2020 Forum*, 2020, pp. 1–18. https://doi.org/10.2514/6.2020-1161, URL https://arc.aiaa.org/doi/abs/10.2514/6.2020-1161.