

# Evolution of aluminium droplets in solid propellant combustion by image analysis using deep learning

Airiau Magdeleine<sup>\*†</sup>, Devillers Robin<sup>\*</sup>, Chan Hon Tong Adrien<sup>\*\*</sup>, Le Besnerais Guy<sup>\*\*</sup>

<sup>\*</sup>DMPE-MPF, ONERA, Université Paris Saclay, F-91123 - Palaiseau, France

<sup>\*\*</sup>DTIS-IVA, ONERA, Université Paris Saclay, F-91123 - Palaiseau, France

Adress

magdeleine.airiau@onera.fr · robin.devillers@onera.fr

<sup>†</sup>Corresponding author

## Abstract

The shadowgraphy technique enables visualising burning aluminium droplets during combustion in combustion gases generated by a solid-propellant sample. To analyze these images, deep learning has already proven its efficiency but requires refined detection. A new method of instance segmentation has been tested with the Mask R-CNN network, which has been adapted to be more efficient. Diameter probability densities and velocity profiles for two pressures for a research composition were estimated from the detections obtained by the network and compared. This study stands out in the existing literature as it provides a substantial amount of data, enabling a rigorous and comprehensive investigation.

## 1. Introduction

Adding aluminium particles in solid-propellant composition increases propulsion performance by about 10% but can also generate troublesome phenomena such as thermo-acoustic instabilities (ITHAC) [1–3], leading to pressure oscillations. Characterization of aluminium combustion above the propellant surface is decisive for understanding the stability of solid-rocket motor (SRM) propulsion. Numerical simulation is to study complex instabilities; however, an accurate aluminium-combustion model with realistic input data is crucial. But regardless of model precision, accurate characterization of aluminium-droplet size and velocity remains limited in realistic combustion conditions. Given the combustion conditions (high pressure, high temperature), these data are particularly complicated to obtain experimentally. ONERA has been using a shadowgraphy [4] set-up to observe small solid-propellant samples in combustion at a high repetition rate. Information-rich images contain multiple objects moving quickly in the gas flow. The complexity of the shadowgraphy images has led to analyze images by deep learning neuronal network, first by semantic segmentation [5, 6], then in this work by instance segmentation using a convolutional pre-trained network. Instance segmentation combines semantic segmentation (providing detailed object shape) and object detection (giving access to localization and classification) by segmenting images only in areas of interest (determined by the network). It gives access to information on the characteristics of each object (shape, size, etc.), their place in the image, and their relationship to other objects (location, distance to neighbours, etc.). The learning phase is conducted with a restricted base of various annotated shadowgraphy images for different experimental conditions (solid-propellant compositions, pressure, acquisition parameters, etc.). For this task, the Mask R-CNN [7], a widely adopted network for instance segmentation, has been selected.

Nevertheless, classical deep learning approaches aim at optimizing global performance, combining detection probability and false alarm rate without considering the objects' size distribution, our goal for aluminium characterization. Under these conditions, the networks tends to favour detecting the most straightforward objects (typically the largest droplets) while neglecting their counter-performance on the most challenging objects (smaller droplets), which can bias the estimated granulometric proportions. To address this issue, the learning process has been modified to enhance robustness to the size distribution, resulting in a modified network named MRCNN-A, as detailed in a separate publication [8]. The results obtained from the image analysis approach and statistical tools enable the precise determination of droplet size distribution. Furthermore, a Kalman filter has been developed as a tracking approach to capture droplet trajectories and estimate their velocity near the burning surface.

To highlight the efficiency of instance segmentation, we aim to establish a benchmark for the performance of instance segmentation in analyzing the burning behaviour of aluminium drops. By comparing our results with state-of-the-art techniques, we can assess the effectiveness and potential advantages of employing instance segmentation

## EVOLUTION OF BURNING ALUMINUM DROPS BY DEEP LEARNING

in this domain. This comparative analysis will provide insights into the capabilities of our method in capturing the intricate dynamics of burning aluminium drops, highlighting any improvements in accuracy, computational efficiency, and predictive capabilities.

Subsequently, the focus shifts towards evaluating physical quantities related to droplet size and velocity. Droplet size plays a crucial role in combustion processes, influencing the surface area-to-volume ratio and hence the rate of heat and mass transfer. Velocity, on the other hand, affects the Al-droplet burning time and mixing characteristics, ultimately shaping the overall combustion behaviour. We investigate these important parameters for two different test conditions for a single propellant composition to understand their influence on the burning characteristics of aluminium drops. By employing MRCNN-A, we can gain a comprehensive understanding of how droplet size and velocity evolve during the combustion process and unravel any underlying trends or correlations.

In summary, this paper presents an efficient image analysis approach employing deep learning techniques to characterize the evolution of aluminium droplets in solid-rocket propellants. It addresses the challenges associated with accurately determining droplet size distribution and showcases the effectiveness of the proposed methodology. Additionally, incorporating a tracking approach using a Kalman filter further enhances the analysis of droplet trajectories and velocity estimation, complementing studies already launched for tracking particles in shadowgraphy images [6,9]. The number of monitoring sessions in this work is significantly higher than in the literature, making the resulting analyses more robust.

## 2. Past work on shadowgraphy analysis

### 2.1 Shadowgraphy images

Shadowgraphy makes it possible to observe solid propellants using an optical set-up to visualize the combustion process. The setup used in this article has already been used to study solid propellants [4, 10]. Images are obtained at high frame rates for relatively small propellant samples, averaging < 1.0 g. The propellants studied are AP/HTPB composites for research. A typical shadowgraphy image is shown in Figure 1




Figure 1: Exemple of shadowgraphy image for a burning solid propellant

The image is divided into two distinct regions: a dark area at the bottom corresponds to the solid propellant sample. In contrast, the lighter area represents the hot gas from propellant combustion. The latter shows the burning aluminium droplets with the alumina cap resulting from droplet combustion and inert aluminium particles that have not yet ignited. The bottom zone of the solid propellant appears more uniform, but its surface often displays complex shapes and patterns. Numerous surface protuberances are generally associated with aluminium particle-related phenomena, such as aluminium melting, ignition, droplet agglomeration and surface ejection [11]. The main objective of the present work is to analyze and detect aluminium droplets as they burn, from the moment they leave the surface to the moment they disappear from the image. However, this task may prove difficult due to the complexity of the images resulting from the extreme conditions.

## 2.2 Image analyses for shadowgraphy images

Different image analysis techniques have been tested on shadowgraphy images to detect aluminium drops. One of these techniques is the MSER (Maximally Stable Extremal Regions) algorithm, as described in [12]. The MSER algorithm is used for blob detection and region-based object recognition. It identifies stable regions in an image with consistent intensity or color properties across different scales. This approach provides robustness to variations in the image and noise, making it suitable for detecting aluminium drops in shadowgraphy images [5, 6, 10].

Another image analysis technique applied to shadowgraphy images is semantic segmentation [6, 13]. Semantic segmentation involves dividing an image into different regions and assigning semantic labels to each pixel. The goal is to classify and understand the contents of an image at the pixel level, providing a detailed understanding of the image by assigning a label to each pixel. This technique allows for precise localization and delineation of objects or regions within the image. One popular architecture used for semantic segmentation is U-Net, as described in [14]. U-Net utilizes an encoder-decoder structure with skip connections, which helps capture local and global context information while preserving fine-grained details, leading to accurate pixel-level predictions.

These two methods have been used for droplet detection [5, 6] and confronted [13]. The network was tested in several classes for 45 training and 15 test images. This results in the detection performance being much better for semantic segmentation than MSER (70% accuracy for U-Net vs 35% for MSER). However, the algorithm tended to confuse the objects in the images, causing confusion between the different classes tested. A study on more images and fewer classes would allow us to refine the segmentation.

Several tracking methods have been developed to track particulates in shadowgraphy images. In the study [6], the velocity of droplets in the gas flow was compared with numerical simulations, while in [9], the velocity of detached droplets was compared with that of detached agglomerates in images of solid propellant loaded with inert particulates. Another approach involves coupling shadowgraphy and Al-PLIF (Aluminum-Planar Laser-Induced Fluorescence) images to enable simultaneous tracking in both types of images, resulting in improved accuracy [10]. However, these tracking methods are limited in their ability to track many objects due to their detection performance (false detection by droplet identification methods).

## 3. Instance segmentation for droplet detection

This section presents another segmentation method for drop detection, instance segmentation, and the neural network chosen to perform it. It also presents our modification to this network to make it more robust to the physical constraints of aluminium drops burning in a solid propellant.

### 3.1 Instance segmentation

Coupling object detection and semantic segmentation, instance segmentation combines the advantages of both methods. It performs segmentation in bounding boxes providing detailed results but only on the targeted detection areas. It then gives access to information on the characteristics of each object (shape, size, etc.), their place in the image, and their relationship to other objects (location, distances to neighbours, etc.). In our case, the key advantage of instance segmentation is its ability to distinguish individual droplets in cases where they overlap or are closely packed together. By assigning a unique label and generating a separate mask for each droplet, instance segmentation allows for a granular understanding of their characteristics and trajectories during the combustion.

### 3.2 Mask R-CNN

To perform instance segmentation of the image, we used the Mask R-CNN [7], which stands for Mask Region-based Convolutional Neural Network, a deep learning model developed, for instance segmentation tasks which gives an accurate mask for each detected object within an image. This means that in addition to identifying objects within an image, it can precisely outline those objects' boundaries, providing detailed segmentation. However, classical deep learning approaches aim to optimize global performance metrics, combining detection probability and false alarm rate without considering objects' size and granulometric distribution. Under these conditions, the networks might favour detecting the easiest objects (typically the largest drops) while neglecting their counter-performance on the most difficult objects (small drops), which can bias the estimated granulometric proportions. The learning process has therefore been modified to make it more robust to the size distribution. The goal is to accept a possible less efficient detection of the network but which has a physical meaning closer to the ground truth for the droplet size distribution, our goal. Until now, the learning process was focused on three tasks: detection, segmentation, and classification,

## EVOLUTION OF BURNING ALUMINUM DROPS BY DEEP LEARNING

followed by post-processing to evaluate the size distribution. Now this last aspect concerning granulometry is studied inside the network. For this purpose, a term has been added to the loss function of the network measuring the deviation between the size distributions of the ground truth and the prediction. This deviation is estimated by the Kullback-Leibler divergence [15] (explained in more detail later in the study). A computer vision-oriented article [8] has been published for more details on the method that we called MRCCN-A. The method is explained in depth in the other publication where it is applied to other dataset

#### 4. Image postprocessing for recovering physical parameters

This section overviews the image analysis tools employed to derive accurate physical parameters. It introduces techniques for evaluating the accuracy of drop detections and tools for measuring the physical quantities related to the size of aluminium drops based on these detections. The section concludes by discussing the tracking methods utilized to monitor the motion of the drops within the gas flow.

##### 4.1 Detection-performance evaluation

The study primarily focuses on the combustion of aluminium drops, distinguishing between two classes on images: the background and the drops themselves. The annotation of images through manual labelling provides a ground truth (GT), which serves to train and evaluate the learning methods of the network.

Segmentation into these two classes can be likened to a classification task, determining the presence or absence of the object of interest. As such, relative classification metrics can be utilized. The GT provides a spatial map of the object locations, enabling a comparison with the predictions. When an object is correctly located in the prediction image, it is classified as a True Positive (TP). If an object is erroneously detected where none exists, it is a False Positive (FP). Conversely, if an object in the GT is undetected, it is a False Negative (FN).

To assess detection performance, Precision and Recall serve as evaluation criteria. Precision measures the ratio of correctly detected relevant objects of interest to all detected objects of interest, evaluating the network's ability to detect the desired objects accurately:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (1)$$

On the other hand, Recall examines the ratio of correctly detected relevant objects of interest to the total number of relevant objects present in the image, evaluating the completeness of the detection:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (2)$$

In practice, it is important to balance Precision and Recall. If the network considers all objects in the image as objects of interest, it may achieve high Recall but have low Precision. Conversely, if it only detects a few objects of interest among the total, Precision may be high but Recall low. Typically, Recall is plotted against Precision to gain insight into the metrics, but relying solely on these criteria is insufficient for network evaluation. To capture the trade-off between Precision and Recall, a combined measure is computed, known as the F-Score. It provides a comprehensive assessment of Precision and Recall by calculating their harmonic mean:

$$F = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3)$$

A higher F-Score indicates better detection quality of the network.

##### 4.2 Droplet size-performance evaluation

###### 4.2.1 Probability density function with Kernel Density Estimation

To estimate the probability density function (PDF) of the droplet's diameter, we employ Kernel Density Estimation (KDE) [16, 17]. This non-parametric technique utilizes observed data points to smooth the data and derive an estimated density. It achieves this by placing a kernel function on each data point and summing up these kernel functions to obtain the density estimate.

Let  $x_1, \dots, x_n$  be an independent and identically distributed sample (of  $n \in \mathbb{Z}$  objects) from a random variable  $X$  with density  $f$ . The density estimator is given by:

$$\hat{f}_h(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x-x_i}{h}\right) \quad (4)$$

Here,  $K$  represents a kernel function that follows a probability law density, and  $h$  is a smoothing parameter. The Kernel Density Estimation technique enables flexible and smooth estimation of the underlying probability density function. It provides valuable insights into the data distribution, even in cases where the underlying distribution is unknown or complex.

#### 4.2.2 Kullback-Leibler divergence

The Kullback-Leibler divergence [15], also known as KL divergence, is utilized to quantify the dissimilarity between two probability distributions. While it does not conform to the properties of a distance measure, such as triangular inequality or symmetry, it effectively assesses the difference between probability distributions.

Consider a dataset  $P = (p_1, \dots, p_n)$  comprising  $n \in \mathbb{N}$  observations, and let  $Q = (q_1, \dots, q_n)$  be an approximation of  $P$  with the same size. The Kullback-Leibler divergence of  $P$  with respect to  $Q$  at a specific point  $x$  is computed using the following equation:

$$KL(P||Q) = \sum_{i=1}^n p_i(x) \log\left(\frac{p_i(x)}{q_i(x)}\right) \quad (5)$$

This divergence is a valuable tool for comparing probability distributions of measured and reference sizes during the learning process.

#### 4.3 Droplet tracking

Droplet tracking was set up using Multi-object tracking (MOT) to establish droplet velocity profiles as the combustion proceeds. It is a computer vision task that detects and tracks multiple objects in successive images. It aims to follow object trajectories despite occlusions and appearance changes accurately. The main steps in MOT include object detection, data association to establish correspondences between detections, trajectory estimation, occlusion handling, track maintenance, and performance metrics evaluation. In our case, object detection is performed by MRCNN-A, trajectory estimation by Kalman Filter, and the other task by post-processing.

The Kalman algorithm is a recursive filtering algorithm used for estimating the state of a dynamic system in the presence of noise and uncertainty. It is widely applied in various fields, including robotics, control systems, and tracking. The Kalman Filter has key features that include optimal state estimation based on noisy measurements, updating and refining the state using a weighted average of predicted state and measurements, estimating the uncertainty with an error covariance matrix, and achieving optimal filtering by dynamically adjusting weights based on uncertainties. The Kalman Filter is particularly useful in situations with measurement noise and uncertainty, allowing for accurate state estimation even in the presence of these disturbances, as in our study.

### 5. Results

This section presents a comparative analysis of deep learning techniques with state-of-the-art methods in studying burning aluminium drops. Furthermore, it explores the evaluation of droplet size and velocity for two different test conditions, shedding light on their impact on combustion behaviour. The physical quantities obtained were normalized to reference values (diameter  $d_{ref}$ , velocity  $v_{ref}$ , height above sample  $H_{ref}$ ). The study is performed on annotated shadowgraphy images of several types of propellants at different test conditions (pressure, image resolution, etc.).

#### 5.1 MRCNN vs U-Net

##### 5.1.1 Repeatability of the networks

To study the repeatability of the networks, we consider learning as a random variable  $X$ . For each network, ten random draws of  $X$  are performed. Each drawing consists of the random separation of the database into 45 images (more than 300 droplets) for learning and 15 (more than 100 droplets) for testing. The learning conditions being the same, the draws are independent: they effectively lead to 10 different learning processes. The learning conditions for U-Net are identical to those implemented in the previous studies [13]. The comparison of the instance segmentation with the state of the art is based solely on the method's performance. The Precision as a function of Recall for the ten draws for the two networks is shown in Figure 2. The scatterplot for U-Net is much more dispersed than that for Mask R-CNN, so

## EVOLUTION OF BURNING ALUMINUM DROPS BY DEEP LEARNING

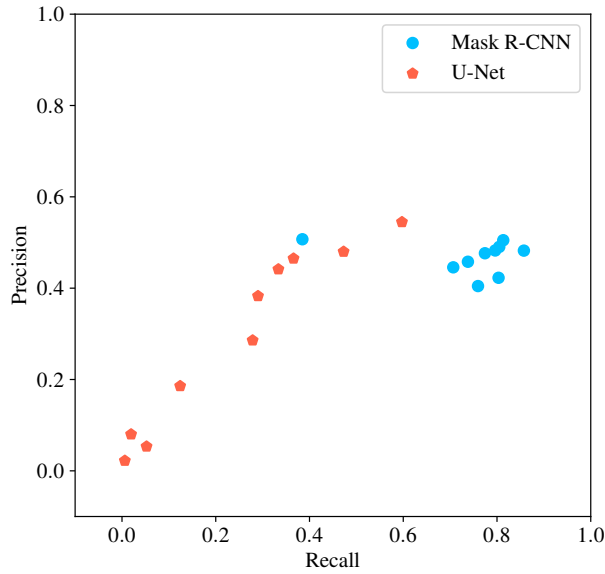


Figure 2: Comparison between Recall based on Precision for 10 training per network

the network is much less stable than Mask R-CNN. On average, performance seems much better with Mask R-CNN. In fact, even if U-Net's Precision is equivalent to Mask R-CNN in a few cases, Mask R-CNN's Recall is always much better (more real objects detected). Regarding detection performance, the network appears more effective than U-Net, as confirmed by Table 1, where the F-Score for each case is represented. On average, the F-Score calculated for Mask R-CNN is much higher than that for U-Net. Instance segmentation seems to perform better than semantic segmentation,

Table 1: Comparison between F-score for each run for each network

	U-Net	Mask R-CNN
1	0.38	0.64
2	0.35	0.60
3	0.01	0.60
4	0.26	0.55
5	0.05	0.55
6	0.03	0.59
7	0.45	0.58
8	0.14	0.36
9	0.52	0.66
10	0.29	0.62
<b>Average</b>	<b>0.25</b>	<b>0.58</b>
<b>Standard deviation</b>	<b>0.17</b>	<b>0.08</b>

taking into account the individuality of objects.

### 5.1.2 Comparison of estimated physical quantities

The PDFs of the detected drops by combining the 10 draws are shown in Figure 3a for U-Net and 3b for Mask R-CNN and their GT to estimate which method is representative of the true physical quantities. If the GT PDF for the two networks are different, this is due to the random selection of images. The U-Net network significantly underestimates drop size, as shown by the peak at low diameter values certainly attributed to false detections. The prediction made by Mask R-CNN is closer to the GT, even if it shows less efficient detection at the level of large drops.

A much lower KL divergence for Mask-RCNN (0.08) than for U-Net (0.35) confirms this. Instance segmentation, therefore, appears to be much more relevant and efficient in detecting and for the quantitative values of physical quantities.

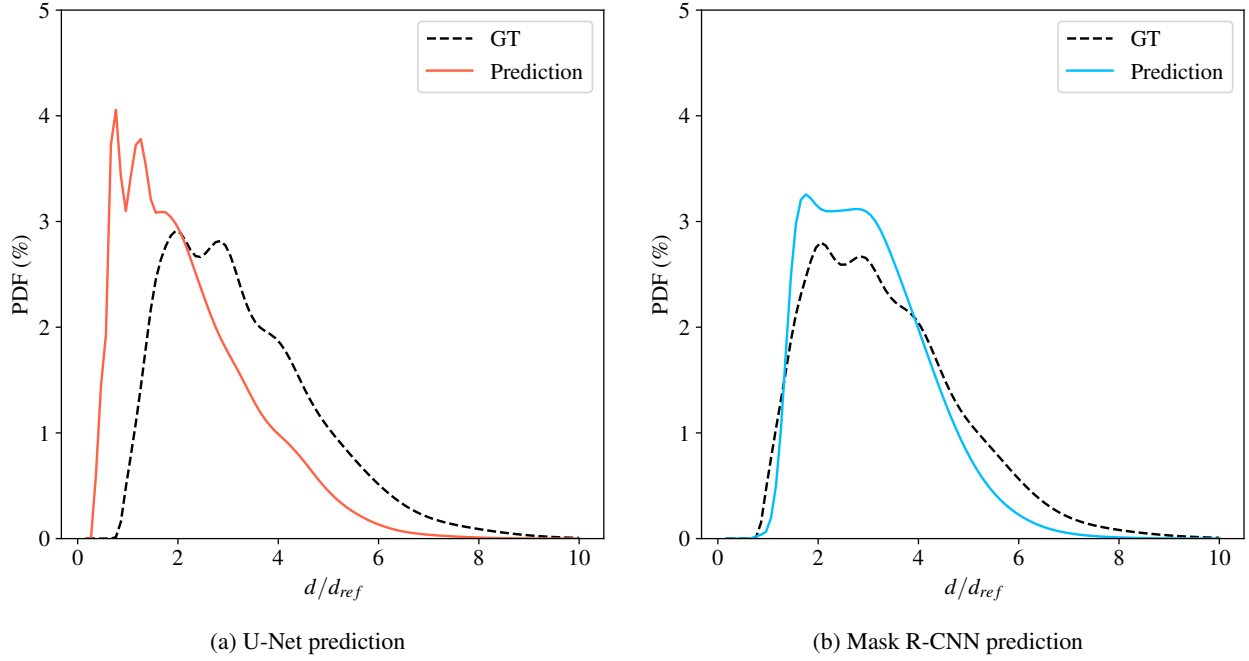


Figure 3: Comparison between probability density function of droplet's diameter between each pressure

### 5.1.3 Comparison of time performance

The processing time for each method is a crucial factor when evaluating its performance and practical use. During the training process, much time is spent annotating the required GT images. On average, it takes around 5 minutes to annotate each image. Subsequently, the trained model is validated using test images. If the model is sufficiently versatile, i.e. if it has been trained on various experimental conditions with different settings, it can be directly applied to complete sequences of shadowgraphy images corresponding to an entire test. As a result, there is no need to train the method for each individual test, nor is there a need for image annotation each time. This makes the method highly adaptable and less time-consuming once trained.

Table 2 illustrates the projected GPU processing times for training a model on 45 images and applying it to sequences comprising 3000 images using the MSER, U-Net, and Mask R-CNN methods. The estimated durations for the MSER method are derived from prior research experiences [5, 6, 13]. The main aspect is the need to adjust MSER settings for each sequence in order to achieve reasonable detection performance via a simple try-and-test approach. The term "Object detection" refers to the use of the CNN model or the MSER algorithm in order to detect droplets in the image.

Table 2: Comparison between time processing and detection quality for each method

		MSER	U-Net	Marsk R-CNN	
Detection quality	Repeatability	+	+	+++	
	Performance	+	++	+++	
Learning process	Annotation	0	5h	5h	
	Train	0	1/2h	1/4h	
Processing Time	Parameter adjustment	2h	0	0	
	Image sequence analyse	Object detection	4h	1h	1/2h
		Post-processing	1h	1/4h	1/4h
		Total time	1 sequence	7h	1h1/4
	5 sequences	35h	6h1/4	3h3/4	

The results indicate that Mask R-CNN surpasses other methods not only in terms of detection accuracy but also in terms of processing time; the network is slightly faster than U-NET on our hardware setting. With Mask R-CNN, it is possible to perform droplet detection and analysis in under an hour once the experimental tests are completed. Once

## EVOLUTION OF BURNING ALUMINUM DROPS BY DEEP LEARNING

the model is trained, droplet detection would be possible in parallel on the experimental measurements, which offers excellent flexibility.

## 5.2 Estimation of physical quantities

This section compares two different test conditions for a single propellant composition to show the efficiency of the instance segmentation approach. The models were derived from MRCNN-A on a database of 250 shadowgraphy images from different test conditions (various propellant compositions, various pressure and various spatial resolution). Figure 4 showcases the outcome of applying the model trained with the MRCNN-A network on one image that has not been used for learning. Each identified drop is accurately segmented and enclosed within a bounding box. Above each box is shown a score, indicating the likelihood of the detected object being a drop, is displayed. The network itself generates this score.




Figure 4: The MRCNN-A network is employed to detect drops, leveraging its trained model

The learned model was then applied to sequences of shadowgraphy images from two testings, one at 9 bar and the other at 11 bar. The two sequences each contain more than 3,000 images. The following Table 3 summarises the characteristics obtained from the images. This is the first time there has been such a large-scale study of aluminium combustion using image analysis.

Table 3: Characteristic details obtained from image sequences by applying the model

Pressure	Number of images	Number of detections	Number of tracks
9 bar	3874	52938	1240
11 bar	3363	53186	1285

In this study, every detection and track has been meticulously validated and carefully sorted through automated post-processing methods to ensure that only accurate data is considered. This meticulous filtering process eliminates any false detection or unexpected trajectories. The resulting dataset contains an extensive collection of reliable trajectories and accurate detections, thus providing a robust foundation for the statistical investigations of various physical properties described below. Consequently, the statistical analyses conducted on these physical quantities accurately represent the real phenomena occurring above the surface of the propellant during combustion.

### 5.2.1 Visualisation of droplets' trajectory

Figure 5 illustrates the trajectories of three distinct drops, each with a different size, leaving the propellant surface. These drops can be classified as follows: a small drop (#5887), a medium-sized drop (#6315), and a large drop (#7962). Although all drops generally exhibit a linear trajectory with minimal horizontal variation, their behaviour differs depending on size. The primary factor influencing their behaviour is the drop size: larger drops take a longer time to



## EVOLUTION OF BURNING ALUMINUM DROPS BY DEEP LEARNING




Figure 5: Comparison of positions of 3 droplets of different sizes: small (#6887), medium (#6315) and large (#7962)

detach from the surface (as depicted in Figure 5d), resulting in a seemingly slower overall speed. This can be observed by the greater number of successive positions documented towards the end of the trajectory, as demonstrated in Figure 5f. Interestingly, it is worth noting that automatic tracking of smaller drops poses a greater challenge due to their lower number of occurrences in the captured images. Consequently, the tracking data for smaller drops are comparatively less informative, leading to less precise predictions of their next positions.

### 5.2.2 Droplet size

From each of these trajectories, statistical analyses were conducted to investigate the behaviour of the droplets; an averaged diameter is attributed to the track. PDFs for different experimental conditions are depicted in Figure 6. Notably, the combustion of aluminium in the propellant at 9 bar results in the formation of larger droplets compared

## EVOLUTION OF BURNING ALUMINUM DROPS BY DEEP LEARNING




Figure 6: Probability density function of droplet's diameter for 9 bar and 11 bar pressures

to higher pressure conditions. This observation is consistent with existing literature, which suggests that the size of aluminium droplet agglomerates tends to decrease with increasing pressure due to the facilitated ignition of aggregates [18]. Also noteworthy is that the droplets appear slightly less dispersed at 11 bar than at 9 bar. It would be interesting to explore the influence of higher pressures on this phenomenon for instance, by analyzing tests up to 30 bar.

### 5.2.3 Droplet velocity

By tracking the droplet's motion during combustion, it becomes possible to establish velocity profiles as a function of its diameter by considering the distance from the surface. This relationship is illustrated in Figure 7, with the left side corresponding to 9 bar and the right side to 11 bar.




Figure 7: Comparison between the dispersion of velocity profiles for pressures 9 and 11 bar. The colorscale corresponds to the normalized diameter.

The variation in speed observed in Figure 5, e.g. smaller droplets moving faster, is corroborated here. For both

pressures, it is evident that larger droplets, depicted by yellow/orange colours, exhibit lower velocities compared to smaller droplets (purple/mauve colour). While the drag force  $F_D$  is proportional to the droplet diameter  $D$  ( $F_D \propto D$ ), the droplet inertia is proportional to the cube of the droplet diameter ( $m \propto D^3$ ). The droplet acceleration  $a$  is hence decreasing with its diameter increasing ( $a \propto \frac{1}{D^2}$ ) [9]. This is visible with the steeper slope on velocity profiles for small droplets. The graph can be easily divided into distinct zones for large, medium, and small droplets, which could be further explored in future investigations.

It is also possible to see that the velocity is lower as the pressure increases: for drops of the same diameter, their velocity will be lower on average at 11 bar than at 9 bar. For example, drops with a normalised diameter of 12 (in dark orange) have a normalised speed of 1-1.5 for 9 bar compared with 0.5-1 for 11 bar. This is consistent with the decrease in gas velocity with increasing pressure. This shows the capacity of the present analysis method to detect a velocity for a relatively small pressure variation (20%).

It is feasible to observe the initial velocities  $v_{init}$  of the droplets, which refers to their speed upon leaving the surface of the propellant. The initial velocity is a key parameter simulation since it is an input parameter for droplets seeded into the flow. These velocities are depicted in Figure 8 through their probability density functions (PDFs). The initial velocities are slightly lower at 11 bar compared to 9 bar.




Figure 8: Probability density function of droplet's initial velocity for 9 bar and 11 bar pressures

More precisely, it is interesting to see if the initial velocity depends strongly on diameter. Figure 9 displays the distribution of droplets across these two quantities. The X and Y axes represent, respectively, initial velocity and diameter. The droplet's points are grouped into rectangular bins based on their values along the X and Y axes. The colour of each bin represents the percentage of droplets found within that particular bin.

It can be compared to the one presented in the article [19], where the drop diameter is plotted against the velocity of the drop upon leaving the burning surface. Both graphs show a similar trend, indicating that as the particle size increases, the initial velocity decreases. However, the authors of [19] had noted a wider dispersion of values compared to what is observed here. This difference can be attributed to the number of velocities studied, with 176 analyzed in the article [19] compared to more than 1200 velocities per pressure shown in the accompanying figures.

In accordance with the PDFs (Figure 8), the initial velocities are higher for 9 bar than for 11 bar. However, the dispersion of  $v_{init}$  as a function of diameter seems much smaller for 9 bar than for 11 bar, where the values cluster around particular diameters. The dependence between droplet diameter and initial velocity, therefore, appears to be greater for higher pressures. Further studies could be carried out to investigate this correlation in particular by looking at the dispersion of  $v_{init}$  by size class.

## 6. Conclusion and perspectives

This paper provides evidence for the effectiveness of deep learning-based instance segmentation in characterizing aluminium combustion of solid propellant. When compared to the current state of the art, instance segmentation

## EVOLUTION OF BURNING ALUMINUM DROPS BY DEEP LEARNING




Figure 9: Comparison between the dispersion of droplet's initial velocity for 9 bar and 11 bar pressure in function of their diameter

exhibits significant superiority, not only in detection but also in the accurate estimation of relevant physical quantities. This approach enables the development of a robust and reliable droplet prediction model, which can be applied to entire sequences of shadowgraphy images, representing a novel advancement in the field.

By utilizing instance segmentation, it becomes possible to derive physical quantities from existing literature. What sets these quantities apart is their derivation from a substantial number of measurements, ensuring the stability and statistical reliability of the subsequent analyses. Overall, this research demonstrates the potential of instance segmentation by deep learning in advancing the understanding of combustion processes and obtaining valuable insights from extensive measurement data.

In the future, it would be valuable to expand the study by including a wider range of pressures. This would enable a comprehensive comparison and analysis of the combustion process across various pressure conditions. Additionally, extending the investigation to different types of propellants would allow for a thorough exploration of their dissimilarities during combustion.

Moreover, the analysis of physical quantities could be further developed to explore additional phenomena, such as the acceleration of the droplets. By delving into these aspects, a deeper understanding of the underlying physics and dynamics of the combustion process could be achieved. This could contribute to advancements in the field and provide insights into the behaviour of droplets and propellants under varying conditions.

## 7. Acknowledgments

The study presented here was made possible thanks to the assistance of a PhD grant provided by CNES (Centre National d'Etudes Spatiales) and ONERA (Office National d'Etudes et de Recherches Aérospatiales). A special thanks to Aurélie Covasso and Laure-Emilie Martin, who helped us with the images annotation.

## References

- [1] F. Godfroy S. Gallier. Aluminum combustion driven instabilities in solid rocket motors. *Journal of propulsion and power*, 25(2):509–521, 2009.
- [2] A. Genot, S. Gallier, and T. Schuller. Thermo-acoustic instabilities driven by fuel droplet lifetime oscillations. *Proceedings of the Combustion Institute*, 37(4):5359–5366, 2019.
- [3] O. Orlandi, M. Plaud, F. Godfroy, S. Larrieu, and N.Cesco. Aluminium droplets combustion and srm instabilities. *Acta Astronautica*, 158:470 – 479, 2019.

- [4] R. W. Devillers, C. Erades, D. Lambert, and J. Bellessa. Mesure et suivi de particules, agglomérats et gouttes en combustion. In *14th CFTL*, 2014.
- [5] M. Nugue, R. Devillers, G. Le Besnerais, and N. Cesco. Particle detection & size evaluation in solid propellant flames via experimental image analysis to improve two-phase flow simulation in rocket motors. In *Space Propulsion*, ROME, France, May 2018.
- [6] M. Nugue. *Outils pour l'étude conjointe par simulation et traitement d'images expérimentales de la combustion de particules d'aluminium utilisées dans les propergols solides*. Theses, Université Paris-Saclay, October 2019.
- [7] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017.
- [8] M. Airiau, A. Chan-Hon-Tong, R. W. Devillers, and G. Le Besnerais. Regressing image sub-population distributions with deep learning. *Sensors*, 22(23), 2022.
- [9] S. Gallier T. G. Decker, R. W. Devillers. Agglomeration in Solid Propellants loaded with inert particles - Study of physical phenomena using shadowgraphy image processing. In *EUCASS 3AF 2022*, Lille, France, June 2022.
- [10] G. Vilmart M. Nugue G. Le Besnerais J. Pichillou R. W. Devillers, N. Dorval. Aluminum particle tracking on experimental shadowgraphy and al-plif images to provide velocity data for two-phase flow simulations of solid rocket motors. In *Proceedings of the Space Propulsion Conference, Seville, Spain*, pages 14–18, 2018.
- [11] T. G. Decker, R. W. Devillers, and S. Gallier. Experimental visualization of aluminum ignition and burning close to a solid propellant burning surface. In *European Combustion Meeting 2023*, 2023.
- [12] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing*, 22(10):761–767, 2004. British Machine Vision Computing 2002.
- [13] R. W. Devillers, M. Nugue, A. Chan Hon Tong, G. Le Besnerais, and J. Pichillou. Experimental analysis of aluminum-droplet combustion in solid- propellant conditions using deep learning. In *EUCASS 2019*, MADRID, Spain, July 2019.
- [14] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. *LNCS*, 9351:234–241, 2015.
- [15] S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79 – 86, 1951.
- [16] M. Rosenblatt. Remarks on some non parametric estimates of a density function. *The Annals of Mathematical Statistics*, 27(3):832–837, 1956.
- [17] E. Parzen. On estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33:1065–1076, 1962.
- [18] J. Duterque. Experimental study of the agglomeration of aluminum in solid propellant motors. *ONERA*, 1996.
- [19] J. Yuan, J. Liu, Y. Zhou, J. Wang, and T. Xv. Aluminum agglomeration of ap/htpb composite propellant. *Acta Astronautica*, 156:14–22, 2019.