

# Experimental analysis of aluminum-droplet combustion in solid-propellant conditions using deep learning

*R.W.Devillers\*, M. Nogue\*\*, A. Chan-Hon-Tong\*\*\*, G. Le Besnerais\*\*\*\*, J. Pichillou\*\*\*\*\**

*\* ONERA, Multiphysics for Energetics Department,  
BP 80100 – 91123 Palaiseau Cedex France, Email: robin.devillers@onera.fr*

*\*\* ONERA, Multiphysics for Energetics Department,  
BP 80100 – 91123 Palaiseau Cedex France*

*\*\*\* ONERA, Information Processing and System Department,  
BP 80100 – 91123 Palaiseau Cedex France, Email: adrien.chan\_hon\_tong@onera.fr*

*\*\*\*\* ONERA, Information Processing and System Department,  
BP 80100 – 91123 Palaiseau Cedex France, Email: guy.le\_besnerais@onera.fr*

*\*\*\*\*\* CNES, The Launcher Directorate,  
52 rue Jacques Hillairet – 75612 Paris Cedex France, Email: julien.pichillou@cnes.fr*

## Abstract

Shadowgraphy is good technique to image aluminum combustion for small solid-propellant samples. But classical detection methods struggle to provide accurate data on aluminum droplets with such different aspects. A new semantic-segmentation approach is tested using deep learning to classify every image pixel over various classes. The UNET network was trained on 45 annotated images for combustion at 1.0 and 2.0 MPa. Detection performances were compared to those of the classical Maximally-Stable Extremal Regions algorithm. The semantic-segmentation approach was able to distinguish droplets from the surrounding flame with better detection performances (Precision, Recall).

## 1. Introduction

Solid propellants are commonly used for space and military applications because they provide high efficiency levels for a relatively low cost. In typical compositions powder-like oxidizers such as ammonium perchlorate (AP) are gathered into a polymer binder, usually hydroxyl-terminated polybutadiene (HTPB). Aluminum is a widely-used performance additive (with weight fractions up to 20%) to improve and adjust effective impulse. However the combustion of aluminum droplets can have a dramatic impact on performance stability [1][2][3] with a strong aluminum-size dependence on the stable or unstable behavior. The diameter for the droplets that effectively leave the burning surface to enter the rocket flow is one of the commanding parameters. But quantifying precisely aluminum-droplet diameter is not straightforward since the aluminum droplets leaving the surface result from the agglomeration of various aluminum particles initially seeded into the energetic material [4][5].

Focusing shadowgraphy is used at ONERA to provide aluminum-particle visualization at high repetition rate close to the surface of burning propellant samples [6][7][8][9]. Every combustion test provides more than 1000 shadowgraphy images so automatic detection methods are used to provide statistics on various droplet parameters such as size and velocity. The present work aims at comparing two detections methods: a classical detection algorithm used for aluminum-droplet detection since 2016 [8][9] (named MSER [10] and described later) and a Convolutional Neural Network trained with a deep learning approach (UNET [13]). The present paper evaluates detection performances for both tested approaches for two targeted regions types associated to aluminum combustion, namely the aluminum-droplet alone and its surrounding flame.

## 2. Experimental apparatus

### 2.1. Shadowgraphy set-up

The present shadowgraphy approach [7][8] uses a focused beam (Figure 1 (a)) as light source in order to increase light intensity as well as reduce the depth of view. Increasing the light-source brightness is a key aspect since solid propellants in general are highly bright and burning aluminum droplets are even brighter. An ideal shadowgraphy set-up to study burning aluminum droplets would use a light bright enough to generate images with dark droplets on a brighter background. But in many cases droplet emission leads to images with brighter droplets on a darker background, as shown in the example image in Figure 1 (b). Typical experimental shadowgraphy images show the burning surface of the solid-propellant sample as a dark stripe at the bottom. Hot gases flow vertically upwards with various burning aluminum droplets moving up as well as vertical dark stripes corresponding to alumina plumes. The goal is to automatically select all visible aluminum droplets without including inaccurate objects such as smoke areas. The final goal is even more ambitious since proper droplet-combustion analysis requires distinguishing the aluminum droplet itself from the surrounding flame. The optical set-up can be adjusted to generate various spatial resolutions for the images. There's a classical tradeoff between "detailed droplets over a small probe window" and "less detailed droplets over a large probe window". Two spatial-resolution conditions were analyzed in the present article.

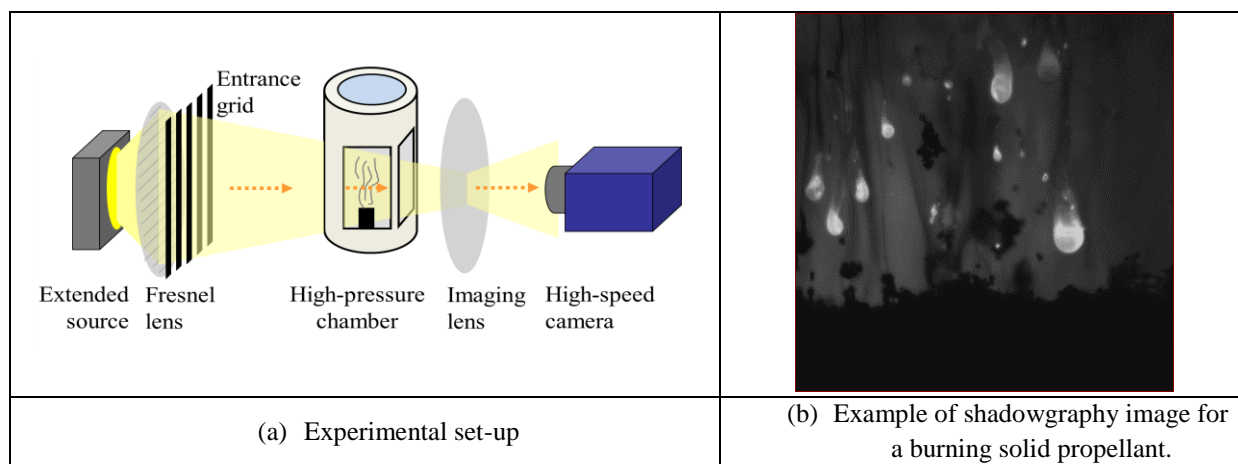


Figure 1. Shadowgraphy set-up for solid-propellant flame visualization and experimental-image example.

### 2.2. Solid propellants and burning conditions

Combustion tests were performed in a close pressurized chamber used to study small samples. The chamber is initially pressurized with nitrogen between 1.0 to 2.0 MPa in the present investigations. Solid-propellant samples are small with total mass below 200 mg, leading to total combustion duration below 2 s. A CO<sub>2</sub> laser at 10.6  $\mu\text{m}$  is used to heat the solid-propellant surface up to ignition. In the present work a single propellant composition was tested since the main goal is the evaluation of the image-analysis approach. The solid-propellant research composition is a typical AP/HTPB composition with an aluminum mass fraction of 18%. The aluminum-droplet size ranged approximately from 10 to 200  $\mu\text{m}$  in diameter during combustion. Three combustion tests were analyzed here as described in Table 1. As stated earlier, two spatial resolutions were analyzed: in the case of the less detailed spatial resolution pixels were approximately twice larger than in the more detailed case. Hypothetically, if the same droplet was seen in the two cases, it would appear 1.8 smaller in the shadowgraphy image for the less spatially-resolved image. The test for the more detailed spatial resolution was labelled 104 whereas two combustion series were tested for the lesser resolution, labeled 502 and 503. The table also includes details of initial pressure. The question is important from an image-analysis point of view since an increase in pressure leads to larger density of objects on the experimental images.

Table 1 Experimental conditions for the various image series.

Al droplet diameter [μm]	Normalized pixel size [a.u.]	Series label #	Pressure [MPa]	Aspect ratio	Image aspect
10 ; 200	1.0	104	2.0	Square	Darker images
	1.8	502	1.0	Square	Bright background with vertical smoke plumes
	1.8	503	1.0	Vertical rectangle	Bright background with vertical smoke plumes

### 2.3. Image characteristics for the analyzed series

As presented above, the choice of dealing with various combustion tests was motivated by the idea of applying the image analysis to data with different characteristics. In order to illustrate more systematically image conditions, various image characteristics were introduced to account for the three types of regions in the flame portion of images: the targeted burning aluminum droplets, disturbing alumina smoke and the overall background corresponding to hot combustion gasses. The aspect of the burning droplets was characterized with two parameters: the signal level, labeled “Al flame level”, and the area size, labeled “Al flame size”. A large & bright aluminum flame area should be less difficult to detect than a small & dark one. The hot-gas area was evaluated by its signal level (“Gas level”) and its homogeneity (“Gas homog.”). Homogeneous backgrounds should make it easier to segment specific droplets over the hot-gas background rather than non-targeted objects. Finally, the potential impact from alumina-smoke regions was quantified by the area proportion on the image corresponding to smoke (“Smoke %”). For comparison sake, all five characteristics can be evaluated over a normalized scale ranging from 1.0 to 5.0 based on their range for the tested images: 5.0 for the most favorable condition and 1.0 for the least favorable.

In the present case, the five characteristics were evaluated by roughly segmenting burning droplets and smoke regions on all tested images. Hence, pixels from the images were divided into 4 classes: burning solid propellant, hot gases, burning aluminum droplets and alumina smokes. The classification is by no means perfect but it enables a rapid evaluation of image characteristics. It is illustrated in Figure 2 for two images extracted from series #104 and #502, i.e. for different spatial resolutions and different initial pressure conditions; the categories are displayed over a 4-level gray color scale. It shows rather clearly that smoke regions are different for the two image conditions and might lead to different kind of aluminum-detection bias. The resulting images characteristics are plotted in Figure 3 as radar graphs. Image conditions are rather different for test #104 on the one side and tests #502 and #503 on the other side. Aluminum-droplet are obviously larger for #104 with the refined spatial resolution but the shadowgraphy setting also led to darker background levels and very bright droplets. Tests #502 and #503 led to narrower contrast between background and droplet regions (brighter background, darker flames) but also a larger proportion of smokes stripes, certainly due to the larger overall field of view. It will be interesting to check if the differences lead to detection bias in the automatic detection algorithms.




Figure 2. Classification of shadowgraphy image regions into 4 categories. An example is given for the test series #104 and another one for test series #502.




Figure 3. Image characteristics for the three analyzed image series.

### 3. Aluminum-droplet detection

The present section introduces the image analyses that were applied in order to detect automatically burning aluminum droplets in shadowgraphy images. Two approaches were tested: a classical image-analysis approach (MSER) and a deep learning approach based on convolutional neural networks.

#### 3.1. Detection-performance evaluation and test images

Reference images were used in order to evaluate the detection efficiency for each tested analysis approach. Shadowgraphy images were annotated manually, e.g. regions on images were classified in various categories of

interest. The annotated images constitute a Ground Truth (GT) to with automatic processing methods can be compared. Up to seven object categories have been studied in previous shadowgraphy-image analysis [11] as listed below. They include aluminum particles at various stages of their life (inert, burning, and aggregated), Al<sub>2</sub>O<sub>3</sub> residuals (as cap on burning Al droplets or as smokes) as well as the solid-propellant sample and the hot gas regions. We are targeting specifically two categories here focused on Al-droplet combustion: the burning aluminum droplet (in blue) and its surrounding flame (in red).




Figure 4. Example of shadowgraphy image with the associated Ground Truth obtained by manual annotation.

For each of the 3 studied image series (Table 1), 20 images were manually labelled at various stages over the full combustion test, leading to 60 Ground Truth images in total. Some of them were used to adjust analysis parameters: the processing input parameters were chosen in order to provide good automatic detection performances for this set of images, called “training images” or “adjustment images”. In the present case, 15 GT images per test series were chosen as “training / adjustment image” (45 in total). The remaining 15 GT images (3 x 5) were used as “test images”: the idea is to check the detection performances for them and see if the performances are still good. A possible bias could come from “learning by heart” the small “train image” sample without any robustness, e.g. tuning processing parameters very precisely over the limited set of images but with poor performances for other images. Dividing the full Ground Truth image set into two subsets is a common approach in machine learning methods to check for processing robustness.

Table 2. Number of objects used for processing adjustments (“training / adjustment”) and performance evaluation (“Test”) for the categories of interest.

Category	Training / Adjustment			Test		
	Images	Regions	Pixels	Images	Regions	Pixels
Burning Al droplet alone	45	327	1.86E+05	15	96	5.01E+04
Al flame + droplet	45	878	1.81E+06	15	279	5.34E+05
Al flame alone	45	878	1.27E+06	15	279	3.59E+05

The number of images used for training and test might seem rather limited and far from statistical representativeness. It is important to comment on the number of targeted objects for each category as presented in Table 2. More than 100 objects were annotated in both categories, which looks like a reasonable statistical sample. Moreover, as will be explained in more details when the deep learning method is described, semantic segmentation corresponds to a classification of every pixel. The number of annotated pixels for each category is also specified in Table 2. More than 10000 pixels are used for training for each category. It illustrates that in theory only a few images need to be annotated in order to train a neuronal network for semantic segmentation. This is not a minor detail to consider since GT annotation is rather time consuming...

Detection efficiency was evaluated with the usual quantifiers by comparing automatic detections with Ground Truth images [8][9][11]. Detection of an object present in the GT image is a True Positive (TP) whereas a detection that

does not correspond to any GT object is a False Positive (FP). A GT object that is not detected is a False Negative (FN). Various detection-performance quantifiers are evaluated from TP, FP and FN:

- *Recall (Rec)*: proportion of real objects that are effectively detected (e.g. TP / GT);
- *Precision (Pr)*: proportion of detections that correspond to real objects (e.g. the proportion of false detections is 1-Pr), calculated as TP / (TP+FP);
- *Intersect over Union (IoU)*: intersect of detection and GT divided by the union of detection and GT.  $IoU = 1$  if a detection matches perfectly the associated GT object. IoU was applied here only for the real objects that were effectively detected.

The idea is to obtain a good *Rec* level in order to detect all the real objects but without detecting many false object (e.g. with good *Pr* level). The *IoU* parameter is a way to evaluate the algorithm capacity to provide detections with a size similar to the associated GT object. *Rec*, *Pr* and *IoU* were evaluated for the two categories of interest introduced previously, “droplet” and “droplet+flame”.

### 3.2. MSER algorithm (Maximally Stable Extremal Regions)

Object detection was first obtained via a Maximally Stable Extremal Regions algorithm (MSER). It was first introduced to associate objects on images with different scales [10] and was used in the past to detect objects on shadowgraphy images for solid-propellant combustion [8][9]. It extracts regions with homogeneous signal level by testing region stability over a range of threshold levels. No fixed threshold level is used, which is convenient for images with complex objects to detect. MSER performances were particularly interesting in the case of research propellant compositions seeded with inert particles [8] since the contrast between dark particles with the bright background was strong and the inert particles corresponded to rather homogeneous regions on the images. However MSER was not as efficient to detect burning-aluminum droplets as shown in [9]. The present aim is to refine MSER droplet detection here and see if the algorithm has good potential for aluminum-combustion studies.

Table 3. MSER parameters.

Target regions	MSER parameters	Circle fit	No alumina smoke criteria	Typical t targeted objects
Droplets	Tighter threshold constraint	Yes	Yes	Small / Round / Homogeneous
Droplet + Flame	Lesser threshold constraints	No	Yes	Small or large / No fixed shape More homogeneous

Previous MSER use for aluminum-combustion studies merely targeted overall “burning droplet regions” without discriminating the droplet from the surrounding flame. Two MSER settings were tested here in order to see if the MSER algorithm could target specifically droplets or the surrounding flame. The droplet parameters were set-up in order to restrain MSER thresholding criteria and provide smaller, rather homogeneous regions. The other parameter set was chosen as less restrictive in order to extract wider regions with lesser homogeneity constraints, hence targeting the overall “droplet + flame” regions in the image. Additionally a couple of simple post-processing steps were added to the mere MSER algorithm. First, morphologic criteria were included in order to remove detected alumina smoke regions (i.e. large vertical region with rather dark signal levels). Second, a simple circle fit was applied to the extracted MSER regions in order to provide only circular regions. The circle fit was used when droplets were targeted but not for “droplet + flame” detections. The chosen MSER settings are gathered in Table 3.

### 3.3. UNET network

Semantic segmentation [12] based on deep convolutional neural network (CNN) is now a marked task. In this task, the CNN has to produce a map with a probability density in each pixel, while, in image classification, CNN just produce a simple probability density for the whole image. In classification, CNN mostly contains convolution and pooling layers which incrementally purify and concentrate the information discarding spatial location. In semantic segmentation, CNN combine an encoder (with an architecture close to a classification CNN) and a decoder which restores spatial-resolution using up-pooling strategy to merge purified information from the encoded layer which

spatially aware information from first layers of the network. The simplest way for this merging is to upsample incrementally low spatial-resolution layer, and to concatenate it with the corresponding resolution layer from the encoder. This idea has been introduced by UNET [13].




Figure 5. UNET network architecture [13].

Due to the specificity of present shadowgraphy images, naive application of UNET leads to poor performance. Yet, only minor preprocessing and architectural changes were needed to reach interesting results. The various tested changes are listed in the following table. The most important change is the use of batch normalization after each convolution (tested cases d) and e)). Batch normalization introduced in ResNet [14] forces mean/variance to be 0/1 for the current batch of processed data. Such layer requires high memory GPU, but, allows a better stability of the weight dynamic, and, thus of the global training. Independently, providing to the network the propellant surface also makes training faster (test case e)). In our experiments, these modifications have much more effect than classical unbalance management tricks that correspond to giving more weights to outnumbered classes (tricks still tested as cases b) and c)). The surface detection was obtained with an active contour approach [15] that had already been tested previously [16]. The high contrast level between the dark solid-propellant sample and the brighter hot-gas region makes it relatively easy for more classical segmentation methods to provide good performances, as had been presented in [16].

Table 4. Input parameters used for the fine tuning process with UNET.

In the table, « im. » corresponds to the experimental image.

Label	Weight of droplet class	Batch normalization	Image channels (RGB)		
			Ch.1	Ch.2	Ch.3
a	1	no	im.	im.	im.
b	20	no	im.	im.	im.
c	10	no	im.	im.	im.
d	1	YES	im.	im.	im.
e	1	YES	im.	im.	segmented surface

## 4. Results

### 4.1. MSER detection trends

Object detections are compared to GT in Figure 6: in (a), detection contours are displayed over the experimental image whereas in (b) all pixels in the image frame are displayed as TP / FP / FN. This is shown for the two sets of MSER input parameters. When droplets alone are targeted various regions detected by MSER correspond to flame areas even if no droplet visible (e.g. several regions corresponding purely to FP regions in red). The MSER algorithm



is not selective enough to avoid segmenting a flame region when no droplet is visible inside. On the other side, when MSER is set to target whole “droplet+flame” regions, the segmentation misses some of the targeted areas as shown with the FN areas in blue. The trends are illustrated in more details in Figure 7 with a close up on a specific Al droplet surrounded by flame. The GT contours are shown in red and blue over the image in the top right frame (with a bright green frame). The bottom two images show TP / FP / FN pixels for each of the two MSER settings. The “droplet” setting shows a large portion in red, i.e. FP, showing that the detected area is much larger than needed to match perfectly GT. The “droplet+flame” setting on the other hand displays a large blue portion, i.e. FN, showing that the segmented area is much smaller than the GT. It seems hard to obtain a proper MSER setting to target one or another region of the burning Al droplet and the surrounding flame.




Figure 6. MSER detections compared to GT for the two sets of inputs parameters.




Figure 7. Detection with the MSER sets of parameters: close up on one example of aluminum droplet.

The GT frame overlay the annotated regions to the shadowgraphy image in the green frame.

The MSER detection frame (black rectangle) presents the comparison to GT for the two sets of MSER inputs.



The reason for the difficult MSER segmentation is illustrated in Figure 8 by showing two examples of burning droplets present on shadowgraphy images. For each example, images show the experimental image with the contour provided by the GT annotation for “droplet” and “flame”. At the left-hand side intensity histograms are shown for the “droplet” (bleu bars) and “flame” regions (red curve). The two droplet examples show clear differences in aspects. In terms of relative signal levels, in (a) the droplet region is darker than the flame whereas in (b) the droplet is very bright, certainly because a large portion of alumina cap is visible. In terms of relative size, the (b) droplet is small compared to the flame area to segment but the (a) flame is not much larger than the associated droplet. The associated intensity histograms illustrate the range of possible aspect for the droplet / flame duet. The MSER algorithm is mostly sensitive to signal level and homogeneous regions in the images. Distinguishing a droplet from a flame certainly requires including other kinds of considerations such as shapes, contrast, morphology...

A more general graph is presented in Figure 9 to show that the two droplets from Figure 8 are not out-layer examples. It displays all annotated GT “droplet+flame” regions characterized by two characteristics. First a size characteristic, the relative area size for the two regions (droplet area / “droplet+flame” area). Second, a signal level characteristic, the relative signal level for two regions (droplet signal / flame signal). The two examples from Figure 8 are shown with red circles in the cloud graph. The data show that the droplet size ranges from 10% to 70% of the total “droplet+flame” region to detect. In terms of signal, droplet and flame remain close most of the time since the signal ratio remains between 0.8 and 1.2 for most cases. No clear criterion is visible in terms of relative size and relative signal levels to add a simple post-processing step to MSER detection that could improve distinction between droplet regions and flame regions.

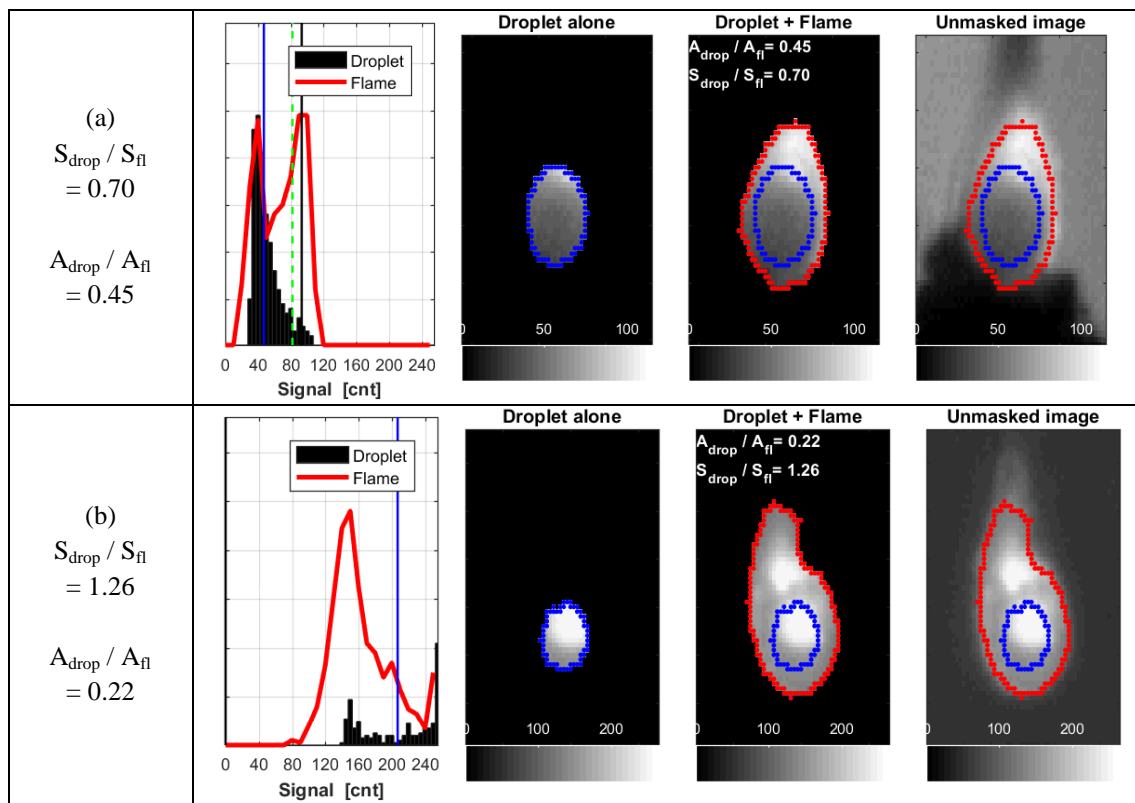


Figure 8. Examples of aluminum-droplet aspect on shadowgraphy images. Droplet and flame regions delimited by manual annotation. The histograms give the signal-level distribution for each region.

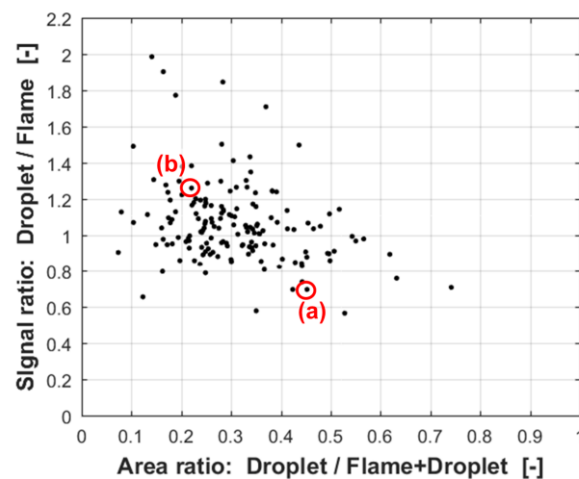


Figure 9. Area and signal characteristics for all annotated droplet/flame regions.

Area ratio: ratio of droplet area and the complete area (droplet+flame).

Signal ratio: ratio of droplet signal level to flame signal level.

#### 4.2. UNET semantic segmentation trends

The present section comments on semantic segmentations trends for the various tested settings on the input images fed to UNET. More quantitative comments on detection performances are given in the following section.

First, as shown on Figure 10, the direct application of UNET to non-modified input images is rather inefficient for droplet detection. The only detected classes are “solid-propellant sample”, “hot-gas region” and “AI flame”, i.e. the main classes. All other classes do not seem learnt by UNET, certainly because, as minority classes, they correspond to fewer pixels on the “train” images. It illustrates that UNET and CNN in general are not fully efficient in a basic plug-and-play approach. Settings and input images need to be adjusted keeping in mind the specific targets and the specific image conditions.

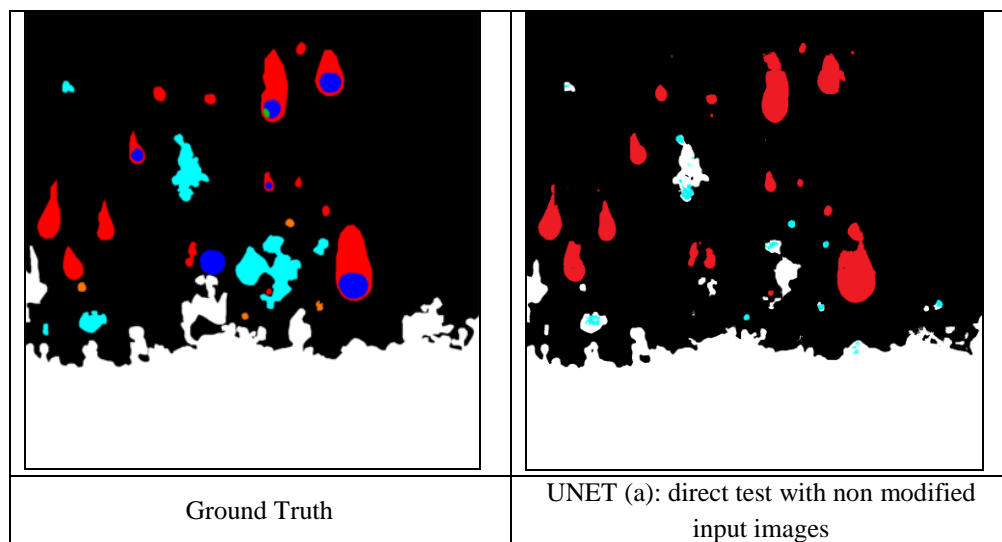


Figure 10. GT (left-hand side) compared to the semantic segmentation provided by UNET with parameters (a) (right-hand side)

The first adjustment axe is increasing the weight for minority classes in order to see if UNET is able to learn them. Segmented images are shown on Figure 11 for adjusted weight levels for the droplet class (x20 and x10). Droplet regions are now visible with this two corrected weight levels. But the segmentation is highly sensitive to the chosen weight correction: UNET seems to detect droplets everywhere when droplet weight is increased by x20 whereas

some droplets are missed with the x10 weight change. Adjusting manually class weight is not a robust approach to generate a segmentation model suitable over a wide range of shadowgraphy-image conditions. But the two weight-adjusted conditions at least show that regions from the droplet class can be segmented with UNET.

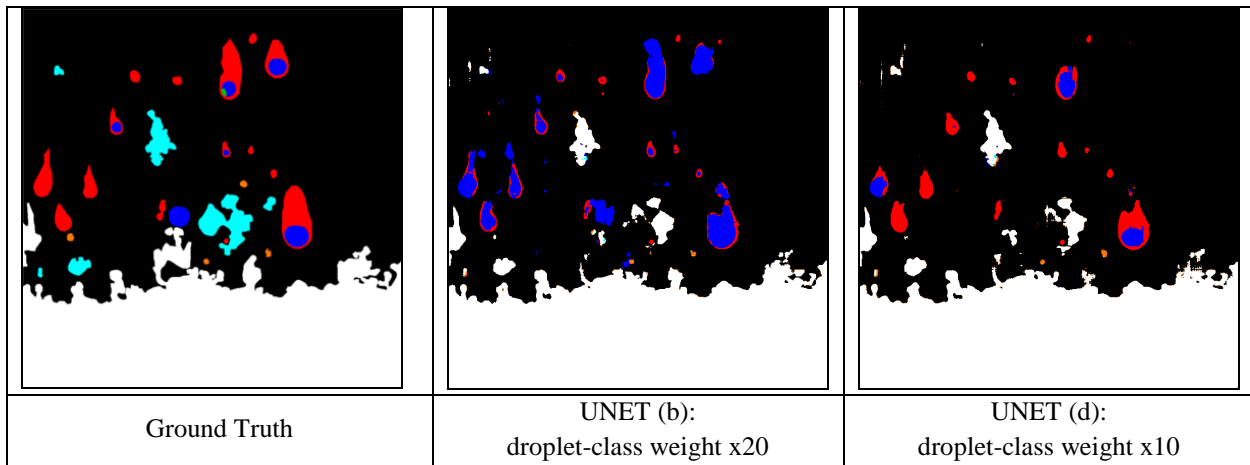


Figure 11. GT (left-hand side) compared to semantic segmentations provided by UNET with two adjusted weight levels for the droplet class (parameters (b) and (c)).

As a more robust adjustment approach, batch normalization has a strong improvement effect as shown in Figure 12 (UNET (d)). Many AI droplets are segmented when batch normalization is used and they show much better correspondence to GT than the previous weight adjustments. Moreover the other object classes are rather appropriately segmented, showing the good level of the overall segmentation. But various droplet regions still show holes or incomplete shapes compared to GT.

Segmentation is even more efficient when the propellant-sample region is segmented previously and “propellant sample / gas region” binary segmentation is included as the third channel. Droplet and flame regions show a much more efficient segmentation. It is also important to see that the “inert Al aggregate” class (shown in cyan) is then segmented better, which is not the case for the other tested conditions. The segmentation quality is really interesting and promising. The UNET condition (e) is used for detection-performance evaluation in the following paragraph.

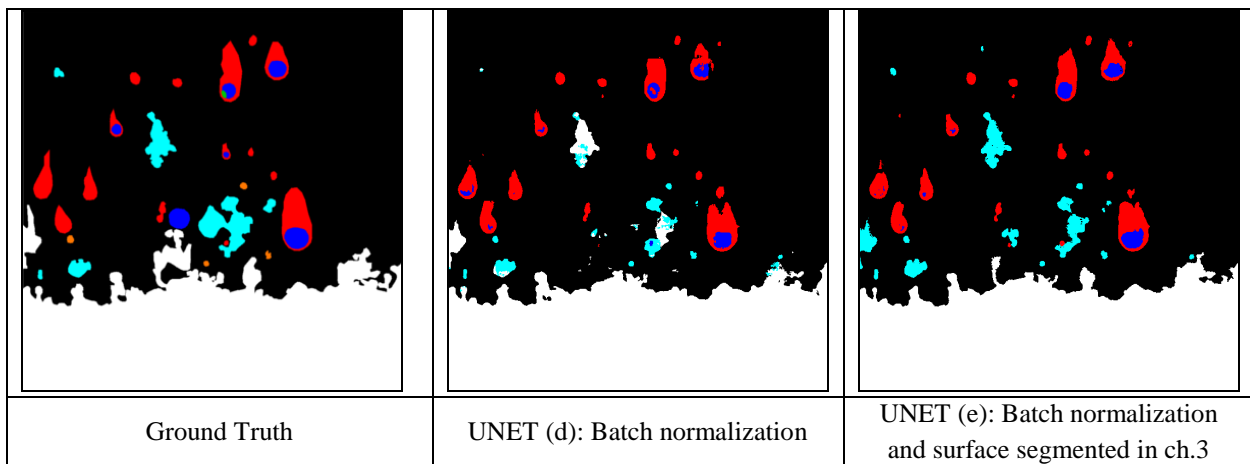


Figure 12. GT (left-hand side) compared to semantic segmentations provided by UNET with two adjusted parameters (parameters (d) and (e)).

#### 4.3. Droplet and flame detection performances

The detection performances are quantified here over the set of 15 “train” images as explained in section 3.1. Precision, Recall and IoU are displayed in Figure 13 for the two detection methods (MSER and UNET). The blue

symbols correspond to “droplet” areas whereas the complete “droplet+flame” regions are displayed in red. As expected from the comments in the previous two sections, UNET provides much better detection performances. The Recall values remain relatively close for UNET and MSER: 65 to 70% for droplets alone and 85 to 90% for the “droplet+flame” regions. Both methods seem quasi similarly capable of segmenting the real targeted regions present in the GT images. However the Precision performances are very good with UNET (> 70%) but extremely poor with MSER (< 35%): UNET segments very few aberrant regions from the images whereas MSER detects any type of homogeneous regions in the images, droplets, flames or meaningless local bright spots...

Moreover the limits of the MSER algorithm with shadowgraphy images is also illustrated by evaluating the detections performances for the 3 images series used to form the “train” set. They are labeled as “... - series” in the legend in Figure 13. The MSER detection performance remains consistent over the 3 images series to segment the “droplet+flame” areas: same Recall, Precision around 35%. However, MSER is less consistent from one image series to another when it comes to segmenting only droplets: Recall values range from 55% to 80% and Precision values vary from 10% to 22%. The MSER input parameters would certainly need to be adjusted for every image series, which is very time consuming. And in any case, such a fine MSER-parameter adjustment might not provide great performances to segment droplets alone since their aspect vary a lot even over a single image. Deep learning approaches really seem very promising for future aluminum-combustion analysis from shadowgraphy images.




Figure 13. Detections performances for the two targeted classes: “droplet alone” (bleu symbols) and “droplet+flame regions” (red symbols).

## 5. Conclusion & perspectives

Experimental data remain sparse for aluminum-droplet combustion in relevant solid-propellant combustion. The present work aims at evaluating automatic detection of aluminum droplets and the surrounding flame on shadowgraphy images acquired during solid-propellant combustion at 1.0 and 2.0 MPa. Two methods were tested with a classical detection approach (MSER) compared to deep learning approach with a CNN (UNET). The detection performances were evaluated by using two sets of annotated images forming the Ground Truth: 45 images from 3 measurement series to adjust or train the segmentation methods and 15 images from the same 3 series to test their performances. Detection performances were significantly better with UNET, particularly in terms of Precision: Precision was larger than 70% with UNET but below 35% with MSER. The targeted objects show a wide range of signal levels and size, which makes it difficult for MSER to recall many real aluminum regions without segmenting many aberrant regions as well. Deep Learning approaches are very promising to provide large and precise sets of data regarding aluminum-droplet size and velocity in a near future.

## Acknowledgments

We acknowledge the financial support of CNES for the development of our shadowgraphy analysis.

## References

- [1] Guéry, J.F., Ballereau, S., Godfroy, F., Gallier, S., Orlandi, O., Della Pieta, P., Robert, E. and Cesco, N. 2008. Thrust Oscillations in Solid Motors. In *Proc. 44<sup>th</sup> AIAA 2008-4979*
- [2] Gallier, S. and Godfroy, F.. 2009. Aluminum Combustion Driven Instabilities in Solid Rocket Motors. *Journal of Propulsion and Power*, Vol. 25, No. 2.
- [3] Casalis, G., Boyer G. and Radenac, E. 2011. Some recent advances in the instabilities occurring in long Solid Rocket Motors. In *Proc. 47<sup>th</sup> AIAA 2011-5642*.
- [4] Cohen, N.S. 1983. A pocket model for aluminium agglomeration in composite propellants, *AIAA Journal*, Vol. 21, No. 5.
- [5] Gallier, S. 2009. A stochastic pocket model for aluminum agglomeration in solid propellants. *Propellants Explos. Pyrotech.*, Vol. 34(2), pp. 97-105.
- [6] Cauty, F. and Erades, C. 2012. Tracking of aluminum particles burning in solid propellant combustion gases by focusing schlieren technique. In *Proc. 15<sup>th</sup> International Symposium on Flow Visualization*.
- [7] Devillers, R.W., Erades, C., Lambert, D. and Belessa, J. 2014. Mesure et suivi de particules, agglomérats et gouttes en combustion au-dessus de la surface d'un propergol en combustion. *14<sup>th</sup> CFTL*, Marseille, France.
- [8] Nogue, M., Devillers, R. W., Le Besnerais, G. and Cesco, N. 2016. Particle detection & size evaluation in solid propellant flames via experimental image analysis to improve two-phase flow simulation in rocket motors. *Space Propulsion conference*.
- [9] Devillers, R.W., Dorval, N., Vilmart, G., Nogue, M., Le Besnerais, G. and Pichillou, J. 2018. Aluminum particle tracking on experimental shadowgraphy and al-plif images to provide velocity data for two-phase flow simulations of solid rocket motors, *Space Propulsion 2018*, Sevilla, Spain.
- [10] Matas, J., Chum, O., Urban, M. and Pajdla, T. 2002. Robust Wide Baseline Stereo from Maximally Stable Extremal Regions. *Proc. Of British Machine Vision Conference*, 384-396
- [11] Nogue, M., Devillers, R.W., Chan Hon Tong, A., Le Besnerais and G., Pichillou, J. 2017. Classification automatique d'images de propergol solide en combustion par utilisation de réseaux de neurones convolutifs, *Conférence GRETSI 2017*, Juan-les-Pins, France.
- [12] LeCun, Y., Bengio, Y., and Hinton, G. 2015. Deep learning, *Nature*, 521(7553), 436-444.
- [13] Ronneberger, O., P. Fischer and T. Brox. 2015. U-net : Convolutional networks for biomedical image segmentation», dans International Conference on Medical image computing and computer-assisted intervention, *Springer*, p. 234–241. 136, 137
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren and Jian Sun. 2015. Deep Residual Learning for Image Recognition, *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016*, Las Vegas, NV, USA, June 27-30, 2016, 770–778
- [15] Chan, T.F. and Vese, L.A. 2001. Active contours without edges, *IEEE TIP* , vol. 10, no. 2, pp. 266 –277
- [16] Devillers, R.W., Le Besnerais, G., Nogue, M. and Cesco, N. 2017. Experimental analysis of solid-propellant surface during combustion with shadowgraphy images: new tools to assist aluminum-agglomeration modelling, *7th European Conference for Aeronautics and Space Sciences (EUCASS)*, Paper #327