



Local anomaly detection in hybrid rocket combustion tests

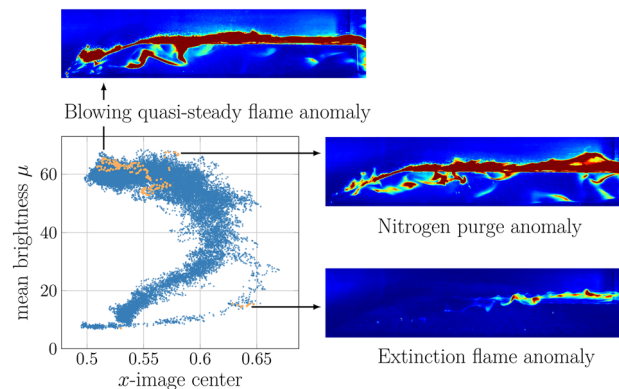
A. Rüttgers¹ · A. Petrarolo²

Received: 19 February 2021 / Revised: 26 May 2021 / Accepted: 2 June 2021 / Published online: 12 June 2021
© The Author(s) 2021

Abstract

Local anomaly detection was applied to image data of hybrid rocket combustion tests for a better understanding of the complex flow phenomena. Novel techniques such as hybrid rockets that allow for cost reductions of space transport vehicles are of high importance in space flight. However, the combustion process in hybrid rocket engines is still a matter of ongoing research and not fully understood yet. Since 2013, combustion tests with different paraffin-based fuels have been performed at the German Aerospace Center (DLR) and the whole process has been recorded with a high-speed video camera. This has led to a huge amount of images for each test that needs to be automatically analyzed. In order to catch specific flow phenomena appearing during the combustion, potential anomalies have been detected by *local outlier factor* (LOF), an algorithm for local outlier detection. The choice of this particular algorithm is justified by a comparison with other established anomaly detection algorithms. Furthermore, a detailed investigation of different distance measures and an investigation of the hyperparameter choice in the LOF algorithm have been performed. As a result, valuable insights into the main phenomena appearing during the combustion of liquefying hybrid rocket fuels are obtained. In particular, fuel droplets entrained into the oxidizer flow and burning over the flame are clearly identified as outliers with respect to the main combustion process.

Graphic abstract



✉ A. Rüttgers
alexander.ruettgers@dlr.de

A. Petrarolo
anna.petrarolo@dlr.de

¹ High-Performance Computing Department, German Aerospace Center (DLR), Institute for Software Technology, Linder Hoehe, 51147 Cologne, Germany

² Satellite and Orbital Propulsion Department, German Aerospace Center (DLR), Institute of Space Propulsion, Im Langen Grund, 74239 Hardthausen, Germany

List of symbols

C	Threshold value (–)
$I(x, y)$	Grayscale pixel intensity at (x, y) (–)
N	Image resolution/ number of pixels (–)
$N_k(a)$	Images with distance $\leq k$ -distance(a) to a (–)
a, b	Images (–)
c_1, c_2, c_3	Stabilization variables (–)
$d(\cdot, \cdot)$	Distance/dissimilarity measure (–)
k	Number of neighbors in LOF algorithm (–)
n	Number of data points (images) in single test (–)

\bar{x}, \bar{y}	Image barycenter coordinates (–)
γ, ν	Hyperparameters of SVM algorithm (–)
μ_a	Mean brightness of image a (–)
σ_a^2	Variance of image a (–)
$\sigma_{a,b}$	Covariance of images a and b (–)

1 Introduction

Despite being born at the same time of solid and liquid propulsion systems, hybrid rocket engines are still not a mature propulsion technology. Only in recent years, there has been a renewed interest in hybrid propulsion due to its characteristic safety, cost and environmental advantages. Due to the fact that the propellants are stored in two different states of matter, hybrid motors are safer than solids. This also contributes to the reduction in the total costs of the engine. Moreover, they are characterized by controllable thrust, including shutoff and restart capability. With respect to liquid engines, they are mechanically simpler and, consequently, cheaper (Karabeyoglu et al. 2011). Finally, their performance is in between those of solid and liquid engines. Unfortunately, due to the characteristic diffusive flame mechanism (the propellants are not premixed, but they need to gasify and mix with each other before being able to react), this technology presents some disadvantages, like poor regression rate performance for conventional polymeric fuels (resulting in low thrust level), time-varying parameters and scaling difficulties. However, most of these drawbacks can be solved by employing a correct design process. In particular, in order to increase the burning rate, the so-called liquefying hybrid rocket fuels, such as paraffin-based ones, can be used. These fast burning fuels are characterized by low viscosity and surface tension and they experience a different combustion mechanism with respect to conventional polymeric fuels (Karabeyoglu et al. 2001). During the combustion, they are able to form a thin liquid layer on the fuel surface, which may become unstable because of the high-speed oxidizer gas flow. The instability waves of the fuel melt layer (Kelvin–Helmholtz waves) may break up and cause the entrainment of fuel droplets into the gas stream. This enables an additional fuel mass transfer and, consequently, an increase in the burning rate (Karabeyoglu et al. 2002). The entrainment mechanism works like a spray injection along the length of the motor, which increases the effective fuel burning area and reduces the blocking effect given from the pyrolysis of the fuel. Unfortunately, this phenomenon has not been well understood yet and still needs to be fully characterized.

For a better understanding of the entrainment process and its relation to the regression rate, optical investigations on the combustion behavior of paraffin-based fuels burning with gaseous oxygen have been performed. The combustion

tests have been recorded with a high-speed video camera, and at the end of the test campaigns, a large amount of data has been collected. Therefore, an automatic analysis method is needed in order to get a faster and deeper insight into the combustion process. An overview of the different flow phases can be obtained by an automatic clustering of the dataset (Rüttgers et al. 2020). This gives essential information on the mean burning behavior. However, in order to completely understand the combustion mechanism, it is important to detect entrained droplets and irregular flow and flame structures, such as Kelvin–Helmholtz waves, appearing during the burning time. The identification of the point in time when such phenomena are showing up could help in understanding the operating conditions leading to an increase in the regression rate. Depending on the fuel composition and oxidizer mass flow rate, a different amount of entrained fuel droplets, as well as different turbulent scales and structures characterizing the flow and the flame, are expected. Even if these turbulent structures only exist within a short period of time, they might strongly affect the overall burning behavior. Therefore, in order to identify specific combustion phenomena, it is necessary to detect “anomalies” in the dataset. In this study, the *local outlier factor* algorithm is applied to the combustion data.

The remainder of this article is organized as follows: First, a short introduction on the *local outlier factor* (LOF) algorithm is given in Sect. 2. In Sect. 3, the experimental setup that is used to obtain the combustion dataset is described. Finally in Sect. 4, the relevance of the chosen algorithm for the current dataset is justified by comparing it to other outlier detection algorithms. Furthermore, a detailed investigation of different crucial parameters used in the LOF algorithm is presented. Lastly, an analysis of the experimental combustion data, based on the outcome of the anomaly detection algorithm, is given. The analysis provides detailed insights into the “anomalies” appearing during the combustion process. In particular, LOF is able to identify specific processes, such as the entrainment of fuel droplets into the gas stream and the flame fluctuations during the transients. The detection of the point in space and time where these phenomena appear during the combustion is important to better understand the whole burning process in liquefying hybrid rocket fuels.

2 Mathematical formulation

Various definitions of outliers exist in the literature, mostly in the area of statistics. One established definition of an outlier (Hawkins 1980) states that it is an observation deviating so much from the other observations that it could be generated by a different mechanism. This definition, however, bases on a global view on the dataset. Consequently, it

primarily covers anomalies that can be described as global outliers. In our application on high-speed image data, the dataset also contains flow features that are only outliers with respect to their local neighborhoods, i.e., the time frame shortly before and after the anomaly. A specific example in this application are represented by melted fuel droplets that only cover small areas of the images and do not significantly change the average image brightness. In the literature, outliers deviating only with respect to their neighborhoods are often denoted as *local outliers* and an algorithm that allows for detecting local outliers is required.

In this manuscript, *local outlier factor* (LOF) (Breunig et al. 2000) is employed as one viable algorithm to detect local outliers. A further justification for using LOF is given in Sect. 4.1 for a simplified 2D representation of the original dataset. In contrast to global outliers, labeling an observation as a local outlier or not depends on the specific application so that a local outlier can be an inlier in other situations. Consequently, local outliers are outliers only up to a certain degree. The LOF algorithm assigns an outlier factor to each image, i.e., a floating point value > 0 , where the size of the numerical value indicates the degree of being a potential outlier.

In the following, the basic definitions of LOF are stated. The core concepts of LOF are the k -distance of an object and the reachability distance (Breunig et al. 2000). Let $d(a, b)$ be a distance measure, while a and b denote images in the dataset. Then, the k -distance(a) is the distance of a to its k th nearest neighbor. The complete set of the $k \in \mathbb{N}$ nearest neighbors includes all images with a distance $\leq k$ -distance(a) and is written as $N_k(a)$. Note that the cardinality of $N_k(a)$, denoted by $\text{card}(N_k(a))$, can be larger than k if several images have the same distance with respect to a . Next, the *reachability distance* from a to b is defined as

$$\text{reach-dist}_k(a, b) = \max \{k\text{-distance}(b), d(a, b)\}. \tag{1}$$

Note that Eq. (1) does not specify a mathematical distance since it is not symmetric. The motivation for Eq. (1) is to define a dissimilarity measure that is robust against statistical fluctuations. These fluctuations occur for very similar images with small distances. For these images, the actual distance is replaced by k -distance(b). The size of the hyperparameter k controls the smoothing effect. For $k = 1$, the usual (unsmoothed) distance measure $d(\cdot, \cdot)$ is recovered. The *local reachability distance* (lrd) of an image a is defined as

$$\text{lrd}_k(a) = \frac{\text{card}(N_k(a))}{\sum_{b \in N_k(a)} \text{reach-dist}_k(a, b)}. \tag{2}$$

Equation (2) defines the inverse of the average reachability distance of image a . Interestingly, if image a has at least k

duplicates, i.e., the reach-dist of all duplicates is zero, Eq. (2) tends to infinity. In our application, this potentially occurs before and after the test, when all images are black and have zero distance with respect to each other. For this reason, these images without information have been removed from the dataset. Finally, the *local outlier factor* (LOF) of a is defined as

$$\text{LOF}_k(a) = \frac{\sum_{b \in N_k(a)} \text{lrd}_k(b)}{\text{card}(N_k(a)) \cdot \text{lrd}_k(a)}. \tag{3}$$

LOF specifies a degree to which a is an outlier based on comparing the local reachability distances of the neighbors. Three different situations can occur in Eq. (3):

$$\text{LOF}_k(a) = \begin{cases} < 1, & \text{if density is higher than neighbors} \\ \approx 1, & \text{if density is similar to neighbors} \\ > 1, & \text{if density is lower than neighbors.} \end{cases}$$

Therefore, if $\text{LOF}_k(a) > 1$, a is a potential outlier. On the other hand, $\text{LOF}_k(a) \leq 1$ indicates potential inliers.

In the following, the advantages and disadvantages of LOF will be discussed. The main disadvantage of LOF is that the ratio defined in Eq. (3) is hard to interpret by the user. Moreover, the threshold value C to separate inliers (images a with $\text{LOF}_k(a) < C$) from outliers (images b with $\text{LOF}_k(b) > C$) is problem dependent and has to be separately determined for each combustion test. The authors of the LOF algorithm have suggested several approaches to allow for a better interpretation of the outlier factor. One possible extension is called *local outlier probability* (LoOP) that scales the LOF ratio to a value range $[0, 1]$ (Kriegel et al. 2009) so that a single threshold value for inlier/outlier separation can be better determined. However, since all combustion tests in this manuscript are evaluated in detail with specific threshold values, this extension is not required here. A further disadvantage is that the correct choice of the hyperparameter k , which controls the smoothness of the distance measure, has to be found. However, since the LOF outcome varies smoothly with k , it is sufficient to determine the correct order of magnitude of the hyperparameter in most applications. In Sect. 4.3, the choice of k for this specific application is considered. On the other hand, LOF has several advantages. The main advantage of LOF is that the algorithm is able to detect local outliers that would be ignored by global anomaly detection algorithms. Furthermore, LOF only requires that $d(\cdot, \cdot)$ is a dissimilarity and not a distance function, i.e., the triangle inequality is not required. Finally, LOF is able to cope with regions of different densities. If, for instance, several high-density regions in a dataset exist, LOF might label a point with a medium density in these clusters as

an outlier. This coincides with our application since we expect different combustion phases, such as an ignition, a steady combustion and an extinction phase, which all form high-density clusters. Even then, LOF is still able to detect outliers in all different phases.

The quality of the algorithm for finding anomalies primarily depends on the quality of the dissimilarity measure $d(\cdot, \cdot)$. An adequate choice simplifies a separation of abnormal images from the images showing regular flow conditions. Furthermore, since each pixel of the image dataset represents a further dimension, the algorithm has to deal with high-dimensional data. In the literature, a lot of effort is spent on anomaly detection for high-dimensional data (Zimek et al. 2012). This is due to the fact that distance or dissimilarity measures are often less effective in high-dimensional problem spaces. For this reason, two different dissimilarity measures will be used in this manuscript: an Euclidean norm (discrete L^2 -norm) and a dissimilarity measure that bases on a structural similarity index measure (SSIM) (Wang et al. 2004). According to the literature, L^1 and L^2 are the most adequate L^p -norms in high-dimensional spaces (Hinneburg et al. 2000). On the other hand, SSIM has been specifically developed for image comparisons and it is one of the most used algorithms in imaging science. Let a and b denote images in the dataset, the SSIM index of these images is defined as

$$\text{SSIM}(a, b) = \underbrace{\frac{2\mu_a\mu_b + c_1}{\mu_a^2 + \mu_b^2 + c_1}}_{\text{luminance } l} \underbrace{\frac{2\sigma_a\sigma_b + c_2}{\sigma_a^2 + \sigma_b^2 + c_2}}_{\text{contrast } c} \underbrace{\frac{\sigma_{ab} + c_3}{\sigma_a\sigma_b + c_3}}_{\text{structure } s} \quad (4)$$

with μ_a as the mean brightness of a (analogously for image b), σ_a^2 as variance of a , $\sigma_{a,b}$ as the covariance of a and b and c_1 , c_2 and c_3 as variables to stabilize the division with small denominators. c_1 , c_2 and c_3 only depend on the dynamic range of the pixel values and are identical in all tests since images with an 8-bit dynamic range are always considered. As illustrated in Eq. (4), SSIM is the product of three comparisons that investigate the difference in luminance l , in contrast c and in structure s of a and b . SSIM is a similarity measure ranging from zero to one. A dissimilarity function, the structural dissimilarity DSSIM, is obtained by a linear transformation according to

$$\text{DSSIM}(a, b) = (1 - \text{SSIM}(a, b))/2. \quad (5)$$

Since DSSIM bases on a threefold comparison of the images, it has a larger computational complexity than an Euclidean norm. In Sect. 4.2, the performance of both measures to find anomalies in the test dataset will be compared.

Note that both distance measures do not contain knowledge from combustion theory and detect all possible image changes. Using domain knowledge, i.e., knowledge on

how an image that shows an anomaly differs from an inlier images, more adapted distance measures, which are able to better separate inliers from outliers, could be constructed.

3 Combustion tests

The combustion tests were performed at the German Aerospace Center (DLR), Institute of Space Propulsion in Lampoldshausen, at the test complex M11. An already existing modular combustion chamber, used in the past to investigate the combustion behavior of solid fuel ramjets (Ciezki et al. 2003), was adjusted and used for the test campaigns at atmospheric pressure. A side view of the whole combustion chamber setup is shown in Fig. 1. The optically accessible combustion chamber is 450 mm long, 150 mm wide and 90 mm high. The flow straighteners with the pre-chamber are in total 450 mm long and the post-chamber is 150 mm long. The oxidizer main flow is entering the combustion chamber from the left, after having passed two flow straighteners, in order to get homogeneous flow conditions in the burner. The mass flow rate is adjusted by a flow control valve and it is measured with a Coriolis flow meter. A high-frequency static pressure sensor is mounted in the combustion chamber. Ignition is done via an oxygen/hydrogen torch igniter from the bottom of the chamber. At the end of the combustion, after having closed the main oxidizer valve, nitrogen is used for purging. A test sequence is programmed before the test and is run automatically by the test bench control system. More details about the test bench and test settings are given in Kobald et al. (2015), Petrarolo and Kobald (2016), Petrarolo et al. (2018).

In the framework of this research, all tests were run at atmospheric pressure and with an oxidizer mass flow ranging from 10 to 120 g/s. Combustion tests were performed using a single-slab paraffin-based fuel with a 20° forward facing ramp angle (see Fig. 2), in combination with gaseous oxygen. Two different fuel compositions were analyzed in this study: As baseline, pure paraffin 6805 from the manufacturer Sasol Wax was used; furthermore, 5% of a commonly available polymer was added to it in order to change




Fig. 1 Side view of the atmospheric combustion chamber setup, adapted from Thumann and Ciezki (2002)



Fig. 2 Fuel slab configuration used in this research, before (top) and after (bottom) combustion test

Table 1 Test matrix

Test no.	Fuel		\dot{m}_{Ox} (g/s)		CH* filter
	6805	6805+5% polymer	50	120	
203		✓	✓		
276	✓			✓	
284	✓		✓		✓

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. The frame rate was set to 10,000 frames per second, thus being able to catch all the important combustion and flow phenomena. The shutter time of the camera was adjusted for each test, according to the test conditions (especially brightness) and position of the camera. Tests were also performed using a CH* chemiluminescence imaging technique, with a band-pass filter centered around 431 nm placed in front of the camera. The excited CH* molecules emit photons around this wavelength, when they relax back to a lower energy state. Since high CH* concentration exists only in the main reaction zone, the resulting images provide a good indication of the instantaneous flame sheet location and topology.

In this study, three combustion tests have been analyzed. The test matrix is presented in Table 1.

4 Results and discussion

4.1 Choice of anomaly detection algorithm

A large number of algorithms for anomaly detection exist in the literature (Schwabacher et al. 2009). In Sect. 2, basic principles of *local outlier factor* (LOF) (Breunig et al. 2000),

an algorithm for local outlier detection, are described. This section justifies the choice of LOF for a two-dimensional benchmark problem, more precisely for a low-dimensional representation of the experimental image dataset. The advantage of this approach is that the outcome of the different anomaly algorithms can be better visualized and compared if only two features of each image in the test set are analyzed. However, the choice of these two features is essential to retain the important information of the uncompressed image dataset (8 GB per test). Furthermore, to avoid any misunderstanding, it is noted that the final results in Sect. 4.4 base on an analysis of the uncompressed dataset and not on the low-dimensional representation used here.

The features used for this benchmark are the mean image brightness μ of all N -pixel images and the x -component of the normalized image barycenters. Since the high-speed videos have 8-bit to capture information on the brightness of each pixel, the mean brightness μ ranges from 0 to 255. μ essentially contains information of the combustion intensity. The computation of the image barycenters is performed as follows: Let $I(x, y)$ denote the grayscale pixel intensity at each pixel of the two-dimensional image in coordinate direction x/y , then the image barycenter (\bar{x}, \bar{y}) is defined as

$$(\bar{x}, \bar{y}) = \left(\frac{M_{10}}{M_{00}}, \frac{M_{01}}{M_{00}} \right) \tag{6}$$

with

$$M_{ij} = \sum_x \sum_y x^i y^j I(x, y). \tag{7}$$

The values for the normalized image barycenters range from 0 to 1 in both coordinate directions and give information about the horizontal and vertical position of bright pixel values. Since the combustion chamber is horizontally oriented, as shown in Fig. 1, the flow field primarily moves in the x -direction. Therefore, the x -component \bar{x} of Eq. (6) will be used for the algorithm benchmark since it contains more information about the flame movement, while the y -component \bar{y} will be neglected for the sake of simplicity. Based on the two features (\bar{x}, μ) , Fig. 3 shows a low-dimensional representation of the three tests considered in this manuscript. This representation will be used to compare the outcome of different anomaly detection algorithms.

In the following, four algorithms for outlier detection are applied to the dataset shown in Fig. 3. In addition to LOF, *Elliptic Envelope*, *One-Class Support Vector Machines* (One-Class SVM) and *Isolation Forest* are considered. For all these algorithms, the implementation from the scikit-learn (Pedregosa et al. 2011) is used. Briefly, *Elliptic Envelope* fits a multivariate Gaussian distribution to the data that has the form of an ellipse in 2D. A *One-Class SVM* fits a hyperplane in a high-dimensional space to the data and

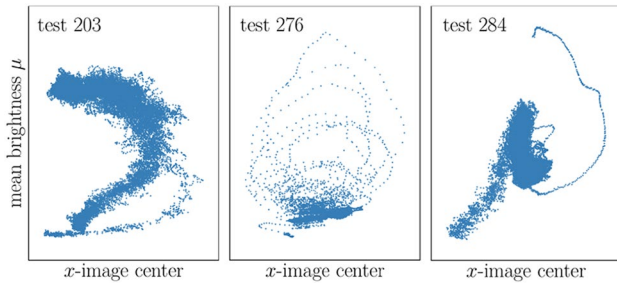


Fig. 3 Two-dimensional representation of each test that bases on the x -component of the normalized image barycenters \bar{x} and on the mean image brightness μ

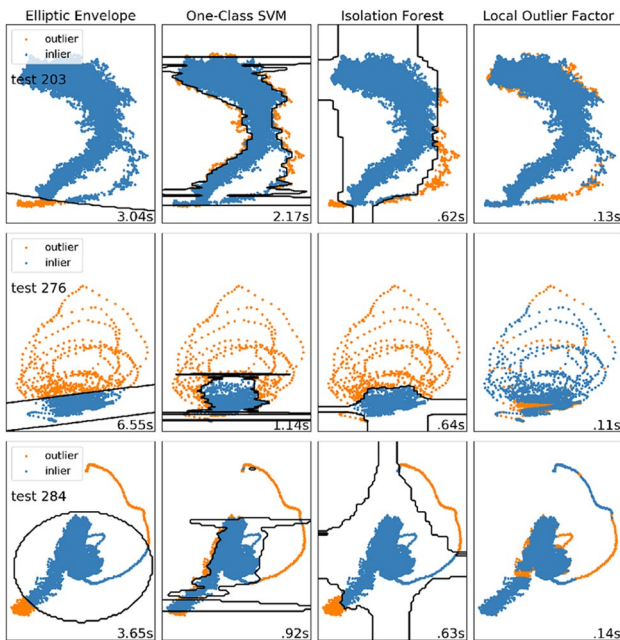


Fig. 4 Performance comparison of four different anomaly detection algorithms on a low-dimensional representation of the test datasets. The runtime of the algorithms in seconds is printed on the lower right side of each benchmark test

decides if additional data belong to the pre-trained class (inlier) or not (outlier). Finally, Isolation Forest detects outliers by an isolation procedure applied to the data.

All algorithms have hyperparameters in the implementation from scikit-learn that need to be adapted to the specific problem to allow for an efficient outliers detection. The *contamination*, i.e., the expected percentage of outliers, is a hyperparameter for the algorithms Elliptic Envelope, Isolation Forest and LOF. Figure 4 visualizes the outcome of all algorithms for a contamination of 0.02. Note, however, that the percentage of outliers in case of the LOF algorithm is used only for determining the threshold value C separating inliers from outliers. Therefore, the contamination hyperparameter only affects the post-processing of the LOF

numerical values. In contrast to the other algorithms, a One-Class SVM has a hyperparameter ν that specifies an upper bound on the fraction of training errors. Basically, ν can be used in a similar way to the contamination hyperparameter and it is also set to 0.02. Furthermore, a SVM has a further hyperparameter γ which affects the variance of a Gaussian kernel function. The best output in the test has been achieved with $\gamma = 0.1$. All other hyperparameters are used with their default values in scikit-learn. As mentioned before, the important hyperparameter for LOF is k defined in Eq. (3) that affects the number of neighbors that is considered by the algorithm (see Sect. 2). The default value $k = 20$ from scikit-learn was also used for the following algorithm benchmark.

Furthermore, Fig. 4 has a black boundary line between inliers and outliers, the so-called decision boundary. A decision boundary can be computed for the algorithms Elliptic Envelope, SVM and Isolation Forest but not for LOF since it is a local algorithm. The decision boundary is important to directly classify new data points as inliers or outliers. However, since the number of images in each combustion test is fixed (30,000 images) and there is no follow-up data, a decision boundary is not required in this application.

Since the data points for the benchmark shown in Fig. 3 are unlabeled, i.e., the problem is unsupervised, it cannot be precisely stated if the labeling of outliers visualized in Fig. 4 is correct. Furthermore, due to the usage of a high-speed video camera, there are no data points that are clearly isolated from the others, i.e., neighboring images have a small distance, and consequently, all data points form a large single cluster. The expected outcome is that potential outliers might occur at the side arms on the right-hand side of test 203 and test 284. These frames represent the extinction phase, where the conditions in the combustion chamber (temperature, pressure, oxidizer mass flow) change very much in a short period of time. In this phase, it is very likely to find outliers, such as a sudden change in the flame brightness, shape and position or the appearance of burning fuel droplets far from the flame. Moreover, some outliers could be found in the small arc on the right-hand side of test 284. These frames belong to the steady combustion phase where usually the burning conditions are constant, and therefore, no anomalies are expected. Consequently, it is most likely that they will show a disturbance to the main burning process. Concerning test 276, potential outliers might occur in the arcs on the upper part of its 2D representation. They represent the ignition phase of the combustion, where a lot of liquid fuel droplets are burning in the whole combustion chamber volume. The combination of a very high oxidizer mass flow with a very low fuel viscosity contributes to increase the entrainment mechanism, as explained in Sect. 1. This makes test 276 quite different from the others (as clearly visible in the two-dimensional representation in Fig. 3):

The strong variation in burning conditions during the ignition phase overtakes all the other combustion phases. Even the steady-state phase, where the majority of the frames are included and which usually dominates, seems to be less important in this case. Some outliers are also expected to be found in the small arc on the bottom of test 276, which represents the extinction phase.

As shown in Fig. 4, the algorithms label different points as outliers. In test 203, the outcome of the Elliptic Envelope algorithm strongly differs from the other algorithms. From our point of view, SVM, Isolation Forest and LOF seem to perform better here. In test 276, the algorithms Elliptic Envelope, SVM and Isolation Forest separate the arcs on the upper part of the 2D representation from the main combustion phase. The situation differs with LOF, which labels different regions on the arcs and in the main combustion phases as potential outliers. This illustrates the local approach of the LOF algorithm. Here, all algorithms predict plausible outlier candidates. Finally, in test 284, the lower left part of the point set is labeled as outliers by the Elliptic Envelope and the Isolation Forest algorithms. Since the point density in the lower left corner is high, it is unlikely to detect outliers there. On the other hand, SVM and LOF primarily detect outliers on the arc on the right-hand side of the dataset. Furthermore, LOF also labels a minor arc on the right-hand side as potential outlier. Again, this is due to the local approach of the algorithm. Therefore, it is likely that LOF performs better in test 284.

The comparison of the algorithms for the benchmark in Fig. 4 suggests the following order of plausibility: LOF and One-Class SVM indicate very plausible outliers, Isolation Forest seems to perform slightly worse and finally Elliptic Envelope predicts some unlikely outliers. The primary reason for favoring LOF over SVM is that LOF essentially has a hyperparameter that can be determined more easily. The numerical output of LOF is only affected by the number of neighbors k and only changes moderately if k is changed from, for instance, $k = 20$ to $k = 30$. In contrast to this, the SVM has two hyperparameters ν and γ that are difficult to determine. Moreover, the output of the SVM becomes very unrealistic if ν and γ are set to unfavorable values. A second reason for choosing LOF is that the runtime of the algorithm, printed on the lower right side of each benchmark test in Fig. 4, is comparatively low. This becomes important in the analysis of the full image dataset discussed in Sect. 4.4. Finally, in the image set that has been considered, the LOF outlier scores were comparatively robust with respect to downsampled images, i.e., a moderate downsampling of the image resolution from 1024×336 pixels to 512×168 pixels or even 256×84 pixels hardly effected the LOF results in a previous study for test 276.

4.2 Investigation of different distance metrics

In the following, the performance of two distance metrics, an Euclidean distance metric and a dissimilarity function DSSIM, defined in Eq. (5), that bases on structural similarity (SSIM), will be compared. SSIM is a similarity measure that ranges from 1 (both images a and b are identical) to -1 . The values of SSIM become negative if the covariance $\sigma_{a,b}$ used in Eq. (4) is negative. It is further noted that SSIM is zero if one image is black (all pixel values are zero) and the other image is white (all pixel values are 255). As a result, DSSIM ranges from zero (a and b are identical) to one.

On the other hand, an Euclidean distance metric delivers unbounded positive values. This complicates the comparison with DSSIM. For this reason, the Euclidean distance measure has been normalized to the range $[0, 1]$ according to

$$\|a - b\|_{2,\text{normalized}} = \frac{1}{\tilde{C}} \left(\sum_{i=1}^N (a_i - b_i)^2 \right)^{1/2} \quad (8)$$

with $\tilde{C} = (N \cdot 255^2)^{1/2}$. Here, the summation a_i and b_i is over the number of pixels (8-bit integer values) in the images a and b . Since the upper part of the camera usually does not contain information from the test, the images are cropped in height by 30% to reduce the computing time of the LOF algorithm. Note that the normalization constant \tilde{C} in Eq. (8) is the Euclidean distance between an 8-bit black image (all pixel values are zero) and an 8-bit white image (all pixel values are 255) with a resolution of N pixels. Furthermore, the normalization with \tilde{C} in Eq. (8) is primarily used to better compare the distance metrics and does not affect the LOF outcome in Eq. (3).

For a better comparison, Fig. 5 illustrates the performance of both distance measures for detecting satellite droplets and its robustness against vertical movement. For this purpose, an image from test 284 at $t = 2.2171$ s was modified. Figure 5a shows that the distance from this image to itself is zero in both distance measures. Next, Fig. 5b, c contains 1 and 20 additional satellite droplets, manually added to the image with a raster graphics editor. In both cases, the Euclidean distance and DSSIM indicate a small distance, since only a small number of pixels has been modified compared to the original image. Nevertheless, in Fig. 5c the distance calculated by the Euclidean metric is 0.04, which is roughly 14 times larger than the distance measured by DSSIM. This might indicate that the Euclidean metric is more receptive for the detection of satellite droplets than DSSIM. Second, this might also indicate that the Euclidean metric leads to a stronger separation of the dataset which results in larger LOF values. Next, Fig. 5d, e shows the original image shifted in the vertical direction by one pixel (d) and by 100 pixels (e). All new pixels that have been added due to the y -shift

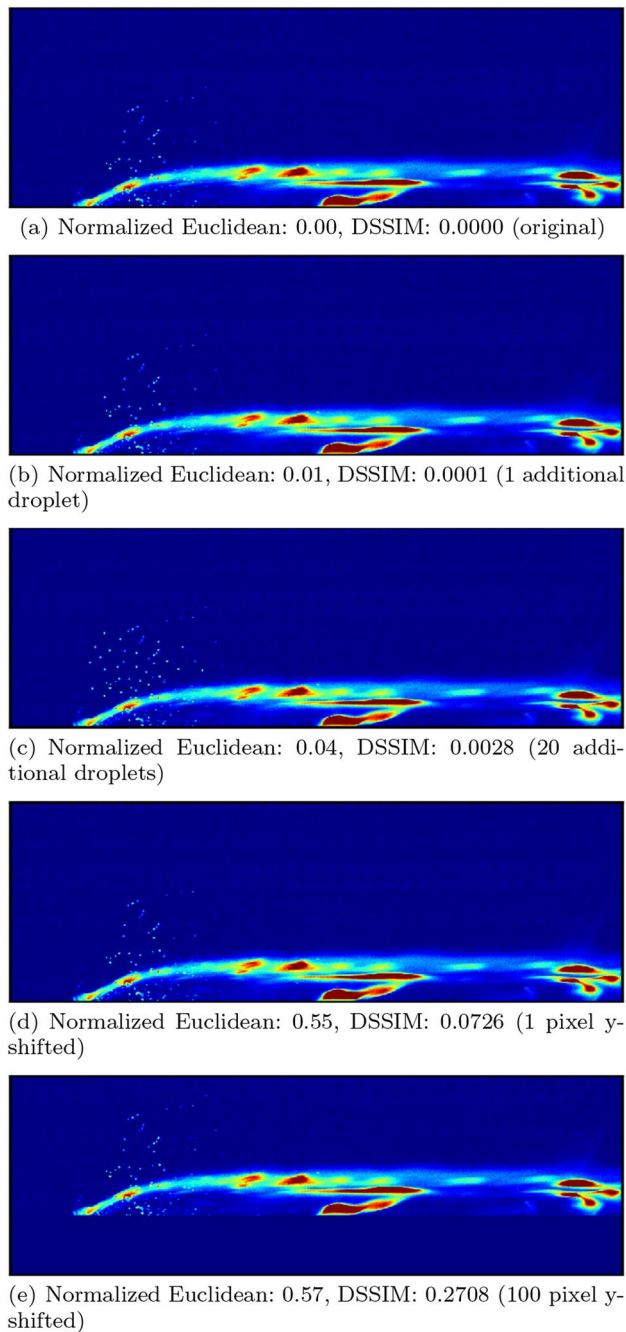


Fig. 5 Efficiency of the normalized Euclidean metric and the structural dissimilarity distance function (DSSIM) in detecting satellite droplets and the robustness against vertical movement

are filled with zeros (black color). In both cases, a large distance is computed by both distance measures since the y-shift affects all pixel values. Again, the Euclidean metric shows larger distances than DSSIM but the distances in (d) and (e) are almost identical. Consequently, although a 1 pixel y-shifted is hard to detect with the naked eye, it already strongly affects the Euclidean distance prediction. Therefore,

an Euclidean distance measure is prone to camera shake that might occur during a combustion test. If camera shake occurs, the images have to be aligned with subpixel accuracy in a preprocessing step. However, the high-speed camera does not vibrate in our test setup shown in Fig. 1 since it is solidly fixed on the ground. On the other hand, DSSIM computes a small distance in Fig. 5d which indicates that, according to this distance measure and similar to human recognition, the image is almost identical to the original image. Therefore, DSSIM seems to be more adequate in the case of camera shake.

Finally, we compare the runtime of an Euclidean distance measure and of DSSIM. The LOF algorithm requires a distance matrix using, for instance, one of the dissimilarity measures $d(\cdot, \cdot)$ described before. The square matrix is of size $n \times n$ and contains the pairwise distances of all $n = 30,000$ images in the dataset. Since the matrix is symmetric with zero diagonal about $n^2/2 - n \approx 4.5 \times 10^8$ image comparisons have to be computed. Basically, this requires the largest percentage of the computing time. On a Linux workstation with 16 cores and 128 GB main memory, the computing time for an Euclidean distance matrix is in the order of 3–4 minutes. The situation is different for DSSIM. In this manuscript, an implementation of SSIM from scikit-image (van der Walt et al. 2014) has been used but it has been optimized for better performance. The image comparison performed in Fig. 5b requires about 34 s on a single core using the reference implementation from scikit-image and about 20 s after optimization. To further reduce the computing time, the matrix computation was parallelized with Dask (Dask Development Team 2016). After parallelization, the computation of the full distance matrix of one test with $n = 30,000$ images took about 4 days on one node with 56 cores of a DLR cluster at the Institute for Software Technology. Note that the parallel efficiency is almost optimal since the matrix entries are independent from each other such that there is no communication overhead.

As a conclusion, both distance measures will be used in this study since they offer different advantages. On the one hand, an Euclidean distance has a lower computational complexity and might be better to capture satellite droplets. On the other hand, DSSIM is more robust to camera shake and is more adapted to human recognition, which could potentially lead to better results, but it has a higher computational complexity. Since camera shake does not occur in our test setup, the robustness of DSSIM is not required in our application.

4.3 Analysis of the number of neighbors

The outcome of the LOF algorithm is essentially affected by one hyperparameter k , which sets the number of neighbors that is considered (see Sect. 2). Therefore, k controls the smoothing of the distance measure to avoid statistical

fluctuations. However, since LOF becomes a global algorithm if k tends to the number of data points n , the size of k should be limited.

For the analysis of the hyperparameter k , Fig. 6 plots statistical parameters for the LOF outcome using the two-dimensional representation of the tests 203, 276 and 284 (see Fig. 3) as input. More precisely, the 2D point set derived from the full dataset is used as input for the LOF(k) algorithm, where k ranges from 1 to 50. Figure 6 displays the

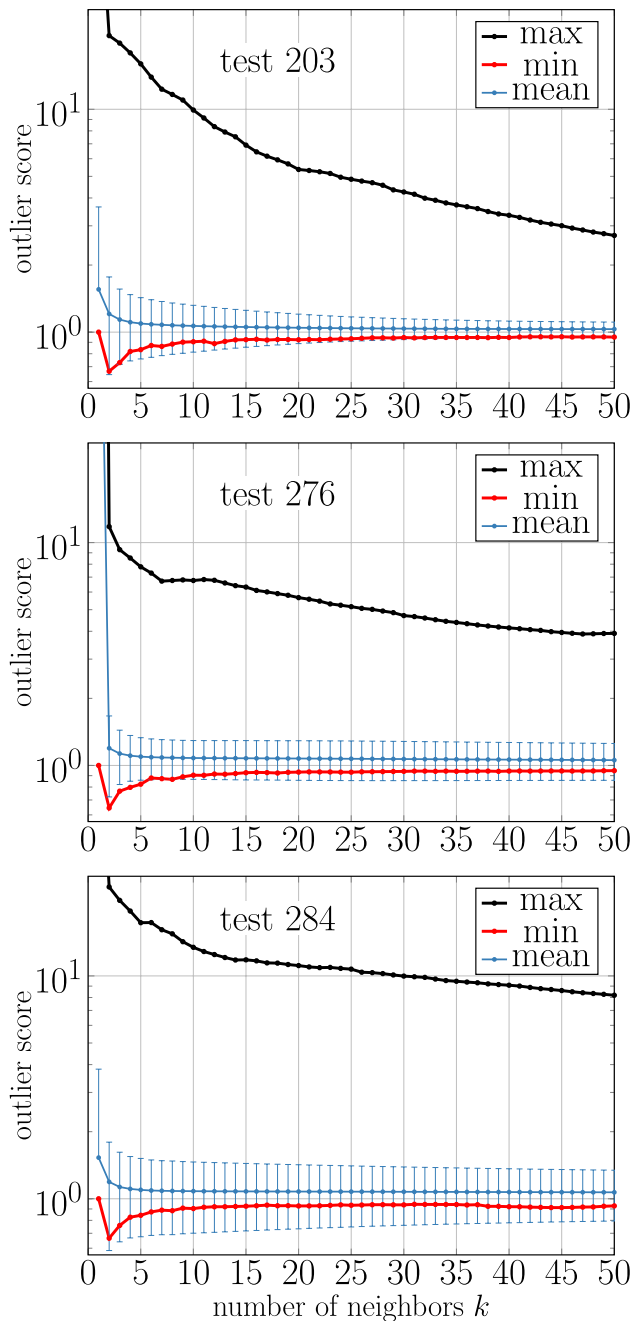


Fig. 6 Effect of the number of neighbors on the outlier score using the low-dimensional representation as an input

maximum (black line), the minimum (red line) and the mean (blue line) with standard deviation (blue vertical lines) of all $n = 30,000$ outlier scores. In general, the result is similar in all tests. For $k < 10$, some outlier scores become very large in the order of $\mathcal{O}(10)$ and the variance of the scores is high. On the other hand, for $k > 10$ the variance of the result is strongly reduced and tends to decrease further with increasing k . Interestingly, test 284 has larger maximum outlier scores compared to the other tests. Therefore, k should be larger than 20 to avoid strong statistical fluctuations. In contrast to this, the authors of the LOF algorithm recommend to set k smaller or equal to the maximum number of neighboring samples that could still be labeled as outliers (Breunig et al. 2000). In this manuscript, a set of about 40-50 data points could still show an anomaly that should be detected. But, if this set consists of more than about 50 points, the anomaly becomes a cluster, i.e., a separate combustion phase.

Therefore, from our point of view, the range of sensible values for k ranges from 20 to about 50 in our application. In the literature, it is recommended to use a heuristic such that only a coarse parameter range of k is required. For each image in the dataset, the heuristic approach computes the maximum outlier score over a given range of different k -values. More precisely, using a range for $k \in \mathbb{N}$ from 20 to 50 the heuristic is to employ

$$\text{LOF}(a) = \max_k \{ \text{LOF}_k(a) | 20 \leq k \leq 50 \} \tag{9}$$

as LOF result for all images a in the test. Consequently, since the maximum outlier score over a range of k -values is calculated, Eq. (9) tends to amplify anomalies. This might increase the number of potential outliers in the full image dataset but might also increase the number of false positive results, i.e., inliniers are falsely detected as outliers. However, since the primary intention of this manuscript is to detect all potential anomalies, the heuristic from Eq. (9) will be used in the following section.

4.4 Analysis of the combustion

The three most representative tests will be analyzed in this work. Test 203 is a typical combustion test with a high video quality and it is considered as the baseline to which all the other tests will be compared. Test 276 has been chosen for the presence of a high number of droplets during the whole ignition transients. Test 284 has been recorded with a band-pass filter placed in front of the lens of the high-speed camera and its analysis gives information on the influence of the CH* filter on the combustion tests (see Table 1 for a summary). The results obtained from the combustion visualizations give many insights into the hybrid combustion process and allow to detect the main phenomena influencing

the burning behavior in this kind of engines. In particular, flame blowing and fuel droplets entrainment are detected as anomalies by the LOF algorithm. It is important to underline that, as explained in Sect. 4.3, the number of neighbors k strongly influences the results, since it controls how “local” the identified outliers will be. In the extreme case of a very high k value, a whole combustion phase might be identified as an anomaly. For this analysis, as previously said in Sect. 4.3, the k value has been varied between 20 and 50, and for each image, the maximum LOF value according to Eq. (9) has been considered. Interestingly, a previous analysis of the distribution of k values on the right-hand side of Eq. (9) showed that all different values of k occur at least 600 times as argmax in the 30,000 images data set but the boundary region, i.e., k close to either 20 or 50, had a higher frequency to deliver the maximum LOF value. Furthermore, during the post-processing phase, the user has to wisely choose the threshold value C separating inliers from outliers. Since the algorithm detects specific phenomena deviating from the main combustion process as well as random background noise as anomalies, it is up to the user to decide which outliers are important for the analysis. One viable approach to determine C is a manual analysis of all LOF score peaks by a domain expert. Then, C can be chosen such that the separation coincides with the manual analysis.

4.4.1 Test 203

As previously said, test 203 is a typical combustion test and it will be considered as the baseline. Looking at Figs. 7 and 8, the three main combustion phases (ignition, steady state and extinction) can be recognized respectively in the three different colored areas in the DSSIM distance matrix and in the three peaks in the LOF outlier score. Each of these areas potentially contains some outliers, which can be recognized with different distance metrics.

The threshold values for this test are $C = 1.5$ for the Euclidean metric and $C = 1.2$ for DSSIM. These values have been determined with an analysis of the outlier score distribution shown in Fig. 8. Similar to Sect. 4.1, Fig. 9 shows a simplified representation of test 203 with different threshold values C . To avoid any misunderstanding, it is noted that this low-dimensional representation is only used here for a better illustration and that the labeling of the outliers results from a LOF analysis of the full image dataset. As expected, the lower the value of C , the more data points are labeled as outliers. If the threshold C is very high (third row in Fig. 9), the Euclidean metric ($C = 1.7$) primarily detects anomalies on the side arm at the bottom (extinction phase) of the 2D representation (colored in red). On the other hand, DSSIM ($C = 1.3$) detects the strongest anomalies (highest outlier score) in an earlier combustion phase on the top of the 2D representation (colored in orange). Interestingly, for low

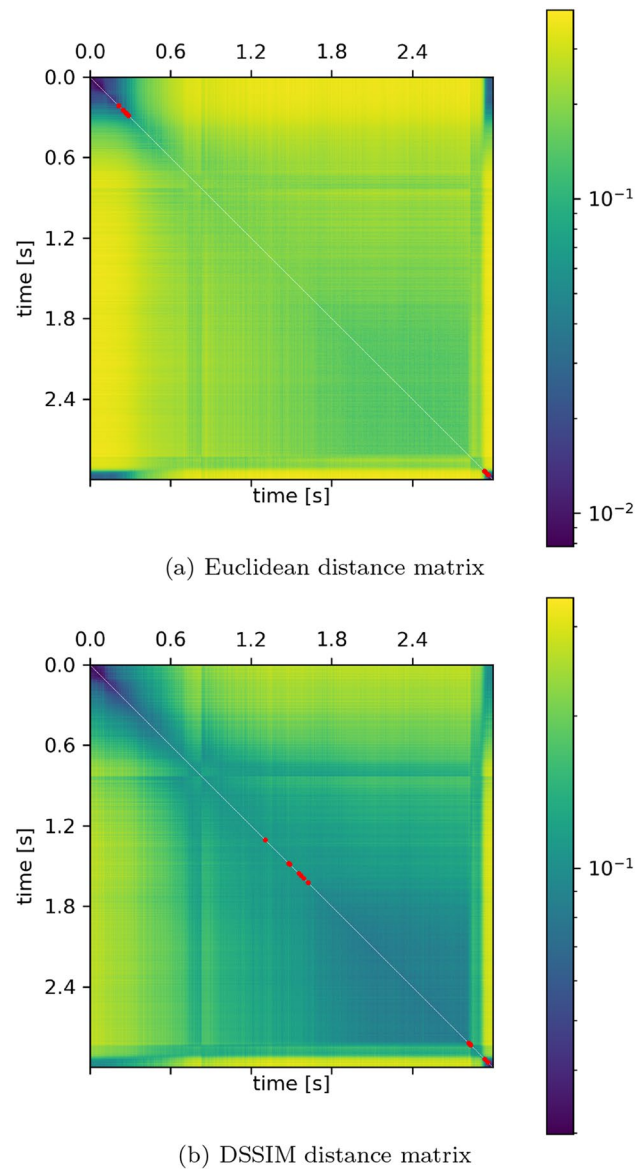


Fig. 7 Euclidean and DSSIM distance matrix of test 203. Potential anomalies according to the threshold values 1.5 and 1.2 as shown in Fig. 8 are highlighted in red on the matrix diagonal

threshold values (first row in Fig. 9) both distance metrics label similar points in the side arm at the bottom as outliers but detect different outliers in the area where μ is large.

The Euclidean metric, which is more sensitive to local pixel luminosity change, such as burning fuel droplets (see Sect. 4.2), is able to detect strong outliers in the two transition phases. During ignition, the flame slowly increases its brightness and moves forward on the fuel slab surface. Furthermore, some burning droplets are visible over the flame. On the other hand, during extinction, the flame conditions are changing so quickly that in less than 0.2 s the flame is completely extinguished. The Euclidean metric is able to recognize satellite droplets as anomalies, as well as




Fig. 8 LOF outlier score over time of test 203. The threshold values C are 1.5 (Euclidean metric) and 1.2 (DSSIM)

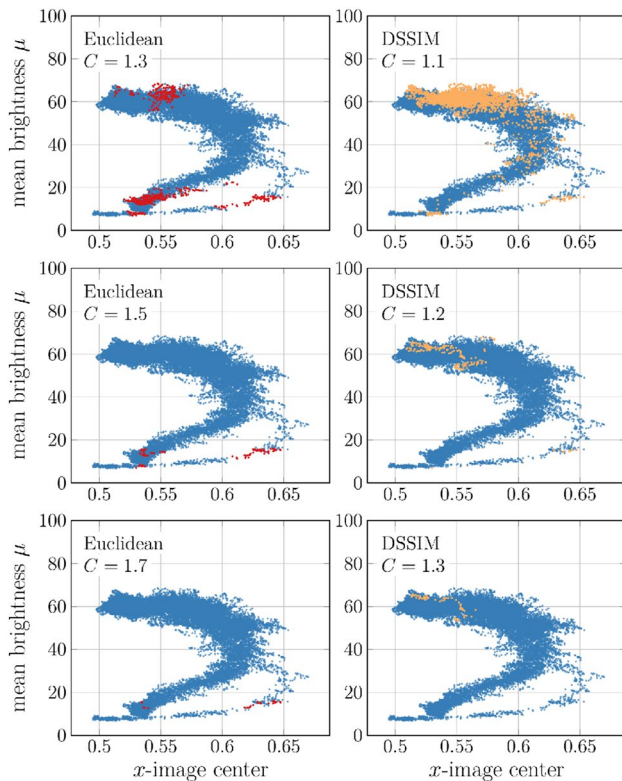


Fig. 9 Visualization of the anomalies found in test 203 in a low-dimensional (2D) representation of the image dataset. For both distance measures, Euclidean distance (left) and DSSIM (right), different threshold values C are compared

horizontal flame movements and vertical flame fluctuations (especially when the flame is not yet completely developed, and therefore, the frame brightness is still quite low).

On the other side, DSSIM, which is more sensitive to a general frame luminosity variation, is able to detect weak outliers in the steady-state phase (also visible as smaller peaks in the Euclidean LOF score values) and stronger ones in the extinction phase. During steady state, no dominant anomalies arise in the flame and in the combustion conditions. The only deviations from the main burning process are due to flame blowing and fluctuations along the fuel slab surface. This brightness shifting in the vertical direction is recognized by both metrics (see Fig. 5), but, due to the chosen lower C value, it is considered as anomaly only in the DSSIM. Moreover, DSSIM is also able to recognize the fast flame brightness variation during the extinction phase. In particular, it recognizes the start of the nitrogen purge at around 2.8 s, directly visible from the disturbance caused to the main flow on the bottom left of the fuel ramp (see Fig. 10d). Moreover, DSSIM as well as the Euclidean metric detects the moment when the main flame actually feels the disturbance and the variation of the burning conditions (at around 2.9 s, cf. Fig. 10e). On the other hand, the slow changes during ignition are not detected as strong anomalies by DSSIM. In this case, the global frame luminosity does not change much among neighboring frames. It is most likely that, by increasing the number of analyzed neighbors, some outliers could be found in the ignition phase also by the DSSIM distance measure. Once more, it is important to remind that it is up to the user, depending on the analyzed problem, to decide whether only data with a high LOF score have to be considered as outliers or whether also “weaker” outlying data identify important anomalies. For instance, a smaller peak in the Euclidean LOF score (see Fig. 8) is visible at 0.8281 s. By looking at the video, it is possible to realize that this peak detects a side flamelets on the fuel slab, being just a noise for the main flame. Consequently, it will not be considered as outlier in this analysis. The same is also valid for the small peak at 0.22515 s in the DSSIM outlier score.

4.4.2 Test 276

Test 276 well represents the entrainment mechanism, due to the combination of a high oxidizer mass flow rate with a low fuel viscosity. The chosen threshold value for this test is $C = 1.5$ for both metrics. During the ignition phase, before the flame is fully settled on the slab, a very high number of burning fuel droplets is clearly visible (see Fig. 11). In this phase, the combustion is characterized by huge and fast variations of flow conditions, which are clearly recognized as anomalies by both metrics. Since the scores of these outliers are very high, the rest of the dataset

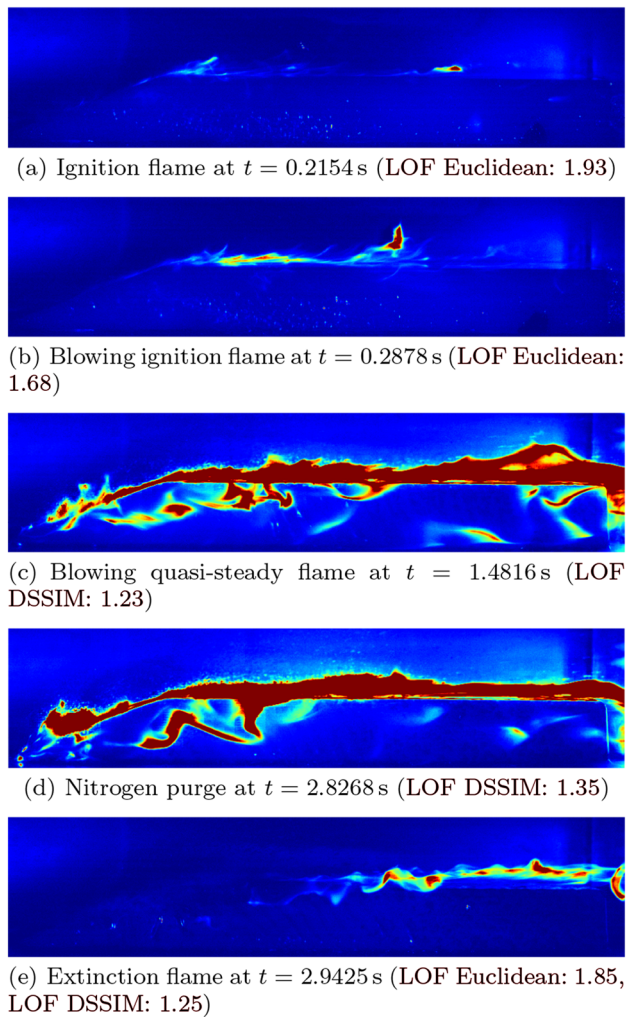


Fig. 10 Representative selection of anomalies found in test 203

appears to be “flattened” and no other anomalies are found. This is especially visible with DSSIM in Figs. 12 and 13: The DSSIM distance matrix is more homogeneous than the Euclidean one (the dissimilarity function has similar values) and the DSSIM outlier scores are much smaller than those coming from the Euclidean metric. Therefore, in this case, a change in the threshold value C does not have a noticeable influence on the results, especially with DSSIM (see Fig. 14). Some anomalies during the steady state are recognized only by the Euclidean metric with $C = 1.4$ (see Fig. 14), but they do not represent important flow phenomena. Indeed, the smaller peaks at around 0.95 s in the Euclidean LOF scores (see Fig. 13) identify some soot formations and light reflections on the window and, therefore, represent a background noise during the combustion process.

This test demonstrates that LOF is a good algorithm to detect the point in space and time where burning fuel droplets appear during the combustion. In this way, the burning

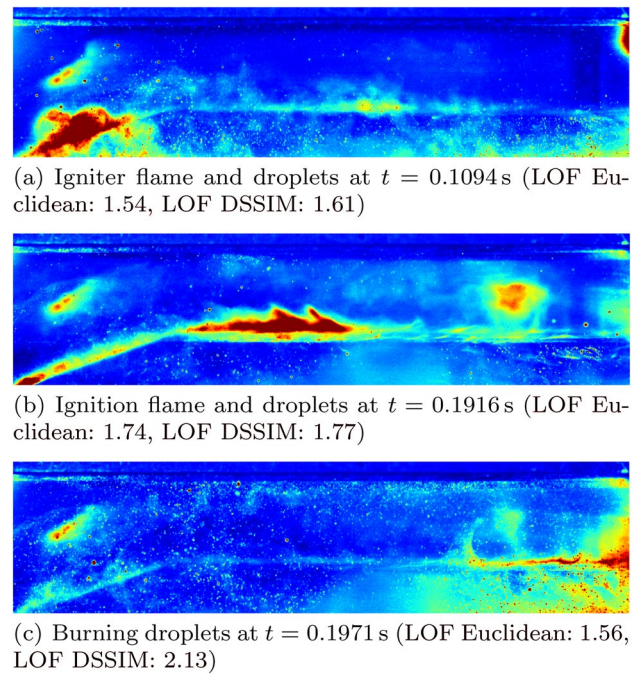


Fig. 11 Representative selection of anomalies found in test 276

conditions that favor the appearance of the entrainment process can be clearly identified.

4.4.3 Test 284

Test 284 has been recorded with a band-pass filter centered around 431 nm placed in front of the high-speed camera. In this way, it is possible to detect the electronically excited species CH^* , which, together with C_2^* and OH^* , are one of the species primarily contributing to the flame luminescence (Devriendt et al. 1996; Schefer 1997). All three species show a close correspondence across the main reaction zone and are thus equally suitable as markers for the flame zone location. In particular, the concentrations of CH^* increase rapidly to a maximum within the flame and then decay rapidly downstream of the reaction zone (Schefer 1997). Therefore, the CH^* images of test 284 give a good representation of the main flame location. The distance matrix (Fig. 15) and the LOF outlier score (Fig. 16) representations are quite different from those of the baseline test (cf. Figs. 7, 8). A higher number of irregular smaller peaks and fluctuations can be detected in test 284, especially in the Euclidean LOF scores. The LOF outlier score values do not change much during the burning process, so that no difference among combustion phases is visible. In particular, the DSSIM LOF scores seem to be quite homogeneous in time. This is due to the presence of the CH^* filter that levels the combustion luminosity out. This is also demonstrated by the mean brightness values in the two-dimensional representation of the dataset: The range




Fig. 12 Euclidean and DSSIM distance matrix of test 276. Potential anomalies according to the threshold value 1.5 as shown in Fig. 13 are highlighted in red on the matrix diagonal

of luminosity values of test 284 is more limited compared to test 203 and 276 (cf. Fig. 17 with Figs. 9, 14). The threshold values for this test are $C = 1.5$ for the Euclidean metric and $C = 1.2$ for DSSIM.

The main anomalies of test 284 are found with both metrics during the steady state in a particular point in time: At around 2.2 s, a lot of droplets detach from the ramp of the fuel slab (see Fig. 18b) and they go burning over the flame. This event corresponds to the arc on the right-hand side of the two-dimensional representation of the dataset at the approximate position $x \in [0.5, 0.6]$ and $\mu \in [19, 25]$, which is clearly identified as anomaly in Fig. 17. Further anomalies are detected also during the ignition phase, mainly by the

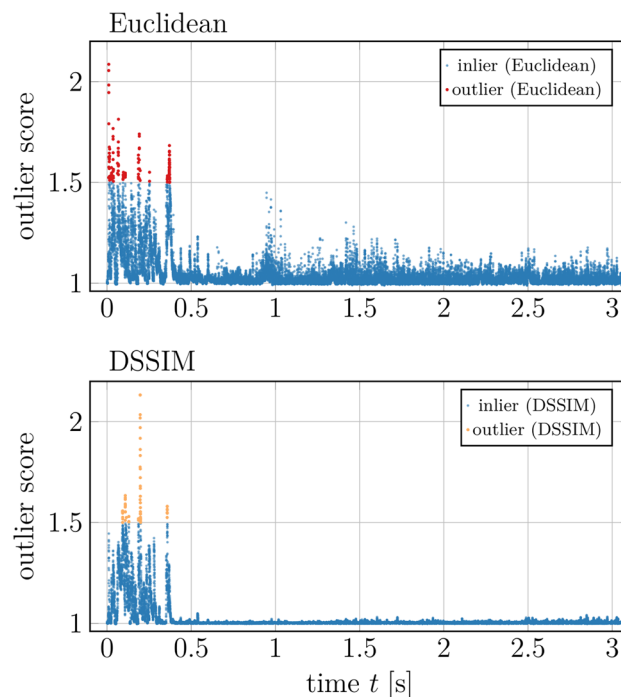


Fig. 13 LOF outlier score of test 276

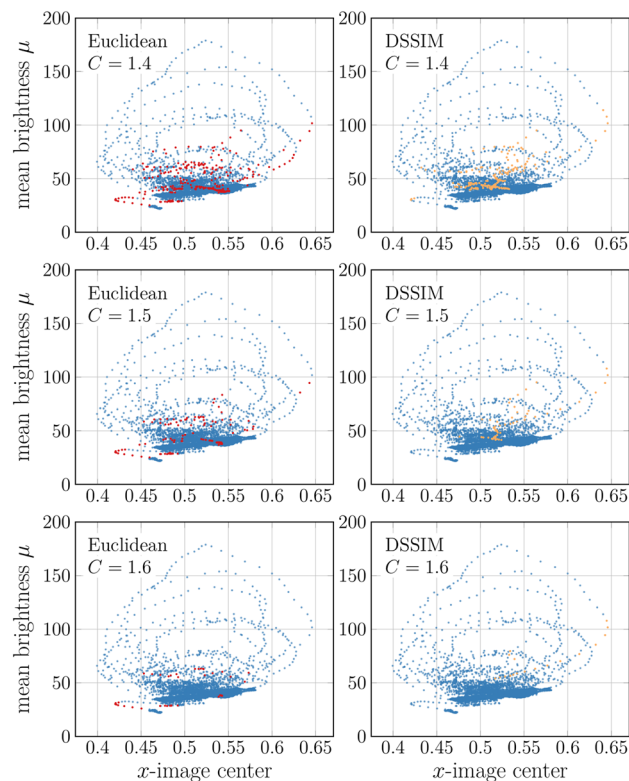


Fig. 14 Visualization of the anomalies found in test 276 in a low-dimensional (2D) representation of the image dataset. For both distance measures, Euclidean distance (left) and DSSIM (right), different threshold values C are compared

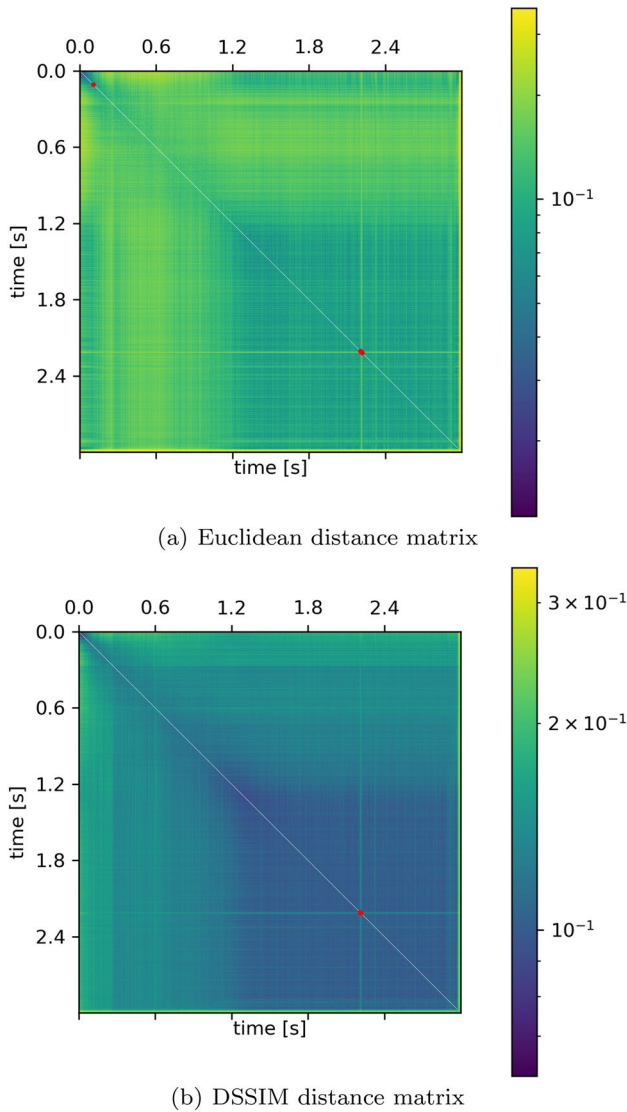


Fig. 15 Euclidean and DSSIM distance matrix of test 284. Potential anomalies according to the threshold values 1.5 (Euclidean) and 1.2 (DSSIM) as shown in Fig. 16 are highlighted in red on the matrix diagonal

Euclidean metric. Some other lower LOF scores peaks (it is important to remember that the threshold value C separating the inliers from the outliers is a user-dependent parameter) are identified by the Euclidean metric during the steady-state and extinction phases. They primarily detect disturbances to the main flow and background noise, such as side flamelets and soot formations. Some smaller flame fluctuations and a few droplets leaving the fuel surface are also identified, but they are too weak to be detected by DSSIM. As expected, by decreasing the C threshold value, more anomalies are detected by both metrics, especially during the steady state (see Fig. 17), but they mainly represent background noise. It could seem strange that the arc on the right-hand side of the




Fig. 16 LOF outlier score of test 284

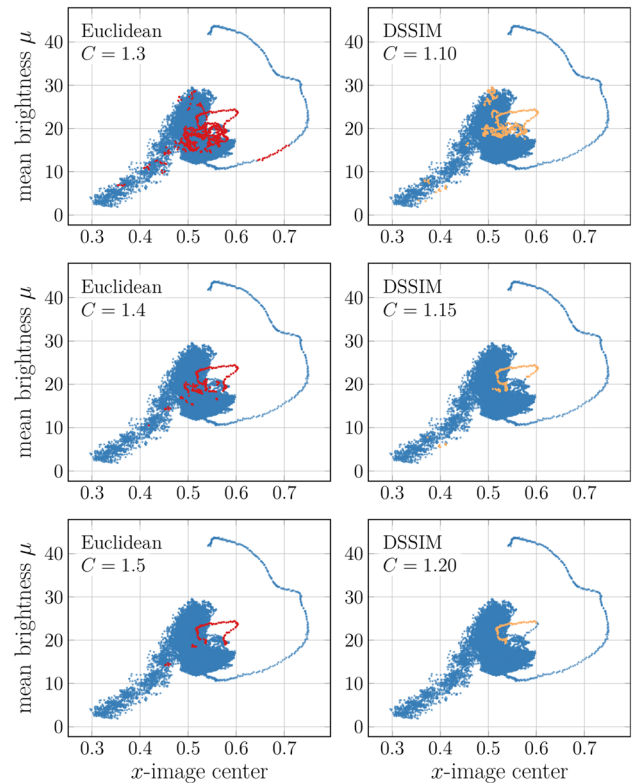


Fig. 17 Visualization of the anomalies found in test 284 in a low-dimensional (2D) representation of the image dataset. For both distance measures, Euclidean distance (left) and DSSIM (right), different threshold values C are compared

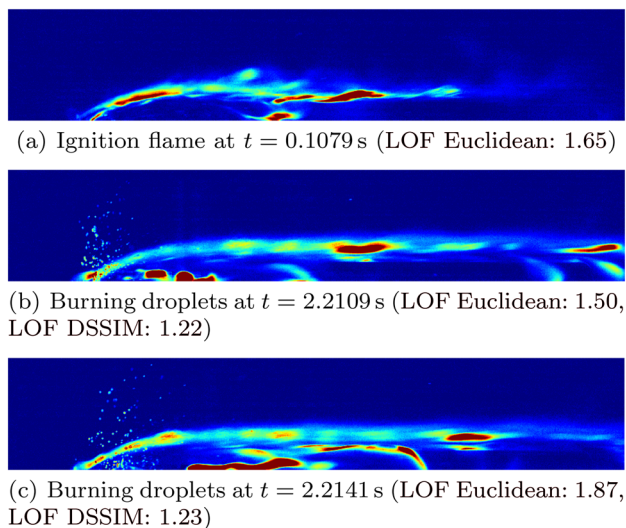


Fig. 18 Representative selection of anomalies found in test 284

low-dimensional representation of the dataset at the approximate position $x \in [0.5, 0.75]$ and $\mu \in [10, 45]$, which represents the extinction phase, is not identified as outlier. This is due to the number of considered neighbors. As explained in Sect. 4.3, the number of neighbors is a parameter chosen by the user that can strongly influence the results. As expected, the lower the value of this parameter, the more the “local” are the outliers. In particular, this parameter can change depending on how fast or slow the dataset changes with time. It is most likely that, in this case, the number of considered neighbors is too low to detect more outliers during the extinction phase. It is expected that, by increasing the value of this parameter, more outliers can be identified. Nevertheless, as shown in Fig. 17, the Euclidean metric with $C = 1.3$ is able to identify some anomalies during the last combustion phase, which represent the extinction flame.

5 Conclusion

Local outlier factor was chosen as one viable algorithm to detect local anomalies in the dataset and applied to hybrid combustion data. To the best of our knowledge, this is the first application of local outlier detection in the specific application of hybrid rocket combustion. The analysis demonstrated that this algorithm is able to reveal several interesting phenomena appearing during the combustion and clearly indicates the potential of unsupervised learning techniques for the study of large datasets. In this work, it was proved that the detection of local anomalies is able to recognize specific phenomena characterizing the combustion of liquefying hybrid rocket fuels. In particular, fuel droplets get entrained into the oxidizer flow and

burning over the flame are clearly identified as outliers with respect to the main combustion process. Furthermore, the first appearance of the flame over the fuel slab surface during ignition, the extinguishing flame after the nitrogen purge, as well as flame blowing and fluctuations during the steady combustion are recognized as weaker disturbances to the regular flow. Depending on what the user is interested in investigating, some parameters can be tuned in the algorithm. The number of neighbors k that is considered by the algorithm influences how “local” the outcome is: By increasing k , the detected outliers range from short-time disturbances to whole combustion phases (thus becoming almost a clustering algorithm). Moreover, in the post-processing phase, the user can choose the threshold value separating the inliers from the outliers, thus deciding which phenomena have to be considered as important anomalies and which ones are only background noise (like soot formations and light reflections on the window, in this case). In this study, it was also shown that, depending on the dissimilarity measure used in the algorithm, the detected anomalies may vary. In particular, the Euclidean metric gives good results in finding local luminosity changes, even of few pixels (see Fig. 5), and therefore, it is particularly suitable to detect burning fuel droplets. On the other hand, the structural dissimilarity measure (DSSIM) is more sensitive to a general frame luminosity variation, and therefore, it is able to detect only strong variations from the main flow, such as a huge amount of droplets or a large flame movement. These characteristics makes the DSSIM distance metric more robust against undesired camera movements, such as shaking, during the test. Strong anomalies, like flows of burning droplets or a fast change in the combustion flame, are detected by both distance measures.

The intention of this work is to demonstrate that unsupervised learning techniques can be successfully applied to large combustion datasets in order to identify specific flow phenomena. In the future, other outlier detection algorithms can be applied to the complete imaging dataset and compared to the LOF algorithm. Moreover, the authors will further develop unsupervised, as well as supervised, learning techniques to apply to combustion data. The objective is to get detailed information on the important phenomena characterizing the hybrid combustion process, such as the entrainment. The next step will be to train the algorithm how to automatically recognize particular structures, such as Kelvin–Helmholtz waves, vortices and droplets, in space and time.

Acknowledgements This research was carried out in part under the *Helmholtz AI* voucher “Investigation and implementation of advanced distance measures for automatic clustering of hybrid rocket combustion video data” and the Projects ATEK and Big-Data-Platform by the German Aerospace Center (DLR).

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Breunig MM, Kriegel HP, Ng RT, Sander J (2000) LOF: identifying density-based local outliers. In: Proceedings of the 2000 ACM SIGMOD international conference on Management of data, pp 93–104
- Ciezki HK, Sender J, Clauß W, Feinauer A, Thumann A (2003) Combustion of solid-fuel slabs containing boron particles in step combustor. *J Propuls Power* 19(6):1180–1191. <https://doi.org/10.2514/2.6938>
- Dask Development Team (2016) Dask: library for dynamic task scheduling. <https://dask.org>
- Devriendt K, Hook HV, Ceusters B, Petters J (1996) Kinetics of formation of chemiluminescent CH by the elementary reactions of C₂H with O and O₂: a pulse laser photolysis study. *Chem Phys Lett* 261:450–456
- Hawkins DM (1980) Identification of outliers, vol 11. Springer. <https://doi.org/10.1007/978-94-015-3994-4>
- Hinneburg A, Aggarwal CC, Keim DA (2000) What is the nearest neighbor in high dimensional spaces? In: 26th International conference on very large databases, pp 506–515
- Karabeyoglu A, Cantwell B, Altman D (2001) Development and testing of paraffin-based hybrid rocket fuels. In: 37th AIAA/ASME/SAE/ASEE joint propulsion conference and exhibit, American Institute of Aeronautics and Astronautics, Salt Lake City, Utah, <https://doi.org/10.2514/6.2001-4503>
- Karabeyoglu A, Altman D, Cantwell BJ (2002) Combustion of liquefying hybrid propellants: part 1, general theory. *J Propuls Power* 18(3):610–620. <https://doi.org/10.2514/2.5975>
- Karabeyoglu A, Stevens J, Geyzel D, Cantwell B, Micheletti D (2011) High performance hybrid upper stage motor. In: 47th AIAA/ASME/SAE/ASEE joint propulsion conference and exhibit, American Institute of Aeronautics and Astronautics, <https://doi.org/10.2514/6.2011-6025>
- Kobald M, Petrarolo A, Schleichtrien S (2015) Combustion visualization and characterization of liquefying hybrid rocket fuels. In: 51st AIAA/SAE/ASEE joint propulsion conference, American Institute of Aeronautics and Astronautics, <https://doi.org/10.2514/6.2015-4137>
- Kriegel HP, Kröger P, Schubert E, Zimek A (2009) Loop: local outlier probabilities. In: Proceedings of the 18th ACM conference on information and knowledge management, pp 1649–1652
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E (2011) Scikit-learn: machine learning in Python. *J Mach Learn Res* 12:2825–2830
- Petrarolo A, Kobald M (2016) Evaluation techniques for optical analysis of hybrid rocket propulsion. *J Fluid Sci Technol* 11(4):JFST0028. <https://doi.org/10.1299/jfst.2016jfst0028>
- Petrarolo A, Kobald M, Schleichtrien S (2018) Understanding Kelvin–Helmholtz instability in paraffin-based hybrid rocket fuels. *Exp Fluids* 59:62. <https://doi.org/10.1007/s00348-018-2516-1>
- Rüttgers A, Petrarolo A, Kobald M (2020) Clustering of paraffin-based hybrid rocket fuels combustion data. *Exp Fluids* 61(1):1–17
- Schefer RW (1997) Flame sheet imaging using CH chemiluminescence. *Combust Sci Technol* 126(1–6):255–279. <https://doi.org/10.1080/00102209708935676>
- Schwabacher M, Oza N, Matthews B (2009) Unsupervised anomaly detection for liquid-fueled rocket propulsion health monitoring. *J Aerosp Comput Inf Commun* 6(7):464–482. <https://doi.org/10.2514/1.42783>
- Thumann A, Ciezki HK (2002) Combustion of energetic materials, chap. Comparison of PIV and Colour-Schlieren measurements of the combustion process of boron particle containing solid fuel slabs in a rearward facing step combustor, vol 5. Begell House Inc., <https://doi.org/10.1615/IntJEnergeticMaterialsChemProp.v5.i1-6.770>
- van der Walt S, Schönberger JL, Nunez-Iglesias J, Boulogne F, Warner JD, Yager N, Gouillart E, Yu T, the scikit-image contributors (2014) scikit-image: image processing in Python. *PeerJ* 2:e453. <https://doi.org/10.7717/peerj.453>
- Wang Z, Bovik AC, Sheikh HR, Simoncelli EP (2004) Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process* 13(4):600–612
- Zimek A, Schubert E, Kriegel HP (2012) A survey on unsupervised outlier detection in high-dimensional numerical data. *Stat Anal Data Min ASA Data Sci J* 5(5):363–387

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.