University of Bristol - Explore Bristol Research
General rights

CrossMark

RESEARCH ARTICLE

Automated mask generation for PIV image analysis based on pixel intensity statistics

Alessandro Masullo[1] · **Raf Theunissen[1]**

Abstract The measurement of displacements near the vicinity of surfaces involves advanced PIV algorithms requiring accurate knowledge of object boundaries. These data typically come in the form of a logical mask, generated manually or through automatic algorithms. The automatic detection of masks usually necessitates special features or reference points such as bright lines, high contrast objects, and sufficiently observable coherence between pixels. These are, however, not always present in experimental images necessitating a more robust and general approach. In this work, the authors propose a novel method for the automatic detection of static image regions which do not contain relevant information for the estimation of particle image displacements and can consequently be excluded or masked out. The method does not require any a priori knowledge of the static objects (i.e., contrast, brightness, or strong features) as it exploits statistical information from multiple PIV images. Based on the observation that the temporal variation in light intensity follows a completely different distribution for flow regions and object regions, the method utilizes a normality test and an automatic thresholding method on the retrieved probability to identify regions to be masked. The method is assessed through a Monte Carlo simulation with synthetic images and its performance under realistic imaging conditions is proven based on three experimental test cases.

✉ Alessandro Masullo
  a.masullo@bristol.ac.uk

[1] Department of Aerospace Engineering, University of Bristol, Bristol, UK

1 Introduction

Particle image velocimetry (PIV) is a non-intrusive technique for measuring flow velocities by analysing successive images of seeding particles captured by specialised cameras. Images are subsequently interrogated by means of statistical operators, such as cross correlation, to infer flow velocities representative of the particle image displacements captured within the image subsections (Theunissen 2012). The most suitable image for PIV analyses constitutes homogeneously distributed bright particle images on a completely dark background. Unfortunately, this ideal scenario is not always possible as the flow region of interest can be occupied by objects producing artefacts (light reflections, shadow areas, local variations in seeding density, etc.) hindering the measurement quality (Westerweel et al. 2005). If the required particle image displacement measurement is in the near vicinity of a surface, special care must be taken in adequate surface preparation during the experimental phase (Paterna et al. 2013), images must be properly pre-processed by means of, e.g., background subtraction (Mendez et al. 2017), or advanced image interrogation processes are to be adopted (Ronneberger et al. 1998; Gui et al. 2003; Usera et al. 2004). Such advanced PIV analysis routines typically involve the exclusion of object regions within cross-correlation windows as to minimise distortions in the cross-correlation map, consequently limiting associated displacement bias errors (Theunissen et al. 2008). This, in turn, requires knowledge of the imaged object boundaries to enable seeded flow regions to be discerned from non-seeded object regions. The generation of such logical masks thus constitutes a critical step, independent from the PIV algorithm used.

Logical masks are typically generated manually when the geometry of the object is very simple and clearly detectable in the PIV images. However, the manual approach becomes unreliable and too time-consuming when complex aerodynamic shapes are involved or several image sets need to be analysed with many different masks. Some examples of automated mask generation have been reported in literature: in Dussol et al. (2016), strong light reflections of the object were manually detected to generate the initial mask, which was dynamically moved using the a priori knowledge of the object motion. In Sanchis and Jensen (2011), the air–water interface was successfully detected exploiting the image intensity gradients through the Radon transform. Ergin et al. (2015) dealt with the motion of a moving body by implementing a tracking method and moving to an object-fixed coordinate system, allowing the use of a manual mask to the moving object. Deen et al. (2010) applied a combination of contrast enhancement and intensity thresholding to identify bubbles in two-phase flows. However, shadows caused by the obstruction of the laser light also constitute unwanted image regions, yet these are not defined by high intensity boundaries. Moreover, experimental settings often produce complex images where objects contours are dark, vague, and blurred. Objects thus become hardly detectable even for a human observer, limiting the application of typical masking algorithms. Alternatively, the spatial coherence between image pixels is involved to filter and identify object regions by means of low-pass filtering (Willert 1999), median filtering (Scholz and Kähler 2004), etc. Due to the finite extent of the filter kernels, such operations tend to introduce inaccuracies in the localisation of the object boundaries and are influenced by noise in the grey-scale distributions. Delnoij et al. (1999) discriminated objects on the basis of inherent differences in their motion. This, however, demands objects to be identified a priori. Similarly, Honkanen and Nobach (2005) utilised the variation in flow velocities to completely remove stagnant objects. In case of temporally well-resolved PIV recordings, frequency analyses of the time-evolution in pixel intensities allow the identification of objects (Sciacchitano and Scarano 2014). This methodology is, however, only conducive for time-resolved PIV recordings involving dynamic objects.

In the current work, the authors propose a general, robust methodology for the automatic detection of static image regions containing no information relevant in the cross-correlation operation. Without any special requirement for the image regions to be excluded (shape, luminosity, or contrast), the method takes advantage of the statistical information contained within each individual pixel across several independent PIV images. A statistical test distinguishes unseeded image regions from image sections of interest with pixel accuracy. The underlying principle is the understanding that the temporal intensity variation of pixels constituting the unseeded partitions is mainly due to camera noise, producing a normal-distributed intensity histogram (Westerweel 2000). Pixels of seeded image regions on the other hand are characterized by the strong alternation of bright and dark values, producing highly skewed histograms (Westerweel 2000). A normality test is, therefore, performed on the time history of each pixel of the image to distinguish objects from particle images. Unlike filtering operations, the categorisation is performed pixel-wise and, therefore, does not incorporate neighbouring pixels. As a result of an automatic thresholding on the p value of the normality test, the proposed method is fully autonomous and independent of user input. The technique is fully assessed on synthetic images and the application to several experimental cases is proposed to show its suitability for a variety of real conditions.

2 Methodology

2.1 Normality test

Considering a sequence of N PIV images, for each pixel, the evolution in intensity across the image sequence can be extracted. Pixel intensities pertaining the imaged flow region will be characterized by strong alternations of high and low values due to the passage of tracer particles. Corresponding intensity probability density functions consequently exhibit a strong skewness (Fig. 1). Conversely, pixels associated with the imaged object display intensity variations principally due to camera noise, which is typically normally distributed (Westerweel 2000). Based on this observation, the authors propose a method to automatically categorise pixels within PIV images, by means of a pixel-wise normality tests. Given that the proposed methodology is based on probability density functions, validity is independent of the temporal resolution in the acquired images.

Normality tests are a tool to compute the likelihood of a set of data to be modelled by a Gaussian probability density function (PDF). Existing literature offers several tests for normality such as D'Agostino's K-squared test (D'agostino et al. 1990), the Kolmogorov–Smirnov test (Fasano and Franceschini 1987), the Pearson's Chi-squared test (Plackett 1983), etc. Lead by simplicity of implementation, the method of choice in the current work is the Jarque–Bera test, for which the heuristic can be easily evaluated on the basis of data skewness and kurtosis (Jarque and Bera 1987).

For a set of observations x (i.e., the set of intensity values of same pixel) with mean μ and standard deviation σ, the skewness s and kurtosis k are defined as:

Fig. 1 Illustration of the principle used for the mask detection: the probability distributions of pixel intensities have very distinguishable shapes depending on whether pixels belong to flow regions or objects
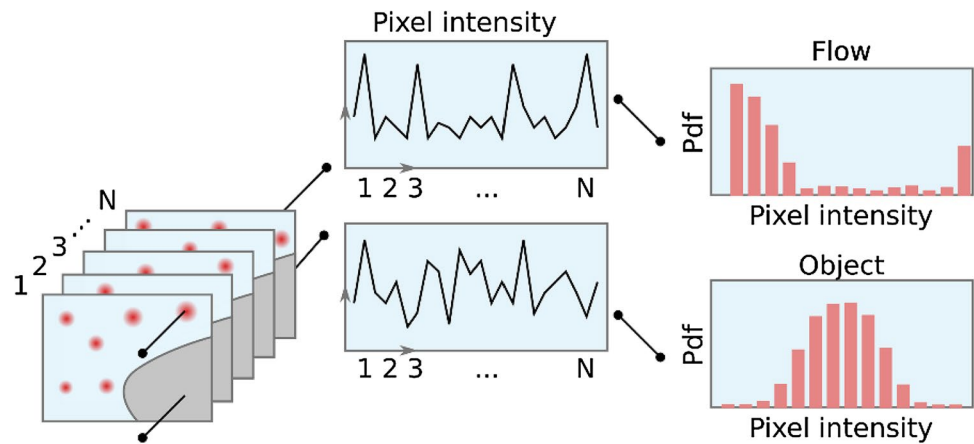
where $E(t)$ represents the expected value of t. For a normally distributed data set, the expected value of the kurtosis will equal 3 with a skewness of 0. The Jarque–Bera statistic quantifies the data set discrepancies from these expected ideal values:

$$JB = \frac{N}{6}\left(s^2 + \frac{(k-3)^2}{4}\right) \quad (2)$$

where N is the size of the sample x (i.e., number of PIV images). The statistic expressed by (2), under the null hypothesis of x being normally distributed, is asymptotically distributed as a Chi-squared with two degrees of freedom (Jarque and Bera 1987). Thanks to this assumption, the p value for the Jarque–Bera test can be evaluated as:

$$p = 1 - \chi^2_{(2)}(JB) \quad (3)$$

where the general Chi-squared with ν degrees of freedom is defined as:

$$\chi^2_{(\nu)}(x) = \int_0^x \frac{t^{(\nu-2)/2} e^{-t/2}}{2^{\nu/2} \Gamma(\nu/2)} dt \quad (4)$$

with Γ being the gamma function:

$$\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt. \quad (5)$$

Given a p value for each pixel of the image, the null hypothesis of normality can be accepted (noise-dominated region) or rejected (flow region) based on the comparison with a cut-off level. This level is typically set to 0.05 (Jarque and Bera 1987).

$$s = \frac{E(x-\mu)^3}{\sigma^3}$$
$$k = \frac{E(x-\mu)^4}{\sigma^4} \quad (1)$$
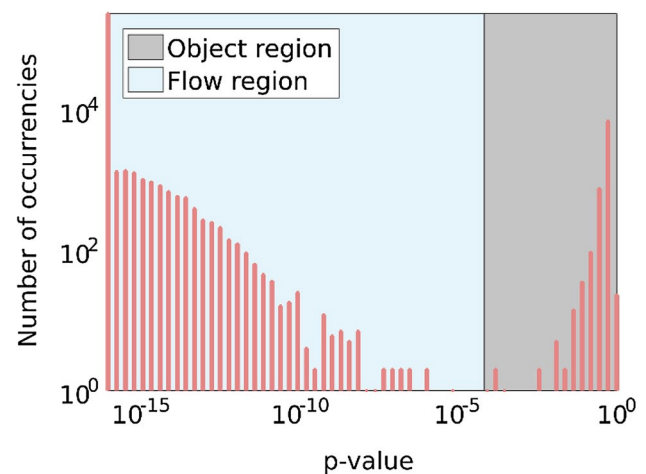
Fig. 2 Histogram of the p values for a set of experimental images (micro-channel flow, see. Sect. 5). The distribution of p values presents the typical bimodal distribution due to the presence of flow (low p values) and object (high p values). The automatic threshold discerns flow regions from the object by maximising the separability of the histogram

The aim of the presented work is not to test the normality of boundary pixels in an absolute sense, but rather to discern them from regions of strong non-normality due to the presence of particle images. Empirical studies lead the authors to conclude that flow regions are typically characterized by extremely low p values, whereas regions void of any signal usable for cross correlation usually present p values which can be several orders of magnitude higher (Fig. 2). Once the p value is evaluated for each image pixel, an automatic threshold exploiting the peculiar bimodality of the probability in p values across the image can be adopted to discern low p values (image regions of interest) from high p values (statistically irrelevant regions) to generate a binary mask. In the current work, the Matlab function *otsuthresh* was used as implementation of Otsu's

Fig. 3 Example of signal pre-processing, applicable in case of decaying laser intensity. **a** Synthetic Gaussian noise combined with decreasing average light intensity and **b** histogram; **c** high pass filtered signal with **d** restored Gaussian histogram
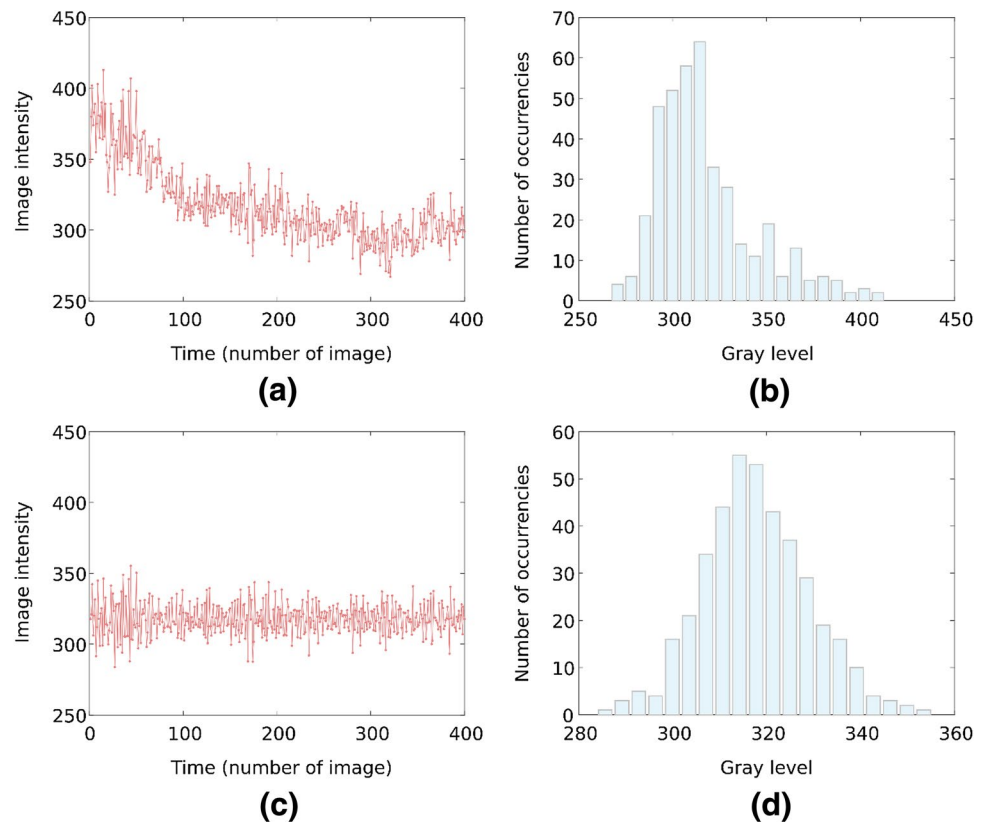
method (Otsu 1979), together with a median filter[1] of a fixed 3 by 3 kernel to enhance the bimodality of the histogram (Gonzalez et al. 2013). The reader should note that this automatic implementation of the threshold ensures a robust distinction of flow from object areas even when related pixels are not perfectly normally distributed, as flow regions will continue to have a much lower distribution of p values due to their more skewed pdf as corroborated by the theoretical intensity pdf suggested by Westerweel (2000).

2.2 Pre-processing

The assumption of normally distributed noise can be violated when the laser light intensity varies across the image sequence, producing artificially skewed histograms of pixel intensity. An example of this behaviour is shown in Fig. 3a and b, where a normally distributed pixel intensity was combined with a decreasing average light intensity. The histogram produced by this light distribution could easily be mistaken for a particle image, as it is not normally distributed. However, even exceptional cases like this do not prevent the application of the proposed mask detection algorithm, since a simple additional pre-processing

step can be implemented to equalize the light intensity. For example, a high pass filter can be applied to the time history of the intensities of the individual pixels to reduce the effect of the light variation and restore the Gaussian shape of the histogram. The high pass filtered signal is presented in Fig. 3c, together with its normally distributed histogram in Fig. 3d.

Another common scenario where the hypothesis of normally distributed noise could be violated is in case of a double pulsed laser. Due to the possible discrepancies in manufacturing of the laser cavities, the light intensity of the two consecutive PIV images might be different, producing two sets of images that have a skewed or bimodal intensity distribution. In such a case, the authors suggest to apply the mask detection algorithm independently on the two image sets, producing two different masks of which the logical union can be used as final mask.

2.3 Minimum number of images required

Besides a sufficient number of images to ensure validity of the normality statistics, a secondary condition for the proposed methodology to work is that each pixel of the flow region should be occupied by a particle image at least once in the entire PIV sequence. The number of images based on the first condition depends on the image quality and consequently cannot be estimated a priori. The second condition,

---

[1] The median filter is only applied for the threshold estimation, therefore the pixel accuracy of the mask holds true.
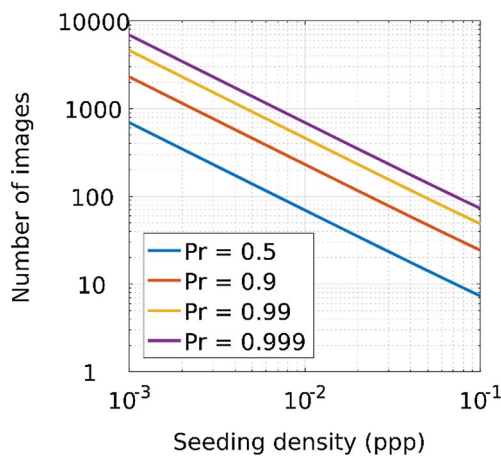
Fig. 4 Number of independent PIV images necessary to allow the correct functioning of the proposed masking algorithm against seeding density. The plots are shown for different probabilities as expressed by Eq. (7)

on the other hand, can be exploited to estimate an absolute minimum number of images necessary to allow the correct functioning of the mask detection. The probability of having n particles in a control volume V follows a Poisson distribution (Adrian 1983):

$$\Pr\{n \text{ particles in } V\} = \frac{s^n}{n!} e^{-s} \tag{6}$$

where s is the seeding particles concentration. The probability of a pixel of being occupied by a particle image at least once in a sequence of N images is described by a binomial distribution. From Eq. (6), for $n = 1$, follows:

$$\Pr\{1 \text{ particle per pixel}\} = 1 - \left(1 - s \cdot e^{-s}\right)^N. \tag{7}$$

Equation (7) can be used to estimate a necessary number of images N, for a given probability Pr. Equation (7) is graphically presented in Fig. 4 and is indicative of the applicability of the method.

3 Post-processing

The application of the normality test explained in the previous section leads to the generation of a mask based solely on the time history of each pixel and as such independent of the spatial coherence between image pixels. The spatial information can be used in a post-detection phase to enhance the quality of the mask, especially when insufficient images are available. In such cases, spurious random pixels might be excluded producing pixel-sized holes in the binary mask generated. To increase the robustness of the mask, morphological operations can be included to simplify

the binary image preserving the characteristics of the shape and eliminating irrelevancies (Haralick et al. 1987). While the reader may require alternative operations when dealing with special geometries, the morphological operations of closing and filling proposed in this work are fundamental. All morphological operations can be performed using the MATLAB routines *bwmorph* and *imfill*. The closing process allows to smoothen contours and to eliminate small holes, whereas the process of filling allows to close every area of the binary mask that cannot be reached by filling in the background from the edge of the image. The majority process allows a final cleansing of the mask by recursively setting a pixel to 1 when more than half of the pixels in the immediate neighbourhood are 1. Figure 5 illustrates the achievable improvements by means of the morphological operations where an example mask is artificially perturbed with noise and holes (Fig. 5a). After the closing process, pixel-sized holes have disappeared (Fig. 5b). The bigger hole inside the mask is removed as a result of the filling process (Fig. 5c) and the recursive application of majority (Fig. 5d) filters out smaller pixel-sized structures without altering the main shape of the mask.

The reader should note that the adoption of morphological operations, despite being beneficial from a general point of view, may cause unwanted effects on the smallest scale of the geometry (8 by 8 pixels). Operations like majority would reduce the size of every sharp corner of the mask by one pixel.[2] Although the advantages of using morphological operations may well outweigh smaller issues on this scale, its use should, therefore, be considered as a facultative addition to the main methodology.

4 Numerical assessment

To assess the validity of the proposed masking methodology and investigate the minimum number of images required for the approach to be reliable, Monte Carlo simulations involving synthetic PIV images were performed. Cylinders of three different diameters of 100, 200, and 400 pixels, occupying 3, 13, and 50% of the entire image area, respectively, served as object. Pixel intensities within the cylinders were randomly drawn from a Gaussian distribution with a mean normalized intensity of 0.12 and standard deviation of 0.03. Images were 500 by 500 pixels in size, discretized in 16 bits, with randomly

---

[2] The majority operations works in cluster of 8 pixels: the sharp edge of a rectangle composed by ones, surrounded by zeros, is a pixel surrounded by 4 ones (the rectangle) and 5 zeros. Therefore, it will be replaced by the majority which is zero.

Fig. 5 Example of application of the proposed sequence of morphological operations to (**a**) an artificially perturbed mask. **b** Application of closure removes single-pixel-sized holes from the mask; **c** filling operation closes any major size hole in the mask; and **d** the majority operation allows a final clean from any smaller scale perturbation

(a)

(b)

(c)

(d)
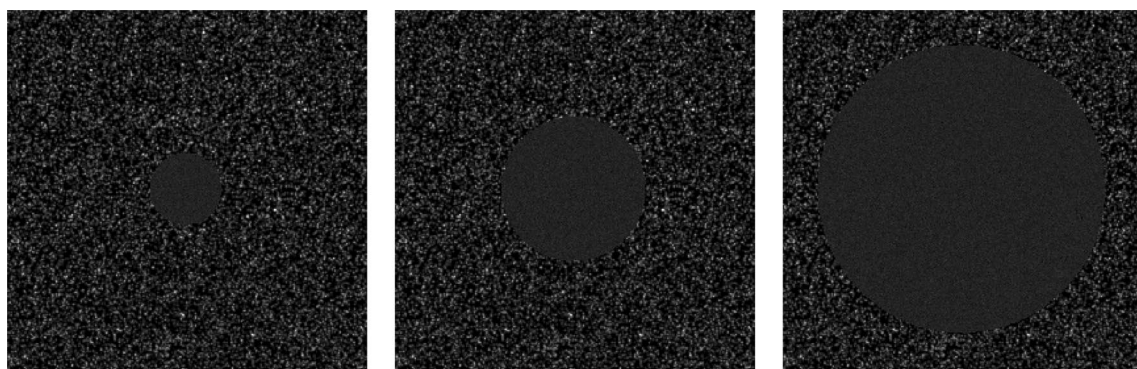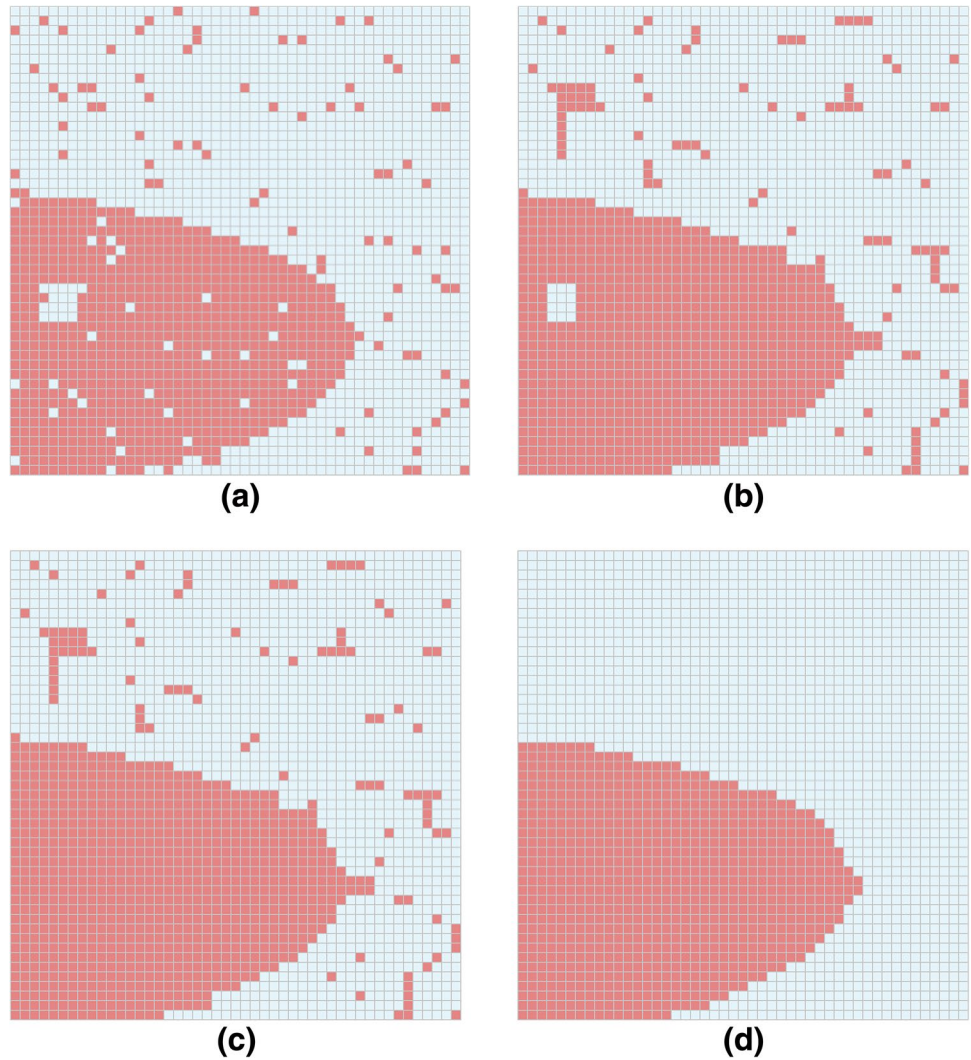
Fig. 6 Three samples of the cylinder images used for the synthetic analysis. Cylinders of 100 px (*left*), 200 px (*centre*), and 400 px (*right*) in diameter were generated using Gaussian noise

Fig. 7 Results of the Monte Carlo simulation for the mask detection of the three synthetic cylinders. The plots present the error of the mask detected versus the number of PIV images used for that detection error
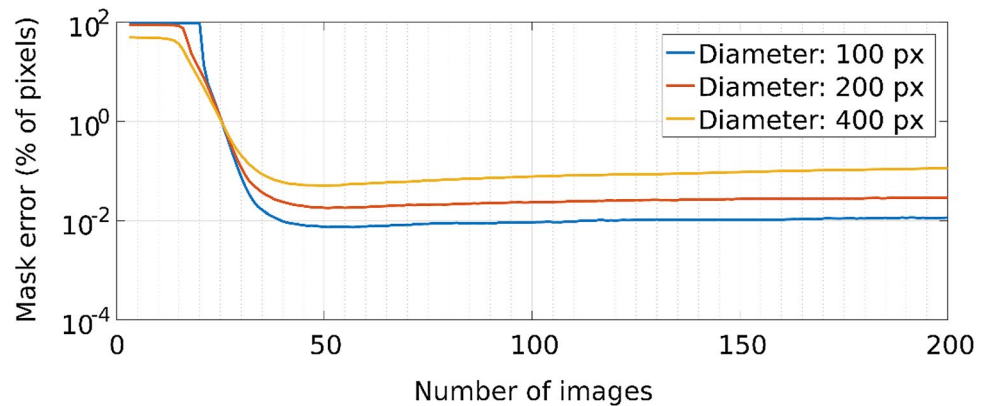
distributed particles in the flow regions having a diameter of 3 pixels and a seeding density of 0.1 ppp (Fig. 6).

For each cylinder, 200 independent images were generated a total of 100 times to obtain statistically meaningful average results. For each size of the set, the error ε in the estimation of the mask area was evaluated as:

$$\varepsilon = \frac{100}{WH} \sum_{i=1}^{W} \sum_{j=1}^{H} \left| M_{i,j} - \hat{M}_{i,j} \right| \qquad (8)$$

where W and H are the width and height of the image in pixels, M is the ideal mask, and $\hat{M}$ is the estimated mask. Results for the error ε are presented in Fig. 7 and show a similar behaviour for all cylinders.

When fewer than 20 images are used for the mask detection, the outcome of the algorithm is erroneous due to an insufficient number of independent occurrences to distinguish normally distributed pixels from highly skewed pixels in flow regions. However, increasing the number of images analysed, the mask error quickly drops to 10^{-1} or less for a set of 40 images. After this value, the error reaches a constant level as increasing the number of images does not affect the normality test anymore. This number of images is further confirmed by the theoretical estimation of Eq. (7), where 40 images with a seeding density of 0.1 ppp yield a probability of 98% of pixels being occupied at least by a particle image throughout the PIV sequence.

A detail of the resulting mask detected for a cylinder of 100 pixels in diameter using 200 PIV images is shown in Fig. 8, demonstrating a near perfect agreement with the imposed mask.

5 Experimental case

To validate the performances of the automatic mask detection in experimental conditions, three experimental test cases are proposed: (a) a micro-channel flow; (b) a transparent acrylic aerofoil, and (c) a pipe junction simulating
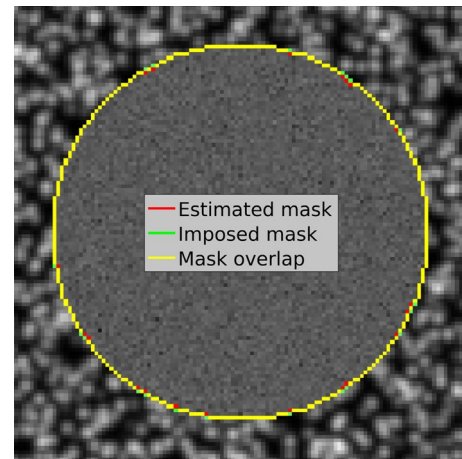

Fig. 8 Result of the mask detection algorithm for the cylinder of 100 pixel of diameter using 200 images, showing a detail of the edges of the detected mask, together with the imposed mask and the respective overlap

a lung. The first test case comprises 600 PIV images of a micro-channel flow driven by 200 bar of pressure, constituting Case A of the 4th International PIV Challenge (Kähler et al. 2016). The resulting mask is presented in Fig. 9a, where a sample of the original PIV images is superimposed with the estimated mask shown in red. In addition to the mask, Fig. 9b, c, respectively, shows a histogram of the grey levels in time for pixels belonging to the masked area and flow region. The plots are in agreement with the observation that masked regions are characterized by normally distributed histograms, whereas seeded image regions present a strongly skewed distribution in grey levels. For the histograms presented in Fig. 9b and c, values of skewness are, respectively, 5.1×10^{-3} and 2.0.

The second test case comes from an experiment performed in the low turbulence wind tunnel of the University of Bristol. PIV experiments were performed using a 200 mJ Nd:Yag laser and 4MP FlowSense EO camera on

Fig. 9 a Automatically detected mask for the micro-channel overlapped with a sample of the set of images (contrast enhanced for clarity), **b** histogram of the *grey* levels for a pixel belonging to the mask and **c** seeded flow portion

(a)

(b)

(c)

(a)

(b)

(c)
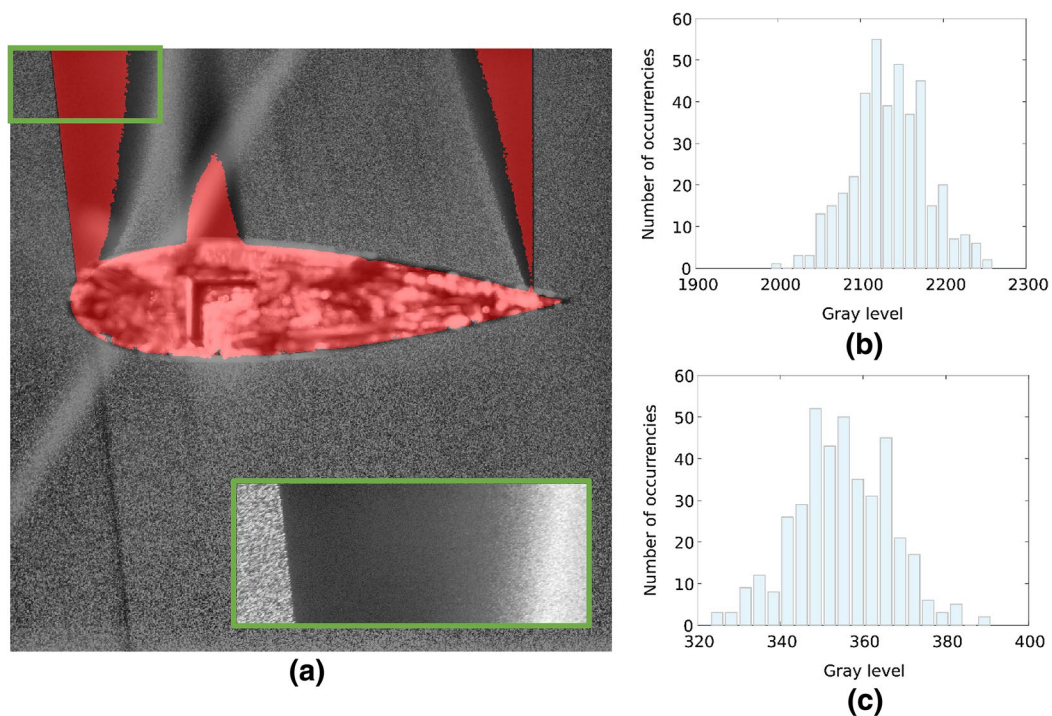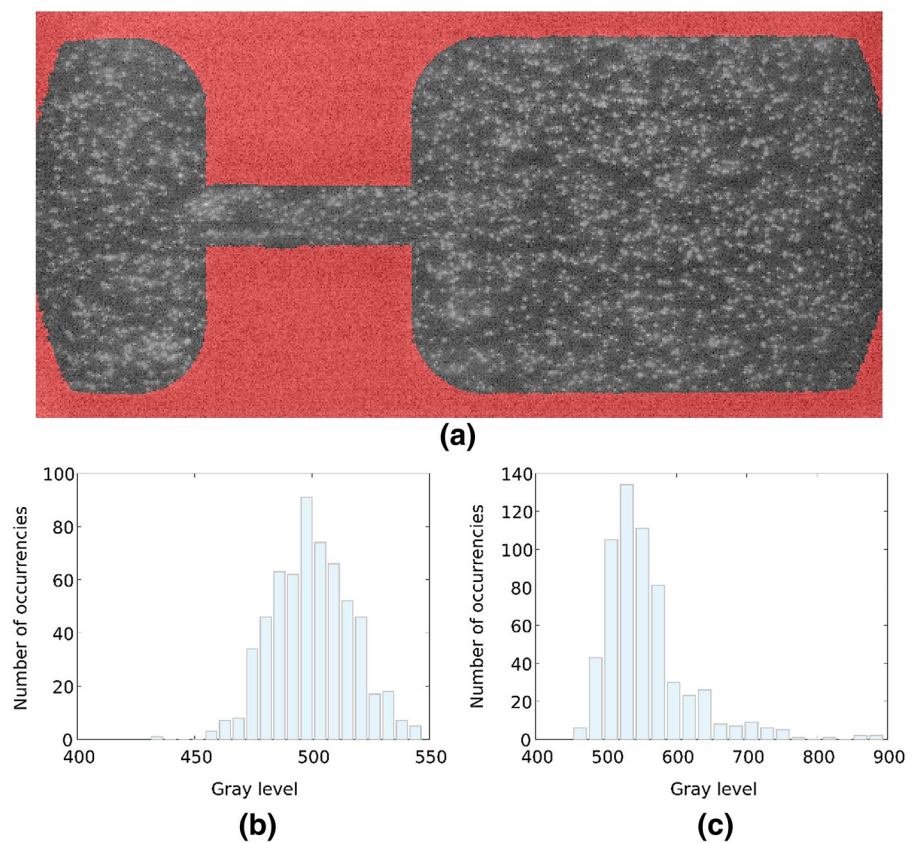
Fig. 10 Automatically detected mask for the transparent aerofoil and detail of the shadow in the top-left corner (**a**); histogram of the *grey* levels for a pixel belonging to the aerofoil (**b**) and the shadow (**c**), both belonging to the mask

an NACA0024 airfoil at a chord-based Reynolds number of 1.4×10^5. The airfoil was manufactured out of acrylic to reduce light reflections in vicinity of the surface. A total of 400 images were acquired in 16 bits image depth. The final mask is presented in Fig. 10a and illustrates how both aerofoil and light shadow area are detected by the masking algorithm. Both regions are characterized by the lack of information due to the absence of particles and present normally distributed grey levels. Figure 10b, c shows two histograms for a pixel belonging to the aerofoil and a pixel belonging to the shadow, respectively; as already mentioned, both the histograms are normally distributed and contribute to the masked region; however, the histogram for the aerofoil region (Fig. 10b) has a mean value which is higher than the shadow area (Fig. 10c), confirming that the suitability of the method is independent from the brightness and the contrast of the masked region.

The reader should notice that the shadow areas produced by the aerofoil are not entirely detected as masks; due to the transparent material used to manufacture the object, the casted shadows are not as dark as those produced by an opaque object, but they gradually fade from dark to bright (see detail of Fig. 10a). This implies that particle images also disappear progressively as they cross the shadow: those areas where the particle image intensity is on par with the random noise are detected as mask, while the regions that may contain valuable information about the displacements are maintained as flow regions, producing the pattern observed in Fig. 10a.

The final test case involves 534 images of a flow bifurcation simulating a lung at $Re = 2.24$ (Theunissen and

Fig. 11 Results of the automated mask detection for the lung test case, overlapped with a sample of the set of images (**a**); histogram of the *grey* levels for a pixel belonging to the object (**b**) and to a noisy area within the channel (**c**), both detected as mask
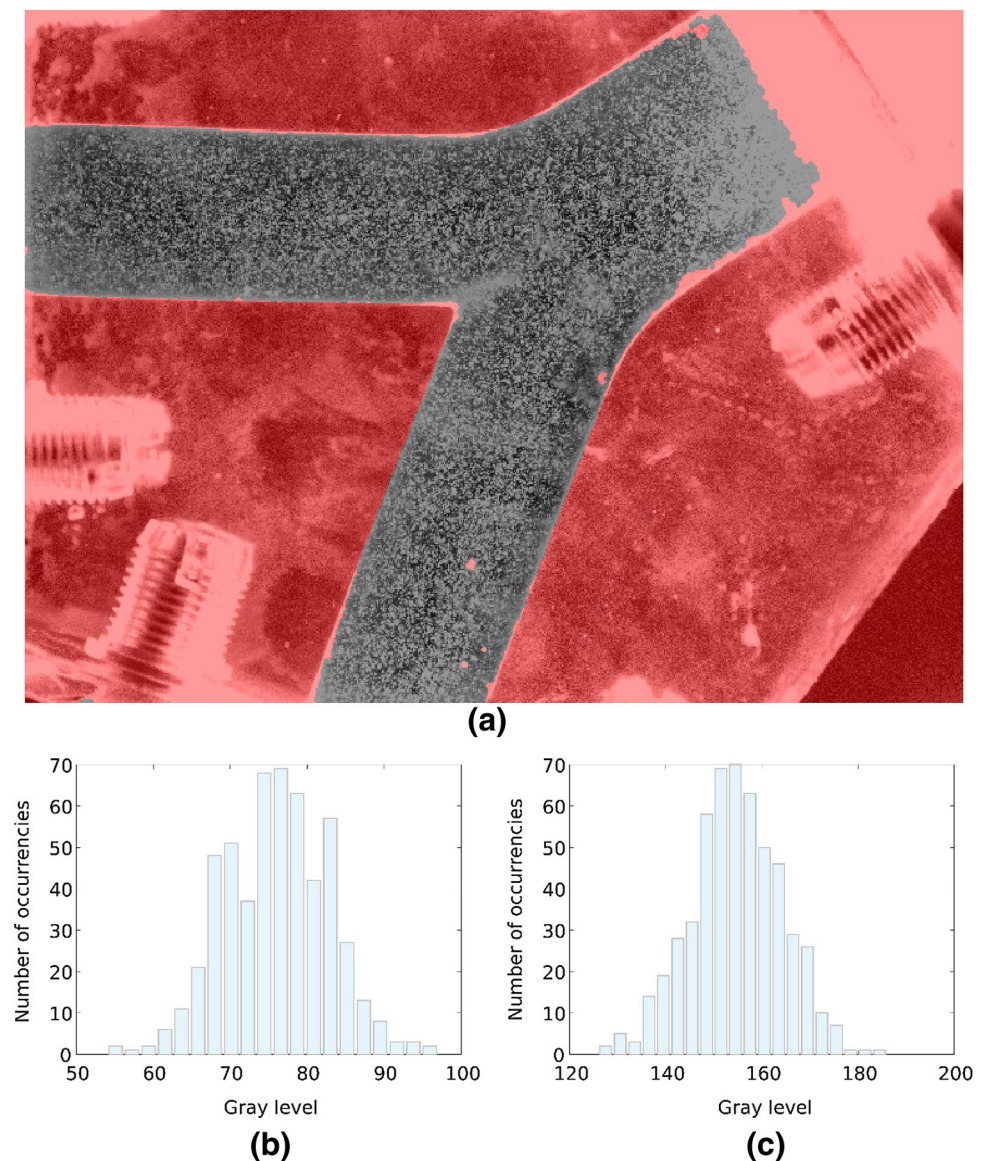
Riethmuller 2008). The 8 bit discretization renders the test case challenging as sporadic noisy areas within the seeded flow region do not produce significant skew in the PDFs. Figure 11b, c, respectively, shows the histograms of pixels belonging to the boundary region and noisy flow area. Both present a normal distribution. This strong image noise in the flow region, coupled with the bright intensity of the particles and the low bit discretization, produces small regions within the flow area to be masked out. Overall, the geometry of the object is correctly retrieved and the occurrence of these additional masked regions is restricted to very small domains. As mentioned in the previous section, these regions can be easily cleaned through the application of additional morphological operations. It is, however, important to note that the mask is based on statistical information of the PIV images and the exclusion of those masked regions from a PIV analysis does not affect the cross-correlation results as those pixels do not contain information from a statistical point of view.

6 Limitations

The automatic mask detection proposed in this paper requires objects in PIV images to be stationary. With the analysis based on the time history of each pixel, vibrations or deformations of the regions to be masked would combine time history of objects with time history of flow regions, hampering the correct classification of the pixels. Conditions like rigid motion and vibrations can be easily tackled by implementing a tracking method and moving the analysis to an object-fixed coordinate system (Ergin et al. 2015). The presence of ghost particles on reflective surfaces constitutes a more complex case. If the experimental settings were such that particle images were reflected within the objects to be masked, the mask detection will not be able to discern real particles from reflected ones. These regions will subsequently be characterized on the basis of pixels containing (reflected) particle images and will consequently be excluded from the masked region.

7 Conclusions

A novel method for the automatic detection of static masks for PIV is proposed. The method does not require any specific feature in terms of brightness or contrast of the images as it exploits pixel-wise statistical information regarding image intensity extracted across several PIV images. Statistically irrelevant image regions are discerned from seeded regions of interest based on the assumption that variations of light intensity in the former are normally distributed in time, while the latter are skewed by the alternation of

background and bright particles. The Jarque–Bera test is used to test the normality of the pixel intensity in time and an automatic threshold based on Otsu's method is adopted to categorise image pixels based on p values. Numerical assessments by means of Monte Carlo simulations on synthetic images were performed to investigate the minimum number of images needed for the normality test to be effective and results showed that at least 40 images are necessary for the method to reliably identify the masked regions. The experimental application on three different sets of PIV images additionally confirmed the applicability of the method for different conditions of background, particle dimensions, seeding densities, and levels of quantization, showing visual agreement of the detected masks with the objects in the images.

References

Adrian RJ (1983) Laser velocimetry. In: Goldstein R (ed) Fluid mechanics measurements. Hemisphere, Washington, DC, pp 155–244

D'agostino RB, Belanger A, D'Agostino Jr RB (1990) A suggestion for using powerful and informative tests of normality. Am Stat 44:316–321

Deen NG, Willems P, Van Sint Annaland M et al (2010) On image pre-processing for PIV of single-and two-phase flows over reflecting objects. Exp Fluids 49:525–530. doi:10.1007/s00348-010-0827-y

Delnoij E, Westerweel J, Deen NG et al (1999) Ensemble correlation PIV applied to bubble plumes rising in a bubble column. Chem Eng Sci 54:5159–5171. doi:10.1016/S0009-2509(99)00233-X

Dussol D, Druault P, Mallat B et al (2016) Automatic dynamic mask extraction for PIV images containing an unsteady interface, bubbles, and a moving structure. Comptes Rendus Mec 344:464–478. doi:10.1016/j.crme.2016.03.005

Ergin FG, Watz BB, Wadhwa N (2015) Pixel-accurate dynamic masking and flow measurements around small breaststroke-swimmers using long-distance MicroPIV. In: 11th International symposium on particle image velocimetry—PIV15. Santa Barbara, California, September 14–16

Fasano G, Franceschini A (1987) A multidimensional version of the Kolmogorov-Smirnov test. Mon Not R astr Soc 225:155–170

Gonzalez RC, Woods RE, Eddins SL (2013) Digital image processing using MATLAB. McGraw Hill Education, New Delhi

Gui L, Wereley ST, Kim YH (2003) Advances and applications of the digital mask technique in particle image velocimetry experiments. Meas Sci Technol 14:1820–1828. doi:10.1088/0957-0233/14/10/312

Haralick RM, Sternberg SR, Zhuang X (1987) Image analysis using mathematical morphology. IEEE Trans Pattern Anal Mach Intell 9:532–550. doi:10.1109/TPAMI.1987.4767941

Honkanen M, Nobach H (2005) Background extraction from double-frame PIV images. Exp Fluids 38:348–362. doi:10.1007/s00348-004-0916-x

Jarque CM, Bera AK (1987) A test for normality of observations and regression residuals. Int Stat Rev/Rev Int Stat 55:163–172. doi:10.2307/1403192

Kähler CJ, Astarita T, Vlachos PP et al (2016) Main results of the 4th international PIV challenge. Exp Fluids 57:1–71. doi:10.1007/s00348-016-2173-1

Mendez MA, Raiola M, Masullo A, Discetti S, Ianiro A, Theunissen R, Buchlin J-M (2017) POD-based background removal for particle image velocimetry. Exp Therm Fluid Sci 80:181–192. doi:10.1016/j.expthermflusci.2016.08.021

Otsu N (1979) A threshold selection method from gray-level histograms. IEEE Trans Syst Man Cybern 9:62–66. doi:10.1109/TSMC.1979.4310076

Paterna E, Moonen P, Dorer V, Carmeliet J (2013) Mitigation of surface reflection in PIV measurements. Meas Sci Technol. doi:10.1088/0957-0233/24/5/057003

Plackett RL (1983) Karl Pearson and the Chi Squared Test. Int Stat Rev/Rev Int Stat 51:59–72. doi:10.2307/1402731

Ronneberger O, Raffel M, Kompenhans J (1998) Advanced evaluation algorithms for standard and dual plane particle image velocimetry. In: Proceedings of the 9th International Symposium on Applied laser techniques to fluid mechanics, pp 13–16

Sanchis A, Jensen A (2011) Dynamic masking of PIV images using the Radon transform in free surface flows. Exp Fluids 51:871–880. doi:10.1007/s00348-011-1101-7

Scholz U, Kähler CJ (2004) Automated Image Processing and Segmentation for mask generation in PIV. In: Proceedings of the 12th international symposium on applications of laser techniques to fluid mechanics. Lisbon

Sciacchitano A, Scarano F (2014) Elimination of PIV light reflections via a temporal high pass filter. Meas Sci Technol 25:84009. doi:10.1088/0957-0233/25/8/084009

Theunissen R (2012) Theoretical analysis of direct and phase-filtered cross-correlation response to a sinusoidal displacement for PIV image processing. Meas Sci Technol 23:65302. doi:10.1088/0957-0233/23/6/065302

Theunissen R, Riethmuller ML (2008) Particle image velocimetry in lung bifurcation models. Particle image velocimetry: new developments and recent applications. Springer, Berlin Heidelberg, pp 73–101

Theunissen R, Scarano F, Riethmuller ML (2008) On improvement of PIV image interrogation near stationary interfaces. Exp Fluids 45:557–572. doi:10.1007/s00348-008-0481-9

Usera G, Vernet A, Ferré JA (2004) Considerations and improvements of the analysing algorithms used for time resolved PIV of wall bounded flows. In: Proc 12 th Int Symp Appl Laser Tech to fluid Mech Lisbon, Port July 11–15

Westerweel J (2000) Theoretical analysis of the measurement precision in particle image velocimetry. Exp Fluids Suppl 29(Suppl 1):S003–S012. doi:10.1007/s003480070002

Westerweel J, Stanislas M, Okamoto K et al (2005) Main results of the second international PIV challenge. Exp Fluids 14:170–191. doi:10.1007/s00348-005-0951-2

Willert C (1999) Stereoscopic digital particle image velocimetry for application in wind tunnel flows. Meas Sci Technol 8:1465–1479. doi:10.1088/0957-0233/8/12/010