13th International Symposium on Particle Image Velocimetry – ISPIV 2019
Munich, Germany, July 22-24, 2019

On the feasibility of selective spatial analysis for temporal adaptivity based on confidence statistics

M. Edwards[1*], R. Theunissen[2], C.B. Allen [1]

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

[2] VITO NV, Environmental Modelling Unit, 2400 Mol, Belgium

[*] m.edwards@bristol.ac.uk

Abstract

Particle image velocimetry is frequently used to obtain time-averaged solutions for quasi-stationary flows. Although temporal variance should be small, fluctuations are unavoidable and the magnitude of such fluctuations will naturally vary over the domain. Larger fluctuations in displacement will necessitate more images to be analysed in order to achieve desired levels of confidence in the solution. Spatial displacement gradients will also vary over the Field of View (FOV) and, due to cross-correlation modulation effects, influence the perceived temporal fluctuations, exacerbating the need for more image pairs in these regions. Conventional image processing routines analyse images in their entirety, implying that further analysis of image pairs may be solely driven by small, isolated, non-converged regions of the FOV. This represents a potentially significant inefficiency and is particularly undesirable if time restrictions are imposed, for example, processing images alongside wind tunnel utilisation, or when processing higher-order statistics, wherein convergence may require $N \approx 20,000$ image pairs in some regions. Naive approaches, i.e. imposing a sampling mask immediately surrounding non-converged regions, truncate information and introduce numerical artifacts to the solution. This paper investigates the distance at which a sampling mask may be imposed around some region of interest, such that the solution is unaffected within the region of interest itself. Via experimental analysis, it is shown that a 44% reduction in the required number of correlations for an ensemble solution can be achieved, compared to conventional image processing routines.

1 Introduction

The pixel count of PIV cameras continues to rise, enlarging the attainable spatial resolution and/or Field of View (FOV). While PIV allows for instantaneous displacement measurements over a given FOV, it is also extensively used to acquire time-averaged solutions of a wide variety of flows including for example, rotorcraft Jenkins et al. (2009), formula 1 cars Michaux et al. (2018), or cyclists Jux et al. (2018). To achieve these temporally-averaged solutions, two approaches are commonly adopted. The first is to analyse each pair of images to obtain N displacement fields, where N is the number of image pairs, and subsequently ensemble-average the displacement fields to obtain the mean solution. The alternative is to take the ensemble-averaged correlation map at each vector location to retrieve a single displacement field equivalent to the time-averaged solution Meinhart et al. (2002). This approach is frequently used in micro-PIV or when experimental conditions, such as seeding, are poor Wereley et al. (2012). While both have their advantages and disadvantages, a primary benefit to the former is the ability to easily obtain statistical quantities about the flow velocities, such as the temporal standard deviation, skewness, or kurtosis.

Such statistics may be of interest to the experimentalist but must be reliable, i.e. converged. Satisfactory convergence of the mean displacement field may be obtained using relatively few, $N < 100$, image pairs for relatively steady or straightforward flows, although may be several times larger for more unsteady flows. Conversely, Ullum et al. (1998) note that reliable convergence of higher-order statistics may require a dramatic increase in the required number of image pairs towards $N = 20,000$, posing significant demands on computational resources. While the computational cost to obtain a single mean displacement field may not be significant, even simple parameter sweeps can rapidly increase the number of mean displacements to be analysed. Furthermore, access to wind tunnels (WT) is often restricted due to their running costs, avail-

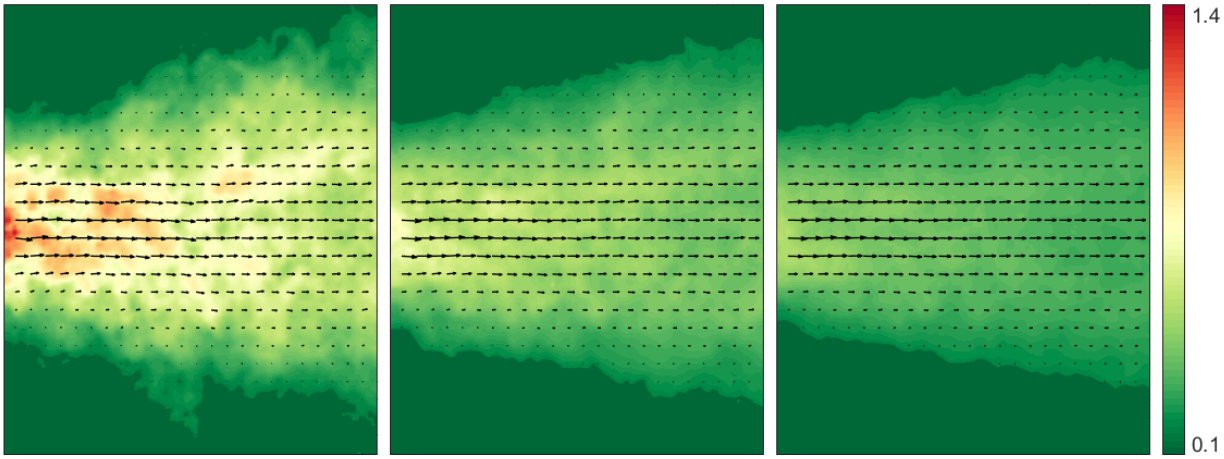

(a) (b) (c)

Figure 1: Evolution of the confidence interval of the mean displacement for an axisymmetric turbulent jet, provided for the second PIV challenge Stanislas et al. (2005), using (a) 20 image pairs, (b) 50 image pairs, (c) 100 image pairs. The CI is capped at 0.1 for the purpose of illustration.

ability, or regulatory bodies, as is the case in Formula 1[1] Federation Internationale de l'Automobile (2019). As such, the experimentalist must often decide which parameter configurations are likely to be the most informative in advance, in order to maximise wind tunnel (WT) utilisation. Reductions in computational requirements to obtain mean displacement fields may conversely enable the user to adjust model-parameter configurations on the fly, based upon preliminary results obtained concurrently with WT run-time, thus offering a considerable benefit towards maximising efficiency of WT usage.

When calculating ensemble averaged solutions, the underlying displacement field (i.e. flow) is required to be quasi-stationary, meaning that the flow velocity should be reasonably constant, with small temporal fluctiations. The magnitude of such temporal fluctuations will naturally vary over the domain due to the underlying flow behaviour. Furthermore, flow features captured within the field of view typically vary spatially, presenting a range of displacement gradients. Strong spatial gradients in displacements are well-known to be detrimental to the accuracy of the correlation outcome, thereby introducing increased uncertainty in the form of artificial displacement fluctuations Westerweel (2008). In addition, it has been previously demonstrated that cross-correlation response is non-linear and is biased towards regions of curvature Theunissen and Edwards (2018). While this produces a constant bias to the mean, it also influences the measured temporal variance, as the amount of bias on a per-image basis will vary depending upon the realisation of particle images within the correlation window. Due to these factors, the number of image pairs required for local statistical convergence will vary over the field of view. To exemplify, the captured FOV may contain regions of little temporal variance (e.g. laminar flow), requiring only few velocity samples to reach statistical convergence in the temporal average, whereas more turbulent regions (actual or artificial) will demand a large number of realisations. This implies that in some image regions statistical convergence might have been reached while additional data analyses (and subsequently computational effort) are driven by isolated, non-converged, image portions. Spatially adaptive sampling strategies may ameliorate the situation by reducing locally the number of correlation windows Theunissen et al. (2007); Edwards and Theunissen (2019), although such approaches do not consider the local convergence and will continue to place correlation windows, albeit fewer of them, in regions which may already be satisfactorily converged. This paper assesses whether it is possible to automatically determine such regions of local convergence, to some user-defined level, as well as whether these regions may be excluded from subsequent analyses, without influencing the solution in the non-converged regions, thereby reducing computational cost without sacrificing accuracy beyond some tolerable level.

The level of convergence in a particular statistic can be assessed by considering its Confidence Interval (CI). Doing so requires calculating the standard deviation of the statistic in question, which can be obtained formally, using standard expressions of variance, or may be approximated using techniques such

[1]Formula 1 teams are restricted to a certain combined total hours WT run-time or equivalent CPU time for CFD. It is thus the prerogative of the team to decide how to split their allocated hours between experimental and numerical.

as bootstrapping Hill (1986); Theunissen et al. (2008). The CI is then obtained by scaling the statistics' standard deviation by a value determined by the desired confidence, typically 95% in engineering applications, and the frequency distribution of the statistic modelled. Where the standard deviation of the statistic population is known, or for sufficiently large N, i.e. typically > 1000, the scaling term is obtained from the standard normal distribution. To obtain the 95% confidence scaling, the distance from the origin to the upper and lower 2.5% tails must be found, which in this case is ± 1.96, and is denoted by z^*. If the number of samples is small or the standard deviation is the sample standard deviation, then the Student's t-distribution should be used to obtain the scaling term, denoted by t. Use of these terms to obtain confidence intervals of the mean, CI_u is shown in Equation 1, where σ_u is the population standard deviation, S_u is the sample standard deviation, and t and z^* are as above, where $t, S_u \rightarrow z^*, \sigma_u$ as $N \rightarrow \infty$. Other frequency distributions, such as the Chi-Squared distribution, may be more appropriate to model particular statistics, however, discussion of such lies beyond the scope of this paper. With the CI calculated, comparison to a threshold value may indicate convergence. For PIV applications, a suitable threshold value for the confidence interval of the mean could be 0.1px, in line with the commonly assumed uncertainty of PIV to be about this level Westerweel and Scarano (2005).

$$CI_{\bar{u}} = z^* \cdot \frac{\sigma_u}{\sqrt{N}} \tag{1a}$$

$$CI_{\bar{u}} = t \cdot \frac{S_u}{\sqrt{N}} \tag{1b}$$

Considering the turbulent jet provided in the second PIV challenge Stanislas et al. (2005), Figure 1 shows how the confidence interval typically decreases as the number of vector fields contributing to the mean increases. Regions in which the confidence is sufficient, c.f. regions in Figure 1 where the CI $<=$ 0.1, can be considered to have reached statistical convergence. Conversely, regions not yet converged necessitate additional samples (i.e. displacement fields) to be taken into account. A simple approach to this problem would be to impose a mask some distance from the Region of Interest (ROI), similarly to image masking, thereby isolating regions of non-convergence and reducing the computational effort (cf. number of cross-correlations). However, this approach artificially truncates information along the flow-mask boundary, which has repeatedly been found to influence correlations near such an interface, unless using advanced processing techniques Masullo and Theunissen (2017); Gui et al. (2003); Ronneberger et al. (1998). The challenge thus lies in defining a mask around the ROI, which does not influence the solution within the ROI, relative to a full conventional solution, while minimising the area of the image to be analysed. The next section will address how the authors achieve such a selective analysis, followed by a demonstration of the attainable benefits in the results section.

2 Methodology

2.1 Calculating the Region of Interest

As detailed previously, the ROI must first be identified by locally assessing the convergence of the solution over the field of view. Although this methodology can be applied to any statistic of interest, utilising bootstrapping where necessary, it shall herein be assumed that 'convergence' refers to convergence of the mean solution, unless otherwise stated. To determine the CI of the mean, the variance of the solution is needed. In a straightforward manner, this necessitates storage of each displacement field and extending the database with every additional displacement field. This imposes vast memory requirements to store each displacement field, such that they can be re-referenced to obtain the new mean and variance each iteration. However, efficient and stable updating algorithms exist, see eq. (2), which can calculate the mean and variance without the need to store all contributing data Chan et al. (2017).

$$M_N = M_{N-1} + (u_N - M_{N-1})/N \tag{2a}$$

$$S_N = S_{N-1} + (u_N - M_{N-1}) \cdot (u_N - M_N) \tag{2b}$$

Here u_N signifies a component of the N^{th} displacement field, M_N is the mean of all N samples (u_i with $i = 1 \ldots N$) and $S_N/(N-1)$ is the variance of N data values. Computing the ROI therefore consists of first updating the variance using (2b) for each pixel, then multiplying by the scaling term $z*$ or t, and finally comparing this to the user-defined CI threshold. Pixels exceeding the threshold thus represent the ROI to analyse in image pair $N+1$.

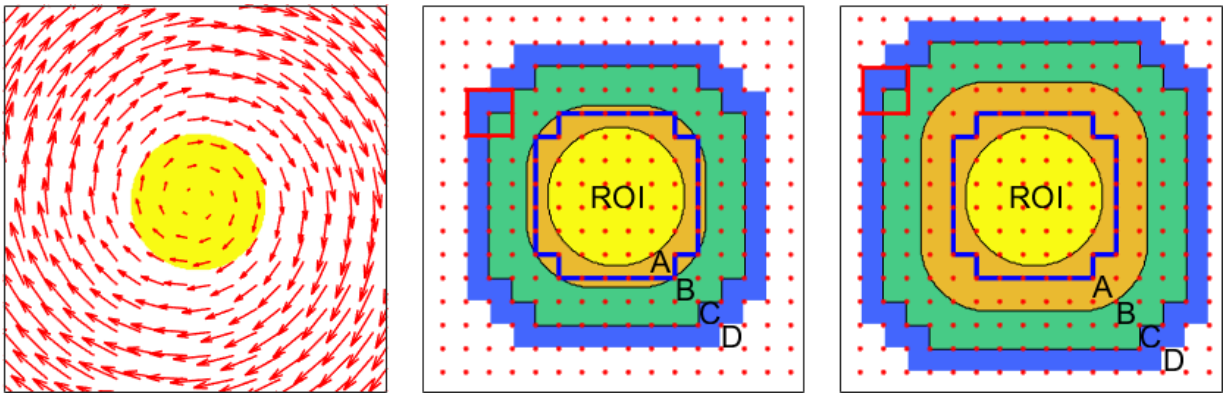

(a) (b) (c)

Figure 2: Depiction of the analysis support radius for an arbitrary region of interest (yellow), showing; the interpolation kernel *A*, the vector validation kernel *B*, image deformation kernel *C*, and finally the region with no influence over the ROI, *D*, for both bi-linear, (b), and bi-cubic, (c), interpolation kernels. The blue line encloses the region equivalent to a full solution following the proposed methodology, and is at least as large as the requested ROI.

2.2 Masking process

Having determined the pixels which constitute the ROI, the next step is to determine a sampling mask which minimises the number of correlations, without influencing the solution within the ROI, relative to a full solution (FS) considering the entire image. Working backwards from the final iteration, the sample mask must extend far enough beyond the ROI to include vectors whose interpolation kernel would overlap the ROI, and thus influence it. For bi-linear interpolation this corresponds to the four enclosing vectors, i.e. a 2x2 stencil, whereas for bi-cubic interpolation the 4x4 stencil of vectors surrounding the pixel must be considered. This region can be obtained by dilating the ROI, in both x and y, by an amount h (bi-linear) or $2h$ (bi-cubic), where h is the spacing between samples, and is indicated by region A in Figures 2(b) and 2(c). According to this methodology, each pixel within a 2x2 stencil would invoke the same four sample locations. Hence if any pixel in such a region is defined as ROI, then it follows that all pixels within the stencil will be equivalent to the FS, and can contribute to the ensemble statistics. The same analysis holds true for cubic interpolation, with any pixel in a given 2x2 stencil invoking the same 4x4 interpolation stencil. The total 'valid region' which can contribute to the ensemble statistics, is therefore often larger than the ROI itself and is depicted in Figures 2(b) and 2(c) by the blue line. Such behaviour is beneficial to the method, since converged pixels may still receive additional sample data, further improving confidence in this location, at no additional computational cost.

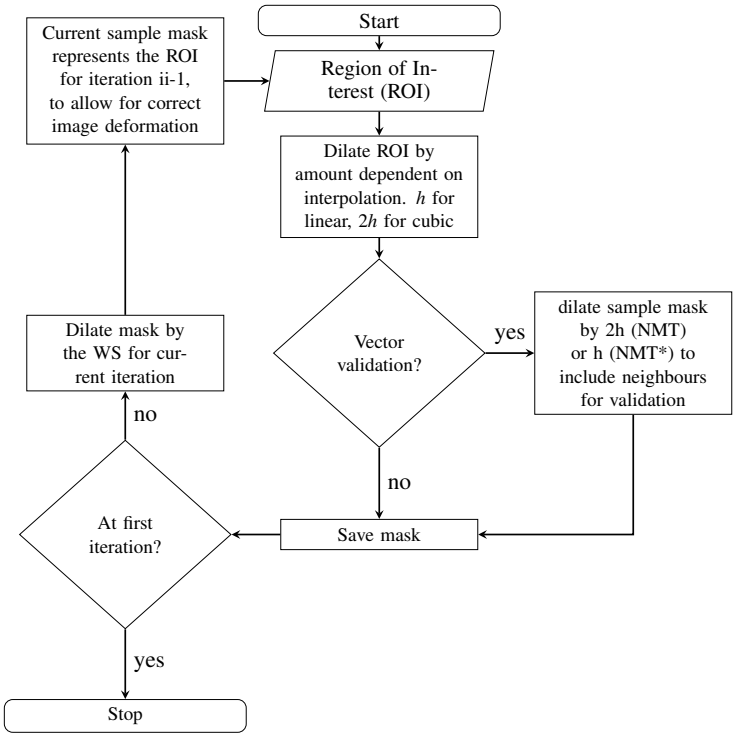

Figure 3: Flow chart showing masking process

Prior to interpolation, vectors are typically subject to validation, using schemes such as the normalised

median theshold test (NMT) Westerweel and Scarano (2005) to detect outliers. The NMT compares each vector to its eight surrounding neighbours from a 3x3 grid. If the vector disagrees with its neighbours beyond some pre-defined threshold, then it is replaced by the analysis routine. The replacement vector is typically calculated as the mean of the valid eight neighbours, or zero if all are invalid. The consequence of this process is that for outlier *detection* equivalent to a FS, the interpolation region must be dilated by another sample spacing, h, to include neighbouring vectors. However, to reproduce the equivalent outlier *replacement*, the validity of all neighbouring vectors must be known, necessitating that the mask be dilated instead by $2h$.

While dilation by $2h$ is required to guarantee numerical equivalence to the FS, dilating by h is generally sufficient for approximate results; vectors at the edge of the sampling mask are guaranteed to have at least three immediate neighbours, and typically five, which can be used for an approximate vector validation attempt. The frequency of incorrect classification from the reduced number of neighbours is relatively low. Drastically anomolous vectors are still reliably detected using just five neighbours, while the reduced neighbour validation typically only fails when the scrutinised vector is only subtly different to its neighbours, in turn leading to only small changes in the replaced vector. Herein, dilation by $2h$ for vector validation, maintaining numerical equivalence, will be herein referred to using NMT, whereas dilation by h, accepting some numerical differences, will be referred to as NMT*. Differences of ensemble mean values using NMT compared to NMT* are typically $< 1\%$, i.e. at least two orders of magnitude smaller than the imposed 'acceptable' threshold, yet, require fewer correlations to obtain.

For a single pass analysis the aforementioned approach, i.e. correlating all vectors within region B, would be sufficient to ensure equivalence within the ROI. However, many PIV analyses make use of iterative routines to obtain improved spatial resolution Scarano (2002). This necessitates the consideration of image deformation in the selection of the sampling mask, such that the correlation response of all the windows within region B are equivalent to a FS. For this methodology to hold true at the final iteration, K, the image within every pixel of every correlation window within 'B' must be deformed following iteration $K-1$ in accordance with the FS. This region can be found by convoluting each 'active' sample location with the current window size (WS), as indicated by region 'C' in Figure 2. Since image deformation, following the methodology of Unser et al. (1993), requires the predictor for iteration $K-1$ at each of these pixels to be correct, region C effectively serves as the ROI for iteration $K-1$. The process is repeated until the first iteration is reached and a sample mask is obtained for each iteration. The described approach is presented as a flowchart in Figure 3, and with each specific sampling region graphically highlighted in Figure 2 for clarity. A standard image analysis routine can then be trivially modified to correlate only sample locations which lie within this region, and ignore everything else, i.e. region 'D', with no negative effects on the solution in the ROI.

It is relatively trivial to modify vector validation and structured interpolation codes to consider only the 'active' region. In doing so, the computational cost for these processes is proportional to the number of vector locations. Image deformation requires reconstructing the image based on the predictor values over the image. This reconstruction operates pixel-wise and hence the code can be easily modified to exclude

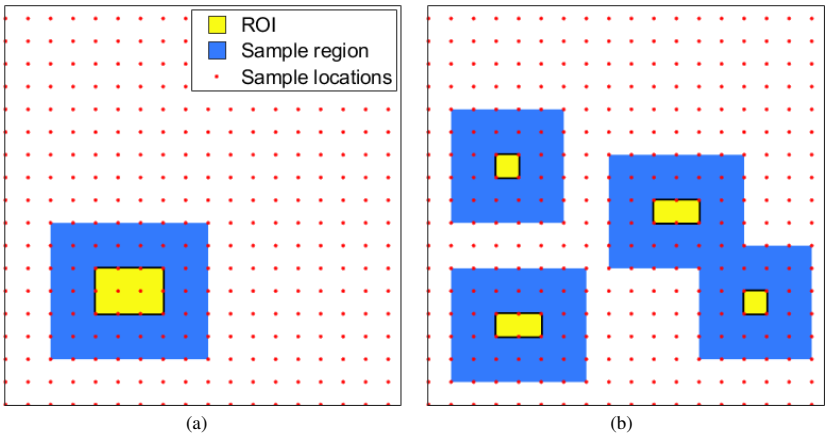

Figure 4: Influence of the composition of the ROI on the number of samples required for (a) a single contiguous ROI, requiring 56 samples for six 'blocks' of ROI, and (b) several distributed ROIs, requiring 150 samples for six 'blocks' of ROI.

Table 1: Values used for each setting as part of the parameter sweep

Parameter	Default	Parameter sweep values
Initial Window Size	97	N/A
Final Window Size	25	(15, 35, 55)
Window Overlap Ratio	0.5	(0, 0.5, 0.75, 0.85)
Number of main iterations	3	(2, 3, 4)
Number of refinement iterations	1	(0, 1, 2, 3, 4)
Interpolation method	Linear	(Linear, Cubic)
Vector validation method	NMT	(none, NMT, NMT*)

pixels outside of the sampling mask from the deformation process, resulting in a computational cost scaling with the number of pixels in the ROI. The reconstruction is applied to filtered versions of the raw images, and such filtering is not be easily reduced to arbitrary sub-regions. However, this step is only required once per image and yields a marginal computational overhead. Hence the methods computational costs scales roughly linearly with increasing ROI extent. The method therefore reduces overall computational costs relative to a conventional analysis, for ROI's which cover $\approx <= 95-97\%$ of the available domain, due to the small overhead of calculating the sampling masks.

3 Results

3.1 Numerical investigation

Changing PIV interrogation parameters such as window size, window overlap ratio (WOR), choice of interpolation kernel and so on, will naturally influence the extent of the resulting sampling masks. For example, cubic interpolation requires dilation by $2h$, instead of h for linear interpolation, therefore demanding more correlations to be performed for a given ROI (Figure 2). To understand the extent to which each parameter affects the total number of correlations required for a given ROI size, a parameter study was conducted. The shape of ROI in this study was kept as a central square with varying area. As illustrated in Figure 4, undoubtedly the composition of the ROI will also influence the total number of correlations. While this effect is difficult to quantify, a generally valid tendency is that fewer, larger, contiguous ROIs, (which are typically encountered in practical situations), will be more effective at reducing the required number of correlations than many scattered regions.

The default parameter configuration for this study is shown in Table 1. For each parameter, a range of typical values were substituted, keeping everything else fixed, and the resulting number of correlations recorded and compared to the total number of correlations required for the full analysis with the same settings. Additionally, each configuration was applied to a range of ROI sizes. Results are presented in Figure 5. In each iteration, the sampling mask dilates by $\approx \eta h$, including the dilation by an amount $WS/2$ (i.e. the correlation window size around each sample) prior to iteration $K-1$. Parameter η is a function of the analysis settings, e.g. $\eta \approx 4.5$ when invoking cubic interpolation with NMT validation. With each additional iteration the sampling mask in earlier iterations increases in size, correlating a larger proportion of the overall image and reducing the efficiency of the method, see Figures 5(d) and 5(e). Settings such as interpolation or vector validation method influence the scaling term η, and thus increase or decrease the extent of the sampling mask around the ROI by some fixed proportion, shown in Figures 5(f) and 5(c). Finally, the WOR, and final WS have a much more significant impact on the overall efficiency as can be deduced from Figures 5(b) and 5(a), due to their influence on the vector spacing itself. Reducing h, either by reducing the WS or by increasing the overlap ratio, does not significantly influence the number of correlations performed[2], i.e. η is unchanged. However, a smaller spacing increases the total number of sample locations which must all be correlated by the FS, and thus a smaller *fraction* of correlations are performed by the proposed method, improving its efficiency. The corollary of the above is that more intensive interrogation parameters, i.e. smaller final WS, larger overlap ratio, with simple vector validation schemes benefit the most from the methodology presented here.

[2]beyond the fact that there may exist more sample locations within the ROI itself

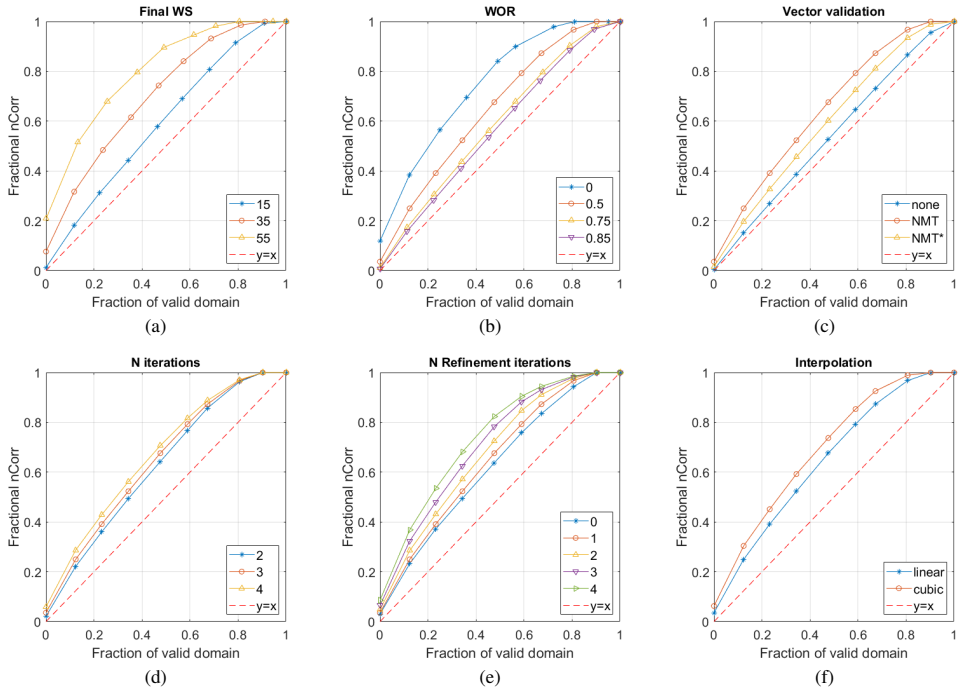

Figure 5: Comparison of the speedup effect of a locally selective WIDIM approach for changing (a) final WS. (b) WOR. (c) Vector validation approach. (d) Number of Iterations. (e) Number of Refinement Iterations. (f) Interpolation method.

3.2 Experimental analysis

To verify the methodology in a real application, it was applied to the flow behind a blunt trailing edge, investigated by Elsahhar et al. (2018) for its aeroacoustic properties. In this study, the effect of trailing edge bluntness on the generation of wake-vortex noise was investigated using Particle Image Velocimetry. Shown in Figure 6 is the reference flow field for this particular case. Since a moderate proportion of the FOV is relatively uniform freestream flow, these images serve as exemplary candidates for the proposed method. Figure 7 shows one of the (in total) 115 raw images, and highlights the strong reflection present in nearly all of the images. An ensemble-minimum background subtraction was applied to all images, removing the majority of the reflections (Figure 7(b)). The reflection is not identical in every image, and hence remnants remain in the images, varying throughout the ensemble. While more advanced pre-processing techniques could certainly be employed to remove such remnants, it does not adversely affect the current investigation and hence these techniques have not been investigated further.

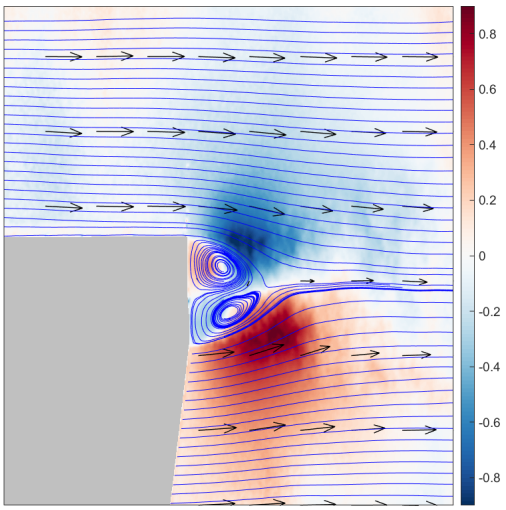

Figure 6: Ensemble averaged, reference flow field obtained by averaging all 115 measured displacement fields in their entirety

The ensemble of images were analysed according to 3 different schemes. First, a reference conventional solution was established using conventional WIDIM Scarano (2002) adopting the default parameters as in Table 1. However, the discrete displacement data (following the cross-correlation of interrogation windows) was interpolated pixel-wise using bi-cubic kernels rather than the bi-linear scheme listed in Table 2. Subsequently, the presented methodology was used to obtain the ensemble mean at a 95% confidence, of being within 0.1px and 0.05px, using the same parameter configuration

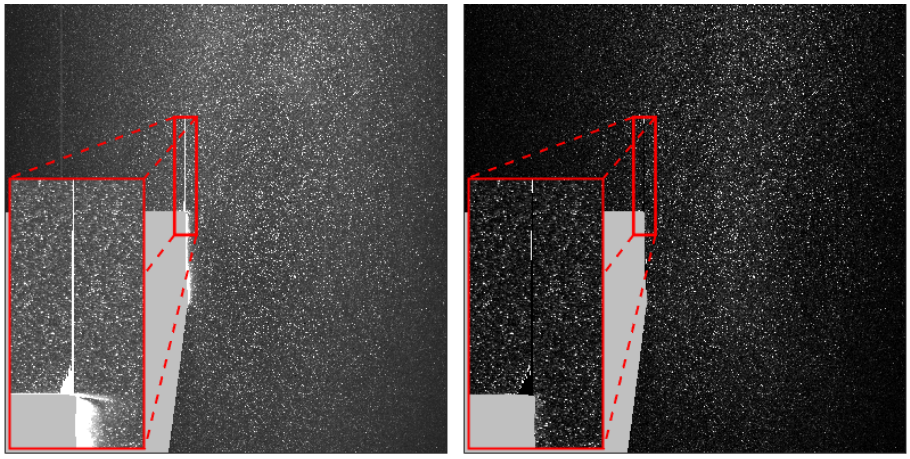

(a) (b)

Figure 7: Exemplary PIV image from Elsahhar et al. (2018). (a) raw image with zoom inset showing reflection, (b) pre-processed image with inset showing diminished reflection. Contrast enhanced for clarity.

as the reference solution. Due to the relatively small number of samples (115) and the unknown standard deviation, (1b) was used to obtain the pixel-wise confidence interval after each timestep, using the Student's t-distribution to calculate t. Regions where the temporally adaptive method uses all 115 images, e.g. regions where the CI exceeds the threshold, such as in the wake region of Figure 6, the solution is numerically equivalent to the reference solution. Outside of these regions, where less than the full 115 images have been considered, the solution is expected to differ from the reference to some degree, depending on the threshold level. Figure 8 shows how much the temporally adaptive solutions differ from the reference WIDIM solution.

The number of samples contributing to the ensemble mean, over the entire domain, for each threshold setting is shown in Figure 9. In both cases, the freestream requires relatively few samples to reach convergence compared to the wake. This is in line with the expressions for the confidence levels in (1. As the local standard deviation increases, the number of independent samples required to reach a pre-defined confidence level increases quadratically. It can also be seen that the stricter tolerance requires more samples in the fringes of the wake, relative to the more relaxed threshold. Although the majority of the image

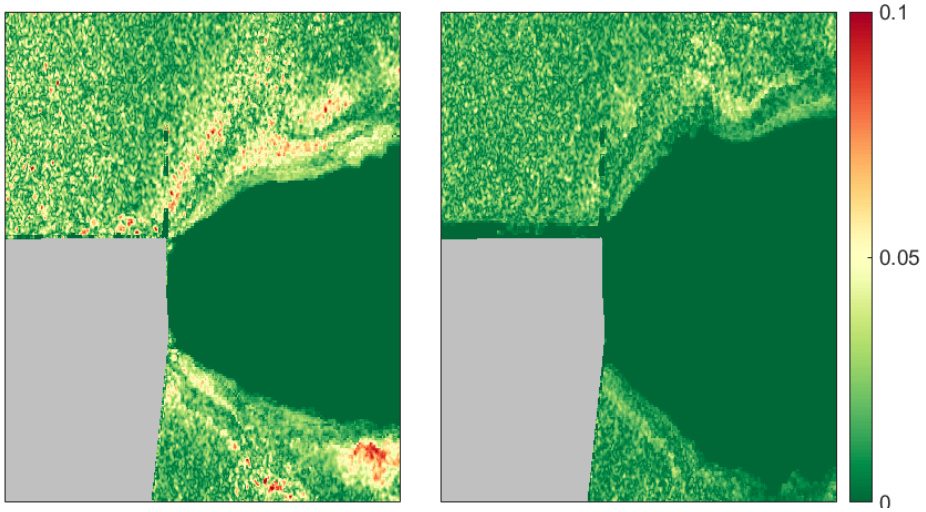

(a) (b)

Figure 8: Effect of changing confidence intervel threshold on the magnitude difference to a reference full solution, for (a) 0.1px threshold and (b) 0.05px threshold.

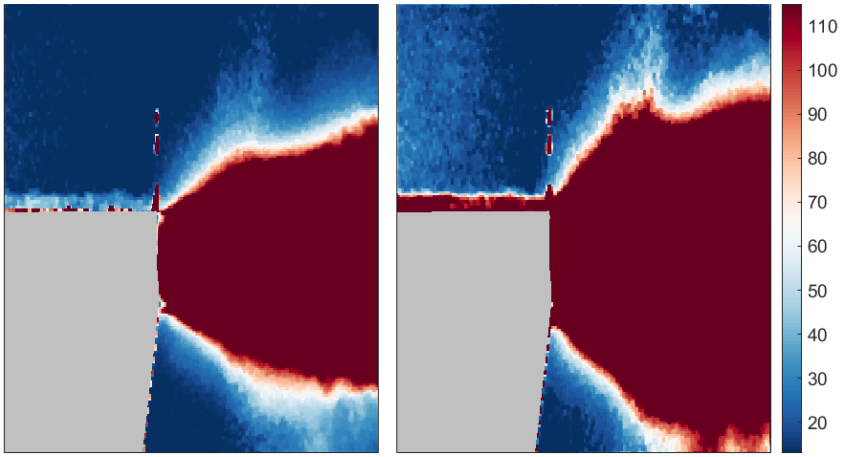

(a) (b)

Figure 9: Number of contributing samples for a threshold of (a) 0.1px, and (a) 0.05px.

reflection was removed, sections of this remained which varied through time. When the relection is present, the displacement will be biased towards 0, and may in fact register as zero displacement, depending on whether this is identified as an outlier. As such, in these locations the recorded displacement will be varying between the 'true' displacement, a displacement biased towards zero, and, occasionally, zero displacement. This results in significantly increased standard deviation in this region and thus necessitates further sampling according to the current criteria. Comparing the total number of correlations for the entire ensemble, Table 2, the strength of the proposed method is made apparent. A reduction of 44% and 28% in the number of correlations is achieved for the 0.1px and 0.05px thresholds, respectively. While overheads exist beyond the correlations, such as interpolation and image deformation, these typically represent a small proportion of the overall computational cost. The number of total correlations is further reduced by using the NMT* approach as described previously, reducing the number of correlations by $\approx$5% in both cases.

Table 2: Comparison of the number of correlations performed for each of the approaches

	Number of correlations
Reference Full Solution	11,867,770 (1.00)
0.05px threshold	8,583,713 (0.72)
0.1px threshold	6,689,502 (0.56)

4 Conclusions

A method has been proposed which automatically reduces the interrogation area subject to satisfactory convergence of arbitrary flow statistics. Convergence is determined by calculating the confidence interval for any arbitrary statistic and comparing to a user-defined threshold value. An automatic Region of Interest (ROI) is obtained which is used to determine a set of sampling masks, one for each iteration. Such a mask identifies which locations influence the solution within the ROI, such that if all are correlated the final solution within said ROI remains numerically identical to an equivalent conventional solution.

This approach can significantly reduce the computational cost required to obtain domain-wide ensemble statistics converged to a pre-defined level. The method was applied to the flow behind a blunt trailing edge, being investigated for its aero-acoustic properties. For typical PIV settings, the number of correlations to obtain the ensemble mean was reduced by 44%. The method is particularly of use when rapid mean solutions are required, i.e. analysing images concurrently with wind tunnel utilisation, or when processing higher-order statistics.

References

Chan TF, Golub GH, and Leveque RJ (2017) Algorithms for Computing the Sample Variance : Analysis and Recommendations. *American Statistician* 37:242–247

Edwards M and Theunissen R (2019) Adaptive incremental stippling for sample distribution in spatially adaptive PIV image analysis. *Measurement Science and Technology*

Elsahhar W, Showkat Ali SA, Theunissen R, and Azarpeyvand M (2018) An Experimental Investigation of the Effect of Bluff Body Bluntness Factor on Wake-vortex Noise Generation

Federation Internationale de l'Automobile (2019) FIA Sporting Regulations

Gui L, Wereley ST, and Kim YH (2003) Advances and applications of the digital mask technique in particle image velocimetry experiments. *Measurement Science and Technology* 14:1820–1828

Hill T (1986) Institute of Mathematical Statistics is collaborating with JSTOR to digitize, preserve, and extend access to Statistical Science. ® www.jstor.org. *Statistical Science* 10:354–363

Jenkins LN, Yao Cs, Bartram SM, Harris J, Allan B, Wong O, and Mace WD (2009) Development of a Large Field-of-View PIV System for Rotorcraft Testing in the 14- x 22-Foot Subsonic Tunnel. *Aiaa*

Jux C, Sciacchitano A, Schneiders JF, and Scarano F (2018) Robotic volumetric PIV of a full-scale cyclist. *Experiments in Fluids* 59:1–15

Masullo A and Theunissen R (2017) On the applicability of numerical image mapping for PIV image analysis near curved interfaces. *Measurement Science and Technology* 28

Meinhart CD, Wereley ST, and Santiago JG (2002) A PIV Algorithm for Estimating Time-Averaged Velocity Fields. *Journal of Fluids Engineering* 122:285

Michaux F, Mattern P, and Kallweit S (2018) RoboPIV: How robotics enable PIV on a large industrial scale. *Measurement Science and Technology* 29

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

Scarano F (2002) Interative image deformation methods in PIV. *Measurement Science and Technology* 13:R1–R19

Stanislas M, Okamoto K, Kähler CJ, and Westerweel J (2005) Main results of the Second International PIV Challenge. *Experiments in Fluids* 39:170–191

Theunissen R, Di Sante A, Riethmuller ML, and Van Den Braembussche RA (2008) Confidence estimation using dependent circular block bootstrapping: Application to the statistical analysis of PIV measurements. *Experiments in Fluids* 44:591–596

Theunissen R and Edwards M (2018) A general approach to evaluate the ensemble cross-correlation response for PIV using Kernel density estimation. *Experiments in Fluids* 59:1–17

Theunissen R, Scarano F, and Riethmuller ML (2007) An adaptive sampling and windowing interrogation method in PIV. *Measurement Science and Technology* 18:275–287

Ullum U, Schmidt JJ, Larsen PS, and McCluskey DR (1998) Statistical analysis and accuracy of PIV data. *Journal of Visualization* 1:205–216

Unser M, Aldroubi A, and Eden M (1993) B-Spline Signal Processing : Part I-Theory 41

Wereley ST, Gui L, and Meinhart CD (2012) Advanced algorithms for microscale particle image velocimetry. *AIAA Journal* 40:1047–1055

Westerweel J (2008) On velocity gradients in PIV interrogation. *Experiments in Fluids* 44:831–842

Westerweel J and Scarano F (2005) Universal outlier detection for PIV data. *Experiments in Fluids* 39:1096–1100