

Validation of Multi-Frame PIV Image Interrogation Algorithms in the Spectral Domain

Steven J. Beresh,^{1,*} Douglas R. Neal,² and Andrea Sciacchitano³

¹ Sandia National Laboratories, Albuquerque, NM, U.S.A.

² LaVision Inc., Ypsilanti, MI, U.S.A.

³ Delft University of Technology, Delft, The Netherlands

* sjberes@sandia.gov

Abstract

Multi-frame correlation algorithms for time-resolved PIV have been shown in previous studies to reduce noise and error levels in comparison with conventional two-frame correlations. However, none of these prior efforts tested the accuracy of the algorithms in spectral space. Even should a multi-frame algorithm reduce the error of vector computations summed over an entire data set, this does not imply that these improvements are observed at all frequencies. The present study examines the accuracy of velocity spectra in comparison with simultaneous hot-wire data. Results indicate that the high-frequency content of the spectrum is very sensitive to choice of the interrogation algorithm and may not return an accurate response. A top-hat-weighted sliding sum-of-correlation is contaminated by high-frequency ringing whereas Gaussian weighting is indistinguishable from a low-pass filtering effect. Some evidence suggests the pyramid correlation modestly increases bandwidth of the measurement at high frequencies. The apparent benefits of multi-frame interrogation algorithms may be limited in their ability to reveal additional spectral content of the flow.

Introduction

Conventional PIV image interrogation cross-correlates two frames at a small time separation, an implementation that has been familiar since the advent of PIV [1]. But the use of only two frames in a correlation places limits on the accuracy and the dynamic range of the velocity measurement (e.g., [2]). One means of improving this situation is to correlate over multiple frames. Three or even four pulses may be recorded in rapid sequence and used to generate a more accurate correlation [2], which has been demonstrated by Ding and Adrian [3]. In this implementation, a triple- or quadruple-pulsed system would generate a single vector field, repeated at the repetition rate possible by the light source and camera framing rate. A three- or four-frame camera analogous to the two-frame interline-transfer CCD camera does not presently exist.

The development of time-resolved PIV (TR-PIV) extends this concept by making long, continuous sequences of images available for interrogation algorithms. As TR-PIV has become widespread, several different multi-frame correlation algorithms have emerged. One approach is to continue conventional two-frame correlation, but to optimize the temporal separation for each vector for maximum accuracy and dynamic range [4]. The most straightforward extrapolation from the triple- or quadruple-frame correlation is the sliding sum-of-correlation (SoC) [5,6], which averages over some N images in correlation space to produce a more accurate correlation. The key difference from the triple- or quadruple-frame approach is that the latter produces a single vector field from the several exposures whereas SoC “slides” over the continuous sequence of images to produce a similar continuous sequence of vector fields. A modification of the SoC method adds Gaussian weighting over the sliding window to improve the correlation accuracy, analogously to the use of such weighting spatially to mitigate the effects of discontinuities at the edges of interrogation windows [7,8].

Even more sophisticated algorithms have been invented. The pyramid correlation from Sciacchitano *et al* [8] offers greater flexibility and claims superior accuracy to SoC. It operates similarly to SoC but it varies the number of frames and the stride between them on a vector-by-vector basis to individually optimize the accuracy. Lynch and Scarano [9] follow the fluid trajectory to translate and rotate interrogation windows comprising the multi-frame correlation. The same concept of following the local trajectory has been found in additional algorithms for TR-PIV [10,11], and these bear some resemblance to the class of Lagrangian Particle Tracking algorithms such as Shake-the-Box employed in three-dimensional Particle Tracking Velocimetry [12,13].

Each advancement of multi-frame correlation analysis has rigorously examined the accuracy of the method in comparison to conventional two-frame correlation and documented the overall improvement in a statistical sense. Distributions of accuracy and fluctuations summed over the entirety of all vectors in the data set show an improvement of each successive multi-frame interrogation algorithm. However, none of these prior studies is concerned with the accuracy of the algorithms in spectral space. This is unfortunate because a key attraction of TR-PIV is its ability to return spectral content of the flowfield. Just because these algorithms improve the accuracy and dynamic range of the statistical populations of measured vectors does not imply that these improvements are equally distributed across the frequency content of the TR-PIV signal and therefore does not guarantee an improvement to the measured spectrum.

TR-PIV recently has shown application to the measurement of turbulence decay properties at high frequencies [14-20], whose interpretation may be sensitive to any dependence of interrogation accuracy upon frequency content. When Beresh *et al* extended their 400 kHz pulse-burst PIV [18] to MHz rates [21], the intent was to employ multi-frame interrogation on the oversampled data to reduce measurement noise and thereby raise the effective frequency limits of the measurement of the turbulence spectrum. However, application of different multi-frame algorithms returned different high-frequency spectral content and it was ambiguous which should be considered most correct. This motivates the present investigation to examine the performance of accepted TR-PIV multi-frame interrogation algorithms with regards to the spectral content they return. To accomplish this, the data set of Neal *et al* [22] has been chosen because it includes simultaneous hot-wire anemometry (HWA) data to compare to its TR-PIV measurements. Though originally used to assess the efficacy of several uncertainty quantification techniques [23], it can be repurposed here to assess multi-frame image interrogation instead. The present effort is designed to answer whether multi-frame TR-PIV image interrogations accurately measure spectral content.

Experimental Configuration

A brief description of the Neal *et al* experiment is provided here, with the full details to be found in Ref. [22]. Measurements were acquired in a rectangular jet of dimensions width $w=72.8$ mm and height $h=10.2$ mm with exit velocity of $U_0=5$ m/s. The jet is initially laminar then transitions and becomes fully turbulent in its far-field. The measurements used here are taken from two positions: a downstream position of $x/h=19.0$ and $y/h=0.0$ at which the flow is fully turbulent, and a nearer position of $x/h=4.0$ and $y/h=0.0$ at which the laminar jet is beginning to become unstable. An additional case was studied at $x/h=4.0$ and $y/h=0.4$ where large-scale vortex roll-up occurs at the edge of the jet, but these results were found to echo the unstable laminar position and therefore are omitted here.

Two TR-PIV measurements were made simultaneously, both using a single-cavity Nd:YLF laser as the light source with a maximum repetition rate of 10 kHz. One was termed the measurement system, or PIV-MS, and used a single LaVision HighSpeedStar 5 CMOS camera to image a wide field of view. This system produced noisy results and therefore may be more revealing of higher-accuracy analysis. The second TR-PIV system, termed the high dynamic range system, or PIV-HDR, provided higher-quality results. This system used high magnification to view a smaller region of the jet and arranged two CMOS cameras such that one was normal to the laser sheet and could be used for two-component PIV analysis while the second was set at an angle to use as a stereo pair to the first. The present analysis uses only one camera in 2C mode. The data were acquired at 10 kHz with 20,000 samples. The original data analysis of Neal *et al* used four iterations of 32×32 pixel interrogation windows with 75% overlap for the HDR data and four

iterations of 16×16 pixel windows with 75% overlap for the MS data. The same settings are used herein. Power spectral density (PSD) functions of the velocity fluctuations were computed using the Welch periodogram method, segmented into blocks of 1000 vector fields with 50% overlap between segments.

In addition, measurements from a hot-wire anemometer were acquired with the probe head positioned immediately downstream of the PIV field of view, within 1-2 mm. The single-wire probe was aligned with the streamwise direction and had a sensitive region of 1 mm. Data were acquired at 30 kHz for a record length of 1,200,000 samples.

Results

The PSD's of the PIV-HDR data and the hot-wire data in the fully turbulent region were given by Neal *et al* in their Fig. 20 [22]. It is reproduced here as Fig. 1 with their original TR-PIV spectrum labeled "MST" and overlaid with reprocessed data labeled "two-frame." The latter have been interrogated using Davis 8.4.1 and a conventional two-frame correlation stepped across the time series using the same four iterations of 32×32 pixel windows with 75% overlap as in the MST data. This serves as a baseline for subsequent comparisons. Whereas the MST spectrum was derived from a single vector sequence, the newly analyzed data averaged spectra from a 6 vector by 7 vector block immediately upstream of the hot-wire probe to improve convergence, reducing the apparent "thickness" of the plotted spectrum. As will be seen in subsequent figures, this is necessary to identify more subtle differences between varied interrogation approaches. The flow characteristics change imperceptibly over the vectors within this block and therefore there is no significant loss of information by averaging across them. Figure 1 demonstrates that the baseline image interrogation used in the present work correctly returns the same spectrum as the original analysis in Neal *et al*. The agreement is near perfect for the HDR data set but the new interrogation returns a slightly elevated spectrum in comparison to the MST spectrum. The reason for this latter discrepancy has not been identified but it is not an impediment to any of the multi-frame analysis to follow because it all originates from this baseline.

The first multi-frame algorithm to be tested is the sliding sum-of-correlation (SoC) without any weighting of the selected images – that is, a top-hat filter is applied across the cross-correlation functions of the N images used. The resulting spectra are shown in Fig. 2 alongside the baseline two-frame correlation of Fig. 1 and the hot-wire spectrum. Two sets of spectra are shown, one for the fully turbulent case and another for the unstable laminar case.

While the noise floor of the two-frame correlation is reduced when using the SoC, this alone does not guarantee that the multi-frame correlations lead to more accurate high-frequency spectra as judged by comparison with the hot-wire data. The $N=3$ SoC spectrum for the HDR data falls nearest to the hot-wire spectrum in Fig. 2a but still betrays a distortion near where the hot-wire begins to roll off to its noise floor (about 2-4 kHz). As N is increased, the spectrum does not move towards closer agreement with the hot-wire even though the correlation theoretically should become more accurate. In fact, the distortion becomes more prominent. It is speculated that this increasing distortion is a form of acceleration error, which is a known source of error in PIV analysis [2,24]. If the flow experiences acceleration between successive frames, then the displacements between those frames will vary as well. As multiple correlations are summed, their peaks will not coincide and therefore correlation broadening or peak splitting will occur. This adds additional noise or even bias to the correlation, and this effect may increase with the use of

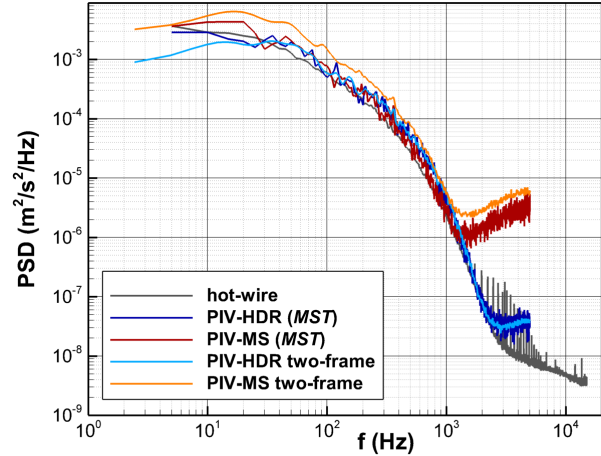


Fig. 1: Comparison of spectra from Neal *et al* [22] in the fully turbulent region, both hot-wire and TR-PIV, to a newly analyzed PIV spectrum using conventional two-frame correlation.

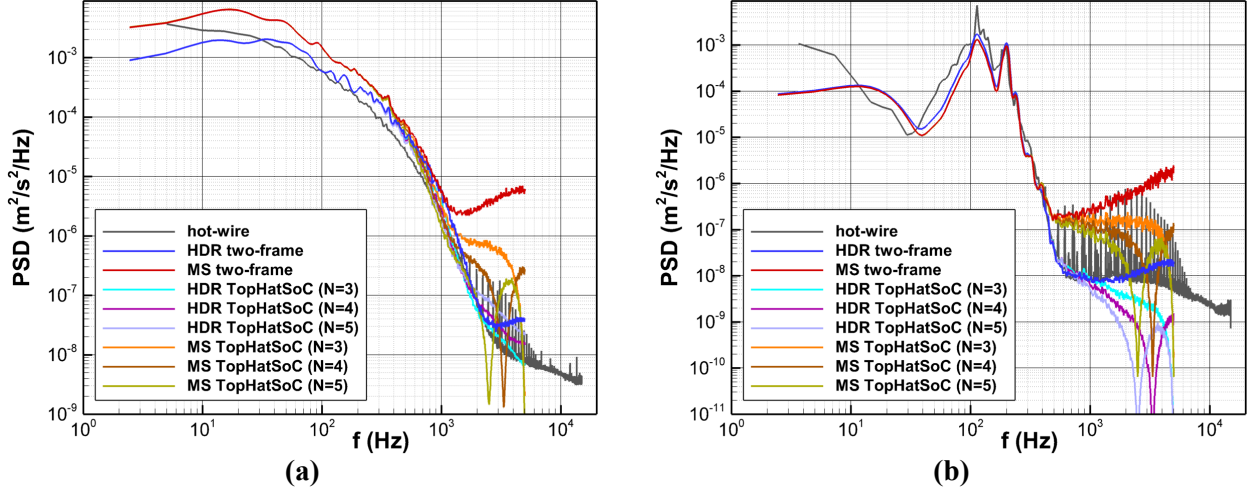


Fig. 2: Comparison of two-frame spectra from Fig. 1 to multi-frame correlation using a sliding sum-of-correlation (SoC) of N images and a top-hat filter. (a) fully turbulent region; (b) unstable laminar region.




Fig. 3: Comparison of two-frame spectra from Fig. 1 to multi-frame correlation using a sliding sum-of-correlation (SoC) of N images and a Gaussian filter. (a) fully turbulent region; (b) unstable laminar region. The arrow in (b) is explained in the text.

additional frames as the displacements become more diverse. In this manner, the distortion observed in the HDR data of Fig. 2a worsens as N increases.

The noise levels for the MS case in Fig. 2a never come into agreement with the hot-wire data no matter how many images are used. The distortions to the HDR spectra of Fig. 2a are not observed here, probably because the unstable laminar flow does not experience acceleration error as does the turbulent case. Instead, the top-hat filter induces high-frequency undulations in the amplitude that clearly are non-physical. This effect is most clearly observed for larger values of N and is stronger for the unsteady laminar case in Fig. 2b in which the noise floor extends for a larger portion of the spectrum. (Note the y -axis in Fig. 2b is expanded to fully reveal the high-frequency spectral undulations.) Even for the HDR case, no continuation of the decay slope emerges from SoC interrogation; instead, the noise floor is filtered and then begins ringing as in the MS case. These artifacts are similar to the response of a top-hat filter in digital signal processing [25] and observed spatially in iterative PIV analysis with image deformations if appropriate weighting is not applied [26]. Though the amplitude of the high-frequency noise is much reduced from the two-frame correlation, the true spectral character of the flow is not revealed. These spectra produced using

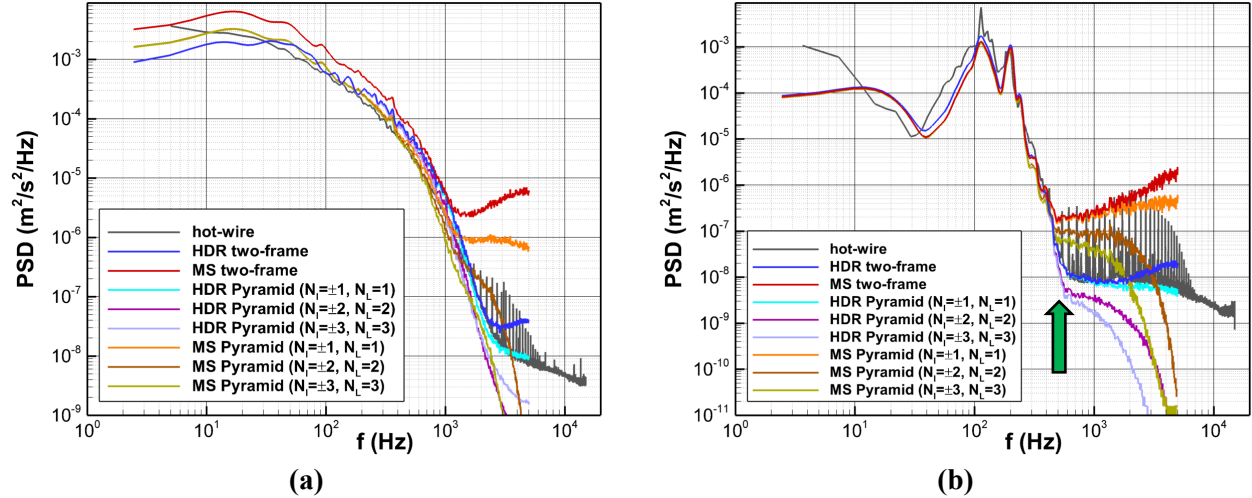


Fig. 4: Comparison of two-frame spectra from Fig. 1 to multi-frame correlation using the pyramid correlation. (a) fully turbulent region; (b) unstable laminar region. The arrow in (b) is explained in the text.

a top-hat filter indicate clearly that an appropriate weighting function is needed to prevent spectral distortions.

When the sliding sum-of-correlation is weighted by a Gaussian filter, strikingly different spectra emerge as shown in Fig. 3. Note that here N is provided as a symmetric range whereas N was absolute when considering the top-hat filter in Fig. 2. Here, as N is increased, the spectra roll off at earlier frequencies and show classic signs of low-pass filtering. This is easiest to identify for the HDR data. The $N=\pm 1$ spectrum for the turbulent case in Fig. 3a shows increased agreement with the hot-wire data but offers only a mild reduction in the noise floor (it actually appears to match the noise of the hot-wire, likely a coincidence). The laminar case in Fig. 3b shows a similarly small reduction in noise but little evident improvement in the spectrum. As N is increased to ± 2 and ± 3 , the HDR spectra in each plot lose high-frequency content and diverge from the hot-wire spectrum. Figure 3b suggests a noise floor that is filtered by the SoC algorithm to produce a spectral shape consistent with a low-pass filter rather than a continuation of the instability slope. The MS data are similar as the value of N is increased; the noise does diminish but the spectral shape simply suggests enhanced filtering. Although the spectra slopes are much smoother and monotonic than those returned by the top-hat filter in Fig. 2, it does not follow that they must be more physically realistic. The application of the Gaussian weighting eliminates the instability of the top-hat filter but instead it acts as a strong low-pass filter on the data. This removes energy in the velocity fluctuations that should be present in the high-frequency range of the spectra. In particular, the high-frequency roll-off associated with turbulent dissipation (Fig. 3a) or the growing instability (Fig. 3b) is not better captured by the multi-frame interrogation.

Finally, the pyramid correlation was tested using five different configurations, varying the maximum number of images N_I and the maximum number of pyramid levels N_L . Three are shown in Fig. 4 for each case in which $N_I = N_L$; reducing N_I without reducing N_L did not meaningfully alter the results. The spectra are remarkably alike to those produced using the Gaussian-weighted SoC and similarly show signs of low-pass filtering. For both flow cases and both PIV measurement systems, an increase in the number of images used corresponds to a reduction in the noise floor but also a characteristic roll off in the high-frequency region of the spectra. One exception is found, however, as the number of images and levels increases to $N_I = N_L = 3$ for the turbulent case in Fig. 4a, the filtering effect tapers off and closely matches the $N_I = N_L = 2$ spectrum. This probably indicates that the optimization of the pyramid correlation reduces the number of images and levels that actually are used in the correlation despite the maximum permissible. Both these

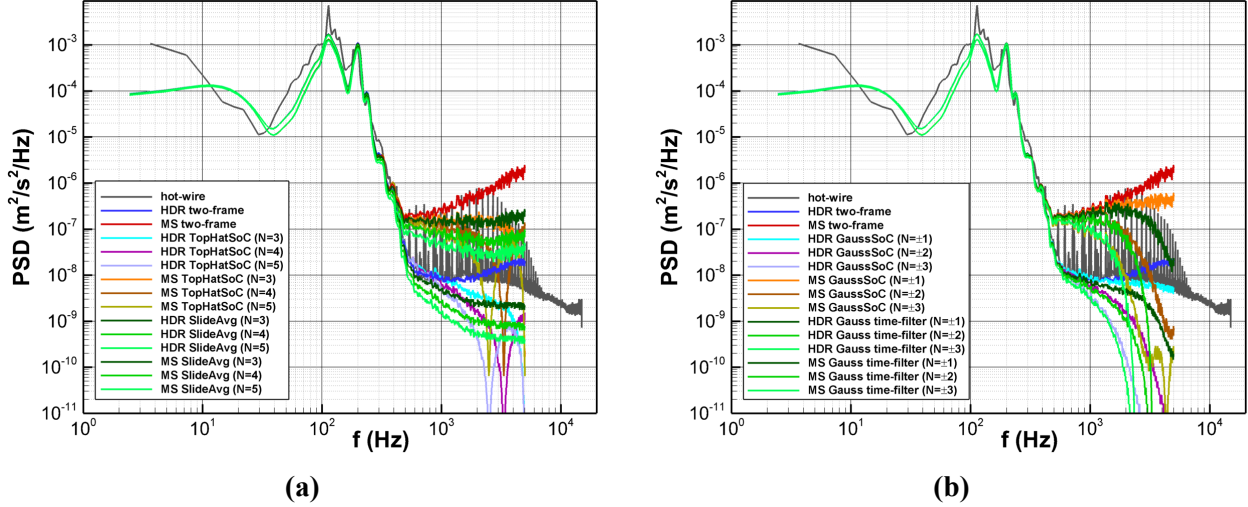


Fig. 5: Comparison of sliding SoC spectra to time-filtering of the two-frame time series in the unstable laminar region. (a) top hat SoC and sliding time average; (b) Gaussian weighted SoC and sliding time average.

spectra agree more closely with the hot-wire spectrum without displaying the characteristic sign of low-pass filtering increasing with number of images.

A closer examination of the spectra in Fig. 4 suggests that the implementation of the pyramid correlation may reveal a little more of the valid physics before the low-pass filtering effect manifests. This is easiest to recognize for the laminar case in Fig. 4b. As the noise levels drop with increasing N_i and N_L , the falling slope of the instability appears to extend to modestly higher frequencies. The characteristic shape of low-pass filtering is applied to the high-frequency noise rather than the fluid dynamics. This is clearest for the HDR data in which the inherent noise is lesser and is indicated with the arrow in the frequency range of 450 – 550 Hz. This contrasts with the Gaussian-weighted SoC in Fig. 3b, in which the different cases branch off the spectrum at the same point, revealing no additional range of the instability slope (also indicated by an arrow). A similar extension of the high-frequency slope also may be visible for the turbulent case in Fig. 4a as N_i increases, but the expected slope is more difficult to judge than in the laminar case. Nonetheless, there is reason to believe that as the pyramid correlation reduces the noise floor, some additional fluid dynamics are revealed without compromising the frequency response.

The top hat and Gaussian filters also can be applied in temporal space rather than correlation space. A sliding average with a kernel of $N=3, 4,$ or 5 images (which is $2, 3,$ or 4 vector fields) to match that of the top hat SoC was applied to the two-frame time-series vectors and then transformed into frequency spectra. These results are shown in Fig. 5a for the unstable laminar flow, overlaid on the spectra from Fig. 2b. No ringing is produced in this case but neither is any additional information added to the high-frequency slope of the instability peak; the correlation noise simply is filtered differently. Therefore, while a sliding average reduces correlation noise, it does not reveal any additional frequency content. More informative is the Gaussian time filter in Fig. 5b. Here, another sliding average is performed on the two-frame time series but now it is weighted by a Gaussian in the same manner as the weighting in correlation space from Fig. 3b. These spectra are overlaid on the Gaussian SoC spectra. Some high-frequency ringing is observed in the time-filtered spectra but is removed for clarity in the plot. The overlay of the two analyses shows that the Gaussian time filter produces very nearly the same spectra as the Gaussian SoC. Since the time-filtering is simply a form of a low-pass filter and would not extract additional accuracy from a correlation, Fig. 5b supports the observation from Fig. 3 that the Gaussian SoC multi-frame correlation merely acts as a low-pass filter and not a means of improving spectral accuracy.

The effect of the multi-frame correlation algorithms on the time-series data is shown in Fig. 6. For this, only the unstable laminar case is examined as it is easier to distinguish flow fluctuations from noise

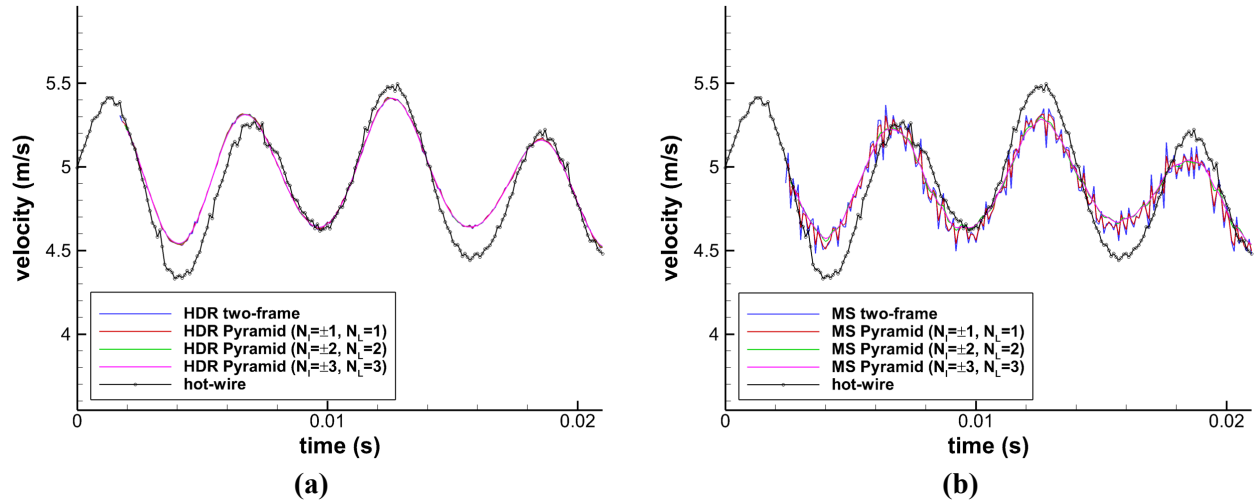


Fig. 6: Time series of data calculated using the pyramid correlation for the unstable laminar region, compared to simultaneous hot-wire data. (a) HDR measurement system; (b) MS measurement system.

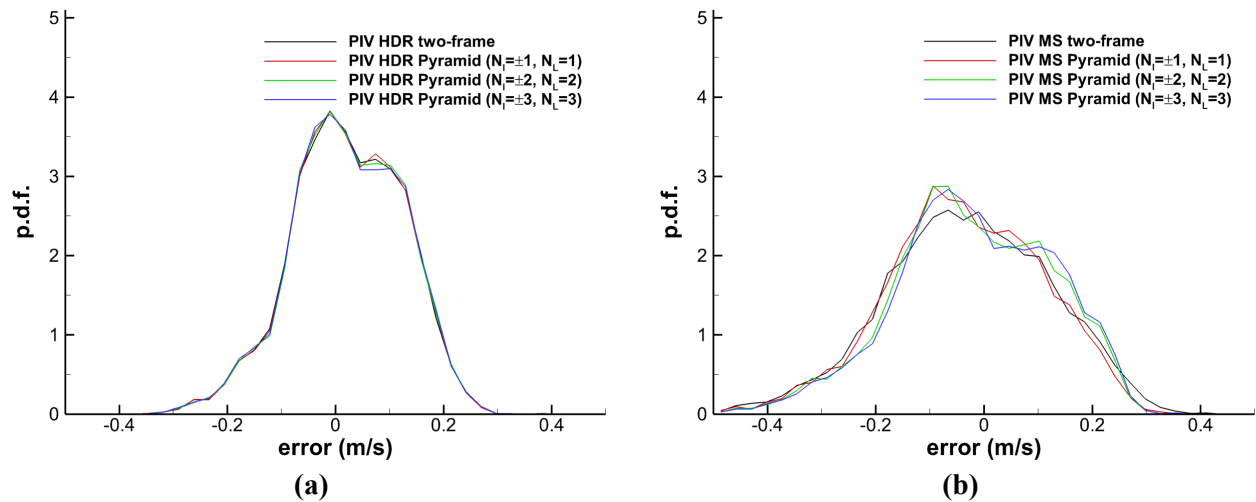


Fig. 7: Probability density functions of the error of the PIV data compared to the hot-wire data for the unstable laminar region. (a) HDR measurement system; (b) MS measurement system.

fluctuations for non-turbulent conditions. Only the pyramid correlation is shown as the Gaussian-weighted SoC is visibly indistinguishable from it in the time domain. (Even the high-frequency undulations of the top-hat filter are too low in magnitude to be visible in the time series.) Fig. 6a shows the performance of the HDR data and Fig. 6b the MS data; both can be compared to the hot-wire traces. A large-scale deviation of both PIV configurations from the hot-wire data is seen. It is not clear why the hot-wire did not return low-frequency characteristics in accordance with both PIV systems but the mid- and high-frequencies of present interest match well. This also can be seen in the spectra for the unstable laminar case such as Fig. 4b, in which the hot-wire spectrum shows a much larger magnitude than the PIV spectra for frequencies below 10 Hz. Possible explanations include the small distance between the hot-wire probe position and the PIV measurement location (though convection time has been accounted) and the differing spatial resolutions.

The PIV time traces in Fig. 6 show reduced fluctuations when a larger number of images are used in the correlation. But it is difficult to evaluate whether the increased smoothness of time-traces in Fig. 6 is a result of reduced noise for increased accuracy or simply a result of low-pass filtering of the signal content.

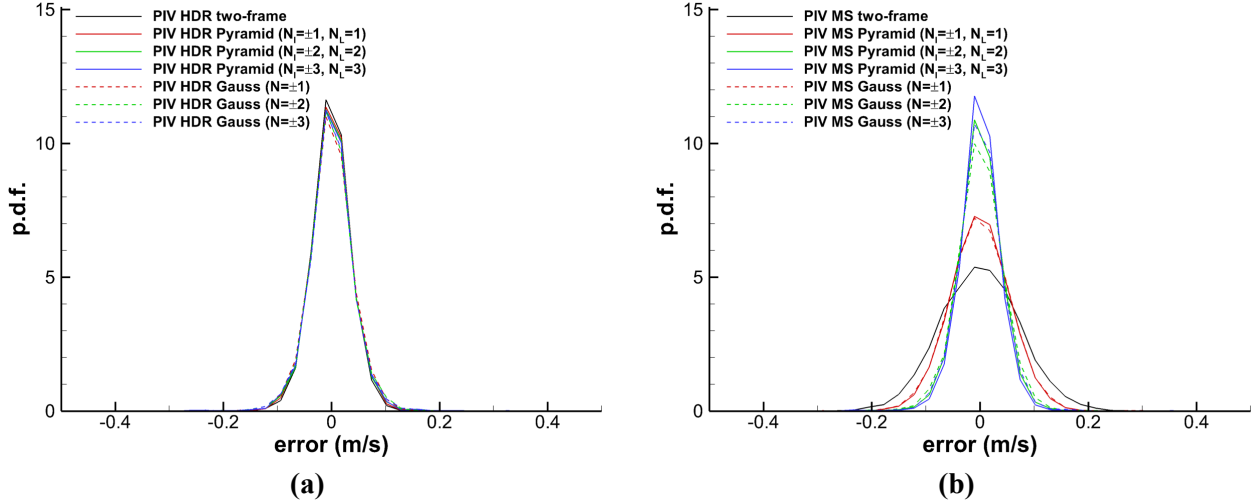


Fig. 8: Probability density functions of the error of the PIV data compared to the hot-wire data for the unstable laminar region. Data have been high-pass filtered at 300 Hz to focus on the high-frequency region of the spectrum. Gaussian-weighted SoC and pyramid correlation analyses are compared. (a) HDR measurement system; (b) MS measurement system.

Another means to evaluate the multi-frame correlation accuracy is to histogram the error of the PIV measurement via comparison with the simultaneous hot-wire data and look for reduced errors when using multi-frame algorithms. This is accomplished in Fig. 7, in which the instantaneous differences between the PIV data and the hot-wire data are formed into probability density functions. These provide the distribution of error when velocities are calculated by the various interrogation algorithms and then are compared to the hot-wire data as the true value (as in Sciacchitano *et al* 2015 [23]).

The error distributions in Fig. 7a do not indicate any improvement in the accuracy for the HDR data set by using the pyramid correlation or increasing the number of images within it. The analogous distributions for the MS data set in Fig. 7b show more variation between the different analysis algorithms but no evidence of a narrower distribution indicative of reduced error when the pyramid correlation is employed. Perhaps this should not be surprising since the accuracy improvements sought in the spectra occur at relatively high frequencies at which the energy in the flow is substantially reduced. Therefore, if such improvements occur, they would be inconsequential to the p.d.f.'s in Fig. 7 that are found from the gross flow field behavior dominated by large-scale low-frequency motion. Also note that the separation of each error distribution from a zero peak is due to the low-frequency discrepancy between the PIV and the hot-wire data. This is not a concern to the present analysis since the focus is on higher frequency ranges.

A better approach to assessing whether the multi-frame algorithms can reveal higher-frequency spectral content must focus on only the high-frequency fluctuations. This was accomplished by high-pass filtering the data before error analysis at a cutoff frequency of 300 Hz, which can be seen in the spectra of the unstable laminar flow to be a reasonable separation between the large-scale motion that all analyses reproduce effectively and the higher frequency content that is subject to analysis artifacts. Figure 8 shows the p.d.f.'s for both the HDR and MS systems and superposes results using both the Gaussian-weighted SoC and the pyramid correlation. As can be seen in Fig. 8a, the noise levels of the HDR system are still too low to differentiate between analysis cases; differences with the hot-wire data remain dominant. Differences in the spatial location of each measurement may allow too much evolution of the flow in this frequency range to draw an effective comparison. The MS system is more effective in revealing the decrease of error levels as more images are used in multi-frame algorithms. However, the difference between the Gaussian-weighted SoC and the pyramid correlation are secondary to the number of images used. Nothing in this plot can be used to determine that one method is superior to the other or to differentiate whether reduced discrepancies with the hot-wire data are due to filtering of noise or actual improved physical accuracy.

Discussion and Future Work

The results provided here demonstrate that the variety of available multi-frame correlation algorithms do not return equivalent results. When the number of images used is large, it should be expected that a low-pass filtering of the velocity spectrum is created. But the differences go beyond this straightforward fact. In the presence of acceleration, an increase of the number of images used for interrogation can also increase the variability of the particle displacements. Hence, when correlations are summed from varying displacements, the peaks become subject to broadening or splitting, thereby increasing correlation noise and/or bias.

But more dominant is the impact of numerical factors in the multi-frame correlations. The sliding sum-of-correlation using a top-hat weighting scheme (i.e., a simple sum of N correlations) induces a numerical instability that leads to strong high-frequency ringing. This is eliminated when instead using a Gaussian weighting on correlation summation, but the results in frequency space are essentially indistinguishable from low-pass filtering of a two-frame time series. The sum of correlation reduces noise and increases overall precision, but this does not accurately reveal any spectral content at frequencies exceeding that of the two-frame correlation. Stated differently, the noise is filtered but not independently from the signal.

The pyramid correlation shows somewhat more promise for revealing high-frequency physics. While low-pass filtering of the noise spectrum is evident, some signal separation occurs as well. This reveals a modest improvement of the signal bandwidth in the mid-to-high-frequency range.

The present results suggest that the multi-frame algorithms return results that are a mixture of low-pass filtering due to the increased number of images and the detection of spectral characteristics from the fluid dynamics. It is non-trivial to distinguish the filtering effects that arise from signal processing considerations from the legitimate physical content sought by these algorithms. Efforts beyond the scope of the present paper are required before this fundamental question can be illuminated. Comparisons of simultaneously acquired TR-PIV and reference hot-wire data would, in principle, be helpful in this regard. The coherence function should reveal over which frequencies valid flow information may be separated from measurement noise by multi-frame interrogation. These have not been reported herein because challenges matching the two signals in the relevant mid-to-high-frequency range interfered with finding any significant distinguishing characteristics between multi-frame algorithms. Unfortunately, the present work also has suggested limits of comparison between PIV and hot-wire data owing to differences in spatial resolution and measurement location. Valid comparisons may only be possible on a statistical basis rather than the instantaneous comparison required to provide a more definitive understanding of multi-frame correlation behavior.

This paper represents the first known study of the accuracy of velocity frequency spectra when interrogated by multi-frame correlation methods. Results establish that the high-frequency outcome of the velocity spectrum is very sensitive to choice of the interrogation algorithm and may not return an accurate response. A low-pass filtering effect due to use of more images may explain the apparent reduction in error rather than an increased precision of the measurement. The apparent benefits of multi-frame interrogation algorithms may be limited in their ability to reveal additional spectral content of the flow. Additional study on this matter is intended in coming years.

Acknowledgements

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

References

- [1] Adrian, R. J., "Twenty Years of Particle Image Velocimetry," *Experiments in Fluids*, Vol. 39, No. 2, 2005, pp. 159-169.
- [2] Westerweel, J., Elsinga, G. E., and Adrian, R. J., "Particle Image Velocimetry for Complex and Turbulent Flows," *Annual Review of Fluid Mechanics*, Vol. 45, 2013, pp. 409-436.
- [3] Ding, L., and Adrian, R. J., "N-Pulse Particle Image Velocimetry-Accelerometry for Unsteady Flow-Structure Interaction," *Measurement Science and Technology*, Vol. 28, No. 1, 2017, pp. 014001.
- [4] Hain, R., and Kähler, C. J., "Fundamentals of Multiframe Particle Image Velocimetry (PIV)," *Experiments in Fluids*, Vol. 42, No. 4, 2007, pp. 575-587.
- [5] Meinhart, C. D., Wereley, S. T., and Santiago, J. G., "A PIV Algorithm for Estimating Time-Averaged Velocity Fields," *Journal of Fluids Engineering*, Vol. 122, No. 2, 2000, pp. 285-289.
- [6] Scarano, F., Bryon, K., and Violato, D., "Time-Resolved Analysis of Circular and Chevron Jets Transition by Tomo-PIV," *15th International Symposium on Applications of Laser Techniques to Fluid Mechanics*, Lisbon, Portugal, 2010.
- [7] Astarita, T., "Analysis of Weighting Windows for Image Deformation Methods in PIV," *Experiments in Fluids*, Vol. 43, No. 6, 2007, pp. 859-872.
- [8] Sciacchitano, A., Scarano, F., and Wieneke, B., "Multi-Frame Pyramid Correlation for Time-Resolved PIV," *Experiments in Fluids*, Vol. 53, No. 4, 2012, pp. 1087-1105.
- [9] Lynch, K., and Scarano, F., "A High-Order Time-Accurate Interrogation Method for Time-Resolved PIV," *Measurement Science and Technology*, Vol. 24, No. 3, 2013, pp. 035305.
- [10] Jeon, Y. J., Chatellier, L., and David, L., "Fluid Trajectory Evaluation Based on an Ensemble-Averaged Cross-Correlation in Time-Resolved PIV," *Experiments in Fluids*, Vol. 55, No. 7, 2014, pp. 1766.
- [11] Yegavian, R., Leclaire, B., Champagnat, F., Illoul, C., and Losfeld, G., "Lucas-Kanade Fluid Trajectories for Time-Resolved PIV," *Measurement Science and Technology*, Vol. 27, No. 8, 2016, pp. 084004.
- [12] Discetti, S., and Coletti, F., "Volumetric Velocimetry for Fluid Flows," *Measurement Science and Technology*, Vol. 29, No. 4, 2018, pp. 042001.
- [13] Schanz, D., Gesemann, S., and Schroeder, A., "Shake-the-Box: Lagrangian Particle Tracking at High Particle Image Densities," *Experiments in Fluids*, Vol. 57, No. 5, 2016, pp. 70.
- [14] Vetel, J., Garon, A., and Pelletier, D., "Denoising Methods for Time-Resolved PIV Measurements," *Experiments in Fluids*, Vol. 51, No. 4, 2011, pp. 893-916.
- [15] Oxlade, A. R., Valente, P. C., Ganapathisubramani, B., and Morrison, J. F., "Denoising of Time-Resolved PIV for Accurate Measurement of Turbulence Spectra and Reduced Error in Derivatives," *Experiments in Fluids*, Vol. 53, No. 5, 2012, pp. 1561-1575.
- [16] Galmiche, B., Mazellier, N., Halter, F., and Foucher, F., "Turbulence Characterization of a High-Pressure High-Temperature Fan-Stirred Combustion Vessel using LDV, PIV, and TR-PIV Measurements," *Experiments in Fluids*, Vol. 55, No. 1, 2014, pp. 1636.
- [17] Miller, J. D., Jiang, N., Slipchenko, M. N., Mance, J. G., Meyer, T. R., Roy, S., and Gord, J. R., "Spatiotemporal Analysis of Turbulent Jets Enabled by 100-kHz, 100-ms Burst-Mode Particle Image Velocimetry," *Experiments in Fluids*, Vol. 57, No. 12, 2016, pp. 192.
- [18] Beresh, S. J., Henfling, J. F., and Spillers, R. W., "Postage-Stamp PIV: Small Velocity Fields at 400 kHz for Turbulence Spectra Measurements," *Measurement Science and Technology*, Vol. 29, No. 3, 2018, pp. 034011.
- [19] Willert, C. E., "'Profile PIV' – More than Just an Optical Hotwire: New Potentials of PIV in Boundary Layer Research," *18th International Symposium on Flow Visualization*, Zurich, Switzerland, June 26-29, 2018.
- [20] Manovski, P., Gulotta, P., Giacobello, M., de Silva, C., Hutchins, N., and Marusic, I., "Velocity and Pressure Measurements of a Turbulent Boundary Layer at 80 kHz using 2D PIV," *3rd Workshop on Data Assimilation and CFD Processing for PIV and LPT (virtual meeting)*, 2020.
- [21] Beresh, S. J., Spillers, R. W., Soehnel, M. M., and Spitzer, S. M., "Extending the Frequency Limits of 'Postage-Stamp PIV' to MHz Rates," AIAA Paper 2020-1018.
- [22] Neal, D. R., Sciacchitano, A., Smith, B. L., and Scarano, F., "Collaborative Framework for PIV Uncertainty Quantification: The Experimental Database," *Measurement Science and Technology*, Vol. 26, No. 7, 2015, pp. 074003.
- [23] Sciacchitano, A., Neal, D. R., Smith, B. L., Warner, S. O., Vlachos, P. P., Wieneke, B., and Scarano, F., "Collaborative Framework for PIV Uncertainty Quantification: Comparative Assessment Methods," *Measurement Science and Technology*, Vol. 26, No. 7, 2015, pp. 074004.
- [24] Boillot, A., and Prasad, A. K., "Optimization Procedure for Pulse Separation in Cross-Correlation PIV," *Experiments in Fluids*, Vol. 21, No. 2, 1996, pp. 87-93.
- [25] Smith, S. W., *The Scientist and Engineer's Guide to Digital Signal Processing*, 1997, <http://www.dspguide.com/pdfbook.htm>.
- [26] Astarita, T., "Analysis of Weighting Windows for Image Deformation Methods in PIV," *Experiments in Fluids*, Vol. 43, No. 6, 2007, pp. 859-872.