CrossMark

RESEARCH ARTICLE

PIV anisotropic denoising using uncertainty quantification

B. Wieneke[1] ⓘ

Abstract Recently, progress has been made to reliably compute uncertainty estimates for each velocity vector in planar flow fields measured with 2D-or stereo-PIV. This information can be used for a post-processing denoising scheme to reduce errors by a spatial averaging scheme preserving true flow fluctuations. Starting with a 5×5 vector kernel, a second-order 2D-polynomial function is fitted to the flow field. Vectors just outside will be included in the filter kernel if they lie within the uncertainty band around the fitted function. Repeating this procedure, vectors are added in all directions until the true flow field can no longer be approximated by the second-order polynomial function. The center vector is then replaced by the value of the fitted function. The final shape and size of the filter kernel automatically adjusts to local flow gradients in an optimal way preserving true velocity fluctuations above the noise level. This anisotropic denoising scheme is validated first on synthetic vector fields varying spatial wavelengths of the flow field and noise levels relative to the fluctuation amplitude. For wavelengths larger than 5–7 times the spatial resolution, a noise reduction factor of 2–4 is achieved significantly increasing the velocity dynamic range. For large noise levels above 50% of the flow fluctuation, the denoising scheme can no longer distinguish between true flow fluctuations and noise. Finally, it is shown that the procedure performs well for typical experimental PIV vector fields. It provides an effective alternative to more complicated adaptive PIV algorithms optimizing interrogation window sizes and shapes based on seeding density, local flow gradients, and other criteria.

✉ B. Wieneke
  bwieneke@lavision.de

1 LaVision GmbH, Göttingen, Germany

1 Introduction

The fundamental challenge to any PIV processing scheme is to select the optimal spatial resolution—mainly determined by interrogation window size and overlap factor—for a given image quality and information density or signal-to-noise ratio. In many cases, this is not uniform across the image or varying from image to image. Typically, one tries to find some compromise in interrogation window size and other processing parameters which work reasonably well everywhere. Instead, it would be advantageous to locally adapt the spatial resolution.

For this purpose, several adaptive PIV techniques have been developed taking into consideration local seeding densities, flow gradients, or physical constraints like walls, locally adjusting the interrogation window position, size, and shape (Scarano 2004; Theunissen et al. 2007, 2008, 2010; Wieneke and Pfeiffer 2010). These techniques have shown to reduce the systematic and random noise level significantly, in particular close to object surfaces.

For time-resolved PIV (TR-PIV), one can (additionally) use the temporal information for noise reduction applying various multi-frame correlation techniques (Hain and Kähler 2007; Sciacchitano et al. 2012; Jeon et al. 2014; see also recent PIV Challenge: Kähler et al. 2016).

Once a velocity field has been obtained, various post-processing noise reduction schemes can be applied, among them standard spatial top-hat or Gaussian smoothing filters, polynomial regression fits, or POD analysis (Raiola et al. 2014). For TR-PIV, more options are available, since the

noise is usually less correlated in time than in space. For example, Oxlade et al. (2012) apply a spectral white noise subtraction recovering the true velocity spectra. Vétel et al. (2011) use an optimal temporal Wiener filter with better results than convolution- or wavelet-based filters.

Other denoising approaches utilize physical constraints. In particular, for volumetric data and divergence-free incompressible flows, one can use the criterion of non-zero divergence for uncertainty quantification (Sciacchitano and Lynch 2015) as well as reducing the noise by making the flow field divergence-free (de Silva et al. 2013; Schiavazzi et al. 2014; Azijli and Dwight 2015; Wang et al. 2016). Atkinson et al. (2014) estimate the noise of volumetric velocity data directly from the velocity fields and apply an appropriate spatial filter to achieve more accurate turbulent power spectra. For time-resolved volumetric data, one can apply Navier–Stokes or some simplified version as a regularizer. For example, Schneiders et al. (2015) apply a time-marching vortex-in-cell procedure to average multiple flow fields with more accurate results than retrieved with spatial/temporal moving average and polynomial regression filters.

The main issue for any denoising scheme is to select the appropriate spatial or temporal filter kernel size, which should be as large as possible to average out the random noise and short enough not to reduce the amplitude of true flow fluctuations (truncation errors). With recent work on PIV uncertainty quantification (Charonko and Vlachos 2013; Timmins et al. 2012; Sciacchitano et al. 2013; Wieneke 2015; Sciacchitano and Lynch 2015), it is now possible to estimate the uncertainty of each computed vector. This can help to guide any denoising scheme for optimal noise reduction while preserving the true flow fluctuations.

The anisotropic denoising filter described here is based on local polynomial approximation (LPA) with locally adaptive kernel size well known in the digital signal processing community (e.g., Katkovnik 2005; Browne et al. 2007) for 1D- and 2D-data (images). While adaptive LPA estimates the local noise level from the residual between the fitted and the original data, direct knowledge of the PIV uncertainties is employed here to guide the size and shape of the final 2D filter kernel size. This anisotropic denoising scheme is finally validated using synthetic and experimental vector field data.

2 Method

The denoising scheme described here is restricted to planar velocity fields with u, v, and possibly w components together with uncertainty values U_u, U_v, U_w on a 1-sigma level, i.e., the true velocity value u_{true} is expected to lie within $u \pm U_u$ with a probability of 68%. Denoising is done independently for each vector of the flow field.

At the beginning, for each vector component (u, v, and w), a second-order 2D-polynomial function is fitted to a 5×5 vector neighborhood around the center vector. The uncertainty of the vectors in the 5×5 neighborhood is averaged and taken as a reference in the following. Vectors just outside the center 5×5 region (white squares in Fig. 1) are tested if they should be added to the filter kernel. Vector a and, at the same time, vector d on the opposite will be added if both adjacent inner vectors b and c are part of the filter kernel, and if all components (u, v, w) of vectors a and d are within an uncertainty band around the fitted polynomial function, as shown in Fig. 2.

If any of the conditions fails, then both vectors, a and d, will be discarded, and subsequently, all vectors further out in this direction are prevented to become part of the smoothing kernel. With such a symmetrical vector addition, the center of mass of the filter kernel remains at the center vector location for all kernel shapes. The resulting 2D-array of valid vectors is then made more compact and regular by deleting vectors if they have ≤ 2 neighbors and adding vectors if they have ≥ 5 neighbors, even if they would be rejected otherwise.

The uncertainty band (Fig. 2) is given by $\pm S$ times the uncertainty (provided on a 1-sigma level), where S is a user-selected filter strength, as shown later typically set to around 2.5–3.5. A narrow band of ± 1-sigma would be too tight, since with a probability of 32%, a vector falls outside this range preventing the growth of the filter kernel. The procedure stops when no more vectors are added or when a user-selected maximum kernel size is reached. At the end, the often quite irregular shape of the filter kernel is converted to a closest ellipse (Fig. 3). The difference in performance with and without ellipse fitting is only minor.

Then, LPA is executed on the vector field inside the filter kernel, and the center vector is replaced by the value of the polynomial function at the center location. Since the spatial derivatives of the flow field are readily available from the fitted polynomial function, they are stored, e.g., for subsequent vorticity or divergence computation.

Finally, the procedure computes a new uncertainty for each vector component using the uncertainty propagation rules outlined in Sciacchitano and Wieneke (2016). A simplified version is used here by taking the reference uncertainty divided by $\text{sqrt}(N_{\text{eff}} - 6)$, where N_{eff} is the number of independent vectors in the final filter kernel and 6 is the number of parameters (degrees of freedom) of the second-order 2D-polynomial function. Roughly, N_{eff} is the total number of vectors in the filter kernel divided by the number of vectors within the size of the interrogation window. For example, with an interrogation window size of 32×32 pixel and 75% overlap, there will be 16 vectors within the window. If one would smooth over these 16 vectors, there will be effectively very little reduction of

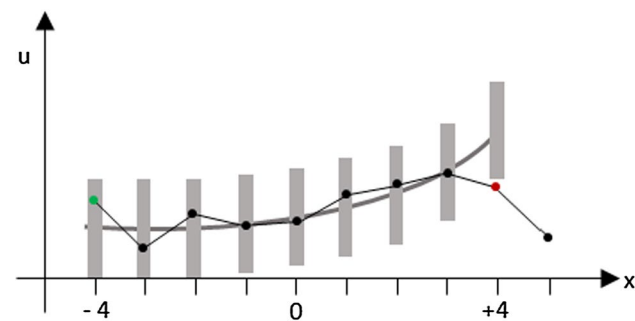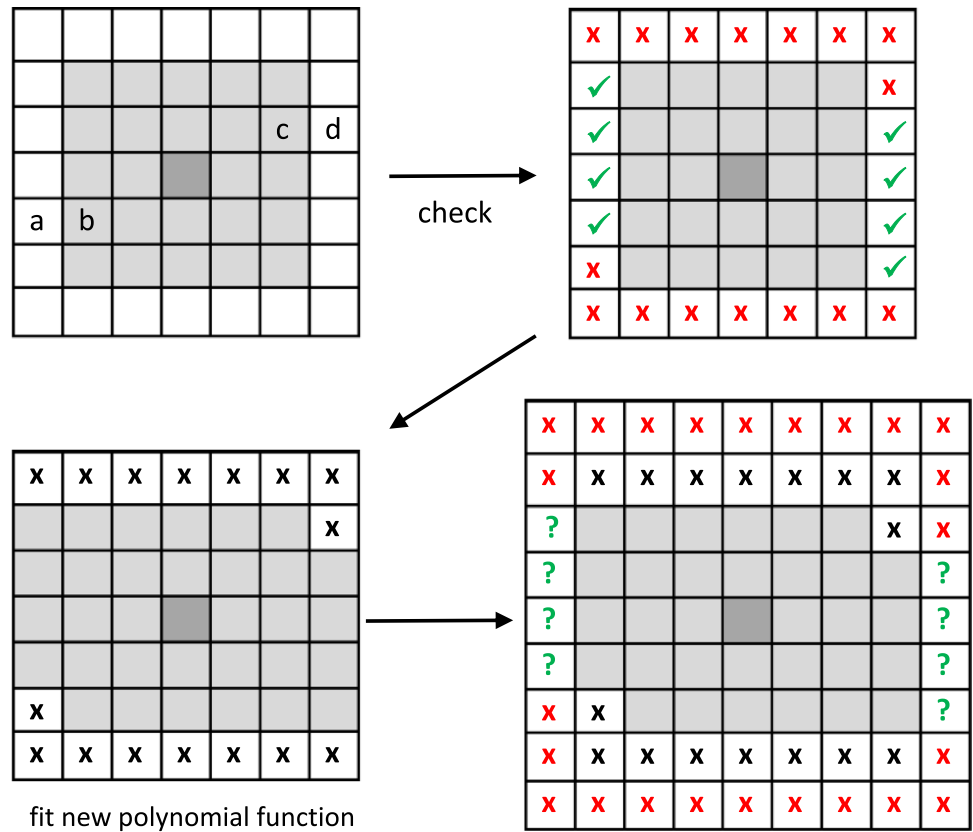**Fig. 1** Adding a new shell to the filter kernel

fit new polynomial function

Fig. 2 Polynomial function computed around the center vector at $x = 0$ with data between -3 and $+3$ and extrapolated to ± 4. *Red vector* at $+4$ is rejected, since it lies outside the *gray* uncertainty band
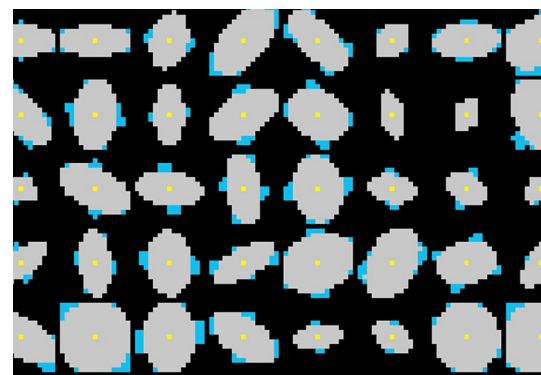

Fig. 3 Final filter kernel (*blue*) is converted to an ellipse (*gray*) shown for every 13th vector

the uncertainty and noise, since the errors of all vectors are closely correlated.

It is also necessary to update the spatial resolution of the vector field, which is related to the spatial autocorrelation coefficients between neighboring vectors. Due to variable filter size and shape, this is different for each vector in magnitude and direction, similar to the adaptive PIV techniques with varying interrogation window sizes and shapes. A fully correct treatment is complicated and would require the storage of many additional correlation values for each vector for subsequent

uncertainty propagation. Again, a simplified version is adopted here setting the spatial resolution to the average linear dimension of the filter kernel. It needs to be shown, if this is sufficient for accurate uncertainty quantification when the directional dependence of the effective spatial resolution becomes important, e.g., for the vorticity field.

The proposed denoising scheme takes typically a few seconds of processing time on a standard PC. It can easily be extended to volumetric data and to the time domain.

3 Synthetic data validation

The denoising scheme is first tested on a synthetic vector field with a wide range of spatial wavelengths L and signal-to-noise ratios (velocity dynamic range). The vector field contains 200×75 vectors with a grid spacing d of 4 pixel. The (true) flow field contains vortices of various sizes with spatial wavelengths of 512 pixel on the left and 32 pixel on the right of the image with a constant (true) amplitude of 1 pixel (Fig. 4 top). Vortices constitute a more challenging case than simple shear flows, where the filter kernel shape can be strongly elongated along the shear.

Any PIV algorithm has a finite spatial resolution equivalent to a characteristic filter length L_{sr} reducing small-scale fluctuations. Here, it is assumed that L_{sr} is 16 pixel, equivalent to 4 vectors, which is similar to using 16×16 pixel interrogation windows with 75% overlap. The filter length L_{sr} as the inverse of the spatial resolution is defined here as the sum of the auto-correlation coefficients between the errors of neighboring vectors (Sciacchitano and Wieneke 2016). If PIV would be a simple single-pass linear top-hat filter averaging the displacement information within an interrogation window of $L \times L$ pixel, then L_{sr} would be equal to L, as can be easily verified. The vector field is filtered here with a Gaussian filter function ($\propto \exp(-x^2/2\sigma^2)$) of equivalent filter length $L_{sr} = \sigma$ sqrt(4π). This leads to a significant reduction in amplitude for small wavelengths, e.g., about 50% for $L/L_{sr} = 2$ (Fig. 4 middle and bottom). The exact spatial frequency response depends on details of, e.g., the predictor–corrector scheme of a multi-pass PIV algorithm. This smoothing is included to put the noise level and its reduction by the anisotropic denoising scheme in perspective to the unavoidable amplitude reduction of small wavelengths due to the limited spatial resolution of the PIV algorithm itself.

Fig. 4 Synthetic vector field with constant fluctuation amplitude (*top*) and after taking the finite spatial resolution of a PIV algorithm into account showing u component (*middle*) and vorticity (*bottom*)
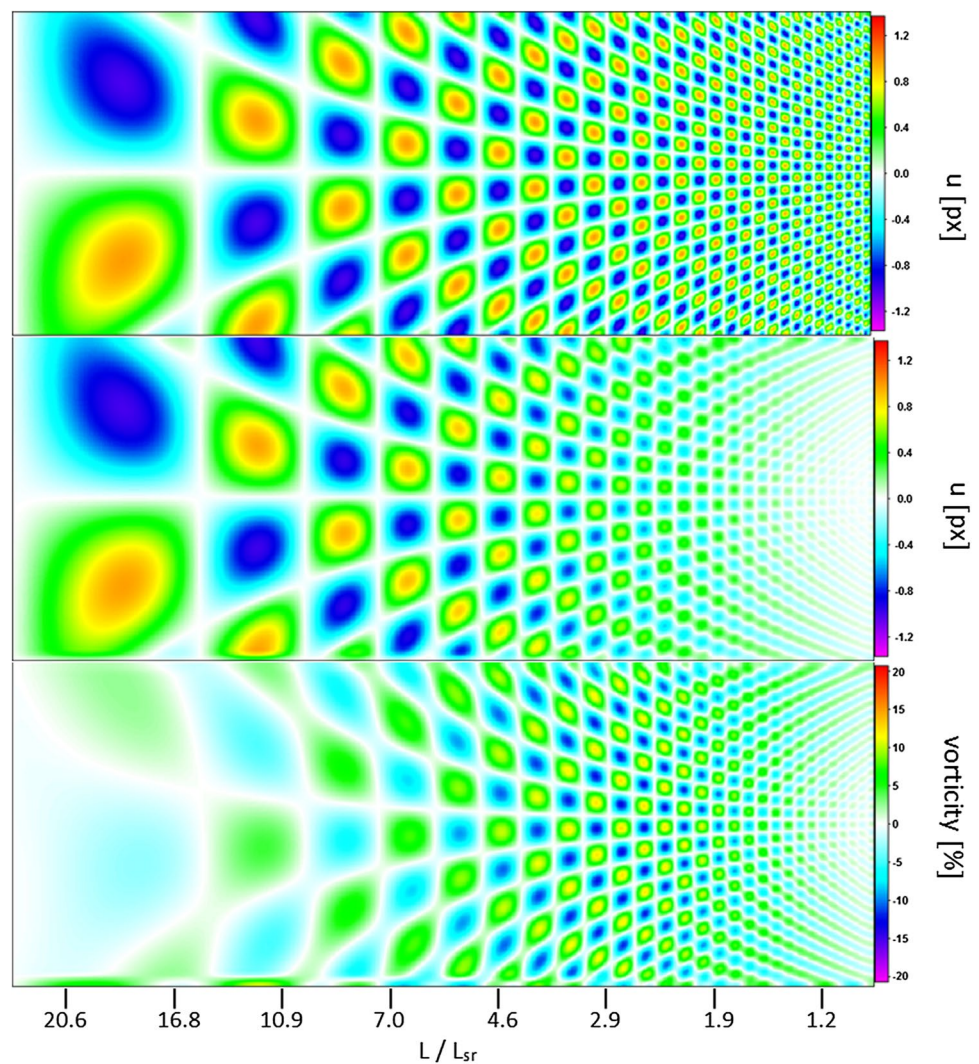
Noise levels of 0–100% are added to each vector component, again subject to the PIV spatial filtering, which leads to locally correlated noise components between neighboring vectors. This becomes important when applying locally confined averaging, where the noise is very little reduced, since it is locally correlated. Viewed another way, local averaging/denoising must be done over a kernel size larger than L_{sr} to become effective. The maximum filter kernel size is set to 41×41 vectors. Final computed kernel sizes are typically in the range of 5–15 vectors in each direction.

Figure 5 shows the u component with increasing noise level from top to bottom of the original vector field and after denoising with strength S of 1, 2, 2.5, 3, and 4 (from left to right). For zero noise level (top), the denoising scheme does not change the vector field apart of slightly decreasing the spatial resolution, i.e., increasing L_{sr} from 16 to 19, due to the initial 5×5 polynomial regression, which is always done. This is barely visible on the top-right, where the amplitude of small-scale fluctuations is decreased slightly.

For low noise levels, the procedure accurately recovers the round shape of vortices for almost all wavelengths. Large-scale vortices with larger possible kernel sizes are recovered even at 100% noise level. For smaller wavelengths further to the right, the denoising procedure is able to reduce the noise as long as the true fluctuations are larger than the errors. Beyond that, the algorithm cannot distinguish between true and noisy fluctuations anymore. Here, given a strong enough filter, the vector field is simply averaged over large regions. The algorithm assumes that everything is noise (see bottom-right of the noise plot in Fig. 6). The optimal filter strength seems to be between 2.5 and 3.0, strong enough to eliminate noise over potentially large regions for larger wavelength while not smoothing over true fluctuations.

The performance of the denoising scheme is quantified in Fig. 7 plotting the local rms of the noise for a filter strength of $S = 3$ as a function of wavelength for the different noise levels of 0–100% (0–1 px). For larger wavelength $L/L_{sr} > 10$, the noise is reduced by a factor of 2, up to a factor of 4 in some cases and larger wavelengths. For large noise levels >50%, only wavelengths $L/L_{sr} > 15$ are recovered, which is not surprising, since even visually it is difficult to detect smaller vortices in the noisy vector field. Smaller wavelengths are simply smoothed-over as the algorithm is unable to distinguish between true vortices and noise. Therefore, the overall noise level is decreased, but hidden flow structures are also removed.

For the noise-free vector field, the error increases for small wavelengths ($L/L_{sr} = 1$–3) by about 5% of the true amplitude due to the 5×5 polynomial regression, which, as mentioned before, leads to 15% lower spatial resolution. Still, one has to keep in mind that for these wavelengths, the amplitude reduction due to the spatial filtering effect of PIV is anyway above 50%.
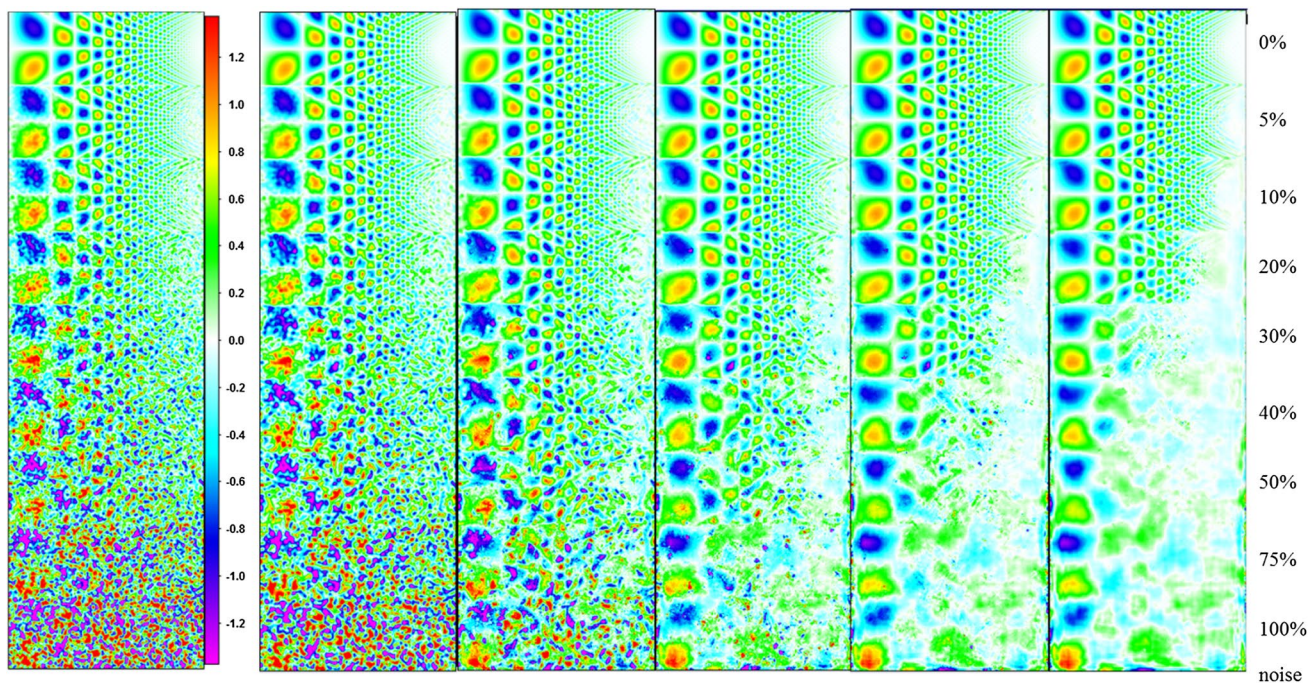

Fig. 5 Original vector field and after denoising with strength $S = 1$, 2, 2.5, 3, and 4 from *left* to *right*. Noise level 0–100% from *top to bottom*. *Color* = u component
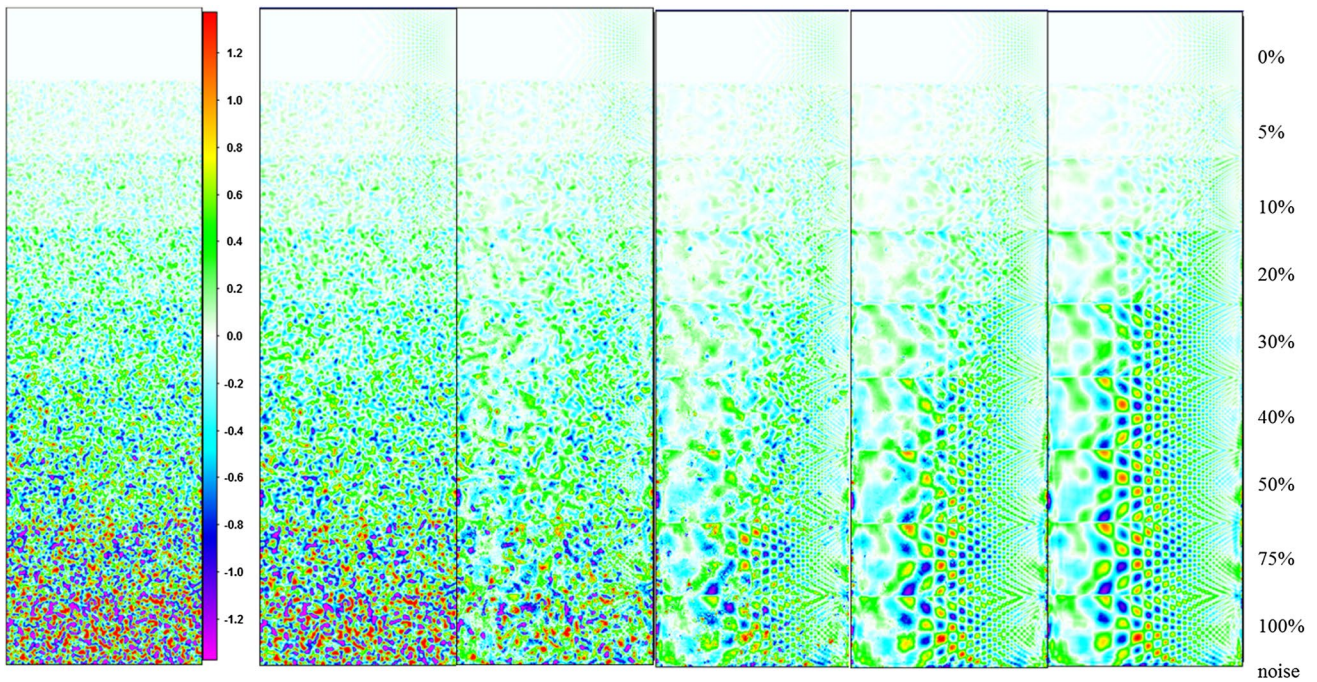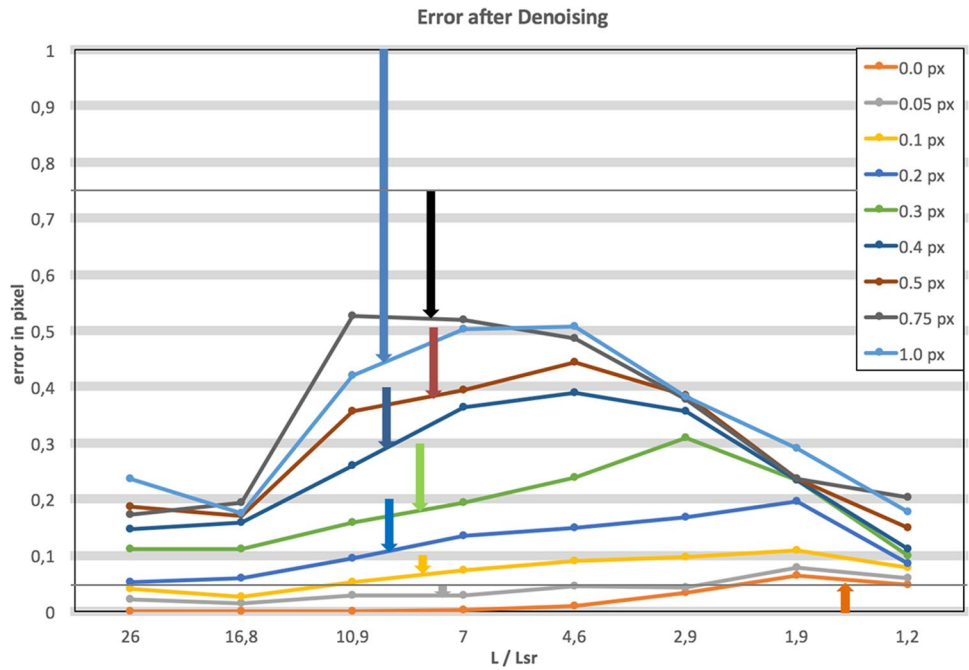
Fig. 6 Original noise and after denoising (equal to Fig. 5 minus Fig. 4 middle) with strength $S = 1$, 2, 2.5, 3, and 4 from *left to right*. Noise level 0–100% from *top to bottom*. *Color = u* component of noise

Fig. 7 Remaining noise level after denoising as a function of spatial wavelengths for different original noise levels of 0–1 px

The denoising scheme has been compared to a standard 2nd-order polynomial regression filter with a fixed kernel size of 5×5 to 11×11 vectors and a top-hat smoothing filter over 9×9 vectors for the case of 20% (0.2 px) noise level (Fig. 8). For large wavelengths, the polynomial regression filter reduces the noise level with successively larger filter kernels. The top-hat 9×9 filter performs better than polynomial fit of 11×11, since it is roughly equivalent to a polynomial filter of 20×20 vectors. For the intermediate range of L/L_{sr} between 2 and 7, the polynomial regression filter even increases the noise level, since the reduction of random noise is less than the

extra noise added due to increased truncation errors, i.e., smoothing the true flow fluctuations. Clearly, the anisotropic denoising filter outperforms all other schemes due to locally adapting the kernel size to the wavelength of the true flow fluctuations.

Fig. 8 Comparison of anisotropic denoising with the second-order polynomial regression filter with 5 × 5 to 11 × 11 vector kernel and 9 × 9 top-hat smoothing filter. Original noise level = 0.2 pixel

4 Experimental assessment

The denoising scheme is first tested on image 50 of PIV challenge 2003 case A (Stanislas et al. 2005). Standard PIV processing with 16 × 16 px window size reveals small-scale vortices inside the jet with a few remaining outliers even after vector validation and leaves a high noise level outside the jet (Fig. 9). It is not clear if some of the small-scale vortices are actually real or noise. Increasing the interrogation window size reduces the noise level, but, at the same time, smears out small-scale vortices leading to lower peak vorticity levels. The user has to choose some compromise, e.g., selecting 24 × 24 or 32 × 32 pixel windows.

Figure 10 shows the result of the denoising scheme with a filter strength S of 3.5 and a maximum kernel size of 15 × 15 vectors. Compared to 32 × 32 windows without denoising, if using 24 × 24 with denoising, one arrives at higher peak vorticity levels and smaller
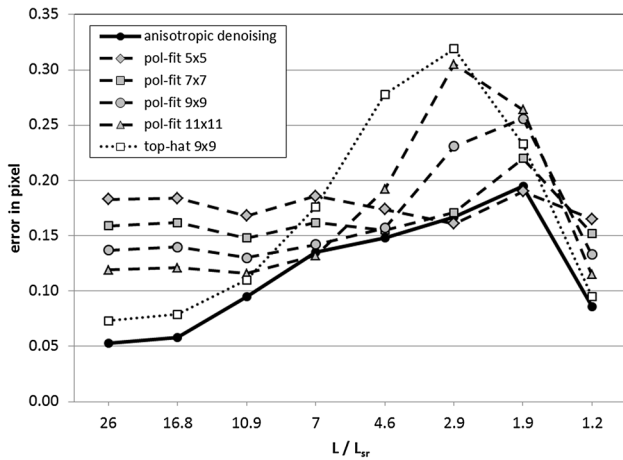
Fig. 9 PIV Challenge 2003, case A, image 50 processed with interrogation window sizes of 16 × 16, 24 × 24, 32 × 32, and 48 × 48 pixel (from *top-left* to *bottom-right*) with 75% overlap. *Color* = vorticity (a.u.)
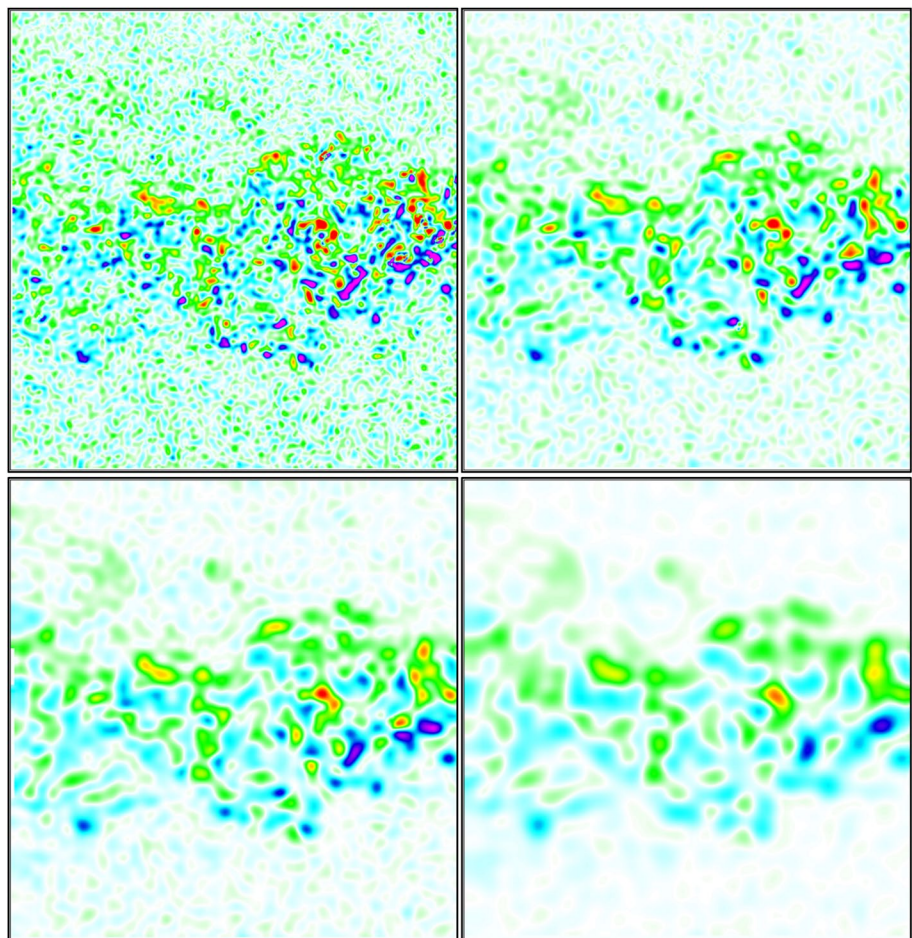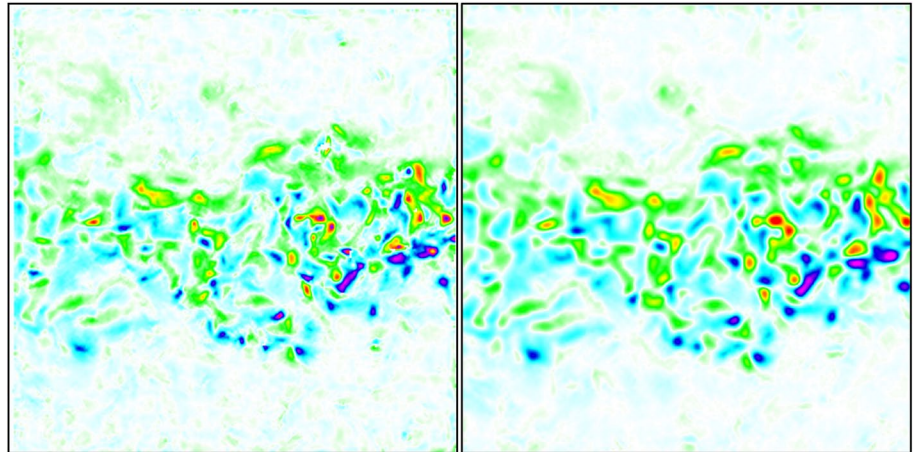
Fig. 10 Denoising applied on vector fields with IW = 16 × 16 (*left*) and 24 × 24 (*right*). S = 3.5

resolved flow features. At the same time, the noise levels in regions of low flow gradients outside the jet are strongly reduced, which would be otherwise only achievable, e.g., with 48 × 48 or larger windows.

Another experimental example is a cut through a ring vortex in air (Fig. 11) with 24 × 24 interrogation windows and 75% overlap. Since the jet is mostly laminar, the small-scale granular structure of the vorticity pattern is purely noise related. Denoising clearly reduces the noise level while preserving and highlighting the true flow structures.

Finally, time-resolved PIV data from a turbulent pipe flow (van Doorne and Westerweel 2007) are analyzed using the new spatial denoising technique and/or applying at each location a temporal polynomial regression filter of second order over 11 time steps. As shown in Fig. 12 and in the associated movie in the supplementary data, with an interrogation window size of 32 × 32 pixel, the vector fields at each time step (top-left) are clearly quite noisy in space and flickering also over time. This qualitative judgement is based on the knowledge that here, small-scale spatial variations lack physical justification.

Within the 2 ms between time steps and corresponding 3–5 pixel displacement of particles, the temporal flickering is mostly due to noise.

After applying the anisotropic denoising scheme, the flow field (top-right) is much smoother spatially, but still flickering significantly over time. When applying the temporal filter on the raw vector fields, the result is now temporally smooth over time at each spatial location, but still spatially noisy (bottom-left). Applying both schemes (here first spatial denoising, then temporal filter, bottom-right), the result is finally spatially and temporally smooth with only small reduction in spatial and temporal resolution compared to the raw vector fields. This indicates a significant reduction of the noise level, about equally by the temporal and spatial denoising scheme.

5 Summary

A rather simple denoising scheme is presented as a postprocessing step on vector fields using the uncertainty value for each vector as a reference. Vectors will be

Fig. 11 Ring vortex processed with 24 × 24 windows and 75% overlap without (*left*) and with denoising (*right*, S = 3.5). Seeding and PIV processing are only applied inside the vortex. *Color* = vorticity (a.u.)
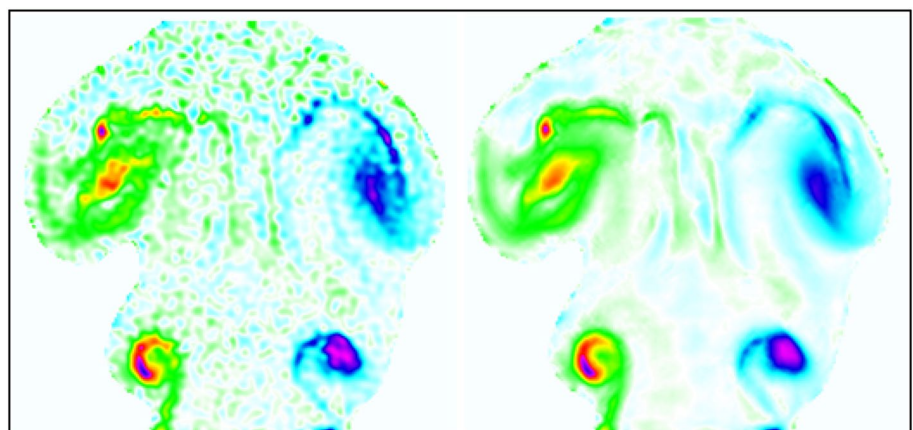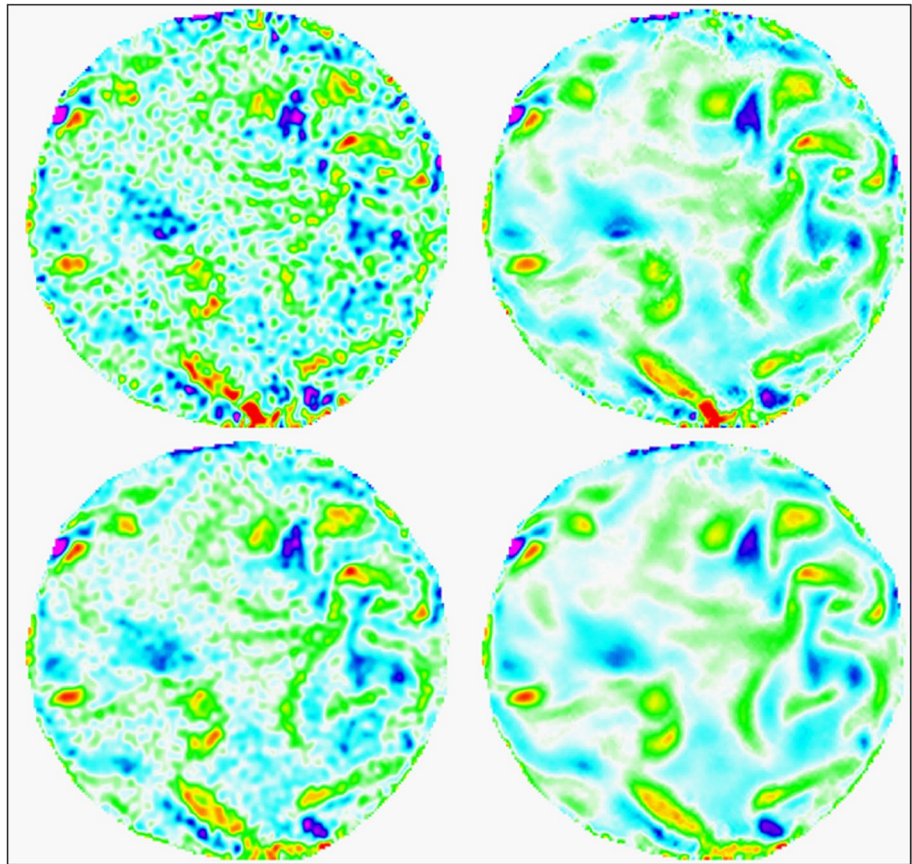
Fig. 12 Turbulent pipe flow (van Doorne and Westerweel 2007), raw vector field (*top-left*) after anisotropic denoising (*top-right*, $S = 3.5$), and for both with polynomial regression filter of the second-order over 11 time steps (*bottom*). *Color* = vorticity (a.u.). See movie in supplementary data

progressively added in all directions to the filter kernel if the velocity value is inside the uncertainty band around the locally fitted function. This is repeated until the true flow field can no longer be approximated by the second-order polynomial function. The final shape and size of the filter kernel automatically adjusts to local flow gradients in an optimal way preserving true velocity fluctuations above the noise level.

This anisotropic denoising scheme is validated on synthetic vector fields with varying spatial wavelengths and noise levels showing a significant increase of the velocity dynamic range for wavelengths about four times larger than the spatial resolution. For noise levels above 50%, the procedure is no longer able to distinguish between true flow fluctuations and noise except for large wavelengths of $L/L_{sr} > 10$.

The procedure has been shown to work well for typical experimental flow fields. It can be easily extended to volumetric data and to the time domain.

References

Atkinson C, Buchmann N, Amili O, Soria J (2014) On the appropriate filtering of PIV measurements of turbulent shear flows. Exp Fluids 55(1):1–15

Azijli I, Dwight RP (2015) Solenoidal filtering of volumetric velocity measurements using Gaussian process regression. Exp Fluids 56:198

Browne M, Mayer N, Cutmore TRH (2007) A multiscale polynomial filter for adaptive smoothing. Dig Sig Proc 17:69–75

Charonko JJ, Vlachos PP (2013) Estimation of uncertainty bounds for individual particle image velocimetry, measurements from cross-correlation peak ratio. Meas Sci Technol 24:1–16

de Silva CM, Philip J, Marusic I (2013) Minimization of divergence error in volumetric velocity measurements and implications for turbulence statistics. Exp Fluids 54:1557

Hain R, Kähler CJ (2007) Fundamentals of multiframe particle image velocimetry (PIV). Exp Fluids 42:575–587

Jeon YJ, Chatellier L, David L (2014) Fluid trajectory evaluation based on an ensemble-averaged cross-correlation in time-resolved PIV. Exp Fluids 55(7):1766

Katkovnik V (2005) Multiresolution local polynomial regression: a new approach to pointwise spatial adaptation. Dig Sig Proc 15:73–116

Oxlade AR, Valente PC, Ganapathisubramani B, Morrison JF (2012) Denoising of time-resolved PIV for accurate measurement of turbulence spectra and reduced error in derivatives. Exp Fluids 53(5):1561–1575

Raiola M, Discetti S, Ianiro A (2014) On the suppression of PIV measurement noise with a POD based filter. 17th Int Symp on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal, 07–10 July

Scarano F (2004) A super-resolution particle image velocimetry interrogation approach by means of velocity second derivatives correlation. Meas Sci Technol 15(2):475

Schiavazzi D, Coletti F, Iaccarino G, Eaton JK (2014) A matching pursuit approach to solenoidal filtering of three-dimensional velocity measurements. J of Comp Phys 263:206–221

Schneiders J, Dwight RP, Scarano F (2015) Tomographic PIV noise reduction by simulating repeated measurements. 11th Int Symp on PIV, Santa Barbara, CA, USA

Sciacchitano A, Lynch K (2015) A posteriori uncertainty quantification for tomographic PIV data. 11th Int Symp on PIV, Santa Barbara, CA, USA

Sciacchitano A, Wieneke B (2016) PIV uncertainty propagation. Meas Sci Technol 27(8):084006

Sciacchitano A, Scarano F, Wieneke B (2012) Multi-frame pyramid correlation for time-resolved PIV. Exp Fluids 53(4):1087–1105

Sciacchitano A, Wieneke B, Scarano F (2013) PIV uncertainty quantification by image matching. Meas Sci Technol 24:1–16

Stanislas M, Okamoto K, Kähler CJ, Westerweel J (2005) Main results of the second international PIV challenge. Exp Fluids 39(2):170–191

Theunissen R, Scarano F, Riethmuller ML (2007) An adaptive sampling and windowing interrogation method in PIV. Meas Sci Technol 18(1):275–287

Theunissen R, Scarano F, Riethmuller ML (2008) On improvement of PIV interrogation near stationary interfaces. Exp Fluids 45:557–572

Theunissen R, Scarano F, Riethmuller ML (2010) Spatially adaptive PIV interrogation based on data ensemble. Exp Fluids 48(5):875–887

Timmins BH, Wilson BW, Smith BL, Vlachos PP (2012) A method for automatic estimation of instantaneous local uncertainty in particle image velocimetry measurements. Exp Fluids 53:1133–1147

van Doorne CWH, Westerweel J (2007) Measurement of laminar, transitional and turbulent pipe flow using stereoscopic-PIV. Exp Fluids 42:259–279

Vétel J, Garon A, Pelletier D (2011) Denoising methods for time-resolved PIV measurements. Exp Fluids 51(4):893–916

Wang C, Gao Q, Wang H, Wei R, Li T, Wang J (2016) Divergence-free smoothing for volumetric data. Exp Fluids 57:15

Wieneke B (2015) PIV uncertainty quantification from correlation statistics. Meas Sci Technol 26(7):074002

Wieneke B, Pfeiffer K (2010) Adaptive PIV with variable interrogation window size and shape, 15th Int Symp on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal, 05–08 July