

Reynolds-stress measurement using stereoscopic ensemble PIV

M. T. Taylor^{1,*}, J. M. Lawson¹, B. Ganapathisubramani¹

1: Aerodynamics and Flight Mechanics Research Group, University of Southampton, Southampton SO17 1BJ, United Kingdom

*Corresponding author: M.T.Taylor@soton.ac.uk

Keywords: Stereoscopic PIV, ensemble correlation, single-pixel PIV, turbulence.

ABSTRACT

Single-pixel ensemble correlation decouples the spatial resolution of particle image velocimetry from the interrogation window size. This is a significant advantage for wall-bounded flows, where high spatial resolution is required to resolve steep gradients near boundaries. The operating principle is to measure mean flow and Reynolds stresses from measurements of the location and width of the time-averaged correlation peak. To date, this technique has been limited to 2D, two-component PIV measurements. The extension to stereoscopic PIV is not straightforward. This is because it requires a measurement of the velocity variance which can only be inferred from cross-correlations between camera pairs, and because the measurement of correlation-peak width is confounded by factors such as the laser sheet thickness, the camera point spread function and image interpolation error. We address this with two developments. Firstly, we introduce a Fourier-domain correlation-peak fitting procedure which accounts for the variation in the correlation peak width due to the imaging properties and image interpolation errors. Secondly, we introduce the “correlation-of-correlations” method to measure the velocity variance between camera pairs. This allows us to recover Reynolds stresses for low signal-to-noise ratio images in synthetic experimental data of a turbulent channel flow.

1. Introduction

Particle image velocimetry (PIV) is the dominant non-intrusive velocimetry technique in experimental fluid mechanics, contributing to many advances in our understanding of turbulent and complex flows (Westerweel et al., 2013). Unlike pointwise methods, such as hot wire anemometry, PIV concurrently resolves the velocity field across many spatial locations by cross-correlating tracer particles within discrete interrogation windows of paired image exposures. Given a statistically converged series of frame pairs, mean velocities and Reynolds stresses can be recovered, which are fundamental to the characterisation of turbulent flows (Pope, 2000).

To calculate accurate velocities and statistics, a robust correlation peak is required. To achieve this, there must be a sufficient number of matched particle-image pairs within the interrogation

window (Westerweel et al., 2013). However, raising the particle density to meet this requirement cannot continue indefinitely, due to practical limitations such as multiple scattering, laser speckle and biases in PIV interrogation (Raffel et al., 2018). Therefore, a minimum interrogation window size exists for any experiment. This minimum also increases with image noise, which raises the correlation noise floor to be comparable to peak height (Scharnowski & Kähler, 2016).

PIV interrogation acts as a low pass spatial filter, attenuating flow features which are smaller than the interrogation window. This leads to the underestimation of Reynolds stresses and over-smoothing mean gradients (Kähler et al., 2012). Wall-bounded turbulence is the canonical setting in which this matters. Since the boundary layer thickness is typically matched to the PIV image size, only 10% of the spatial resolution is available to resolve the log layer. For smooth walls, the viscous and buffer layers span only a few wall units $y^+ = yu_\tau/\nu$, where u_τ is the friction velocity and ν the kinematic viscosity. Above moderate Reynolds numbers, the minimum window required for a robust peak corresponds to $> 10y^+$, so the near-wall peak and the steep mean shear at the wall cannot be resolved. Methods that improve the spatial resolution of PIV interrogation are therefore needed.

Single-pixel ensemble correlation, introduced by Westerweel et al. (2004), removes window size constraints by averaging correlation planes through a time series of frame pairs before peak detection. Signal-to-noise then decouples from the window dimension since particle accumulation is through time rather than spatially. The interrogation window size can shrink to a single pixel given sufficient data. Scharnowski et al. (2012) extended the operator from mean to second-order statistics by demonstrating that the ensemble averaged correlation map is the velocity probability density function convolved with the particle-image auto-correlation. Therefore, the broadening of the cross-correlation map above the particle autocorrelation recovers the normal stresses, and the orientation of the correlation map recovers the shear stress.

These developments have been focused on planar PIV, which measures only two velocity components in a two-dimensional plane. Flows of practical interest, including turbulent boundary layers with vortical structure, separated flows, and swirling jets, require statistics of all three components. Stereoscopic PIV infers the third component from the parallax error introduced by different perspectives of particles moving within the laser sheet (Prasad, 2000). Combining single-pixel ensemble correlation with a stereoscopic setup could therefore improve the spatial resolution of stereoscopic PIV, allowing high-resolution characterisation of all three components of the mean flow and Reynolds stresses.

A trivial extension of the method of Scharnowski et al. (2012) to two-camera stereoscopic imaging does not exist. Because the two cameras' lines of sight lie along opposite oblique directions, an out-of-plane displacement in the object plane appears in the image plane as a positive apparent in-plane shift in one camera and a negative shift in the other. For the mean flow, subtracting the two measured displacements isolates the out of plane displacement. However, Reynolds stresses are a

second moment, and squaring the displacement removes the ability to distinguish the out-of-plane and streamwise normal stresses.




Figure 1. A particle at the sheet undergoes the true displacement (u_x, u_z) . Each camera records only its image-plane projection: Cam 1 measures $d_1 = u_x + u_z \sin \theta$ and Cam 2 measures $d_2 = u_x - u_z \sin \theta$. The in-plane part u_x (blue) is common to both views. The out-of-plane part $u_z \sin \theta$ shifts the recorded image forward for Cam 1 and backward for Cam 2. Subtracting the two recovers u_z , but squaring them for a variance leaves both with $\sin^2 \theta u_z^2$, fusing R_{xx} and R_{zz} .

To close the deficiency, we present a stereoscopic correlation of correlations framework. This correlates the standard frame pair correlations of each camera and the resulting cross camera plane is ensemble averaged across all frame pairs. This ensemble carries the joint statistics of the two cameras' displacements. Because the views agree on in-plane motion and differ only in their out-of-plane sensitivity, the spread of this cross camera correlation is directly related to the out of plane fluctuations. This paper outlines current progress towards utilising this correlation of correlation information for the separation of the streamwise and out-of-plane normal stresses. To our knowledge, this is the first operator to correlate per-frame ensemble correlation planes across cameras for the explicit purpose of recovering the joint two-camera covariance and closing the stress rank deficiency.

2. Methodology

2.1. Planar single-pixel ensemble correlation

For an ensemble of N image pairs (A_i, B_i) , the per-window cross-correlation, evaluated at the correlation lag ξ , is averaged across the ensemble before peak detection:

$$\langle R_{AB} \rangle(\xi) = \frac{1}{N} \sum_{i=1}^N (A_i \star B_i)(\xi). \quad (1)$$

For turbulent flow under a Reynolds decomposition of the velocity field, $\mathbf{u} = \bar{\mathbf{u}} + \mathbf{u}'$, the measured displacement of each time instance inherits the same split, $\mathbf{d}_i = \boldsymbol{\mu} + \mathbf{d}'_i$, where $\boldsymbol{\mu}$ is the mean displacement and $\mathbf{d}'_i$ a zero-mean fluctuation governed by the turbulence of the flow. After taking

an ensemble of the instantaneous correlation peaks distributed around a zero mean, a correlation plane centred on $\boldsymbol{\mu}$ and broadened by the scatter of $\mathbf{d}'_i$ is formed. The mean displacement is used as the predictor field deformation so that the residual displacement in each frame is the fluctuation $\mathbf{d}'_i$ alone.

The expectation of the correlation peak over the ensemble is then the convolution of the particle-image autocorrelation with the displacement probability density function (Westerweel et al., 2004; Scharnowski et al., 2012). Scharnowski et al. (2012) model the velocity PDF as a rotated ellipse parametrised by two principal widths and an orientation. They show that the ensemble correlation peak is then itself a rotated elliptical Gaussian whose covariance is the sum of the particle image and displacement contributions,

$$\langle R_{AB} \rangle(\boldsymbol{\xi}) \propto G(\boldsymbol{\xi} - \boldsymbol{\mu}; \mathbf{S} + \boldsymbol{\Sigma}), \quad (2)$$

where $\mathbf{S}$ is the image autocorrelation covariance, and $\boldsymbol{\Sigma}$ the 2×2 displacement covariance. The classical recovery route fits a rotated elliptical Gaussian to the ensemble cross-correlation peak and a second one to the autocorrelation, then subtracts the principal widths axis-by-axis, $\sigma_{u'u'}^2 = \sigma_{AB,x}^2 - \sigma_{AA,x}^2$ and likewise for the spanwise and shear components.

2.2. K-space transfer-function fit

The spatial domain variance subtraction is subject to the combined fitting uncertainty of both the auto- and cross correlation's independently estimated variances. We avoid the subtraction altogether by fitting the direct transfer function between the autocorrelation and cross correlation in the Fourier domain in a single fit. Taking the Fourier transform of Eq. (2) and forming the transfer function with the geometric mean of the auto-correlations gives

$$T(\mathbf{k}) = \frac{\mathcal{F}\{\langle R_{AB} \rangle\}(\mathbf{k})}{\sqrt{\mathcal{F}\{\langle R_{AA} \rangle\}(\mathbf{k}) \cdot \mathcal{F}\{\langle R_{BB} \rangle\}(\mathbf{k})}} = \exp(-2\pi^2 \mathbf{k}^\top \boldsymbol{\Sigma} \mathbf{k}) \exp(-2\pi i \mathbf{k} \cdot \boldsymbol{\mu}). \quad (3)$$

The denominator $\sqrt{\mathcal{F}\{\langle R_{AA} \rangle\} \cdot \mathcal{F}\{\langle R_{BB} \rangle\}}$ is the measured Fourier transform of the per-frame autocorrelation. Where Scharnowski et al. (2012) fit a rotated elliptical Gaussian to the autocorrelation and then subtract its principal widths axis-by-axis, here we simply map the transfer function by its exact spectrum. The particle-image shape is implicitly handled, whether circular, elliptical, or fully astigmatic. Therefore, the Fourier approach holds for imperfect particle images that are stretched such as in the case of supersonic flow. The transfer function remains Gaussian, carrying $\boldsymbol{\Sigma} = (\Sigma_{xx}, \Sigma_{yy}, \Sigma_{xy})$ and a complex linear phase carrying $\boldsymbol{\mu}$. The fit is implemented in PIVtools as a bounded nonlinear least-squares problem on the stacked real and imaginary parts of $T(\mathbf{k})$ (Taylor et al., 2026).

In multipass PIV, the two images of each pair are resampled by the symmetric $\pm \mathbf{d}/2$ predictor before correlation, which requires interpolating each image at sub-pixel positions. This reduces in

plane pair loss by keeping the correlation plane centred on μ . Astarita & Cardone (2005) show that this interpolation is not a free choice, and that the scheme used to reconstruct the deformed image sets the bias, random error of the recovered displacement, and compute cost. As such, PIVtools exposes a bicubic and a Lanczos-3 kernel so this accuracy against-cost balance can be struck per case (Taylor et al., 2026). Any such kernel has finite support, so it implicitly acts as a low-pass filter, attenuating the high-wavenumber content of each interrogation.

The interpolation kernel leaves a coloured noise floor in the autocorrelation planes that form the denominator of Eq. (3). Its shape is set by the fractional pixel shift $\mathbf{f} = (f_x, f_y)$ at which the re-sampling lands at each window centre. We denote it $P_{\text{noise}}(\mathbf{k}; \mathbf{f})$, the squared magnitude of the discrete-time Fourier transform of the interpolation kernel weights. The same kernel also modulates the recovered displacement field, an effect quantified for image-deformation PIV through the modulation transfer function of Astarita (2007).

The measured reference is therefore not the bare particle spectrum but that spectrum lifted by the coloured floor,

$$\mathcal{F}\{\langle R_{AA} \rangle\}(\mathbf{k}) = P_{\text{noise}}(\mathbf{k}; \mathbf{f}) [\hat{S}(\mathbf{k}) + N_0], \quad (4)$$

with $\hat{S}$ the particle-image power spectrum and N_0 the pre-resampling white noise level. Substituting Eq. (4) into the denominator of Eq. (3) gives the transfer function that is actually fitted,

$$T(\mathbf{k}) = \frac{\mathcal{F}\{\langle R_{AB} \rangle\}(\mathbf{k})}{\sqrt{P_{\text{noise}}(\mathbf{k}; \mathbf{f}) [\hat{S} + N_0]_A \cdot P_{\text{noise}}(\mathbf{k}; \mathbf{f}) [\hat{S} + N_0]_B}}. \quad (5)$$

The floor enters the denominator alone. The cross-correlation numerator is built from the two independent exposures, so its interpolation noise is zero-mean and averages out over the ensemble, whereas the autocorrelation is a self-product and retains a positive coloured floor. The multiplicative kernel response P_{noise} on the signal cancels against the numerator in the geometric-mean ratio, leaving the additive N_0 as the residual bias. Because P_{noise} is fixed a priori by the chosen kernel and fractional shift $\mathbf{f}$, only the amplitude N_0 is estimated, jointly with Σ and μ .

2.3. Stereoscopic projection and the rank-five singularity

The extension of the planar displacement PDF estimation is considered for two cameras at symmetric angles $\pm\theta$ about the world z -axis, viewing a common illuminated xy -plane as shown in Figure 1. Both cameras lie in the world xz -plane and are rotated only about the y -axis, so neither optical axis carries a y -component and the wall-normal displacement projects identically onto both sensors, decoupling u_y from u_x and u_z . Each camera's measured in-plane displacement projects

from the 3-D velocity (u_x, u_y, u_z) as

$$\begin{aligned} d_1(t) &= u_x(t) + \sin \theta \cdot u_z(t), \\ d_2(t) &= u_x(t) - \sin \theta \cdot u_z(t), \\ d_{y,c}(t) &= u_y(t) \quad \text{for } c \in \{1, 2\}. \end{aligned} \quad (6)$$

Each camera's cross-correlation plane yields the in-plane covariance $(\Sigma_{xx}^c, \Sigma_{yy}^c, \Sigma_{xy}^c)$, giving six measurements in total. Substituting the projection (Eq. 6) into the variance and covariance definitions:

$$\begin{aligned} \Sigma_{xx}^1 &= R_{xx} + 2 \sin \theta R_{xz} + \sin^2 \theta R_{zz}, \\ \Sigma_{xx}^2 &= R_{xx} - 2 \sin \theta R_{xz} + \sin^2 \theta R_{zz}, \\ \Sigma_{yy}^c &= R_{yy}, \\ \Sigma_{xy}^1 &= R_{xy} + \sin \theta R_{yz}, \\ \Sigma_{xy}^2 &= R_{xy} - \sin \theta R_{yz}. \end{aligned} \quad (7)$$

The six measurements collapse to five linearly independent combinations. In this camera arrangement, R_{yy} is measured directly, the half-sum and half-difference of the σ_{xy}^c giving R_{xy} and R_{yz} , the half-difference of the σ_x^{c2} giving R_{xz} , and the half-sum

$$A \equiv \frac{1}{2}(\Sigma_{xx}^1 + \Sigma_{xx}^2) = R_{xx} + \sin^2 \theta R_{zz}. \quad (8)$$

The system has rank five in the six-component Reynolds tensor $\mathbf{R} = (R_{xx}, R_{yy}, R_{zz}, R_{xy}, R_{xz}, R_{yz})^\top$. Its null space is defined by $\mathbf{n} = (\sin^2 \theta, 0, -1, 0, 0, 0)^\top$. Therefore, an increase of R_{zz} by ε , accompanied by a decrease of R_{xx} by $\sin^2 \theta \varepsilon$, leaves all six per-camera measurements identically unchanged. The stress coupling is structural and independent of calibration disparity and persists for any camera arrangement.

2.4. Correlation-of-correlations recovers the joint statistic

Closing the rank deficiency requires a sixth measurement, independent of the five per camera combinations, that separates R_{xx} from R_{zz} . The two per camera variances are discarded as the single projection is squared such that $+\sin \theta u_z$ and $-\sin \theta u_z$ becomes $+\sin^2 \theta u_z^2$. However, by taking the correlation of each camera's frame pair correlations, the sign of the covariance is retained. Within a single frame, a camera's per-window AB plane is approximately Gaussian about the instantaneous displacement $d_c(t)$,

$$R_{AB}^c(\boldsymbol{\xi}; t) \approx G(\boldsymbol{\xi} - d_c(t); S). \quad (9)$$

Cross-correlating the two cameras' frame pair correlation planes centres the result on the inter-camera difference $d_2(t) - d_1(t) = -2 \sin \theta u_z(t)$, which carries only the out-of-plane motion, with a width set by the two particle-image variances,

$$C(t; \Delta \xi) \equiv R_{AB}^1(\cdot; t) \star R_{AB}^2(\cdot; t) \approx G(\Delta \xi - (d_2(t) - d_1(t)); 2S). \quad (10)$$

Averaging this map across the ensemble broadens its peak by the variance of the difference,

$$\langle C \rangle(\Delta \xi) \approx G(\Delta \xi - \langle d_2 - d_1 \rangle; 2S + \text{Var}(d_2 - d_1)). \quad (11)$$

The broadening above the particle-image width is therefore $\text{Var}(d_2 - d_1) = \Sigma_{11,xx} + \Sigma_{22,xx} - 2\Sigma_{12,xx}$. The per-camera diagonals $\Sigma_{11,xx}$ and $\Sigma_{22,xx}$ are already fixed in Section 2.3, so the measured broadening delivers $\Sigma_{12,xx}$, and with it the separation of Eq. (13).

$$\Sigma_{12,xx} \equiv \langle d_1 d_2 \rangle = \langle (u_x + \sin \theta u_z)(u_x - \sin \theta u_z) \rangle = R_{xx} - \sin^2 \theta R_{zz}. \quad (12)$$

Here R_{zz} enters with the opposite sign to A , which restores a rank six formulation.

$$R_{xx} = \frac{1}{2}(A + \Sigma_{12,xx}), \quad R_{zz} = \frac{1}{2\sin^2 \theta}(A - \Sigma_{12,xx}), \quad (13)$$

2.5. Dewarping prerequisite

Instantaneous PIV in PIVtools correlates in the image plane, following the standard guidance that preserves sharp particle images and, as Jankee & Ganapathisubramani (2020) show for stereoscopic PIV, keeps each camera's sub-pixel displacement statistics intact. However, the correlation of correlations operator cannot directly follow this route. The output of correlating the two cameras' per-frame correlation planes against one another is only physically meaningful when both planes share a single coordinate system with common units. A camera at $\pm\theta$ has magnification differing from a normal view by approximately $\cos \theta$, so the per-camera pixel² spreads would not be directly comparable across cameras.

Therefore we dewarp both camera images onto a common world- xy grid *before* correlation. The trade-off is that interpolation noise is admitted into the correlation step in exchange for unit consistency across cameras. Importantly, the image to object plane operator and predictor field deformation are jointly applied in one operation, producing a known coloured attenuation and therefore a known value of $P_{\text{noise}}(\mathbf{k}; \mathbf{f})$.

2.6. Finite laser sheet thickness

The dewarping maps both cameras onto the calibration plane, $z = 0$, yet a stereoscopic laser sheet has finite thickness. Therefore, a particle lying ahead of or behind $z = 0$ is mapped to the

wrong world position and because the two cameras view from $\pm\theta$, the error is antisymmetric. This cross-camera disparity smears the correlation plane through the mechanisms that Wieneke (2005) observes in self-calibration, where the disparity peak is an extended streak whose length scales with the sheet thickness. This static spread adds to the variance $\text{Var}(d_2 - d_1)$ of Eq. (11), so the cross-camera plane $\langle C \rangle$ is broadened above the genuine out-of-plane signal. We remove it with a per-frame autocorrelation correlation-of-correlations $\langle C_{AA} \rangle = \langle (C_{1,AA-BB} \star C_{2,AA}) \rangle$, which carries the cross camera disparity but no inter-frame motion. Thus, it isolates the disparity induced spread.

3. Synthetic Images

Both the planar single-pixel ensemble model of Section 2.1 and the stereoscopic correlation-of-correlations method are validated against synthetic images of a turbulent channel flow at friction Reynolds number $\text{Re}_\tau = u_\tau \delta / \nu = 1000$. The underlying velocity fields are taken from the Johns Hopkins Turbulence Databases (JHTDB) channel-flow direct numerical simulation of Graham et al. (2016), accessed through the JHTDB public web-services interface (Li et al., 2008). The parameters of the two datasets are summarised in Table 1.

Particles are seeded uniformly throughout the measurement volume and the velocity is interpolated at a single time instant through fourth-order Lagrange spatial interpolation with no temporal interpolation. The field of view corresponds to a square $h \times h$ region spanning the bottom half of the channel, where h is the channel half-height. The two exposures of each pair follow from a symmetric first-order Taylor step in time, and the streamwise homogeneity of the channel is exploited by drawing each query from a random x offset within the domain $[0, 8\pi] h$ while y is held fixed, improving the statistical independence of the sampled turbulence realisations.

Synthetic images are generated with the EUROPIV Synthetic Image Generator (SIG) of Lecordier & Westerweel (2004). Particle images were generated with a characteristic image diameter of 3px and a maximum intensity of 116 counts and 8-bit depth. The planar dataset was aligned with the streamwise (x) and wall-normal (y) directions in a near-telecentric imaging configuration. The stereoscopic dataset simulates two cameras at $\pm 45^\circ$ from the sheet normal in a side-scatter arrangement.

Table 1. Synthetic-image parameters for the planar noise-model validation and the stereoscopic laser-sheet-thickness sweep. h is the channel half-height and U_b the bulk velocity.

Parameter	Planar Noise Test	Stereo
Image pairs	4,000	1,000
Sensor resolution	2048×2048 px	1440×2080 px
Cameras	single, in-plane	two, $\pm 45^\circ$ side-scatter
Laser sheet thickness	$0.0078 h$ (16 px)	$\{0.7, 0.8, 0.9, 1.0, 1.1, 1.2\}$ mm
Particles per image	2.2×10^4 ($N_{\text{ppp}}=0.005$)	8.5×10^4
Background noise	Gaussian, $\mu=80$, $\sigma=16$	none
Interframe Δt	$6.48 \times 10^{-3} h/U_b$	$6.48 \times 10^{-3} h/U_b$

To identify the disparity sensitivity to laser-sheet thickness, six laser sheet thicknesses T are tested as listed in Table 1. Each thickness is realised by two coupled changes. First, the SIG sheet model is set to a Gaussian intensity profile of full width T , so that a particle at out-of-plane position z is weighted by $\exp[-8(z/T)^2]$ and excluded beyond $\pm T/2$. Second, the stored particle field is rescaled in the out-of-plane direction by the factor T/T_0 , with $T_0 = 1.2$ mm the original render thickness, while the frame-to-frame out-of-plane displacement $\delta z = z_B - z_A$ is preserved,

$$z'_A = \frac{T}{T_0} z_A, \quad z'_B = z'_A + (z_B - z_A). \quad (14)$$

Preserving δz in Eq. (14) holds the true out-of-plane velocity, and hence the full Reynolds-stress tensor, fixed across the sweep. By rescaling only the absolute z we project the same two-dimensional particle-image at every thickness.

4. Results

4.1. Planar Ensemble PIV

Figure 2 compares the three in-plane Reynolds-stress components recovered from the noisy planar dataset, processed with and without the coloured-noise model of Section 2.2 against the DNS reference. The ensemble pipeline used a four-pass scheme progressing from 64×64 through 32×32 and 16×16 to 16×8 px at 50% overlap.

Without the model, the recovered profile's error shows periodic deviations from the reference. The overestimations grow with the mean particle displacement d and are largest at odd integer values of d . Since image deformation is applied symmetrically, this corresponds to half-integer values of the displacement $d/2$ applied when interpolating the images. The streamwise normal stress overshoots its inner peak by about 20% and overestimates the whole near-wall region. The wall-normal stress is significantly overestimated at odd integer displacement. With the noise model enabled, the oscillations are removed even at the lower window size of 16×8 px, resolving down to $y^+ \approx 2$. The same figure includes a standard 32×32 instantaneous reconstruction of the identical

noisy images. This is the smallest window that returns valid vectors for the instantaneous method; this larger window size is unable to meaningfully resolve the near-wall peak since the first vector is at $y^+ \approx 12$.

At even integer pixel deformation, the image interpolation becomes exact and the image noise contribution is white. This is why, even without the model, the stresses trend towards the ground truth at these deformations. Away from these values, image noise becomes filtered during image interpolation, artificially broadening the correlation peak. An uncorrected fit mis-attributes this as displacement variance. This effect is therefore a leading-order error source for ensemble PIV with noisy images, and modelling it directly is required to make robust high-resolution statistics usable.




Figure 2. Planar Reynolds stresses for a noisy dataset recovered by ensemble correlation, with the coloured-noise model enabled (orange squares) and disabled (blue circles), against the directly binned DNS reference (black lines with 95% bootstrap band). The vertical lines mark the mean streamwise displacement $d/2$ applied to correct both images for the mean flow. Green triangles show a standard 32×32 instantaneous reconstruction of the same noisy images.

4.2. Stereoscopic Ensemble PIV

To extend planar ensemble PIV to stereoscopic ensemble PIV, we must measure the variance of the parallax displacement between cameras. This is achieved using the correlation-of-correlations technique. Figure 3 shows the correlation of correlations reference plane $\langle C_{AA} \rangle$ and the cross-camera plane $\langle C \rangle$ at the thinnest and thickest laser sheets for a representative 32×32 px window, with a cut through each peak. We observe that increasing the laser sheet thickness increases the correlation peak width. This manifests as a wide, non-Gaussian tail, which a single Gaussian peak fit cannot adequately model. This tail arises from image interpolation and the finite laser sheet calibration error rather than from added sensor noise. This effect is not present in the y direction, since the disparity between perspectives does not affect the y direction for this camera configuration, and because there is no mean displacement in the y direction.




Figure 3. Cross-camera 32×32 px correlation planes at the thinnest ($T = 0.7$ mm, top) and thickest ($T = 1.2$ mm, bottom) laser sheets for a single window at $h/2$. The left two columns are the matched-pair reference $\langle C_{AA} \rangle$ (COC AC) and the cross-camera plane $\langle C \rangle$ (COC) as normalised two-dimensional maps, and the right two columns are a horizontal cut through each peak with an attempted single Gaussian fit.

Table 2 reports $\Sigma_{\text{diff},xx} \equiv \Sigma_{11,xx} + \Sigma_{22,xx} - 2\Sigma_{12,xx} = \text{Var}(d_2 - d_1)$ of Eq. (11), obtained from our peak fitting method. This quantity corresponds to the difference in width between the correlation peaks found in the $\langle C \rangle$ and $\langle C_{AA} \rangle$ planes caused by the parallax displacement. It measures the variance of the parallax displacement seen between cameras, and is proportional to the out-of-plane Reynolds stress R_{zz} .

Across the range of laser sheet thicknesses tested, we are able to robustly identify this velocity variance. Our transfer function method does not require the correlation peaks to be Gaussian. Instead, we simply seek to model the relationship between $\langle C \rangle$ and $\langle C_{AA} \rangle$. Thus, the non-Gaussian tail evident in $\langle C \rangle$ and $\langle C_{AA} \rangle$ is tolerated and an effective out-of-plane signal is recovered despite it.

Table 2. Out-of-plane reference $\Sigma_{\text{diff},xx}$ recovered by the cross-camera k -space transfer fit, evaluated in the band $y = -74.5 \pm 12$ mm, across the laser-sheet-thickness sweep.

Sheet thickness T [mm]	0.8	0.9	1.0	1.1	1.2
$\Sigma_{\text{diff},xx}$	0.617	0.620	0.622	0.625	0.628




Figure 4. (a) The per-camera half-sum $A = \overline{u'u'^+} + \sin^2 \theta \overline{w'w'^+}$ of Eq. (8), at $T = 1.2$ mm and a 32×32 px cross-camera window, against the directly binned DNS A (line with 95% bootstrap band). A is built from the per-camera variances alone; its agreement with the DNS localises the $\overline{w'w'^+}$ deficit to the partition of A performed by the cross-camera correlation of correlations. (b) The three normal Reynolds stresses in wall units for the same window against the DNS reference (solid lines with 95% band).

Figure 4(a) shows the per-camera half-sum $A = \overline{u'u'^+} + \sin^2 \theta \overline{w'w'^+}$ of Eq. (8), recovered from the cross-camera reconstruction at $T = 1.2$ mm and a 32×32 px interrogation window, against the DNS reference. A lies within the statistical uncertainty of the ground truth. Because A is built from the per-camera variances alone, its agreement establishes that the per-camera measurement and reconstruction are correct. The out of plane and streamwise stresses are then decoupled through the cross-camera term $\Sigma_{12,xx}$ which is the closure of Eq. (13). Therefore, any error in $\Sigma_{12,xx}$ tends to cause $\overline{u'u'^+}$ to be overestimated and $\overline{w'w'^+}$ to be underestimated, or vice versa.

Figure 4(b) shows the three normal stresses. The wall-normal stress $\overline{v'v'^+}$ is measured in plane and projects identically onto both cameras, independently of the out-of-plane reconstruction. $\overline{v'v'^+}$ is

in close agreement with the ground truth. The out-of-plane stress $\overline{w'w'}^+$ reproduces the correct profile shape but is systematically underestimated and $\overline{u'u'}^+$ is correspondingly overestimated. Therefore, the measured $\Sigma_{\text{diff},xx}$ from the correlation of correlations is under predicted, which has been attributed to the correlation peak broadening seen in Figure 3. The transfer fit cannot yet cleanly separate this effect from turbulence signal, biasing the recovered $\Sigma_{12,xx}$ high and R_{zz} low.

5. Conclusions

We have introduced a new peak fitting method for ensemble PIV. This method infers the turbulent velocity variance by describing a transfer function in Fourier space which maps the autocorrelation function of an image to the cross-correlation function between images, accounting for the effects of image interpolation. With noisy synthetic PIV images, a naïve fit shows that this image interpolation effect can introduce errors of up to 50% in the velocity variance. Incorporating interpolation effects into the peak fitting algorithm recovers all three in-plane Reynolds stresses to within 3% of the ground truth at a window size of 16×8 px. For context, conventional PIV only yielded adequate measurements at a window size of 32×32 in this test case.

Building on this, we have formulated a method to measure all three normal Reynolds stresses from stereoscopic PIV images. This is achieved using the correlation of correlations technique, which measures the variance in the parallax of the displacement observed from different perspectives. Across a range of laser sheet thicknesses, we recover a consistent parallax-variance measurement from synthetic PIV data, though it is biased low by the peak broadening. While in-plane stresses are measured to within the statistical uncertainty of our data, the out-of-plane stress shows the correct trends but is systematically underestimated. This is because of effects due to calibration and interpolation errors, which confound the accurate determination of changes in the correlation-of-correlations peak width that are due to the flow.

Closing this gap requires separating the out of plane signal from these deterministic effects. Promising routes are a model for the cross-camera correlation-of-correlations, analogous to the coloured-noise model used for the planar fit, and performing cross-correlations in undistorted images to avoid introducing interpolation artefacts. Both are left for future work.

Acknowledgements

The authors acknowledge funding from EPSRC grants EP/Y036409/1 and EP/W026090/1.

References

- Astarita, T. (2007). Analysis of weighting windows for image deformation methods in PIV. *Experiments in Fluids*, 43, 859–872. doi: 10.1007/s00348-007-0314-2
- Astarita, T., & Cardone, G. (2005). Analysis of interpolation schemes for image deformation methods in PIV. *Experiments in Fluids*, 38(2), 233–243. doi: 10.1007/s00348-004-0902-3
- Graham, J., Kanov, K., Yang, X. I. A., Lee, M., Malaya, N., Lalescu, C. C., ... Meneveau, C. (2016). A Web services accessible database of turbulent channel flow and its use for testing a new integral wall model for LES. *Journal of Turbulence*, 17(2), 181–215. doi: 10.1080/14685248.2015.1088656
- Jankee, G. K., & Ganapathisubramani, B. (2020). Comparison between object and image plane cross-correlation for stereoscopic PIV in the presence of pixel locking. *Experiments in Fluids*, 61, 89. doi: 10.1007/s00348-020-2916-x
- Kähler, C. J., Scharnowski, S., & Cierpka, C. (2012). On the resolution limit of digital particle image velocimetry. *Experiments in Fluids*, 52, 1629–1639. doi: 10.1007/s00348-012-1280-x
- Lecordier, B., & Westerweel, J. (2004). The EUROPIV synthetic image generator (S.I.G.). In M. Stanislas, J. Westerweel, & J. Kompenhans (Eds.), *Particle image velocimetry: Recent improvements (EUROPIV 2)* (pp. 145–161). Berlin: Springer. doi: 10.1007/978-3-642-18795-7_11
- Li, Y., Perlman, E., Wan, M., Yang, Y., Meneveau, C., Burns, R., ... Eyink, G. (2008). A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence. *Journal of Turbulence*, 9(N31). doi: 10.1080/14685240802376389
- Pope, S. B. (2000). *Turbulent flows*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511840531
- Prasad, A. K. (2000). Stereoscopic particle image velocimetry. *Experiments in Fluids*, 29(2), 103–116. doi: 10.1007/s003480000143
- Raffel, M., Willert, C. E., Scarano, F., Kähler, C. J., Wereley, S. T., & Kompenhans, J. (2018). *Particle image velocimetry: A practical guide* (3rd ed.). Berlin: Springer. doi: 10.1007/978-3-319-68852-7
- Scharnowski, S., Hain, R., & Kähler, C. J. (2012). Reynolds stress estimation up to single-pixel resolution using PIV-measurements. *Experiments in Fluids*, 52, 985–1002. doi: 10.1007/s00348-011-1184-1
- Scharnowski, S., & Kähler, C. J. (2016). On the loss-of-correlation due to PIV image noise. *Experiments in Fluids*, 57, 119. doi: 10.1007/s00348-016-2203-z

- Taylor, M. T., Lawson, J. M., & Ganapathisubramani, B. (2026). PIVtools: A comprehensive open source software for particle image velocimetry. *SoftwareX*. (Preprint, submitted 2026-05-08)
- Westerweel, J., Elsinga, G. E., & Adrian, R. J. (2013). Particle image velocimetry for complex and turbulent flows. *Annual Review of Fluid Mechanics*, 45, 409–436. doi: 10.1146/annurev-fluid-120710-101204
- Westerweel, J., Geelhoed, P. F., & Lindken, R. (2004). Single-pixel resolution ensemble correlation for micro-PIV applications. *Experiments in Fluids*, 37, 375–384. doi: 10.1007/s00348-004-0826-y
- Wieneke, B. (2005). Stereo-PIV using self-calibration on particle images. *Experiments in Fluids*, 39, 267–280. doi: 10.1007/s00348-005-0962-z