

Quantification and adjustment of pixel-locking in Particle Image Velocimetry

R. J. Hearst · B. Ganapathisubramani

the date of receipt and acceptance should be inserted later

Abstract A quantification metric is provided to determine the degree to which a particle image velocimetry data set is pixel-locked. The metric is calculated by integrating the histogram equalization transfer function and normalizing by the worst-case-scenario to return the percentage pixel-locked. When this metric is calculated for each position in the vector field, it is shown that pixel-locking is non-uniform across the field. Hence, pixel-locking adjustments should be made on a vector-by-vector basis rather than uniformly across a field, although the latter is the common practice. A methodology is provided to compensate for the effects of pixel-locking on a vector-by-vector basis. This includes applying a Gaussian filter directly to the images, processing the images with window deformation, ensuring the vector fields are in pixel displacements, applying histogram equalization calculated at each vector coordinate, and mapping the adjusted vector fields to physical space.

Pixel-locking is the tendency of a particle image velocimetry (PIV) velocity field to be biased towards integer values of pixel displacements. It is a result of insufficient resolution to accurately determine the sub-pixel displacement. This phenomenon is also frequently referred to as 'peak-locking'. Pixel-locking is an intrinsic error imbedded in an image at acquisition. To a certain degree, the effects of pixel-locking can be mitigated by the sub-pixel estimator during processing, but in severe cases the choice of estimator has no influence (Christensen, 2004).

Pixel-locking can adversely effect our interpretation of the underlying physics of flows. Christensen (2004) demonstrated that the only statistic that was insensitive to pixel-locking was the mean velocity profile. Pixel-locking also

corrupts our interpretation of flow phenomena such as uniform momentum zones (UMZs); see Meinhart and Adrian (1995) for a full description of this flow phenomenon. UMZs manifest as peaks in the instantaneous probability density function (PDF) of a velocity flow field, however, so does pixel-locking. Kwon et al. (2014) demonstrated in fully developed channel flow that PIV data tended to predict the same UMZ modal velocities in every image. However, direct numerical simulations of the same flow found that there was no bias towards particular velocities for UMZs, demonstrating that the coalescence of UMZ modal velocities in PIV is unphysical.

Recent advances in PIV uncertainty have focussed on local effects and non-uniformity of uncertainty across the image plane. Sciacchitano et al. (2015) found that the most accurate estimate of the uncertainty in PIV came from the correlation statistics approach of Wieneke (2015). However, this approach was found to decrease in its accuracy of the estimated uncertainty by 10% in the presence of pixel-locking (Sciacchitano et al., 2015). This suggests a need for a method of both quantifying pixel-locking and determining if its effects are global or local.

Two common techniques to reduce the effects of pixel-locking are deliberately defocussing the image (either prior to acquisition or during pre-processing) and applying a histogram equalization (HE) algorithm during post-processing. The use of HE in PIV is detailed by Roth and Katz (2001), and will be discussed in detail as we progress through this letter. The goals of this work are to: (i) present a quantification methodology to assess how pixel-locked a data set is, (ii) identify that pixel-locking varies through an image and each vector position should be treated independently, and (iii) present a simple algorithm to compensate for pixel-locking on a vector-by-vector basis.

A sample data set of planar PIV is used to facilitate the present discussion. The captured flow was a turbulent

boundary layer with a turbulent free-stream. This data was deliberately chosen because there was significant turbulence throughout the image plane, and as such none of the pixel-locking phenomena shown here are dependent on the flow structure. The turbulent boundary layer was formed over a plate in a $0.9 \text{ m} \times 0.6 \text{ m} \times 4.5 \text{ m}$ wind tunnel. The free-stream turbulence was generated by an active grid based on the original design of [Makita \(1991\)](#). The turbulence intensity in the free-stream was 12.7% and was measured with a hot-wire probe. 2000 images were captured with a LaVision ImagerProLX CCD 16 mega-pixel camera fitted with a Nikon Nikkor 200 mm lens. The field-of-view was approximately $1.2\delta \times 0.8\delta$ and was illuminated by a Litron Lasers Nano L200 15PIV Nd:YAG laser (523 nm wavelength, 200 mJ/pulse 15 Hz repetition rate).

The images were processed with LaVision DaVis 8.2.2, although the present technique is not dependent on this software. During pre-processing, a sliding background subtraction filter was applied to the raw images to reduce noise. This was followed by applying a 3×3 Gaussian blur kernel to each image to synthetically smooth the particles over a greater number of pixels. It is important to exercise caution when synthetically blurring images, because the uncertainty associated with the sub-pixel estimator will increase if particles are broadened over too many pixels. The images were then processed with five passes of decreasing window size from $96 \text{ pixels} \times 96 \text{ pixels}$ to $16 \text{ pixels} \times 16 \text{ pixels}$. Window deformation was used on final passes, which also mitigates pixel-locking. Vector fields were deliberately generated with units of millimeters and meters per second because we sought an algorithm that could be applied to any completely processed vector field set rather than requiring specific processing output, e.g., unmapped and unscaled pixel displacements, however, the following analysis could also be conducted on such a field.

First, the degree of pixel-locking in the data set is assessed through a histogram equalization process. Histograms of the sub-pixel displacement of the total \tilde{u} - and \tilde{v} -velocity distributions combined are compiled for each vector (i.e., one histogram per vector coordinate in the vector field); the local velocity is denoted $\tilde{u} = U + u$, where U is the local mean velocity. The streamwise and wall-normal velocity components are combined because there should be no bias of pixel-locking to any particular direction. An example of such a histogram for the position $(x/\delta, y/\delta) = (1.09, 0.74)$ in the present field is given in Fig. 1(a). This is the histogram of the velocity in pixel displacement about any integer value at this position across all 2000 images. Fig. 1(a) reveals a bias towards the integer value (0 in the figure). A transfer function, $s(r)$, can then be generated that will reorganize the histogram to have no bias toward any particular sub-pixel

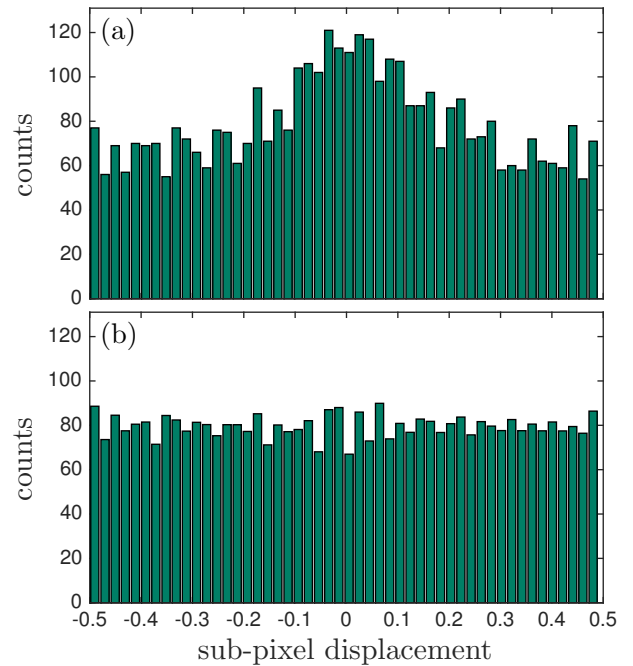


Fig. 1: Histogram of the sub-pixel displacement (a) without and (b) with histogram equalization at $(x/\delta, y/\delta) = (1.09, 0.74)$.

value ([Roth and Katz, 2001](#)),

$$s(r) = r_{\max} \int_{-0.5}^r \mathcal{P}(r) dr, \quad (1)$$

where r is the sub-pixel displacement, and \mathcal{P} is the probability density function of the sub-pixel displacement distribution. A transfer function, $s(r)$, is generated for each vector in the field. An example of the transfer function for $(x/\delta, y/\delta) = (1.09, 0.74)$ is given in Fig. 2. If there was no pixel-locking for this vector position, then the transfer function would lie along the one-to-one line. When $s(r)$ is applied to the velocity field, the histogram about integer values no longer shows a bias, Fig. 1(b). This approach assumes a deterministic relationship between pixel-locking and the true velocity field.

In the past it has been common to generate a single transfer function for the entire flow field ([Roth and Katz, 2001](#)) or one for each row of vectors in the wall-normal direction. These assume either that the pixel-locking is uniform across the entire image or for a given wall-normal position. Here, we have computed separate transfer functions for each position in the vector field. As such, we can determine the distributions of pixel-locking throughout the image. As a metric, we employ

$$Q = \frac{1}{0.25} \int_{-0.5}^{0.5} |s(r) - r| dr. \quad (2)$$




Fig. 2: Transfer function at $(x/\delta, y/\delta) = (1.09, 0.74)$.

This quantity is the absolute value of the area under the curve of the transfer function less the case with no pixel-locking. The coefficient represents the worst-case pixel-locking scenario, and thus Q is the percentage of pixel-locking at each position; 0 represents no pixel-locking and 1 represents the case where only integer values are measured. In practice, (1) and (2) are evaluated as sums of discrete values and as such can be affected by the bin size of the histograms. In this study we employ 0.02 pixel bins as a balance between producing a smooth $s(r)$ and the actual accuracy of the measurement.

A map of Q in our flow field is provided in Fig. 3. For reference, $Q = 0.11$ at $(x/\delta, y/\delta) = (1.09, 0.74)$, the position investigated in Fig. 1 and 2. Fig. 3 clearly identifies that pixel-locking is non-uniform across the image plane, with variations in the degree of pixel-locking of up to 10%. The non-uniformity is likely a combination of several factors including non-uniform magnification, and sensor, lens, and laser sheet impurities. This suggests that applying a HE methodology uniformly across the flow field would introduce a bias of its own, and hence HE adjustments should be carried out vector-by-vector rather than en masse.

Given that the effects of pixel-locking are local, we now provide a simple algorithm that can be used on an existing data set to adjust for pixel-locking on a vector-by-vector basis. A sample instantaneous velocity field is given in Fig. 4(a)¹. The histogram of \tilde{u} in this field is given in Fig. 5(a). We wish to know if the peaky structure (e.g., peaks at $\tilde{u}/U_\infty = 1.03, 1.08$ and 1.12) is a result of the measurement technique or a consequence of the physics. If a calibration has been used to give the velocity field in m/s, the first step is to rescale the

¹ The high level of turbulence in the free-stream makes it possible for the instantaneous velocity to exceed U_∞ by a significant amount.




Fig. 3: Map of the degree to which each vector is pixel-locked.

velocity field back to the image plane in units of pixel displacements. This can be accomplished with the following set of equations (Adrian and Westerweel, 2011):

$$\Delta x = \frac{\partial f(x', y')}{\partial x'} \Delta x' + \frac{\partial f(x', y')}{\partial y'} \Delta y', \quad (3)$$

$$\Delta y = \frac{\partial g(x', y')}{\partial x'} \Delta x' + \frac{\partial g(x', y')}{\partial y'} \Delta y', \quad (4)$$

where Δx and Δy are the image plane pixel displacements, $\Delta x'$ and $\Delta y'$ are the object plane displacements, and $x = f(x', y')$ and $y = g(x', y')$ are the mapping functions between planes. The calibration process produces the mapping functions, and hence the above can be used to take mapped velocities back to their native state. An alternative is to process the particle images to obtain unmapped, unscaled pixel displacements, which reduces the total number of operations.

Histogram equalization is then applied as described above on a vector-by-vector basis, producing a series of transfer functions with one for each vector position in the field. After these transfer functions have been applied to the vector fields, the data is mapped and scaled to the object plane with the calibration and a new histogram equalized velocity field is produced, Fig. 4(b). The histogram of the adjusted field is given in Fig. 5(b) where two of the unadjusted image's peaks appear to have been locked to integer values of the pixel displacement. It is also significant that the new velocity field is not visibly different from the original field. A disparity map between the original and the adjusted instantaneous \tilde{u} fields is provided in Fig. 4(c) and reveals that in general the changes are contained within $\pm 0.3\%$ for this data set. This is important because it indicates that the influence of pixel-locking has been damped while the overall flow organization remains unchanged.




Fig. 4: Instantaneous U -velocity field (a) without and (b) with histogram equalization applied vector-by-vector, and (c) the disparity map between them.




Fig. 5: Histograms of the instantaneous U -velocity field (a) without and (b) with histogram equalization applied vector-by-vector.

In the implementation of the above protocol, pixel-locking adjustments are applied only where they are necessary (i.e., where the peaks are at integer pixel values). A peak may occur at a non-integer value for very small pixel displacements (e.g., near the wall), or in approximately laminar flows. These non-integer peaks are due to the flow phenomenology rather than pixel-locking. To avoid distorting these peaks, a simple conditional statement that does not apply the transfer function at positions where the sub-pixel displacement peak is misaligned from the integer value by more than ± 0.05 pixels appears to preserve the nature of the flow.

Ultimately, this processing algorithm can be broken down into the follow steps:

- (i) De-noise the images.
- (ii) Synthetically blur the images with a Gaussian filter.
- (iii) Process the images using window deformation.
- (iv) De-map the vector fields back to the image plane (if they were not processed as unmapped, unscaled pixel displacements).
- (v) Calculate histograms across all images for each vector in the field and identify pixel-locked vectors.
- (vi) Compute the transfer functions for pixel-locked vectors and apply it.
- (vii) Map the vector fields to the object plane.

A quantification metric was also provided to determine the degree of pixel-locking. The metric is the absolute value

of the area under the histogram equalization transfer function deviating from the no pixel-locking case. This metric was used to demonstrate that pixel-locking is not uniform across an image. Therefore, identifying and adjusting for pixel-locking with histograms computed based on entire vector fields may be erroneous. Consequently, any equalization process should be applied on a vector-by-vector basis. This finding, and the quantification metric provided in (2) could be integrated into the local uncertainty estimates discussed by Sciacchitano et al. (2015), to improve their accuracy in pixel-locked data.

The local variations in pixel-locking are particularly significant when a vector field is processed from multiple simultaneous images, especially stereoscopic PIV. Pixel-locking is imbedded in each image separately, and thus the final field cannot be corrected in the fashion described above. In principle, one would need to adjust for pixel-locking in each image separately, either by computing the velocity fields individually and using a similar process to that described herein before combining the separate vector fields to produce the stereo field, or using another form of adjustment directly on the images before performing the correlations. Future work should thus focus on identifying and compensating for pixel-locking in such cases.

Acknowledgements This project was funded by ERC (grant agreement no. 277472) and EPSRC (grant reference no. EP/I037717/1). RJH acknowledges the financial support of the Natural Sciences and Engineering Research Council of Canada (NSERC). The authors are grateful to Ms. E. Dogan who acquired the data used for this analysis, Dr. R. de Kat for enlightening conversations on pixel-locking, and Mr. T. Medjnoun for validating the methodology with independent data sets. RJH would also like to thank Dr. C. Vanderwel for her PIV support and guidance.

References

- R. J. Adrian and J. Westerweel. *Particle Image Velocimetry*. Cambridge University Press, 2011.
- K. T. Christensen. The influence of peak-locking errors on turbulence statistics computed from PIV ensembles. *Exp. Fluids*, 36:484–497, 2004.
- Y. S. Kwon, J. Philip, C. M. de Silva, N. Hutchins, and J. P. Monty. The quiescent core of turbulent channel flow. *J. Fluid Mech.*, 751:228–254, 2014.
- H. Makita. Realization of a large-scale turbulence field in a small wind tunnel. *Fluid Dyn. Res.*, 8:53–64, 1991.
- C. D. Meinhart and R. J. Adrian. On the existence of uniform momentum zones in a turbulent boundary layer. *Phys. Fluids*, 7(4):694–696, 1995.
- G. I. Roth and J. Katz. Five techniques for increasing the speed and accuracy of PIV interrogation. *Meas. Sci. Technol.*, 12:238–245, 2001.
- A. Sciacchitano, D. R. Neal, B. L. Smith, S. O. Warner, P. P. Vlachos, B. Wieneke, and F. Scarano. Collaborative framework for PIV uncertainty quantification: comparative assessment methods. *Meas. Sci. Technol.*, 26(7):074004, 2015.
- B. Wieneke. PIV uncertainty quantification from correlation statistics. *Meas. Sci. Technol.*, 26(7):074002, 2015.