Review of Scientific Instruments

RESEARCH ARTICLE | JULY 07 2023

Testing image-velocimetry methods for turbulence diagnostics

Y. W. Enters ; S. Thomas ; M. Hill ; I. Cziegler

Rev. Sci. Instrum. 94, 075101 (2023)
https://doi.org/10.1063/5.0133453

Review of Scientific Instruments | ARTICLE | pubs.aip.org/aip/rsi

Testing image-velocimetry methods for turbulence diagnostics

Cite as: Rev. Sci. Instrum. 94, 075101 (2023); doi: 10.1063/5.0133453
Submitted: 3 November 2022 • Accepted: 27 May 2023 •
Published Online: 7 July 2023

View Online Export Citation CrossMark

Y. W. Enters,[1,2,a)] S. Thomas,[1] M. Hill,[1] and I. Cziegler[1]

AFFILIATIONS

[1]York Plasma Institute, School of Physics, Engineering and Technology, University of York, York, United Kingdom
[2]Culham Centre for Fusion Energy, Culham Science Centre, Oxfordshire, United Kingdom

[a)]Author to whom correspondence should be addressed: ye525@york.ac.uk

ABSTRACT

Two image-based velocity-inference techniques, cross-correlation time-delay estimation (CCTDE) and dynamic time warping (DTW), were tested. These techniques are conventionally used in the study of plasma dynamics, but they can be applied to any data where features propagate across the image field-of-view. Differences between the techniques were investigated, which showed that the shortcomings of each technique are complemented well by the strengths of the other. Thus, the techniques should be used in conjunction with each other for optimal velocimetry. For ease of use, an example workflow that applies the results in this paper to experimental measurements is provided for both techniques. The findings were based on a thorough analysis of the uncertainties for both techniques. Specifically, the accuracy and precision associated with inferred velocity fields were systematically tested using synthetic data. Novel findings are presented that strongly improve the performance of both techniques, some of which are as follows: CCTDE was able to operate accurately under most conditions with an inference frequency as short as 1 per 32 frames, as opposed to the typical 1 per ≥256 frames used in the literature; an underlying pattern in CCTDE accuracy depending on the magnitude of the underlying velocity was found; spurious velocities due to the barber pole illusion can now be predicted *prior* to CCTDE velocimetry through a simple analysis; DTW was more robust against the barber pole illusion than CCTDE; DTW performance with sheared flows was tested; DTW was able to reliably infer accurate flow fields from data with as low as 8×8 spatial channels; and however, if the flow direction was unknown prior to DTW analysis, DTW could not reliably infer any velocities.

I. INTRODUCTION

Turbulence is a notoriously complicated and pervasive phenomenon in physics. Experimentally inferring turbulent flow fields is a key component in improving our understanding of turbulent systems. This paper focuses on testing techniques that infer velocity fields without relying on introducing tracer particles to the system. The context of this paper is around tokamak plasmas, and the tested velocimetry methods are primarily used in this field. Nevertheless, the applications of this paper can extend to any data in which features propagate across the field-of-view. Some examples include artificial speech recognition,[1] low Reynolds number particle image velocimetry (PIV),[2] and 3D lidar velocimetry.[3]

In tokamak plasmas, it is known that the transport levels across the confining field are enhanced compared to neoclassical predictions due to turbulent fluctuations.[4,5] On top of that, turbulence is a ubiquitous phenomenon in tokamak plasmas and has been linked to a multitude of related processes, such as filament evolution in the edge and SOL,[6] streamers,[7] and zonal flow evolution.[8]

One of the standard approaches of particle image velocimetry (PIV), especially at its inception, was to introduce high contrast tracer particles to enable visualization of the fluid's underlying velocity fields.[9] In the context of magnetic confinement fusion (MCF), no tracer particle exists, which can survive the extreme conditions in the plasma. Instead, fluctuations of the plasma properties, such as density or temperature fluctuations, can be used as tracer "particles" and tracked across a field of view to infer the underlying velocity fields. These fluctuations decorrelate over time, which can lead to complications if the decorrelation-timescale is comparable to or faster than the underlying velocity. In this paper, Taylor's hypothesis for frozen-in turbulence was used to assume that decorrelation effects were negligible. This can be done without loss of generality as known corrections can be introduced to account for decorrelation.[10]

In MCF plasmas, PIV is typically applied to density fluctuation data obtained from diagnostics, such as beam emission spectroscopy[11] or gas-puff imaging.[12] The majority of published literature is centered around two techniques: cross-correlation time-delay estimation (CCTDE) and dynamic time warping (DTW).[13–15] A fundamental difference between DTW and CCTDE is that DTW finds velocity fields based on variations in the spatial information in the images, whereas CCTDE relies on varying temporal information in the time-series. Note that DTW is otherwise known as "orthogonal dynamic programming" and CCTDE can be used to describe a host of different varieties of the same core technique. The specific implementation of CCTDE in this paper is discussed in Sec. III A. In theory, both of these techniques can infer highly accurate velocity fields, and some case-specific testing of the techniques does exist in current literature. However, extensive testing reported in a comprehensive fashion has been lacking, although it would enable the reliable use of the techniques for a broad range of experimental data scenarios.[13,14,16,17] In this paper, such tests were performed while varying the following parameters: the characteristic density-fluctuation spatial scale, tilt angle, characteristic in-plane velocities, the number of spatial channels of the diagnostic, and the signal-to-noise ratio. The following aims are defined for this research:

- to quantify the accuracy and precision of the measured velocity fields for both the CCTDE and DTW techniques for a broad range of conditions,
- to test both techniques to the limits of their operational parameter space,
- to provide examples of the reliable application of velocimetry techniques using the results in this paper, and
- to compare the techniques with each other under controlled conditions.

Section II discusses the methods used to generate the synthetic density fluctuation data and how the velocity fields were imposed on the density fluctuations. Sections III and IV introduce and test the CCTDE and DTW techniques, respectively. These sections outline the methods used to test their uncertainties, present the results, and discuss the findings. In Sec. V, DTW and CCTDE were compared to each other according to the aims listed above. Section VI presents an example application of the results. An example workflow is proposed, which minimizes the drawbacks from both methods. Conclusions are drawn in Sec. VII.

II. SYNTHETIC DENSITY FLUCTUATION DATA

Synthetic data were generated with the intent of representing the typical structures observed in turbulence diagnostics, such as GPI[12] and BES.[11] These diagnostics measure a time-series of two-dimensional images. The same synthetic dataset was analyzed by both CCTDE and DTW so that the two techniques could be straightforwardly compared.

Two types of fluctuation structures were generated. The first type contained isolated density-fluctuations, often observed in plasmas near marginal stability. This will be referred to as "isolated density feature" (IDF) data. In this paper, these isolated density features are sometimes referred to as "blobs." It should be noted, however, that the term "blob" in this paper does not specifically refer to the scrape-off layer filaments observed in MCF plasmas.[18] The second type contained "turbulent density fields" (TDF), representing the density structures observed in fully developed turbulent plasmas.

The three main user-controlled inputs for the synthetic data generation were the in-plane velocity field imposed on the density structures, the characteristic spatial scales associated with the structures themselves, and the signal-to-noise ratio of the data. Additionally, the angle of the fluctuations could be altered to investigate the effect of the barber pole illusion (see Sec. III D) and separately, sheared flows could be introduced in the TDF data. All variables in this paper are given in terms of machine units. For example, velocities are given in pixels per frame and length-scales are given in pixels. This ensures that the results can be applied to a wide range of diagnostics.

A. Isolated density feature data

Synthetic data were generated to represent isolated density features propagating through the field-of-view. Images of the IDF were generated using a two-dimensional Gaussian function,

$$z(x,y) = A\,\exp\left(-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2}\right), \tag{1}$$

where A and (x_0, y_0) are the amplitude and centroid location of the blob, respectively, and (σ_x, σ_y) are the standard deviations of the Gaussian blob shape, which could be tuned to vary the spatial size of the blobs.

However, true Gaussian features of the form in Eq. (1) are infinite in spatial extent, which is non-physical and undesirable for testing purposes. In order to localize the density features, 25% of the maximum intensity was subtracted from all images, and all resultant negative intensities were set to zero. Finally, the intensities were multiplied by 4/3 such that the original maximum amplitude was recovered. The result was a modified 2D Gaussian shape with a spatial extent that could be calculated in each direction using the relation

$$\Delta\lambda_{x,y} = 2\sqrt{2\ln(1/0.25)}\sigma_{x,y} = 3.33\sigma_{x,y}, \tag{2}$$

where $\Delta\lambda_{x,y}$ corresponds to the full size of the blob in pixels in the x- or y-direction. Throughout this paper, it is also often referred to as the blob y-size or x-size. An example of a synthetic IDF image can be found in Fig. 1(a).

Synthetic time-series were generated by creating images of these blobs and moving the blob centroid according to the imposed velocity with each frame. All images spanned 128 pixels × 128 pixels, and the blob centroid location was initialized below the image field-of-view (negative y-direction). The starting location was restricted such that blobs did not clip the edges of the x-side of the field-of-view. Only one blob was present in the images at any one time, and the time-series generation was terminated upon the blob completely leaving the field-of-view. Noise was optionally generated and added

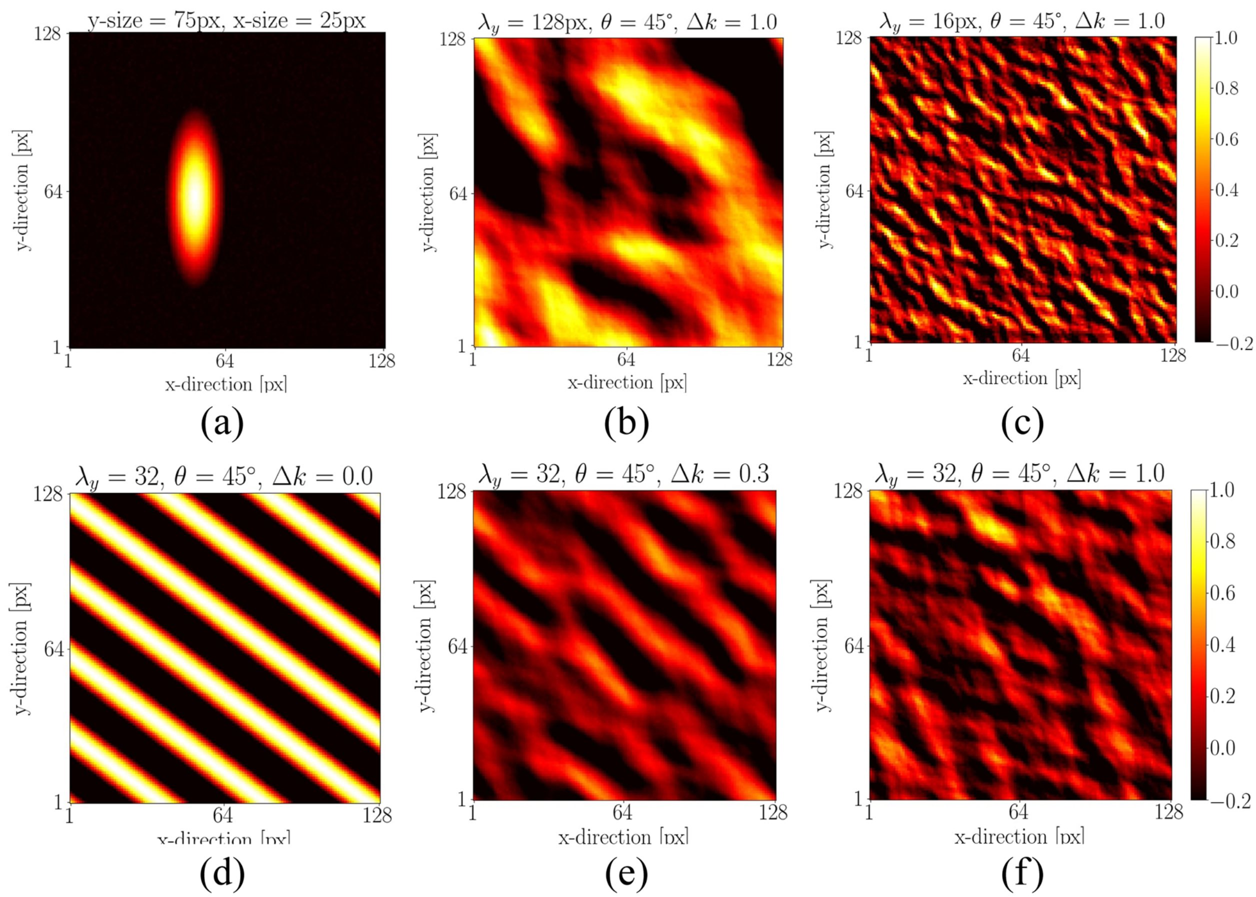

FIG. 1. Example images of synthetic data showing an isolated blob (a) and turbulent density fields with large (b) and small (c) spatial scales. (d)–(f) show the effect of varying Δk. All images were generated with SNR = 100, and color bars were normalized to maximum intensity in time-series. (a) IDF example. (b) Large TDF fluctuations. (c) Small TDF fluctuations. (d) $\Delta k = 0$. (e) Increasing Δk. (f) Increasing Δk further.

to each frame in the form of normally distributed, pixel-size noise in order to mimic diagnostic electronic noise. The signal-to-noise ratio (SNR) was defined as the ratio between the maximum blob amplitude and the rms of the noise added,

$$\mathrm{SNR_{blob}} = \frac{\mathrm{A}}{\mathrm{rms(noise)}}. \quad (3)$$

Synthetic data were generated for a range of blob y-sizes [1–100 px], blob velocities in the y-direction [0.1–60 px/frame], and signal-to-noise ratios [1 − Inf]. All data were generated using an arbitrary blob x-size of 25 px, which was large enough to be registered by the velocimetry techniques but small enough to have little clipping. The x-size did not affect velocimetry performance in initial tests (unless it was close to the pixel size or frame size). The velocity of the blobs was set to zero in the x-direction in all tests performed in this paper. The lack of generality here is discussed further in Sec. III E.

B. Turbulent density fields

Density fields observed from turbulence diagnostics often display a complex structure that cannot be directly reproduced through a simple function like in Sec. II A, Eq. (1). However, when represented in wavenumber space, turbulent density fluctuations can often be approximated as a broad, singly peaked function, such as a Gaussian or a Lorentzian.[14] Through this observation, the distributions were generated first in reciprocal space and then inverse Fourier transformed, producing the real-space TDF images.

Specifically, arrays with a Lorentzian distribution were generated of the form

$$P(k_x, k_y) = \frac{\Delta k^2}{(k_x - k_{x0})^2 + (k_y - k_{y0})^2 + \Delta k^2}, \quad (4)$$

where P is the distribution amplitude, (k_{x0}, k_{y0}) was the centroid location of the Lorentzian, and Δk is the function width. All elements were given a randomized phase, and the real component of the inverse-Fourier transform was taken as the final TDF image.

The Lorentzian centroid location, (k_{x0}, k_{y0}), could be altered to change the spatial size and angle of the density features, as seen in Figs. 1(b) and 1(c). All k_{y0} values in this paper were normalized to correspond to the number of full wavelengths in the y-direction per width of the image frame (wavelengths per 128 px). For ease of interpretation, the real-space length-scale of the fluctuations was also defined: $\lambda_{y0} = 128\ \text{px}/k_{y0}$. k_{x0} was not varied independently and instead was normalized to k_{y0}. The tilt angle of the density features, clockwise from horizontal, was then defined as $\theta = \text{Arctan}(k_{x0}/k_{y0})$.

The Lorentzian width, Δk, could be increased from zero to increase how "broken up" the structures appeared in the real-space image [see Figs. 1(e) and 1(f)]. Δk was also normalized to k_{y0}, varying from 0 to 2 in increments of 0.3, where 1.3 was the typical value observed in density fluctuation diagnostics.[14]

In order to generate a time-series with an imposed velocity field, the density fields were generated to be spatially larger than the final image dimensions. This allowed entire columns of pixels to be rigidly shifted up and down according to the imposed velocity field, which took the form

$$\boldsymbol{v}_{\text{imposed}} = v_0\ \hat{\boldsymbol{y}} + v_1\ \cos\left(k_{v,y}\ x\right)\ \hat{\boldsymbol{y}}, \tag{5}$$

where $k_{v,y}$ was the wavenumber of the velocity field sinusoid. All imposed velocity fields pointed purely in the y-direction and $k_{v,y}$ was varied from 1 to 8 wavelengths per 128 px. This method, although more computationally expensive, avoids the use of periodic boundary conditions, which could cause aliasing in the velocimetry analysis. The time-series were generated with image dimensions of 128 px × 128 px and were 512 frames long. Velocity fields with $v_0 = 0.1$–60 px/frame and $v_1 = 1$–15 px/frame were used. Furthermore, normally distributed, pixel-size noise was added on top of the images with $\text{SNR}_{\text{rms}} = 1$ – infinity and was generated individually for each image frame,

$$\text{SNR}_{\text{rms}} = \frac{\text{rms(signal)}}{\text{rms(noise)}}. \tag{6}$$

The SNR was defined as the ratio between the signal rms and the noise rms. Note that SNR_{blob} and SNR_{rms} represent significantly different definitions, often separated by more than an order of magnitude. Conversion factors vary solely on the blob size and can be found in Fig. 14.

III. CROSS-CORRELATION TIME DELAY ESTIMATION

A. A review of the technique

Cross-correlation time-delay estimation is a technique that can be used to estimate a velocity between two spatially separated time-signals. It is a specific implementation of the general cross-correlation-based PIV approach.[19] The technique is based on finding the time-delay, τ_m, at which the cross-correlation between the two signals is maximized. If the maximum amplitude of the cross-correlation function (CCF) is close to one, it is assumed that identical fluctuations are present in both signals and that they have traveled between the two spatial locations in a time τ_m. By repeating this process for a range of spatial locations, a velocity field can be constructed. In this paper, the two-point CCTDE was investigated. This version of CCTDE could be considered as fundamental to most other CCTDE variations. It should therefore be straightforward to extrapolate the results in this paper to the more modern and elaborate line method[20] and hybrid method.[16] The procedure of the two-point technique is outlined as follows:

- Consider a time-series of spatially resolved images.
- Two spatially separated pixels were selected, and their time-signals were cross-correlated using the function

$$\text{CC}(\tau) = \begin{cases} \dfrac{N-1}{N+\tau-1} \dfrac{\sum_{n=1}^{N-\tau}[f(n+\tau)-\bar{f}][g(n)-\bar{g}]}{\sqrt{\sum_{n=1}^{N}[f(n)-\bar{f}]^2[g(n)-\bar{g}]^2}}, & \tau < 0, \\ \dfrac{N-1}{N-\tau-1} \dfrac{\sum_{n=1}^{N-\tau}[f(n)-\bar{f}][g(n+\tau)-\bar{g}]}{\sqrt{\sum_{n=1}^{N}[f(n)-\bar{f}]^2[g(n)-\bar{g}]^2}}, & \tau \geq 0, \end{cases} \tag{7}$$

  where N is the length of the time-series and τ is the time-delay between the two signals, f and g. The bar denotes the mean of f and g, and the time-delay was given a range from $-N$ to N frames in all tests. The prefactor in the expressions is necessary for unbiased estimators by normalizing the time-series based on their length, which varies with τ.
- The time-delay, τ_m, at which the CCF peak occurred was determined. If the correlation peak was above 0.5, a velocity was then inferred by taking $v = \Delta\ell/\tau_m$, where $\Delta\ell$ is the spatial separation between the two signals.
- This procedure was repeated for all pixel pairs in both orthogonal directions to produce two spatially resolved velocity fields.

An important user-defined parameter was the characteristic separation distance, $\Delta\ell$, which is the distance between chosen pixels to be analyzed. Three pixels were required per velocity measurement. The first pixel, which was designated to have its velocity inferred, was called the reference. The other two pixels were defined to be located

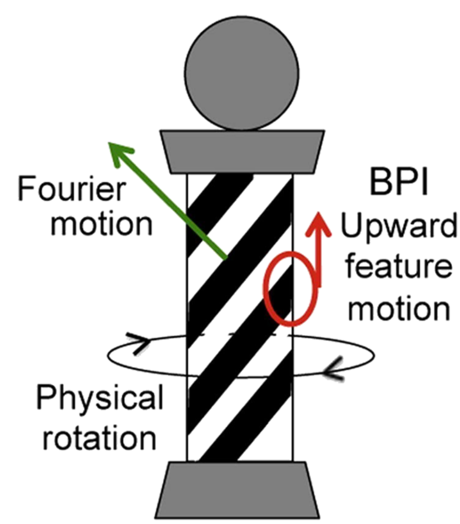

FIG. 2. Apparent up/down motion of barber pole as it spins. Reproduced with permission from Sun *et al.*, Vision Res. **111**, 43 (2015). Copyright 2015 Author(s), licensed under a Creative Commons Attribution 4.0 License.[30]

in the positive x- and y-directions from the reference by $\Delta\ell$ px. This CCTDE method is referred to as the "two-point" method and will be the primary focus for testing in this section. The effects of varying this separation distance, $\Delta\ell$, and the length of the time-series, N, are discussed in Sec. III C.

One notable source of spurious velocity measurements is known as "the barber pole illusion," after the apparent up/down motion of a barber's pole as it spins,[21,22] as can be seen in Fig. 2. This effect can also be observed typically when large, tilted density features propagate through the frame. Much like a barber pole, tilted density features may appear to be moving in a different direction than their true underlying velocity. It is difficult and indeed sometimes impossible in these cases for any velocimetry method to distinguish between apparent motion and true motion. The extent to which CCTDE was found to be susceptible to the barber pole illusion is investigated in Sec. III D.

B. Analysis of the measured velocity fields

To compare results, the velocity fields inferred by CCTDE were condensed into metrics assessing the overall accuracy and precision. First, the "percentage velocity-deviation field," Δv_{meas}, of the measured velocity field, v_{meas}, from the imposed velocity field, v_{imp}, was calculated,

$$\Delta v_{\text{meas}} = 100\% \cdot \left(\frac{v_{\text{meas}} - v_{\text{imp}}}{v_{\text{imp}}}\right). \tag{8}$$

The metric for accuracy used in, e.g., Fig. 3 was then taken to be the mean of Δv_{meas}. In the case of isolated density features, the average was performed over the area of the blob. In the turbulent density fluctuation case, the average was performed over the entire field. Similarly, the precision was quantified by the root-mean-square of Δv_{meas}, with the same areas as above.

C. Results—isolated density features

This section assesses the performance of CCTDE when presented with data containing isolated density features. The velocimetry accuracy parallel to the imposed velocity was tested. Density features with spatial scales ranging from 1 px to 100 px were investigated. Imposed velocities were varied from 1/10 px/frame to 30 px/frame. Additionally, the effect of $\Delta\ell$, length of time-series N, noise levels, and noise filtering is discussed.

As seen in Fig. 3, negligible dependence of the velocity measurement accuracy was found on the spatial size of the blobs. This was expected from the CCTDE method especially in the case of negligible noise. One exception to this observation can be seen in Fig. 3(b) at blob y-size <10 px and $0.5 < v_0 < 1.0\ \Delta\ell$/frame (at $\Delta\ell = 20$ px). In this region, the blobs were small and fast enough to entirely skip over one of the measurement locations, and thus, no accurate velocity could be measured. Additionally, varying the length of the time-series, N, was found to have no effect on the velocimetry accuracy. This means that accuracy does not need to be considered with single blobs when setting N.

Figure 3 also shows that the accuracy with which CCTDE measures the velocity of the blobs is strongly dependent on the underlying velocity. This dependence is expected from the two-point CCTDE method and is further discussed in Sec. III E. Initial tests showed that using a "line" method instead—which is equivalent to performing the two-point method simultaneously for a range of $\Delta\ell$ and selecting the measurement with the highest correlation—relieved the strong accuracy dependence on imposed velocity magnitude. Despite this advantage of the line method, further testing was outside the scope of this paper, but it is further discussed in Sec. III E.

The effect of varying $\Delta\ell$ was investigated next. As might be expected, increasing $\Delta\ell$ enables the accurate measurement of faster velocities. Additionally, the accurately measurable space expanded near slow velocities (~1 px/frame), as can be seen in Fig. 3. Despite the aforementioned benefits, maximizing $\Delta\ell$ may not always be beneficial in reality, which is discussed further in Sec. III E. Another noteworthy observation on varying $\Delta\ell$ is that the accuracy follows a pattern which is constant with normalized velocity, as shown in Figs. 3(a) and 3(b). Once again this was expected and is further discussed in Sec. III E.

Results show an often overlooked fact: the finite length of the time-series imposes a limit on the *minimum velocity* that can be measured. This threshold is given analytically as

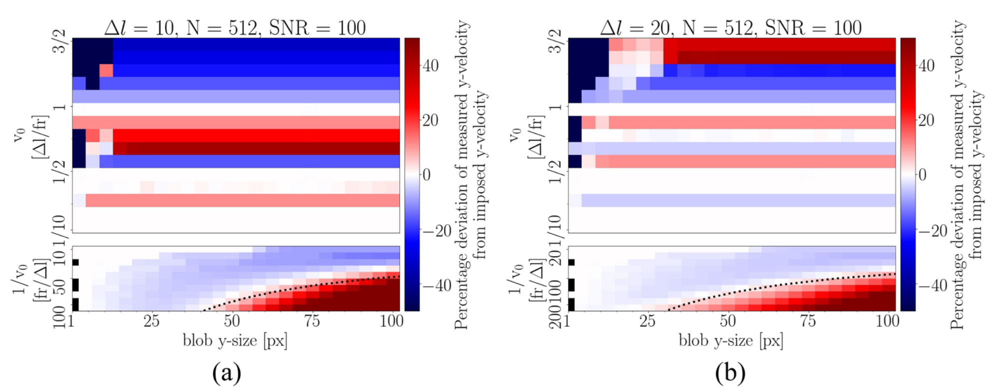

FIG. 3. The accuracy of CCTDE with isolated density features. Velocities are normalized to $\Delta\ell$. The bottom subplots show velocities below 1 px/frame. The dashed line denotes the predicted minimum measurable velocity that is imposed due to signal clipping by the finite length of the time-series, N. Accurate velocity fields (white) were typically measured with a standard deviation of <1%; precision was not a proxy for accuracy. (a) $\Delta\ell = 10$. (b) $\Delta\ell = 20$.

$$v_{\min} = \frac{\lambda_{y0} + \Delta\ell}{N}, \tag{9}$$

where N is the length of the time-series, λ_{y0} is the blob y-size, and $v_{\min}$ is the minimum velocity in the y-direction that can be measured at a given combination of N, $\Delta\ell$, and λ_{y0}. This expression was overplotted as an example in Fig. 3 and found to accurately predict the minimum measurable velocity in all cases. Threshold was found to correspond approximately to the 25% velocity deviation mark. Additionally, the precision of the velocity inference decreased with decreasing N.

The effect of increasing noise levels on the performance of CCTDE is summarized in Fig. 4. It was observed that the accuracy of CCTDE remained largely unaffected by noise from $\mathrm{SNR}_{\mathrm{blob}}$ = inf to 10. Accuracy in this range looked like Fig. 3(b), with only marginal increases of standard deviation observed by $\mathrm{SNR}_{\mathrm{blob}} = 10$. Decreasing $\mathrm{SNR}_{\mathrm{blob}}$ from 10 to 4 showed significant signs of accuracy degradation by the noise, which is shown in Fig. 4(a) in the form of unreliable accuracies and a shrinking reliably accurate parameter space. Decreasing the $\mathrm{SNR}_{\mathrm{blob}}$ further from 4 to 1 caused the accurately measurable parameter space to shrink steadily and the measured velocity fields showed increased standard deviation. By $\mathrm{SNR}_{\mathrm{blob}} = 1$, no accurate velocities could be inferred, as shown in Fig. 4(b). Importantly, the standard deviation of the inferred velocity was *not necessarily a good predictor* for the accuracy.

Next, the ability to recover accurate CCTDE measurements from noisy data using frequency filtering was assessed. The synthetic data used for Fig. 4 were passed through a two-way low-pass Butterworth filter prior to CCTDE analysis. A cutoff frequency of 0.3 times the Nyquist frequency was chosen due to effective filtering of the noise while leaving the signal relatively unaffected. These initial tests showed marginal improvements in accuracy; the $\mathrm{SNR}_{\mathrm{blob}} = 4$ case after filtering showed accuracies comparable to the unfiltered $\mathrm{SNR}_{\mathrm{blob}} = 100$ case, although filtering data with $\mathrm{SNR}_{\mathrm{blob}} = 3$ showed only partial recovery of the accuracy (not shown). No improvement was seen for $\mathrm{SNR}_{\mathrm{blob}} = 1$. Whether optimization of the noise filtering techniques could be used to fully recover the accuracy for $\mathrm{SNR}_{\mathrm{blob}} = 3$ (and lower) could not be determined using these initial tests.

It was also found in all tests in this subsection that the blob velocity orthogonal to the direction of blob propagation was accurately measured to be zero on average for imposed velocities >1/20 $\Delta\ell$/frame. Standard deviations of up to 20% relative to the measured parallel velocity were observed in this range. Orthogonal velocity inferences at $v_0 < 1/20\ \Delta\ell$/frame were found to be susceptible to high statistical variation of the measurements. The standard deviation often exceeded that of the parallel velocity measurements, thus significantly reducing the reliability of the overall measurements in this region of velocity space.

D. Results—turbulent density fluctuations

This section assesses the performance of CCTDE when presented with TDF data. The characteristic spatial scale, λ_{y0}, was varied from 1 px to 100 px, and the imposed velocity was varied from 1/10 px/frame to 30 px/frame. The angle of the density features was given a range from 0° to 75° clockwise from horizontal. Δk was given a range from $0k_{y0}$ to $2k_{y0}$ in $0.3k_{y0}$ increments. $\Delta\ell$ was varied from 1 px to 20 px.

As expected, varying $\Delta\ell$, $\mathrm{SNR}_{\mathrm{rms}}$, and N had an effect comparable to what was found in Sec. III C. In summary, decreasing N was known to reduce precision and impose a minimum measurable velocity. The corresponding expression defined in Eq. (9) held for the turbulent density data. $\mathrm{SNR}_{\mathrm{rms}}$ was found to have a negligible effect on accuracy for $\mathrm{SNR}_{\mathrm{rms}} > 1$, although reductions in precision were observed.

It was found that the barber pole illusion had a negligible impact on the measurement accuracy parallel to the flow direction. Spurious perpendicular velocity measurements were widely observed when $\Delta k \leq 0.3k_{y0}$. When Δk is this low, density features often have a spatial extent comparable to the image field-of-view. Furthermore, negative velocities were found, which were due to a combination of barber poling and aliasing. A subtle example of aliasing can be seen as the blue regions in Fig. 5(a). These spurious velocities could not consistently be predicted, and experimental CCTDE analysis in the region of $\Delta k \leq 0.3k_{y0}$ is unreliable without further in-depth testing.

In the region of $\Delta k \geq 0.6k_{y0}$, it was found that there is typically a well-defined threshold for when the barber pole illusion becomes

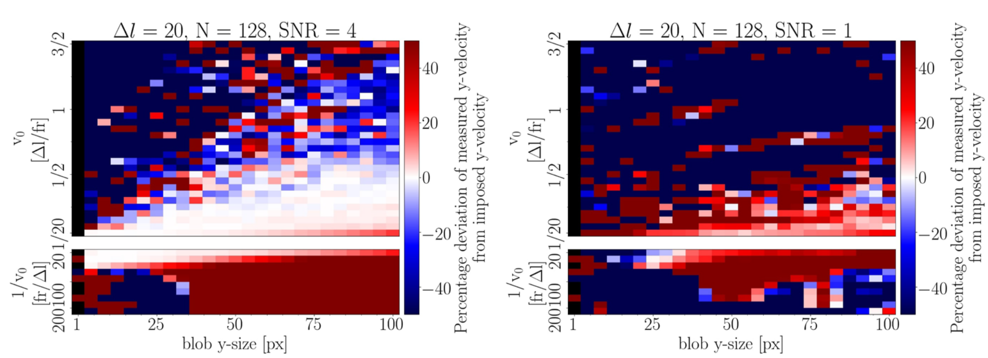

(a) (b)

FIG. 4. The accuracy of CCTDE with isolated density features. The $\mathrm{SNR}_{\mathrm{blob}}$ was varied to assess the dependence of method accuracy on noise. Noise had little effect on accuracy for $\mathrm{SNR}_{\mathrm{blob}} > 10$ (not shown). The standard deviation of the measured velocity fields varied from <20% in (a) to >100% in (b). (a) SNR = 4. (b) SNR = 1.

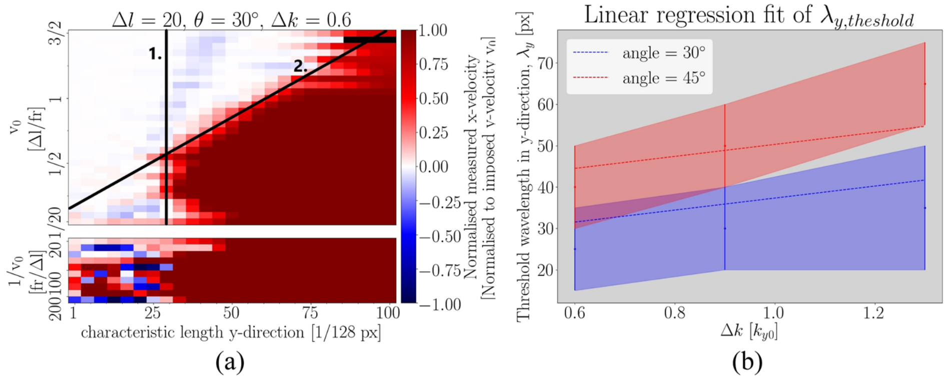

FIG. 5. Defining a threshold at which barber poling becomes significant: (a) thresholds drawn approximately at 25% deviation from v_0; (b) examples from the multivariate linear regression of threshold line 1. Data-points denote empirical measurements with estimated error bars, and the dashed lines denote the regression equation fits. *R*-squared was 0.9. (a) Barber pole approximate thresholds. (b) Barber pole threshold regression fit.

significant. An example can be found in Fig. 5(a), where two distinct lines are drawn approximately at the 25% velocity deviation mark. Threshold 1 in Fig. 5(a) was found to depend solely on the spatial parameters: $\Delta k, \Delta \ell$, θ, and λ_{y0}, although no analytical expression for the threshold could be determined from first principles. Instead, a multivariate linear regression approach was used in an attempt to define an empirical expression for the threshold.

Estimates of the characteristic length, λ_{y0}, at which threshold line 1 occurred had to be taken in order to perform a regression. Angles of 30° and 45°, $\Delta k = 0.6$–$1.3, k_{y0}$ and $\Delta \ell = 5$–20 px were used for estimation of the λ_{y0} thresholds. The threshold was determined as follows: the λ_{y0} at which the percentage velocity deviation reached 25% were traced across v_0. The five-point moving average was taken thereof, which consistently resulted in a trace with a bi-linear form [approximately corresponding to lines one and two in Fig. 5(a)]. The change in gradient between the two linear regions was not sharp and an approximate gradient transition region was defined by visual inspection. The λ_{y0} threshold was defined as the lowest λ_{y0} of the gradient transition region [approximately equivalent to the intersection of lines 1 and 2 in Fig. 5(a)]. The error margin of the λ_{y0} threshold was defined to be equal to the half-width of the gradient transition region.

A weighted multivariate linear regression was performed with the determined λ_{y0} thresholds as the dependent variable and $\Delta k, \Delta \ell$, and $\tan(\theta)$ as the independent variables (which individually showed approximate linear dependence prior to regression). The following expression was constructed from the regression to approximate the λ_{y0} threshold:

$$\lambda_{y0,\text{threshold}} = 30 \cdot \tan(\theta) + 15 \cdot \Delta k + 2.2 \cdot \Delta \ell - 40. \tag{10}$$

For this regression of 24 datapoints, all coefficients were found with a standard error of ~15% and an R^2 of 0.9. Δk was normalized to k_{y0} and $\Delta \ell$ was given in px. A good prediction (within errors) was also found when extrapolated to parameters $\theta = 15°$–$60°$, $\Delta k = 0.6$–$2.0 k_{y0}$, and $\Delta \ell = 5$–20 px.

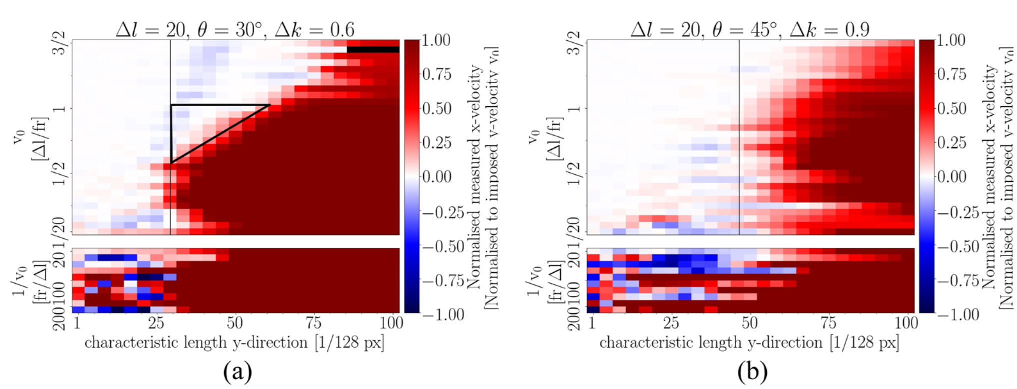

FIG. 6. Two example plots showing the λ_{y0} threshold as predicted by the regression equation. Fit within error margins was found for most (~90%) cases tested. Triangular region in (a) shows accurate area that can be gained by quantifying threshold 2. This area does not exist in (b). (a) BP predicted threshold. (b) BP predicted threshold.

With increasing λ_{y0}, the perpendicular velocity increases and saturates at a well-defined value. This saturated value could be calculated by the following expression:

$$v_\perp = v_\parallel / \tan(\theta). \tag{11}$$

The $v_\perp$ calculated in Eq. (11) is strictly in case that there is no underlying velocity in the x-direction. It has not been explicitly tested if the barber-pole-induced x-velocity is affected by true underlying x-velocities.

Threshold 2, shown in Fig. 5(a), was not investigated in detail because it would have little impact; quantifying threshold 2 would not add a significant area in parameter space that can be measured accurately with confidence. The potential additional area is exemplified by the small triangular region in Fig. 6(a), and no added benefit would be seen in Fig. 6(b).

E. CCTDE discussion and summary

1. Accuracy dependence on underlying velocity

The measurement accuracy of two-point CCTDE depends strongly on the underlying velocity of the fluctuations. The accuracy varied in a predictable pattern, which can be seen in Fig. 3. This pattern was theoretically expected from the two-point CCTDE method because only velocities equal to factors of $\Delta\ell$ can be measured accurately. In-between velocities will be approximated to the nearest factor of $\Delta\ell$. The salient issue is that the underlying velocity is unknown prior to CCTDE analysis, which complicates the prediction of the measurement accuracy. Next to this, *the standard deviation of the inferred velocity fields was not a good predictor of measurement accuracy*. This is why it is essential that two-point CCTDE be used in conjunction with other velocity estimation techniques.

2. A lack of precision may be beneficial

In initial tests, it was found that the line-CCTDE method accuracy did not show the same dependence on velocity. It may not be necessary, however, to employ this more computationally expensive method. This is due to a most welcome case where experimental noise provides an unexpected beneficial effect: turbulent velocity fluctuations and optical jitter cause fluctuations in the apparent velocity. This can lead to the inferred velocity "flip-flopping" between factors of $\Delta\ell$ throughout the velocity field. The result is an average velocity field with improved accuracy but at a reduced temporal resolution. This effect is especially pronounced for underlying velocities halfway between two factors of $\Delta\ell$, and the flip-flopping can also be encouraged by decreasing N and thereby reducing the measurement precision. In practice, it is hypothesized that this beneficial averaging is commonplace, and it is possibly the reason behind the lack of accuracy dependence on underlying velocity in previous literature.[13,16] That said, this effect cannot be fully relied upon and cross-comparison with other velocity estimates is advised.

3. Accuracy dependence on SNR

The extent to which noise degrades the performance of CCTDE was found to be negligible for $\mathrm{SNR_{rms}} > 1$, in both IDF and TDF data. These results appear to contradict previous literature (e.g., Ref. 13), which cites much stronger dependence on SNR, and define a measurement limit around SNR = 10. However, due to the lack of a rigorous description of the data generation parameters and ambiguity in the SNR definition in the aforementioned work, no meaningful comparisons could be made.

4. Choosing appropriate $\Delta\ell$

As may be expected, increasing $\Delta\ell$ increases the maximum accurately measurable velocity (up to 1 $\Delta\ell$/frame). It simultaneously expands the accurately measurable velocity space at low velocities near 1 px/frame (see Fig. 3). However, maximizing $\Delta\ell$ may not be desirable if the decorrelation timescale is significant compared to the underlying velocity. In this case, $\Delta\ell$ could be reduced to limit the impact of decorrelation. $\Delta\ell$ can be independently defined in different directions without significantly complicating the analysis because velocity inferences in different directions are independent of each other. In this paper, $\Delta\ell$ was kept symmetric in all tests.

5. The effect of reducing length of the time-series, N

It is desirable to decrease N for a number of reasons: the computational cost of the technique is reduced, it may result in beneficial averaging due to reduced precision, and the temporal frequency of the velocity inference is increased. The reduction in precision in CCTDE has been covered in existing literature.[23] One consideration that must be made is that the expected velocities do not lie below the minimum velocity limit defined in Eq. (9) for the chosen N. This equation can be applied to both IDF and TDF data, although some extra care needs to be taken in determining a representative λ_{y0} for TDF data. Assuming that the minimum velocity limit is not surpassed, it has been observed in this paper that CCTDE can be operated effectively with N as low as 32 in all cases tested with little effect on measurement accuracy, whereas previous literature applying CCTDE typically uses $N \geq 256$.[13]

6. A simplified method to avoid barber pole illusions

The prevalence of spurious velocity measurements due to the barber pole illusion was quantified in Sec. III D. In this research, it can be predicted if the barber pole illusion has a significant impact on velocimetry *prior* to CCTDE analysis. The prediction requires estimation of the spatial parameters associated with fluctuations in the data, Δk and λ_{y0}. By referring to Eq. (10), spurious velocities can typically also be avoided by increasing $\Delta\ell$. An example of this simple check is shown in Sec. VI. These results are in contrast to previous literature that focused on correcting spurious velocities post-velocimetry.[24] This correction process is typically rather laborious and requires assumption of the shape of the underlying structures.

7. Precision in the low-velocity regime

It was found that for imposed velocities, $v_0 < 1$ px/frame, the statistical variation of the velocity measurements becomes significant. The standard deviation becomes comparable to the average inferred velocity. This is why extra care is advised in ensuring statistical convergence in this low velocity regime.

IV. DYNAMIC TIME WARPING

A. A review of the technique

Dynamic time warping is a technique that falls under the broader "optical flow" approach in PIV, where brightness conservation along flow trajectories is assumed. A brief review of the broader optical flow research is given in the introduction of Cai *et al.*[25] The aim of DTW is to find an optimized displacement field from one image to another.[15,17,26,27] The spatial transformation is found through an iterative process that calculates a displacement field according to the largest intensity features first and then makes increasingly small corrections in subsequent iterations. The displacement field optimization in each iteration is based on the minimization of the intensity difference between the two images. DTW is considered to be a promising technique in image velocimetry because it is theoretically able to find accurate displacement fields even in ambiguous scenarios, such as turbulent flow fields, and because it promises to deliver a velocity field at the diagnostic sampling rate.

The operation of the algorithm is explained in detail in Quenot's 1998 paper[15] and will briefly be summarized here. Two images for which an optimized transformation is to be found are loaded into the DTW algorithm. The two images are then divided into strips that overlap each other by half in the short direction, where the number of pixels along the short direction is known as the strip width. The optimized displacements were found along the long direction, or "slicing direction," of the strips. The optimization is based on the minimization of the intensity difference between the two strips.[15] Splitting the image into strips introduces a natural ordering of the pixels and imposes a continuity constraint on the displacement search. By re-combining the strips into the original images, a full displacement field is built up with pixel displacements *along the slicing direction* of the strips. This process is repeated in the orthogonal direction, and a 2D displacement field is inferred. In this first iteration, the algorithm has inferred a displacement field weighted toward intensity features with a spatial size comparable to the first strip width. Subsequent iterations are used to make increasingly small corrections to the displacement field by using increasingly small strip widths. In this paper, the strip width was reduced by a factor of $\sqrt{2}$ in each iteration. This is not necessarily an optimized value but follows convention used by Quenot *et al.*[15] The exact code used throughout this paper can be found in version 1.0.0 of S. Thomas' GitLab repository.[28]

The performance of DTW was tested for a range of synthetic data input parameters, such as spatial scale of the density fluctuations, underlying velocity of the density fluctuations, signal-to-noise ratio, shear flow of the imposed velocity field, and the number of spatial channels in the data. The impact of some of these parameters have been investigated in existing literature,[14,17] which are used as a point of comparison in this paper.

B. Results—isolated density features

This section assesses the performance of DTW when presented with data containing isolated density features. The accuracy and precision of the inferred velocity fields were determined in the direction parallel to- and perpendicular to the imposed velocity. Fluctuation spatial scales were varied from 1 px to 100 px, and imposed velocities were varied from 0.1 px/frame to 60 px/frame. SNR_{blob} was varied from 1 to infinity. DTW was typically operated with seven iterations with the following strip widths: [32, 22, 16, 12, 8, 6, 4] px, unless otherwise stated. The initial slicing direction was set parallel to the known imposed velocity, which was a crucial step and is discussed in more detail in Sec. IV D.

As can be seen in Fig. 7(a), DTW can measure velocities highly accurately across the vast majority of parameter space at SNR_{blob} = 100. The main exception occurs at $v_0 \leq 2$ px and blob sizes > 20 px, where there is consistent overestimation of the velocity. Other than this region of overestimation, the accuracy of DTW does not depend on the spatial size of the blobs or the underlying velocity. All perpendicular velocity measurements were zero. When SNR_{blob} is decreased, the accuracy of DTW rapidly degrades. The result can be seen in Fig. 7(b) where DTW cannot recover any velocities accurately at SNR_{blob} = 2. Upon further investigation, it was determined that the displacement field was often accurate after the first iteration of DTW, and subsequent iterations would distort the originally accurate field. This effect was due to an operational quirk of DTW in situations where there are areas with no signal and only noise, which is further discussed in Sec. IV D. Follow-up tests with only one DTW iteration at strip width = 32 px and varying

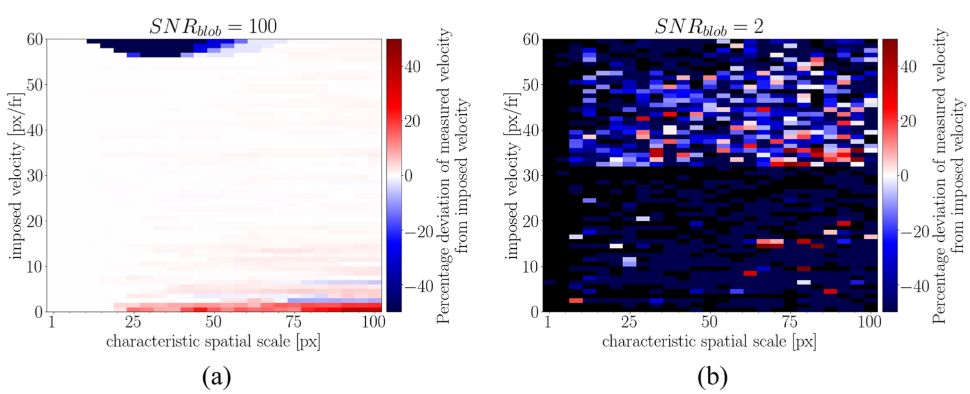

FIG. 7. The effect of pixel-size noise on DTW velocity inference accuracy. Input data were unsmoothed and seven DTW iterations were used. (a) SNR_{blob} = 100. (b) SNR_{blob} = 2 and 7 DTW iterations.

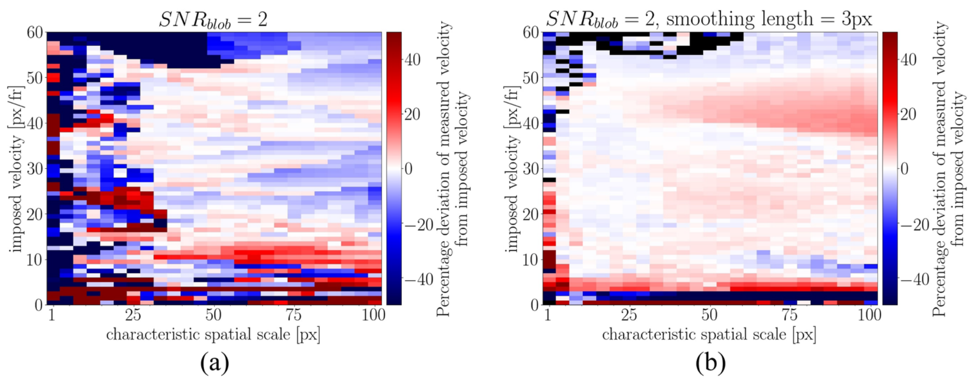

FIG. 8. Two approaches for recovering from noise in IDF data. The result of using only one iteration is shown in (a). The result of smoothing the images pre-analysis is shown in (b). (a) $\text{SNR}_{\text{blob}} = 2$, 1 DTW iteration. (b) $\text{SNR}_{\text{blob}} = 2$, pre-smoothed.

$\text{SNR}_{\text{blob}} = 100 - 1$ were performed. Highly improved accuracies were found with typically <10% velocity deviation, as can be seen in Fig. 8(a), although full recovery of accuracy could not be achieved using this approach.

The effectiveness of spatially smoothing out the noise, before DTW analysis, was investigated by passing the input data images through nested 1D convolution filters with a Gaussian kernel. The optimum Gaussian FWHM was found to lie around 1–3 px where the noise was smoothed out effectively while leaving the underlying signal relatively unaffected. The smoothed data were then analyzed using DTW. It can be seen in Fig. 8(b) that smoothing is also highly effective across the majority of parameter space in recovering the accuracy of the DTW algorithm such that most velocities could be measured to within 5% accuracy. At low velocities, $v_0 < 5$ px/fr, the velocity could not accurately be inferred, as can be seen in Fig. 8(b).

C. Results—turbulent density features

The velocity measurement accuracy and precision were determined in both the direction parallel to and perpendicular to the imposed velocity. The DTW operational parameters used were identical to those in Sec. IV B unless otherwise stated. The characteristic spatial scale, λ_{y0}, was varied from 1 px to 100 px, and the imposed velocity in the y-direction was varied from 1 px/frame to 60 px/frame. All the tests performed in this section used TDF data with the parameters $\Delta k = 1.3k_{y0}$ and $\theta = 45°$.

Initial tests at $\text{SNR}_{\text{rms}} = 100$ showed that DTW could accurately infer velocities across the entire range of spatial sizes and imposed velocities, much like the IDF results shown in Fig. 7(a). Upon incrementally decreasing SNR_{rms} down to 1, the accuracy and precision were found to generally decrease, although tests at $\text{SNR}_{\text{rms}} = 1$ still showed accurately measurable regions in parameter space, which can be seen in Fig. 9(a). The standard deviation of the velocity fields in Fig. 9(a) was typically around 10%. The problems associated with using numerous iterations that were observed in Sec. IV B were not reproduced with TDF data. Smoothing was applied using the same approach as described in Sec. IV B, which also showed substantial improvements to the measurement precision. At $\text{SNR}_{\text{rms}} = 1$ and a 2 px smoothing length, standard deviation was reduced to within 5% of the average. This revealed a consistent overestimation of velocities below 10 px/frame, as shown in Fig. 9(b).

The effect of reducing the number of spatial channels available in the input data was investigated. The original 128 × 128 channel images were down-resolved into new N_{ch} by N_{ch} channel images. This was done by splitting the original images into arrays of size Δd by Δd, where $\Delta d = 128/N_{ch}$. Thus, each channel in the new N_{ch} by N_{ch} image corresponded to one Δd by Δd array in the original image. The intensities of the channels in the N_{ch} by N_{ch} images were calculated by taking weighted averages of the respective Δd by Δd arrays. The weighted average was calculated using a 2D Gaussian kernel with a FWHM equal to Δd and was centered on the center of the Δd by Δd array. Exploratory tests showed negligible dependence on the shape of the averaging kernel. Nevertheless, the Gaussian shape was chosen to approximate the increased sensitivity in the center of the channels found in diagnostics such as BES.[29] Synthetic data with $N_{\text{ch}} = [4, 8, 16, 32, 64]$ were generated and pixel-size noise was re-introduced with $\text{SNR}_{\text{rms}} = [1, 2, 3, 4, 10, 100]$. Although DTW could technically run using data with down to 4 × 4 spatial channels, the accuracy of the velocity inferences was found to rapidly degrade with decreasing N_{ch} and SNR_{rms}.

The down-resolved data were re-interpolated into the original 128 × 128 channel grid using a bivariate cubic spline interpolation before DTW analysis.[14] At $\text{SNR}_{\text{rms}} = 100$, the accuracy was only marginally affected in the range $N_{\text{ch}} = 128$–16. Measured velocities typically deviated less than 10% of the imposed values. As can be seen in Fig. 10, the accuracy deteriorated strongly going from <10% deviation at $N_{\text{ch}} = 16$ to ≥20% deviation at $N_{\text{ch}} = 8$. Additionally, there was a marked decrease in the amount of reliable parameter space shown in Fig. 11. Decreasing N_{ch} also showed a strong decrease in the precision of the DTW velocity fields, which could be exacerbated by decreasing SNR_{rms}. The combined result of these two detrimental effects can be seen in Fig. 11(a). It shows unreliable velocity measurements across all parameter space at $N_{\text{ch}} = 8$ and $\text{SNR}_{\text{rms}} = 2$, which are not uncommon values observed in BES diagnostic measurements. However, upon averaging 32 measurements from successive pairs of frames, the precision could be improved to a standard deviation of <10%, and considerable regions

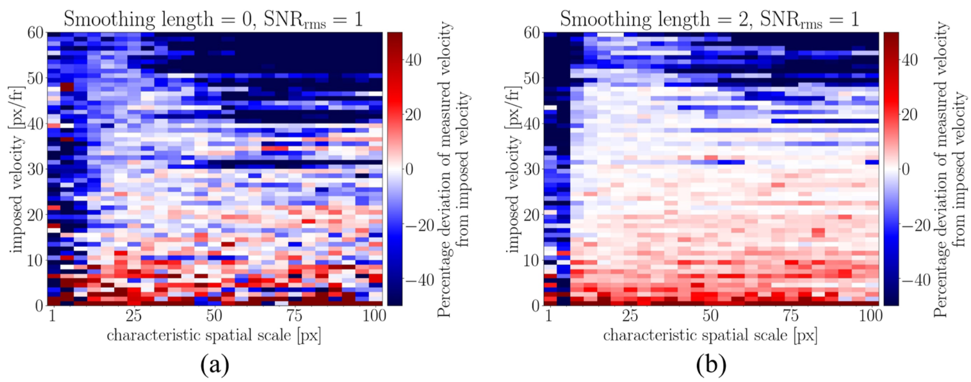

FIG. 9. The effect of smoothing TDF input data on DTW performance with noise. (a) No smoothing. (b) With smoothing.

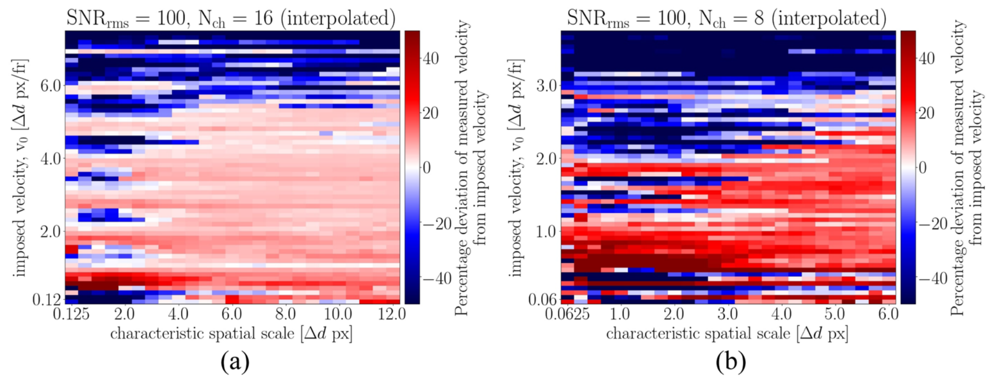

FIG. 10. The effect of reducing the number of available channels on DTW performance. Images contained N_{ch} by N_{ch} channels. All images were re-interpolated onto 128 × 128 channels before DTW analysis. (a) 16 × 16 channels. (b) 8 × 8 channels.

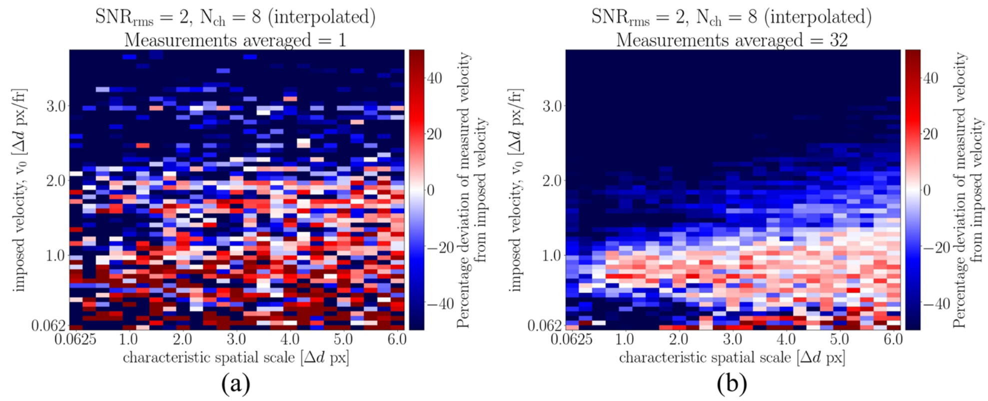

FIG. 11. Assessing DTW performance at N_{ch} = 8 under noisy conditions. Reduced precision shown in figure (a) is shown to converge upon averaging multiple measurements in (b). (a) 8 × 8 channels and SNR = 2, one measurement. (b) 8 × 8 channels and SNR = 2, 32 measurements averaged.

of parameter space were recovered in which accurate measurements could be made, as can be seen in Fig. 11(b). Averaging over multiple frames improved the precision of the velocity field inference, which revealed an area in parameter space where velocities could accurately be inferred. This area, shown in Fig. 11(b), consistently showed velocity inferences to within 10% of the true values. It is highlighted here that a high temporal frequency is one of the main desirable characteristics for DTW. Averaging significantly reduces DTW temporal resolution and can become comparable to CCTDE temporal resolution, as shown in Fig. 11(b).

For 128 × 128 channel images with $\Delta k \geq 0.3k_{y0}$, the inferred velocity was zero in the direction perpendicular to the imposed velocity. At $\Delta k < 0.3k_{y0}$, spurious velocities were found due to the barber pole illusion.

Exploratory tests were performed to assess the impact of sheared velocity fields on the performance of DTW. TDF data wereе used with a characteristic fluctuation spatial scale of 7 px. The velocity field parameters v_1 and $k_{v,y}$ [see Eq. (5)] were varied from 1 to 15 px/frame and from 1 to 8 wavelengths per 128 px, respectively. In the literature, the same $k_{v,y}$ range had been investigated, but v_1 was not varied in those tests.[14] Independently increasing the amplitude and the wavenumber showed detrimental effects on the accuracy and precision in both cases. The existence of thresholds for the onset of these detrimental effects was observed depending on v_1 and $k_{v,y}$. Thresholds were not quantified, but it was found that the accuracy and precision stayed constant when the maximum shear amplitude, $\partial_x(v_y)|_{\max} = v_1 k_{v,y}$, was kept constant. This observation held, while v_1 and $k_{v,y}$ were both varied.

D. DTW summary and discussion

1. Accuracy dependency on SNR

The accuracy of DTW decreases with increasing noise levels, as shown in Fig. 7. This was observed in 128 px × 128 px images for both IDF and TDF data. This loss of accuracy could largely be recovered through spatial smoothing of the images before velocimetry, as seen in Fig. 9. One exception is observed in Fig. 9(b), where a consistent overestimation of up to 20% is observed at low velocities, $v_0 < 10$ px/frame. The accuracy at characteristic spatial scales below 5 px could not be recovered using spatial smoothing. The precision of velocity fields also varied with signal-to-noise levels. Velocity fields with 1% standard deviation were observed at $\mathrm{SNR_{rms}} = 100$. The standard deviation increased to 10% at $\mathrm{SNR_{rms}} = 1$. The reduced precision due to noise could be mostly, but not completely, counteracted by spatial smoothing of the images before DTW analysis, as can be seen in Fig. 9.

2. An operational quirk with IDF data

This was seen in Sec. III C and thought to be due to regions in the images where no signal, and only noise, was present. With a lack of signal outside of the blob area, DTW would transform the images here according to the noise. This distortion outside the blob area would affect the displacement field inside the blob area due to continuity constraints in the DTW method. This is why care is advised for DTW velocimetry of IDF-like data. Two approaches in minimizing detrimental effects were shown in Fig. 8.

3. Accuracy dependency on the number of spatial channels

It is important to note that DTW accuracy decreases rapidly below $N_{\mathrm{ch}} = 32$, even when noise levels are negligible. The beneficial effect of interpolating the images onto a higher grid prior to DTW is ubiquitous *but at best marginal* for the accuracy of velocimetry. Despite this improvement, DTW increasingly and consistently overestimated velocities upon decreasing N_{ch}, which can be seen in Fig. 10. Spatial scales below approximately three times the channel size were unreliable, especially if v_0 was not an integer. The introduction of noise was investigated at $N_{\mathrm{ch}} = 8$. Decreasing $\mathrm{SNR_{rms}}$ rapidly decreased the precision of DTW. At $\mathrm{SNR_{rms}} = 2$, this effect could be considered fatal, as can be seen in Fig. 11(a). Upon averaging multiple consecutive inferences, the precision improved and accurate velocity inferences were revealed in Fig. 11(b). Specifically, a consistent overestimation of ~10% was found at $v_0 = 1$ px/frame. This was well below the measurement limit of SNR = 10 defined in previous literature but still relevant to realistic experimental situations.[14]

4. Sheared flow fields

Exploratory tests investigated the ability of DTW to infer sheared velocity fields by varying the velocity field parameters v_1 and $k_{v,y}$ seen in Eq. (5). It was found that both the accuracy and precision were reduced by increasing *the maximum shear* past a threshold. Interestingly, the accuracy and precision did not vary locally with local shear amplitude. Instead, the global maximum shear was found to be the parameter that governs the accuracy and precision. This is a result that generalizes previous tests by Kriete *et al.*,[14] who found that the shear flow wavenumber is an important parameter that affects DTW accuracy and precision. The reduction in accuracy presented itself as a reduction in amplitude of the measured velocity sinusoid, although the shape and wavelength were conserved, which was consistent with previous findings.[14]

5. The limited effect of the barber pole illusion

All measurements in the direction perpendicular to the imposed velocity were accurately inferred to be zero at $\Delta k \geq 0.3$. Spurious perpendicular velocities due to the barber pole illusion occurred at $\Delta k < 0.3$. The presence of spurious velocities are hypothesized when the density features are tilted, extend past the diagnostic field of view, and if spatial variations within the density features are negligible compared to the noise levels.

6. The main complication of using DTW

In all tests, the first slicing direction was set parallel to the direction of the imposed velocity. *Choosing the perpendicular direction instead was found to be catastrophic for DTW velocity inference.* The threshold at which misalignment between the initial slicing direction and flow direction becomes an issue was not tested, although increasing the first strip width was hypothesized to accommodate larger misalignments. Nevertheless, it is critical that the direction of the velocity field must be estimated prior to DTW analysis in a real experimental setting. Additionally, sub-pixel velocities could not be accurately inferred by DTW. This shortcoming can easily be circumnavigated by increasing the temporal spacing between images that are analyzed.

V. A COMPARISON OF THE TWO TECHNIQUES

A fundamental difference between DTW and CCTDE is that DTW finds velocity fields based on variations in the spatial information in the images, whereas CCTDE relies on varying temporal information in the time-series. This means that DTW is theoretically able to infer velocity fields with a frequency equal to the framerate of the diagnostic, whereas CCTDE typically operates with a measurement frequency at least an order of magnitude slower. This inherent drawback of CCTDE is in contrast to the fact that it is less reliant on spatial information, which may result in higher accuracy and precision than DTW when the number of spatial channels is reduced.

The CCTDE inference frequency has historically been at least 2 orders of magnitude slower than the diagnostic frame-rate.[13] In these investigations, it was found that CCTDE can be operated with $N = 32$ in most cases, which represents at least an order of magnitude improvement compared to previous literature. Additionally, DTW was found to require averaging of multiple subsequent inferences with noisy, low-spatial-channel data (e.g., $\mathrm{SNR}_{\mathrm{rms}} = 2, N_{\mathrm{ch}} = 8$). This reduced the effective velocity-inference frequency of DTW, and in some cases, the inference frequency could be comparable between the two techniques.

Both techniques were prone to inferring spurious velocities due to the barber pole illusion. Neither technique could reliably infer accurate velocity fields at $\Delta k < 0.3$. At $\Delta k \geq 0.3$, DTW typically inferred accurate velocity fields. CCTDE could also infer accurate fields in this region, assuming that it passed a simple check using Eq. (10).

Shear in the velocity fields negatively impacted the DTW accuracy once a threshold in the maximum local shear was surpassed. Shear did not affect two-point CCTDE because the velocity inferences at each spatial location are independent of each other. For the same reason, CCTDE accuracy is unaffected by reducing the number of spatial channels, N_{ch}. On the other hand, reducing N_{ch} had detrimental effects on DTW accuracy and precision.

The accuracy of both techniques varied with the underlying velocity field. The *direction* of the velocity must be known before setting the initial slicing direction of DTW. If this was done incorrectly, DTW inferences were unreliable. The *magnitude* of the velocity must be known to infer the accuracy of CCTDE. The salient issue is that the velocity is unknown prior to velocimetry—once again highlighting that a combination of velocimetry techniques must be used for accurate velocity inferences.

VI. AN EXAMPLE APPLICATION OF THE RESULTS AND EXPERIMENTAL WORKFLOW

Note that the simple workflow presented here is not intended to be comprehensive, but rather an example that can be expanded upon and tailored to specific scenarios. Velocimetry was performed on two synthetic TDF time-series, simply named A and B. The signal-to-noise ratio was known prior to analysis at $\mathrm{SNR}_{\mathrm{rms}} = 100$. The spatial parameters and underlying velocity of the data were unknown prior to analysis but could be recovered from metadata at a later point. Despite the use of synthetic data in this section, the example workflow is directly relevant to velocimetry with experimental data.

Before any velocity inferences were made, some preparatory analysis was performed to assess the spatial characteristics of the fluctuations in the data. This was done by taking the magnitude of the 2D spatial Fourier transform of each image in the time-series and manually setting the DC peaks to zero. The resultant Fourier images were summed together, and the result was normalized to the maximum amplitude. An example of a Fourier spectrum can be seen in Fig. 12(a), which showed a clear, single peak. In order to quantify the spatial parameters, one dimensional slices of the Fourier spectra were taken through the peak. These slices were then fitted through a least-squares routine to a Lorentzian of the form shown in Eq. (4). An example of the fit can be seen in Fig. 12(b). A summary of the final estimates for the spatial parameters is shown in Table I. The minimum $\Delta\ell$ required to avoid spurious velocity measurements due to the barber pole illusion was also shown here [as calculated by Eq. (10)].

At $\mathrm{SNR}_{\mathrm{rms}} = 100$, which was well above the $\mathrm{SNR}_{\mathrm{rms}} \geq 1$ limit determined in Secs. III and IV, the impact of noise was thought to be negligible and spatial smoothing or filtering was assumed to be unnecessary. Correlation parameters were calculated through standard techniques[10] and showed that decorrelation effects were negligible. Given that no estimate of the underlying velocity was known at this point, a time-series length of $N = 128$ frames was chosen for CCTDE analysis. This corresponded to a minimum accurate velocity less than 1 px/frame with $\Delta\ell \leq 30$ px and $\lambda_{y0} \leq 50$ px, as determined by Eq. (9). As seen in Fig. 3, the accuracy is expected to vary strongly with velocity. Thus, $\Delta\ell$ had to be varied over a number of velocity inferences to enable approximate quantification of the velocity. A trial range of $\Delta\ell = 10$–30 px was chosen after preliminary testing.

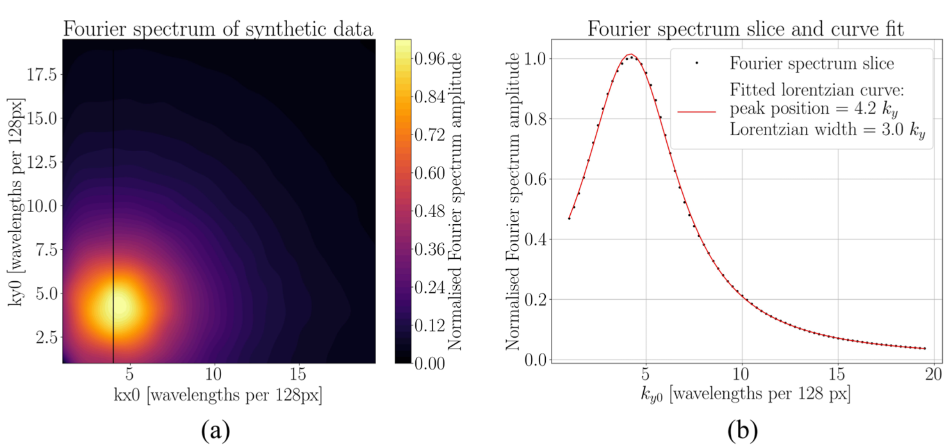

FIG. 12. Example 2D spatial Fourier spectra of time-series A (a). 1D slice of the Fourier spectrum in the y-direction with Lorentzian fit (b). Slice location shown in (a). (a) Spatial Fourier spectrum of synthetic data. (b) Slice of spatial Fourier spectrum with Lorentzian fit.

TABLE I. Table summarizing the estimated spatial parameters associated with the density fluctuations in the time-series. Minimum $\Delta\ell$ was calculated from Eq. (10). All estimated values have uncertainty margins around 10%.

Time-series	k_{x0}	k_{y0}	θ (deg)	Δk	$\Delta\ell_{min}$
A	4.2	4.2	45	$0.7k_{y0}$	13 px
B	2.8	2.7	45	$0.7k_{y0}$	21 px

TABLE II. Summarizing CCTDE inferred velocities for time-series A and B. $\Delta\ell$ was varied from 10 px to 30 px. Correlation amplitudes of the inferences are included. NaN velocities given when the time-delay was zero.

Time-series	$\Delta\ell$	v_x	v_x corr.	v_y	v_y corr.
A	10	NaN	0.3	10	0.5
	15	NaN	0.3	15	1.0
	20	NaN	0.3	20	0.5
	25	NaN	0.3	12.5	0.5
	30	NaN	0.3	15	0.5
B	10	10	0.5	10	0.7
	15	15	0.4	15	1.0
	20	NaN	0.4	20	0.9
	25	NaN	0.3	25	0.6
	30	NaN	0.3	15	0.9

Using these parameters, CCTDE was used to estimate velocities in both time-series A and B. Estimates were made using the standard procedure in Sec. III and are summarized in Table II. The inferred velocities varied strongly with $\Delta\ell$, as was expected. Example velocity fields are not shown here because no spatial variations were observed. For time-series A, no velocities were recorded in the x-direction and the velocity in the y-direction was found to have the highest correlation at 15 px/frame. In time-series B, non-zero velocities in the x-direction were found at $\Delta\ell \leq 15$. These were hypothesized to be spurious velocities due to the barber pole illusion, as predicted by Table I. This hypothesis could be cross-checked with DTW results at a later point. The velocity in the y-direction had multiple values with high correlation, and the y-velocity was estimated to lie around 15–20 px/frame.

The CCTDE results suggest that both the velocity fields point purely in the y-direction. Additionally, no sheared flows were observed and no spurious velocities due to the barber pole illusion should be expected from DTW at $\Delta k = 0.7k_{y0}$. DTW accuracy does not depend strongly on fluctuation spatial scale or underlying velocity at $\text{SNR}_{\text{rms}} = 100$. Thus, DTW could be applied without expected complications. Strip widths of [32, 22, 12, 8, 6, 4, 2] px were used for a total of seven iterations, and the initial slicing direction was set in the y-direction. As can be seen in Fig. 13, DTW was able to measure velocity fields with standard deviations of less than 10% from the mean velocity. For time-series A, the average velocities were 0 px/frame in the x-direction and 15 px/frame in the y-direction. For time-series B, the average velocities were 0 px/frame in the x-direction and 17 px/frame in the y-direction. These results generally agree with the CCTDE estimates and supports the hypothesis that the previous non-zero velocities in the x-direction were indeed due to the barber pole illusion.

CCTDE was re-run for both time-series A and B while varying $\Delta\ell$ from 15 px to 20 px in 1 px increments. For time-series A, the velocity in the y-direction with the highest correlation remained unchanged at 15 px/frame. For time-series B, the velocity in the y-direction with the highest correlation was 17 px/frame. By utilizing both velocimetry methods and cross-checking results, reliable velocity field inferences were obtained for both time-series A and B. The fluctuations in both time-series were found to have no significant velocity in the x-direction. The velocity in the y-direction for time-series A was found to be 15 ± 0.5 px/frame and the velocity in the y-direction for time-series B was found to be 17 ± 0.5 px/frame.

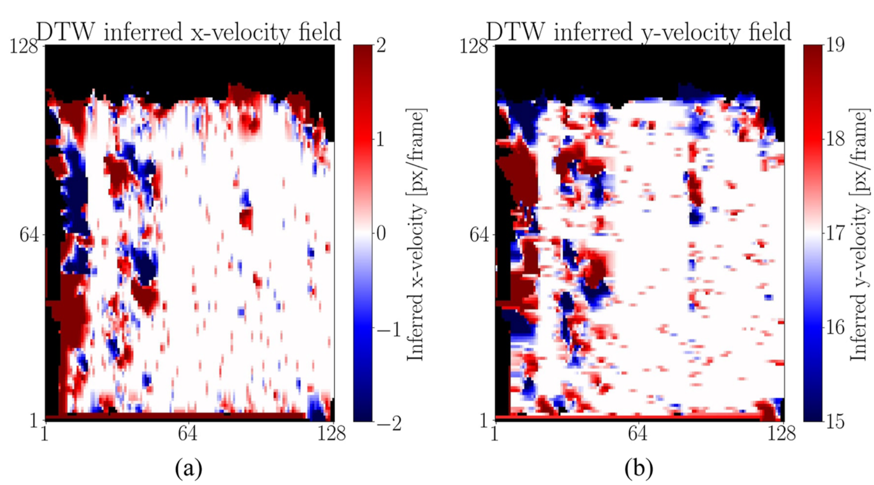

FIG. 13. Example DTW velocity field inferences of time-series B. Velocity in the x-direction averaged out to zero and velocity in the y-direction averaged to 17 px/frame. Color bars centered on the averages. All velocity fields were measured with a standard deviation of <10% from the average. (a) DTW-inferred velocity field in the x-direction. (b) DTW-inferred velocity field in the y-direction.

Uncertainties were defined as the half-step size in $\Delta\ell$ for the final CCTDE inferences. Prescribed velocities were recovered from the metadata at this point, which found y-velocities of 15 px/frame and 17 px/frame for time-series A and B, respectively. The velocity in the x-direction was 0 px/frame in both cases.

In this section, it was shown how fluctuation parameters were used to set the operational parameters of CCTDE, and they accurately predicted spurious velocities due to the barber pole illusion. The direction of the CCTDE-inferred velocity fields was then used to set the slicing direction in the DTW analysis. By cross-checking results from both techniques and refining the CCTDE inferences, consistent results between the two techniques were found. It is emphasized here that the methods can complement each other and should be used together to produce reliable results.

VII. DISCUSSION AND CONCLUSIONS

Two common image velocimetry techniques, CCTDE and DTW, were tested extensively to quantify the dependencies of their accuracy and precision on key parameters in the underlying fields. Synthetic data were used to represent a range of fluctuation structures observed in turbulence diagnostics, namely, ranging from isolated density fluctuation structures to fully developed turbulent density fields. Additionally, specific scenarios were investigated, such as the barber pole illusion, sheared velocity fields, and variation of the number of spatial channels. It was shown in Secs. III and IV that the accuracy of both techniques can exhibit strongly nonlinear behavior. It is therefore ill-advised to extrapolate results from any such tests, including previous literature, beyond the investigated parameter range. Nevertheless, the scope of this paper covers the typical data parameters for most plasma turbulence diagnostics.

Decorrelation effects were not imposed on the synthetic data. For DTW, this is not thought to be an issue unless the decorrelation timescale is comparable to the diagnostic measurement frequency, which is the realm where any velocimetry method is doomed to fail anyway. CCTDE can be more strongly affected by decorrelation effects, especially if $\Delta\ell$ is increased. Nevertheless, decorrelation effects can often be mitigated by analyzing the decorrelation timescale, as discussed in Ref. 10.

This study only investigates velocities pointing in the orthogonal directions, thus preventing the investigation of, for example, rotational flows. However, the impact of such flows on two-point CCTDE and DTW is likely limited. If rotational flows are spatially larger than the spatial resolution, CCTDE will be unaffected. If rotations are smaller, the diagnostic will not record them anyway. It should also be noted that some more elaborate CCTDE methods can infer diagonal velocities.[16] It has been shown previously that DTW can perform accurately with rotational velocity fields.[15]

It was found that CCTDE's accuracy strongly depends on the underlying velocity, which can unknowingly introduce inaccuracies if not carefully considered. A number of options that address this issue were discussed in Sec. III E.

The signal-to-noise ratio and density fluctuation spatial scale did not have a significant impact on CCTDE accuracy in the majority of cases. Precision loss due to noise could mostly be counteracted via filtering.

The length of the time-series could be reduced down to 32 frames without significant impact on the CCTDE accuracy. This represents an inference frequency, which is an order of magnitude faster compared to the typical ≥256 frames.

Spurious CCTDE velocities due to the barber pole illusion were quantified and could largely be avoided through a simple analysis *before* velocimetry. This complication, which is highly prevalent in plasma turbulence studies, could previously not be circumvented without extensive additional analysis.[24]

The accuracy of DTW, with 128 × 128 channel images, was not strongly dependent on the spatial size of the fluctuations or the underlying velocity. Noise was detrimental to accuracy, but these effects could largely be counteracted via spatial smoothing of the images. Sheared flows were detrimental to accuracy, and a threshold in the maximum shear was observed. These tests confirmed the expected result that DTW is typically a reliable technique for such high-spatial-resolution images. A major caveat with DTW is that *the direction of the underlying velocity field must be known prior to analysis*. The flow direction is used to set the initial slicing direction of DTW, which results in complete failure of the velocimetry if set incorrectly.

A more challenging test for DTW was to observe its performance when the number of spatial channels is reduced. In this regime, it was found that the accuracy and precision were strongly reduced in images with 16 × 16 channels or fewer. Additionally, noise had a strongly enhanced detrimental effect on the method precision and accuracy. Through averaging consecutive measurements, accurate velocity fields could be recovered in 8 × 8 channel images with $SNR_{rms} = 2$. This shows that DTW can confidently be applied to high-time-resolution plasma diagnostics, such as beam emission spectroscopy;[11] it is crucial to check with the results in this paper that accurate velocity results can be inferred. No spurious velocities due to the barber pole illusion were observed using DTW at $\Delta k \geq 0.3k_{y0}$.

In conclusion, the accuracy and precision of both CCTDE and DTW were quantified under a broad range of conditions. Improvements on multiple fronts were made to the operational range of both techniques. It was found that two of CCTDE's main drawbacks were as follows: the velocimetry accuracy can depend strongly on the **magnitude** of the underlying velocity and CCTDE is prone to spurious inferences due to the barber pole illusion. Conversely, DTW accuracy was found to strongly depend on the **direction** of the underlying velocity and the number of spatial channels in the data. Notably, DTW was much more robust against spurious inferences due to the barber pole illusion. In general, it is recommended that both CCTDE and DTW can be used in conjunction with each other for accurate velocity field inferences. To this end, a basic example workflow was presented, which utilizes both techniques and should significantly improve confidence in velocity estimates.

ACKNOWLEDGMENTS

The authors would like to thank A. Malcolm-Neale for helpful discussions and H. S. Hacı Ali Oğlu for unwavering support. This work was supported by the Engineering and Physical Sciences Research Council, Award No. EP/R034737/1.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

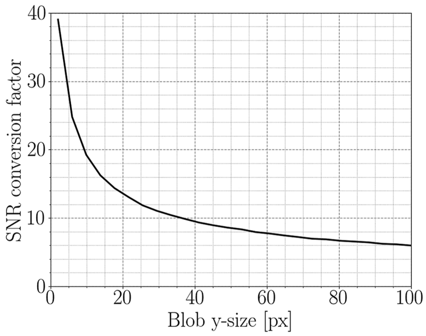

FIG. 14. Conversion factor between SNR_{blob} and SNR_{rms} for IDF data. Factor defined as SNR_{blob}/SNR_{rms}. Blob x-size fixed at 25 px.

Author Contributions

Y. W. Enters: Conceptualization (supporting); Data curation (lead); Formal analysis (lead); Methodology (equal); Software (equal); Validation (lead); Visualization (lead); Writing – original draft (lead); Writing – review & editing (lead). **S. Thomas**: Conceptualization (supporting); Methodology (supporting); Software (equal); Writing – review & editing (supporting). **M. Hill**: Formal analysis (supporting); Writing – review & editing (supporting). **I. Cziegler**: Conceptualization (lead); Funding acquisition (lead); Methodology (equal); Project administration (lead); Supervision (lead); Writing – review & editing (supporting).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

APPENDIX: CONVERTING BETWEEN SNR DEFINITIONS

The conversion factor between SNR_{blob} and SNR_{rms} for IDF data is shown in Fig. 14.

REFERENCES

[1] M. E. Rahaman, S. M. S. Alam, H. S. Mondal, A. S. Muntaseer, R. Mandal, and M. Raihan, "Performance analysis of isolated speech recognition technique using MFCC and cross-correlation," in 2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT), 2019.

[2] S. T. Wereley and C. D. Meinhart, "Recent advances in micro-particle image velocimetry," Annu. Rev. Fluid Mech. **42**, 557–576 (2010).

[3] N. S. Prasad and A. Radhakrishnan Mylapore, "Three-beam aerosol backscatter correlation lidar for wind profiling," Opt. Eng. **56**, 031222 (2017).

[4] J. W. Connor and H. R. Wilson, "Survey of theories of anomalous transport," Plasma Phys. Controlled Fusion **36**, 719–795 (1994).

[5] A. J. Wootton, B. A. Carreras, H. Matsumoto, K. Mcguire, W. A. Peebles, C. P. Ritz, P. W. Terry, and S. J. Zweben, "Fluctuations and anomalous transport in tokamaks," Phys. Fluids B **2**, 2879–2903 (1990).

[6] N. Vianello, D. Carralero, C. K. Tsui, V. Naulin, M. Agostini, I. Cziegler, B. Labit, C. Theiler, E. Wolfrum, D. Aguiam, S. Allan, M. Bernert, J. Boedo, S. Costea, H. De Oliveira, O. Fevrier, J. Galdon-Quiroga, G. Grenfell, A. Hakola, C. Ionita, H. Isliker, A. Karpushov, J. Kovacic, B. Lipschultz, R. Maurizio, K. McClements, F. Militello, A. H. Nielsen, J. Olsen, J. J. Rasmussen, T. Ravensbergen, H. Reimerdes, B. Schneider, R. Schrittwieser, E. Seliunin, M. Spolaore, K. Verhaegh, J. Vicente, N. Walkden, W. Zhang, ASDEX Upgrade Team, TCV Team, and EUROfusion MST1 Team, "Scrape-off layer transport and filament characteristics in high-density tokamak regimes," Nucl. Fusion **60**, 016001 (2020).

[7] A. Ghizzo, D. Del Sarto, X. Garbet, and Y. Sarazin, "Streamer-induced transport in the presence of trapped ion modes in tokamak plasmas," Phys. Plasmas **17**, 092501 (2010).

[8] P. H. Diamond, S.-I. Itoh, K. Itoh, and T. S. Hahm, "Zonal flows in plasma—A review," Plasma Phys. Controlled Fusion **47**, R35–R161 (2005).

[9] I. Grant, "Particle image velocimetry: A review," Proc. Inst. Mech. Eng., Part C **211**, 55–76 (1997).

[10] M. F. J. Fox, A. R. Field, F. van Wyk, Y.-c. Ghim, A. A. Schekochihin, and MAST Team, "Experimental determination of the correlation properties of plasma turbulence using 2D BES systems," Plasma Phys. Controlled Fusion **59**, 044008 (2017).

[11] W. Mandl, R. C. Wolf, M. G. von Hellermann, and H. P. Summers, "Beam emission-spectroscopy as a comprehensive plasma diagnostic-tool," Plasma Phys. Controlled Fusion **35**, 1373–1394 (1993).

[12] S. J. Zweben, J. L. Terry, D. P. Stotler, and R. J. Maqueda, "Invited review article: Gas puff imaging diagnostics of edge plasma turbulence in magnetic fusion devices," Rev. Sci. Instrum. **88**, 041101 (2017).

[13] C. Holland, G. R. Tynan, G. R. McKee, and R. J. Fonck, "Investigation of the time-delay estimation method for turbulent velocity inference," Rev. Sci. Instrum. **75**, 4278–4280 (2004).

[14] D. M. Kriete, G. R. McKee, R. J. Fonck, D. R. Smith, G. G. Whelan, and Z. Yan, "Extracting the turbulent flow-field from beam emission spectroscopy images using velocimetry," Rev. Sci. Instrum. **89**, 10E107 (2018).

[15] G. M. Quénot, J. Pakleza, and T. A. Kowalewski, "Particle image velocimetry with optical flow," Exp. Fluids **25**, 177–189 (1998).

[16] J. M. Sierchio, I. Cziegler, J. L. Terry, A. E. White, and S. J. Zweben, "Comparison of velocimetry techniques for turbulent structures in gas-puff imaging data," Rev. Sci. Instrum. **87**, 023502 (2016).

[17] G. R. McKee, R. J. Fonck, D. K. Gupta, D. J. Schlossberg, M. W. Shafer, C. Holland, and G. Tynan, "Turbulence velocimetry of density fluctuation imaging data," Rev. Sci. Instrum. **75**, 3490–3492 (2004).

[18] D. A. D'Ippolito, J. R. Myra, and S. J. Zweben, "Convective transport by intermittent blob-filaments: Comparison of theory and experiment," Phys. Plasmas **18**, 060501 (2011).

[19] M. Raffel, C. E. Willert, and F. Scarano, *Particle Image Velocimetry* (Springer Cham, 2018).

[20] S. J. Zweben, J. L. Terry, M. Agostini, W. M. Davis, A. Diallo, R. A. Ellis, T. Golfinopoulos, O. Grulke, J. W. Hughes, B. LaBombard, M. Landreman, J. R. Myra, D. C. Pace, and D. P. Stotler, "Comparison of edge turbulence imaging at two different poloidal locations in the scrape-off layer of Alcator C-Mod," Phys. Plasmas **20**, 072503 (2013).

[21] Y.-c. Ghim, A. R. Field, D. Dunai, S. Zoletnik, L. Bardóczi, A. A. Schekochihin, and M. Team, "Measurement and physical interpretation of the mean motion of turbulent density patterns detected by the beam emission spectroscopy system on the mega amp spherical tokamak," Plasma Phys. Controlled Fusion **54**, 095012 (2012).

[22] J. P. Guilford, "Illusory movement from a rotating barber pole," Am. J. Psychol. **41**, 686–687 (1929).

[23] B. Tal, A. Bencze, S. Zoletnik, G. Veres, and G. Por, "Cross-correlation based time delay estimation for turbulent flow velocity measurements: Statistical considerations," Phys. Plasmas **18**, 122304 (2011).

[24] N. Fedorczak, P. Manz, S. C. Thakur, M. Xu, G. R. Tynan, G. S. Xu, and S. C. Liu, "On physical interpretation of two dimensional time-correlations regarding time delay velocities and eddy shaping," Phys. Plasmas **19**, 122302 (2012).

[25]S. Cai, E. Mémin, P. Dérian, and C. Xu, "Motion estimation under location uncertainty for turbulent fluid flows," Exp. Fluids **59**, 8 (2018).
[26]H. Sakoe and S. Chiba, "Dynamic-programming algorithm optimization for spoken word recognition," IEEE Trans. Acoust., Speech, Signal Process. **26**, 43–49 (1978).
[27]G. M. Quenot, "The orthogonal algorithm for optical-flow detection using dynamic-programming," in *ICASSP-92: 1992 International Conference on Acoustics, Speech, and Signal Processing* (IEEE Computer Society, 1992), Vol. 1–5, pp. C249–C252.
[28]S. Thomas and I. Cziegler (2023). "2D dynamic time warping (DTW) algorithm for python," Zenodo, V.1.0.0, Dataset https://doi.org/10.5281/zenodo.7649022.
[29]Y.-c. Ghim(Kim), A. R. Field, S. Zoletnik, and D. Dunai, "Calculation of spatial response of 2D beam emission spectroscopy diagnostic on MAST," Rev. Sci. Instrum. **81**, 10D713 (2010).
[30]P. Sun, C. Chubb, and G. Sperling, Vision Res. **111**, 43 (2015).