

Coupling Coherent Point Drift and Affine Least-Squares Transformation to build trajectories in TR-PTV

L. Thomas^{1,*}, B. Mercier¹, B. Tremblais², L. David¹,

1: Institut P¹, CNRS, Univ. de Poitiers, ISAE-ENSMA - UPR 3346, SP2MI, Poitiers, France

2: Université de Poitiers, Univ. Limoges, CNRS, XLIM, Poitiers, France

*Corresponding author: lionel.thomas@univ-poitiers.fr

Keywords: 2D2C PIV, 3D3C PIV, Coherent Point Drift.

ABSTRACT

The robustness and accuracy of state-of-the-art time-resolved Particle Tracking Velocimetry algorithms is strongly supported by the use of the temporal coherence of particle trajectories. However, to exploit this coherence, the particle needs to be tracked for a sufficiently long period. This condition cannot be achieved if the time of flight of the particles across the measurement volume is not long enough, for example because the limited illumination energy requires the thickness of the volume to be thin. The authors propose to address this difficulty by providing an algorithm based on the spatial coherence of the trajectories rather than the temporal one. This algorithm results from the coupling of two methods; the Coherent Point Drift (CPD), which has proven to be accurate in pairing particles from two instants, but with an insufficient recall for time-resolved application, and the Affine Least Square Transformation (ALST), which allows transporting particles that were missed by the CPD. This algorithm is integrated into a global process that includes an Iterative Particle Reconstruction (IPR) algorithm to extract a list of particles from the images. First tests performed on a synthetic test case consisting of images of particles advected by a turbulent flow highlighted the potential of the method for tracking short tracks. With six frames, the recall is found to be 99.6 % with the present method against 98.3% with the DaVis 10.2 implementation of Shake-The-Box (STB). However, the error is much larger with a false positive rate of 5.1% with our method but only 0.2% with (STB). The cause of this high error rate is analysed and it is speculated that it is partly explained by the tendency of the IPR to detect the same real particle multiple times, which is a concern that has to be addressed in future work.

1. Introduction

Particle Tracking Velocimetry (PTV) encounters two critical bottlenecks: detecting particles in space from a set of images, and correlating detected particles across successive frames.

Schanz et al. (2016) showed that these two problems can be solved together by extrapolating the particle positions based on the knowledge of their previous positions. However, this requires that the first instants of the tracks are somehow initialized.

Emerging methods attempt to circumvent the reliance on a priori knowledge by quantifying the probability of particle pairing across frames. (Zhang et al. (2015); Nie et al. (2021); Yang & Heitz (2021)). In particular, the authors of this contribution are developing a framework (Mercier et al. (2023, 2024b)) based on the Coherent Point Drift (CPD) method introduced by Myronenko & Song (2010). The CPD allows to derive an estimator of the particle displacement field, even in the context of a data set corrupted by a large number of ghost particles, and with hardly any tuning parameters. However, although the error rate in pairing particles using CPD is very low, the recall varies between 50% and 95% depending on the configuration, which is not sufficient to ensure reconstructing long tracks (Mercier et al. (2024b)). The novelty of this specific contribution is the introduction of a new step, referred to as "Affine Least-Squares Transformation" (ALST) by the authors, that allows the accurate advection of any particle based on a local deformation gradient tensor $\mathbf{F}(\mathbf{x}, t)$ derived from the pairs of particles identified by the CPD in the vicinity of the particle under consideration, following the idea of, for instance, Benkley et al. (2023) in the context of the finite mechanical strain analysis. These two steps of the framework are currently implemented and have been coupled with an Iterative Particle Reconstruction (IPR) method by means of a more classical particle prediction method.

The CPD, ALST, particle position prediction and IPR are introduced in the next sections. Finally, a qualitative and quantitative evaluation of the algorithm's performance on a synthetic test case is given, with a comparison to the results obtained by Shake-The-Box implemented in Davis 10.2.

2. Algorithm description

The workflow of the algorithm that is described in this contribution is displayed in figure 1. In short, particles are detected from images at the instants t_0 (in blue in the figure 1) and t_1 with the IPR scheme. Particles from t_0 and t_1 are paired by means of the CPD. Pairs are used to estimate the local deformation gradient tensor $\mathbf{F}(\mathbf{x}, t)$ in the vicinity of each particles detected in t_0 . These particles are then advected to t_1 from the local deformation described by $\mathbf{F}(\mathbf{x}, t)$. If after advection, the particle is within a given distance from a particle detected at t_1 , the particle is validated (in green in the figure 1). If no match is found between the advected particle and a detection at t_1 , the particle will be added to the list of detection but will be flagged as virtual (in orange in the figure 1). The operation is repeated between t_1 and t_2 . A virtual particle can still be converted to "validated" if a match is found in the subsequent time-steps, however, if it remains virtual too long, the tracks is ended. This approach differs from the classical Shake-The-Box since the prediction of the next position of a particle is made from the local spatial coherence of the trajectories rather than the

temporal coherence of a single track. This is expected to improve the quality of the reconstruction of short tracks, for example if the particle tends to move across the measurement volume in a short time. It is also expected to avoid track breaks when a particle is not detected at a given time by introducing virtual particles.

The positions in the next time step of all the individual particles are then estimated more classically to inform the IPR of the position where a particle is likely to be.




Figure 1. Schematic of the workflow described in the paper.

3. Coherent Point Drift (CPD)

The Coherent Point Drift (CPD) method, first introduced in 2006 by Myronenko et al. (2006); Myronenko & Song (2010), is a robust point-set registration algorithm used to align two sets of points in space. Originally developed for computer vision applications, CPD has been successfully adapted to fluid mechanics contexts as demonstrated in Mercier et al. (2023, 2024b). In double-frame velocimetry, CPD aligns particles detected from the images associated with frame 1 (source points) with those from frame 2 (target points). The coordinates of these particles are respectively represented as $\mathbf{Y} \in \mathbb{R}^{M \times D}$ and $\mathbf{X} \in \mathbb{R}^{N \times D}$, where $D = 2$ or 3 is the dimension of the problem.

CPD performs this alignment by determining a geometrical transformation $\mathcal{T}(\mathbf{Y}, \boldsymbol{\theta})$ that maximizes the correspondence between the transformed source points $\tilde{\mathbf{Y}} = \mathbf{Y} + \mathcal{T}(\mathbf{Y}, \boldsymbol{\theta})$ and the target points \mathbf{X} . The method handles rigid, affine, and non-rigid deformations, with the non-rigid variant being the most effective for tracking particles in turbulent flows. This transformation is modeled using a Gaussian Mixture Model (GMM), where the centroids are aligned with the source points, and the Gaussians are weighted by parameters $\boldsymbol{\theta}$. The optimization involves minimizing a cost function, details of which are given in Myronenko & Song (2010). The method had to be adapted

to particle tracking in several ways, which are described in Mercier et al. (2024b).

4. Affine Least-Squares Transformation

The Affine Least-Squares Transformation assumes that the transformation $\zeta_{t_0}^{t_1}$ that defines the displacements field of the particles at time t_0 to the time t_1 is only a function of the local deformation gradient tensor, $\mathbf{F}(\mathbf{x}_0, t)$, and the velocity $\mathbf{U} = [U, V, W]'(\mathbf{x}_0)$ at a reference position \mathbf{x}_0 given

$$\zeta_{t_0}^{t_1}(\mathbf{x}, t) = \mathbf{U}_{\mathbf{x}_0} + (\mathbf{F}(\mathbf{x}_0, t) - \mathbf{I})(\mathbf{x} - \mathbf{x}_0) \quad (1)$$

where \mathbf{I} is the identity matrix. The partial derivatives that constitute the tensor \mathbf{F} and the local velocity can be estimated from a set of known displacements in the vicinity of the considered reference position \mathbf{x}_0 , here corresponding to the position of the particle to be advected. The displacements $[U_p, V_p, W_p]'$ of the neighbor particle pair p are obtained during the application of the CPD. The problem can be written in the form of a linear system, and is solved in a least-squares sense with

$$\tilde{\mathbf{x}} = \operatorname{argmin}(\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2) \quad (2)$$

where \mathbf{A} contains the positions of the particle p $[X_p Y_p Z_p]' = [x_p y_p z_p]' - \mathbf{x}_0$, \mathbf{b} is filled with the samples $[U_p, V_p, W_p]'$, and \mathbf{x} contains the elements of $\mathbf{F}(\mathbf{x}_0, t)$ and $\mathbf{U}_{\mathbf{x}_0}$ where u_x is $\partial u / \partial x$.

More details about the method can be found in Mercier et al. (2024a).

5. Particle position prediction

Following the reasoning that the proposed method should be robust even with short tracks, we propose an alternative approach to the Wiener filter (for instance in (Schanz et al., 2016)) that requires a sufficient number of prior positions on a given track to work optimally. In our approach, we try to predict the next position of the particle not only from its previous positions, but also from the behaviour of the neighbouring particle. Taking into account the local topology of the flow has indeed proven to be beneficial by Khojasteh et al. (2021).

The present method consists in estimating the next coordinate x_{n+1} from the last m known coordinate by

$$x_{t+1} = \sum_{i=0}^{m-1} a_i x_{t-i} \quad (3)$$

with the coefficients a_i obtained by minimizing the cost function

$$J = \sum_{j=1}^n \left(\left(\sum_{i=1}^m a_i x_{t-i,j} \right) - x_{t,j} \right)^2 \quad (4)$$

Table 1. Value of the least square coefficients computed in a turbulent flow for the direction x , aligned with the mean advection direction, and the transverse direction y . Coefficients are given for three levels of noise that simulate particle detection errors. The noise is normalised to the average particle displacement.

	a_2	a_1	a_0
noise 0%			
x	-0.0242	-0.9489	1.9732
y	0.1610	-1.1175	1.9564
noise 5%			
x	-0.5726	0.1532	1.4194
y	-0.2991	0.1827	1.1164
noise 10%			
x	-0.6340	0.2949	1.3391
y	-0.0969	0.2981	0.7988

calculated from the previous coordinates of the same track, and the $n - 1$ closest. This procedure is given here for one dimension, but can be repeated independently for other dimensions of the problem. The choice of m depends on the number of time steps that are available, but it must remain strictly less than this number. The value of the coefficients evolves depending on the magnitude of the acceleration driving the flow, or on the amount of noise in the particle localisation. Values for the three coefficients configuration are given in the Table 1. Coefficient a_i where determined from a large number of tracks of particles advected by a synthetic homogeneous isotropic turbulent flow with a mean velocity in the x direction, and a turbulence intensity of 5% with respect to the mean velocity. Timing parameters, and turbulence scales are defined to match that of a typical PIV experiment in a turbulent flow. Three sets of coefficients are given in the table 1 for three noise levels of 0%, 5% and 10% of the mean particle displacement between two samples. The 0% case shows that the flow direction has a direct effect on the coefficients, as only two coefficients dominate in the flow direction and the three are needed to extrapolate the particle position optimally in the transverse direction. Coefficients are more balanced when 5% noise is added to the true particle position. Surprisingly, at 10% noise, the coefficients in the x direction remain fairly balanced, which is expected since balancing the weight of each observation overall reduces the amount of propagated noise, but the a_2 coefficient in the y direction is greatly reduced. This result is not a singularity, it is repeatable and amplified at higher noise levels. It is the result of a lack of long term consistency of the trajectory in transverse direction due to the small scales of the turbulence. This is confirmed by the fact that a_2 increases again to a value closer to a_1 when the sampling time is reduced.

The Table 2 reports the error obtained for different noise levels and prior positions m used for the prediction. The prediction is made from a bundle of $n = 20$ neighbour tracks, and the error is

averaged over a large number of independent predictions. The error are normalised by the average displacement of particles between two successive instants.

Table 2. Value of the error in the predicted position of a particle with respect to its true position, for different amounts of noise simulating particle triangulation errors, and different number of points used for the prediction. The noise is normalised to the average particle displacement.

Noise	2	3	4	5
0%	4%	4%	4%	3%
5%	11%	9%	8%	8%
10%	18%	14%	13%	13%
15%	25%	19%	19%	17%

In practice, we observe that under the present conditions, increasing m from 2 to 3 noticeably reduces the prediction error, but there is no further substantial improvement for values of m larger than 3. It is also observed that there is a floor of 4% to 3% in the prediction error, even though the particle position is not affected by any noise. This is caused by the scales of the turbulence that cannot be resolved, and that contribute to the apparent randomness of the particle motion in the same way as a triangulation error.

6. Iterative Particle Reconstruction (IPR)

The Iterative Particle Reconstruction (IPR) workflow is schematized in figure 2. The first step is to identify the local peaks in the images (after proper filtering in case of noise) or the signal ratio \widehat{N}_s which is the ratio between the number of non black pixels to the total number of pixels. This estimation of the number of particles per pixel (\widehat{ppp}) can be used with an estimate of the particle diameter \widehat{d} and the volume depth ΔZ to evaluate the real number of particles density in the images, ppp and the detection error \widehat{e}_d . The detection error is necessary to estimate the triangulation error, \widehat{e}_t that is the particle localization error in 3D. A triangulation step returns a list of possible 3D particles. The reprojection error $\widehat{e}_r = \widehat{C}_m + \widehat{e}_d$ can be estimated. \widehat{C}_m is the camera model systematic error. It is in principle corrected during self-calibration, but it can be also adjusted during IPR when there are camera vibrations. The Point Spread Function (PSF) can also be evaluated or refined at this step. After model correction, the reprojection error \widehat{e}_r can be used to estimate more precisely the detection error \widehat{e}_d , and then the triangulation error \widehat{e}_t . The triangulation is performed in an original way (see figure 3).

The volume is divided into boxes. Each 3D box is projected on the images, and the particles detected in the images that fall into the boxes projection, are added to the box. After that, if there is more than one particle in each image, the 3D box is equally divided into eight smaller boxes that are added to the list of 3D boxes to be processed, and the box is removed. Otherwise, if there is

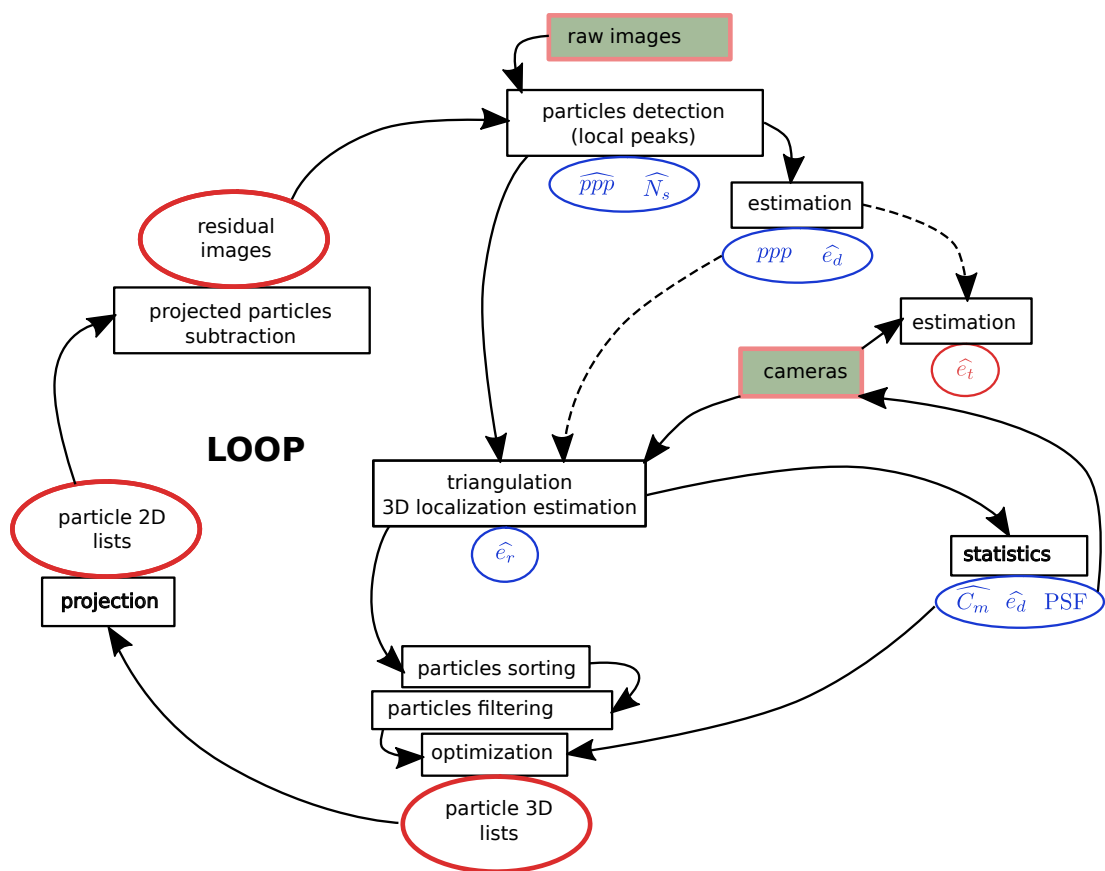


Figure 2. Schematic of IPR workflow. The process uses the particle images and the camera models in order to generate 2D and 3D particle lists, and to estimate the triangulation error.




Figure 3. Illustration of the triangulation process. The boxes are projected on the images in order to determine the particles in the images that could be a projection of a particle inside the box.

exactly one particle for each image in the box, the particle is triangulated using Newton method; if there are at least three particles available, an intersection is also evaluated. If the particle position is outside the box (taking into account a possible triangulation error), the box is removed. The process goes then to the next box. The advantage of this triangulation process is that there is no reference image, and all camera models can be used.

After the triangulation step, the particles are sorted according to their reprojection error, filtered and optimized (taking into account the PSF). Only particles associated with a reprojection error smaller than a threshold that increases with iterations are retained for the next step. This next step consists in the projecting the particle on the images and subtracting them from the original image, giving the so-called residual images. However, care must be taken to avoid subtracting ghost particles, as this would affect the detection of true particles in subsequent iterations. To mitigate this concern, only the particles that are detected and have a position that matches the expected position of an already tracked particle are considered. Other particles are considered to be ghosts and are discarded from the list. This screening step is only performed for a certain number of iterations, after which all particles are taken into account, allowing new tracks to be created with particles not previously expected.

The particle reconstruction process is then repeated with a new reprojection error threshold until either the predetermined maximum threshold is reached or the process stops adding new particles.

7. Performance of the method

The method is now tested on a synthetic experiment that consists of images of randomly distributed particles advected in a controlled manner between two time steps. The images were generated using a high-fidelity in-house PIV image generator, described in Acher et al. (2022), with a density of 0.05 particles per pixel. The particles are massless and are advected by integrating their trajectory over three-dimensional velocity fields for a given time. The setup is made of four high speed cameras, of resolution 1600×1200 pixel, and the volume thickness is 20 mm.

The velocity field has the characteristics of a turbulent flow. It is generated by the method proposed by Martinez-Sanchis et al. (2021). The mean velocity field has a range of values between 0 and 1 m/s along the vertical direction of the measurement volume, which is about 220 mm. The turbulence intensity is 5% relative to 1 m/s, the integral length scale is 100 mm and the Kolmogorov scale is 0.05 mm. The particles are 0.5 mm apart and are advected for 0.5 ms. The mean displacement is therefore close to the mean distance between particles in the high speed region of the flow. Two detail views highlight specific errors in the reconstruction.

The performance of the present method (TR-CPD) is assessed through statistical analysis of the false negative and false positive reconstructed tracks in the table 3. A comparison with the reference method Shake-he-Box (STB) from Davis 10.2 is also provided. Focusing first on the case where

tracks are at least 5 points long, it appears that the false negative rate, calculated as the number of missed tracks divided by the total number of tracks available, is overall better with TR-CPD than with STB. Indeed, this rate is between 0.3% and 0.4% for TR-CPD, while for STB it is 1.7% before complete convergence, and 0.5% after convergence. This is a consequence of advecting virtual particles when they are not triangulated from the images. The downside of this process is that the rate of false positives is much higher with the present method than with STB. The two indicators show more balanced figures when tracks with at least 7 points are considered. The false negative rate increases up to 2.4%, whereas the false positive rate drops down to 1.0%. This indicates that tracks tend to be closed before they reach the end of the domain, which directly increases the false negative rate. It also means that incorrect tracks are not tracked over long period of times, so the false positive rate drops.

Table 3. Value of the rate of false negative, and false positive in tracks obtained from the present method (TR-CPD), and Shake-The-Box (STB). Statistics are calculated for 14×10^3 tracks. Enumeration of time steps starts at 1.

Min track length:		5 points		7 points
Time step:		4	10	10
False negative rate	TR-CPD	0.4%	0.3%	2.4%
	STB	1.7%	0.5%	0.7%
False Positive rate	TR-CPD	5.1%	5.1%	1.0%
	STB	0.2%	0.2%	0.1%

A qualitative view of the reconstructed tracks is shown in the figure 4 to investigate the cause of the observed large false positive rate and premature track closure. Two characteristic failures are highlighted in details A and B of the figure. Detail B shows the post common problem that is observed. A particle is detected many times during the iterative particle reconstruction. This leads to a duplication of the track, which explains the number of false positives. Also, this is not stable over time as the number of spurious detection of a single particle is somewhat random, which overall tends to shorten the lengths of the tracks. Disabling the detection of two particles that are too close to each other is likely to be a solution to this problem. The consequence of such a solution, which is used in the IPR implemented in Davis 10.2, is emphasised in detail A of the figure 4. There are two tracks that are too close to be detected independently by STB. These events may contribute to the false negative rate of the STB method. The benefit of resolving these two tracks independently in terms of analysing the physics of a flow would be negligible, as both tracks exhibit very similar trajectories.




Figure 4. Example of track reconstruction in a synthetic turbulent flow. — : true tracks; -o- : Shake-The-Box; - - - : present method. Thick blue dots show virtual particles. Two detail views highlight specific defects in the reconstruction.

8. Conclusions

This study presents a novel framework for time-resolved particle tracking velocimetry that integrates coherent point drift (CPD) to address the particle pairing problem, and introduces affine least squares transformation (ALST) to increase the robustness of track reconstruction. CPD and ALST are combined with an iterative particle reconstruction (IPR) scheme that iteratively triangulates particle positions from a set of images. The aim of this approach is to bypass the reliance on long-term temporal coherence by exploiting local spatial coherence, thereby improving the reconstruction quality of short tracks and reducing track breaks.

The performance of the method is evaluated using a synthetic test case simulating turbulent flow and a typical experimental setup. The data show that the proposed method achieves a negative false rate comparable to the conventional Shake-The-Box (STB) method, particularly for tracks with a minimum length of 5 points, although the false positive rate is higher with the proposed method. It is also found that tracks are shorter overall with the present method than with STB. A qualitative analysis highlights the role of multiple particle detection in this problem.

Overall, the introduced framework shows encouraging performance, especially in the context of short particle transit times in the measurement volume. Future work will focus on refining the detection algorithm to reduce false positives and further validating the approach with experimental

data.

Acknowledgments

The authors would like to thank the ERDF of the Nouvelle Aquitaine Region for their financial support to the Grinfil project (Convention P-2020-BAFE-33) as well as the ANR through the France relance plan (Convention ANR-21-PRRD-0001-01) .

References

- Acher, G., Thomas, L., Tremblais, B., & David, L. (2022, sep). A new camera model combining an analytical model and a discrete correction to overcome refractive index variation challenges. *Measurement Science and Technology*, 33(12), 125204.
- Benkley, T., Li, C., & Kolinski, J. (2023). Estimation of the deformation gradient tensor by particle tracking near a free boundary with quantified error. *Experimental Mechanics*, 63(7), 1255–1270.
- Khojasteh, A. R., Heitz, D., Yang, Y., & Fiabane, L. (2021). Particle position prediction based on lagrangian coherency for flow over a cylinder in 4d-ptv. In *Proceedings of the 14th international symposium on particle image velocimetry, illinois, usa* (pp. 1–9).
- Martinez-Sanchis, D., Sternin, A., Sternin, D., Haidn, O., & Tajmar, M. (2021). Analysis of periodic synthetic turbulence generation and development for direct numerical simulations applications. *Physics of Fluids*, 33(12), 125130. Retrieved from <https://doi.org/10.1063/5.0071002>
doi: 10.1063/5.0071002
- Mercier, B., Gomez, Q., Thomas, L., Tremblais, B., & David, L. (2023). Proof-of-concept study of coherent point drift registration for particle pairing in particle tracking velocimetry. In *15th international symposium on particle image velocimetry* (pp. 1–6).
- Mercier, B., Thomas, L., Tremblais, B., & David, L. (2024a). A particle based approach for improved resolution piv and tomo-piv based on the coherent point drift and the affine least-squares transformation. In *21st international symposium on applications of laser and imaging techniques to fluid mechanics, lisbon, portugal*.
- Mercier, B., Thomas, L., Tremblais, B., & David, L. (2024b). A robust pairing method for two-pulse particle tracking velocimetry based on coherent point drift. *Measurement Science and Technology*.
- Myronenko, A., & Song, X. (2010). Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12), 2262-2275.

- Myronenko, A., Song, X., & Carreira-Perpiñán, M. (2006). Non-rigid point set registration: Coherent point drift. In B. Schölkopf, J. Platt, & T. Hoffman (Eds.), *Advances in neural information processing systems*. MIT Press.
- Nie, M., Pan, C., Wang, J., & Cai, C. (2021). A hybrid 3d particle matching algorithm based on ant colony optimization. *Experiments in fluids*, 62(4).
- Schanz, D., Gesemann, S., & Schröder, A. (2016). Shake-the-box: Lagrangian particle tracking at high particle image densities. *Experiments in Fluids*, 57(5), 70.
- Yang, Y., & Heitz, D. (2021). Kernelized lagrangian particle tracking. *Experiments in Fluids*, 62(12), 252.
- Zhang, Y., Wang, Y., Yang, B., & He, W. (2015, jun). A particle tracking velocimetry algorithm based on the voronoi diagram. *Measurement Science and Technology*, 26(7), 075302.