



HAL
open science

Double-frame tomographic PTV at high seeding densities

Philippe Cornic, Benjamin Leclaire, Frédéric Champagnat, Guy Le Besnerais, Adam Cheminet, Cédric Illoul, Gilles Losfeld

► **To cite this version:**

Philippe Cornic, Benjamin Leclaire, Frédéric Champagnat, Guy Le Besnerais, Adam Cheminet, et al.. Double-frame tomographic PTV at high seeding densities. *Experiments in Fluids*, 2020, 61 (2), 10.1007/s00348-019-2859-2 . hal-02443793v2

HAL Id: hal-02443793

<https://hal.science/hal-02443793v2>

Submitted on 17 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Double frame Tomographic PTV at high seeding densities

Philippe Cornic^{1†} · Benjamin Leclaire^{2†} · Frédéric Champagnat¹ · Guy
Le Besnerais¹ · Adam Cheminet³ · Cédric Illoul² and Gilles Losfeld²

Received: date / Accepted: date

Abstract A novel method performing 3D PTV from double frame multi-camera images is introduced. Particle velocities are estimated by following three steps. Firstly, separate particle reconstructions with a sparsity-based algorithm are performed on a fine grid. Secondly, they are expanded on a coarser grid on which 3D correlation is performed, yielding a predictor displacement field that allows to efficiently match particles at the two time instants. As these particles are still located on a voxel grid, the third, final step achieves particle position refinement to their actual subvoxel position by a global optimization process, also accounting for their intensities. As it strongly leverages on principles from tomographic reconstruction, the technique is termed Double-Frame Tomo-PTV (DF-TPTV). Standard synthetic tests on a complex turbulent flow show that the method achieves high particle and vector detection efficiency, up to seeding densities of around 0.08 particles per pixel (ppp). On these tests, it also shows a higher robustness to noise and lower root-mean-square errors on velocity estimation than similar state-of-the-art methods. Results from an experimental campaign on a transitional round air jet at Reynolds number 4600 are also presented. Average seeding density varies in time from

0.06 to 0.03 ppp during the considered run, with different densities and signal-to-noise ratios being observed with time in the jet and ambient air regions, supplied by two different seeding systems. The strong polydisperse nature of the seeding, as well as the coexistence of two spatial zones of significantly different particle densities and signal-to-noise ratios, are observed to be the most influential sources of limitation for DF-TPTV performance. However, the method still successfully reconstruct a large amount of particles, and, associated to an outlier rejection scheme based on temporal statistics, truthfully reconstructs the instantaneous jet dynamics. Further quantitative performance assessment is then provided by introducing statistics performed by bin averaging, upon assuming statistical axisymmetry of the jet. Mean and fluctuating axial velocity components in the jet near-field are compared with reference results obtained from planar PIV at higher seeding density, with an interrogation window of size comparable to that of the bins. Results are found to be in excellent agreement with one another, confirming the high performance of DF-TPTV to yield reliable volumetric vector fields at seeding densities usually considered for tomographic PIV processing.

Keywords 3D PTV; tomographic PTV; jet

†: P. Cornic and B. Leclaire contribute equally to this paper and are co-first authors.

P. Cornic
E-mail: philippe.cornic@onera.fr

¹ DTIS, ONERA, Université Paris Saclay
F-91123 Palaiseau - France

² DAAA, ONERA, Université Paris Saclay
F-92190 Meudon - France

³ LMFL-Kampe, Univ. Lille, CNRS, ONERA, Arts et Métiers
ParisTech, Centrale Lille, FRE 2017
F-59000 Lille, France

1 Introduction

Over the past two decades, volumetric optical velocimetry techniques have undergone tremendous advances, paving the way towards instantaneous and unsteady characterization of three-dimensional (3D) complex flows at a large variety of spatial scales. Tomographic Particle Image Velocimetry (TomoPIV) has been introduced as a first technique yielding dense

3D vector fields (Elsinga et al., 2006). While providing a wealth of information on a regular vector grid, allowing convenient physical analyses, one of its major limitations has been quickly identified as its more important degree of spatial filtering than in planar PIV. Indeed, due to 3D imaging constraints, the maximum acceptable seeding density maintaining an acceptable accuracy is known to be lower than for planar PIV (Kähler et al., 2016; Scarano, 2012), imposing to perform cross-correlation with large Interrogation Volumes (IV), thereby smoothing the smallest turbulent scales.

Three-dimensional Particle Tracking Velocimetry (3DPTV) methods, on the other hand, have been introduced much earlier than TomoPIV (Maas et al., 1993; Malik et al., 1993; Nishino et al., 1989; Willert and Gharib, 1992). While having since been applied to a variety of flow situations (including difficult viewing conditions such as in micro-scale flows, see e.g. Park and Kihm, 2006) these methods have long been characterized by a different trade-off, as the accuracy of the velocity estimation is rather linked to the particle image size (much smaller than the typical IV size in TomoPIV), but at the cost of a low seeding density, typically lower than 0.001 particles per pixel (ppp). However, recent years have seen major improvements in 3DPTV’s performance by acquiring multiple images in a time-resolved (TR) mode.

Multiframe 3DPTV, exploiting temporal consistency over a large horizon (typically 10 time steps or more) has led to obtaining reliable particle trajectories and accurate particle location, velocity and material acceleration (Jux et al., 2018; Malik et al., 1993; Schanz et al., 2016; Schröder et al., 2009, 2011). Associated with TomoPIV for 3D detection (Schröder et al., 2009, 2011) or with iterative stereo matching techniques like the Iterative Particle Reconstruction (IPR, Wieneke, 2013) or Shake-The-Box (STB, Schanz et al., 2016), temporal consistency has been the key factor behind recent 3DPTV successes at seeding densities up to ppp ≈ 0.05 . However, these approaches require TR measurements, which suffer a lower signal-to-noise ratio due to limited energy per pulse, and therefore lose in accuracy in situations where good seeding or contrast quality are difficult to achieve. Furthermore, acquisition rate in a regular pulse TR mode is limited to several kHz , higher frequencies requiring to decreasing even more the light intensity. These techniques therefore still suffer from severe limitations in the context of high-speed flows, except if more complex and costly setups can be assembled (such as a pair of interlaced double-pulse lasers enabling four pulse acquisition, see e.g. Novara et al., 2016).

In contrast, we consider here dual frame 3DPTV using a novel processing pipeline, whose design leads in standard synthetic cases to high performance for ppp up to 0.08, while still yielding reliable and accurate results in difficult experimental conditions. As the approach uses several algorithmic steps of the TomoPIV processing, we termed it Dual Frame Tomo-PTV (DF-TPTV). The paper is organized as follows: firstly, section 2 reviews the main recent proposed algorithms that also take as their objective to perform 3DPTV on conventional double frame acquisition to measure velocity information, and underlines the main factors defining the performance in this context. Section 3 describes the proposed DF-TPTV method, organized in three main steps: sparse tomographic 3D particle reconstruction, temporal matching of particles, and vector refinement. Section 4 then characterizes the behavior of DF-TPTV on a large range of seeding density values using synthetic data generated from a turbulent channel flow direct numerical simulation (DNS) (Graham et al., 2016; Li et al., 2008), in both ideal and noisy situations. Section 5 presents an experimental application to a cylindrical air jet. Detection performance of DF-TPTV is first assessed, to illustrate its performance in a real experiment, here with quite challenging conditions of polydispersity, and significant differences in seeding density and signal-to-noise ratios between the jet itself and the outer ambient flow regions, supplied by two different seeding systems. Additional synthetic tests are introduced, mimicking this experiment, confirming that these two characteristics can have a primary influence on its efficiency. The ability of DF-TPTV to yield reliable vector fields is then scrutinized both on instantaneous results, and by estimating statistical quantities which are compared with results from a classical planar PIV system. Section 6 is devoted to conclusions and perspectives.

2 Related works

Recently, high accuracy measurements have been obtained from 3DPTV by Fuchs et al. (2016) and Agüera et al. (2016), with performance illustration by computing ensemble statistics through spatio-temporal binning of 3DPTV vectors. However, these methods differ significantly in their processing steps.

Agüera et al. (2016) first conduct 3D particle detection by a classical 2D particle detection in the images and stereoscopic triangulation, and then solve the temporal matching in two steps. Particles at the first time instant are displaced using a ”predictor” motion field obtained by correlation on low-resolution TomoPIV volumes. They are then matched with particles at the sec-

ond time instant by nearest-neighbor association. The method appears to be limited to low seeding densities ($\text{ppp} = 0.003 - 0.005$) setting hard constraints either on the size of the bin, or on the number of acquired snapshots, in the context of statistical estimation by bin averaging. The authors use as a matter of fact large bins and propose a technique to mitigate the influence of a mean velocity gradient inside each bin that otherwise would bias the computation of second-order statistics.

Fuchs et al. (2016) reconstruct a volume by TomoPIV (using either MLOS or MART) and detect 3D particles by fitting 3D Gaussian to voxel intensities. A selection of particles is then applied: a detection is confirmed if its projections in PIV images can be associated unambiguously to a unique particle image in each frame. This rule eliminates almost all ghosts at the cost of losing a significant number of true particles. Selected detections are then triangulated and propagated to the next time instant by means of a previously estimated displacement field. They finally use a matching process which takes into account the spatial regularity of the motion field so as to reduce outlier vectors as much as possible. Their method has been applied to the computation of mean flow velocity and Reynolds stresses of a turbulent boundary layer, with results equivalent to 3DPTV on TR data (Fuchs et al., 2016).

Both the mentioned methods, Agüera et al. (2016) and Fuchs et al. (2016), are characterized by a relatively low maximum density of estimated 3D vectors, as a consequence of their choices of 3D reconstruction method. The epipolar stereo matching method used in Agüera et al. (2016) can only work at very low seeding densities ($\text{ppp} = 0.003 - 0.005$); TomoPIV methods used in Fuchs et al. (2016) yield a high percentage of ghost particles leading the authors to choose a drastic selection rule.

In the following, we show that, in contrast, by using in particular sparse TomoPIV reconstruction, (Cornic et al., 2015a) we have been able to contain the proportion of ghosts in the 3DPTV process while maintaining a higher number of estimated vectors.

3 The Double Frame Tomo-PTV technique

The technique involves 3 stages, sketched in figure 1: initial particle reconstructions on a voxel grid at the two instants, 3D matching of particles yielding a first estimate of the displacement vectors, and subvoxel refinement of the particle positions and thus of the displacements.

Note that the very first version of the technique was introduced in Cornic et al. (2014), with good perfor-

mances at densities up to $\text{ppp} \approx 0.03$, as attested by the 4th PIV challenge (Kähler et al., 2016). It was further improved in Cornic et al. (2015b). The following presentation corresponds to the stabilized and optimized version of the method, featuring in particular much simplification compared to our previously communicated work.

3.1 Initial particle reconstructions

The first step is a fast and efficient localization and intensity reconstruction of particles on a 3D voxel grid based on sparsity principles, comprehensively described in Cornic et al. (2015a). It is applied to obtain separate and initial particle reconstructions at the two time instants, in the form of indices and intensities of voxels supposed to contain a particle. The initial step is a traditional MLOS operation (Atkinson and Soria, 2009) on a grid with a voxel-to-pixel (v/p) ratio of 0.5. The number of potential particles is further reduced by retaining only voxels corresponding to local intensity maxima ("LocM" strategy). The tomographic reconstruction then relies on the Particle Volume Reconstruction (PVR) imaging model, which seeks to explain the images of a particle with a cluster of few non-zero voxels (Champagnat et al., 2014). The PVR system $Y = WE$ relates pixel intensities Y to voxel intensities E through a weight matrix W made of Point Spread Function (PSF) samples. A sparse solution is defined through the following constrained minimization problem:

$$\min_E \|Y - WE\| \text{ subject to } \|E\|_0 \leq S, \quad (1)$$

where $\|E\|_0$ is the number of non-zero entries of E . We use CoSaMP (Needell and Tropp, 2009), a sparsity-based algorithm to solve this problem over voxels yielded by the LocM selection. The main parameter is S , the upper bound on the number of non-zero voxels in the reconstruction, which in practice is taken as the expected number of particles in the volume. The overall reconstruction algorithm is termed LocM-CoSaMP. As shown by Cornic et al. (2015a), it has a high efficiency to preserve real particles and remove ghosts, which, as discussed earlier, is a critical asset in the context of dual frame 3DPTV.

3.2 3D particle temporal matching

This step consists in identifying the same physical particles in the two instants, with the aim of reducing as much as possible the number of ghosts in the individual reconstructions. As proposed in previous works, we




Fig. 1: Method overview (illustrated in 2D): Particle reconstructions at t and $t + dt$ (yielding sparse/spiky particle representations in the voxel space), temporal matching, and subvoxel refinement of the matched particles. The temporal matching comprises three steps: estimation of an initial displacement field by 3D correlation on a coarse grid of the initially expanded particles, application of the interpolated displacement field to particles at time t to predict their position at time $t + dt$, and local matching with particles at time $t + dt$.

proceed in two steps, first using a predictor motion and then a nearest-neighbor association restricted to a small region around the predicted position. In a preliminary version of the algorithm (Cornic et al., 2014), 2D displacement fields in the four images were used as motion predictor, but this was found to be insufficient, in particular for turbulent flows. A 3D motion predictor derived from the correlation of the two reconstructed volume was introduced in Cornic et al. (2015b) and also used in Agüera et al. (2016).

In practice, LocM-CoSaMP reconstructions are first post-processed before correlation. Non-zero voxels are expanded with Gaussian filtering on a low-resolution, voxel-to-pixel ratio $v/p \geq 2$ grid, on which the 3D correlation will be performed. Gaussian filtering is necessary in the grid transfer process, as the PVR model, used in LocM-CoSaMP, is designed to yield spiky particle reconstructions, extending to a minimum number of voxels (instead of the more traditional "blob" paradigm of 3D PIV, see Champagnat et al., 2014). 3D correlation is obtained using FOLKI-3D (Cheminet et al., 2014), a 3D extension of FOLKI-PIV (Champagnat et al., 2011). Once particles are propagated, a search region of matches of three-voxel radius (expressed here in $v/p = 1$ units) is in practice sufficient.

3.3 Subvoxel refinement

After matching, the obtained particles are still located on a voxel grid, so that a final step performing subvoxel localization is required. Contrary to techniques such as the Iterative Particle Reconstruction (IPR, Wieneke, 2013) or STB (Schanz et al., 2016), that process all the particles sequentially, subvoxel refinement is here performed through a global optimization so as to fully account for the interactions between the particles in the images. The objective function to minimize is the sum of squared differences (SSD) between the recorded images and the images corresponding to the projections of the obtained 3D particles:

$$J(\mathbf{X}_p, E_p) = \sum_j \sum_x \left\| Y_j(x) - \sum_p E_p h(x - F_j(\mathbf{X}_p)) \right\|^2, \quad (2)$$

where Y_j are the recorded images, x a given pixel coordinate in an image, F_j is the projection function in image Y_j yielded by the calibration and h is the PSF. This SSD is thus performed over all pixels of all images. Without loss of generality and to alleviate the notations, we suppose that h is constant and the same for all images.

The non linear least squares criterion J is minimized over the 3D positions and intensities $\{\mathbf{X}_p, E_p\}$ of the particles, i.e. potentially a huge number of variables. To cope with this issue, we used the L-BFGS algorithm (Nocedal and Wright, 2006)¹. The gradient of J can be efficiently computed using projection and back-projection steps. We perform this optimization independently at time steps t and $t + dt$, on matched particles only. In other words, we move independently the two ends of each 3D vector, as illustrated in Fig. 1. At the end of the process, DF-TPTV thus produces estimated 3D displacement vectors with real, subvoxel coordinates.

4 Assessment on synthetic data

Being a PTV algorithm, the method must be assessed for performance both by the number of vectors produced, compared to the number of tracers present in the observation volume, and by their precision, in terms of RMS error on the 3D positions and displacements. This is the purpose of the simulation study presented in this section. We here intend to characterize the behavior of DF-TPTV with respect to traditional parameters such as seeding density and noise. In doing so, we also provide some elements of comparison with performances reported in the literature.

4.1 Synthetic data generation and performance metrics

Similar to e.g. case D of the 4th PIV Challenge (Kähler et al., 2016), we use here one of the flow cases of the Johns Hopkins Turbulence Database, namely the turbulent channel flow DNS (Graham et al., 2016; Li et al., 2008). We define x, y and z as the streamwise, spanwise and wall-normal directions, respectively. A domain $100 \times 110 \times 20 \text{ mm}^3$ has been simulated, taken in contact with the lower wall. Figure 2 shows iso-contours of the Q-criterion color-coded by the local velocity norm of the velocity snapshot used for particle displacement, illustrating its complex turbulent structure. As the database uses dimensionless coordinates, a scaling has been chosen, such that the size of one voxel roughly corresponds to the viscous length scale of the database, i.e. that the first voxel away from the wall corresponds to one wall unit. This leads to a voxel size of $50 \mu\text{m}$, and a volume extension of $2000 \times 2200 \times 400$ wall units.

A traditional, four cameras observation setup is simulated. With the world coordinate origin located on




Fig. 2: Flow snapshot used for synthetic data generation (iso-contours of Q-criterion color-coded by velocity norm), obtained from the Johns Hopkins Turbulence Database (channel flow case, Graham et al., 2016; Li et al., 2008). Friction Reynolds number is equal to $Re_\tau = 1000$. The observed volume has a thickness of 400 wall units starting from the wall, and a length and width of 2000 and 2200 wall units, respectively.

the wall of the channel, at the middle of the illuminated zone in the streamwise and spanwise directions, these cameras are located at the corners of the base of a right square pyramid, whose apex is at $(0, 0, 0)$ and height coincides with the z axis. All four cameras have a 2016×2016 pixel sensor, with a pixel pitch of $11 \mu\text{m}$ (e.g. similar to the PCO Dimax S4), and are equipped with a lens of focal length $f = 200 \text{ mm}$. Their roll position with respect to the optical axis, and their Scheimpflug angle, are computed by assuming that they are all in focus at mid-thickness of the illuminated volume. Imaging of the particles is supposed to be diffraction limited. The apparent diameter of a particle's image is $d_\tau = 2.4$ pixels, resulting from a Gaussian PSF of standard deviation $\sigma = 0.6$ pixel integrated over the pixel surface. Several cases have been considered, both without and with different levels of Gaussian noise.

Laser illumination is modeled as a constant intensity throughout the illuminated volume, whose extension is infinite in the y direction, and spans the range $[-50, 50]$ mm and $[0, 20]$ mm in the x and z directions, respectively. For all synthetic experiments, the upper bound S on the LocM-CoSaMP sparse reconstruction (see Eq. 1) is chosen equal to the true number of particles considered in the simulation and seen simultaneously by the four cameras. It ranges from 33,280 (ppp = 0.01) to 332,287 (ppp = 0.1). As will be explained in section 5, in a real experiment parameter S is chosen on the basis of an estimation of the image seeding density.

¹ we used the L-BFGS implementation of L. Stewart, available at <http://www.cs.toronto.edu/~liam/software.shtml>

The set of particles seen by all the cameras is used as the ground truth for measuring the performance of DF-TPTV in terms of particle detection. To quantify it, we adopt the metrics of Champagnat et al. (2014) and Cornic et al. (2015a). A detection is referred to as a true positive (TP) if it lies within 1 voxel chessboard distance of a true particle. Otherwise it is a false positive (FP), a.k.a a ghost. A non-detection, or false negative (FN), is reported for a true particle with no detection within 1 voxel distance. From these basic quantities two indices of performance are defined:

$$\text{Recall} = \frac{\#TP}{\#TP + \#FN} \text{ and Precision} = \frac{\#TP}{\#TP + \#FP}, \quad (3)$$

where # stands for number of. Recall is the number of detected particles over the total number of true particles in the volume, and is thus the detection rate. Precision is the fraction of true particles among the detected particles. The best achievable performance is given by Recall = 1 (#FN = 0, every particle is detected) and Precision = 1 (#FP = 0, no ghosts).

4.2 Influence of seeding density

Figure 3 (left) shows the Precision and Recall of the reconstruction at initial time t , as a function of the seeding density expressed in particles per pixel. As LocM-CoSamp is tuned to retrieve the exact number of particles, one has $S = \#TP + \#FN = \#TP + \#FP$, so that Recall and Precision are equal. Performance is observed to decrease first slowly with the seeding density; value $\text{ppp} = 0.08$ then sets a break after which the performances decline faster. Note that such a concentration is already above the usually acknowledged optimum for TomoPIV, i.e. roughly around 0.05 ppp . To enable a more detailed comparison with the reconstruction phase of TomoPIV, we also show in this figure the Precision and Recall corresponding to MLOS-SMART (Atkinson and Soria, 2009), using 25 iterations. As is classically the case with such methods, MLOS-SMART operates at a very different trade-off between both quantities, showing a nearly maximum Recall throughout the range of ppp , with Precision dropping dramatically, and indeed reaching some quasi saturation at very low values from $\text{ppp} \approx 0.05$. Returning to comparison with another 3D PTV approach, by Fuchs et al. (2016), we observe that their selection rules led them to a quasi negligible percentage of ghosts, at the cost of low Recall (detection rate). According to Fuchs et al. (2016)'s Fig. 1(a), their Recall is indeed around 70% at $\text{ppp} = 0.05$, while we here obtain a 95% Recall at this ppp . Even though

the dataset considered by Fuchs et al. (2016) and the present one are different and therefore cannot be compared directly, these figures tend to support the fact that the two methods rely on a different operating point between Precision and Recall. In other words, the reconstruction step of DF-TPTV appears more balanced between the percentage of ghosts and the detection rate compared to Fuchs et al. (2016). We will see below that the fact of tolerating a higher rate of ghosts allows here to retrieve a higher percentage of true vectors while maintaining the ratio of outliers vectors as low as a few percent.

In figure 3 (right), we now account for the accuracy in particle localization, by plotting the mean and RMS position errors yielded by DF-TPTV, again compared to that of MLOS-SMART, for the same variation of ppp . For homogeneity, in the case of MLOS-SMART, these quantities are computed by only keeping the true particles in the reconstructions, and performing a three-dimensional Gaussian fit to determine their subvoxel positions. The curves for DF-TPTV again show a gradual increase in the error with growing ppp , until $\text{ppp} = 0.08$ where a change in slope occurs. However for both the mean and RMS values, DF-TPTV yields systematically smaller error values than MLOS-SMART, except at the highest density, $\text{ppp} = 0.1$, as a result of this change of slope. Over a large range of densities starting from the lower values and up to $\text{ppp} \approx 0.05 - 0.06$, the errors of the former are between two and three times smaller than those of the latter. Even more interestingly, the mean location errors obtained with DF-TPTV are lower than that reported in the literature for state-of-the-art particle reconstruction or double frame 3D PTV techniques, *e.g.*, IPR Wieneke (2013) or Fuchs et al. (2016), and comparable to the 3D positional errors of STB obtained for the first result available from the initialization phase, which relies on the processing of the first four snapshots of a sequence (Schanz et al., 2016). To illustrate the repartition in location error among the components, figure 4 displays the location error probability density function (pdf) in the xy and xz planes, for $\text{ppp} = 0.05$. It may be seen that the pdf decays very fast in the vicinity of the origin and is elongated in the z direction in the xz plane, showing that higher error values are obtained in the z component. This is quite classical due to the geometry of the optical setup, and is observed for all seeding densities considered here (not shown for conciseness).

Precision and Recall, that have been introduced above for characterizing the reconstructed particles, can also be computed on velocity vectors. Here, the ground truth is made of all vectors defined by a true particle visible at the two time instants by all cameras. A true




Fig. 3: Performances of DF-TPTV for particle reconstruction with respect to seeding density (ppp), compared with MLOS-SMART. For the latter, the particles' positions are extracted using a 3D Gaussian fit. Left: Precision and Recall of the particles reconstruction step. Right: total mean and RMS position error of true particles.




Fig. 4: Position error probability density function for DF-TPTV at ppp=0.05. Left: xy -plane. Right: xz -plane.

positive vector (TPv) stems from the detection of the same particle at both instants and the correct matching of the two detections. The fraction of true positive vectors TPv among all retrieved vectors is the Precision, while the fraction of TPv among all true vectors is the Recall. Figure 5 (left) shows vector Precision and Recall of DF-TPTV over a range of ppp up to 0.1. Upon comparing for each ppp the values of Recall in figure 3 (left), one observes that the obtained vector Recall is approximately equal to the square of the reconstruction Recall. This is ascribed both to the nature of the displacement considered here, and to the monodisperse particle distribution. These two conditions jointly en-

sure that, at the two instants, the probabilities for each particle of being reconstructed are independent, leading to this value of the vector Recall. This would not be the case, either for a monodisperse distribution and a very smooth displacement field (such as e.g. a pure translation), for which epipolar geometry would induce that more or less the same particles at t and $t + dt$ would be missed, or for a polydisperse distribution, in which the probability of each particle to be reconstructed would be linked to its intensity. Note that in the latter case, as will be observed in section 5.4, this leads to a vector Recall higher than the square of the particle Recall. Overall, this result shows that vector Recall is mainly




Fig. 5: Performances of DF-TPTV for vector reconstruction with respect to seeding density (ppp). Left : vector Precision and Recall. Right: RMS errors on velocity, RMS_v and RMS_v^* (see the text for definitions). RMS_v^* can be computed, and is shown here, over all vectors (blue) and only true positive vectors (red). RMS_v is computed over true positive vectors only.

limited by the missing detections rather than by the matching step, justifying the use of a predictor motion field based on 3D correlation. An additional sign of the efficiency of the 3D matching is that, contrary to the Precision of particle detection, vector Precision stays close to 1 when the ppp increases, supporting its ability to filter out most of the ghost particles. The vector Recall can be translated straightforwardly to the *effective amount of particles by pixel* metric introduced by Fuchs et al. (2016), by multiplying it with the seeding density. For instance, at ppp equal to 0.05, figure 5 (left) indicates a vector Recall equal to 0.9, thus corresponding to a $0.045 = 0.05 \times 0.9$ effective amount of particles by pixel. As a comparison, Fuchs et al. (2016) report a maximum effective amount of 0.032 (see their Fig. 2b) — however bearing in mind the slightly different characteristics of their dataset. The percentage of ghost vectors among the total number of true vectors can be computed as

$$\%ghost = \frac{(1 - \text{Precision}) \times \text{Recall}}{\text{Precision}}. \quad (4)$$

For the considered $ppp = 0.05$, figure 5 (left) shows that $\text{Recall} = 0.9$, $\text{Precision} = 1 - 0.01$, consequently the percentage of ghosts is equal to 0.9%, i.e. indeed a very low proportion of the vector field. In this figure, we also introduce the matching rate, defined as the number of vectors obtained by DF-TPTV divided by the total number of particles in the flow. This quantity is seen to nearly collapse with the vector Recall, showing only a slight deviation from it at the highest seeding densi-

ties. In fact, it can be shown that, in this synthetic case where the (particle) Recall and Precision are equal, this matching rate is equal to the ratio between the vector Recall and the vector Precision and thus does not introduce much new information. We introduce it here mostly as a reference for comparison, since we will analyze the matching rate in our experiments in section 5.

Vector fields obtained by DF-TPTV are further assessed using the RMS error on the velocity vectors. In that respect we introduce two quantities, denoted RMS_v and RMS_v^* , where subscript v stands for velocity, and which are defined as:

$$RMS_v = \sqrt{\frac{1}{P} \sum_{p=1}^P \sum_{i=1}^3 (u_i(x_p^{GT}, y_p^{GT}, z_p^{GT}) - \tilde{u}_i(x_p, y_p, z_p))^2}, \quad (5)$$

$$RMS_v^* = \sqrt{\frac{1}{P} \sum_{p=1}^P \sum_{i=1}^3 (u_i(x_p, y_p, z_p) - \tilde{u}_i(x_p, y_p, z_p))^2}, \quad (6)$$

where P is the number of detected particles, (u_1, u_2, u_3) is the ground truth 3D displacement and $(\tilde{u}_1, \tilde{u}_2, \tilde{u}_3)$ is the DF-TPTV estimation of the 3D displacement. Coordinates (x_p, y_p, z_p) are the spatial locations of the particles detected by DF-TPTV, while $(x_p^{GT}, y_p^{GT}, z_p^{GT})$ denotes their ground truth locations. RMS_v is built by considering the set of true particles only, and therefore only assesses whether the velocity of these particles

reconstructed by DF-TPTV has been estimated accurately, whereas RMS_v^* also includes the fact that, due to its positional error, a true particle might have been reconstructed at the wrong spatial location, possibly inducing an erroneous spatial portrait of the flow. Figure 5 (right) shows these two quantities as a function of the ppp. In the case of RMS_v^* , we include assessment both over only the TP vectors (red curve) and over all the vectors (blue curve); examining only the TP yields a lower bound RMS error of what would be obtained after outlier rejection, if perfect. Firstly, it appears that the curves of RMS_v , and of RMS_v^* considering the TP only, collapse over the whole range of ppp. This shows that errors on particle positions are much lower than the spatial scale of the finer flow gradients. As this will be systematically the case, also with the case of the next section including noise, we will restrict to RMS_v^* in the following. Trends are logically similar to that observed for particle localization error, with a gradual increase of RMS_v up to $ppp = 0.08$, and a more pronounced performance loss for higher densities. It can also be seen that, although there are only few wrong vectors compared to the number of good ones, they have a noticeable influence on the RMS when the ppp increases. This will motivate the introduction of outlier rejection in the processing of experimental data, in section 5.

4.3 Influence of noise

We now provide a first series of tests to assess the robustness of DF-TPTV to noise, by considering the same image settings as in Schanz et al. (2016), to allow comparison. Datasets generated are similar to the previous ones, except that several levels of Gaussian noise, of standard deviation σ expressed relative to the mean peak image intensity of the particles I_{avg} , are considered. Similar to Schanz et al. (2016), three levels, from very low to very high noise, are chosen: $\sigma = 0.03$, 0.1 and 0.2. Mean level of noise is set to $\sigma I_{avg}/2$. As in Schanz et al. (2016), no image preprocessing has been applied.

In a similar way as in the nominal case presented above, we show particle detection performances (Precision and Recall for both particles and vectors, in figure 6), mean and RMS errors on particle positions (figure 7), and RMS errors on particle velocities (figure 8). In all these, we recall the result without noise, to allow grasping the effect of the lowest noise level. For all these quantities, the noise logically leads to loss of performance, in direct relation to its magnitude. The overall shape of the curves is conserved, with a discrepancy compared to the ideal case growing in value as the ppp grows. This fact is also expected, as noise leads to more




Fig. 8: Sensitivity of DF-TPTV RMS velocity error to noise, for varying seeding density. Separate analyses including all reconstructed vectors (TP+FP) and only TP vectors.

ambiguities in particle reconstruction, and should make situations of overlapping particle images even more difficult to handle. While the lowest noise level, $\sigma = 0.03$, only has minor differences with the ideal case, considering higher values leads to increasing the drop of performance (change of slope of the curves) for earlier values of the ppp, namely $ppp \approx 0.07$ for $\sigma = 0.1$, and $ppp \approx 0.06$ for $\sigma = 0.2$. Noteworthy, it may be seen in figure 6 (right) that whatever the noise up to $ppp=0.06-0.07$, the vector Precision lies very close to 1. It then begins to decay fastly only in the noisiest case.

Values at $ppp = 0.01$ and 0.05 of the mean error on particle positions shown in figure 7 (left) can be compared to values obtained by Schanz et al. (2016) corresponding to the first point of the temporal development of STB processing which, as mentioned above, is obtained from considering the first four frames. Comparison shows that, except for the lowest noise level $\sigma = 0.03$ at $ppp = 0.05$, mean errors obtained by DF-TPTV are all lower, even though the algorithm only exploits the information contained in two frames. The discrepancy is in particular quite pronounced for the intermediate and high levels, errors for DF-TPTV being even lower than converged errors of STB, i.e. when considering for the latter a large number of snapshots.




Fig. 6: Sensitivity of DF-TPTV detection performance to different levels of noise, for varying seeding density. Precision and Recall for particle reconstruction (left) and vector reconstruction (right).




Fig. 7: Sensitivity of DF-TPTV particle localisation performance to different levels of noise, for varying seeding density. Left: mean position error. Right : RMS position error.

5 Experimental results on a round air jet

5.1 Experimental setup

Demonstration on experimental data is performed by considering a round air jet of exit diameter $D = 12 \text{ mm}$. With $V_j = 5.8 \text{ m.s}^{-1}$ the velocity at the centre of the exit plane, the flow Reynolds number based on D and V_j is equal to $Re = 4600$. Flow conditions in the exit plane are transitional in the boundary layer, due to the presence of a small circular step at the nozzle wall 20 mm upstream from the exit. In the following, only dimensionless quantities, built using D and V_j as reference length and velocity, are considered. The centre of the jet exit plane is taken as the origin O of the coordi-

nate system; y denotes the direction aligned with the jet axis, here vertical, and x and z the horizontal axes (see figure 9).

The near field of this jet, up to $y/D \approx 7.3$, is measured using two PIV systems. 3D measurements are acquired in a parallelepiped with its largest edges in the x and y directions, with an approximate thickness of 16 mm , centered around O in the z direction (see figure 9). Illumination is achieved using a Quantel Twin Ultra Nd-Yag laser delivering 120 mJ per pulse, and observation by two Dantec HiSense and two LaVision Imager ProX 4 Mpixels cameras (2048×2048 pixel), set up in a cross like configuration. A mirror is placed in order to reflect back the volumetric illumination and compen-




Fig. 9: Transitional jet experiment, featuring 3D PTV measurement in a parallelepiped of 16 mm thickness in the z direction (with volumetric illumination on), observed by four cameras, and planar PIV measurement in a jet-longitudinal yz plane (camera in perspective observation, therefore used as a reference for the streamwise velocity component only).

sate for the different scattering condition among the cameras.

In order to have a reference on flow quantities for performance assessment, and also to provide an independent seeding density evaluation during the tests, an additional planar PIV measurement system is set up. It is made of one Litron Laser Dual Nd-YAG 532nm laser and one Dantec HiSense 11M (4000×2672 pixel) camera. The laser sheet, of estimated thickness 1.5 mm , is located in the $x = 0$ plane, orthogonal to the main axis of the tomographic laser. Due to the presence of the volumetric illumination, the camera optical axis cannot be placed perpendicular to the laser sheet, and is thus set with a slight perspective. The camera is equipped with a Scheimpflug mount, and calibration is applied in order to compensate for perspective distortion in the observation. Due to this single-camera observation, this system is used for flow comparisons only on the streamwise, v component, and its camera is therefore labelled as "2D1C camera" in figure 9. Acquisitions of the 3D and planar systems are intertwined, in the sense that during a run, 3D and planar snapshots are acquired alternatively, with a separation time of 0.25s. This results

in a respective acquisition frequency for both systems of 2Hz. Besides, both also operate with the same inter-frame time of $50\mu s$.

Seeding is achieved using two different aerosol generators producing DEHS droplets, one whose particles are injected into the jet settling chamber (thus seeding the jet), and one used to seed the experimental room (and thus the outer shear layers and the entrained flow). Whereas the former continuously injects particles during the run, the latter is operated by initially saturating the (quite large) room, and waiting for homogenization to begin acquisition, without further injection later on. As shown in figure 10, the consequence of this setup and operating mode is twofold. Firstly, the jet and ambient air regions exhibit same seeding density only for a limited time horizon, around snapshot number 200. For earlier snapshots, the ambient air is more densely seeded than the jet, and conversely for later snapshots, due to a constant decay in density there during the 8 minutes lasting run. Secondly, due to the larger area in the images occupied by the ambient air, the overall seeding density also decreases during the run, from slightly less than 0.06 ppp at the beginning of the run,




Fig. 11: Evolution of seeding density during the run: in the entire images (global), and in the regions corresponding to the jet and to the outer ambient air (see yellow dashed line depicting the considered boundary in figure 10). Also included is the ppp corresponding to the number of particles (sparsity S) set in the particle reconstruction step of DF-TPTV.

until roughly 0.03 ppp. Note that, to determine this seeding density, we count the average number of local maxima (on a 3×3 pixel neighbourhood) present in the raw images. In that process, a threshold for detection is applied, which is chosen as an upper bound of the noise level in the images; the latter can be estimated from a non-illuminated zone. Overall, it is known that such a procedure for estimating seeding density remains accurate for low ppp, where only few particle images overlap, and leads to a possibly significant underestimation at higher densities (see, e.g., Warner and Smith, 2014). We thus emphasize here that the present numbers are conservative with respect to the actual seeding density. To enable separate analysis between jet and ambient air regions in the following and account for their different ppp, we have chosen a simple volumetric delineation, whose projection in the images is sketched schematically in figure 10. The jet region has been considered as a truncated cone, of section diameter equal to the jet diameter at the exit, and twice the jet diameter at the most downstream illuminated location.

5.2 Processing parameters

Calibration and self-calibration of the 3D system are done according to a pinhole model (Cornic et al., 2016), leading to a voxel size ($v/p = 1$) of $53.8\mu\text{m}$. Standard pre-processing steps are applied to the images of the

3D cameras before applying the DF-TPTV method, i.e. historical minimum subtraction and thresholding by identifying the maximum noise level from a non-illuminated zone (see figure 10). Additionally, to compensate for differences in dynamics between the camera images (due to the use of different camera models, and different scattering conditions in spite of the mirror), normalization was also necessary to obtain comparable signal-to-noise ratios between the cameras. It is performed as a gamma correction (using the value $\gamma = 0.5$) mapping each camera image from its native range, $[\text{LB UB}]$, to the common range for all $[0 \ 1000]$. Here, LB and UB respectively stand for lower and upper (intensity) bounds, and correspond, respectively, to the noise level mentioned above, and to the intensity level above which 0.5% of the intensity levels lie.

To obtain velocity fields from these pre-processed data, the DF-TPTV method requires choosing the number of particles to retrieve (sparsity parameter S in the LocM-CoSaMP particle reconstruction step). This number should be logically linked to the seeding density estimated in the images. In practice, we choose to set S to a value inferior to the number of particles arising from image density estimation. Concretely, this is simply done by setting a higher intensity value for the threshold used to consider local maxima in the images as actual particles. As shown in figure 11, it leads to targeting a number of reconstructed particles representing between 60% and 80% of the counted particles, depending on the instant in the run. We indeed observed that, at least for the present experiments, even though it amounts to reconstructing less vectors than could be possibly extracted from the data, such a choice leads to an increase in the matching rate between particles at the two time instants, and to less wrong vectors in the final fields. This was verified by processing a few snapshots with different values of S . This choice will also be further supported by the analysis derived in section 5.4. We would like to emphasize here that the value that we chose for S compared to the (possibly already underestimated) number of particles obtained from ppp determination by counting probably depends of the nature of the present experimental data, making it difficult to indicate it as a systematic guideline. Another justification to this point is that we think that other possible choices for S might also be possible, depending on the remaining processing steps. In particular, in the idea of maximizing the final amount of measurements, one could instead try to choose a higher S , closer to the number of particles corresponding to the ppp, which could possibly lead to more outliers in the vector fields, and then set the outlier rejection scheme to a higher threshold value in order to remove them. Other tuning




Fig. 10: Sample particle images from one of the 3D cameras, average seeding density estimated to 0.06 ppp (left) and 0.03 ppp (right). Contrast has been enhanced for display purposes. In the left subfigure, the dashed yellow line schematically depicts the boundary that has been used in the volume to perform separate diagnostics between the inside of the jet, considered as a truncated cone, and the outer ambient flow.

parameters of the DF-TPTV algorithm, i.e. pertaining to matching and subvoxel refinement, are left to their default values, mentioned in section 4.

In order to filter the results from remaining outliers, we introduce a rejection post-processing. To do so in a most adapted and efficient way, following e.g. Griffin et al. (2010), we choose to rely on temporal statistics computed by bin averaging, that will be introduced to compute mean flow fields (see section 5.5). As already noted in the PTV literature, such an approach, when available, is more efficient for turbulent flows than relying on comparisons to a spatial neighborhood. Contrary to Griffin et al. (2010) however, we here rely on a simpler combination of univariate statistical rejection rules, as we choose to reject a given vector if any of its (u, v, w) components deviates from more than three standard deviations from its mean (the latter two referring to that of the bin to which the vector belongs).

Finally, note that variations in illumination within the volume led in practice to restrict the results to $-0.61 \leq z/D \leq 0.40$, i.e. a 13 mm thickness, in order to exclude edge effects where light intensity was lower. More details on this are given in section 5.5.1. Instantaneous fields will also be restricted to this zone for consistency. Also, note that the most upstream location of the volumetric vector fields yielded by DF-TPTV is located slightly above the jet exit, i.e. $y/D \approx 0.3$. As can be seen in figure 10, in order to avoid intense light

reflections on the nozzle, the laser volume was indeed positioned slightly above its exit plane.

5.3 Example of instantaneous results

As a first experimental illustration of the DF-TPTV method performance, we show in figure 12 the instantaneous vector field obtained for the first pulse of the run, at highest seeding density, estimated to around 0.06 ppp (see also corresponding particle images of one of the cameras in figure 10 left). We show both raw results and results post-processed via the outlier rejection method presented above, in order to illustrate its effect. Out of the 30,787 vectors obtained in the retained $-0.61 \leq z \leq 0.40$ zone, outlier rejection filtered roughly 2.7% of them, leading to a useful set of 29,947 vectors. This post-processing seems indeed to be efficient and adapted, rejecting a large number of spurious vectors while keeping the physical ones, as can be observed most evidently in the external flow.

Figure 12 also shows that the expected instantaneous structure of this transitional jet is successfully retrieved by the DF-TPTV. Starting from the jet exit, one first observes uniform axial velocity in the jet core ($|x|/D < 0.4$, y/D ranging from 0.3 to around 1.5), together with a thin shear layer on the lateral edges. Note that in this jet potential core zone, as the volume spans over $-0.61 \leq z \leq 0.40$, one also observes the start of




Fig. 12: Instantaneous vector field of the transitional $Re = 4600$ round air jet yielded by the DF-TPTV method, color-coded by the normalized streamwise component v/V_j . Left: raw result, right: result with outlier filtering (see text for details). Full field (top) and close-up on the region delineated by the dashed rectangle.

the shear layer in nearly all azimuthal directions, which translates here into the presence of vectors also with lower axial velocities (of around $v/V_j \approx 0.5$, appearing in green). Following the jet evolution in the downstream direction, one then observes typical toroidal vortical structures in the shear layer, due to the Kelvin-Helmholtz instability, coinciding with accelerations in the jet core. As also expected, in each region separating

two successive of these vortices, the jet column is seen to expand, leading to flow deceleration on the axis as a result of mass conservation. These typical dynamical features can be observed in more detail in the close-up also shown in figure 12.

To complete this first assessment of instantaneous results, we finally compare the vector fields yielded by DF-TPTV together with those obtained by traditional




Fig. 13: Left: instantaneous vector field yielded by Tomo-PIV processing of the snapshot shown in figure 12, color-coded by the normalized streamwise component v/V_j and using the same colormap. Right: Vectors corresponding to the difference between DF-TPTV and TomoPIV processing, color-coded by the normalized correlation coefficient corresponding to each vector obtained by TomoPIV.

TomoPIV processing. For the latter, we perform tomographic reconstructions of the two instants using an MLOS first guess followed by 25 SMART iterations, as has been done in section 4.2. We then cross-correlate the obtained particle distributions with FOLKI-3D (Cheminet et al., 2014), to get the final vector fields. The algorithm relies on an iterative scheme with volume deformation, whereby cubic BSpline interpolation is performed. An interrogation volume (IV) size of $33 \times 63 \times 33$ voxels ($0.15D \times 0.28D \times 0.15D$) has been chosen, in order to guarantee a minimum number of tracers per IV, with the anisotropy accounting for the present physics, in particular the thin shear layers nearest to the jet exit. Note that the spatial resolution is in practice lower than the size of the IV, as a Gaussian weighting is used, with $\sigma = 8 \times 15.5 \times 8$ voxels ($0.036D \times 0.069D \times 0.036D$). In Figure 13 left, we represent the values of the velocity vectors yielded by TomoPIV, at the spatial locations of the particles found by DF-TPTV. This is performed by interpolating the velocity fields yielded by FOLKI-3D. Note that this interpolation can be considered quite accurate; indeed, the specific algorithmic architecture

of FOLKI-3D relies on computing the displacement *for each voxel* in the 3D grid, i.e. with a maximum overlap. We here take advantage of this possibly dense result to avoid truncation during this interpolation process. In the figure, the same representation and colormap as for DF-TPTV have been used. In Figure 13 (right), we display the field of difference vectors between DF-TPTV and Tomo-PIV, which we color using the normalized correlation (or correlation coefficient) corresponding to each vector obtained by Tomo-PIV. The value of this coefficient can usually be considered as an indicator of the measurement accuracy, a value close to 1 being most often associated to a large number of particles (or a low level of noise) and to weak displacement gradients within the IV. As the seeding is quite homogeneous in the considered snapshot, one indeed notices that the values of this correlation coefficient reflect the presence of flow gradients, since high values are observed in the jet core and the outer flow, while the shear layers and the most downstream zone of the jet exhibit lower values. Of specific interest to the present study is that there is also a coincidence between the value of the cor-




Fig. 14: Evolution with snapshot number of the matching rate between the two instants in a snapshot, in the entire flow field (global) and in the regions corresponding to the jet and to the ambient air. Also shown is the estimation of the vector recall during the run.

relation coefficient and the norm of the difference vector between the DF-TPTV and TomoPIV processings, the higher values of the coefficient being more often associated to small discrepancy values, and conversely. This trend indicates that DF-TPTV certainly improves the displacement estimation in the high gradient zones, where TomoPIV is hindered by the necessarily large IVs, leading to a smoothed and less reliable estimation. This observation will be completed and confirmed when considering statistical results, in section 5.5.

5.4 Detection performance assessment

As a first, more quantitative assessment in this experimental context, we now investigate some performance metrics pertaining to vector detection by DF-TPTV. Figure 14 presents the evolution during the run of the matching rate, defined as the proportion of particles matched between the two instants of a snapshot. This matching rate is indicated for the global volume, as well as separately for the jet and ambient air regions, since they are characterized by different seeding densities. Also shown in the figure is the estimation of the vector recall, which can be determined from other shown quantities, i.e. as the ratio between the number of particles S chosen for the reconstructions (corresponding to the orange ppp curve in figure 11), divided by the estimated actual number of particles (corresponding to the solid blue curve in this figure), multiplied by the global matching rate. To appreciate the change in behaviour of the algorithm in these realistic conditions

compared to more idealistic synthetic ones, even with noise, these obtained values can be compared to those of figure 6 (right). A first striking, but possibly not totally unexpected feature is that the levels of both the matching rate and the vector recall are quite significantly lower in these experimental results than in the synthetic conditions considered in section 4. This is especially the case for the vector Recall. Considering for instance a situation corresponding to the first snapshot, for $ppp \approx 0.06$, this quantity is comprised between 0.68 and 0.88 depending on the noise level, while in the experiment it is equal to 0.43. The matching rate obtained for the noisiest synthetic case is also close to 0.7, while it is of 0.55 in the experiment. A first explanation for this discrepancy, at least regarding the vector Recall, can be ascribed to our choice of reconstructing less particles than estimated from the images, as explained in section 5.2. However, this explanation does not account for observed differences with the theoretical expected behaviour: both experimental curves for the matching rate and the vector Recall should increase with the snapshot number. Indeed, as shown in figure 11, the global seeding density decreases during the run, which, according to synthetic results, should lead to an increase in performance. When considering separately the jet and the ambient flow in figure 14 (right), another possible cause appears, being due to the difference in seeding between these regions. It turns out that the matching rate indeed increases with time in the ambient flow, where the ppp decreases. However, according to the synthetic results presented in section 4, this increase should be more pronounced; also, there is no understandable reason why the matching rate in the jet drops with time whereas the seeding density there is nearly constant - after a slight decrease in the earliest instants.

To better understand these numbers and trends, we characterized more in depth the experimental images, further than just by their seeding density. Upon also considering the intensities of these particles, i.e. of local maxima in the images, we could build the evolution during the run of their Signal-to-Noise Ratio (SNR), by dividing the average value of these maxima by the noise level estimated during pre-processing. This quantity is plotted, again separately for the jet and ambient flow regions, in figure 15. Its evolution appears closely linked to that of the seeding density in both regions, remaining roughly constant for the jet and decreasing steadily in the ambient flow. This is confirmed by the visual impression in figure 10, and is most probably the effect of multiple light reflections by the particles, leading to brighter particle images at high seeding density, and lower intensities for these at low seeding density.




Fig. 16: Left: histogram of local maxima intensities within a sample experimental image. Right: Compared Cumulative Density Functions (CDF) of this image with two synthetic images, corresponding to a monodisperse distribution, and to a polydisperse one chosen to mimic the experimental distribution. In both these figures, intensities are those obtained after the image pre-processing steps described in section 5.2.




Fig. 15: Evolution with snapshot number of Signal-to-Noise Ratio (SNR) in the jet and outer ambient regions.

This figure also shows that for the latest instants in the run, SNR in the ambient region reaches quite critical values (around 3), showing that true particles only weakly stand out from the noise, and also that the corresponding noise level there, around 35%, is higher than the maximal one considered in section 4. Examining also the distribution in the image of these local maxima showed, as seen in figure 16 left for a sample snapshot, that the particle image intensities were strongly polydisperse, and of a different nature than the Gaussian

distributions traditionally included in synthetic tests when dealing with polydispersity.

To verify to what extent each of these features could explain the results of DF-TPTV in this experiment, we built new synthetic test cases mimicking part or all of these characteristics. Firstly, geometrical setup was chosen identical to the experiment, with the same dimensions of the reconstructed volume, as well as positions, optics and sensor sizes of the cameras. Displacement values of the particles were obtained by an interpolation of the 3D vector field yielded by the TomoPIV processing of the first snapshot shown in the previous subsection. Geometrical subdomains corresponding to the jet and the ambient flow were also defined in the same way as in the experiment. Volumetric seeding density was chosen to be different between the jet and ambient regions, being set to yield an image density of 0.055 ppp in the former, and to various values corresponding to the range 0.0054 – 0.055, depending on the test case. This allows to span entirely the spatially inhomogeneous nature and the range of densities observed during the experimental run. Accordingly to what has been observed in the experimental images, the SNR in the ambient region has been adjusted depending on the chosen seeding density. Figure 17 shows both the seeding density characteristics considered, as well as the corresponding SNRs in the jet and ambient regions. Finally, in order to assess specifically the effect of polydispersity, we consider both cases of, on the one hand, particles with identical intensities (monodisperse), and on the other hand a polydisperse distribution chosen




Fig. 17: Evolution of relevant quantities with snapshot number in the synthetic test reproducing the experiment. Left: image seeding density as a function of the volumetric seeding density in the ambient flow (recall that the density in the jet is kept constant), for the entire images (global), and for the zones corresponding to the jet and to the ambient flow. Right: Signal-to-Noise Ratio (see text for definition) in the jet and the ambient flow, for both mono- and polydisperse particle distribution cases.

so as to reproduce the experimental histograms of local intensity maxima. Figure 16 compares the cumulated histograms, or cumulated density functions (CDFs), of both these synthetic images together with that of the experimental image, confirming that the polydisperse distribution is indeed much closer than the monodisperse distribution.

Performance metrics for both the monodisperse and polydisperse synthetic images are shown in figures 18 and 19. The former pertains to particles, and the latter to vectors, with corresponding Recall and Precision values. As was the case for the synthetic datasets of section 4, the same trends as a function of the global ppp can be observed for both particle and vector results. A first interesting observation is that, for all levels of seeding density (and thus SNR), polydispersity leads to a significant drop in Recall, and in a lesser extent in Precision. It nearly appears as a constant offset independent of the ppp for the Recall, and as a discrepancy of increasing magnitude as the ppp grows for the Precision. Regarding the latter quantity, one observes that a value of roughly 0.9 is obtained at the highest seeding density, therefore for a still intermediate SNR, while in the synthetic cases of section 4, this value was closer to 1 even at the highest noise levels considered. We explain this discrepancy by the fact that in the present jet experiment, more ghost vectors are formed in the ambient flow, which is characterized by nearly vanishing displacements. This nearly zero displacement firstly probably leads to ghost particles being reconstructed at

similar locations between the two instants, which can be then matched as their (zero or close to zero) displacement is compatible with the physical displacement there. Focusing now on vector Recall, one observes that the value obtained at $\text{ppp} = 0.055$, which corresponds to experimental conditions at the beginning of the run, is equal to 0.51, while in the experiment it was found equal to 0.43. As in the synthetic tests of section 4, the minimum value obtained, corresponding to the highest noise level, was equal to 0.68, this confirms that strong polydispersity is a first important factor of influence on the method's efficiency. Note that for this considered point we can single out this factor, since at this ppp the image is close to spatially homogeneous in terms of ppp and SNR.

Still focusing on figure 19 (left), another interesting observation can be drawn, pertaining to the choice of how to tune the sparsity parameter S , i.e. the number of particles to be reconstructed. In this figure, the green curves correspond to our choice of restricting the number of particles compared to what can be estimated in the images, as described in section 5.2, whereas the orange curves are the performance metrics obtained when considering all possible particles. One can see that, even if the vector Recall is higher in the latter case, vector Precision is lower, which leads to more wrong vectors in the final fields, confirming our experimental observations, and our parameter choice.

Considering in figure 19 (left) the evolution of vector Recall as the ppp decreases, one observes that this




Fig. 18: Precision and Recall of particle detection for the synthetic cases mimicking the experiment. Left: global performance for the entire domain, comparison between several synthesis parameters. Right: separate analysis in the jet and ambient flow domains for the polydisperse synthesis, the closest to the experiment.




Fig. 19: Precision and Recall of vector detection for the synthetic cases mimicking the experiment. Left: global performance for the entire domain, comparison between several synthesis parameters. Right: separate analysis in the jet and ambient flow domains for the polydisperse synthesis, the closest to the experiment.

quantity first increases, then reaches a maximum at $\text{ppp} \approx 0.028$, and then decreases. When considering the same quantities separately in the jet and ambient regions, it appears that this maximum and decrease are due to image conditions, and presumably the low SNR associated with the low ppp values, in the ambient region. Indeed, for the lowest ppp, a stronger decrease in the vector Recall in this region than at a global level is observed, while this quantity reaches a plateau within the jet. Overall, this global decrease in vector Recall for the lowest ppp is consistent with what is observed

in the experiments, although in the present synthetic cases, it does not reach lower value than at the highest ppp, as is the case in the experiments. However, the trend is strongly different than what was observed in the more idealistic synthetic cases of section 4.

As we now show in figure 20, the exact explanation for it is not directly a very low SNR value in the images in a global sense, but rather the coexistence in the images of two regions of significantly different ppp and SNR. This figure compares the already shown synthetic case closest to the experiment (green curves) with two




Fig. 20: Precision and Recall of particle (left) and vector (right) detections for several synthetic cases, highlighting the role of spatial inhomogeneity. The green curves are identical to that of figures 18 and 19 and are the closest to the experiment, exhibiting all features considered.

other situations. One of them (red curves) differs only by its homogeneous character, i.e. the jet and ambient regions have identical ppp and SNR, the latter being chosen equal as that of the ambient region, with thus an SNR still varying as a function of ppp; as a result of this choice, these curves can be extended towards lower ppp values compared to the green ones. Comparison between these two curves shows that, in the homogeneous case, vector Recall in particular does experience a maximum and then a decrease as the ppp (and SNR) decrease, but for much lower values of the seeding density, around $\text{ppp} \approx 0.01$. To confirm that this decrease is linked to the corresponding very low SNR reached at these densities, we show what would be obtained with the same homogeneous distribution of particles, but at a higher SNR of 5, corresponding to the beginning of the experimental run (orange curves). Indeed, in this more favorable case, the maximum and drop disappear, and vector Recall reaches nearly maximum value of 1 for the lowest ppp considered. We thus see that, while very low global SNR values lead to a drop of performance, the one which is observed when images exhibit two regions characterized by significantly different levels of ppp and SNR is more influential.

5.5 Statistical results

For further performance assessment, we now consider mean and fluctuating velocities in the jet obtained by statistical averaging, which we compare to the same quantities yielded by planar PIV and TomoPIV. We

decompose each velocity component as, for instance on the axial one, $v = V + v'$, where V denotes the mean and v' the instantaneous fluctuation.

5.5.1 Runs and averaging characteristics

As 3D techniques and planar PIV operate at different optimal seeding densities, a dedicated run at higher seeding density was performed in order to obtain the reference mean flow with planar PIV. We thus here present the compared characteristics of the common run used for 3D-PTV and TomoPIV, and of that performed with for planar PIV, as well as the respective methods for obtaining statistics, and their spatial resolution. Table 1 sums up associated relevant quantities.

For DF-TPTV, which yields scattered vector data, we resort to bin averaging, as traditionally done in PTV methods (Agüera et al., 2016; Jux et al., 2018; Kähler et al., 2016). In the present jet flow context, we choose a specific form of bins in order to increase the number of samples. We exploit the assumption of statistical axisymmetry to define them as annuli of increasing radius in a cross-sectional plane. More precisely, introducing the radius $r = \sqrt{x^2 + z^2}$, a bin centered at radius r , with radial resolution $BS_r(r)$ and streamwise resolution BS_y , is defined as the volume:

$$\left[r - \frac{BS_r(r)}{2}; r + \frac{BS_r(r)}{2} \right] \times \left[\frac{y}{D} - \frac{BS_y}{2}; \frac{y}{D} + \frac{BS_y}{2} \right]$$

Another specificity is that, as depicted in figure 21 (right), these bins are not all strictly annuli, but rather

Table 1: Characteristics of mean flows yielded by DF-TPTV, planar PIV and Tomo-PIV: spatial resolution (interrogation window/volume or bin size) and number of samples available for averaging. For Tomo-PIV, interrogation volumes with Gaussian weighting have been used, as specified by (GW) and detailed in section 5.3. Planar PIV results correspond to a separate run performed at higher seeding density.

Parameter	DF-TPTV	Planar PIV	Tomo-PIV
Bin/IW/IV size in the radial / horizontal (z, x) direction (D)	From 0.09 ($r = 0$) to 0.01 ($r \geq 0.4$)	0.014	0.15 (GW)
Bin/IW/IV size in the vertical (y) direction (D)	0.084	0.084	0.28 (GW)
Number of snapshots	1044	3000	1044
Number of samples for averaging	From $\approx 2200 - 2300$ ($r = 0$) to $\approx 1000 - 1200$ ($r \geq 0.4$)	3000	1044




Fig. 21: Geometrical definition of bins used for statistical averaging of the DF-TPTV results. Left: bin radial resolution, $BS_r(r)$ (expressed in jet diameter units), as a function of r/D . Right: Layout of in a cross-sectional plane. Bins, which have been subsampled for clarity, are depicted as black rings whose thickness is equal to their radial resolution $BS_r(r)$. The red circle denotes the jet nozzle.




Fig. 22: Intensity partition of reconstructed particles in cross-sectional plane $y/D = 1.4$, averaged over square bins of $0.025D \times 0.025D$ size. The red circle denotes the jet nozzle.

truncated annuli for the majority of them, as vectors outside of the range $-0.61 \leq z \leq 0.40$ have been excluded. This choice is justified by the repartition of light intensity in the volume, which was slightly asymmetrical with respect to the jet centre. As figure 22 shows, intensity drops quite rapidly for $z/D > 0.4$, which leads to spurious end effects in the motion estimation there within DF-TPTV. This was observed in turn to yield less reliable results, in spite of the rejection (which was also observed to exclude more samples); indeed, values of the fluctuating velocity were observed to be spuriously higher when including this region in the estimation.

In order to gather enough samples in each bin for convergence, while introducing as little spatial smoothing as possible, we choose $BS_r(r)$ to decrease from $0.09D$ for $r/D = 0$, to $0.01D$ for $r/D \geq 0.4$, as seen in figure 21 (left). Note that, although BS_r decreases with r in the jet core and then stabilizes in the shear layer and the outer flow, bins are nearly of same volumetric extent in the whole field, as a result of the radial geometry. Value $BS_r(r > 0.4) = 0.1$ has been observed to be the smallest reachable, i.e. preserving enough samples for a satisfactory level of convergence. Also, we observed that the value of maximum RMS velocity in the shear layer is nearly insensitive to moderate variations of BS_r around this value. Also, in order to allow fine sampling within the shear layer, we consider a high overlap between the bins, equal to at least 85%. The streamwise resolution BS_y is set to $0.084D$, which corresponds to the interrogation window (IW) size of the planar PIV processing (see below).

For each bin, we compute the mean and variance of each velocity component as (here exemplified on the streamwise component):

$$\overline{v'^2} = \frac{1}{M-1} \sum_{t=0}^N \sum_{i=1}^{N_t} (v(t)_i - V)^2 \quad (7)$$

where

$$V = \frac{1}{M} \sum_{t=0}^N \sum_{i=0}^{N_t} v(t)_i \text{ and } M = \sum_{t=0}^N N_t \quad (8)$$

where $v(t)_i$ denotes the velocity of the i^{th} vector measured in the considered bin at time t , N_t is the total number of vectors in the bin at time t and N the number of snapshots.

For consistency, images acquired in the run for planar PIV are processed with rectangular, top-hat IWs of 31×5 pixels ($0.084D \times 0.014D$, respectively in the y and z directions). With such a parameter, the IW size is as close as possible to the bin size of DF-TPTV, although resolution in z is slightly higher. For this planar PIV

run, 3000 snapshots were acquired, in order to reach full convergence of mean and fluctuating velocities.

Finally, similar to the instantaneous results shown in section 5.3, we also report here, for comparison, statistical results stemming from a TomoPIV processing of the same image data as for DF-TPTV, leading to 1044 flow snapshots. As mentioned in section 5.3, rectangular IVs have been used, with a size along y twice bigger as that along x and z , to account for the flow anisotropy. Also, the overall dimension has been chosen so as to guarantee a minimum number of tracers throughout the run, especially at the latest instants where the outer seeding is low. As a Gaussian weighting is used within the IVs, the standard deviation of which being equal to one fourth of the size in each direction, the effective spatial resolution is expected to be lower than the mere IV size, as recalled in table 1. Still, this resolution will be significantly larger than that obtained with DF-TPTV or with planar PIV.

5.5.2 Compared velocity profiles in the jet near field

To compare mean and RMS velocities yielded by DF-TPTV, TomoPIV and planar PIV, we restrict to the jet near field, here $y/D \leq 2$. As a matter of fact, in this zone the turbulence rate remains moderate and can be estimated with a satisfactory accuracy with DF-TPTV, which is characterized by the lowest number of samples per bin.

Figure 23 shows the number of samples obtained by DF-TPTV in each bin, as well as the percentage of vectors rejected, as a function of r/D , at $y/D = 1.40$. As a result of the difference in average seeding density between the jet and the outer flow (see also section 5.1), up to maximum 2350 samples per bin are obtained in the jet core, progressively decreasing in the outer shear layer and the ambient flow, here to 1100 – 1200. Rejection also has a slightly different behaviour in the jet core and in the external flow. Around 3.5 to 4% of the initial vectors are discarded in the core, while the average rejection rate in the outer flow is rather of the order of 2.5%, with a slight increase in the outermost region. This is again consistent with the difference in seeding between these two zones, which translates immediately in an increase in the average measurement error (see in particular figures 5 (right) and 8), therefore inducing more rejection where seeding is denser. Similar trends and orders of magnitude are observed at all other stations in y/D shown hereafter.

Figure 24 shows profiles of the mean and fluctuating streamwise velocity in the radial or horizontal direction, depending on the technique. As the volumetric and planar measurements have been performed with differ-

ent calibrations, slight discrepancies in their respective frame of reference have been detected. To compensate for these and ease comparison, a shift in abscissa has been performed on the curves. Its value has been determined so that the value of the mean velocity at the centre of the shear layer matches for both measurements. In other words, denoting $r_{1/2}$ (resp. $z_{1/2}$) the position such that V is equal to $V(0)/2$, this shift has been chosen such that $r_{1/2} = z_{1/2}$. Comparing first DF-TPTV with planar PIV, consistently with the choice for bin or IW size, it is observed that the mean velocity profiles match almost perfectly. Only very minor discrepancies appear, with the profiles of planar PIV characterized by slightly larger shear layers, this being in line with the slightly larger horizontal resolution (see table 1). On the other hand, as expected due to the necessarily large IVs used, spatial filtering is obvious on the TomoPIV results, the shear layer appearing significantly larger at the most upstream stations, the discrepancy with DF-TPTV and planar PIV decreasing progressively with y and the natural widening of the shear layer.

Regarding the fluctuating streamwise velocity or turbulence rate, curves for DF-TPTV and planar PIV are observed either to collapse, or to differ by up to roughly 0.02. On the other hand, TomoPIV results appear logically even more biased by the spatial filtering due to the large IVs, with a thickening of the peak of fluctuating velocity in the shear layer at the most upstream stations, evolving towards a global underestimation of this quantity as one progresses downstream. Focusing for the rest of this analysis on the more interesting comparison between DF-TPTV and planar PIV, their maximum discrepancies are observed in a localized zone in the outer shear layer ($r/D, z/D \approx 0.6 - 0.7$ for $y/D = 1.7$, and with a lower magnitude at $y/D = 1.4$ and 2), and also within the shear layer for $y/D = 0.8$ and 1.1. It should be noted that they cannot be ascribed to the partially converged character of $\sqrt{v'^2}/V_j$ estimated by DF-TPTV. Indeed, monitoring of its convergence with respect to the number of samples yields an estimate of the statistical uncertainty of roughly 0.001 – 0.002. Turning now to the detailed explanation of these curves, a first, most directly understandable observation is the level of turbulence rate in the jet core ($r/D, z/D < 0.3 - 0.4$) and in the outer flow ($r/D, z/D > 0.8$) for the most upstream locations, say up to $y/D = 1.1$. Indeed, flow in these zones should be strictly laminar and therefore $\sqrt{v'^2}/V_j$ should vanish. Levels observed there are thus measurement noise and can be directly compared, in the case of DF-TPTV, to the levels of the RMS error on velocity from synthetic data, RMS_v , presented in section 4 (see figures 5 (right) and 8). One observes, in the present experiment, values

close to 0.025 in the core, and 0.02 in the outer flow, which correspond to 0.13 and 0.1 voxel, respectively. The higher value in the jet core is again consistent with the denser average seeding there; on a more general point of view, the fact that these values are quite low (especially when comparing to values of RMS_v^* for situations with noise in figure 8), shows the effectiveness of the outlier rejection scheme chosen here, and overall, the very good efficiency and robustness of the DF-TPTV approach itself. Finally, it is interesting to observe that DF-TPTV exhibits a noise level which is very similar to that of planar PIV processed with FOLKIPIV, as long as the bin size and the IW size are taken equal. In the curves, these noise levels are seen to be nearly equal, or slightly higher for DF-TPTV (note that the IW size for planar PIV is in fact very slightly larger than the bin size, as shown in table 1). This fact also helps to understand why the agreement between results yielded by the two methods is very close in a general way. Possible reasons for zones with discrepancies could be a partial lack of axial symmetry of the jet, making the estimation by bin averaging of DF-TPTV less accurate, or a different sensitivity of DF-TPTV and of planar PIV to flow gradients.

6 Conclusion

We proposed here a novel technique for performing 3D PTV from traditional double frame images, termed double frame tomographic-PTV (DF-TPTV). Its main specificity is that it takes advantage of the sparse nature of the tomographic PIV problem. It first produces spiky particle reconstructions located on a voxel grid, leveraging on the PVR model (Champagnat et al., 2014) and the sparsity-based algorithm LocM-CoSaMP (Cornic et al., 2015a). Reconstructed particles are matched using a low resolution predictor yielded by 3D correlation before being accurately localized through a global optimization procedure. Good performances have been obtained over a large range of seeding densities ($ppp \in (0, 0.08)$) on traditionally generated synthetic images, using a DNS data of a turbulent channel flow. DF-TPTV has then been demonstrated on a round air jet experiment at $Re = 4600$, with a global seeding density evolving from 0.06 to 0.03 ppp through time, with different densities and signal-to-noise ratios in the jet and ambient flow on a large period of time in the run. Analysis of the detection performance, together with dedicated supplementary synthetic tests designed to mimic closely this experiment, have shown that strong polydispersity, as well as this coexistence of spatial zones characterized by significantly different levels of signal-to-noise ratio and ppp, are parameters of influence for the

method's performance. In particular, the spatial inhomogeneity has been found of stronger importance than a low but homogeneous signal-to-noise ratio in the images. Overall, the DF-TPTV technique allowed to produce both reliable instantaneous velocity vector fields and accurate ensemble statistics, upon introducing an additional outlier rejection step. Statistical results have been obtained by a specific bin averaging process exploiting the jet average axisymmetry, and were found in excellent agreement with reference measurements by planar PIV at comparable spatial resolution.

Overall, it is worthwhile emphasizing that DF-TPTV is a particle tracking technique that relies on the same amount of information as TomoPIV, not only in terms of hardware (double frame acquisition, either at low or high frequency), but also of seeding densities. Indeed, it has been reported to yield reliable results for up to 0.08 ppp in a synthetic context, and to 0.06 ppp in the presented experiments; in the latter experimental case, due to the quite challenging conditions mentioned above, optimal performance has been obtained by lowering the number of particles to reconstruct, leading to vector fields corresponding to an effective ppp of around 0.045. A consequence is that, within a given experiment, there is no need to perform separate runs at lower densities in order to perform DF-TPTV as well as TomoPIV processing.

A possible direct perspective to this work could be to further improve the method's robustness to low signal-to-noise ratios, so as to exploit the maximum from a given volumetric illumination, often less intense on the volume edges, as was the case in the present experiment. Future research paths will also target a more drastic increase in the capability of DF-TPTV to characterize complex turbulent flows, on an instantaneous point of view. This will be tackled by proposing new data assimilation schemes, performing physically-sound instantaneous interpolation between the obtained scattered vectors.

References

- Agüera N, Cafiero G, Astarita T, Discetti S (2016) Ensemble 3D PTV for high resolution turbulent statistics. *Measurement Science and Technology* 27(12):124011
- Atkinson C, Soria J (2009) An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Experiments in Fluids* 47:553–568
- Champagnat F, Plyer A, Le Besnerais G, Leclaire B, Davoust S, Le Sant Y (2011) Fast and accurate PIV computation using highly parallel iterative correlation maximization. *Experiments in Fluids* 50:1169–1182
- Champagnat F, Cornic P, Cheminet A, Leclaire B, Le Besnerais G, Plyer A (2014) Tomographic PIV: particles versus blobs. *Measurement Science and Technology* 25(8):084002
- Cheminet A, Leclaire B, Champagnat F, Plyer A, Yeghavian R, Le Besnerais G (2014) Accuracy assessment of a Lucas-Kanade based correlation method for 3D PIV. In: 17th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal
- Cornic P, Champagnat F, Plyer A, Leclaire B, Cheminet A, Le Besnerais G (2014) Tomo-PTV with sparse tomographic reconstruction and optical flow. In: 17th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal
- Cornic P, Champagnat F, Cheminet A, Leclaire B, Le Besnerais G (2015a) Fast and efficient particle reconstruction on a 3D grid using sparsity. *Experiments in Fluids* 56(3):62
- Cornic P, Leclaire B, Cheminet A, Champagnat F, Losfeld G, Illoul C, Le Sant Y, Plyer A, Le Besnerais G (2015b) Two time steps tomo-PTV with sparse tomographic reconstruction versus tomo-PIV. In: 11th International Symposium on PIV - PIV15, Santa Barbara, USA
- Cornic P, Illoul C, Cheminet A, Le Besnerais G, Champagnat F, Le Sant Y, Leclaire B (2016) Another look at volume self-calibration: calibration and self-calibration within a pinhole model of scheinpflug cameras. *Measurement Science and Technology* 27(9)
- Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006) Tomographic particle image velocimetry. *Experiments in fluids* 41(6):933–947
- Fuchs T, Hain R, Kähler CJ (2016) Double-frame 3D-PTV using a tomographic predictor. *Experiments in Fluids* 57(11):174
- Graham J, Kanov K, Yang XIA, Lee M, Malaya N, Lalescu CC, Burns R, Eyink G, Szalay A, Moser RD, Meneveau C (2016) A web services accessible database of turbulent channel flow and its use for testing a new integral wall model for LES. *Journal of Turbulence* 17(2):181–215
- Griffin J, Schultz T, Holman R, Ukeiley LS, Cattafesta L (2010) Application of multivariate outlier detection to fluid velocity measurements. *Experiments in Fluids* 49(4):305–317
- Jux C, Sciacchitano A, Schneiders JFG, Scarano F (2018) Robotic volumetric PIV of a full-scale cyclist. *Experiments in Fluids* 59(4):74
- Kähler C, Astarita T, Vlachos P, Sakakibara J, Hain R, Discetti S, La Foy R, Cierpka C (2016) Main results

- of the 4th international PIV challenge. *Experiments in Fluids* 57(6):97
- Li Y, Perlman E, Wan M, Yang Y, Meneveau C, Burns R, Chen S, Szalay A, Eyink G (2008) A public turbulence database cluster and applications to study lagrangian evolution of velocity increments in turbulence. *Journal of Turbulence* 9:N31
- Maas HG, Gruen A, Papantoniou D (1993) Particle tracking velocimetry in three-dimensional flows. Part 1. Photogrammetric determination of particle coordinates. *Experiments in Fluids* 15:133–146
- Malik NA, Dracos T, Papantoniou DA (1993) Particle tracking velocimetry in three-dimensional flows. *Experiments in Fluids* 15(4):279–294
- Needell D, Tropp JA (2009) CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied and computational harmonic analysis* 26:301–321
- Nishino K, Kasagi N, Hirata M (1989) Three-dimensional particle tracking velocimetry based on automated digital image processing. *Journal of fluids engineering* 111(4):384–391
- Nocedal J, Wright S (2006) *Numerical Optimization (Second Edition)*. Springer
- Novara M, Schanz D, Reuther N, Kähler CJ, Schröder A (2016) Lagrangian 3D particle tracking in high-speed flows: Shake-the-box for multi-pulse systems. *Experiments in Fluids* 57(8):128
- Park J, Kihm K (2006) Three-dimensional micro-ptv using deconvolution microscopy. *Experiments in Fluids* 40(3):491
- Scarano F (2012) Tomographic PIV: principles and practice. *Measurement Science and Technology* 24(1):012001
- Schanz D, Gesemann S, Schröder A (2016) Shake-the-box: Lagrangian particle tracking at high particle image densities. *Exp Fluids* pp 57–70
- Schröder A, Geisler R, Sieverling A, Wieneke B, Henning A, Elsinga G, Scarano F, C P (2009) Lagrangian aspects of coherent structures in a turbulent boundary layer flow using TR-tomo PIV and PTV. In: 8th International Symposium on Particle Image Velocimetry - PIV09, Melbourne
- Schröder A, Geisler R, Staack K, Elsinga G, Scarano F, Wieneke B, Henning A, Poelma C, Westerweel J (2011) Eulerian and Lagrangian views of a turbulent boundary layer flow using time-resolved tomographic PIV. *Experiments in Fluids* 50(4):1071–1091
- Warner S, Smith B (2014) Autocorrelation-based estimate of particle image density for diffraction limited particle images. *Measurement Science and Technology* 25(6):065201
- Wieneke B (2013) Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology* 24(2):024008
- Willert C, Gharib M (1992) Three-dimensional particle imaging with a single camera. *Experiments in Fluids* 12(6):353–358




Fig. 23: Percentage of rejected vectors per bin (left), and number of vectors per bin after rejection (right) as a function of r/D , in cross-sectional plane $y/D = 1.4$, yielded by the DF-TPTV method.




Fig. 24: Streamwise mean and fluctuating velocity profiles, along a jet radius, for various cross-sectional locations y/D in the near-field. Comparison between DF-TPTV, which uses bin averaging with statistical axisymmetry assumption (r/D abscissa), together with planar PIV (z/D direction) and Tomo-PIV (y/D direction). To ease readability and comparison, a horizontal coordinate shift has been applied, see the text for more details.