

Lagrangian Particle Tracking: a link between localisation error and fraction of missed particles

P. Cornic^{1*}, F. Champagnat¹, B. Leclaire²

¹ ONERA – The French Aerospace Lab, F-91761 Palaiseau, France

² ONERA - The French Aerospace Lab, F-92190 Meudon, France

* philippe.cornic@onera.fr

Abstract

This paper aims at analysing the behaviour of particle localisation error in 3D Lagrangian Particle Tracking (LPT) techniques, with a particular emphasis on general properties, independent of a specific algorithm. Based on the hypothesis that in LPT algorithms, errors on the image formation models are solely due to random noise, we show/prove the existence of a best achievable root mean square error (RMSE) on particle localisation, that, for a setup at a given seeding density, depends only on the noise level. We provide a procedure to estimate this lower bound, and show that it can only be reached if there are no missed detections; further on, we establish a link between localisation error and fraction of missed particles. We illustrate the consistency of this model on the results of the recent First Challenge on LPT (see ISPIV21 papers by Leclaire et al. and Sciacchitano et al.)

1 Introduction

Lagrangian Particle Tracking (LPT) is gradually becoming the new standard for high seeding density 3D velocity measurements, both for time-resolved series (see e.g. Schanz et al. (2016)) and double-frame data (then also referred to as 3D PTV, see e.g. Fuchs et al. (2016), Yang et al. (2019), Lasinger et al. (2019), Cornic et al. (2020)). Whereas 3D PIV, relying on cross-correlation, induces important spatial averaging, and thus smoothing of corresponding turbulent scales, residual error on LPT vectors is much smaller, of the order of a fraction of the particle image size. To cope with high seeding densities, where algorithmic performance is usually observed to drop, 3D PTV algorithms are often faced with a dilemma in the choice of the operating point, between retrieving as many vectors as possible – at the risk of getting ghost vectors a.k.a. false positive (FP), and emphasising reliability i.e. obtaining as few FP as possible – at the cost of a large number of missed true vectors a.k.a. false negative (FN). Recent literature shows that no clear consensus yet exists, the two possible choices being made depending on the algorithms.

Our purpose is twofold. The first one is to show that there is a best achievable root mean square error (RMSE) on particle localisation, depending, for a given flow and camera setup, essentially on the image noise and the seeding density. We provide a way to estimate it. The second one is to show that the choice consisting in reducing as much as possible the number of ghosts, leading potentially to less true vectors retrieved, might have some unexpected consequences on the particles localisation error, and consequently on the RMSE on the velocity vectors. More precisely, we find that a low velocity RMSE is achievable only provided that very few vectors are missed; a coupling which, to the best of our knowledge, has not yet been documented in the literature. While we demonstrate this result in the framework of the recently introduced DF-TPTV algorithm (Cornic et al., 2020), performance metrics obtained match with results of the most advanced methods at the 1st Challenge on LPT (Sciacchitano et al., 2020), (Leclaire et al., 2021) thus confirming its global scope.

The paper is organised as follows. Section 2 briefly describes the synthetic data generation, intended to reflect as much as possible the one of the 1st Challenge on LPT (Leclaire et al., 2020). Section 3 briefly states the common principles of tomographic LPT algorithms and introduces the Best Achievable Root mean square error on particle localisation (BAR). It highlights the requirement of a suitable preconditioning technique to perform gradient-based optimisation to refine all particle position and intensity. Section 4 presents the RMSE on particle localisation model depending for a given setup and seeding density on only




Figure 1: (a): cameras setup. (b): crop 256×256 of a synthetic image (ppp=0.12).

two parameters and deals with the consistency of the model predictions with the results of the 1st Challenge on LPT.

2 Synthetic data generation

The ground truth of 1st LPT challenge data is strictly dedicated to the evaluation of the LPT challenge and forthcoming submissions to the benchmark site (Leclaire et al., 2021). It is not available for other purposes, including this work. The starting point of this work being to get a better understanding of the LPT challenge results, we built a dedicated synthetic two-pulse dataset that mimics as much as possible the images available to the participants. We have used the particles trajectories of the publicly available dataset of the 1st Challenge on Data Assimilation https://w3.onera.fr/first_lpt_and_da_challenge/.

The physical situation of this case is a wall-bounded turbulent cylinder wake flow (resolved by LES), with a momentum thickness Reynolds number of the boundary layer of around 4,500. We used the same optical setup - four virtual cameras are set along an arc of circle as shown figure 1(a) - and image settings as for the LPT case of the challenge, as disclosed during the presentation of the challenge results during the 3rd CFD for PIV workshop (Leclaire et al., 2020). In particular in terms of particles' polydispersity with uniform intensity distribution in $[500, 1700]$, diffraction limited particle image, photon noise (20 counts mean and 8 counts standard deviation), and shot noise. We named this dataset HMDB standing for home made data base.

3 Determination of best achievable root mean square error on particle localisation

Most state-of-the-art LPT algorithm based on tomographic principles share the same organisation principle as sketched figure 2. An image-based detection task is performed that yields a rough localisation that is refined through an optimisation process. Recorded images are subtracted with predicted images to form the residue images the process loops on. This loop may include in addition a prediction from previous time steps.

3.1 Particle accurate localisation

Many different tomo LPT algorithms, Schanz et al. (2016), Yang et al. (2019) and Cornic et al. (2020) resort to minimising the discrepancy between recorded images and predicted images given by a particle imaging model depicted in figure 3(a). This is also the case for IPR Wieneke (2013). This boils down to minimising a criterion which is a function of 3D coordinates X_p and intensities E_p of each particle p . We note: $\mathbf{X}, \mathbf{E} = \{X_p, E_p; p = 1 \dots P\}$ the set of all the particle 3D coordinates and intensities, P is the number




Figure 2: schematic of tomo LPT algorithms. A rough detection followed by an optimisation step looping on residue images.




Figure 3: (a): imaging model. (b): true location and estimated location due to the noise.

of particles. The criterion writes:

$$\mathcal{C}(\mathbf{X}, \mathbf{E}) = \sum_c \sum_x \left\| I_c(x) - \sum_p E_p h_c(x - \mathcal{F}_c(X_p)) \right\|^2, \quad (1)$$

where \mathcal{F}_c is the projection function for camera c , h_c denotes the point spread function and x is a pixel coordinate. We assume that \mathcal{F}_c and h_c have been calibrated with enough accuracy so that they can be considered as perfectly known.

3.2 Best achievable root mean square error on particles localisation

On synthetic experiments, the recorded images I_c are computed from X_p^* and E_p^* , the true location and intensity of particle p corrupted with additive noise n :

$$I_c(x) = \sum_p E_p^* h_c(x - \mathcal{F}_c(X_p^*)) + n(x) \quad (2)$$

Let $\widehat{\mathbf{X}}, \widehat{\mathbf{E}}$ denote the minimizer of \mathcal{C} . Due to the noise, the value of the criterion at true locations and intensities, $\mathcal{C}(\mathbf{X}^*, \mathbf{E}^*)$ is not zero and furthermore $(\mathbf{X}^*, \mathbf{E}^*)$ is not a minimizer of \mathcal{C} . Due to noise, $\widehat{\mathbf{X}}$ drifts away from \mathbf{X}^* leading to discrepancies between $\widehat{\mathbf{X}}$ and \mathbf{X}^* . For low noise level, the perturbation $\Delta \mathbf{X}(n)$ is a linear function of the noise realisation n (Fessler, 1996). A similar result holds for $\widehat{\mathbf{E}}$.

The Best Achievable Root mean square error on particle localisation (BAR) is the expected value of the norm $\|\Delta \mathbf{X}(n)\|$ for the distribution \mathcal{N} of n . It writes:

$$BAR = \mathbb{E}_{\mathcal{N}} \left[\frac{\|\Delta \mathbf{X}(n)\|}{\sqrt{P}} \right]. \quad (3)$$

Here it is important to note that the BAR is defined as a “true” RMSE. It writes : $\sqrt{\frac{1}{P} \sum_{p=1}^P \|\hat{X}_p - X_p^*\|^2}$

and differs from a mean positional error: $\frac{1}{P} \sum_{p=1}^P \sqrt{\|\hat{X}_p - X_p^*\|^2}$, whose value is lower. Estimating the BAR would require to generate many images with different noise realisations n_i according to the law of \mathcal{N} and then average the obtained $\|\Delta \mathbf{X}(n_i)\|$. As typical residual image noise (i.e. once possible light reflections have been removed by background subtraction) is uncorrelated in space, and as our synthetic experiments involve thousands of particles already at low seeding density, estimating the BAR can be achieved on the basis of a single snapshot, which will lead enough particle position realisations to enable convergence of the expected value:

$$\mathbb{E}_{\mathcal{N}} \left[\frac{\|\Delta \mathbf{X}(n)\|}{\sqrt{P}} \right] \approx \frac{\|\Delta \mathbf{X}(n = n_{l_c})\|}{\sqrt{P}}, \quad (4)$$

where $n = n_{l_c}$ stands for the noise realisation of the recorded image.

3.3 optimisation

Computing $\Delta \mathbf{X}$ boils down to finding the minimizer of C close to $(\mathbf{X}^*, \mathbf{E}^*)$. A straightforward way to do so is to use an optimisation algorithm suitable for non linear high dimension problems – here four times the number of particles because we also must account for the particles intensity – starting from the first guess $(\mathbf{X}^*, \mathbf{E}^*)$. With tens of thousand of particles, computing the Hessian matrix of C or its inverse is impractical and Newton’s methods are definitely out of reach. Even quasi-Newton method like BFGS algorithm with a computational complexity of $O(16P^2)$ turns out to be impractical and one has to turn to L-BFGS with complexity $O(4P)$ or improved gradient descent like eg. conjugate gradient (CG). But straightforward application of latter techniques fail in practice due ill conditioned criterion C .

3.4 Preconditioning

C suffers from a conditioning problem that originates in the fact that positional gradients are several order of magnitude greater than intensity gradients. When using L-BFGS or CG, it results in intensity values that virtually do not change through the iterations and impede the accuracy of localisation. This can be illustrated on a toy example with only one particle and no noise. The first guess has a localisation error on the 3 axes of $[0.74, -0.45, 1]$ voxel (vx), whose size is equal to the backprojection of a pixel ie. a voxel to pixel ratio v/p equal to 1, and an intensity error of 10%. Initial residue images ($I_c - I_{guess}$) are shown figure 4(a). Negative values, appearing dark blue, show how much shifted the guessed particle is.

It may be seen figure 4(b) that BFGS which determines the descent direction by preconditioning the gradient with curvature information is not prone to the conditioning problem and correctly estimates the intensity. This is not the case for L-BFGS: the intensity stay still through the iterations and leads to a localisation error 2.5 times higher as shown figure 4(b).

A good conditioning is when the criterion has more or less the same curvature in all the directions, ie. when its Hessian is close to the identity matrix (up to a scale factor). Let’s note Θ the vector of particles positions and intensities. Preconditioning amounts to applying a change of variable $\Theta' = M\Theta$ such that the condition number of $\tilde{C}(\Theta') = C(M^{-1}\Theta')$ is better than that of $C(\Theta)$. It can easily be shown that: $\nabla^2 \tilde{C}(\Theta') = M^{-t} \nabla^2 C(M^{-1}\Theta') M^{-1}$. Using the Cholesky decomposition: $\nabla^2 C(M^{-1}\Theta') = LL^t$ and setting $M = L^t$ yields: $\nabla^2 \tilde{C}(\Theta') = \mathbb{I}$ (Nocedal and Wright, 2006). Green curve in figure 4(a) shows that using this preconditioning for Θ equal to the first guess, the optimisation through L-BFGS reaches the same performances as BFGS with a good estimation of intensity. But this approach is impractical with tens of thousands of particles since we cannot neither compute/store $\nabla^2 C(\Theta)$ nor its Cholesky decomposition. We have implemented a solution that considers an approximate block diagonal Hessian, this amounts to apply the same change of variable M , as in the single particle case, to the P particles. In practice, we compute a “mean particle Hessian” M by averaging the Hessian over a hundred of randomly chosen particles in the first guess.

3.5 BAR for the home made data base

Figure 5(a) presents the BAR (black solid line) as a function of the ppp (particle per pixel, tightly related to seeding density) for the dataset described section 2. We have verified that the BAR value does not depend

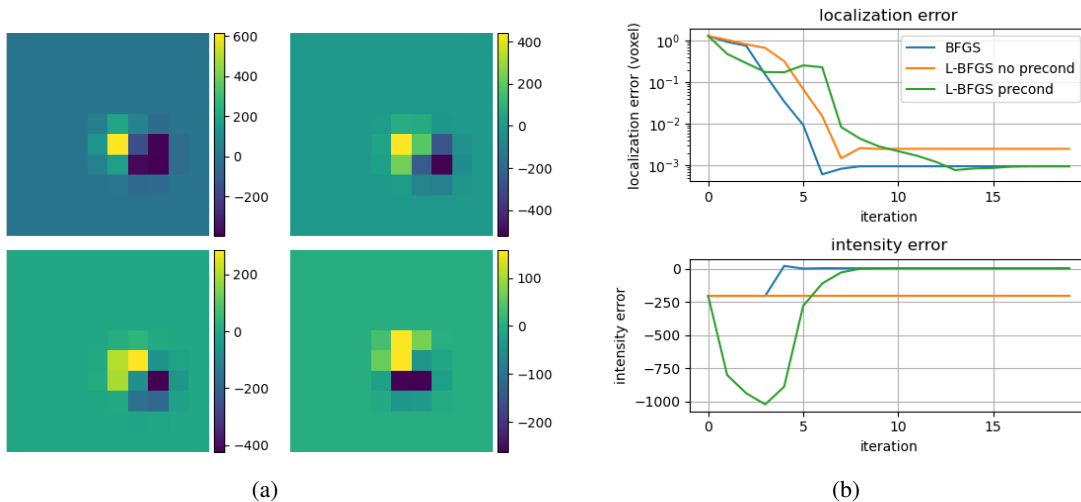


Figure 4: Toy example. (a): initial residue images. It may be seem how much shifted the guessed particle is. (b): localisation (up) and intensity (bottom) error as a functions of optimisation iterations. The voxel (v_x) size is equal to the backprojection of a pixel. Blue: BFGS, orange: L-BFGS without preconditioning, green: L-BFGS with preconditioning.

on the noise realisation. It is also of practical interest to know if the BAR can be reached initialising the optimisation process from an other first guess than \mathbf{X}^* , \mathbf{E}^* mimicking the rough localisation step sketched on figure 2. The solid curves present the final RMSE resulting from an optimisation using the above described preconditioning technique from various first guess with initial RMSE, named ϵ , respectively equal to 0.5, 0.75 and 1 voxel. In this experiment, the various initial RMSE are obtained by initialising the particles on the closest node of a voxel grid as shown in the schematic figure 5(b). The greater the grid step, the greater ϵ . The initial intensities are set by solving in a least squares sense the tomographic system (Cornic et al., 2015) with the particles being on the nodes of the grid.

The dashed coloured curves show the final RMSE without using preconditioning. It may be seen that the greater the ppp and ϵ , the farther the result is from the BAR. Conversely, using pre-conditioning with the same first guesses (see solid coloured curves), it may be seen that the final RMSE virtually does not depend on the first guess. Whatever a reasonable rough initialisation, the BAR can virtually be reached for all ppp provided the preconditioning technique is used.

4 Model for predicting an algorithm’s RMSE on particle position based on its detection performances

DF-TPTV (Cornic et al., 2020) is a recently introduced two frame LPT algorithm relying on tomographic and sparsity principles. It may be seen in the left hand side figure 6(a) that its performances in terms of RMSE get farther from the BAR as the ppp grows. This algorithm was originally designed to yield as few ghosts as possible. Indeed, the Precision figure 6(b) which measures the fraction of true particles among detected ones is almost 1 up to ppp equal 0.08. It indicates that ghosts are not involved in the RMSE soaring. The choice of having few ghosts has for this algorithm the consequence of missing particles as shown by the dropping Recall (fraction of retrieved particles) as the ppp grows. This suggests that RMSE and Recall might be anti-correlated. A qualitative explanation could be as sketched figure 7(a) and 7(b). When the ppp is low, the particles do not significantly overlap. A missed particle (ie. not detected) do not perturb the optimisation of the detected one (red) whose image through the point spread function can be matched in a least square sense with the recorded one (green) as shown figure 7(a). This is no longer the case when the ppp gets higher and particles overlap. The missed particle interferes with the optimisation that results in a compromise solution as displayed on figure 7(b). Accurate location is no longer possible through optimisation. When the particles overlap a missed particle results from the optimisation point of view in a much stronger “noise” than background noise.




Figure 5: (a): black: BAR as a function of the ppp for the home made dataset (uniform intensity distribution in $[500, 1700]$, 20 counts mean and 8 counts standard deviation photon noise); coloured curves: RMSE resulting from various first guess (see text) and optimisation technique. (b): The particles are displaced on the grid node for initialisation.

4.1 Modelling the False Negative effect on RMSE

False Negative (FN) are the missed particles whose percentage is given by $1 - Recall$. An interesting question to deal with is whether the RMSE of DF-TPTV - and more generally of any LPT algorithm relying on tomographic principles - can be approximated by simply optimising \mathcal{C} from a first guess with the same amount of adequately selected FN. Figure 8(a) shows that optimising with the same amount of randomly selected FN from a first guess with $\epsilon = 1$ results, except for low ppp, in overestimating the RMSE.

It is not surprising because, at least for algorithm relying on tomographic principles, there are particles that are more liable than others to be missed. One then might think that a random selection is not appropriate. Indeed, with missed particles being chosen among the least bright, one can see that we get a quite good match between model prediction and DF-TPTV results as shown figure 8(b).

This result is remarkable because it establishes a link between the root mean square error on particle localisation and the percentage of missed particles.

4.2 RMSE on particle localisation model

Section 3.5 has shown that without missed particles, ϵ the RMSE of the first guess, if reasonable, has almost no influence on the final RMSE. But this no longer holds in presence of missed particles. For ppp ranging from 0.005 to 0.12, figure 9 shows the RMSE as a function of FN percentage, missed particles being chosen among the least bright. To get this bunch of curves, we optimised criterion \mathcal{C} from initial first guesses with a given rate of FN among the least bright and various initial RMSE. The box plot spans $\epsilon = 0$ to $\epsilon = 1$ voxel. The coloured solid curves are for $\epsilon = 0.5$. Given a certain amount of FN, the foot of the box plot is the best one can do in terms of RMSE. It may be seen that for the lowest ppp in blue, with no significant overlap, FN rate and ϵ have virtually no influence on the RMSE. It is consistent with the quantitative explanation given section 4. The greater the ppp, the greater the influence of the FN rate and ϵ . The later being visible on the extent of the box plot. The BAR can be retrieved from the foot of the box plot for a FN percentage equal zero.

4.3 Consistency with 1st LPT challenge results

This RMSE on particle localisation error model is very simple. It involves, for a given ppp, only 2 parameters. Originally built to study the performances of DF-TPTV, the question of its scope naturally arises.




Figure 6: (a): DF-TPTV RMSE and BAR as a function of the ppp . (b): DF-TPTV Precision and Recall as a function of the ppp .




Figure 7: Green particles are the recorded one, the detected (red) particle's location has to be optimised. (a): The particles do not overlap. (b): overlapping particles, the optimisation results in a compromise solution.

To answer this question we asked the participants' and organisers' permission to compare the results of the 1st LPT challenge (Sciacchitano et al., 2020) with the model predictions. It should be stressed that even if they share the same principles, the algorithms of each team are different and we do not know exactly how they work. It should also be highlighted, as already mentioned, that the LPT challenge's ground truth is not available for the BAR and RMSE model. Five teams competed in the two pulse LPT challenge, namely: DLR, LAVISION, IRAE, ETHZ and ONERA. The RMSE prediction given their detection performances is reported for all the teams except ETHZ, for which the model is not directly applicable. DLR and LAVISION teams submitted results for all the proposed ppp (ie. up to 0.16) while IRAE and ONERA stopped at $ppp=0.08$. Figures 10(a), 10(b), 10(c), 10(d) show for the four teams, their RMSE, evaluated with the ground truth of the 1st LPT challenge, the BAR and the model prediction as a box plot given their detection performances yielded using the HMDB. Please note that the y-scale are not the same and depend on the team's performance.

The model predicts quite well the DLR's performance that is very close to the BAR, betraying detections without missed particles. The model's predictions are not too far from LAVISION results, and suggest, if we trust the model, that given their detection performances, slightly better RMSE performance is achievable. In the same manner, IRAE's performance is also captured reasonably well by the model that suggests achieved RMSE could also be slightly better. As far as ONERA is concerned, we used an iterative version of DF-TPTV to compete in the challenge, named I-DF-TPTV v0. The model predicts far better RMSE performance than the achieved one, given the detection performances, suggesting there was something wrong with the optimisation step of our algorithm, that did not make use of preconditioning. Furthermore, the understanding of the key role of FN lead to significant changes in the algorithm structure, now v1. Figure 11 compares




Figure 8: DF-TPTV performance in terms of RMSE (blue) (a): RMSE predicted by optimisation with the same amount of randomly selected FN. (b): RMSE predicted by optimisation with the same amount of FN being chosen among the least bright.

v0 and v1 performances. Note that v0 is evaluated on the challenge data while v1 is evaluated on the home made data base. But in the latter, evaluation and model prediction does not use the same set of particles. It may be seen that v1 performance is now consistent with the model prediction, and much closer to the BAR than it used to be for v0.

5 Conclusions

We have highlighted that for a given synthetic setup there is a Best Achievable Root mean square error on particle localisation (BAR) that depends on image noise and ppp. We have shown that the BAR can virtually be reached starting from a reasonable rough first guess provided there are no missing particles (FN) and that a suitable optimisation algorithm using an appropriate preconditioning technique is used. We have brought up that missing particles, in view of the optimisation involved in the accurate localisation steps, results – at ppp larger than 0.025 – in a much stronger noise than background image noise. For a given ppp, we proposed a model that relies on only 2 parameters: the rate of missed particles among the least bright and the first guess RMSE on particle localisation. One might have noticed that ghosts do not take part in the model. Contrary to FN they are not easy to simulate since intensity balance is also required. But we have indications that the model holds even with a large amount of ghosts associated with FN. Of course, this model is too simple to be perfect, but it matches quite well with 1st LPT challenge results, at least for 3 teams. Although we have demonstrated this model in a two frame context we think that to a large extent it might apply to multi pulses process. Finally, it is well known that synthetic images are a pale reflection of experiments images, but the model suggests that if an accurate localisation is of critical importance, one should set a ppp such that the used LPT algorithm misses as few particles as possible.

Acknowledgements

The authors wish to thank the organisers and participant teams of the 1st LPT Challenge for agreeing on the use of their results. We hope this work will be useful to improve the performances of their methods.

References

- Cornic P, Champagnat F, Cheminet A, Leclaire B, and Le Besnerais G (2015) Fast and efficient particle reconstruction on a 3D grid using sparsity. *Experiments in Fluids* 56:62
- Cornic P, Leclaire B, Champagnat F, Le Besnerais G, Cheminet A, Illoul C, and Losfeld G (2020) Double-frame tomographic PTV at high seeding densities. *Experiments in Fluids* 61:1–24




Figure 9: RMSE on particle localisation as a function of the FN percentage (missed particles being among the least bright) and initial RMSE ϵ of the first guess. A box plot spans $\epsilon = 0$ to $\epsilon = 1$ voxel. Uniform intensity distribution in $[500, 1700]$, 20 counts mean and 8 counts standard deviation photon noise.

Fessler JA (1996) Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Transactions on Image Processing* 5:493–506

Fuchs T, Hain R, and Kähler CJ (2016) Double-frame 3D-PTV using a tomographic predictor. *Experiments in Fluids* 57:174

Lasinger K, Vogel C, Pock T, and Schindler K (2019) 3D fluid flow estimation with integrated particle reconstruction. *International Journal of Computer Vision* pages 1–16

Leclaire B, Mary I, Liauzun C, Péron S, Sciacchitano A, Schröder A, Cornic P, and Champagnat F (2021) First Lagrangian particle tracking and data assimilation challenge: datasets description and planned evolution to an open online benchmark. in *14th International Symposium on Particle Image Velocimetry – ISPIV 2021 August 1–5*

Leclaire B, Mary I, Liauzun C, Péron S, Sciacchitano A, and Schröder A (2020) Lpt and da challenge: datasets generation.. in *In 3rd Workshop on Data Assimilation and CFD Processing for PIV and Lagrangian Particle Tracking*. online meeting


Nocedal J and Wright S (2006) *Numerical optimization*. Springer Science & Business Media

Schanz D, Gesemann S, and Schröder A (2016) Shake-the-box: Lagrangian particle tracking at high particle image densities. *Exp Fluids* pages 57–70


Sciacchitano A, Leclaire B, and Schröder A (2020) Main results of the lpt challenge. in *3rd Workshop on Data Assimilation and CFD Processing for PIV and Lagrangian Particle Tracking*. online meeting

Wieneke B (2013) Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology* 24:024008


Yang Y, Heitz D, and Mémin E (2019) Lagrangian particle image velocimetry. in *13th International Symposium on Particle Image Velocimetry, ISPIV 2019*. Munich, Germany




(a)



(b)



(c)



(d)

Figure 10: BAR, RMSE on particle localisation and model prediction (bar plot) in voxel units. BAR and model prediction are computed using HMDB. RMSE are from the 1st challenge results (a): DLR (b): LAVISION. (c): INRAE (d): ONERA. Note the y-scale are not the same.




Figure 11: BAR, I-DF-TPTV v0 and v1 RMSE as a function of the ppp. Model's prediction for I-DF-TPTV v1.