Tomographic PIV: particles vs blobs

Frédéric Champagnat[1], Philippe Cornic[1], Adam Cheminet[2], Benjamin Leclaire[2] and Guy Le Besnerais[1]

[1] Department of Information Processing and Modelization, ONERA frederic.champagnat@onera.fr

[2] Department of Fundamental and Experimental Aerodyanmics, ONERA

ABSTRACT

We present an alternative approach to tomo-PIV that seeks to recover single voxel particles rather than blobs of extended size. Our approach follows the classical MLOS-SMART method with a weight matrix built with the system's Point Spread Function coefficients. As such an approach requires only a few voxels to explain the image appearance, it yields much more sparse reconstructed volumes. We carefully justify our method for building the weight matrix and show on synthetic PIV images with a large exploration of generating conditions that it always leads to better results than MLOS-SMART.

1. Introduction

In numerous PIV [1] and tomo-PIV experiments, particles have a very small physical size, so that their images on the camera sensors are mostly controlled by the aperture diffraction, and are thus essentially a characteristic of the imaging system. Optimization in plane PIV developments has led to seeking a resulting point spread function (PSF) having a 3×3 pixel square shape, in order to guarantee a good subpixel accuracy in the displacement estimation. During the development of tomo-PIV, the same rule-of-thumb has been maintained so that conventional MART/SMART-based tomo-PIV methods [7][2] have aimed at reconstructing "particle blobs", i.e. voxel aggregates of several voxels width, rather than the actual particles, which would lie in a single voxel (and, usually, be much smaller than the voxel size).

In difficult experimental conditions such as when viewing through curved walls, in the presence of defocusing, or in compressible flows, it is known that the PSF shape can vary dramatically, leading to accuracy losses in the estimation. Recently, this has led Schanz et al. [9] to propose an approach in which these changes are taken into account. However, their algorithms do not fully recover the particle vision, and still reconstruct blobs.

In this paper we introduce a simple and accurate method for recovering more closely the particle's physical reality. The proposed method is based on the image formation model from optics. The main difference between our approach and tomo-PIV is best captured in Figure 1 : a particle whose geometrical image is much smaller than the PSF produces an Airy-like image sketched on the left. The orthodox tomo-PIV (such as introduced by [7]) models this image as the integration of a 3D bell-shaped blob in the volume, along a thin pixel-sized pencil elongated in the direction of the line of sight.
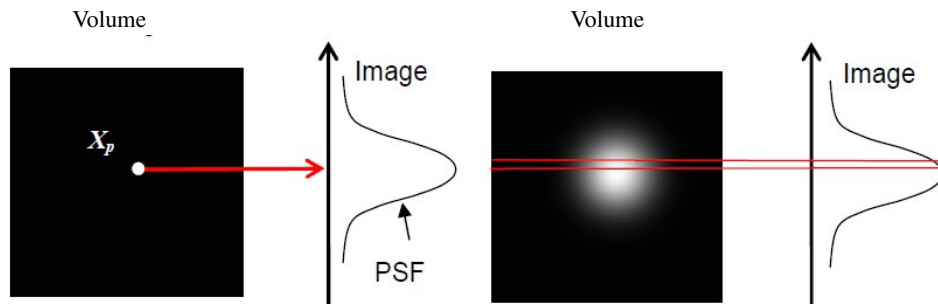

Figure 1: Image formation from a volumic projection: particle model (left) and blob model (right).

In our approach the physical image model is carefully discretised in order to yield a rigorous representation of image data in terms of voxelised particle intensity and a weight matrix. This weight matrix is shown to be built from appropriate samples of the PSF. This result is original and enables to incorporate PSF calibration data in a straightforward manner. Analysis of the different degrees of freedom in the discretisation step allows us to explore an original path in tomo-PIV, that of voxel grids with higher density than the classic voxel-to-pixel ratio $v/p = 1$.

Using the proposed weight matrix a MART/SMART-based reconstruction can be used to yield nearly one-voxel wide particles. This is a first step in a global framework, where the subvoxel accuracy will be preserved by adapted refinement schemes in future developments.

In addition to being more physically sound, we show that this particle approach is indeed more accurate than the conventional methods for a large range of PSF sizes and thus will potentially maintain a good accuracy in the difficult experimental conditions mentioned above (*e.g.,* interfaces, compressible flows, illuminated volumes larger than the cameras depth of fields, leading to the presence of defocused particles in the images). Besides, the method is directly adapted to sparsity-based techniques such as [8, 3, 6]. In a companion communication [6], we show in particular that this enables a dramatic reduction of the problem size, and thus of the processing time.

This paper is outlined as follows. Section 2 recalls the basic image formation model for particle images. Then the main contribution of this paper is introduced in Section 3, that is, a careful discretisation of the image model that justifies the classic expression of image data in terms of an appropriate weight matrix, and a voxelized intensity field. The usefulness of such a representation is demonstrated on a set of numerical experiments whose conditions are described in Section 4, the experiments themselves being gathered in Section 5.

2. Imaging model from optics

In our approach, which we call Particle Volume Reconstruction (PVR), the weighting matrix W is built directly by following the physical model shown in Figure 1 (left). A particle with index p at point $\boldsymbol{X}_p$ in 3D-space has a geometrical image $F(\boldsymbol{X}_p)$ in the focal plane (in practice, the 3D to 2D projection function F results from the camera calibration procedure), and the particle intensity E_p is shared between the impacted pixel and its immediate neighbours according to weights given by the so-called Point Spread Function $h(\boldsymbol{x})$. For the sake of simplicity, we assume hereafter that the PSF is homogeneous over the illuminated volume. This is not a strong assumption, as all the following derivations can be readily adapted to the case of known variation of the PSF (which is the case of astigmatism generated by water/air interface or defocusing). The intensity distribution in the image plane for a number of particles P then reads

$$I(\boldsymbol{x}) = \sum_{p=1}^{P} E_p h(\boldsymbol{x} - F(\boldsymbol{X}_p)), \tag{1}$$

where $\boldsymbol{x} = (x,y)$ denotes any location in the image plane. In practice, camera sensors consist of an array of pixels $\boldsymbol{k} = (k_1,k_2)$, each gathering a single intensity value $I(\boldsymbol{k})$.

For all the numerical experiments in this paper we assume a PSF made of a Gaussian Optical Tranfer Function with standard deviation σ_{psf} averaged on the pixel surface, assuming a 100% fill factor:

$$h(x,y) = \frac{1}{4} \left(\text{erf}\left(\frac{x+0.5}{\sqrt{2}\sigma_{\text{psf}}} \right) - \text{erf}\left(\frac{x-0.5}{\sqrt{2}\sigma_{\text{psf}}} \right) \right) \left(\text{erf}\left(\frac{y+0.5}{\sqrt{2}\sigma_{\text{psf}}} \right) - \text{erf}\left(\frac{y-0.5}{\sqrt{2}\sigma_{\text{psf}}} \right) \right). \tag{2}$$

3. Discretisation of the imaging model

Based on (1), reconstruction of the number of particles, their 3D positions and intensities is a difficult nonlinear problem. A more workable approach is to build an intermediate representation of particles on a 3D grid associated to a voxel space representation, encoding simultaneously the particles location and intensity. An important parameter of such a grid is its scale Δ, *i.e.,* the voxel physical size. Without loss of generality any node of this grid (the center of a voxel) can be indexed with a triplet of signed integers $\boldsymbol{n} = (n_x, n_y, n_z)$ such that the 3D position of the node is $\boldsymbol{n}\Delta$.

We look for a matrix-vector counterpart of (1)

$$\boldsymbol{I} = \text{W}\boldsymbol{E}, \tag{3}$$

where $\boldsymbol{I}$ is a vector that collects image pixels indexed by $\boldsymbol{k}$ introduced in Section 2, $\boldsymbol{E}$ is a vector of intensities located at discrete notes $\boldsymbol{n}\Delta$ and the weight matrix W has entries denoted by $\text{W}_{\boldsymbol{kn}}$.

Discretisation of (1) requires a volumic interpolation of the Projection+PSF function $\boldsymbol{X} \mapsto h(\boldsymbol{x} - F(\boldsymbol{X}))$ based on grid node values near $\boldsymbol{X}$, *i.e.,*

$$h(\boldsymbol{x} - F(\boldsymbol{X})) \approx \sum_{\boldsymbol{n}} h(\boldsymbol{x} - F(\boldsymbol{n}\Delta))\beta(\boldsymbol{X} - \boldsymbol{n}\Delta). \tag{4}$$

β can be any interpolation kernel, but for the sake of particle reconstruction we advocate that this kernel should be as short as possible; thus in this paper we will use a trilinear interpolation.

Using (4) for each particle position $\boldsymbol{X}_p$, and plugging into (1)

$$
\begin{aligned}
I(\boldsymbol{x}) &\approx \sum_{p=1}^{P} E_p \left(\sum_{\boldsymbol{n}} h(\boldsymbol{x} - F(\boldsymbol{n}\Delta))\beta(\boldsymbol{X}_p - \boldsymbol{n}\Delta) \right), \\
&\approx \sum_{\boldsymbol{n}} h(\boldsymbol{x} - F(\boldsymbol{n}\Delta)) \left(\sum_{p=1}^{P} E_p \beta(\boldsymbol{X}_p - \boldsymbol{n}\Delta) \right),
\end{aligned}
\tag{5}
$$

where the last equality is obtained by swapping the summations.

The last expression enables us to introduce a discrete 3D field

$$\widetilde{E}_{\boldsymbol{n}} = \sum_{p=1}^{P} E_p \beta(\boldsymbol{X}_p - \boldsymbol{n}\Delta). \tag{6}$$

The 3D field $\widetilde{E}_{\boldsymbol{n}}$ can be seen as a discrete approximate representation of the original Dirac-like particle field. Conversely, given $\widetilde{E}_{\boldsymbol{n}}$, one can recover the intensity and position of particles with subvoxel accuracy if they do not overlap within this 3D representation.

Using (5) at pixel index $\boldsymbol{k}$ and introducing (6) we get:

$$I(\boldsymbol{k}) \approx \sum_{\boldsymbol{n}} h(\boldsymbol{k} - F(\boldsymbol{n}\Delta))\widetilde{E}_{\boldsymbol{n}}. \tag{7}$$

This is the sought matrix-vector representation (3). The entries of the weight matrix W have a straightforward expression in terms of the PSF and the geometric projection function:

$$W_{\boldsymbol{kn}} = h(\boldsymbol{k} - F(\boldsymbol{n}\Delta)). \tag{8}$$

Such a property is in particular in contrast with the approach of [9], in which the computation of W needs the PSF and a postprocessing which is unnecessary here.

The accuracy of the approximation in (7) depends on the accuracy of the original approximation (4). The accuracy of (4) depends in turn on the v/p ratio. Typically, a PSF $\sigma_{\mathrm{psf}} = 1$ can be accurately sampled using $v/p = 1$, but for a well focused particle with $\sigma_{\mathrm{psf}} = 0.4$ it is not sufficiently accurate, in this case $v/p = 0.5$ is more appropriate. Further examples of this dependence will be given in Section 5.3.

Another way to control the accuracy of approximation (4) is through the shape of the interpolation kernel β and its size in voxel unit. In this respect, nearest neighbour interpolation is too coarse and would require too small a v/p. The trilinear interpolation kernel is better, but higher-order polynomial interpolation would achieve a better order of approximation. But another factor to consider is the shape of volume $\widetilde{E}$: by (6), the kernel controls the smoothness of the discrete particle volume. Selecting the kernel thus depends on the subsequent use of the reconstructed volume. If the reconstructed volume is to be correlated between successive time steps, β should be smooth and spread over at least $3 \times 3 \times 3$ voxels in order to yield low peak locking, this is typically the logic promoted by Schanz *et al.* [9] although there is no explicit β kernel defined in their paper. Conversely, if the emphasis is put on detection, the opposite direction can be taken: choose the kernel with the most compact support so as to yield spiky volumes. This direction is favored by the fact that the system $\boldsymbol{I} = W\boldsymbol{E}$ is most of the time inverted with multiplicative algorithms like MART or SMART that promote sparsity [8]. The kernel of trilinear interpolation is thus a good sparsity/accuracy trade-off. In that case, specific strategies can be implemented to refine the particle position to subvoxel accuracy; to us this should be best done when estimating the displacement, and is thus not addressed here.

4. Simulation setting

We now turn to an evaluation of the efficiency of representation (7) based on synthetic image data and reconstruction algorithm. The goal of this section is to specify the principles of the numerical experiments used for this evaluation.

4.1 Synthetic setup

All our simulations involve four cameras, which are positioned on a single side of the laser volume at the vertices $(\frac{\pm 1}{2}, \frac{\pm 1}{2}, \frac{1}{\sqrt{2}})$ of a square of 1 meter side. They are positioned at 1 meter from the centre of the reconstructed volume, the latter defining the origin $(0, 0, 0)$ and point at it. The pin-hole model is assumed for the cameras (without Scheimpflug adapter for simplicity) and calibration is supposed to be perfectly known. The focal length is 100 mm, thus the magnification factor M is equal to 0.1, and the pixel size is $10\mu m$ with 100% fill factor. Thus a voxel-to-pixel ratio $v/p = 1$ leads to voxels of 0.1 mm side. The image size is set to 512×512; hence, the field of view is fixed for each simulation. The dynamic range of the images is chosen equal to 8 bit.

The laser volume is modelled as a 20 mm thick parallelepiped. The reconstructed volume, also 20 mm thick, is defined as the smallest parallelepiped containing the portion of the illuminated volume seen by all the cameras.

The tracers particles are uniformly distributed in the light sheet volume. The density is controlled by the particle per voxel count (ppv). Horizontal and vertical extension of the sheet are larger than the field of view covered by all the cameras. It is important to notice that this leads to the fact that all the illuminated particles cannot be seen by all cameras, which is systematically the case in real datasets, and is however not often taken into account in synthetic experiments. Our companion papers [6, 5] give further details on this important issue. The scattered light is proportional to the square of the particle physical diameter d_p. Note that Mie's scattering is not taken into account in this study. Consequently, the intensity of a particle depends only on its diameter and on its depth. It is given by $I_0 d_p^2 e^{-z^2/2\sigma_L}$, where I_0 is a constant, and σ_L is the standard deviation of the laser sheet profile modelled as a Gaussian. In this study the dependence in z is very weak as we consider the laser sheet as an almost perfect top hat with $\sigma_L = 0.05$. The particle diameters are supposed small enough (a few microns) to neglect the size of their geometric image Md_p. d_p is drawn in $[mind_p maxd_p]$ according to a Gaussian law with mean md_p and standard deviation $stdd_p$. Unless otherwise specified: $I_0 = 300$, $mind_p = 0.5$, $maxd_p = 2.5$, $md_p = 1.5$, $stdd p = 0.15$.

The images are synthesized according to (1) with a PSF given by (2). Unless otherwise specified, we take $\sigma_{\mathrm{psf}} = 0.6$. With this value, a particle has a 4×4 pixel image pattern. In a given experiment, σ_{psf} does not vary with the volume depth, but a range values of σ_{psf} is considered between 0.4 and 1.4. This is a first-order account of the impact of defocusing. Unless otherwise specified, a Gaussian noise with mean 5 and standard deviation 2 is added to the images. Its amplitude is thus about 10% relative to the maximum particle intensity.

4.2 Reconstruction Algorithm

The Simultaneous Multiplicative Algebraic Reconstruction Technique (SMART) solves linear systems under non negative constraints [4], and is a popular choice for tomographic PIV reconstruction since Atkinson *et al.* [2]. It is considered here because of its parallel structure (compared to MART which is highly sequential). Applying SMART in order to solve (3) yields the following update equation (in logarithmic form in order to emphasize parallelism):

$$\log E^{k+1} = \log E^k + \mu \widetilde{W} \left(\log I - \log W E^k \right) \tag{9}$$

$\widetilde{W}$ is obtained by normalising W over the columns, and μ is a relaxation parameter. In our simulations μ is set to 1 in order to guarantee convergence [4].

Our simulations consider a reference algorithm referred to as tomo-SMART and a new algorithm derived from the present framework referred to as PVR-SMART. Both algorithms start by an MLOS step [2] that aims to reduce the number of voxels that have to be considered in the reconstruction. MLOS also reduces the memory requirement for storage of matrix W, if required. Images are thresholded before MLOS, all pixels with graylevel lower than 4 are set to zero. The MLOS volume is also thresholded: only the voxels greater than 4 are retained for further refinement. Then, 25 iterations of SMART are performed. In both cases, the iteration is defined by (9), but the matrix W is different. tomo-SMART builds W by computing the volume intersected by a cylinder and a sphere centered on the voxel as described in [2]. PVR-SMART computes W using PSF samples and geometric projection functions with (8). Two PVR-SMART versions can be considered, depending on the scale of the voxel selected in the discretisation of the image formation model, see discussion below (7). The usual choice is $v/p = 1$, but for small size PSF we anticipate that a finer discretisation is better, so we also consider $v/p = 0.5$. This latter setting has no counterpart for the tomo-SMART algorithm, so with tomo-SMART we have always $v/p = 1$.

The MLOS, tomo-SMART and PVR-SMART algorithms have been coded in CUDA in order to take advantage of their intrinsic parallelism. Taking advantage of the brute force of GPU, the matrix W is not stored, its entries are computed on the fly for projection and backprojection, and recomputed at the next SMART iteration. Using such an implementation, the processing of four 512×512 images with PVR-SMART with $v/p = 0.5$ yields a reconstructed volume of $1254 \times 1182 \times 401$ voxels in a few minutes.
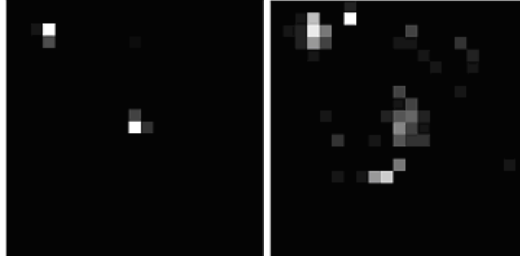

Figure 2: Zoom on a slice of reconstructed volume using PVR-SMART with $v/p = 1$(left) and tomo-SMART (right). $\sigma_{psf} = 0.9$. PVR-SMART clearly shows more spiky results than tomo-SMART.

Figure 2 shows samples extracted from PVR-SMART and tomo-SMART reconstructions for $\sigma_{psf} = 0.9$. The proposed method clearly shows more spiky results than the conventional one, although we do not explicitly use sparsity-enhancing techniques like [8, 3, 6]. MART/SMART techniques are indeed sparsity techniques that are already exploited by using a weight matrix W such as that proposed here.

4.3 Performance metrics

The difference between the particle approach considered here and the classical tomographic approach initially introduced by [7] raises the question of which ground truth to consider to build performance criteria for the algorithms. Indeed, for a same set of physical particles, the former will aim at providing the list of voxels containing a particle together with its corresponding intensity, while the latter will reconstruct volumetric blobs, of approximately the back-projected particle image size, centred around these physical particles. In this tomographic framework in particular, [7] builds the ground truth by expanding locally the physical particles to a 3D gaussian blob, usually of the order of 2-3 voxels size, in the idea of having a volumetric distributions well adapted to the subsequent correlation step yielding the 3D displacement field. A natural quality measure is then the Q criterion, which indicates the degree of correlation between the reconstruction and this ground truth.

In the case of PVR-SMART, the Q criterion cannot be applied directly, as the particle reconstructions generated by PVR-SMART have nearly one voxel size. Therefore, to compute this criterion in the sequel, we expand the PVR-SMART reconstruction using the same method as for building the ground truth, *i.e.,* filter the voxel volume using the same 3D Gaussian kernel.

We also introduce two metrics adapted to measure the detection performance of PVR-SMART *w.r.t.*tomo-SMART. The reconstructed volume is thresholded and only the S brightest voxels are retained as detections. S controls a trade-off between the number of true particles detected and ghosts, the logic of setting this parameter will be developed in the next section. In order to characterize detection performance, we compute *Precision* and *Recall*, two well known measures in pattern recognition and information retrieval. A detection (*i.e.,* a thresholded voxel of the reconstruction) is a True Positive (TP) if it is in the neighbourhood of a true particle. Unless otherwise specified, the neighbourhood is a $3 \times 3 \times 3$ voxel cube centred on the voxel of a true particle. A detection is a False Positive (FP), *i.e.,*

a ghost, if it is not in the neighbourhood of a true particle. A particle is recorded as False Negative (FN) if there is no detection in its neighbourhood. Precision then gives the fraction of true particles among all detected particles, and Recall is defined as the number of true positive divided by the total number of particles, *i.e.,*:

$$\text{Precision} = \frac{\#TP}{\#TP + \#FP} \text{ and Recall} = \frac{\#TP}{\#TP + \#FN}, \tag{10}$$

where # stands for "number of". The best achievable performance is given by Recall = 1 (#FN = 0, every particle is detected) and Precision = 1 (#FP = 0, all the detections are true particles). In practice both measures depend on a common parameter (*e.g.,* the threshold or S) and are correlated (*i.e.,* Precision decreases as Recall increases). An example of Recall and Precision curves is given in Figure 3 in the next section.

Note that, whatever the quality criterion considered, and in all the simulations below, the ground truth will consist of the particles that are seen by all the cameras exclusively, which is consistent with the fact that all reconstructions are initialized with MLOS.

5. Results

The goal of our simulations is to show the increase in performance of PVR-SMART *w.r.t.*tomo-SMART in a variety of experimental conditions, controlled by two parameters N_{ppp} and σ_{psf}. First we explore the ideal case where the PSF size σ_{psf} is known by the experimentalist, but allow this factor to vary. Then we test the robustness of PVR-SMART *w.r.t.* inaccurate knowledge of this parameter.

5.1 Selection of detection threshold

Before reviewing these results we need to fix the parameter S that has an impact on the detection performance metrics Precision and Recall. Figure 3 shows the evolution of the metrics *w.r.t.* a quantity called *normalised sparsity* (NS), which is the ratio of the number S over the total number of particles. Normalised sparsity is merely a monotone function of the detection threshold. N_{ppp} is set to 0.07 and $\sigma_{\text{psf}} = 0.6$. The left part of Figure 3 shows Precision and Recall for each method. One can observe the antagonist behaviour of Precision and Recall *w.r.t.* NS : the objective of getting 1, the best score for both quantities, cannot be achieved, and thus a trade-off has to be chosen. On the other hand the Q factor implicitly balances the two quantities, as shown in the right part of Figure 3: it reaches a maximum between NS equal to 1 and 1.5, which slightly favors Recall *w.r.t.*Precision.
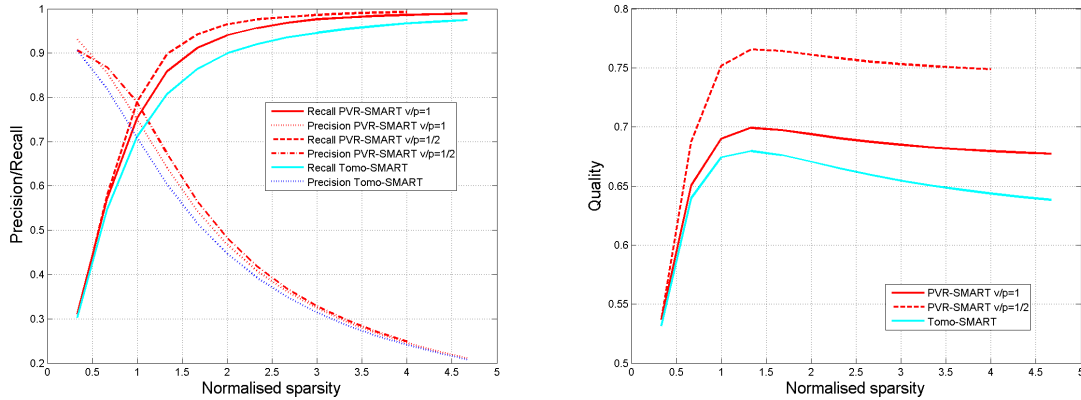

Figure 3: Performance measures *w.r.t. normalised sparsity* (NS), which is the ratio of the number S of brightest voxels retained over the total number of particles. $N_{ppp} = 0.07$ and $\sigma_{\text{psf}} = 0.6$. Left : Precision and Recall have an antagonist monotonic behaviour *w.r.t.* NS. Right : Q has an optimum *w.r.t.* NS.

It is interesting to note that the ranking of the algorithms is identical for Q and Recall whatever the NS, *i.e.,*, by decreasing score value: PVR-SMART $v/p = 0.5$, PVR-SMART $v/p = 1$, and tomo-SMART. In particular, a good discrimination between the algorithm performance is obtained for NS larger than 0.75 and lower than 2.5. Thus any value of NS between 1 and 1.5 can be chosen as reference, we will choose Recall for $NS = 1$ to be a good landmark for assessing relative the performance of the algorithms, in conjunction with the Q criterion.

5.2 Behaviour w.r.t.N_{ppp}

The image particle density N_{ppp} is known to be one of the major factor that drives reconstruction performance in tomo-PIV. Our first experiment thus deals with the sensibility *w.r.t.* this experimental parameter.

Figure 4 presents a comparison of a classical tomographic reconstruction (tomo-SMART) and PVR (PVR-SMART) for $\sigma_{\text{psf}} = 0.6$. This PSF size corresponds to a good balance between contrast and peak-locking. Apart from the predictible decrease of performance for Q and Recall as N_{ppp} grows, two facts can be noticed:

1. PVR-SMART behaves better than tomo-SMART all over the range of N_{ppp};

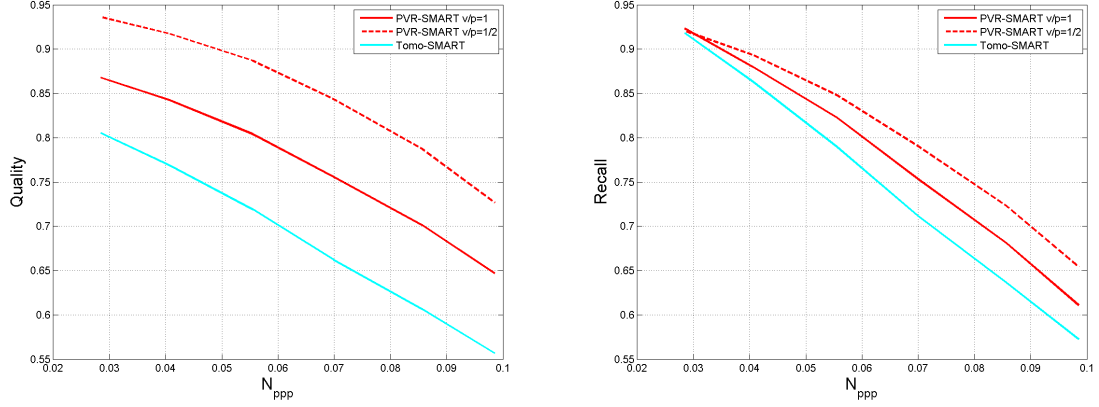2. PVR-SMART with $v/p = 0.5$ behaves better than $v/p = 1$.

Figure 4: Comparison between conventional tomographic reconstruction and the proposed PVR approach, using the SMART algorithm, $\sigma_{\text{psf}} = 0.6$. Left: Q criterion; Right: fraction of detected true particles, or Recall, for $NS = 1$. Note the performance gain of PVR-SMART with $v/p = 0.5$ compared to $v/p = 1$.

The better performance of $v/p = 0.5$ compared to $v/p = 1$ is an original but expected result, which we anticipated in the discussion that followed the discretisation of the image formation model. The reason is that using $v/p = 1$ in the case $\sigma_{\text{psf}} = 0.6$ clearly introduces aliasing errors in the discrete image model. As a natural way to defeat aliasing is to sample more densely the reconstructed volume, $v/p = 0.5$ thus works very well in this respect, though at the expense of an increase of memory and computations. An interesting question is now to determine if this increase in reconstruction performance obtained by refining the discretisation (reducing v/p) is robust, *i.e.*, occurs in all experimental situations.

5.3 Behaviour w.r.t.σ_{psf}

Our second numerical experiment takes a step in answering this question. Figure 5 presents the results for tomo-SMART, PVR-SMART $v/p = 1$ and $v/p = 0.5$ for a large range of σ_{psf} values.
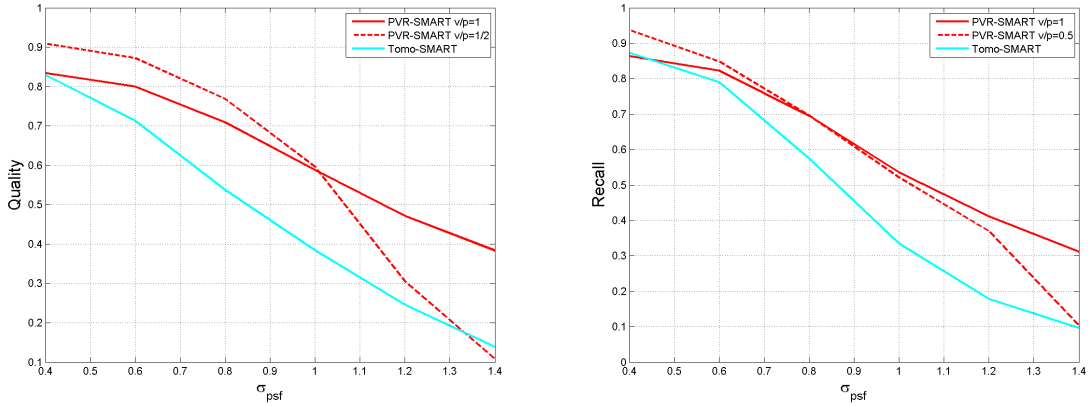

Figure 5: Evolution of Q criterion (left) and Recall (right, $NS = 1$) for various σ_{psf}. $N_{ppp} = 0.07$.

The left and right part of Figure 5 are quite consistent: $v/p = 0.5$ is best for low σ_{psf}, then $v/p = 1$ takes over; for the Q criterion (resp. Recall) the transition takes place near $\sigma_{\text{psf}} = 1$ (resp. $\sigma_{\text{psf}} = 0.85$). We explain this behaviour by *overfitting*. Setting $v/p = 0.5$ gives indeed more degrees of freedom than $v/p = 1$ to explain image data. This freedom is useful to fit the image of a particle that projects on two or more pixels, especially if the PSF is sharp and induces aliasing. But as the PSF gets smoother this freedom is unnecessary and even counter-productive: it makes possible a situation where different volume configurations can explain the same image which is equivalent to say that the linear system has a lower condition number), and the SMART algorithm is apparently unable to solve the corresponding ambiguities.

So our answer is no: reducing v/p is not always the best choice, as it can reduce performance by over-fitting the PSF. v/p should be adapted to the PSF. In any case, Figure 5 shows that a v/p can always be found such that PVR-SMART performs consistently better than tomo-SMART.

5.4 Behaviour w.r.t. inaccurate PSF knowledge

In practice the PSF is known approximately as it results from a calibration, and depending on the context a family of models can be fitted to experimental data as demonstrated in [9]. This means that in practice PVR-SMART will use a PSF model which is different from the true PSF. We propose a first-order study of this problem by assuming that the true PSF is given by (2) with $\sigma_{psf} = 0.6$, but let the parameter σ_{psf} used for reconstruction vary between 0.5 and 0.9. The results are gathered in Figure 6.
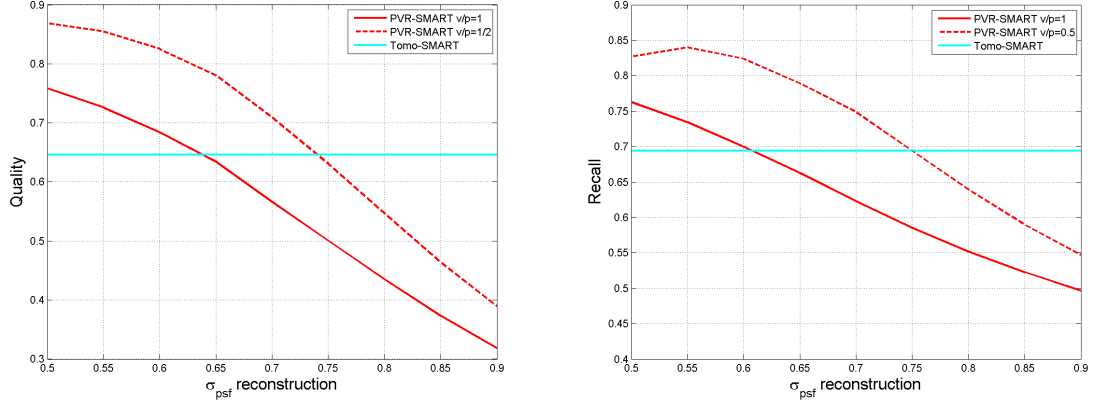

Figure 6: Evolution of Q criterion (left) and Recall (right) as a function of the value for σ_{psf} used in PVR-SMART. The true σ_{psf} equals 0.6 and N_{ppp} equals 0.07.

One observes that PVR-SMART is fairly tolerant to inaccurate knowledge of σ_{psf}, in particular $v/p = 0.5$ has a large domain of validity around 0.6. Moreover, it seems always better to underestimate σ_{psf}.

6. Discussion and concluding remarks

This paper deals with the first step of tomo-PIV, volume reconstruction. We have presented an alternative approach to the classical tomographic reconstruction, that seeks to recover single voxel particles rather than blobs of extended size, an approach referred to as "Particle Volume Reconstruction" (PVR). The proposed method is based on the image formation model from optics. The image model is carefully discretised in order to yield a rigorous representation of image data in terms of voxelised particle intensity and a weight matrix built with PSF samples. The degrees of freedom of this representation are analysed and we show the usefulness of adapting one of them, the volume-to-pixel ratio, depending on the PSF size. In particular we underline the interest of using a volume-to-pixel ratio lower than one in order to deal with sharp PSFs corresponding to well focused, low f-number optics. Another interesting degree of freedom is a kernel that controls the smoothness of the reconstructed volume. The proposed PVR method promotes detection of particles through compact kernels that enforce sparsity, but other choices could be studied further that yield smoother reconstructions and promote correlation.

PVR-SMART has been shown to perform better than tomo-SMART on a large domain of generating conditions. Robustness to inaccurate knowledge of the PSF size has also been checked. It is worthwhile mentioning that the proposed image data representation has also been tested successfully in association with other sparsity-based algorithms, see our companion paper [6].

As PVR directly builds the weight matrix from samples of the PSF, PSF calibration techniques as proposed by Schanz *et al.* [9] can readily be used. Future work will aim at developing the post-processing necessary to perform velocimetry from the output of the PVR method. Since our first implementation PVR is oriented to detection at the voxel scale, particle location refinement is indeed necessary before tracking the particle along time.

References

[1] R. J. Adrian and J. Westerweel. *Particle Image Velocimetry*. Cambridge University Press, 2010.

[2] Callum Atkinson and Julio Soria. An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Experiments in Fluids*, 47:553–568, 2009.

[3] I. Barbu, C. Herzet, and E. Mémin. Sparse models and pursuit algorithms for PIV tomography. In *Forum on Volumic Reconstruction 2011*, 2011.

[4] Charles L. Byrne. *Applied Iterative Methods*. AK Peters, 2008.

[5] A. cheminet, B. Leclaire, F. Champagnat, P. Cornic, and G. Le Besnerais. On factors affecting the quality of tomographic reconstruction. In *Proceedings of PIV13*, 2013.

[6] P. Cornic, F. Champagnat, A. Cheminet, B. Leclaire, and G. Le Besnerais. Computationally efficient sparse algorithms for tomographic PIV reconstruction. In *Proceedings of PIV13*, 2013.

[7] G.E. Elsinga, F. Scarano, Wieneke B., and B.W. van Oudheusden. Tomographic particle image velocimetry. *Experiments in Fluids*, 41:933–947, 2006.

[8] S. Petra and C. Schnörr. TomoPIV meets compressed sensing. *Pure Mathematics and Applications*, 20(1-2):49 – 76, 2009.

[9] Daniel Schanz, Sebastian Gesemann, Andreas Schröder, Bernhard Wieneke, and Matteo Novara. Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Meas. Sci. Technol.*, 24:1–15, 2013.