HAL
archives-ouvertes.fr

A new sparsity based particle image reconstruction approach for particle detection

Soufiane Ait Tilat, Frédéric Champagnat, Cédric Herzet

▶ **To cite this version:**

Soufiane Ait Tilat, Frédéric Champagnat, Cédric Herzet. A new sparsity based particle image reconstruction approach for particle detection. ISPIV 2019 - 13th International Symposium on Particle Image Velocimetry, Jul 2019, Munich, Germany. pp.1-10. hal-02389648

HAL Id: hal-02389648
https://hal.inria.fr/hal-02389648

Submitted on 2 Dec 2019

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.

13th International Symposium on Particle Image Velocimetry – ISPIV 2019
Munich, Germany, July 22-24, 2019

A new sparsity based particle image reconstruction approach for particle detection

Soufiane Ait Tilat[1*], Frédéric Champagnat[1], Cédric Herzet[2]

[1] ONERA the French Aerospace Lab, Information Processing and Systems Department, Palaiseau, France

[2] Inria Rennes-Bretagne Atlantique, SIMSMART, Rennes, France

[*]soufiane.ait_tilat@onera.fr

Abstract

The rationale behind this paper is to propose a methodogy to identify and localize the particle image centers under high density conditions. Cheminet proposed a reconstruction method based on the inversion of an imaging model using a non negative least squares (NNLS) approach with extremely small grid sizes, followed by a local maxima step. This strategy enables to approximate accurately the particles position and handle overlapping. We address here the issues encountered when reconstructing with NNLS in noisy conditions. We also proposed two approaches based on l_1-norm regularization with non-negativity constraints : The first one, called BP, aims to enhance NNLS results in the presence of noise and the second one, called C-BP, aims to reduce the grid size adopted for recovery while keeping the same performances as BP.

1 Introduction

Accurate detection and localization of individual particles in PIV recordings is a critical step of Particle Tracking Velocimetry (Ouellette et al., 2006). Indeed, uncertainties in the determination of particle positions greatly worsen velocity errors. For low seeding densities, particle detection and localization can be merely achieved by using local maxima detection followed by 1D Gaussian fit to the pixels directly adjacent to it (Ouellette et al., 2006). However, as the seeding density increases, the localization accuracy achievable with such a method degrades drastically due to particle image overlapping. One approach to address this issue is to use an imaging model that explicitly accounts for overlapping: PIV recordings are described as a weighted sum of shifted identical patterns (Champagnat et al., 2014):

$$i(n) = \sum_{p=1}^{P} e_p h(n - x_p) \tag{1}$$

where $i \in R^N$ denotes the image, $n \in R^2$ a pixel location, $h\colon R^2 \to R_+$ is the so-called point spread function (PSF) modelling the diffraction pattern, P the number of particles, $e_p \in R$ the intensity of each particle and $x_p \in R^2$ the (subpixel) continuous position. Our aim is to identify the position of each particle along with its associated amplitude. This is a difficult problem because of the nonlinear dependence of i on x_p, furthermore, the number of particles P is unknown.

There are two main approaches to tackle this problem : the first one preserves the continuous aspect of the particles position and leads to iteratively add or remove particles and refine their position in order to fit a least squares criterion based on (1). Two instances are the Iterative Particle Reconstruction (IPR) (Wieneke, 2013) and the Sliding Frank-Wolfe algorithm (SFW) (Denoyelle et al., 2018). Although relying on physically sound model because particles live on a continuum, they lead to some practical implementation issues due to non convexity of the least squares criterion. The second approach leads to simpler methods. The continuous space of particle image locations x_p is disretized on a 2D grid. This leads to a linear model. Inversion of this model requires prior information. For instance, Cheminet et al. (2018) considered a least squares approach under positivity constraint. When the grid size is relatively coarse, this strategy has two limitations, namely the presence of multiple detections in the vicinity of the true particles and the inability to handle overlapping. Cheminet et al. (2018) showed that these limitations can be overcome using a non negative least squares (NNLS) approach with extremely small grid sizes, followed by a local maxima step

to extract only relevant points in the refined subspace. Undoubtedly, refining the grid allows to more closely approximate x_p and handle overlapping, but demands more computation. It should be emphasized, that this method is sensitive to noise and requires high thresholds intensity level to rule out false detections.

Our contributions are twofold: first we introduce a constraint enforcing the sparsity of solution, that is we intend to find the solution of the observation equation (1) involving the fewest number of particles. Our approach is based on the use of l_1-norm penalization with a non-negativity constraint in order to enhance the results of Cheminet et al. (2018). This amounts to a Basis Pursuit approach (BP) (Chen et al., 1998). Second, we introduce a continuous basis pursuit method (C-BP) (Ekanadham et al., 2011) in order to preserve the sub-grid localization accuracy while maintaining a relatively coarse grid. The central idea is to approximate the particle image by a linear combination of the PSF and its derivative via a first-order Taylor expansion. Then detection and localization derive from a convex optimization problem with appropriate linear constraints. We have generalized the original 1D approach of Ekanadham et al. (2011) and developed a solver adapted to the linear constraints induced by the 2D C-BP approach. The real challenge here is to recover faithfully the position of particles under high density conditions while keeping a reasonable algorithmic cost. We address this trade-off in the present paper.

The structure of the paper is as follows. Section 2 introduces the image formation model. Section 3 describes the rationale of using l_1 regularization and the algorithm used for reconstruction. Section 4 presents an extension of Continuous Basis Pursuit detection techniques to particle imaging. Finally, all details about the synthetic setup, the reconstruction steps, the performance metrics and the simulation results are covered in the last section.

2 Image model discretization and NNLS inversion

We adopt here the particle approach proposed in (Champagnat et al., 2014). The image formation model is given by equation (1). For all the simulations performed in this paper, we assume that the PSF is a convolution between a Gaussian function with standard deviation σ_{psf} accounting for diffraction and a gate function for spatial integration over the detector's surface. We seek to recover the particle position x_p and its associated amplitude e_p. The classical way to do that is to discretize the subpixel position x_p to a finite set of possible nodes, and consider only the nearest ones to encode the particle locations. Let's consider a grid with M nodes and Δ its scale. For a given particle p, localized in x_p, we associate the closest node, indexed m. The position of the latter is denoted $m\Delta$, $m \in \mathbb{N}^2$. So the approximation can be expressed as follows :

$$h(n-x_p) = h(n-(m\Delta+\varepsilon_m)) \simeq h(x_{nm}). \tag{2}$$

where $x_{nm} = n - m\Delta$. The order of approximation is $O(\Delta)$. Evaluating (2) at the pixel n, we get:

$$i(n) \simeq \sum_{m=1}^{M} e_m h(x_{nm}). \tag{3}$$

Equation (3) is then equivalently written in the matrix form :

$$i \simeq \mathrm{H}e \quad \mathrm{H} = (h(x_{nm}))_{nm} \tag{4}$$

To invert the system one must account for two considerations :

1. The observed image can be altered by different types of noise.
2. The use of a thin grid to construct H results in an under-determined problem of a very high dimensionality. Indeed, refining the grid leads to deal with a matrix H having more columns M than rows N. Consequently, the system is ill conditionned.

In order to single out a proper solution, Cheminet proposed, in (Cheminet et al., 2018), to minimise the quadratic error between the approximated model and the observations under positivity constraint (see equation (5)).

$$\arg\min_{e \geq 0} \|i - \mathrm{H}e\|_2^2 \tag{5}$$

and use the active set algorithm of (Lawson and Hanson, 1974) to recover the associated solutions. As discussed in the same paper, NNLS combined with a local maxima extraction has an interesting level of

sparsity and a high capacity to distinguish between two particles in overlapping conditions. In noise-free simulations, when the grid is thin and the sub-pixel distance between particles is not extremely small to be considered critical, Cheminet et al. (2018) shows that overlapping situations are resolved and detections are localized in the vicinity of the true particles.

Nonetheless, this method is extremely sensitive to noise and requires a mandatory thresholding step to discard non relevant detections introduced by non-negative values of the added noise. A qualitative illustration of the NNLS solution, in the context of noisy data, is represented in Figure 1 before thresholding (left) and after thresholding (right). The observed image corresponds to a density of 0.10 particle per pixel. σ_{psf} is set to 0.6. All particles have the same intensity value and are placed randomly in the studied region. More details are discussed in section 5.1. We denote by (o) the real particles and ($\times$) the observed detections.

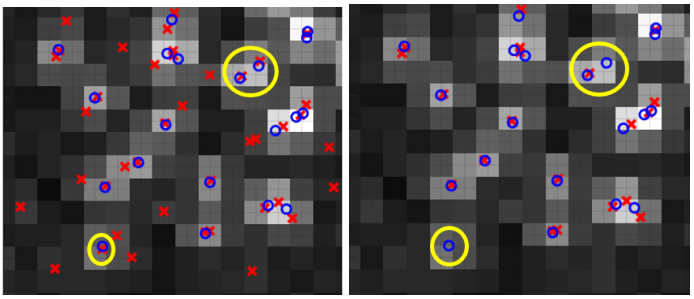

Figure 1: NNLS solution before thresholding (left) and after thresholding (right) in the presence of noise. The initial image contains many false detections. The thresholding process enables to discard almost all undesirable detections produced by noise but removes some true detections with low intensity values.

Obviously, the reconstructed solution is impacted by the background noise. This is evidenced by the important number of false detections obtained in areas where there are no particles. The latter can be removed using a high threshold level at the cost of eliminating true particles (see yellow circles).

3 Basis Pursuit method

As argued above, the positivity constraint is clearly insufficient to deal with noisy images. This is mainly due to the inherent structure of NNLS which has no parameter to tune to increase the sparsity, when the noise level varies. Our first approach to overcome this issue is to introduce l_1 regularization term as presentend in equation (6). We refer to this methodology as Basis Pursuit (BP).

$$\arg\min_{e\geq 0} \|i - \mathrm{H}e\|_2^2 + \lambda \|e\|_1 , \tag{6}$$

where λ is a tuning parameter that controls parsity. Figure 2 shows the result obtained by BP on the same image presented previously. We consider the same grid size Δ adopted to solve the NNLS problem. The sub-figure on the left shows the initial result without any prior thresholding and the sub-figure on the right shows the ultimate result obtained after a thresholding step. Clearly, comparing to the NNLS results, the number of false detections generated by the background noise are subtsantially eliminated. Although, thresholding improves the quality detection, it sometimes removes some true detections (see yellow circles in Figure 2). This last situation is more frequent with NNLS. Hence, the interest of using l_1 regularization.

3.1 BP implementation

Thanks to the positivity constraint, the BP problem (7) is indeed quadratic and can be written as follows :

$$\arg\min_{e\geq 0} \|i - \mathrm{H}e\|_2^2 + \lambda \sum_{m=1}^{M} e_m \tag{7}$$

. It can be solved using the Active set algorithm (Nocedal and Wright, 1999). Indeed, the Matlab function `lsqnonnneg` is an Active set algorithm based on (Lawson and Hanson, 1974), it is used by Cheminet et al.

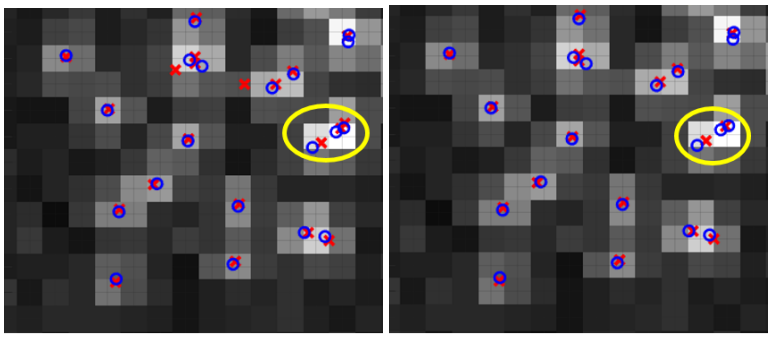

Figure 2: BP solution before thresholding (left) and after thresholding (right) in the presence of noise. The initial image contains a little number of false detections. As in the case of NNLS, the thresholding step may remove some true detections.

(2018) to solve NNLS problems. We propose here a modification of this algorithm to solve the BP problem (see Algorithm 1). The input variables are the image i, the matrix H and the sparsity parameter λ. The output variable is the intensity e. We denote by r the residual, Q the support of e, *i.e.* the non zero coordinates of e. $\bar{Q}$ is the complement of Q, H_Q the submatrix of H whose columns index are in Q and e_Q the solution restricted to Q.

Algorithm 1 Active set algorithm for Basis Pursuit methods

Input: i , H, λ
Output: e

$Q \leftarrow \emptyset,\ e \leftarrow 0,\ r \leftarrow i$
$\mu_{\bar{Q}} \leftarrow H^t_{\bar{Q}}\, r - \lambda$
while $\bar{Q} \neq \emptyset$ & $\max \mu_{\bar{Q}} > 0$ **do**
  $l \leftarrow \arg\max \langle \mathrm{H}^t_{\bar{Q}}, r \rangle$ {returns the index corresponding to the maximal of the inner product between r and $\mathrm{H}_{\bar{Q}}$}
  $Q \leftarrow Q \cup \{l\}$
  $x_Q \leftarrow \arg\min_{e_Q} \left\| i - \mathrm{H}_Q\, e_Q \right\|_2^2 + \lambda \sum_{m \in Q} e_m$
  $x_{\bar{Q}} \leftarrow 0$
  while $\min x_Q \leq 0$ **do** {Inner loop to remove negative components from x}
    j$\leftarrow \arg\min_{m \in Q}(-\frac{e^k_m}{e^k_m - x^k_m})$ {Finds indices where x is negative. k is the current iteration of the inner loop.}
    $\alpha \leftarrow \min\{1, \min_{m \in Q}(-\frac{e^k_m}{e^k_m - x^k_m})\}$
    $e \leftarrow e + \alpha(x - e)$ {Chooses new e with no negative components. }
    $Q \leftarrow Q - \{j\}$ {The index j is removed from the current support Q}
    $x_Q \leftarrow \arg\min_{e_Q} \left\| i - \mathrm{H}_Q\, e_Q \right\|_2^2 + \lambda \sum_{m \in Q} e_m$
    $x_{\bar{Q}} \leftarrow 0$
  end while
  $e \leftarrow x$
  $r \leftarrow i - \mathrm{H}\, e$
end while

At each iteration, The algorithm selects the atom having the highest correlation with the current residual. It adds the associated index to the support Q then solves the penalized least squares criterion with respect to coordinates index in Q. If all coefficients in x_Q are non-negative, e is set equal to x. When we add a new variable to the support Q, there is a chance that in the new penalized least squares problem, some of the solution coefficients will turn negative. If that is the case, an inner loop is repeated until all violating variables are eliminated from Q. Overall, the most expansive operation is the penalized least squares step. Note that the case $\lambda = 0$ corresponds exactly to Matlab implementation of NNLS.

4 Continuous Basis Pursuit Method

As described earlier, NNLS and its regularized version BP requires a very thin grid to recover the particle distribution accurately. Nevertheless, in this case one requires the manipulation of a matrix with a very high dimension which leads to important memory and computational burden. In order to circumvent this issue, we adopted an alternative discretization. The central idea is to smoothly approximate the particle image by a linear combination of the PSF and its derivative via a first-order Taylor expansion. As shown in Figure 3, the particle image can be faithfully approximated by a linear combination of the same particle image localized in $m\Delta$ and its derivative at the same grid node.

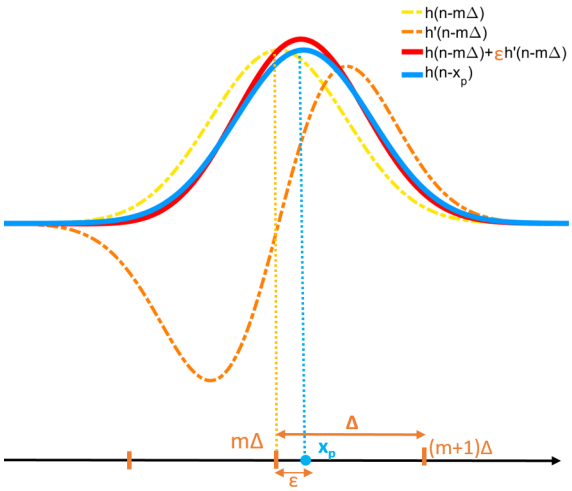

Figure 3: First-order Taylor expansion leads to an approximation of a shifted PSF function by a weighted sum of the PSF and its derivative. Δ is the grid size. m is the nearest node to the particle image center x_p.

4.1 Problem formulation

Using the same notations as in section 2, the Taylor approximation is written as follows :

$$\begin{aligned} h(n-x_p) &= h(n-(m\Delta+\varepsilon_m)) \\ &\simeq h(x_{nm}) - \varepsilon_{1m}.\frac{\partial h}{\partial x}(x_{nm}) - \varepsilon_{2m}.\frac{\partial h}{\partial y}(x_{nm}) \end{aligned} \tag{8}$$

.

In a 2D case, the Taylor approximation involves two additional space variable : ε_{1m} and ε_{2m}, encoding the horizontal and the vertical discretization error. The latter are constrained in a 2D interval depending on the grid scale Δ :

$$|\varepsilon_{1m}| \leq \frac{\Delta}{2} \qquad |\varepsilon_{2m}| \leq \frac{\Delta}{2} \tag{9}$$

Using (8) and (9) for each particle position x_p, and plugging into (1), we get :

$$i(n) \simeq \sum_{m=1}^{M} e_m h(x_{nm}) - d_{1m}\frac{\partial h}{\partial x}(x_{nm}) - d_{2m}\frac{\partial h}{\partial y}(x_{nm}) \tag{10}$$

with

$$|d_{km}| = e|\varepsilon_{km}| \leq \frac{\Delta}{2}\, e_m \qquad k \in \{1,2\}$$

As pointed out in (Ekanadham et al., 2011) the auxiliary constraint $|d_{km}| \leq \frac{\Delta}{2}\, e_m$ has two goals. First, it guarantees that the Taylor approximation is only used when it is accurate. Second, it discourages neighboring nodes from explaining the same particle image. This leads to the following matrix form :

$$i \simeq \mathrm{H}e + \mathrm{H}_1 d_1 + \mathrm{H}_2 d_2, \quad \begin{cases} \mathrm{H} = [h(x_{nm})]_{nm} \\ \mathrm{H}_1 = [-\frac{\partial h}{\partial x}(x_{nm})]_{nm} \\ \mathrm{H}_2 = [-\frac{\partial h}{\partial y}(x_{nm})]_{nm}. \end{cases} \tag{11}$$

In order to solve this system with the basis pursuit method, we introduce an l_1 penalization and a non negativity constraint, we obtain the continuous BP (C-BP) problem :

$$\begin{aligned} &\underset{e,d_1,d_2\in\mathcal{E}}{\arg\min}\ \|i-(\mathrm{H}e+\mathrm{H}_1 d_1+\mathrm{H}_2 d_2\|_2^2+\lambda\|e\|_1 \\ &\mathcal{E}=\left\{e_m\geq 0 \quad |d_{km}|\leq \tfrac{\Delta}{2}e_m \quad k\in\{1,2\} \quad m\in\{1..M\}\right\} \end{aligned} \tag{12}$$

4.2 Parametrization as Basis Pursuit

We propose hereafter a simple way to reparameterize the C-BP problem as a classical basis pursuit problem under positivity constraints. This is possible because $\mathcal{E}$ is a convex cone. Indeed it can be seen that $\mathcal{E}$ can be rewritten as follows:

$$\mathcal{E}=\{(e,d_1,d_2)|\,(e_m,d_{1m},d_{2m})\in\mathcal{E}_m, m\in\{1..M\}\},$$

where

$$\mathcal{E}_m=\left\{(a,b,c)\in\mathbb{R}^3|\,|b|\leq\alpha\frac{\Delta}{2}a,\ |c|\leq\frac{\Delta}{2}a\right\}.$$

The shape of each basic cone $\mathcal{E}_m$ is depicted on Figure 4. Indeed, one can deduce a generative description with positive coefficients that fully represents $\mathcal{E}_m$, see Figure 4. This procedure was suggested by Duval and Peyré (2017) in the 1D case. The extension to the 2D case is defined in equation (15) and is obtained by setting :

$$\alpha=\frac{\Delta}{2},\quad \mathrm{H}_\alpha=[\mathrm{H}\mathrm{H}_1\mathrm{H}_2]\mathrm{P}_\alpha,\quad \mathrm{P}_\alpha=\begin{pmatrix}1 & 1 & 1 & 1\\ \alpha & -\alpha & \alpha & -\alpha\\ \alpha & \alpha & -\alpha & -\alpha\end{pmatrix}\otimes\mathbb{I}_M \tag{13}$$

and

$$u\geq 0 \quad \text{such that} \quad [{}^te,{}^td_1,{}^td_2]^t=\mathrm{P}_\alpha u \tag{14}$$

$\otimes$ is the Kronecker product and $\mathbb{I}_M$ the unit matrix of size M.

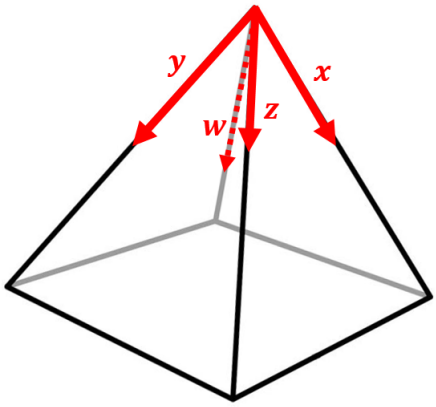

Figure 4: $\mathcal{E}_m$ is a pyramidal cone generated by the linear combinations of 4 vectors with positive coefficients. These vectors correspond to the columns of P_α.

Finally, the problem to solve is :

$$\underset{u\geq 0}{\arg\min}\,\|i-\mathrm{H}_\alpha u\|_2+\lambda\|u\|_1 \tag{15}$$

To sum up, the C-BP approach has four steps : First, we discretize (1) using first-order Taylor interpolation. This yields matrices H, H_1 and H_2. Second, we compute the matrix H_α using (13). Third, we solve (15) using the active set algorithm described in section 3.1. Finally, we map back from the obtained solution u to the initial parameters e, d_1 and d_2 which encode the intensity and the particle position using (14).

5 Numerical Results

5.1 Synthetic setup

To illustrate the benefits of our methods, we generated a set of synthetic images of varying particle density according to the generative model (1). Particles are monodisperse and their image centers are uniformly distributed in a 32×32 pixel image. We set $\sigma_{psf} = 0.6$ (particle image diameter equal to 2.4). With this value, a particle has a 4×4 pixel image pattern. The density is controlled by the particle-per-pixel count (N_{ppp}). A zero-mean Gaussian distribution representing camera noise is added to the images. Its standard deviation is about 5% relative to the maximum particle image intensity. 30 images are generated to guarantee the convergence of the metrics used in our simulations.

5.2 2D reconstruction steps and performance metrics

In order to handle overlappings and achieve an accurate localization of particles, one needs to have a very small grid step size. However, when the size of the grid increases, we obtain a small cluster of non-zero weights in the vicinity of each true particle, more precisely on the vertices of the discretization cell that surrounds it. A way to mitigate this issue is to perform a post processing to rule out irrelevant detections. In all our simulations, we use a grid step size Δ equal to $\frac{1}{20}$ pix to invert NNLS and BP and Δ equal to $\frac{1}{5}$ pix for C-BP. Other relevant parameters are selected as described in the next section 5.3. Figure 5 illustrates the overall steps followed in this process, for instance in the case of BP : Step 1 corresponds to BP Inversion. The associated solution is represented in the refined subspace. Step 2 corresponds to a thresholding process to eliminate detections generally produced by noise. Step 3 corresponds to a local maxima extraction. The coordinates of each maxima are taken as the particle location estimate.

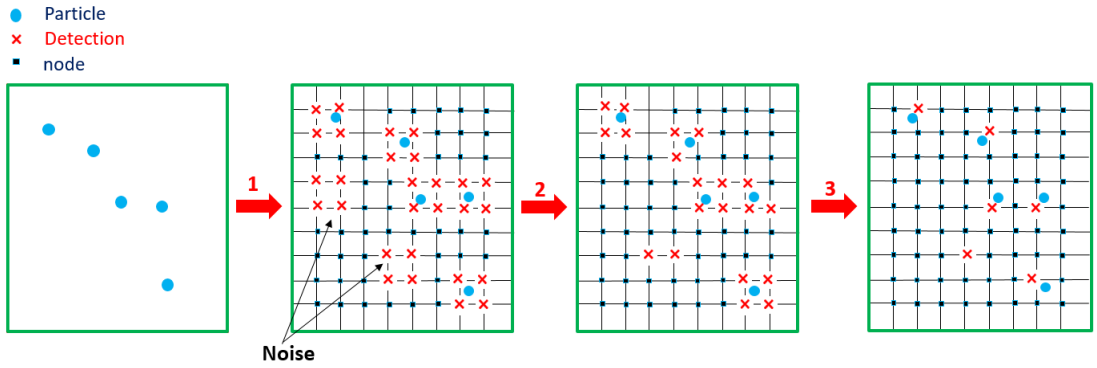

Figure 5: 2D reconstruction steps using BP. Step 1 : BP recovery. Step 2 : thresholding. Step 3 : local maxima extraction.

Our three methods yield a list of nonzero nodes occupied by particles. A detection, i.e. a nonzero node, is a True Positive (TP) if it is in the neighbourhood of a true particle. The neighborhood of any particle is the disk centered around its geometrical projection x_p with a radius r lower than 0.5 pixel. A detection is a False Positive (FP), 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. As mentionned in (Champagnat et al., 2014), to measure the detection performance of the three studied algorithms, we define two metrics : Precision and Recall.

$$Precision = \frac{\#TP}{\#TP + \#FP} \qquad Recall = \frac{\#TP}{\#TP + \#FN} \tag{16}$$

where $\#TP$, $\#FP$ and $\#FN$ are the number of True Positives, False Positives and False Negatives detections respectively. We also used the Fscore indice to select the optimal value of the important parameters (see section 5.3). It is defined as the harmonic average of the precision and recall :

$$Fscore = 2\,\frac{Precision \times Recall}{Precision + Recall} \tag{17}$$

The selection of the highest value of Fscore corresponds to a compromise between Precision and Recall. The localization performance is assessed using the averaged distance between estimated TP positions $\tilde{x_p}$ and ground truth positions x_p :

$$\text{Mean error localization} = \frac{1}{\#TP}\sqrt{\sum_{p=1}^{\#TP}(x_p - \tilde{x_p})^2} \tag{18}$$

5.3 Parameters selection and performance results

We perform a Monte Carlo experiment as described in section 5.1 for each value of the sparsity coefficient λ and the thresholding value S to calculate the associated detection metrics. This procedure is performed for λ varying from 0 to 0.12 and S from 0 to 0.9. Figure 6 shows three maps corresponding to the three detection measures defined in the previous section: Precision, Recall and Fscore. Notice that the first row corresponds to the NNLS case : λ = 0. As expected, an increase of the sparsity coefficient λ and the threshold level S enables to enhance the Precision index but deteriorates the Recall. Conversely, a decrease of these two parameters values, improves the recall but to the detriment of the precision quality. A fair compromise between Recall and Precision is found at the maximum of the Fscore Map. The red circle corresponds to the optimal couple (λ, S) in the case of BP and the green one corresponds to the optimum in the case of NNLS. Similar simulations were performed on C-BP to set the same parameters.

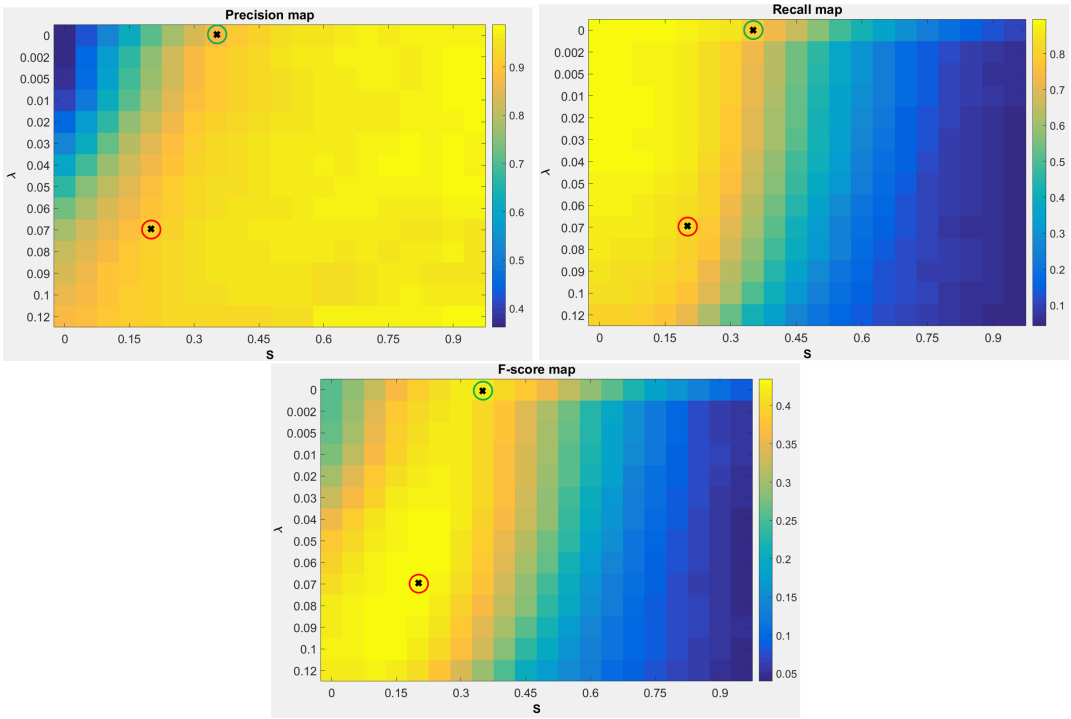

Figure 6: Precision map, Recall map and Fscore map, obtained with the BP method. The optimal point correspond to a sparsity coefficient equal to 0.07 and a threshold S equal to 0.2 (Red circle). The optimal NNLS threshold is deduced from the same map and is equal to 0.35 (Green circle).

The comparison of the performance metrics obtained by the two proposed Basis pursuit methods and NNLS with optimal parameters, is shown in Figure 7, as a function of the particle density, N_{ppp}. It is observed that NNLS Precision is 5% lower than BP precision for a grid size $\Delta = 1/20$ pix and for low densities N_{ppp}. This gap slightly decreases for high level densities. Recall curves show the same tendency but the difference is not quite significant. These results validate the fact that BP with optimal settings manages to eliminate more False positives produced by noise and to retain slightly more low intensity detections after the thresholding process. We can also notice that the C-BP curve obtained with Δ equal to

$\frac{1}{5}$ pix is almost identical to the BP curve obtained with Δ four times smaller. This is an important result of C-BP because one can use a grid size 4 times coarser and get exactly the same detection performance as BP. This corresponds to a matrix with a number of columns 4 times smaller and hence a gain in memory and computing time.

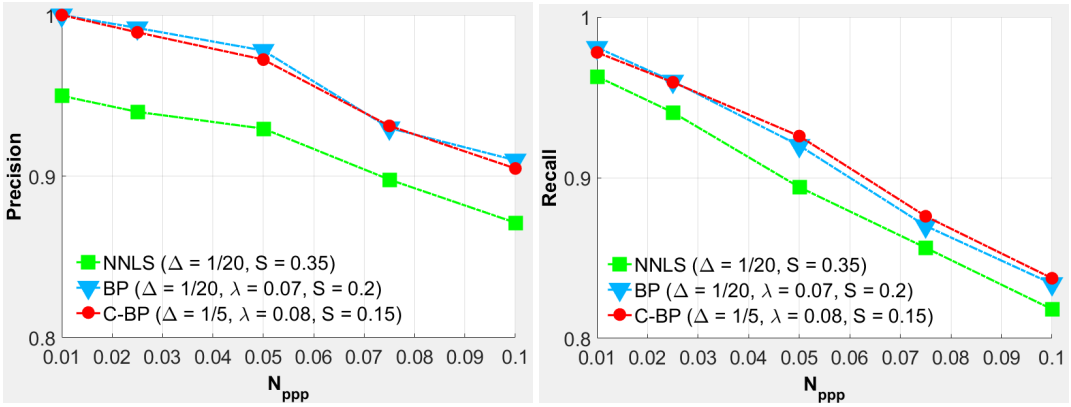

Figure 7: Comparison between NNLS (Green), BP (Blue) and C-BP (Red) in terms of detection performance for different densities N_{ppp} and different grid sizes Δ. BP outperforms NNLS in terms of Precision and Recall for the same Δ. C-BP leads to similar detection results as BP with a grid 4 times coarser.

In terms of localization performance, The overall results are almost the same. Figure 8 shows the mean localization error obtained with the three studied methods as a function of N_{ppp}. We keep the same parameter values : $\Delta = \frac{1}{20}$ pix for BP and NNLS. $\Delta = \frac{1}{5}$ pix for C-BP. The three curves are quasi identical including the C-BP one. This is also an interesting feature of the C-BP method. Even for a discretization step, four times coarser, we get very small localization errors.

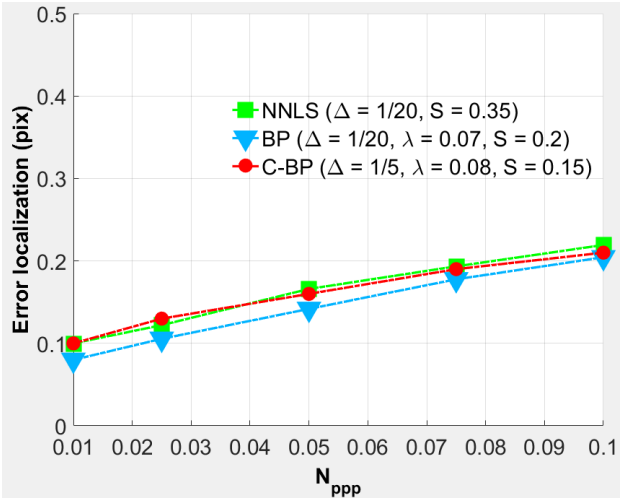

Figure 8: Mean Localization error obtained with NNLS, BP and C-BP as a function of the particle density N_{ppp}. The overall results of the three methods are quite similar. C-BP used with $\Delta = \frac{1}{5}$ pix has the same localization error as NNLS and BP used with $\Delta = \frac{1}{20}$ pix.

6 Conclusion

We presented an improvement of the NNLS method to filter out wrong detections produced in the presence of noise. Our approach was based on the use of l_1-norm penalization with a non-negativity constraint. We

dubbed it BP and presented a simple implementation using an active set algorithm. We also proposed a second approach C-BP which can be seen as a continuous form of the latter. One of its most salient features is that it relies on a coarse grid. Thanks to the accurate approximation of the continuous PSF funtion, C-BP yields similar quality detection and almost the same localization error with a grid size 4 times coarser. We believe our C-BP strategy can be extended to the $3D$ context if it is combined with a screening procedure (Fercoq et al., 2015). The latter consists on simple and safe tests (before or during the optimization process) to eliminate all the unused nodes and consequently a dramatic reduction of the matrix size used for recovery.

Acknowledgements

This work was supported by the project ANR 15-CE23-0021 BECOSE.

References

Champagnat F, Cornic P, Cheminet A, Leclaire B, Le Besnerais G, and Plyer A (2014) Tomographic PIV: Particles versus blobs. *Measurement Science and Technology* 25:084002

Cheminet A, Krawczynski J, and Druault P (2018) Particle image reconstruction for particle detection in particle tracking velocimetry. *Measurement Science and Technology* 29:125202

Chen S, Donoho D, and Saunders M (1998) Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing* 20:33–61

Denoyelle Q, Duval V, Peyré G, and Soubies E (2018) The sliding Frank-Wolfe algorithm and its application to super-resolution microscopy. *arXiv e-prints* arXiv:1811.06416

Duval V and Peyré G (2017) Sparse spikes super-resolution on thin grids II: the continuous basis pursuit. *Inverse Problems* 33:095008

Ekanadham C, Tranchina D, and Simoncelli P (2011) Recovery of sparse translation-invariant signals with continuous basis pursuit. *IEEE Transactions on Signal Processing* 59:4735–4744

Fercoq O, Gramfort A, and Salmon J (2015) Mind the duality gap: safer rules for the lasso. *Proceedings of the 32nd International Conference on Machine Learning* 37:333–342

Lawson CL and Hanson RJ (1974) *Solving Least Squares problems*. pages 149–199. Prentice-Hall, Saddle River, NJ, USA

Nocedal J and Wright S (1999) *Numerical Optimization*. Springer-Verlag, New-York

Ouellette N, Xu H, and Bodenschatz E (2006) A quantitative study of three-dimensional Lagrangian particle tracking algorithms. *Experiments in Fluids* 40:301–313

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