

A robust method for estimating seeding density in PIV images

M. Hebey^{1*}, F. Champagnat², B. Leclaire¹

¹DAAA, ONERA, Institut Polytechnique de Paris, 92190, Meudon, France

²DTIS, ONERA, Université Paris-Saclay, 91120, Palaiseau, France

* maximilien.hebey@onera.fr

ABSTRACT

We introduce a new method for estimating the seeding density of particle images, using a statistical model of their pixel intensities based on Filtered Poisson Processes. We show that computing the ratio of two different statistical quantities from the images allows to obtain an estimator proportional to the true seeding density. This is in particular the case up to high density values, where multi-scattering effects can be present. To estimate the proportionality coefficient between estimated and true densities, which involves contributions from the particle image shapes and from their distribution of intensities, we introduce an *in situ* calibration process at low densities, considering estimation by counting local intensity maxima as the reference. To guarantee an optimal accuracy of this calibration step, new characterizations and guidelines for optimal and easily automated use of this reference counting have been derived. Performance assessment of the method is first presented using synthetic cases corresponding to both 2D and 3D virtual observation setups, the latter being introduced to assess the impact of spatial inhomogeneity linked to perspective viewing. They confirm the soundness of both the model and the calibration process, the excellent accuracy of the method up to very high seeding densities, as well as its robustness to both uncorrelated and correlated spatial noise (for the latter, once a historical minimum image has been subtracted). The method is then tested on experimental images from the 3D Rayleigh-Bénard convection experiment of Bosbach et al. (2021). Results confirm its very good performance across a wide range of densities, as well as its robustness to multiple scattering effects. Further analysis of the results also shows that the image cumulants, as well as our estimator, can be used as valuable diagnostic tools to monitor changes in the distribution of particle intensities.

1. Introduction

Knowledge of the image seeding density in Particle Image (resp Tracking) Velocimetry (PIV, resp. PTV) experiments is of foremost importance. It can indeed help adjusting the interrogation region size and assessing the uncertainty in PIV (Timmins et al., 2012), and constitutes a crucial landmark in the context of PTV (e.g. adjustment of seeding in an experiment to guarantee final accuracy), all the more as the maximum accessible density of these methods increases continuously. Due to particle image overlap, estimation of this quantity (in ppp, particles per pixel, this abbreviation being often used to denote the density itself) is an increasingly difficult task as the ppp grows. Warner & Smith (2014) suggest an estimator based on an empirical relationship between the particle image density and the following parameters: the autocorrelation peak magnitude, the particle image diameter, the particle image intensity, and the interrogation region size. The estimator is robust up to 0.08 ppp. More recently, Sankaran et al. (2024) suggested a Voronoi-cell based estimator to identify

the local density. Whilst being very innovative, this method is subject to a heavy calibration process and most suited to peaked particles with diameters inferior to 2 pixels. In a more theoretical spirit, Thomas et al. (2025) suggested a mathematical analysis of overlapping linking the real *ppp* and the number of detectable peaks in an image. Our method aims at drastically widening the range of accessible *ppp* by estimating a simple quantity, derived from image statistics, that linearly depends on it. To do so, similarly to our recent contribution on the subject (Hebey et al., 2025), we model the particle image formation process as a probabilistic model for the pixel intensity random variable. While our initial work requested to prescribe a distribution for the particles intensities, which can be complex in practice (as this distribution might be significantly different from an experiment to another), we here introduce a new and simpler method which does not rely on such a prescription, and will therefore be robust and accurate in a large number of experimental contexts. Besides, it relies on a limited number of computations so that, after an initial calibration phase, it can be used to provide a real-time estimation of the *ppp*.

We first describe the underlying model and proposed method in section 2. Sections 3 and 4 then provide assessments, respectively on synthetic and experimental data, the latter consisting of a 3D PTV experiment performed by DLR, described in Bosbach et al. (2021) and Weiss et al. (2024) and which data are partially available on the website <https://onera-flow-benchmarks.fr>.

2. Proposed method

2.1. Outline

This section presents the derivation of our method for estimating seeding density. Being defined relative to a digital image, this quantity is traditionally expressed in particles per pixels (*ppp*). In the mathematical expressions, we will denote it by λ , while in the text we will sometimes also refer to it directly as "the *ppp*", for simplicity.

A first necessary step for the estimation of λ is to introduce a model for PIV/PTV image formation, which is done in section 2.2. The probabilistic counterpart of this model as a Filtered Poisson Process (FPP) is then exposed in section 2.3, which allows to derive a set of equations linking λ to statistical quantities accessible directly from the images, but also to other quantities stemming from the optical setup. As these quantities can be complex to estimate in practice, section 2.4 then shows how some of these equations can be combined to obtain, *solely from image statistics*, an estimate $\tilde{\lambda}$ of λ up to a proportionality constant C , i.e. such that $\tilde{\lambda} = C \lambda$. As constant C has no reason to be equal to unity in practice, an additional operation is necessary to determine its value, which we introduce as an *in situ* calibration step. This is done in section 2.5.

2.2. PIV image formation model

We first introduce the following model for the pixel intensities of a PIV/PTV image:

$$I(\mathbf{x}) = I_P(\mathbf{x}) + N(\mathbf{x}) = \sum_{p=1}^P E_p h(\mathbf{x} - F(\mathbf{X}_p)) + N(\mathbf{x}) \quad (1)$$

where $I(\mathbf{x})$ is the grey level at pixel of coordinate $\mathbf{x}$, resulting from contributions of particles $I_P(\mathbf{x})$ and of background noise $N(\mathbf{x})$. As shown by the expression of $I_P(\mathbf{x})$, P particles with individual intensities E_p are considered. Their positions in the object space are denoted by $\mathbf{X}_p$, so that their projected position on the camera sensor writes $F(\mathbf{X}_p)$, with F the camera projection function obtained from calibration. Function h denotes the so-called Point Spread Function (PSF, also known as Optical Transfer Function, OTF), modelling the effect of diffraction-limited optics, leading to particle images spreading with an Airy shape over a few pixels in both directions, whose diameter is often denoted by d_τ (Raffel et al., 2018). In practice, h is modelled as a Gaussian (e.g. Keane & Adrian (1992), or more recently Schanz et al. (2013)) or as an erf function (e.g. Stanislas et al., 2004; Champagnat et al., 2014), the latter choice modelling the integration over the pixel surface, possibly accounting for a fill ratio smaller than unity. Given this definition of h , E_p corresponds to the intensity of particle p collected at the camera aperture, before diffraction sets in. It thus gathers all possible effects due to inhomogeneities in laser sheet intensity, particle size, as well as Mie scattering, and possibly multi-scattering at high concentrations (Raffel et al., 2018). While early PIV references modelled some of these effects (such as intensity variations within the laser sheet thickness, as e.g. Westerweel, 2000), more recent approaches tend to rather consider the present choice, either for density estimation (Warner & Smith, 2014), or in the framework of 3D particle reconstruction (Wieneke, 2013; Champagnat et al., 2014), possibly due to the difficulty of accurately modelling each individual effect. Whatever the choice done, equations such as (1) also underlie synthetic image generators, used for PIV/PTV algorithm assessment (see, e.g. Stanislas et al., 2004).

As outlined by e.g. Adrian & Yao (1985), Westerweel (1997, 2000) and Stanislas et al. (2004), model (1) considers that particle images are additive, which holds strictly for incoherent light (such as LED light), whereas it neglects potential effects of interference and speckle at high densities for coherent (laser) light. In the latter case, it can therefore be expected to progressively deviate from reality for increasing source density N_S , this quantity being defined as the fraction of illuminated pixels in the image, and thus equal to $N_S = \lambda \pi d_\tau^2 / 4$. In the above cited references, an acknowledged typical value for significant speckle is identified as $N_S \approx 1$, which is also a classical upper bound for tuning seeding in an experiment, and will be our upper bound as well here.

In the background noise term $N(\mathbf{x})$, a systematically-encountered contribution is the camera thermal noise, which is known to be accurately modelled by a Gaussian distribution at each pixel $\mathbf{x}$, with no spatial correlation (Westerweel, 2000). Spatially correlated noise, if any, stems from the optical setup, e.g. laser light reflections, illumination of background objects by the particles, or even as a result of a variation of seeding density within the image. A common practice to mitigate such a correlated noise is to subtract a historical minimum to the images. In this study, we will tackle these noise contributions in accordance with their systematic or case-related nature: $N(\mathbf{x})$ will be considered as spatially uncorrelated Gaussian noise, while the effect of historical minimum subtraction as well as of residual spatial inhomogeneities will be considered as model error, and assessed using synthetic tests.

Finally, since h is a deterministic function, model (1) as such does not include photon shot noise. In the same way, since its expression does not depend on the particle position $\mathbf{X}_p$, it cannot account for images with defocus nor astigmatism. Contrary to what this notation suggests, our method is in fact fully compatible with these situations: presenting its underlying model in the form of (1) is a deliberate choice to maintain simplicity and not overload the key results. Shot noise, as well as effects linked to varying h (defocus and astigmatism) will be introduced subsequently as

extensions, as they lead to somewhat heavier expressions.

2.3. Probabilistic counterpart and results

2.3.1. PIV image as a filtered Poisson process

Equation (1) usually translates into bright particle images over a dark background, which can enable to locate them individually, e.g. as local maxima (LM) of intensity. While counting these maxima to determine λ is a quite frequently used method (hereafter referred to as the LM method, similar to Warner & Smith, 2014), it is also known to be limited to very small ppp values; indeed, as density and/or diameter d_τ increases, images will overlap so that counting will lead to a strongly underestimated value of λ . The key ingredient of the present study to overcome this limitation is to interpret (1) probabilistically, as a filtered Poisson process (FPP), which allows to rather rely on image statistics to estimate λ . Such an approach is in fact close in spirit to seminal works on PIV fundamentals, such as those of Adrian et al. (1991), Keane & Adrian (1992) and Westerweel (1997), which have used statistics to derive the main characteristics of a PIV image, and of the cross-correlation process. The present analysis can be viewed as an extension to these works, as it exploits further properties of the FPP, in order to model PIV images with significant overlap of particle images.

Switching to this probabilistic context is rather straightforward, as the generic equation for a 2D homogeneous FPP reads (Rice, 1944)

$$I_{FPP}(\mathbf{x}) = \sum_p E_p h(\mathbf{x} - \mathbf{\Pi}_p) \quad (2)$$

where $(\mathbf{\Pi}_p)_{p \geq 0}$ is sequence of independent identically distributed 2D random variables generated by a Poisson process of parameter λ , and $(E_p)_{p \geq 0}$ a sequence of independent identically distributed random variables, each following a probability law of density f_E . In this expression, h is a function that only needs to be continuous and compactly supported. Setting firstly the noise term $N(\mathbf{x})$ aside, comparison of equations (1) and (2) shows that a PIV image can be interpreted as a 2D FPP directly, provided that projected positions $F(\mathbf{X}_p)$, which are the counterparts of $(\mathbf{\Pi}_p)$, can be considered as uniformly distributed over the camera sensor. In such a case, λ will then identify as the image seeding density. Conversely, considering the framework of the FPPs to determine λ thus only imposes to restrict to homogeneous subzones of the image, which is not a very strong limitation. Contrary to the above cited early analyses on the mathematical modelling of PIV, a slight difference here is that we consider directly the *image* density as obeying a Poisson process, whereas, as noted by e.g. Adrian et al. (1991), this rather applied to the volumetric concentration C in a homogeneously seeded flow. In the sequel, by considering in particular optical setups from 3D experiments, in which cameras can be placed at angles quite different from a traditional normal viewing, we will confirm that perspective viewing does not prevent to directly extend the Poisson hypothesis from the 3D to the 2D projected world.

Considering then a more realistic situation including noise is also quite simple, since this can be done additively, similar to the image formation equation:

$$I(\mathbf{x}) = I_{FPP}(\mathbf{x}) + N(\mathbf{x}) \quad (3)$$

so that each pixel's gray level will then be viewed as a superposition of an FPP and of Gaussian noise. As I_{FPP} is the strict probabilistic counterpart of I_P , we will preferentially use the latter notation in the sequel, for simplicity.

2.3.2. Campbell's equations applied to PIV

The interest of switching to filtered Poisson processes is that some of its associated theorems allow to link the seeding density λ to statistical quantities which can be computed directly from the images. These results are known as Campbell's equations (Rice, 1944), which we here generalize for 2D variables, and involve:

- the (non-centered) moments of order $n \geq 1$, defined by $m_n = \mathbb{E}(Y^n)$ for a random variable Y , where $\mathbb{E}$ denotes the expectation. If the probability law f_Y of Y is known, then $m_n = \int y^n f_Y(y) dy$.
- the cumulants κ_n of order $n \geq 1$, defined as the coefficients of the Taylor expansion of the logarithm of the moment-generating function:

$$\log(\mathbb{E}(e^{tI})) = \sum_{n \geq 0} \frac{\kappa_n}{n!} t^n$$

In practice, the cumulants can be easily determined from the moments by

$$\kappa_n = m_n - \sum_{k=1}^{n-1} \binom{n-1}{k-1} \kappa_k m_{n-k}$$

Expectation $\mathbb{E}$ is often estimated using a spatial or temporal average, provided that ergodicity holds. Since we already assume homogeneous subzones, it will thus only be necessary to also assume that subsequently acquired images share the same properties, should more than one image be necessary to reach convergence.

In the context of particle images modelled by equation (3), Campbell's equations provide a set of relations between the cumulants of the images and of the noise, respectively denoted by $(\kappa_n)_{n \geq 1}$ and $(N_n)_{n \geq 1}$, the moments of the distribution of particles intensities E_p , denoted by $(M_n)_{n \geq 1}$, and the seeding density λ , as:

$$\kappa_n = \lambda H_n M_n + N_n \quad (4)$$

where $H_n = \int h^n(\mathbf{x}) d\mathbf{x}$ are the 2D integrals of successive powers of the PSF.

Although this is not strictly considered as Campbell's equations, a result with a similar structure also holds for the auto-covariance of the images, hereafter denoted by κ_δ and defined as

$$\kappa_\delta = \mathbb{E}(I(\mathbf{x})I(\mathbf{x} + \delta)) - \mathbb{E}(I(\mathbf{x})) \mathbb{E}(I(\mathbf{x} + \delta)) \quad (5)$$

with δ a 2D offset in pixel units. Indeed, considering an offset $\delta \neq 0$, one has:

$$\kappa_\delta = \lambda M_2 H_\delta \quad (6)$$

with $H_\delta = \int h(\mathbf{x})h(\mathbf{x} + \delta)d\mathbf{x}$ the corresponding autocorrelation of the PSF. Since background noise is uncorrelated in space, choosing $\delta \neq 0$ precisely eliminates any possible contribution of it, which is particularly advantageous in the context of our method. As this equation might be seen as a way to estimate λM_2 , which is independent of δ , then in practice, one value of δ needs to be considered (of course smaller than the width of the PSF since otherwise (6) would become trivial). In the following, we will systematically consider $\delta = (1, 0)$.

The above equation extends in fact results which have been introduced in the early days of PIV, to the case of images with background noise. Denoting the auto-correlation of a noiseless image by R_I , Keane & Adrian (1992) and Westerweel (1997) have indeed shown it to be linked to the image source density N_S , in an equivalent way. Upon introducing a frequently used decomposition in the form $R_I = R_C + R_F + R_P$ (as for instance summarized in the textbook by Raffel et al., 2018), where R_C is the convolution of the mean intensities, R_F is the contribution of overlapping particles, and R_P to the auto-correlation of particles, then auto-covariance κ_δ is equal to $R_I - R_C = R_F + R_P$, and accounts for both the auto-covariance of particles and of the overlaps.

On the other hand, Campbell's equations (4) seem to have been less investigated in the context of PIV. Su et al. (2017) introduced a similar FPP model and Campbell's equations to model images of speckle patterns and then the estimation of displacement by Digital Image Correlation (DIC). To the best of our knowledge, the only work considering a cumulant of a PIV image with a link to source density N_S , is that of Westerweel (2000), dealing with PIV measurement error. In this study, a simplified version of equation at order 4 is introduced, since the image fourth-order cumulant κ_4 is linked to N_S using a linear fit from synthetic images. In the article, this fit is then reported to lack accuracy when considering different types of images, in particular with a different energy repartition in the light sheet thickness. In light of Campbell's equations (4), this appears quite logical, as the slope of the fit should precisely translate properties of particle intensities. This early result hints in fact to a difficulty of determining seeding density λ using only one of equations (4) and (6), even if they are linked in a simple - linear - way with λ . Proposing a way to still use Campbell's equations and overcome these difficulties will precisely be the objective of subsections 2.4 and 2.5.

2.3.3. Extension to defocus, astigmatism and shot noise

As mentioned above, PIV images include an amount of shot noise, and the occurrence of more or less pronounced defocus and/or astigmatism can be quite common. In this paragraph, we will show that when these effects are accounted for, the key results derived in section 2.3.2 either still hold exactly, or in a slightly modified version only.

The case of defocus and astigmatism, in which the PSF function varies as a function of the particle position in object space $\mathbf{X}_p$, is the simplest one. This function models the shape of the particle image due to diffraction, which, with effects such as defocus or astigmatism, can vary in size or turn to an ellipse depending on the particle position. This range of shapes can be described by a single parametric family (such as done e.g. in Schanz et al., 2013), so that one can formally denote

this varying PSF by $h(\mathbf{x}, \sigma)$, with σ the shape controlling parameter. If the distribution of σ obeys the same probability law f_Σ for all pixels of the image zone considered for determining λ , then it is quite easy to show that in equations (4) and (6), H_n will simply have a different definition, as:

$$\begin{cases} H_n = \int_{\mathbb{R}} \tilde{H}_n(s) f_\Sigma(s) ds \\ \tilde{H}_n(s) = \int_{\mathbb{R}^2} h^n(\mathbf{x}, s) d\mathbf{x} \end{cases} \quad (7)$$

Consequently, our method is also compatible with defocus blur and astigmatism, under the condition that these effects must have uniform spatial properties within the zone considered for ppp estimation (i.e., same amount of defocus or type of image deformations within the whole zone).

Shot noise is linked to the fact that the intensity collected by the sensor results from the arrival of successive photons, the number thereof for a given exposure time thus being variable, and in practice modelled using a probability law. To account for it, we will introduce in equation (1) a modified contribution of the particles intensities, denoted by $\tilde{I}_P(\mathbf{x})$. Another necessary quantity to introduce is the camera gain g , which identifies the number of photons which are necessary to increase by one the value of a pixel's gray level. Probability theory then allows to show that (see Kuba & Panholzer, 2016):

$$\tilde{\kappa}_n = \lambda \sum_{j=1}^n \binom{n}{j} \frac{H_j M_j}{g^{n-j}} \quad (8)$$

with $\tilde{\kappa}_n$ identifying the n-th cumulant of $\tilde{I}_P$. Note that, similar to quantities H_n , M_n and N_n , knowledge of new parameter g will not be necessary for our method to work, as will be described in the next subsections. The situation is even simpler regarding image auto-correlation; indeed, since shot noise is not spatially correlated, using similar notations, one immediately has

$$\tilde{\kappa}_\delta = \kappa_\delta \quad (9)$$

Therefore, upon adding the contribution of background noise $N(\mathbf{x})$, the image probabilistic model obtained in section 2.3.2 will simply be modified as

$$\tilde{\kappa}_n = \lambda \sum_{j=1}^n \binom{n}{j} \frac{H_j M_j}{g^{n-j}} + N_n \equiv \lambda K_n + N_n \quad (10)$$

$$\tilde{\kappa}_\delta = \lambda M_2 H_\delta \quad (11)$$

To account for full image complexity, we will from here on use equations (10) and (11) instead of (4) and (6). We will simply write κ_n resp. $\tilde{\kappa}_\delta$ in place of $\tilde{\kappa}_n$ resp. κ_δ to alleviate notations.

2.4. Density estimator from image statistics

As mentioned in section 2.3.2, using one of equations (10) and (11) to determine the ppp λ from image statistics requires at least knowledge of the PSF (or PSF distribution in case of defocus or astigmatism), as well as of one of the moments M_n of the distribution of particle intensities. Note

that in fact, and as probably also remarked by Westerweel (2000) due to their order choice, knowledge of exact properties of background noise N is not necessary: indeed, since it is assumed Gaussian, it is known that $N_n = 0$ for $n \geq 3$, i.e. noise contributions disappear in Campbell's equations above from order 3 on. While approaches for calibrating the PSF can be quite simple when it is uniform in the images, an accurate estimation of its distribution in the case of astigmatism or defocus can require quite advanced methods (see, e.g. Schanz et al., 2013). The situation is even worse regarding moments M_n of the distribution of particle intensities: as mentioned in section 2.2, this quantity results from the combination of several physical and potentially complex characteristics of the experiments, i.e. the distribution of particle sizes, the variation of illumination intensity in the light sheet, Mie scattering or even multi-scattering as the volumetric concentration becomes large. Whereas Warner & Smith (2014) proposed an empirical estimation of the effects of these parameters (excepted for individual Mie scattering and multi-scattering), their estimation was shown to still exhibit a sensitivity to the distribution of particles intensities, as shown by estimating λ on images of the same flow viewed two different cameras. In the same spirit, we found that our first version of the method (Hebey et al., 2025), which considered a physical model of the distribution of particle intensities, was adapted only to a fraction of possible experimental contexts.

Given these observations, we will try and use Campbell's equation so as to reduce as much as possible the number of setup-dependent parameters to estimate, and in a way that the remaining parameter(s) can be estimated easily. To follow, we will show that it is in fact possible to combine two of these equations to get an estimate $\tilde{\lambda}$ of seeding density λ , only from image statistics, so that a proportionality holds between both, i.e.

$$\tilde{\lambda} = C \lambda, \quad (12)$$

What we mean by proportionality here is that, for a given experimental setup, C should remain quasi-constant as the seeding density λ changes while all other parameters of influence (e.g. laser power, magnification, aperture number, seeding material and device) remain unchanged. Since C will be formed of a combination of quantities H_n , H_δ and M_n , it has indeed no reason to have the same value from a setup to another. Then, if a relationship such as (12) can be shown to apply, the only remaining task is to estimate C , which can be done via an adapted calibration, provided a reference is available. Our proposition for this is detailed in the next subsection. To follow, we will present how we construct $\tilde{\lambda}$.

To find the right form for $\tilde{\lambda}$, a first key question to consider is thus whether quantities H_n , H_δ and M_n (or which among them) can be truthfully considered as independent of the ppp λ . If all optical parameters are considered fixed between calibration and experiment (and do not change during the experiment), then all H_n and H_δ will remain constant, since they are linked to the camera aperture and characteristics. Concerning particle intensity moments M_n however, the answer will be mostly negative, unless a quite strong restriction is posed on the maximum seeding density considered. Indeed, as exposed e.g. in Raffel et al. (2018) and recalled in section 2.2, an illuminated particle will also act as a secondary source of illumination for its neighbors, leading to significant multi-scattering as λ increases. Within our model (1), this means that, considering a virtual situation where particles are progressively added to increase λ , multi-scattering will imply that intensity E_p of a given particle will change depending on λ , meaning that moments M_n will change accordingly. While this phenomenon is too complex to be accounted for precisely in the present study, we advocate that a first crude model of it could be of multiplicative nature. Considering then that particle intensities E_p are multiplied by a constant A when comparing given low

and high density cases, robustness to multi-scattering will translate by C being independent or at most only very slightly dependent on A . Together with the other above remarks, this leads us to build $\tilde{\lambda}$ as follows:

$$\tilde{\lambda} = \frac{\kappa_\delta^3}{\kappa_3^2} = \frac{M_2^3 H_\delta^3}{K_3^2} \lambda = C \lambda \quad (13)$$

While the choice of κ_δ and κ_3 is linked to the above remarks on noise N , the idea of considering a ratio of these two statistical quantities is directly linked to the sought robustness to multi-scattering. Indeed, from (10), one has

$$K_3 = \left(H_3 M_3 + 3 \frac{H_2 M_2}{g} + 3 \frac{H_1 M_1}{g^2} \right)$$

so that when comparing low and high seeding densities based on the multiplicative hypothesis with constant A , denoted by $\tilde{\lambda}_{low} = C_{low} \lambda$ and $\tilde{\lambda}_{high} = C_{high} \lambda$, one will have

$$C_{low} = \frac{M_2^3 H_\delta^3}{\left(H_3 M_3 + 3 \frac{H_2 M_2}{g} + 3 \frac{H_1 M_1}{g^2} \right)^2} \quad \text{and} \quad C_{high} = \frac{A^6 M_2^3 H_\delta^3}{A^6 \left(H_3 M_3 + 3 \frac{H_2 M_2}{A g} + 3 \frac{H_1 M_1}{A^2 g^2} \right)^2}$$

showing that, within this simple model, a variation in C is only due to shot noise. As it is reasonable to assume that A will be at most equal to a few units, one thus sees that a condition for C_{low} and C_{high} to be nearly identical is that $H_3 M_3 \gg \frac{H_2 M_2}{g}$ and $H_3 M_3 \gg \frac{H_1 M_1}{g^2}$. This can hold provided that the intensities of particles are large enough, to guarantee that $H_3 M_3 \gg H_2 M_2$ and $H_3 M_3 \gg H_1 M_1$, and that g is not much smaller than 1. In practice, g is typically of order unity, as inferred from datasheets of typical PIV cameras. Besides, a smaller g will lead to brighter particles and conversely, which is favorable in light of the present discussion, as it means that it is not common to have both a small g and weak intensities (meaning that $H_3 M_3$, $H_2 M_2$ and $H_1 M_1$ would be comparable). Synthetic tests, which we choose not to present here for conciseness, confirmed that, for a considered maximum multiplication $A = 3$, a reasonable maximum value for the difference between C_{low} and C_{high} is of 3%. Consistently with the above remark, such a value is similar between a higher-gain camera ($g = 2.5$) and low particle intensity values, and a lower-gain camera ($g = 0.5$) and higher intensities. Consequently, it is reasonable to assume that

$$C_{low} \approx C_{high} \approx C \approx \frac{M_2^3 H_\delta^3}{H_3^2 M_3^2} \quad (14)$$

confirming robustness to multi-scattering, provided that it approximately obeys a multiplicative behaviour. Experimental results, presented in section 4, will show that this hypothesis on multi-scattering seems indeed to be valid, at least at first order.

2.4.1. Effect of spatially correlated background noise

In experiments, it is common to obtain PIV images with a certain degree of spatially correlated background noise, for instance due to objects in the background illuminated by the particles. Usually, image pre-processing is performed, such as historical minimum subtraction, to suppress the

largest possible amount of it, and reduce the final measurement noise. An interesting question for our method is thus to determine if it can be applied on the raw images, or if similar pre-processing should be considered; additionally, knowing the impact of residual noise, in case of partially successful pre-processing, is also an important point. From a theoretical point of view, this background noise translates by the addition of a constant term in the right-hand-sides of both the equations for the third cumulant and the auto-covariance. We denote these additional terms respectively by P_δ and P_3 . Our estimator $\tilde{\lambda}$ then becomes:

$$\tilde{\lambda} = \frac{(M_2 H_\delta \lambda + P_\delta)^3}{(K_3 \lambda + P_3)^2}$$

Introducing $\Gamma_3 = \frac{P_3}{K_3}$ and $\Gamma_\delta = \frac{P_\delta}{M_2 H_\delta}$ for conciseness, we get:

$$\tilde{\lambda} = C\lambda \cdot \frac{\left(1 + \frac{\Gamma_\delta}{\lambda}\right)^3}{\left(1 + \frac{\Gamma_3}{\lambda}\right)^2}$$

It is straightforward to show that this expression reaches a minimum at $\tilde{\lambda}_{crit} = 2\Gamma_\delta - 3\Gamma_3$, so that $\tilde{\lambda}$ is no longer linear with λ . This quick analysis thus confirms that the same pre-processing for background subtraction as used before the PIV or PTV interrogation should be performed before applying our method. If this pre-processing step still leaves some residual noise, it is still reasonable to assume that its intensity is low compared to the intensities of particles. This allows us to assume that $\frac{\Gamma_\delta}{\lambda}, \frac{\Gamma_3}{\lambda} \ll 1$, so that a first order development gives:

$$\tilde{\lambda} = C\lambda + C \cdot (3\Gamma_\delta - 2\Gamma_3)$$

Consequently, remaining background noise after pre-processing might translate in the appearance of a zero-offset in our estimator $\tilde{\lambda}$, compared to the true value λ . This observation will in particular be confirmed by the synthetic analysis in section 3.3.

2.5. Calibration

The remaining step to complete our method is thus to propose a way for calibrating proportionality coefficient C linking our statistical estimate $\tilde{\lambda}$ and the ground truth density λ . As mentioned above, C is setup-dependent so that calibration must be performed in the exact same experimental conditions (other than the seeding density) as those corresponding to the productive images. Besides, the method yielding the reference values necessary for calibration should be of course as accurate as possible, and also as simple as possible to operate. Based on these constraints, it seems reasonable to pick the LM method as the reference, and thus to perform the calibration at low density. Similar to e.g. Warner & Smith (2014), ppp estimate λ_{LM} of the LM method at ground truth density λ is obtained by considering that each pixel whose gray level is above user-selected threshold I_{LM} , and which is a local maximum over a 3×3 stencil, is a particle. Note that the present choice of this stencil leads to restrict to the most typical particle image sizes, of maximal diameter $d_\tau \approx 4 - 5$ pixels. Indeed, for larger particles and with such a stencil size, the presence

of image noise could lead to a particle yielding several local intensity maxima, which is not the desired behavior of the LM method.

Our dedicated calibration process will be composed of the following steps:

1. Acquisition of a few images at several different low densities, ranging from λ_{min} to λ_{max} . λ_{max} should be as low as possible; if the experiment does not allow it, then a correction of the LM method will be necessary (see step 3. and below). The number of different values acquired should be enough to guarantee a satisfactory fit at step 2., thus 4-5 points minimum are required. If due to the optical setup the density varies within the images, then λ_{min} and λ_{max} will refer to the (nearly) homogeneous subzone that is selected for estimation. The number of images acquired at each density (or total number of pixels) should be in line with values advised for convergence of the cumulants, as studied e.g. in section 3.2
2. Computation of estimates $\tilde{\lambda}$ and λ_{LM} from these images or subzones, and linear fit of $\tilde{\lambda}$ vs. λ_{LM}
3. C is then obtained as the slope of this fit, possibly using a pre-tabulated correction coefficient to account for the possible difference between λ_{LM} and λ , as detailed below.

As already suggested above, if λ_{max} is sufficiently low, then images will exhibit nearly no overlap. Provided as well that their SNR is not too low, then threshold I_{LM} will be easy to select, so that assumption $\lambda_{LM} \approx \lambda$ will hold with good precision. Then, C will be directly obtained as the slope of the fit of $\tilde{\lambda} = C \lambda = C \lambda_{LM}$.

In many situations however, these hypotheses will be difficult to meet, either because the density cannot be decreased enough in the experiment, or due to a moderate or low SNR of the images making the choice of I_{LM} more difficult. To tackle them as well, we here propose to model these deviations of the LM method by a single corrective coefficient, as mentioned at step 3. The main idea is to locally approximate the behavior of λ_{LM} vs. λ in an affine way, i.e. supposing that in the range $[\lambda_{min}, \lambda_{max}]$, one has

$$\lambda_{LM} \approx a_{LM} \lambda + b_{LM}, \quad (15)$$

From this linear approximation, the link between $\tilde{\lambda}$ and λ_{LM} can then be expressed as:

$$\tilde{\lambda} = \frac{C}{a_{LM}} (\lambda_{LM} - b_{LM}) \quad (16)$$

so that the fit between $\tilde{\lambda}$ and λ_{LM} will yield C/a_{LM} , showing that in practice, only the slope a_{LM} of the local linear model is to be known, in situations where it can be expected to be different from 1.

As mentioned above, situations with $a_{LM} < 1$ will arise if λ_{max} is not low enough and/or for low image SNR, making it difficult to select appropriate threshold I_{LM} . In such cases, it is then important to know a_{LM} to guarantee maximum accuracy of the method. However before aiming at estimating it, it is worthwhile emphasizing that an optimal threshold value I_{LM}^{opt} , making sure to obtain the highest possible a_{LM} (closest to 1 for the given images), can be determined practically and easily. Figure 1, representing λ_{LM} for different I_{LM} , first confirms the dependency of λ_{LM} to




Figure 1. Seeding density estimation by local maxima counting, considering two different thresholds: the optimal threshold $I_C = 180$ (blue), and a lower value $I_C = 140$ (red).

I_{LM} . While it most directly shows its impact on zero-intercept b_{LM} , local linear fits confirm that it also leads to slightly different values of a_{LM} .

This quantity in fact exhibits a maximum as a function of I_{LM} , at corresponding value hereafter denoted by I_{LM}^{opt} . Its existence can be understood by the fact that picking a too large I_{LM} can lead to missing particles, leading to lowering a_{LM} . Conversely, decreasing it too much will lead to identifying noise patterns as particles, translating in a non zero offset b_{LM} . Then, considering two different values of λ , it can be understood that increasing λ will lead to add as many local maxima as the additional particle number; however, some of them will coincide with pre-existing local maxima due to noise, finally leading to lower a_{LM} . Since the slope of the fit of $\tilde{\lambda}(\lambda_{LM})$ at step 2 provides C/a_{LM} , and C is completely independent of the LM method, then an easy practical way to identify I_{LM}^{opt} is to perform the $\tilde{\lambda}(\lambda_{LM})$ fit for different values of I_{LM} , and monitor the obtained values of C/a_{LM} . As confirmed in Figure 2 for the two synthetic test cases that will be introduced in section 3, I_{LM}^{opt} will then be associated to the minimum of C/a_{LM} . Note that in figure 2a pertaining to case 1, the fact that an asymptote is rather observed for large I_{LM} is due to the quite favorable SNR of these images: the maximum sampled value for I_{LM} on this figure is still lower than the minimal particle intensity.

The above reasoning relies on the assumption that the properties of the images used for calibration do not vary over the range $[\lambda_{min}, \lambda_{max}]$. In section 4.3, dealing with experimental data for which such variations are observed, we will introduce a dedicated renormalization step allowing to overcome this difficulty.

For a tabulated estimation of a_{LM} to be efficient in bringing the necessary correction of step 3, it




Figure 2. Evolution of $\frac{C}{a_{LM}}$ (obtained by fitting $\tilde{\lambda}$ as a function of λ_{LM}) vs. intensity threshold I_{LM} for the synthetic test cases introduced in section 3.

should be done from a synthetic situation reproducing as closely as possible the values of parameters influencing it. As e.g. noted by Warner & Smith (2014) and Thomas et al. (2025), one of the most important ones is the particle image diameter d_τ . Indeed, larger particle images will lead to particle image overlaps occurring at lower density λ . As mentioned above, the calibration range $[\lambda_{min}, \lambda_{max}]$, and possibly the image SNR as well, might have an influence. Regarding the latter point, this is because, even at I_{LM}^{opt} and at low λ_{max} , and as will be seen in section 3, one might obtain a_{LM} slightly lower than 1. However, while estimating d_τ and $[\lambda_{min}, \lambda_{max}]$ can be done quite accurately, determining the image SNR is much more difficult. Another important point is to be able to provide the associated uncertainty on a_{LM} . Indeed, still in the case of significant λ_{max} , the very estimation of λ_{min} and λ_{max} might have a degree of inaccuracy, due to its possible discrepancy with $\lambda_{min,LM}$ and $\lambda_{max,LM}$, obtained by maxima counting, which are the only quantities accessible to the user. The same holds for the estimation of d_τ , depending on the fitting method used. At the time of writing of this paper, performing this tabulation and related uncertainty quantification is still in progress. Therefore, for the cases to be presented in the next sections, we will instead simply present how we determine each individual value of a_{LM} to be used in the calibration correction, the automatization of which would lead to obtaining tabulated values. Note that for simplicity, we will perform this under the hypothesis that $\lambda_{min,LM} = \lambda_{min}$ and $\lambda_{max,LM} = \lambda_{max}$ for the synthetic test cases, leaving estimation of the impact of deviations in these values to future work.

3. Assessment on synthetic data



3.1. Test case description

To test our model, we generated synthetic particle images, of size 1024×1024 pixels, corresponding to both 2D and 3D virtual observation setups. The former will be considered in sections 3.2 and 3.3, and the latter in section 3.4. Particle intensities E_p are modelled with a χ^2 (i.e. squared Gaussian) distribution. This corresponds to assuming that the distribution of particle sizes is Gaussian, while their values lie in the range for which Mie scattering intensity is on average proportional to the square of the particle diameter (see Raffel et al., 2018). Using these intensities, particle images are then obtained by considering a Gaussian + pixel integration (erf) PSF, parameterized by standard deviation σ_{PSF} . In order to account for slight PSF variations due to defocus blur, we chose σ_{PSF} to obey a Gaussian distribution with average 0.6 pixel (leading to a particle image diameter $d_\tau = 2.4$ pixel, thus respecting the 2 – 3 pixel rule-of-thumb for avoiding peak-locking in PIV, or accurately detecting the peak position in PTV) and standard deviation 0.05 pixel. As justified in section 2, we superimpose noise modelled as a Gaussian distribution of mean μ_{noise} and standard deviation σ_{noise} . Two different cases will be considered corresponding to two different SNR ratios. Figure 3 shows both the noise distribution and the distribution of the maximum intensity of particles' projections. Both plots are normalised for readability purposes. The right column shows details of the images at 0.005 ppp, both saturated at 300 counts. As one can see, synthetic case 1 corresponds to a situation with no overlap between the particle's intensity distribution and thermal, whereas synthetic case 2 is a more degraded situation with such overlap.



In this series of synthetic test cases, we will assume that the virtual experimental setup allows to acquire images at very low density, and consider $\lambda_{min} = 0.001$ ppp and $\lambda_{max} = 0.005$ ppp. Within this range, we consider 5 regularly spaced values, thus with an interval of 0.001 ppp between each. Estimation of seeding density by our method will be performed on images with λ ranging from 0.01 to 0.25 ppp. With the considered $d_\tau = 2.4$ pixel image diameter, this maximal density approximately corresponds to a source density of 1, in line with the limitation of our model mentioned in section 2.2.

3.2. Calibration and estimation at varying signal-to-noise ratio

In this first case, particles positions are randomly drawn within the extent of the images, thus corresponding to a 2D homogeneous seeding and also exactly matching the underlying hypothesis of our method. As our method relies on statistical quantities, a first natural question is that of the number of images, or rather of pixels, to consider in order to obtain a sufficiently converged value of $\tilde{\lambda}$, be it for the calibration or the actual estimation. Figure 4 shows the convergence of this quantity, normalized by its theoretical value (which is known in this synthetic context), for a selection of densities. We consider convergence suitable as soon as the absolute error with respect to the converged value is less than 1%. From this criterion, we observe that a value of 15MPix (resp. 20 MPix) of data are sufficient for case 1 (resp. case 2). Note that the fact that some of these convergence curves do not converge towards a unit value, at the lowest densities, is not a concern here, as it is presumably linked to the difficulty of generating synthetic images at a precise density value when it becomes very low.



(a) Case 1



(b) Case 2

Figure 3. Left: Histograms of pixel intensities corresponding to noise, and to maxima of particle images. Right: Detail of a sample image of the 2D dataset at 0.005 ppp, saturated at 300 counts for visualization purposes.

We then apply our calibration process in the dedicated $\lambda_{min} = 0.001 - \lambda_{max} = 0.005$ ppp range. As shown in Figure 2 in section 2.5, the optimal threshold values for the LM method are found




Figure 4. Convergence of normalized estimator $\tilde{\lambda}$ as a function of the number of pixels considered.

as $I_{LM}^{opt} = 168$ for case 1 and $I_{LM}^{opt} = 144$ for case 2. Interestingly, for case 1, one obtains $I_{LM}^{opt} \approx \mu_{noise} + 4\sigma_{noise}$. Since the distribution of noise is Gaussian, it is found that the optimal setting for counting local maxima amounts to excluding nearly all of the possible impact of noise. On the contrary, when overlap occurs, the optimal threshold is shifted towards lower values. Turning to the estimation of a_{LM} in the $[\lambda_{min}, \lambda_{max}]$ range, Figure 5 confirms that this quantity exhibits a slight dependency on the SNR, with $a_{LM} = 0.965$ for case 1 and $a_{LM} = 0.944$ for case 2.

We then present the results of our density estimation on these two cases in figure 6. Note that in these figure, we label this final estimation by "calibrated estimation", to distinguish from the uncalibrated estimator, $\tilde{\lambda}$, which we will also represent in specific situations. Since estimating the image SNR is difficult in practice, we use for both cases the value $a_{LM} = 0.965$ obtained in the reference case 1, which has a perfect SNR (no particle image with intensity lower than the maximum intensity of noise). As can be seen, the model presents a remarkably linear behaviour with respect to λ for both cases. The maximal error is of 0.7% for case 1 and 2.2% for case 2. These values are consistent with both the convergence error of the estimator and the error in the value of a_{LM} for case 2.

3.3. Spatially-varying noise and background subtraction

The spatially uncorrelated nature of noise is one of the building hypotheses of our model. As shown in the mathematical analysis in section 2.4.1, when however a significant coherent background noise is present, our estimator can experience a deviation of our model from its expected linear behavior with the true seeding density. If it can be brought to a residual level, i.e. at lower level than the particles' intensities, then the estimator becomes linear again, this residual noise simply adding a zero-intercept. To confirm the validity of this simple analysis and confirm the




Figure 5. Illustration of the calibration process on the synthetic datasets: ppp estimation by local maxima counting, λ_{LM} , as a function of the true ppp, for the two cases. Calibration is performed using 5 regularly spaced densities between 0.001 and 0.005 ppp (triangles), yielding the local slope of the local maxima estimation, a_{LM} (see inset).




Figure 6. Results on the 2D synthetic datasets. Plot of our calibrated estimation, together with estimation by counting local maxima.

necessity of adapted image pre-processing, we generate images with the same parameters as case 1, and add to them a similar background as in Mendez et al. (2017) (see sample image in figure 7). We underline that, contrarily to Mendez et al. (2017), no pixels are saturated in the image.




Figure 7. Sample synthetic image with spatially-correlated background noise

Values of our uncalibrated estimator $\tilde{\lambda}$ for the raw images, i.e. considering no background subtraction, are provided in figure 8. The behaviour of the estimator is consistent with the one predicted in section 2.4, with a departure from a linear behavior and the presence of minimum as a function of λ . From this plot, it appears that calibrating at low densities on these images would thus lead to strong errors. To mitigate this effect, in line with pre-processing before PIV or PTV interrogation, we subtract a historical minimum image. To do so, we generate 50 images on which to compute the minimum. The estimator is then computed on 15 images, in agreement with the convergence study of section 3.2. In this case, we found the optimum threshold for the LM method to be $I_{LM}^{opt} = 145$. This optimum is approximately 20 counts lower than the one found in case 1 in section 3.2: this is due to the fact that, in addition to the background, the minimum also removes part of the noise. Figure 10 shows the plot of the calibrated solution with a subtracted minimum computed on 20 images (plot 9a) and 50 images (plot 9b). The latter presents the ideal case of an almost fully converged historical minimum, and figure 9b indeed shows that the calibrated solution is comparable to 6a. However, computing a historical minimum on 50 images may not always be suited, as intensity can vary with time in experimental cases and thus bias the minimum. The plot in figure 9a shows the result using a historical minimum computed on a more limited sequence of 20 images. The error is logically greater at high density, with a maximum at 3.13%. This is due to the fact that, at these densities, the historical minimum is not fully converged and removes some of the particles.




Figure 8. Behavior of the uncalibrated estimator $\tilde{\lambda}$ as a function of the true ppp λ , in the case of 2D synthetic images with spatially-correlated background noise.






Figure 9. Results on 2D synthetic images with spatially coherent background noise, to which a historical minimum has been subtracted. Plot of our calibrated estimation, together with estimation by counting local maxima.

3.4. Images from 3D measurements: perspective viewing and spatially varying seeding density

We now turn to images obtained from a three-dimensional PIV/PTV setup, in order to get closer to the experimental test cases which will be introduced in section 4, which correspond to 3D PTV



(a) Sample image at 0.05 ppp. The red rectangle identifies the central subzone used for evaluation.



(b) Estimation results on the subzone identified in Figure a. Plot of our calibrated estimation, compared with estimation by counting local maxima.

Figure 10. Presentation and results of the 3D synthetic dataset.

campaigns. Since seeding is still considered to be homogeneous in the object space, perspective viewing associated to the multi-camera setup will introduce a deviation from our underlying assumption of 2D homogeneity in the image plane. For this test case, we considered one camera of sensor size 2048×2048 pointing towards a nearly cubic particle volume. The extend in the Z direction of the 3D illuminated volume is indeed of the same order of magnitude as the X and Y directions, in order to maximize the potential impact of 3D epipolar effects. Aside from this geometrical difference, all other parameters were kept to the same values as case 1 of the 2D datasets.

Figure 10a confirms that perspective viewing and the finite extent of the illuminated domain in all directions lead to an inhomogeneous particle density in the image, with edges showing a lower concentration than the central region. To test our method, we then focus on a central 1024×1024 subzone, where density is homogeneous and maximal. The intensity parameters being the same as case 1, I_{LM}^{opt} is logically unchanged. The plot of the calibrated estimation is presented in 10b. The maximal error is of 0.9%, which is commensurable to the error obtained in section 3.2. This plot highlights that, even though 3D viewing constitutes a deviation from the theoretical 2D model, this deviation is negligible as long as the chosen zone still exhibits a seemingly homogeneous seeding density.

4. Assessment on experimental data

In this section, we assess the performance of our method on images from a Rayleigh-Bénard convection (RBC) experiment conducted at DLR Göttingen (Bosbach et al., 2021; Weiss et al., 2024). Images from this case have been selected to build case 3 of the 2nd Challenge on Lagrangian Particle Tracking (Sciacchitano et al., 2026). In the following, we use these challenge image sets as well as additional images at other seeding densities, communicated to us privately by D. Schanz from DLR. The interest of this dataset is twofold: the part of it that has been used within the Second LPT Challenge can be downloaded freely (see https://onera-flow-benchmarks.fr/en/Second_LPT_3_Datasets), and a quite accurate estimation of the total number of particles, and therefore of the image seeding density, is available. As described in particular in Bosbach et al. (2021), seeding has been performed with a specifically designed system allowing to generate long-lifetime Helium-Filled Soap Bubbles (HFSB). The RBC cylindrical tank, of diameter and height both equal to 1.1 m, is illuminated using an LED-array, and observed by a six-camera system. This high number of cameras allows to obtain high-quality LPT processing results using DLR's Shake-The-Box (STB) algorithm, the obtained number of tracks then forming an accurate reference for the number of particles. As done for the Second LPT challenge, this number of particles is then divided by the number of pixels of the active size of the images, to obtain the seeding density. While this estimation has been performed from the six-camera acquisition, all datasets considered here only include the four central cameras, as in the Second LPT Challenge. In the sequel, these will be numbered from 0 to 3.




Figure 11. Sample image of the experiment (camera 3), corresponding to 0.04 ppp. The red square identifies the subzone chosen for density evaluation with our method.

Table 1 summarizes the characteristics of the images used. Available seeding densities range from 0.005 to 0.14 ppp. For a majority of them, 50 images corresponding to a time-resolved (TR) acquisition are available. We also include images that have been chosen for the two-pulse (TP) cases of the 2nd LPT Challenge; given that only 2 images are available, it can be expected that statistics will

Average seeding density (ppp)	Number of images	Acquisition type	Available for download
0.005	2	TP	X
0.01	50	TR	X
0.015	2	TP	X
0.02	50	TR	
0.025	50	TR	
0.03	2	TP	X
0.04	50	TR	X
0.045	2	TP	X
0.05	50	TR	
0.06	2	TP	X
0.07	50	TR	X
0.075	2	TP	X
0.08	50	TR	
0.09	50	TR	
0.1	50	TR	X
0.11	50	TR	
0.12	50	TR	
0.13	50	TR	
0.14	50	TR	


Table 1. Seeding densities and number of images of the experimental Rayleigh-Bénard convection dataset from DLR Göttingen (Bosbach et al., 2021; Weiss et al., 2024). Images from case 3 of the 2nd LPT Challenge are available for download (X) at https://onera-flow-benchmarks.fr/en/Second_LPT_3_Datasets.

be less converged at the corresponding densities. An example of experimental image is presented in figure 11. Due to the cylindrical nature of the experimental container, the density is not uniform across the image. In line with our method's assumptions, and as in section 3.4, we choose a sub-zone where density and intensities of particles are seemingly homogeneous. This zone, of size 1300×1300 pixels, is identified by the red square in figure 11. As it is located in the centre of the image, it is expected to exhibit a higher density than the global seeding density estimation from STB. To determine its exact value, we used the self-calibration performed during our participation to the 2nd Challenge to synthetically reproduce the experimental setup. Virtual particles were distributed homogeneously within the cylinder and synthetic images generated on the self-calibrated cameras, allowing to estimate, for each camera, a multiplicative factor between the average density and the density in our subzone. This factor was respectively found equal to 1.15, 1.14, 1.15 and 1.16 for cameras 0, 1, 2 and 3. For consistency's sake, in the text we will however continue referring to the data at a given density by using the average density, while figures will use the exact subzone density.


4.1. Image pre-processing

As shown by figure 11, the subzone considered for evaluation lies in a central region, quite far away from the cylinder bottom and top faces. However, a certain degree of background noise from the upper face can be perceived. Consistently with the analyses shown in sections 2.4.1 and 3.3, we thus subtract a historical minimum as a first pre-processing step. Comparing the corresponding final estimation with an estimation performed on the raw images (not shown here for conciseness) confirms the importance of this step. For the TR densities, the historical minimum was computed over the 50 available images. For the TP densities, we used historical minimums computed on 20 images provided by DLR within the 2nd LPT challenge, which are accessible at https://onera-flow-benchmarks.fr/en/Second_LPT_3_Datasets.

By nature, the third order cumulant is sensitive to pixels with highest intensities. For the estimated value of this cumulant to not be biased by rare and unconverged events, it is important to exclude intensity values which would both be of very high value and rarely encountered in the images. Our second post-processing step thus consists in identifying these pixels based on a threshold, and set their intensity values to 0. In such a process, the choice of this threshold is also of importance. For this experimental dataset, we monitor the evolution of maximal intensities within the image sequences of the TR densities. An example of this is presented for camera 3 at 0.07 ppp, in figure 12a. This plot quite clearly allows to distinguish, among the highest intensities, which ones are frequent and will participate to the cumulant's convergence, and which ones might bias its estimation due to their rareness. To ensure consistency across the whole image sequence, we then build the list formed by the intensity maxima of each image, and then choose our threshold as the minimum value of this list. This ensures that every intensity taken into account in the distribution is present in all images. For the TP densities, as less information is available, we set the threshold to be 1.02 times the threshold of the TR sequence whose density is immediately below the considered TP sequence. The reason for picking a multiplicative factor slightly above unity is due to the observation, appearing in figure 12b (plotting the evolution of the threshold as a function of seeding density), that the threshold monotonically increases with the ppp. This observation suggests that multi-scattering is present in the images at sufficiently high ppp, which will be further investigated in the sequel.



(a) Maximal image intensity at each instant of the TR sequence at 0.07 ppp



(b) Evolution of the threshold used to filter out rare high intensity events, as a function of seeding density

Figure 12. Study of the maximum threshold for pixels' intensity (camera 3)

4.2. Convergence

In this section, we study the convergence of the model with respect to the number of images. To do so, we plot the evolution of the uncalibrated estimator $\tilde{\lambda}$ as a function of the number of images for the four cameras. For clarity's sake, we restrict to densities 0.01, 0.04, 0.07, 0.1 ppp, which were deemed to be representative of the range of densities studied. Values of $\tilde{\lambda}$ are normalized using a mean computed on the last five values. The results are presented in figure 13.

Given the plots across all four cameras, we see that $\tilde{\lambda}$ converges under 1% error for 20 images, the convergence being achieved with less images at higher densities. We also observe that the convergence seems quicker for cameras 2 and 3. We underline that the number of corresponding MPix, namely 33.8MPix, despite being higher than the one required in synthetic cases (see section 3.2), is within the same order of magnitude. Hereafter, we will thus be presenting results computed on 20 images for the TR points. Moreover, figure 13 shows that the TP densities, where we only have 2 images, are not fully converged, so that the results on these densities are expected to be untrustworthy.

4.3. Calibration

As mentioned in section 2.5, determining the optimal threshold value used for counting local maxima during calibration supposes that a single value is adapted to all ppp within range $[\lambda_{min}, \lambda_{max}]$. This suggests in particular that the noise properties of all images remains similar, which can be a strong assumption if λ_{min} and λ_{max} are too significantly different. This will be precisely the case with the present experimental data, where we consider for calibration the four TR datasets at $\lambda_{min} = 0.01, 0.02, 0.025$ and $\lambda_{max} = 0.04$ ppp, corresponding to ppp values in the studied subzone




Figure 13. Convergence of uncalibrated estimator $\tilde{\lambda}$ for selected densities for all cameras

from 0.011 to 0.046. Note that even if the lowest density of the dataset is even lower, equal to 0.005 ppp, we do not consider it here as only 2 images are available. To cope with such situations where the noise values are different around $\lambda = \lambda_{min}$ and $\lambda = \lambda_{max}$, we propose a renormalization step, to be now described, which will allow to find a single value for optimal threshold I_{LM}^{opt} over the whole $[\lambda_{min}, \lambda_{max}]$ range.

To derive it, we first analyse the evolution at each density of value λ_{LM} obtained by local maxima counting, as a function of threshold I_{LM} . The plots for the two synthetic cases of section 3 and of the present experimental case are showed in figure 14 for densities varying between 0.005 and 0.06 ppp. Plots are normalized by their maximal value to allow comparison.

As expected, λ_{LM} remains constant when threshold I_{LM} is below the minimal intensity of noise. As it then sweeps across noise values, a steep decrease is observed, then stabilizing or showing a much weaker decrease, for values larger than the maximal intensity of noise. Stabilization is




Figure 14. Evolution of density λ_{LM} estimated by local maxima counting as a function of intensity threshold I_{LM} , for both synthetic cases of section 3 and the present experimental case.

observed if the images exhibit a clear separation between particle intensities and noise (no particle with intensity lower than the maximum intensity of noise, synthetic case 1 in figure 14a), whereas a continued slow decrease happens otherwise (figure 14b for synthetic case 2). Interestingly, the shape of curves in figure 14c thus hints to the fact that the distribution of particle intensities overlaps with noise in the present experimental case.

Given the shape of the plots in figure 14, a natural next step to identify relevant values for I_{LM} is to represent the second derivatives of the obtained curves as a function of I_{LM} . This is done in figure 15, where the opposite of this second derivative is shown, and curves have been normalized by their maximal value. As one can see, the plots all collapse in both synthetic cases, whereas they show much more variability for the experimental case. Each plot shows a maximum and a minimum, respectively reached at thresholds which we denote by I_{max} and I_{min} . As synthetic cases are characterized by common noise parameters and different intensity distributions, we assume that the values of these extrema characterize the noise distribution. Thus, an adapted normalization using associated intensity value I_{norm} could be obtained by forming I_{norm} from a combination of I_{max} and I_{min} . We will here choose $I_{norm}(\lambda) = I_{max}(\lambda) + 3 \cdot (I_{max}(\lambda) - I_{min}(\lambda))$, because for both synthetic cases we find $I_{norm} \approx \mu_{noise} + 4 \cdot \sigma_{noise}$, which corresponds to the optimal threshold I_{LM}^{opt} of synthetic case 1 (see section 2.5), and hence confirms the significance of this quantity relative to the noise level. Our renormalization will thus consist in multiplying the images ratio $\frac{I_{norm}(\lambda_{min})}{I_{norm}(\lambda)}$.

After this renormalization, the optimal value I_{LM}^{opt} found in the present experimental case is respectively of 138, 126, 148 and 143 counts for cameras 0 to 3. An example of determination of this value is shown for camera 3 in figure 16a.

Due to our choice for the values of λ_{min} and λ_{max} , which are much higher than in the synthetic




Figure 15. Opposite of the second derivative of density λ_{LM} estimated by local maxima counting as a function of intensity threshold I_{LM} , for both synthetic cases of section 3 and the present experimental case.




Figure 16. Method calibration in the experimental case. Data from camera 3.

cases of section 3, it can be expected that density λ_{LM} provided by local maxima counting will exhibit a significant discrepancy with respect to true density λ . This will translate into a value of a_{LM} possibly significantly lower than 1, making it essential to know the value of this quantity for calibration correction (step 3 described in section 2.5). To determine it, we simply generated synthetic

images using all parameters of synthetic case 1, except for the particle image diameter d_τ , which was set to $d_\tau = 1.6$ px, as obtained from a Gaussian fit performed on local intensity maxima at the lowest ppp. This resulted in $a_{LM} = 0.738$, as shown by Figure 16b. As will be seen in the results to come, this determination of a_{LM} performed on synthetic images with very different characteristics than the present experimental images (except for d_τ) will lead to very accurate results. This confirms the soundness of our assumption, made in section 2.5, that a_{LM} can be tabulated offline for use in a variety of practical situations.

4.4. Results




Figure 17. Calibrated results in the experimental case. The color code is the following : camera 0, camera 1, camera 2, camera 3. Plot of our calibrated estimation, together with estimation by counting local maxima. Two-pulse (TP) densities, for which only 2 images are available, are singled out with circle symbols.

The results obtained by our method, after it has been calibrated, are presented in 17. In order to assess the accuracy of the calibration process, we also present results obtained by considering the available reference ppp estimation (obtained from Shake-The-Box processing on the six-camera dataset) during the calibration process, instead of the estimation by local maxima counting. More precisely, at step 2. of the calibration process, $\tilde{\lambda}$ is fitted against this reference instead of λ_{LM} , thus avoiding the need of any correction unnecessary at step 3.

As with synthetic cases, we observe that the method shows a nearly linear behavior, confirming our physical assumptions and models. Exceptions to this are the two highest seeding densities, where a deviation is observed. Aside from this anomaly, for which we will provide an interpretation in section 4.5, our method is seen to yield an estimation of seeding density with very high accuracy up to roughly 0.14 ppp. Besides, comparison between the two sets of curves, which are

sensibly similar, confirms that our calibration process relying on local maxima counting at low density guarantees a good performance. This tends to confirm the hypothesis that the local maxima counting method presents trends of a universal behavior. If we except the two highest densities, the largest error is of 0.016 ppp for cameras 1 to 3, and of 0.022 ppp for camera 0. These values correspond to densities equal to 0.13 ppp for cameras 2 and 3, and 0.14 ppp for cameras 0 and 1. Moreover, as it is possible, in spite of the established strong performance of the STB algorithm, that its estimation of the ppp might still come with some uncertainty, we deem this error to be an upper bound of the real error.

Note that, in order to provide an estimate of the method's random uncertainty, Figure 17 also includes, in the form of error bars, the variability of the estimation found by considering different sequences of 20 consecutive images within the 50 instant TR datasets. For readability, these error bars were only displayed for the three highest densities only, as they were almost undistinguishable otherwise; indeed, we found a maximal relative variability of the result of 2% across all four cameras.

4.5. Further use of the image cumulants




Figure 18. Evolution of normalized values of κ_δ, κ_3 and $\tilde{\lambda}$ as a function of the reference seeding density

To explain the deviation observed in our estimation at the two highest densities of this dataset, it is worth examining not only the uncalibrated estimator $\tilde{\lambda}$, but, as well, its two contributors, the auto-covariance κ_δ and the third-order cumulant κ_3 . Their evolutions as a function of the ppp, normalized by their first value, are presented in figure 18.

This plot shows that the last two values of both κ_δ and κ_3 show a sudden decrease with respect to the previous ones, while $\tilde{\lambda}$ shows an increase. This hints towards a change in the intensity

distribution, so that the calibration, which does not include these points, is logically no longer valid. If we delve further in the interpretation, the decreases in both κ_δ and κ_3 are associated to decreases in the non centered moments M_2 and M_3 of the distribution of particle intensities, indicating its less spread and less asymmetrical behavior. One can confirm this observation by examining the histograms of pixels' intensities, in a log-log plot to highlight their differences. This is done for selected ppp values in figure 19. As one can observe, the general shape of the histogram of the last two densities is different from the other ones, as they show more intensities in the [100, 1000] counts range, followed by a steeper decrease for intensities greater than 2000 counts. This also shows that considering image histograms in the present log-log form is a valuable tool for monitoring a change in the distribution of particle intensities, and consequently a loss of accuracy of the calibrated estimator.

Moreover, figure 18 also corroborates our idea of using a ratio of cumulants instead of a raw cumulant to build estimator $\tilde{\lambda}$, and thus the relevance of considering the effect of multi-scattering as a multiplication of intensity distribution. Indeed, $\tilde{\lambda}$ keeps its linear behavior with a constant slope throughout (until 0.14 ppp), while κ_δ and κ_3 first tend to show a linear behavior until 0.025 – 0.03 ppp, and then deviate from it, with a certain amount of curvature appearing. These observations are to be confronted to characterizations of the average particle intensities reported by Bosbach et al. (2021) in the description of their experiment. Due to the nature of the chosen seeding (helium-filled soap bubbles) and to the closed flow under consideration, these authors first injected a large quantity of bubbles, and performed their acquisitions in a continuous manner after some settling time. Images at progressively decreasing ppp were thus obtained, as HFSB progressively burst. Concomitantly, large but quite quickly decreasing average intensities, which the authors ascribed to secondary scattering effects, were observed in a first phase (before $t = 600$ s), to then to reach a quasi-plateau beyond $t = 600$ s. Given the mean residence time of the used HFSB of 326 ± 5 s, one can link $t \geq 600$ to a global density $\lambda \leq 0.025$ ppp. This consequently confirms that our estimator $\tilde{\lambda}$ is robust to multi-scattering occurring at high seeding densities ($0.025 - 0.03 \leq \text{ppp} \leq 0.14$, the upper bound being due to a change in the property of particle intensities in the present experiment as explained above), but also to a switch towards a regime where these effects become much less pronounced or even absent ($\text{ppp} \leq 0.025$). As for the log-log histograms in figure 19, this quick analysis additionally confirms that cumulants can also be a diagnostic tool of the evolution of the particles' intensity distribution in the fluid.

5. Conclusions

Viewing a PIV or PTV particle image as a filtered Poisson process allows to build an estimator linearly depending on the seeding density. This estimator is based on the computation of the statistical moments of the image's pixel intensities and has been designed to be robust to a multiplicative dilatation of the particles' intensity distribution, that we assume to be a first order model of multiscattering. A calibration process based on the local maxima counting (LM) method has been suggested. For now, this calibration relies on the restrictive hypothesis that the upper and lower bounds of the density range used for calibration are known. Synthetic tests have shown robustness up to very high densities (up to 0.25 ppp). Additional hardships such as coherent background subtraction or effects of 3D viewing have shown to be handled well by our estimator. Finally, confronting our method to experimental data has confirmed the excellent performances of our estimator up to high experimental densities. Further analysis of these data also hinted us




Figure 19. Log-log histograms of pixels' intensities for various images. The labelled densities correspond to the global density of the image.

at the soundness of modelling multiscattering as a global multiplication of particles' intensities at first order.

The main perspective of our work is to delve further into the local maxima counting method in order to offer a more general calibration method that relies on fewer restrictive hypothesis.

Acknowledgements

This work is supported by the French Government defence procurement and technology agency (DGA) and by the French National Research Agency (ANR) under project BANG, contract ANR-22-CE30-0025. The authors would like to especially thank Daniel Schanz, from DLR Göttingen, for providing the images of the Rayleigh-Bénard convection experiments. Adam Cheminet, Philippe Cornic and Jean Le Bris are also acknowledged for fruitful scientific discussions.

References

- Adrian, R. J., et al. (1991). Particle-imaging techniques for experimental fluid mechanics. *Annual review of fluid mechanics*, 23(1), 261–304.
- Adrian, R. J., & Yao, C.-S. (1985). Pulsed laser technique application to liquid and gaseous flows and the scattering power of seed materials. *Applied optics*, 24(1), 44–52.

- Bosbach, J., Schanz, D., Godbersen, P., & Schröder, A. (2021). Spatially and temporally resolved measurements of turbulent rayleigh-bénard convection by lagrangian particle tracking of long-lived helium-filled soap bubbles. In *Proceedings of 14th international symposium on particle image velocimetry* (Vol. 1, pp. 1–12).
- Champagnat, F., Cornic, P., Cheminet, A., Leclaire, B., Le Besnerais, G., & Plyer, A. (2014). Tomographic PIV: particles versus blobs. *Measurement Science and Technology*, 25(8), 084002.
- Hebey, M., Leclaire, B., & Champagnat, F. (2025). Robust estimation of particle image density up to high values through probabilistic modelling. In *16th international symposium on particle image velocimetry*.
- Keane, R. D., & Adrian, R. J. (1992). Theory of cross-correlation analysis of piv images. *Applied scientific research*, 49(3), 191–215.
- Kuba, M., & Panholzer, A. (2016). On moment sequences and mixed poisson distributions. *Probability Surveys*, 13, 89–155.
- Mendez, M. A., Raiola, M., Masullo, A., Discetti, S., Ianiro, A., Theunissen, R., & Buchlin, J.-M. (2017). Pod-based background removal for particle image velocimetry. *Experimental Thermal and Fluid Science*, 80, 181–192.
- Raffel, M., Willert, C. E., Scarano, F., Kähler, C. J., Wereley, S. T., & Kompenhans, J. (2018). *Particle image velocimetry: a practical guide*. springer.
- Rice, S. O. (1944). Mathematical analysis of random noise. *The Bell System Technical Journal*, 23(3), 282–332.
- Sankaran, A., Hain, R., & Kähler, C. J. (2024). Particle image based simultaneous velocity and particle concentration measurement. *Measurement Science and Technology*, 35(6), 065206.
- Schanz, D., Gesemann, S., Schröder, A., Wieneke, B., & Novara, M. (2013). Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Measurement Science and Technology*, 24(2), 024009.
- Sciacchitano, A., Godbersen, P., Grauer, S., Grille Guerra, A., Leclaire, B., Schanz, D., & Schröder, A. (2026). Main results of the second ipt challenge. In *Proceedings of the 22nd international symposium on the application of laser and imaging techniques to fluid mechanics*.
- Stanislas, M., Westerweel, J., & Kompenhans, J. (2004). Proceedings of the europiv 2 workshop held in zaragoza, spain, march 31-april 1, 2003. *Particle Image Velocimetry: Recent Improvements*.
- Su, Y., Zhang, Q., & Gao, Z. (2017). Statistical model for speckle pattern optimization. *Optics express*, 25(24), 30259–30275.
- Thomas, L., Tremblais, B., Mercier, B., & David, L. (2025). Estimation of particle number and detection error from simple peak detection for ptv and ipt. In *16th international symposium on particle image velocimetry–ispiv 2025*.

- Timmins, B. H., Wilson, B. W., Smith, B. L., & Vlachos, P. P. (2012). A method for automatic estimation of instantaneous local uncertainty in particle image velocimetry measurements. *Experiments in fluids*, 53, 1133–1147.
- Warner, S. O., & Smith, B. L. (2014). Autocorrelation-based estimate of particle image density for diffraction limited particle images. *Measurement science and technology*, 25(6), 065201.
- Weiss, S., Schanz, D., Erdogan, A. O., Schröder, A., & Bosbach, J. (2024). On lagrangian properties of turbulent rayleigh–bénard convection. *Journal of Fluid Mechanics*, 999, A90.
- Westerweel, J. (1997). Fundamentals of digital particle image velocimetry. *Measurement science and technology*, 8(12), 1379–1392.
- Westerweel, J. (2000). Theoretical analysis of the measurement precision in particle image velocimetry. *Experiments in Fluids*, 29(Suppl 1), S003–S012.
- Wieneke, B. (2013). Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology*, 24(2), 024008.