

A new camera model combining an analytical model and a discrete correction to overcome refractive index variation challenges

G. Acher¹, L. Thomas¹, B. Tremblais², L. David¹

Abstract

When performing Tomo-PIV experiments, the calculation of accurate camera models is a key point for successful measurements. In severe configurations, where optical interfaces are involved and refractive index variations occur along the line of sight, analytical models can fail to accurately represent the projection and back-projection functions. This inability can lead to significant inaccuracies in particle location and volume reconstruction, which has a considerable impact on the calculation of velocity fields. In order to overcome these limitations, an innovative camera model based on the combination of an analytical model, such as a pinhole model or a polynomial model, with discrete corrections is proposed. In this method, the analytical projection and back-projection are adjusted with a discrete correction stored in two adaptive grids that save both memory and computation time. These correction grids require calibration which is performed with triangulation procedures similar to those used in misalignment corrections. The calculation and operation of the model are described in this paper. The performance of the camera model is evaluated on simulated and experimental setups based on a large depth-of-field calibration performed in a glass water tank. The presence of multiple optical interfaces and fluids, resulting in large light deviations, makes it difficult to compute a high accuracy camera model. In this configuration, the proposed technique successfully reduces the triangulation error from 1 pixel to less than 0.01 pixel. The usefulness of the model is demonstrated in a Tomo-PIV experiment where the deflection of light through the water tunnel walls prevents classical analytical functions from accurately modelling the projection and back-projection functions. This corrected model can also solve discontinuity problems in the projection functions and can be used when there are interfaces in the measurement volume. It opens new perspectives in the study of fluid-structure interaction when transparent solids are involved.

Keywords: Multiple camera calibration, Optical interface, Triangulation optimization, Discrete camera model, Tomo-PIV

1 Introduction

The processes of calculation and calibration of camera models are crucial operations when implementing an experiment with tomographic particle image velocimetry (Tomo-PIV) or Lagrangian particle tracking (LPT) measurements. Numerous publications have stressed the importance of accurate calibration (Scarano [2012]), analysed the effect of error in the camera model (Thomas et al. [2014]) and even addressed the problem of improving the calibration for volume reconstruction (Wieneke [2008] and Wieneke [2018]).

¹Institut Pprime, UPR3346, CNRS - Université de Poitiers - ISAE-ENSMA, France

²Institut XLIM, UMR 7252, France

However, in many experimental setups, even if accurate calibration procedures -and even de-alignment corrections- are implemented, the projection errors remain large enough to prevent the performance of Tomo-PIV calculations with the desired accuracy. This lack of accuracy is due to the inability of the chosen analytical camera model to represent the complexity of the transformation between the physical reference frame and the image plane. Such a configuration is reported in Earl et al. [2016], which studied the mixing of a heated jet in a cross-flow. The study points out that the self-calibration technique does not achieve adequate accuracy throughout the volume. In order to overcome this difficulty, the volume reconstruction is divided into two areas and calculated with two different sets of camera models. Many other experimental setups, such as the measurement of compressible flows, or reactive flows (Lecordier et al. [2012]), or even the measurement of constant density flows through complex or thick windows, pose this type of problem.

The problem described here is caused by local variations in the refractive index in the medium being measured, often due to local fluctuations in the density of the fluid, or to index variations along the optical path between the measurement volume and the camera. When the refractive index variations are sufficiently small in intensity and gradients, these problems can be solved by using high-order camera models such as high degree polynomials (Soloff et al. [1997]). However, when the gradients are too steep, the camera model is not able to accurately represent the solution because continuous models cannot accommodate such very local variations.

Very little development has been done on camera models in the fluid dynamics community since the advent of tomographic PIV and the models still used for Tomo-PIV were developed for Stereo-PIV with the aim of mapping a "thick plane" (see Kwon and Casebolt [2006], Weng et al. [1992], Soloff et al. [1997] and Scarano et al. [2005]). Although under many favourable experimental conditions these models are suitable, in configurations with pronounced distortions or high depth measurement volume they lack local accuracy. The best course of action is always to strive for favourable experimental conditions, but this is not always possible due to reduced optical access or the nature of the research and/or the facility used. In order to overcome these limitations, Belden [2013] proposes a ray-tracing camera model that is adaptable to many configurations but requires a detailed knowledge of the shape of the interface (although most parameters are calibrated). Lecordier et al. [2012] uses advanced image corrections in combination with high-order polynomials to deal with distortions caused by a cylindrical wall for internal combustion studies. More recently, Muller et al. [2020] have used images of targets placed at random positions to calibrate underwater cameras. Paolillo and Astarita [2022] have developed a ray-tracing camera model to make measurements inside a transparent cylinder immersed in water, while in Paolillo and Astarita [2021], the accuracy of the previous model has been compared to the case of a 7th order polynomial model. They show that they are equivalent.

As a general solution to this problem, we present in this paper a new camera model with increased accuracy. This camera model, based on the combination of an analytical model (e.g. linear or polynomial) and a local discrete correction, can model projection operations with local high gradient variations and possibly capture discontinuities if the correction grid is locally refined enough. It adapts to configurations where strong light deflections occur and can be useful in many cases where interfaces are involved or where refractive index changes in the measurement field occur. The new model is an innovative way of dealing with these light deflection challenges and offers a new approach to problems such as fluid-structure interaction

studies where transparent models are involved.

The paper is structured as follows. The first two sections are devoted to the description of the operating principles of the camera model and its calibration procedure. Then, the performance of the camera model is evaluated on a numerical and experimental test case in which the model is used to perform an accurate calibration in a water tank through a glass window with a large depth of field. Finally, the section 5 presents a Tomo-PIV experiment that was performed using the developed model to overcome the problem of distortions due to the thick walls of the water tunnel.

2 Working principles of the camera model

The proposed camera model is based on the combination of a classical analytical model (pinhole, polynomial or other analytical model) and a local discrete correction. The choice to use a combination of the analytical and discrete approaches was made to take advantage of the ability of the discrete approach to model strong local variations in the projection model while using the analytical approach to model global variations in the projection function in order to keep data storage reasonable. The use of discrete correction allows the modelling of steep gradients and even discontinuities if the mesh is locally refined enough. Thus, it is suitable for configurations where strong light deflections are caused by the presence of an optical interface along the optical path between the measurement volume and the camera or even within the measurement volume.

The projection and back-projection corrections are stored in a dedicated data structure. They are calculated during a calibration procedure described in the next section. These corrections depend on the three dimensions of the physical space, which is described by the coordinate system (X, Y, Z) . The coordinates (u, v) are the coordinates of the pixels in the image. The projection of a physical point in space is calculated as the sum of the analytical projection (X, Y, Z) and a pixel correction (du, dv) depending on the location of the volume point to be projected (eq. 1). Since the correction is stored as a physical volume variable, the calculation of the back-projection of an image point onto a $Z = cst$ plane is performed in two steps: first, the analytical back-projection is computed by returning a point $(X', Y', Z') = \Psi^{-1}(u, v, Z)$ and then the correction $(dX, dY, 0)$, computed in terms of the back-projected point (X', Y', Z') , is added (eq. 2). No dZ correction is added in the backprojection operation since it is always performed on a fixed Z coordinate.

$$\begin{cases} (u, v) = \Psi(X, Y, Z) + (du(X, Y, Z), dv(X, Y, Z)) & (1) \end{cases}$$

$$\begin{cases} (X, Y, Z) = \Psi^{-1}(u, v, Z) + (dX(\Psi^{-1}(u, v, Z)), dY(\Psi^{-1}(u, v, Z)), 0) & (2) \end{cases}$$

3 Model calibration

The calibration of this newly developed camera model is performed in two distinct phases. The first phase corresponds to a standard calibration with a manufactured calibration chart used to calculate the analytical model on which our projection model is based. Classical methods are well adapted to carry out this operation (see Calluau [2003], Zhang [1999]); this aspect will therefore not be detailed in this article. Once the analytical model has been calculated and potentially improved by

advanced techniques, such as volume auto-calibration (see Wieneke [2008], Wieneke [2018]), the second phase, corresponding to the calibration of the corrections, begins.

The discrete correction is intended to compensate for the lack of precision of the analytical model to accurately describe the imaging process. In order to calculate the correction (du, dv) for the projection function Ψ , both the position of the particle in the image and its position in the real world must be known. Once these are evaluated, the correction can be calculated as the difference between the projection of the particle with the analytical model and its position in the image (eq. 3).

$$(du, dv) = (u, v) - \Psi(X, Y, Z) \quad (3)$$

Therefore, like many proposed and used projection correction algorithms, this iterative process relies on triangulation of the imaged particles to assess the most likely location of the particles in the volume of interest. The procedure is briefly described here in a few lines, and will be explained in more detail in the following paragraphs. All imaged particles are located in each image acquired simultaneously. From this detection of particles in the image planes, the most probable locations of the real-world particles are determined by triangulation. The real-world particles are then projected onto the CCD plane of each camera to calculate the reprojection error. A statistically representative reprojection error is then calculated within a local volume box to determine the correction to be applied during projection. Corrections for back-projection are evaluated in the same way. At the end of the procedure, the local corrections are determined and stored for both projection and back-projection and the process can be restarted with an improved triangulation as the camera models have been improved with the correction. The main steps of the discrete model calibration are shown in Figure 1.

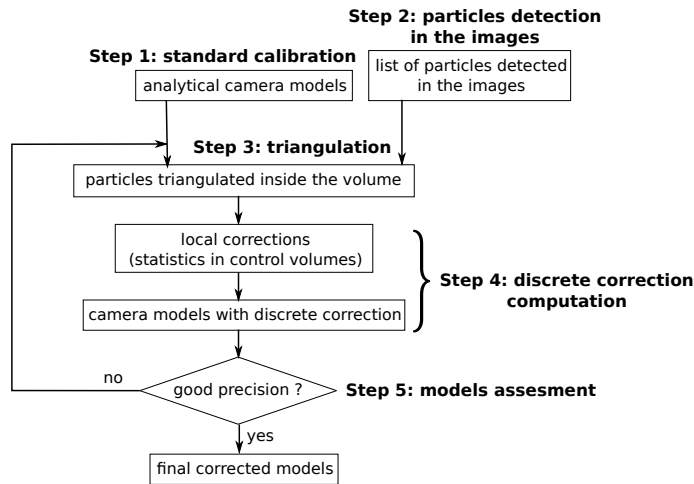


Figure 1: Algorithm of the camera model calibration

The calibration procedure will now be detailed with a four-camera configuration, as this is a standard in Tomo-PIV measurements, but it can be applied for a number of cameras greater than or equal to two.

Step 1: The first step in the calibration of the system is the calibration of the analytical model. This can be done using any of the calibration techniques developed for volume calibration in camera vision, as in Soloff et al. [1997]. In this study, this first step was carried out by precisely translating a fabricated calibration plate along the Z-axis in a range corresponding to the depth of the volume of interest. All markers,

whose positions are known in the physical reference frame, from the characteristics of the plate, the translation step and the initial positioning, are detected, for each camera, in the image sequences. Knowing the coordinates in the physical reference frame and in the images of each marker, the chosen analytical model is calculated by best fitting to the data. If a strong misalignment between the cameras is observed during the analysis of the recorded particle image series, these analytical models can be corrected by applying correction algorithms such as volume auto-calibration (Wieneke [2008]). It should be noted that it is unnecessary to store the correction of the analytical model as the discrete correction can contain it (Wieneke [2008]). Nevertheless, there are some arguments for keeping the analytical model or not: the analytical model is useful to interpolate or extrapolate the camera model, but it can be less effective if the analytical model is sophisticated and requires a high computational cost.

Step 2: The basic model is calculated, the discrete correction has to be calibrated. To this end, in step 2, all particle centres in the image quadruplets are detected by performing a local maximum detection and an anisotropic Gaussian fit on a near neighbourhood (typically 3 pixel \times 3 pixel). These particles, detected once and for all in the correction calibration process, are then used to determine the most likely location of the particles in the volume of interest by the triangulation procedure.

Step 3: The triangulation procedure performed in step 3 allows the epipolar particles contained in the images to be matched and the physical coordinates of the source particle to be determined. In order to have a reliable triangulation and to avoid, as much as possible, the formation of ghost particles, it is strongly recommended to perform the triangulation with a low density of particles, in accordance with the expected projection error. Ideally, images acquired with low seeding should be used, but it is possible to calculate the triangulation on a percentage of the brightest particles if such low particle density images are not available.

In the case of perfectly accurate camera models, the four lines of sight derived from the centre of the epipolar particles in the images intersect at a single point in the measurement volume, which is the centre of the source particle. In practice, due to two main sources of error, the lines of sight never cross perfectly at a single point but approach each other in a restricted volume containing the real particle. Indeed, due to poor calibration or their inability to model the complexity of the optical projection, camera models are never perfect in an experimental context. Discrete correction aims to correct the latter problem. In addition, the detection of the position of particles in the image (by peak fitting) is also prone to errors and can lead to degraded line-of-sight intersections.

Therefore, the implementation of the triangulation procedure used in this study is based on the following principle: a particle is unambiguously located in the measurement volume if one and only one line of sight from each camera image intersects a small three-dimensional box in the measurement volume. Based on this principle, triangulation is implemented as follows: First, a list of adjacent three-dimensional boxes, whose union represents the volume of interest, is created. These boxes represent the cells of a regular grid dividing the domain. They have no methodological interest but are used to speed up the calculation of the triangulation and allow an easier parallelization of the calculation. One of these boxes is shown in the diagram in Figure 2.

Each box is processed independently. For each box, all lines of sight (from the




Figure 2: Illustration of the triangulation method used in this study. The volume is divided into boxes. For each box, and for each camera, the particles whose line of sight intersect the box are identified. If the box contains at least one line of sight for each camera, it is conserved, otherwise, it is eliminated from the list of boxes. When a conserved box contains more than one line of sight for at least one camera, it is divided into eight smaller boxes which are added to the list.

image particles) that intersect that box are identified. Three different configurations are then considered. The first case is where the box is intersected by more than four lines of sight and at least one line of sight from each camera. In this case, it is possible that a particle is inside this box but an ambiguity about the epipolar particles remains (case of the main box in figure 2). In order to remove this ambiguity, the box is divided into 8 sub-boxes and the identification of the intersecting lines of sight is rerun for each box. This procedure of refining the box mesh is repeated until one of two cases appears: either, for a camera image, there is no line of sight intersecting the box (orange dashed boxes in Figure 2), then there is no particle in that box and the box is rejected; or there is a single line of sight from each camera image and there is probably a particle inside. In this case, the particle is located by minimising the average reprojection error across the cameras (eq. 5). If the estimated position is inside the box and corresponds to a quality threshold set by the user (based on the expected triangulation quality), the particle is validated and stored with the box and the epipolar particles in the image (black line boxes with particles in figure 2). Otherwise, if the particle is outside the box (red dotted boxes in figure 2) or if the quality threshold is not reached, the particle and the box are ignored since the particle is respectively located in another box or is probably a ghost particle.

This triangulation process is repeated for as many quadruplets of images as necessary to obtain enough particles to converge the statistics in the next step. A representation of a triangulated list is shown in Figure 3 with associated lines of sight and content boxes.

step 4: The fourth step is the calculation of the discrete correction *per se*. The triangulated particles in physical space are projected onto the four image planes of the camera to determine a reprojection error for each particle. It is calculated as the difference between the position of the detected particle in the image and the projection of the triangulated particle (eq. 5). Similarly, the back-projection error is calculated as the difference between the triangulated particle and the back-projection of the detected image particle onto the Z-plane in which the triangulated

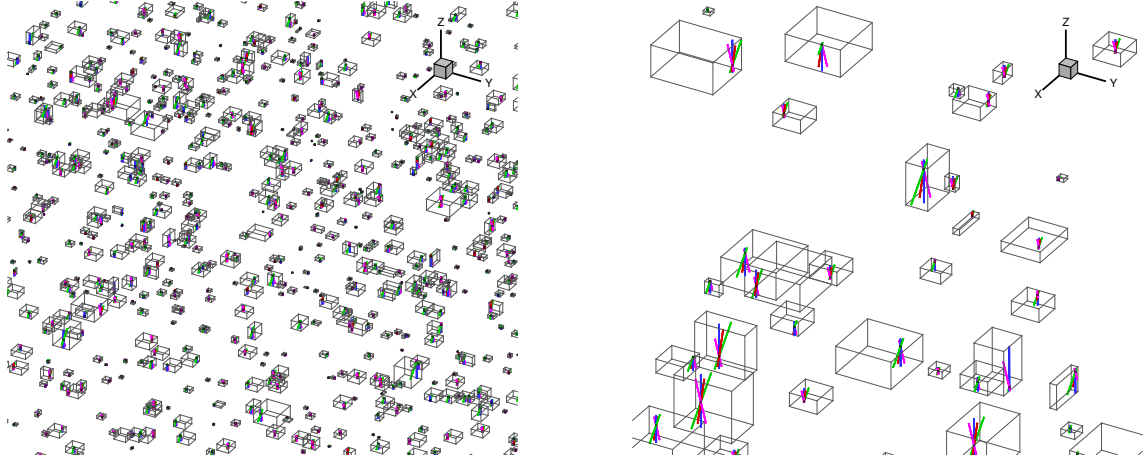


Figure 3: Triangulated particles with their detection box and the corresponding lines of sight

particle is located. The previous steps calculate the reprojection error and the back-projection error for all particles in the measurement volume and thus correct the projection and back-projection operations. As these corrections are distorted by the detection error of the particles in the images, representative corrections can only be determined by a statistical procedure to minimise this error. Therefore, for each selected location, the correction is calculated as a statistically representative value within a three-dimensional box centred on the evaluation point. In this work, the median value is used as the representative value of the correction rather than using the mean value, which allows strongly biased values to be ignored. Other statistical procedures such as clustering could also be implemented but were not investigated in this study. The size of the local domains in which the statistics are computed is chosen as a compromise between the convergence of the statistics favoured by a large number of particles, i.e. large boxes, and boxes small enough to minimise spatial filtering.

The corrections to the projection and back-projection functions are stored in two adaptive grids which are shown to improve the efficiency of the calculations while minimising data storage. They are implemented in such a way as to ensure that the grid is locally refined enough to meet the required accuracy (if possible). The adaptive grid, containing the corrections, is initialised with the size of the volume of interest and a given number of cells. Each cell behaves as an octree structure. The correction is determined at each node using a statistically representative correction calculated in a three-dimensional search area centred on the node. Once the first level is initialized, the value at the centre of each cell is both calculated in the same way as the nodes, and interpolated using the value of the nodes. Trilinear interpolation is implemented in this study but other interpolation schemes could be considered. If the interpolation error (defined as the distance between the interpolated value and the calculated value) is smaller than the target error, the correction is valid and the process stops for the cell. Otherwise, the cell is split in 2 in each direction, giving rise to 8 sub-cells in which the nodes are recalculated and the centre is tested for interpolation accuracy. The process continues until the target accuracy or a maximum cut level is reached. This procedure allows the correction grids to be calculated with a controlled error. The correction grids for projection and back-projection are calculated independently in a similar way.

Step 5: At this stage, the whole model (analytical and discrete parts) has been initialized. This step is a check step to determine whether the camera models have reached a satisfactory accuracy and whether the process should stop or whether another iteration through steps 3 and 4 is required. Often, several iterations are required to achieve appropriate accuracy. Where further iterations are required, the process characteristics can be adjusted. As the corrections are improved with each iteration, the camera models become more and more reliable. As a result, the triangulation error and the proportion of ghost particles decrease in the triangulation process. It can therefore be performed with a higher particle density, which leads to a highly seeded domain, allowing the size of the statistical box domains to be reduced. Indeed, as in the misalignment correction, for the first few iterations the particle density is strongly constrained by the (lack of) accuracy of the model and thus by the ambiguity in the search for matching particles in the images. The more iterative the process, the more the models are adjusted and the more the corresponding particles are recovered. In the last iterations, almost no ghost particles remain and the triangulation process is mainly limited by the quality of the particle detection in the images (which is supposed to be compensated by the statistical analysis of the data).

This calibration procedure with the iterative correction calculation process and the use of the octree data structure allows to calculate camera models with a desired and controlled accuracy. At the beginning of the calibration, the projection and back-projection corrections are set to zero. Thus, the transformations depend solely on the accuracy of the analytical model calculated for the first triangulation. It is therefore very important to take sufficient care in the calibration of the analytical model. Furthermore, the accuracy of the complete model depends strongly on the accuracy of the base model. This will be discussed in more detail in the model evaluation section.

4 Assessment of the camera model

The developed camera model has been evaluated on numerical and experimental set-ups. For this quality assessment, the experimental setup is rather simple as it is based on measurements in a water tank. A numerical dataset is then derived from this setup in order to accurately assess the accuracy of the model and to discriminate between different sources of error.

4.1 Experimental and numerical configuration

The configuration studied is based on an experiment designed for this purpose and is therefore very academic. It consists of a scanning Tomo-PIV measurement inside a water tank. The configuration is shown in figure 4. The water is seeded with polyamide particles of diameter $50\ \mu\text{m}$ and the depth of the measurement volume is fixed at $115\ \text{mm}$. The light source consists of a laser diode whose beam is moved in depth by an oscillating mirror and enlarged by a cylindrical divergent lens. The measuring device consists of four cameras (8-bit JAI CV-M2) arranged in a pyramidal configuration and a light sheet that scans the measurement space. The cameras are mounted on Scheimpflug tilted plates with lenses of $50\ \text{mm}$. The orientations of the cameras are collected in the table 1. In the end, 29 cycles of 250 shots each were recorded, giving a scan step close to $0,5\ \text{mm}$.




Figure 4: Photograph of the experimental configuration for the camera model assessment

Table 1: Orientation of the cameras

	Camera 0	Camera 1	Camera 2	Camera 3
α_x	0	0	15	14
α_y	25	-29	27	-28
α_z	0	0	6	7

On the basis of this experiment, a numerical simulation of the Tomo-PIV experiment was carried out. The characteristics of the simulation were precisely derived from this experiment. The camera position is directly extracted from the experiment and the refraction at the glass wall of thickness 6 mm is taken into account when projecting the particles onto the camera planes, using ray-tracing algorithms. The size of the particles in the images is 3-4 pixel, which corresponds to the filtered experimental images. Gaussian noise was superimposed on the images. Two levels were studied ($\sigma = 0,01$ and $\sigma = 0,05$, for image grey levels between 0 and 1) but due to the rather close behaviour of the model between the two sets, only the higher noise level is presented. Finally, the calibration images of the calibration target were synthesised under the same conditions as the experiment, in order to reproduce the whole experimental calibration. No bias was induced at this stage.

4.2 Scalar errors definition

Analytical and corrected camera models are evaluated on two main criteria: the quality of particle localisation and the quality of triangulation. These two criteria can be quantified by two scalars each representing the quality in the volume and the quality in the image plane. Only the scalars associated with the image plane will be presented because they have the advantage of being at the "pixel" scale and therefore more universal.

The projection error scalar is calculated in the volume as follows. A Cartesian grid representing the measurement volume is defined. Each node, assimilated to a particle, is projected with the "true" model and the model studied on the images. The projection error is defined as the distance between these two projected points

(eq. 4):

$$err_{proj} = \sqrt{(u - u_{ref})^2 + (v - v_{ref})^2} \quad (4)$$

where u , v , u_{ref} and v_{ref} are respectively the horizontal and vertical coordinates of the node projected with the studied model and the real model. Obviously, the calculation of this scalar requires the knowledge of the real projection and can only be used in the context of the simulated experiment.

The second scalar, which is experimentally accessible, is the reprojection error. It characterises the proximity of the lines of sight in the volume. Ideally, the lines of sight from epipolar particles cross perfectly inside the measurement volume, however, in the experimental setup, these lines of sight only pass close to the particle location and never cross perfectly. The calculation of the reprojection error makes it possible to quantify this distance, which is a crucial feature for the tomographic reconstruction. For the calculation, the location of the particle in the volume is triangulated from the four epipolar particles in the images, and then the triangulated particle is reprojected onto the images. The distance between the original particle in the image and the projection of the triangulated particle is defined as the reprojection error (eq. 5).

$$err_{reproj} = \sqrt{(u_{reproj} - u_{ref})^2 + (v_{reproj} - v_{ref})^2} \quad (5)$$

Unlike the projection error, the calculation of the reprojection error is based only on the models studied, once for the triangulation and a second time for the reprojection. It can therefore be used in both experimental and numerical analysis, assuming that the location of the particle in the image is accurately determined during a detection phase. In practice, the location of the particles in the image is subject to a non-negligible error during detection and the reprojection relies on a statistical procedure to extract the most reliable information. Indeed, the modelling errors are constant biases whereas the detection is supposed to introduce random errors allowing statistical filtering. For this study, the reprojection error at a point in the volume is evaluated as the median reprojection error in a small control volume centred on the point of interest. The size of this volume is set to $10 \text{ mm} \times 10 \text{ mm} \times 10 \text{ mm}$ with a grid spacing of $2,5 \text{ mm}$ in all three spatial directions.

4.3 Camera model error database

The calibration of the installation was carried out for the numerical and experimental configurations with a pinhole model and a polynomial model of degree 3 in the x- and y-directions and 2 in the z-direction (Soloff et al. [1997]). This calibration is calculated from the images of the calibration target. As is customary in Tomo-PIV experiments, a misalignment correction procedure is then performed on these camera models. These models form the starting point for the analysis of the discrete correction calculation and its contribution to the overall camera models.

Firstly, the error of these models is quantified using the available error scalars. Secondly, the discrete correction is calculated for both analytical models on the experimental and synthetic images (with the highest noise). Finally, the correction is calculated using intermediate data that is only available in the numerical setup. This investigation, which is of no experimental interest because it cannot be carried out in the experimental configuration, makes it possible to identify the sources of error in the calibration process. In the first case, the discrete correction is calculated based on the "real" position of the particles in the volume, reducing the calibration

of the correction to its essence, without triangulation bias. In the second case, the correction is calculated using the actual position of the particles in the images. It brings the triangulation back into the process but does not take into account the error in the location of the particles in the images due to the detection procedure and thus provides a solution with triangulation on perfect data. The comparison with the previous case allows us to characterise the triangulation bias due to the inaccuracy of the camera models, discussed in the paragraph "Impact of triangulation". Finally, the comparison of the last case with the experimental case provides information on how the statistical procedure used in the correction calibration deals with the data dispersion caused by all the localisation errors, discussed in the paragraph 4.8.

4.4 Experimental and numerical data comparison

The analysis of the camera model errors is mainly based on the synthetic experiment, as the projection and reprojection errors can be calculated and the primary steps of the correction calibration can be broken down. The representativeness of the synthetic data can be ensured by comparing the reprojection error of the analytical model calculated from the numerical and experimental data (see Figures 6a (centre) and 5a (left) for the pinhole model and Figures 7a (centre) and 5b (left) for the polynomial model). Although the error maps do not match perfectly between the experimental and numerical cases, their topology and magnitude are definitely comparable. Therefore, the discussions based on the synthetic data are relevant for the experimental case which is used to ensure the compatibility of the discrete correction model with the real experiments. The similarity between the experimental and numerical cases with respect to the reprojection errors also indicates that the calculation of the error, carried out using a statistical procedure for the experimental case, is relevant, while for the numerical case the error is calculated at the exact point using ground truth knowledge.

4.5 LOS intersection quality

To evaluate the newly developed model, the first criterion that can be used is the quality of the proximity of the lines of sight, called the quality of the intersection of the LOS (intersection being the ideal case). This quality can be quantified directly by the reprojection error and is available for both numerical and experimental data. This error is plotted for the corrected model in Figure 6d (centre), 7d (centre), 5a (centre) and 5b (centre) for the numerical case with the basic pinhole and polynomial model respectively and for the experimental case also with the basic pinhole and polynomial model. These corrected model errors are to be compared with the corresponding error of the basic analytical model, shown in Figures 6a (centre), 7a (center), 5a (left) and 5b (left). For the pinhole-based model, the error, before the application of the discrete correction, is of the order of a few tenths of a pixel to more than one pixel (Figures 6a (centre) and 5a (left)). The error falls to less than 0.005 pixels throughout the volume, with the exception of a marginal band at the edge of the domain for the digital data (Figure 6d (centre)) and to less than 0.05 pixels for the experimental data (Figure 5a (centre)). The discrepancy between the numerical and experimental data is due to the rather low intensity of the particles in the experimental images. For the models based on a polynomial function, the error goes from 0.05 pixels (Figures 7a (centre) and 5b (left)) to 0.005 pixels in the whole domain for the numerical test case (Figure 7d (centre)) and 0.005-0.01 for

the experimental case (Figure 5b (centre)).

The fine scratches are caused at the edges by the local statistical procedure used to determine the correction. First of all, the statistics on the edges are biased by an inhomogeneous distribution in the local volume and the centre of mass of the distribution is no longer the centre of the domain. Moreover, the domain being partially empty, the statistical convergence of the correction is impacted. Secondly, the correction map which is filtered between iterations by outlier detection based on a comparison with neighbouring values may suffer more replacements caused by the reduced number of neighbours. Finally, as the domain boundaries are the areas where the errors are the most important, some particles may be considered as erroneous during the triangulation if the threshold is not well adjusted, and thus the bias described above may be reinforced. Special treatment of boundaries could be implemented if they are important in the case under study.

The reduction of the LOS intersection error, by discrete correction, to a few hundredths of a pixel at most has been demonstrated for all cases studied. Keeping this error at low levels is crucial for Tomo-PIV since the reconstruction algorithms rely on this pattern matching to distribute the intensity of the particles in the volume. It is also of major importance for 3D PTV/LPT as it ensures better particle matching and reduces ambiguities that tend to produce ghost particles (both in PIV and PTV).

4.6 Localization quality

The accuracy of the model is then studied through the projection error. This error characterises the ability of the model to describe the geometric correspondence between the measurement volume and the camera image plane. This analysis can only be carried out with the digital database since knowledge of the ground truth is necessary. The projection error is represented in figures 6a (left) and 7a (left) for the basic analytical model and in figures 6d (left) and 7d (left) for the models corrected for the pinhole and polynomial base models respectively. The initial errors are of the order of one pixel for the pinhole model and 0.1 pixel for the polynomial model (Figures 6a (right) and 7a (right)). Although the reprojection error has been considerably reduced compared to the basic models, the projection error remains of the same order of magnitude despite a modest improvement. This observation is corroborated by the fact that the correction maps presented in Figures 6d (right) and 7d (right) are comparable to the reprojection error maps of the base model in norm (Figures 6a (centre) and 7a (centre)) but not to its projection error (Figures 6a (left) and 7a (left)) which would be the ideal solution. The discrete correction is therefore unable to significantly improve the accuracy of the model. That is, it will significantly improve the quality of the intersection of the lines of sight at a position in the volume that is only slightly improved. In order to understand the inability of the correction to significantly improve the location of the particles, other correction processes relying on intermediate data, only accessible in the numerical data set, have been performed.

4.7 Impact of triangulation

The calibration of the discrete correction relies heavily on triangulated particles for its calculation. In order to assess the ability of the corrected model to accurately describe the reality of the geometric transformation between the volume and the im-

age plane, the correction was calculated with a set of so-called perfectly triangulated particles. This set is composed of the precise coordinates of the particles seeding the measurement volume and their corresponding projection location on the camera image planes. In this case, the bias of image particle detection and particle triangulation is eliminated. Only the statistical procedure of the correction calculation (step 4 of the procedure) with perfect data remains. The projection and reprojection errors are shown in Figure 6b (left) and (centre) for the basic pinhole model and in Figure 7b (left) and (centre) for the polynomial model. For both models, both errors fall to values below 0.005 pixels, demonstrating that, unlike the analytical models considered alone, the model defined as the combination of an analytical model and a discrete correction has the ability to accurately describe the projection (and back-projection) but is hampered by the imprecision of the triangulated data used for its calibration.

In order to isolate the impact of the triangulation procedure, step 3 (triangulation) is reintegrated into the calculation with perfectly localised image particles. In this case, the bias due to the detection of particles in the images is removed and the comparison with the previous scenario allows to discuss the impact of the triangulation itself. The projection and reprojection errors are plotted for this case in Figure 6c (left) and (centre) for the basic pinhole model and Figure 7c (left) and (centre) for the basic polynomial model. In this case, the reprojection error drops drastically (as before) but the projection error remains of the same order of magnitude as the basic model but with a slight improvement. Moreover, the correction calculated in this case (Figures 6c (right) and 7c (right)) strongly resembles the norm reprojection error of the base model (Figure 6a (centre) and 7a (centre)) whereas, in the previous case, the correction (Figures 6b (right) and 7b (right)) was almost similar to the projection error of the base model (Figures 6a (left) and 7a (left)) as expected. Therefore, the triangulation procedure induces a bias responsible for the lack of improvement in the accuracy of the projection. This is due to a lack of accuracy in locating the particles in the volume during the triangulation process. This bias is inherent in the triangulation procedure with inaccurate models since it calculates the most likely location by minimising a distance either in the volume or in the image planes and cannot find the actual location of the particles. This triangulation bias in the procedure leading to a projection bias is therefore not specific to the proposed model but has important consequences for all procedures relying on triangulation such as misalignment correction.

4.8 Impact of particle detection - statistical procedure

After having characterised the impact of the triangulation procedure on the accuracy of the model, the bias due to the detection of particles by the image is studied. This step is isolated by comparing the results of the previous scenario (triangulation with perfectly detected particles) and the complete procedure. No significant divergence is observed between the two cases, neither on the projection error (Figures 6c (left) and 6d (left) for the basic pinhole model and Figures 7c (left) and 7d (left) for the basic polynomial model), nor on the reprojection error (Figures 6c (centre) and 6d (centre) for the basic pinhole model and Figures 7c (centre) and 7d (centre) for the basic polynomial model). Due to inaccurate detection, the reprojection error is slightly higher in the experimental type data set, but is effectively mitigated by the statistical procedure of the correction calculation. The detection error mainly affects the statistical convergence of the correction. The more accurate the detection, the

more accurate the triangulation and the fewer particles will be needed to obtain an accurate correction calculation. A better detection quality will therefore result in a reduced computational load (if the box size is kept constant) or in a more accurate correction (if the number of particles is kept constant and the box size is reduced).




Figure 5: Maps of the reprojection error norm of the base model (left column), reprojection error norm of the corrected model (central column) and computed discrete correction norm (right column) for a pinhole base model (first row) and a polynomial base model (second row)

5 Contribution of the model in a Tomo-PIV experiment

Finally, the proposed model has been successfully used in an experimental study briefly presented here to highlight the advantages of such a camera model. The configuration chosen to carry out this study is that of an open-ended wing, placed at an incidence of 30^{circ} in the TunHyde hydrodynamic tunnel at the Pprime Institute. With a square cross-section of 230 mm on each side, the test section of the hydrodynamic tunnel is composed of a stainless steel structure and large 30 mm thick PMMA (polymethylmethacrylate) windows, allowing wide optical access over a length of 800 mm. It can be used in open configuration with flow velocities up to 2 m/s and in closed configuration (under load) up to 8 m/s. For all studies, the flow velocity is set to 1.25 m/s, the chord-based Reynolds number to 10^5 and the




Figure 6: Maps of the projection error norm (left column), reprojection error norm (central column) and computed discrete correction norm (right column) for a pinhole base model

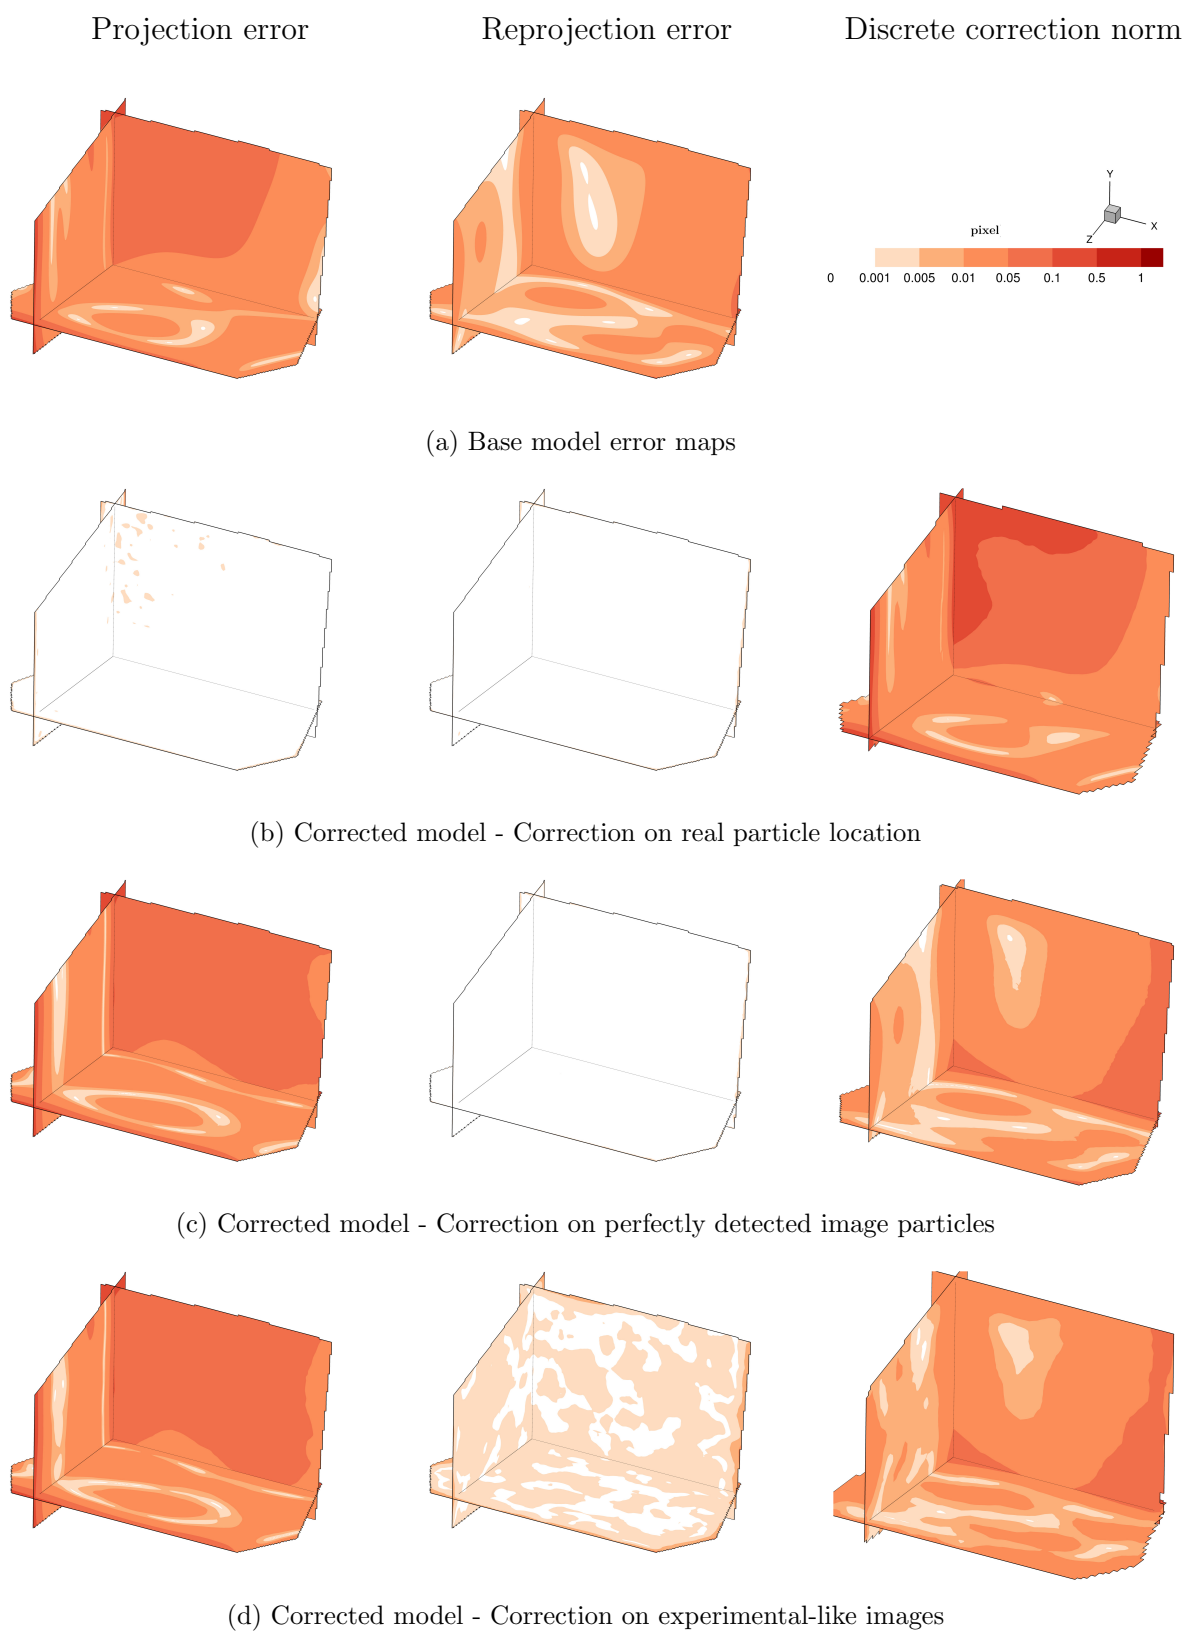


Figure 7: Maps of the projection error norm (left column), reprojection error norm (central column) and computed discrete correction norm (right column) for a polynomial base model

blockage rate to 11 %. In the study configuration thus defined, illustrated by the sketch in Figure 8, Tomo-PIV velocity measurements were to be performed. The overall study frame, shown in Figure 3.3, is chosen to form an orthonormal reference frame, with the X axis oriented along the flow axis, the Z axis oriented along the increasing span of the wing and the Y axis coinciding with the vertical axis oriented upwards. The origin of the reference frame is positioned at the quarter chord of the wing profile, at the root (coincident with the side window). To allow optical measurements, the flow is seeded with polyamide particles of size $20\ \mu\text{m}$ (Vestosint 2158). These particles have the advantage of scattering light well and, above all, of having a density close to that of water ($\rho = 1016\ \text{kg} \cdot \text{m}^{-3}$), which favours their movement with the flow and their suspension.




Figure 8: Sketch of the experimental configuration: camera positions and hydraulic tunnel. Cameras 3, 4, 5 and 6 were used for tomoPIV measurements.

For these three-dimensional measurements, the velocity should be measured primarily on the upper surface and in the wake of the wing, as this is where the vortex structures appear most. For these reasons, it was decided to modify the classical (optimal) configuration of the square-based pyramid distribution of cameras by moving one of the cameras to the centre, in the axis of the profile. This configuration allows for at least three cameras to be useful for wake reconstruction. The cameras used in the experiment are two Photron APX-RS high-speed cameras at positions 3 and 4 and two Photron Fastcam SA-Z at positions 5 and 6. All the cameras used have a square sensor with a resolution of $1024\ \text{pixel} \times 1024\ \text{pixel}$. Cameras 1 and 2 were used for other measurements. The two APX-RS cameras are mounted with a Sigma lens 105 mm, and the SA-Z cameras with a Nikon lens 105 mm. Scheimpflug mounts are used on the three tilted cameras to place the focal plane on the median plane of the measurement volume. For space reasons, the front camera (number 6) is positioned slightly backwards, and a focal length multiplier ($\times 1.4$) is added to the lens to maintain an optimal measurement field. The angular positioning of the cameras is shown in the table 2. The measurement volume is illuminated by two fast lasers: a TerraPIV Continuum, from below and a MesaPIV from above to ensure sufficient light intensity and to limit profile shadowing in the measurement field. The laser sheet are generated using the Lavision volume generator, whose assembly of converging and diverging cylindrical lenses creates a thick sheet collimated on the Z axis and diverging in the (X, Y) plane. The edges of the sheets are also

Table 2: Camera angles used for the tomoPIV measurements

	Camera 3	Camera 4	Camera 5	Camera 6
$\alpha_x/^\circ$	-18	18	-18	0
$\alpha_y/^\circ$	-16	16	16	0

screened to ensure that they are sharp and avoid the presence of faintly illuminated particles, which are detrimental to the reconstruction. The illuminated polyamide tracers reflect the laser light and become visible to the cameras during the laser shot. The particle images are recorded simultaneously by the four cameras of the system at a frequency of 2 kHz and then filtered to limit the impact of measurement noise on the reconstruction. The filtering operation consists of several successive steps. First, the average of the time series is subtracted from each image in order to eliminate reflections as much as possible. Then, a local normalisation is performed to homogenise the intensity in the images and limit the impact of shadows created by the wing. Next, Gaussian filtering is applied, followed by (weak) thresholding to remove low-level noise.

The system is calibrated by recording images from a calibration chart of points by the four cameras. The test pattern is positioned in the volume according to external references and shifted along the Z axis to 13 positions in the range $[-15\text{ mm}; 15\text{ mm}]$ using a micrometric shift table. During positioning, particular attention is paid to checking the alignment (tangency) between the laser sheet and the calibration chart. The camera model chosen for the study is a polynomial model with discrete corrections. First of all, in order to correct the displacements of the optical system that occur systematically between the calibration stage and the recording of the series of images, a misalignment correction stage is carried out on each set of measurements. Indeed, a misalignment (of small amplitude) was observed between acquisitions, probably due to interventions on the test section and around the camera between each configuration (change in the thickness of the sheet, change in the profile of the wing in the section, calibration of the CMOS sensor blacks). However, no significant variations were observed during the time series themselves: the application of a calibration correction between each acquisition is therefore sufficient. There is no vibration. This correction of model misalignment (self-calibration) is performed using the procedure described in section 3 applied without discrete correction. The reprojection error is presented in figures 9a and 9b before and after the alignment correction. The error is presented for the horizontal component dU of the first camera (numbered 3). The vertical component and the errors of the other cameras related to the calibration errors are of the same order of magnitude as the one presented, or even lower, or almost non-existent for some of them. The overall movement is clearly visible with the polynomial modelling since it results in a (quasi)-constant error (about 0,7 pixel) superimposed on the intrinsic error of the model, which is very small in the centre of the domain (about 0,1 pixel), but can be large towards the top and bottom (about 0,25 pixel). In order to improve the models, a discrete volume correction is added to them, according to the procedure described in section 3. The final size of the camera model on disk was about 10 Mb per camera. It should be noticed that the size on disk of the camera model depends on the complexity of the experimental setup. The number of iterations used to estimate this correction was 6. The residual errors after correction are shown in Figure 9c. The discrete correction drastically reduces the reprojection errors on the edges, but also inside the domain. The global error is about 0,01 pixel everywhere. The

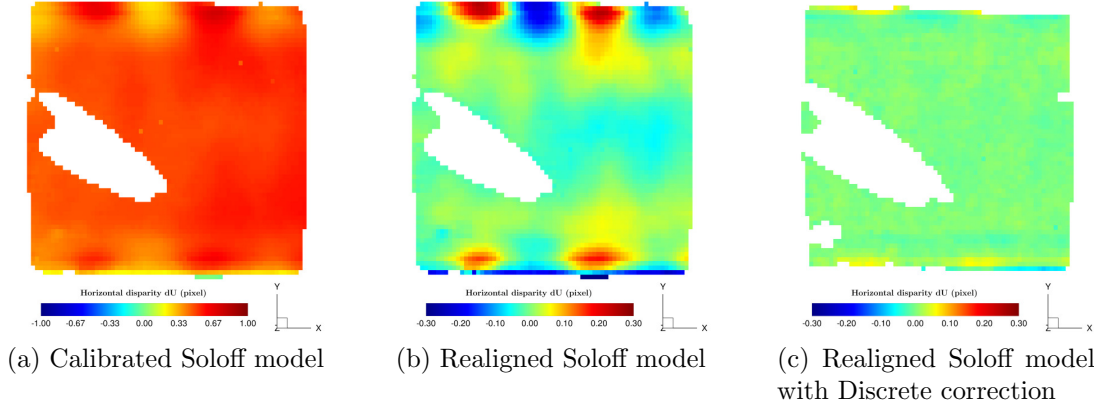


Figure 9: Maps of the horizontal component of the reprojection error for Camera 3, in the middle volume plane. The camera model is a polynomial (Soloff) model (degrees (3,3,2) in (X,Y,Z) , with a discrete correction. (a) Model calibrated without correction (b) Model realigned without discrete correction (c) Model realigned with discrete correction

particular shape of these errors corresponds to a regular pattern whose wavelength corresponds to the distance between the clamping screws of the viewing window. These errors cannot be corrected by simply increasing the degree of the polynomial, given their localised nature.

6 Discussion and Conclusion

The discrete camera model presented in this paper is a very versatile model. However, the quality of the model depends on different conditions:

- with enough images of (detectable) particles, it is possible to reduce the reprojection error with no real limit, even if it is more difficult on the edges of the domain.
- the projection error is not reduced using a discrete correction: this depends mostly on the analytical model.
- the analytical model can be obtained using a standard calibration with a target (like in this article) but it could also be obtained through a raytracing model. In this case, it could reduce the computational cost.
- the discrete correction need a good initialization in order to be computable. This initialization can be provided by the analytical model, or by a discrete model obtained by the calibration procedure.
- the discrete correction can contain the analytical model, but analytical models are useful to interpolate or extrapolate the camera model.
- it is necessary to use a discrete correction in complex optical configurations as no analytical model can take into account optical imperfections.
- in situations where no calibration target can be used, the only way is to calibrate ex-situ and to find an analytical model to reduce the calibration error. In this case, the localization errors should be checked carefully.

In conclusion, this paper describes a new technique for improving the accuracy of camera models. It is based on the combination of an analytical camera model, which can be based on a physical model, and a discrete correction. The analytical model and its correction can be calibrated in a conventional way using a test pattern, and improved to correct for projection errors. The discrete model requires a large amount of data, especially if an accurate model is required. This is the price to pay for minimising the errors associated with the camera model (reprojection errors). This large amount of data can be obtained by using many test pattern positions, but also particle images. A new triangulation procedure has been used in this paper, based on a search in physical space and not starting from the images as is more traditionally done in the literature. It allows for simple parallelization, and makes it possible to work from the lists of particles, without using the images, to match the particles.

The camera calibration method has been validated on synthetic images, on a dedicated experiment (compared to scanning) and illustrated on a real experiment. Accuracy can be achieved to an arbitrary level and an accuracy of $SI0.01pixel$ can be achieved with sufficient data. The method can be used to correct optical defects in windows of any shape. If the distortions are too large, the procedure will have to be adapted, for example by reducing the number of particles in each image to improve the match. Similarly, an appropriate analytical model will have to be chosen to reduce localisation errors. Another important advantage of the method is the possibility to describe discontinuous functions with a discrete correction. It is possible to handle the presence of optical interfaces in the measurement volume, provided that one is able to match the particles, which should be possible for fixed interfaces, but certainly more complicated for moving ones. This new camera model should be tested in more complex configurations, with curved interfaces, or in fluid-structure interaction configurations with transparent solids (with or without particles inside).

7 Acknowledgments

The grant of Gwenaël Acher's PhD thesis is supported by the Direction Générale de l'Armement. The current work has been partially funded by the CPER FEDER of the Region Nouvelle Aquitaine.

References

- J. Belden. Calibration of multi-camera systems with refractive interfaces. *Experiments in Fluids*, 54(2):1463, Feb 2013. ISSN 1432-1114. doi: 10.1007/s00348-013-1463-0. URL <https://doi.org/10.1007/s00348-013-1463-0>.
- D. Calluau. *Développement d'une méthode de mesures tridimensionnelles par PIV Stéréoscopique: application à l'étude de l'écoulement naissant et établi autour d'un parallélépipède*. PhD thesis, Poitiers, 2003.
- T. Earl, Y. J. Jeon, B. Lecordier, and L. David. F2dpr: a fast and robust cross-correlation technique for volumetric piv. *Measurement Science and Technology*, 27(8):084007, 2016. URL <http://stacks.iop.org/0957-0233/27/i=8/a=084007>.
- Y. Kwon and J. B. Casebolt. Effects of light refraction on the accuracy of camera calibration and reconstruction in underwater motion analysis. *Sports Biome-*

- chanics*, 5(2):315–340, 2006. doi: 10.1080/14763140608522881. URL <https://doi.org/10.1080/14763140608522881>. PMID: 16939159.
- B. Lecordier, C. Gobin, C. Lacour, A. Cessou, B. Tremblais, L. Thomas, and L. David. Tomographic piv study of lifted flames in turbulent axisymmetric jets of methane. *16th Int Symp on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal*, July 2012. URL http://ltces.dem.ist.utl.pt/LXLASER/lxlaser2012/upload/377_abstract_ywqbvr.pdf.
- K. Muller, C. K. Hemelrijk, J. Westerweel, and D. S. W. Tam. Calibration of multiple cameras for large-scale experiments using a freely moving calibration target. *Experiments in Fluids*, 61(1), 2020. ISSN 0723-4864. doi: 10.1007/s00348-019-2833-z.
- G. Paolillo and T. Astarita. On the PIV/PTV uncertainty related to calibration of camera systems with refractive surfaces. *Measurement Science and Technology*, 32(9):094006, jun 2021. doi: 10.1088/1361-6501/abf3fc. URL <https://doi.org/10.1088/1361-6501/abf3fc>.
- G. Paolillo and T. Astarita. Perspective camera model with refraction correction for optical velocimetry measurements in complex geometries. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3185–3196, June 2022. ISSN 1939-3539. doi: 10.1109/TPAMI.2020.3046467.
- F. Scarano. Tomographic piv: principles and practice. *Measurement Science and Technology*, 24(1):012001, 2012.
- F. Scarano, L. David, M. Bsibsi, and D. Calluaud. S-piv comparative assessment: image dewarping+misalignment correction and pinhole+geometric back projection. *Experiments in Fluids*, 39(2):257–266, 2005. ISSN 0723-4864. doi: 10.1007/s00348-005-1000-x.
- S. M. Soloff, R. J. Adrian, and Z.-C. Liu. Distortion compensation for generalized stereoscopic particle image velocimetry. *Measurement Science and Technology*, 8(12):1441, 1997. URL <http://stacks.iop.org/0957-0233/8/i=12/a=008>.
- L. Thomas, B. Tremblais, and L. David. Optimization of the volume reconstruction for classical tomo-piv algorithms (mart, bimart and smart): synthetic and experimental studies. *Measurement Science and Technology*, 25(3):035303, 2014. URL <http://stacks.iop.org/0957-0233/25/i=3/a=035303>.
- J. Weng, P. Cohen, and M. Herniou. Camera calibration with distortion models and accuracy evaluation. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 14:965–980, 10 1992. ISSN 0162-8828. doi: 10.1109/34.159901. URL doi.ieeecomputersociety.org/10.1109/34.159901.
- B. Wieneke. Volume self-calibration for 3d particle image velocimetry. *Experiments in Fluids*, 45(4):549–556, Oct 2008. ISSN 1432-1114. doi: 10.1007/s00348-008-0521-5. URL <https://doi.org/10.1007/s00348-008-0521-5>.
- B. Wieneke. Improvements for volume self-calibration. *Measurement Science and Technology*, 29(8):084002, jul 2018. doi: 10.1088/1361-6501/aacd45. URL <https://doi.org/10.1088/1361-6501/aacd45>.

Z. Zhang. Flexible camera calibration by viewing a plane from unknown orientations. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 1, pages 666–673 vol.1, Sept 1999. doi: 10.1109/ICCV.1999.791289.