Exp Fluids (2016) 57:13
DOI 10.1007/s00348-015-2100-x

RESEARCH ARTICLE

A direct approach for instantaneous 3D density field reconstruction from background-oriented schlieren (BOS) measurements

F. Nicolas[1] · V. Todoroff[1] · A. Plyer[2] · G. Le Besnerais[2] · D. Donjat[1] · F. Micheli[1] · F. Champagnat[2] · P. Cornic[2] · Y. Le Sant[3]

Received: 16 September 2015 / Revised: 25 November 2015 / Accepted: 3 December 2015

Abstract We present a new numerical method for reconstruction of instantaneous density volume from 3D background-oriented schlieren (3DBOS) measurements, with a validation on a dedicated flexible experimental BOS bench. In contrast to previous works, we use a direct formulation where density is estimated from measured deviation fields without the intermediate step of density gradient reconstruction. Regularization techniques are implemented to deal with the ill-posed problem encountered. The resulting high-dimensional optimization is conducted by conjugate gradient techniques. A parallel algorithm, implemented on graphics processing unit, helps to speed up the calculation. The resulting software is validated on synthetic BOS images of a 3D density field issued from a numerical simulation. Then, we describe a dedicated 3DBOS experimental facility which has been built to study various BOS settings and to assess the performance of the proposed numerical reconstruction process. Results on various datasets illustrate the potential of the method for flow characterization and measurement in real-world conditions.

✉ D. Donjat
david.donjat@onera.fr

[1] DMAE, ONERA, 2 Avenue Edouard Belin, 31400 Toulouse, France

[2] DTIM, ONERA, Chemin de la Hunière, 91123 Palaiseau, France

[3] DAFE, ONERA, 8 rue des Vertugadins, 92190 Meudon, France

1 Introduction

1.1 BOS and 3DBOS

Density field visualization such as strioscopy has been used extensively to understand fluid mechanics phenomena such as compressibility or thermal effects. Yet, the mostly qualitative information issued from this measurement does not give enough information to refine computational fluid dynamics codes. Background-oriented schlieren (BOS) is one way to obtain quantitative measurement of density gradients. Density fluctuations in a fluid come with inhomogeneities of the optical index. The BOS technique is based on observation of ray deviations through a medium of inhomogeneous optical index. The BOS optical setting is remarkably simple: it only requires that the flow under study is placed in between a camera and some textured background on which the camera is focused. Comparing the images of the background with or without the flow using digital image correlation techniques reveals displacements which are the projections of the light ray deviations. First BOS methods have been proposed in the early 2000 years (Dalziel et al. 2000; Raffel et al. 2000; Meier 2002), a recent review of BOS techniques can be found in Raffel (2015).

We focus on 3DBOS, i.e., the numerical reconstruction of 3D density fields from a set of deviation fields measured by BOS. Please note that, here, 3DBOS means quantitative estimation of the density field volume, usually by providing density estimates on each vertex of some 3D voxel grid recovering the flow. This is not to be confused with the use of multiple BOS measurements for *3D localization* of structures such as vortices within a large experimental volume (Bauknecht et al. 2015).

Published online: 26 December 2015

1.2 Related works

When considering an axisymmetric flow, the 3D density field can be computed from one BOS measurement along a viewing direction orthogonal to the flow axis by inversion of an Abel integral equation. This technique has been applied on the study of supersonic flow around cones (Venkatakrishnan and Meier 2004; Sourgen et al. 2004, 2012), on jet experiments (Venkatakrishnan 2005; Venkatakrishnan and Suriyanarayanan 2009) and also for the study of rotor blades vortices (Kindler et al. 2007). All other instances of 3DBOS require the tomographic processing of several BOS deviation fields obtained along various viewing directions—they will be called "BOS projections". These projections can be recorded at several time instants by rotating a unique BOS setting (camera + background) with respect to the flow. Such approaches usually rely on the assumption of a stationary flow (Ota et al. 2010, 2011; Sourgen et al. 2012), although they have been used recently for periodic unstationary flows by synchronizing the measurement with the flow period (Cabaleiro 2013; Christian et al. 2014). A relatively large number of BOS projections can then be obtained. For instance, Ota et al. (2010) and Sourgen et al. (2012) collect 19 BOS projections every 5° along a quarter of a circle around the jet.

Synchronized recording of a flow by several BOS settings opens the way to 3D reconstruction of nonaxisymmetric unstationary 3D flows. There are few references in this direction. The main piece of work originates from Ihrke and Magnor (2004), Ihrke (2007), Atcheson et al. (2008) and was motivated mainly by the problem of realistic rendering of transparent and turbulent media in computer graphics. In this context, Atcheson et al. (2008) published the first experimental demonstration of 3DBOS using 12 cameras mounted in a half-circle configuration around a hot air flow produced by a gas burner. Zeb et al. (2011) studied heat convection with a three-camera system, leading to a reconstruction with a rather low resolution and large artifacts. Alhaj and Seume demonstrate a coplanar 8-cameras system within a linear cascade wind tunnel (Alhaj and Seume 2010). The present paper, which follows several conference communications made since 2012 (Todoroff et al. 2012, 2014; Le Sant et al. 2014), presents a quantitative 3DBOS method on coplanar and noncoplanar camera settings.

Let us briefly review the numerical approaches which have been put forward to reconstruct 3D density volume from a set of BOS measured deviation fields. As will be described more precisely in Sect. 2, deviations result from the integration of density gradients of the volume along each ray going from the background pattern to a pixel of the BOS camera. Most of the previous references address the reconstruction problem in two steps (Atcheson et al. 2008; Ota et al. 2010, 2011; Sourgen et al. 2012; Leopold et al. 2013): tomographic reconstruction [or computed tomography, CT (Kak and Slaney 2001)] of density gradients, then integration of the gradients to estimate the density volume. The classical approach for the first step of CT is Filtered Backprojection (FBP) using Fast Fourier Transform, which is fast and efficient if a large number of projections is available. FBP has been used in 3DBOS studies of stationary flows where several views can be done sequentially (Goldhahn and Seume 2007; Sourgen et al. 2012) with the same camera or even in axisymmetric cases (Venkatakrishnan and Suriyanarayanan 2009; Schröder 2009). Yet, FBP becomes impractical when dealing with limited number of projections, or with occultations (Kak and Slaney 2001). To tackle these problems, most recent studies Atcheson et al. (2008), Ota et al. (2010, 2011), and Leopold et al. (2013) propose methods based on the iterative resolution of a discretized system derived from the modelization of the ray propagation (see Sect. 2)—such methods are known as "algebraic reconstruction techniques" (ART) in CT.

The integration step is usually done after tomographic reconstruction of the three components of density gradients by solving a Poisson equation (Atcheson et al. 2008; Ota et al. 2011), whereas other authors propose to simply integrate numerically along each direction of space (Leopold et al. 2013).

Finally, let us mention references which rely on a coplanar camera configuration and a parallel ray hypothesis so as to recast the 3D reconstruction problem as a stack of 2D tomographic reconstructions, which are solved by a 2D FBP (Goldhahn and Seume 2007; Alhaj and Seume 2010) or ART (Zhang et al. 2015). These approaches use cameras with very limited field of view. In order to obtain independent 2D tomographic problems stack by stack, they systematically neglect the out-of-plane component of the deviation. Here we will consider general camera configurations (larger field of view, noncoplanar configurations) and handle 3D deviations.

In contrast to previous works, we have proposed in Todoroff et al. (2012, 2014) a one-step approach applicable to general configurations, where the 3D density field is directly estimated from BOS deviation fields. A discretized volume around the flow is defined and related to measured deviations through a global linear system under a paraxial assumption. Volume reconstruction then reduces to numerical inversion of this system. However, as well known in CT, the resulting linear system is not only underdetermined but also badly conditioned. Trying to solve it with usual least-squares approaches leads to artifacts related to the missing projections and to amplification of the measurement noise. Regularization (Tikhonov and Arsenin 1977; Idier 2010) is a way to define a solution with better properties while fitting the data. There are very few papers which explicitly

address the problem of regularizing the BOS reconstruction. In their tomographic reconstruction of density gradient components, Ihrke and Magnor (2004) and Atcheson et al. (2008) solve the system by a Conjugate Gradient (CG) algorithm with a support constraint on the sought volume to limit artifacts. Besides, they stop the algorithm before convergence to benefit from the good behavior of CG, which restores large-scale structures before small ones and thus regularizes the solution (Ihrke and Magnor 2004). Actually, most ART approaches of CT rely on such early stopping criterion to avoid noise amplification.

Here we propose to adopt a classical regularization framework (Idier 2010) where the estimated density volume is defined as the minimizer of a compound criterion made up of a data fidelity term and a regularization term. As we search directly for the density volume rather than for an auxiliary quantity such as the gradient, we can use the regularization term to enforce desirable properties of the density. We present results obtained with a classical Tikhonov smoothness regularization (Tikhonov and Arsenin 1977). The main tuning parameter is the regularization parameter which balances both terms of the criterion. We propose to select the regularizing parameter by a L-curve approach (Hansen 1992).

Another important aspect of 3DBOS is the computational burden associated with the reconstruction of large volumes. The problem at hand is the large-scale iterative optimization of a convex, but possibly nonquadratic, criterion. As in Ihrke and Magnor (2004), we use a CG algorithm and restrict the number of variables as much as possible by selecting a tight mask (or support) for our reconstruction. In the CG algorithm, the most intensive part is to project and backproject data between the image planes and the discretized volume. Those two operations are highly parallel. Following several works in CT (Gross et al. 2009; Pan et al. 2010), we propose here to implement the optimization on graphics processing units (GPU), which are massively parallel architectures providing high computational resource for a low cost.

Finally, the previous literature can also be analyzed in terms of the test cases used for validation. We focus here on jet flow experiments and propose a validation on both simulated and real data. As one is forced to work with a very limited number of projections in 3DBOS, an important aspect of the validation is the choice of the camera configuration around the jet. Most references limit their study to data acquired with a planar configuration, where all camera centers belong to a unique plan, usually orthogonal to the main axis of the flow under study. In contrast, we provide results on noncoplanar configurations, not only from simulation but also from real data provided by a dedicated "geode-shaped" 3DBOS test bench built at ONERA DMAE (Toulouse, France). Indeed, nonplanar configurations may lead to better results than planar ones for flows with arbitrary geometry or to restore details which are not aligned with the main axis of the jet. Moreover, these configurations are sometimes required when working in experimental facilities where optical accesses around the jet are constrained. Their study is then of primary interest to prepare the deployment of BOS in wind tunnels, which is our main concern for the future.

1.3 Contributions and outline of the paper

The main contribution of the paper is a numerical framework for a one-step reconstruction 3D density volume from a limited number of deviation fields. This framework is described in Sect. 2 and validated on simulated data in Sect. 3. The second major contribution is an original experimental bench for the empirical study of noncoplanar configurations for 3DBOS, described in Sect. 4. Section 5 describes the processing chain going from recorded images to deviation fields and also gives some details on geometrical calibration of the cameras. In Sect. 6, we present reconstructed density volumes from various datasets and discuss the results. Finally, concluding remarks and perspectives are given in Sect. 7.

2 A regularized framework for direct 3DBOS reconstruction

The BOS technique is based on observation of ray deviations through a medium of inhomogeneous optical index n. The deviation angle of a ray is equal to the integral of the optical index gradient along the optical path. Optical index is then related to density via $n - 1 = G\rho$, where G is the Gladstone–Dale constant which is a function of the light wavelength and of the chemical composition of the flow. Globally, considering one ray going through the flow under study, the deviation ε writes

$$\varepsilon = \frac{G}{n_0} \int_{\text{ray}\subset\text{flow}} \nabla\rho(s)\,\mathrm{d}s, \tag{1}$$

where the integration is done on the part of the light ray which is inside the flow and assuming that the optical index of the medium surrounding the flow is constant and equal to n_0. This problem is generally nonlinear as the integration path depends on the unknown and variable density of the flow. Similar to most other references (except for Ihrke 2007), we assume the paraxial hypothesis and integration is done along the (straight) unperturbed ray.

The linear problem associated with the BOS reconstruction can be formulated as follows. Let us suppose that K projections, each one made up of $I \times J$ deviation values, have been measured around a given volume. The problem then writes as a set of three coupled linear systems:

$$\varepsilon_u(i,j,k) = \frac{G}{n_0} \int_{\text{ray}(i,j,k)} \frac{\partial \rho}{\partial u} \mathrm{d}s \quad \begin{cases} 1 \leq i \leq I, \\ 1 \leq j \leq J, \\ 1 \leq k \leq K \end{cases} \quad \text{with } u \in \{x, y, z\}. \tag{2}$$

As we use a paraxial approximation, the deviations $\varepsilon_u(i,j,k)$, $u \in \{x,y,z\}$ result from integrations of the corresponding density gradient along the straight, unperturbed ray which crosses the k^{th} camera at pixel (i, j). In practice, one first chooses the world frame coordinate system, which defines the (x, y, z) directions in (2), and the geometry of the reconstructed volume. Then, the cameras are calibrated in the world reference frame, which defines the 3D equation of each line segment associated with the observed deviations. This process is called the geometrical calibration of the experiment. It is discussed in Sect. 4.3.

2.1 Discretization

The goal of the discretization process is to derive a matrix approximation of the system (2) such as:

$$\varepsilon = \begin{bmatrix} \varepsilon_x \\ \varepsilon_y \\ \varepsilon_z \end{bmatrix} = \boldsymbol{A}\rho = \boldsymbol{T} \begin{bmatrix} \boldsymbol{D}_x \\ \boldsymbol{D}_y \\ \boldsymbol{D}_z \end{bmatrix} \rho \tag{3}$$

where, for $u \in \{x,y,z\}$, $\boldsymbol{D}_u$ is a finite difference matrix along the u axis and ε_u collects all measured deviations along the u axis. $\boldsymbol{T}$ is a tomographic projection matrix: each line is made up of positive weights which describe the contribution of each volume element to a particular deviation. The structure of $\boldsymbol{T}$ depends on the choice of the basis functions which are used to discretize the volume, i.e., to approximate it by a finite set of values. As usual in algebraic approaches of CT, we choose a piecewise constant discretization of the volume so that $\boldsymbol{T}$ is sparse, i.e., there are very few nonzero coefficients in each row of $\boldsymbol{T}$. Knowing the geometrical calibration of the experiment, the computation of $\boldsymbol{T}$ is straightforward. However, because this matrix is huge, it is not stored during the estimation but computed on the fly.

2.2 Regularization

Measurement of deviations is subject to errors and the observation system should then be written as

$$\varepsilon = \boldsymbol{A}\rho + \boldsymbol{b}. \tag{4}$$

where $\boldsymbol{b}$ accounts for measurement noise. There are several issues with Eq. (4). First, $\boldsymbol{A}$ is neither square nor regular, so inversion should be done in the least-squares sense, i.e., by searching for a minimizer of $\|\varepsilon - \boldsymbol{A}\rho\|^2$— or some other weighted least-squares criterion if information is known about variable variance of the measurement noise. Second, as we aim to measure instantaneous flows, the noise cannot be diminished by averaging several images. Finally, $\boldsymbol{A}$ inherits the ill-conditioning of the tomographic matrix $\boldsymbol{T}$. It means that the effect of the noise will be amplified in the least-squares solution. Fidelity to the data is not sufficient to obtain a reliable solution, and some regularization process should be applied (Idier 2010). One way to do so is to define the solution as the minimizer of a compound criterion:

$$\mathscr{J}(\rho) = \|\boldsymbol{A}\rho - \varepsilon\|^2 + \lambda \mathscr{R}(\rho), \tag{5}$$

where the first term is the classical least-squares criterion and the second one is the regularization term, which encodes the a priori information on the volumic distribution of ρ. The balance between both terms is set by the regularization parameter $\lambda > 0$. In this work, we consider first-order Tikhonov regularization (Tikhonov and Arsenin 1977) implemented in our discretized setting by choosing the L2 norm of the density spatial gradient as the regularization term:

$$\mathscr{R}(\rho) = \|\bar{D}\rho\|^2 = -\rho^T \triangle \rho, \tag{6}$$

where $\bar{D}$ is the upwinded discrete gradient operator and $\triangle$ is the discrete Laplacian operator with appropriate boundary conditions (see Sect. 2.3.2).

This leads to a well-behaved quadratic criterion. It enforces smoothness of the solution and reduces the effect of the noise but tends to oversmooth the density discontinuities that may be present in the flow. Note that we have also considered Total Variation (TV) regularization, aimed at preserving discontinuities or edges in the reconstructed volume (Todoroff et al. 2012, 2014). In the case of convective flow with relatively small density gradients, a TV approach does not appear necessary to recover a good solution.

One issue with the use of a compound criterion such as (5) is to determine the optimal value of the regularization parameter. In the case of quadratic regularization, the L-curve has been proposed to obtain the optimal regularization parameter (Hansen 1992; Idier 2010). The L-curve is obtained by plotting the gradient norm (6) as a function of the data term $\|A\rho - \varepsilon\|^2$ for several density estimates obtained with various regularization parameters λ. It can be shown that this curve is usually L-shaped (hence its name) and that a good choice for the parameter is the one corresponding to the point of highest curvature. Indeed, at this point, both terms of the criterion (5) are close to their minimum. Note that the L-curve has been used by Ihrke and Magnor in tomographic reconstruction of flames (Ihrke and Magnor 2004) (and later by Atcheson et al. 2008 for 3DBOS) to determine a stopping condition in the CG optimization of a least-squares criterion.

2.3 Optimization

Let us first give some typical dimensions of the problem at hand: we consider at least 10 image acquisitions with typically 5×10^5 displacement vectors per image, leading to a total number of collected deviation of the order of 10^6. The size of the reconstructed volume is of the order of 10^7 voxels. Our long-term goal is to be able to reconstruct gigavoxel volumes. In all cases, the optimization problem is of very large dimension, which means that only first-order descent methods such as gradient descent and conjugate gradients (CG) are affordable. In this section, we present the chosen CG algorithm, constraints in force during the optimization and some details about GPU implementation.

2.3.1 Conjugate gradients algorithm

Conjugate gradients are a classical iterative strategy for solving large linear systems (Wright and Nocedal 1999). They have been used in 3DBOS by Ihrke and Magnor (2004), and Atcheson et al. (2008) for solving (4) in the least-squares sense, and it is also a classical algorithm in algebraic approaches to CT (Gross et al. 2009). We use it here to optimize criterion $\mathscr{J}$ of Eq. (5), with the quadratic regularization term of Eq. (6).

The GC algorithm is an iterative descent technique, i.e., the estimated volume is updated at iteration k according to:

$$\rho_k = \rho_{k-1} + \alpha_k \boldsymbol{d}_k \tag{7}$$

where $\boldsymbol{d}_k$ is the descent direction which is defined recursively by:

$$\boldsymbol{d}_k = -\boldsymbol{g}_k + \beta_k \boldsymbol{d}_{k-1}, \quad \text{with } \boldsymbol{d}_0 = -\boldsymbol{g}_0, \tag{8}$$

where

$$\boldsymbol{g}_k = \nabla \mathscr{J}(\rho_k) = 2\left(\boldsymbol{A}^T(\boldsymbol{A}\rho_k - \varepsilon) - \lambda \triangle \rho_k\right) \tag{9}$$

β_k is selected in order to achieve a sequence of orthogonal descent directions (Wright and Nocedal 1999):

$$\beta_k = \frac{\|\boldsymbol{g}_{k+1}\|^2}{\|\boldsymbol{g}_k\|^2}. \tag{10}$$

And α_k achieves steepest descent in the direction $\boldsymbol{d}_k$:

$$\alpha_k = -\frac{\boldsymbol{g}_k^T \boldsymbol{d}_k}{\boldsymbol{d}_k^T(\boldsymbol{A}^T\boldsymbol{A} - \lambda\triangle)\boldsymbol{d}_k}. \tag{11}$$

We have observed empirically that GC indeed outperforms other gradient-based methods for minimizing criterion (5).

2.3.2 Boundary conditions, valid rays and constraints

In practice, the reconstructed volume is in the shape of a parallelepiped which is significantly larger than the studied flow. It is then interesting to define a tighter working volume by defining a mask indicating the active voxels and the ones which are fixed during the optimization. This process is obviously useful for limiting the computing cost. It also helps to avoid artifacts in the reconstructed volumes as demonstrated in Ihrke and Magnor (2004). In practice, 2D masks are defined by the user on the deviation images and backprojected into a 3D mask: this process is described more precisely in Sect. 5.3. Usually, this mask is elongated along the principal axis of the flow. Boundary conditions are applied on the faces of the mask. They should be carefully chosen to avoid artifacts and differ depending on the faces of the mask, as illustrated in Fig. 1. If the face is known a priori to cut the flow (gray faces), it is associated with a free condition, while if it is known to be outside the flow support, a constant boundary condition is enforced with value ρ_0. This concept is also used to select valid rays. When browsing a valid ray, no gradient density should be encountered outside the mask. As shown in Fig. 1, a valid ray can simply be defined as a ray going through the mask by crossing only "ρ_0 boundaries" (i.e., no gray faces, as the green ray in Figure). This ensures that integration that is inside the mask, described in Sect. 2.3.3, is correct. Deviation data associated with nonvalid rays, and the corresponding lines of matrix $\boldsymbol{A}$, are deleted from the criterion (5).

Setting appropriate boundaries to an outer value ρ_0 implicitly sets the reference value of ρ inside the reconstructed flow. Indeed, the data term in (5) is invariant to the addition of a constant. The only term which constrains the average level of ρ is the regularization term through the constant boundary conditions.

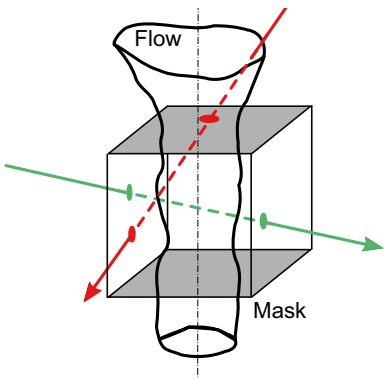

Fig. 1 Mask and boundary conditions. *Gray faces* are a priori known to cut the flow, while other faces are outside the flow. Valid rays go through the mask without crossing *gray faces* (*green ray*), and other ones are rejected from the optimization (*red ray*)

2.3.3 Parallel implementation on GPU

Here we briefly describe the parallel implementation of CG on GPU. GPU are massively parallel architectures, which involve a large number of elementary computation cores (e.g., 2688 cores in a recent Titan GPU) and a simple and adaptable hierarchy of registers, shared and global memories. Another interesting feature of GPUs is a texture memory which allows very fast interpolation. Originally designed for graphic applications only, they are since 2009 used for general programming, thanks to the release of high-level languages and associated SDK (CUDA, OpenCL). Implementation of CG algorithm on GPU amounts to identify the most computationally demanding operations, find a way to parallelize them and carefully control global memory access which is often the bottleneck of GPU. In our case, the costly operations are related to the observation matrix $\boldsymbol{A}$, which is too large to be stored.

However, the CG only requires matrix-vector products such as $\boldsymbol{A}\rho$ and $\boldsymbol{A}^T\varepsilon$. Their parallel implementations are derived from projection and backprojection operations in CT which have been studied in several works (Gross et al. 2009; Gac et al. 2010). For both, we adopt a pixel-wise ray tracing scheme (Gac et al. 2010), which means that the volume is traveled along the rays associated with the pixels of each image as sketched in Fig. 2. More precisely, the equation of the ray associated with a given pixel is determined by using the calibration parameters and the first intersection point P between the ray and the volume is located. From this point, the ray is finely discretized (the step s is typically chosen as 1/10 of the voxel size) and, for each spatial step, values are read (respectively, written) in the crossed voxel to compute $\boldsymbol{A}\rho$ (resp. $\boldsymbol{A}^T\varepsilon$). In a sequential implementation, one would then consider the next pixel, and so on for all available images. Here, thanks to the massively parallel architecture of the GPU, thousands of rays can be cast simultaneously, which drastically increases the efficiency of the calculation. The same GPU core also executes the spatial derivations using neighboring values of the density volume which can be stored in fast access local buffers. While multiple rays could be cast and averaged to account for integration on the pixel area, we have found that calculation of only one ray per pixel already provides sufficient accuracy and saves much computations.

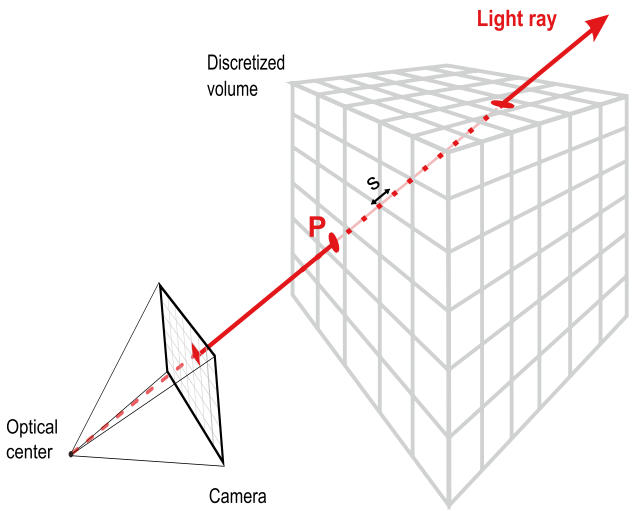

Fig. 2 Illustration of the pixelwise ray tracing operation

3 Validation of the one-step inversion method on simulated data

3.1 Test case

We consider a simulated flow coming from a calculation of coplanar jet issued from the CoJen project (Vuillot et al. 2008). The flow is generated by a coplanar double stream nozzle in static atmosphere. The inlet conditions are transonic. For primary flow, the stagnation temperature equals to 850 K with a pressure ratio of 1.4. For secondary flow, the stagnation temperature reaches 335 K and the pressure ratio equals to 1.69. A 3D view and slices of the volume are presented in the first row of Fig. 10: the density is mostly lower than the ambient one since it was heated in the turbine. We can notice vortex ring formation in the mixing layer between the static surrounding air and the $M = 0.9$ secondary flow. The density field is fully 3D with no axisymmetry, and it is then a challenging case to reconstruct with a limited number of views. The CFD computation is interpolated on a Cartesian grid of $300 \times 300 \times 300$ cubic voxels of edge length 3.6 mm.

Simulation of deviation data begins with the choice of the camera configuration. The basic configuration (denoted "coplanar180") uses 12 cameras equally spaced along half circle as presented in Fig. 3 and in Table 1. For each of the 12 cameras, its position is depicted by a green circle and the 4 rays defining the limits of the CCD are plotted in different colors.

Other 3D configurations are evaluated below. Each virtual camera is formed by 500×500 pixels, has a focal length of 14 mm, a field of view of 24° and points toward the center of the reconstructed volume at a distance of 3 m. This leads to voxel/pixel ratio of 1.4. For each camera, integrated deviations for the three directions x, y, z are computed by tracing rays within the CoJen simulated volume. Figure 4 shows the simulated deviation field for camera 5 of Fig. 3. Note that the axial component ε_x is not negligible, which means that the problem cannot be formulated as a stack of 2D tomographic

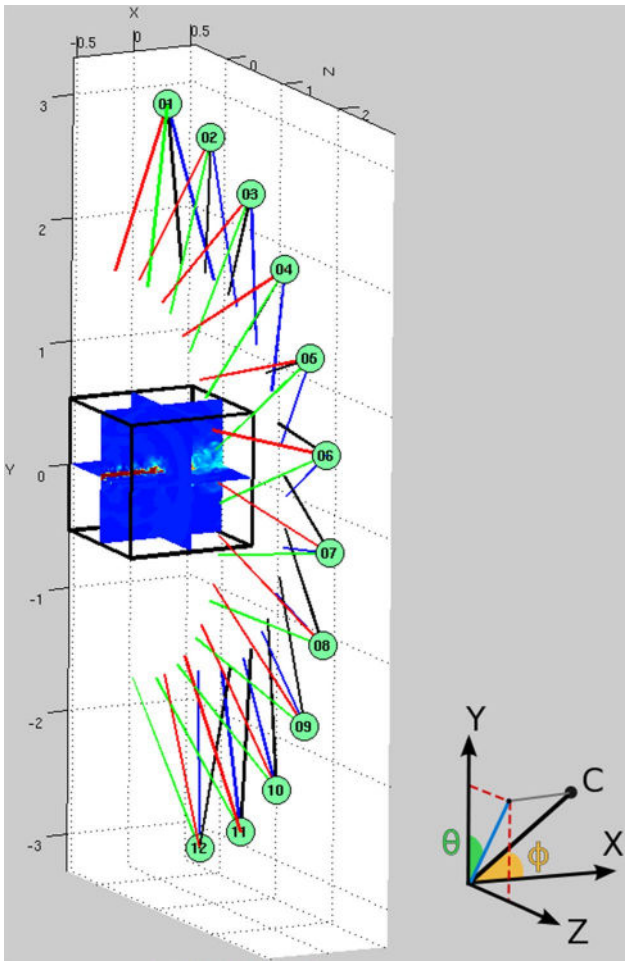

Fig. 3 Reference camera configuration for 3DBOS: "coplanar180" case

problems, as proposed in Goldhahn and Seume (2007), Alhaj and Seume (2010), and Zhang et al. (2015).

To account for imperfect measurements, we add a white and homogeneous Gaussian noise to the deviations. We choose a noise variance $\sigma = 1.0 \times 10^{-5}$ rad resulting from a displacement noise of 0.1 pixel, which is a typical order for random errors of digital image correlation methods. Of course, this is only a rough first-order model of the real errors occurring in estimated deviation fields, which are due to image noise, illumination effects, limitations of the optical flow method, etc. However, these effects will be accounted for in the real experiments. Our goal here is simply to assess the robustness of the method with respect to noisy data, which is crucial when dealing with ill-conditioned systems.

3.2 Choice of the regularization parameter

As it is well known, approaches based on the minimization of a compound criterion such as (5) rely on a "good" choice of the regularization parameter λ. Effects of under- or over-regularization are illustrated in Fig. 5. Reconstruction with too low regularization are noisy, while over-regularized ones are too smooth with clearly underestimated density levels.

We adopt here the L-curve strategy of Hansen (1992) to automatically choose the regularization parameter. Figure 6 shows that, on the tested configuration, the point of highest curvature of the L-curve corresponds to a value ($\lambda = 1 \times 10^{-4}$), which also minimizes the mean-squared error (MSE) to the true synthetic density volume.

Table 1 Cameras angular positions

Camera	Coplanar180		Shifted180		Cluster30	
	θ	ϕ	θ	ϕ	θ	ϕ
01	1°	90°	1°	100°	0°	76.3°
02	16°	90°	16°	80°	0°	90°
03	31°	90°	31°	100°	0°	103.3°
04	46°	90°	46°	80°	10°	76.3°
05	61°	90°	61°	100°	10°	90°
06	76°	90°	76°	80°	10°	103.3°
07	91°	90°	91°	100°	20°	76.3°
08	106°	90°	106°	80°	20°	90°
09	121°	90°	121°	100°	20°	103.3°
10	136°	90°	136°	80°	30°	76.3°
11	151°	90°	151°	100°	30°	90°
12	166°	90°	166°	80°	30°	103.3°

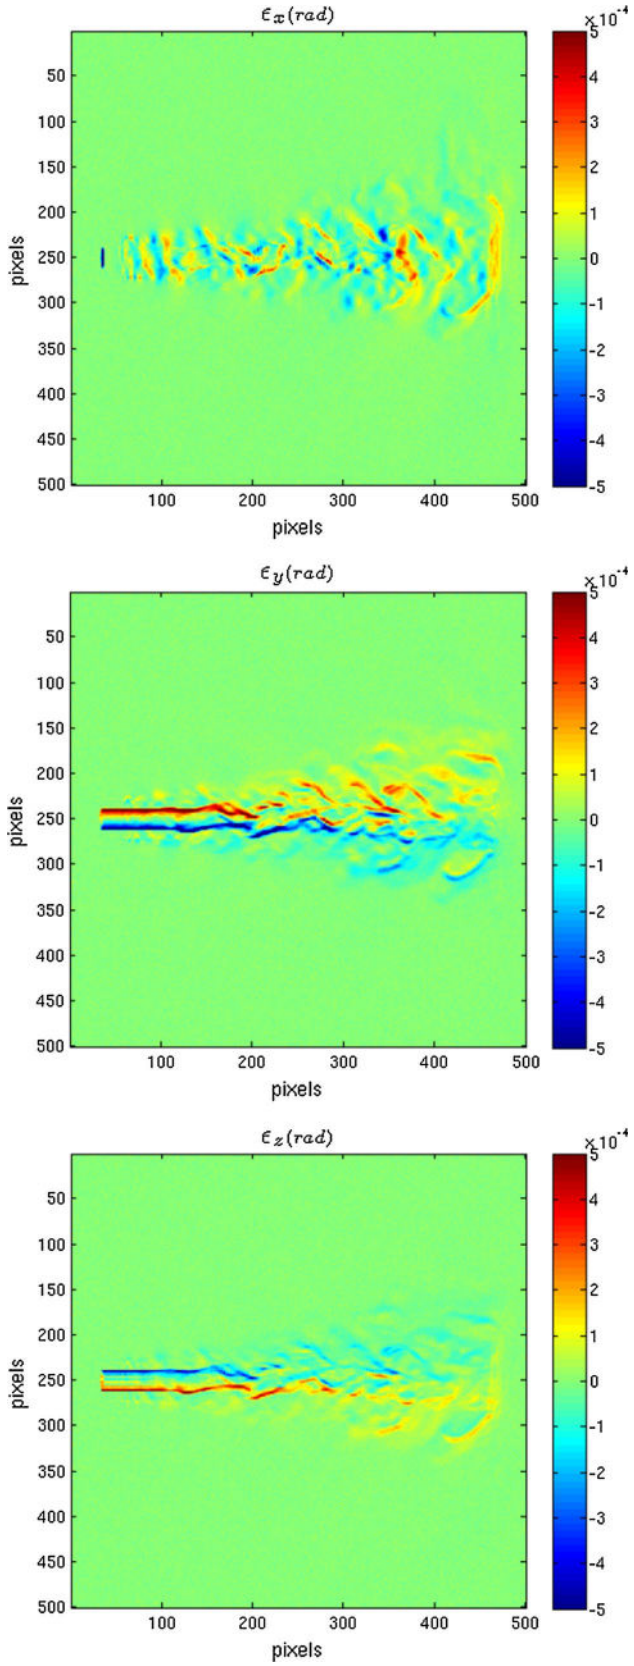

Fig. 4 Simulated deviation field in the geometry of camera 5 of the coplanar180 configuration shown in Fig. 3. From *top* to *bottom*, the x, y and z components of the deviation, *color scale* in radians.

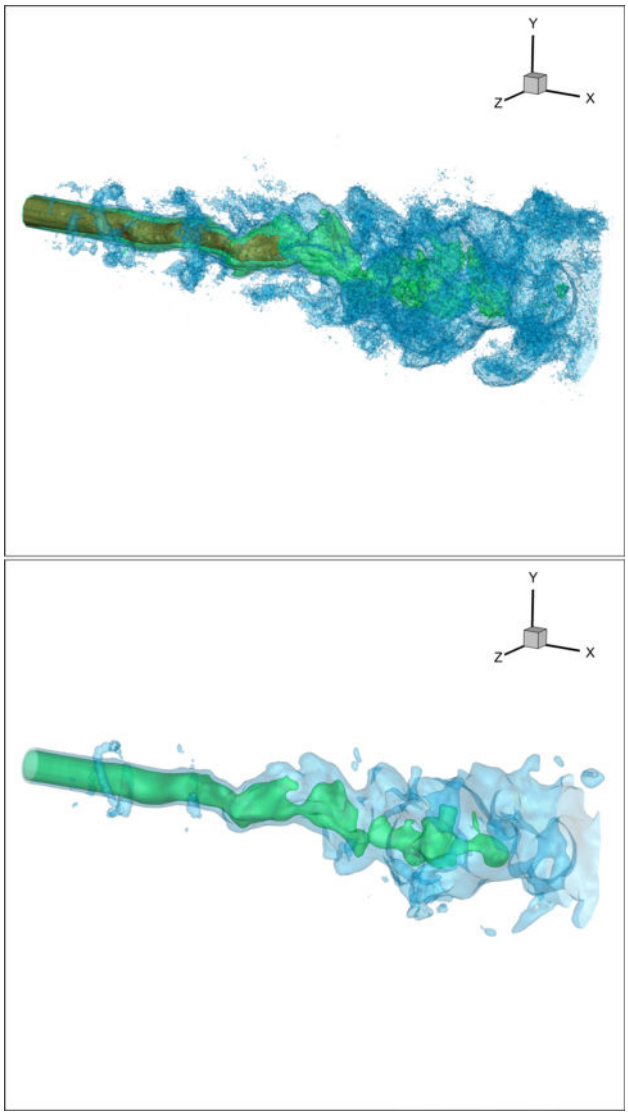

Fig. 5 Choice of the regularization parameter. *Top* under-regularized solution ($\lambda = 1.0 \times 10^{-6}$). *Bottom* over-regularized solution ($\lambda = 1.0 \times 10^{-3}$)

The reconstruction obtained with the optimal regularization parameter according to the L-curve criterion is shown on the second row of Fig.10. Iso-surfaces, XY and YZ slices of the reconstructed volume are represented, which should be compared with the views of the true volume located on the top row of Fig.10. The iso-surface representation shows that the morphology and levels of the true volume are well reconstructed. However, low density regions inside the jet are smoothed, particularly in the expanded part on the right, as can be seen on the XY slices in the second column. Artifacts related to the limited number of projections are visible in the YZ slice on the right. Those finger patterns are associated with the camera direction.

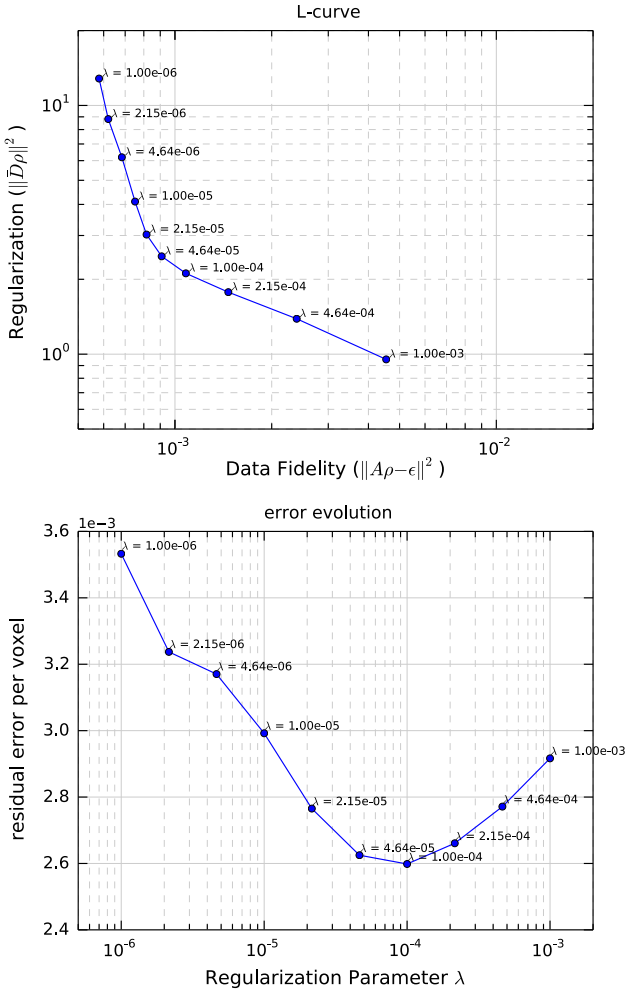

Fig. 6 Choice of the regularization parameter. *Top* L-curve. *Bottom* MSE with respect to the true volume. Regularization parameter values are written along the curve

3.3 Influence of the mask

As introduced in Sect. 2.3.2, we use a 3D mask to improve the efficiency of the regularization. In this section, we briefly discuss the influence of the mask on the 3D reconstruction. In practice, the 3D mask is derived from 2D masks defined interactively by the user on the displacement images, as described in Sect. 5.3. The resulting reconstruction, denoted by "interactive mask", is shown in the second row of Fig. 10. On this simulated example, a tighter mask can be obtained by directly thresholding the true density volume. However, the resulting reconstruction is very close to the one deriving from the interactive mask and is not displayed.

Besides, we present on the third row of Fig. 10 a reconstruction obtained without any mask. The reconstruction without the mask captures the general form of the flow but severely underestimates the density variations in the center. In the slices presented in the two rightmost columns, the energy of the reconstruction appears to be spread over the whole volume. This spreading phenomenon in the absence of a mask (or more generally, of a "support constraint") is well known in tomography, and it has also been noticed by Ihrke et al. in their work on 3D density reconstruction (Ihrke 2007). Density profiles extracted from those reconstructions are compared to the CFD reference on Fig. 7. Whereas the reconstruction without any mask smooths all the density gradients, the interactive mask result shows a good behavior, being able to capture most of the density variations. Huge variations like the fall toward density 0.45 kg/m^3 are restored, though the recovered gradients are smoother than the originals. One can also notice the

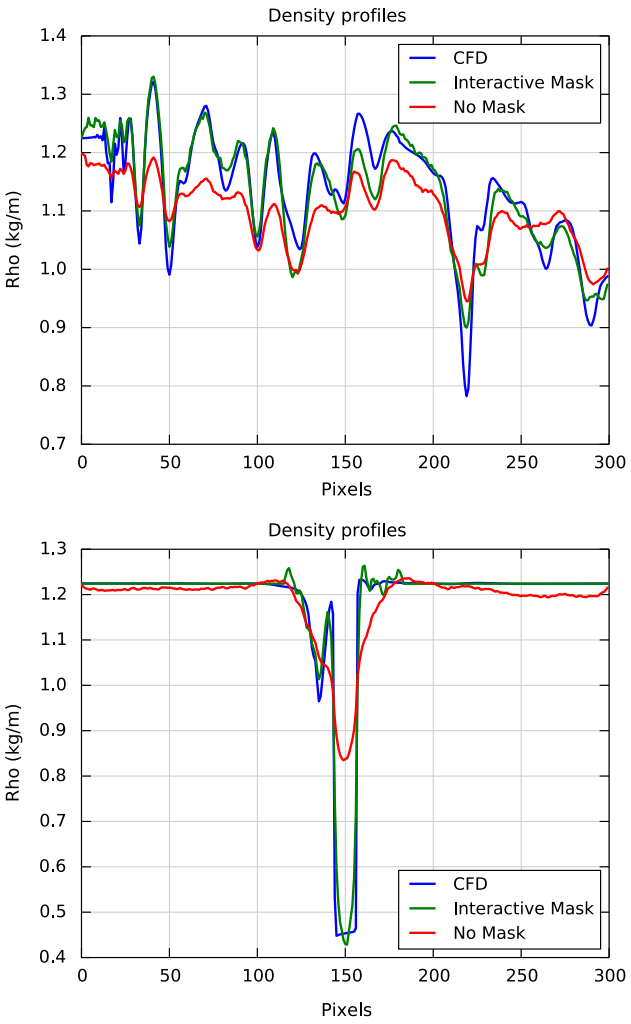

Fig. 7 Profiles of density along slices extracted from the volumes presented in the three first lines of Fig. 10. The positions of the profiles are indicated in the CFD XY Slice of Fig. 10. *Top* Slice 1 ($Y = -15$ voxels). *Bottom* Slice 2 ($X = -105$ voxels)

presence of overshoots in density levels at the edge of the 3D mask on the lower plot of Fig. 7. These results provide empirical evidence that using a mask not only accelerates the computations but also significantly improves the reconstruction.

3.4 Geometrical configuration of the cameras

Not surprisingly, Fig. 8 shows that, for the coplanar configuration, the use of a higher number of cameras yields a better reconstruction. However, even for a complex volume such as the one studied here, the marginal gain also decreases. Moreover, it is unrealistic to assume that 50 cameras could be used in a 3DBOS setting. Therefore, we limit our study to 12 cameras and investigate various geometrical configurations in this section.

The coplanar180 is compared with a variant where cameras are shifted along the main axis of the jet and rotated so as to all point toward the same point in the middle of the flow. The tilt angle of the camera is 10° (to be compared with the vertical field of view 24°). This configuration called "shifted180" is shown in Fig. 9 (left). Another configuration, depicted in Fig. 9 (right) is the "cluster30" where all cameras are confined in an angular region of 30° × 30°, which can represent situations where optical accesses are very limited. The detailed orientation of the cameras for all tested configurations can be found in Table 1.

Results with noncoplanar configurations are presented in the lower part of Fig. 10. All results are obtained with an optimal 3D mask and using the regularization parameter derived from the L-curve strategy. The result obtained with the "shifted180" configuration appears quite similar

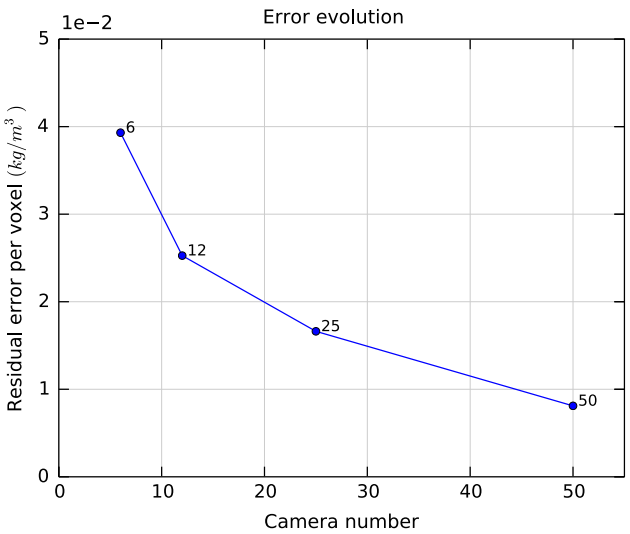

Fig. 8 Residual error per voxel as a function of the number of cameras used in the reconstruction process

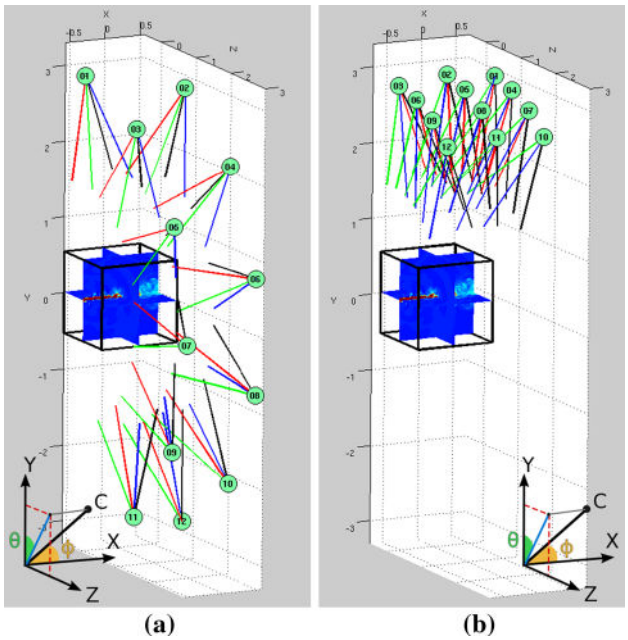

Fig. 9 Noncoplanar camera configurations. **a** Shifted180 **b** Cluster30

to the "coplanar180", even if its MSE is slightly higher. Indeed, for flows having a geometrical main axis (such as jet flows), the coplanar configuration, where all cameras belong to a plane orthogonal to the main axis, appears optimal. But the result indicates that varying the cameras' viewpoint around this configuration does not significantly alter the quality of the reconstruction. In contrast, the "cluster30" result shows that when the variety of available viewpoints is too restricted the reconstruction is very poor and information is spread out along camera axis. In such situations, one would probably reduce the number of cameras and consider 2D BOS only, 3D reconstruction being possible only for axisymmetric objects.

3.5 Conclusion of the simulation study

The proposed one-step density reconstruction has been assessed on simulated deviation data. Using 12 cameras providing various angles on the object, a good reconstruction can be obtained. In our quadratic regularization setting, the L-curve strategy is efficient to determine the correct regularization parameter. While several parameter studies have been conducted, we have only presented the effect of the geometrical configuration of the camera which is of primary interest for practical use of 3DBOS in experimental facilities. As could be anticipated, for the studied jet, the best reconstruction is obtained by a coplanar configuration in the main axis orthogonal plane and when covering as much as possible the circle of view around the object.

4 Experimental bench for instantaneous 3DBOS: the Geode

The Geode facility dedicated to the study of instationary flow by 3DBOS has been designed and built in ONERA DMAE. This bench is a small rhombicuboctahedron, 8 of its faces are triangular and 18 are square. The diameter Z of the Geode is 3.20 m which makes it possible to reconstruct large flows. As it can be seen in Fig. 11, half of the Geode is used to support the cameras, while the other half supports the backgrounds.

4.1 Camera setup

This installation is equipped with 12 JAI cameras BM500 GE. The cameras are mounted all around the structure, and each of them is fixed on a ball joint allowing the camera to rotate in every direction. The design of this 3DBOS test bench allows setting up various camera configurations. All cameras are mounted with 23 mm focal length Schneider C-mount optics. These optics allow for the visualization of a common cubic volume with edge length of about 50 cm placed at the center of the geode. The main characteristics of the cameras and main optical parameters are summarized in Table 2. During the experiment, a uniform illumination is provided by four 500 W halogen spots. The amount of light being limited, all the tests were conducted with an aperture corresponding to a f-number $f_{\#} = 2$. The cameras focus is set on the backgrounds which results in a depth of field of 26 cm.

The cameras are synchronized with a TTL generator which can handle 24 independent signals. To be able to record 10 fps with 5 Mpx images, on each camera, we used a C5G124-24 *Enterasys* switch connected to 20 Gb/s *Intel* network card. All images are stored in the RAM memory during the acquisition and then transferred on SSD disks. Following this process, we perform an acquisition of 900 images per camera in 4 min 30 s:1 min 30 s for the acquisition and 3 min for the storage.

4.2 Background

Multiple backgrounds patterns have been tested. The wavelet background (Atcheson et al. 2008), whose main advantage relies on its multiscale random content, did not perform well in areas with large density gradients. Random distribution of points is a simple and efficient way to design backgrounds; however, they can lead to local inaccuracies in the estimated displacement fields because the number of dots inside a correlation window is not constant. To overcome this issue, we have designed a new background pattern named "semi-random background" (Fig. 12b). It is obtained by small-scale random perturbations of dots around a regular grid pattern: the regular pattern guarantees a large number of dots inside the correlation window, while random perturbation prevents matching ambiguities during correlation.

In the facility, the backgrounds have been printed on aluminum plates and cut at the exact size of the triangle and square surfaces. They are held with small screws with a spring nut positioned in the slot of the BOSCH beams.

4.3 Geometric setting and calibration

The geometrical configuration is sketched in Fig. 13. Each camera observes the background plate in front of it, where the focus is set, and the flow under study is placed halfway between the cameras and the background. From the dimensions of the bench and the camera parameters, the diameter of the circle of confusion due to defocus at the flow location can be estimated at 6 mm—approximately half the diameter of the aperture $f/f_{\#}$. It means that the smallest observable spatial scales of the gradients of ρ are of the order of 6 mm. In order to mitigate this effect, some authors (Bichal and Thurow 2014) propose to focus the cameras halfway between the background and the flow. Such a configuration would lead to a better circle of confusion (about 4 mm), but this would be compensated by a reduced deviation-to-displacement sensitivity (factor 1.5) and a lower accuracy of estimated displacement vectors due to a blurred background image. So the overall gain of this alternate configuration is difficult to foresee. A better overall sensitivity could probably be obtained but would require an optimization of the experiment which falls outside the scope of the present study.

3DBOS reconstruction as described in this paper requires synchronous observation of the flow by a set of cameras, typically 12 cameras in our bench. To compute and optimize criterion (5), the equation of each ray associated with some pixel of the image plane of a camera should be calibrated in the coordinate frame fixed to the reconstructed volume (called the working frame in the sequel). This operation requires a geometrical calibration of the multicamera setting in the same coordinate frame, i.e., the identification of all internal parameters of the cameras (focal lengths, principal points, distortion parameters) and of all camera poses (position and orientation).

Multicamera calibration is a well-documented problem, and developments specific to the 3DBOS have been discussed in Le Sant et al. (2014). Here, as in several other experiments, we use as a calibration body (CB), a simple white plane with regularly spaced black dots. Three extra dots are added to capture the orientations of the CB. The basic principle of our multicamera calibration is to record images of a CB moved around the common field of view of the cameras and to minimize the reprojection errors of all visible features of the CB with respect to calibration parameters. However, there are three specific difficulties in the calibration of a 3DBOS bench.

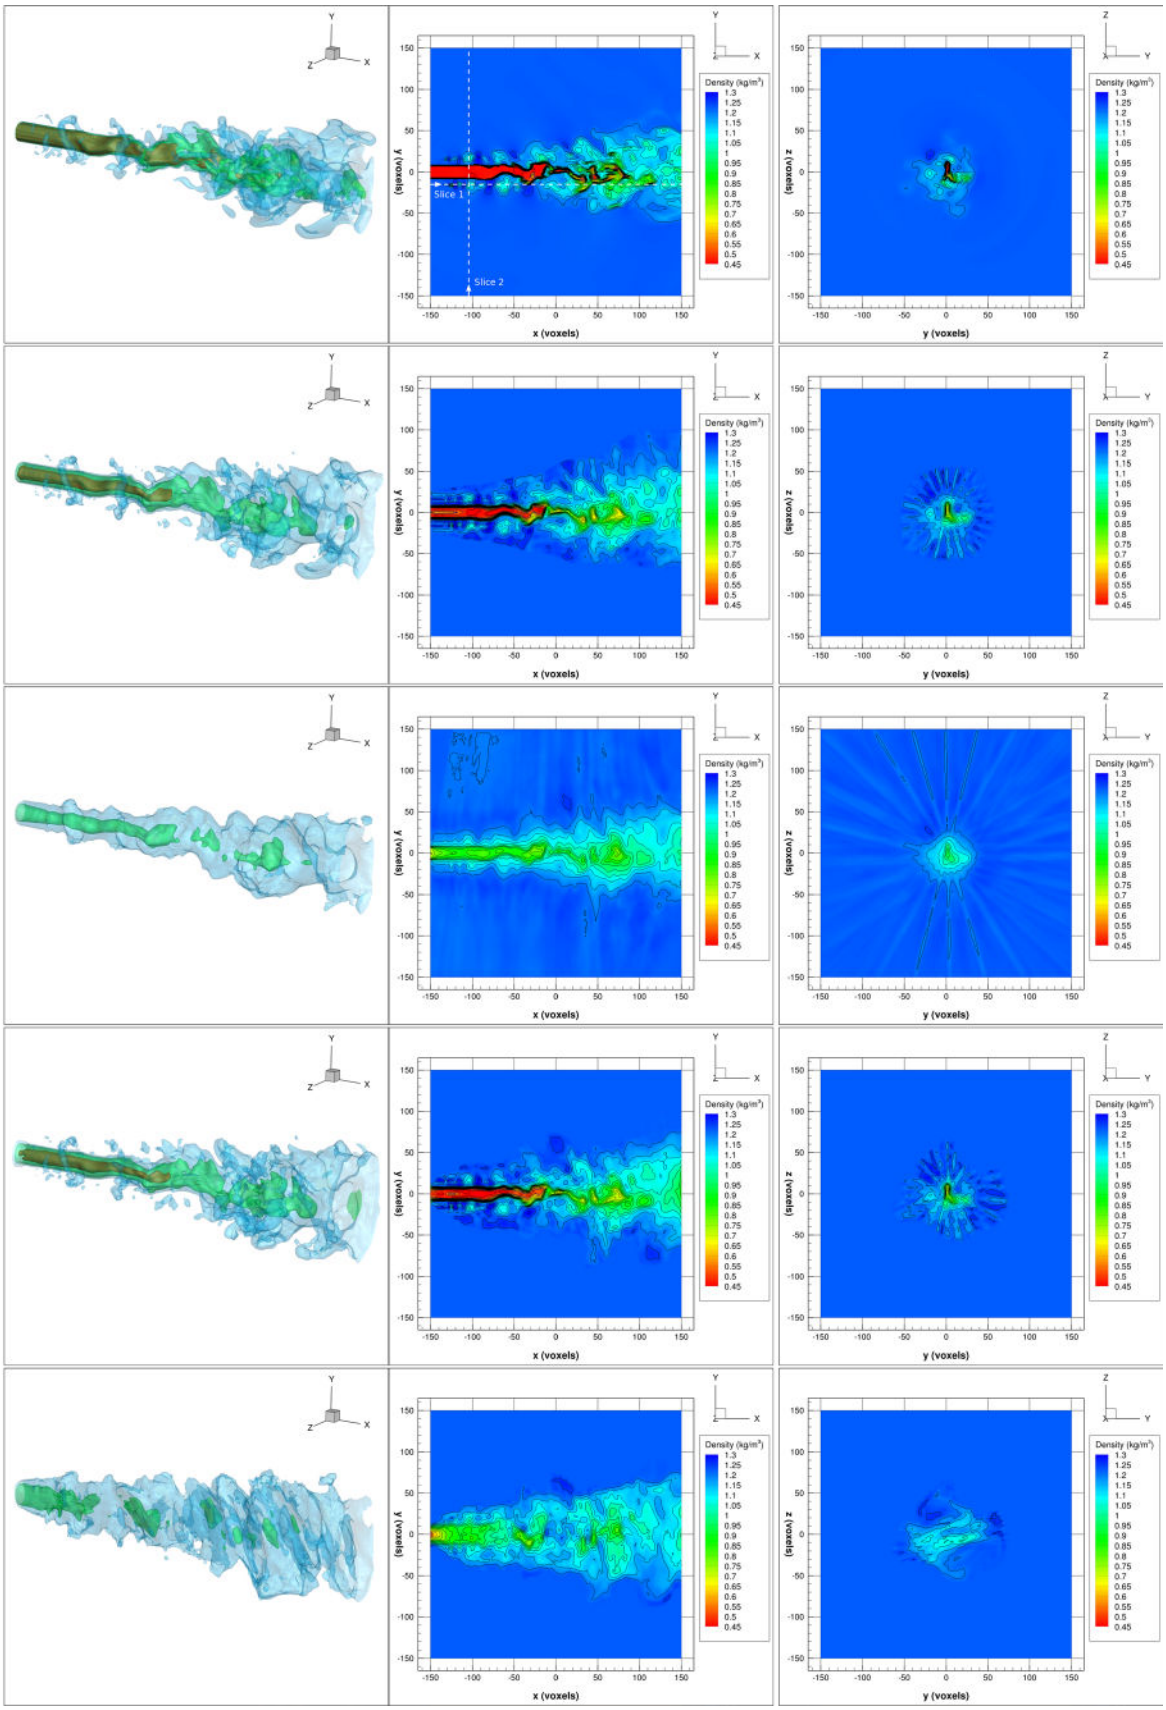
CFD
Interactive mask
No mask
Shifted 180
Cluster 30
Slice 1
Slice 2
x (voxels)
y (voxels)
z (voxels)
Density (kg/m³)

◀**Fig. 10** Ground truth density volume from CFD (*first row*) compared with various reconstructed volumes: using the coplanar 180 camera configuration and an interactively user-defined mask (*second row*) or no mask (*third row*); using the shifted 180 configuration (*fourth row*) and the cluster 30 configuration, both with a user-defined mask. All reconstruction are obtained with $\lambda = 1.0 \times 10^4$. From *left* to *right*: 3D representation of 3 iso-density surfaces 0.5 (*red*), 0.9 (*green*), 1.1 (*blue*); (*X*, *Y*) slice; (*Y*, *Z*) slice

First, depending on the camera configuration, it is often not possible that the CB is viewed from all cameras. However, in all tested configurations, all cameras are located in the same half-space and the CB is always seen at least by 3 cameras. In this case, it is still possible to calibrate all camera parameters in a consistent way by exploiting a chain of correlation between views having CB features in common.

The second issue is illustrated in Fig. 13: the region where the CB can be viewed by several cameras is the workspace (where the flow will be put), located halfway between the cameras and the background. In the Geode, illumination was limited and we were forced to work with a large aperture, leading to severely blurred images of the CB. In Le Sant et al. (2014), we have compared several calibration strategies including in-focus images of the CB together with blurred images of the CB located in the workspace. It has been found that using in-focus images is not necessary and that an accurate calibration can be obtained with blurred images of the CB located in various positions inside the workspace. In practice, the calibration process requires about 50 images per camera which are recorded, while the CB is moved by hand into the workspace.

The third difficulty is related to the workload of the calibration. Indeed, there are several hundreds of images to process which is at least an order of magnitude higher than for other optical measurement methods such as PIV or TomoPIV. A real-time calibration is not required, but the

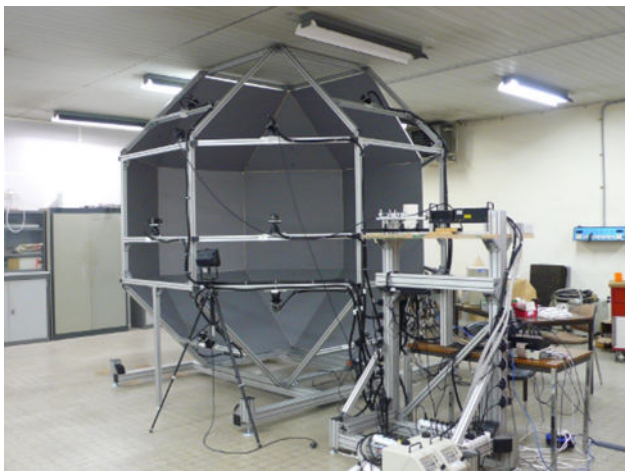

Fig. 11 The Geode, a 3DBOS experimental facility

Table 2 Camera JAI BM500 GE characteristics and acquisition parameters

Sensor	CCD, 2/3”
Pixel resolution	2456 × 2058
Pixel size	3.45 μm × 3.45 μm
Acquisition frequency	10 images/s
Exposure time	750 μs
Aperture	$f_\# = 2$

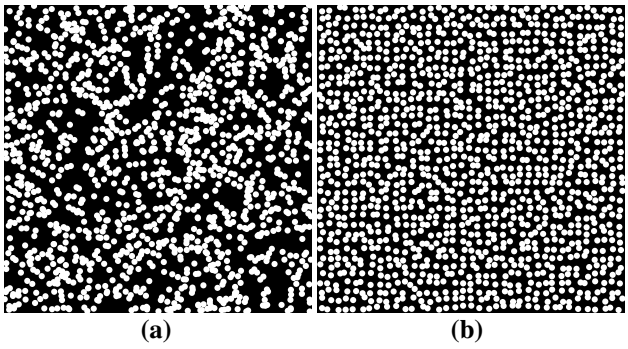

Fig. 12 The semi-random background guarantees a fixed number of dots per correlation window with random arrangement inside the window. In contrast, the random background (Poisson distribution) may occasionally have too few dots inside a correlation window. Both backgrounds presented here have the same number of dots. **a** Random background. **b** Semi-random background

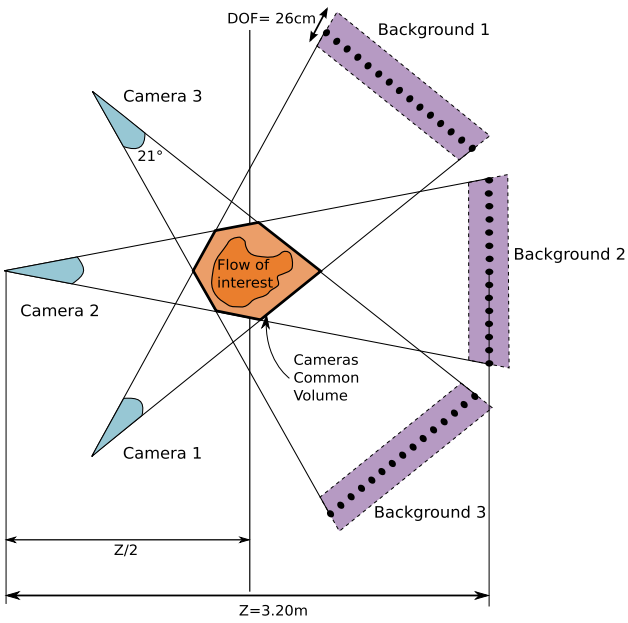

Fig. 13 Sketch of the geometrical configuration of cameras in a typical 3DBOS setting. The cameras are focused on the backgrounds

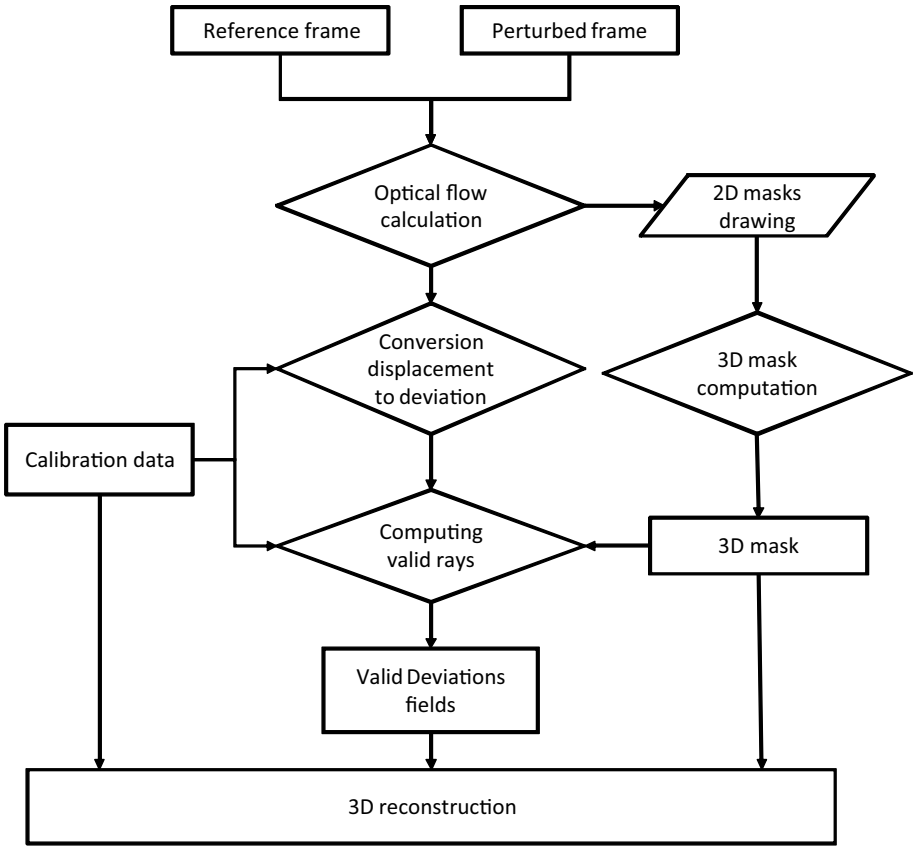

Fig. 14 Processing chain relating the recorded frames to the deviation fields which are fed to the 3D reconstruction software

calibration must be carried out within a reasonable delay (say below 1h). This issue has been solved using a mixed CPU-GPU implementation [see Le Sant et al. (2014)].

5 From images to deviation fields

In this Section we describe the operations needed to measure deviation fields in a typical multicamera 3DBOS experiment. We discuss about the generic experimental methodology and processing chain which have been developed to conduct 3DBOS reconstruction in several experimental facilities at ONERA. Numerical values such as camera parameters or actual distance between camera and background panels related to the particular geode 3DBOS bench are given in Sect. 4.

We list below the steps of the acquisition and processing chain, which are also summarized in Fig. 14. More details are given in the referenced Sections.

1. Multicamera calibration (Sect. 4.3)
2. Acquisition of background images with and without the flow
3. Correction of geometrical distorsions of the images
4. Estimation of image 2D displacement fields by optical flow calculation (Sect. 5.1)
5. Conversion of 2D displacements into 3D deviations (Sect. 5.2)
6. 3D mask creation and selection of valid rays (Sect. 5.3)

5.1 Displacements estimation

Images of a textured background are recorded without flow (reference frame) and with the flow (perturbed frame) in between the camera and the background. The image displacement field between these two frames is computed using the FOLKI software developed by ONERA (Plyer et al. 2014), which is described within the PIV context in Champagnat et al. (2011). FOLKI relies on the Lucas-Kanade paradigm of iterative registration of local interrogation windows. It provides a dense displacement field (i.e., one vector per image pixel) although, as discussed in Champagnat et al. (2011), its spatial resolution is limited by the size of the interrogation window, which is the main parameter of the method. The norm of an experimental 2D displacement field is shown in Fig. 16. Thanks to its highly parallel structure, the GPU implementation of FOLKI is computationally very efficient: as shown in Fig. 7 of Plyer et al. (2014), 5 Mpixels images are processed in 30 ms on a Titan GPU. It is of high interest in 3DBOS where several large images (e.g., 12 images of 5 Mpixels) have to be correlated at each acquisition time.

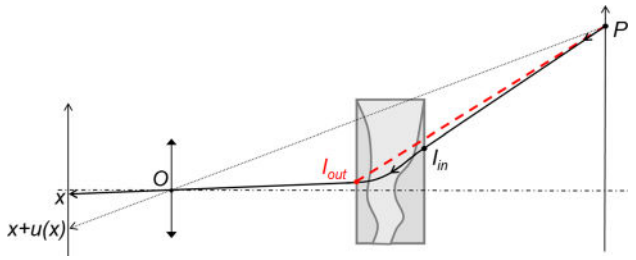

Fig. 15 Principle of BOS measurement. A point P of the background is imaged in position x with the flow and in $x + u(x)$ without

5.2 Converting image displacement into deviations

Figure 15 recalls the principle of BOS measurement. A detail at position x in the perturbed image frame is associated with point P in the background pattern. Point P was previously imaged at position $x + u(x)$ in the reference frame. Conversely, from the estimation of the displacement $u(x)$, and knowing the camera calibration, one is able to locate P. We assume that some external envelope of the flow, represented as a gray square in Fig. 15, is known. This envelope is a surface separating the flow, i.e., the region of space where the optical index can vary, and the outer region where the optical index is assumed to be constant. In practice, we simply use the boundaries of the 3D volume of reconstruction as an envelope. The deviation angle ε of Eq. 1 is the difference between the directions of the ray (I_{out}, O) outgoing from the envelope and of the ingoing ray (P,I_{in}):

$$\varepsilon = \boldsymbol{a}_{\mathrm{I_{out}O}} - \boldsymbol{a}_{\mathrm{PI_{in}}}$$

where $\boldsymbol{a}_{\mathrm{AB}}$ stands for the unit vector associated with the segment (A, B). However, it requires the knowledge of the point I_{in} where the ray coming from P impacted the envelope. This position is not known accurately, in contrast to the position of point I_{out} which can be obtained from the ray associated with pixel x. The deviation is then approximated by

$$\varepsilon \approx \boldsymbol{a}_{\mathrm{I_{out}O}} - \boldsymbol{a}_{\mathrm{PI_{out}}}. \qquad (12)$$

5.3 3D mask

The 3D mask is helpful to increase the efficiency of the optimization. The 3D mask computation is done interactively from camera displacement fields. For each camera, the user draws a 2D mask on the image of the estimated displacement norm (see Fig. 16). Each voxel of the volume is projected into the image plane of each camera and incremented if the projection is inside the corresponding 2D mask. Voxels are retained if the obtained total number is above some threshold, which means that they are inside the 2D mask of a given number of cameras. Finally, valid deviations are selected. For each candidate deviation, the corresponding ray is casted toward the volume. The faces of the 3D mask crossed by the ray are identified. If one of the crossed faces cuts the flow as shown in Fig. 1, the corresponding ray and its associated deviation data are removed. This process allows to update the 2D masks as shown in the lower image of Fig. 16.

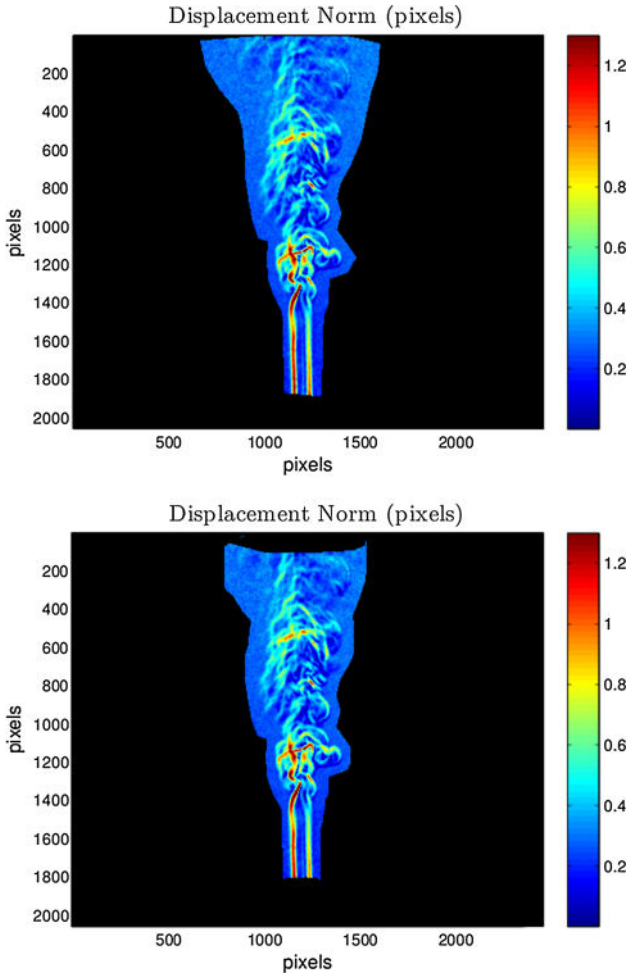

Fig. 16 *Upper row* Norm of the displacement field estimated by FOLKI and 2D mask drawn by the user. *Lower row* Final 2D mask after removing nonvalid rays

6 Experimental results and discussion

We present here results obtained on four convective flows: a candle plume, a hot jet generated by a heat gun, and two helicoidal plumes obtained by rotation of a gas burner and of an ember. These tests allow assessing our method on various density volumes with weak gradients and complex

Table 3 Experimental results: size of the reconstructed volumes and computation times

	Volume grid	Computation time
Candle	207 × 422 × 213 (18.6 M)	2.2 h
Hot Jet (inst.)	470 × 344 × 343 (55M)	5.9 h
Hot Jet (mean)	336 × 472 × 361 (57.3 M)	6.3 h
Rotating gas burner	511 × 449 × 566 (130M)	7.8 h
Rotating ember	560 × 421 × 420 (99 M)	10.6 h

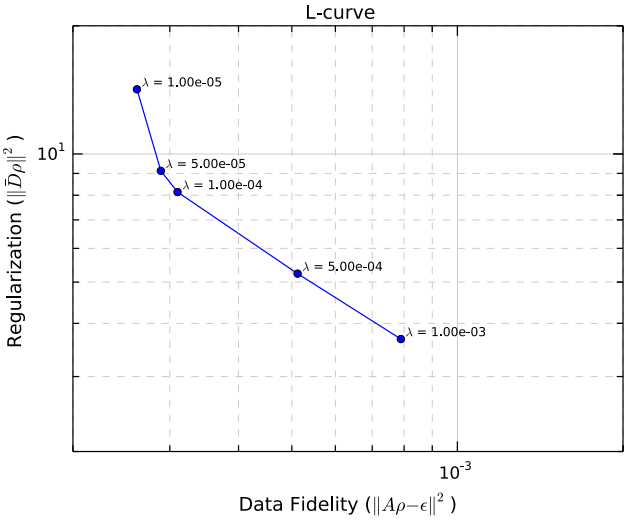

Fig. 17 L-curve of the 3D reconstruction of a candle plume for a coplanar configuration

3D structures. For each case, the flow is recorded using the Geode with 12 cameras put in different 3D configurations. Each camera acquires 500 frames allowing the reconstruction of both instantaneous and mean density fields. Mean fields are derived from averaged displacement images. Moreover, atmospheric pressure and temperature are measured during the tests in order to evaluate the reference density outside the flow. Figures of volume sizes and processing times are gathered in Table 3: the largest reconstructed volume is about 130 megavoxel, and its reconstruction takes approximately 8 h with 500 iterations.

6.1 Plume of a candle

The plume flow is simply generated by a burning candle. Combustion of liquid wax transported by capillary flow through the porous wick results in a flame with a highly nonlinear temperature profile in which local temperatures in the candle flame can exceed 1400 °C. Heat transfer from the candle flame generates a hot plume led by natural convection. Figure 18 illustrates the 3D reconstruction of the convection flow observed a few centimeters above the flame of a candle using a coplanar configuration of the cameras. The size of mesh is about 18 megavoxels for a resolution of 1 mm. The regularization parameter $\lambda = 5 \times 10^{-5}$ was automatically chosen according to the L-curve shown on Fig. 17.

The reconstructed air flow is consistent with the evolution of a slowly ascending hot plume classically observed above a candle's flame. In particular, the flow transition is clearly observed. Neglecting the combustion products, the reconstruction shows a density range comprised between 0.7 and 1.2 kg/m^3 (a proper comparison on temperature profiles can not be done here because of the influence of the different chemicals species). After the destabilization of the flow, the turbulence induces a very quick increase in the density. In order to check the quality of the reconstruction, ray tracing has been performed through the reconstructed volume and the generated deviations were compared to the measured ones. In Fig. 19, we provide measured/generated 2D deviation fields and also two profiles of both fields along lines passing across the flow, one in the core of the flame, the other in the turbulent region. A very good fit between measured deviations and generated ones can be observed: only a slight smoothing of small-scale structures is noticeable in the turbulent part. One can conclude that the proposed reconstruction method exploits the deviation data almost fully.

6.2 Hot plume

The reconstruction of an instantaneous hot air plume generated by a heat gun is presented in Fig. 20. The hot gun provides a hot jet with an inlet velocity of 5 m/s and an estimated temperature of about 650 K. The Reynolds number based on inlet diameter is about 3700. The jet was scanned according to a coplanar distribution of the cameras. Setting the exposure time to 750 µs leads to a maximal spatial displacements around 5 mm, and this is consistent with the 6-mm-diameter circle of confusion induced by defocus at the flow location, see Sect. 4.3.

Despite its complex nonaxysymmetric structure, the jet is well reconstructed and appears larger and more turbulent than the candle plume. In Fig. 21, the mean density field, obtained by processing averaged deviation images (500 samples), is presented. Boundary artifacts occur at the inlet

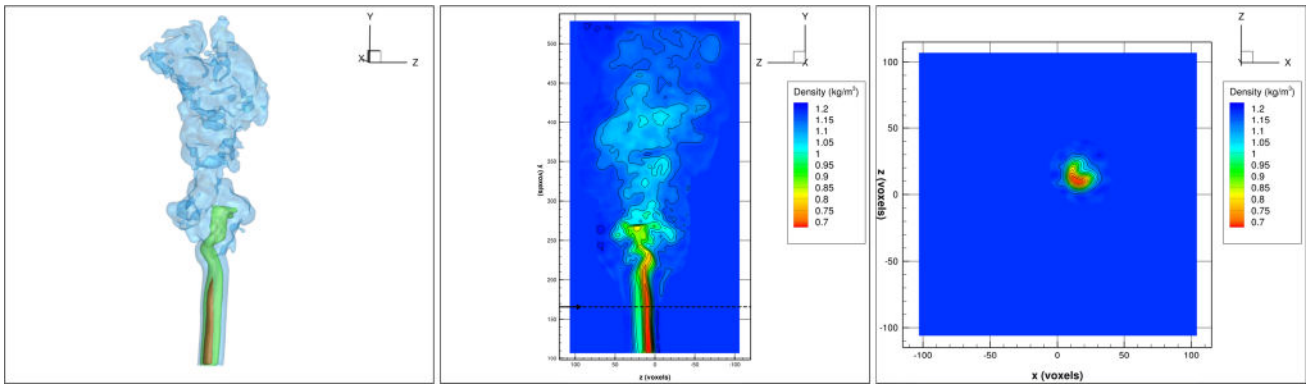

Fig. 18 Optimal reconstruction ($\lambda = 5 \times 10^{-5}$) of a candle plume for a "coplanar180" configuration: 3D representation of 3 iso-density surfaces 0.7, 0.9, 1.1, respectively, represented in *red*, *green* and *blue*; (Y, Z) slice; (X, Y) slice

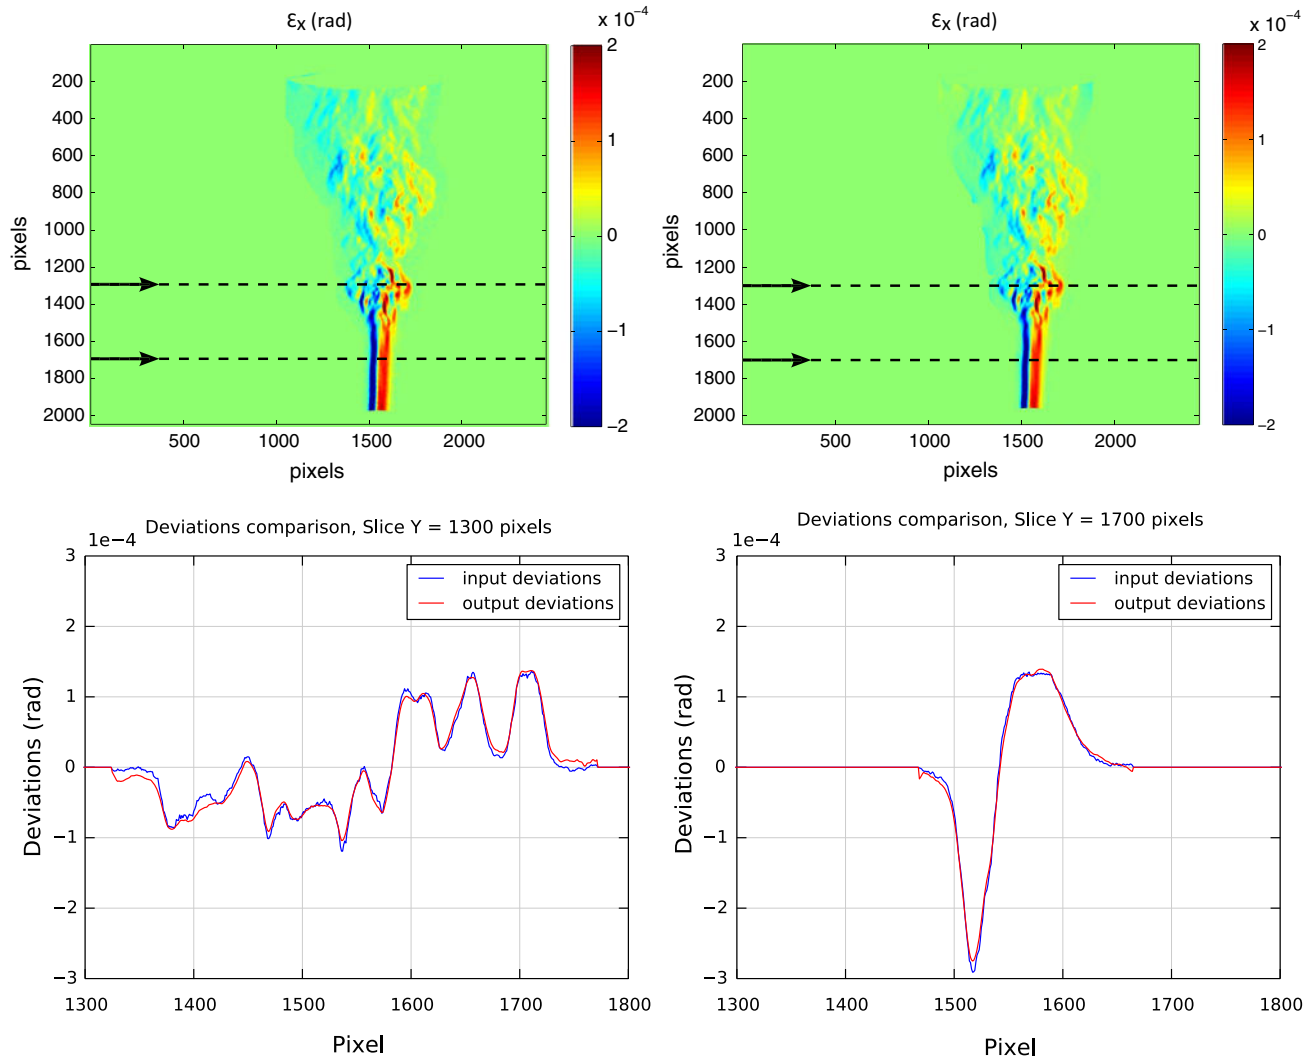

Fig. 19 Comparison between input ε_x deviations (*top left*) and computed deviations through the reconstructed volume (*top right*). Slices at $Y =$ 1300 pixels (*bottom left*) and $Y = 1700$ pixels (*bottom right*)

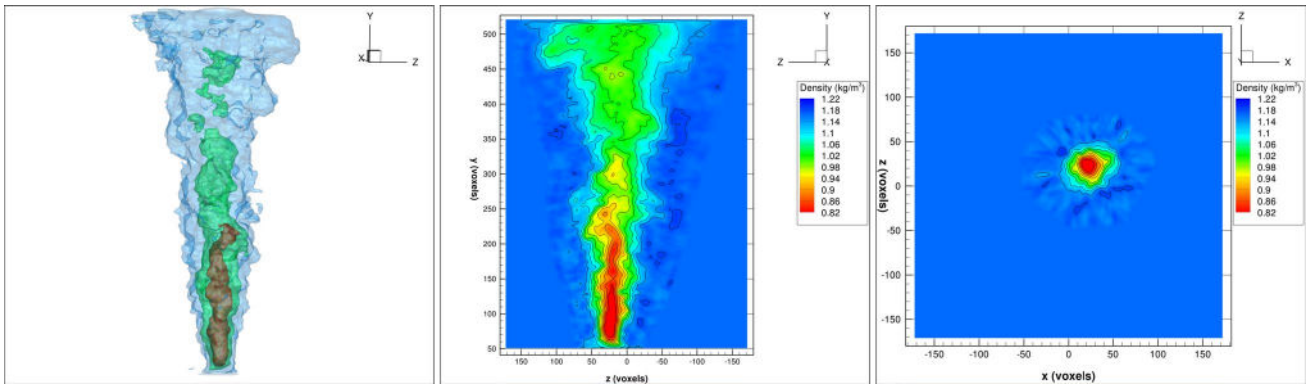

Fig. 20 Optimal reconstruction ($\lambda = 1 \times 10^{-4}$) of a hot jet for a "Shifted180" configuration: 3D representation of 3 iso-density surfaces 0.95, 1.05, 1.15, respectively, represented in *red*, *green* and *blue*; (Y, Z) slice; (X, Y) slice

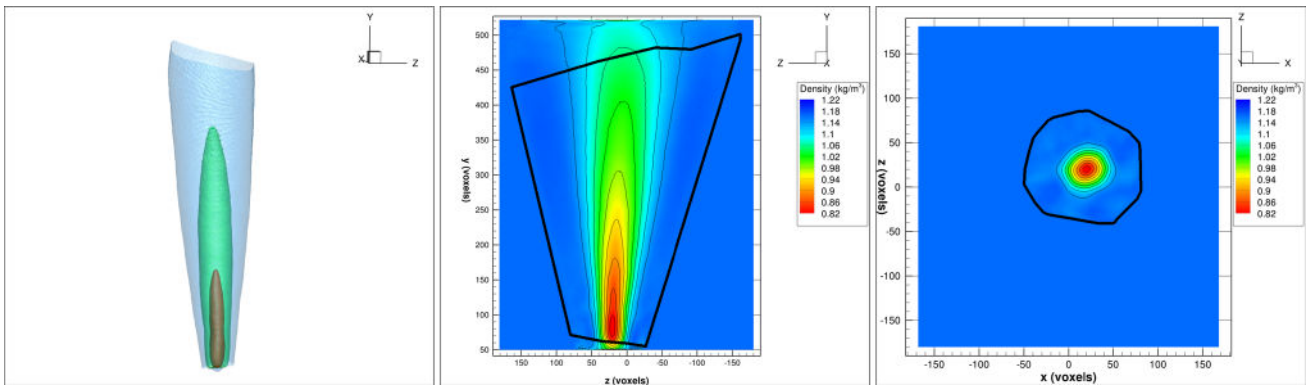

Fig. 21 Optimal reconstruction ($\lambda = 1 \times 10^{-4}$) of the mean hot jet for a "shifted180" configuration: 3D representation of 3 iso-density surfaces 0.95, 1.05, 1.15, respectively, represented in *red*, *green* and *blue*; (Y, Z) slice; (X, Y) slice

and outlet edges of the domain. These misleading alterations are a consequence of the reduction in available data near the boundaries. Indeed, the black polygon indicates the limit of the common domain seen by all the 12 cameras. Outside, the lack of measurement alters the reconstructed solution. Inside the common polygon, the 3D reconstruction is consistent.

The same flow has also been recorded using the non-coplanar "shifted180" configuration of cameras shown in Fig. 9. Results are very close to the ones obtained with the coplanar ones as could be expected from the simulation study. Differences are not significant and are most likely due to variations of SNR in the recorded images.

The density amplitudes obtained in the mean field reconstruction are compared to profiles acquired with thermocouples in Fig. 22. We sample the hot air plume every 2 mm with a type K thermocouple using a timeout interval of 60 s in order to reduce the impact of the heat gun temperature fluctuations. Neglecting pressure variations, the density can be derived from the temperature measurement using the ideal gas relation. Whereas the agreement is fairly good, the comparison between both profiles shows a slight underestimation of the BOS result for the profile nearest to the base of the jet ($Y = 230$ mm). Such an underestimation of high gradients is classical with quadratic regularization (it has been also observed in Fig. 7 in the simulation study), and it could be mitigated by using a tighter mask.

6.3 Rotating flows

In order to significantly increase the complexity of flow topology, different heat sources have been fixed on a rotating arm with a very slow rotation speed. Tests have been conducted for both a butane torch and an ember. Results are presented, respectively, in Figs. 23 and 24. In the first case, the turbulent plume describes a helicoidal path and disappears quickly due to the high diffusion induced by the rotation. The reconstructed density reveals a temperature field at less than 373 K, illustrating the high sensitivity of the proposed 3DBOS method which appears well adapted to such convective flows with very weak density gradients.

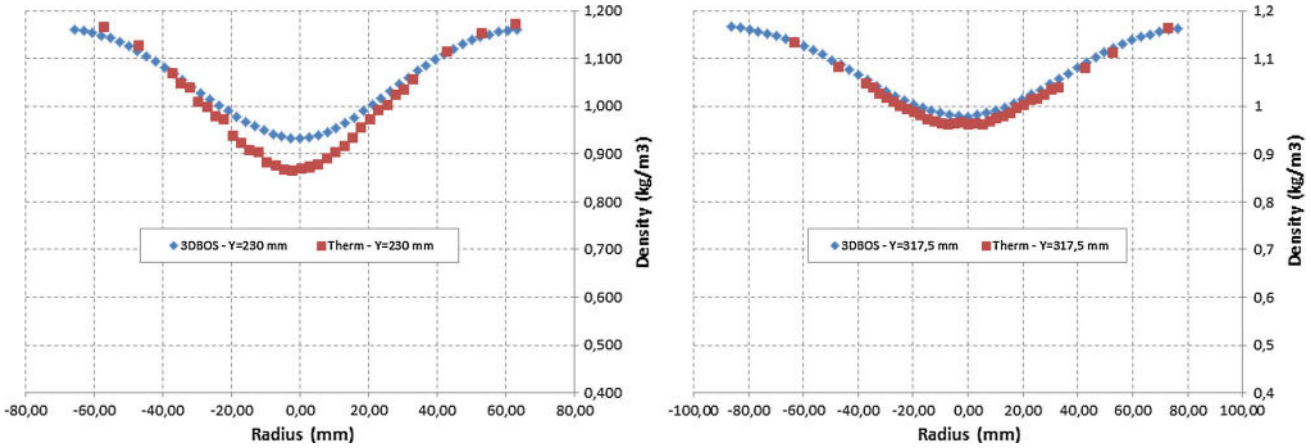

Fig. 22 Comparison between thermocouple measurements and density field reconstructed by 3DBOS for two Y positions from the inlet plan of the hot jet

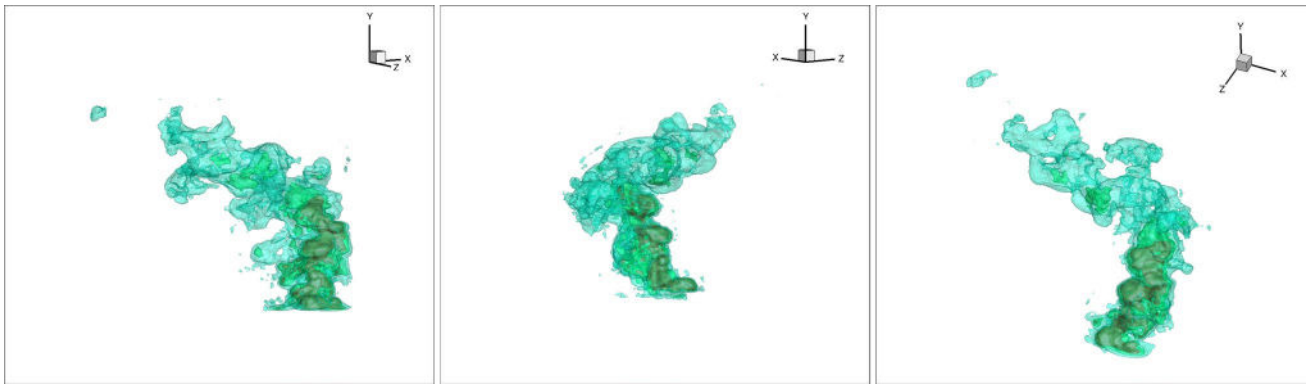

Fig. 23 Reconstruction of an ember helicoidal plume using the "coplanar180" configuration. Regularization parameter $\lambda = 1 \times 10^{-4}$ is chosen according to the L-curve strategy. The reconstructed volume is represented from three different viewpoints using 3 iso-density surfaces 1.155, 1.165, 1.175, respectively, drawn in *red*, *green* and *blue*

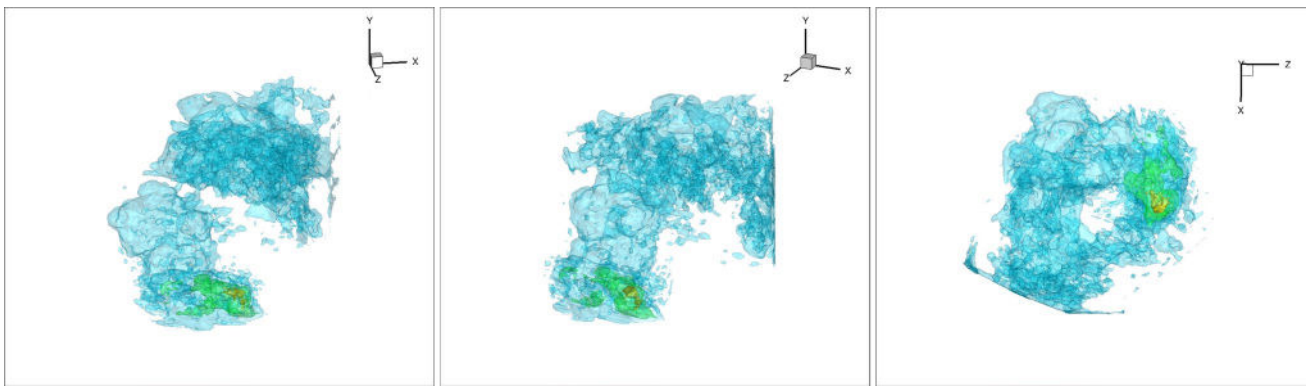

Fig. 24 Reconstruction of a gaz burner helicoidal plume for a "coplanar180" configuration. Regularization parameter $\lambda = 1 \times 10^{-4}$ is chosen according to the L-curve strategy. The reconstructed volume is represented from three different viewpoints using 3 iso-density surfaces 0.97, 1.07, 1.17, respectively, drawn in *red*, *green* and *blue*

In the second case, the inlet flow is larger and more intense. The helicoidal path is also very well reconstructed with larger density gradients: the central hole in the flow structure is apparent, showing again the good sensitivity of the method. Some artifacts associated with the boundary of the common field of views of the cameras appear on the side of the reconstructed domain. Let us emphasize that this example is of very large dimensions, with close to 130 megavoxels reconstructed by casting more than 500 millions of rays every iteration. This huge number of rays explains the longer computational time needed on this particular test case.

7 Conclusion

In this paper, we have presented a new reconstruction method to determine the 3D instantaneous density fields from several BOS images taken simultaneously. This method relies on a regularized framework for the estimation of 3D density volume directly from a limited number of deviation fields. This approach allows to account for prior information on the quantity of interest: the density volume. To solve the resulting huge optimization problem, a CG algorithm implemented on GPU hardware was proposed. This approach is efficient, with volumes of more than 130 megavoxels reconstructed from more than 100 million deviations in a few hours. We believe that further optimization of the software will allow us to reach volumes of one gigavoxel in the near future.

We have realized a bench dedicated to the study of three-dimensional instantaneous flows by 3DBOS, including the case of noncoplanar configurations of cameras. 3D reconstructions of convection flows characterized by weak density gradients and by a transitional or turbulent behavior have been presented. Not only the topology of the flow was correctly captured, but, for very slow flows, consistent estimations of the temperature field at the origin of density variations were obtained.

We intend to conduct further studies about sensitivity and accuracy of 3DBOS for various configurations of the cameras with these software and experimental resources. Another research direction concerns the extension of our approach to flows with high density gradients, so as to eventually address the compressible domain. Preliminary results using an approximated Total Variation regularization to enhance the density discontinuities in the reconstruction were presented in Todoroff et al. (2012, 2014). Going further will require to let aside the paraxial approximation and tackle the nonlinear inversion problem. Finally, ongoing work focuses on the deployment the proposed 3DBOS method in different ONERA wind tunnels, together with other optical measurement techniques such as PIV.

References

Alhaj O, Seume JR (2010) Optical investigation of profile losses in a linear turbine cascade. In: ASME turbo expo 2010: power for land, sea, and air. American Society of Mechanical Engineers, pp 1503–1513

Atcheson B, Ihrke I, Heidrich W, Tevs A, Bradley D, Magnor M, Seidel HP (2008) Time-resolved 3d capture of non-stationary gas flows. In: ACM Transactions on graphics (Proceedings of SIGGRAPH Asia) 27(5):132

Bauknecht A, Ewers B, Wolf C, Leopold F, Yin J, Raffel M (2015) Three-dimensional reconstruction of helicopter blade tip vortices using a multi-camera BOS system. Exp Fluids 56(1):1–13

Bichal A, Thurow B (2014) On the application of background oriented schlieren for wavefront sensing. Measur Sci Technol 25(1):015,001

Cabaleiro Aea (2013) Single camera time-resolved 3d tomographic reconstruction of a pulsed gas jet. J Vis

Champagnat F, Plyer A, Le Besnerais G, Davoust S, Le Sant Y (2011) Fast and accurate PIV computation using highly parallel iterative correlation maximization. Exp Fluids 50:1169–1182

Christian Wolf C, Gardner AD, Ewers B, Raffel M (2014) Starting process of a pulsed jet as seen by schlieren measurements. AIAA J 1–6

Dalziel SB, Hughes GO, Sutherland BR (2000) Whole-field density measurements by synthetic schlieren. Exp Fluids 28(4):322–335

Gac N, Vabre A, Mohammad-Djafari A, Rabanal A, Buyens F, et al. (2010) Gpu implementation of a 3d bayesian ct algorithm and its application on real foam reconstruction. In: The 1st CT meeting proceedings, pp 151–155

Goldhahn E, Seume J (2007) The background oriented schlieren technique: sensitivity, accuracy, resolution and application to a three-dimensional density field. Exp Fluids 43(2–3):241–249

Gross D, Heil U, Schulze R, Schoemer E, Schwanecke U (2009) Gpu-based volume reconstruction from very few arbitrarily aligned X-ray images. J Sci Comput 31(6):4204–4221

Hansen PC (1992) Analysis of discrete ill-posed problems by means of the L-curve. SIAM Rev 34(4):561–580

Idier J (2010) Bayesian approach to Inverse problems, vol 35. Wiley, New York

Ihrke I (2007) Reconstruction and rendering of time-varying natural phenomena. Ph.D. thesis, Max-Planck-Institut fur Informatik

Ihrke I, Magnor M (2004) Image-based tomographic reconstruction of flames. In: Boulic DPR (ed) Eurographics/ACM SIGGRAPH symposium on computer animation

Kak AC, Slaney M (2001) Principles of computerized tomographic imaging. Society for Industrial and Applied Mathematics

Kindler K, Goldhahn E, Leopold F, Raffel M (2007) Recent developments in background-oriented schlieren methods for rotor blade tip vortex measurements. Exp Fluids 43:233–240 (Cite par Sourguen 2012 comme un article BOS axisym par Abel)

Le Sant Y, Todoroff V, Bernard-Brunel A, Le Besnerais G, Micheli F, Donjat D (2014) Multi-camera calibration for 3dbos. In: Application of laser techniques to fluid mechanics. Lisbon

Leopold F, Ota M, Klatt D, Maeno K (2013) Reconstruction of the unsteady supersonic flow around a spike using the colored background oriented schlieren technique. J Flow Control Measur Vis 1(2):69–76

Meier G (2002) Computerized background-oriented schlieren. Exp Fluids 33(1):181–187

Ota M, Hamada K, Kato H, Maeno K (2011) Computed-tomographic density measurement of supersonic flow field by colored-grid background oriented schlieren (CGBOS) technique. Meas Sci Technol 22:104–111

Ota M, Hamada K, Maeno K (2010) Quantitative 3d density measurement of supersonic flow by colored grid background oriented

schleiren (CGBOS) technique. In: 27th international congress aeronautical sciences (ICAS2010)

Pan Y, Whitaker R, Cheryauka A, Ferguson D (2010) Regularized 3d iterative reconstruction on a mobile c-arm CT. In: Proceedings of the first CT meeting, Salt Lake City

Plyer A, Le Besnerais G, Champagnat F (2014) Massively parallel Lucas Kanade optical flow for real-time video processing applications. J Real-Time Image Process 1–18

Raffel M (2015) Background-oriented schlieren (BOS) techniques. Exp Fluids 56(3):1–17

Raffel M, Richard H, Yu Y, Meier G (2000) Background oriented stereoscopic schlieren (BOSS) for full scale helicopter vortex characterization. In: In 9th international symposium on flow visualization. Heriot-Watt University, Edinburgh

Schröder A, Over B, Geisler R, Bulit A, Schwane R, Kompenhans J (2009) Measurements of density fields in micro nozzle plumes in vacuum by using an enhanced tomographic background oriented schlieren (BOS) technique. In: 9th international symposium on measurement technology and intelligent instruments. Saint-Petersburg

Sourgen F, Haertig J, Rey C (2004) Comparison between background oriented schlieren measurements (BOS) and numerical simulations. In: 24th AIAA aerodynamic measurement technology and ground testing conference, Portland OR, USA, paper AIAA, vol 2602

Sourgen F, Leopold F, Klatt D (2012) Reconstruction of the density field using the colored background oriented schlieren technique (CBOS). Opt Lasers Eng 50:29–38

Tikhonov A, Arsenin V (1977) Solutions of ill-posed problems. Winston

Todoroff V, Le Besnerais G, Donjat D, Micheli FAP, Champagnat F (2014) Reconstruction of instantaneous 3d flow density fields by a new direct regularized 3dbos method. In: 17th international symposium on application of laser techniques to fluid mechanics, Lisbon

Todoroff V, Plyer A, Le Besnerais G, Champagnat F, Donjat D, Micheli F, Millan P (2012) 3D reconstruction of th density field of a jet using synthetic BOS images. In: 15th international symposium on flow visualization, Minsk

Venkatakrishnan L (2005) Density measurements in an axisymmetric underexpanded jet by background-oriented schlieren technique. AIAA J 43(7):1574–1579

Venkatakrishnan L, Meier G (2004) Density measurements using the background oriented schlieren technique. Exp Fluids 37(2):237–247

Venkatakrishnan L, Suriyanarayanan P (2009) Density field of supersonic separated flow past an afterbody nozzle using tomographic reconstruction of bos data. Exp Fluids 47:463–473

Vuillot F, Lupoglazoff N, Rahier G (2008) Double stream nozzles flow and noise computations and comparisons to experiments. In: 46th AIAA aerospace sciences meeting and exhibit

Wright S, Nocedal J (1999) Numer Opt, vol 2. Springer, New York

Zeb MF, Ota M, Maeno K (2011) Quantitative measurement of heat flow in natural heat convection using color-stripe background oriented schlieren (CSBOS) method. J JSEM 11(Special Issue), s141–s146

Zhang B, Wu Z, Zhao M (2015) Deflection tomographic reconstructions of a three-dimensional flame structure and temperature distribution of premixed combustion. Appl Opt 54(6):1341–1349