

Assessment of three-dimensional turbulent density measurements from tomographic background-oriented schlieren (BOS)

Shoaib Amjad^{1*}, Julio Soria¹, Callum Atkinson¹

¹Monash University, Laboratory for Turbulence Research in Aerospace and Combustion (LTRAC),
Department of Mechanical and Aerospace Engineering, Melbourne, Australia

*shoaib.amjad@monash.edu

Abstract

The accuracy of three tomographic BOS reconstruction methods are assessed by generating synthetic BOS displacements from the instantaneous density field of a heated jet at Reynolds number $Re = 5000$ and inlet temperature 75K above ambient computed via direct numerical simulation. Synthetic displacements are produced by ray tracing through the jet's refractive index field, obtained using the Gladstone-Dale relation. The displacements form the basis of tomographic reconstruction of the refractive index gradient fields. Reconstruction is performed using filtered back-projection (FBP), iterative algebraic reconstruction technique (ART) and a sequential FBP-ART method. Eighteen views around the flow are used, which mimics a typical experimental setup. FBP features strong reconstruction artefacts even with 18 simultaneous independent views. The reconstruction performs well in the laminar jet core, but accuracy decreases as the flow transitions to turbulence. The sequential FBP-ART uses FBP as the first guess to ART iterations. Unlike previous investigations, improvements were found to marginal, and could degrade overall accuracy in the turbulent region where FBP gives strong artefacts. Using a pre-filter for FBP before ART accounting for the jet spreading downstream is expected to produce a more favourable comparison.

1 Introduction

Density measurements in turbulent heat transfer and mixing provide insight in how to control these complex flows, which appear in a variety of natural and industrial. Schlieren techniques can be used for these measurements, by relating the flow density gradients to the propagation of light rays. This is done using the Gladstone-Dale relation, which connects the flow density ρ to its refractive index n by the Gladstone-Dale constant G which weakly varies with the wavelength of light λ passing through flow,

$$n - 1 = \rho G(\lambda). \quad (1)$$

The background-oriented schlieren (BOS) technique (Raffel et al., 2000), specifically its tomographic implementation (Venkatakrisnan and Meier, 2004; Atcheson et al., 2008), is a variation of schlieren that can provide quantitative instantaneous three-dimensional density measurements. BOS captures path-integrated information of the density gradient field by placing a camera looking through the flow at a background pattern. The deflection of light rays due to changes in wave speed through the flow's inhomogeneous refractive index are detected as displacements in the background. In tomographic BOS (TBOS), multiple cameras are placed around

the flow to simultaneously capture the instantaneous path-integrated displacements. This provides a basis for tomographic reconstruction of the instantaneous turbulent density gradients, from which the density field itself is obtained.

Measurement accuracy depends strongly on the ability to resolve background displacements and the reconstruction of the gradient field. Analytical reconstruction techniques such as filtered back-projection (FBP) can introduce erroneous reconstruction artefacts when only few cameras are used (tens in BOS compared to hundreds in medical imaging). Iterative methods such as algebraic reconstruction technique (ART) cope better with fewer views and allow artefact mitigation using filtering, but this may simultaneously reduce the resolution of small-scale turbulence features. ART inherently requires more computation and multiple solutions can exist. Alternatively, ray deflections can be directly related to an iterated density field solution rather than the gradients (Nicolas et al., 2016; Grauer et al., 2018) but this approach suffers similarly to ART. An assessment of the degree to which these techniques can resolve the multi-scale turbulence is necessary. Studies using FBP (Venkatakrisnan and Meier, 2004; Goldhahn and Seume, 2007) and ART (Atcheson et al., 2008; Lang et al., 2017) show that TBOS provides good agreement of mean temperature (from the ideal gas law on the density field) with thermocouples, and more generally that large-scale features are well-resolved. Small-scale resolution is limited by high spatial frequency noise, or spatial averaging (which is also inherent in the BOS setup). Using an FBP solution as the first guess to ART iterations (Hartmann and Seume, 2016) is a promising development. ART is directed towards a solution with an analytical basis, converging in fewer iterations and FBP artefacts can be filtered out while preserving spatial resolution.

An assessment of a combined FBP-ART method for experimental 3D turbulent density measurements is presented. The density field of a heated free jet from direct numerical simulation (DNS) underpins the assessment. Ray tracing is used in conjunction with Snell's law to produce the exact BOS displacements that are used in tomographic reconstruction. This way, the uncertainties in BOS geometry and displacement calculation methods is separated from the tomographic reconstruction. The performance of FBP, ART and FBP-ART are gauged with downstream length, which is indicative of a greater range of spatial scales to be reconstructed.

2 Background-oriented schlieren (BOS) technique

Briefly, the principles and compromises of background-oriented schlieren are explained. Figure 1 shows a typical BOS setup with a camera focused on a background pattern, with the flow placed an arbitrary distance between the two. A light ray propagating towards the camera will be deflected due to the variable refractive index along its path compared to a uniform refractive index field. Deflection is measured by taking an image of the background without the flow, and again with the flow; comparing the two images provides a two-dimensional displacement field $\Delta\mathbf{X}$.

In gasses, the variation in refractive index is weak enough to approximate the ray path as the result of a single deflection angle ε (assuming the flow's depth along the optical axis is small compared to the camera-background distance Z_B), where Z_D is the object-background distance and f is the focal length. The angle can be expressed in terms of the recorded displacement,

$$\Delta\mathbf{X} = f \left(\frac{Z_D}{Z_B - f} \right) \varepsilon. \quad (2)$$

To increase the signal-tonoise ratio, it is desirable to use longer focal lengths, and place the flow close to the camera. However, the flow is then typically out of focus. A spatial averaging (geometric blurring) δ is introduced (Raffel, 2015), that depends on the aperture size d_a . At the object Z_A ,

$$\delta \approx d_a \left(\frac{Z_D}{Z_B} \right). \quad (3)$$




Figure 1: BOS optical setup and geometric parameters

Reducing δ in the image requires short focal lengths and small apertures. This demands powerful lighting, and emphasises the accuracy of displacement calculation methods, as they become smaller too. BOS displacements can be found using cross-correlation or optical flow methods (Atcheson et al., 2009). Cross-correlation introduces spatial averaging of its own, in interrogation windows; in practice, δ is reduced below the window size to avoid further resolution loss.

The deflection angles for each view ($\epsilon_{x'}$, $\epsilon_{y'}$, $\epsilon_{z'}$) are directly related to the path-integrated refractive index gradients (Eq. 4 with equations in y' and z' having similar form), which provides a basis for tomographic reconstruction of the three-component refractive index gradient field when multiple views are used.

$$\epsilon_{x'} = \frac{1}{n_0} \int \frac{\partial n}{\partial x'} dz' \quad (4)$$

From the reconstructed refractive index gradients, the refractive index field (and hence density field) in a global coordinate system is obtained by solving a Poisson equation (Eq. 5) where the source term q is obtained from the derivative of the reconstructed gradients at each point.

$$\frac{\partial^2 n}{\partial x^2} + \frac{\partial^2 n}{\partial y^2} + \frac{\partial^2 n}{\partial z^2} = q \quad (5)$$

3 Implementing tomographic BOS

Background displacements from each camera must be related to density gradients in global coordinates. First, the displacements are used to find the deflection angles at each camera. Short timescales in turbulent flows mean they must be captured simultaneously. These are then used to reconstruct a global refractive index gradient field.

In filtered back-projection, assuming parallel light rays from each camera (e.g. for a pinhole model), each pixel in each camera contains a projection: the refractive index gradients ∇n integrated along the ray reaching the pixel. This is represented as system of equations describing the gradients in terms of the ray vector \mathbf{x}' (where Δx is the cell size and n_0 is the ambient refractive index),

$$\frac{n_0}{\Delta x} \boldsymbol{\varepsilon}' = \mathbf{x}'^T \sum_{ray} \nabla n \quad (6)$$

For FBP, the cameras are placed in a plane, so each projection is at a certain angle and distance from the object centre. From the Fourier slice theorem, these projections ‘fill in’ the object’s frequency-space representation at that same angle and distance. As a result, higher frequencies are more sparsely populated than lower frequencies. This can also result in erroneous reconstruction artefacts being introduced in regions of high spatial frequency. A filter (typically a ramp filter) is required to emphasise the high-frequency information to increase resolution, but this may enhance the artefacts as well. Also, using more projections allows the object to be described in more detail. The inverse Fourier transform of the frequency space (back-projection) reconstructs the object.

In contrast, algebraic reconstruction technique assigns each point (voxel) in the reconstructed domain a weighting for each light ray. The weighting represents the contribution of that point to an individual ray’s deflection. Thus, each recorded projection is represented as the sum over all the light rays of each volume point and its weighting. From an initial guess, the volume points are updated to minimise the difference in the calculated projection to the recorded projection. Filtering can be used in between iterations to target artefact formation. This approach allows flexibility in geometry and can incorporate prior information into the initial guess. The $(k+1)^{th}$ update for the refractive index gradients causing the deflection of the i^{th} ray through volume point j is obtained using (where L_i is the length of the ray, w_{ij} is the weighting of the j^{th} point to the i^{th} ray and λ is a relaxation factor),

$$\frac{\lambda_i n_0 w_{ij}}{L_i} [\boldsymbol{\varepsilon}' - \boldsymbol{\varepsilon}'^k] = \mathbf{x}'^T (\nabla n_j^{k+1} - \nabla n_j^k). \quad (7)$$

The weighting factor of each point to each ray is calculated according to Atkinson and Soria (2009), by finding the volume intersection of each voxel with a ray. It is approximated as $w_{ij} = \max(0, 1 - \frac{b}{r})$ where b is the distance between the ray and voxel centres and r is the equivalent radius of a voxel.

ART typically starts from a uniform initial guess, in this case zero gradient. Prior information can be used instead to enhance convergence. This study will test both zero gradient and the FBP initial solutions to determine if FBP affords an improvement.

Modifications to ART include: a user-defined mask, which sets the reconstructed gradients to zero beyond a certain radius (as the ambient region around the jet can be clearly distinguished); using gradual unmasking (Liao, 2007) to update the gradients up to a certain threshold each iteration which is progressively relaxed to target artefact formation; a gradually-relaxed Gaussian filter targeting reconstruction artefacts towards the edges; and using a Hamming window on the correction across a ray for a given iteration so that the correction is weighted towards the centre of the volume where the jet is located. In each iteration, the cameras and rays are also selected in a random order to fill in the updated reconstruction, improving convergence (Kak and Slaney, 2001).

Following gradient reconstruction, the Poisson equation is solved using a finite-difference discretisation and an algebraic multigrid approach (Olson and Schroder, 2018). Because the first-order gradients are available directly from the reconstruction, the second-order term of the right-hand side of the Poisson equation contains a smaller truncation error than the left-hand side for the same discretisation scheme. This encourages the use of higher-order discretisation for the right-hand side term than the left and to this end a 4th order central difference is used for the left-hand side and 6th order for the right. For the simulated case presented in section 4, two-dimensional transverse slices will be analysed along the jet’s axis. Only the transverse gradient components are used in the Poisson equation, the longitudinal gradient is typically more than an order of magnitude smaller than the transverse gradients and hence its impact on measurement accuracy is minimal

particularly in comparison to the blurring. The boundaries in these slices are far outside the jet and hence Dirichlet boundary conditions are used specifying the boundaries to have a known refractive index value of $n_0 = 1.000226$. The solution is iterated until residuals reach 10^{-16} order.

4 Synthetic BOS experiment

Simulated BOS displacement fields are created for assessing the reconstruction methods by ray tracing through a known refractive index field to a multiple-camera BOS setup. The refractive index field is obtained by using Eq. 1 on the density field from the DNS of a heated jet (Karami et al., in press). Shown in Fig. 2, the jet is 75K above ambient at exit, with a Reynolds number of 5000 and Mach number 0.5; it has been interpolated to a uniform Cartesian grid with spacing of 0.0204 times the nozzle diameter D for the ray tracing and reconstruction algorithms. The gradients of refractive index are found using 6th-order central differences. Ten downstream transverse planes will be investigated, covering the laminar, transitional and turbulent regions, from $0.2864D$ to $9.2812D$ in increments of approximately one nozzle diameter. The size of each plane is 476×476 grid points, reaching far outside the jet so that Dirichlet boundary conditions are valid in the Poisson equation.

The optical setup for the cameras is chosen to mimic a typical experimental setup while conforming to the compromise between increasing displacements but reducing geometric blurring. To this end, each virtual camera uses a 25mm lens at $f/22$ aperture. Cameras are placed in a y - z plane transverse to the jet axis, in a semicircle. A pinhole model is used to locate the cameras in a global space. The optical centre of each camera is located 275 mm from the jet, and the background 300 mm further behind the jet. The resolution at the object plane is 0.04125 mm/px, based on a standard camera with 3.75 μm pixels. Choosing the jet diameter to be $D = 10$ mm, each grid point thus represents a 5 pixel window while the blur is approximately 15 pixels; blur is simulated by applying a uniform filter to each point with its immediate neighbours.

Raytracing is performed by propagating a ray from its entry to the volume and updating its direction at several steps until it exits the volume. The number of steps along the volume is chosen until the final displacement has converged. Snell's law is used to update the ray's direction at each step, by considering the angle of the refracted ray based upon its entry angle and the refractive index change from the local refractive index gradient at that step. The ray's direction is held constant once it leaves the reconstruction domain, and hence forms the background displacement. Eighteen cameras are used for each slice; this is chosen as it is a practically packageable number of cameras (Amjad et al., 2018). The cameras are equally spaced and located in a 180° arc transverse to the jet axis.




Figure 2: slice through $z/D = 0$ of the DNS heated jet density field. Lines - - show transverse slices to be reconstructed, from $x/D = 0.2864$ to 9.2812 in increments of D .

5 Results and discussion




Figure 3: Error in the Poisson solution of the true transverse gradients. Error in Poisson solution of 6th order central difference gradients with respect to the true refractive index field \blacktriangle ; error in Poisson solution of blurred 6th order central difference gradients with respect to the true refractive index field (zero-reconstruction error case) \blacksquare .

The ultimate accuracy of tomographic BOS is achieved for zero reconstruction error; in this case, the only error is due to spatial averaging from geometric blur and the truncation error associated with discretising the Poisson equation. Figure 3 compares the Poisson solution of the true refractive index gradients (from central difference) to the true refractive index fields themselves with downstream length, both with and without blurring. Errors are presented relative to the difference in refractive index between the centreline and ambient at $x/D = 0.2864$, $\Delta n_p(x = 0.2864) = 3.82 \times 10^{-5}$. The error introduced from the blurring is more prominent than the error from the truncation error and disregarding the longitudinal gradients in the Poisson solution, being nearly double. Geometric blurring is responsible for a high error in regions with a broad range of spatial scales present, as only the larger scales are discerned.

In Figure 4, the error from the FBP and ART reconstructions are shown, after solving the Poisson equation and compared to the Poisson solution of the blurred refractive index gradients. This allows the reconstruction error to be examined in isolation, being an additional error to blurring. FBP introduces the largest error, mainly associated with reconstruction artefacts in regions of high spatial frequency. The number of cameras is insufficient for a high-quality instantaneous reconstruction using FBP alone. Low frequency features appear to be well-captured however, as evidenced by the lower peak error in the laminar region. Artefacts appear towards the edges of the jet, but the core is well-captured. FBP solutions should be filtered at the edges before using ART. To this end, FBP is first multiplied by a Gaussian mask before ART, decreasing artefacts and emphasising the FBP prediction at the domain centre. It is seen that if this is not done, FBP-ART gives a much worse prediction than ART alone, closer to that of using FBP by itself. At present, the same-sized Gaussian mask and reconstruction cut-off mask were used for all downstream slices. The mask size was made to fit around the jet slice furthest downstream, where it occupies most of the domain. Using an adaptive mask that grows with downstream length would undoubtedly produce a closer comparison. Presently, this means that the jet itself is more affected by the mask downstream, rather than just the artefacts.

ART with uniform initial refractive index gradient field converges with increasing iterations; by 80 iterations the error reduces significantly. The present study agrees well with previous results on the reconstruction of synthetic density fields where turbulence is represented by sinusoidal variations with increasing frequency for smaller scales (Atkinson et al., 2017a; Atkinson et al., 2017b). The peak gradients appear to be underpredicted, unlike FBP. The increase in RMS error with downstream length may be attributed to the proliferation of smaller scales, which was seen in synthetic turbulence for high frequencies (Atkinson et al., 2017b). There is little improvement in mean or peak error from 40 to 80 iterations; marginally sacrificing ultimate reconstruction accuracy significantly reduces computing time.




Figure 4: Error in Poisson solution of reconstructed gradients and refractive index field at equivalent spatial resolution to geometric blurring, with downstream distance. FBP only \blacktriangle ; ART 10 iterations \square ; ART 40 iterations \blacksquare ; ART 80 iterations \blacksquare ; FBP-ART 10 iterations \circ ; FBP-ART 40 iterations \circ ; FBP-ART 80 iterations \bullet . ART/FBP-ART cases use reconstruction mask, relaxation $\lambda_i = 0.5$, gradual unmasking gradient threshold 10^{-6} m^{-1} per iteration, inversely iteration weighted Gaussian filter, Hamming window and random camera and ray order. FBP-ART uses a Gaussian-filtered FBP before ART.

For fewer iterations, FBP-ART overall presents a less accurate solution than ART alone, especially in the turbulent regions where artefacts are prominent. Note that the choice and size of the filter, especially with regards to the domain and jet size, strongly affect this result. In the laminar region, the solution is marginally better as FBP captures the top-hat profile very well; this trend is maintained with more iterations. The first few FBP-ART iterations are essentially spent on mitigating the FBP artefacts. With increasing iterations, ART overtakes FBP-ART with marginally higher mean and peak accuracy. FBP itself does not perform well for high-frequency variations, and hence does not represent a better (or much worse) initial guess for ART in the turbulent regions. The strength and location of pre-filtering appears to be very important, it can be seen that for 10 iterations the strong filtering of the FBP solution in the FBP-ART, and the ART by itself, actually produced lower peak accuracy at $x/D = 3.2846$, where the peak gradients were severely affected by the Gaussian filter. Because a constant filter size is used for the given domain, as the jet grows downstream the filter affects not only affects the artefacts but also the jet.

The relaxation parameter may be adjusted to enhance convergence. Figure 5 shows that for a given number of iterations, the solution accuracy is not independent of this parameter. When it is set too low, the solution converges slowly, but if it is too high the solution may diverge. The mean and peak errors appear to be influenced differently by this parameter. The best peak accuracy occurs for $\lambda_i = 0.5$, but the mean error lowers with higher relaxation.

Figure 6 compares the best cases obtained for each reconstruction method for $x/D = 4.2841$. Although FBP is very efficient and captures the shape of jet core well in this region, the reconstruction artefacts are severe. The finer detail in the jet core does not compare well with true solution and fluctuates much more. Further downstream, the core can become indistinguishable from the artefacts as the gradients. Hence, FBP on its own is not appropriate for turbulent density measurements using BOS. When the artefacts are filtered out, the FBP-ART solution is very similar to ART alone. The detail within the jet core is more like the true solution. Interestingly, the ART solution appears to have some slight artefacts just outside the jet core, which is absent in FBP-ART, which may be due to the filtering used in the FBP-ART. Overall, the comparison between the two is very similar. The effect of blurring further removes fine-scale detail within the jet core, for all cases because it is inherent in BOS. The mean error due to reconstruction, however, in the best case is still twice that of blurring alone. With a more appropriate choice of pre-filter, it is expected that FBP-ART could be more favourable than ART alone. If the domain size were to be adjusted to compensate for the larger jet downstream, along with an adaptive pre-filter, FBP-ART could converge faster than ART while maintaining equivalent accuracy.

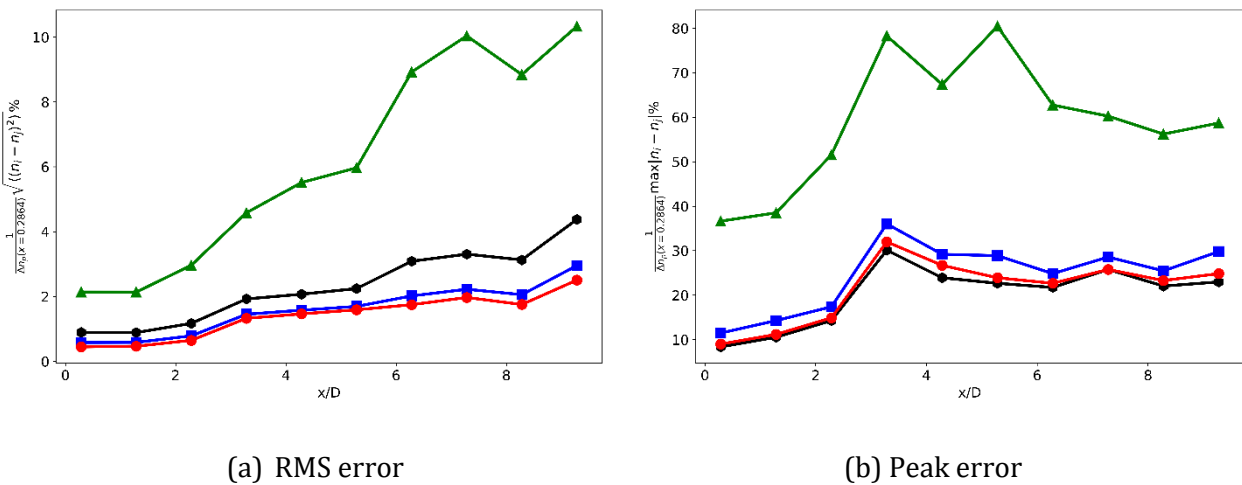


Figure 5: Error in Poisson solution of reconstructed gradients and blurred refractive index field with downstream distance. Relaxation parameter $\lambda_i = 0.1$ \blacktriangle ; $\lambda_i = 0.5$ \bullet ; $\lambda_i = 1.0$ \blacksquare ; $\lambda_i = 1.5$ \bullet . All cases use Gaussian-filtered FBP-ART with 10 iterations, reconstruction mask, gradual unmasking, inversely iteration weighted Gaussian filter, Hamming window and random camera and ray order.

6 Conclusion

Tomographic BOS is a versatile experimental method for three-dimensional density measurements in turbulent jets. The largest sources of uncertainty in this method include the inherent out-of-focus blurring and errors introduced from imperfect tomographic reconstruction.

Filtered back-projection is an efficient reconstruction method but will introduce significant reconstruction artefacts that reduce measurement accuracy, particularly in turbulent flow. These artefacts are associated with having insufficient views of the flow during the reconstruction. Algebraic reconstruction technique copes better with having fewer views and can deliver a more accurate reconstruction. As the number of ART iterations I_s is increased, the solution converges. However, for the simulated jet reconstruction tested, there was only a slight improvement between 40 and 80 iterations. This indicates that a trade-off between computation time and accuracy can be made with minimal impact on measurement accuracy.

Finer scale details are better captured by ART than FBP, and strong reconstruction artefacts are avoided. Using FBP as the initial solution to ART gives a marginal improvement in accuracy in the

laminar regions (where FBP artefacts are not as prominent) but can degrade the solution slightly compared to ART alone in the turbulent regions of the jet. The FBP solution must be filtered first before using it as the initial solution to ART, otherwise the FBP-ART solution is far less favourable. The choice of filter is extremely important, with an appropriate filter that accounts for the jet's growth downstream it is expected that FBP-ART can deliver similar accuracy to ART from uniform initial conditions in fewer iterations. At present, with a uniform FBP pre-filter for the entire jet length, ART from zero initial conditions was only marginally better than the FBP-ART. Modifications were incorporated that target the growth of artefacts at the volume boundaries. The mean error in the most accurate reconstruction case was still twice that of solving the Poisson equation of the blurred gradients directly, which represents a perfect reconstruction.

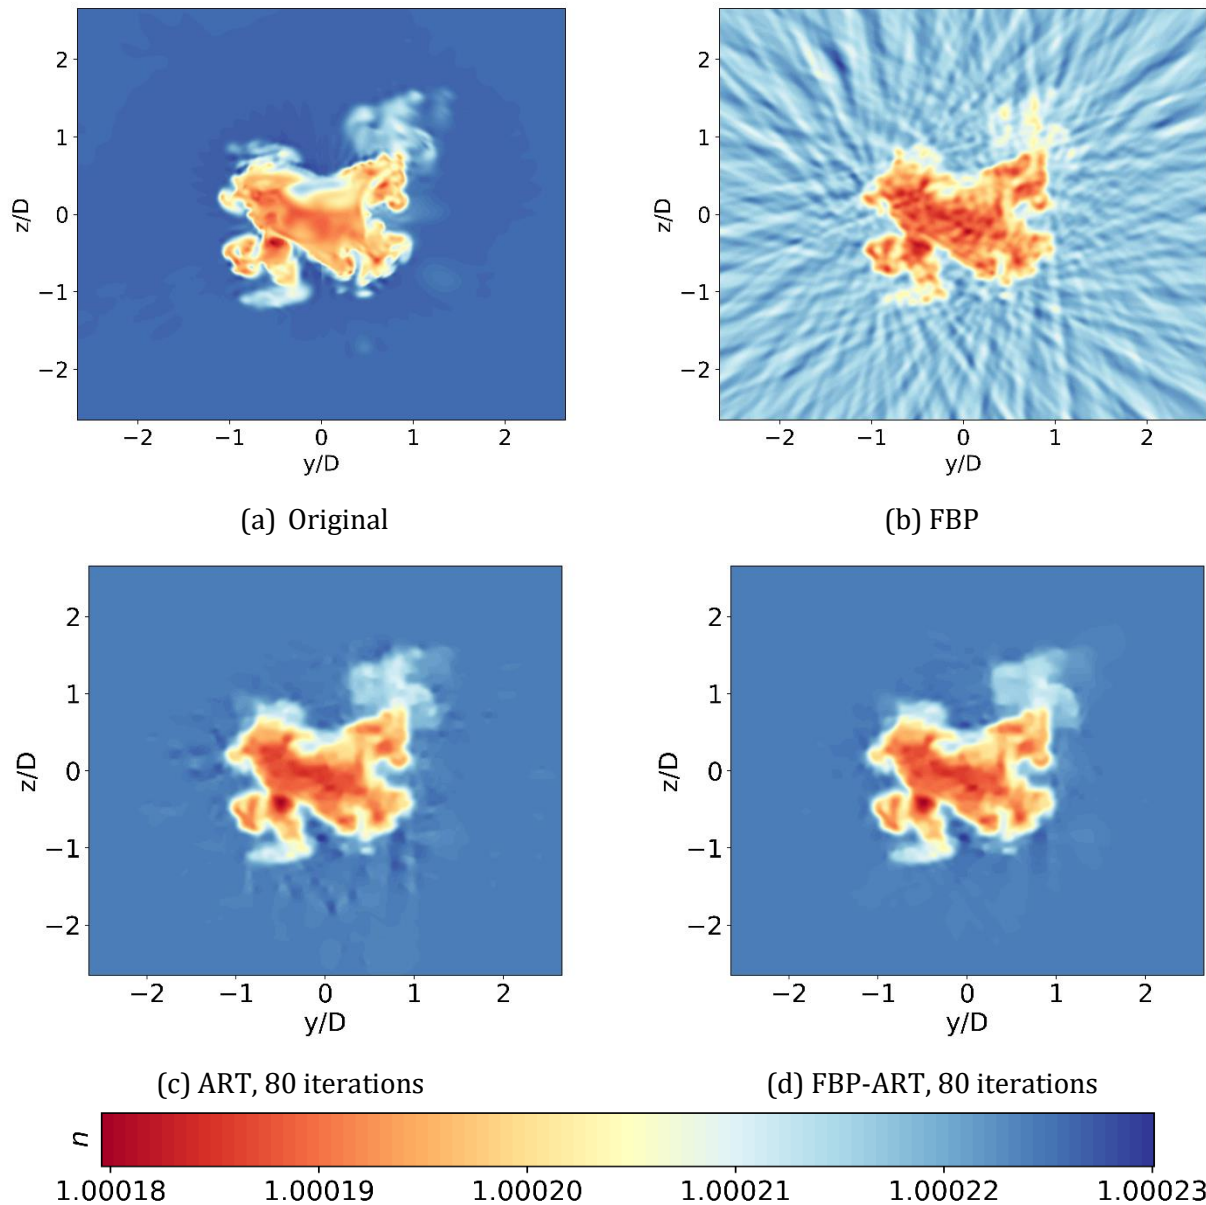


Figure 6: comparison of reconstruction methods with original refractive index field at $x/D = 4.2841$.

Acknowledgements

The authors gratefully acknowledge the support of the Australian Research Council (ARC). Dr Atkinson was supported by an ARC Discovery Early Career Researcher Award (DECRA) while undertaking this work.

References

- Amjad S, Soria J, and Atkinson C (2018) Time-averaged three-dimensional density and temperature field measurement of a turbulent heated jet using background-oriented schlieren. In *21st Australasian Fluid Mechanics Conference, Adelaide, December*
- Atcheson B, Ihrke I, Heidrich W, Tevs A, Bradley D, Magnor M, and Seidel H-P (2008) Time-resolved 3D capture of non-stationary gas flows. *ACM Transactions on Graphics* 27:5
- Atcheson B, Heidrich W, and Ihrke I (2009) An evaluation of optical flow algorithms for background oriented schlieren imaging. *Experiments in Fluids* 46:3
- Atkinson C, and Soria J (2009) An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Experiments in Fluids* 47:4
- Atkinson C, Buchner AJ, Amjad S, and Soria J (2017a) Evaluation and application of tomographic background oriented schlieren and PIV measurements to a heated jet. In *12th International Symposium on Particle Image Velocimetry, Busan, South Korea, July*
- Atkinson C, Amjad S, and Soria J (2017b) A tomographic background-oriented schlieren method for 3D density field measurements in heated jets. In *11th Asia-Pacific Conference on Combustion, Sydney, Australia, December*
- Goldhahn E, and Seume, J (2007) The background-oriented schlieren technique: sensitivity, accuracy, resolution and application to a three-dimensional density field. *Experiments in Fluids* 43:2
- Grauer SJ, Unterberger A, Rittler A, Daun KJ, Kempf AM, and Mohri K (2018) Instantaneous 3D flame imaging by background-oriented schlieren tomography. *Combustion and Flame* 196
- Hartmann U, and Seume JR (2016) Combining ART and FBP for improved fidelity of tomographic BOS. *Measurement Science and Technology* 27:9
- Kak AC, and Slaney M (2001) *Principles of Computerized Tomographic Imaging*. Classics in Applied Mathematics. SIAM
- Lang HM, Oberleithner K, Paschereit CO, and Seiber M (2017) Measurement of the fluctuating temperature field in a heated swirling jet with BOS tomography. *Experiments in Fluids*, 58:7
- Liao, HY (2007) A gradually unmasking method for limited data tomography. In *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Arlington, Virginia, USA*
- Nicolas F, Todoroff V, Plyer A, Le Besnerais G, Donjat G, Micheli F, Champagnat F, Cornic P, and Le Sant Y (2016) A direct approach for instantaneous 3D density field reconstruction from background-oriented schlieren (BOS) measurements. *Experiments in Fluids*, 57:13
- Olson LN, and Schroder JB (2018) PyAMG: Algebraic multigrid solvers in Python v4.0. Release 4.0
- Raffel M, Richard H, and Meier GEA (2000) On the applicability of background oriented optical tomography for large scale aerodynamic investigations. *Experiments in Fluids* 28:5
- Raffel M (2015) Background-oriented schlieren (BOS) techniques. *Experiments in Fluids* 56:3
- Venkatakrishnan L, and Meier GEA (2004) Density measurements using the background-oriented schlieren technique. *Experiments in Fluids* 37:2