

Development of a tomographic background-oriented schlieren method for instantaneous three-dimensional density field measurements of heated jets

Shoaib Amjad

A thesis submitted for the degree of Doctor of Philosophy at Monash University in 2021



Laboratory for Turbulence Research in Aerospace and Combustion Department of Mechanical and Aerospace Engineering Monash University Clayton, Victoria, Australia

Copyright notice

© Shoaib Amjad (2021).

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

Declaration

This thesis is an original work of my research and contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Signature: Print name: Date:

Abstract

A complete description of the evolution of variable density flows requires measurement of both the velocity and density field. At present, there are very few well-developed techniques for studying the density field in variable density jets, let alone the threedimensional structure of turbulence. In this work, we develop the tomographic backgroundoriented schlieren (TBOS) technique for high-quality 3D density field measurements in heated jets with low temporal and spatial blurring. TBOS is a refraction-based method, where density gradients are related to the refractive index through the Gladstone-Dale relation. Based on path-integrated apparent displacements in a background pattern from path variations in light rays travelling through the flow, a tomographic reconstruction of the three-dimensional refractive index gradients can be obtained. A Poisson equation of the reconstructed gradients is solved to obtain the 3D refractive index field, from which the density and temperature can be obtained. This thesis presents the systematic development of a TBOS method.

TBOS experiments suffer from a compromise between measurement sensitivity and defocus blur. The latter obscures the resolution of small scales and is detrimental to turbulence measurements. Temporal blurring must also be considered in the measurements. Furthermore, the tomographic reconstruction of the refractive index gradient field utilises multiple cameras placed around the flow. Reconstruction accuracy depends on the wavelength of the fluctuations, the number and position of cameras around the volume, and the ability of the reconstruction algorithm to cope with both measurement noise and limited camera numbers. These issues are addressed in a three-part investigation.

Using a synthetic density field, the first part of the investigation quantifies the performance of the filtered back-projection (FBP) reconstruction, the iterative algebraic reconstruction technique (ART), simultaneous algebraic reconstruction technique (SART), and a sequential usage of two techniques (FBP+ART), modified with a series of intermediate filtering, windowing and reconstruction damping techniques that are used to improve the reconstruction quality and accuracy. Background displacements, which are used as input to the reconstruction algorithms, are produced by ray tracing through the synthetic refractive index field. It is shown that all techniques under-resolve fluctuations with wavelengths less than 4 voxels. FBP introduces strong reconstruction artefacts due to an inadequate number of projections in a typical BOS setup. An optimised ART is developed, which is the most accurate of all techniques tested and converges in 100 iterations. The quality of FBP+ART is sensitive to appropriate mask selection, the absence of which can lead to a degradation in reconstruction quality and a proliferation of reconstruction artefacts. The optimal FBP+ART is marginally less accurate than the optimal ART due to under-prediction of peak gradients, but it converges in only 20 iterations. When the density field is obtained from the reconstructed gradients, a gradient integration scheme that attenuates higher frequencies is recommended. The high-frequency information from the reconstruction is unusable and likely to propagate noise unnecessarily. In present cases, a Poisson equation is solved to obtain the density field, and it was found that the filtering characteristics of a second-order central difference of both the volume and gradients resulted in the optimal balance between truncation error and propagation of reconstruction noise.

In the second part of the investigation, the density field of a heated jet obtained via direct numerical simulation is used to validate a proposed 15-camera experimental setup. The simulation contains the near-field laminar-to-turbulent transition of a jet with a Reynolds number based on nozzle diameter D of $Re_D = 10,000$ and exit-to-ambient density ratio $\rho_e/\rho_{\infty} = 0.8$. This case is ideal for further testing of the reconstruction algorithms on a realistic flow. Temporal and defocus blurring are found to be detrimental to the measurement quality and are the dominant sources of error in a typical experiment. Defocus blur at the measurement object δ should be limited to $\delta/D \leq 11$ % to preserve smaller scales, based on the correlation between the blurred and true fields. The temporal blurring is controlled by the exposure time of the BOS images, and should not exceed 10% of the flow's characteristic timescale t_c . If both sources of blurring are controlled, very detailed and accurate reconstructions of the flow can be obtained using the optimised ART.

The last part of the work concerns experimental TBOS measurements. To address the temporal blurring issue, a pulsed laser-speckle TBOS technique is devised. A high-power, short-pulse laser is used for both illumination and creation of the background pattern in a 15-camera setup using a laser-speckle pattern. The high-power laser ensures adequate exposures in nanoseconds, so temporal integration is minimised in each measurement. A novel method of identifying the ideal compromise between measurement sensitivity, defocus blur and speckle size through considered selection of the focal length, aperture and focussing distance is presented. Reconstructions of the near-field of a heated jet show excellent resolution of 3D flow structures, and good agreement with thermocouple measurements. A comparison is made with DNS measurements to show the influence of Mach number and boundary conditions on the measurement. A demonstration of TBOS for use in turbulence modelling measurements is presented, which shows that the peak scalar dissipation in the near-nozzle measurement domain occurs in the shear layer associated with vortex roll-up. An analysis of the 3D density potential core structure reveals that the potential core undergoes stretching and fragmentation during the turbulence transition.

Publications and awards from this work

Journal publications

 S. Amjad, S. Karami, J. Soria, and C. Atkinson. Assessment of three-dimensional density measurements from tomographic background-oriented schlieren (BOS). *Measurement Science and Technology*, 31(11):114002, 2020.

Publications based on chapters 4 and 6 of this thesis are forthcoming.

Conference proceedings

- 1. S. Amjad, J. Soria, and C. Atkinson. Three-dimensional measurement of density and temperature in a turbulent heated jet. In 11th Australasian Heat and Mass Transfer Conference, Melbourne, 2018.
- 2. S. Amjad, J. Soria, and C. Atkinson. 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, 2018.
- S. Amjad, S. Karami, J. Soria, and C. Atkinson. Assessment of three-dimensional turbulent density measurements from tomographic background-oriented schlieren. In 13th International Symposium on Particle Image Velocimetry (ISPIV 2019), Munich, 2019.
- 4. S. Amjad, J. Soria, and C. Atkinson. Instantaneous three-dimensional density measurements using tomographic background-oriented schlieren. In *(9th Australian Conference on Laser Diagnostics, Adelaide, 2019.)*
- S. Amjad, S. Karami, J. Soria, and C. Atkinson. Assessment of three-dimensional turbulent density measurements from tomographic background-oriented schlieren (part 2). In 73rd Annual Meeting of the American Physical Society's Division of Fluid Dynamics, 2020.
- 6. S. Amjad, J. Soria, and C. Atkinson. Three-dimensional density measurements of a heated jet using laser-speckle tomographic background-oriented schlieren. In 14th International Symposium on Particle Image Velocimetry (ISPIV 2021), 2021.

Awards and achievements

- Invited paper Assessment of three-dimensional turbulent density measurements from tomographic background-oriented schlieren (BOS) for special issue of Measurement Science and Technology journal, Special Section on the 13th International Symposium on Particle Image Velocimetry (ISPIV 2019).
- Best Student Presentation at the 9th Australian Conference on Laser Diagnostics, 2-4 December 2019. Adelaide, Australia.

Acknowledgements

Professional acknowledgments

The support of the Australian Research Council (ARC) for this work through a discovery grant is gratefully acknowledged. This research was supported by an Australian Government Research Training Program (RTP) Scholarship.

The research benefited from computational resources provided through the National Computational Merit Allocation Scheme (NCMAS), supported by the ARC. The computational facilities supporting this project included the Australian NCI Facility, the partner share of the NCI facility provided by Monash University through an ARC LIEF grant, Pawsey Supercomputing Centre and the Multi-modal Australian ScienceS Imaging and Visualisation Environment (MASSIVE) HPC facility (www.massive.org.au).

Manufacturing and construction services, design advice, and electronics support for the experimental rig was provided by the Monash University Mechanical and Aerospace Engineering Technical Services Group (MAE-TSG).

Laboratory facilities and resources belong to the Laboratory for Turbulence Research in Aerospace and Combustion (LTRAC) in the Department of Mechanical and Aerospace Engineering, Faculty of Engineering, Monash University (Clayton, Victoria, Australia).

Personal acknowledgments

It has been quite a journey. The focus is on acquiring new knowledge and trying to prove yourself, but it pushes you and causes you to grow in so many ways. I don't think I am the same person now as I was at the start of 2018, hopefully for the better. I have had the best people around me to guide, support and nurture me. This is my attempt to thank everyone, but I know that words cannot do justice to how much I am grateful to them.

This may be the longest acknowledgments section you've ever seen in your life, by far. No one can claim to have completed a journey all on their own while being entirely self-sufficient. I certainly don't. I apologise if I missed anyone. Without further ado...

Firstly, I must heartily thank my supervisors, Dr Callum Atkinson and Prof Julio Soria. I have learnt so much from them in the past few years, and they have been excellent role models for me. It was a great pleasure to work with, and learn from, the both of them. Dr Atkinson created the initial version of the tomographic reconstruction code that was further developed and tested in this work, obtained initial results in chapter 4, and drafted portions of the text in chapter 4. Both Dr Atkinson and Prof Soria contributed to the research direction, and proofreading and revising the thesis. Sometimes I felt like young Luke being tutored by Masters Obi-Wan Kenobi and Yoda. Wise teachers, and excellent people, who have left a lasting impression on me. My co-author Dr Shahram Karami deserves an enormous amount of thanks for his hard work and patience. I thank him for running the heated jet DNS used in chapter 5, and for helping me with, and offering me advice and insights on, DNS and the behaviour of jets. Dr Karami drafted a portion of the text on the DNS code in chapter 5.2, and contributed to proofreading and revising the publication, Amjad et al. [4], upon which chapter 5 is based.

I greatly appreciate the time and feedback given to me by my milestone panel members, Prof Richard Manasseh, Dr Isaac Pinar, and Prof Murray Rudman. It was a pleasure to share my work with you.

Without the MAE workshop, my experimental setup would only be drawings on paper. They helped bring my vision to life, and provided me with excellent advice on design, fabrication, and electronics. The fruits of this labour are presented in chapter 6. My thanks go to: Santhosh Babu, Paul Campbell, Nat Derose, Keith Erbs, Chris Pierson, Andrew Smith, Mark Symonds and Hugh Venables. Hugh was the first workshop staff member I met when I was a bright-eved undergraduate, and he patiently answered my many questions and provided helpful feedback. Nat has been great mentor to me, teaching me so much about design and manufacturing. I am always eager to hear his great stories, and I appreciate the time he has given me that has undoubtedly made me a better engineer. Mark is the electronics guru who patiently works with mechanical engineers, knowing that we haven't the foggiest about what is required, but always comes up with a solution that amazes us. Chris has been so helpful in organising everything I needed to construct my experiment, and I really admire his hard work. Santhosh manufactured and constructed my experimental rig. I have enjoyed working with him, and I appreciate the effort he put in to making sure everything was perfect. The workshop folks are probably the busiest and hardest-working people I know, and I think they do not get even half as much recognition as they deserve.

Dr Daniel Edgington-Mitchell started me on this path many years ago. I witnessed his charismatic teaching as an undergraduate, which made me want to know more about his work. His encouragement and advice introduced me to research and helped me find a passion.

My PhD pals are a treasure. My friends were great company on this journey. If I tried to list them all, it might be longer than the rest of my thesis. But I will try! I couldn't have met a nicer and more caring bunch of folks. The humour, camaraderie, and kind-heartedness we share make the journey worthwhile. Dr Shevarjun Senthil and I have become the best of friends through this journey. His encouragement and advice means so much to me, and it is difficult to put my immense gratitude to him into words. I got to know Harry Scott and Daniel Jovic well through our tutoring duties, and they are stand up guys who I had a lot of fun working with. Mohamed Tolba taught me a lot about the 'bigger picture' in life, and I really appreciate his friendship and wisdom. Sanjay Mohan has been a good friend throughout the journey, and he is one of the hardestworking people I know. Ahmed Mahil is one of the most vibrant people I have had the pleasure of meeting. I would also like to thank my friends Asif Ahmed, Dariush Ashtiani, Dat Bui, Michael Eisfelder, Rhiannon Kirby, AJ Kusangaya, Keith Lai, Sean Lawrence, Chi Nguyen, Muhammad Shehzad, Bihai Sun and Dominic Tan. My school friends also deserve credit for encouraging me to get this far.

Dr Samuel Grauer is a friend and well-wisher that I have met along the way. Sam encouraged me along the journey, and helped me understand the practical aspects of the BOS experiment. Dr Jim Kostas is someone I met relatively late in my journey, and he has been so kind and has provided me with wise career advice.

Many thanks to my thesis examiners, Prof Markus Raffel and Dr Gerrit Elsinga. Their glowing comments were much appreciated, and their thorough reading of the thesis and suggestions were much appreciated.

My schoolteachers worked so hard so that I could succeed. I really, sincerely acknowledge and appreciate all of their efforts. I apologise if I was ever difficult, recalcitrant, or mischievous. It will be difficult for me to name everyone. But I want to at least name some of my high school teachers which set me on the path to studying engineering: Adele Booth, Ron Cerbasi, Cathy Craig, Victoria Heffernan, Christine McCann, Carol Mulraney, and Jon Neall. I find it funny that I didn't think I would ever have to use the photography concepts I learnt with Ms Booth in my life, yet it basically forms the basis of the experimental methodology in this thesis.

Finally, thanks to my family, both immediate and extended. My family also reinforced my faith, and my religion has kept me strong, persistent, and grateful. My parents and my sister have put in so much effort to help me and to make sure I got this far. My father, Amjad Hafizullah, and my mother, Rubina Amjad, deserve more thanks than words could ever do justice. My sister, Fareha Amjad, is my hero. She is hard-working, smart, creative, loyal, good and generous. I wouldn't be anywhere without my family.

Thank you. I couldn't have done it without any of you.

Alhamdulillah

Contents

Li	st of	tables	i
\mathbf{Li}	st of	figures	xiii
N	omen	nclature	xiv
1	Intr	oduction	1
	1.1	Aims and overview of the thesis	2
2	Bac	kground	4
	$2.1 \\ 2.2$	A description of variable density jets	4
	2.3	tration measurements in variable density jets	$\begin{array}{c} 11 \\ 13 \end{array}$
3	The	e tomographic background-oriented schlieren (TBOS) technique	21
	3.1	Principles and overview of the TBOS measurement process	21
	3.2	Background design and displacement calculation methods	26
		3.2.1 Background design	27
		3.2.2 Displacement calculation methods	28
	3.3	Tomographic reconstruction of the refractive index field	31
		3.3.1 Filtered back-projection (FBP)	33
		3.3.2 Algebraic reconstruction technique (ART)	34
		3.3.3 Other iterative reconstruction techniques for BOS	36
		3.3.4 Integration of the reconstructed gradients	37
4	A p	arametric study of TBOS methods in a fluctuating density field	39
	4.1	Aims and overview of the chapter	39
	4.2	Implementation of a tomographic BOS reconstruction	40
		4.2.1 Filtered back-projection (FBP)	40
		4.2.2 Algebraic reconstruction technique (ART)	41
		4.2.3 Calculation of the refractive index field	42
	4.3	Numerical validation procedure	43
	4.4	A synthetic fluctuating density field test case: the heated jet phantom	44
		4.4.1 Convergence of the ray tracing method	47
	4.5	Results and discussion	48
		4.5.1 Finite difference schemes to solve the Poisson equation	48
		4.5.2 FBP reconstruction	52
		4.5.3 ART reconstruction schemes and enhancements	54

	4.6	4.5.4 4.5.5 4.5.6 Summ	Hybrid FBP+ART reconstructions	61 63 69 71
5	Nur	nerical	l validation of a proposed experimental setup and assessmen	ıt
	of t	hree-d	imensional TBOS density measurements of a heated jet sim	1-
	ulat	ion		74
	5.1	Aims a	and overview of the chapter	74
	5.2	Detail	s of the heated jet direct numerical simulation (DNS)	75
	5.3	Valida	tion of proposed experimental setup	75
	5.4	Result	s and discussion	77
		5.4.1	Impact of defocus blurring and other spatial averaging on TBOS	
		549	measurements	77
		5.4.2	Impact of temporal blurring on 1 BOS measurements	80
		0.4.3	Comparison of FBP, ARI and FBP+ARI reconstructions of the	00
		511	Effect of displacement field poise on APT reconstruction	02 85
	55	5.4.4 Summ	ary and conclusions	86
	0.0	Summ		00
6	\mathbf{Exp}	erime	ntal three-dimensional density field measurements of a heate	d
	jet 1	using l	aser-speckle TBOS	88
	6.1	Aims a	and overview of the chapter	88
	6.2	Laser-	speckle BOS	88
	6.3	Exper	imental setup and method	90
		6.3.1	Camera and laser configuration and control	90
		6.3.2	Heated jet	94
	6.4	A syst	ematic selection methodology for camera lens focal length, aperture	
		and fo	cussing distance	96
	6.5	Result	s and discussion	101
		6.5.1	Comparison of mean temperature field with thermocouple measure-	105
		6 F 0	ments	105
		0.5.2	Comparison of density field statistics with heated jet DNS	107
		0.5.3	Insights on the scalar variance transport equation for temperature	110
	66	0.3.4 Summ	Potential core behaviour in the heated jet	111
	0.0	Summ		114
7	Sun	nmary	and conclusions	117
Α	Mat	terial r	related to chapter 4	120
	A.1	Transf	fer function of the discretised 1D Poisson equation	120
	A.2	Sensit	ivity study of ART inversely iteration-weighted Gaussian filtering	
		standa	rd deviation	120
	A.3	Sensit	ivity study of ART progressively tightened Gaussian mask	121
	A.4	Sensit	ivity study of FBP+ART progressively tightened Gaussian mask pa-	
		ramete	ers	123
	A.5	Conve	rgence of ART schemes compared to FBP+ART	124

в	Mat	erial related to chapter 5	126
	B.1	Definition of the blur kernel	126
	B.2	Convergence of the ray tracing method	127
	B.3	Anisotropic diffusion modification to the Poisson solver	128
\mathbf{C}	Mat	erial related to chapter 6	131
	C.1	Additional experimental setup photographs	131
	C.2	Matching non-dimensional jet parameters to experimental conditions	132
	C.3	Additional experimental density field visualisations	134
Bibliography 13-			134

List of Tables

- 4.1 Central finite-difference schemes used to solve the Poisson equation. In each scheme, each dimension has the same order of accuracy. Left-hand side is abbreviated as LHS, and the right-hand side is abbreviated as RHS. The 'No. points' column refers to the number of points considered by the finite-difference equation in each dimension. The entries in the 'Abbreviation' column will be used to refer to these schemes henceforth.
- 4.2Influence of ART reconstruction techniques for 16 camera reconstruction and $\lambda_{x,z} = L/14$ after 100 iterations with relaxation parameter $\lambda_i = 1.0$. RMS errors $\sqrt{\langle (..., -..., s)^2 \rangle}$ and peak errors max |..., -..., s| are given for the reconstructed gradient fields ∇n and Poisson-solved reconstructed refractive index fields n, where subscripts \dots_r and \dots_s denote the reconstructed and true synthetic fields, respectively. $R_{\dots r \dots s}$ denotes the correlation coefficient between the reconstructed and synthetic fields. Errors and correlation coefficients are calculated within a radius that is twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$. The Poisson equation is solved with the (3, 5) discretisation scheme. Gradients are presented both normalised by the peak difference between the centreline and outer flow refractive index Δn_n divided by domain size L, and as a percentage of the peak gradient in the spatial average synthetic field ∇n_{peak} . Errors in the Poisson solved refractive index field n are given as a percentage of Δn_p . Case K is selected for further investigation (bold text, red ticks).
- Combinations of camera lens focal length, aperture and focussing distance 6.1(independent variables) for laser speckle BOS evaluated by sensitivity to displacements, blur at the object and average speckle size (dependent variables, bold). For all options $Z_A = 500$ mm; camera pixel size is $l_{pix} =$ 3.45 µm. Focus distances are chosen such that blur $\delta/D \approx 11\%$ for all options. The closest cross-correlation window size for focal length f = 25mm is 16 pixels, while for f = 50 mm it is 32 pixels. Text colour indicates: red is unfavourable (speckles too small/large for current pixel size); orange is sub-optimal (field of view too small); black is selected. 99 6.2100A.1 One-dimensional Poisson equation finite difference (FD) schemes and cor-

49

List of Figures

2.1	Evolution of a submerged free jet with downstream length x (not to scale), adapted from Abdel-Rahman [1] and Ball et al. [11]. Flow is from left to right. The jet is initially laminar, emerging from a contoured nozzle with exit diameter D . The growth of instabilities transitions the flow to a turbulent state, encouraging mixing with the ambient fluid. Pertinent flow features are labelled in the top half of the diagram. Axial regions are labelled in the bottom half. Badial regions are labelled on the right.	5
2.2	Cutaway normalised density field of a heated jet from direct numerical sim- ulation, with Reynolds number based on nozzle diameter $Re_D = 10,000$, exit Mach number $Ma = 0.6$ and exit density ratio $\rho_e/\rho_{\infty} = 0.8$. Flow is	0
	from left to right. Details of the simulation are given in chapter 5	6
2.3	A non-exhaustive family tree of density, temperature and concentration measurement techniques in variable density gas flows. Shading indicates	1 1
<u>ົ</u> ງ /	The deflection of the requirements in this project.	11
$\frac{2.4}{2.5}$	Experimental setup for collimated light shadowgraphy Light can be colli-	14
2.0	mated from a point source using lenses or a parabolic mirror.	15
2.6	Experimental setup for density gradient methods: a) Töpler schlieren, b) moiré deflectometry, c) speckle photography, d) structured light refractog-	
	raphy, e) background-oriented schlieren (BOS)	17
2.7	Experimental setup for digital holographic interferometry. The holograms can be recorded without imaging optics, e.g. camera lens	20
3.1	Schematic of BOS ray deflection and nomenclature. In traditional BOS, the camera (composed of image plane, lens, and aperture) is focussed on the background $(l = +Z_D)$. In the laser speckle BOS implemented in chapter 6, the camera may be focussed anywhere in the range $-(Z_A + f) < l < +\infty$. The global coordinate system origin is located at the centre of the refractive volume. The optical axis of camera 1 is oriented in the global \dot{C}	0.0
3.2	<i>z</i> -axis	22
	instead has a diameter of d_i at the image plane. \ldots \ldots \ldots \ldots	24

ii

3.3	Process for density measurements using TBOS. After the refractive index field $n(x, y, z, t)$ has been reconstructed, the density field can be obtained using the Gladstone-Dale relation. Numbers correspond to the steps listed	
	on page 25	26
3.4	Types of backgrounds used in previous studies: a) random dots, b) hori-	
3.5	zontal lines, c) wavelet noise, d) laser speckles, e) checkerboard. \ldots A 16 × 16 pixel cross-correlation interrogation window of a laser speckle background pattern with a heated jet. Axes coordinates are in pixels. a) Reference image interrogation window, b) displaced image interrogation window, c) correlation plane (zero displacement origin marked with yellow cross). A clear, white peak is visible in the correlation plane (marked with red cross), as well as measurement noise, indicating a displacement of 1.35	27
3.6	pixels to the right and 1.67 pixels downwards (magnitude 2.15 pixels) Principle of tomographic reconstruction. A ray passing through the field of interest g generates a projection on the sensor plane P_{θ} . The rays need not be perpendicular to the sensor plane. Adapted from Kak and Slaney	30
3.7	[62]	32
3.8	to lower frequencies is evident as the increasing distance between adjacent data points from different projections	33 34
4.1	Possible TBOS investigations (shaded) using ray-tracing of the synthetic and DNS heated jet fields (True solution). This work focuses on investi- gating spatial averaging (chapter 5), temporal averaging (chapter 5) and reconstruction methods and Poisson equation (chapters 4 and	
4.2	5), shown with darker shading	44
4.3	are shown for simplicity; up to 22 cameras positioned in a 180° arc are tested. Convergence of displacement, as a function of sub-grid steps along the ray,	45
	at pixel with highest displacement for a camera oriented at $\theta = 45^{\circ}$ to the volume with $\lambda_{x,z} = L/14$. The volume has 65 grid points in each direction.	47

4.4 RMS error between the synthetic refractive index fields and the refractive index fields calculated by the Poisson solver, n_s and n_r , respectively, within twice the jet half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$, as a function of fluctuation wavelength $L/\lambda_{x,z}$ and additive Gaussian noise level $\sigma_{noise}/\nabla n_{peak}$. The error is calculated within the jet core for varying kernel sizes in the discretisation used in the left-hand side (multigrid) and right-hand side (calculating Laplacian from the reconstructed refractive index gradient field), respectively: a) 3 and 3 points, b) 3 and 5 points, c) 5 and 3 points, d) 5 and 5 points. Each data point is averaged over 100 samples of added random noise.

50

51

52

53

- 4.6 Bode magnitude plot of the analytical (A) and finite-difference Poisson equation transfer functions H as a function of spatial frequency ω . The finite-difference schemes are denoted by the number of left- and right-hand side points used in the discretisation kernel as per table 4.1, e.g. (3, 5) for 3 points on the left-hand side and 5 points on the right-hand side. The Nyquist frequency is $\omega_x = 0.5 \text{ voxel}^{-1}$, or $\lambda_{x,z} \approx L/32...$
- 4.7 Contour plots, and profiles through x = 0 (black dotted line is the original synthetic field and red dashed line is the reconstruction), for the reconstructed refractive index gradient $\partial n/\partial x$ for 16 cameras and $\lambda_{x,z} = L/8$: a) synthetic field, b) FBP. Bottom row is $\lambda_{x,z} = L/14$: c) synthetic field, d) FBP.
- 4.8 Contour maps of the RMS (a) and peak (b) errors within twice the halfwidth $r \leq 2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ between the synthetic and FBP reconstructed refractive index fields n_s and n_r , respectively, as a function of wavelength and camera number, normalised by the peak change in the synthetic refractive index field from the Poisson solution. The minimum RMS and peak errors are 0.3% and 0.7%, respectively. The maximum RMS and peak errors are 35% and 97%, respectively.

- 4.11 RMS error in the refractive index gradients ∇n (a) and the refractive index fields n (b) for 16 camera reconstruction and $\lambda_{x,z} = L/14$ in the region twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$, as a function of ART iterations for: ART $\lambda_j = 0.2 \circ$; ART $\lambda_j = 0.5 \Box$; ART $\lambda_j = 1.0 \triangle$; ART $\lambda_j = 4.0 \diamond$; SART $\lambda_j = 1.0 \blacktriangle$; SART $\lambda_j = 4.0 \diamond$. In all cases the Poisson equation is solved using 3- and 5-point kernels for the left- and right-hand side calculation, respectively. Reconstructions use case K in table 4.2. . .
- 4.12 Contour maps (top row) of the RMS (a) and peak (b) errors between the synthetic and optimised ART reconstructed refractive index fields, n_s and n_r , respectively, with 100 iterations within twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ as a function of wavelength $L/\lambda_{x,z}$ and camera number $n_{cameras}$, normalised by the peak change in the synthetic refractive index field from the Poisson solution n_{peak} . The minimum RMS and peak errors for ART are 0.6% and 1.5%, respectively. The maximum RMS and peak errors are 3.0% and 11.2%, respectively. Colourbar is consistent with figure 4.8 for comparison.
- 4.13 RMS error (top row) and peak error (bottom row) in the reconstructed refractive index gradients ∇n (left column) and the refractive index fields n (right column) for 16 camera reconstruction and $\lambda_{x,z} = L/14$ in the region twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$, as a function of the number of ART iterations for the cases shown in the table below. In all cases the Poisson equation is solved using 3- and 5-point kernels for left-and right-hand side discretisation, respectively. All ART and FBP+ART reconstructions are performed using randomly-ordered cameras and pixels with Hamming windowed corrections and relaxation $\lambda_j = 0.5$. ART and FBP+ART use a sharp cut-off mask with $r_{mask} = 30$ voxels ($3.3\sigma, 2.8r_{1/2}$). Progressively tightened Gaussian mask decreases from $r_m = 35$ voxels ($3.9\sigma, 3.3r_{1/2}$) to $r_{m,final} = 30$ voxels. ART case A is the same as case K from table 4.2 with $\lambda_j = 0.5$.
- 4.14 Contour plots, and profiles through x = 0 (black dotted line is the original synthetic field and red dashed line is the reconstruction), for the reconstructed refractive index gradient $\partial n/\partial x$ for 16 camera reconstruction and $\lambda_{x,z} = L/14$: a) synthetic field; b) FBP and 100 ART iterations with gradual unmasked and Hamming windowed correction (case C in the table of figure 4.13); c) FBP multiplied by progressively tightened Gaussian mask and 100 ART iterations with Hamming windowed corrections (case E), and d) FBP multiplied by progressively tightened Gaussian mask and 100 ART iterations with inversely iteration-weighted Gaussian filter and Hamming windowed corrections (case G).

60

58

62

- 4.15 Average absolute error $|..._r ..._s|_{bin}$ in the reconstructed refractive index gradients ∇n normalised by ∇n_{peak} (a), and the refractive index fields nnormalised by n_{peak} (b) for 16 camera reconstruction and $\lambda_{x,z} = L/14$, as a function of normalised radial position r/σ with a bin size of $\sigma/4$. Shown are: FBP •, ART case A from figure 4.13 ×, and FBP+ART case E from figure 4.13 \triangle (the latter two both correspond to the same markers in figure 4.13). Vertical dotted line indicates the usual $2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ limit that the RMS error is calculated within. Sharp cut-off mask radius is located at $r_{mask} = 3.3\sigma$. Initial width of the progressively tightened Gaussian mask is $r_m = 3.9\sigma$, and the final width is $r_{m,final} = 3.3\sigma$
- 4.16 Power spectral density of the analytical $\partial n/\partial x|_{x=0,z}$ for different imposed frequency fluctuations ω , from L/2 to L/32, in increments of L/2 ($\omega = 0.03 \text{ voxel}^{-1}$), left to right. Each frequency (individual lines) produces a distinct, sharp peak.

66

66

67

- 4.17 Peak power (left column), and relative peak power (right column), of power spectral density of $\partial n/\partial x|_{x=0,z}$ for different imposed frequency fluctuations ω , from L/2 to L/32, in increments of L/2 ($\omega = 0.03 \text{ voxel}^{-1}$). L/32 corresponds to the Nyquist frequency $\omega = 0.5 \text{ voxel}^{-1}$. Top row corresponds to A = 0.125, while the bottom row corresponds to the usual A = 0.25. Red horizontal dashed line indicates 50% relative power criterion. Shown are peak powers of the spectra for: the analytical fields ∇ , FBP •, selected ART case \times , and selected FBP+ART case \triangle (the latter two both correspond to the same markers in figure 4.13).
- 4.18 Contour plots, and profiles through x = 0 (black dotted line is the original synthetic field and red dashed line is the reconstruction), for the reconstructed Poisson solved refractive index $n_0 - n$ for 16 camera reconstruction and $\lambda_{x,z} = L/14$: a) synthetic field; b) FBP; c) selected ART 100 iterations (case A in figure 4.13); d) selected FBP+ART 100 iterations (case E in figure 4.13).

4.20	Average absolute errors in the refractive index gradients ∇n (a) and re- fractive index fields n (b) as a function of radial position r/σ for noise level $\sigma_{noise}/\Delta X_{peak} = 5\%$ for 16 camera reconstruction and $\lambda_{x,z} = L/14$. Shown are: FBP •, selected ART case ×, and selected FBP+ART case \triangle (the latter two both correspond to the same markers in figure 4.13). Faint lines indicate the error with no added noise (figure 4.15). The Poisson solution uses 3- and 5-point kernels for the left- and right-hand sides, re- spectively. Each point is averaged over 100 samples; error bars indicate the 95% confidence level and are approximately the size of the markers. The red vertical line indicates the $2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ limit up to which the RMS error is calculated	71
	placements with $\sigma_{noise}/\Delta A_{peak} = 5\%$ added holse for: FBP (top row); selected AR1 (middle row); selected FBP+ART (bottom row).	72
5.1	Slice through $z/D = 0$ of the DNS heated jet density ratio field ρ/ρ_{∞} at one snapshot. Flow is from left to right. Dotted lines : show evenly spaced	
5.2 5.3	transverse slices for reconstruction from $x/D = 0.28$ to 9.28 The proposed 15-camera experimental setup	76 76
5.4	fields.	78
5.5	as previous figure) as a function of blurring δ/D and axial distance x/D . RMS (a) and peak (b) errors between the Poisson solution of the transverse gradients and the true DNS density field as a function of axial distance x/D. Errors are normalised by the peak refractive index difference near the nozzle $\Delta n_p(x/D = 0.28)$ and averaged over the 104 samples. Marker • represents the Poisson solution of the true gradients with no blurring and (3, 5) discretisation. Markers \bigstar and \blacksquare represent the Poisson solution of the gradients with blurring $\delta/D = 6\%$ for grid resolutions of $d_{an} = \delta/15$	79
5.6	and $\delta/5$, respectively. Error bars are approximately the size of the markers and indicate a 95% confidence level	80 80

5.7RMS (a) and peak (b) errors between the Poisson solution of the reconstructed gradients and the Poisson solution of the blurred true gradients as a function of axial distance x/D. Errors are normalised by the peak refractive index difference near the nozzle $\Delta n_p(x/D=0.28)$ and averaged over the 104 samples. Markers correspond to: FBP ★; ART 10 iterations from null initial conditions \Box ; ART 40 iterations \Box ; ART 100 iterations \blacksquare ; FBP+ART 10 iterations O; FBP+ART 40 iterations O; FBP+ART 100 iterations •. All optimised ART/FBP+ART cases use the sharp cut-off mask, relaxation parameter $\lambda_i = 0.5$, inversely iteration-weighted Gaussian filter, Hamming windowed corrections, random camera and ray order and Gaussian mask. FBP+ART initial FBP solution is filtered with the Gaussian mask. Error bars are approximately the size of the markers and indicate a 95% confidence level. 83 Absolute error as a function of radius (normalised by nozzle radius R) and 5.8axial length (normalised by nozzle diameter D) for a) FBP, and b) ART 100 iterations, both averaged over the 104 samples. Markers \bullet show the mean jet radius \bar{r} calculated based on displacement magnitude threshold $\left| \vec{\Delta X} \right| < 0.1$ pixels. Markers + show the calculated sharp cut-off mask radius r_m . Radial bin size is equal to 0.4R. 83 Cross-section contours of normalised 'excess' density for one snapshot at 5.9x/D = 5.2, and corresponding centreline profiles at z/D = 0, for: a) Poisson solution of the blurred gradients; b) FBP; c) optimised ART with 100 iterations; d) optimised FBP+ART with 100 iterations. 84 5.10 RMS (a) and peak (b) errors for the ART reconstruction with 40 iterations, averaged over the 104 samples, as a function of axial distance x/D. Markers \bigstar show the ideal case with no displacement field noise (identical to figure 5.7); markers \bullet show reconstruction accuracy with random uniform noise added to the displacements in the range ± 0.1 pixels. Error bars are approximately the size of the markers and indicate a 95% confidence 86 5.11 Contours of normalised 'excess' density for one snapshot (on the same colour scale as figure 5.9) at x/D = 5.2, and corresponding centreline profiles at z/D = 0, for: a) the Poisson solution of the reconstructed ideal blurred gradients (no displacement field noise), and b) Poisson solution of the reconstructed gradients from noisy displacements. Dotted line ... indicates Poisson solution of blurred gradients (figure 5.9a). 87 Configurations for laser-speckle BOS introduced by Meier and Roesgen 6.1[81]. a) 'Single-pass' mode, where the expanded laser illuminates a surface without passing through the refractive volume (flow) under study. b) 'Double-pass' mode, where the illumination is introduced in-line with the camera's optical axis. Notice that in both of these modes, the camera is not necessarily focussed on the speckle surface $(l \neq Z_D)$ like in traditional BOS (cf. figure 3.1). c) 'Interferometry mode', like 'double-pass' mode except that the camera is focussed on the refractive volume (l = 0). . . . 91

6.2	a) Experimental setup with 15 cameras modified for laser-speckle TBOS,	
	b) schematic of laser expansion. Optical axis of camera 1 is aligned with	
	the global z-axis, and x is the jet axis. The laser beam is guided into the	
	beam splitter at the correct orientation using an articulated arm (ILA 5150	
	Articulated Mirror Arm, not pictured). Bottom: photograph of expanded	
	laser beams illuminating the backgrounds. The cylinder placed above the	
	nozzle is used to check that the beams do not directly cross the recon-	
	struction volume In the real experiment the laser optics and mirror arm	
	are not directly attached to the rig's 'table top' upper surface as shown	
	in (a) and (b) but rather are magnetically secured to the underside of a	
	small bonch that fits onto the 'table top' and over the top of cameras. The	
	banch is visible in the bettern right corner of the bettern photo. Additional	
	apported photographs given in Appendix C 1	02
63	A severe example of shot to shot variations in intensity in the speckle im	92
0.3	A severe example of shot-to-shot variations in intensity in the speckie ini-	
	ages. No now is present in either image (ambient environment). The	0.9
C 4	O l'hatin the speckles do not seem to be anected, only their intensity.	93
0.4	Calibration target, designed by Himpel et al. [56], imaged from three ad-	
	jacent cameras. The inverted 1-snaped structure is used to determine the	
	orientation of each camera relative to camera 1. The cameras are mounted	
	sideways to align the longer edge of the sensor with the streamwise direc-	
	tion, so the 1-shaped structure appears to (subtly) move up and down in	0.4
0 F	the images of adjacent cameras, not left and right.	94
0.5	Cutaway schematic of converging nozzle with matched-cubic profile [93].	
	Dimensions are in millimetres. The matched cubic contour itself is 50 mm	0 5
0.0	in length from entrance to exit.	95
6.6	Thermocouple temperature traces of nozzle centreline exit temperature	
	(x/D = 0.3) as a function of time for a) warmup and b) cooldown (after	
	one hour of stable heating), when the air mass flow rate is 1.7×10^{-6} kg s ⁻¹	
	and the heater voltage is 120 V. Samples are recorded at intervals of $dt = 2$	0.0
0.7	seconds	96
0.7	I hermal image of the insulated nozzle when the jet is at an exit tempera-	0.0
0.0	ture of 91°C. Image captured using Seek Thermal Compact camera	96
6.8	a) Sensitivity $\Delta/\tan\varepsilon$ (mm), and b) blur at object in pixels d_i/l_{pix} , as a	
	function of focussing distance l (mm) for $Z_A = 500$ mm. Legend for b):	
	solid line $-f = 25$ mm at $f/8$; dotted line $\cdots f = 25$ mm at $f/10$; dash-dot	00
<i>c</i> 0	line $f = 50 \text{ mm at } f/10$	98
0.9	Recorded laser speckles images in a 100 \times 100-pixel area for apertures: a)	
	J/8, b) $J/11$ and c) $J/10$. Red squares show a 10 × 10-pixel area. Physical	
	pixel size is $3.45 \times 3.45 \mu\text{m/pixel}$. Brightness and contrast of images b) and	
	c) have been enhanced for clarity. Bottom: speckle patterns recorded by	00
C 10	each camera at $f/11$ superimposed on the experimental setup (illustration).	99
6.10	a) Horizontal displacements for camera 1, and b) vertical displacements	
	for camera 1 using the median of multiple reference images for one time-	
	step (physical pixel size is $3.45 \times 3.45 \mu\text{m/pixel}$, and spatial resolution at	
	the focus plane is 140 µm/pixel). Flow is from bottom to top. Bottom	
	image: measured displacement magnitude (0 to 1.5 pixels) captured by	
	each camera at a given instant superimposed on the experimental setup	102
	(Illustration $)$	102

- 6.11 Reconstructed refractive index gradient magnitude $|\nabla n|$ at one time-step. Longitudinal slices (flow is from bottom to top): a) y/D = 0; b) z/D = 0. Transverse cross-section slices (on the same colour bar as the above longitudinal slices): c) x/D = 0.3; d) x/D = 1.3; e) x/D = 1.8; f) x/D = 2.6; g) x/D = 3.4; h) x/D = 4.3.

103

6.19 Histogram illustrating frequency of samples as a function of number of potential core fragments observed using the $(\rho_{\infty} - \rho)/(\rho_{\infty} - \rho_e) > 0.9$ criterion. a) All potential core fragments, b) fragments with a volume $V \ge 0.05\overline{V}_{pc}$.

114

- 6.20 Box-and-whisker plots (first and third rows) and scatter plots (second and fourth rows) of the potential core fragments' volumes (normalised by D^3) and the corresponding centres of mass as a function of position, respectively. Only potential core structures/fragments with $(\rho_{\infty} - \rho)/(\rho_{\infty} - \rho_e) >$ 0.9 with a volume $V \ge 0.05\overline{V}_{pc}$ are considered. a) Fields with 1 structure, b) fields with two structures, c) fields with three structures, d) fields with four structures. Box-and-whisker plots: orange line – is the median, box represents the interquartile range IQR (Q_1 to Q_3), whiskers extend to $1.5 \times$ IQR beyond Q_1 and Q_3 , black circles \circ are outliers, red dashed line - - is the volume of the mean potential core \overline{V}_{pc} . Scatter plots: largest structure (main potential core) \circ , second-largest structure \Box , third-largest structure \triangle , fourth-largest structure \diamond ; red cross \times shows the centre of mass of the mean potential core (translated slightly off-axis for clarity).
- A.1 a) Average absolute error of the ART gradient field reconstruction as a function of normalised radial position r/σ with a bin size of $\sigma/4$. The lightness of the lines (purple to yellow in colour, dark grey to light grey in greyscale) shows σ_{GF} increasing from 0.1 voxels to 2 voxels in increments of 0.1 voxels ($0.011 \leq \sigma_{GF}/\sigma \leq 0.22$ in increments of $0.011\sigma_{GF}/\sigma$) and from 2 voxels to 5 voxels in increments of 1 voxel ($0.22 < \sigma_{GF}/\sigma \leq 0.56$ in increments of $0.11\sigma_{GF}/\sigma$). b) Normalised RMS error between the reconstructed gradients and synthetic field as a function of σ_{GF}/σ . Results are presented for a 16-camera reconstruction and $\lambda_{x,z} = L/14$. ART uses the same modifications as case A (selected ART) in figure 4.13. Vertical dotted line indicates the $2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ limit for the RMS error. . . .

- A.3 a) Average absolute error of the masked FBP gradient field reconstruction as a function of normalised radial position r/σ with a bin size of $\sigma/4$; the lightness of the lines shows $r_{m,initial}$ increasing from 0.5σ to 5.5σ in increments of 0.5σ (purple to yellow in colour, dark grey to light grey in greyscale), and the black dashed line is the FBP solution from figure 4.15. b) Normalised RMS error between the masked FBP reconstructed gradients and synthetic field as a function of $r_{m,initial}$ normalised by σ . Results are presented for a 16-camera reconstruction and $\lambda_{x,z} = L/14$. Vertical dotted line indicates the $2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ limit for the RMS error. 123
- A.4 a) Average absolute error of the FBP+ART gradient field reconstruction as a function of normalised radial position r/σ with a bin size of σ/4; the lightness of the lines shows r_{m,final} increasing from 0 to 3.5σ in increments of 0.5σ (purple to yellow in colour, dark grey to light grey in greyscale).
 b) Normalised RMS error between the FBP+ART reconstructed gradients and synthetic field as a function of r_{m,final} normalised by σ. Results are presented for a 16-camera reconstruction and λ_{x,z} = L/14. FBP+ART uses the same modifications as case E (selected FBP+ART) in figure 4.13 with 100 iterations. Initial mask size r_m is fixed at r_m = 35 voxels = 3.9σ. Vertical dotted line indicates the 2r_{1/2} = 2σ√2 ln 2 limit for the RMS error. 124
- A.5 RMS error (top row) and peak error (bottom row) in the reconstructed refractive index gradients ∇n (left column) and the refractive index fields n (right column) for 16 camera reconstruction and $\lambda_{x,z} = L/14$ in the region twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$, as a function of the number of ART iterations for the cases shown in the table below. In all cases the Poisson equation is solved using 3- and 5-point kernels for left-and right-hand side discretisation, respectively. All ART reconstructions are performed using randomly-ordered cameras and pixels with Hamming windowed corrections and relaxation $\lambda_j = 0.5$. ART uses a sharp cutoff mask with $r_{mask} = 30$ voxels $(3.3\sigma, 2.8r_{1/2})$. Progressively tightened Gaussian mask decreases from $r_m = 35$ voxels $(3.9\sigma, 3.3r_{1/2})$ to $r_{m,final} = 30$ voxels. ART case A is the same as case K from table 4.2 but with $\lambda_j = 0.5$.
- B.1 Illustration of blur from one camera at a point in the reconstruction volume.127
- B.2 Convergence of the calculated displacements (in pixels) as a function of number of steps taken through the slice, relative to 5,000 steps, for a camera at x/D = 9.28. The slice contains 462 grid points in each direction. 127
- B.3 RMS (left) and peak (right) errors for the ART reconstruction with 40 iterations in the noise-free (top) and noisy cases (bottom) averaged over the 104 samples, as a function of anisotropic diffusion parameter α . Noisy samples are the same as those used in section 5.4.4. Markers correspond to: x/D = 0.28, x/D = 5.28, and x/D = 9.28. Error bars are approximately the size of the markers and indicate a 95% confidence level. 129

C.1	Multiple-camera heated jet rig used for laser speckle TBOS experiment.	
	In normal use, the entire rig is barricaded with black boards, as seen at	
	the back of shot, to prevent stray laser light from injuring users. Not	
	pictured: electronic control system for camera triggering, computer, and	
	heater voltage control are located to the left: laser is located to the right:	
	compressed air input, inline air heater, settling chamber located beneath	
	rig	131
C.2	The settling chamber and nozzle are shrouded in thermal insulation (par-	
	tially unwrapped to show the layers). They are mounted to a height-	
	adjustable platform which is positioned using linear bearings on guide	
	rails (one on each corner of the platform). The platform is moved us-	
	ing a hydraulic jack located below (not pictured). The compressed air line	
	and inline air heater connect to the settling chamber from below the plat-	
	form. The temperature at the settling chamber inlet is measured using a	
	thermocouple	132
C_{3}	Close up of lens groups showing how they are magnetically attached to the	102
$\bigcirc.0$	underside of the banch a) Lager arm outlet plate hearm calittee and long	
	underside of the bench. a) Laser and outlet, plate beam spitter and lens	199
a 4	group 1. b) High-energy mirror and lens group 2.	133
C.4	Perpendicular cross-sections through the density field at one instant: a)	
	x - y plane at $z = 0$, b) $x - z$ plane at $y = 0$	134
C.5	Transparent contour view of density field at one instant	135
C.6	Cross-section of density field in $x - y$ plane at $z = 0$ captured in eight	
	successive instances. The measurement is not time-resolved.	135

Nomenclature

- * Convolution
- α Anisotropic diffusion tuning parameter
- $\alpha,\,\beta,\,\gamma\,$ Rotation of camera about global $z,\,y$ and x axes, respectively
- χ Scalar dissipation rate
- $\Delta \varphi$ Phase difference
- ΔI Change in intensity
- Δn_p Peak refractive index difference
- Δx Grid spacing/voxel width
- $\Delta Z'$ Distance from refractive volume to recording plane
- δ Defocus blur size at object plane
- δ_{IW} Interrogation window size projected to object plane
- \dot{m} Mass flow rate
- γ Specific heat ratio
- $\hat{\delta}$ Dirac delta function
- κ Thermal diffusivity
- Λ Eigenvalue array
- λ Wavelength of light
- $\lambda_1, \lambda_2, \lambda_3$ Eigenvalues
- λ_j Relaxation factor
- λ_r Regularisation parameter
- $\lambda_{x,z}$ Wavelength of density field phantom
- \mathbb{R} Real number space
- \mathcal{F} Fourier transform

\mathcal{W}	Hamming window function
tr ()	Trace
μ	Viscosity
∇n	Spatial refractive index gradients
∇	Gradient operator
ω	Spatial frequency
\otimes	Outer product operation
$\overline{\overline{D}}$	Anisotropic diffusion tensor
$\overline{\overline{J_{\sigma}}}$	Gaussian-filtered structure tensor
_	Ensemble average
\overline{V}	Eigenvector matrix
\overline{V}_{pc}	Average potential core volume
Т	Transpose of vector/matrix
k	Value at k^{th} iteration
$_{0},$	$_\infty$ Value at ambient conditions
c	Value at jet centreline
e	Value at nozzle exit
recon	Reconstructed field
Φ_{ν}	Viscous dissipation rate
П	Probability function
ψ	Active scalar
ρ	Density
σ	Standard deviation of Gaussian function
σ_{GF}	Standard deviation of Gaussian filter kernel
σ_{noise}	Standard deviation of added Gaussian noise
au	Scalar dissipation timescale
θ	Angle of projection in $x - z$ plane
$\tilde{\epsilon}$	Dissipation rate for DNS
$ ilde\eta_k$	Kolmogorov scale for DNS

XV

- $\tilde{\nu}$ DNS dynamic viscosity
- $\varepsilon_{x'}$ Deflection angle of light ray in x' direction
- $\vec{\Delta X}, \Delta$ Displacement of light ray on recording plane, i.e. image displacement
- $\vec{\eta}$ Correlation space position vector
- $\vec{\hat{i}}$ Incident ray vector
- $\vec{\hat{n}}$ Unit vector of ∇n
- \vec{t} Refracted ray vector
- \vec{x} Position vector of point on light ray
- \vec{c} Optical centre vector
- \vec{D} Finite difference operator
- \vec{K} Intrinsics matrix
- \vec{P}_o Reference background point coordinates
- \vec{P}_r Refracted background point coordinates
- \vec{R} Rotation matrix
- \vec{r} Refracted ray vector
- \vec{v} Velocity
- \vec{X}, \vec{X}_o, X, Y Image plane coordinates
 - Favre average
- ζ Position along light ray
- " Favre fluctuation
- ' Reynolds fluctuation
- A Amplitude
- *a* Height of image
- A_n Nozzle exit area
- *b* Shortest distance between ray and voxel centre
- C Concentration
- c Speed of sound
- c_p Specific heat at constant pressure
- D Nozzle exit diameter

- D_{ψ} Diffusion coefficient for active scalar
- d_{Σ} Overall geometric blur
- d_a Aperture diameter
- D_C Diffusion coefficient for concentration field
- d_d Diffraction-limited spot size
- d_i Defocus blur size at image plane
- d_s Average speckle diameter
- d_{gp} DNS equivalent grid point size
- d_{IW} Interrogation window size at image plane
- $d_p i$ Particle size in cross-correlation images
- dx Unequal DNS grid spacing
- f Focal length of lens
- Fr Froude number
- G Gladstone-Dale constant
- g Function to be reconstructed
- I Intensity
- k Iteration number, or tuning parameter (appendix B.3 only)
- k Thermal conductivity
- K_{σ} Gaussian filter kernel
- L Domain length
- *l* Focus distance
- L_i Length of i^{th} ray
- L_x DNS domain length in streamwise direction
- l_{px} Pixel size
- Ma Mach number
- N Lens f-number, i.e. f/N
- N Number of grid points
- *n* Refractive index
- N_{it} Total number of iterations

 N_{sm} Number of identified particles in cross-correlation interrogation window O_X, O_Y Image centre in X and Y directions, respectively PPath-integrated projection PPressure Q Frequency filtered Fourier transform of path-integrated projection Poisson equation right-hand side source term qRadiative heat transfer q_{rad} RCorrelation coefficient, or nozzle radius, or ideal gas constant Radial coordinate, or equivalent radius of a voxel (chapter 4 only) r $r_m, r_{m,initial}, r_{m,final}$ Mask radius, initial and final Half-width $r_{1/2}$ Reyolds number ReRiRichardson number SFourier transform of path-integrated projection sPath direction along light ray S_1, S_2 Intensity signal for cross-correlation displacement calculation method S_{ψ} Source term for active scalar S_C Source term for concentration field TTemperature t Time, or ray distance in tomography (chapter 3.3 only) Jet characteristic time t_c Jet through-time t_j Gradual unmasking threshold t_o Exposure time t_{exp} UAverage velocity VVolume Weighting of i^{th} ray to j^{th} voxel w_{ij} x, r, θ Cylindrical coordinate system, x is the jet's stream-wise coordinate

x,y,z~ Global Cartesian coordinate system; z-axis oriented along camera 1 line of sight, x-axis aligned with jet axis

x',y',z' Local Cartesian coordinate system of a light ray; z'-axis is oriented ray path

- x'_r A direction orthogonal to ray path z', i.e. x' or y'
- Z_A Lens-to-object distance
- Z_B Lens-to-background distance
- Z_D Object-to-background distance
- Z_I Image plane-to-lens distance
- (T)BOS (Tomographic) background-oriented schlieren
- ART Algebraic reconstruction technique
- DNS Direct numerical simulation
- FBP Filtered back-projection technique
- LES Large-eddy simulation
- PIV Particle-image velocimetry

Chapter 1 Introduction

Turbulent, variable density jets are fluid flows that arise in a wide variety of natural and industrial processes, such as convective cooling, combustion, pollutant dispersion and spray drying. Jets consist of high-momentum fluid issuing into an ambient environment from a nozzle or orifice. Variable density effects influence the development and characteristics of jets, in which fluids of different chemical composition are mixed, or are subject to thermal or compressibility effects. Due to the impossibly high computational demand of direct numerical simulation of flows of practical interest, the governing equations must be modelled to predict and control these flows. A complete description of variable density flows requires modelling of the turbulent fluctuations of the density and velocity fields and their interactions. The models must be informed by, and evaluated against, high-quality experimental measurements of the density and velocity fields.

Velocity field measurement techniques have seen rapid development over the past few decades from point measurements using hot-wire anemometry and doppler velocimetry to two- and three-dimensional velocity field measurements (including time-resolved capability) using particle-image velocimetry (PIV), for a wide variety of flows. Density measurement techniques have not enjoyed the same rapid advancement. It remains uncommon to see a three-dimensional quantification of the density field, let alone the existence of a standard, well-developed technique for these measurements. Optical methods present a promising avenue to conduct non-intrusive measurements of the density field in turbulent flows, which take advantage of the relation between a fluid's density field and its refractive index field. Background-oriented schlieren (BOS) is one such technique based on this relation, which utilises a camera to image a textured background pattern while looking through the flow of interest [112]. When imaged from multiple perspectives, the apparent distortions of the background images due to refractive index gradients can be used as the basis for a tomographic reconstruction of the 3D refractive index field, from which the density and temperature fields may be obtained.

The accuracy of tomographic BOS (TBOS) measurements depends strongly on both the ability to resolve background displacements and the reconstruction of the gradient field. For the former, a prominent dilemma in the BOS experimental setup is the compromise between measurement sensitivity and defocus blur. The sensitivity to background pattern displacements is increased by increasing the distance between the flow and the background, and by using longer focal length lenses. However, this also increases the defocus blurring in the measured object, which obscures the flow features. Further compromises between image illumination and temporal blurring arise when one attempts to solve the original dilemma. These issues have not been satisfactorily addressed in previous work, and a key objective of this research is to develop a methodology to find the ideal compromise between measurement sensitivity and defocus blurring while mitigating temporal blurring. Regarding the tomographic reconstruction, there are special considerations in fluid mechanics measurements such as TBOS and TPIV that are not found in other fields such as medical imaging. For example, the short timescales of the flow requires that all of the projections are recorded simultaneously. This requires that each projection correspond to an individual camera, limiting the number of projections that can be used due to physical packaging constraints. Ultimately, the limited number of views results in an imperfect reconstruction which will under-resolve high spatial frequency flow features and produce undesirable reconstruction artefacts. Furthermore, unique to TBOS, the reconstructed quantity is a gradient that can be positive or negative, unlike intensity-based reconstructions. The current work presents the development of an optimised reconstruction algorithm for TBOS suitable for limited-view tomography. The resulting findings are used to develop an experimental setup capable of high-quality 3D density measurements of heated jets with low temporal and defocus blurring.

1.1 Aims and overview of the thesis

The progression of the current work is divided into three sequential parts, and the aims of each part are listed below.

- 1. Development of TBOS reconstruction methods using a heated jet phantom. Aims:
 - (a) Optimise the reconstruction methods to improve accuracy and measurement quality.
 - (b) Determine the appropriate number of cameras for high-quality reconstruction.
 - (c) Identify the range of spatial scales resolvable by the reconstruction methods relative to the Nyquist frequency of the optical system, as a function of camera number.
- 2. Validation using a realistic flow: the heated jet direct numerical simulation (DNS). Aims:
 - (a) Validation of a proposed 15-camera experimental TBOS setup.
 - (b) Identify the limit on temporal blurring due to exposure time for high-quality measurements relative to the jet's characteristic time scale.
 - (c) Identify the limit on defocus blurring for high-quality measurements relative to the jet's nozzle diameter as a function of camera aperture.
- 3. Demonstration of an improved experimental technique: pulsed laser-speckle TBOS. Aims:
 - (a) Implement pulsed laser-speckle TBOS as a novel method to mitigate temporal blurring.
 - (b) Develop a systematic procedure to enable optical setup optimisation identifying the ideal compromise between measurement sensitivity and defocus blurring.

(c) Demonstration of 3D measurements of the potential core in a heated jet to identify transitional behaviour.

In chapter 2, the background to the present work is discussed. The characteristics, complexity and open questions on heated jets are introduced. The principles and qualities of different measurement techniques for density, temperature, and concentration in variable density jets are compared. The current work focuses on refraction-based density measurements, which can produce the desired 3D density field measurements. A density gradient method known as BOS is selected for further development. BOS is chosen because of its simple and versatile experimental setup, which can be expanded to the necessary multiple-camera setup for 3D density field measurements.

Chapter 3 discusses BOS in greater detail. The principles of the technique are outlined, including the compromise between measurement sensitivity and defocus blur. The background design and displacement calculation methods are presented. Finally, the tomographic reconstruction of 3D density fields from BOS image displacements are considered, including past approaches.

Part 1 of the research is presented in chapter 4, which is focussed on developing and validating four tomographic reconstruction methods using a synthetic density field phantom (test case). A ray tracing procedure is developed to create synthetic BOS displacements from the phantom to use as input to the reconstruction. This allows a systematic assessment of the accuracy and measurement quality of four reconstruction methods. Modifications are introduced to enhance convergence, improve the prediction of gradients, and reduce reconstruction artefacts associated with limited-view tomography. The reconstruction error and range of resolvable spatial scales are quantified.

Part 2 is presented in chapter 5. Having established the optimal reconstruction method, the reconstruction method is tested using the data from DNS of a heated jet with the virtual version of a proposed 15-camera experimental setup. The study gauges the accuracy of the technique on a realistic flow to validate the proposed setup. Furthermore, the study establish limits on the allowable defocus blur and temporal blurring in TBOS measurements for high-quality 3D flow measurements. Because this approach starts by considering the impact of temporal and defocus blurring relative to the flow structures, the present work demonstrates a significant step forward in addressing the compromise between measurement sensitivity and defocus blurring in TBOS.

The last part of the research is presented in chapter 6. In part 3, the guidelines established in part 2 are used to develop the 15-camera experimental setup for 3D density measurements in a heated jet. This chapter demonstrates the implementation of a novel pulsed laser-speckle tomographic background-oriented schlieren method. A laser-speckle pattern is used as the background for all cameras. This modification mitigates temporal blurring in the measurements and furthers the use of TBOS to the study of turbulent flow structures and their statistics. A systematic procedure is developed to control defocus blurring in the measurement and maximise measurement sensitivity. This creates an extremely useful set of guidelines to optimise future TBOS experiments. Finally, measurements of the near-nozzle region of a heated jet are presented, including validation against thermocouple measurements and insights on the breakup of the potential core of a heated based on 3D experimental data.

Chapter 7 concludes the work with a summary of the key findings. Remarks on the potential applications and limitations of the technique are presented, and recommendations for future work are made.
Chapter 2

Background

This chapter first introduces the characteristics of variable density jets, and how their governing equations and behaviour differ from the canonical incompressible jet. Variations in density can be due to heating, high-speed compressibility effects, or the mixing of various species. The current project investigates heated jets emanating from round nozzles. Important findings from previous studies are outlined, including the statistics and coherent structures in these flows, and open research questions on heated jets are introduced.

Subsequently, classes of techniques for experimental density, temperature and concentration measurements in variable density flows are compared. These measurements are critical to further the understanding of the structure of heated jets. The range of techniques available span 1D (point), 2D and 3D measurements.

The discussion narrows towards refraction-based density measurement techniques, which are preferred in the current work due to their non-intrusive nature. The backgroundoriented schlieren technique is a member of a sub-class, density gradient methods, which is selected for further development.

2.1 A description of variable density jets

A jet is a free shear flow in which fluid emanates from a nozzle or orifice into ambient fluid. The ambient fluid may have its own motion, such as counter-, cross- or co-flow to the jet's axis, or be quiescent, which is known as a submerged jet [2]. The canonical jet is a well-studied flow with incompressible fluid emanating from a round nozzle into a quiescent fluid with identical properties. The evolution of the jet flow is broadly characterised by the Reynolds number based on nozzle diameter,

$$Re_D = \frac{\rho_e U_e D}{\mu_e},\tag{2.1}$$

which describes the relative importance of inertia to viscous effects, where ρ is the fluid density, U is the average velocity, D is the nozzle diameter, μ is the kinematic viscosity and subscript $_{e}$ refers to conditions at the nozzle exit. Except for very low Reynolds numbers where the developing flow is completely laminar, the shear between the jet and ambient fluid causes turbulent motions and induces enhanced mixing. These motions are instigated by the development of shear layer instabilities.

Although variable density jets, which are the subjects of this study, introduce some changes to the jet's behaviour, a description of the canonical case is still helpful in un-



Figure 2.1: Evolution of a submerged free jet with downstream length x (not to scale), adapted from Abdel-Rahman [1] and Ball et al. [11]. Flow is from left to right. The jet is initially laminar, emerging from a contoured nozzle with exit diameter D. The growth of instabilities transitions the flow to a turbulent state, encouraging mixing with the ambient fluid. Pertinent flow features are labelled in the top half of the diagram. Axial regions are labelled in the bottom half. Radial regions are labelled on the right.

derstanding the evolution of jets. Figure 2.1 illustrates the zones of development and flow features in the canonical jet, which can also be understood in terms of their axial distance x from the source relative to the nozzle diameter, i.e. x/D. The potential core is the region of unmixed fluid near the source, typically in the region $0 \le x/D \le 7$ [11]. The fluid in this region may be laminar or turbulent, depending on the upstream conditions inside the nozzle. For jets of practical interest, a turbulent region exists far from the nozzle where ambient fluid is entrained into the jet through the action of turbulent eddies. The transition from a laminar potential core to a turbulent far-field can occur from $7 \le x/D \le 70$ and is initiated through the development of shear layer instabilities, which result in vortex roll-up (Kelvin-Helmholtz instability) and pairing. The instabilities extract energy from the flow and their break-down near the end of the potential core gives rise to the transition region of the jet. Beyond the transition region exists the far-field, which is characterised by chaotic motions across a range of scales. The jet achieves self-similarity in the mean and fluctuating profiles of velocity and transported scalar in the far-field. The effects of buoyancy become more important in the far field of a variable density jet, as the jet's momentum is steadily decreased through the action of turbulent motions. Radially, the jet can be divided into a centreline region, the shear layer, and outer region. The shear layer is most interesting as the flow structures and mixing with the ambient environment originate here.

In their reviews on the canonical jet, Abdel-Rahman [1] and Ball et al. [11] classify previous work broadly into two categories: study of jet's statistics (especially the farfield), and study of coherent flow structures (especially in the near- to intermediatefield). The former is useful for model development and validation, and these studies have been conducted over nearly a century. In the far-field, where the jet exhibits selfsimilarity of mean velocity and concentration profiles, the characteristics of the flow can be estimated through the Reynolds-averaged Navier-Stokes (RANS) equations and assuming a turbulence model such as the $k - \varepsilon$ model. These characteristics include the mean velocity and concentration self-similarity scaling, decay rate, spreading rate



Figure 2.2: Cutaway normalised density field of a heated jet from direct numerical simulation, with Reynolds number based on nozzle diameter $Re_D = 10,000$, exit Mach number Ma = 0.6 and exit density ratio $\rho_e/\rho_{\infty} = 0.8$. Flow is from left to right. Details of the simulation are given in chapter 5.

and turbulent stresses, among other properties [103]. Point-based measurements of the far-field velocity are numerous, due in part to the limited measurement techniques, and agree well with the modelled far-field characteristics [11]. The classical view of far-field jet statistics originating with Kolmogorov asserts that the small-scale turbulence in the far-field is isotropic. This allows the jet to be treated as a point source of momentum at a virtual origin producing a universal far-field. More recent work questions this view, and the inflow and boundary conditions are found to have an impact on the jet's stream-wise evolution, which has been reviewed by George [41] and shows that differences in inflow boundary conditions, such as velocity profile and turbulence intensity, can significantly alter the development of the jet's near-field turbulent structures, with evidence that the far-field characteristics are influenced as well. Ball et al. [11] and George [41] summarise the key findings of recent studies on the influence of initial conditions on the far-field characteristics of the jet. It is seen that there are at least three forms of self-similar behaviour in the far-field which arise from the initial conditions: self-similarity of all statistical moments and scales, self-similarity of a limited number of moments and scales, and local self-similarity as the profiles scale with local quantities. George [40] notes that development of the shear layer depends on the growth of a sequence of instabilities and coherent structures near the nozzle, eventually influencing the jet's far-field growth. The dependence of the jet's far-field characteristics on the initial conditions and near-nozzle development complicate the development of turbulence models and prediction, which should also consider the influence of inlet conditions and near-nozzle structures.

The jet may also have a different density to its ambient environment due to the added fluid having a different temperature or molecular composition. Changes in density may also be due to compressibility effects, when the local Mach number Ma > 0.3. The variations in density can modify the jet's growth, as will be discussed. The spatial and temporal evolution of temperature and molecular composition (concentration) of the jet due to mixing with the environment can be described through the scalar transport equation. Both these effects may be present simultaneously. The density variation is characterised through the exit-to-ambient density ratio, ρ_e/ρ_{∞} . The jet is light if $\rho_e/\rho_{\infty} < 1$, and heavy if $\rho_e/\rho_{\infty} > 1$.

The density and scalar fields of variable density jets can be extremely complex, like the velocity field. If the scalar transport has little or no effect on fluid properties such as density, it is considered a passive scalar. Examples include weakly-heated jets, or concentration of an injected dye or tracer particles. However, if the scalar influences the fluid properties, it is an active scalar. This includes strongly heated jets, where the active scalar is temperature (and a thermodynamic equation-of-state is also required to relate temperature and density) and mixing of different chemical species such as helium in air, where the active scalar is species concentration. In this study, compressibility effects associated with high-speed flows (Mach number Ma > 0.3) are not considered. These flows possess their own unique characteristics, especially the supersonic jet, and their modelling is significantly different to low-speed flows [25].

A section of the density field in a heated jet obtained via direct numerical simulation (DNS) is shown in figure 2.2, which demonstrates the complex transition to turbulence through the growth of instabilities. Small- and large-scale features can be seen far from the inlet. Chassaing et al. [25] provide an extensive description of the theory of variable density jets, and a summary of previous experiments. They note that the spreading rate, turbulence intensity and peak Reynolds stresses are inversely proportional to the density ratio. A related flow to the variable density jet, the plume, is produced by a source of buoyancy, with no point-source of momentum such that the flow is not inertia-dominated [74]. As mentioned previously, buoyancy also becomes important in the Boussinesq regime of the heated jet, where the flow resembles a plume.

Modelling a variable density flow is significantly more complex than an incompressible flow. In addition to the continuity and momentum equations, an equation for concentration of each additional species is required [25],

$$\frac{\partial(\rho C)}{\partial t} + \nabla \cdot (\rho C \vec{v}) = \nabla \cdot (\rho D_C \nabla C) + \rho S_C, \qquad (2.2)$$

where ρ is the density field, C is the concentration field, t is time, \vec{v} is the velocity field, D_C is the mass diffusivity of the species, and S_C is a source term. Heat transfer is described by the energy equation,

$$\frac{\partial \left(\rho c_p T\right)}{\partial t} + \nabla \cdot \left(\rho c_p T \vec{v}\right) = \frac{\partial P}{\partial t} + \vec{v} \cdot \nabla P + \Phi_{\nu} + \nabla \cdot \left(k \nabla T\right) + q_{rad}, \tag{2.3}$$

where T is the temperature field, c_p is the specific heat at constant pressure, P is the pressure field, Φ_{ν} is the viscous dissipation rate, k is the thermal conductivity and q_{rad} represents radiative heat transfer.

Peters [101] explains that the evolution of any active scalar ψ , assuming low speed flow such that pressure variations and viscous dissipation can be neglected, can be expressed in the general form,

$$\frac{\partial \left(\rho\psi\right)}{\partial t} + \nabla \cdot \left(\rho\vec{v}\psi\right) = \nabla \cdot \left(\rho D_{\psi}\nabla\psi\right) + \rho S_{\psi},\tag{2.4}$$

where D_{ψ} is a diffusivity and S_{ψ} is a source term, e.g. heat release from chemical reaction (which is not relevant to a non-reacting jet). By introducing the density-weighted decomposition of the scalar, also called Favre averaging,

$$\psi = \tilde{\psi} + \psi'' \tag{2.5}$$

where

$$\widetilde{\psi} = \frac{\overline{\rho\psi}}{\overline{\rho}},\tag{2.6}$$

and the over-bar refers to the ensemble average, the density-weighted scalar variance is defined as $\widetilde{\psi''^2}$. Mixing in the flow can be described by the transport equation for scalar variance [101, 144],

$$\underbrace{\frac{\partial\left(\overline{\rho}\psi^{\prime\prime2}\right)}{\partial t}}_{1} + \underbrace{\nabla\cdot\left(\overline{\rho}\widetilde{\psi}\widetilde{\psi}^{\prime\prime2}\right)}_{2} + \underbrace{\nabla\cdot\left(\overline{\rho}\widetilde{\psi}^{\prime\prime}\widetilde{\psi}^{\prime\prime2}\right)}_{3} = \underbrace{\nabla\cdot\overline{\left(\rho D_{\psi}\nabla\psi^{\prime\prime2}\right)} + 2\overline{\psi^{\prime\prime}\nabla\cdot\left(\rho D_{\psi}\nabla\widetilde{\psi}\right)}}_{4} - \underbrace{2\overline{\rho}\widetilde{\psi^{\prime\prime}\psi^{\prime\prime}}\cdot\nabla\widetilde{\psi}}_{5} - \underbrace{2\overline{\rho D_{\psi}\left(\nabla\psi^{\prime\prime}\cdot\nabla\psi^{\prime\prime}\right)}}_{6}, \tag{2.7}$$

where the terms refer to:

- 1. Unsteadiness of the scalar variance.
- 2. Mean convection.
- 3. Turbulent transport.
- 4. Molecular diffusion, which should be negligible for a turbulent flow.
- 5. Turbulent production.
- 6. Fluctuating scalar dissipation rate due to turbulence.

The terms on the right-hand side require modelling informed by experimental measurements. For example, the second-last term, $-2\rho D_{\psi} (\nabla \psi'' \cdot \nabla \psi'')$, is unclosed, and affects the flow's mixing by dictating the dissipation of turbulent motions. The transport equation of the Favre-averaged Reynolds stress $\vec{v''v''}$ also provides insight on the mechanisms of mixing in these flows. Panchapakesan and Lumley [100] find positive production of turbulence kinetic energy (TKE) tr $(\vec{v''v''})$ in an incompressible air jets, which is nearly twice as high in the far-field of a helium-air jet ($\rho_e/\rho_{\infty} = 0.14$). In contrast, Charonko and Prestridge [24] found negative production of turbulent kinetic energy near the centreline in the far-field of a sulphur hexafluoride (SF₆)-air jet ($\rho_e/\rho_{\infty} = 4.2$), but positive production in the shear layer with a similar magnitude to a constant density air jet. A large turbulent flux is observed near the centreline, which is absent in the air jet, which the authors state is responsible for transporting TKE from the shear layer to the centreline. The mechanisms underlying the spreading and mixing behaviour of variable density jets are quite poorly understood, which in turn affects the quality of flow modelling. Additional experimental measurements of the density and concentration fields are invaluable to furthering our understanding of these flows.

Even in low-speed jets, the vortex roll-up and potential core break-down can lead to large fluctuations in density in the near- to intermediate-region of the jet. Using the Reynolds average notation,

$$\psi = \overline{\psi} + \psi', \tag{2.8}$$

Chassaing et al. [25] notes some observations on the correlations of fluctuating velocity, density and scalar in low-speed (isobaric) jets:

- In lighter jets, where $\rho_e/\rho_{\infty} < 1$, the mean velocity-density fluctuation correlations $\overline{\rho' u'_i}$ are negative. In heavy jets, where $\rho_e/\rho_{\infty} > 1$, these correlations are positive. In the radial velocity correlation, this signifies that the motion of lighter fluid into heavier fluid is statistically dominant.
- The density-concentration fluctuation correlation $\overline{\rho'C'}$ has the same sign as $\overline{\rho'u'_i}$, but the density-temperature fluctuation correlation $\overline{\rho'T'}$ is always negative.
- The higher order correlation of density and temperature $\overline{\rho T'^2}$ is always negative (note that instantaneous density is used). The correlation $\overline{\rho T'u'_i}$ is positive in light jets and negative in heavy jets. The concentration correlation $\overline{\rho C'u'_i}$ is always positive.

Although these observations provide some insight on the jet's mixing, experimental data are required to meaningfully develop models of the jet's mixing. Ultimately, simultaneous density and velocity measurements are required to investigate the density-velocity correlations. The present work takes a step towards this by developing a density measurement technique to complement existing velocity measurement techniques.

In addition to statistical characterisation of the jet, the study of coherent structures can provide insights on the underlying flow physics. Various definitions of coherent structures have been proposed, defined based on statistical correlation of flow quantities or through modal analysis supporting qualitative inspection [11]. Coherent structures can be identified through flow visualisation and 2D or 3D velocity or concentration measurements. Until recently, these visualisations were limited to qualitative inspections of the flow's development, e.g. through dye or smoke injection. True- or high-fidelity simulations such as DNS and large-eddy simulation (LES) are also employed in the study of coherent structures. In variable density flows, optical techniques based on the relationship between a fluid's density and refractive index can be used for both qualitative and quantitative flow visualisation, which will be discussed in more detail in the next section.

The flow structures which develop in the near- to intermediate-regions of the jet are complex, three-dimensional and can be quite well-organised in time and space. The Reynolds number and initial conditions also play a large role in determining the nature of these structures. For example, Mi et al. [86] observed that a fully developed pipe jet of identical Reynolds number to contoured nozzle jet does not display a laminar potential core and possesses much smaller and less well-organised vortical structures. The pipe jet consequently achieves self-similarity much earlier. As discussed previously, shear layer instabilities cause vortex rings to form, which pair and merge at the end of the potential core. Shearing may cause the paired rings to be oriented off-axis [78]. The end of the potential core is 'pinched' by the merging rings, and dramatically disintegrates into chaotic multi-scale turbulence. Previous experiments have demonstrated that trains of vortex rings may form along the shear layer, with the rings pairing alternately, with the end of the potential core bobbing up and down as a result [154]. These trains breed streamwise-oriented vortices between them, known as braids, which may organise into counter-rotating pairs spaced evenly around the circumference of the vortex rings. Liepmann and Gharib [78] found small mushroom-shaped radial ejections in transverse slice visualisations through the jet, which they associated with fluid being pushed out radially between the braids. However, the chaotic nature of the vortex ring pairing ensures that the mushrooms do not grow very large before they are destroyed at the end of the potential core.

When the density ratio $\rho_e/\rho_{\infty} \lesssim 0.7$ and $\rho_e/\rho_{\infty} \gg 1$, the jet's structures can be significantly different to the incompressible jet. The light jet is a well-studied case, yet the formation of the near-nozzle region structures is still poorly understood. The heavy jet is not even well-studied, with only Charonko and Prestridge [24] recently investigating a heavy jet $(\rho_e/\rho_{\infty} = 4.2)$ and finding more mixing than in an incompressible jet very near the nozzle but hindered spreading further downstream. As the density ratio is lowered, the formation of structures in the near-field shear layer becomes more organised for an initially laminar jet [117]. This contributes to a mild increase in spreading compared to the incompressible case along the length of the jet. Below a certain density ratio, many studies have noted the formation of spectacular side jets, which emit a large quantity of fluid perpendicular to the jet's axis. The side jets form sporadically, are short-lived, and are spaced evenly along the jet's azimuth. The side jets are believed to be related to the formation of an oscillating instability originating from a self-excited low-density potential core. The exact value of the density ratio which corresponds to the onset of the instability is the subject of debate. Monkewitz and colleagues estimated a density ratio $\rho_e/\rho_{\infty} < 0.72$ using linear stability theory, which was also observed experimentally [58, 92]. Kyle and Sreenivasan [70] disagree and report that their experiments required $\rho_e/\rho_{\infty} < 0.6$ before side jets were observed as they posit a strong dependence on the initial momentum thickness of the jet's velocity profile as well, which reconciles the reports of increased mixing from only $\rho_e/\rho_{\infty} < 0.6$ in many older experiments, e.g. Corrsin and Uberoi [27]. Russ and Strykowski [117] confirmed that the side jets require an initiallylaminar flow, indicating that the well-organised shear layer structures are required for the formation of the side jets, and not only a low density ratio. Further investigation of a forced, incompressible jet DNS by Brancher et al. [17] showed that the side jets may be related to mushroom structures being allowed to persist and grow for a longer time by the well-organised vortex ring train. However, the reason for their sporadic rapid growth is not well understood.

The transition to turbulence in jets is complex. Three-dimensional measurements of the density and velocity field are needed not only for the study of flow structures, but for generating experimental data to validate flow and heat transfer models. Thus, it is imperative to develop and validate an experimental technique which can be used to study the development of the jet structure, which is the primary aim of this work.



Figure 2.3: A non-exhaustive family tree of density, temperature and concentration measurement techniques in variable density gas flows. Shading indicates the path to the selected techniques in this project.

2.2 Overview of experimental techniques for density, temperature and concentration measurements in variable density jets

As discussed in the previous section, variations in density in low-speed flows can be due to thermal or concentration effects. Numerous techniques have been developed to indirectly measure density fields through the temperature or concentration fields in gases.

These techniques can broadly be described as contact or non-contact methods, additive or non-additive methods. Figure 2.3 illustrates the classification of the measurement techniques used for density, temperature and concentration field measurement in gases. The family of techniques that are the focus of this project, refractive index gradient methods, are highlighted. The density field in gases is also related to the refractive index field, n. A family of techniques has been developed which rely on measuring the disturbances to light propagating through the flow of interest. These techniques are non-contact and non-additive, with absolutely no disturbance to the flow. Qualitative flow visualisation, and quantitative density field measurements, can be obtained through measurement of the refractive index, or its spatial gradients, based on the deflection of light. This gives rise to further subcategories: shadowgraphy, density gradient methods (including background-oriented schlieren), and phase difference methods (which include holography). In all of these techniques, a camera is used to record the passage of light rays through the volume. This allows the measurement of a large and continuous field of view. These refraction-based methods will be discussed in further detail in the next section. But one should also briefly consider other prominent measurement techniques, which were the backbone of many experimental studies underpinning our understanding of variable density jets, and still have a place in fluid diagnostics.

The oldest methods for quantitative measurements have involved the use of electrical or mechanical probes inserted into the flow which produce point measurements. These are obviously contact methods, but require no additional substances introduced to the flow. Electrical probes are mostly used to measure temperature, and include the thermocouple, thermistors, resistance temperature detectors (RTDs) and cold wire systems. Thermocouples utilise the Seebeck effect, and the temperature is inferred from the voltage generated at the junction between two dissimilar metals. Thermistors and RTDs are electrical elements, whose resistance depends strongly on the temperature, and are actively powered by a constant current circuit. Cold wire systems are similar in principle to thermistors and RTDs but operate using a constant voltage circuit (like hot-wire anemometry) with a very small temperature difference between the wire and ambient temperature. The 'point-wise' nature of these measurements really depends on the size of the thermocouple junction or element size, and there is an associated spatial averaging. These techniques are precise to the order of a tenth of a Kelvin, but the probes record only their own temperature, rather than the flow temperature [118]. There are further disadvantages associated with using probes for turbulence measurements. Physical disturbance of the flow means that it is difficult to detect coherent flow structures with probe arrays. Electrical probes often suffer from slow response times (milliseconds or seconds) which further confounds efforts to detect structures, there is a thermal inertia associated with the sensing element, and they may heat themselves due to electrical resistance (particularly thermistors). The sensing element of the probe also conducts heat to the associated circuitry, which are termed 'end-losses'. It is possible to correct for thermal inertia and end-losses in cold wire systems, such that the measurement rate can be on the order of a few kilohertz [26, 83]. Concentration measurements using probes are far less common than temperature measurements. Way and Libby [148] described a hot wire-film probe for concentration measurements in helium-air jets. The thermal fields of the wire and film overlap, and the extent of the film's field increases with the concentration of helium which is detected by the wire. This probe was used by Panchapakesan and Lumley [100] for their extensive measurements of the velocity and concentration field characteristics in helium-air jets. Although cost-effective, non-intrusive techniques have gained traction over probes in modern turbulence diagnostics due to the disadvantages of probe methods.

Non-contact methods based on optical phenomena include absorption/emission methods, scattering techniques, and the refraction-based methods. Point-based measurements can be obtained non-intrusively by utilising scattering phenomena, while fluorescence methods can provide planar or volumetric measurements of temperature or concentration in variable density gas flows. The development and application of these methods is broad and mature, and comprehensive reviews of scattering and fluorescence techniques are provided by Kowalewski et al. [69] and Lee [75]. The response and acquisition times of these techniques are far shorter than probes and are typically much shorter than flow characteristic time scales, so they are ideal for the study of turbulent flows. Rayleigh scattering describes the scattering of light by particles much smaller than the wavelength of the incident light. It is an elastic phenomenon, meaning that the scattered light is of the same frequency as the incident light. Richards and Pitts [113] utilised Rayleigh scattering in their characterisation of the far-field in isothermal helium-air, methane-air and propane-air free jets. The authors noted that the technique required a clean facility, free of dust, to prevent the measurements being contaminated by the more-dominant Mie scattering. A similar technique based on inelastic scattering, Raman scattering, has been used to determine the composition of flames; combined Rayleigh-Raman scattering techniques can be used to determine both the flow temperature and composition fields. Mie scattering involves the elastic scattering of light by particles much larger than the wavelength of the light. Sautet and Stepowski [119, 120] utilised Mie scattering in their investigation of the near-nozzle development of hydrogen-air jets to obtain the scalar concentration profiles. Furthermore, combined Rayleigh and Mie scattering have been used to provide simultaneous point concentration and velocity measurements [88, 89]. Although the present discussion of non-contact methods is limited to optical methods, note that acoustic methods have also been used to provide point-wise temperature and velocity measurements in heated gaseous flows [52].

A prominent example of absorption/emission methods for concentration measurements is laser-induced fluorescence (LIF), which is based on the emission of light (fluorescence) from a substance molecule after excitation, and subsequent relaxation, from a laser source of a particular wavelength. The fluorescence is emitted at another, longer wavelength. LIF is used for fluid concentration measurements by contaminating a flow with a fluorescent substance, e.g. acetone, and guiding a laser to illuminate a section of the flow. The intensity of the fluorescence wavelength is dependent on the local concentration of the substance through the Beer-Lambert law and is captured by a camera. The concentration field can then be related to the density. The laser can be formed into a sheet to illuminate a 2D plane in the flow, called planar laser-induced fluorescence (PLIF), or used for volumetric illumination, called volumetric laser-induced fluorescence (VLIF). In the latter case, to obtain 3D measurements of intensity/concentration, it is necessary to use emission tomography from multiple cameras viewing the flow simultaneously. PLIF is a well-established diagnostic for variable density turbulent flows. For example, Charonko and Prestridge [24] paired PLIF with particle image velocimetry to simultaneously measure velocity and concentration in air and sulphur hexafluoride (SF_6) jets. This was used to yield the profiles of the Favre-averaged Reynolds stresses, and investigate the budget of terms in the Favre-averaged turbulent kinetic energy transport equation.

2.3 Refraction-based density measurements

Optical techniques for fluid density measurement are based on relations between a fluid's refractive index and its density. Merzkirch [84] describes the derivation of the Clausius-Mosotti relation between the refractive index n of a fluid (liquid or gas) to its density ρ , based on the propagation of electromagnetic waves, i.e. light, through a medium of variable density and its effect on the polarity of fluid molecules. In gases, the Clausius-Mosotti relation simplifies to the Gladstone-Dale relation,

$$n-1 = \rho G, \tag{2.9}$$

where G is the Gladstone-Dale constant which depends on the species composition of the gas, and weakly on temperature and the wavelength λ of the light itself. Assuming



Figure 2.4: The deflection of the ray within the volume has been exaggerated for clarity.

the Gladstone-Dale constant of the gas is known, the density field of the gas can be determined by measuring the refractive index field. Hence, it is established the refractive index field allows the measurement of the density field. Now consider how the refractive index field itself can be measured. The propagation of light rays through a volume with varying refractive index is described by the Fermat principle [16],

$$\frac{\mathrm{d}}{\mathrm{d}s} \left(n \frac{\mathrm{d}\vec{\mathbf{x}}}{\mathrm{d}s} \right) = \nabla n, \qquad (2.10)$$

where \vec{x} is the position vector, ds is a differential length along a light ray and ∇n is the (spatial) refractive index gradient. Sharma et al. [125] and Hewak and Lit [55] use the variable substitution dt = nds to obtain,

$$\frac{\mathrm{d}}{\mathrm{d}t} \left(n^2 \frac{\mathrm{d}\vec{\mathbf{x}}}{\mathrm{d}t} \right) = \frac{1}{n} \nabla n, \qquad (2.11)$$

and subsequently use the variable substitution du = ds/n to rewrite the Fermat principle in the following form:

$$\frac{\mathrm{d}^2 \vec{\mathbf{x}}}{\mathrm{d}u^2} = n \cdot \nabla n. \tag{2.12}$$

Equation 2.12 is numerically integrated using a Runge-Kutta method to find the deviation of light rays passing through a medium of variable refractive index due to changes in wave speed [125]. This is known as ray tracing. The deviation of rays can also be understood in terms of the optical wavefront, as light rays are perpendicular to the wavefront. Huygen's principle states that the wavefront is a source of secondary spherical waves [16], which dictates that the rays (normal vectors to the wavefront) are oriented towards the region of higher refractive index [124]. In experiments, the refractive index field and its gradients are not known. They are sought to measure the fluid density field. From the Fermat principle, it is seen that detecting changes in the trajectory of light rays passing through the flow can be used to measure the refractive index field and its gradients.

Equation 2.12 demonstrates the effect of the variable refractive index field on the trajectory of a ray passing through the medium. By comparing the difference in trajectory or travel time between this ray and a ray that does not encounter any refractive index gradients, it is possible to measure the strength of the refractive index gradients.



Figure 2.5: Experimental setup for collimated light shadowgraphy. Light can be collimated from a point source using lenses or a parabolic mirror.

Consider a three-dimensional volume of variable refractive index situated in otherwise ambient conditions (*n* is constant outside the volume), illustrated in figure 2.4. In ambient conditions (no refractive index gradients), light rays travel in a constant direction. A light ray is launched through the volume, initially oriented in the z' direction. The variable refractive index field will result in the ray following a curved path through the volume, but restricting the discussion to weak refractive index gradients such as those in heated gases and a thin volume relative to the ray's path, the ray leaves at roughly the same position in x' and y' but travelling in a different direction. The weak refraction implies that the slope of the exiting ray, $\partial x'/\partial z'$ and $\partial y'/\partial z'$, are small and it can be approximated as a small deflection of angle ε' from the boundary of the refractive object. Equation 2.12 simplifies to describe the ray's curvature:

$$\frac{\partial^2 x_r}{\partial z'^2} = \frac{1}{n} \frac{\partial n}{\partial x_r},\tag{2.13}$$

where x_r denotes the components x' and y'. Finally, a photodetector, photographic film or camera sensor records the ray's position at a plane on the opposite side of the volume. Lenses may be used, in which case standard geometric optics can be applied to determine the modified ray path from the lens to the sensor. The ray curvature equation is the basis of the refraction-based measurement techniques. However, which aspect of the light ray's deflection is measured depends on the technique.

Perhaps the most straightforward refraction-based technique is the shadowgraph. The simplest shadowgraphs can be formed with only a point or collimated light source, the refractive object to be measured, and a screen on which to cast the object's shadow. This is illustrated in figure 2.5. The deflection of the light rays from their ambient path leaves a shadow at their ambient position on the screen, with the ray instead illuminating a point ΔX away. A sharper shadow is cast using collimated light. The shadow is related to the gradients of the deflection angle ε' , which is the second derivative of refractive index $\partial^2 n/\partial x_r^2$, and the change in intensity I at a point on the screen is given by:

$$\frac{\Delta I}{I} = \Delta Z' \int \left(\frac{\partial^2}{\partial x'^2} + \frac{\partial^2}{\partial y'^2}\right) \ln n \, \mathrm{d}z'. \tag{2.14}$$

A somewhat-related technique is silhouette photography [130], which places illumination in line with a camera's optical axis using a semi-silvered mirror. The illumination passes once through the flow, and retroreflective material then casts the light back through the volume towards the camera, casting a shadow with detailed flow structures. The shadowgraph and silhouette photograph are useful and versatile tools for flow visualisation, but the nature of shadow formation and the double integration means that it is difficult to develop these techniques to produce quantitative measurements of the refractive index field. As shown in the family tree in figure 2.3, these techniques are considered to stand apart from other refraction-based methods which can more readily be extended to quantitative measurements.

Instead, consider measurement of a ray's deflection angle. The deflection of the light ray can be found by integrating the ray curvature equation along the ray's line of sight,

$$\tan \varepsilon'_{x_r} = \int \frac{1}{n} \frac{\partial n}{\partial x_r} \,\mathrm{d}z'. \tag{2.15}$$

In the simple setup with no imaging optics, e.g. camera lens, shown in figure 2.4, it is seen that $\tan \varepsilon' = \Delta X / \Delta Z'$. For the weak refractive index gradients associated with heated flows, one may use the small angle approximation $\tan \varepsilon' \approx \varepsilon'$. Equation 2.15 shows that the deflection angle, or displacement, is related to the path-integrated refractive index gradients. This integral equation is in the form of a tomographic reconstruction inverse problem. Solving these problems allows the reconstruction of an object based on its projections and they will be discussed in detail in the next chapter. For now, let it suffice that solving the inverse problem, with the deflection angles as the input, allows the three-dimensional reconstruction of the refractive index field, which is our aim. Tomographic reconstruction of complex flow structures requires projections of the object from multiple views, and the short time scales of turbulent flows mean that these views must be recorded simultaneously for instantaneous density field measurements. The exact number of views required for a high-quality reconstruction will be investigated in later chapters, but obviously the selected method should lend itself well to a multiple-camera setup. There are many implementations of deflection angle measurement, which are also known as density gradient methods. These methods will be the focus of this project. Now consider the most prominent density gradient methods, to evaluate which methods are conducive to a multiple-camera setup for quantitative 3D density field measurements.

Perhaps the best known density gradient method is schlieren, or Töpler schlieren after its inventor. Although there are many variations of the Töpler schlieren system, the simplest places the refractive object to be measured in a beam of collimated light, after which a converging lens is placed. A knife-edge or grid is placed at the focus of the lens, such that half the beam is blocked, after which the imaging system is located, as shown in figure 2.6a. In practice, parabolic mirrors are used to collimate or focus the light source, as large mirrors can be obtained more easily than large lenses and, hence, a much greater field of view can be achieved. The deflection of the rays allows them to either pass over the knife-edge, or be obstructed by the knife-edge. Correspondingly, a bright or dark region is formed on the screen, and the change in intensity is given by:

$$\frac{\Delta I}{I} = \frac{f_1}{a} \int \frac{1}{n} \frac{\partial n}{\partial x_r} \,\mathrm{d}z',\tag{2.16}$$

where a is the height of the image in the image plane. Non-coherent light is preferred, as diffraction effects from the knife-edge are minimised. Note that the orientation of the knife-edge is important in determining which component of the refractive index gradient is measured; by rotating the knife-edge through 90° about the optical axis, the other component of the refractive index gradient may be captured as well. Like shadowgraphy, it is difficult to obtain quantitative information on the refractive index gradients from Töpler schlieren because only changes in intensity are recorded. However, in some simple



Figure 2.6: Experimental setup for density gradient methods: a) Töpler schlieren, b) moiré deflectometry, c) speckle photography, d) structured light refractography, e) background-oriented schlieren (BOS).

cases such as 2D flow, it is possible to obtain density and temperature measurements from schlieren photographs. Zamyatina et al. [156] were able to obtain the time-averaged axisymmetric temperature field in a round heated jet from schlieren photographs in the 1960s. Merzkirch [84] and Settles [124] detail several modified experimental setups. Notable modifications include double-pass and focussing schlieren, which increase sensitivity and resolution, respectively, and colour schlieren, which provides spectacular colour images of the density gradients by replacing the knife-edge with a colour filter. Elsinga et al. [35] obtained path-integrated quantitative measurements of the density field of the supersonic flow over a two-dimensional wedge by calibrating the colour shift in colour schlieren to corresponding density values. The colour filter can be designed to record ray deflections in perpendicular directions simultaneously, which is not possible with standard Töpler schlieren. They note that this method had a high noise level and smaller dynamic range compared to background-oriented schlieren.

As it is seen from shadowgraphy and Töpler schlieren, recording only changes in intensity in a bright- or dark-field due to light ray deflections is generally not conducive to quantitative measurements of the refractive index gradients. The solution is to organise the light into a structured pattern before it reaches the image sensor. In ambient conditions, the structured light will be imaged as a pattern. Then the ray deflections correspond to measurable shifts in the pattern $\Delta \vec{X}$, which may be related to the deflection angle as discussed previously. Now consider four different methods of generating the pattern of structured light: moiré deflectometry, speckle photography, structured light refractography, and background-oriented schlieren.

Moiré deflectometry uses an experimental setup very similar to Töpler schlieren, with the omission of the knife-edge at the beam focus. Instead, two fine rulings of alternating opaque and transparent lines known as Ronchi gratings are placed away from the beam focus as shown in figure 2.6b. The rulings are oriented at a slight angle to one another as seen from the camera, which produces a fringe pattern called moiré in the image in the absence of any refractive object. When the refractive object is introduced, the deflection of light rays accordingly shifts the fringes. The method was introduced by Kafri [61], and subsequently used to measure the time-averaged axisymmetric temperature field in flames [12, 67]. The authors noted the similarity of the moiré fringes to those produced in interferometry (discussed later), except that the technique was much more robust to vibrations and that the sensitivity could be easily adjusted by altering the distance between the Ronchi gratings. Both schlieren and moiré deflectometry methods can be difficult to set up due to the need to collimate the light source.

Speckle photography requires a laser to be used as the light source. By expanding the beam to cover the required field of view, and passing through a ground-glass diffuser as shown in figure 2.6c. Transmission of the coherent light through the diffuser's rough surface results in random deflection and interference [37]. The interference pattern is imaged by the camera as a random pattern of bright and dark spots, which is called a speckle pattern. With no refractive object present, the speckle pattern is stable. When the refractive object is introduced, the ray deflection causes an apparent shift in the speckles $\Delta \vec{X}$. The original method of recording specklegrams on photographic plates and interpreting deflections based on Young's fringes [85] has been superseded by recording the speckles directly onto an image sensor and evaluating the speckle shifts directly. This makes the speckle photography much less prone to errors, as well as being more practical and straightforward to process. The theory of speckle formation will be revisited in chapter 6, as this work combines speckle photography with background-oriented schlieren. Instead of producing a random interference pattern, a laser beam can also be shaped into precise patterns such as an array of dots or rings, lines, concentric circles, or waves. These can be achieved through an arrangement of optics, including lenses, diffraction gratings and fly-eye condensers [79], as shown in figure 2.6d. This technique is called structured light refractography, and the ray deflections are measured as shifts in the laser pattern, like speckle photography [114].

Background-oriented schlieren (BOS) is both the most recent, and simplest, density gradient method that will be discussed here. Introduced by Raffel et al. [105], the BOS setup requires only a printed background pattern, ordinary non-coherent lighting, and a camera. The camera images the background pattern, and ray deflections from the refractive object result in an apparent shift of background features compared to ambient conditions. The simplicity of the setup lends itself well to a multiple-camera setup required for tomographic setup, and, as such, BOS is chosen for further development in this project. This is in stark contrast to more complicated setups such as calibrated colour schlieren that could also deliver quantitative measurements. Compared to the other density gradient methods, the resolution of BOS is limited due to the camera being focussed on the background pattern rather than the flow. As shown in the next chapter, the sensitivity of the method to refractive index gradients requires that the background pattern be far from the flow, so there are conflicting requirements. Developing the ideal compromise between measurement resolution and sensitivity in BOS is a major theme of this work.

Lastly, consider a separate class of refraction-based density measurements, interferometry. Interferometric techniques measure the phase difference $\Delta \varphi$ in coherent light passing through the (weakly) refractive object compared to ambient conditions,

$$\Delta \varphi = \frac{2\pi}{\lambda} \int n - n_0 \,\mathrm{d}z,\tag{2.17}$$

where λ is the wavelength of light and n_0 is the ambient refractive index value. The phase difference, which cannot be measured with any of the previously discussed techniques, is captured by introducing a secondary beam that does not pass through the flow. These methods have a much higher sensitivity than the previously discussed methods. Up to fractions of a wavelength in the optical path difference can be detected [121]. Although there are many types of interferometers, the discussion is restricted to holographic interferometers, which are the most widely used. An overview of other interferometers used in fluid mechanics is given by Merzkirch [84].

The holographic interferometer relies on splitting a laser beam into two branches, as shown in figure 2.7. One branch passes around the flow in ambient conditions (reference beam), while the other (signal beam) passes through the region where the refractive object will be. A fringe pattern is formed on the holographic plate or digital sensor by the interference of the two beams. An initial interference pattern (hologram) is recorded with no flow present, and subsequent holograms can be recorded with the signal beam attenuated by the flow. The process of viewing holograms recorded on holographic plates is detailed by Briers [18], but this process has been almost entirely superseded by digital recording. The phase of the initial and subsequent holograms is reconstructed using the Fresnel approximation [121], and from the difference between the two, one obtains the path-integrated phase distribution of the flow.

Holographic interferometry has been applied to 3D flow measurements, but the complexity of the experimental setup has generally limited these investigations to singlecamera measurements of the time-averaged axisymmetric flow field of jets. It is difficult



Figure 2.7: Experimental setup for digital holographic interferometry. The holograms can be recorded without imaging optics, e.g. camera lens.

to extend these techniques to a multiple-view system for tomographic reconstruction of a complex flow field [147]. However, there are some notable studies. Timmerman [139, 140] used a rapid-switching double-pulsed laser to create a complex two-view system, while Doleček et al. [33] reconstructed the phase-averaged density fields of a starting jet through synchronisation of the exposures with the flow's frequency. The sensitivity of these systems to vibration and dust also hinders their wider adoption. This was partially overcome by Guo et al. [50], who combined a moiré system with a 4f optical correlator to generate a laser interferometric system with one beam and presented the axisymmetric time-averaged temperature field in a flame from tomographic reconstruction.

Note that in the interest of brevity, the density gradient methods and interferometry are presented as entirely separate techniques. However, there can be considerable overlap in experimental setups, such as holographic schlieren, and moiré interferometry. Briers [18] presents a detailed discussion of the relationship between schlieren, speckle photography, moiré and holography. These overlaps are often invented to improve measurement resolution or overcome experimental difficulties such as sensitivity to unwanted vibrations. Undoubtedly, more variations of refraction-based density measurements will emerge. But this project develops background-oriented schlieren due to its simple experimental setup and potential for conducting 3D density measurements through multipleview tomographic reconstruction.

Chapter 3

The tomographic background-oriented schlieren (TBOS) technique

3.1 Principles and overview of the TBOS measurement process

In the last chapter, it was shown that BOS has arguably the simplest experimental setup of all density gradient methods. It requires only a printed background pattern facing a camera, with the flow located in between the two. The background pattern may be illuminated by ordinary, non-coherent light. Now the intricacies of the BOS measurement process will be explained.

The principles of BOS displacements are shown in figure 3.1 using a simple camera model. The camera is focussed on the background pattern, which is located at a distance Z_B from the camera's lens. The flow is placed an arbitrary distance between the two, a distance Z_D from the background and a distance Z_A from the lens. A light ray propagating towards the camera will be deflected due to the variable refractive index along its path, compared to the path followed in a uniform refractive index field. As discussed in chapter 2, the curved path of the deflected light can be approximated by a single deflection event with refraction angle $\vec{\varepsilon}$ about the centre of the volume. This approximation is made due to the typically weak refractive index gradients in gaseous flows, and assuming that the depth of the flow is small compared to Z_B [104]. The deflection means that the ray appears to come from the position $\vec{P_r}$ instead of its true position $\vec{P_0}$. The background displacements can be determined to sub-pixel accuracy (≈ 0.1 pixels) using cross-correlation methods for PIV, or optical flow techniques, which will both be discussed in detail later in the chapter. On the image plane, following Richard and Raffel [112] by assuming a thin lens and in-focus background, the ray's displacement can be related to its deflection angle using geometric optics by

$$\vec{\Delta X} = \left(\frac{Z_D f}{Z_B - f}\right) \tan \vec{\varepsilon},\tag{3.1}$$

where f is the lens focal length.

Weak refraction allows the use of the small angle approximation, $\tan \vec{\varepsilon} \approx \vec{\varepsilon}$. Increasing the distance between the volume and the background will increase the displacement of the ray, however it will also reduce the magnification of the background meaning that



Figure 3.1: Schematic of BOS ray deflection and nomenclature. In traditional BOS, the camera (composed of image plane, lens, and aperture) is focussed on the background $(l = +Z_D)$. In the laser speckle BOS implemented in chapter 6, the camera may be focussed anywhere in the range $-(Z_A + f) < l < +\infty$. The global coordinate system origin is located at the centre of the refractive volume. The optical axis of camera 1 is oriented in the global z-axis.

the background features' sizes must often be tailored for each BOS experimental setup.

As discussed in the previous chapter, the deflection angles are the path-integrated refractive index gradients encountered by a light ray,

$$\tan \varepsilon_{x_r} = \frac{1}{n_0} \int \frac{\partial n}{\partial x_r} \,\mathrm{d}z',\tag{3.2}$$

where n_0 is the ambient refractive index. Equations 3.1 and 3.2 demonstrate that the sensitivity to the refractive index gradients is increased by moving the camera closer to the refractive volume, increasing the distance between the camera and background $(Z_D/Z_B \rightarrow 1)$, and using longer focal length lenses (maximising f). The maximum displacement will be obtained when the lens is as close to the measurement volume as possible and the background is as far as possible, approaching a maximum displacement of $\Delta X \approx f \vec{\varepsilon}$ for a given deflection angle.

Notice that the camera is focussed on the background, but the flow may be located much closer to the camera to increase the measurement sensitivity. One must consider the impact of the defocus blurring that is introduced into the measurement owing to this, which is illustrated in figure 3.2. The light that falls on the sensor will be collected from multiple rays, which fall within a cone with a diameter δ at the object plane (refractive volume), which is related to the aperture diameter d_a by:

$$\delta \approx d_a \left(\frac{Z_D}{Z_B}\right). \tag{3.3}$$

This represents the resolution of the gradient field that can be resolved by the optical setup and imposes a restriction on the background distance, since in the limit of infinite background distance and negligible distance between the lens and the volume, $\delta \rightarrow d_a$. Defocus blurring, i.e. circle of confusion d_i , in the image plane of a point at the refractive volume Z_A , when the camera is focussed on the background Z_B is given by,

$$d_{i} = d_{a}Z_{I} \left| \frac{1}{f} - \frac{1}{Z_{A}} - \frac{1}{Z_{I}} \right|$$
(3.4)

where

$$Z_I = \left(\frac{1}{f} - \frac{1}{Z_B}\right)^{-1},\tag{3.5}$$

following Rowlands [116]. Minimising defocus blurring requires smaller apertures and shorter focal length lenses, as well as moving the background/focus plane closer to the measurement object [44]. The diffraction limit also needs to be considered, as this dictates the smallest features that can be captured in the background pattern. According to Raffel [104], the diffraction limit of background features on the image plane is given by

$$d_d = 2.44\lambda \frac{f}{d_a} \left(\frac{Z_I}{Z_B} + 1\right),\tag{3.6}$$

where λ is the wavelength of the light. The overall blur size on the image plane is

$$d_{\Sigma} = \sqrt{d_i^2 + d_d^2} \tag{3.7}$$

For a typical BOS setup, $d_i > d_d$ so that $d_{\Sigma} \approx d_i$, and the priority should be to minimise d_i . But the requirements for minimising blur are contrary to those required for maximising sensitivity. Given that displacement calculation methods such as digital cross-correlation PIV analysis introduce spatial averaging to the measurement over a search area (interrogation window) d_{IW} , the effective spatial averaging in the measured background distortions will be the larger of the defocus blurring or interrogation window. It is therefore sufficient to make blur smaller than the window size. This will be explored further in chapter 5.

Although discussed in the seminal paper by Richard and Raffel [112], the careful selection of appropriate focal length, aperture, and distances to strike a compromise between blur and sensitivity does not seem to have been a priority in some later studies using BOS or TBOS, e.g. [96], but considered in others, e.g. [72]. The careful selection of the setup dimensions, lens focal length and aperture is required to minimise defocus blurring to a reasonable range while maintaining sensitivity so that flow features are faithfully captured.

Decreasing the lens aperture is often a practical solution to reduce blur while maintaining sensitivity. But this requires that the light source must be very bright. To complicate the matter, minimising temporal blurring in the measurement to freeze the turbulent flow features in each snapshot requires short exposure times. This places more emphasis on a bright light source. So, the sensitivity/defocus blur compromise in BOS means that obtaining high-quality 3D density measurements of gaseous flows with low temporal and spatial integration is very challenging.

Mitigating temporal blur, while finding a balance between defocus blur and measurement sensitivity is a recurring theme in this thesis. Chapter 5 contains an evaluation of the impact of increasing defocus blurring and temporal blurring on the ability to capture turbulent flow features in a heated jet. A reasonable limit on defocus blurring in TBOS measurements is proposed to ensure smaller scale turbulence is captured. In chapter 6, a set of guidelines is created which can be used to design a BOS experimental setup with an ideal compromise between defocus blurring and measurement sensitivity.

Returning to the discussion on the deflection of light rays, the path-integrated information on the refractive index gradients provided by equation 3.2 is insufficient to provide three-dimensional information on the refractive index gradients themselves for an arbitrary three-dimensional turbulent flow. The introduction of additional cameras to



Figure 3.2: Two ways of considering the origin of defocus blur in the BOS measurement. In both cases, the camera is focussed on the background image $(l = Z_D)$, which dictates the distance Z_I . a) A cone of light from a point at the background has a finite diameter δ when it passes through the refractive volume (object plane). b) A cone of light from a point at the object plane (thick lines) is not brought to a focus (sharp point) at the image plane due to Z_I being chosen for the background plane Z_B , and instead has a diameter of d_i at the image plane.

monitor the flow from different viewpoints enables the tomographic reconstruction of the three-dimensional refractive index gradient fields ∇n in a global coordinate system. A detailed presentation and discussion of the tomographic reconstruction of the refractive index gradients is provided later in the chapter. Given a reconstructed three-dimensional refractive index gradient field, the refractive index field itself and hence, the density field, in a global coordinate system is obtained by solving a Poisson equation [6],

$$\nabla^2 n = \frac{\partial^2 n}{\partial x^2} + \frac{\partial^2 n}{\partial y^2} + \frac{\partial^2 n}{\partial z^2} = q, \qquad (3.8)$$

where the right-hand side of the equation is populated by taking the derivatives of the reconstructed gradient fields,

$$q \equiv \frac{\partial}{\partial x} \left(\frac{\partial n}{\partial x} \right)_{recon} + \frac{\partial}{\partial y} \left(\frac{\partial n}{\partial y} \right)_{recon} + \frac{\partial}{\partial z} \left(\frac{\partial n}{\partial z} \right)_{recon}.$$
 (3.9)

Subsequently, the Gladstone-Dale relation can be applied to obtain the measured density field.

A camera calibration model is required to relate the image plane coordinates (X, Y) to the global coordinate system centred on the measurement volume (x, y, z). This can be done using a pinhole model [141], as follows:

$$\zeta \begin{bmatrix} X \\ Y \\ 1 \end{bmatrix} = \vec{K}\vec{R} \begin{bmatrix} x \\ y \\ z \end{bmatrix} - \vec{K}\vec{R}\vec{c}$$
(3.10)

$$\vec{K} = \begin{bmatrix} \frac{-f}{l_{px}} & O_X & 0\\ 0 & \frac{-f}{l_{px}} & O_Y\\ 0 & 0 & 1 \end{bmatrix},$$
(3.11)

$$\vec{R} = \begin{bmatrix} \cos\alpha & -\sin\alpha & 0\\ \sin\alpha & \cos\alpha & 0\\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos\beta & 0 & \sin\beta\\ 0 & 1 & 0\\ -\sin\beta & 0 & \cos\beta \end{bmatrix} \begin{bmatrix} 1 & 0 & 0\\ 0 & \cos\gamma & -\sin\gamma\\ 0 & \sin\gamma & \cos\gamma \end{bmatrix},$$
(3.12)

where f is the focal length of the lens, l_{px} is the image plane pixel size, O_X and O_Y are the centre of the image plane in pixels, α , β and γ are rotations of the camera with respect to the global z, y and x axes, respectively and \vec{c} is the vector representing the optical centre of the camera relative to the origin of the measurement volume. These matrices are populated through a calibration procedure, e.g. [157]. A ray between the optical centre of the camera through a given point on the image plane can similarly be computed from:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \vec{c} + \zeta \vec{R}^{-1} \vec{K}^{-1} \begin{bmatrix} X \\ Y \\ 1 \end{bmatrix}, \qquad (3.13)$$

where ζ represents the relative position along this ray. This enables the incident angle and the path of the ray through the volume to be determined.

The process of obtaining three-dimensional density measurements from BOS images is given in figure 3.3 and summarised below. Note that there are several differing approaches reported in previous studies, which will be discussed in subsequent sections, but they generally involve the following steps to obtain the three-dimensional refractive index field:

- 1. Recording apparent deflections of background patterns using cameras looking through the flow. The apparent deflections are caused by the observed variations in the flow's refractive index. The optical setup is often a compromise between maximising sensitivity to the flow's refractive index gradients, and minimising defocus blurring. The choices of background pattern are discussed later in the chapter.
- 2. Calculating the background pattern displacements. These are related to the pathintegrated refractive index gradients observed by the cameras using equations 3.1 and 3.2.
- 3. Tomographic reconstruction of the three-dimensional refractive index gradients. Solving the ill-posed tomographic reconstruction problem relating the deflection angle of light rays seen by all cameras to the refractive index gradients in a global coordinate system. The cameras are calibrated to the global coordinate system using a multiple-camera calibration procedure, e.g. [56, 157].



Figure 3.3: Process for density measurements using TBOS. After the refractive index field n(x, y, z, t) has been reconstructed, the density field can be obtained using the Gladstone-Dale relation. Numbers correspond to the steps listed on page 25.

4. Solution of a Poisson equation, or integration of the gradients, to obtain the threedimensional refractive index field itself. The Poisson equation is usually discretised using finite differences. From here the Gladstone-Dale relation can be used to obtain the density field. The temperature field could also be obtained for an ideal gas with a known pressure field and ideal gas constant.

Each of these steps will be explored further in the remaining sections of this chapter.

3.2 Background design and displacement calculation methods

Aside from the optical setup, the choice and design of background pattern is crucial to accurate TBOS measurements. The displacement calculation method, too, must suit the type of background pattern used. In this section, several approaches used in previous studies will be discussed.



Figure 3.4: Types of backgrounds used in previous studies: a) random dots, b) horizontal lines, c) wavelet noise, d) laser speckles, e) checkerboard.

3.2.1 Background design

The studies of Raffel et al. [105] and Richard and Raffel [112] aimed to develop a density measurement technique with a simple experimental setup. They were inspired by laser speckle photography [85], but in creating BOS they desired a simpler experimental setup that was robust to vibrations and readily scalable to large-scale industrial studies. Modelling the background after a particle image velocimetry (PIV) image, these studies used background patterns of random white dots painted on a black surface to study the flow around helicopter blades. Independently, Dalziel et al. [29] developed a 'synthetic schlieren' method using a horizontal line background pattern. In general, background patterns with periodic features have fallen out of favour, because flow structures located in between the background's period (or an integer multiple of it) cannot be detected. This risk is minimised by using random patterns. A comparison of these backgrounds in shown in figure 3.4. The dot background pattern is easy to generate and can be designed to conform to guidelines for PIV cross-correlation methods (discussed in the next section), so that high-quality displacement fields are obtained. Coloured patterns have also been used in place of monochrome backgrounds. Overlapping coloured dots [99, 131] allow the information density of the background pattern to be increased by minimising featureless space. The displacements from each dot colour are processed separately by filtering the other colours out.

Atcheson et al. [7] proposed using wavelet noise patterns instead of dots. They argue that this type of background carries several advantages over dot background patterns. Foremost is that wavelet noise patterns contain features at multiple scales, so it is possible to use a single background pattern for experiments with different magnifications rather than designing an optimised background for each setup. In their tests on synthetic data, the authors also found that the wavelet noise pattern resulted in displacement fields with higher accuracy than dot patterns when using optical flow for displacement calculations. They also suggest that this type of background, combined with optical flow, is immune to 'peak-locking', where the cross-correlation displacements incorrectly tend to integer values. Despite these advantages, the dot background pattern has proliferated, mainly because strong density gradients can totally smear out the features in wavelet backgrounds, making displacement calculation impossible.

Meier and Roesgen [82] returned to the roots of BOS in laser speckle photography by using the diffuse reflection of a laser off a rough surface to produce a laser speckle background pattern. Instead of forming specklegrams, the laser speckle background is used in the same way as a dot pattern. The characteristics of the speckles can be adjusted using the lens aperture, and bear the unique property that the speckle is always in focus regardless of the camera's focussing distance. The sensitivity of the measurement can also be increased by focussing closer than the object [22, 81, 87]. This decouples the optical setup from the physical dimensions of the setup. The laser speckle background is favoured in the current study and is discussed in more detail in chapter 6.

In general, BOS measurements cannot be performed in real time because the displacement calculation is often much slower than the image acquisition. A special mention is made of the checkerboard BOS developed by Wildeman [151]. This method uses a Fourier demodulation algorithm to calculate the displacements of a checkerboard pattern with similar accuracy to, and much faster than, cross-correlation and optical flow methods. With a GPU implementation, the author reports that real-time BOS measurements up to 190 Hz are easily achieved, while an optical flow pipeline could only operate at 1 Hz.

3.2.2 Displacement calculation methods

Background image displacements have been calculated using digital cross-correlation PIV analysis. Optical flow and particle-tracking methods are also used.

Digital cross-correlation methods have traditionally been used on dot- or speckle-type BOS backgrounds, because they mimic the particle images found in PIV. A brief description of the principles of cross-correlation is relayed here with detailed explanations provided by Soria [129] and Raffel et al. [106]. Cross-correlation methods calculate the displacement between the reference and distorted images in small regions known as interrogation windows. The position of each window is fixed at the same location in both the reference and distorted images, and only the particles inside the window are displaced. It is assumed that the displacement of particles within the interrogation window between the two images is uniform, i.e. that the gradient of displacement in the window is zero. This, of course, is very unlikely to be true for any flow of interest, but it can be a reasonable approximation for small window sizes.

The reference and displaced interrogation windows contain information on the intensity of N_{sm} distinct particles with an assumed Gaussian intensity distribution I (which may be less than the true number of particles in the window N). The reference and displaced signals containing the position and intensity information of the particles are written as

$$S_1(\vec{x}, t) = \sum_{i=1}^{N_{sm}} I_i\left(\vec{X}, t; d_{pi}\right), \qquad (3.14)$$

$$S_2(\vec{x}, t + \Delta t) = \sum_{i=1}^{N_{sm}} I_i\left(\vec{X}, t + \Delta t; d_{pi}\right), \qquad (3.15)$$

respectively, where \vec{x} is the location of the interrogation window and d_{pi} is the diameter of the i^{th} particle. In BOS, t refers to the reference image and $t + \Delta t$ is the distorted image, rather than an actual time difference. The cross-correlation between the two signals can be computed using Fourier transforms,

$$R_{S1,S2}\left(\vec{\eta}\right) = \mathcal{F}^{-1}\left[\mathcal{F}\left[S_1\left(\vec{X},t\right)\right]^* \mathcal{F}\left[S_2\left(\vec{X},t+\Delta t\right)\right]\right],\tag{3.16}$$

where $[]^*$ is the complex conjugate and $\vec{\eta}$ represents the position in the space of possible particle displacements. The displacement of each particle is given by $R_{ii}\left(\vec{\eta} - \Delta \vec{X}\left(\vec{X}_i(t)\right)\right)$ (which can be used for particle tracking), and the most likely displacement of particles in the interrogation window is given by a peak in the correlation plane,

$$\Delta \vec{X} = \underset{\eta}{\arg\max} R_{S1,S2}\left(\vec{\eta}\right). \tag{3.17}$$

The particle intensity distributions can be modelled by a Gaussian function, and a Gaussian function can be fitted to the correlation function to determine the window displacement to a sub-pixel level. An example of the interrogation window and correlation plane from experimental speckle images is shown in figure 3.5. As well as a strong peak just off-centre corresponding to a displacement magnitude of 2.15 pixels, there are nearby fluctuations which are associated with measurement noise.

Some rules-of-thumb have emerged for optimising the accuracy of the cross-correlation method for PIV and measurements [39, 43, 106, 145, 152] that may be pertinent to BOS measurements as well, including:

- The displacements should not exceed half the window size, otherwise many particles will exit the window. The displacements in BOS are typically much smaller than those in PIV, e.g. 1–2 pixels in an 8- or 16-pixel square window, so this is usually not a problem.
- The particle size should be at least 2 pixels in diameter to avoid peak-locking, which results in the displacements incorrectly tending to integer values.
- There should be at least 7–10 particles per window. They should cover as much of the window as possible, i.e. they should not be clustered.
- A smaller window size results in less spatial averaging but a lower signal-to-noise ratio.

Obviously, not all guidelines for optimising the digital PIV cross-correlation analysis will be relevant to BOS, especially those concerning out-of-plane displacements, which do not exist in BOS measurements.

The principles of optical flow are unrelated to those of cross-correlation. Optical flow methods are based on the assumption that the intensity of moving structures is constant



Figure 3.5: A 16×16 pixel cross-correlation interrogation window of a laser speckle background pattern with a heated jet. Axes coordinates are in pixels. a) Reference image interrogation window, b) displaced image interrogation window, c) correlation plane (zero displacement origin marked with yellow cross). A clear, white peak is visible in the correlation plane (marked with red cross), as well as measurement noise, indicating a displacement of 1.35 pixels to the right and 1.67 pixels downwards (magnitude 2.15 pixels).

from shot to shot [13, 31]. The intensity I of an image region between two shots must follow

$$I\left(\vec{X},t\right) = I\left(\vec{X} + \Delta \vec{X}, t + \Delta t\right).$$
(3.18)

Like cross-correlation for BOS, t and $t + \Delta t$ do not refer to time, but instead refer to the reference and displaced images, respectively. From a Taylor series expansion of the left-hand side of equation 3.18 with second- and higher-order terms neglected, and after some algebraic manipulation, Davies [31] derives the fundamental equation of optical flow:

$$\frac{\partial I}{\partial X}\frac{\Delta X}{\Delta t} + \frac{\partial I}{\partial Y}\frac{\Delta Y}{\Delta t} + \frac{\partial I}{\partial t} = 0.$$
(3.19)

The intensity gradients $\partial I/\partial X$ and $\partial I/\partial Y$ are obtained from the BOS images, but this equation has two unknown displacement components ΔX and ΔY . To obtain a unique solution, additional equations are required. Atcheson et al. [7] compared three implementations of optical flow using synthetic BOS images: Lucas-Kanade [80], Horn-Schunck [57] and Brox [20]. The Lucas-Kanade algorithm considers an interrogation window around each pixel and solves an overdetermined system of optical flow equations in a least-squares sense. The Horn-Schunck algorithm adds regularisation to the optical equation to enforce smoothness while minimising a cost function. The Brox algorithm assumes constancy in other properties related to the brightness such as the Laplacian and seeks to minimise an energy function that incorporates this additional information.

Grauer and Steinberg [47] took advantage of the fact that the optical flow displacements can be represented in terms of the gradients of intensity to incorporate the optical flow displacement calculation into the camera calibration and BOS tomographic reconstruction. This creates a 'unified' one-step reconstruction pipeline, where the BOS images are used as input, and the density field is obtained with no intermediate intervention from users.

Although Atcheson et al. [7] report higher accuracy using the optical flow algorithms on synthetic BOS images than cross-correlation, one must consider that all the optical flow algorithms contain a tuning parameter that must be optimised to obtain the correct displacements. The authors note that this is not a straightforward task for experimental data. This uncertainty, and the constant intensity assumption, lead to the selection of cross-correlation over optical flow in this investigation. In particular, Atcheson et al. [7] discouraged the use of fluorescent lighting due to flickering and the associated shotto-shot variations in intensity. The present work implements laser speckle BOS with a high-power pulsed laser, for reasons that will become clear in the following chapters, and this also produces shot-to-shot variations in illumination intensity. Therefore, although it may have higher accuracy than cross-correlation in some situations, optical flow is unsuitable for the current application as the experimental images would violate its basic constant intensity assumption.

Rajendran et al. [108] developed a particle-tracking method which utilised the known location of dots in the reference background image to improve the accuracy and dynamic range of the displacement calculation compared to cross-correlation alone. The particle identification and centroid location is done using an image segmentation algorithm, which is then refined using an additional cross-correlation in a tight window around each particle. The authors report that the optimised method has a noise floor three times smaller than cross-correlation and a four-fold higher dynamic range, in terms of sensitivity to minimum and maximum detectable refractive index gradient, from tests using synthetic images. However, the authors also disclose that the method is not ideal for background patterns with high-density, overlapping or greyscale dot patterns (rather than black dots on a white background). The performance of this method with regards to speckle patterns was left as an open question. For this reason, the present study will use the traditional digital cross-correlation PIV analysis.

3.3 Tomographic reconstruction of the refractive index field

In general, the aim of tomographic reconstruction is to solve for a distribution g(x, y, z) given path-integrated projections P of the field along several lines running through the field [62]. For a 2D field with 1D projections, each line (ray) is described by the angle θ and distance t from a global origin, and could represent the pixels in a sensor:

$$P_{\theta}(t) = \int_{s} g(x, z) \,\mathrm{d}s, \qquad (3.20)$$

or if the rays are perpendicular to the sensor plane (they need not be),

$$P_{\theta}(t) = \int \int_{\mathbb{R}^2} g(x, z) \hat{\delta}(x \cos \theta + z \sin \theta - t) \, \mathrm{d}x \, \mathrm{d}z \tag{3.21}$$

where ds is a differential length along the line and $\hat{\delta}$ is the Dirac delta function. This is illustrated in figure 3.6. The extension to three-dimensional data is similar, but a 3D reconstruction can also be achieved by stacking multiple 2D planes of the field g. Only the projections P are available, but g is desired; this is called an 'inverse problem'. Notice that equation 3.20 has a similar form to equation 3.2, if $P_{\theta} \equiv \tan \varepsilon_{x_r}$ and $g \equiv \frac{1}{n_0} \frac{\partial n}{\partial x_r}$. Obtaining the three-dimensional refractive index gradients, or refractive index field itself, from path-integrated BOS displacements is therefore an inverse problem that



Figure 3.6: Principle of tomographic reconstruction. A ray passing through the field of interest g generates a projection on the sensor plane P_{θ} . The rays need not be perpendicular to the sensor plane. Adapted from Kak and Slaney [62].

can be achieved through tomographic reconstruction. The projections P are the Radon transform of the field g. It is mathematically possible to obtain g using the inverse Radon transform given an infinite number of perfect projections, but this is obviously not practical under real constraints such as a limited camera number and measurement noise. In practice, other algorithms are used which account for these limitations. There are essentially two classes of reconstruction algorithms: analytical methods, and iterative methods. Analytical methods are related to the inverse Radon transform and use the Fourier slice theorem to relate the projections to the desired field. Iterative methods, broadly, discretise the field into a basis such as pixels (in 2D) or voxels (in 3D), and seek to describe the contribution of the voxels to the line integral in a linear system of equations such that they produce measured projections. The following sections will discuss the principles of a popular analytical method, filtered back-projection (FBP), and an analytical method, algebraic reconstruction technique (ART), in a general sense. In chapter 4, implementations of these techniques specific to TBOS will be presented.

Note that this project is focussed on developing reconstruction methods for full threedimensional density measurements. As such, methods for axisymmetric BOS measurements will not be discussed in detail. Significant inroads have been made for axisymmetric BOS measurements such as the time-averaged density field of a round jet by Tan et al. [136] and Xiong et al. [153]. These measurements can utilise classical axisymmetric reconstruction methods closely related to the inverse Radon transform such as reduced back-projection, the onion-peeling method and three-point Abel transform [30]. In recent years, specialised axisymmetric BOS reconstruction techniques have been developed and validated, such as the adaptive Fourier-Hankel Abel algorithm [136], and an 'indirect' approach which first relates the BOS deflection angles to the path-integrated density which is then subjected to the Abel transform [153].



Figure 3.7: Illustration of Fourier slice theorem. The Fourier transform S of the projections of angle θ with discrete data points t fills a line in the 2D frequency domain of the object field g. Under-sampling of higher frequencies relative to lower frequencies is evident as the increasing distance between adjacent data points from different projections.

3.3.1 Filtered back-projection (FBP)

Filtered back-projection relies on the Fourier slice theorem, which states the Fourier transform of a projection, denoted as S, is equivalent to a slice through the Fourier transform of the object field f, i.e.

$$S_{\theta}(w) = \int_{\mathbb{R}} P_{\theta}(t) \exp\left[-2\pi i w t\right] dt$$

= $\int \int_{\mathbb{R}^2} g(x, z) \exp\left[-2\pi i w \left(x \cos \theta + y \sin \theta\right)\right] dx dy,$ (3.22)

where the first line of the equation is the Fourier transform of the projection, the second line is the Fourier slice theorem, and w is a Fourier-space frequency variable. This implies that the field g can be recovered by taking the inverse 2D Fourier transform of all the projections. Unfortunately, reconstruction of the true field requires an infinite number of projections in θ and along t in each θ . Using a limited number of projections results in an incomplete 2D frequency space representation of the object, as shown in figure 3.7. The lower frequencies (small u and v) are sampled more than the higher frequencies, which results in some blurring of the object. Intermediate values require interpolation, which implies that there is a higher error at higher frequencies.

Attempting to reconstruct the object from the noisy, limited-view data will result in a corrupted reconstruction. Filtered back-projection is a modification to this process that attempts to deal with these shortcomings by first filtering the data in frequency space,

$$Q_{\theta}(t) = \int_{\mathbb{R}} S_{\theta}(w) |w| \exp\left[2\pi i w t\right] \,\mathrm{d}w, \qquad (3.23)$$

before inversion, where |w| is a frequency-space filter that regularises the frequency information. Often a simple ramp function is used as the filter. The field is finally obtained by integrating over all the filtered projections,

$$g(x,z) = \int_0^{\pi} Q_\theta \left(x \cos \theta + y \sin \theta \right) \, \mathrm{d}\theta. \tag{3.24}$$



Figure 3.8: Discretisation of the reconstruction volume into voxels for ART. The i^{th} ray passes through all voxels marked with a star, i.e. these voxels have a non-zero weight w_{ij} for this ray (dashed outline). Red stars (red outline) are voxels containing the field, while grey stars (black outline) mark voxels that do not contain the field. The latter can be prevented from being updated during the reconstruction process through masking to speed up computations. Like FBP, the rays need not be perpendicular to the sensor plane.

The integration range assumes that all the projections were captured in a semicircle, but this does not have to be the case. The last step is known as back-projection and dictates that each point t in the projection has an equal contribution to all of the reconstructed points along the ray leading to point t. A common description of back-projection is that the projection is smeared back along the object [62], which blurs the field further.

The entire FBP reconstruction can be performed very efficiently by making use of Fast Fourier Transforms. TBOS reconstructions of the density field via FBP have been performed by Goldhahn and Seume [45], Schröder et al. [122], Venkatakrishnan and Meier [143] and Hartmann and Seume [53]. Goldhahn and Seume [45] used a 'one-step' reconstruction where the deflection angles are used to reconstruct the density field directly, eliminating the need to solve the Poisson equation. Although FBP is orders of magnitude faster than an equivalent iterative reconstruction, a key observation is the formation of reconstruction artefacts around the object, e.g. the jet core, due to the inadequate number of views and measurement noise, and the associated errors at higher frequencies as discussed above. The reconstruction improves as the number of views is increased, but the artefacts can easily be confused with genuine flow features. This will become apparent in the next chapter when the BOS implementation of FBP is introduced.

3.3.2 Algebraic reconstruction technique (ART)

Iterative tomographic reconstruction is very different from the analytical methods, as the goal is to solve a system of equations that relate the volume points to the recorded projections. The main advantages of ART over FBP are flexible camera placement and a more accurate solution in limited-view tomography [62]. ART discretises the field g into j discrete cells, as shown in figure 3.8. The ray corresponding to a projection i is also assumed to have a finite width as it traverses the volume. Equation 3.20 can be rewritten as

$$P_i = \sum_j w_{ij} g_j, \tag{3.25}$$

where w_{ij} is the weight (contribution) of the j^{th} voxel to the i^{th} ray. There are many choices of weights, including: binary box function $w_{ij} = 1$ if any part of the ray crosses the voxel or 0 otherwise [54], radially symmetric $w_{ij} = \max(0, 1-b)$ where b is the distance between the ray centre and voxel centre [6], linear radial basis $w_{ij} = \max(0, 1-b/r)$ where r is the radius of the voxel [9], and a trilinear basis function [59]. As each ray will only pass through a few voxels, the weight matrix w is typically sparse. The tomographic reconstruction has been recast as the solution to the system of equations, but an illposed one. According to Hadamard [102], a system of equations is well-posed if the solution exists, is unique, and is stable. Tomographic reconstruction cannot meet the second condition as the system is under-determined (the number of voxels far exceeds the number of rays), and the presence of measurement noise can create an unstable solution.

ART solves the system of equations using Kaczmarz's algorithm [54, 62],

$$g_j^{k+1} = g_j^k + \lambda_j \frac{P_i - \sum_j g_j^k w_{ij}}{\sum_j w_{ij}^2} w_{ij}, \qquad (3.26)$$

where superscript ^k refers to the iteration number and λ_j is a relaxation factor. The aim is to update the field to minimise the difference between projections through the reconstructed field and the actual recorded projections. ART requires an initial guess, which may arbitrarily be set to a zero field or incorporate information about the characteristics of the field known prior to reconstruction, e.g. values must be non-negative in intensitybased measurements (such as in medical imaging). One must be careful in setting the prior characteristics, for example positive and negative values of the gradient are valid in TBOS refractive index gradient reconstruction. The weighting matrix w_{ij} means that only voxels intersected by a ray are updated, by smearing the second term in the above equation along all of the intersected voxels like back-projection. As discussed by Tanabe [137] and Kak and Slaney [62], ART converges to a solution which is nearest to the initial solution, i.e. the solution is not unique. This implies that there is scope to improve the accuracy of the reconstruction through careful choice of the initial solution, which is considered in the next chapter. ART reconstructs lower spatial frequency features in earlier iterations, and higher frequencies are updated in later iterations. This means that the general (large-scale) features of a field can be obtained with only a few iterations. In the presence of measurement noise, the ART algorithm has a regularising effect on the solution as the iterations progress [34].

TBOS reconstructions using ART have been conducted by Atcheson et al. [6] and Lang et al. [72]. Atcheson et al. [6] used synthetic displacements of a heated jet CFD simulation created via ray tracing as the basis for ART reconstruction, and reported good agreement between their simulated BOS and ground-truth CFD. Lang et al. [72] also considered the effect of camera number and placement and voxel size using a synthetic case. Their experimental temperature field measurements in a swirling jet matched thermocouple measurements in the shear layer to within 1 K, although the centreline temperature was under-predicted by up to 16 K due to excessive defocus blurring reducing the measured gradients.

A noteworthy modification to the basic ART scheme is simultaneous ART (SART) [5]. During each iteration, ART makes corrections to each voxel multiple times based on the number of rays crossing the voxel. SART corrects each voxel once per iteration based on the information from all projections:

$$g_j^{k+1} = g_j^k + \lambda_j \frac{\sum_i \left[w_{ij} \frac{P_i - g_j^k w_i}{\sum_j w_i} \right]}{\sum_i w_{ij}}$$
(3.27)

This yields a faster convergence than ART while improving its noise suppression characteristics. This survey of ART algorithms is not exhaustive, and other approaches such as multiplicative ART (MART) are used extensively in other fields of study such as tomographic PIV due to superior particle reconstruction ability [8]. But these approaches may not be suitable for TBOS. For example, the denominator in the MART algorithm,

$$g_j^{k+1} = g_j^k \left(\frac{P_i}{\sum_j w_{ij} g_j^k}\right)^{\lambda w_{ij}},\tag{3.28}$$

may be zero due to the addition of positive and negative gradient contributions along a ray. As such, the next chapter concentrates on implementations of modified standard ART and SART for TBOS measurements.

3.3.3 Other iterative reconstruction techniques for BOS

As discussed, the tomographic reconstruction is an ill-posed problem. As well as using a 'one-step' reconstruction where the density field is obtained directly from the BOS deflections, Nicolas et al. [96] approached TBOS from the point of Tikhonov regularisation which aims to solve a well-conditioned system in place of the ill-conditioned system [138]. They rewrite the system of equations,

$$\vec{\varepsilon} = \vec{w} \vec{D} \rho, \tag{3.29}$$

where \vec{w} is a weight matrix as before and \vec{D} is a finite difference operator which allows the density to be related to the refractive index gradients and, hence, the deflections. The reconstruction is the weighted least-squares solution to a system describing the observed deflections including a regularisation parameter λ_r which enforces smoothness on the density field,

$$\mathcal{J}(\rho) = ||\vec{w}\vec{D}\rho - \varepsilon||^2 + \lambda_r \mathcal{R}(\rho), \qquad (3.30)$$

$$\mathcal{R}(\rho) = ||D'\rho||^2 = -\rho^{\mathrm{T}} \nabla^2 \rho.$$
(3.31)

The system is solved using conjugate gradient minimisation method [32]. The optimal value of λ_r is selected using an L-curve criterion by plotting $||D'\rho||^2$ as a function of $||\vec{w}\vec{D}\rho - \varepsilon||^2$ with several λ_r . The point on the curve with highest curvature indicates a solution with minimal error compared to a ground-truth case. The L-curve must be constructed for each camera setup. Nicolas et al. [96] present detailed 3D reconstructions of the density field of a candle plume and heat gun, and this initial study was supported by further measurements in a supersonic jet [97]. It is an interesting approach to TBOS

reconstructions, but one that offers no apparent advantage in accuracy over the ART reconstructions. However, the 'one-step' reconstruction may be more convenient and faster than an equivalent two-step reconstruction, by virtue of having fewer computations. Furthermore, Grauer and Steinberg [47] reported that their 'unified' method, which also incorporates displacement calculation into the reconstruction, was more than 60% faster than their previous one-step reconstruction which still had a separate displacement calculation step.

The studies by Grauer [48, 49] use a Bayesian framework tomographic reconstruction. This approach uses a one-step method like Nicolas et al. [96] with a simultaneous iterative reconstruction technique (SIRT) reconstruction. Bayesian statistics are used to update the reconstruction by marrying the observed data with prior knowledge of the process using Bayes's equation

$$\Pi(\vec{g}|\vec{P}) = \frac{\Pi(\vec{P}|\vec{g})\Pi_{prior}(\vec{g})}{\Pi(\vec{P})} \propto \Pi(\vec{P}|\vec{g})\Pi_{prior}(\vec{g}), \qquad (3.32)$$

where $\Pi(\vec{g}|\vec{P})$ is the posterior probability density function (PDF), $\Pi(\vec{P}|\vec{g})$ is the likelihood PDF, $\Pi(\vec{P})$ is the evidence PDF and $\Pi_{prior}(\vec{g})$ is the prior information PDF. The reconstructed field is chosen to be $\vec{g}_{max} = \arg \max \Pi(\vec{g}|\vec{P})$. The prior information PDF is formulated to incorporate Tikhonov regularisation for a smooth solution, and total variation regularisation so that sharp variations in density, e.g. at shocks and flame fronts, are still permitted in the smoothed solution. The likelihood and evidence PDFs contain the linear system to be solved by a least squares approach. This approach was able to reconstruct the large-scale features in a flame simulation, but appeared to struggle with smaller scale features.

3.3.4 Integration of the reconstructed gradients

After tomographic reconstruction of the three 3D refractive index gradient components ∇n , the refractive index field itself is obtained by solving a Poisson equation (equation 3.8). Both sides of the equation can be discretised using finite differences, e.g. second-order accurate central differences [72] with first-order discretisation at the boundaries [6] or the use of ghost points as the ambient value is known so Dirichlet boundary conditions can be employed. The impact of different discretisation schemes will be examined in the next chapter. Atcheson et al. [6] suggested an anisotropic diffusion scheme where the right-hand side source term q is multiplied by a diffusion tensor. This has the effect of regularising noise while not over-smoothing edges. This scheme will be explained in more detail, and evaluated, in chapter 5.

The discretised equation is to be solved at all points, leading to a linear system of equations of the type $\vec{A}\vec{x} = \vec{b}$ that is sparse and positive definite [6]. Various iterative techniques have been used to solve for the refractive index field, including successive over-relaxation [143] and conjugate gradients [6, 72]. Demmel [32] provides a comparison of the complexity for solving the Poisson equation on a grid with N points with various direct and iterative methods. The standard Gauss-Seidel direct method requires $\mathcal{O}(N^2)$ operations to solve the system. The successive over-relaxation and conjugate gradient iterative methods both require $\mathcal{O}(N^{3/2})$ operations. The fastest direct method is block cyclic reduction, which requires $\mathcal{O}(N \log N)$ operations. The iterative multigrid methods are the optimal iterative solver for linear systems, delivering the solution for a system of N points in just $\mathcal{O}(N)$ operations. In this project, a multigrid method is preferred for this reason.

Instead of solving the Poisson equation, Rajendran et al. [107] proposed a weighted least squares integration of the refractive index gradients. The grid points are weighted based on the local measurement uncertainty in the density gradients. This uncertainty itself is related to the uncertainty in the displacement calculation, which can be estimated using various methods developed for PIV [14, 123] and their earlier work for BOS [110, 111]. The integration is then formulated as an optimisation problem which requires minimising a cost function based on the difference between the reconstructed gradients and the finite-difference gradients of the solved refractive index field. This is a promising method of obtaining the refractive index field that can limit the spread of measurement noise from a localised region, e.g. due to uneven illumination. In such a situation, the Poisson equation may instead propagate the noise over a larger volume. However, the authors note that although their experimental demonstration reduced the uncertainty in the measurement compared to the Poisson solver, they were not able to show that the error itself was reduced.

Chapter 4

A parametric study of TBOS methods in a fluctuating density field

Part 1 of the research road map. Development of reconstruction methods using a heated jet phantom.

4.1 Aims and overview of the chapter

The initial stage of this project is focussed on developing, optimising, validating, and comparing several tomographic reconstruction algorithms using a synthetic density field phantom (test case). Introduced in section 4.2, this investigation will compare four 'classical' methods: filtered back-projection (FBP), algebraic reconstruction technique (ART), simultaneous ART (SART) and a hybrid FBP+ART method. The phantom is modelled on the self-similar region of a heated jet modulated with sinusoidal variations to mimic turbulence fluctuations, which allows the range of scales resolved by each method to be determined. Using the ray tracing procedure discussed in section 4.3, synthetic background displacements are generated for a virtual camera setup with a variable number of cameras. The synthetic displacements are used as input to the reconstruction methods to assess their accuracy while modifications are made to the basic algorithms to increase the reconstruction quality.

Results are presented in section 4.5. Using the optimised methods, this study determines the appropriate number of cameras for an experimental TBOS setup. As well as optical limits to the range of flow scales that can be resolved in the BOS measurement, the reconstruction algorithm will have a limit on the scales that can be measured. This study investigates the range of scales that can be reconstructed faithfully up to the Nyquist frequency (dictated by grid spacing), assuming that the defocus blurring is small relative to reconstruction grid spacing. Fluctuations of increasingly higher spatial frequency are imposed on the density field to determine the performance of the TBOS methods by comparing the true and reconstructed power spectra of the phantom. This chapter will also examine the characteristics of different discretisation schemes in terms of resolvable scales and propagating measurement noise in the reconstruction.
4.2 Implementation of a tomographic BOS reconstruction

To perform a TBOS reconstruction from an arrangement of cameras, it is necessary to relate the background displacements computed in each camera to the density gradients in the global coordinate system. The global position and orientation of each camera is determined using a standard camera calibration procedure [157], which populates the pinhole matrices from equations 3.11 and 3.12. Background displacements in each camera's image plane can be calculated using the methods discussed in the previous chapter.

Following figure 3.1, the deflection of the light ray from the centre of each interrogation window, $\vec{X_o}$, is determined by using the camera model (equation 3.13) to project the nonrefracted ray from the the image plane to the corresponding point on the background $\vec{P_o}$. The point where this ray crosses the centre of the measurement volume is denoted by $\vec{I_o}$. This ray is represented by a vector $\vec{x'} = I_o \vec{P_o}$ in global coordinates, with local axes orthogonal to this ray denoted as $\vec{y'}$ and $\vec{z'}$. A second ray is projected from the displaced point in the image plane $\vec{X_o} + \Delta \vec{X}$ to the background point $\vec{P_r}$. The deflection of the non-refracted ray is therefore the angle between the vector $\vec{x'}$ and the vector $\vec{r} = I_o \vec{P_r}$. Using the small angle approximation, the estimated angles $\varepsilon_{x'}, \varepsilon_{y'}$ and $\varepsilon_{z'}$ about the local axes become

$$\varepsilon_{x'} = \vec{r} \cdot \vec{x'}, \tag{4.1}$$

$$\varepsilon_{y'} = \vec{r} \cdot \vec{y'}, \tag{4.2}$$

$$\varepsilon_{z'} = \vec{r} \cdot \vec{z'}. \tag{4.3}$$

In this project, tomographic reconstruction of the refractive index gradient fields $\frac{\partial n}{\partial x}(x, y, z)$, $\frac{\partial n}{\partial y}(x, y, z)$ and $\frac{\partial n}{\partial z}(x, y, z)$ fields that correspond to the measured background displacements and associated ray deflections will be performed using 'classical' methods, either Fourier-slice based filtered back-projection (FBP) [62] or iterative algebraic reconstruction techniques (ART) [54]. As discussed in the previous chapter, FBP is a popular technique in the medical field where a single sensor is often rotated in a single plane around a stationary object, which allows the recording and reconstruction of hundreds of slices through a volume. The time scales of turbulent flows require that the projections are recorded simultaneously using individual cameras. Owing to cost, and the need to physically package these systems in experimental facilities, the number of projections is necessarily limited. ART typically behaves better in the case of limited projections and requires no fundamental restrictions on camera spacing and placement, however its iterative nature makes it far more computationally expensive. This study will consider the use of both methods for TBOS density measurements in heated jets.

4.2.1 Filtered back-projection (FBP)

Filtered back-projection makes use of the inverse Radon transform to relate a sinogram, representing the sum of a scalar quantity through a volume along the line of sight of a sensor as a function of the position and angle of the sensor, to the corresponding distribution of the scalar within the volume [62]. Assuming each camera records parallel projections, and the sensor resolution remains constant across all angles, this transformation can be performed rapidly using Fast Fourier Transforms. Presently, the reconstruction of each refractive index gradient is performed independently based on sinograms that represent

the sum of each component of the refractive index gradient, $\sum_{ray} \nabla n$, along the camera's axis at each position along the camera's interrogation windows. In this case, a camera corresponds to a different projection angle. The sum of the gradients at each window are determined by solving the following system of equations

$$\begin{bmatrix} \varepsilon_{x'} \\ \varepsilon_{y'} \\ \varepsilon_{z'} \end{bmatrix} = \frac{1}{n_0} \begin{bmatrix} \vec{x'}^T \\ \vec{y'}^T \\ \vec{y'}^T \\ \vec{z'}^T \end{bmatrix} \sum_{ray} \nabla n \Delta x$$
(4.4)

where n_0 is the refractive index outside the measurement volume and Δx is the voxel width. Reconstruction is performed following the common practice of taking the inverse Radon transform of ramp-filtered sinograms. A cylindrical reconstruction domain is applied which corresponds to the common view of all cameras, outside which the density gradients are set to zero.

4.2.2 Algebraic reconstruction technique (ART)

Algebraic reconstruction is based on representing a series of projections or views of a volume P_i in terms of a weighted contribution from a discretised point in the volume I_j such that the final reconstruction should satisfy,

$$P_i = \sum_j w_{ij} I_j, \tag{4.5}$$

where w_{ij} is the weighted contribution of the *j*-th point in the volume to the *i*-th ray seen by one of the cameras. These algorithms iteratively update I_j based on the difference between the estimated projections $\sum_j w_{ij}I_j$ and the measured projections P_i of the volume. In TBOS, P_i represents the measured deflections of the refracted light rays, which needs to be formulated in terms of the refractive index gradients. This can be done following equation 3.2, such that the deflection of the *i*-th ray in the local ray coordinate system that results from the *k*-th iteration of the reconstruction of the gradient field in the volume is given by the integral of the weighted dot product of the vector of the local axis in global coordinates with the gradient vector of the refractive index field ∇n

$$\varepsilon_{x'i}^k = L_i \frac{\sum_j w_{ij} \vec{x'} \cdot \nabla n_j^k}{n_o \sum_j w_{ij}}, \qquad (4.6)$$

$$\varepsilon_{y'i}^k = L_i \frac{\sum_j w_{ij} \vec{y'} \cdot \nabla n_j^k}{n_o \sum_j w_{ij}}, \qquad (4.7)$$

$$\varepsilon_{z'i}^k = L_i \frac{\sum_j w_{ij} \vec{z'} \cdot \nabla n_j^k}{n_o \sum_j w_{ij}}, \qquad (4.8)$$

where L_i is the length of the path followed by the ray through the measurement volume. Following Atkinson and Soria [9], the weighted contribution of each point in the volume j to a ray i from a given interrogation window depends on the volume intersection of the voxel and the projection of the window, which can be approximated by representing the voxel as a sphere of equivalent volume and the projection of each ray as a cylinder with equivalent cross-sectional area. This allows the contribution to be parameterised in terms of each radius and the shortest distance between the ray and the centre of each voxel. This volume can be calculated analytically [71], or can be approximated and more rapidly computed by a linear radial basis:

$$w_{ij} = \max\left(0, 1 - \frac{b}{r}\right),\tag{4.9}$$

where b is the shortest distance between the ray and the centre of the voxel and r is the radius of the voxel. Using the full analytically derived volume intersection had negligible effect on the reconstruction accuracy when compared to the radial basis approximation above.

The iterative correction of the density gradients at each point in the volume is performed one ray at a time by relating the difference between the measured ray deflections to the projection, $\varepsilon_{x'i} - \varepsilon_{x'i}^k$, to the required correction in each of the refractive index gradient components $\nabla n_j^{k+1} - \nabla n_j^k$ following

$$\frac{\lambda_j n_o w_{ij}}{L_i} \begin{bmatrix} \varepsilon_{x'i} - \varepsilon_{x'i}^k \\ \varepsilon_{y'i} - \varepsilon_{y'i}^k \\ \varepsilon_{z'i} - \varepsilon_{z'i}^k \end{bmatrix} = \begin{bmatrix} \vec{x'} \\ \vec{y'} \\ \vec{y'} \\ \vec{z'} \end{bmatrix} \left(\nabla n_j^{k+1} - \nabla n_j^k \right), \tag{4.10}$$

where λ_j is a relaxation parameter associated with the reconstruction. This correction is based on a standard additive ART correction [54]. While multiplicative algorithms are favoured in limited-view tomographic PIV and medical imaging due to their lower noise [8], the reconstruction of density gradients differs from intensity-based reconstructions because the gradient is not restricted to a positive value. The summed quantity may also be zero from a contribution of positive and negative gradients along a ray. This precludes the use of reconstruction algorithms which involve normalisation with denominators that could be zero in this application.

Various methods are explored to improve the reconstruction quality including: random ordering of projections; Gaussian filtering of the density fields after each iteration, with filtering relaxed to zero as the final iteration is approached; Hamming windowed correction to penalise the generation of artefacts at the boundaries of the reconstruction; gradual unmasking that restricts corrections below a threshold that tends to zero with progressing iterations [77]; and using FBP as an initial solution to the iterative ART [53]. This work also tests the simultaneous correction of the volume from all rays using simultaneous ART, or SART [5], where the left-hand side of equation 4.10 is now averaged over all rays (and the right-hand side is unchanged),

$$\frac{1}{\Sigma_i w_{ij}} \Sigma_i \left[\frac{\lambda_j n_o w_{ij}}{L_i} \begin{bmatrix} \varepsilon_{x'i} - \varepsilon_{x'i}^k \\ \varepsilon_{y'i} - \varepsilon_{y'i}^k \\ \varepsilon_{z'i} - \varepsilon_{z'i}^k \end{bmatrix} \right] = \begin{bmatrix} \vec{x'}^T \\ \vec{y'}^T \\ \vec{z'}^T \end{bmatrix} \left(\nabla n_j^{k+1} - \nabla n_j^k \right).$$
(4.11)

Details and results of the improvements are covered in section 4.5.3.

4.2.3 Calculation of the refractive index field

Refractive index fields n(x, y, z) can be calculated from the reconstructed components of ∇n by solving a Poisson equation (equation 3.8). The right-hand side term q (equation 3.9) is populated by taking the derivatives of the reconstructed gradient fields. The present work uses a finite difference discretisation to solve this equation numerically from the reconstructed gradient data. The solution to this equation involves a second-derivative finite difference discretisation of the refractive index field (left-hand side), and a first-derivative discretisation of the reconstructed gradients (right-hand side). The accuracy of the solution is therefore dependent on the order of the finite difference scheme used for both sides of the equation, which do not have to be the same. Examples of the finite

difference discretisation for different order schemes can be found in Appendix A.1. The gradients ∇n are reconstructed without the need for any finite difference type gradient calculation and hence, without any associated truncation error (but possibly with defocus blurring from the optical setup). This should result in a lower truncation error from the right-hand side of the Poisson equation than the left-hand side, for the same kernel size, because the input gradients themselves are not computed from finite differences. As higher order difference schemes are used this problem will be diminished, however relatively low order schemes are often required in the analysis of experimental data, in order to reduce the influence of any reconstruction noise, in analogy to what is done with PIV data by Foucaut and Stanislas [38]. The effect of altering the kernel size and the order of the difference scheme will be investigated in section 4.5.1.

A universal Dirichlet boundary condition of n_0 is applied with q = 0 at all boundaries far outside the flow, with an appropriate number of ghost points used for a given discretisation scheme order. The measurement domain is large enough to capture the full extent of all gradients present in the flow. The present chapter is restricted to examining a 2D slice transverse to the jet axis.

The discretised Poisson equation is solved using an algebraic multigrid method [98]. As discussed in section 3.3.4, multigrid methods are the optimal iterative solver for linear systems. The multigrid solver runs until the iterative residual to n is $\mathcal{O}(10^{-16})$.

4.3 Numerical validation procedure

To assess the ability of the different reconstruction methods to resolve small scale density fluctuations in turbulent heated jets, synthetic BOS displacements were created by tracing rays from each camera using the pinhole model through a known refractive index distribution.

The direction of the ray is updated at each step using Snell's law where the refracted vector \vec{t} is given by:

$$\vec{t} = \frac{n_1}{n_2}\vec{i} + \left(\frac{n_1}{n_2}\left(\vec{i}\cdot\vec{n}\right) - \sqrt{1 - \left(\frac{n_1}{n_2}\right)^2\left(1 - \left(\vec{i}\cdot\vec{n}\right)^2\right)}\right)\vec{n},\tag{4.12}$$

$$n_2 = n_1 + \nabla n \cdot \hat{i}, \tag{4.13}$$

where \tilde{i} is the unit vector of the incident ray, \tilde{n} is the unit vector of ∇n and n_1 is the local refractive index of the incoming ray [90]. Upon exiting the measurement volume, the final ray direction is maintained and projected onto the background plane, from which the displacement from the projected un-refracted ray is obtained. The use of Snell's law, rather than the integration of density gradients along a fixed ray, allows for the inclusion of the variations in the ray path which could affect the accuracy of the TBOS reconstruction in the presence of strong localised refractive index variations. This analytical refractive index distribution could also be used to explore the influence of the exact refractive index distribution and the influence of the circle of confusion on fluctuation scales down to Nyquist, however as shown in subsequent sections fluctuation begins to be poorly resolved even at twice the Nyquist wavelength.

Figure 4.1 presents the investigation methodology for chapters 4 and 5. Starting from a known density field, the synthetic displacements are generated using ray tracing.



Figure 4.1: Possible TBOS investigations (shaded) using ray-tracing of the synthetic and DNS heated jet fields (**True solution**). This work focuses on investigating **spatial averaging** (chapter 5), **temporal averaging** (chapter 5) and **reconstruction methods** and **Poisson equation** (chapters 4 and 5), shown with darker shading.

The synthetic displacements are used to investigate the error of each step in the TBOS measurement process in isolation from the other steps. The present chapter focuses on optimising the tomographic reconstruction. The next chapter will investigate the impact of defocus blurring, temporal blurring, spatial averaging from solving the Poisson equation using a test case more representative of experimental measurements, a heated jet DNS. Previous studies have also investigated the displacement calculation methods using synthetic images in a similar fashion [7, 108, 109], and so this aspect will not be covered in this project. The accuracy presented here assumes a perfect calibration, without distortion or displacement error, and hence, the best-case input to the tomographic reconstruction.

4.4 A synthetic fluctuating density field test case: the heated jet phantom

This study uses a phantom (test case) based on the far-field of an air jet exiting from a round orifice with a time-averaged Gaussian temperature distribution. The jet allows the implementation of simple boundary conditions in the Poisson equation because the



Figure 4.2: Schematic of the synthetic TBOS setup, with jet flowing into the page. Cameras are positioned circumferentially to the jet axis. Only two cameras are shown for simplicity; up to 22 cameras positioned in a 180° arc are tested.

flow far outside is quiescent with known properties. The jet phantom has a peak change in refractive index of $\Delta n_p = n_0 - n_{centreline} = 1.5 \times 10^{-4}$ from the ambient refractive index $n_0 = 1.000293$, corresponding to a centreline temperature of approximately 286°C at standard atmospheric temperature and pressure according to the following relation for ideal gases derived from the ideal gas law and Gladstone-Dale relation,

$$n = 1 + (n_0 - 1) \frac{P}{P_0} \frac{T_0}{T},$$
(4.14)

where P is pressure, T is absolute temperature and subscript $_0$ refers to the ambient (reference) value. It is assumed that the jet does not deviate from ambient pressure, which should hold for a low-speed heated jet [25]. The synthetic jet will be modelled on the properties of a low-density jet at a downstream distance of x/D = 25, where it is fully turbulent. The geometric arrangement was chosen such that the interrogation window of $d_{iw} = 16$ pixels would correspond to 16 voxels (0.66 mm) when projected to the volume centre using

$$\delta_{IW} \approx d_{IW} \left(\frac{Z_A - f}{f}\right),$$
(4.15)

and would remain larger than the BOS resolution estimated by equation 3.3, so that the resolution is limited by the displacement interrogation window size. The measurement domain of $65 \times 3 \times 65$ voxels ($42.9 \times 1.98 \times 42.9$ mm) was sized based on a jet diameter of D = 2 mm and $3.75 \times 3.75 \ \mu\text{m}^2$ pixel cameras equipped with f = 25 mm focal length lenses at an aperture of $d_a = f/22$. The optical centre of the camera is positioned 275 mm from the centre of the volume and 575 mm from the background. The spatial resolution at the jet is $0.04125 \ \text{mm/px}$. The synthetic fields were generated at the same resolution to remove any effects of spatial filtering from the analysis, i.e. the volume discretisation Δx matches the optical resolution of the BOS system δ and δ_{IW} . Synthetic background

images were created for different camera numbers, with cameras evenly spaced in a 180° arc about the jet axis in the transverse x-z plane, as shown in figure 4.2. This camera configuration is practical to set up and lends itself well to FBP. It also allows a more efficient comparison of the different reconstruction methods by limiting our region of interest to a thin slice normal to the jet axis.

Because a primary goal of this analysis is to assess the ability of TBOS to resolve spatial density fluctuations across multiple wavelengths, the Gaussian refractive index distribution was modulated using sinusoids in the x and z directions as given by:

$$n(x,y,z) = n_o - \Delta n_p \exp\left[-[x^2 + z^2]/(2\sigma^2)\right] \left[1 + A\sin\left(\frac{2\pi}{\lambda_x}x\right)\sin\left(\frac{2\pi}{\lambda_z}z\right)\right], \quad (4.16)$$

where λ_x and λ_z are wavelengths which are varied to represent density fluctuations of differing length scales. Unless otherwise stated, the wavelengths in both directions are set identically, i.e. $\lambda_x = \lambda_z = \lambda_{x,z}$. For the current spatial resolution of 0.66 mm/voxel, the wavelength corresponding to the Nyquist frequency ($\omega = 0.5 \text{ voxels}^{-1}$) is $\lambda_{x,z Nyquist} = 1.32 \text{ mm}$, or $\lambda_{x,z Nyquist} \approx L/32$, where L = 42.9 mm is the domain length in x and z. The standard deviation of the jet's Gaussian profile is $\sigma = 9 \text{ voxels} = 5.94 \text{ mm}$, which is chosen to suit the desired downstream distance of $x/D \approx 25$, based on an approximate relationship between spreading of the half-width $r_{1/2}$ and downstream distance, $r_{1/2}/D = 0.13(x/D)$ from the helium jet of Panchapakesan and Lumley [100], and the relationship between the half-width of a Gaussian and its standard deviation, $r_{1/2} = \sigma \sqrt{2 \ln 2}$. A is the amplitude of the modulation, which unless otherwise stated was set to 0.25. This corresponds to a density field with fluctuations up to 25% of the mean value which decay towards the jet boundaries, like passive scalar fluctuations in a fully developed turbulent round jet [100].

As BOS projections do not record the integrated density gradients along a single ray, but rather the integrated density gradients (denoted with a tilde) along all rays within the finite aperture limit resolution δ of the system, the averaged density gradient is given by,

$$\frac{\partial \tilde{n}}{\partial x}(x,z) = \frac{-\Delta n_p}{\delta^2} \int_{x-\delta/2}^{x+\delta/2} \int_{z-\delta/2}^{z+\delta/2} \exp\left[\frac{-[x^2+z^2]}{2\sigma^2}\right] \cdot \left[\frac{2\pi A \cos\left(\frac{2\pi x}{\lambda_{x,z}}\right) \sin\left(\frac{2\pi z}{\lambda_{x,z}}\right)}{\lambda_{x,z}} - \frac{x\left(A \sin\left(\frac{2\pi x}{\lambda_{x,z}}\right) \sin\left(\frac{2\pi z}{\lambda_{x,z}}+1\right)\right)}{\sigma^2}\right] dz dx, \quad (4.17)$$

$$\frac{\partial \tilde{n}}{\partial z}(x,z) = \frac{-\Delta n_p}{\delta^2} \int_{x-\delta/2}^{x+\delta/2} \int_{z-\delta/2}^{z+\delta/2} \exp\left[\frac{-[x^2+z^2]}{2\sigma^2}\right] \cdot \left[\frac{2\pi A \cos\left(\frac{2\pi x}{\lambda_{x,z}}\right) \sin\left(\frac{2\pi z}{\lambda_{x,z}}\right)}{\lambda_{x,z}} - \frac{z\left(A \sin\left(\frac{2\pi x}{\lambda_{x,z}}\right) \sin\left(\frac{2\pi z}{\lambda_{x,z}}+1\right)\right)}{\sigma^2}\right] dz dx, \quad (4.18)$$

The above equations the density gradient are obtained by taking the analytical partial derivatives of equation 4.16, and integrating across the voxel to represent the spatial averaging of the true local gradients. As δ is chosen to match the voxel size, equations 4.17 and 4.18 correspond to the spatially-averaged gradients at each point in the domain.



Figure 4.3: Convergence of displacement, as a function of sub-grid steps along the ray, at pixel with highest displacement for a camera oriented at $\theta = 45^{\circ}$ to the volume with $\lambda_{x,z} = L/14$. The volume has 65 grid points in each direction.

4.4.1 Convergence of the ray tracing method

The ray tracing method used to generate the synthetic background displacements must be sufficiently accurate to ensure that no significant bias is introduced into the subsequent reconstruction and analysis. In the current ray tracing scheme, each interrogation window in the displacement field corresponds to one ray using the pinhole model described in section 4.3. The path, and hence, the displacement, of a ray converges as the number of steps along the ray through the refractive index volume is increased, because the sampling of the refractive index variations is increased. The number of steps along a ray is increased beyond the number of refractive index grid points by interpolating sub-grid values with three-dimensional linear interpolation. The convergence of rays can be measured in terms of the residual of displacements as a function of step number.

Typical experimental displacement calculation methods, like cross-correlation, offer subpixel accuracy of displacements to the order of 10^{-1} pixels. Additionally, the reconstruction methods were set to cut-off the displacements to the same order, so any further accuracy is not accounted for in the reconstruction. To concretely define convergence of a ray's displacement for our purpose, the ray path is therefore considered to have converged if the displacement residual is smaller than the 10^{-1} pixels criterion. A convergence study was undertaken to determine the required number of steps to reduce the residual beyond 10^{-1} pixels.

The number of steps is chosen based on a study with a camera oriented at $\theta = 45^{\circ}$ to the volume with $\lambda_{x,z} = L/14$. Figure 4.3 shows the variation in displacement at the pixel with the highest recorded displacement. It is observed that by 125 steps, nearly twice as much as the linear dimension of the volume, the displacement has converged to the order of 10^{-3} pixels, which is well beyond the accuracy with which the background displacement can be determined by cross-correlation or optical flow. Therefore, it is specified that at least 125 steps are used in this investigation.

4.5 Results and discussion

The results presented here focus on optimising the reconstruction methods for turbulence measurements in the fully developed region of a heated jet, and assessing the range of turbulence fluctuations that can be faithfully resolved by these methods. This study investigates the effect of camera number, noise, turbulence fluctuation wavelength, and modifications to the reconstruction schemes, to comprehensively tune the TBOS schemes for the desired measurements and assess their capabilities and limitations.

In section 4.5.1, the behaviour of the Poisson solver is examined with respect to the spatial frequency of fluctuations. Even if the gradients are perfectly reconstructed, any error imparted by the Poisson solver will be present in the final measurement. Thus, it is crucial to understand the behaviour of the solver itself and its ability to handle noise in the reconstructed gradients. The Poisson-solved gradients, for four different discretisation schemes, will be compared to the original refractive index field as a function of fluctuation wavelength and added noise. A simplified, one-dimensional model of the Poisson equation is also used to conduct a von Neumann stability analysis, to compare the transfer function of the discretised solver to the analytical transfer function as a function of input fluctuation frequency.

In section 4.5.2, the behaviour of the FBP reconstruction, with no added displacement field noise, is examined as a function of camera number and wavelength. This is used to determine if FBP is suitable for turbulence measurements in the self-similar region of the heated jet, and if so, the minimum number of cameras required for this. The assessment of the reconstruction methods is conducted with noiseless, perfect background displacement fields using 6 to 22 evenly-spaced cameras, with fluctuation wavelengths from $\lambda_{x,z} = 4$ (twice the Nyquist frequency) to 32 voxels, or ~ L/16 to ~ L/2.

Section 4.5.3 evaluates ART in a similar manner; it is shown that quality of the ART reconstruction depends greatly on the modifications made to the basic reconstruction method. Section 4.5.4 assesses the FBP+ART method, with appropriate modification to remove FBP artefacts. A direct comparison of the reconstruction error as a function of position, and the range of scales that can be resolved by the optimised algorithms, is presented in section 4.5.5. The investigation is concluded in section 4.5.6 with an examination of the effect of displacement field noise, with varying strength, on the reconstruction. These findings on the nature of noise propagation through the reconstruction are combined with our understanding of the Poisson solver to recommend an optimal discretisation scheme.

Unless otherwise stated, the reconstruction accuracy is measured in the region up to twice the jet half-width, i.e. $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ (note $2r_{1/2} \approx 2.35\sigma = 21.2$ mm). To measure the quality of the reconstruction schemes, this study will typically use the RMS error $\sqrt{\langle (..._r - ..._s)^2 \rangle}$ and peak error max $|..._r - ..._s|$ in the refractive index gradients ∇n or refractive index field n between the synthetic field (denoted by subscript $_s$) and reconstructed field (denoted by subscript $_r$). Errors in the gradients will be normalised by the local peak gradient ∇n_{peak} , and errors in the refractive index field will be normalised by $\Delta n_p \equiv n_{peak}$.

4.5.1 Finite difference schemes to solve the Poisson equation

In the absence of any random displacement calculation or reconstruction errors (a perfect reconstruction with infinite views), the error in the TBOS reconstruction is a combination

Table 4.1: Central finite-difference schemes used to solve the Poisson equation. In each scheme, each dimension has the same order of accuracy. Left-hand side is abbreviated as LHS, and the right-hand side is abbreviated as RHS. The 'No. points' column refers to the number of points considered by the finite-difference equation in each dimension. The entries in the 'Abbreviation' column will be used to refer to these schemes henceforth.

Abbreviation	LHS		RHS	
	Order	No. points	Order	No. points
(3, 3)	2	3	2	3
(3, 5)	2	3	4	5
(5, 3)	4	5	2	3
(5, 5)	4	5	4	5

of the systematic error from blurring and spatial averaging of the reconstructed gradients, and the truncation of terms in the finite difference discretisation used to solve the Poisson equation. As discussed in section 4.2.3, the refractive index field is recovered by calculating the second derivatives of the reconstructed density field for the right-hand side of the Poisson equation (equations 3.8 and 3.9), that is then solved numerically with a chosen discretisation of the left-hand side. The influence of the finite difference schemes used to evaluate the right-hand side of the Poisson equation and the kernel size used in the discretised solution to the equation, was investigated by using spatially averaged analytical refractive index gradients (from equations 4.17 and 4.18) as input to the Poisson solver. A description of the schemes tested is provided in table 4.1, and the finite-difference equations for each scheme (in 1D for brevity) are provided in the Appendix A.1 in table A.1.

As the finite difference discretisation can amplify noise in the input fields, it is also of interest to see how the final measurement would be degraded by noise in the reconstruction. This is investigated in figure 4.4, by comparing the RMS error between the synthetic refractive index fields and the refractive index fields calculated by the Poisson solver, n_s and n_r , respectively, as a function of fluctuation wavelength and added Gaussian noise for a variety of left- and right-hand side discretisation orders. The standard deviation of added noise to the gradient fields is expressed as a percentage of the peak true gradient, i.e. $\sigma_{noise}/\nabla n_{peak}$, from 0% (no added noise) to 10%. The results presented at each data point are averaged over 100 samples of added random noise. At low additive noise levels $(\sigma_{noise}/\nabla n_{peak} < 4\%)$, both schemes using a 5-point kernel for the right-hand side discretisation are similarly accurate. It would be expected that the staggered discretisation scheme (3-point left-hand side and 5-point right-hand side) would better account for the smaller truncation error in the gradient fields, as the gradients are returned directly from the reconstruction and do not need to be computed from finite differences, and so produce a smaller error than the scheme using a higher-order on both sides (5, 5). This difference appears to be negligible in practice for the tested cases, and both schemes behave very similarly across the range of wavelengths and noise levels. The (3, 5) and (5, 5) kernels give the lowest RMS and peak errors across the range of noise levels tested, calculated within $r \leq 2r_{1/2}$ as in figure 4.4 and figures 4.5a and b; as discussed below, this is not true if the whole domain is considered (i.e. ambient points included). The schemes using a 5-point right-hand side kernel exhibit a sharper increase in error as a function of noise,



Figure 4.4: RMS error between the synthetic refractive index fields and the refractive index fields calculated by the Poisson solver, n_s and n_r , respectively, within twice the jet half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$, as a function of fluctuation wavelength $L/\lambda_{x,z}$ and additive Gaussian noise level $\sigma_{noise}/\nabla n_{peak}$. The error is calculated within the jet core for varying kernel sizes in the discretisation used in the left-hand side (multigrid) and right-hand side (calculating Laplacian from the reconstructed refractive index gradient field), respectively: a) 3 and 3 points, b) 3 and 5 points, c) 5 and 3 points, d) 5 and 5 points. Each data point is averaged over 100 samples of added random noise.

to the point that at a noise level of $\sigma_{noise}/\nabla n_{peak} = 10\%$, these schemes have a similar error to the schemes using a 3-point left-hand side.

A closer look at the trends of the RMS and peak errors of the solved fields is shown in figure 4.5, for one wavelength $\lambda_{x,z} = L/14$. The behaviour of the different discretisation schemes is quite different depending on whether the jet core (within twice the half-width $2r_{1/2}$) or the whole domain is considered. This is because outside the jet the density gradients approach zero, and so the addition of a constant noise level has a larger effect there. Within the jet (figures 4.4a and b), the higher-order right-hand side schemes clearly deliver the lowest RMS and peak errors up to the highest noise level tested. For $\sigma_{noise}/\nabla n_{peak} < 4\%$, the staggered (3, 5) scheme delivers a marginally better than the (5, 5) scheme, for the reason discussed above. Beyond this, the (5, 5) scheme has a less sharp increase in error with increasing noise level.

These trends change when the ambient points are considered as well (figures 4.4c and d). At low noise levels ($\sigma_{noise}/\nabla n_{peak} \leq 2\%$), the lowest RMS and peak errors are still



Figure 4.5: RMS (left) and peak (right) errors between the synthetic refractive index fields and the refractive index fields calculated by the Poisson solver, n_s and n_r , respectively, for fluctuation wavelength $\lambda_{x,z} = L/14$, calculated up to: twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ (top row), and in the whole domain (bottom row). Errors are shown as a function of additive Gaussian noise level $\sigma_{noise}/\nabla n_{peak}$. The error is calculated for varying kernel sizes in the discretisation used in the left-hand side (multigrid) and righthand side (calculating Laplacian from the reconstructed refractive index gradient field), respectively: 3 and 3 points \circ ; 3 and 5 points \Box ; 5 and 3 points \triangle ; 5 and 5 points \diamondsuit . Each data point is averaged over 100 samples of added random noise; error bars indicate the 95% confidence level and are approximately the same size as the markers.

delivered by the 5-point gradient (right-hand side) discretisation schemes. But it becomes apparent that using a higher discretisation order on noisy reconstructed data (3, 5 and 5, 5 schemes) creates a higher measurement error from as early as $\sigma_{noise}/\nabla n_{peak} = 3\%$. At higher random noise levels, the lowest-order discretisation scheme (3, 3) avoids sharp amplification of the gradient field noise in the final refractive index measurement. The reverse-staggered (5, 3) scheme delivers the lowest error of all at $\sigma_{noise}/\nabla n_{peak} = 10\%$, but its poor performance within the jet may preclude its use. The ambient regions could typically be filtered in an experimental measurement, which may allow the use of higher-order right-hand (3, 5 and 5, 5) schemes in the whole domain. This indicates that the optimal Poisson discretisation scheme depends on the local signal-to-noise ratio throughout the domain. Of course, this is not practical, so it is recommended to use a lower order scheme to avoid propagating noise.



Figure 4.6: Bode magnitude plot of the analytical (A) and finite-difference Poisson equation transfer functions H as a function of spatial frequency ω . The finite-difference schemes are denoted by the number of left- and right-hand side points used in the discretisation kernel as per table 4.1, e.g. (3, 5) for 3 points on the left-hand side and 5 points on the right-hand side. The Nyquist frequency is $\omega_x = 0.5$ voxel⁻¹, or $\lambda_{x,z} \approx L/32$.

To gain further insight on the behaviour of the different discretisation schemes, consider the one-dimensional Poisson equation,

$$\frac{\mathrm{d}^2 n}{\mathrm{d}x^2} = q \equiv \frac{\mathrm{d}}{\mathrm{d}x} \left(\frac{\mathrm{d}n}{\mathrm{d}x}\right)_{recon} \tag{4.19}$$

As explained in Appendix A.1, this can be considered as a single-input single-output (SISO) system between the reconstructed gradients and the desired refractive index field. Although obviously much simpler than the 2D or 3D Poisson equations, through von Neumann stability analysis, the transfer function of this system can effectively illustrate the effect of discretisation on the calculated solution.

Figure 4.6 shows the Bode magnitude plot of the analytical and finite-difference transfer functions for the one-dimensional Poisson equation. It is immediately obvious that for lower spatial frequency fluctuations (longer wavelengths), all of the finite-difference schemes are able to match the behaviour of the analytical transfer function. Similarly, as the frequency of the fluctuations approaches the Nyquist frequency ($\omega = 0.5 \text{ voxel}^{-1}$, $\lambda_{x,z} \approx L/32$), all of the discretisation schemes strongly diverge from the analytical transfer function, which indicates a strong attenuation of these higher frequencies. The staggered (3, 5) scheme provides a closer match to the analytical transfer function across the frequency range than the other schemes, but this would also allow higher-frequency noise to have a greater effect in the final measurement.

This approach can be used to select the most suitable discretisation scheme, if the range of frequencies resolved by the reconstruction, and the nature of the noise, can be estimated. These will be examined in sections 4.5.5 and 4.5.6, and the investigation will return to this analysis to select the most appropriate scheme.

4.5.2 FBP reconstruction

In the limit of infinite views, FBP should provide both an accurate and efficient reconstruction of the refractive index gradients. However, in the case of limited simultaneous viewing angles, this approach can generate significant reconstruction artefacts outside the



Figure 4.7: Contour plots, and profiles through x = 0 (black dotted line is the original synthetic field and red dashed line is the reconstruction), for the reconstructed refractive index gradient $\partial n/\partial x$ for 16 cameras and $\lambda_{x,z} = L/8$: a) synthetic field, b) FBP. Bottom row is $\lambda_{x,z} = L/14$: c) synthetic field, d) FBP.

jet. The artefacts can clearly be seen in the 16 camera reconstruction with a fluctuation wavelength of $\lambda_{x,z} = L/14$ in figure 4.7. The nature of the artefacts seems to depend on the spatial frequencies present in the reconstruction, and they can become indistinguishable from the flow itself. This makes filtering the artefacts out more challenging in general, although for the self-similar region of the heated round jet, it is known that the density field should conform to a Gaussian envelope which can be used for filtering the flow.

The performance of FBP seems to be very sensitive to the number of cameras used. Given enough cameras, the refractive index field in the jet's core can be predicted very well by FBP, as seen in figures 4.8, assuming the background displacements are calculated exactly. This seems to require at least 12 cameras and using any less than this causes an exponential increase in the RMS and peak errors. Using any more than 12 cameras,



Figure 4.8: Contour maps of the RMS (a) and peak (b) errors within twice the halfwidth $r \leq 2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ between the synthetic and FBP reconstructed refractive index fields n_s and n_r , respectively, as a function of wavelength and camera number, normalised by the peak change in the synthetic refractive index field from the Poisson solution. The minimum RMS and peak errors are 0.3% and 0.7%, respectively. The maximum RMS and peak errors are 35% and 97%, respectively.

the rate of decrease of RMS and peak error is far less dramatic, indicating diminishing returns (see also figure 4.12e and f). For the range of wavelengths tested, using more than 16 cameras does not significantly increase the accuracy of the FBP reconstruction.

For a given number of cameras, the FBP RMS and peak errors steadily decrease with wavelength, i.e. increase with spatial frequency (also see figure 4.12c and d). For the longest wavelengths tested, $L/\lambda_{x,z} \leq 4$, FBP is the most accurate reconstruction method in the region of interest, but the ART methods introduced in the following section are superior at shorter wavelengths ($L/\lambda_{x,z} > 4$). This study will return to exploring the performance of FBP with increasing wavelengths in section 4.5.5. If the jet's core is wellpredicted by FBP, then it may be a reasonable basis for a reconstruction which is further enhanced by mitigating artefacts. This will be explored further in section 4.5.4, where the FBP reconstruction is used as the initial solution to an iterative ART reconstruction to possibly shorten the convergence of ART.

4.5.3 ART reconstruction schemes and enhancements

The performance of ART is strongly related to the modifications made to the basic reconstruction scheme. The modifications aim to improve the quality of the reconstruction in three categories: rate of convergence, accuracy of predicted gradients, and reduction of artefacts. Examples of the ART reconstruction and the influence of different reconstruction schemes are shown in figure 4.9, for 16 cameras with a fluctuation wavelength of $\lambda_{x,z} = L/14$ after 100 iterations with a relaxation parameter $\lambda_j = 1.0$. The convergence of these schemes is shown in figure 4.10. The basic ART method in figure 4.9b contains visible artefacts like FBP, while under-predicting the gradients to a much greater degree than FBP. As combinations of modifications are progressively added from figures 4.9c to 4.9f, the artefacts are diminished, and the strength of the predicted gradients are improved by concentrating the corrections towards the centre of the volume.

First, consider a constant 16-camera reconstruction with $\lambda_{x,z} = L/14$, and the mod-

Table 4.2: Influence of ART reconstruction techniques for 16 camera reconstruction and $\lambda_{x,z} = L/14$ after 100 iterations with relaxation parameter $\lambda_j = 1.0$. RMS errors $\sqrt{\langle (\dots_r - \dots_s)^2 \rangle}$ and peak errors max $|\dots_r - \dots_s|$ are given for the reconstructed gradient fields ∇n and Poisson-solved reconstructed refractive index fields n, where subscripts \dots_r and \dots_s denote the reconstructed and true synthetic fields, respectively. $R_{\dots_r \dots_s}$ denotes the correlation coefficient between the reconstructed and synthetic fields. Errors and correlation coefficients are calculated within a radius that is twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$. The Poisson equation is solved with the (3, 5) discretisation scheme. Gradients are presented both normalised by the peak difference between the centreline and outer flow refractive index Δn_p divided by domain size L, and as a percentage of the peak gradient in the spatial average synthetic field ∇n_{peak} . Errors in the Poisson solved refractive index field n are given as a percentage of Δn_p . Case K is selected for further investigation (bold text, red ticks).

Case	$\sqrt{\langle (\nabla \cdot) \rangle}$	$\overline{\left(n_r - \nabla n_s\right)^2}$	$\max \nabla$	$7n_r - \nabla n_s$	$R_{\nabla n_r \nabla n_s}$	$\sqrt{\left\langle \left(n_r - n_s\right)^2\right\rangle}$	$\max n_r - n_s $	$R_{n_r n_s}$
	$\frac{\Delta n_p}{L}$	$\nabla n_{peak}\%$	$\frac{\Delta n_p}{L}$	$\nabla n_{peak}\%$		$\Delta n_p\%$	$\Delta n_p \%$	
А	3.09	16.68	14.58	78.77	0.757	3.07	14.93	0.990
В	3.09	16.68	14.57	78.74	0.757	3.06	14.86	0.990
\mathbf{C}	2.76	14.90	12.08	65.30	0.810	2.74	12.55	0.992
D	2.76	14.90	12.15	65.66	0.810	2.76	12.75	0.992
Ε	2.75	14.85	12.01	64.93	0.812	2.72	12.58	0.992
F	1.99	10.76	7.75	41.89	0.905	1.97	8.76	0.996
G	1.99	10.76	7.78	42.02	0.906	2.00	8.96	0.996
Η	1.97	10.64	7.77	41.99	0.909	1.84	8.35	0.997
Ι	1.94	10.46	7.24	39.12	0.913	1.83	8.15	0.997
J	1.59	8.70	6.59	35.95	0.944	1.16	4.36	0.999
Κ	1.24	6.78	4.52	24.65	0.966	1.00	3.35	0.999
L	1.62	8.85	6.74	36.77	0.940	1.18	3.72	0.999
Μ	1.27	6.92	4.58	24.97	0.964	1.09	3.42	0.999

Case	Random order	Sharp-cutoff mask	Inversely iteration-weighted Gaussian filter	Gradual unmasking	Hamming -windowed correction	Progressively tightened Gaussian mask
А						
В	\checkmark					
\mathbf{C}	\checkmark	\checkmark				
D	\checkmark	\checkmark	\checkmark			
Ε	\checkmark	\checkmark		\checkmark		
\mathbf{F}	\checkmark	\checkmark			\checkmark	
G	\checkmark	\checkmark	\checkmark		\checkmark	
Η	\checkmark	\checkmark		\checkmark	\checkmark	
Ι	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	
J	\checkmark	\checkmark			\checkmark	\checkmark
Κ	\checkmark	\checkmark	\checkmark		\checkmark	\checkmark
L	\checkmark	\checkmark		\checkmark	\checkmark	\checkmark
Μ	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark

ifications which may be made to enhance the rate of convergence. ART algorithms are known to become unstable at higher values of λ_j , however the value at which this occurs depends on the modifications (to be discussed) to the algorithm that are used. So this

study will first devleop an optimal set of modifications that improves accuracy, before examining the effects of varying the relaxation parameter λ_j . Unless otherwise stated, the results will use a relaxation value of $\lambda_j = 1.0$. Given enough iterations to reach a converged state with a given set of modifications, the choice of λ_j is not critical so long as the solution does not diverge. Kak and Slaney [62] describe another well-established method of improving the rate of convergence by randomising the order in which rays and cameras are considered during an iteration, rather than considering them sequentially. This allows the volume to be 'filled-in' faster. The effect of convergence of *randomised ray and camera ordering* (case B) compared to sequential ordering (case A) is presented in figure 4.10, and the impact on the accuracy of the final measurement is shown in table 4.2. It is demonstrated that randomising the camera and ray orders each iteration marginally accelerated the convergence of the reconstruction during the early iterations yet had negligible effect on the final reconstruction when 10 or more iterations were performed. Cases B onward retain the random camera and ray ordering.

Now consider modifications targeting artefact reduction and increasing the strength of the predicted gradients. Given that the artefacts appear outside the jet, the simplest means of reducing artefacts is to add a *sharp cut-off mask* which masks the volume so that the reconstructed gradients are set to zero beyond a specified radius r from the centre of the reconstructed volume;

$$\nabla n_j^k = 0, \text{ where } r > r_{mask}, \tag{4.20}$$

and $r_{mask} = 30$ voxels $(3.3\sigma, 2.8r_{1/2})$ in this case. Care must be taken not to cut off areas of the domain that could legitimately be inhabited by significant flow features. There are various methods to implement this. The mask size can be automatically detected from the displacement fields decreasing below a certain threshold, e.g. 0.1 pixels, which signifies the edge of the jet. This is like the 'visual hull' approach used by Atcheson et al. [6], and others, which prevents any gradient reconstruction in regions clearly outside the flow. As the sharp cut-off mask restricts the reconstruction algorithm from spreading the reconstructed gradients across the whole domain, the strength of the predicted gradients at the centre of volume naturally increases. This results in a modest reduction in the RMS and peak errors, as seen by comparing cases C and B in table 4.2 and figure 4.10.

An inversely iteration-weighted Gaussian filtering of the reconstructed gradients is introduced, with a standard deviation of 0.5 voxels ($\sigma_{GF}/\sigma = 0.06$ where σ_{GF} is the standard deviation of Gaussian filtering) expressed as,

$$\nabla n_j^k = \nabla n_j^k \big|_{unfiltered} + \left(1 - \frac{k}{0.5N_{it}}\right) \left[\nabla n_j^k \big|_{filtered} - \nabla n_j^k \big|_{unfiltered}\right], \text{ where } k \le 0.5N_{it},$$

$$(4.21)$$

where N_{it} is the total number of iterations to be performed. The Gaussian filtering of each point is performed on a 5×5 kernel with its neighbours on the same transverse plane. Neighbours outside of the domain are assigned a value of zero, as they are far outside the jet and the refractive index gradients are negligible there. The inversely iterationweighted Gaussian filtering aims to aggressively reduce artefacts as soon as they appear in the earliest iterations, backing off as the iteration process continues. A sensitivity study on the effect of standard deviation is provided in the Appendix A.2, which shows that a value close to $\sigma_{GF} = 0.5$ voxels ($\sigma_{GF}/\sigma = 0.06$) is optimal for lower RMS and peak errors. With no other modifications save for the random ordering and sharp cut-off mask, the inversely iteration-weighted Gaussian filtering presents only a marginal improvement



Figure 4.9: Contour plots, and profiles through x = 0 (black dotted line is the original synthetic field and red dashed line is the reconstruction), of the ART reconstructed refractive index gradient $\partial n/\partial x$ for 16 camera reconstruction and $\lambda_{x,z} = L/14$ after 100 iterations with $\lambda_j = 1.0$: a) synthetic field; b) case B in table 4.2; c) case C; d) case D; e) case I; and f) case K (best case).



Figure 4.10: RMS error (a) and peak error (b) between the reconstructed gradients ∇n_r and synthetic gradients ∇n_s for 16 camera reconstruction and $\lambda_{x,z} = L/14$ in the region twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$, as a function of ART iterations with relaxation parameter $\lambda_j = 1.0$ for various ART modification schemes. Shown are: case A in table 4.2 \circ ; case B \Box ; case C \triangle ; case I \diamond ; and, case K \bigtriangledown . Cases are consistent with gradient field visualisations presented in figure 4.9.

in accuracy (compare cases D and C in table 4.2). However, as will be discussed later, the inversely iteration-weighted Gaussian filter is a crucial addition to the optimal scheme (case K), where its presence results in a large decrease in the measurement error.

Following Liao [77], a gradual unmasking method is tested, which appears to reduce artefact formation in limited-view ART. This approach is based on recovering the most intense features (highest gradient magnitude in this case) first in early iterations, then allowing lower intensity features to develop. Gradual unmasking aims to prevent artefacts forming through an update to the reconstruction at each iteration using a progressively decreased threshold given by:

$$abla n_j^k = 0, \text{ where } \left| \nabla n_j^k \right| < t_o \left(1 - \frac{t_o k}{N} \right),$$

$$(4.22)$$

where t_o is an initial threshold level set at $t_0 = 0.00044 \frac{\Delta n_p}{\Delta x}$. This modification, at best, marginally improves measurement accuracy (e.g. comparing cases E and D in table 4.2), but it may also slightly increase the error in certain cases (e.g. cases L and J). Due to its negligible impact, it does not feature in the optimal ART scheme. Still, due to its potential in reducing artefacts, this modification will be revisited when in the discussion of the hybrid FBP+ART method in section 4.5.4.

As the jet is centred in the domain, artefacts form away from the centre. Based on this, a *Hamming windowed correction* is proposed:

$$\nabla n_j^{k+1} = \nabla n_j^k + \mathcal{W}(\nabla n_j^{k+1} - \nabla n_j^k), \qquad (4.23)$$

$$\mathcal{W} = 0.54 - 0.46 \cos\left(\frac{2\pi\zeta}{L_i}\right),\tag{4.24}$$

where ζ is the relative position along a ray through the reconstruction volume of length L_i . This modification acts to weight the correction in the gradient fields towards the

centre of the volume, consistent with the location of the jet in this case. Out of all of the modifications tested, this is the most effective in reducing the artefacts and improving the prediction of the gradients (especially in conjunction with the progressively tightened Gaussian mask which will be introduced later). This is clearly demonstrated in table 4.2 through the decrease in RMS and peak errors, and increase in correlation coefficient (discussed below) between the reconstructed and original fields (compare cases F and C).

Lastly, consider a *progressively tightened Gaussian mask* which acts to further concentrate the corrections towards the centre of the volume,

$$\nabla n^k = \nabla n^{k-1} \left[\exp\left(\frac{-r}{2(r_m/3)^2}\right) \right],\tag{4.25}$$

where k is the iteration number and r_m is the mask width. The mask width r_m starts with an initial value of $r_{m,initial} = 35$ voxels $(3.9\sigma, 3.3r_{1/2})$ and decreases linearly with respect to iteration number to $r_{m,final} = 30$ voxels $(3.3\sigma, 2.8r_{1/2})$ at the final iteration. A sensitivity study on mask size is presented in Appendix A.3; the mask can overrestrict the reconstruction, and so an over-sized mask is preferred (both $r_{m,initial}, r_{m,final} > 3\sigma$). This modification also significantly reduces the measurement errors, in addition to the Hamming windowed corrections (compare cases F and J in table 4.2). Adding the inversely iteration-weighted Gaussian filter as well (case K) achieved the best ART reconstruction. Case K manages to capture the large, sharp gradients in the $\lambda_{x,z} = L/14$ field, as shown in figure 4.9, although the convergence does appear to be slower than the methods which do not use the progressively tightened Gaussian mask (see figure 4.10).

It is important to note that unlike standard intensity-based reconstructions (e.g medical imaging or tomographic PIV), the TBOS reconstruction can produce positive and negative gradients at the opposite sides of the volume which may cancel each other out when a ray is projected through them. These gradients may satisfy the observed ray deflection, but result in the spurious gradients outside the jet observed in figure 4.9b (for example). As the basic reconstruction scheme only seeks that the iterated gradient field satisfies the observed displacements, these would be a totally valid (but completely unphysical) solution, which enforces the need for these modifications. This is the rationale for using the correlation coefficient as an additional metric for the reconstruction quality in table 4.2, as it provides a measure of similarity between the reconstruction and original field. Note that the Poisson solver acts to filter out some of these spurious fluctuations as well, so they do not propagate fully from the reconstructed gradients to the final refractive index field. Although the improvements in RMS and peak errors and correlation coefficient within $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ may be slight between successive modifications presented in table 4.2, it is found from figure 4.9 that artefacts can be removed completely, and the field is predicted very well, with the prescribed modifications (e.g. case K) compared to the standard sequential-order ART or unmodified random-order ART (case B).

Having selected the optimal reconstruction scheme, it is prudent to test the effect of relaxation factor λ_j on the optimised ART (case K), and using SART instead of ART. The SART iterative correction [5] is based on simultaneously satisfying all observed projections rather than considering rays sequentially, given by equation 4.11. The aim is to provide a reasonable reconstruction in just one iteration. Figure 4.11 shows the RMS and peak errors for both ART and SART approaches, using the case K modifications. Experiments using SART show slower convergence, owing to the attempt to simultaneously satisfy every projection of the volume at once. SART convergence can be accelerated via the use



Figure 4.11: RMS error in the refractive index gradients ∇n (a) and the refractive index fields n (b) for 16 camera reconstruction and $\lambda_{x,z} = L/14$ in the region twice the halfwidth $r \leq 2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$, as a function of ART iterations for: ART $\lambda_j = 0.2 \circ$; ART $\lambda_j = 0.5 \Box$; ART $\lambda_j = 1.0 \triangle$; ART $\lambda_j = 4.0 \diamond$; SART $\lambda_j = 1.0 \blacktriangle$; SART $\lambda_j = 4.0 \blacklozenge$. In all cases the Poisson equation is solved using 3- and 5-point kernels for the left- and right-hand side calculation, respectively. Reconstructions use case K in table 4.2.

of a larger relaxation parameter, λ_j , although care must be taken to ensure the solution does not become unstable. Both ART and SART approaches were found to be stable if $\lambda_j \leq 4$, but as shown the accuracy of the gradient field and final Poisson solution were not independent of this parameter. It seems that 100 iterations is sufficient for ART to converge when $0.5 \leq \lambda_j \leq 4.0$. SART behaves like a slower-converging ART, so a larger λ_j is required for it to converge in 100 iterations. Negligible difference was observed in the final reconstruction quality between these two approaches. As it converges very slowly compared to ART for values of λ_j that do not cause the solution to diverge, from here on, SART will be not considered.

In summary, ART was found to be most accurate when using randomly-ordered cameras and vectors, a sharp cut-off mask, Gaussian filtering, Hamming windowed corrections and the progressively tightened Gaussian mask (case K in table 4.2). These modifications increase the quality and accuracy of the reconstruction by concentrating the iterative corrections towards the centre of the volume. Although this ART scheme tends to converge more slowly than other tested schemes, it is observed that the relaxation parameter λ_j has the greatest effect on convergence. Setting $0.5 \leq \lambda_j \leq 4.0$ allows the reconstruction to converge within 100 iterations while avoiding the possibility of a diverging solution.

Using case K as the optimal method, the variation in accuracy for ART with camera number and fluctuation wavelength is shown in figure 4.12. Unlike FBP, ART copes much better with fewer cameras, and does not show as sharp an increase in RMS and peak error with decreasing camera numbers. The performance of ART within the jet when using less than 12 cameras is far superior to FBP. As little as 6 cameras could be used. Apart from wavelengths $L/\lambda_{x,z} < 6$, the modified ART is superior to the FBP reconstruction for the 16 camera reconstruction. The modified ART does not under-predict the peaks in the gradient field as much as the FBP, and it does not suffer from artefacts.

Ultimately, the accuracy of a TBOS reconstruction depends not only on the reconstruction methodology but also on the number and orientation of the background views, the strength of the density gradients and the relative distance between the volume and the background. A range of 6 to 18 cameras is considered practical, as fewer cameras results in significant errors and the benefit associated with the use of more cameras diminishes significantly, not to mention the cost and geometric challenges associated with the use of higher camera numbers. Note that this study has only considered an even camera spacing across a 180° arc in a plane transverse to the jet axis. Nicolas et al. [96] have tested the effect of camera placement for their iterative method, and found that spreading the cameras out is necessary for a high-quality reconstruction. Clustering the cameras closely together does not provide the reconstruction algorithm with enough information on the distribution of refractive index gradients. It does not matter if the cameras are evenly spread out across a 180° arc, or across a full 360° range, as the displacements are path-integrated and, hence, the same information is captured.

4.5.4 Hybrid FBP+ART reconstructions

Given that ART requires an initial guess, which is often set arbitrarily to zero, it is logical to try to use FBP as the initial guess to ART (FBP+ART) to try to improve the rate of convergence. It is clear that FBP is less accurate than ART, so care must be taken to avoid corrupting the ART solution due to the FBP input. There is no expectation for FBP+ART to be more accurate than the optimised ART, but any improvement in the rate of convergence is most welcome. In the current study, it has been shown that FBP captures the fluctuations in the core of the jet quite well, although the peak gradients are under-predicted (figure 4.7). This indicates, like Hartmann and Seume [53], that FBP may provide ART a good baseline for resolving the turbulent fluctuations in the jet in fewer iterations than starting from arbitrary initial conditions. The artefacts which are introduced by FBP outside the jet, which inflate its calculated RMS error, should be removed completely in the optimised ART. This study demonstrate a few combinations of the modifications discussed previously in section 4.5.3 on the FBP+ART reconstruction, with the aim of finding an optimised FBP+ART method that converges quicker than the optimised ART, with comparable accuracy.

A comparison between the FBP, optimised ART and a few variations of FBP+ART is presented in figure 4.13 for 16 camera measurements with a fluctuation wavelength of $\lambda_{x,z} = L/14$. All ART iterations are performed using random camera and pixel ordering and the sharp cut-off mask; λ_j is now set to a very conservative 0.5 to minimise the risk of a diverging solution brought about by the growth of FBP artefacts. FBP+ART is tested with different combinations of inversely iteration-weighted Gaussian filtering, gradual unmasking and the progressively tightened Gaussian mask to gauge their effects on reconstruction error and artefact removal. The progressively tightened Gaussian mask acts on the FBP reconstruction before the FBP+ART iterations begin:

$$\nabla n^{k=0} = \nabla n_{FBP} \left[\exp\left(\frac{-r}{2(r_m/3)^2}\right) \right], \qquad (4.26)$$

where k = 0 refers to the initial guess for the ART iterations, ∇n_{FBP} are the FBP reconstructed gradients and r_m is the effective radius. A sensitivity study of initial and final mask size on the FBP+ART reconstruction is presented in Appendix A.4, where the change in error of the input FBP solution to FBP+ART as $r_{m,initial}$ is varied is examined, and also the change in the final FBP+ART error as $r_{m,initial}$ is kept constant and $r_{m,final}$ is varied. Based on this study, the initial radius of the Gaussian mask is chosen to be $r_{m,initial} = 35$ voxels $(3.9\sigma, 3.3r_{1/2})$. The mask width is decreased linearly



Figure 4.12: Contour maps (top row) of the RMS (a) and peak (b) errors between the synthetic and optimised ART reconstructed refractive index fields, n_s and n_r , respectively, with 100 iterations within twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ as a function of wavelength $L/\lambda_{x,z}$ and camera number $n_{cameras}$, normalised by the peak change in the synthetic refractive index field from the Poisson solution n_{peak} . The minimum RMS and peak errors for ART are 0.6% and 1.5%, respectively. The maximum RMS and peak errors are 3.0% and 11.2%, respectively. Colourbar is consistent with figure 4.8 for comparison.

Line plots (middle row) show RMS (c) and peak errors (d) as function of wavelength for: 8camera FBP • and optimised ART \times ; 16-camera FBP • and optimised ART \times ; 20-camera FBP • and optimised ART \times . FBP data is from figure 4.8.

Line plots (bottom row) show RMS (e) and peak errors (f) as a function of camera number for wavelength $L/\lambda_{x,z} = 14$ only: FBP \diamond , and optimised ART \circ .

Optimised ART corresponds to case K from table 4.2 with $\lambda_j = 0.5$.

with respect to iteration number, such that $r_{m,final} = 30$ voxels $(3.3\sigma, 2.8r_{1/2})$ by the final iteration (identical settings to the progressively tightened Gaussian mask for ART). This is aimed at removing the reconstruction artefacts introduced by the FBP, while leaving the gradients at the centre of the jet unaffected.

Along with the visual comparison of the FBP+ART reconstructions in figure 4.14, it is clear that FBP+ART requires that the ART introduces some form of filtering (progressively tightened Gaussian mask or inversely iteration-weighted Gaussian filtering) to avoid maintaining, and even strengthening, the FBP artefacts. If this is not done, FBP+ART can have a higher error than either FBP or ART alone (case C in figure 4.13 and figure 4.14b), due to the presence of artefacts *and* weak gradients. Although gradual unmasking is used, it does not prevent the growth of artefacts, and this case essentially combines the worst of both methods.

The remaining FBP+ART cases (cases D-H) in figure 4.13 use either the progressively tightened Gaussian mask or inversely iteration-weighted Gaussian filtering, or both. These methods all manage lower RMS and peak errors after 100 iterations than FBP alone, but to varying degrees. They also show different convergence behaviours. Cases using the inversely iteration-weighted Gaussian filter (cases D, G and H) possess the slower convergence style similar to the selected ART case (case A in figure 4.13, which is the same as case K from table 4.2 but using $\lambda_i = 0.5$). The two cases which use the progressively tightened mask, and forgo the inversely iteration-weighted Gaussian filter (cases E and F), show rapid convergence, with negligible improvements after 20 iterations. Although the errors of schemes E and F are higher than best ART case A (and indeed the very similar FBP+ART case G which uses the same settings), the improved rate of convergence for cases E and F may favour their use for analysis of large data sets because a marginal loss in accuracy may be worth the large savings in computation time. Again, the gradual unmasking results in a slightly higher error for these two otherwise identical cases but has little effect on modifying the convergence rate. The error of case F is higher than case E because the gradual unmasking encroaches on the smaller gradients near the $2r_{1/2}$ limit up to which the errors are calculated.

Further evidence of the FBP+ART method's fast convergence is observed by comparing figure A.5 in Appendix A.5 and figure 4.13. Figure A.5 shows the convergence of the ART method with the different modification schemes from figure 4.13. It is clear that FBP+ART schemes E and F of figure 4.13 converge faster than any tested ART scheme in figure A.5.

A comparison of the FBP+ART reconstructed gradient fields in figure 4.14 shows that the higher error in the fast-converging FBP+ART case E (figure 4.14c) compared to the most-accurate FBP+ART case G (figure 4.14d) is due to under-prediction of the highest gradients. Note the extreme similarity between these two reconstructions and the best two ART reconstructions shown in figures 4.9e and f. In subsequent comparisons featuring the FBP+ART method, case E is used. Although it is somewhat less accurate than case G, its fast convergence may be extremely valuable in practice.

4.5.5 Comparison of the optimised FBP, ART and FBP+ART methods

Now it is possible to directly compare the optimal FBP, ART and FBP+ART methods. From the many reconstruction schemes tested, FBP, ART case A from figure 4.13, and FBP+ART case E from figure 4.13 (the latter two with 100 iterations) are selected for



Figure 4.13: RMS error (top row) and peak error (bottom row) in the reconstructed refractive index gradients ∇n (left column) and the refractive index fields n (right column) for 16 camera reconstruction and $\lambda_{x,z} = L/14$ in the region twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$, as a function of the number of ART iterations for the cases shown in the table below. In all cases the Poisson equation is solved using 3- and 5-point kernels for left- and right-hand side discretisation, respectively. All ART and FBP+ART reconstructions are performed using randomly-ordered cameras and pixels with Hamming windowed corrections and relaxation $\lambda_j = 0.5$. ART and FBP+ART use a sharp cutoff mask with $r_{mask} = 30$ voxels $(3.3\sigma, 2.8r_{1/2})$. Progressively tightened Gaussian mask decreases from $r_m = 35$ voxels $(3.9\sigma, 3.3r_{1/2})$ to $r_{m,final} = 30$ voxels. ART case A is the same as case K from table 4.2 with $\lambda_j = 0.5$.

Marker	Case	Туре	Inversely iteration-weighted Gaussian filter	Gradual unmasking	Progressively tightened Gaussian mask
×	А	ART	\checkmark		\checkmark
	В	FBP	n/a	n/a	n/a
0	\mathbf{C}	FBP+ART		\checkmark	
	D	FBP+ART	\checkmark	\checkmark	
Δ	Ε	FBP+ART			\checkmark
\diamond	F	FBP+ART		\checkmark	\checkmark
	G	FBP+ART	\checkmark		\checkmark
\bigtriangledown	Η	FBP+ART	\checkmark	\checkmark	\checkmark



Figure 4.14: Contour plots, and profiles through x = 0 (black dotted line is the original synthetic field and red dashed line is the reconstruction), for the reconstructed refractive index gradient $\partial n/\partial x$ for 16 camera reconstruction and $\lambda_{x,z} = L/14$: a) synthetic field; b) FBP and 100 ART iterations with gradual unmasked and Hamming windowed correction (case C in the table of figure 4.13); c) FBP multiplied by progressively tightened Gaussian mask and 100 ART iterations with Hamming windowed corrections (case E), and d) FBP multiplied by progressively tightened Gaussian mask and 100 ART iterations with inversely iteration-weighted Gaussian filter and Hamming windowed corrections (case G).

further examination. To address the aims of this study, it is required to judge the accuracy of the reconstructions across the domain and determine the range of scales that are wellresolved for each method. Consider the average error of the reconstructed gradients and Poisson-solved refractive index fields as a function of normalised radius r/σ in figure 4.15. Results are presented for fluctuations of a single wavelength for clarity and brevity, but similar behaviour is observed for other wavelengths down to $\lambda_{x,z} = L/16$. Within $r \leq \sigma$ (recall that the half-width is $r_{1/2} = \sigma \sqrt{2 \ln 2} \approx 1.17\sigma$), ART maintains the lowest



Figure 4.15: Average absolute error $|..._r - ..._s|_{bin}$ in the reconstructed refractive index gradients ∇n normalised by ∇n_{peak} (a), and the refractive index fields n normalised by n_{peak} (b) for 16 camera reconstruction and $\lambda_{x,z} = L/14$, as a function of normalised radial position r/σ with a bin size of $\sigma/4$. Shown are: FBP •, ART case A from figure $4.13 \times$, and FBP+ART case E from figure $4.13 \triangle$ (the latter two both correspond to the same markers in figure 4.13). Vertical dotted line indicates the usual $2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ limit that the RMS error is calculated within. Sharp cut-off mask radius is located at $r_{mask} = 3.3\sigma$. Initial width of the progressively tightened Gaussian mask is $r_m = 3.9\sigma$, and the final width is $r_{m,final} = 3.3\sigma$.



Figure 4.16: Power spectral density of the analytical $\partial n/\partial x|_{x=0,z}$ for different imposed frequency fluctuations ω , from L/2 to L/32, in increments of L/2 ($\omega = 0.03$ voxel⁻¹), left to right. Each frequency (individual lines) produces a distinct, sharp peak.

error of all, while FBP and FBP+ART errors possess very similar characteristics and are noticeably higher than ART due to under-prediction of the peak gradients. After peaking in error within $r < \sigma$, the ART and FBP+ART errors steadily decline with increasing radius, as the modifications to these methods (sections 4.5.3 and 4.5.4) are able to remove artefacts outside the jet and return to ambient conditions. Conversely, the FBP increases again from $r > \sigma$, especially up to twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$, which contributes to the higher indicated RMS errors. The error is especially high in this region because the gradients in the jet approach zero, but the FBP artefacts actually start from here outwards.



Figure 4.17: Peak power (left column), and relative peak power (right column), of power spectral density of $\partial n/\partial x|_{x=0,z}$ for different imposed frequency fluctuations ω , from L/2 to L/32, in increments of L/2 ($\omega = 0.03 \text{ voxel}^{-1}$). L/32 corresponds to the Nyquist frequency $\omega = 0.5 \text{ voxel}^{-1}$. Top row corresponds to A = 0.125, while the bottom row corresponds to the usual A = 0.25. Red horizontal dashed line indicates 50% relative power criterion. Shown are peak powers of the spectra for: the analytical fields ∇ , FBP •, selected ART case \times , and selected FBP+ART case \triangle (the latter two both correspond to the same markers in figure 4.13).

Having assessed the errors of the reconstructions for larger scales, now consider the performance of these methods for the full range of scales, up to the Nyquist frequency $\omega = 0.5 \text{ voxel}^{-1}$ ($\lambda_{x,z} \approx L/32$). Under-prediction of the gradients has been identified as the primary issue that is detrimental to reconstruction performance. Consider a new metric to assess a method's ability to resolve fluctuations of a particular frequency. As the frequency of the fluctuation is increased, the power spectrum of a line through $\partial n/\partial x(x = 0, z)$, i.e. $|\mathcal{F}[\frac{\partial n}{\partial x}|_{x=0,z}]|^2$ (where \mathcal{F} is the Fourier transform), will show a clear peak at each frequency, illustrated in figure 4.16. Given that the reconstructions have been observed thus far to under-predict the gradients, but still correctly identify the frequency, the power spectrum of the reconstructions should also show a peak at each frequency, but smaller in magnitude.

The left column of figure 4.17 compares the peak of each spectrum of the analytical and reconstructed fields for increasing fluctuation frequency. The amplitudes of turbulence fluctuations are inversely proportional to their spatial frequency, and so it is wise to check

if the reconstruction methods can resolve the higher frequencies when their amplitude is smaller than the usual A = 0.25. Consequently, A = 0.125 has also been tested in 4.17. For both amplitudes, all three of the methods are able to match the peaks of the analytical field very well up to a frequency of $\omega \leq 0.1$ voxel⁻¹ ($\lambda_{x,z} \geq 10$ voxels or $\lambda_{x,z} \leq L/6.5$), indicating that none of the methods struggle with large-scale features. When $\omega > 0.1$ voxel⁻¹, the ART and FBP+ART methods show a reduced ability to match the peak power of the analytical gradients, which manifests as under-predicted gradients. All of the methods show a sharp decline in peak power after approximately $\omega = 0.22$ voxel⁻¹ ($\lambda_{x,z} \approx L/14$) up to the Nyquist frequency, which shows that these scales are poorly resolved.

The relative ability of a method is measured in the right column of figure 4.17 simply as the ratio of the method's peak power to the analytical field peak power for each frequency. The criterion for a frequency to be considered 'strongly under-resolved' is defined as when this ratio is below 0.5 (50% strength). It is indeed clear that the ART and FBP+ART methods under-resolve the fluctuations at frequencies that are far lower than for FBP, at both tested amplitudes A = 0.125 and the usual A = 0.25. ART and FBP+ART show similar behaviour in the relative peak power for smaller frequencies, but diverge from each other from $\omega > 0.1$ voxel⁻¹. FBP+ART shows a sharper decline in relative peak power than ART until $\omega = 0.3$ voxel⁻¹. Despite this, both methods reach the 50% strength criterion near $\omega = 0.22$ voxel⁻¹ ($\lambda_{x,z} = 4.5$ voxels $\approx L/14$). FBP crosses the threshold at $\omega = 0.3$ voxel⁻¹ ($\lambda_{x,z} = 3.3$ voxels $\approx L/20$) onwards. Even though FBP performs up to a higher frequency than the other two methods, recall that the artefacts become stronger too, and it may be impossible to distinguish between artefacts and genuine flow features. This is particularly challenging considering that the gradients can legitimately be positive or negative, unlike an intensity-based reconstruction where only a strictly positive quantity is sought. This analysis also shows that the range of resolvable scales is not a significant factor in choosing between ART and FBP+ART; the reconstruction errors and rate of convergence provide more distinction between the two methods.

It is apparent that the reconstruction of the small scales can suffer well before they approach the Nyquist frequency of the optical system, for the amplitudes tested. The resolution of higher frequencies may improve with a greater number of cameras, but the number of cameras required to resolve a significantly larger range of frequencies (i.e. up to $\omega = 0.5 \text{ voxel}^{-1}$) would be impractically high. This suggests that the best way to resolve small-scale features is to increase the magnification of the flow, so that they appear larger relative to the size of a voxel. But as discussed in section 3.1, defocus blurring may increase as well, which would also prevent these scales from being captured.

Referring back to figure 4.6, these frequency characteristics also makes selection of the appropriate Poisson discretisation scheme easier. It is not necessary for the discretisation scheme to maintain a high gain, matching the analytical transfer function, up to the highest frequencies (i.e. beyond $\omega \gtrsim 0.25$ voxel⁻¹). The reconstructions themselves resolve these frequencies poorly, and they are likely to become polluted with displacement field noise, so it is even helpful if the Poisson scheme can attenuate the information at these frequencies. A lower-order scheme will complement the characteristics of the reconstruction methods.

Returning now to the well-resolved $\lambda_{x,z} = L/14$, consider the predicted refractive index fields. The refractive index fields that result from the application of the Poisson solution (with the 3, 5 scheme) to the FBP, ART and FBP+ART reconstructed gradient fields are shown in figure 4.18. The predictions inside the jet are very similar for the three



Figure 4.18: Contour plots, and profiles through x = 0 (black dotted line is the original synthetic field and red dashed line is the reconstruction), for the reconstructed Poisson solved refractive index $n_0 - n$ for 16 camera reconstruction and $\lambda_{x,z} = L/14$: a) synthetic field; b) FBP; c) selected ART 100 iterations (case A in figure 4.13); d) selected FBP+ART 100 iterations (case E in figure 4.13).

methods. Across the whole domain, ART and FBP+ART are almost indistinguishable and match the true field very closely, but the FBP reconstruction has persistent artefacts in the ambient region. Because both the reconstruction and discretised Poisson equation resolve the $\lambda_{x,z} = L/14$ wavelength quite well, the measured refractive index fields strongly resemble the analytical field.

4.5.6 Effect of displacement field noise

Up to now, this study has only considered the accuracy of the reconstruction using perfect background displacements. In reality, the displacements are estimated, which introduces random noise that is propagated through the reconstruction and the Poisson



Figure 4.19: RMS error (top row) and peak error (bottom row) in the refractive index gradients ∇n (left column), and refractive index fields n (right column) for 16 camera reconstruction and $\lambda_{x,z} = L/14$ calculated within twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$, as a function of the standard deviation of added noise normalised by the peak displacement $\sigma_{noise}/\Delta X_{peak}$. Shown are: FBP •, selected ART case ×, and selected FBP+ART case \triangle (the latter two both correspond to the same markers in figure 4.13). The Poisson solution uses 3- and 5-point kernels for the left- and right-hand sides, respectively. Each point is averaged over 100 samples of added noise; error bars indicate the 95% confidence level and are approximately the size of the markers.

solution. For example, when using standard PIV algorithms to estimate the deflection of the background, one would expect a random error on the order of 0.1 pixels [128], which for the current maximum background displacement of 2 pixels corresponds to a 5% random error. To investigate the effect of the displacement field random error on the reconstruction, Gaussian-distributed noise is added to the displacements. The added noise is expressed as a percentage of the true peak displacement $\sigma_{noise}/\Delta X_{peak}$. This was done for each camera. The results presented are averaged over 100 instances of added noise to the reconstructions.

The influence of this noise on the $\lambda_{x,z} = L/14$ reconstruction from 16 cameras with different fluctuations wavelengths are shown in figure 4.19. The RMS and peak errors in the gradients and Poisson-solved refractive index field within $r \leq 2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ increase steadily with increasing noise level. The RMS error in all three methods increase at similar rates, but the peak error in ART appears to grow at a faster rate than FBP



Figure 4.20: Average absolute errors in the refractive index gradients ∇n (a) and refractive index fields n (b) as a function of radial position r/σ for noise level $\sigma_{noise}/\Delta X_{peak} = 5\%$ for 16 camera reconstruction and $\lambda_{x,z} = L/14$. Shown are: FBP •, selected ART case \times , and selected FBP+ART case \triangle (the latter two both correspond to the same markers in figure 4.13). Faint lines indicate the error with no added noise (figure 4.15). The Poisson solution uses 3- and 5-point kernels for the left- and right-hand sides, respectively. Each point is averaged over 100 samples; error bars indicate the 95% confidence level and are approximately the size of the markers. The red vertical line indicates the $2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ limit up to which the RMS error is calculated.

or ART with increasing noise level. The outer noise in the FBP+ART appears to be due to stronger artefacts in the initial FBP solution, which are retained through the ART iterations. As an aside, similar behaviour was observed for other wavelengths down to $\lambda_{x,z} = L/14$, and the RMS and peak error decreased only marginally with increasing camera number up to 22 cameras.

From figure 4.20, the increase in absolute error as a function of radius seems quite evenly distributed across the domain. Although ART shows the highest increase in error towards the centre of the domain $(r < \sigma)$, the error in the ART reconstruction is still considerably lower than that of FBP and FBP+ART. FBP+ART shows the greatest increase in the region $1 < r/\sigma < 2$. A look at the noisy gradient field reconstructions, and Poisson-solved refractive index fields, in figure 4.21 shows that the jet's features are affected negligibly, although the edge of the jet develops a patchy appearance. A reasonably small noise level such as $\sigma_{noise}/\Delta X_{peak} = 5\%$ should not be too troublesome in terms of affecting the quality of measurements of turbulence measurements within twice the half-width, but additional filtering of the outer regions may be required.

4.6 Summary and conclusions

This study focussed on the development of a robust tomographic reconstruction algorithm for TBOS. The FBP, ART and FBP+ART methods were compared using the reconstruction of a heated jet phantom. The phantom has a variable spatial fluctuation wavelength, which allows the range of resolvable scales to be characterised. A range of 2 to 22 cameras in the virtual setup were spaced equally in a 180° arc in a plane transverse to the jet axis. This configuration was chosen because it is convenient to implement in an experi-



Figure 4.21: Contour plots, and profiles through x = 0 (black dotted line is the original synthetic field and red dashed line is the reconstruction), for reconstructed $\partial n/\partial x$ (left column) and $n_0 - n$ (right column) with $\lambda_{x,z} = L/14$ and 16 camera reconstruction from displacements with $\sigma_{noise}/\Delta X_{peak} = 5\%$ added noise for: FBP (top row); selected ART (middle row); selected FBP+ART (bottom row).

mental setup. The optical parameters for this chapter were chosen to remove the impact of defocus blurring and concentrate the analysis on the behaviour of the reconstruction algorithms.

It was shown that FBP suffers greatly as the number of cameras used for reconstruction is reduced, with RMS errors and peak errors increasing exponentially. It is not practical to use FBP with less than 12 cameras to measure fluctuations across the tested wavelengths. The main disadvantage of FBP is the introduction of reconstruction artefacts to the measured field, which can be indistinguishable from flow features.

A range of modifications were introduced to the ART algorithm which focussed on artefact reduction and reducing the under-prediction of the reconstructed gradients. The sharp cut-off mask, inversely iteration-weighted Gaussian filter and progressivelytightened Gaussian mask had the largest impact on reducing the reconstruction error and mitigating the artefacts. The optimised ART algorithm provide reasonable reconstructions with as few as 6 cameras. The reconstruction time using the optimised ART is orders of magnitude higher than FBP, but the high-quality reconstruction means that it is the preferred method in this study.

The FBP+ART method has shown mixed results. Without appropriate modifications, FBP+ART encourages the growth of the FBP artefacts. This results in a solution with a higher error than either FBP or the basic ART alone. The rationale for using FBP+ART is to improve the rate of convergence compared to ART, but this was achieved only for certain modifications. The fastest-converging FBP+ART reconstructions featured the progressively tightened Gaussian mask and converged in a fifth of the iterations required for the optimised ART, but with marginally higher error.

The error in all reconstructions increases steadily with the strength of Gaussian noise added to the displacement fields. The largest impact of noise is found outside the flow, in the ambient regions, which lead to a patchy appearance there. For the typical displacement field noise level of ± 0.1 pixels (5% of peak displacements in the current case), the impact on the final jet density and temperature measurements is expected to be negligible.

For all reconstruction methods, the degree to which the fluctuations are poorly resolved increases with fluctuation frequency. This study introduces the definition of underresolved fluctuation to mean the peak power of the reconstructed field being less than 50% of the peak power of the analytical field. The optimised ART and FBP+ART reconstruction methods under-resolve fluctuations with a spatial frequency $\omega \gtrsim 0.22$ voxel⁻¹ ($\lambda_{x,z} < 4.5$ voxels), or a little more than half the Nyquist frequency of the optical system. FBP under-resolves frequencies $\omega \geq 0.3$ voxel⁻¹ ($\lambda_{x,z} < 3.3$ voxels), but the generation of strong artefacts precludes its use.

The behaviour of the Poisson solver used to obtain the refractive index field is strongly dependent on the discretisation scheme used for both the refractive index field (left-hand side) and the discretisation of the reconstructed gradients (right-hand side). As it is known that the optimal ART under-resolves frequencies $\omega \gtrsim 0.22$ voxel⁻¹, the Poisson discretisation scheme that is selected only needs to transmit information up to this frequency. As the higher frequencies are likely to be more affected by noise, it is actually advantageous if they are attenuated. For these reasons, a lower-order discretisation of the gradients (right-hand side) should be used, e.g. the (3, 3) scheme. If it is required to obtain more information on scales smaller than 4.5 voxels, the only reasonable solution is to modify the optical characteristics of the TBOS system, such that higher magnification is achieved with small defocus blurring.

Chapter 5

Numerical validation of a proposed experimental setup and assessment of three-dimensional TBOS density measurements of a heated jet simulation

Part of the research presented in this chapter has been already been published in the article by Amjad et al. [4], which is also listed in the **Publications and awards from this** work.

Part 2 of the research roadmap. Validation using a realistic flow: the heated jet DNS.

5.1 Aims and overview of the chapter

The previous chapter focussed on developing and characterising tomographic reconstruction algorithms for TBOS density measurements using a density phantom. Now, this study will numerically validate the optimised reconstruction algorithm for a proposed 15-camera experimental setup on a significantly more complex and realistic test case, the density field of a heated jet direct numerical simulation (DNS) following the investigation methodology in section 4.3.

As well as gauging the accuracy of the reconstructed fields, the effect of increasing defocus and temporal blurring on the measurement will be quantified. These results are used to establish practical limits for defocus blurring and temporal blurring in heated jet measurements, in terms of the jet characteristic length scale and characteristic time scale, respectively. The limits provide a guideline for experiment design in the next chapter.

Finally, a numerical assessment of FBP, ART and a combined FBP+ART method for instantaneous 3D turbulent density measurements is presented. The performance of these algorithms is gauged with downstream location, as the jet transitions from a laminar state at the exit to fully turbulent at the end of the domain. This will demonstrate the spatial scales that can be resolved with each technique. The effect of displacement field noise on the experimental reconstruction is also evaluated.

5.2 Details of the heated jet direct numerical simulation (DNS)

As in the previous chapter, simulated BOS displacement fields are created to assess the reconstruction methods by ray tracing through a known refractive index field to a multiple-camera BOS setup. This time, the refractive index field is obtained using the Gladstone-Dale relation on the density field of a heated free jet simulation. An inhouse high-fidelity multi-block DNS parallel code (ECNSS) that has been tested and validated in previous studies is used to solve the compressible Navier-Stokes equations [63, 64, 66, 132, 133]. Karami et al. [65] give full details on the numerical method.

The configuration is a heated free jet with an exit-to-ambient density ratio of $\rho_e/\rho_{\infty} =$ 0.8 (an exit temperature of 361 K assuming an ideal gas with constant pressure and an ambient temperature of 300 K), with a Reynolds number based on nozzle diameter Dof $Re_D = 10,000$ and bulk nozzle exit velocity U_e , Mach number Ma = 0.6 and an initial momentum thickness of 0.02D. The jet issues from an orifice in an isothermal flat plate at ambient temperature. The simulation domain extends to 30D in the streamwise direction x (with a sponge region from x/D > 20) and 15D in the radial direction r. Convective boundary conditions are used at all boundaries, except the flat plate. There are $N_x \times N_r \times N_{\theta} = 1280 \times 448 \times 144$ computational grid points in the streamwise, radial and azimuthal directions, respectively. A uniform grid is employed in the azimuthal direction. In the axial direction, stretch grid points are used, with two stretching rates; a fine grid is used in the region 0 < x/D < 20. In the radial direction, a fine grid is employed in the centre of the mixing layer with a polynomial stretch of the grid towards the centre of the jet and the far-field. The Kolmogorov scale is defined as $\tilde{\eta}_k = (\tilde{\nu}^3/\tilde{\epsilon})^{1/4}$ and is used to assess the turbulence resolution in this test case. The simulation was run for 3 jet through-times, $t_j = L_x/U_e$ where L_x is the domain length, to obtain a statistically stationary solution. This study uses 104 snapshots from the last $0.5t_j$ in the present study with an interval of $0.35t_c$, where t_c is the characteristic time scale of the jet. The minimum $\tilde{\eta}_k/dx$ occurs in the shear layer, where a resolution of $\tilde{\eta}_k/dx > 0.5$ is maintained, which is sufficient for DNS [103]. The integration time-step was adjusted to maintain a constant Courant-Friedrichs-Lewy (CFL) number of 1.5.

For subsequent analysis, the density field was interpolated onto a uniform Cartesian grid with a spacing of 0.0204D in all directions. A slice through z/D = 0 of the DNS heated jet density field shown in figure 5.1. The gradients of the refractive index are found with 6th-order central differences. The domain is large enough to assume zero density gradients at the boundaries.

Ten downstream transverse planes are investigated, covering the laminar, transitional, and turbulent regions, from 0.28D to 9.28D 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 applicable in the Poisson equation.

5.3 Validation of proposed experimental setup

A virtual optical setup is created around the DNS heated jet for ray tracing, shown in figure 5.2. This setup is proposed for construction, but the expected quality of the reconstructions should be assessed first. There are 15 cameras placed in a full circle in a y - z plane circumferentially around the jet axis. Each camera has a corresponding


Figure 5.1: Slice through z/D = 0 of the DNS heated jet density ratio field ρ/ρ_{∞} at one snapshot. Flow is from left to right. Dotted lines : show evenly-spaced transverse slices for reconstruction from x/D = 0.28 to 9.28.



Figure 5.2: The proposed 15-camera experimental setup.

background pattern on the opposite side of the jet. Consistent with the previous chapter, and previous studies on TBOS [72, 96], this number and placement delivers reconstruction accuracy towards the upper limit of the TBOS technique for a practical number of cameras. Each virtual camera uses a f = 25 mm focal length lens. The usable range of apertures will be decided based on the findings of this chapter to limit defocus blur. The choice of aperture will also influence the illumination used in the experiment, because if very small apertures are required then powerful pulsed lighting must be used to both limit temporal blur and form well-lit images.

The front of each camera lens is $Z_A = 490$ mm from the jet axis, and the background is a further $Z_D = 500$ mm behind the jet axis. The resolution at the object plane (jet) is $68 \,\mu\text{m/px}$, based on a camera with $3.45 \,\mu\text{m}$ pixels. The jet diameter is 10 mm, meaning each grid point in the DNS represents a 3 pixel 'inherent' window, or grid size $d_{gp} = 3$ pixels, while the blur is 12 pixels from equation 3.3. Blur is simulated by applying a box filter to each point with its immediate neighbours. The reader is referred to Appendix B.1 for further discussion on the definition of the blur kernel.

The number of steps for ray tracing was chosen to be at least the number of grid points along the diagonal of a slice, which reduces the residual in the displacements to less than 10^{-3} pixels (two orders of magnitude smaller than the 0.1 pixel resolution of typical displacement calculation methods). Further details on the convergence of the ray tracing method are provided in Appendix B.2. The ray's direction is held constant once it leaves the volume, until it intersects with the imaging plane and creates the background displacement.

5.4 Results and discussion

A truly accurate TBOS measurement can only be achieved with perfect calibration, zero reconstruction error, no defocus blurring, no temporal blurring, and a perfect Poisson solution discretisation that introduces no truncation error and hence, no spatial averaging. The latter four sources of error will be investigated independently. These are repeatable errors in the TBOS measurement. Additionally, uncertainty in the BOS measurement is present due to random errors from noise in the calculated displacement field. This will also be investigated.

5.4.1 Impact of defocus blurring and other spatial averaging on TBOS measurements

Figure 5.3 shows the correlation coefficient and RMS error introduced to the jet because of blurring compared to the true DNS density field, assuming a perfect reconstruction and Poisson solution are achieved. Blur is nondimensionalised by the nozzle diameter to make the results generalisable for a heated jet. Errors are presented relative to the difference in refractive index between the centreline and ambient values at x/D = 0.28, which is $\Delta n_p = 3.82 \times 10^{-5}$. The black dotted lines compare the effect of the aperture stop f/N on the measurement quality from f/16, f/11 and f/8. The measurement quality decreases considerably with apertures larger than f/11 for this setup, with low correlation (R < 0.9) in the laminar region. The red dashed line shows the spatial averaging introduced to the displacements by a 16-pixel cross-correlation interrogation window, which for this setup is comparable to using f/11. The limiting factor in any BOS measurement may be either the window size or blurring, depending on which is larger.

Furthermore, figure 5.4 qualitatively compares the blurred field in the laminar, transitional, and turbulent regions as blur is increased. The laminar region is more strongly affected for a given blur, with lower correlation and higher RMS error with the true field. The jet has a sharp transition (i.e. strong gradient peaks) between its core and the ambient surroundings in this region and is affected to a greater extent than the gentler gradient fluctuations further downstream. Sharp gradients are smeared with increasing blurring, which is especially evident close to the nozzle, which causes the density to be under-predicted and the jet to appear larger. As the jet transitions to turbulence, the blurring causes the small scales to be lost. Beyond $\delta/D = 6\%$ blurring, some pockets of high- and low-density fluid at x/D = 5.2 and x/D = 9.2 inside the jet appear to be mixed. But beyond $\delta/D = 11\%$ blurring, the small-scale features are lost and the jet is



Figure 5.3: Correlation coefficient (a) and RMS error (b) between the true and computed refractive index fields for one snapshot as a function of axial distance x/D and blurring δ non-dimensionalised by nozzle diameter D. Dotted lines \cdots indicate (from bottom to top) blurring for the current optical setup with apertures varied from f/16, f/11, and f/8. Red dashed line - - indicates blurring from 16-pixel interrogation windows on the displacement fields.

severely smeared, with the correlation coefficient R < 0.9.

Therefore, $\delta/D = 11\%$ is established as the limit of acceptable defocus blur for heated jet measurements to preserve large scale structures. Note that preserving smaller scale structures depends not only on defocus blur, but also the grid resolution. During the design of a BOS experiment, the researcher should check that this $\delta/D = 11\%$ criterion is not exceeded by any camera, and can adjust the focal length f, aperture stop f/N, or relative distances Z_A/Z_B and Z_D/Z_B accordingly. Often, the experiment is constrained by the available physical space, and using adjustable-focal length lenses for a multiplecamera setup is expensive, so the only practical option is to use smaller apertures. This will limit the light collected by the camera during each image exposure, so it is tempting to use a longer exposure. But this leaves the measurement vulnerable to temporal blurring, which will be investigated in the next section.

Returning to spatial averaging, it is expected that the error introduced to the solution through defocus blurring will often be much greater than the filtering introduced by the Poisson solver discretisation. The previous chapter's study set the spatial resolution of the virtual optical setup such that the blurring was smaller than the grid size and hence, any filtering effects related to discretisation. This may not always be the case, so it is wise to check if the effect of defocus blurring or discretisation-related filtering is larger when the two are of comparable physical size. Figure 5.5 compares the following two cases against the true DNS density field: the Poisson solution of the true DNS gradients, and the Poisson solution of the blurred gradients. To decouple the effect of the growing jet size for the constant domain size on the RMS error, the RMS error is only calculated within the mean local jet diameter that is found from the displacements (shown in figure 5.8). Furthermore, it can be concluded from figure 5.5 that, for a given blur (e.g. $\delta/D = 6\%$), there is negligible benefit to the measurement accuracy in using a grid resolution that is smaller than the blur (e.g. $d_{gp} = \delta/5$ and $\delta/15$ in figure 5.5) and when the blur is larger



Figure 5.4: Cross-section contours of normalised 'excess' density for one snapshot (same as previous figure) as a function of blurring δ/D and axial distance x/D.

than the spatial averaging due to the Poisson solution. Together with our findings in the previous chapter, it is demonstrated that it is acceptable to use lower-order discretisations, and perhaps even preferable to avoid propagating measurement noise, because the defocus blurring is larger.

It is clearly demonstrated, that in the proposed setup, the defocus blur is the dominant source of spatial averaging, not the filtering associated with discretising the Poisson equation. Significant improvement to the measurement can only be achieved if the dominant source of spatial averaging is reduced. Thus, it is a priority to restrict the defocus blurring in our experiment. The largest errors are occur in regions of sharp gradients; this is especially prominent in the shear layer of the transitional region.



Figure 5.5: RMS (a) and peak (b) errors between the Poisson solution of the transverse gradients and the true DNS density field as a function of axial distance x/D. Errors are normalised by the peak refractive index difference near the nozzle $\Delta n_p(x/D = 0.28)$ and averaged over the 104 samples. Marker \bullet represents the Poisson solution of the true gradients with no blurring and (3, 5) discretisation. Markers \bigstar and \blacksquare represent the Poisson solution of the gradients with blurring $\delta/D = 6\%$ for grid resolutions of $d_{gp} = \delta/15$ and $\delta/5$, respectively. Error bars are approximately the size of the markers and indicate a 95% confidence level.



Figure 5.6: a) Instantaneous density ratio ρ/ρ_{∞} as a function of the characteristic time scale of the jet $t_c \bullet$ for all available samples at x/D = 5.2, y/D = 0, z/D = 0.5 (i.e. r/D = 0.5), with the running average given by red crosses \times . b) Autocorrelation of the density ratio signal at the same location, with the red dashed line indicating $t_c = 1$.

5.4.2 Impact of temporal blurring on TBOS measurements

While it is true that increasing the light intensity allows the dependence between sensitivity and defocus blur to be uncoupled, with reference to equations 3.1 and 3.4, the issue of temporal blurring and exposure time is then brought into play. Often the only practical way of controlling defocus blurring is to reduce the lens aperture. This can severely restrict the amount of light that is collected by the sensor during an exposure for the BOS images because each successive stop on a standard aperture scale halves the aperture area. Decreasing the aperture to very small diameters to reduce the defocus blurring then places a premium on the illumination intensity while the requirements for a short exposure time to freeze the flow remain unchanged. To ensure adequate illumination, one can either use longer exposures or more powerful lighting. A high-quality measurement must mitigate both defocus and temporal blurring.

It can be expensive to change lighting systems in an experimental setup, while using a longer exposure is more convenient. But long exposures increase the amount of temporal blurring (temporal integration) in the measured density field. It is not possible to conduct a truly 'instantaneous' measurement, as the camera must integrate the flow measurement over some interval, no matter how small. Thus, it would be wise to see how the quality of the 'instantaneous' measurement is degraded as the temporal integration range increases.

The temporal integration can be expressed non-dimensionally as a function of the jet's characteristic time scale, $t_c = D/U_j$. Increasing temporal blurring will first obscure the smallest scales. As the temporal integration increases, larger and larger scales will be affected. In figure 5.6, the density field is probed at x/D = 5.2, r/D = 0.5 (the shear layer) as a function of time scale t_c . Figure 5.6a illustrates how rapidly the signal can change as a function of time. For example, in this figure the density ratio fluctuates from $\rho/\rho_{\infty} = 1.02$ at $t_c = 8.5$ to $\rho/\rho_{\infty} = 0.79$ at $t_c = 10$.

The autocorrelation of the signal in figure 5.6b confirms the previous observation, and it shows that the correlation of the measured signal decreases very rapidly within an integration time of $t_c = 1$. After an exposure of $t_c = 1$, the density field will have little resemblance to that at the beginning of the integration, such that the correlation decreases to R < 0.5. Therefore, the temporal integration must be limited to a fraction of the jet's characteristic time, say $0.1t_c$, to preserve the measurement quality.

Furthermore, the temporal blurring must not exceed the time taken for flow structures to convect at a speed of U_j over the length of a grid point d_{gp} , which is $t_{gp} = d_{gp}/U_j$. This time will often be smaller than the $0.1t_c$ requirement. Overall, the exposure time must be less than the smaller of $0.1t_c$ and t_{qp} .

Consider a typical value of t_c for experimental measurements in a heated jet. For example, the DNS heated jet has a Reynolds number $Re_D = \frac{\rho U_j D}{\mu} = 10,000$ and $\rho_e/\rho_{\infty} =$ 0.8. Creating an air jet experimentally at these conditions in a standard laboratory environment, with D = 10 mm, one obtain an exit velocity of $U_j = 22 \text{ m s}^{-1}$ (calculations are covered in detail in Appendix C.2), such that $t_c = 455 \text{ µs}$ (0.1 $t_c = 45 \text{ µs}$). The grid convection time for this setup is $t_{gp} = 9.3 \text{ µs}$. Therefore, the maximum acceptable exposure time should be 9.3 µs.

Combined with the use of small apertures, this almost certainly necessitates using very powerful lighting, such as an over-driven pulsed LED or pulsed laser.

The previous studies on non-time-averaged TBOS used the following light sources and exposure times:

- Atcheson et al. [6] used 800 W halogen stage lights. The exposure time was not specified.
- Nicolas et al. [96] used 500 W halogen spotlights. The exposure time was 750 µs.
- Lang et al. [72] used 1000 W halogen lights. The exposure time was 2 ms.

- Nicolas et al. [97] used a pulsed laser with a 10 ns exposure time, although they did not use the laser speckle method introduced in the next chapter. This was required to mitigate temporal blurring, i.e., freeze the flow.
- Grauer et al. [48] used 200 W LED floodlights. The exposure time was 300 µs for adequate illumination.

Most of these studies achieve adequate lighting, at the expense of a longer exposure time and hence, larger temporal blurring. Only the laser illumination method used by Nicolas et al. [97] would be appropriate for the present study. Preliminary experimental tests of over-driven LED lighting shows that the required illumination and short exposure time can be achieved in a narrow field of view. Thus, each camera would require its own LED, which is prohibitively expensive and also difficult to package when using the design like that of Buchmann et al. [21]. So, this study will consider spreading a high-power, pulsed laser beam as a possible solution to this problem in the next chapter.

5.4.3 Comparison of FBP, ART and FBP+ART reconstructions of the heated jet

The error introduced by the tomographic reconstruction is in addition to that introduced by blurring and the Poisson solution spatial averaging. The last chapter focussed on the development of the optimised ART and FBP+ART algorithms for high-quality reconstructions of the synthetic density field phantom. Now, it is must shown that the same high-quality measurements can be expected for the DNS heated jet and, by extension, the experimental TBOS measurements. Figure 5.7 presents the RMS and peak errors from FBP, ART and FBP+ART Poisson-solved reconstructions, in the absence of displacement field noise. The errors are relative to the Poisson solution of the blurred, true DNS gradients. This allows the reconstruction error to be isolated from the error due to blurring. Further, figure 5.8 compares the distribution of error in the FBP and ART reconstructions as a function of radius and downstream position. The mean jet radius \overline{r} at each axial station is identified based on the radial position where the mean displacement from all cameras is less than 0.1 pixels, $\left| \vec{\Delta X} \right| < 0.1$ pixels. The radius of the sharp cut-off mask r_m at each instant at each axial station is identified as the maximum radial position from all cameras where the displacement magnitude is less than 0.1 pixels with a 10% margin. Figure 5.9 provides a direct comparison of the blurred, reconstructed, Poisson-solved density fields at x/D = 5.2 for FBP, ART and FBP+ART against the density field from the blurred, Poisson-solved true DNS gradients.

Figure 5.7 shows that the FBP introduces the largest error of the three methods across the jet's length, as expected. Figure 5.8 shows that the FBP can capture the jet's laminar core (up to x/D = 2), but it introduces strong reconstruction artefacts immediately outside this area, as seen by the increase in the error. The streaky nature of these artefacts can be seen in figure 5.9. The artefacts still fall within the detected mask, and so are not removed. Figure 5.8 also shows that without the sharp cut-off mask, the FBP artefacts continue to the domain edges and that the overall error of the FBP inside the jet is higher than for ART. The error for FBP appears more evenly distributed inside the jet further downstream. As expected, the number of cameras that are typical of a BOS experimental setup are therefore insufficient for high-quality FBP reconstructions of the instantaneous turbulent density field of a heated jet.



Figure 5.7: RMS (a) and peak (b) errors between the Poisson solution of the reconstructed gradients and the Poisson solution of the blurred true gradients as a function of axial distance x/D. Errors are normalised by the peak refractive index difference near the nozzle $\Delta n_p(x/D = 0.28)$ and averaged over the 104 samples. Markers correspond to: FBP \bigstar ; ART 10 iterations from null initial conditions \Box ; ART 40 iterations \blacksquare ; FBP+ART 10 iterations \odot ; FBP+ART 40 iterations \blacksquare ; FBP+ART 10 iterations \odot ; FBP+ART 40 iterations \bullet ; FBP+ART 100 iterations \bullet . All optimised ART/FBP+ART cases use the sharp cut-off mask, relaxation parameter $\lambda_i = 0.5$, inversely iteration-weighted Gaussian filter, Hamming windowed corrections, random camera and ray order and Gaussian mask. FBP+ART initial FBP solution is filtered with the Gaussian mask. Error bars are approximately the size of the markers and indicate a 95% confidence level.



Figure 5.8: Absolute error as a function of radius (normalised by nozzle radius R) and axial length (normalised by nozzle diameter D) for a) FBP, and b) ART 100 iterations, both averaged over the 104 samples. Markers \bullet show the mean jet radius \overline{r} calculated based on displacement magnitude threshold $|\Delta X| < 0.1$ pixels. Markers \bullet show the calculated sharp cut-off mask radius r_m . Radial bin size is equal to 0.4R.



Figure 5.9: Cross-section contours of normalised 'excess' density for one snapshot at x/D = 5.2, and corresponding centreline profiles at z/D = 0, for: a) Poisson solution of the blurred gradients; b) FBP; c) optimised ART with 100 iterations; d) optimised FBP+ART with 100 iterations.

The ART and FBP+ART methods fare much better than FBP for the same number of cameras. Figures 5.7 and 5.8 show that the RMS error is approximately halved along the length of the jet, while figures 5.8 and 5.9 show that reconstruction artefacts are not introduced outside the jet. From figure 5.8, ART errors decay in the shear layer even within the sharp cut-off mask due to the Hamming windowed corrections. Figure 5.9 qualitatively demonstrates how prominent the artefacts from FBP can be. The peak errors of ART therefore appear closer to the jet centreline, unlike FBP. ART generally predicts all regions of the jet core more accurately than FBP and converges quickly. From figure 5.7, the RMS solution error is reduced by only approximately 0.5% from 10 to 40 iterations (average across the jet length), and even less from 40 to 100 iterations (0.1% average across the jet length). This agrees well with our findings in the previous chapter, and shows that a marginal sacrifice in reconstruction accuracy can greatly reduce the computational cost.

Because FBP does not perform well in the jet core, it is expected that FBP+ART could hardly be an improvement on ART from null initial gradients. Indeed, without filtering the FBP gradient fields before using them as the initial solution to ART using a Gaussian mask (equation 4.26), the FBP+ART solution can have a much higher error than ART which is only marginally lower than FBP, consistent with the previous chapter. With masking and progressively relaxed Gaussian filtering in place, FBP+ART is only marginally worse than the ART from null gradients. Figure 5.7 shows that FBP+ART converges at a similar rate to ART, discussed previously, and that the RMS error across the jet length is like ART. The peak errors of FBP+ART vary slightly more from ART, with a tendency to have a marginally higher error than ART (average 0.1% higher) at 10 iterations and slightly lower peak error (average 0.5% higher) at 100 iterations. The idea of using FBP as the initial solution to reduce ART underprediction by preserving regions of high (or overpredicted) [53] is not realised in practice due to the strong artefacts throughout the domain (figure 5.9), exacerbated by the Gaussian mask which is necessary to reduce the growth of the FBP artefacts. This contrasting result to the previous studies on FBP+ART can be explained in terms of the flow under study. The present flow is significantly more complex than the one used by Hartmann and Seume [53], with a wide range of scales to be captured, and is not reconstructed adequately by FBP. It confirms that for FBP+ART to have any advantage over ART, the FBP solution itself must be adequate. Overall, it is concluded that FBP is a more disruptive initial solution to ART for the instantaneous density gradient field reconstruction, marginally slowing the ART convergence down compared to starting from null gradients.

5.4.4 Effect of displacement field noise on ART reconstruction

Figures 5.10 and 5.11 show the effect of adding noise to the displacement field. In line with the typical accuracy of cross-correlation methods [128], uniform random noise in the range ± 0.1 pixels was added to each window in the displacement field. For the given optical setup, the peak displacements were approximately 1 pixel, with a mean of approximately 0.7 pixels from all cameras across the length of the jet, thus the added noise is approximately 15% of the typical displacements. Without any additional filtering or conditioning of the displacements the accuracy in the laminar region suffers most, from figure 5.10. This is due to corruption of the sharp displacement peaks which contribute to the 'top-hat' profile of the laminar region. In figure 5.11, the overall appearance of the jet at x/D = 5.2 is retained despite the noise. Gaussian filtering, Hamming win-



Figure 5.10: RMS (a) and peak (b) errors for the ART reconstruction with 40 iterations, averaged over the 104 samples, as a function of axial distance x/D. Markers \bigstar show the ideal case with no displacement field noise (identical to figure 5.7); markers \blacklozenge show reconstruction accuracy with random uniform noise added to the displacements in the range ± 0.1 pixels. Error bars are approximately the size of the markers and indicate a 95% confidence level.

dowed corrections and gradual unmasking marginally improve the solution, particularly by smoothing out sharp local fluctuations. However, the input displacements themselves are affected (which ART compares its solution against) and so the solution remains noisy throughout the iteration process. The features of the jet core appear to be preserved, but outside the jet core patchy regions are observed. These may be due to the comparatively large-magnitude noise in these low-displacement regions. The spatial averaging that is imparted by the solution of the Poisson equation also helps to reduce the impact of the noisy gradient field on the reconstructed density field. These observations are consistent with the previous chapter and indicate that the typical levels of displacement field noise will have a negligible effect on the measurement quality. Appendix B.3 also investigates potential improvements in accuracy from noisy displacements using an anisotropic diffusion scheme in the Poisson solver, as used by Atcheson et al. [6]. No improvement was observed by using this scheme.

5.5 Summary and conclusions

This study aimed to validate an experimental TBOS setup, and assess the impacts of defocus blur, temporal blur, and reconstruction error on TBOS measurements of a heated jet. The method of ray tracing through a heated jet DNS density field allows these sources of error in the TBOS measurement to be examined in isolation from one another and the uncertainty due to random error.

Spatial averaging in the TBOS measurement process results from two sources, one inherent from the defocus blurring, and the other a second-order effect from discretising the Poisson solver. This study shows that the impact of blurring is much greater than the spatial averaging introduced through the discretisation of the Poisson equation, for the proposed experimental setup. Both sources are seen to produce higher RMS and



Figure 5.11: Contours of normalised 'excess' density for one snapshot (on the same colour scale as figure 5.9) at x/D = 5.2, and corresponding centreline profiles at z/D = 0, for: a) the Poisson solution of the reconstructed ideal blurred gradients (no displacement field noise), and b) Poisson solution of the reconstructed gradients from noisy displacements. Dotted line \cdots indicates Poisson solution of blurred gradients (figure 5.9a).

peak errors in the region x/D < 4 and gradually decrease thereafter. The errors due to blurring are hardly affected by using a finer grid resolution. The sharp gradient peaks contributing to the 'top-hat' profile in the laminar region are smoothed due to spatial averaging and this leads to an under-predicted solution. The correlation coefficient may decrease to below R = 0.9 in this region when the blur $\delta/D > 11\%$.

Temporal blurring can also severely degrade the measurement quality. The temporal integration time must be limited to just a fraction of the jet's characteristic timescale, as well as being less than the grid convection time. In practice, this necessitates microsecond exposures, and very strong lighting for the experimental setup.

The optimised ART performs well on the complex DNS heated jet density field. The ART solution converges quickly, with little change observed between 40 to 100 iterations. This indicates that a compromise between computation time (which is significantly greater than for FBP) and accuracy can be made with minimal impact on measurement accuracy due to the small change in error. The best ART reconstruction adds only a modest amount to the blurring error, at most an additional 4.5% RMS and 25% peak of the blur error in the transitional region from x/D = 3 to 6, for the proposed setup. Thus, efforts to improve the accuracy of TBOS should focus on reducing defocus blurring rather than improving the reconstruction quality.

This study establishes an optimised reconstruction method, a validated experimental setup concept for heated jet measurements, and practical limits for defocus and temporal blurring. Next, these will be applied to obtain experimental 3D density measurements.

Chapter 6

Experimental three-dimensional density field measurements of a heated jet using laser-speckle TBOS

Part 3 of the research roadmap. Demonstration of an improved experimental technique: pulsed laser-speckle TBOS.

6.1 Aims and overview of the chapter

The previous two chapters focussed on developing an optimised TBOS reconstruction algorithm, validating an experimental setup concept, and establishing the limits of defocus and temporal blurring. This chapter reports on a study that applies these findings to construct an optimal experimental setup for TBOS 3D density measurements in a heated jet and its analysis. The most pressing need is to restrict the temporal blurring, and so this study considers the laser-speckle BOS technique, where a laser is used for background illumination. In contrast to the limited number of previous studies on laserspeckle BOS, the present work proposes using a high-power, pulsed laser, which makes the temporal integration negligible. Lasers are a coherent light source, and this can be taken advantage of to both create a background pattern and illuminate the setup. An important contribution of this chapter is to present a systematic selection methodology for the camera lens focal length, aperture and focussing distance which restricts defocus blur to permissible levels while finding a good compromise with the measurement sensitivity. Finally, experimental density and temperature measurements of a heated jet are reported, which are validated against thermocouple measurements and compared with a heated jet DNS. The behaviour of the jet's potential core is examined, which can provide insight on the unique behaviour of variable density jets.

6.2 Laser-speckle BOS

Lasers are a coherent light source, meaning that the light may be treated as waves of identical frequency that may interfere. Speckle patterns are generated when coherent light is diffusely reflected off a surface with roughness of similar size to the light's wavelength. The phase differences of the diffusely reflected light propagating through free space results in interference. The incident light on an observation plane will be viewed as a pattern of light and dark patches known as speckles due to the interference. The characteristics of the speckle pattern depend on whether the observation plane is also located in free space, in which an 'objective speckle pattern' is formed, or if optical elements such as lenses are used to record the light on an imaging plane, in which a 'subjective speckle pattern' is formed [146].

The speckle pattern that is formed depends on the roughness distribution of the surface that diffusely reflects the laser. Rather than seeking to describe the pattern exactly, it is more common (and practical) to describe the pattern statistically instead. Goodman [46] provides a detailed explanation and derivation of speckle statistics, and finds that the intensity of speckles on discrete camera pixels follows a gamma distribution. Although quantisation effects and aliasing should be considered if the aim is to faithfully capture the speckle pattern [127], it is not critical to fully-resolve the speckle pattern to conduct accurate speckle measurements which require recording the displacements of speckles [76].

Laser speckle BOS is a promising solution to the sensitivity/defocus blurring dilemma. Although used extensively in other fluid diagnostic techniques [37], laser speckles were introduced more recently to BOS by Meier and Roesgen [81], who replaced the traditional printed background pattern with the speckle pattern generated by illuminating a rough surface from the side with an expanded laser ('single-pass mode'). One can consider laser speckle BOS to be a simpler, more versatile implementation of the older laser speckle photography technique [85]. Meier and Roesgen [81] were able to exactly replicate the BOS displacements using the speckle background as they had obtained using a printed background. The method was developed further, taking advantage of the unique properties of speckle imaging. Firstly, that the speckles are always sharply imaged, regardless of the camera focussing on the speckle surface or closer or further. In this way, the sensitivity and blurring of the setup is decoupled from the most restrictive element of the BOS setup, the overall dimension Z_B , and equation 3.1 for the image displacement vector $\Delta \vec{X}$ and equation 3.5 for the focal plane distance Z_I can now be written as,

$$\vec{\Delta X} = \left(\frac{fl}{l + Z_A - f}\right) \tan \vec{\varepsilon},\tag{6.1}$$

and

$$Z_{I} = \left(\frac{1}{f} - \frac{1}{Z_{A} + l}\right)^{-1},\tag{6.2}$$

respectively. Here l is the focus distance along the camera's optical axis measured relative to the centre of the refractive volume, i.e., l = 0 means that the camera is focussed on the measurement volume, l > 0 refers to focussing further away, l < 0 refers to a focussing closer than the volume, and $l = Z_D$ means that the camera is focussed on the background like traditional BOS. The camera may be focussed closely up to $l > -(Z_A + f)$, or further up to $l < \infty$ [81]. The equation for on the image plane (equation 3.4) still has the same form, but the equation for the blur cone at the object plane (equation 3.3) for l > 0 can be rewritten as

$$\delta \approx d_a \left(\frac{l}{Z_A + l}\right). \tag{6.3}$$

The perceived size of the speckles on the image sensor is only dependent on the fnumber (N) being used, due to diffraction-limited imaging. A smaller aperture (larger N) results in larger speckles, and Meier and Roesgen [81] derive a rough relationship for the 'average' speckle diameter d_s on the image sensor based on an imaging model described by Goodman (2007) [as cited in 81],

$$d_s \approx 4\lambda N/\pi,\tag{6.4}$$

but speckles of many sizes coexist due to the random interference of light off the rough surface. Thus, an additional compromise is added to the selection of the aperture, which is to optimise the speckle size.

Meier and Roesgen [81] introduced alternative experimental setups that allow the sensitivity of the measurement to be increased, by illuminating the surface in line with the camera's optical axis, similar to silhouette photography [130]. As shown in figure 6.1, the light passes through the refractive volume twice, once forward to the speckle surface, and then back to the camera ('double-pass'). The authors also introduced an 'interferometry mode' where the inline illumination is again used, but the camera is focussed on the refractive volume itself. This only generates a displacement from the light's forward pass through the volume; the advantage is that defocus blur is removed. While this would solve the sensitivity/defocus blurring dilemma, in practice it may be unfeasible to achieve inline illumination for the typical number of cameras needed for tomographic BOS due to the cost of the required optics such as beam-splitters and mirrors (which may cost more than the cameras themselves). Therefore, only the 'single-pass' mode is considered in this study.

Meier and Roesgen [81] use a low-power continuous wave laser to generate the speckles. The current work will investigate the use of a high-power, short-pulse laser to generate both sufficient illumination and short exposure times for the small-aperture, low temporal integration images for high-quality turbulence measurements. With appropriate selection of the focal length, and varying the aperture and focus distance, it will be possible to obtain well-illuminated images with enough sensitivity for a high signal-to-noise ratio for the displacement calculation method, with blurring smaller than the interrogation window size.

6.3 Experimental setup and method

6.3.1 Camera and laser configuration and control

Following the last chapter, a 15-camera setup is utilised to conduct TBOS density measurements. The cameras are 3 megapixel Daheng MER302-56U3M USB3-interface monochrome machine vision cameras with a pixel size of $3.45 \times 3.45 \,\mu\text{m}$ and 10-bit image recording. The cameras are placed circumferentially around the reconstruction volume, and each face a plane speckle surface (acrylic sheet with adhesive paper film) oriented towards the camera that is 100 mm wide \times 130 mm high. The cameras and opposing backgrounds are roughly equidistant 500 mm from the centre of the reconstruction volume. In order to illuminate all of the surfaces with a single laser, the experimental setup deviates slightly from the proposed setup in chapter 5 by placing all of the cameras on one half of the circumference, and all of the speckle surfaces on the other, shown in figure 6.2. This will have no adverse effect on the measurement quality, because BOS measurements are path-integrated, and the same information is captured by the 15 cameras although half of them are now on the opposite side (tests on the DNS jet confirm this).

A 120 mJ Nd:YAG dual-cavity PIV laser ($\lambda = 532$ nm, New Wave Solo 120) is used to generate the speckle patterns. Only a single cavity is required to conduct the BOS



Figure 6.1: Configurations for laser-speckle BOS introduced by Meier and Roesgen [81]. a) 'Single-pass' mode, where the expanded laser illuminates a surface without passing through the refractive volume (flow) under study. b) 'Double-pass' mode, where the illumination is introduced in-line with the camera's optical axis. Notice that in both of these modes, the camera is not necessarily focussed on the speckle surface $(l \neq Z_D)$ like in traditional BOS (cf. figure 3.1). c) 'Interferometry mode', like 'double-pass' mode except that the camera is focussed on the refractive volume (l = 0).

measurements, as the reference images are taken before the flow is started, and images containing the flow require only a single exposure. The pulse time of the laser is 10 ns, which controls the effective exposure time of the images. The cameras themselves use an exposure time of 400 µs, but do not capture anything when the laser is not firing. The 10 ns exposure is expected to be much smaller than the characteristic time scale of the flow for any practical flow configuration, meaning that temporal blurring is negligible.



Figure 6.2: a) Experimental setup with 15 cameras modified for laser-speckle TBOS, b) schematic of laser expansion. Optical axis of camera 1 is aligned with the global z-axis, and x is the jet axis. The laser beam is guided into the beam splitter at the correct orientation using an articulated arm (ILA 5150 Articulated Mirror Arm, not pictured). Bottom: photograph of expanded laser beams illuminating the backgrounds. The cylinder placed above the nozzle is used to check that the beams do not directly cross the reconstruction volume. In the real experiment, the laser optics and mirror arm are not directly attached to the rig's 'table top' upper surface as shown in (a) and (b), but rather are magnetically secured to the underside of a small bench that fits onto the 'table top' and over the top of cameras. The bench is visible in the bottom-right corner of the bottom photo. Additional annotated photographs given in Appendix C.1.

The cameras and laser are triggered externally using a BeagleBone board to synchronise the timing [36]. A custom-made trigger box relays the signal to all 15 cameras. The MATLAB Image Acquisition Toolbox is used to facilitate image acquisition from the cameras to computer memory via USB3 connection. As the images are quite small, only one computer with additional USB3 ports installed is required. To avoid lost frames, the image acquisition frequency is limited to 5 Hz while the laser operates at 15 Hz.

A plate beam splitter (Thorlabs BSS10) and mirror (Thorlabs NB1-K12) are used to create two branches of the laser beam, as shown in figure 6.2b, which are separately expanded using groups of lenses. Spherical concave lenses (f = -25 mm, Thorlabs LD2297-A) are used to expand each beam to ensure vertical coverage and initial horizon-



Figure 6.3: A severe example of shot-to-shot variations in intensity in the speckle images. No flow is present in either image (ambient environment). The locations of the speckles do not seem to be affected, only their intensity.

tal spreading. The beams subsequently pass through one or two cylindrical concave lenses (f = -9.69 or -15 mm) for additional horizontal spreading (Thorlabs LK1836L1-A or LK1753L1-A, respectively). One branch illuminates 8 speckle surfaces, the other illuminates 7 speckle surfaces. As discussed in section 6.2, it is difficult to enable 'double-pass' or 'interferometry' mode on all the cameras. To ensure that all cameras record in 'single-pass' mode, the expanded laser branches are skirted around the measurement volume on their way to the speckle surfaces, illustrated in figure 6.2b. A cylinder, representing the measurement volume, is placed over the nozzle during alignment of the optics to ensure that the light from the branches does not directly pass through the volume. Only diffuse reflected light from the speckle surfaces may pass through the volume in 'single-pass' mode.

It is expected that there will be shot-to-shot variations in illumination intensity. A severe example of this is given in figure 6.3. These variations do not affect the position of the speckles, which would be a larger issue, when the laser is fully warmed-up before use, which is estimated to take at least 1.5 hours. But the variations will prevent smaller, dimmer speckles from being recorded in occasional shots. Although the impact of this is small, because the observed displacement with no flow present is less than 0.1 pixels, to mitigate this, it is specified that each camera must record multiple (30) reference images. Then each instantaneous flow image from a camera is compared with each of that camera's reference images. The median of these 30 displacements is used as the instantaneous displacement field for that camera for the reconstruction at that time-step.

As discussed in section 3.1, camera calibration is required in tomographic BOS for three reasons: to locate the origin of the global coordinate system; to establish the rotation (pose) and translation of the cameras relative to the global coordinate system; and, to accurately locate the background distance and orientation from each camera. This investigation utilises the multiple-camera calibration toolbox developed by Himpel et al. [56]. The advantage of this toolbox is that it uses a calibration target with dots. Because the global origin is often out of focus in BOS (the cameras are focussed further



Figure 6.4: Calibration target, designed by Himpel et al. [56], imaged from three adjacent cameras. The inverted T-shaped structure is used to determine the orientation of each camera relative to camera 1. The cameras are mounted sideways to align the longer edge of the sensor with the streamwise direction, so the T-shaped structure appears to (subtly) move up and down in the images of adjacent cameras, not left and right.

away or closer to increase sensitivity), it can be difficult to accurately locate the corners of a traditional checkerboard pattern, which become blurred. The centroid of dots can still be located accurately when they are out of focus, encouraging their use for TBOS calibration. Figure 6.4 shows the calibration target viewed simultaneously by three adjacent cameras. Notice that the dots are not sharply imaged. Using at least 10 images per camera with a clear view of the target, reprojection errors around 0.25 pixels in the calibration are obtained. The two-element radial and tangential distortion coefficients k_i and p_i , respectively, are also obtained. A standard polynomial relation between lens distortion and radial distance from the image optical centre is assumed, where the radial distortion affects the image coordinates (X, Y) according to,

$$X_{dist} = X \left(1 + k_1 r^2 + k_2 r^4 \right), \tag{6.5}$$

$$Y_{dist} = Y \left(1 + k_1 r^2 + k_2 r^4 \right).$$
(6.6)

Likewise, the tangential distortion which is due to misalignment between the imaging plane and lens plane, is represented by,

$$X_{dist} = X + \left(2p_1 XY + p_2 \left(r^2 + 2X^2\right)\right), \tag{6.7}$$

$$Y_{dist} = Y + \left(p_1 \left(r^2 + 2y^2\right) + 2p_2 XY\right).$$
(6.8)

These models are used in the OpenCV package to undistort the images with a linear image interpolation scheme before displacement calculation.

6.3.2 Heated jet

The heated jet is generated using a compressed air supply and inline heater (Omegalux AHP-7562). The air mass flow rate is regulated using a pressure regulator and mass flow meter (Alicat M-500SLPM-D), while the jet temperature is controlled by adjusting the heater voltage with a Variac device (up to 240 V at 3 A). The maximum mass flow rate of the system is 500 standard litres per minute, approximately $10^{-2} \text{ kg s}^{-1}$. After passing



Figure 6.5: Cutaway schematic of converging nozzle with matched-cubic profile [93]. Dimensions are in millimetres. The matched cubic contour itself is 50 mm in length from entrance to exit.

through an aluminium settling chamber equipped with three levels of meshes, the jet emits through an aluminium nozzle, with the profile shown in figure 6.5. The nozzle has an exit diameter of D = 10 mm and an entrance-to-exit area contraction ratio of 13.69:1. The internal contour of the nozzle is a matched-cubic profile derived using the method devised by Morel [93].

For each flow configuration, the approximate nozzle centre exit temperature (at x/D = 0.3) is recorded to determine the appropriate warm-up and cool-down times for the jet to reach a stable temperature after the heater is activated. Figure 6.6 shows the warm-up and cool-down temperature traces for the jet when the mass flow rate is set to $1.7 \times 10^{-3} \text{ kg s}^{-1}$ and the heater voltage to 120 V, producing a stable exit temperature of approximately 91 °C. The stable exit temperature is repeatable to within ± 1 °C for the same mass flow rate and heater voltage, seen on four separate runs. The laser is warmed up for 2 hours (half an hour beyond the stable jet temperature) to ensure a stable speckle pattern is maintained. In this experiment, the reference images are taken after the jet has cooled down, because the warm-up time of the jet coincides with the warm-up time of the laser. The jet is allowed to cool for 30 minutes after the heater is turned off before the reference images are taken, which allows the entire rig to reach ambient temperature.

As the air is heated, the settling chamber and nozzle will be heated too. This could affect the boundary conditions of the domain and induce convective currents around the hot nozzle. To mitigate this, several layers of thermal wool insulation (RS Pro Superwool) are wrapped around the heater, settling chamber and nozzle. Thermal images of the fully warmed-up nozzle, for the conditions described in figure 6.6, show that the outermost layers of insulation reach a temperature of approximately 40 °C, as seen in figure 6.7. Thermocouple temperature profiles near the jet, discussed in section 6.5.1, show that the radial boundary condition is still at ambient temperature, indicating the insulation is effective.



Figure 6.6: Thermocouple temperature traces of nozzle centreline exit temperature (x/D = 0.3) as a function of time for a) warmup and b) cooldown (after one hour of stable heating), when the air mass flow rate is 1.7×10^{-3} kg s⁻¹ and the heater voltage is 120 V. Samples are recorded at intervals of dt = 2 seconds.



Figure 6.7: Thermal image of the insulated nozzle when the jet is at an exit temperature of 91 °C. Image captured using Seek Thermal Compact camera.

6.4 A systematic selection methodology for camera lens focal length, aperture and focussing distance

Appropriate selection of lens focal length, aperture and focussing distance requires careful thought. The BOS setup should have the ideal compromise between sensitivity and defocus blur, which have conflicting requirements. The proposed pulsed laser speckle TBOS method allows a short exposure time so that the temporal blurring constraint is met. Even though the laser is very intense, it can still be difficult to achieve adequate brightness in the images when using an aperture small enough to limit defocus blurring.

Thus, the smallest aperture setting which allows a reasonable illumination is constrained. To keep blurring to a reasonable level, one must then adjust the focus distance (laser speckle TBOS) or change the distance between flow and background (traditional BOS). This has the effect of reducing sensitivity, as seen from equation 3.1 and 6.1. So, in practice, there will always be some compromise in sensitivity if it is desired to keep temporal and defocus blurring in check. Incorporating the laser speckle technique into the BOS optical setup provides an additional tool to optimise the measurement with various experimental constraints.

From previous TBOS studies [48, 72, 96], it is not clear how this can be found, as most have shown little regard for finding the ideal compromise. This work outlines a systematic approach that will be useful for researchers. The results of section 5.4.2 show that mitigating the temporal blurring is most critical to high-quality measurements, and thus it cannot be compromised. As the short exposure time constraint has been addressed using the pulsed laser, there are three remaining constraints: spatial (defocus) blurring, sensitivity or signal-to-noise ratio, and appropriate speckle size. The acceptable values for each of these constraints must now be decided.

Digital cross-correlation PIV analysis is used to obtain the displacements, which introduces blurring in each displacement field over an interrogation window. Too large a window will increase the spatial averaging at each displacement vector, and too small a window will increase the random noise. Testing shows that a 16-pixel window appears to offer good compromise, although this ultimately depends on the background speckle density (the characterisation of which is not straightforward). Therefore, the defocus blurring can have a maximum size of 16 pixels without further reducing the effective spatial resolution. Figure 6.8 shows the variation in sensitivity and blur for different focal lengths and apertures. Sensitivity is expressed as $\Delta/\tan\varepsilon$, where Δ and ε are one component of the displacement and deflection angle, respectively, from equation 6.1, as a function of focus distance l for fixed $Z_A = 500$ mm for easy comparison of different focal lengths and focussing distances (the lines begin at the minimum possible focus distance for each focal length); obviously, a higher sensitivity is desired. Defocus blur is calculated using equations 3.4 and 6.2 also as a function of focussing distance for fixed $Z_A = 500$ mm; blur is normalised by the pixel size to easily identify when the 16-pixel constraint is crossed. Blur is also expressed as a percentage of nozzle diameter at the spatial resolution at the measurement volume, so that the $\delta/D \leq 11\%$ criterion is met, which must also be considered. Figure 6.8 shows that higher sensitivity can be achieved in the l < 0region (focussing closer than the object), but this also causes the blur to grow much more rapidly than the l > 0 region. Due to this, only the l > 0 region will be considered.

Although the expected speckle size can be calculated using equation 6.4, it is still wise to check if the speckles are small enough to fit within the 16-pixel windows, and large enough to avoid peak-locking (> 2 pixels). Speckle size is not related to the surface characteristics and depends only on the lens aperture and camera pixel size. Figure 6.9 compares the speckles for 3 apertures. Using f/8 generates speckles that are too small and using f/16 generates speckles that are too large for the current pixel size and interrogation window dimensions. Hence, the ideal aperture is in between these two, and f/11 is selected as the optimal aperture for this setup.

A comparison of possible choices is shown in table 6.1, where the focus distance is chosen to maximise the sensitivity until the blur reaches $\delta/D = 11\%$. Based on the constraints, the optimal selection is to use focal lengths of f = 25 mm at f/11 with l = 450 mm for $Z_A = 500$ mm. The optimal f = 25 mm and f = 50 mm choices have



Figure 6.8: a) Sensitivity $\Delta/\tan\varepsilon$ (mm), and b) blur at object in pixels d_i/l_{pix} , as a function of focussing distance l (mm) for $Z_A = 500$ mm. Legend for b): solid line – f = 25 mm at f/8; dotted line $\cdots f = 25$ mm at f/16; dash-dot line $\cdots f = 50$ mm at f/16.

very similar sensitivity, and the blur δ/D was chosen to be the same. To choose between the two, consider the field of view. The advantage of the f = 25 mm lens is a larger field of view over the f = 50 mm lens. The illuminated sensor size of the cameras is quite small, so it is desirable to increase the field of view in the jet so that both the near-nozzle laminar region and transition to turbulence can be captured simultaneously. For any of the given setups in table 6.1, it would also be possible to increase sensitivity by accepting an increase in defocus blurring by moving the focus distance further away.

Note that the diffraction-limited spot size (equation 3.6) and total blur (equation 3.7) were not considered. For the given range of apertures, the diffraction-limited spot size is expected to be much smaller than the defocus blur. For example, option 6 (f = 50 mm, f/16, l = 95 mm) in table 6.1 has the largest diffraction-limited spot size, $d_d = 6.6$ pixels. Option 1 (f = 25 mm, f/8, l = 260 mm) has the smallest diffraction-limited spot size $d_d = 3.1$ pixels. They have total blur sizes of $d_{\Sigma} = 17.1$ and 16.3 pixels, respectively, given that $d_i = 16$ pixels, so defocus blur is far more important than the diffraction-limited spot size.

Note that the selection methodology for traditional BOS with a printed background will be slightly different, because the speckle size is no longer a consideration. The aperture only affects illumination and blur. It may be worthwhile considering the total blur d_{Σ} rather than only the defocus blur d_i , because diffraction-limited blur d_d can become around the same size as, or larger than, the defocus blurring at small apertures. A multi-objective optimisation method may then be necessary to find the optimal balance between sensitivity, defocus blur and diffraction-limited blur.

The final experimental parameters are given in table 6.2. The spatial resolution at the object and focus planes is determined by placing a ruler at the respective plane, but they can also be determined using geometric optics if the camera's pixel size l_{pix} is known. The spatial resolution at the object can be estimated by $(Z_A l_{pix})/Z_I$, and likewise for the background resolution using $(Z_B l_{pix})/Z_I$. Although the cameras have a full sensor size of 2048 × 1536 pixels, only a 640 × 600 pixel section is used at present. This is due to difficulties in expanding the laser with the available optics. This smaller sensor



Figure 6.9: Recorded laser speckles images in a 100×100 -pixel area for apertures: a) f/8, b) f/11 and c) f/16. Red squares show a 16×16 -pixel area. Physical pixel size is $3.45 \times 3.45 \,\mu\text{m/pixel}$. Brightness and contrast of images b) and c) have been enhanced for clarity. Bottom: speckle patterns recorded by each camera at f/11 superimposed on the experimental setup (illustration).

Table 6.1: Combinations of camera lens focal length, aperture and focussing distance (independent variables) for laser speckle BOS evaluated by sensitivity to displacements, blur at the object and average speckle size (dependent variables, bold). For all options $Z_A = 500$ mm; camera pixel size is $l_{pix} = 3.45 \,\mu\text{m}$. Focus distances are chosen such that blur $\delta/D \approx 11\%$ for all options. The closest cross-correlation window size for focal length f = 25 mm is 16 pixels, while for f = 50 mm it is 32 pixels. Text colour indicates: red is unfavourable (speckles too small/large for current pixel size); orange is sub-optimal (field of view too small); black is selected.

Independent variables			Dependent variables			
Focal length f (mm)	Aperture f/N	Focus distance $l \ (mm)$	$\begin{array}{c} \mathbf{Blur} \\ d_i/l_{pix} \\ (\mathrm{px}) \end{array}$	$\begin{array}{c} \mathbf{Blur} \\ \delta/D \\ (\%) \end{array}$	$\begin{array}{c} \mathbf{Sensitivity} \\ \Delta/\tan\varepsilon \\ (\mathrm{mm}) \end{array}$	Speckle size d_s/l_{pix} (px)
25	8	260	16.0	10.7	8.8	1.57
25	11	450	16.0	10.7	12.2	2.16
25	16	1145	16.0	10.7	17.7	3.14
50	8	100	32.0	10.4	9.1	1.57
50	11	148	32.0	10.4	12.4	2.16
50	16	249	32.0	10.4	17.8	3.14

Parameter	Value
Number of cameras	15
Nominal angular spacing of cameras	12°
Nominal Z_A	$515 \mathrm{mm}$
Nominal l	490 mm
Focal length and aperture	25 mm at f/11
Pixel size, active sensor size	$3.45 \times 3.45 \mu\mathrm{m}, 640 \times 600 \mathrm{pixels}$
Spatial resolution at measurement volume (Z_A)	$72\mu\mathrm{m/pixel}$
Spatial resolution at focus plane $(Z_A + l)$	$140\mu\mathrm{m/pixel}$
Effective exposure time	10 ns at 5 Hz
Nozzle diameter D	10 mm
Field of view at object plane	$4.6D \times 4.3D$

Table 6.2: Experimental setup parameters

size means that a larger field of view of the f = 25 mm focal length lens is even more desirable, allowing the domain from the nozzle outlet up to the turbulence transition to be captured.

Image displacements are calculated using an in-house digital cross-correlation PIV analysis code. The in-house code is well-known and has been developed, validated and applied in many studies over nearly three decades, e.g. [10, 126, 128]. Square 16-pixel interrogation windows are used with a grid spacing of 8 pixels in each direction (50% window overlap), employing Hart's correlation-based correction [51], and median value validation of 2 pixels [150].

The selection methodology for laser speckle TBOS measurements of heated jets can be summarised as:

- 1. Determine the maximum allowable blur in relation to the jet's characteristic length scale, i.e. δ/D . The nozzle diameter (characteristic length) D is often fixed. For the 11% limit established in the previous chapter, the maximum allowable blur will be $\delta = 0.11D$. For the current setup, $\delta = 0.11 \times 10 = 1.1$ mm.
- 2. Select a range of available focal lengths and aperture combinations. For the current setup, f = 25 mm and f = 50 mm focal lengths are available. Longer focal lengths allow higher magnifications but increase the blur, while the field of view may be severely restricted depending on the camera sensor size. The minimum focussing distance also increases, meaning that larger apertures cannot be used as the lens cannot focus closely enough to negate the compounded increase in blur. Apertures of f/8, f/11 and f/16 are considered. The largest aperture f/8 is chosen based on testing which shows that the speckles become very small beyond this, and so are unsuitable for cross-correlation. The smallest aperture f/16 is chosen based on testing which shows that the speckles become dim for f = 25 mm (as the amount of light entering the aperture is restricted) and very large. Recall that equation 6.4 can be used to find the approximate speckle size, which can be normalised by l_{pix} .
- 3. Calculate the focussing distance l which sets $\delta = 0.11D$ for each combination of focal length and aperture. This can be done by rearranging equation 6.3 (and restricting

l > 0 as discussed previously):

$$l = \frac{\delta Z_A}{d_a - \delta} \tag{6.9}$$

where $d_a = f/N$, e.g. $d_a = 25 \div 11 = 2.3$ mm for f = 25 mm at f/11. This requires $d_a > \delta$, which is true for typical BOS experiments.

- 4. Calculate the blurring on the image plane d_i using equations 6.2 and 3.4. Normalise by the pixel size l_{pix} .
- 5. Determine the closest cross-correlation interrogation window size to the choices of d_i/l_{pix} . Common sizes are 8, 16, 24, 32, 48, 64, 96, 128 and 256 pixels; a rule of thumb is to choose an integer power of two, but other sizes are possible. For each focal length and aperture combination, the focussing distance l should be adjusted so that blur d_i approximately matches the nearest window size.
- 6. Select the optimal aperture for laser speckles. At this point, it is possible to rule out apertures which produce speckles that are not optimal for the chosen cross-correlation interrogation window size. The speckles may be too small or too large, too dense or too sparse.
- 7. Calculate the relative sensitivity $\Delta/\tan\varepsilon$ for the remaining choices using equation 6.1. For each focal length, only consider the remaining choices with the highest sensitivity.
- 8. Lastly, select between the focal lengths by considering the field of view. The spatial resolution at the object can be estimated by $(Z_A l_{pix})/Z_I$. Together with the sensor size, one can obtain the field of view at the object. The field of view can also be non-dimensionalised by D. If a large field of view is desired, use a smaller focal length. For example, using a sensor of 600 pixels with $Z_A = 500$ mm and D = 10 mm, the field of view at the object for f = 25 mm at f/11 with l = 450 mm is 4.6D. For f = 50 mm at f/11 with l = 148 mm, the field of view at the object is only 2D. Therefore, the f = 25 mm lens is selected.

6.5 Results and discussion

The experimental measurements from one case are presented here. The jet mass flow rate is set to 1.7×10^{-3} kg s⁻¹ with an exit temperature of 90 °C based on thermocouple measurement. Using the calculations described in Appendix C.2, these conditions approximately correspond to an exit velocity of $U_j = 22 \text{ m s}^{-1}$, Reynolds number based on jet diameter of $Re_D = 10,000$ and exit-to-ambient density ratio of $\rho_e/\rho_{\infty} = 0.8$. It was not possible to characterise the jet's exit velocity profile, velocity field and turbulence intensity in the present study. The jet's characteristic time scale is estimated to be $t_c = D/U_e = 450 \text{ µs}$ and the grid convection time is $t_{gp} = 3.3 \text{ µs}$ where the grid size is taken to be the spatial resolution at the object of 72 µm/pixel. One could instead use the 16-pixel interrogation window size as d_{gp} , which instead gives $t_{gp} = 52.4 \text{ µs}$. The smallest of these values is taken in accordance with the guidelines of the previous chapter, the maximum allowable exposure time is 3.3 µs. As the laser pulse is only 10 ns, temporal blurring is negligible. The Reynolds number and density ratio of this case are chosen to



Figure 6.10: a) Horizontal displacements for camera 1, and b) vertical displacements for camera 1 using the median of multiple reference images for one time-step (physical pixel size is $3.45 \times 3.45 \,\mu\text{m/pixel}$, and spatial resolution at the focus plane is $140 \,\mu\text{m/pixel}$). Flow is from bottom to top. Bottom image: measured displacement magnitude (0 to 1.5 pixels) captured by each camera at a given instant superimposed on the experimental setup (illustration).

match the heated jet DNS from the previous chapter, and the experimental BOS measurements will be compared with the DNS. In the DNS the Mach number is Ma = 0.6 whereas the experiment has an estimated Mach number of $Ma \approx 0.06$, which will have the most impact in the jet's near-field.

From each camera, 5,100 flow images are collected. Each instantaneous image is correlated against 30 reference images from that camera and the median is used as the input to the instantaneous reconstruction to compensate for shot-to-shot variations in illumination intensity. An example of the obtained median instantaneous displacement fields based on multiple reference images is shown in figure 6.10. The displacements caused by the jet are clearly visible, and spurious displacements caused by variations in laser speckle intensity are negligible. The displacements fall within the sub-pixel range, with the peak displacement magnitude measured to be approximately 1.5 pixels.

The reconstructed refractive index gradient magnitude at one instant is shown in figure 6.11. The reconstruction domain uses a Cartesian coordinate system, such that the optical axis of camera 1 is aligned with the global z-axis and with x being the jet axis. There are $50 \times 59 \times 59$ voxels in (x, y, z), with a grid spacing of 0.84 mm/voxel in each direction. The reconstruction is performed using the optimised 'best case' ART algorithm denoted as case K in table 4.2. Using 100 iterations and the discussed modifications, no reconstruction artefacts are visible, and the flow features are distinct from the ambient environment. Preliminary tests show that the FBP reconstruction is obtained in less than 1 second per sample, but the reconstruction is unusable due to strong artefacts throughout the domain. The 100 iteration ART reconstruction was completed in approximately 40



Figure 6.11: Reconstructed refractive index gradient magnitude $|\nabla n|$ at one time-step. Longitudinal slices (flow is from bottom to top): a) y/D = 0; b) z/D = 0. Transverse cross-section slices (on the same colour bar as the above longitudinal slices): c) x/D = 0.3; d) x/D = 1.3; e) x/D = 1.8; f) x/D = 2.6; g) x/D = 3.4; h) x/D = 4.3.

minutes on a workstation using a single core, but this large increase in the time taken for reconstruction is necessary for high quality measurements. It is possible to optimise the ART implementation further by utilising parallel processing or GPU computing, but this is beyond the scope of the current work.

A reconstructed density field at one instant is shown in figure 6.12. Additional visualisations are provided in the Appendix C.3. The laminar region, potential core instabilities and transition to turbulence are captured with striking detail. The flow field is qualitatively consistent with the description of potential core evolution by Yule [154], where



Figure 6.12: Reconstructed density field ρ (kg m⁻³) at one time-step. Domain length is 0.3 < x/D < 4.6. Top: longitudinal slice at y/D = 0, flowing from bottom to top. Transverse slices (on same colourbar as the longitudinal slice): a) x/D = 0.3; b) x/D = 1.3; c) x/D = 2.3; d) x/D = 3.2; e) x/D = 4.2. Also see Appendix .

the downstream growth of instability wave in the shear layer results in the formation of vortex rings which merge at the end of the potential core. These subsequently 'pinch' and break off from the potential core, from which point the flow becomes turbulent. The transverse slices clearly show the growth and development of the secondary instabilities described by Liepmann and Gharib [78], where mushroom-like structures are ejected radially outwards. These are secondary instabilities with high streamwise vorticity which engulf ambient fluid between them. This contributes to the spreading of the jet, like the well-studied incompressible round jet [11], with the exception that the potential core in this jet is likely much shorter than an incompressible jet of equivalent Reynolds number [73]. This similarity is expected as the density ratio is not low enough to trigger the spectacular side jets and the associated enhanced mixing [70], as $\rho_e/\rho_{\infty} > 0.6$ presently. Overall, these observations instill confidence that the optimised TBOS experiment is able to faithfully capture the coherent structures related to the jet's transition to turbulence.

The convergence of the mean density and RMS density fluctuations are shown in figure 6.13. Azimuthal averaging with a bin width of $\Delta r/D = 0.084$ is employed for radial profiles, lending a smooth profile. For both the centreline and radial profiles shown, both the mean density and RMS density fluctuations converge quickly. There is negligible variation beyond using 2,000 samples.

6.5.1 Comparison of mean temperature field with thermocouple measurements

The mean temperature field obtained from TBOS is validated against thermocouple measurements of the radial temperature distribution at several axial stations. The Type-K thermocouple wire has a bead diameter of approximately 0.5 mm, cf. TBOS voxel side length 0.84 mm. The wire is sheathed in a rigid insulating ceramic coating, which itself is located within an aluminium sleeve connecting the thermocouple to a spring-loaded linear traverse. The traverse provides radial movement with a resolution of 0.01 mm. The thermocouple bead traverses along the y-axis, with z = 0. Thus the thermocouple measurements correspond to measurements in a radial direction to the jet r/D. Note that there is a slight offset of the thermocouple bead from the y/D = 0 position by 1 mm due to limitations in the reach of the thermocouple and traverse, i.e. thermocouple measurements start from r/D = 0.1. This offset was measured by fitting a machined tip snugly into the nozzle to identify its centreline and measuring the radial distance to the thermocouple bead through caliper measurements. Movement in x of the thermocouple is performed without a traverse, and the axial location of the bead from the nozzle exit plane is measured using calipers.

Thermocouple measurements are performed at four axial locations with a radial increment of $\Delta r/D = 0.05 \ (0.5 \text{ mm})$: x/D = 0.3, x/D = 1.3, x/D = 2.3, x/D = 3.2, x/D = 4.2. These axial locations allow comparisons along the full length of the BOS domain. During measurement, the thermocouple is traversed to a particular radial station and held for 30 s to allow the temperature to stabilise; the temperature at each station is recorded for 30 s at 1 Hz and the average of these measurements is presented. It is observed that the indicated thermocouple temperature could fluctuate in a range ± 0.8 °C depending on location (thermal inertia was not compensated for). Locations near the centreline and outer region fluctuate in a narrower range than the shear layer, as expected.

The comparison of measurements is shown in figure 6.14. The mean TBOS temper-



Figure 6.13: Convergence of BOS mean density statistics. First row shows values at centreline r/D = 0: a) centreline mean density ratio ρ/ρ_{∞} , and b) centreline normalised RMS density fluctuation $\sqrt{\rho'^2}/\rho_{\infty}$. Second row shows radial profiles at x/D = 2.3 with azimuthal averaging using a bin width of $\Delta r/D = 0.084$: c) mean density ratio, and d) normalised RMS density fluctuation. Results are shown for 500 to 5,000 samples in increments of 500 samples. The darker the marker colour, the more samples are used. Pale yellow is 500 samples and black is 5,000 samples.

ature is obtained by averaging across all 5,100 samples, and further by averaging in the azimuthal direction. It is assumed that the ideal gas law with constant pressure will approximate the jet temperature field based on the measured density field, given that the pressure difference and fluctuations in a low-speed jet, where Ma < 0.3, are practically negligible [11]. Across the measured length of the jet, the mean TBOS temperature field matches the thermocouple temperature very well in the outer region and shear layer, with a maximum difference in temperature between TBOS and the thermocouple of approximately 2 °C. TBOS underestimates the jet temperature towards the centreline, which is exaggerated near the nozzle at x/D = 0.3, where the maximum observed difference rises to 7 °C. Like Lang et al. [72], this can be attributed chiefly to the defocus blurring, which reduces the strength of measured density gradients. The limited sensitivity will also increase the noise floor for the measured gradients which also means that smaller gradients near the jet centreline are not correctly measured. Still, the close agreement

between the thermocouple and BOS measurements are encouraging and demonstrate the potential of our optimised TBOS to be a versatile flow diagnostic technique.

6.5.2 Comparison of density field statistics with heated jet DNS

Although the experimental measurement and DNS case are approximately matched in terms of Reynolds number and density, the Mach numbers and boundary conditions are significantly different. As discussed in chapter 2, differences in inflow boundary condition and velocity profile can significantly alter the development of the jet's near-field, with evidence that the far-field characteristics are influenced as well, in contradiction to the long-standing belief in a universal far-field condition [11, 68, 70, 86, 94, 115, 117].

The experimental heated jet emerges from a contoured heated nozzle with free space around the nozzle. However, the DNS jet emerges from an orifice in flat boundary. The boundary is isothermal with no-slip conditions. The flat boundary is expected to increase the entrainment from the jet's surrounds at the nozzle exit, as well as pushing the development of instabilities and the vortex roll-up further downstream [115]. The nozzle exit boundary layer thickness also strongly influences the development of instabilities, the jet spreading and the point of turbulent transition [70, 117]. The nozzle contour itself plays a part in determining the exit boundary layer thickness, and it is worth noting that in the extreme case of a fully developed pipe flow at the nozzle exit, the large scale structures such as vortex roll-up are not observed and the transition to turbulence is more sudden [86]. For the current experiment, it is expected that the nozzle contraction ratio of 13.69 is sufficient to generate a thin exit boundary layer, similar to those used in the DNS hyperbolic tangent exit profile [65]. But the heated nozzle in the experiment may also create convective currents around the laminar core at the jet exit. Furthermore, the experimental setup may contain upstream or external disturbances producing instabilities that affect the development of the jet.

The DNS Mach number of Ma = 0.6 is an order of magnitude higher than the experiment, and well beyond the range for low-speed flows, Ma < 0.3, where compressibility effects are negligible. Increasing the Mach number restricts the jet's near-field spreading [73, 155], but this does not necessarily mean that the point of turbulent transition is closer to the nozzle, e.g. as measured by the increase in the RMS density fluctuation. The two can be independent.

The comparison of the experiment and DNS in figure 6.15 does, indeed, show the significance of the Mach number and inflow conditions on the near-field development of the heated jet. Matching the density ratio and Reynolds number is not sufficient to produce 'identical' flows. However, it is not possible to determine in what proportion these differences are attributable to compressibility effects, differences in boundary conditions, or upstream conditions such as instabilities. As expected, the spreading of the DNS jet is restricted close the nozzle, when x/D < 2, illustrated in the density ratio field ρ/ρ_{∞} . This is also seen in figure 6.16 the normalised half-width $r_{1/2}/D$ based on normalised 'excess' density $(\rho_{\infty} - \overline{\rho})/(\rho_{\infty} - \overline{\rho}_c)$, where subscript $_c$ indicates the centreline value. Although the shear layer in the DNS is thinner than the experimental jet for any x/D in the domain, the RMS density fluctuations are much more intense than the experiment, with the peak normalised RMS density fluctuation $\sqrt{\rho'^2}/\rho_{\infty}$ in the DNS approximately double that of the experiment. This may be at least partially attributable to the spatial filtering in the experimental data. The experiment also shows mild fluctuations outside the jet close to the nozzle, which could be associated with convective currents generated by the hot



Figure 6.14: a) Mean temperature field from TBOS (flow is from left to right) as a function of radial position r/D and axial position x/D, with red dashed lines - - showing locations of thermocouple measurements. Comparison of temperature between BOS \circ and thermocouple \times at: b) x/D = 0.3; c) x/D = 1.3; d) x/D = 2.3; e) x/D = 3.2; f) x/D = 4.2. Error bars represent a 95% confidence level and are approximately the same size as the markers.



Figure 6.15: Comparison of BOS and DNS mean density statistics. First row shows mean density ratio ρ/ρ_{∞} for a) BOS, and b) DNS; second row shows normalised RMS density fluctuation $\sqrt{\overline{\rho'^2}}/\rho_{\infty}$ for c) BOS, and d) DNS, as a function of radial position r/D and axial position x/D. Comparison of BOS \circ and DNS -- e) centreline mean density ratio, and f) centreline normalised RMS density fluctuation, as a function of axial position x/D. Error bars on the centreline mean density ratio indicate a 95% confidence level and are approximately the same size as the markers.



Figure 6.16: Comparison of BOS \circ and DNS - - halfwidth $r_{1/2}/D$ based on local normalised 'excess' centreline density $(\rho_{\infty} - \overline{\rho})/(\rho_{\infty} - \overline{\rho}_c)$, where subscript $_c$ indicates local centreline value.

nozzle. Furthermore, although the centreline density ratio shows good agreement (the effect of the TBOS underestimating the gradients is clearly visible), the centreline RMS density fluctuation shows a sharp increase from x/D > 3, while the experiment continues to show a linear increase in this region. This signifies that the end of the laminar potential core has been reached in the DNS, while the end of the potential core in the experimental jet is not contained in the measured domain.

6.5.3 Insights on the scalar variance transport equation for temperature

With reference to equation 2.4, the temperature equation for non-reacting low-speed flows can be written as,

$$\frac{\partial \left(\rho T\right)}{\partial t} + \nabla \cdot \left(\rho \vec{v} T\right) = \nabla \cdot \left(\rho \kappa \nabla T\right), \qquad (6.10)$$

where \vec{v} is the velocity vector, $\kappa = k/(\rho c_p)$ is the thermal diffusivity, k is the thermal conductivity, c_p is the specific heat at constant pressure [101]. Interestingly, the temperature equation can be used to obtain the velocity field by solving an optimisation problem if the instantaneous scalar field and its temporal evolution are described [23, 28, 134, 135]. For the subsonic heated jet, the density field can be obtained from TBOS, and the temperature field obtained as well through the ideal gas law by assuming a constant pressure field. Though beyond the scope of this work, the temporal evolution of the density fields can also be obtained by using high-speed imaging TBOS or by using both laser cavities in the current setup and staggering the camera exposures, which will facilitate simultaneous 3D velocity, density, and temperature measurements.

As discussed in section 2.1, examination of the scalar variance can provide insights

on the flow's mixing. The temperature variance transport equation is

$$\underbrace{\frac{\partial\left(\overline{\rho}T''^{2}\right)}{\partial t}}_{1} + \underbrace{\nabla\cdot\left(\overline{\rho}\widetilde{v}T''^{2}\right)}_{2} + \underbrace{\nabla\cdot\left(\overline{\rho}\widetilde{v}''T''^{2}\right)}_{3} = \underbrace{\nabla\cdot\overline{\left(\rho\kappa\nabla T''^{2}\right)} + 2\overline{T''\nabla\cdot\left(\rho\kappa\nabla\widetilde{T}\right)}}_{4} - \underbrace{2\overline{\rho}\widetilde{v}''T''\cdot\nabla\widetilde{T}}_{5} - \underbrace{2\overline{\rho\kappa}\left(\nabla T''\cdot\nabla T''\right)}_{6},$$
(6.11)

with the numbering scheme following equation 2.7. The dissipation term $-2\overline{\rho\kappa}(\nabla T'' \cdot \nabla T'')$ can be modelled by introducing the scalar dissipation rate,

$$\chi = \kappa \left(\nabla T'' \cdot \nabla T'' \right). \tag{6.12}$$

The dissipation term can be rewritten as $-2\overline{\rho\chi}$. Using the definition of Favre averaging, $\overline{\rho\chi} = \overline{\rho\chi}$, the dissipation term is simply $-2\overline{\rho\chi}$. The magnitude of χ corresponds to the degree of small-scale turbulent mixing in the flow. It can also be written in terms of Reynolds-averaged quantities,

$$\widetilde{\chi} = \frac{1}{\overline{\rho}} \frac{k}{c_p} \left(\nabla T' - \nabla \left(\frac{\overline{\rho' T'}}{\overline{\rho}} \right) \right) \cdot \left(\nabla T' - \nabla \left(\frac{\overline{\rho' T'}}{\overline{\rho}} \right) \right).$$
(6.13)

Additionally, the dissipation time scale can be defined as

$$\tau = \frac{\widetilde{T''^2}}{2\widetilde{\chi}},\tag{6.14}$$

where the density-weighted scalar variance can be written in terms of Reynolds-averaged quantities as

$$\widetilde{T''^2} = \overline{T'^2} - \frac{\left(\overline{\rho'T'}\right)^2}{\overline{\rho}^2} + \frac{\overline{\rho'T'^2}}{\overline{\rho}}.$$
(6.15)

Figure 6.17 illustrates the measured scalar dissipation rate and mixing time scale for the jet obtained from the TBOS measurements. These provide a deeper insight on the mixing between the hot jet and ambient environment than the mean density field, as they are directly related to the actions of the large-scale turbulent structures [19]. The vortex roll-up ingests and mixes fluid from both the jet and the ambient fluid starting from the upstream instability waves. The peak value of the dissipation rate appears in the shear layer near x/D = 1, which is associated with the vortex roll-up and large temperature gradients. The dissipation time scale τ is much longer than the jet's characteristic time scale t_c , indicating that the mixing due to the large scale occurs much more slowly than the motion of the bulk fluid exiting the nozzle.

6.5.4 Potential core behaviour in the heated jet

The near-field behaviour of the jet is described by the emergence of the bulk fluid from the nozzle, the growth of shear layer instabilities into coherent vortex rings which pair and merge, and their eventual breakup into turbulence. The potential core is the central structure in the near field consisting of unmixed fluid, around which the instabilities and transition occur. As discussed, it is well known that the potential core of a heated jet can


Figure 6.17: Favre-averaged a) scalar dissipation rate $\tilde{\chi}$ normalised by T_{∞}^2/t_c , and b) dissipation time scale τ normalised by t_c , as a function of radial position r/D and axial position x/D.

produce events of significantly increased spreading known as side jets. The conditions under which side jets form has been the subject of investigation [70, 91, 92], but the flow evolution leading to their production is still poorly understood. This phenomenon has been considered from the perspective of the velocity field in forced flows which approximate the same behaviour due to the belief that the hot potential core acts as an instability that can be replicated by external forcing in incompressible jets [17, 142], where side jets can be reproduced with regularity. But there is strong evidence that the density distribution in the near field has a significant role in the formation of side jets due its interactions with the velocity field [15], and it seems less predictable in the heated jet [95].

Laser-speckle TBOS will be a useful tool for studying the near field behaviour of heated jets that lead to side jet formation. Although the current jet does not fall within the $\rho/\rho_{\infty} < 0.61$ criterion for heated jet formation proposed by Kyle and Sreenivasan [70], this study will examine the characteristics of the heated jet observed from TBOS measurements to demonstrate this potential application. There is a lack of data on the potential core of heated jets in any case, and the optimised TBOS technique developed in this project can be used to parametrically investigate the behaviour of the potential core, though a more extensive investigation is beyond the scope of this work.

The unmixed fluid in the potential core can be identified based on an 'excess' velocity criterion, $(u_x - U_\infty)/(U_e - U_\infty) > 0.9$, where u_x is the axial velocity and U_∞ is the freestream velocity (e.g. coflow velocity) [42]. Consider a similar definition to define the density potential core, $(\rho_\infty - \rho)/(\rho_\infty - \rho_e) > 0.9$, with the corresponding contours shown in figure 6.18 for a few instantaneous samples. The location of the end of the potential core, based on the location of 'pinching', varies significantly and tends to bob up and down. Time-resolved data is needed to confirm if there is any periodicity to the location of pinching, but these measurements agree well with previous schlieren visualisations that show that the instantaneous potential core length varies significantly [117]. In some instances, the potential core is stretched to a degree where the end is not captured in the domain.

Pinching tends to break the potential core into smaller fragments. Figure 6.19 shows



Figure 6.18: Contours of normalised 'excess density' $(\rho_{\infty} - \rho)/(\rho_{\infty} - \rho_e)$, coloured by normalised density fluctuation ρ'/ρ_{∞} . Flow is from bottom to top. Solid contours indicates potential core, defined as $(\rho_{\infty} - \rho)/(\rho_{\infty} - \rho_e) = 0.9$. Transparent contours indicate $(\rho_{\infty} - \rho)/(\rho_{\infty} - \rho_e) = 0.25$. Various patterns are observed, labelled as: a) stretching, b) stretching with 1 fragment, c) pinching, d) compressed with 1 fragment, e) compressed with 2 fragments, f) compressed splattering.



Figure 6.19: Histogram illustrating frequency of samples as a function of number of potential core fragments observed using the $(\rho_{\infty} - \rho)/(\rho_{\infty} - \rho_e) > 0.9$ criterion. a) All potential core fragments, b) fragments with a volume $V \ge 0.05\overline{V}_{pc}$.

the frequency of samples as a function of number of the potential core fragments observed. Occasionally, multiple fragments are visible, but it is not possible to determine if this is due to the ejection of multiple fragments in quick succession, or the breakup of one fragment into smaller fragments still. When considering only larger structures, e.g. with volume greater than 5% of the mean potential core volume \overline{V}_{pc} in figure 6.19b, most fields tend to have only one structure, which is the potential core itself. This could either indicate that the majority of pinching and fragments form outside the measurement domain, or that the 'excess' density of the majority of fragments has decreased so that they do not fall within the $(\rho_{\infty} - \rho)/(\rho_{\infty} - \rho_e) > 0.9$ criterion. Figures 6.20 compares the distribution of fragments' volumes and their centres of mass. The greater the number of structures, the lower the volume of the main potential core. The fragments, whether on-axis or off, tend to be much smaller than the main potential core, but the multiple fragments themselves are similar in size. The four structure fields, indicative of the 'splattering' behaviour are very infrequent. The centres of mass of fields with two structures shown in figure 6.21 indicate that the off-axis structures tend to be much smaller than the on-axis structures. In fact, no large structures appear off-axis in the two-, three- or four-structure fields, in which case the two-structure fields show potential core itself and a minute off-axis fragment.

6.6 Summary and conclusions

The short exposure time requirement for high-quality BOS images is a challenging aspect of experiment design. As well as having a short pulse time, the lighting system must be practical for illuminating a 15-camera setup. The use of high-power, pulsed PIV laser was chosen as the solution.

Informed by the findings of Meier and Roesgen [81], the present work takes advantage of the coherent laser illumination to create a laser speckle pattern for the BOS background. This decouples the BOS equations from the physical setup dimensions. This is advantageous for optimising the BOS optical setup, as the measurement sensitivity and defocus blur are functions of the camera lens focal length, aperture and focussing



Figure 6.20: Box-and-whisker plots (first and third rows) and scatter plots (second and fourth rows) of the potential core fragments' volumes (normalised by D^3) and the corresponding centres of mass as a function of position, respectively. Only potential core structures/fragments with $(\rho_{\infty} - \rho)/(\rho_{\infty} - \rho_e) > 0.9$ with a volume $V \ge 0.05 \overline{V}_{pc}$ are considered. a) Fields with 1 structure, b) fields with two structures, c) fields with three structures, d) fields with four structures. Box-and-whisker plots: orange line – is the median, box represents the interquartile range IQR (Q_1 to Q_3), whiskers extend to $1.5 \times$ IQR beyond Q_1 and Q_3 , black circles \circ are outliers, red dashed line - - is the volume of the mean potential core \overline{V}_{pc} . Scatter plots: largest structure (main potential core) \circ , second-largest structure \Box , third-largest structure \triangle , fourth-largest structure \diamond ; red cross \times shows the centre of mass of the mean potential core (translated slightly off-axis for clarity).



Figure 6.21: Centres of mass of the 2-structure fields from figure 6.20 as a function of position (magnified for clarity) coloured by volume V normalised by D^3 . a) Largest structure (main potential core), b) fragment. Red cross × shows the centre of mass of the mean potential core (translated slightly off-axis for clarity).

distance.

This study presents a selection methodology to address the compromise between measurement sensitivity and blur. Such a methodology has not been demonstrated before. The methodology is 'blur-led', in that it focuses on mitigating the impact of defocus blur on the measurement first, informed by the limits that are identified in chapter 5, and then sequentially optimising the speckle size, measurement sensitivity and field-of-view. Such a systematic design of the BOS experimental setup is previously unseen, and will allow BOS measurements to move beyond simple qualitative visualisation or quantitative measurements of large-scale features only. This novel contribution can be applied to future BOS experiments to obtain high-quality measurements of the 3D density field across a range of spatial scales by striking an optimal balance between defocus blurring and measurement sensitivity.

The 15-camera experimental setup with laser-speckle TBOS and the optimised reconstruction algorithm was used to study the density field of a heated jet. Flow features in the near- to intermediate-field were identified with excellent resolution. Measurements were validated against thermocouple measurements, which show good agreement in the shear layer. TBOS under-predicts temperature near the jet centreline in the near-field due to defocus blurring and limited measurement sensitivity. A comparison against DNS provides insights on the effect of Mach number and boundary conditions on the development of the near-field structures and potential core length. This study demonstrates that the optimised laser-speckle TBOS can be used for numerical model validation, and to investigate the potential core evolution in variable density jets.

Chapter 7

Summary and conclusions

Prediction and control of variable density jets requires models which are informed by experimental measurements. There are presently a limited number of techniques which can provide 3D density measurements. This study presents the development of a refraction-based density measurement technique known as tomographic background-oriented schlieren with the goal of obtaining faithful 3D measurements with low temporal and spatial blurring. The work was conducted in three parts to systematically develop each aspect of the TBOS measurement process.

In the first part, a synthetic density field phantom with variable frequency fluctuations was used to test and develop four reconstruction algorithms: FBP, ART, SART and FBP+ART. The key findings from this part are:

- For the typical range of cameras in a TBOS setup, 4-22 cameras, FBP reconstructions contain artefacts outside the jet core which are indistinguishable from flow features. At least 12 cameras are required for a reasonable FBP reconstruction.
- The basic ART algorithms requires several modifications to produce high-quality reconstructions. The modifications are aimed at improving convergence, reducing artefacts, and masking and concentrating iterative corrections towards the location of the jet. A reasonable reconstruction from perfect displacements was obtained with 6-8 cameras and diminishing returns thereafter.
- FBP+ART in practice delivers no significant advantage over the optimised ART.
- SART is slower to converge than ART, but produces the same result.
- If the reconstruction grid size is larger than the defocus blurring, then the maximum spatial frequency that can be resolved is limited by the reconstruction algorithm. In this situation, fluctuations with wavelengths smaller than 4 voxels are underresolved.
- Lower-order discretisations are preferred to solve the Poisson equation via finite differences. The higher spatial frequencies are attenuated by either the reconstruction algorithm or defocus blur, so there is no benefit to using higher-order discretisations, with the added risk of amplifying measurement noise.
- The optimised ART is robust to displacement field noise, for the typical level of random noise expected from the digitial cross-correlation PIV analysis.

The second part of the work focussed on validating a proposed 15-camera experimental setup and testing the reconstruction algorithms on a realistic flow, the density field of a heated jet obtained via DNS. The key findings are:

- In a typical experimental setup, where compromises between measurement sensitivity, defocus blur, adequate illumination and exposure time must be considered, the defocus blur is the dominant source of spatial averaging in the measurement.
- In a heated jet, the defocus blur should be limited to $\delta/D \le 11\%$ to preserve smaller scales in the turbulent transition.
- The temporal blurring should not exceed $0.1t_c$, nor the grid convection time t_{gp} , to prevent degradation of the measurement. Temporal blurring is more detrimental to the measurement than defocus blurring in a typical experimental setup.
- A high-quality reconstruction of the near- to intermediate-field flow structures can be obtained with the proposed 15-camera setup and optimised ART. Reconstruction artefacts are not visible.

The final part of the work implemented an experimental technique for low temporal and spatial blurring measurements. The key findings and contributions are:

- A pulsed laser-speckle TBOS method can be used to overcome temporal blurring. This method provides adequate illumination with temporal integration much shorter than the jet's characteristic time scale. The laser-speckle background pattern decouples the measurement sensitivity and defocus blurring from the physical setup dimensions. This allows the compromise between measurement sensitivity and defocus blurring to be addressed by only adjusting the lens focus distance and aperture, which is much more convenient than adjusting the background distance as in traditional BOS.
- Development of a novel method of optimising the optical setup. This strikes an ideal compromise between measurement sensitivity and defocus blurring. This method is based on limiting blur to the $\delta/D \leq 11\%$ threshold, and subsequently determining the appropriate speckle size and highest possible measurement sensitivity.
- Establishment of a 15-camera facility to conduct laser-speckle TBOS measurements on jets. 3D measurements show excellent resolution of near- to intermediate-field flow structures.
- Laser-speckle TBOS measurements compare well with thermocouple temperature measurements in the shear layer.
- Comparison with DNS shows the influence of Mach number and boundary conditions on the near-field development of the jet.
- Laser-speckle TBOS can be used for closure model measurements.
- Measurements of the potential core show that it is subject to stretching and fragmentation.

The numerical and experimental validation procedures inspire confidence in the accuracy of the optimised laser-speckle TBOS. The development phase of the technique can be concluded, at least for the investigation of heated jets. The most significant improvement to the current implementation of the technique would be a reduction in the reconstruction time by implementing parallel processing or GPU computing. A double-pulse PIV laser can be used to obtain temporal information on the density field as well, which could be used to estimate the velocity field through the scalar transport equation. The laser-speckle TBOS method can pivot away from the development phase, and is ready to be applied to the studies of near- to intermediate-field coherent structures, collection of flow statistics, and model validation. The technique can be used in conjunction with velocity measurement techniques such as tomographic PIV or scalar imaging velocimetry [23, 28, 134, 135], to obtain information on the density-velocity correlations in these flows.

The ability to extend these measurements further downstream depends on sensitivity to the density gradients, as they are expected to decrease significantly in the far-field, where the turbulence approaches statistical isotropy. This study presents a method of scaling the experimental setup parameters to measure these regions of the flow, accounting for the increased sensitivity required. Expanding to other flow configurations, such as an annular jet or thermal boundary layer, will require some knowledge of the geometry, to appropriately modify the masking and windowing techniques in the reconstruction algorithm, and for the Poisson equation boundary conditions.

Appendix A

Material related to chapter 4

A.1 Transfer function of the discretised 1D Poisson equation

This discussion is restricted to examining the one-dimensional Poisson equation, as it is a single-input single-output (SISO) system with an easily-derivable single transfer function. Conducting the von Neumann stability analysis, one can see the behaviour of the analytical and finite-difference transfer functions with respect to the fluctuation frequency ω . The one-dimensional Poisson equation is given in equation 4.19. Letting the reconstructed gradient be

$$\left(\frac{\mathrm{d}n}{\mathrm{d}x}\right)_{recon} \equiv m,\tag{A.1}$$

by taking the Laplace transform of the Poisson equation, the analytical system transfer function is found to be

$$\frac{N(s)}{M(s)} = \frac{1}{s} \tag{A.2}$$

where s is the complex frequency variable.

The finite difference (FD) equations for the different discretisation schemes are given in table A.1. To obtain the transfer function (TF) of the finite-difference equations, the Z-transform is used. The transfer functions are also given in table A.1, where z is the complex frequency variable for a discrete signal.

The response of the analytical and discrete transfer functions with respect to the spatial frequency of the signals ω is given the Bode magnitude plot in figure 4.6.

A.2 Sensitivity study of ART inversely iteration-weighted Gaussian filtering standard deviation

It is seen in table 4.2 that using inversely iteration-weighted Gaussian filtering alongside the sharp cut-off mask, Hamming windowed corrections and progressively tightened Gaussian mask can reduce the RMS and peak errors in the reconstructed gradients further (e.g. case J to case K in table 4.2). The moderation provided to the reconstruction by the inversely iteration-weighted Gaussian filtering can be controlled by adjusting the standard deviation of the Gaussian filter σ_{GF} (the filter acts on a 5 × 5 kernel). Figure A.1 presents the change in absolute error as a function of radial position, and the change

Scheme	Equation					
(3, 3)	FD	$\frac{n_{i-1}-2n_i+n_{i+1}}{dx^2} = \frac{-(1/2)m_{i-1}+(1/2)m_{i+1}}{dx}$				
	\mathbf{TF}	$\frac{dx[(1/2)z^2 - (1/2)]}{z^2 - 2z + 1}$				
(3, 5)	FD	$\frac{n_{i-1}-2n_i+n_{i+1}}{dx^2} = \frac{(1/12)m_{i-2}-(2/3)m_{i-1}+(2/3)m_{i+1}-(1/12)m_{i+2}}{(1/12)} + \frac{(1/12)m_{i-2}-(2/3)m_{i-1}+(2/3)m_{i+1}-(1/12)m_{i+2}}{(1/12)}$				
	TF	$\frac{dx[-(1/12)z^{4}+(2/3)z^{3}-(2/3)z+(1/12)]}{z^{3}-2z^{2}+z}$				
(5, 3)	FD	$\frac{-(1/12)n_{i-2} + (4/3)n_{i-1} - (5/2)n_i + (4/3)n_{i+1} - (1/12)n_{i+2}}{dx^2} = \frac{-(1/2)m_{i-1} + (1/2)m_{i+1}}{dx}$				
	TF	$\frac{dx[(1/2)z^3 - (1/2)z]}{-(1/12)z^4 + (4/3)z^3 - (5/2)z^2 + (4/3)z - (1/12)}$				
(5, 5)	FD	$\frac{-(1/12)n_{i-2} + (4/3)n_{i-1} - (5/2)n_i + (4/3)n_{i+1} - (1/12)n_{i+2}}{dx^2}$ = $\frac{(1/12)m_{i-2} - (2/3)m_{i-1} + (2/3)m_{i+1} - (1/12)m_{i+2}}{dx^2}$				
	TF	$\frac{dx[-(1/12)z^4+(2/3)z^3-(2/3)z+(1/12)]}{-(1/12)z^4+(4/3)z^3-(5/2)z^2+(4/3)z-(1/12)}$				

Table A.1: One-dimensional Poisson equation finite difference (FD) schemes and corresponding transfer functions (TF).

in RMS error, as σ_{GF} is varied. This figure indicates that there may be an optimal value of σ_{GF} in case K which minimises both errors. The optimal value may slightly change depending on the wavelength of fluctuations to be measured, but for the tested case as the value of σ_{GF} is increased from the optimal value, the errors will increase too as smaller scales are strongly attenuated (washed out), and then plateau to an error that is smaller than that without filtering. A value of $\sigma_{GF} = 0.5$ voxels $= \sigma_{GF}/\sigma = 0.06$ is chosen for all cases in the investigation.

A.3 Sensitivity study of ART progressively tightened Gaussian mask

The progressively tightened Gaussian mask radius r_m is decreased linearly as ART iterations progress. An oversized mask is used, where $r_{m,initial}$ is a fixed 35 voxels = 3.9σ , and then vary its final value $r_{m,final}$. In figure A.2, the absolute error with respect to radial position and the RMS error are shown as functions of $r_{m,final}$. The errors appear quite sensitive to $r_{m,final}$, which is due to the mask clipping the gradients. If this region is avoided, the mask can be very useful in improving the quality of the reconstruction. For the current reconstruction, a lower error is obtained when $r_{m,final}$ is within 85% of the current value of $r_{m,initial} = 35$ voxels. A value of $r_{m,final} = 30$ voxels = 3.3σ is chosen for all cases in the investigation.



Figure A.1: a) Average absolute error of the ART gradient field reconstruction as a function of normalised radial position r/σ with a bin size of $\sigma/4$. The lightness of the lines (purple to yellow in colour, dark grey to light grey in greyscale) shows σ_{GF} increasing from 0.1 voxels to 2 voxels in increments of 0.1 voxels ($0.011 \leq \sigma_{GF}/\sigma \leq 0.22$ in increments of $0.011\sigma_{GF}/\sigma$) and from 2 voxels to 5 voxels in increments of 1 voxel ($0.22 < \sigma_{GF}/\sigma \leq 0.56$ in increments of $0.11\sigma_{GF}/\sigma$). b) Normalised RMS error between the reconstructed gradients and synthetic field as a function of σ_{GF}/σ . Results are presented for a 16-camera reconstruction and $\lambda_{x,z} = L/14$. ART uses the same modifications as case A (selected ART) in figure 4.13. Vertical dotted line indicates the $2r_{1/2} = 2\sigma\sqrt{2 \ln 2}$ limit for the RMS error.



Figure A.2: a) Average absolute error of the ART gradient field reconstruction as a function of normalised radial position r/σ with a bin size of $\sigma/4$; the lightness of the lines shows $r_{m,final}$ increasing from 0 to 3.5σ in increments of 0.5σ (purple to yellow in colour, dark grey to light grey in greyscale). b) Normalised RMS error between the FBP+ART reconstructed gradients and synthetic field as a function of $r_{m,final}$ normalised by σ . Results are presented for a 16-camera reconstruction and $\lambda_{x,z} = L/14$. ART uses the same modifications as case A (selected FBP+ART) in figure 4.13 with 100 iterations. Initial mask size r_m is fixed at $r_m = 35$ voxels = 3.9σ . Vertical dotted line indicates the $2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ limit for the RMS error.



Figure A.3: a) Average absolute error of the masked FBP gradient field reconstruction as a function of normalised radial position r/σ with a bin size of $\sigma/4$; the lightness of the lines shows $r_{m,initial}$ increasing from 0.5σ to 5.5σ in increments of 0.5σ (purple to yellow in colour, dark grey to light grey in greyscale), and the black dashed line is the FBP solution from figure 4.15. b) Normalised RMS error between the masked FBP reconstructed gradients and synthetic field as a function of $r_{m,initial}$ normalised by σ . Results are presented for a 16-camera reconstruction and $\lambda_{x,z} = L/14$. Vertical dotted line indicates the $2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ limit for the RMS error.

A.4 Sensitivity study of FBP+ART progressively tightened Gaussian mask parameters

The Gaussian mask modification described by equation 4.26 is applied to the FBP reconstructed gradients before the ART iterations begin, i.e. when k = 0. The strength of the mask is dictated by position and the mask's effective radius r_m . As the iterations progress, the size of the mask is decreased, linearly with respect to iteration number, from the initial $r_{m,initial}$ to a final mask size $r_{m,final}$.

First, consider the effect of changing r_m when the Gaussian mask is initially applied to the pure FBP reconstructed gradients. Figure A.3 presents: the radial position absolute error of the masked FBP solution (similar to figure 4.15) for different $r_{m,initial}$; and, the RMS error of the masked FBP solution as a function $r_{m,initial}$. From the absolute error as a function of radius, it can be seen that increasing $r_{m,initial}$ decreases the error somewhat in the region $r > 2r_{1/2}$ as intended, but the error within $r \leq r_{1/2} \approx 1.2\sigma$ increases instead. As the FBP predicts this inner region quite well, the damping predictably worsens the reconstruction here. The FBP reconstruction should be preserved in this region, as this region should help improve the convergence of FBP+ART. As long as $r_{m,initial}$ is not made too small ($0 < r_{m,initial} < 2\sigma$), the masked FBP does not seem to be very sensitive to this parameter. Hence, $r_{m,initial} = 35$ voxels = 3.9σ is chosen.

Now consider the change in error in the FBP+ART reconstructed gradients when $r_{m,initial}$ is a fixed 35 voxels, and its final value $r_{m,final}$ is varied instead. Results are presented in figure A.4, with the FBP+ART reconstruction otherwise using all of the same modifications as the best FBP+ART case in figure 4.13 (case E). The FBP+ART reconstruction appears to be even less sensitive to $r_{m,final}$ than r_m , at least for the value of r_m used here.



Figure A.4: a) Average absolute error of the FBP+ART gradient field reconstruction as a function of normalised radial position r/σ with a bin size of $\sigma/4$; the lightness of the lines shows $r_{m,final}$ increasing from 0 to 3.5σ in increments of 0.5σ (purple to yellow in colour, dark grey to light grey in greyscale). b) Normalised RMS error between the FBP+ART reconstructed gradients and synthetic field as a function of $r_{m,final}$ normalised by σ . Results are presented for a 16-camera reconstruction and $\lambda_{x,z} = L/14$. FBP+ART uses the same modifications as case E (selected FBP+ART) in figure 4.13 with 100 iterations. Initial mask size r_m is fixed at $r_m = 35$ voxels = 3.9σ . Vertical dotted line indicates the $2r_{1/2} = 2\sigma\sqrt{2\ln 2}$ limit for the RMS error.

A.5 Convergence of ART schemes compared to FBP+ART

To aid comparison between ART and FBP+ART, figure A.5 presents the convergence of several ART modification schemes for $\lambda_j = 0.5$, similar to figure 4.13 does with FBP+ART. The case names and markers in both figures correspond to the same modifications. From these two figures, it is seen that FBP+ART really can improve the rate of convergence, as none of the ART cases in figure A.5 possess the 'flat' convergence curve of FBP+ART cases E and F in figure 4.13.



Figure A.5: RMS error (top row) and peak error (bottom row) in the reconstructed refractive index gradients ∇n (left column) and the refractive index fields n (right column) for 16 camera reconstruction and $\lambda_{x,z} = L/14$ in the region twice the half-width $r \leq 2r_{1/2} = 2\sigma\sqrt{2\ln 2}$, as a function of the number of ART iterations for the cases shown in the table below. In all cases the Poisson equation is solved using 3- and 5-point kernels for left-and right-hand side discretisation, respectively. All ART reconstructions are performed using randomly-ordered cameras and pixels with Hamming windowed corrections and relaxation $\lambda_j = 0.5$. ART uses a sharp cut-off mask with $r_{mask} = 30$ voxels (3.3 σ , $2.8r_{1/2}$). Progressively tightened Gaussian mask decreases from $r_m = 35$ voxels (3.9 σ , $3.3r_{1/2}$) to $r_{m,final} = 30$ voxels. ART case A is the same as case K from table 4.2 but with $\lambda_j = 0.5$.

Marker	Case	Type	Inversely iteration-weighted Gaussian filter	Gradual unmasking	Progressively tightened Gaussian mask
×	А	ART	\checkmark		\checkmark
	В	FBP	n/a	n/a	n/a
0	С	ART		\checkmark	
	D	ART	\checkmark	\checkmark	
Δ	Ε	ART			\checkmark
\diamond	\mathbf{F}	ART		\checkmark	\checkmark
∇	G	ART	\checkmark	\checkmark	\checkmark

Appendix B

Material related to chapter 5

B.1 Definition of the blur kernel

Blur is introduced into the measurement through defocusing effects and windowing from displacement calculation methods. The larger of these two is the effective blur. Blur of a volume point from defocusing will have a truncated conical shape due to the diffuse spreading of light rays from the background, while blur due to windowing will have a rectangular shape. Blur from many views will be those shapes, revolved about the point. The blur (whether due to defocussing or cross-correlation windowing) was implemented here using a rectangular blur kernel. The truncated cone shape can also be approximated with the rectangular shape when three factors are considered:

- The shallow expansion angle of the blur cone, due to the large distance between the background and lens aperture.
- All cameras are placed in a plane circumferential to the jet axis, i.e. there is no vertical angle between the cameras and hence, the blur cones.
- Blurring from each view of the flow so that the blur cone essentially becomes revolved around the jet axis.

Figure B.1 can be used to visualise these points. Individual voxels are shown in a transverse slice through the jet (as is used in all analyses here). The point under consideration is marked as an orange cross. The blur cone expands from the background towards the aperture with an angle of 2α . At the point under consideration, the cone has a width of δ .

When the width δ is revolved around that point (axis into the page), it forms a circle. This is approximated with a square-shaped kernel. A circle whose diameter is the same as the side length of a square has an area that is $\pi/4 \approx 0.785$ times that of the square. In order to correct for this discrepancy in area between a true circular kernel and a square kernel with side length δ'/D , the equivalent blur diameter δ/D is given by $\delta/D = (\delta'/D)\sqrt{4/\pi} \approx 1.12(\delta'/D)$ when quantifying the blur in terms of the nozzle diameter, e.g. in the vertical axis of figure 5.3.

To apply blur parametrically to the DNS field, the following method is used. The grid resolution of the DNS field is 0.0204D, the displacement field resolution is 0.068 mm/px and the jet diameter is D = 10 mm. Therefore, the physical size of each grid point is $d_{gp} = 0.0204 \times 10 = 0.204 \text{ mm}$ per grid point (and there are $10 \div 0.204 = 49$ grid points across the nozzle diameter). This means that each grid point is equivalent to



Figure B.1: Illustration of blur from one camera at a point in the reconstruction volume.



Figure B.2: Convergence of the calculated displacements (in pixels) as a function of number of steps taken through the slice, relative to 5,000 steps, for a camera at x/D = 9.28. The slice contains 462 grid points in each direction.

 $d_{gp} = 0.204 \div 0.068 = 3$ displacement field pixels. The blur is implemented over a square kernel where the side length is an odd number of pixels. So, to implement, say, a blur of $\delta' = 15$ pixels, a kernel with a side length of $\delta' \div d_{gp} = 15 \div 3 = 5$ grid points is used. In terms of the nozzle diameter, this is $\delta'/D = 5 \div 49 = 10.2\%$ whereupon the correction of 1.12 is applied so that $\delta/D = 11.4\%$). Even-numbered blurring is taken as the average of the adjacent integer-valued blur kernels.

The test of increasing grid resolution for a constant blur in figure 5.5, e.g. keeping blur at $\delta/D = 6\%$ while increasing the grid resolution from $\delta/5$ to $\delta/15$, was achieved by upscaling the true DNS field by a factor of 3 so that $d_{gp} = 1$ pixel. The blur kernels are then applied on a 1:1 basis between pixels and grid points, e.g. a blur of 15 pixels is applied using a kernel with a side length of 15 grid points. As seen from figure 8, the error changes only marginally when the grid resolution is made finer, but the blur is the same size in terms of nozzle diameter, indicating that the effect of defocus blurring has a greater impact than grid resolution when the blur is larger than the grid resolution.

B.2 Convergence of the ray tracing method

The ray tracing convergence analysis presented in section 4.4.1 is repeated for the DNS field. As an example, shown in figure B.2 is the average variation (residual) in transverse displacement of all rays inside the cut-off mask for a camera at x/D = 9.28 oriented along

the z-axis as the number of steps is increased, compared to the displacement obtained with 5,000 steps through the volume (where the values of refractive index gradient are linearly interpolated for sub-grid values). Other cameras show similar behaviour. The DNS field contains 462 points in each orthogonal direction (x and z) perpendicular to the jet axis.

The maximum variation in the calculated displacement field with increasing iterations is smaller than 10^{-4} pixels. This is considered negligible, as this is approximately 3 orders of magnitude finer than the criterion of 10^{-1} pixels. Based on these findings, the number of steps along the ray is chosen to be at least 1.5 times the number of grid points encountered by the non-refracted ray (in this example, 700 steps).

B.3 Anisotropic diffusion modification to the Poisson solver

The Poisson equation assumes that $\nabla \times \nabla n = 0$, which may not be true for the reconstructed field in which blurring and measurement noise are present. Atcheson et al. [6] noted that this can cause overshoots in the solution of the Poisson equation due to the imperfect reconstruction. To combat this, they introduce an anisotropic diffusion scheme following Agrawal et al. [3] and Weickert [149]. The anisotropic diffusion model used by Atcheson et al. [6] is applied to the right-hand side source term q,

$$q \equiv \nabla \cdot \left(\overline{\overline{D}} \nabla n_{recon}\right),\tag{B.1}$$

where ∇n_{recon} are the reconstructed gradients and \overline{D} is the anisotropic diffusion tensor. This tensor places more weight on information from similar isosurfaces, which can be useful in inhibiting the effects of noisy information. When \overline{D} is the identity tensor, this method naturally becomes the same as the normal Poisson solver. Following Atcheson et al. [6], the anisotropic diffusion tensor is calculated at each point by constructing the structure tensor $\overline{J_{\sigma}}$ at the point,

$$\overline{\overline{J_{\sigma}}} = K_{\sigma} * (\nabla n_{recon} \otimes \nabla n_{recon}^{\mathrm{T}}),$$
(B.2)

where \otimes is the outer product and K_{σ}^* indicates that each component of the structure tensor (e.g. $(\partial n/\partial x)^2$, $(\partial n/\partial x)(\partial n/\partial y)$, etc.) have been individually convolved with a Gaussian blur kernel with a standard deviation of $\sigma = 0.5$ voxels (points outside the domain are given a value of 0). An eigendecomposition is applied to the structure tensor,

$$\overline{\overline{J_{\sigma}}} = \overline{V}\Lambda\overline{V}^{-1}.$$
(B.3)

Then the diffusion tensor at that point is

$$\overline{\overline{D}} = \overline{V}\tilde{\Lambda}\overline{V}^{-1},\tag{B.4}$$

where $\tilde{\Lambda}$ is a modified eigenvalue array,

$$\tilde{\Lambda} = \begin{bmatrix} \alpha & 0 & 0 \\ 0 & \alpha + (1 - \alpha) \exp\left(\frac{-\max|\nabla n_{recon}|}{k(\lambda_0 - \lambda_1)^2}\right) & 0 \\ 0 & 0 & \alpha + (1 - \alpha) \exp\left(\frac{-\max|\nabla n_{recon}|}{k(\lambda_0 - \lambda_2)^2}\right) \end{bmatrix}.$$
 (B.5)



Figure B.3: RMS (left) and peak (right) errors for the ART reconstruction with 40 iterations in the noise-free (top) and noisy cases (bottom) averaged over the 104 samples, as a function of anisotropic diffusion parameter α . Noisy samples are the same as those used in section 5.4.4. Markers correspond to: x/D = 0.28, x/D = 5.28, and x/D = 9.28. Error bars are approximately the size of the markers and indicate a 95% confidence level.

The λ terms, $\lambda_0 \geq \lambda_1 \geq \lambda_2$, are the ordered eigenvalues of $\overline{J_{\sigma}}$, and k and α are tuning constants, and max $|\nabla n_{recon}|$ refers to the maximum of the absolute value of the gradients at that point, i.e. max $(|\partial n/\partial x|, |\partial n/\partial y|, |\partial n/\partial z|)$. Identical to Atcheson et al. [6], $k = 0.5 \times 10^{-5}$ is used. They indicated that using lower values of the tuning constant α provides better noise removal and settled on $\alpha = 0.8$ from their 'ground-truth' CFD data study.

The impact of using anisotropic diffusion for the Poisson solver, for both noise-free (but blurred) and noisy data, is also evaluated. The RMS and peak errors in the 40iteration optimised ART reconstruction are shown in figure B.3 as parameter α is varied from 0.5 to 1.5.

The behaviour of the error appears similar regardless of the presence of significant random noise. The increase in error in all cases is due to the under-prediction of the solved refractive index field increases as α strays from 1.0. Although there may be a small benefit in reducing peak error near the nozzle by using $\alpha > 1$, this study is not able to replicate the findings of Atcheson et al. [6] with the present data and methods to show that using $\alpha < 1$ is beneficial for noise removal.

Appendix C Material related to chapter 6

C.1 Additional experimental setup photographs



Figure C.1: Multiple-camera heated jet rig used for laser speckle TBOS experiment. In normal use, the entire rig is barricaded with black boards, as seen at the back of shot, to prevent stray laser light from injuring users. Not pictured: electronic control system for camera triggering, computer, and heater voltage control are located to the left; laser is located to the right; compressed air input, inline air heater, settling chamber located beneath rig.



Figure C.2: The settling chamber and nozzle are shrouded in thermal insulation (partially unwrapped to show the layers). They are mounted to a height-adjustable platform which is positioned using linear bearings on guide rails (one on each corner of the platform). The platform is moved using a hydraulic jack located below (not pictured). The compressed air line and inline air heater connect to the settling chamber from below the platform. The temperature at the settling chamber inlet is measured using a thermocouple.

C.2 Matching non-dimensional jet parameters to experimental conditions

The round heated jet can be characterised by the exit-to-ambient density ratio ρ_e/ρ_{∞} and the exit Reynolds number based on nozzle diameter D,

$$Re_{D,e} = \frac{\rho_e U_e D}{\mu_e},\tag{C.1}$$

where U_e is the average exit velocity and μ_e is the exit viscosity [25]. However, the experimental setup requires setting the mass flow rate and heater voltage, which both influence the exit properties of the jet. To set these parameters such that the jet matches a given Reynolds number and density ratio, say $\rho_e/\rho_{\infty} = 0.8$ and $Re_{D,e} = 10,000$, the following procedure is used.

- 1. The ambient conditions are assumed to be $T_{\infty} = 20$ °C and P = 101.325 kPa. Measurements show that the lab temperature does not fluctuate more than ± 0.5 °C, while the pressure does fluctuate depending on weather, but these fluctuations should have a negligible impact on the experimental conditions.
- 2. Assuming dry air, the ideal gas law $(R = 0.287 \text{ kJ kg}^{-1} \text{ K}^{-1})$ is used to obtain the ambient density ρ_{∞} . If the target density ratio is $\rho_e/\rho_{\infty} = 0.8$, then the required exit density is $\rho_e = 0.96 \text{ kg m}^{-3}$. Again using the ideal gas law, the target exit temperature is found to be $T_e = 91$ °C. This will be used to set the heater voltage once the required mass flow rate is determined.
- 3. To determine the mass flow rate, the Reynolds number is used to calculate the required exit velocity, $U_e = \frac{Re_{D,e}\mu_e}{\rho_e D}$, where D = 10 mm. Linear interpolation viscosity data from Incropera et al. [60] for ρ_e yields $\mu_e = 214.4 \times 10^{-7}$ Pas. The exit velocity is found to be $U_e = 22.3 \,\mathrm{m \, s^{-1}}$.



Figure C.3: Close up of lens groups showing how they are magnetically attached to the underside of the bench. a) Laser arm outlet, plate beam splitter and lens group 1. b) High-energy mirror and lens group 2.

- 4. Assuming a top-hat velocity profile at the nozzle exit, the mass flow rate is found to be $\dot{m} = \rho_e A_n U_e = 1.7 \times 10^{-3} \text{ kg s}^{-1}$, where $A_n = (\pi/4)D^2$ is the nozzle exit area.
- 5. The exit Mach number is defined as $Ma_e = U_e/c_e$, where $c_e = \sqrt{\gamma RT_e} = 382 \,\mathrm{m \, s^{-1}}$ is the speed of sound at exit conditions with $\gamma = 1.4$. The exit Mach number is found to be $Ma_e = 0.058$, which is well below the DNS Mach number of Ma = 0.6.
- 6. The DNS jet does not consider gravitational effects, therefore buoyancy is not present in the simulation. To be consistent with the simulation, the experimental jet should be momentum-dominated, rather than buoyant, which indicates that the Richardson number in the region of interest (near nozzle region) should be $Ri \ll 1$. The Richardson number based on nozzle diameter at exit conditions is $Ri_D = \frac{g\beta(T_e - T_\infty)D}{U_e^2} = 3.85 \times 10^{-5}$, where $\beta = 1/T_e$ is the expansion coefficient at exit conditions and $g = 9.81 \,\mathrm{m \, s^{-2}}$. Also, the Froude number should be $Fr \gg 1$. The Froude number at the exit is $Fr_D = \frac{U_e}{\sqrt{gD}} = 71$. These indicate that the experimental jet exit is momentum-dominated like the DNS, and gravitational effects are negligible.

Temporal blurring should also be considered. It is insignificant if the effective exposure time of the BOS images t_{exp} is much shorter than the characteristic time scale of the jet $t_c = D/U_e$. For the jet described above, the characteristic time scale is $t_c = 4.5 \times 10^{-4}$ s. The pulse time of the laser illumination is $t_{exp} = 10^{-8}$ s = $2.2 \times 10^{-5} t_c$, so temporal blurring is not a problem.



Figure C.4: Perpendicular cross-sections through the density field at one instant: a) x - y plane at z = 0, b) x - z plane at y = 0.

C.3 Additional experimental density field visualisations

All contour plots use the same colour bar scale and domain as figure 6.12.



Figure C.5: Transparent contour view of density field at one instant.



Figure C.6: Cross-section of density field in x - y plane at z = 0 captured in eight successive instances. The measurement is not time-resolved.

Bibliography

- A. Abdel-Rahman. A review of effects of initial and boundary conditions on turbulent jets. WSEAS Transactions on Fluid Mechanics, 4(5):257-275, 2010.
- G. N. Abramovich. The theory of turbulent jets. MIT Press, Cambridge, 2003. Reprint of 1963 edition.
- [3] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? In A. Leonardis, H. Bischof, and A. Pinz, editors, *Computer Vision - ECCV 2006.*, volume 3951 of *Lecture Notes in Computer Science*, 2006.
- [4] S. Amjad, S. Karami, J. Soria, and C. Atkinson. Assessment of three-dimensional density measurements from tomographic background-oriented schlieren (BOS). *Measurement Science and Technology*, 31(11):114002, 2020.
- [5] A. H. Andersen and A. C. Kak. Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm. *Ultrasonic imaging*, 6 (1):81–94, 1984.
- [6] B. Atcheson, I. Ihrke, W. Heidrich, A. Tevs, D. Bradley, M. Magnor, and H.-P. Seidel. Time-resolved 3D capture of non-stationary gas flows. In ACM transactions on graphics (TOG), volume 27, page 132. ACM, 2008.
- [7] B. Atcheson, W. Heidrich, and I. Ihrke. An evaluation of optical flow algorithms for background-oriented schlieren. *Experiments in Fluids*, 46(3):467–476, 2009.
- [8] C. Atkinson and J. Soria. Algebraic reconstruction techniques for tomographic particle image velocimetry. In 16th Australasian Fluid Mechanics Conference, pages 191–198, 2007.
- [9] C. Atkinson and J. Soria. An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Experiments in Fluids*, 47(4):553–568, 2009.
- [10] C. Atkinson, N. A. Buchmann, O. Amili, and J. Soria. On the appropriate filtering of PIV measurements of turbulent shear flows. *Experiments in Fluids*, 55:1654, 2013.
- [11] C. G. Ball, H. Fellouah, and A. Pollard. The flow field in turbulent round free jets. Progress in Aerospace Sciences, 50:1–26, 2012.
- [12] E. Bar-Ziv, S. Sgulim, O. Kafri, and E. Keren. Temperature mapping in flames by moire deflectometry. *Applied Optics*, 22(5):698–705, 1983.

- [13] S. Beauchemin and J. Barron. The computation of optical flow. ACM Computing Surveys (CSUR), 27(3):433–466, 1995.
- [14] S. Bhattacharya, J. J. Charonko, and P. P. Vlachos. Particle image velocimetry (PIV) uncertainty quantification using moment of correlation (MC) plane. *Measurement Science and Technology*, 29(11):115301, 2018.
- [15] A. Boguslawski, A. Tyliszczak, and K. Wawrzak. Large eddy simulation predictions of absolutely unstable round hot jet. *Physics of Fluids*, 28(2):025108, 2016.
- [16] M. Born and E. Wolf. Principles of optics. Pergamon Press, Oxford, 6th (corrected) edition, 1980.
- [17] P. Brancher, J. M. Chomaz, and P. Huerre. Direct numerical simulations of round jets: vortex inducation and side jets. *Physics of Fluids*, 6(5):1768–1774, 1994.
- [18] J. D. Briers. Holographic, speckle and moiré techniques in optical metrology. Progress in Quantum Electronics, 17(3):167–233, 1993.
- [19] G. L. Brown and A. Roshko. On density effects and large structures in turbulent mixing layers. *Journal of Fluid Mechanics*, 64(4):775–816, 2006.
- [20] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *Proceedings of the 8th ECCV*, pages 25–36, Prague, 2004.
- [21] N. A. Buchmann, C. E. Willert, and J. Soria. Pulsed, high-power LED illumination for tomographic particle image velocimetry. *Experiments in Fluids*, 53(5):1545– 1560, 2012.
- [22] P. Bühlmann. Laser speckle background oriented schlieren imaging for near-wall measurements. PhD thesis, ETH Zürich, 2020.
- [23] S. Cai, Z. Wang, F. Fuest, Y. J. Jeon, C. Gray, and G. E. Karniadakis. Flow over an espresso cup: inferring 3-D velocity and pressure fields from tomographic background oriented Schlieren via physics-informed neural networks. *Journal of Fluid Mechanics*, 915:A102, 2021.
- [24] J. J. Charonko and K. Prestridge. Variable-density mixing in turbulent jets with coflow. *Journal of Fluid Mechanics*, 825:877–921, 2017.
- [25] P. Chassaing, R. Antonia, F. Anselmet, L. Joly, and S. Sarkar. Variable density fluid turbulence, volume 69 of Fluid mechanics and its applications. Springer Netherlands, Dordrecht, 2002.
- [26] S. Cleve, E. Jondeau, P. Blanc-Benon, and G. Comte-Bellot. Cold wire constant voltage anemometry to measure temperature fluctuations and its application in a thermoacoustic system. *Review of Scientific Instruments*, 88(4):044904, 2017.
- [27] S. T. Corrsin and M. S. Uberoi. Experiments on the flow and heat transfer in a heated turbulent jet. Technical Report 1865, NACA Technical Note, 1949.

- [28] W. J. A. Dahm, L. K. Su, and K. B. Southerland. A scalar imaging velocimetry technique for fully resolved four-dimensional vector velocity field measurements in turbulent flows. *Physics of Fluids A: Fluid Dynamics*, 4(10):2191–2206, 1992.
- [29] S. B. Dalziel, G. O. Hughes, and B. R. Sutherland. Whole-field density measurements by 'synthetic schlieren'. *Experiments in Fluids*, 28(4):322–335, 2000.
- [30] C. J. Dasch. One-dimensional tomography: a comparison of Abel, onion-peeling, and filtered backprojection methods. *Applied Optics*, 31(8):1146–1152, 1992.
- [31] E. R. Davies. *Machine vision theory, algorithms, practicalities.* Morgan Kaufmann, San Francisco, 3rd edition, 2005.
- [32] J. W. Demmel. Applied numerical linear algebra. Society for Industrial and Applied Mathematics (SIAM), 1997.
- [33] R. Doleček, P. Psota, V. Lédl, T. Vít, J. Václavík, and V. Kopecký. General temperature field measurement by digital holography. *Applied Optics*, 52(1):A319– A325, 2013.
- [34] T. Elfving, P. C. Hansen, and T. Nikazad. Semi-convergence properties of Kaczmarz's method. *Inverse problems*, 30(5):55007, 2014.
- [35] G. E. Elsinga, B. W. van Oudheusden, F. Scarano, and D. W. Watt. Assessment and application of quantitative schlieren methods: calibrated color schlieren and background oriented schlieren. *Experiments in Fluids*, 36(2):309–325, 2004.
- [36] M. Fedrizzi and J. Soria. Application of a single-board computer as a low-cost pulse generator. *Measurement Science and Technology*, 26:095302, 2015.
- [37] N. Fomin. Speckle photography for fluid mechanics. Springer-Verlag, Berlin, 1998.
- [38] J.-M. Foucaut and M. Stanislas. Some considerations on the accuracy and frequency response of some derivative filters applied to particle image velocimetry vector fields. *Measurement Science and Technology*, 13(7):1058, 2002.
- [39] J. M. Foucaut, B. Miliat, N. Perenne, and M. Stanislas. Characterization of different PIV algorithms using the EUROPIV synthetic image generator and real images from a turbulent boundary layer. In *Particle image velocimetry: recent improvements*, pages 163–185, Berlin, 2004. Springer.
- [40] W. K. George. The self-preservation of turbulent flows and its relation to initial conditions and coherent structures. In W. George and R. Arndt, editors, Advances in Turbulence, pages 1–41. Hemisphere, New York, 1989.
- [41] W. K. George. Asymptotic effect of initial and upstream conditions on turbulence. Journal of Fluids Engineering, 134(6):061203, 2012.
- [42] N. J. Georgiadis and P. J. Papamoschou. Computational investigations of highspeed dual-stream jets. In AIAA/CEAS Aeroacoustics Conference and Exhibit, page 3311, 2003.

- [43] A. B. Gojani and S. Obayashi. Assessment of some experimental and image analysis factors for background-oriented schlieren measurements. *Applied Optics*, 51(31): 7554–7559, 2012.
- [44] A. B. Gojani, B. Kamishi, and S. Obayashi. Measurement sensitivity and resolution for background oriented schlieren during image recording. *Journal of Visualization*, 16(3):201–207, 2013.
- [45] E. Goldhahn and J. Seume. The background oriented schlieren technique: sensitivity, accuracy, resolution and application to a three-dimensional density field. *Experiments in Fluids*, 43(2-3):241–249, 2007.
- [46] J. W. Goodman. Speckle phenomena in optics: theory and applications. SPIE Press, Bellingham, Washington, 2nd edition, 2020.
- [47] S. Grauer and A. Steinberg. Fast and robust volumetric refractive index measurement by unified background-oriented schlieren tomography. *Experiments in Fluids*, 61:80, 2020.
- [48] S. Grauer, A. Unterberger, A. Rittler, K. Daun, A. Kempf, and M. K. Instantaneous 3D flame imaging by background-oriented schlieren tomography. *Combustion and Flame*, 196:284–299, 2018.
- [49] S. J. Grauer. *Bayesian methods for gas-phase tomography*. PhD thesis, University of Waterloo, 2018.
- [50] Z. Guo, Y. Song, Q. Yuan, T. Wulan, and L. Chen. Simultaneous reconstruction of 3D refractive index, temperature, and intensity distribution of combustion flame by double computed tomography technologies based on spatial phase-shifting method. *Optics Communications*, 393:123–130, 2017.
- [51] D. P. Hart. PIV error correction. *Experiments in Fluids*, 29(1):13–22, 2000.
- [52] R. C. Hart, R. J. Balla, and G. C. Herring. Nonresonant referenced laser-induced thermal acoustics thermometry in air. *Applied Optics*, 38(3):577–584, 1999.
- [53] U. Hartmann and J. R. Seume. Combining ART and FBP for improved fidelity of tomographic BOS. *Measurement Science and Technology*, 27(9):097001, 2016.
- [54] G. T. Herman and A. Lent. Iterative reconstruction algorithms. Computers in biology and medicine, 6(4):273–294, 1976.
- [55] D. W. Hewak and J. W. Y. Lit. Numerical ray-tracing methods for gradient index media. *Canadian Journal of Physics*, 63(2):234–239, 1985.
- [56] M. Himpel, B. Buttenschön, and A. Melzer. Three-view stereoscopy in dusty plasmas under microgravity: a calibration and reconstruction approach. *Review of Scientific Instruments*, 82:053706, 2011.
- [57] B. Horn and B. Schunck. Determining optical flow. Artificial Intelligence, 17: 185–203, 1981.

- [58] P. Huerre and P. A. Monkewitz. Local and global instabilities in spatially developing flows. Annual Review of Fluid Mechanics, 22:473–537, 1990.
- [59] I. Ihrke and M. Magnor. Image-based tomographic reconstruction of flames. In ACM SIGGRAPH 2004 Sketches, page 16, 2004.
- [60] F. Incropera, T. Bergman, A. Lavine, and D. Dewitt. Fundamentals of heat and mass transfer. John Wiley, Hoboken, 7th edition, 2011.
- [61] O. Kafri. Noncoherent method for mapping phase objects. *Optics Letters*, 5(12): 555–557, 1980.
- [62] A. C. Kak and M. Slaney. Principles of computerized tomographic imaging. Society for Industrial and Applied Mathematics (SIAM), 2001.
- [63] S. Karami, P. Stegeman, V. Theofilis, P. Schmid, and J. Soria. Linearised dynamics and non-modal instability analysis of an impinging under-expanded supersonic jet. *Journal of Physics: Conference Series*, 1001(1):012019, 2016.
- [64] S. Karami, D. Edgington-Mitchell, and J. Soria. Large eddy simulation of supersonic under-expanded jets impinging on a flat plate. In *Proceedings of the 11th Australasian Heat and Mass Transfer Conference*, page 12, 2018.
- [65] S. Karami, P. C. Stegeman, A. Ooi, and J. Soria. High-order accurate large-eddy simulations of compressible viscous flow in cylindrical coordinates. *Computers & Fluids*, 191:104241, 2019.
- [66] S. Karami, P. Stegeman, A. Ooi, V. Theofilis, and J. Soria. Receptivity characteristics of supersonic under-expanded impinging jets. *Journal of Fluid Mechanics*, 889:A27, 2020.
- [67] E. Keren, E. Bar-Ziv, I. Glatt, and O. Kafri. Measurements of temperature distribution of flames by moire deflectometry. *Applied Optics*, 20(24):4263–4266, 1981.
- [68] J. Kim and H. Choi. Large eddy simulation of a circular jet: effect of inflow conditions on the near field. *Journal of Fluid Mechanics*, 620:383–411, 2009.
- [69] T. Kowalewski, P. Ligrani, A. Dreizler, C. Schulz, and U. Fey. Temperature and heat flux. In C. Tropea, A. L. Yarin, and J. F. Foss, editors, *Springer Handbook of Experimental Fluid Mechanics*, chapter 7, pages 487–561. Springer-Verlag, Berlin, 2007.
- [70] D. M. Kyle and K. R. Sreenivasan. The instability and breakdown of a round variable-density jet. *Journal of Fluid Mechanics*, 249:619–664, 1993.
- [71] F. Lamarche and C. Leroy. Evaluation of the volume of intersection of a sphere with a cylinder by elliptic integrals. *Computer Physics Communications*, 59(2):359–369, 1990.
- [72] H. M. Lang, K. Oberleithner, C. O. Paschereit, and M. Seiber. Measurement of the fluctuating temperature field in a swirling heated jet with BOS tomography. *Experiments in Fluids*, 58(7):58–88, 2017.

- [73] J. C. Lau. Effects of exit Mach number and temperature on mean-flow and turbulence characteristics in round jets. *Journal of Fluid Mechanics*, 105:193–218, 1981.
- [74] J. H. W. Lee and V. H. Chu. *Turbulent jets and plumes: a Lagragian approach*. Kluwer Academic Publishers/Springer, Boston, 2003.
- [75] T.-W. Lee. Optical diagnostics for measurements of species concentrations and temperature. In R. W. Johnson, editor, *Handbook of Fluid Mechanics*, chapter 52, pages 52(1)–52(20). CRC Press, Boca Raton, 2016.
- [76] M. Lehmann. Measurement optimization in speckle interferometry: the influence of the imaging lens aperture. *Optical Engineering*, 36(4):1162–1168, 1997.
- [77] H. Y. Liao. A gradually unmasking method for limited data tomography. In Biomedical Imaging: From Nano to Macro, 2007. ISBI 2007. 4th IEEE International Symposium on, pages 820–823. IEEE, 2007.
- [78] D. Liepmann and M. Gharib. The role of streamwise vorticity in the near-field entrainment of round jets. *Journal of Fluid Mechanics*, 245:643–668, 1992.
- [79] N. Lindlein, A. Bich, M. Eisner, I. Harder, M. Lano, R. Voelkel, K. Weible, and M. Zimmermann. Flexible beam shaping system using fly's eye condenser. *Applied Optics*, 49(12):2382–2390, 2010.
- [80] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 674–679, Vancouver, 1981.
- [81] A. Meier and T. Roesgen. Improved background oriented schlieren imaging using laser speckle illumination. *Experiments in Fluids*, 54:1549, 2013.
- [82] A. H. Meier and T. Roesgen. Improved background oriented schlieren imaging using laser speckle illumination. *Experiments in Fluids*, 54(6):1549, 2013.
- [83] B. Mercier, E. Jondeau, T. Castelain, Y. Ozawa, C. Bailly, and G. Comte-Bellot. High frequency temperature fluctuation measurements by Rayleigh scattering and constant-voltage cold-wire techniques. *Experiments in Fluids*, 60:110, 2019.
- [84] W. Merzkirch. Flow visualization. Academic Press, Orlando, 2nd edition, 1987.
- [85] W. Merzkirch. Density-sensitive whole-field measurement by optical speckle photography. Experimental Thermal and Fluid Science, 10(4):435–443, 1995.
- [86] J. Mi, D. S. Nobes, and G. J. Nathan. Influence of jet exit conditions on the passive scalar field of an axisymmetric free jet. *Journal of Fluid Mechanics*, 432:91–125, 2001.
- [87] Q. Michalski, C. J. Benito Parejo, A. Claverie, J. Sotton, and M. Bellenoue. An application of speckle-based background oriented schlieren for optical calorimetry. *Experimental Thermal and Fluid Science*, 91:470–478, 2018.

- [88] A. F. Mielke. Development of a molecular Rayleigh scattering diagnostic for simulataneous time-resolved measurement of temperature, velocity, and density. PhD thesis, Case Western Reserve University, 2008.
- [89] A. F. Mielke, R. G. Seasholtz, K. A. Elam, and J. Panda. Time-average measurement of velocity, density, temperature, and turbulence velocity fluctuations using rayleigh and mie scattering. *Experiments in Fluids*, 39(2):441–454, 2005.
- [90] A. Mikš and P. Novák. Determination of unit normal vectors of aspherical surfaces given unit directional vectors of incoming and outgoing rays: comment. *Journal of* the Optical Society of America A, 29(7):1356–1357, 2012.
- [91] P. A. Monkewitz, B. Lehmann, B. Barsikow, and D. W. Bechert. The spreading of self-excited hot jets by side jets. *Physics of Fluids A: Fluid Dynamics*, 1(3): 446–448, 1989.
- [92] P. A. Monkewitz, D. W. Bechert, B. Barsikow, and B. Lehmann. Self-excited oscillations and mixing in a heated round jet. *Journal of Fluid Mechanics*, 213: 611–639, 1990.
- [93] T. Morel. Comprehensive design of axisymmetric wind tunnel contractions. Journal of Fluids Engineering, 97(2):225–233, 1975.
- [94] G. J. Nathan, J. Mi, Z. T. Alwahabi, G. J. R. Newbold, and D. S. Nobes. Impacts of a jet's exit flow pattern on mixing and combustion performance. *Progress in Energy and Combustion Science*, 32(5):496–538, 2006.
- [95] J. W. Nichols, P. J. Schmid, and J. J. Riley. Self-sustained oscillations in variabledensity round jets. *Journal of Fluid Mechanics*, 582:341–376, 2007.
- [96] F. Nicolas, V. Todoroff, A. Plyer, G. Le Besnerais, D. Donjat, F. Micheli, F. Champagnat, P. Cornic, and Y. Le Sant. A direct approach for instantaneous 3D density field reconstruction from background-oriented schlieren (BOS) measurements. *Experiments in Fluids*, 57(13), 2016.
- [97] F. Nicolas, D. Donjat, O. Léon, G. Le Besnerais, F. Champagnat, and F. Micheli. 3D reconstruction of a compressible flow by synchronized multi-camera bos. *Experiments in Fluids*, 58(5):1–15, 2017.
- [98] L. N. Olson and J. B. Schroder. PyAMG: algebraic multigrid solvers in python, 2018. release 4.0.
- [99] M. Ota, K. Hamada, H. Kato, and K. Maeno. Computed-tomographic density measurement of supersonic flow field by colored-grid background oriented schlieren (CGBOS) technique. *Measurement Science and Technology*, 22(10):104011, 2011.
- [100] N. Panchapakesan and J. Lumley. Turbulence measurements in axisymmetric jets of air and helium. part 2. helium jet. *Journal of Fluid Mechanics*, 246:225–247, 1993.
- [101] N. Peters. *Turbulent combustion*. Cambridge University Press, Cambridge, 2000.

- [102] Y. P. Petrov and V. S. Sizikov. Well-posed, ill-posed, and intermediate problems with applications. Koninklijke Brill NV, Leiden, 2005.
- [103] S. Pope. *Turbulent flows*. Cambridge University Press, Cambridge, 2000.
- [104] M. Raffel. Background-oriented schlieren (BOS) techniques. Experiments in Fluids, 56:60, 2015.
- [105] M. Raffel, H. Richard, and G. E. A. Meier. On the applicability of background oriented optical tomography for large scale aerodynamic investigations. *Experiments* in Fluids, 28(5):477–481, 2000.
- [106] M. Raffel, C. Willert, F. Scarano, C. Kähler, S. Wereley, and J. Kompenhans. *Particle image velocimetry: a practical guide.* Springer, Berlin, 3rd edition, 2018.
- [107] L. Rajendran, J. Zhang, S. Bane, and P. Vlachos. Uncertainty-based weighted least squares density integration for background-oriented schlieren. *Experiments in Fluids*, 61(11):239, 2020.
- [108] L. K. Rajendran, S. P. M. Bane, and P. P. Vlachos. Dot tracking methodology for background-oriented schlieren (BOS). *Experiments in Fluids*, 60(11):162, 2019.
- [109] L. K. Rajendran, S. P. M. Bane, and P. P. Vlachos. PIV/BOS synthetic image generation in variable density environments for error analysis and experiment design. *Measurement Science and Technology*, 30(8):085302, 2019.
- [110] L. K. Rajendran, S. P. M. Bane, and P. P. Vlachos. Uncertainty amplification due to density/refractive index gradients in background oriented schlieren experiments. *Experiments in Fluids*, 61(6):139, 2020.
- [111] L. K. Rajendran, J. Zhang, S. Bhattacharya, S. P. M. Bane, and P. P. Vlachos. Uncertainty quantification in density estimation from background-oriented schlieren measurements. *Measurement Science and Technology*, 31(5):54002, 2020.
- [112] H. Richard and M. Raffel. Principle and applications of the background oriented schlieren (BOS) method. *Measurement Science and Technology*, 12(9):1576, 2001.
- [113] C. D. Richards and W. M. Pitts. Global density effects on the self-preservation behaviour of turbulent free jets. *Journal of Fluid Mechanics*, 254:417–435, 1993.
- [114] B. S. Rinkevichyus, O. A. Evtikhieva, and I. L. Raskovskaya. Laser refractography. Springer, New York, 2010.
- [115] G. P. Romano. The effect of boundary conditions by the side of the nozzle of a low Reynolds number jet. *Experiments in Fluids*, 33(2):323–333, 2002.
- [116] A. Rowlands. *Physics of digital photography*. IOP Publishing, Bristol, 2nd edition, 2020.
- [117] S. Russ and P. J. Strykowski. Turbulent structure and entrainment in heated jets: the effect of initial conditions. *Physics of Fluids A: Fluid Dynamics*, 5(12):3216– 3225, 1993.

- [118] G. P. Russo. Aerodynamic measurements: from physical principles to turnkey instrumentation. Woodhead Publishing, Cambridge, 2011.
- [119] J. C. Sautet and D. Stepowski. Single-shot laser Mie scattering measurements of the scalar profiles in the near field of turbulent jets with variable densities. *Experiments* in Fluids, 16(6):353–367, 1994.
- [120] J. C. Sautet and D. Stepowski. Dynamic behavior of variable-density, turbulent jets in their near development fields. *Physics of Fluids*, 7(11):2796–2806, 1995.
- [121] U. Schnars and W. Jüptner. Digital holography: digital hologram recording, numerical reconstruction, and related techniques. Springer-Verlag, Berlin, 2005.
- [122] A. Schröder, B. Over, R. Geisler, A. Bulit, R. Schwane, and J. Kompenhans. Measurements of density fields in micro nozzle plumes in vacuum by using an enhanced tomographic background oriented schlieren (BOS) technique. In *Proceedings of the 9th International Symposium on Measurement Technology and Intelligent Instruments (ISMTII-2009)*, pages 1–6, Saint Petersburg, 2009.
- [123] A. Sciacchitano. Uncertainty quantification in particle image velocimetry. Measurement Science and Technology, 30(9):92001, 2019.
- [124] G. S. Settles. Schlieren and shadowgraph techniques. Springer-Verlag, Berlin, 2001.
- [125] A. Sharma, D. V. Kumar, and A. K. Ghatak. Tracing rays through graded-index media: a new method. Applied Optics, 21(6):984–987, 1982.
- [126] M. Shehzad, B. Sun, D. Jovic, Y. Ostavan, C. Cuvier, J.-M. Foucaut, C. Willert, C. Atkinson, and J. Soria. Investigation of large scale motions in zero and adverse pressure gradient turbulent boundary layers using high-spatial-resolution particle image velocimetry. *Experimental Thermal and Fluid Science*, 129:110469, 2021.
- [127] L. Song and D. S. Elson. Effect of signal intensity and camera quantization on laser speckle contrast analysis. *Biomedical Optics Express*, 4(1):89–104, 2012.
- [128] J. Soria. An investigation of the near wake of a circular cylinder using a videobased digital cross-correlation particle image velocimetry technique. *Experimental Thermal and Fluid Science*, 12(2):221–233, 1996.
- [129] J. Soria. Digital particle image velocimetry. In P. Rastogi, editor, *Digital optical measurement techniques and applications*, chapter 9, pages 347–376. Artech House, Boston, 2015.
- [130] J. Soria, W. K. Chiu, and M. P. Norton. A study of unsteady laminar boundary layer flow on a flat plate using a smoke-wire/silhouette flow visualization technique. *Experimental Thermal and Fluid Science*, 3(3):291–304, 1990.
- [131] F. Sourgen, F. Leopold, and D. Klatt. Reconstruction of the density field using the colored background oriented schlieren technique (CBOS). Optics and Lasers in Engineering, 50:29–38, 2012.

- [132] P. Stegeman, J. Pèrez, J. Soria, and V. Theofilis. Inception and evolution of coherent structures in under-expanded supersonic jets. *Journal of Physics: Conference Series*, 708(1):012051, 2016.
- [133] P. Stegeman, J. Soria, and A. Ooi. Interaction of shear layer coherent structures and the stand-off shock of an under-expanded circular impinging jet. In *Fluid-Structure-Sound Interactions and Control*, pages 241–245, 2016.
- [134] L. K. Su and W. J. A. Dahm. Scalar imaging velocimetry measurements of the velocity gradient tensor field in turbulent flows. i. assessment of errors. *Physics of Fluids*, 8(7):1869–1882, 1996.
- [135] L. K. Su and W. J. A. Dahm. Scalar imaging velocimetry measurements of the velocity gradient tensor field in turbulent flows. ii. experimental results. *Physics of Fluids*, 8(7):1883–1906, 1996.
- [136] D. J. Tan, D. Edgington-Mitchell, and D. Honnery. Measurement of density in axisymmetric jets using a novel background-oriented schlieren (BOS) technique. *Experiments in Fluids*, 56(11):204, 2015.
- [137] K. Tanabe. Projection method for solving a singular system of linear equations and its applications. *Numerische Mathematik*, 17(3):203–214, 1971.
- [138] A. N. Tikhonov, A. Goncharsky, V. V. Stepanov, and A. G. Yagola. Numerical methods for the solution of ill-posed problems. Springer, Dordrecht, 1995.
- [139] B. Timmerman and D. W. Watt. Tomographic high-speed digital holographic interferometry. *Measurement Science and Technology*, 6(9):1270–1277, 1995.
- [140] B. Timmerman, D. W. Watt, and P. J. Bryanston-Cross. Quantitative visualization of high-speed 3D turbulent flow structures using holographic interferometric tomography. Optics & Laser Technology, 31(1):53-65, 1999.
- [141] R. Tsai. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of Robotics* and Automation, RA-3(4):323–344, 1987.
- [142] A. Tyliszczak. Parametric study of multi-armed jets. International Journal of Heat and Fluid Flow, 73:82–100, 2018.
- [143] L. Venkatakrishnan and G. Meier. Density measurements using the background oriented schlieren technique. *Experiments in Fluids*, 37(2):237–247, 2004.
- [144] D. Veynante and L. Vervisch. Turbulent combustion modeling. Progress in Energy and Combustion Science, 28(3):193–266, 2002.
- [145] N. A. Vinnichenko, A. V. Uvarov, and Y. Y. Plaksina. Accuracy of background oriented schlieren for different background patterns and means of refraction index reconstruction. In *Proceedings of the 15th International Symposium on Flow Visualization*, Minsk, 2012.

- [146] M. R. Viotti. Digital speckle pattern interferometry. In P. Rastogi, editor, *Digital optical measurement techniques and applications*, chapter 5, pages 167–216. Artech House, Boston, 2015.
- [147] T. E. Walsh. A comparative study of laser speckle photography and laser interferometry for optical tomography. PhD thesis, Texas A&M University, 1996.
- [148] J. Way and P. A. Libby. Hot-wire probes for measuring velocity and concentration in helium-air mixtures. AIAA Journal, 8(5):976–978, 1970.
- [149] J. Weickert. Anisotropic diffusion in image processing. B. G. Teubner, Stuttgart, 2008.
- [150] J. Westerweel and F. Scarano. Universal outlier detection for PIV data. Experiments in Fluids, 39(6):1096–1100, 2005.
- [151] S. Wildeman. Real-time quantitative schlieren imaging by fast Fourier demodulation of a checkered backdrop. *Experiments in Fluids*, 59(6):97, 2018.
- [152] C. E. Willert and M. Gharib. Digital particle image velocimetry. Experiments in Fluids, 10(4):181–193, 1991.
- [153] Y. Xiong, T. Kaufmann, and N. Noiray. Towards robust BOS measurements for axisymmetric flows. *Experiments in Fluids*, 61(8):178, 2020.
- [154] A. J. Yule. Large-scale structure in the mixing layer of a round jet. Journal of Fluid Mechanics, 89(3):413–432, 1978.
- [155] K. B. M. Q. Zaman. Asymptotic spreading rate of initially compressible jets experiment and analysis. *Physics of Fluids*, 10(10):2652–2660, 1998.
- [156] N. A. Zamyatina, O. I. Navoznov, and A. A. Pavel'ev. Experimental determination of the temperature field of axisymmetric jets by the Töpler schlieren method. *Combustion, explosion, and shock waves*, 3(3):279–282, 1967.
- [157] Z. Zhang. A flexible new technique for camera calibration. Technical report, Microsoft Research, 2008.