

Experimental Investigation of Mass Transport, Dynamics, and Stirring in Isolated Thermal Plumes

by

William H. Newsome

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Geological Sciences)
in The University of Michigan
2011

Doctoral Committee:

Associate Professor Aline J. Cotel, Co-Chair
Associate Professor Carolina R. Lithgow-Bertelloni, Co-Chair
Professor Rebecca A. Lange
Professor Larry J. Ruff
Associate Professor Luis P. Bernal
Scientist Emeritus Stanley R. Hart, Woods Hole Oceanographic Institution
Scientist Emeritus John A. Whitehead, Woods Hole Oceanographic Institution

© William H. Newsome 2011

All Rights Reserved

To Susan, Lavonne, Jackie, Clasine, Carolina, Aline, and Christie

ACKNOWLEDGEMENTS

This work would have been very difficult without the kind support of Rhodri Davies (Department of Earth Science and Engineering, Imperial College London) who ran our numerical models. I am also grateful to the craftsmen both at Woods Hole Oceanographic Institution and the University of Michigan for the experimental apparatus their work helped construct. Two gentlemen on my committee, Stan Hart and Jack Whitehead, have been extremely supportive, and their wisdom has been invaluable in shaping my own thinking regarding mantle geodynamics and the evolution of the solid Earth. Actually, the generosity and enthusiasm of all the committee members have been a constant reminder of why our academic institutions are so special.

And then there are my advisers, Carolina and Aline, who have been indispensable assets and mentors. They have shown me new ways of tackling problems that I previously never appreciated. I am extremely grateful for their confidence in my efforts, willingness to let me explore, and constant encouragement.

Go Blue.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vii
LIST OF TABLES	xx
LIST OF APPENDICES	xxii
ABSTRACT	xxiii
CHAPTER	
I. Introduction	1
II. Stereo-PIV and thermochromic liquid crystals for investigation of mixing and entrainment in thermally driven plumes	6
2.1 Introduction	6
2.2 Experimental setup	9
2.3 Calibration	15
2.3.1 Photogrammetric calibration	15
2.3.2 Thermochromic calibration	19
2.4 Data processing	28
2.5 Results	33
2.6 References	37
III. Hybrid PIV/PTV scheme for data recovery	43
3.1 Introduction	43
3.2 The hybrid PIV/PTV scheme	49
3.3 Application of the hybrid scheme	60
3.4 Conclusions	73

3.5	References	73
IV.	Mass origin and transport in thermal plumes	78
4.1	Introduction	78
4.2	Data sources and outline of discussion	81
4.3	Lagrangian coherent structures of mantle plumes	86
4.3.1	Development	87
4.3.2	Thermal plume LCS	95
4.3.3	Relationship between the plume LCS and features observed with other analysis techniques	103
4.3.4	LCS Applications for Analysis of Plume Mass Trans- port	107
4.3.5	Experimental characterization of plume head leading edge LCS	122
4.4	A metric for the morphology of flow-induced deformation	140
4.5	Conclusions	150
4.6	References	154
V.	Time scale analysis for onset of convection in laminar plumes	161
5.1	Introduction	161
5.2	Experimental setup	164
5.3	Experimental results and analysis	165
5.4	Conclusions	174
5.5	References	174
VI.	Reflections and future work	177
APPENDICES	185
A.	Dynamical systems	186
A.1	Introduction	186
A.2	The language of dynamical systems theory	195
A.3	Fixed points	198
A.4	Lyapunov exponents	206
A.5	Stable and unstable manifolds	214
A.6	Aperiodic systems and finite time dynamics	231
A.7	Summary	235
A.8	References	236
B.	SPIVET	239

B.1	Overview	239
B.2	Licensing	240
B.3	Architecture	240
	B.3.1 PivLIB	243
	B.3.2 TlcLIB	246
	B.3.3 SPIVET steps	247
	B.3.4 FloLIB	247
B.4	Coordinate system and array indexing	248
B.5	Installation	250
B.6	Filesystem layout	251
B.7	Dependencies	251
B.8	Avenues for improvement	253
B.9	References	256

LIST OF FIGURES

Figure

2.1	Experimental setup (top-down view). All critical components are ultimately secured to an optical table (not shown).	10
2.2	Camera separation angle.	13
2.3	Calibration target. Lines are spaced 6.35 mm apart.	17
2.4	Setup for photogrammetric calibration. Note that the front and left-side panels of both the inner and outer tanks have been removed for clarity.	18
2.5	Pinhole camera showing an object point and the corresponding image point. All object points along the illustrated ray will map to the same point on the image plane.	19
2.6	Definition of view angle and z -depth for thermochromic calibration. The light sheet is represented by the gray shaded region.	20
2.7	Demonstration of color correction scheme on a dewarped image from a thermal plume experiment. The heater is located at the bottom center of the images. (a) Uncorrected image. (b) Corrected image.	22
2.8	Dewarped thermochromic calibration image for an isothermal fluid temperature of 26.2 °C. The fiber optic line generator is positioned out of view on the right side of the image. Note that the lateral color variation shown is due entirely to view angle. The small yellow-white particles are secondary tracers.	23
2.9	Processed hue field for Figure 2.8. Note that processing has been restricted to a particular region of interest within the full image.	25

2.10	Sample curves from the thermochromic calibration. Curves are obtained by evaluating Eq. (2.1) with z set to the mid-tank position. Note that a hue of 0.0 corresponds to pure red, while a hue of 0.67 represents pure blue.	27
2.11	Thermochromic calibration error versus bath temperature. Central box encloses the 25-75 th percentiles and the median is denoted with a horizontal line. Whiskers mark 99 th percentile confidence intervals, and outliers are represented as +	27
2.12	High-level overview of data reduction process. Arrows represent generic data flow. Circles represent the hue channel, squares the saturation channel, and diamonds the intensity channel.	28
2.13	Geometry for 3C reconstruction showing a single camera at \mathbf{x}_c	34
2.14	Flow variables for an experimental plume analyzed with the SPIV system. The data are shown 1.8 hr after heater activation. (a) Cross section of flow field at mid-tank. Heater is located at the origin. Pseudocolor background is the component of vorticity perpendicular to the section. The solid black line represents the 25.4 °C temperature contour. Arrows are velocity vectors. (b) 3D perspective view showing the vertical origin of a portion of the mass transported by the plume. The heater is located at the bottom of the image and centered beneath the upwelling. The buoyant mass is convecting heat away from the heater, and the 25.4 °C contour is shown here as a translucent purple solid.	35
2.15	Perspective view showing the vertical origin of mass transported by a failed plume. Heater is located at the bottom of the image and centered beneath the upwelling.	36
3.1	Motion-controlled stereoscopic PIV system used to investigate thermal plumes. The working fluid is contained in the inner tank with a well-controlled isothermal boundary condition provided by a water bath in the outer tank. The light sheet is produced using halogen white light sources and a fiber optic line generator.	44

3.2	<p>Comparison of PIV and PTV methods. Initial particle position is depicted by open circles, while final location is denoted by closed circles. (a) Frame 1 divided into 4 windows for processing with PIV. One window contains 5 particles. (b) Frame 2 with the location of the same particles shown after some elapsed time. (c) Best alignment of the tracer pattern in lower left window of Frame 1 with particles in Frame 2 as determined by an image registration algorithm. The window displacement distance is given by $\Delta\mathbf{x}$. (d, e) Same as (a, b), but for processing with PTV. (f) PTV provides the displacement for each individual particle, $\Delta\mathbf{x}_i$. The initial location of the particles has been repeated (dotted circles) for clarity.</p>	45
3.3	<p>Sample image of tracers illuminated using white light. The tracers are Hallcrest NSL33R25C15W100 thermochromic liquid crystals. Note the similarity in appearance of the individual tracers.</p>	46
3.4	<p>Application of PIV to a rotational flow. Format follows Figures 3.2(a-c). (a) Frame 1 with an excessively large, frame-sized PIV window. (b) Frame 2 showing tracers after some time interval. (c) The flow has rotated the initial pattern of tracers to such a degree that poor results are obtained from image registration algorithms that consider only rigid body translations.</p>	47
3.5	<p>Flow chart for the hybrid PIV/PTV scheme. The variable <code>it</code> represents a counter for a loop that terminates after <code>maxits</code> iterations.</p>	50
3.6	<p>A PTV analysis region (dash-dotted line) constructed around a smaller PIV window (dashed line). The maximum expected particle displacement in the time interval between frames is Δx_{\max}.</p>	51
3.7	<p>Diagram of a one-to-one pairing Φ for a linear correspondence problem between two sets of particles A and B. Outliers are shown as solid circles.</p>	53
3.8	<p>Sample shape context (b) for the particle inside the triangle (a). (c) Particle count legend for (b). The shape context bin boundaries are shown as dotted lines in (a). The $\log r$ bin boundaries visible in (a) are also depicted as dotted lines in (b).</p>	54
3.9	<p>Spring model. (a) Particle a from Frame 1 and its four neighbors. (b) Particle b from Frame 2 (solid circle) and a's neighbors in their warp estimated, Frame 2 configuration (grey circles). Superscripts are used to emphasize the temporal configuration of the α_i.</p>	56

3.10	Thin plate spline example. (a) A thin plate pinned at the four corners and displaced by an armature at 5 interior locations. The warped sheet is colored according to out of plane displacement. The armature is shown as a purple surface embedded in the displaced sheet, and also alone in the inset where it has been colored according to displacement. (b) Pseudo-color plot of plate displacement with selected contours shown in white. Purple circles identify the location of the armature points. (c) Superposition of an initially uniform grid (gray lines) that has been vertically deformed (black lines) using the TPS warp of (a,b). Positive out of plane displacements for the TPS have been equated to positive vertical displacement of the image.	59
3.11	Synthetic 100×100 pixel PIV/PTV image sequence containing 30 tracers. (a) Frame 1. (b) Frame 2. (c) Known displacements. Particles in their initial locations corresponding to (a) are shown as open circles. Final particle locations from (b) are identified with solid circles.	61
3.12	RMS error in pixels for four PTV formulations applied to the image sequence of Figure 3.11. (+) Standard PTV scheme using the parameters of Table 3.1. (\times) No blended cost restart. (\square) Blended cost restart used for each iteration. (\circ) Standard PTV setup, but without regularization.	63
3.13	Final correspondences for PTV scheme without regularization. The two mismatched pairs are identified with a dashed box.	63
3.14	RMS error (\circ) in pixels for the standard PTV formulation applied to the image sequence of Figure 3.11, but with 8 outliers added to Frame 2 (inset). The Frame 2 outliers are circled. The RMS error for the no outlier case of Figure 3.12 has been included for comparison and is denoted with (+).	65
3.15	Synthetic image sequence of Figure 3.11 modified by adding four outliers to each frame (a,b).	65
3.16	Results for the standard PTV formulation applied to the image sequence of Figure 3.15. (a) RMS error (\circ). The error trace for the case with no outliers has been included for comparison and is denoted with (+). (b) Final correspondences. The single mismatched pair is enclosed in the dashed box.	66
3.17	PTV analysis regions from an actual dataset. Particles identified by peak extraction algorithm are labeled. (a) Frame 1 region. (b) Frame 2 region.	67

3.18	RMS error for PTV analysis of Figure 3.17 using various outlier fractions.	68
3.19	RMS error for PTV analysis of Figure 3.17. (a) Effect of various restart intervals, with ∞ indicating no restart. $\beta = 0.5$ was used for all 4 cases. (b) Effect of different β values using the restart interval of 5.	70
3.20	Two frame PIV image sequence.	72
3.21	Displacement vectors from (a) traditional PIV analysis and (b) analysis with the hybrid PIV/PTV scheme. The vectors shown originate from the region of the flow field identified by the boxed region in (c).	72
4.1	Experimental setup. The cubic inner tank (inner dimension of 26.5 mm per side) contains the corn syrup, and the outer tank provides the water bath. The heater is located in the center of the inner tank bottom. The volume of fluid analyzed via SPIV is shown in gray.	83
4.2	Stagnation point flow.	89
4.3	Cross-section of velocity, temperature, and FTLE fields $t = 1050$ s following heater activation for case N6 (a-c) and the corresponding numerical model (d-f). The 25.3 °C temperature contour is shown in all images as a black line for reference. The heater is centered at the origin. (a,d) Pseudocolor background is the velocity magnitude. Velocity vectors are shown as arrows. (b,e) Backward-time FTLE field with integration time $\tau = -1050$ s. (c,f) Forward-time FTLE field with integration time $\tau = 910$ s.	96
4.4	Three-dimensional backward-time LCS at the plume leading edge 1360 s following heater activation for case N6. Corresponding cross-section of the backward-time FTLE field (integration time $\tau = -1360$ s) is presented for reference.	99
4.5	Influence of plume leading edge LCS on nearby fluid elements. Passive tracers initially configured as shown in (a) at $t = 740$ s are advected using the velocity field for case N6. The subsequent state of the tracers is presented for integration times of (b) $\tau = 310$ s, (c) $\tau = 620$ s, (d) $\tau = 930$ s, (e) $\tau = 1170$ s. Note that the initial configuration of the tracers is repeated in the inset for clarity. The leading edge LCS is depicted in gray for each snapshot.	100

- 4.6 Strain rate characterization of instantaneous fluid behavior in the vicinity of the plume leading edge LCS from case N6. Fluid elements normal to the LCS are instantaneously attracted ($\dot{\epsilon}_n < 0$) at locations shown in blue and repelled ($\dot{\epsilon}_n > 0$) where red. Zones where fluid elements are predominantly sheared ($\dot{\epsilon}_t > |\dot{\epsilon}_n|$) are depicted in green. The leading edge LCS is shown at the same times as for Figure 4.5. 102
- 4.7 Comparison of different means of locating the leading edge of a plume. All images are generated using the numerical model (Fluidity) of case N6. The material surface grown from the leading edge LCS is shown in purple. Images (a) through (c) are presented at $t = 2000$ s following heater activation, and contours of the depicted scalar field are shown in white or black for reference. Note that the value of the scalar field for images (b) and (c) has been clipped to the ranges shown for clarity. (a) Magnitude of the temperature gradient. (b) Laplacian of temperature. (c) Magnitude $\partial u_r / \partial r$. (d) Simulated release of two dyed fluid parcels which are colored orange and red at $t = 0$ s. The orange dye is released just above but does not wet the heater surface. The final configuration of the dyed fluid at $t = 2000$ s is also shown, but using a different color. The original orange maps to black at $t = 2000$ s, while the red line similarly maps to the blue. Note that the purple material surface in the figure overlies essentially identical sections of the fluid colored black and blue at $t = 2000$ s. The two individual scenarios are repeated in the insets for clarity. 104
- 4.8 Cross-section of velocity, temperature, and vorticity fields for the two numerical models. The 25.3 °C temperature contour is shown in all figures as a solid black line. (a,b,c) Velocity, temperature, and vorticity fields for case R6 at $t = 1050$ s following heater activation. Selected vorticity contours are shown in (c) at ± 0.02 , ± 0.005 , and ± 0.001 s⁻¹. (d,e,f) Velocity, temperature, and vorticity fields for case G6 at $t = 130$ s following start of injection. Selected vorticity contours are shown in (f) at ± 0.75 , ± 0.38 , and ± 0.12 s⁻¹. 110
- 4.9 FTLE fields for the numerical models of starting plumes generated with a localized heat source (a, b) and via injection of hot syrup (c, d). (a,b) Case R6 backward and forward-time FTLE fields, respectively, at $t = 1050$ s following heater activation. (c,d). Case G6 backward and forward-time FTLE fields, respectively, at $t = 130$ s following start of injection. Integration times used to compute the FTLE fields (ref. Eq. 4.12) are (a) $\tau = -1050$ s, (b) $\tau = 910$ s, (c) $\tau = -130$ s, and (d) $\tau = 130$ s. The 25.3 °C contour is shown in all figures as a solid black line. 112

4.10	Schematic of the Lagrangian coherent structures from numerical models (a) R6 and (b) G6. The tank bottom is illustrated as a hatched volume in both figures. The heater is shown cross-hatched in (a), and the injection nozzle is depicted as a gap in the tank bottom in (b). LCS (or segments of LCS) extracted from the forward-time FTLE field are shown in red, while those extracted from the backward-time FTLE field are shown in blue. Note that the outermost forward-time LCS in (b) has been dotted for clarity. The dashed blue line in (b) connects three backward-time LCS extracted from Figure 4.9(c) to form a material line surrounding injected mass.	113
4.11	Shielding effect of the outermost LCS for the injection plume of case G6. The upper series of plots show the evolution of passive tracers launched at $t = 100$ s (a). The subsequent state of the tracers is presented at (b) $t = 250$ and (c) $t = 300$ s. The lower series of images depict the forward-time FTLE field at the same time as the figure immediately above. The solid black lines in all plots represent the 25.7, 31, and 45 °C temperature contours. The outermost LCS is illustrated in the plots as a dotted black line. Note that the FTLE fields were computed using an integration time of $\tau = 100$ s.	115
4.12	Origin of mass within the reservoir bounded between the outermost LCS and the LCS surrounding injected fluid. (a) Configuration of passive tracers from Figure 4.11(a) at $t = 100$ s, repeated for reference. (b) Tracer distribution at $t = 0$ s computed by advecting tracers in (a) backward in time. The white pockets in (b) represent mass that has already been entrained by the head at $t = 100$ s, whereas the green tracers will be temporarily shed to fill the reservoir as shown in (a).	116
4.13	Head vorticity for the numerical model of case G6 shown at (a) $t = 250$ s and (b) $t = 300$ s following the start of injection. The outer black line is the 25.3 °C contour. Vorticity contours are provided at ± 0.4 , ± 0.2 , and ± 0.1 s ⁻¹	118
4.14	Vertical and radial origin of mass contained within (a) the plume generated via a localized heat source in case R6, and (b) the injection plume of case G6. The image of (a) is provided at $t = 3250$ s following heater activation, while the image of (b) corresponds to $t = 300$ s after the start of injection. Mass contours are shown at 15 mm intervals. Radial contours are represented as solid blue lines while the vertical contours are red. The solid black line and shaded region enclose all material that has a temperature of at least 25.7 °C.	119

4.15	Histogram of vertical and radial origin for mass within the plume heads of cases R6 (a, b) and G6 (c, d) at the times shown in Figure 4.14. Note that only non-injected mass for the G6 case is considered.	120
4.16	Cumulative distribution function for vertical and radial origin of mass within the plume heads of cases R6 (a, b) and G6 (c, d) at the times shown in Figure 4.14. Note that only non-injected mass for the G6 case is considered.	120
4.17	Plume head leading edge height above the tank bottom versus time for the six experimental cases. Best fit lines are shown in black for each case. (Δ) N1. (\circ) N2. (\blacktriangle) N3. (\bullet) N4. (\blacklozenge) N5. (\blacktriangledown) N6. (+) Numerical model of case N6.	123
4.18	Position sensitivity analysis setup for case N6. The best estimate for the location of the leading edge material surface at $t = 420$ s following heater activation is shown as the purple line. The dashed black lines represent the best estimate shifted $\pm 10\%$ vertically. The pseudocolor background is the backward-time FTLE field.	125
4.19	Evolution of positional uncertainty for cases of Table 4.4. (a) Time history of $\delta y(\tau)/\delta y(0)$. The best fit lines (solid) for exponential decay of uncertainty are also provided. (b) Uncertainty relative to plume leading edge vertical position y plotted as a function of plume position. (Δ) N1. (\circ) N2. (\blacktriangle) N3. (\bullet) N4. (\blacklozenge) N5. (\blacktriangledown) N6. (+) R6.	126
4.20	Strain rate normal to the leading edge LCS for cases N6 (\blacktriangledown) and R6 (+). The earliest time at which a ridge in the FTLE field can be reliably identified is depicted by a vertical dotted line for each case. These times correspond to $\tau = 0$ in Figure 4.19(a).	126
4.21	Plume head rise velocity as a function of heater power for cases of Table 4.4 (\bullet). The black line is the best fit determined via least squares. The numerical model (case R6) is shown for reference (\blacklozenge), but is not used in the least squares fit.	130
4.22	Plume head rise velocity as a function of adjusted heater power ($Q - Q_0$) for cases of Table 4.4 (\bullet). The black line is the best fit determined via least squares. The numerical model (case R6) is shown for reference (\blacklozenge), but is not used in the least squares fit. Note that due to the adiabatic lower boundary condition, $Q_0 = 0$ for the numerical model.	132

4.23	Comparison of temperature contours for case N1 obtained from the finite volume thermal model (solid lines) and determined experimentally via liquid crystal thermometry (dashed lines). For each liquid crystal derived contour, corresponding finite volume contours are shown at ± 0.05 °C. Note that the isolated, lowermost dashed red line visible in figures (b) and (c) should be ignored as it is an artifact caused by the fluid temperature exceeding the color play of the liquid crystals (ref. Chapter II). Contours are shown at (a) $t = 1900$ s, (b) $t = 5200$ s, and (c) $t = 8500$ s.	133
4.24	Ellipsoidal fit of plume leading edge material surface for numerical model case R6 shown at 250 s intervals starting from $t = 300$ s. The grown material surface is shown in purple, and the best fit ellipsoid determined via least squares is depicted in blue.	135
4.25	Time evolution of the ellipsoidal axes. See text for explanation of colors. (a) Ellipsoid diameter a . (b) Ellipsoid b -axis. Note that the error bars have been removed for clarity (ref. Figure 4.26). (c) Ratio of ellipsoid diameter to b . (Δ) N1. (\circ) N2. (\blacktriangle) N3. (\bullet) N4. (\blacklozenge) N5. (\blacktriangledown) N6. (+) R6.	136
4.26	Evolution of ellipsoid axis uncertainty for cases of Table 4.5. (a, c) Time history of $\delta a(\tau)/\delta a(0)$ and $\delta b(\tau)/\delta b(0)$, respectively. The best fit lines (solid) for exponential decay of uncertainty are also provided. (b, d) Ellipsoid axis uncertainty relative to axis dimension plotted as a function of plume leading edge position. (Δ) N1. (\circ) N2. (\blacktriangle) N3. (\bullet) N4. (\blacklozenge) N5. (\blacktriangledown) N6. (+) R6.	138
4.27	Ellipsoid axes growth rates versus leading edge rise velocity. Values of \dot{a} from Table 4.5 are plotted as (\circ) and values of \dot{b} are shown as (\bullet). The best fit line determined via least squares for $\dot{a} = \dot{a}(v_l)$ is shown dashed while that for $\dot{b} = \dot{b}(v_l)$ is shown solid. The numerical model, case R6, is represented as an open or solid diamond for reference, but was not used in the fits.	140
4.28	Deformation of the unit cube by an arbitrary flow.	143
4.29	Shape metric for cases (a) N6 and (b) R6. Metric is computed using an integration time $\tau = 2000$ s, at which point the plume has risen 200 mm as measured using the leading edge LCS. The images presented, however, show the configuration of mass at $t = 0$. Two contours for $\Psi = 0.5, 1.5$ are shown as solid black lines.	146

4.30	Deformation induced by the flow of case R6 for four macroscopic, initially spherical inclusions. (a) Location of the spherical fluid parcels at $t = 0$. The shape metric pseudocolor plot is provided for reference. (b) The deformed state of the identically colored parcels at $t = 2000$ s.	147
4.31	Orientation of maximum stretch and evolution of material pockets for case R6. (a) Stretch metric with eventual orientation of maximum principal stretch shown in the position the fluid elements have at the time of heater activation ($t = 0$ s). The maximum principal stretch axis orientation is illustrated using black line segments. Fluid enclosed by the the dotted black line has a maximum stretch orientation perpendicular to the cross-section shown (hence the line segments are invisible). Two contours for $\Psi = 0.5, 1.5$ are shown as solid black lines. (b) Blocks of passive tracers at $t = 0$ colored to illustrate various pockets of fluid in (a) where fluid in each pocket deforms similarly. (c) Position of the tracers in (b) at $t = 2000$ s.	148
5.1	Schematic of the time scale t_0 for onset of convection. Thermal boundary layer (solid line) grows by conduction in Phase I then becomes unstable in Phase II. The resulting plume head reaches a steady rise velocity in Phase III.	164
5.2	Head structure for a laminar plume generated with a constant heater power of 1.8 W. The view is a vertical cross section through the plume with the heater centered at the origin. Steady state heater temperature is 80 °C. The pseudocolor background is the backward-time FTLE field with larger values representing higher rates of separation between initially nearby fluid elements. The outer black contour (dashed) represents a temperature of 0.3 °C above ambient. The umbrella-shaped solid black line is the leading edge LCS.	168
5.3	Plume head leading edge height above the tank bottom versus time for the six experimental cases. Best fit lines are shown in black for each case. (Δ) N1. (\circ) N2. (\blacktriangle) N3. (\bullet) N4. (\blacklozenge) N5. (\blacktriangledown) N6.	169
5.4	Ra_c versus viscosity contrast for the data of Table 5.4.	170
5.5	Fluid temperature 5 mm away from heater perimeter as a function of time. Curve labels correspond to case numbers. The time t_{0f} of peak temperature prior to plume liftoff is marked with an X.	171

5.6	Effect of u_0 on the time scale t_0 . Thermal boundary layer (solid line) grows by conduction in Phase I then becomes unstable in Phase II. The resulting plume head reaches a steady velocity in Phase III. Provided development of two plumes, a and b , are identical in Phases I and II, an increase in the head velocity for case b equates to an increase in t_0 . The pivot is identified with an open circle.	173
A.1	Velocity field of the periodically forced Duffing equations at four time values: (a) $t = 0$, (b) $t = \pi/6$, (c) $t = \pi/3$, (d) $t = \pi/2$. For clarity, the velocity vectors are of equal length and colored by magnitude. .	190
A.2	Three scenarios, one for each row of the figure, demonstrating the evolution of two groups of passive tracers (red and blue) under the periodically forced Duffing equations. (a, c, e) Initial distribution of tracers at $t = 0$ for each scenario with black line segments drawn tangent to the instantaneous velocity field for reference. (b, d, f) Corresponding tracer fields after an elapsed time equivalent to sixteen forcing periods ($t = 32\pi/3$).	191
A.3	Evolution of two groups of passive tracers (red and blue) under the periodically forced Duffing equations. (a) Initial distribution of tracers at $t = 0$ with black line segments drawn tangent to the instantaneous velocity field for reference. (b, c, d) Tracer fields after fourteen ($t = 28\pi/3$), fifteen ($t = 10\pi$), and sixteen ($t = 32\pi/3$) forcing periods, respectively.	192
A.4	Velocity field for the periodically forced Hill's spherical vortex at four time values: (a) $t = 0$, (b) $t = 1/4$, (c) $t = 1/2$, (d) $t = 3/4$. Velocity magnitude is represented by the length of the vectors as well as the pseudocolor background.	194
A.5	Evolution of colored passive tracer blocks subjected to the periodically forced Hill's spherical vortex velocity field. (a) Initial tracer distribution at $t = 0$ with line segments (black) proportional in length and tangent to the instantaneous velocity field provided for reference. (b, c) Tracer distribution after three ($t = 3$) and eight periods ($t = 8$) of the forcing, respectively.	196
A.6	Lipschitz continuity in 1D for $f(\theta) = \sin \theta$. Note that for all θ in \mathbb{R} , $ f(\theta) \leq \theta $. Hence $\sin \theta$ is Lipschitz continuous with $K = 1$	200
A.7	The blue ball can be placed almost anywhere on the hill and it will subsequently roll down it. The very apex of the mound, however, is an equilibrium point, and a ball placed there will remain there for all time provided the ball is never perturbed.	201

A.8	<p><i>Phase portraits</i> of several example fixed points in a 2D state space. For each case, the fixed point is identified by the open circle, and several trajectories are illustrated with arrows depicting the direction of trajectory evolution. (a) Stable node. (b) Unstable node. (c) Saddle point. (d) Center. (e) Spiral node.</p>	204
A.9	<p>Phase portrait for the unforced Duffing oscillator. The elliptic points at $\mathbf{x}^* = (\pm 1, 0)$ are indicated by small open circles, while the hyperbolic point at the origin is shown as a solid circle.</p>	205
A.10	<p>Section of Figure A.5(a) with one of the orange/red blocks magnified in the inset. The group of tracers bounded by the ellipse in the inset are initially very close together, but evolve differently as shown in Figure A.5.</p>	206
A.11	<p>In 2D, the matrix $\nabla_0 \mathbf{F}$ transforms the unit circle (dashed line) into an ellipse whose axes align with the unit vectors \mathbf{u}_1 and \mathbf{u}_2. These vectors are the normalized eigenvectors of the matrix \mathbf{N} (see text).</p>	210
A.12	<p>Phase portraits of the steady (a) Duffing oscillator and (b) Hill's spherical vortex. Hyperbolic fixed points are denoted by solid, black circles. Stable manifolds of the hyperbolic points are shown as dashed, red lines, and unstable manifolds are presented as solid blue lines. Several other trajectories are shown for reference as black lines.</p>	216
A.13	<p>Homoclinic (a) and heteroclinic (b) manifold tangles illustrated on a plane that cuts through the actual 3D manifolds. The hyperbolic points, shown as solid black circles, mark the intersection of a closed trajectory (<i>e.g.</i>, a saddle cycle) in 3D state space with the plane. Similarly, the intersections of the (a) homoclinic and (b) heteroclinic trajectories with the plane of section are depicted by open circles. The subscript denotes the order in which the intersections occur relative to p_0.</p>	219
A.14	<p>Segments of stable and unstable manifolds for the forced Duffing system's saddle cycle. A portion of state space spanning $3\pi/2 \leq \theta < 2\pi$ has been removed to show the complicated tangle that has formed. The intersection of the manifolds with the $\theta = 0$ plane is highlighted, and the corresponding saddle cycle intersection is denoted by the green ball. The stable manifold (red) is shown alone in the inset (a). Similarly, the unstable manifold (blue) is shown inset in (b). The combined structures are presented in (c).</p>	221

A.15	Longer segments of the stable (red) and unstable (blue) manifolds shown intersecting the $\theta = 0$ plane. The saddle cycle intersection is represented by the solid black circle near the origin.	222
A.16	Poincaré section for the forced Duffing oscillator at the $\theta = 0$ plane.	224
A.17	Poincaré section of the forced Duffing oscillator shown in 3D state space along with several example trajectories, each assigned a unique color. The intersection of the trajectories with the plane of section ($\theta = 0$) has been highlighted for clarity.	226
A.18	Segments of the stable (red) and unstable (blue) manifolds for the perturbed Hill’s spherical vortex.	228
A.19	Poincaré section for the perturbed Hill’s spherical vortex. The unstable manifold (solid black line) of the leading edge saddle cycle has been included for reference. Trajectories have been colored for easy identification.	228
A.20	Passive tracer blocks from Figure A.5 superimposed on the stable and unstable manifolds (black) of the saddle cycles for the perturbed Hill’s spherical vortex at times (a) $t = 0$ and (b) $t = 8$	230
B.1	Architectural overview of SPIVET. Primary dependencies on third party software are also shown (see Section B.7).	241
B.2	SPIVET source code directory structure.	252

LIST OF TABLES

Table

2.1	Corn syrup fluid properties. The thermal expansion coefficient and temperature dependence of viscosity were measured in the laboratory. Specific heat capacity and thermal conductivity values were provided by the manufacturer. T has units of °C.	10
2.2	RMS calibration errors for the SPIV setup of Figure 2.1.	15
3.1	Hybrid scheme baseline parameters.	61
3.2	Mismatch count for PTV analysis of Figure 3.17 using various outlier fractions.	68
3.3	Mismatch counts corresponding to the analyses of Figure 3.19.	69
4.1	Corn syrup fluid properties. The thermal expansion coefficient and temperature dependence of viscosity were measured in the laboratory. Specific heat capacity and thermal conductivity values were provided by the manufacturer (LSI Specialty Products).	83
4.2	Heater operating conditions. Flux power values represent the steady state power delivered directly to the syrup as measured with a heat flux gauge affixed to the heater. These values are lower than the total heater power due to parasitic heat losses to the acrylic tank.	84
4.3	Common notation.	86
4.4	Experimental results for the six cases of Table 4.2. The numerical model of case N6 is identified as R6. Note that Q is the steady state power delivered to the syrup, and ΔT is the maximum temperature contrast achieved. The Rayleigh number provided is taken as $Ra = \alpha g Q d^2 / \kappa^2 \mu c_p$, where all quantities are evaluated at ambient temperature.	124

4.5	Ellipsoid axis parameters. Heater power Q , temperature contrast ΔT , and characteristic Rayleigh number Ra are repeated from Table 4.4 for reference.	137
4.6	Characteristic dimensionality assessment provided by the metric Ψ	143
5.1	Notation.	162
5.2	Corn syrup fluid properties. The thermal expansion coefficient and temperature dependence of viscosity were measured in the laboratory. Specific heat capacity and thermal conductivity values were provided by the manufacturer (LSI Specialty Products).	165
5.3	Experimental parameters. Flux power values represent the steady state power delivered directly to the syrup as measured with a heat flux gauge. These values are lower than the total heater power due to parasitic heat losses to the acrylic tank.	166
5.4	Experimental results for the six cases of Table 5.3.	168
5.5	Comparison of liftoff time scale t_{0f} to t_0 for cases N3-N6. Time-averaged plume head diameters, \bar{a} , and Rayleigh critical values computed using t_{0f} are also included.	172
B.1	PivLIB package contents.	243
B.2	TlcLIB package contents.	246
B.3	FloLIB package contents.	248

LIST OF APPENDICES

Appendix

A.	Dynamical systems	186
A.1	Introduction	186
A.2	The language of dynamical systems theory	195
A.3	Fixed points	198
A.4	Lyapunov exponents	206
A.5	Stable and unstable manifolds	214
A.6	Aperiodic systems and finite time dynamics	231
A.7	Summary	235
A.8	References	236
B.	SPIVET	239
B.1	Overview	239
B.2	Licensing	240
B.3	Architecture	240
B.3.1	PivLIB	243
B.3.2	TlcLIB	246
B.3.3	SPIVET steps	247
B.3.4	FloLIB	247
B.4	Coordinate system and array indexing	248
B.5	Installation	250
B.6	Filesystem layout	251
B.7	Dependencies	251
B.8	Avenues for improvement	253
B.9	References	256

ABSTRACT

Experimental Investigation of Mass Transport, Dynamics, and Stirring in Isolated Thermal Plumes

by

William H. Newsome

Co-Chairs: Aline J. Cotel and Carolina R. Lithgow-Bertelloni

Significant differences exist between isotopic signatures of typical mid-ocean ridge basalts and those associated with many ocean islands, with ocean island basalts (OIB) generally exhibiting more variability in trace element concentrations and a bias towards enrichment in more primitive isotopes in some cases. Such observations coupled with other geophysical evidence have been used to suggest that OIB's are surface manifestations of upwellings originating in the deep interior near the core-mantle boundary that interact with distinct geochemical reservoirs as material is transported from the Earth's interior to the surface. Although many have studied the chemistry and dynamics of these mantle plumes, fundamental questions remain.

Lagrangian coherent structures and elements of dynamical systems theory are used to extract key material lines and surfaces in isolated laminar plumes. These structures are shown to provide a taxonomic picture of plumes operating in different regimes, to govern how the plume interacts with the ambient during its ascent, and to have a pronounced effect on the origin of mass transported by the plume. A metric

is developed to provide a means of predicting the morphology of mass transported by a general flow, and the rise velocity of the starting plume is used to investigate time scales for liftoff.

All investigations are conducted using a series of experiments and numerical models where laminar, thermal plumes are generated in a high Prandtl number fluid having strongly temperature dependent viscosity. Experimental data are acquired using a custom-built stereoscopic particle image velocimetry with thermochromic liquid crystals to measure the 3D flow and temperature fields within the tank. A hybrid particle image/particle tracking velocimetry scheme is presented which provides improved robustness to particle pattern deformation when using PIV techniques.

In agreement with others, we find starting plumes rise with an essentially constant velocity for a significant portion of their evolution. During the time necessary to traverse the depth of the tank, the strongest experimental case considered has the capacity to reduce heterogeneity length scales by a factor of at least 1000.

CHAPTER I

Introduction

For the most part, extreme pressures, temperatures, and the vast sizes involved have restricted the Earth sciences community to studying the planetary interior using indirect methods. Volcanic lavas, on the other hand, provide one of the rare opportunities for a somewhat direct observation of Earth's inner workings. As portions of Earth's mantle melt, some trace element species and volatiles preferentially migrate into the melt. These 'incompatible' elements along with radiogenic isotopes provide a significant amount of information regarding the composition and evolution of the overall mantle [2, 6, 12]. When melts make it to the planetary surface, the signature provided by the trace and radiogenic elements become frozen in lavas awaiting analysis by geochemists.

Several lines of geophysical reasoning suggest that the planetary mantle is vigorously convecting on geologic timescales [8], with some numerical models showing that portions of the full mantle may have overturned at least 15 times in the past 4.5 Gyr [11]. One might expect a vigorously convecting planetary layer to be relatively well mixed given sufficient time, and at first glance, part of the geochemical record would appear to support such a finding.

The largest volcanic feature on Earth is the 60,000 km network of spreading centers known as mid-ocean ridges. Geochemical analysis of the lavas produced at these mid-

ocean ridges reveals a remarkable degree of uniformity both in space and time of trace element and radiogenic isotope composition across the full extent of the mid-ocean ridge system [4, 1]. So one might argue that the source for these mid-ocean ridge basalts (MORB), to first order at least, would seem reasonably well-mixed. In truth, however, MORB has been found to possess heterogeneities on both large and small spatial scales [12], but typically not to the same degree as another type of oceanic basalts – those produced at ocean islands.

The most famous and currently most active ocean island volcanic feature is the Hawaiian hotspot, which has left a trail of volcanic islands on the Pacific plate that can be traced back at least 70 Myr in the geologic record [10]. Geochemical analysis of these ocean island basalts (OIB) show a much more pronounced degree of chemical heterogeneity among different hotspots and, for some, even within the same hotspot [1, 12]. But perhaps even more intriguing, geochemists have found that OIB from some hotspots are enriched in incompatible elements and certain radiogenic isotopes, like ^3He , that are not produced in any significant quantity in the interior [6, 12, 2]. The mere existence of these OIB heterogeneities suggests that mantle mixing is not complete. Furthermore, OIB enrichment in incompatibles, particularly the noble gas isotope ^3He , implies that OIB sources are potentially sampling reservoirs of ‘primitive’ material that are not only unmixed with the greater mantle, but have been isolated from thorough mixing for at least several billion years [3]. These and other observations insinuate that the sources for OIB and MORB are perhaps different, or at a minimum, a common source samples and entrains material from distinct chemical reservoirs that vary depending on whether the melt is eventually emplaced at ocean islands or mid-ocean ridges.

One mechanism for explaining the geochemical and geophysical signatures of hotspot volcanism, like that of Hawaii, which has gained much interest is the notion of a mantle plume [5, 6, 7, 9]. In concept, these thermo-chemical plumes would

originate in the deep interior, potentially as deep as the core-mantle boundary. Either by way of source material heterogeneity or otherwise interacting with discrete compositional reservoirs, these plumes would transport and potentially mechanically stir mantle material during ascent to the Earth's surface, thereby producing the characteristic observations associated with hotspot volcanism. While plumes have proven to be a convenient and versatile tool for reconciling many observables, questions still remain regarding the fundamental behavior of these dynamical processes.

Such questions can be grouped into two general issues: a) Overall plume structure and dynamical interaction with the surrounding mantle, and b) the degree of entrainment and mixing in plumes of chemically distinct material from the mantle. To provide additional insight toward resolving these questions, we have constructed a special-purpose stereoscopic particle image velocimetry (SPIV) system that is used to experimentally investigate the structure, temperature, velocity, and degree of entrainment and mixing in thermal plumes. Our experimental plumes are generated in corn syrup, a fluid having strongly temperature-dependent viscosity, via a localized heat source. Once the flow variables are acquired for a given experimental case using the SPIV system, the data are processed to compute specific material lines which govern how a plume interacts with the ambient during its ascent.

The chapters of the thesis are intended to be rather self-contained, and each attempts to thoroughly introduce the material discussed. Chapter II provides an overview of the experimental setup, presents the SPIV system, and introduces the reader to stereoscopic particle image velocimetry in general. We also explain the use of thermochromic liquid crystals for non-invasive measurement of temperature, and provide an analysis of measurement errors that affect both velocity and temperature.

The material in Chapter III takes a deeper look at particle image velocimetry. In particular, we consider measurement errors which arise when velocity gradients distort the shape formed by a cluster of tracer particles. To address the issue, we

present a hybrid scheme that transitions from tracking clusters of particles (*i.e.*, the PIV approach) to tracking individual particles when PIV fails.

The tools of the previous chapters and some additional techniques are employed in Chapter IV to investigate the dynamics and mass transport properties of laminar thermal plumes. The plumes analyzed come from a series of six laboratory experiments and two cases generated via a numerical model. The SPIV system is utilized to measure the velocity and temperature fields of the experimental cases, and then all datasets are processed to extract Lagrangian coherent structures (LCS). The LCS approximately locate material lines and surfaces that are shown to have a significant impact on how the plume interacts with its surroundings and entrains mass. We also develop a metric that provides a predictive measure of how a given flow will transform the dimensionality of fluid elements (*i.e.*, whether spherical inclusions become cigar shaped, sheet like, or remain spherical).

Our analysis of the dynamics of thermal plumes is extended in Chapter V through consideration of the time required for a developing thermal boundary layer to become unstable and thereby produce a plume. The development relies exclusively on the six experimental datasets discussed above.

And finally, some concluding remarks and suggestions for future work are provided in Chapter VI. We also include two appendices for reference. The first is a slightly deeper look at some of the foundational elements of dynamical systems theory. The appendix furnishes additional examples of how certain material lines and surfaces can impact mass transfer for fluid flows. The second appendix provides an overview of our open-source SPIV software package SPIVET.

References

- [1] D. N. Barfod, C. J. Ballentine, A. N. Halliday, and J. G. Fitton. Noble gases in the Cameroon line and the He, Ne, and Ar isotopic compositions of high μ

- (HIMU) mantle. *Journal of Geophysical Research*, 104(B12):29509–29527, 1999.
- [2] D. W. Graham. Noble gas isotope geochemistry of mid-ocean ridge and ocean island basalts: Characterization of mantle source reservoirs. *Reviews in Mineralogy and Geochemistry*, 47:247–317, 2002.
- [3] S. R. Hart. Heterogeneous mantle domains: Signatures, genesis and mixing chronologies. *Earth and Planetary Science Letters*, 90(3):273–296, 1988.
- [4] S. R. Hart, E. H. Hauri, L. A. Oschmann, and J. A. Whitehead. Mantle plumes and entrainment: Isotopic evidence. *Science*, 256(5056):517–520, 1992.
- [5] E. H. Hauri, J. A. Whitehead, and S. R. Hart. Fluid dynamic and geochemical aspects of entrainment in mantle plumes. *Journal of Geophysical Research*, 99(B12):24275–24300, 1994.
- [6] A. W. Hofmann. Mantle geochemistry: The message from oceanic volcanism. *Nature*, 385(6613):219–229, 1997.
- [7] A. M. Jellinek and M. Manga. Links between long-lived hot spots, mantle plumes, D^{''}, and plate tectonics. *Reviews of Geophysics*, 42(3), 2004.
- [8] G. Schubert, D. L. Turcotte, and P. Olson. *Mantle Convection in the Earth and Planets*. Cambridge University Press, 2001.
- [9] P. J. Tackley. Mantle convection and plate tectonics: toward an integrated physical and chemical theory. *Science*, 288:2002–2007, 2000.
- [10] J. A. Tarduno, R. A. Duncan, D. W. Scholl, R. D. Cottrell, B. Steinberger, T. Thordarson, B. C. Kerr, C. R. Neal, F. A. Frey, M. Torii, and C. Carvallo. The Emperor Seamounts: Southward motion of the Hawaiian hotspot plume in Earth’s mantle. *Science*, 301(5636):1064–1069, 2003.
- [11] P. van Keken and S. J. Zhong. Mixing in a 3D spherical model of present-day mantle convection. *Earth and Planetary Science Letters*, 171(4):533–547, 1999.
- [12] A. Zindler and S. Hart. Chemical geodynamics. *Annual Review of Earth and Planetary Sciences*, 14:493–571, 1986.

CHAPTER II

Stereo-PIV and thermochromic liquid crystals for investigation of mixing and entrainment in thermally driven plumes

2.1 Introduction

Particle image velocimetry (PIV) provides a powerful tool for the non-invasive observation of fluid flows. There are other equally non-invasive techniques for the measurement of velocity [1, 69] or temperature [3, 4, 17], however PIV permits the user to capture both simultaneously provided temperature-sensitive tracer particles are utilized. This concurrency provides tight coupling between the velocity and temperature fields while reducing the number of different datasets that must be analyzed. PIV systems and associated software have reached a level of maturity that enable the technique to be utilized on a wide variety of flows, and development is still ongoing [14]. As a natural consequence of PIV's evolution, practitioners can now choose from a number of specific approaches that all draw from the same fundamental principles. In the following presentation, our aim is not to provide a comprehensive examination of the different PIV techniques available. Such information can be best obtained from one the many PIV reviews [2, 14, 51, 76, 71]. Instead we present details of what has proven to be a very workable stereoscopic particle image velocimetry (SPIV) system

constructed specifically for the analysis of creeping flows relevant to convection in planetary interiors. Along the way we provide insight into why particular avenues were chosen from those available.

In short, PIV functions by taking images of a fluid laced with passive tracers. As the fluid moves, so do the tracers, and image processing can then be used to track the particles through space and time. Conceptually PIV is no different from a time lapse sequence of images taken of the night sky. On playing the sequence back, the observer can detect motion of the stars.

The term PIV is generally reserved for the class of techniques that functions by tracking clusters of particles as opposed to individual particles – the purview of particle tracking velocimetry [46, 45]. Appealing to the stargazing analogy once more, tracking of particle clusters is somewhat akin to locating the north star Polaris by first identifying the constellation Ursa Minor. Like stars, the individual passive tracers imaged by a PIV system all look quite similar. However small clusters of tracers tend to move together thereby preserving their relative orientation. And that pattern is simply much easier to identify uniquely during image processing (ref. Chapter III).

The minimal PIV system then consists of the working fluid laced with passive tracers, a light source to illuminate those tracers, and a camera to photograph the tracers as they move through time. A single camera projects the 3D world it images onto a 2D sensor or film [58, 65]. Consequently information is destroyed in the process, and if one wishes to reconstruct how the PIV tracers move in 3D space, multiple cameras [74, 51] or other techniques [77, 35, 32] must be used. The desire to measure all three components of the displacement (velocity) field is the very reason stereoscopic PIV systems like the one detailed here have been developed.

Functionally, each camera of the PIV system acquires a pair of images at slightly different times. The time interval between these image frames is chosen to ensure the magnitude of tracer motion is sufficient for reliable detection, but not so large that

the particle pattern begins to deform under the influence of velocity gradients (ref. Chapter III and [24, 34, 56, 64]). During the image processing step, at least one of the images from the pair is divided into discrete windows (*i.e.*, smaller images). The tracers visible in a particular window then form the pattern, and one searches for a match to the pattern in the other frame via some image registration technique [10]. The distance in pixels that the pattern must be shifted such that it lines up with itself in the second frame is then directly related, via the photogrammetric calibration, to the desired velocity of the fluid. In the end, each analysis window provides a single displacement (velocity) vector. Using smaller windows will yield more vectors from a given image, but each window will also have fewer particles with which to hopefully form an unambiguous pattern.

When selecting passive tracers, practically anything that is neutrally buoyant and of appropriate diameter will work [47, 27, 72]. Thermochromic liquid crystals (TLC's) are some of the more interesting passive tracers in that they preferentially scatter light of a particular color based on their temperature [15, 5, 6]. TLC's have several disadvantages. One of the biggest from a particle tracking perspective is that they only scatter visible light for a small range of temperatures. Consequently, TLC's with temperatures above or below the active range are transparent, and an auxiliary means of seeding the flow must be employed in cases where this is an issue. The system we have constructed was specifically designed to utilize TLC's, and we will provide some guidance on their use.

In what follows, we first provide a detailed look at the hardware used in the experimental setup. Both the SPIV system and experimental apparatus will be discussed to illuminate some of the design choices that were made. We then follow with a discussion of how the system is calibrated. Because we extract velocity and temperature data with our system through the use of TLC's, we must perform photogrammetric and thermochromic calibration. Both procedures will be discussed in turn, and then

we provide an overview of how the velocity and temperature fields are constructed from the raw images. We close with results obtained using the system.

2.2 Experimental setup

As mentioned in the Introduction, our SPIV system was designed to study slow flows, specifically laminar plumes generated with a localized heat source. A graphical depiction of the experimental apparatus is shown in Figure 2.1. The working fluid, LSI Specialty Products' Liquidose 436 corn syrup, is contained in a sealed, acrylic tank (hereafter referred to as the inner tank) with cubic inner dimensions of 26.5 cm per side. The side and top walls are 1 cm thick while the tank bottom is 3 cm thick. A small heater, 2 cm in diameter, is located in the center of the inner tank and attached to a programmable closed-loop controller. The heater is installed such that the top surface of the heater is flush with the inner tank bottom, thereby providing a smooth surface for the flow of syrup across the tank bottom. The inner tank is in turn enclosed in an additional acrylic tank. This outer tank, having an outside rectangular footprint of 36.8 cm per side with 40.6 cm tall walls, provides an isothermal water bath that is maintained by a dual-purpose refrigerating and heating recirculator pump. All six outer-tank walls are 1 cm thick. Properties of the working fluid are provided in Table 2.1.

Both the inner tank and the outer tank are constructed with locating features that permit the assembly to be installed on an optical table in a highly repeatable fashion. This is an important design consideration in general but is especially significant for our setup for a couple of reasons. If the refractive environment (*i.e.*, fluids and geometry of the cameras, tanks, and light source) changes, then the PIV system must be recalibrated. However, photogrammetric and thermochromic calibrations are slow and tedious processes in a highly viscous fluid like corn syrup. Once a calibration is available, it is definitely advantageous to try and ensure it remains relevant between

Property	Value or formulation	Units
Density at 25 °C (ρ)	$\rho = 1450 - 0.447T$	kg m^{-3}
Thermal expansion coefficient (α)	3.1×10^{-4}	K^{-1}
Dynamic viscosity (μ)	$\mu = 1080e^{-0.156T+6.25 \times 10^{-4}T^2}$	Pa s
Specific heat capacity (c_p)	2280	$\text{J kg}^{-1} \text{ }^\circ\text{C}^{-1}$
Thermal conductivity (k)	0.34	$\text{W m}^{-1} \text{ }^\circ\text{C}^{-1}$
Refractive index [40] (n)	$n = 1.500 - 2.095 \times 10^{-4}T$	-

Table 2.1: Corn syrup fluid properties. The thermal expansion coefficient and temperature dependence of viscosity were measured in the laboratory. Specific heat capacity and thermal conductivity values were provided by the manufacturer. T has units of $^\circ\text{C}$.




Figure 2.1: Experimental setup (top-down view). All critical components are ultimately secured to an optical table (not shown).

experiments. Furthermore, our system was constructed with the express intent of comparing datasets from different experimental runs to investigate a parameter range representative of varying conditions in terrestrial planetary interiors. By having the experimental test cell located in a fixed position in space, these datasets are readily compared since they all share the same laboratory coordinate system.

Two types of tracer particles are used to seed the corn syrup: micro-encapsulated thermochromic liquid crystals (TLC's) and a high-gloss, white polymer powder used in the coatings industry for a paint-like process known as powder coating. The use of liquid crystals for PIV work is not new, and their utility for non-invasive temperature measurement and as PIV tracers has been well-recognized [15, 13, 50, 28]. As men-

tioned in the Introduction, liquid crystals selectively scatter certain frequencies of light based on the temperature of the liquid crystal and the view angle [23, 54]. The wavelength of scattered radiation will lie in the visible spectrum for some range of liquid crystal temperatures, depending on the TLC formulation, with lower temperatures scattering predominantly in the red region of the spectrum and higher temperatures scattering in the blue. This range of temperatures within which the TLC's scatter visible wavelengths is commonly referred to as the formulation's color-play and can range from a few degrees Celsius to 20°C. The micro-encapsulated TLC's used in the experiments of the present paper were manufactured by LCR Hallcrest¹ to have an approximate diameter of 50 - 100 μm . The TLC's, provided in a water-based slurry at 33% solids by volume, are added to the syrup at 54 ppm of slurry by volume.

Although TLC's provide a convenient means of non-invasively measuring a temperature field while simultaneously serving as PIV tracers, they do have limitations. The significant view angle dependence previously mentioned must be accounted for in the calibration that maps TLC color to fluid temperature. Additionally, the view angles typically used in PIV work reduce the realized color-play to a fraction of the quoted value [30]. The TLC formulation used here, for example, has a quoted color play of 15°C, however the realized color-play for the setup of Figure 2.1 is only 1.6°C. Outside of the active color-play temperature band, the TLC particles are essentially transparent rendering them useless as optically visible flow tracers when operating in such conditions. Hence, secondary tracer particles that always scatter in the visible spectrum regardless of the fluid temperature are also added to the corn syrup so that camera images have sufficient visible particles for PIV processing. The secondary tracers utilized have a density of 1.9 g/cm³, an approximate diameter of 1 - 40 μm , and are added to the syrup at 2 ppm by volume. These gloss-white secondary tracers are supplied as a dry powder and can be obtained from various sources that market

¹Formulation NSL33R25C15W100.

powder coating media.

There are a couple of considerations when choosing the quantity of TLC's or other tracers to add to the working fluid. The seeding density needs to be high enough to ensure that the analysis window of each image contains a sufficient number of particles (ref. Introduction and [76, 72]), but not so large that the fluid is opaque. Fluid opacity becomes increasingly important as the amount of fluid between the camera and the illuminated plane of interest increases.

Our use of secondary tracers creates additional issues. As noted, the perceived color of the liquid crystals is directly related to the temperature field of interest. Any factor, such as reflections from secondary tracers, that affects image color will skew the extracted temperature values and must be compensated for during image processing. Of course, the use of two tracer types also compounds the opacity problem, so care must be exercised in selecting an appropriate seeding density for the secondary tracers.

Once the fluid is seeded and the tanks mounted on the optical table, the flow field of the inner tank is captured using a rotational stereoscopic PIV (SPIV) setup [51]. The angle between the two 3CCD cameras is 18° , and both cameras are equipped with a 15 mm focal length lens. The cameras have a resolution of 1024×768 pixels and are positioned approximately 55 cm from the light sheet. Given the geometry of the SPIV system, the lenses used provide sufficient depth of field over the full thickness of the light sheet, so the Scheimpflug arrangement [52] is not utilized.²

The angle ψ between the cameras in a rotational stereoscopic setup (ref. Figure 2.2) has a direct impact on the error of the out-of-plane component for the velocity field. If the in-plane uncertainty is σ_i , then the out-of-plane error will vary as $\sigma_o = \sigma_i / \tan(\psi/2)$ [41, 51]. While $\psi = 90^\circ$ may be optimal, practical considerations may dictate a smaller value. In our case, for example, a separation angle larger than

²Cameras constructed using the Scheimpflug arrangement have the image sensor rotated such that the lens plane, object plane (*i.e.*, the light sheet in our case), and image plane all intersect in a common line. This configuration improves focus across the lateral extent of the light sheet.




Figure 2.2: Camera separation angle.

$\psi = 18^\circ$ resulted in the tank structure occluding part of the cameras' field of view.

However in addition to occluding objects, one must also consider depth of field with large camera separation angles. Unlike a pinhole camera, a real camera can only maintain a sharp focus for objects that are within a certain range of distances from the lens [60]. As the cameras in a SPIV system are rotated away from each other, their view of the light sheet becomes more and more oblique. Depending on the lens focal length, lens aperture, the distance to the light sheet, and physical size of the CCD pixels, the extreme edges of the light sheet may blur. To compensate and provide the required depth of field, one must close down the lens aperture or modify the internals of the camera and lens to satisfy the Scheimpflug condition [52, 51].

The light sheet for our setup is provided by two programmable 150 W tungsten halogen light sources (Volpi Intralux DC-1100) connected to a custom, dual-input Volpi fiber optic line generator. A 51 mm focal length cylindrical Fresnel lens collimates the light from the line generator producing an 8 mm thick light sheet. Generally speaking, one wants the thickness of the light sheet to be less than the smallest feature of the flow perpendicular to the sheet. If particles at the back of the light sheet move in one direction while those in front move in another, the observed particle pattern will deform through time making a typical correlation-type PIV algorithm fail. On the other hand, if the out-of-plane velocity component is strong, then the light sheet thickness must be chosen to ensure the same set of particles remains within the sheet

for acquisition of the two necessary image frames [12].

We utilize white light for two reasons. First, white light is important for liquid crystals if one wants to visualize a range of temperatures as opposed to isotherms [17]. Second, white light sources are typically much cheaper than a high power laser. Nevertheless, white light can be problematic. Most white light sources use bulbs which are much larger in size than the beam of a laser. In order to produce a light sheet that is sufficiently small, the user is forced to collimate the light from the source using optics. Collimating white light can be a challenge given the areal extent of the bulb and the multiple wavelengths that may be present [60]. Finally, most white light sources come nowhere near the intensity available from a high power laser, so the user must compensate with longer exposures and wider lens apertures. Xenon arc lamps are available, but their cost rivals that of a laser, and the user must deal with the copious UV and IR radiation these sources produce [49].

As mentioned above, PIV usually employs a thin light sheet and hence captures velocity and temperature data for a limited region of the flow field. To extract information on a larger scale, a different PIV scheme must be utilized [9, 7, 20], or the light sheet and cameras must be physically moved [11, 26]. Given the creeping nature of our flows, we chose the latter configuration. The cameras and the fiber optic line generator are mounted on motion-controlled linear slides powered by Applied Motion Products' Stac6-Si drives. The line generator is swept across the entire tank, typically in 5 mm increments, and the cameras are moved to maintain focus of the tracer particles within the light sheet. The step distance for the cameras is $1/n$ that of the light sheet, where n is the refractive index of the working fluid (n is set to 1.49 for the present work) [44].

The primary advantage of scanning systems such as ours is the relative simplicity of design and construction. Nevertheless, the virtual speed limit the system places on flow dynamics must be carefully considered. Given that the light sheet and camera

Photogrammetric		
In plane	Out of plane	Thermochromic
42 μm	190 μm	0.05 $^{\circ}\text{C}$

Table 2.2: RMS calibration errors for the SPIV setup of Figure 2.1.

must move to photograph different sections of the flow field, each of these regions are imaged at different times. Accounting for the speeds of the camera shutter and motion control hardware, sweeping across the tank to acquire a single sequence of 52 images requires 20 s with our setup. However, our flows of interest are sufficiently slow (maximum velocities on the order of 0.2 mm/s) that we can temporally interpolate the velocity and temperature data from each plane to produce an instantaneous view of the 3D flow field. In those cases where better simultaneity is required, mirror scanners and stationary cameras can be employed [12, 33], or the user can resort to a different volumetric PIV technique [35, 32].

2.3 Calibration

In order to extract useful information from the camera images, two calibrations are performed. The photogrammetric calibration outlined in Section 2.3.1 provides a mapping between the image coordinates of the camera CCD sensor and laboratory coordinates, while the thermochromic calibration of Section 2.3.2 permits computation of temperature values based on the observed color of the liquid crystals. The calibration errors obtained for our setup (detailed in Sections 2.3.1 and 2.3.2) are summarized in Table 2.2 for reference.

2.3.1 Photogrammetric calibration

As mentioned previously, cameras project the 3D world onto a 2D sensor or film and will generally introduce distortions in the process [60]. Therefore, in order to form accurate estimates of flow variables, some means is necessary to relate image

coordinates back to locations in laboratory space. The process of photogrammetric calibration provides this mapping and will enable raw images to be dewarped (*i.e.*, projected onto the mid-plane of the light sheet). Pixels in these dewarped images are then directly related to laboratory units (*e.g.*, mm) by a constant scale factor.

The photogrammetric mapping produced during the calibration procedure is specific to a given refractive environment. As a result, fluids and other refractive materials that will be used during later experiments, must be present in the same configuration during calibration. Seeding of the working fluid with passive tracers is not necessary since the particles are present in trace amounts.

There are numerous approaches that one can take for photogrammetrically calibrating the cameras, including the explicit ray tracing of the full optical environment. However the fitting of a simplified camera model has been utilized with great success in computer vision applications [66, 70, 65] as well as PIV [75, 74], and this technique is employed here. The simplified camera model obviates the need for a detailed optical analysis of the entire experimental system by capturing the essential properties of a projective camera operating in a refractive environment. Hence in our approach, the output of photogrammetric calibration is a model camera formulation that can be used to simulate (and undo) the image formation process of the the real camera.

For our system, photogrammetric calibration proceeds by first imaging a target consisting of intersecting lines with a known spacing as shown in Figure 2.3. The intersection locations are extracted from the raw images of each camera, and a pinhole camera model is fit to the data. Each of these steps is discussed below.

To ensure proper alignment of the calibration target with respect to the light sheet, the target is mounted on a rigid structure that securely engages the top of the inner tank. At the start of the calibration cycle, the target is installed such that it is located in the center of the tank (Figure 2.4), and the cameras are adjusted to image a common region of the target. Once the cameras have been suitably adjusted, the




Figure 2.3: Calibration target. Lines are spaced 6.35 mm apart.

camera orientation and lens focus are locked to prevent future changes that would render the calibration invalid. With the position of the cameras held fixed along the z -axis, the calibration target is translated fore and aft on the target mounting fixture such that it is imaged by both cameras at 5 locations that span the full 8 mm thickness of the light sheet.

After the images of the calibration target are obtained, the image coordinates for the line intersections are determined. We first compute the locations of the line intersections using a linear Hough transform [18], then recompute to higher precision using cross-correlation and a cross-shaped template. The motivation for using a two phase approach is to provide a degree of tolerance to missing intersections. Although great care is taken during photogrammetric calibration to ensure that images of the target have minimal imperfections, periodically a bubble or other debris will obscure a line intersection when the target is installed in the tank. Unfortunately, cross-correlation cannot match the template to obscured intersections. The Hough transform, however,




Figure 2.4: Setup for photogrammetric calibration. Note that the front and left-side panels of both the inner and outer tanks have been removed for clarity.

is robust to this issue since it provides the location and orientation of entire lines, as opposed to smaller features like individual intersections. Any missing intersections from cross-correlation can then be provided by the set of lines extracted via the Hough transform.³

Following extraction of line intersection coordinates from the calibration images, the pinhole camera model of Tsai [66] is fit to the data for each camera. The pinhole camera takes a line extending from the pinhole into laboratory space and projects it onto a single point (*e.g.*, a pixel) in the image plane as depicted in Figure 2.5. The basic pinhole model has no concept of refraction in lenses or other objects, so these effects are modeled as distortions or accommodated by other parameters such as focal length. The full, distortion-compensated camera model consists of 18 parameters, six of which are typically identified as external with the remainder being termed internal. The external parameters define the location and orientation of the camera in three-space, while the internal parameters describe the pinhole camera itself. The fit is performed in two stages using an approach similar to that of Heikkila and Silven [31].

³The Hough transform parameterizes linear image features using a cylindrical coordinate system. A line in the xy -space of the image then maps to a single point in $\rho\theta$ -space of the transform via the parameterized equation for a line given by $x \cos \theta + y \sin \theta = \rho$, where ρ extends from the origin and is normal to the line [18, 29].

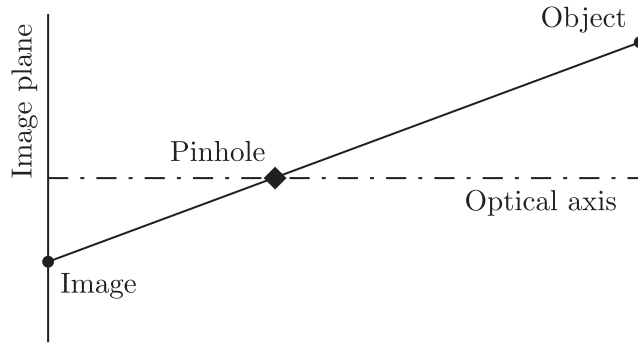


Figure 2.5: Pinhole camera showing an object point and the corresponding image point. All object points along the illustrated ray will map to the same point on the image plane.

The first stage of our calibration procedure simply solves for the distortionless camera parameters using a direct linear transform [31, 65]. The second stage incorporates the distortion model of Weng [70] using the Levenberg-Marquardt optimization algorithm [48] to solve the resulting non-linear least squares problem. For more information regarding use of Tsai’s camera model with PIV applications, the reader is referred to papers by Wieneke [74] and Willert [75].

After the camera models have been calibrated, an estimate of system error can be obtained using the same setup of Figure 2.4 and moving the target a known distance. These images can then be analyzed (ref. Section 2.4) and the extracted displacements compared to the known. Installing the calibration target such that it is positioned in the center of the tank and translating it 8.70 mm along the z -axis yields RMS error estimates of 42 μm for the in-plane displacement components and 190 μm for the out-of-plane components.

2.3.2 Thermochromic calibration

The observed color of thermochromic liquid crystals is a complex function of temperature, view angle, and particulars of the optical environment (*e.g.*, the spectral properties of the working fluid) [23, 54]. Hence, to employ TLC’s in thermometry




Figure 2.6: Definition of view angle and z -depth for thermochromic calibration. The light sheet is represented by the gray shaded region.

applications, an additional calibration must be performed.

The view angle dependence of observed TLC hue requires some means of quantitatively determining the angle between the light sheet propagation vector and the ray from camera to TLC. As long as the camera model of Section 2.3.1 is valid (*i.e.*, barring changes to the laboratory refractive environment, the physical camera orientation, or focus), the pinhole camera external parameters can be used to compute the view angle as depicted in Figure 2.6. Consequently, photogrammetric calibration must be performed prior to thermochromic calibration.

The objective of thermochromic calibration is to construct a polynomial of the form

$$T(z, \theta, h) = \sum_{i=0}^l \sum_{j=0}^m \sum_{k=0}^n c_{i,j,k} z^i \theta^j h^k = \mathbf{c}'\mathbf{b} \quad (2.1)$$

that provides fluid temperature, T , as a function of z -depth, view angle (θ), and observed TLC hue (h). We include z -depth dependence in the polynomial to compensate for the yellow color of our syrup. The coefficients, $c_{i,j,k}$, are determined using a penalized least squares approach detailed below, with the polynomial order chosen

to provide the best fit.

Because photogrammetric calibration models the image formation process within a particular optical setting, we took great care during that process (ref. Section 2.3.1) to ensure all aspects of the refractive environment were accounted for in the setup. Similarly, thermochromic calibration is providing a means to estimate temperature based on color, so characteristics of the experimental setup that affect color must be included in thermochromic calibration. Recall that in addition to TLC's we utilize secondary tracers for our PIV work. Although image processing is used to remove the secondary tracers from raw images prior to hue extraction (ref. Section 2.4), the light scattered from the secondary tracers can still contaminate the observed TLC hue to a certain degree. Hence the corn syrup used during thermochromic calibration is mixed with the same TLC and secondary tracer concentrations that will be used for later experiments.

We also find that RMS error for the thermochromic calibration (see below) can be reduced by a factor of $O(5)$ if raw images are corrected for the color cast of our syrup and tungsten halogen light sources prior to additional processing. Following Farina *et al.* [21], a 3×3 matrix \mathbf{A} can be constructed to transform the red, green, and blue (RGB) values \mathbf{r} of an image pixel to corrected values \mathbf{r}_c

$$\mathbf{r}_c = \mathbf{A}\mathbf{r}. \tag{2.2}$$

The color correction matrix \mathbf{A} is obtained by imaging an X-Rite ColorChecker chart (model MSCCC) located in the middle of the tank filled with syrup and no passive tracers. The ColorChecker is composed of a series of 24 colored squares having known RGB values, and \mathbf{A} can be determined via least squares from an image of the target. A comparison between raw and color-corrected images is presented in Figure 2.7.

Once the inner tank is prepped with syrup, the tank and its contents are allowed




Figure 2.7: Demonstration of color correction scheme on a dewarped image from a thermal plume experiment. The heater is located at the bottom center of the images. (a) Uncorrected image. (b) Corrected image.

48 hours to equilibrate to a specific bath temperature. The uniform temperature fluid is then imaged at five planes spaced evenly along the full z -depth of the inner tank. This process is repeated for bath temperatures ranging from 25.2 to 26.8 °C in 0.2 °C increments, yielding the series of images that will be used to determine the coefficients of Eq. 2.1. A sample image is presented in Figure 2.8 for reference. The process of allowing the syrup to equilibrate to each bath temperature setpoint is the most time consuming part of the operation, however not all thermochromic calibrations will be subject to this constraint. In fluids that are much less viscous, the process of equilibration can be expedited by mechanical stirring. An even simpler solution is available for setups which can minimize heat loss from the tank sidewalls (*e.g.*, a large width to height aspect ratio Rayleigh-Bénard convection apparatus). Namely, a known temperature gradient can be imposed across the fluid layer and a single set of images acquired [13].

As observed by others, we find that smaller view angles (θ) provide improved color-play [21, 28]. Consequently, we only use imagery from the right camera (ref.




Figure 2.8: Dewarped thermochromic calibration image for an isothermal fluid temperature of $26.2\text{ }^{\circ}\text{C}$. The fiber optic line generator is positioned out of view on the right side of the image. Note that the lateral color variation shown is due entirely to view angle. The small yellow-white particles are secondary tracers.

Figure 2.6) for TLC thermometry. Hue, saturation, and intensity (HSI) channels are computed from the raw RGB channels for each of the calibration images using the HSI formulation of Ledley *et al.* [43]. The HSI channels are subsequently dewarped by projecting onto the mid-plane of the light sheet using the photogrammetric calibration of Section 2.3.1 and then broken into 32×32 pixel windows for additional processing. Applying Westerweel and Scarano’s universal outlier detection scheme [73], which itself is an adaptation of the median filter [29], secondary tracers are selectively removed from each of the image windows. Valid TLC particles are then identified with HSI value thresholds, and the hue for the entire window is taken as the average hue of pixels associated with valid TLC tracers. An example of the processed hue field is shown in Figure 2.9.

The other two independent variables in Eq. (2.1) are also readily determined. The z -depth for each image is recorded by the SPIV hardware during acquisition, and hence known. The view angle for each 32×32 window, as depicted in Figure 2.6, is computed using the photogrammetric parameters of the rightmost camera and the laboratory coordinates of the window.

With the hue, z -depth, and view angles available, the coefficients of Eq. (2.1) are determined using a penalized least squares technique which we now explain. The highly nonlinear temperature versus hue characteristics of TLC’s suggest a high polynomial order, n , for the hue dependence in Eq. (2.1). As is well known, however, interpolation with high order polynomials tends to produce oscillations [39]. This behavior can be controlled by introducing a roughness penalty term to the standard least squares fit criterion [53]. Let \mathbf{y} represent the vector of known bath temperatures. The penalized metric to be minimized becomes

$$\text{PENSSE} = \|\mathbf{y} - \mathbf{B}\mathbf{c}\|^2 + \lambda\mathbf{c}'\mathbf{R}\mathbf{c} \quad (2.3)$$




Figure 2.9: Processed hue field for Figure 2.8. Note that processing has been restricted to a particular region of interest within the full image.

where the first term represents the standard sum of squares error, and the second term contains the curvature penalty. In Eq. (2.3), \mathbf{B} is the matrix of monomial basis functions evaluated at the hue, view angle, and z -depth values for each block, \mathbf{c} the unknown polynomial coefficients, λ a smoothness parameter, and \mathbf{R} the roughness penalty matrix given by

$$\mathbf{R} = \int_z \int_\theta \int_h \nabla^2 \mathbf{b} \nabla^2 \mathbf{b}' dz d\theta dh \quad (2.4)$$

Note that in Eq. (2.4) as in Eq. (2.1), \mathbf{b} is a vector of the monomial basis functions. Differentiating Eq. (2.3) with respect to \mathbf{c} yields the following system of equations that can be inverted for the unknown coefficients

$$(\mathbf{B}'\mathbf{B} + \lambda\mathbf{R})\mathbf{c} = \mathbf{B}'\mathbf{y} \quad (2.5)$$

Although sophisticated techniques such as generalized cross validation [53] exist for determining an optimal value of the smoothing parameter, for the present work λ is chosen on a trial and error basis such that the resulting polynomial of Eq. (2.1) is essentially monotonic.

Using the above procedure, a polynomial of order $(l, m, n) = (1, 1, 7)$, and $\lambda = 0.002$, the sample calibration curves shown in Figure 2.10 are produced for our TLC formulation and SPIV setup. These curves are obtained by evaluating Eq. (2.1) with z set to the mid-tank position, and five θ values that span the full field of view. To determine a measure of calibration error, the polynomial of Eq. (2.1) is evaluated at the raw hue, view angle, and z -depths used to compute the polynomial coefficients. The error is then taken as the difference between the computed temperature and the known bath temperature. In this manner, the aggregate RMS error for the calibration is 0.05 °C. A more detailed breakdown of calibration error as a function of bath temperature is presented in Figure 2.11.




Figure 2.10: Sample curves from the thermochromic calibration. Curves are obtained by evaluating Eq. (2.1) with z set to the mid-tank position. Note that a hue of 0.0 corresponds to pure red, while a hue of 0.67 represents pure blue.




Figure 2.11: Thermochromic calibration error versus bath temperature. Central box encloses the 25-75th percentiles and the median is denoted with a horizontal line. Whiskers mark 99th percentile confidence intervals, and outliers are represented as + .

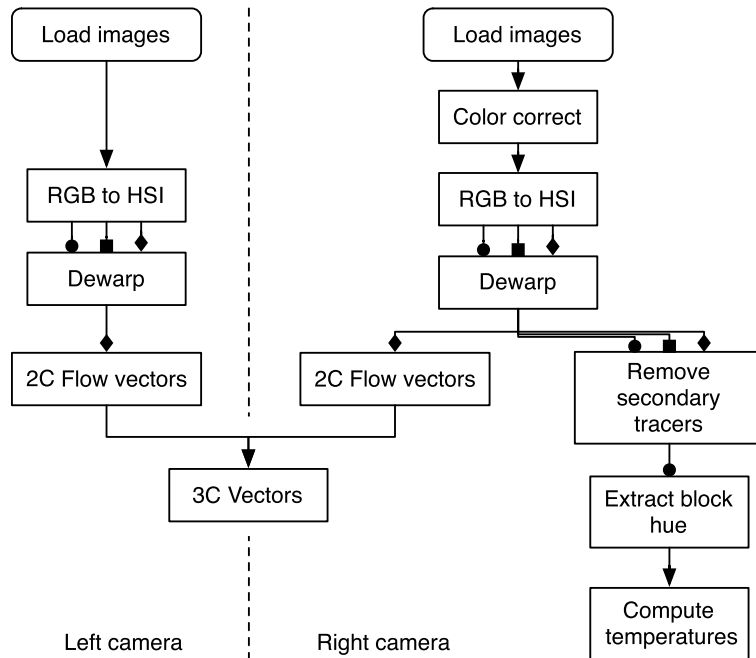


Figure 2.12: High-level overview of data reduction process. Arrows represent generic data flow. Circles represent the hue channel, squares the saturation channel, and diamonds the intensity channel.

2.4 Data processing

Once the system has been calibrated it can be put into service capturing images of experiments, however these images must still be processed to extract velocity and temperature data. A high-level overview of our data reduction process is presented in Figure 2.12. To start, the RGB channels from the left and right cameras are converted to HSI color space, and HSI channels are subsequently projected onto the mid-plane of the light sheet using the photogrammetric calibration of Section 2.3.1. The temperature field is computed using the same image processing steps described in Section 2.3.2, while the scheme utilized for velocity estimation is detailed below.

Extracting all three components (3C) of the velocity vector starts with the compu-

tation of two component (2C) displacement vectors for each camera (*i.e.*, the in-plane components). We compute 2C flow vectors using the PIV technique of Correlation Image Velocimetry (CIV) [25]. In the CIV method, a small image window of size $M \times N$ pixels from one image, I_1 , forms a template (*i.e.*, the pattern) which is iteratively translated over a larger search region of the second image, I_2 . At each location of the template within the search region, a similarity measure is computed based on how well the template matches the underlying search region. The translation that maximizes similarity is then taken as the 2C displacement vector. The process is repeated for several template windows which together divide I_1 into a uniform, two-dimensional array of grid points. Consequently, each grid point provides a single displacement vector.

Our implementation of the CIV scheme employs two different similarity measures for locating the template within the search region. These two metrics are normalized cross-correlation (NCC) and the similarity measure of Barnea and Silverman’s [8] Sequential Similarity Detection Algorithm (SSDA). The primary motivation for using two similarity measures is computational cost. Before discussing the two measures in more detail, however, we note that robustness to intensity variation is an important criteria when working with TLC’s. Reflecting back on Figure 2.7, strong variation in intensity can be seen. If the similarity measure used to perform image registration is sensitive to intensity, as is the case with standard cross-correlation, then false registration can occur [10]. In the case of Figure 2.7, a template window from the dark zone near the heater may be erroneously matched to the green band of TLC’s purely because that region of the search image has a high mean intensity (*i.e.*, not because the particle patterns match). Both normalized cross-correlation and SSDA have built in tolerance for intensity variation. A brief overview of these two measures follows.

Let (i, j) represent the current position of the template from I_1 within the search

region of I_2 . The normalized cross-correlation (NCC) coefficient given by

$$\rho(i, j) = \frac{\sum_{k=1}^M \sum_{l=1}^N [I_1(k, l) - \bar{I}_1][I_2(i+k, j+l) - \bar{I}_2(i, j)]}{\left[\sum_{k=1}^M \sum_{l=1}^N [I_1(k, l) - \bar{I}_1]^2 \sum_{k=1}^M \sum_{l=1}^N [I_2(i+k, j+l) - \bar{I}_2(i, j)]^2 \right]^{1/2}} \quad (2.6)$$

is a widely used similarity measure for PIV and image registration applications with a long history ([10, 61]). The measure is robust to intensity variations within and between the two image regions being registered, and the magnitude of the coefficient provides an absolute measure of match quality. There is, however, one significant drawback to using the NCC coefficient for image registration: the measure is computationally expensive. In Eq. (2.6), I_1 and I_2 represent the two frames of a dewarped pair of images from a single camera. Note that \bar{I}_1 is the average intensity of the template window, and only has to be computed once. On the other hand, \bar{I}_2 , the average of the corresponding $M \times N$ region from the search frame, must be re-evaluated for each location (i, j) of the template.

To reduce the computational overhead associated with the normalized cross-correlation coefficient, the image blocks described above are first registered using a method derived from Barnea and Silverman's [8] Sequential Similarity Detection Algorithm (SSDA). The SSDA algorithm uses a more economical similarity measure that is simply the sum of the absolute differences between the template and search blocks

$$\gamma(i, j) = \sum_{k=1}^M \sum_{l=1}^N |I_1(k, l) - I_2(i+k, j+l)| \quad (2.7)$$

Like NCC, the SSDA measure is robust to mean intensity variations and provides an absolute measure of match quality.

Using these two similarity measures for registration of the template window to the

search region, extraction of 2C vectors proceeds in a four-step fashion as follows. Step 1 of the 2C vector extraction process utilizes the SSDA similarity measure, a large search region, and coarse, 32×32 pixel template window over which the similarity measure is computed. Employing large search regions in Step 1 accommodates large particle translations between frames [25].

Step 2 proceeds to refine the Step 1 results using the same template window size, but employing the normalized cross-correlation measure and a much smaller search region of a few pixels in both dimensions. Step 2 also provides sub-pixel precision by fitting to the NCC measure peak to a Gaussian in each direction [76].

Velocity fields are rarely gradient free, and large displacements are often accompanied by excessive particle pattern deformation between frames. If the particle pattern observed in the template window deforms sufficiently under the influence of the flow, a high-quality registration of the template will not be possible in the preceding steps (ref. Chapter III). The result is typically a poor estimate of the displacement vector. Given that the two similarity measures, NCC and SSDA, provide an absolute estimate of confidence in the registration, low quality vectors computed using the CIV scheme generally correspond to low values of the similarity measures. The goal of Step 3 is to isolate these poor quality vectors and form a better estimate using an altogether different technique. We do this by transitioning to particle tracking velocimetry (PTV) [36, 45, 46] when a low value from the NCC measure of Step 2 suggests the CIV scheme has failed (ref. Chapter III). PTV uniquely identifies and tracks the motion of each individual tracer, so the method is in theory immune to many circumstances (*e.g.*, the afore mentioned particle pattern deformation) that negatively impact the CIV technique. ⁴

Upon entering Step 4, an estimate of the displacement field has been computed for the entire image, but despite the processing done in previous stages, some spurious

⁴PTV can certainly be used in place of CIV to compute all velocity vectors, however we restrict PTV use to low quality vectors due to the technique's high computational cost (ref. Chapter III).

vectors can remain. Step 4 removes these spurious vectors and further refines the displacement field through an iterative loop. At the start of each iteration, the 2C displacement vectors are filtered first using the outlier detection scheme of Westerweel and Scarano [73] and then by Gaussian smoothing [29]. The filtering step ensures that the displacement field estimate is a smoothly varying function of spatial coordinates, so the template window can be deformed by shifting each pixel using the estimated displacements [56]. Once the template window is deformed, Steps 1-3 are repeated, and a correction for each displacement vector is computed.

After 2C vectors have been computed for each camera using the above four-step process, the third component can be determined. The geometry for a single camera is illustrated in Figure 2.13 with \mathbf{x}_1 representing the laboratory coordinates of a particle cluster from frame 1 and \mathbf{x}_2 representing the true coordinates of the cluster at the time of frame 2. The out-of-plane displacement Δz generates an error $\delta\mathbf{x}$ that contaminates the true in-plane displacement $\Delta\mathbf{x}$. This contamination is shown for the x -component of displacement in Figure 2.13. As a consequence of the out-of-plane motion, the process of two component registration described above extracts an in-plane displacement of $\Delta\mathbf{x}' = (\Delta x', \Delta y')$ as though the particle pattern moved within the xy -plane from \mathbf{x}_1 to \mathbf{x}'_2 . Given that two cameras a and b are used in the stereoscopic system, the out-of-plane component can be determined. Two component flow processing provides 4 knowns $\Delta\mathbf{x}'_a$ and $\Delta\mathbf{x}'_b$, and the desired three components of the true displacement $\Delta\mathbf{x}$ can then be obtained via least squares. Let

$$r = \left[(x_0 + \Delta x' - x_c)^2 + (y_0 + \Delta y' - y_c)^2 \right]^{1/2}, \quad (2.8)$$

and

$$\begin{aligned}\tan \eta &= \frac{r}{z_c} \\ \cos \phi &= \frac{x_0 + \Delta x' - x_c}{r} \\ \sin \phi &= \frac{y_0 + \Delta y' - y_c}{r}.\end{aligned}\tag{2.9}$$

The system of equations for $\Delta \mathbf{x}$ is then

$$\begin{pmatrix} \Delta x'_a \\ \Delta y'_a \\ \Delta x'_b \\ \Delta y'_b \end{pmatrix} = \begin{pmatrix} 1 & 0 & \tan \eta_a \cos \phi_a \\ 0 & 1 & \tan \eta_a \sin \phi_a \\ 1 & 0 & \tan \eta_b \cos \phi_b \\ 0 & 1 & \tan \eta_b \sin \phi_b \end{pmatrix} \begin{pmatrix} \Delta x \\ \Delta y \\ \Delta z \end{pmatrix},\tag{2.10}$$

and the 3C velocity vector is computed as

$$\mathbf{u} = \frac{\Delta \mathbf{x}}{\Delta t},\tag{2.11}$$

where Δt is the time interval between images I_1 and I_2 .

2.5 Results

To illustrate the setup's utility for investigating slow flows which are characteristic of planetary interiors, we present selected results from two datasets. Velocity and temperature data are obtained on a 3D grid of $60 \times 78 \times 52$ cells along the x , y , and z axes, respectively. Each cell has dimensions of $(\Delta x, \Delta y, \Delta z) = (0.28, 0.28, 0.50)$ cm yielding a total data domain size of $16.8 \times 21.8 \times 26.0$ cm. For the results presented, the domain is tangent to the bottom of the inner tank and centered over the heater along the x and z axes.




Figure 2.13: Geometry for 3C reconstruction showing a single camera at \mathbf{x}_c .

We are particularly interested in mass transport and mechanical stirring that occurs in laminar plumes as these mechanisms have important implications for the thermochemical evolution of Earth’s mantle and other terrestrial planets [55, 57]. Consequently, as part of our standard post-processing procedures, we typically use the extracted SPIV velocity field to numerically advect passive tracers using a fourth-order Runge-Kutta scheme [63]. The numerical tracer field can then be used to investigate the origin of mass within the plume as well as its stretching history (ref. Chapter IV).

Figure 2.14 presents a snapshot of velocity, vorticity, temperature, and mass origin for a plume dataset generated with the heater delivering 0.72 W into syrup having an initial temperature of 25.2 °C. The derived quantities of vorticity and mass origin have been computed from the time-dependent 3D-3C velocity field obtained via SPIV. To produce Figure 2.14(b), the extracted SPIV velocity field is numerically seeded with 1.7×10^6 tracers. Each of these tracers is assigned a scalar value based on the initial height of the tracer above the tank bottom. The tracers are then computationally advected using the SPIV velocity field as described.

Our second example is of a plume that failed to reach the top of the tank following

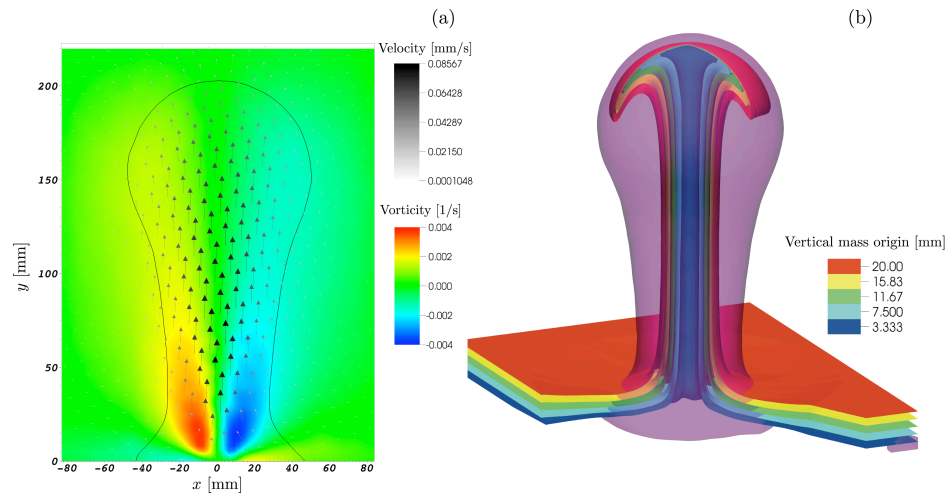


Figure 2.14: Flow variables for an experimental plume analyzed with the SPIV system. The data are shown 1.8 hr after heater activation. (a) Cross section of flow field at mid-tank. Heater is located at the origin. Pseudocolor background is the component of vorticity perpendicular to the section. The solid black line represents the $25.4\text{ }^{\circ}\text{C}$ temperature contour. Arrows are velocity vectors. (b) 3D perspective view showing the vertical origin of a portion of the mass transported by the plume. The heater is located at the bottom of the image and centered beneath the upwelling. The buoyant mass is convecting heat away from the heater, and the $25.4\text{ }^{\circ}\text{C}$ contour is shown here as a translucent purple solid.




Figure 2.15: Perspective view showing the vertical origin of mass transported by a failed plume. Heater is located at the bottom of the image and centered beneath the upwelling.

heater activation. Instead, this plume becomes unstable and eventually collapses back on itself. The reasons for collapse are not fully understood, however it is likely that compositional layering of the syrup occurred as the sealant used to secure the inner tank lid cured over a period of two weeks. While curing, the tank is not completely full of syrup and the syrup at the air interface dehydrates. Once the sealant has cured, the tank is usually stirred for several days by activating the heater and allowing the fluid to convect. The dataset shown here was collected prior to that mixing cycle. Again the syrup was initially at $25.2\text{ }^{\circ}\text{C}$ prior to heater activation for this particular case. The heater was then digitally controlled such that it maintained a temperature of $38\text{ }^{\circ}\text{C}$ following a warm up transient that lasted 18 s. Mass transport for this flow is visualized using the same techniques as for the previous case. The vertical origin of advected mass is shown 4.7 hours following heater activation in Figure 2.15. Instead of producing the mushroom shape seen in Figure 2.14(b), the uppermost leading edge of the plume head has collapsed. As a consequence, hot buoyant fluid is forced to flow out around the downwelling material producing the observed lateral bulge.

Overall, both the SPIV system and data processing software detailed within the present paper have proven very useful for the analysis of slow flows like those discussed above. These creeping flows are of particular interest to the geodynamics community

because they simulate features, such as thermal plumes [19, 59], thought to be relevant to convection in Earth’s mantle [57, 62]. While PIV is not new to the experimental geodynamics community, much of the existing work has been spatially limited to measuring flow variables for a single cross-sectional plane within the flow [16, 37, 68]. For simple flows with a high degree of symmetry (*e.g.*, the plume of Figure 2.14), single-plane data is typically not a significant limitation. Convection within Earth’s mantle, however, is expected to be unsteady and inherently three-dimensional [22, 38, 42, 67]. Therefore, the ability to capture the full three-dimensional, time-dependent flow field is critical to leveraging the full utility of laboratory experiments for the study of mantle flows.

2.6 References

- [1] R. J. Adrian. Laser velocimetry. In R. J. Goldstein, editor, *Fluid Mechanics Measurements*, chapter 5, pages 155–244. Hemisphere, 1983.
- [2] R. J. Adrian. Particle-imaging techniques for experimental fluid mechanics. *Annual Review of Fluid Mechanics*, 23:261–304, 1991.
- [3] M. A. Ahmadshahi, S. Krishnaswamy, and S. Nematnasser. Dynamic holographic-electronic speckle-pattern interferometry. *Journal of Applied Mechanics*, 60(4):866–874, 1993.
- [4] D. Ambrosini, D. Paoletti, and G. S. Spagnolo. Study of free-convective onset on a horizontal wire using speckle pattern interferometry. *International Journal of Heat and Mass Transfer*, 46(22):4145–4155, 2003.
- [5] M. R. Anderson and J. W. Baughn. Liquid-crystal thermography: Illumination spectral effects. Part 1 - experiments. *Journal of Heat Transfer*, 127(6):581–587, 2005.
- [6] M. R. Anderson and J. W. Baughn. Thermochromic liquid crystal thermography: Illumination spectral effects. Part 2 - theory. *Journal of Heat Transfer*, 127(6):588–597, 2005.
- [7] M. P. Arroyo and K. D. Hinsch. Recent developments of PIV towards 3D measurements. In *Particle Image Velocimetry*, volume 112 of *Topics in Applied Physics*, pages 127–154. Springer, 2008.

- [8] D. I. Barnea and H. F. Silverman. A class of algorithms for fast digital image registration. *IEEE Transactions on Computers*, C 21(2):179–186, 1972.
- [9] D. H. Barnhart, R. J. Adrian, and G. C. Papen. Phase-conjugate holographic system for high-resolution particle-image velocimetry. *Applied Optics*, 33(30):7159–7170, 1994.
- [10] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, 1992.
- [11] C. Brücker. Digital-Particle-Image-Velocimetry (DPIV) in a scanning light-sheet: 3D starting flow around a short cylinder. *Experiments in Fluids*, 19:255–263, 1995.
- [12] C. Brücker. 3-D scanning-particle-image-velocimetry: Technique and application to a spherical cap wake flow. *Applied Scientific Research*, 56(2-3):157–179, 1996.
- [13] M. Ciofalo, M. Signorino, and M. Simiano. Tomographic particle-image velocimetry and thermography in Rayleigh-Bénard convection using suspended thermochromic liquid crystals and digital image processing. *Experiments in Fluids*, 34:156–172, 2003.
- [14] D. Dabiri. Digital particle image thermometry/velocimetry: a review. *Experiments in Fluids*, 46(2):191–241, 2009.
- [15] D. Dabiri and M. Gharib. Digital particle image thermometry: the method and implementation. *Experiments in Fluids*, 11:77–86, 1991.
- [16] A. Davaille, M. Le Bars, and C. Carbonne. Thermal convection in a heterogeneous mantle. *Comptes Rendus Geoscience*, 335:141–156, 2003.
- [17] A. Davaille, A. Limare, F. Touitou, I. Kumagai, and J. Vatteville. Anatomy of a laminar starting thermal plume at high Prandtl number. *Experiments in Fluids*, pages 1–16, 2010.
- [18] R. O. Duda and P. E. Hart. Use of the Hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1975.
- [19] R. A. Duncan and M. A. Richards. Hotspots, mantle plumes, flood basalts, and true polar wander. *Reviews of Geophysics*, 29(1):31–50, 1991.
- [20] G. E. Elsinga, F. Scarano, B. Wieneke, and B. W. van Oudheusden. Tomographic particle image velocimetry. *Experiments in Fluids*, 41(6):933–947, 2006.
- [21] D. J. Farina, J. M. Hacker, R. J. Moffat, and J. K. Eaton. Illuminant invariant calibration of thermochromic liquid crystals. *Experimental Thermal and Fluid Science*, 9(1):1–12, 1994.

- [22] C. G. Farnetani and H. Samuel. Beyond the thermal plume paradigm. *Geophysical Research Letters*, 32(7), 2005.
- [23] J. L. Fergason. Liquid crystals in nondestructive testing. *Applied Optics*, 7(9):1729–1737, 1968.
- [24] A. Fincham and G. Delerce. Advanced optimization of correlation imaging velocimetry algorithms. *Experiments in Fluids*, 29:S13–S22, 2000.
- [25] A. M. Fincham and G. R. Spedding. Low cost, high resolution DPIV for measurement of turbulent fluid flow. *Experiments in Fluids*, 23(6):449–462, 1997.
- [26] N. Fujisawa and S. Funatani. Simultaneous measurement of temperature and velocity in a turbulent thermal convection by the extended range scanning liquid crystal visualization technique. *Experiments in Fluids*, 29:S158–S165, 2000.
- [27] N. Fujisawa, T. Nakajima, N. Katoh, and Y. Hashizume. An uncertainty analysis of temperature and velocity measured by stereo liquid-crystal thermometry and velocimetry. *Measurement Science and Technology*, 15:799–806, 2004.
- [28] S. Funatani and N. Fujisawa. Simultaneous measurement of temperature and three velocity components in planar cross section by liquid-crystal thermometry combined with stereoscopic particle image velocimetry. *Measurement Science and Technology*, 13:1197–1205, 2002.
- [29] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Prentice Hall, 2 edition, 2002.
- [30] A. Günther and P. R. von Rohr. Influence of the optical configuration on temperature measurements with fluid-dispersed TLCs. *Experiments in Fluids*, 32:533–541, 2002.
- [31] J. Heikkila and O. Silven. A four-step camera calibration procedure with implicit image correction. In *1997 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1106–1112, 1997.
- [32] K. D. Hinsch. Holographic particle image velocimetry. *Measurement Science and Technology*, 13(7):R61–R72, 2002.
- [33] T. Hori and J. Sakakibara. High-speed scanning stereoscopic PIV for 3D vorticity measurement in liquids. *Measurement Science and Technology*, 15(6):1067–1078, 2004.
- [34] H. T. Huang, H. E. Fiedler, and J. J. Wang. Limitation and improvement of PIV part II: Particle image distortion, a novel technique. *Experiments in Fluids*, 15(4-5):263–273, 1993.
- [35] L. Kajitani and D. Dabiri. A full three-dimensional characterization of defocusing digital particle image velocimetry. *Measurement Science and Technology*, 16:790–804, 2005.

- [36] N. Kasagi and K. Nishino. Probing turbulence with three-dimensional particle-tracking velocimetry. *Experimental Thermal and Fluid Science*, 4:601–612, 1991.
- [37] I. Kumagai, A. Davaille, and K. Kurita. On the fate of thermally buoyant mantle plumes at density interfaces. *Earth and Planetary Science Letters*, 254:180–193, 2007.
- [38] I. Kumagai, A. Davaille, K. Kurita, and E. Stutzmann. Mantle plumes: Thin, fat, successful, or failing? constraints to explain hot spot volcanism through time and space. *Geophysical Research Letters*, 35(16):L16301, 2008.
- [39] C. B. Laney. *Computational Gasdynamics*. Cambridge University Press, 1998.
- [40] N. Laudenbach and U. R. Christensen. An optical method for measuring temperature in laboratory models of mantle plumes. *Geophysical Journal International*, 145(2):528–534, 2001.
- [41] N. J. Lawson and J. Wu. Three-dimensional particle image velocimetry: error analysis of stereoscopic techniques. *Measurement Science and Technology*, 8:894–900, 1997.
- [42] M. Le Bars and A. Davaille. Whole layer convection in a heterogeneous planetary mantle. *Journal of Geophysical Research*, 109(B3), 2004.
- [43] R. S. Ledley, M. Buas, and T. J. Golab. Fundamentals of true-color image processing. In *Proceedings of the 10th International Conference on Pattern Recognition*, volume 1, pages 791–795, 1990.
- [44] P. M. Lutjen, D. Mishra, and V. Prasad. Three-dimensional visualization and measurement of temperature field using liquid crystal scanning thermography. *Transactions of the ASME*, 123:1006–1014, 2001.
- [45] H. G. Maas, A. Gruen, and D. Papantoniou. Particle tracking velocimetry in three-dimensional flows. Part 1. Photogrammetric determination of particle coordinates. *Experiments in Fluids*, 15:133–146, 1993.
- [46] N. A. Malik, T. Dracos, and D. Papantoniou. Particle tracking velocimetry in three-dimensional flows part II: Particle tracking. *Experiments in Fluids*, 15:279–294, 1993.
- [47] A. Melling. Tracer particles and seeding for particle image velocimetry. *Measurement Science and Technology*, 8(12):1406–1416, 1997.
- [48] J. J. Moré. The Levenberg-Marquardt algorithm: Implementation and theory. In G. A. Watson, editor, *Numerical Analysis*, pages 105–116. Springer, 1978.
- [49] J. R. Oliver and F. S. Barnes. A comparison of rare-gas flashlamps. *IEEE Journal of Quantum Electronics*, 5(5):232–237, 1969.

- [50] H. G. Park, D. Dabiri, and M. Gharib. Digital particle image velocimetry/thermometry and application to the wake of a heated circular cylinder. *Experiments in Fluids*, 30:327–338, 2001.
- [51] A. K. Prasad. Stereoscopic particle image velocimetry. *Experiments in Fluids*, 29:103–116, 2000.
- [52] A. K. Prasad and K. Jensen. Scheimpflug stereocamera for particle image velocimetry in liquid flows. *Applied Optics*, 34(30):7092–7099, 1995.
- [53] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis*. Springer, 2nd edition, 2005.
- [54] E. P. Raynes. Electro-optic and thermo-optic effects in liquid crystals. *Philosophical Transactions of the Royal Society of London A*, 309(1507):167–178, 1983.
- [55] N. M. Ribe, A. Davaille, and U. R. Christensen. Fluid dynamics of mantle plumes. In J. R. R. Ritter and U. R. Christensen, editors, *Mantle Plumes: A Multidisciplinary Approach*, pages 1–48. Springer, 2007.
- [56] F. Scarano. Iterative image deformation methods in PIV. *Measurement Science and Technology*, 13(1):R1–R19, 2002.
- [57] G. Schubert, D. L. Turcotte, and P. Olson. *Mantle Convection in the Earth and Planets*. Cambridge University Press, 2001.
- [58] L. G. Shapiro and G. C. Stockman. *Computer Vision*. Prentice Hall, 2001.
- [59] N. H. Sleep. Hotspots and mantle plumes: Some phenomenology. *Journal of Geophysical Research - Solid Earth*, 95(B5):6715–6736, 1990.
- [60] W. J. Smith. *Modern Optical Engineering*. McGraw-Hill, 3rd edition, 2000.
- [61] M. Svedlow, C. D. McGillem, and P. E. Anuta. Experimental examination of similarity measures and preprocessing methods used for image registration. In *Symposium on Machine Processing of Remotely Sensed Data*, pages 9–17, 1976.
- [62] P. J. Tackley. Mantle convection and plate tectonics: toward an integrated physical and chemical theory. *Science*, 288:2002–2007, 2000.
- [63] J. C. Tannehill, D. A. Anderson, and R. H. Pletcher. *Computational Fluid Mechanics and Heat Transfer*. Taylor and Francis, 2nd edition, 1997.
- [64] P. T. Tokumaru and P. E. Dimotakis. Image correlation velocimetry. *Experiments in Fluids*, 19(1):1–15, 1995.
- [65] E. Trucco and A. Verri. *Introductory Techniques for 3-D Computer Vision*. Prentice Hall, 1998.

- [66] R. Y. 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.
- [67] P. van Keken and S. J. Zhong. Mixing in a 3D spherical model of present-day mantle convection. *Earth and Planetary Science Letters*, 171(4):533–547, 1999.
- [68] J. Vatteville, P. E. van Keken, A. Limare, and A. Davaille. Starting laminar plumes: comparison of laboratory and numerical modeling. *Geochemistry Geophysics Geosystems*, 10(12):1–9, 2009.
- [69] G. Voulgaris and J. H. Trowbridge. Evaluation of the acoustic Doppler velocimeter (ADV) for turbulence measurements. *Journal of Atmospheric and Oceanic Technology*, 15(1):272–289, 1998.
- [70] J. Weng. Camera calibration with distortion models and accuracy evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(10):965–980, 1992.
- [71] J. Westerweel. Fundamentals of digital particle image velocimetry. *Measurement Science and Technology*, 8:1379–1392, 1997.
- [72] J. Westerweel. Theoretical analysis of the measurement precision in particle image velocimetry. *Experiments in Fluids*, 29:3–12, 2000.
- [73] J. Westerweel and F. Scarano. Universal outlier detection for PIV data. *Experiments in Fluids*, 39(6):1096–1100, 2005.
- [74] B. Wieneke. Stereo-PIV using self-calibration on particle images. *Experiments in Fluids*, 39(2):267–280, 2005.
- [75] C. E. Willert. Assessment of camera models for use in planar velocimetry calibration. *Experiments in Fluids*, 41(1):135–143, 2006.
- [76] C. E. Willert and M. Gharib. Digital particle image velocimetry. *Experiments in Fluids*, 10:181–193, 1991.
- [77] C. E. Willert and M. Gharib. Three-dimensional particle imaging with a single camera. *Experiments in Fluids*, 12(6):353–358, 1992.

CHAPTER III

Hybrid PIV/PTV scheme for data recovery

3.1 Introduction

Seeding a working fluid with small particles that behave as passive tracers and tracking the resulting motion of those particles by way of digital image processing continues to be a valuable tool for noninvasively extracting velocity data from a flow of interest. A broad class of techniques that generally fall under the categories of particle image velocimetry (PIV) [1, 34, 52, 50] or particle tracking velocimetry (PTV) [23, 28, 29] are typically constructed such that a thin sheet within the flow is illuminated via laser or other light source and imaged with one or more calibrated cameras [18, 37, 44, 48, 49, 51] as depicted in Figure 3.1. If the same illuminated sheet is repeatedly imaged through time, then suitably designed image processing software can track the motion of particles within the sheet. An estimate of the 2D planar velocity \mathbf{u} is defined as

$$\mathbf{u} = \frac{\Delta \mathbf{x}}{\Delta t}, \quad (3.1)$$

where $\Delta \mathbf{x}$ is the distance a particle is observed to move between two images taken in a time interval Δt .

Whether the motion of individual particles or that of clusters is tracked forms the primary distinguishing criteria between PTV methods and those of PIV, respectively




Figure 3.1: Motion-controlled stereoscopic PIV system used to investigate thermal plumes. The working fluid is contained in the inner tank with a well-controlled isothermal boundary condition provided by a water bath in the outer tank. The light sheet is produced using halogen white light sources and a fiber optic line generator.

(see Figure 3.2). The motivation for using a cluster of particles in PIV techniques can best be summarized by the notion that a cluster of particles forming a spatial pattern contains more information than a single particle. Just as in astronomy, locating a cluster of passive tracers (or constellation of stars) can often be done more robustly than a single particle (or star) simply because more constraints are being satisfied by way of spatial relationships when identifying the cluster.

Specifically, correlation-type PIV algorithms function by considering at least two image frames, Frame 1 and Frame 2, taken at different times. Generally Frame 1 is divided into an array of discrete blocks or ‘windows’, typically 16×16 or 32×32 pixels, with each window containing the images of several particles (ref. Figure 3.2(a)). Then a chosen image registration technique (*e.g.*, cross-correlation) is used to locate the windows from Frame 1 in Frame 2 [6, 14, 21, 24]. The relative displacement in pixels required to translate a particular window in Frame 1 to its location in Frame 2 then forms, by way of the camera calibration, the spatial displacement $\Delta \mathbf{x}$ in Eq.

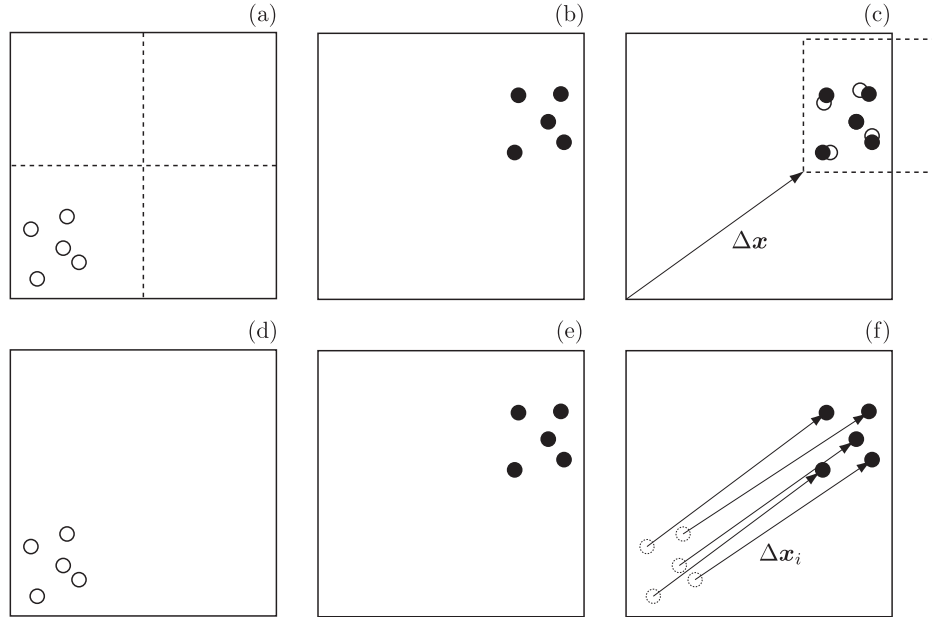


Figure 3.2: Comparison of PIV and PTV methods. Initial particle position is depicted by open circles, while final location is denoted by closed circles. (a) Frame 1 divided into 4 windows for processing with PIV. One window contains 5 particles. (b) Frame 2 with the location of the same particles shown after some elapsed time. (c) Best alignment of the tracer pattern in lower left window of Frame 1 with particles in Frame 2 as determined by an image registration algorithm. The window displacement distance is given by $\Delta \mathbf{x}$. (d, e) Same as (a, b), but for processing with PTV. (f) PTV provides the displacement for each individual particle, $\Delta \mathbf{x}_i$. The initial location of the particles has been repeated (dotted circles) for clarity.

3.1 (ref. Figure 3.2(c)). While using windows containing clusters of particles can lead to robust estimates of fluid velocities using a correlation-type PIV technique, the approach does have some drawbacks.

The resolution of the extracted velocity field is directly dependent on the size of the interrogation window. As the window size is decreased, the number of velocity vectors extracted increases, however smaller windows contain less particles which can lead to erroneous estimates of the velocity. If the seed particles are reasonably uniform in size, the image of each particle looks nearly identical to all the others as demonstrated in Figure 3.3. Therefore without additional constraints, a window from




Figure 3.3: Sample image of tracers illuminated using white light. The tracers are Hallcrest NSL33R25C15W100 thermochromic liquid crystals. Note the similarity in appearance of the individual tracers.

Frame 1 containing just one or even no particles will produce ambiguous results, since it will register perfectly atop any other tracer (or void if the window contained no particles) in Frame 2. Accordingly, the PIV window size must be sufficiently large to contain multiple tracers and the resolution improved, if desired, using overlapping windows or some other technique [14].

Unfortunately, most flows of interest are not characterized by a spatially uniform velocity field, therefore there are upper bounds on the size of the interrogation window as well. If the window is too large, spatial gradients in the velocity field can deform the pattern of particles such that the pattern in Frame 1 can no longer be registered by a simple translation to Frame 2. This scenario is depicted in Figure 3.4. To successfully employ a PIV approach in this case, one must reduce the size of the window or employ a registration technique that considers higher order deformations [13, 19, 36, 43].

In contrast to PIV, particle tracking velocimetry attempts to resolve and follow each individual particle through a sequence of images (ref. Figures 3.2(d-f)). Hence, in principal at least, PTV methods offer the potential for extremely high resolution velocity fields and immunity to the problems associated with particle pattern deformation under strong velocity gradients. Like PIV, PTV functions with at least two




Figure 3.4: Application of PIV to a rotational flow. Format follows Figures 3.2(a-c). (a) Frame 1 with an excessively large, frame-sized PIV window. (b) Frame 2 showing tracers after some time interval. (c) The flow has rotated the initial pattern of tracers to such a degree that poor results are obtained from image registration algorithms that consider only rigid body translations.

image frames, although some approaches (*e.g.*, [17, 29]) require more. To employ PTV, two separate problems must be solved. First, and representing perhaps the greatest challenge, each of the particles in Frame 1 must be matched with the corresponding partner in later frames – this is the so-called correspondence problem in computer vision literature [11, 38]. The difficulty associated with matching tracer pairs is the very reason why PIV methods have a lower limit on the useful window size. As before, the images of individual tracer particles used in PIV/PTV are almost identical. Therefore when presented with two or more image frames taken some time interval apart, identifying the same particle in all frames generally requires some additional information (*e.g.*, the shape a cluster of nearby particles makes). Nevertheless, once particle pairs have been identified, then the second piece of the PTV puzzle can be tackled: estimating the displacement field, or warp, that maps the particles in Frame 1 onto their locations in subsequent frames.

While solving either of the two PTV problems in isolation is nearly impossible, iteratively solving for both simultaneously has generally proved very productive [7, 15, 33, 35]. In the broader computer vision context, such schemes usually proceed

by forming some sort of metric that is minimized to indicate the best match between an object or feature point in Frame 1 and one in Frame 2. Several such metrics have been proposed that measure similarity in appearance [8, 12], preservation of some gross spatial relationship [4, 31, 32], or adherence to prescribed limits on motion [20, 17]. Once an estimate of the correspondence is available, an approximation to the warp can be computed and the correspondence subsequently updated. The methods generally proceed in an iterative fashion until a specified convergence threshold has been reached.

Overall, PIV-type schemes often offer superior accuracy and computational efficiency, at the expense of difficulty handling strong gradients and somewhat reduced resolution in extracted velocity vectors. PTV, on the other hand, can provide improved resolution and more easily accommodate strong velocity gradients, but the method can be computationally expensive due to the iterative nature of simultaneously estimating the point correspondences and warp. Clearly some sort of hybrid technique that leverages the strengths of each method has much to offer, and several such approaches have been investigated in the literature [3, 9, 25, 41, 39]. While the methods referenced are more focused on the use of PTV as a means to improve the resolution of the velocity field following application of conventional PIV, here we present a hybrid PIV/PTV algorithm developed specifically for the purpose of computing velocity vectors in those regions of the flow field where PIV has failed.

The advantages of the presented technique are three-fold. First the hybrid technique relaxes the upper limit on PIV window size from a particle deformation perspective. Hence the user can select a window size with more emphasis on balancing the trade off between a window containing a sufficiently large number of tracers and the desired spatial resolution of the extracted velocity field. Second, the scheme permits better utilization of information available in the raw images. While spurious velocity vectors can be removed and replaced with some statistical measure of the

vectors in a small neighborhood [46, 47], such approaches do not necessarily take full advantage of the available data. Furthermore, if several spurious vectors are extracted in a connected neighborhood, then this group of erroneous vectors can skew a neighborhood average or median value. The presented hybrid PIV/PTV scheme can often construct a reasonable estimate of the offending vector(s) prior to later processing with spurious vector removal tools, thereby improving the fidelity of the extracted velocity field. And finally, the PTV portion of scheme only requires two frames from each camera for estimation of the velocity vectors.

In Section 3.2, the formulation of the hybrid scheme is detailed. Results from application of the method are presented in Section 3.3, and some final remarks are provided in Section 3.4.

3.2 The hybrid PIV/PTV scheme

An overview of the hybrid PIV/PTV scheme is presented in Figure 3.5. The method begins by computing flow vectors with traditional PIV using image registration metrics that provide a normalized measure for goodness of fit. Two examples of such metrics are normalized cross-correlation (NCC) [6, 40] and the similarity measure of Barnea and Silverman’s sequential similarity detection algorithm (SSDA) [2].¹ In both cases, the value of the metric extremum indicates the degree to which a perfect match has been found for the PIV window of Frame 1 within the search region of Frame 2. Consequently, by treating the registration metric as a confidence measure and specifying a suitable threshold, one can flag those extracted PIV vectors that are likely of poor quality. With normalized cross-correlation, for example, the metric value is constrained to vary between 0 and 1, with 1 representing 100% confidence in the match. Spurious vectors would then be associated with low values of the NCC metric, and an appropriate threshold would be selected based on sample images such

¹For more on NCC and SSDA, refer to Chapter II.

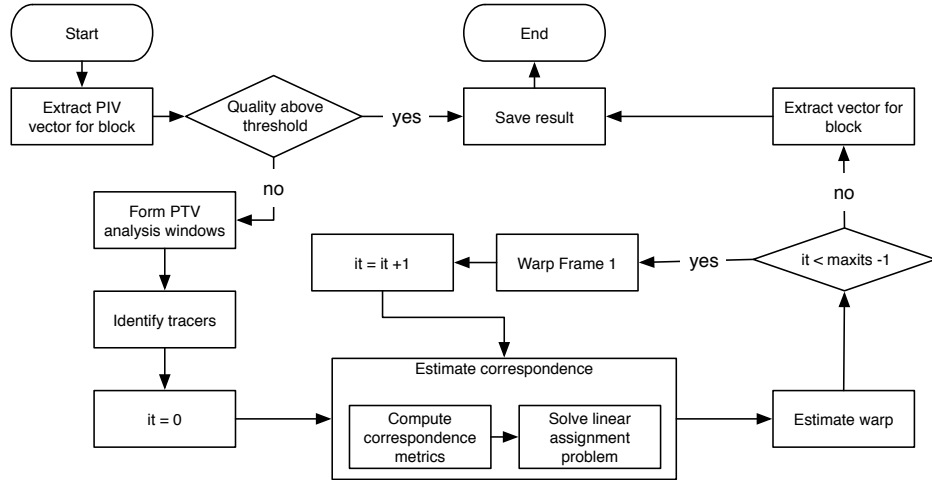


Figure 3.5: Flow chart for the hybrid PIV/PTV scheme. The variable `it` represents a counter for a loop that terminates after `maxits` iterations.

that only problematic vectors are further processed using the hybrid scheme.

If the registration metric for a given PIV window is above a specified threshold, the displacement estimate for the window is stored as a valid PIV vector and the hybrid scheme terminates for that block of Frame 1. If the registration metric is below the threshold, then the PTV procedure is launched by first extracting a sufficiently large sized region of both frames. These PTV windows are aligned between the frames (*i.e.*, the same region in both frames is extracted) and are of dimensions no less than that of the original PIV window plus twice the maximum expected tracer displacement between frames, as depicted in Figure 3.6. Ensuring that the PTV windows are large enough to capture the tracers of the original Frame 1 PIV block in their final locations in Frame 2 is critical. However the window size should not be so large that an excessive number of tracers must be matched.²

Once the PTV analysis region has been extracted from both frames, the tracer particles within must be identified. The objective of this phase is to locate the center

²Keep in mind that the objective of using PTV for the hybrid scheme is to compute a better estimate of the velocity vector for a single PIV window where PIV itself has failed. In the hybrid scheme presented here, the intention is not to apply PTV to the entire image frame.




Figure 3.6: A PTV analysis region (dash-dotted line) constructed around a smaller PIV window (dashed line). The maximum expected particle displacement in the time interval between frames is Δx_{\max} .

of mass of each particle. Several schemes for identifying tracers have been developed specifically for PTV [9, 16, 33, 39], however the two step peak finding algorithm of Kierkegaard [26] has proven very effective in the authors' tests. We employ the method to identify particles by first dividing the PTV analysis region from each frame into a number of small blocks. The block size is chosen to be equal to the typical size of a particle in the images. For our images, a block of 3×3 pixels works well. The maximum intensity value within each block is then retained as a candidate particle centroid provided the corresponding pixel has an intensity that lies above a specified threshold, κ , typically set to the brightest decile of all pixels in the PTV analysis region. During the second phase of the peak finding procedure, a window of size equal to the small blocks (*e.g.*, 3×3 pixels in our case) is centered on top of each candidate particle's location. If the candidate centroid from the prior step remains the pixel having the highest intensity within the second window, then the candidate particle position is identified as a valid tracer centroid. Note that the particle identification scheme as described only locates a tracer with a resolution of one pixel. If the location of the particle's centroid is desired with sub-pixel resolution, then a fit to a Gaussian or other function can be performed following application of the peak finding procedure [52].

One additional note regarding the extraction of tracer particles is in order. Gen-

erally speaking, the remaining steps for computing tracer correspondences are quite sensitive to the quality of the tracer extraction algorithm. Namely, correspondence estimation is dependent on the degree to which the PTV analysis regions from both frames contain the same particles. This sensitivity of the correspondence problem to particle extraction should be intuitively clear. After all, no correspondence algorithm can find a correct match for which none exists. While the PTV scheme described here explicitly compensates for outliers (*i.e.*, particles in one frame that have no match in the other), such provisions are an added complication to an already difficult problem, and occasional misclassification is bound to occur. For this reason, any effort spent tuning the particle extraction algorithm has a direct bearing on the overall quality of the velocity vector estimated using the PTV scheme.

With the set of tracers from the PTV analysis regions of both frames identified, the particle correspondences can be computed. As discussed in the Introduction, correspondence estimation is typically based on the minimization of a similarity measure. The approach taken by the hybrid scheme is to minimize the overall cost for the set of all pairings using cost metrics that measure how well a particle from Frame 2 appears to be the same particle from Frame 1. This optimized one-to-one matching is commonly referred to as the linear assignment problem. In the context of PTV, the linear assignment problem is formulated as follows. Consider two sets of points A and B , with the set A taken from the PTV analysis region of Frame 1 and set B taken from the analysis region of Frame 2. For each particle in A there is at most one corresponding partner in B . If A and B are of different size, then some particles in the larger set are clearly outliers (*i.e.*, these outliers have no matching particle in the other set). Using an appropriate metric (or combination of metrics), a cost $C_{ab} \equiv C(a, b)$ can be constructed for pairing particle a from set A with tracer b from set B . The set of all possible pairing costs then forms a 2D matrix \mathbf{C} that can be made square by appending rows or columns of zeros to represent outliers. The ob-




Figure 3.7: Diagram of a one-to-one pairing Φ for a linear correspondence problem between two sets of particles A and B . Outliers are shown as solid circles.

jective of the linear assignment problem is then to choose the optimal pairing Φ such that the total cost C_T ,

$$C_T(\Phi) = \sum_{a \in A} C(a, \Phi(a)), \quad (3.2)$$

is minimized. The problem is depicted in Figure 3.7 for reference. Note that several techniques for finding Φ given \mathbf{C} are available [22, 27, 30].

The remaining piece of correspondence estimation is then selection of suitable cost metrics, several of which were presented in the Introduction. In the PTV scheme detailed here, two individual measures are actually employed: the shape context of Belongie *et al.* [4] and the spring model of Okamoto *et al.* [32]. Both of these metrics will be discussed briefly before considering how they are utilized together.

For a given tracer in a collection of tracers, the shape context provides a coarse, 2D, log-polar histogram of the location of all other particles in the collection relative to the selected tracer. Again, let a represent a particle in the PTV analysis region of Frame 1, and $\{\alpha_i\}_{i \in (1,n)}$ be the n neighbors of a . The location of a neighbor with respect to a is then

$$\mathbf{r}_i = \mathbf{x}_{\alpha_i} - \mathbf{x}_a, \quad (3.3)$$

and the azimuth of the neighbor becomes

$$\theta_i = \arctan \frac{r_{i,y}}{r_{i,x}}. \quad (3.4)$$

By constructing a 2D histogram with axes spanned by $\log r$ and θ , a statistical expression of the spatial distribution of a 's neighbors is formed, as shown in Figure 3.8.

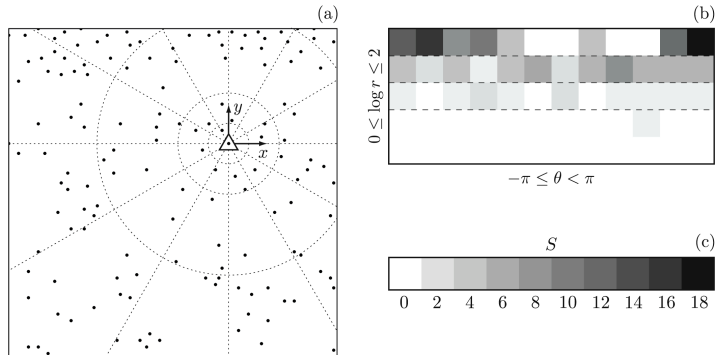


Figure 3.8: Sample shape context (b) for the particle inside the triangle (a). (c) Particle count legend for (b). The shape context bin boundaries are shown as dotted lines in (a). The $\log r$ bin boundaries visible in (a) are also depicted as dotted lines in (b).

If the shape context for particle a from Frame 1 is represented by S_a , and the shape context for particle b from Frame 2 given by S_b , then the pairing cost C_{ab} is taken as

$$C_{ab} = \frac{1}{2} \sum_{k=1}^K \frac{(S_a(k) - S_b(k))^2}{S_a(k) + S_b(k)}, \quad (3.5)$$

where k iterates over the K bins of the shape context. Note that Eq. 3.5 is simply the χ^2 test for S_b against an expected distribution of $(S_a + S_b)/2$ [42].

Clearly, shape contexts can also be constructed for the particles of Frame 2. The implicit assumption for PTV use of the shape context is that while the particle pattern can deform during the time interval between frames, some gross spatial relationship between the particles is likely preserved. Therefore when selecting the bin sizes for r and θ , the user must take care to ensure the bins are large enough to compensate for adequate particle deformation. The authors' testing concurs with the recommendation of Belongie *et al.* [4] that 5 bins for $\log r$ and 12 bins for θ are good starting points.

Due to the necessary coarseness of the histogram bins used for the shape context, the metric used in isolation often permits spurious matches of tracers within the same

bin. To compensate, an adaptation of the spring model cost metric of Okamoto *et al.* [32] is also used in the computation of particle correspondences by the present PTV scheme. Once again, consider a particle a from the PTV analysis region of Frame 1 and its n neighbors, $\{\alpha_i\}_{i \in (1,n)}$, that are now connected to a via virtual springs (ref. Figure 3.9(a)). The notion behind the spring model is that while a cluster of particles will move and deform under the influence of the fluid flow, they should do so together. Each spring has a force constant

$$k_i = \frac{1}{r_i^{(1)}}, \quad (3.6)$$

where $r_i^{(1)}$ is the distance given by the magnitude of Eq. 3.3. The superscript is used merely to emphasize that the relative distance, and hence the spring constant, is being computed for particles as arranged in Frame 1. Accordingly, the closest neighbors of particle a in Frame 1 have the stiffest springs. If the relative distance between particle a and neighbor α_i changes in moving from Frame 1 to Frame 2, then the virtual spring will generate a force

$$f_i = k_i \left| r_i^{(2)} - r_i^{(1)} \right| \quad (3.7)$$

on a . Assuming the location of a 's neighbors are known exactly at the time corresponding to Frame 2, the spring model provides a pairing cost

$$C_{ab} = \sum_{i=1}^n k_i \left| \rho_i - r_i^{(1)} \right|, \quad (3.8)$$

where ρ_i is the distance from the i -th neighbor of a with coordinates evaluated at the time of Frame 2 to particle b in Frame 2 (ref. Figure 3.9(b)). Consequently, the spring model considers the best match to particle a as the particle from Frame 2 with the lowest total force from all a 's neighbors (in their positions corresponding to the time of Frame 2).




Figure 3.9: Spring model. (a) Particle a from Frame 1 and its four neighbors. (b) Particle b from Frame 2 (solid circle) and a 's neighbors in their warp estimated, Frame 2 configuration (grey circles). Superscripts are used to emphasize the temporal configuration of the α_i .

Note that via ρ_i , the formulation of Eq. 3.8, however, depends on information that is not truly known – the exact location of a 's neighbors at the time of Frame 2. In practice once the PTV scheme is initialized, ρ_i is computed using the latest estimate of the warp to map the initial coordinates of a 's neighbors to their estimated locations. For the first iteration of the particle correspondence/warp estimation loop, however, no estimate of the warp exists. Hence the spring model as formulated here is not self initializing. It is for this very reason that the use of shape contexts has been included in the scheme. Namely, the correspondence/warp estimation loop is initialized with shape contexts alone to provide the first estimate of particle pairing and a warp. Afterwards, the PTV scheme transitions to correspondence estimation utilizing the spring model.

Testing conducted with various synthetic and real PIV image sequences indicates that correspondence/warp estimation convergence is often accelerated by way of periodic restarting (*e.g.*, every 5-10 iterations) using a blend of the shape context and spring model costs. The blended cost metric is taken as

$$C_{ab} = C_{ab}^{SC} + 0.5 \frac{\max(C_{ab}^{SC})}{\max(C_{ab}^{SP})} C_{ab}^{SP}, \quad (3.9)$$

where C_{ab}^{SC} is the shape context pairing cost, C_{ab}^{SP} the spring model pairing cost, and

$\max(\cdot)$ is the maximum pairing cost of all pairs using the specified metric. The factor of 0.5 simply gives more weight to the shape context pairing cost and is otherwise arbitrary. Note that during a restart, Frame 1 shape contexts are computed with tracer coordinates that have been warped using the latest warp estimate. Recalculating the Frame 1 shape contexts based on the warp estimate permits the Frame 1 shape contexts to asymptotically converge to their Frame 2 counterparts.

With the metrics used for the correspondence problem in hand, attention can now turn to the formulation used for estimating the warp. As soon as an estimated pairing Φ is available, the associated warp that maps Frame 1 onto Frame 2 using the correspondences can be computed. As illustrated in Figure 3.5, this warp estimate serves two purposes. Ultimately the warp is used to compute an improved velocity vector for the PIV block that produced the low quality result. But within the iterative correspondence/warp estimation loop, the warp estimate is also used to deform Frame 1. Conceptually, the motivation for iteratively deforming Frame 1 is straightforward. Namely, as the warp estimate is improved and the deformed analysis region of Frame 1 begins to better resemble Frame 2, the quality of the correspondences should improve as well.

Following Chui and Rangarajan [7], thin plate splines (TPS) [5, 45] are used to formulate the warp function although other approaches may be utilized. Thin plate splines have the benefit of being naturally multivariate and hence readily adaptable to modeling the warp of a 2D image. The non-regularized thin plate spline is an interpolating function that minimizes the bending energy of an idealized, thin 2D plate subject to small displacements. Consider a horizontal flat plate, and let $\{v_l\}_{l \in (1,m)}$ represent m vertical displacements of the plate at locations $\{\mathbf{p}_l\}_{l \in (1,m)}$ in the xy -plane (ref. Figures 3.10(a,b)). The shape, $f(x, y)$, taken by the plate that minimizes the

bending energy function

$$E = \int_{\mathbb{R}^2} \left[\left(\frac{\partial^2 f}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 f}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 f}{\partial y^2} \right)^2 \right] dx dy \quad (3.10)$$

is given by

$$f(x, y) = a_1 + a_2 x + a_3 y + \sum_{l=1}^m w_l U(r_l), \quad (3.11)$$

with $U(r_l) = r_l^2 \log r_l$ and $r_l = |(x, y) - \mathbf{p}_l|$. The coefficients w_l are subject to three additional constraints

$$\sum_{l=1}^m w_l = \sum_{l=1}^m w_l x_l = \sum_{l=1}^m w_l y_l = 0 \quad (3.12)$$

which together with the requirement that

$$f(x_l, y_l) = v_l, \quad l = 1 \dots m \quad (3.13)$$

produces the linear system of equations

$$\mathbf{L} \begin{bmatrix} \mathbf{w} \\ \mathbf{a} \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ \mathbf{0} \end{bmatrix}, \quad (3.14)$$

where \mathbf{L} is an $(m+3) \times (m+3)$ block matrix

$$\mathbf{L} = \begin{bmatrix} \mathbf{K} + \lambda \mathbf{I} & \mathbf{P} \\ \mathbf{P}^T & \mathbf{0} \end{bmatrix}. \quad (3.15)$$

The system of Eq. 3.14 can be inverted for the w_l along with a_1 , a_2 , and a_3 . The elements of \mathbf{K} are provided by $k_{ij} = U(|(x_i, y_i) - (x_j, y_j)|)$, and the rows of \mathbf{P} are the \mathbf{p}_l expressed as the homogeneous coordinates $(1, x_i, y_i)$ for each row i . The positive scalar λ controls the amount of regularization (*i.e.*, smoothing) used, with $\lambda = 0$




Figure 3.10: Thin plate spline example. (a) A thin plate pinned at the four corners and displaced by an armature at 5 interior locations. The warped sheet is colored according to out of plane displacement. The armature is shown as a purple surface embedded in the displaced sheet, and also alone in the inset where it has been colored according to displacement. (b) Pseudo-color plot of plate displacement with selected contours shown in white. Purple circles identify the location of the armature points. (c) Superposition of an initially uniform grid (gray lines) that has been vertically deformed (black lines) using the TPS warp of (a,b). Positive out of plane displacements for the TPS have been equated to positive vertical displacement of the image.

resulting in an exact fit to the displacements v_l . Note that \mathbf{I} is the $m \times m$ identity matrix. For the warp estimate in the PTV problem, the v_l are set equal to the lateral displacements necessary to map the tracers of the Frame 1 analysis region onto their partners in Frame 2. Hence two TPS models must be constructed for the PTV scheme, one for the x and one for the y displacements. An example demonstrating the use of one TPS model to warp an image along the y -axis alone is presented in Figure 3.10(c).

Given that the initial particle pairing is likely to contain errors, using regularization to smooth the warp is critical for early iterations. Regularization relaxes the requirement that the TPS model must produce an exact fit to the particle displacements, so the interpolating constraint of Eq. 3.13 effectively becomes the looser requirement $f(x_l, y_l) \approx v_l$ [4, 45]. Without regularization, the warp is interpolating and will therefore deform the analysis region of Frame 1 such that the correspondence estimate is perfectly realized. Using an interpolating warp based on an immature ini-

tial pairing will grossly deform Frame 1, and the iterative method may consequently converge to a sub-optimal pairing. Hence it is advantageous to regularize initial estimates of the warp, and use this smoothed warp to gently deform the analysis region of Frame 1 toward the state of Frame 2. Following Chui and Rangarajan [7], the amount of regularization can then be reduced each iteration by a factor of γ , an annealing rate, to reflect the increasing maturity of the correspondence estimate. For the present PTV scheme, the regularization parameter is initially set to $\lambda = r_{max}^2$, with r_{max} taken as the largest dimension of the PTV analysis window. The annealing rate γ is limited to the range $0.8 \leq \gamma \leq 0.99$ to provide a sufficiently, but not overly, slow rate of convergence.

Once a certain number of iterations or some other convergence criteria has been reached, the estimation loop terminates, and the PTV estimate of the velocity vector for the PIV block under consideration can be computed. For the scheme detailed here, the final PTV warp is used to calculate a displacement vector for every pixel in the original PIV block. The mean displacement for the block is then returned by the PTV method as the replacement for the original low-quality PIV vector.

3.3 Application of the hybrid scheme

To develop an appreciation for the impact various aspects of the hybrid PIV/PTV scheme have on overall performance, the method is first applied to a manually constructed field of tracers. Finally the scheme will be applied to actual PIV data. Unless otherwise specified, the parameters specified in Table 3.1 have been used for all computations.

Consider the two synthetic images of 30 tracers shown in Figures 3.11(a, b) below which have been constructed to ensure that both frames contain the same particles. The tracers for Frame 2 have been generated by deforming Frame 1 horizontally using a sinusoidal warp with a vertical wavelength of 54 pixels and amplitude of 5 pixels.

Particle extraction	Intensity threshold percentile (κ)	90
	Pass 1 window size [pixels]	3×3
	Pass 2 window size [pixels]	3×3
Correspondence estimation	Loop iterations	50
	Loop restart interval [iterations]	5
Warp estimation	Starting regularization (λ)	$O(10^4)$
	Annealing rate (γ)	0.83

Table 3.1: Hybrid scheme baseline parameters.

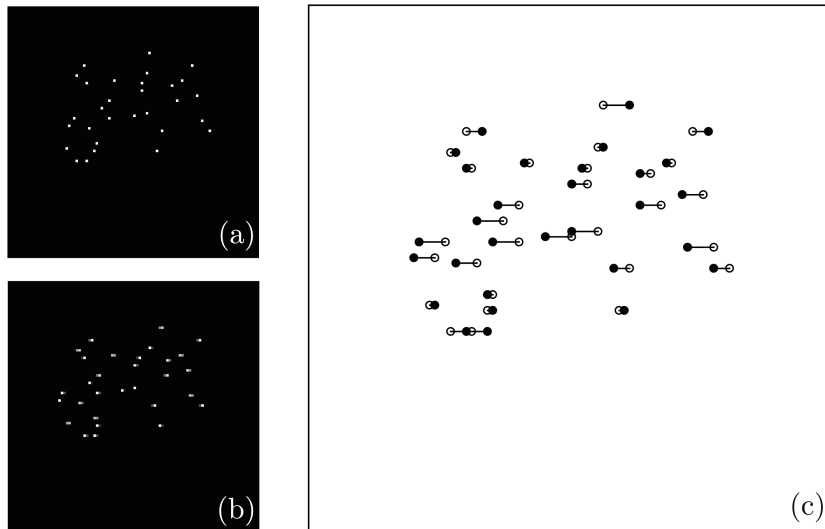


Figure 3.11: Synthetic 100×100 pixel PIV/PTV image sequence containing 30 tracers. (a) Frame 1. (b) Frame 2. (c) Known displacements. Particles in their initial locations corresponding to (a) are shown as open circles. Final particle locations from (b) are identified with solid circles.

The warped particles are depicted in Figure 3.11(b), and the known displacement field between frames is shown in Figure 3.11(c). The displacement field gradient along the vertical axis is of sufficient magnitude to cause the correlation type measures of both SSDA and normalized cross-correlation to fail. Consequently the scenario presents an opportunity to investigate the PTV scheme’s behavior using an idealized dataset.

Given that known pairings are available, an error measure can be constructed. Let a and b represent correctly paired particles from Frame 1 and Frame 2, respectively. Further let the position of particle a computed using the current warp estimate be

$\mathbf{x}_a^{(2)}$, and the known position of the Frame 2 particle be \mathbf{x}_b . The error for each individual pairing is then taken as

$$e_{a,b} = |\mathbf{x}_a^{(2)} - \mathbf{x}_b|. \quad (3.16)$$

The parameters of Table 3.1 have been chosen such that the warp regularization is reduced from an initial value of $\lambda \sim O(10^4)$ to $\lambda \sim O(1)$ by the end of the estimation loop. Therefore elevated RMS errors computed using Eq. 3.16 are indicative of erroneous particle pairings.

Figure 3.12 presents a convergence history using this RMS error for four different variations of the PTV scheme applied to the scenario of Figure 3.11. As shown, the convergence histories are essentially identical regardless of whether the standard set of parameters given in Table 3.1 are used, the estimation loop is never restarted using the blended cost of Eq. 3.9, or the loop is restarted every iteration. In these three cases, the PTV scheme perfectly matches the Frame 1 and Frame 2 particles with no errors. On the other hand, the fourth case which used the standard parameters of Table 3.1 but no regularization demonstrates the importance of warp regularization. After the first iteration, this latter example converges to a local extremum characterized by two mismatched pairings and an elevated RMS error. The final correspondence for the fourth case is shown in Figure 3.13 which contrasts with Figure 3.11(c).

As discussed in the previous section, one of the most important stages of the hybrid scheme's PTV segment is the extraction of valid tracers. When both PTV analysis regions contain the exact same set of tracers, the PTV scheme is incredibly robust as demonstrated above. However, if one or both analysis regions contain outliers having no real match then method performance can deteriorate quickly. Unfortunately, the presence of outliers in most PIV images is a virtual certainty. Some of these lone particles are real tracers moving into or out of the light sheet (or PTV analysis




Figure 3.12: RMS error in pixels for four PTV formulations applied to the image sequence of Figure 3.11. (+) Standard PTV scheme using the parameters of Table 3.1. (x) No blended cost restart. (□) Blended cost restart used for each iteration. (o) Standard PTV setup, but without regularization.




Figure 3.13: Final correspondences for PTV scheme without regularization. The two mismatched pairs are identified with a dashed box.

window) between frames. Other matchless ‘particles’ can be generated by the two step peak extraction method if the image signal to noise ratio is sufficiently low that noise is classified as a valid tracer. Regardless of the origin of such outliers, their essentially guaranteed presence necessitates an investigation of subsequent impact on the PTV scheme detailed here. To this end, the scenario of Figure 3.11 has been modified by introducing matchless particles into either Frame 2 alone or both frames.

When confined to a single frame, outliers result in an excessive number of particles in the affected analysis region thereby producing a rectangular cost matrix C . For the present PTV scheme, however, rectangular cost matrices are made square by appending rows or columns, as discussed in Section 3.2. Of course one must choose a cost with which to fill these added rows or columns. Given that a true match for an outlier doesn’t exist, pairing outliers with particles from the other frame will generally produce very high costs. Therefore, it is advantageous to make the cost matrix square using rows or columns filled with zeros. The linear assignment solution procedure will then be enticed to assign expensive, matchless particles to the zero cost outlier space. Figure 3.14 presents RMS error following PTV application to the previous synthetic image sequence which has been modified by adding 8 outliers to Frame 2. After the initial six iterations, the PTV scheme convergence plot mirrors that for the prior cases having no outliers, and once again particle pairings have been produced with no mismatches.

On the other hand, if the outliers are evenly distributed between frames, then the problem is much more complex. Suppose the synthetic dataset used above is once again modified such that four outliers are added to each frame as illustrated in Figure 3.15. In this scenario the cost matrix is already square. Without additional provisions, the correspondence estimation procedure will assign a match between all particles in the frames, known outlier or otherwise. The solution, then, is to again augment the cost matrix with rows and columns of zeros to provide a low cost route




Figure 3.14: RMS error (\circ) in pixels for the standard PTV formulation applied to the image sequence of Figure 3.11, but with 8 outliers added to Frame 2 (inset). The Frame 2 outliers are circled. The RMS error for the no outlier case of Figure 3.12 has been included for comparison and is denoted with (+).

for pairing outliers. Clearly in the present scenario of synthetic data, the number of auxiliary rows and columns needed to accommodate outliers is known precisely. Padding the cost matrix with four rows and four columns of zeros produces the results presented in Figure 3.16. Here the RMS error at the end of the estimation loop is nearly as low as that with no outliers, and only one incorrect pairing has been made in a region of the image sequence characterized by very small displacements.

In a real dataset processed without manual intervention, however, the true number




Figure 3.15: Synthetic image sequence of Figure 3.11 modified by adding four outliers to each frame (a,b).

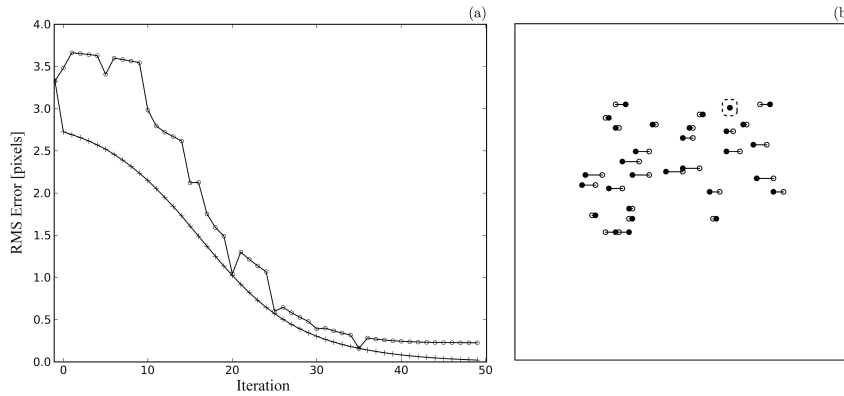


Figure 3.16: Results for the standard PTV formulation applied to the image sequence of Figure 3.15. (a) RMS error (o). The error trace for the case with no outliers has been included for comparison and is denoted with (+). (b) Final correspondences. The single mismatched pair is enclosed in the dashed box.

of outliers is unknown and must be estimated. If the assumed outlier count is too small, then the PTV scheme will be forced to make high cost, likely erroneous pairings. On the other hand, if the chosen outlier fraction is too large, then an excessive number of otherwise pairable particles will instead be diverted to outlier space. In either case, the quality of the final correspondence and warp suffers. Based on the authors' tests using actual PIV images, assuming that 20-30% of all particles are outliers produces the best results as demonstrated next. Consider the raw 100×100 pixel PTV analysis regions of Figure 3.17 which have also been manually processed to identify known pairings for error analysis. These images were extracted from a PIV dataset where the time interval between frames was excessively long permitting particles to move roughly 20 pixels between frames. Of the full set of particles identified by the two step peak extraction scheme, 102 tracers from Frame 1 were manually classified as having a known partner in Frame 2. In the ensuing discussion, a particle pairing will be identified as a mismatch if a particle in one frame is not paired with the correct manually identified partner – that is unless the particles in question are both paired to outlier space, as such assignments will not be classified as mismatches.




Figure 3.17: PTV analysis regions from an actual dataset. Particles identified by peak extraction algorithm are labeled. (a) Frame 1 region. (b) Frame 2 region.

Five separate cases have been considered using the data of Figure 3.17, with the allotted outlier fraction from each image varying from 10% to 50% in 10% increments. In all cases, the parameters of Table 3.1 have been used. The resulting RMS errors are presented in Figure 3.18 with mismatch counts captured in Table 3.2. Although the cases where the per-frame outlier fraction is greater than 10% perform similarly, the scenario accommodating 30% outlier particles produces the best results. The RMS error at loop termination for the 30% outlier fraction is below 2 pixels, but this value is far above the final error achieved with the idealized test cases presented earlier. Note that the PTV scheme correctly pairs all 102 known pairs with a final RMS error of 0.1 pixel if the outliers are manually removed from the set of particles. Consequently, the elevated final error and mismatches presented in Figure 3.18 and Table 3.2 are a direct consequence of the difficulty in automated identification and removal of outliers. Nevertheless, these results are typical and serve to illustrate the challenges associated with processing real data.

Before presenting an example of what the scheme can do to a full image with many windows, one final example using the data in Figure 3.17 is worthwhile. In the




Figure 3.18: RMS error for PTV analysis of Figure 3.17 using various outlier fractions.

Outlier fraction	Mismatches
10%	30
20%	9
30%	4
40%	3
50%	3

Table 3.2: Mismatch count for PTV analysis of Figure 3.17 using various outlier fractions.

Restart interval	Mismatches	β	Mismatches
1	7	0	4
5	4	0.25	4
10	24	0.5	4
∞	48	0.75	4
		1	4

Table 3.3: Mismatch counts corresponding to the analyses of Figure 3.19.

results of Figure 3.12 presented for the idealized dataset discussed earlier, there was no apparent benefit to using one cost metric over the other as long as regularization was used with the warp. For general datasets like that of Figure 3.17, however, such uniform behavior is atypical. Each metric has its own strengths and weaknesses when it comes to classifying matches, and the PTV scheme has been specifically constructed to leverage both. As discussed in Section 3.2, the estimation loop is first initialized using shape contexts alone, then transitions to a cost function employing only the spring model, but is periodically restarted using the blended cost function of Eq. 3.9. Consequently, two parameters must be chosen: the restart interval and the blending factor, β . To provide some guidance in choosing these parameters, RMS errors for several values of the restart interval and β are presented in Figure 3.19 below, with corresponding mismatch counts provided in Table 3.3. These data have been generated by sweeping one parameter and holding the other fixed. An outlier fraction of 30% was specified, and all other controls were as listed in Table 3.1. The results shown, typical of a wide range of tests, suggest a smaller restart interval provides lower errors, while choice of β is less critical. A reasonably good initial selection of parameters is to restart the estimation loop every 5 iterations using $\beta = 0.5$.

The merits of the hybrid PIV/PTV scheme become most apparent when applied to full images. As discussed in the Introduction, several methods exist for detection and removal of spurious vectors created when traditional correlation type PIV meth-

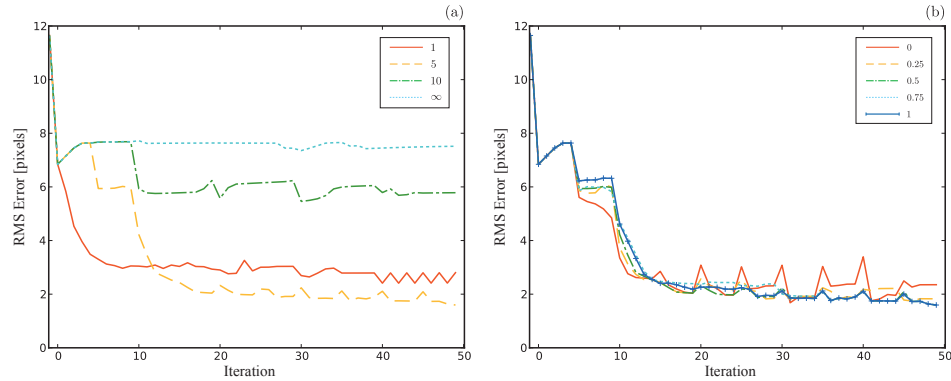


Figure 3.19: RMS error for PTV analysis of Figure 3.17. (a) Effect of various restart intervals, with ∞ indicating no restart. $\beta = 0.5$ was used for all 4 cases. (b) Effect of different β values using the restart interval of 5.

ods fail. Many of these techniques identify spurious vectors by comparing a vector under consideration to others in a small neighborhood. In the authors' experience, spurious vectors are not usually isolated, but instead tend to be produced in clusters. Consequently a scheme that simply compares vectors in a local region will have difficulty classifying a single vector as valid or erroneous if that vector is embedded in a cluster of erroneous values.

Consider the PIV frames shown in Figure 3.20 below. These images were taken using the left (bottom as shown in the figure) camera of the experimental setup shown in Figure 3.1 which was designed to study thermal convection in a fluid with temperature-dependent viscosity.³ The brighter tracers in Figure 3.20 are thermochromic liquid crystals which scatter different colored light based on temperature (ref. Chapter 1 and [10]). Near the heater in the darker bulbous area (bottom center of the image), the fluid temperature is beyond the working range of the liquid crystals, so the liquid crystals are transparent. In the region of hot fluid, the visible tracers are actually a powdered pigment added to the fluid specifically for the purpose of capturing the flow in this area. A number of issues associated with the image sequence make extraction of velocity vectors using PIV alone difficult. Upon cursory

³For more information on the experimental setup, see Chapter 1.

inspection of Figure 3.20, four large scale stationary structures near the bottom of the images are visible: two faint horizontal lines, and two vertical bright objects on either side of the dark upwelling fluid. These image artifacts are caused by light reflecting off various components of the experimental setup and do not represent actual structures influencing the flow. Consequently the reflections have an adverse effect on PIV performance. Other sources of difficulty for PIV application to the image sequence arise from the low tracer density in the hot upwelling region and an excessively long interval between frames.

PIV vectors extracted using the correlation-based PIV scheme of Chapter 1 are presented in Figure 3.21(a) for the flow field region directly over the heater. The PIV scheme was run using a 32×32 pixel window with 75% overlap. Note that prior to PIV processing, the raw image frames were filtered in the frequency domain using a notch filter to remove the parallel horizontal reflections. Even with the image pre-processing, several spurious velocity vectors are still generated in the region surrounding the lower reflected line, as illustrated in Figure 3.21(a). Given that these spurious vectors appear in clumps, post-processing using a local statistical measure will have difficulty in identifying and replacing erroneous values.

The hybrid PIV/PTV scheme, however, provides a marked improvement over correlation-type PIV alone. Setting the intensity threshold percentile $\kappa = 95$, but otherwise using the standard parameters of Table 3.1 produces the vector field of Figure 3.21(b). Close inspection of the vector field computed using the hybrid scheme reveals that it also includes spurious vectors, but generally not in the large clusters observed previously. Hence the hybrid results provide a final displacement field with higher fidelity to the underlying flow kinematics.




Figure 3.20: Two frame PIV image sequence.

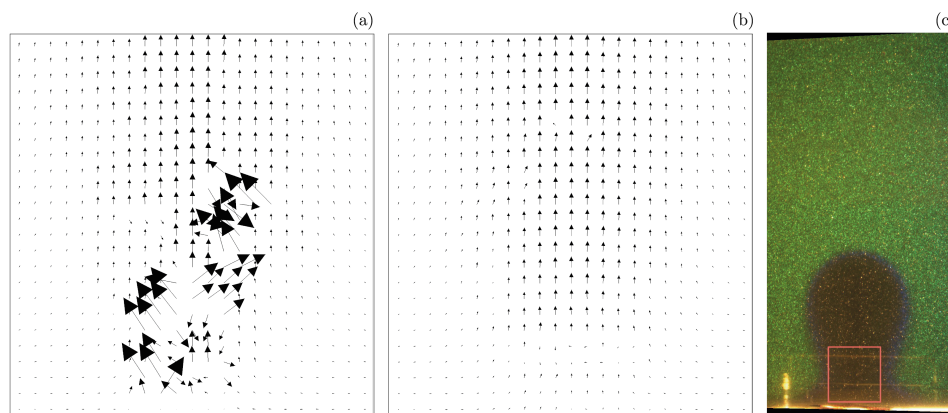


Figure 3.21: Displacement vectors from (a) traditional PIV analysis and (b) analysis with the hybrid PIV/PTV scheme. The vectors shown originate from the region of the flow field identified by the boxed region in (c).

3.4 Conclusions

A hybrid PIV/PTV scheme specifically designed to address those instance where PIV fails has been presented, and the method has been applied to several example scenarios. In many cases where the scheme finds applicability, the optimum solution to improving the quality of extracted velocity vectors is undoubtedly a recollection of the original data using better control over experimental parameters. Unfortunately, repeating experiments with the necessary modifications is not always feasible. In such circumstances, techniques like the hybrid PIV/PTV scheme permit better utilization of the available data.

While the hybrid scheme has been shown to be quite robust in situations where the exact same particles are present in both analysis windows, the method was found to have limitations arising from sensitivity to correct identification of particles and further classification of matchless outliers. Additional improvements are certainly possible in this regard. In particular, the current formulation of the warp estimate includes regularization to enforce smooth variation in the displacement field, however a similar mechanism is not employed as a correspondence cost measure. Penalizing pairings that produce local displacements which strongly depart from the latest warp estimate would likely provide a significant benefit to the identification of outliers. Furthermore, the quantity of outliers is currently preset at the start of the hybrid scheme to a fixed fraction of the total number of identified particles in each frame. Detailed analysis of the nature of pairing costs for valid matches and outlier particles will likely provide a route for automated estimation or refinement of this critical quantity.

3.5 References

- [1] R. J. Adrian. Particle-imaging techniques for experimental fluid mechanics. *Annual Review of Fluid Mechanics*, 23:261–304, 1991.

- [2] D. I. Barnea and H. F. Silverman. A class of algorithms for fast digital image registration. *IEEE Transactions on Computers*, C 21(2):179–186, 1972.
- [3] R. J. M. Bastiaans, G. A. J. van der Plas, and R. N. Kieft. The performance of a new PTV algorithm applied in super-resolution PIV. *Experiments in Fluids*, 32(3):346–356, 2002.
- [4] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(24):509–522, 2002.
- [5] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.
- [6] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, 1992.
- [7] H. L. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2-3):114–141, 2003.
- [8] D. Comaniciu, V. Ramesh, and P. Meer. Kernel-based object tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(5):564–577, 2003.
- [9] E. A. Cowen and S. G. Monismith. A hybrid digital particle tracking velocimetry technique. *Experiments in Fluids*, 22:199–211, 1997.
- [10] D. Dabiri and M. Gharib. Digital particle image thermometry: the method and implementation. *Experiments in Fluids*, 11:77–86, 1991.
- [11] U. R. Dhond and J. K. Aggarwal. Structure from stereo - a review. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(6):1489–1510, 1989.
- [12] P. Fieguth and D. Terzopoulos. Color-based tracking of heads and other mobile objects at video frame rates. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 21–27, 1997.
- [13] A. Fincham and G. Delerce. Advanced optimization of correlation imaging velocimetry algorithms. *Experiments in Fluids*, 29:S13–S22, 2000.
- [14] A. M. Fincham and G. R. Spedding. Low cost, high resolution DPIV for measurement of turbulent fluid flow. *Experiments in Fluids*, 23(6):449–462, 1997.
- [15] S. Gold, A. Rangarajan, C. P. Lu, S. Pappu, and E. Mjolsness. New algorithms for 2D and 3D point matching: Pose estimation and correspondence. *Pattern Recognition*, 31(8):1019–1031, 1998.

- [16] Y. G. Guezennec, R. S. Brodkey, N. Trigui, and J. C. Kent. Algorithms for fully automated three-dimensional particle tracking velocimetry. *Experiments in Fluids*, 17:209–219, 1994.
- [17] Y. A. Hassan and R. E. Canaan. Full-field bubbly flow velocity measurements using a multiframe particle tracking technique. *Experiments in Fluids*, 12(1-2):49–60, 1991.
- [18] J. Heikkila and O. Silven. A four-step camera calibration procedure with implicit image correction. In *1997 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1106–1112, 1997.
- [19] H. T. Huang, H. E. Fiedler, and J. J. Wang. Limitation and improvement of PIV part II: Particle image distortion, a novel technique. *Experiments in Fluids*, 15(4-5):263–273, 1993.
- [20] M. Ishikawa, Y. Murai, A. Wada, M. Iguchi, K. Okamoto, and F. Yamamoto. A novel algorithm for particle tracking velocimetry using the velocity gradient tensor. *Experiments in Fluids*, 29:519–531, 2000.
- [21] K. Jambunathan, X. Y. Ju, B. N. Dobbins, and S. Ashforth-Frost. An improved cross correlation technique for particle image velocimetry. *Measurement Science and Technology*, 6:507–514, 1995.
- [22] R. Jonker and A. Volgenant. A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing*, 38(4):325–340, 1987.
- [23] N. Kasagi and K. Nishino. Probing turbulence with three-dimensional particle-tracking velocimetry. *Experimental Thermal and Fluid Science*, 4:601–612, 1991.
- [24] R. D. Keane and R. J. Adrian. Theory of cross-correlation analysis of PIV images. *Applied Scientific Research*, 49:191–215, 1992.
- [25] R. D. Keane, R. J. Adrian, and Y. Zhang. Superresolution particle imaging velocimetry. *Measurement Science & Technology*, 6(6):754–768, 1995.
- [26] P. Kierkegaard. A method for detection of circular arcs based on the Hough transform. *Machine Vision and Applications*, 5:249–263, 1992.
- [27] H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–87, 1955.
- [28] H. G. Maas, A. Gruen, and D. Papantoniou. Particle tracking velocimetry in three-dimensional flows. Part 1. Photogrammetric determination of particle coordinates. *Experiments in Fluids*, 15:133–146, 1993.
- [29] N. A. Malik, T. Dracos, and D. Papantoniou. Particle tracking velocimetry in three-dimensional flows part II: Particle tracking. *Experiments in Fluids*, 15:279–294, 1993.