

ABSTRACT

Title of dissertation: SCHLIEREN SEQUENCE ANALYSIS
USING COMPUTER VISION

Nathanial Timothy Smith,
Doctor of Philosophy, 2013

Dissertation directed by: Professor Mark J. Lewis
Department of Aerospace Engineering

Computer vision-based methods are proposed for extraction and measurement of flow structures of interest in schlieren video. As schlieren data has increased with faster frame rates, we are faced with thousands of images to analyze. This presents an opportunity to study global flow structures over time that may not be evident from surface measurements. A degree of automation is desirable to extract flow structures and features to give information on their behavior through the sequence. Using an interdisciplinary approach, the analysis of large schlieren data is recast as a computer vision problem. The double-cone schlieren sequence is used as a testbed for the methodology; it is unique in that it contains 5,000 images, complex phenomena, and is feature rich.

Oblique structures such as shock waves and shear layers are common in schlieren images. A vision-based methodology is used to provide an estimate of oblique structure angles through the unsteady sequence. The methodology has been applied to a complex flowfield with multiple shocks. A converged detection success rate be-

tween 94% and 97% for these structures is obtained. The modified curvature scale space is used to define features at salient points on shock contours. A challenge in developing methods for feature extraction in schlieren images is the reconciliation of existing techniques with features of interest to an aerodynamicist. Domain-specific knowledge of physics must therefore be incorporated into the definition and detection phases. Known location and physically possible structure representations form a knowledge base that provides a unique feature definition and extraction. Model tip location and the motion of a shock intersection across several thousand frames are identified, localized, and tracked.

Images are parsed into physically meaningful labels using segmentation. Using this representation, it is shown that in the double-cone flowfield, the dominant unsteady motion is associated with large scale random events within the aft-cone bow shock. Small scale organized motion is associated with the shock-separated flow on the fore-cone surface. We show that computer vision is a natural and useful extension to the evaluation of schlieren data, and that segmentation has the potential to permit new large scale measurements of flow motion.

SCHLIEREN SEQUENCE ANALYSIS USING
COMPUTER VISION

by

Nathanial Timothy Smith

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2013

Advisory Committee:
Professor Mark Lewis, Chairman/Advisor
Professor Rama Chellappa
Associate Professor Christopher Cadou
Associate Professor Kenneth Yu
Professor Ashwani K. Gupta

© Copyright by
Nathaniel Timothy Smith
2013

Dedication

In memory of Dr. Robert H. Korkegi and Dr. Clark W. Hawk, whose influence on me helped to make this possible.

Acknowledgments

I would first like to thank my advisor Prof. Mark Lewis for his mentoring and guidance. Many seemingly esoteric areas of fluid dynamics were demystified beginning with a knock on his door and me asking “got a second?” I am grateful for his support through unexpected difficulties which went above and beyond the call of duty of an advisor.

I would like to thank Prof. Rama Chellappa of Electrical and Computer Engineering who took the time and effort to mentor and aerospace engineering student in the area of computer vision, who also became very acquainted with the phrase “got a second?”, and always responded “yes”.

Thanks to Joe Coblish, Michael Smith, Dan Marren and John Lafferty at Arnold Engineering Hypervelocity Wind Tunnel 9 for providing the data used in this research, and for their helpfulness whenever I had questions.

I would like to thank my committee, my officemates in JMP 2116 through the years, and the students at the Center For Automation Research, all from whom I have learned a great deal.

Finally, I would like to thank my wife, Xue Wu (Ph.D. CS 2008) without whose support this would not have been possible.

Table of Contents

List of Tables	vii
List of Figures	viii
Nomenclature	xiii
1 Introduction	1
1.1 Motivation	1
1.2 Objectives and Scope	3
1.3 Organization of the Dissertation	6
1.4 Previous Work	6
1.4.1 Image Scanning Methods	7
1.4.2 Image Processing-Based Methods	9
1.4.3 Methods Towards Computer Vision	10
1.4.4 Tracking in Sequences	12
1.5 Concepts from Vision	14
1.5.1 Image Interpretation and Understanding	15
1.5.2 The Role of Noise	18
1.5.3 The Role of Scale	21
2 Double-Cone Flows and Schlieren Data	23
2.1 Introduction	23
2.2 The Double-Cone Flowfield and Previous Findings	24
2.3 Hypervelocity Wind Tunnel 9	28
2.4 Tunnel 9 Double-Cone Experiment	30
2.5 The Double-Cone High-Speed Schlieren Sequence	35
3 Flow Isolation Using Segmentation	39
3.1 Motivation	39
3.2 Segmentation	41
3.3 Region Preserving Diffusion-Based Filtering	42
3.4 Manual Segmentation of the Model	46
3.5 The Image as a Graph: Graph Terminology	49

3.6	Graph Cut Segmentation	52
3.6.1	Determining a Physically Meaningful Flow Partition	57
3.6.2	Flow Segmentation Algorithm	61
3.7	Segmentation Results	63
3.7.1	Segmentation Accuracy	67
3.7.2	The Importance of the Labeled Image	69
4	Extracting Oblique Structures and Outer Contours	71
4.1	Motivation	71
4.2	Edge Detection	72
4.2.1	Limitations and Noise	76
4.2.2	Representation of Shock Structure with Edge Maps	77
4.2.3	Variation Through the Sequence	81
4.3	Isolation of the Outer-Shock	84
4.3.1	Outer-Shock Motion History	86
4.3.2	Instances of False Shock Structure	87
4.4	Line Fitting	88
4.4.1	The Hough Transform	89
4.4.2	Grid Sensitivity Studies	91
4.5	Detecting Shock Wave and Contact Surface Angles	95
4.5.1	Scale Selection for Shock Representation	96
4.5.2	Knowledge Representation for Structure Definitions	97
4.5.3	Algorithm for Measuring Oblique Features	99
4.6	Results for the Shock and Contact Surface Angle	104
4.6.1	Analysis of Detected Failure Modes	107
4.6.2	Combining Evidence Across Grids and Convergence	108
5	Feature Tracking of Large Scale Interest Points	111
5.1	Motivation	111
5.2	Feature Concepts	112
5.2.1	Scale in Features	114
5.2.2	Tracking	116
5.2.3	Specific Challenges in Schlieren Data	117
5.3	Grayscale Features	118
5.3.1	Corner Detection Evaluation	120
5.3.2	Performance in The Regions of Interest	121
5.4	Curvature as a Feature	124
5.4.1	Curvature Evolution and Scale Space	126
5.4.2	The Curvature Scale Space Corner Detector	126
5.5	Contour Extraction and Processing	128
5.5.1	Contour and Curvature Evolution	129
5.5.2	Kernel Construction and Convolution	132
5.6	Algorithm Design	133
5.6.1	Large Scale Curvature as a Canonical Descriptor	135
5.6.2	Model Tip Identification	136

5.6.2.1	Localization	138
5.6.2.2	Detection Algorithm	139
5.6.3	Shock Intersection Identification	141
5.6.3.1	Intersection Curvature Primitives	143
5.6.3.2	Localization	147
5.6.3.3	Detection Algorithm	149
5.6.4	Handling of Missing Contours	152
5.7	Results	153
5.7.1	Model Detection and Motion	154
5.7.2	Shock Intersection Detection and Motion	156
6	Flowfield Motion Analysis	164
6.1	Introduction	164
6.2	Outer-Shock Motion Analysis	165
6.3	Full Flowfield Motion Analysis	168
6.4	Fore and Aft-Cone Flowfield Relations	171
6.4.1	Fore-Cone Region Motion Analysis	172
6.4.2	Aft-Cone Region Motion Analysis	175
6.4.3	Flowfield Area Correlations	176
6.5	Flowfield Spectral Content	180
6.5.1	The Full Flowfield	181
6.5.2	The Fore-Cone Flow	187
6.5.3	The Aft-Cone Flow	192
6.6	Flowfield Motion Summary	196
7	Conclusions and Future Work	198
7.1	Conclusions	198
7.2	Original Contributions	202
7.3	Directions for Future Work	203
A	Extensions of Salient Point Identification Methodology	206
B	Performance Estimate of Salient Point Detection	209
	Bibliography	212

List of Tables

3.1	Flowfield segment distributions for each graph-connectivity [pixles]. . .	69
4.1	Hough transform peak votes for corner ROI $\mathbf{I}_e(t_o + 1500)$	94
4.2	Test cases for shock angle β_s	105
4.3	Test cases for contact surface angle θ_{cs}	105
4.4	Shock failure modes.	107
4.5	Contact surface failure modes.	108
5.1	Distribution properties for the triple point location.	161
6.1	Full segment centroid fluctuation \mathbf{x}'_c distributions.	171
6.2	Flowfield segment area distributions for the full, fore-cone, and aft-cone [pixels].	172
6.3	Fore-cone segment centroid fluctuation \mathbf{x}'_c distributions.	174
6.4	Aft-cone segment centroid fluctuation \mathbf{x}'_c distributions.	175

List of Figures

1.1	Typical line scanning boundary detection method.	8
1.2	Line scanning to determine shock location	9
1.3	Boundary detection method using Sobel and morphology.	11
1.4	Method of determining shock locations using edge maps.	12
1.5	Oblique shock detection method.	13
1.6	Expert knowledge: Medical image interpretation.	16
1.7	Examples of illusory contours believed to be completed with prior knowledge of geometry.	17
1.8	Corruption of a step image with Gaussian noise and its derivatives.	20
1.9	Mixing layer scales visualized with the wavelet transform	22
2.1	Double-cone flowfield schematic.	25
2.2	Numerical and experimental pressure and heat transfer distributions on the double-cone surface from Nompelis.	27
2.3	Tunnel 9 operational envelope.	29
2.4	Tunnel 9 test section.	29
2.5	The double-cone schematic and model installed in the test section.	31
2.6	Unsteady CFD results from Coblish <i>et al.</i>	34
2.7	Time averaged computational results from Coblish <i>et al.</i>	35
2.8	Two consecutive images and their absolute difference.	38
3.1	Consecutive images from the sequence.	40
3.2	Schematic of image interpretation.	40
3.3	Bilateral filter response on a step edge.	45
3.4	Effect of the bilateral filter.	47
3.5	Non-region preserving Gaussian filtering for image $\mathbf{I}(t_o + 4960)$, $\sigma_D = 4$	48
3.6	Motion energy approximation image.	49
3.7	Example of a weighted undirected graph.	50
3.8	Graph partitioning illustrated on a cluster.	53
3.9	The normalized cut performed on a synthetic image.	54
3.10	Illustration of the minimum graph cut.	55
3.11	Illustration of inclusive and exclusive points of the freestream seg- mentation.	60

3.12	Evaluation of labeled Images.	64
3.13	Two images freestream segmentation removed.	65
3.14	Automatically detected segmentation failures.	66
3.15	Histogram distribution of flowfield segment area with graph connectivity.	68
3.16	Convergence of flowfield segment area with graph connectivity.	69
3.17	Image labeling scheme and segmentations for image $\mathbf{I}(t_o)$	70
4.1	Shock structure image understanding schematic.	72
4.2	Edge images \mathbf{I}_e at multiple scales σ_c with naive thresholding for image $\mathbf{I}(t_o)$	79
4.3	Edge images \mathbf{I}_e at multiple scales σ_c with thresholding $T_L = 0.2, T_H = 0.5$ for image $\mathbf{I}(t_o)$	80
4.4	Representations edge maps at constant scale $\sigma_c = 1$ with hysteresis through the run sequence.	83
4.5	Determining a usable contour from an edge map.	85
4.6	An unsuccessful contour isolation due to line breakage.	85
4.7	Motion history of isolated outer-shock structure	87
4.8	Contour images \mathbf{I}_c indicating false structure.	88
4.9	The Hough transform parameter space.	91
4.10	Corner ROI of image $\mathbf{I}(t_o + 1500)$	92
4.11	Hough transform accumulator grid sensitivities for corner ROI image $\mathbf{I}_e(t_o + 1500)$	94
4.12	Scale selection and raw primal sketches for Corner ROI $\mathbf{I}(t_o)$	97
4.13	Illustration of knowledge base incorporation.	98
4.14	Histograms of detected shock angle β_s across grid resolution.	106
4.15	Histograms of detected contact surface angle θ_{cs} across grid resolution.	106
4.16	Shock angle convergence across $\Delta\theta$ spacing.	110
4.17	Contact surface angle convergence across $\Delta\theta$ spacing.	110
5.1	A one-dimensional signal and its scale space representation.	116
5.2	Image $\mathbf{I}(t_o)$ showing all Harris corner responses.	121
5.3	Harris responses in the cone tip ROI.	122
5.4	Harris responses in the shock intersection ROI.	123
5.5	Isolated outer-shock contour showing feature locations and sign conventions.	134
5.6	Curvature κ for $\mathbf{I}_c(t_o)$ at two scales.	136
5.7	Curvature κ for $\mathbf{I}_c(t_o + 2158)$ at two scales.	136
5.8	Contour evolution near the cone tip for images $\mathbf{I}(t_o)$ and $\mathbf{I}(t_o + 1678)$	142
5.9	Shock intersection primitive structures in the raw image.	144
5.10	Shock intersection primitive structures in the contour image.	144
5.11	Shock intersection curvature primitives.	145
5.12	Unphysical intersection representation violating predicate logic.	147
5.13	Contour evolution near the shock intersection for images $\mathbf{I}(t_o)$ and $\mathbf{I}(t_o + 1678)$	151

5.14	Model tip localization results.	154
5.15	Row and column distributions of the model tip.	156
5.16	Two-dimensional histogram of the model tip.	157
5.17	Shock intersection localization results.	158
5.18	CSS localization of the contour primitives.	158
5.19	Failure instance undetected by the algorithm.	159
5.20	Row and column distributions of the shock intersection point.	161
5.21	Two-dimensional histogram of shock intersection motion.	163
5.22	Joint probability distribution of shock intersection motion.	163
6.1	Motion history of isolated outer-shock structure.	165
6.2	Outer-shock motion standard deviation.	166
6.3	Outer-shock motion kurtosis.	167
6.4	Cropped flowfield segment showing centroid location for $\mathbf{I}(t_o)$	168
6.5	Full flowfield segment centroid $\mathbf{x}'_{c_{FF}}$ fluctuation motion history. . . .	170
6.6	Full segment centroid motion overlaid on outer-shock motion history. . . .	170
6.7	Cropped flowfield segment split into fore and aft-cone regions by the detected shock intersect for $\mathbf{I}(t_o)$	172
6.8	Fore-cone flowfield segment centroid $\mathbf{x}'_{c_{FC}}$ fluctuation motion history. . . .	173
6.9	Fore-cone segment centroid motion overlaid on outer-shock motion history.	174
6.10	Aft-cone flowfield segment centroid $\mathbf{x}'_{c_{AC}}$ fluctuation motion history. . . .	176
6.11	Aft-cone segment centroid motion overlaid on outer-shock motion his- tory.	176
6.12	Standardized flowfield segment area change correlations.	179
6.13	Normalized autocorrelations for area and \mathbf{x}_c fluctuations.	183
6.14	Power spectral density of full segment area fluctuation A'_{FF}	185
6.15	Power spectral density of full segment row centroid fluctuation $x'_{c_{FF}}$	186
6.16	Power spectral density of full segment col. centroid fluctuation $x'_{c_{FF}}$	187
6.17	Normalized autocorrelations for fore-cone area and \mathbf{x}_c fluctuations. . . .	189
6.18	Power spectral density of fore-cone segment area fluctuation A'_{FC}	190
6.19	Power spectral density of fore-cone row centroid fluctuation $x'_{c_{FC}}$	191
6.20	Power spectral density of fore-cone col. centroid fluctuation $x'_{c_{FC}}$	192
6.21	Normalized autocorrelations for aft-cone area and \mathbf{x}_c fluctuations. . . .	193
6.22	Power spectral density of aft-cone segment area fluctuation A'_{AC}	194
6.23	Power spectral density of aft-cone row centroid fluctuation $x'_{c_{AC}}$	195
6.24	Power spectral density of aft-cone col. centroid fluctuation $x'_{c_{AC}}$	195
A.1	Waverider test image.	206
A.2	Waverider contour image.	207
A.3	Waverider contour curvature $\kappa(\sigma, s)$	208
A.4	Bow shock identification with modified CSS method. Fine scale CSS point labeled \triangle , final assignment in \square	208

B.1	Test shape and effects of noise corruption. Each image is a 32 x 32 pixel closeup.	210
B.2	Test shape and deformed isolated contours. Each image is a 32 x 32 pixel closeup.	210

List of Algorithms

1	Determining a physically meaningful segmentation.	62
2	Determination of separation shock and contact surface angles.	103
3	Identifying the cone tip for known contours.	140
4	Identifying the shock intersection for known contours.	150

Nomenclature

\otimes	convolution in one or two-dimensions
\in	inclusive
\forall	for all
$\{\emptyset\}$	the null set
A	flowfield area
CSS	Curvature Scale Space
\mathbf{D}	contour evolution distance
d	vector of graph node degrees
∂S	graph perimeter
\mathbf{E}	set of graph edges
e	graph edge indexed by node
ff	flowfield label
fs	freestream label
$\mathbf{G}(\mathbf{N}, \mathbf{E})$	graph
$g(\mathbf{x}, \sigma)$	Gaussian in one or two-dimensions
\mathbf{H}	Harris matrix
\mathbb{H}	Hough transform
h_G	isoperimetric constant
I	image intensity
\mathbf{I}	image
\mathbf{I}_{BF}	bilateral filtered image
\mathbf{I}_c	contour image
\mathbf{I}_{diff}	motion energy approximation image
\mathbf{I}_e	edge image
\mathbf{I}_L	labeled image
\mathbf{I}_σ	image smoothed at scale σ
k	set of line segments
L	graph Laplacian indexed by node
l	arbitrary label, or line segment
l_{min}	minimum line length
\mathbf{N}	set of graph nodes
\mathcal{N}	normal distribution
\mathcal{N}_p	image neighborhood about a pixel
n_i	i^{th} graph node
\mathbf{P}, \mathbf{Q}	sets of image points
\mathbf{p}, \mathbf{q}	image points
ROI	Region Of Interest
\mathbf{R}	Harris response matrix
\mathbf{R}_{xx}	autocorrelation function
\mathbb{R}	set of real numbers
r	normal distance from origin to line segment
r_p	Pearson's correlation coefficient

S	set
\bar{S}	set complement
s	arc length parameter
T	threshold
T_H	high threshold
T_L	low threshold
t	integer time index
t_{max}	end of image sequence
t_o	start of image sequence
Vol_s	graph volume
$W_{\mathbf{p}}$	normalizing factor for bilateral filter
X, Y	smoothed image coordinates
\mathbf{x}	image label vector
\mathbf{x}_c	flow segment centroid
x, y	image coordinate system
x_m, y_m	row and column coordinates of line segment midpoint
\mathbb{Z}	set of integers
α_c	graph cut value
α_I	graph intensity affinity weighting
β_s	separation shock angle
Γ	image contour
Δr	Hough transform distance discretization, pixels
$\Delta\theta$	Hough transform angle discretization, degrees
ϵ	convergence parameter, degrees
$\hat{\eta}$	noise estimate
θ	arbitrary angle, degrees
θ_{cs}	contact surface angle, degrees
θ_l	lower bound angle, degrees
θ_u	upper bound angle, degrees
κ	curvature
μ	statistical mean
$\frac{\mu_3}{\sigma^3}$	skewness
$\frac{\mu_4}{\sigma^4}$	kurtosis
σ	scale parameter
σ_c	Canny scale parameter
σ_D	image domain scale parameter for bilateral filter
σ_R	image range scale parameter for bilateral filter
σ_{std}	statistical standard deviation
τ	lag time
ω	graph weight matrix

Superscripts

' fluctuating component

Subscripts

AC aft-cone segment area

FC fore-cone segment area

FF full flowfield segment area

Chapter 1: Introduction

1.1 Motivation

The ability to capture schlieren data sets at higher frame rates has increased rapidly throughout the past decade. This study utilizes an optical data set sampled at 10 kHz continuously for 0.5 seconds. Studies have employed techniques that sample a small number of frames at in the MHz range [1–3]. Ultra high-speed cameras such as the Shimadzu HPV-2 camera are capable of imaging in this range for 100 consecutive frames. Other commercially available devices, for instance Vision Research’s Phantom v1610 camera, can acquire images at resolutions of 384 x 384 at a rate of 72.9 kHz, until the camera memory is full. Application of high-speed imaging techniques in wind tunnel facilities with long run times are producing increasingly larger data sets. These data may contain tens of thousands of images, a quantity that will only increase with continued scientific camera development. Large optical data sets offer the potential for insight into phenomena of the global flowfield that traditional surface measurements may not, particularly in unsteady aerodynamics.

Beyond human inference, there exist few methods to extract and quantify the information these data sets contain. Existing methods for interpretation of these sequences differ substantially, and require a high degree of supervision and

user interaction (some examples are provided in section 1.4). Most importantly, researchers are not currently maximizing the ability to understand all of the information in these data sets, because there is too much information to process. Current methods of analysis do not extract the richness of information that the data may contain. To fully interpret and understand these data, new analysis tools from outside the fluid mechanics community must be introduced and incorporated. Techniques developed for computer vision offer a framework for information extraction and maximizing data interpretation. This work proposes an interdisciplinary approach to the analysis and understanding of large schlieren data sets by recasting the sequence understanding as a computer vision problem.

With only a handful of images, there is no need to automate the measurement of flow structures. Once images of unsteady phenomena number in the thousands, manual data extraction methods are no longer viable. For large data sets, a degree of automation is desirable for analysis. Given a long sequence of images from which extraction and description of features is required, methods developed in computer vision can be useful for automating and assisting the inference process. Settles [4] suggests that with the advancement of image processing (computer vision), "...can we do 'character recognition' to reveal specific features of a schlieren object?" The present research is a step in this direction: the application of computer vision algorithms to a 5,000 frame schlieren sequence.

For this work, a schlieren data set with several thousand frames characterized by noise and small resolution, is used as a testbed for the methodology. The unsteady double-cone data sequence is used in this research for several reasons: the data set is

unique in that it consists of over 5,000 frames, the sequence contains multiple coupled complex flow phenomena, the images have contrast sufficient for many features allowing interpretation by the eye, and a high degree of unsteadiness is observed. These factors make this data set ideal for use as a testbed for the development and application of computer vision-based measurements for schlieren sequences.

1.2 Objectives and Scope

Computer vision techniques are proposed as a new means of analysis for understanding schlieren data. The approach here is an interdisciplinary effort merging the fields of Computer Vision and flow visualization analysis. Our main goal is to extend present analysis capabilities, establishing a vision-based analysis framework. This can remove non-repeatability associated with human analysis and to reduce human subjectivity of measurement. To accomplish these goals, the schlieren data analysis task is recast as a computer vision problem. It must be determined which methods and algorithms from computer vision are best suited for analysis of these data. Computer vision techniques must be adapted for their application in areas where image structure knowledge is required for interpretation (schlieren, radar and medical imaging, *etc.*) Computer vision *cannot* fully replace the expert. Computer vision or image understanding can, however, be used as a tool to quantitatively describe what a viewer can infer and observe from examining a sequence. Vision-based methods can assist the human, while providing automation and an engineering estimate that does not suffer from human measurement unrepeatability.

The first step of the analysis is to reduce the image content to physically meaningful representations. Flow image measurements are given in terms in pixels, yet the images are more usefully interpreted as objects or regions that correspond to large scale flow phenomena of interest. The process of converting images into labeled regions is termed segmentation. From these segmented representations, it is later investigated if the segments themselves that correspond directly to a measurement of flow unsteadiness, can provide new quantitative information.

Linear or near-linear structures in compressible flows include shock waves, contact surfaces, and shear layers are ubiquitous in many schlieren and shadowgraph sequences. The first objective of this work is to develop a vision-based approach for a robust and automated means to extract linear features in the sequence. The goal is a means of reliable extraction of unsteady angle measurements of linear shock waves and shear layers in the double-cone sequence. A rule-based classifier for extraction and measurement of oblique structure angles is constructed using vision methods that enforce physical knowledge of aerodynamics in the sequence. Known location and likely bounds are used to assist structure identification. In addition to the schlieren sequence in this research, the approach outlined here provides a framework to analyze common flows where unsteady shock, contact surface and shear layer motion are dominant features.

Curved and intersecting shock wave structures are also dominant features in schlieren and shadowgraph imagery. Interest points on shock structures include triple points, inflection points, and intersections whose motions can provide insight into unsteady flow mechanisms. For example, regions of inflection can identify the

maximum forward point on a curved or bow shock, which can help characterize standoff distance and shock strength. All of these are phenomena that can be identified in some way by changes or inflections of shock curvature. The second major portion of this work is the application of feature detection algorithms from computer vision to measure large scale “features” which can be defined by shock contours. The features of interest here are the motions of a sharp model tip, and a shock intersection point. These salient features are defined, localized, and tracked through the sequence using fused feature identification algorithms. A general feature detection framework is presented that, by design, captures structures of interest based on curvature properties that physically define shock features. A framework has been developed based on the Curvature Scale Space (CSS) representation that can be applied to a wide class of schlieren images when shock contours can be isolated.

This research focuses on large scale unsteady global flow structures as opposed to velocimetry or measurement of density *etc.* Each of the flow structures examined are large scale phenomena common in many schlieren data sets. Instances where linear shock, shear layer, bowshock or shock intersection identification are important factors to identify in unsteady flows include: the engine start and unstart processes, hysteresis in shock reflection patterns, shock-wave/boundary-layer interactions, *etc.* With this in mind, vision-based methods are introduced in a manner such that the algorithms may be generally applied to a wide range of problems, and to higher quality and resolution data sets.

1.3 Organization of the Dissertation

The remainder of this Chapter provides a review of current image analysis methods for schlieren and shadowgraph images. Concepts important in computer vision are also briefly discussed; elaboration and additional concepts that are task-specific are introduced where needed. A brief discussion of double-cone flows, the double-cone schlieren sequence, and a description of the facility where it was acquired are presented in Chapter 2. Previous findings of the flowfield are also discussed. Chapter 3 deals with pre-processing and labeling of the image sequence. Shock contour extraction and a means of measuring oblique structures are dealt with in Chapter 4. A method of providing shock motion histories of the data sequence is presented. Feature extraction is developed to identify, localize and track structures of interest on the shock contours in Chapter 5. Chapter 6 presents an analysis of information deduced from the flowfield using results from previous chapters. The dissertation concludes with a summary of the findings, a discussion of the impact of the methodologies applied and developed here, and suggestions for future work.

1.4 Previous Work

Schlieren imaging is typically used for qualitative analysis by the observer. Fluid structures are readily detected by the human eye as discontinuities in image gradients. The visualization of density gradients provides a means to interpret the flow physics that these images contain. This is an example of “expert knowledge”;

image content has no meaning without knowledge of the scene. Quantitative methods, on the other hand, typically focus on density and velocity profile information to understand smaller scale flow information, not global aspects such as shock motion. A general overview of these quantitative techniques is discussed in reference [4]. While schlieren and shadowgraphy are well-established techniques, advances to the techniques continue to occur (*e.g.* references [5] and [6]) in image acquisition and quantitative analysis, in part due to recent developments in modern cameras, optics, and computational capabilities.

1.4.1 Image Scanning Methods

Common methods to determine shock standoff distances, shock location, or shock angles involve line scanning of a row intensity trace across a location where a shock is known to exist [7–9]. Intensity traces of schlieren images are among the most common methods to determine shock information. Kastengren, Dutton and Elliott [7] developed a method to measure recompression shock unsteadiness in bluff body wakes, illustrated in Figure 1.1. To determine the recompression shock boundary in supersonic flow imaged using planar laser scattering, images were first divided into 16 vertical strips, 16 pixels in width each to reduce resolution. The pixel strips were averaged, resulting in 16 averaged image columns through the shock boundary (shown in Figure 1.1a). To determine the shock boundary location, each vertical strip was scanned for the largest difference in averaged intensity; this point, shown in Figure 1.1b, was defined as the shock location. Shock angles were determined

from linking four adjacent boundary points used to define slope. This method of down-sampling the image destroys, rather than preserves, image structure. The averaging filter similarly distorts image structure.

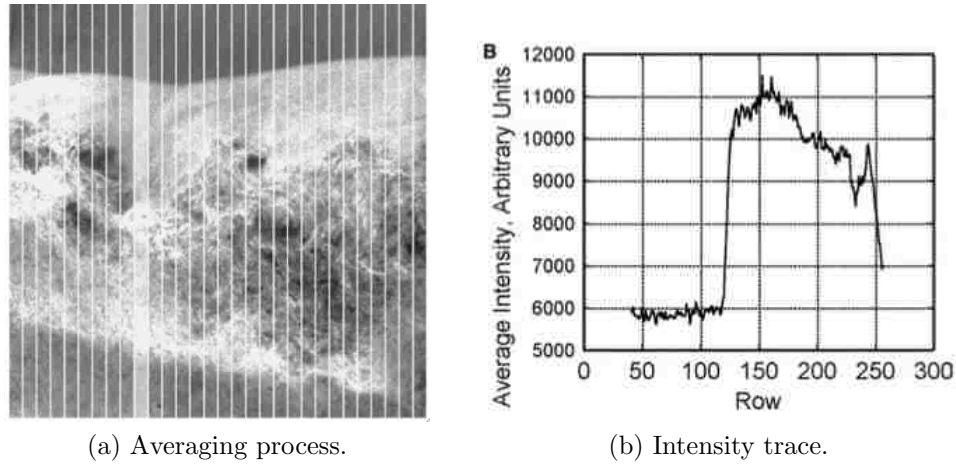


Figure 1.1: Typical line scanning boundary detection method from reference [7].

Bruce and Babinsky [8] studied the location of an oscillating normal shock in a duct. Shock location was determined, again, using a line scanning method. The line scanning method is illustrated in Figure 1.2, which shows the normal shock, a streamwise intensity trace, and the determined position throughout the test time. A similar technique was employed by Timmerman *et al.* [9] to determine shock position and angles in shadowgraphy images.

Line scanning methods, discussed above, are in fact a zeroth order approximation to edge detection, a computer vision technique. In such methods, locations of gradient change alone are considered, not flow features or definitions. Therefore, properties of the flow such as shocks or shear layers are not extracted. Intensity changes in schlieren images can be caused by a multitude of factors; line scanning methods therefore do not provide a firm definition of a fluid structure. These meth-

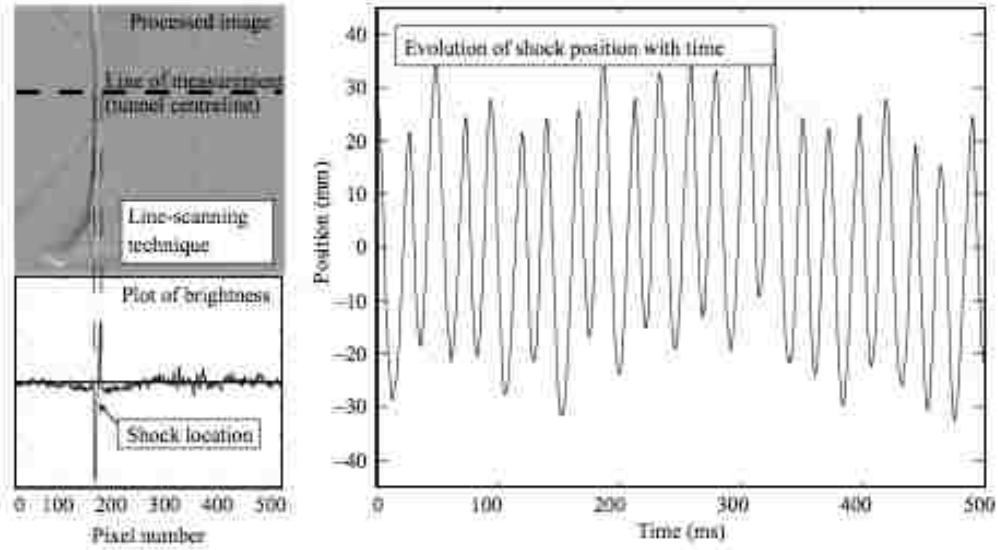


Figure 1.2: Line scanning to determine shock location from Bruce and Babinsky [8].

ods offer insight into the intensity variations of images, not the flow physics that the images contain, although the two are generally related. A tempting approach is to use edge maps, or intensity trace locations to measure shock angles by scanning image lines for two edge points, and calculating an angle. These methods may be faced with the following problems: incompleteness and streaking of edge maps, shock motion that pitches outside of the region being scanned, and the fact that the gross structure contained in an individual frame may change substantially within the sequence.


1.4.2 Image Processing-Based Methods

Beloki Perurena *et al.* determined water jet boundary penetration height and width in a Mach 6 crossflow by direct thresholding of the images [10]. Probability density functions of image intensities were used to infer information on mixing of the water jet with the freestream air. Kouchi *et al.* [11] developed a method to determine

the jet boundary in a supersonic cross flow based on the Sobel [12] operator. This method was recommended as a general boundary detection method, later used and further clarified in reference [13]. One hundred consecutive schlieren images were acquired at 250 kHz framing rate. Contrast adjustments were made to each image to standardize the intensity. “Fluctuation” images highlighted the boundary structure. Adjusted and fluctuating images are shown in Figure 1.3a. A gradient image was calculated using the Sobel operator, which was thresholded to create a binary image. Small objects corresponding to “noise” were removed. Morphology was used to fill the binary images to provide a jet boundary in each image. Figure 1.3b shows the binary image and the detected boundary overlaid on the schlieren image. The Sobel operator performs poorly in textured regions and is therefore an undesirable method for determining image gradients and boundaries in particular. Morphology is simply an operation that fills holes in a binary image by enforcing a structuring element; consideration of grouping and segmentation rules such as spatial proximity and object or gradient similarity are not modeled. Such methods therefore typically give poor segmentation results as they cannot guarantee accurate boundary detection.

1.4.3 Methods Towards Computer Vision

Estruch *et al.* [14] developed a method to measure shock unsteadiness using “background” subtraction (not to be confused with the computer vision term) of a tare image from the test sequence, followed by edge detection using the Canny [15] algorithm on a shock-wave/boundary-layer interaction schlieren sequence. The





(a) Standardized and fluctuating image.

(b) Binary image and boundary

Figure 1.3: Boundary detection method using Sobel and morphology operations from Kouchi *et al.* [11].

technique was later employed to study shock motion in front of a protuberance at high Mach numbers [16]. Point tracking on the edge map was used as a means of shock tracking, and a frequency analysis was performed directly on the edge map. Figure 1.4 shows the schlieren images with marked “tracking points” alongside the point motion history through the test time. Although this method gives a motion estimate, the direct correspondence of a shock to an edge point was assumed. Additionally, properties such as shock inclination angle were not extracted.

Efforts to develop software for wind tunnel schlieren image analysis at the High Enthalpy Shock Tunnel Göttingen [17] have recognized the need for automation and information extraction from large imaging data sets. Shock standoff distance from a cylinder, shock angle measurement, and model displacement tracking were performed using the Hough Transform [18,19] (a technique also used in the present research) and its variants through consecutive frames. The measured shock angle in a scramjet combustor is shown in Figure 1.5. The schlieren image (Figure 1.5a) is



(a) Schlieren image with marked tracking points.


(b) Shock location through sequence.

Figure 1.4: Method of determining shock locations using edge maps from reference [14].


converted into the gradient image shown in Figure 1.5b. The “inlet shock” (labeled in Figure 1.5c) is then fitted using the Hough transform. The results compared well against traditional pressure instrumentation measurements. However, the extraction required significant user involvement. This and the lack of a firm definition can contribute to feature subjectivity, and, without a classifier, feature extraction can be limited to specific data sequences.

1.4.4 Tracking in Sequences


While feature tracking methodologies are common in such applications as autonomous vehicle navigation, three-dimensional modeling from videos, and video alignment, there is little information found on using these techniques for schlieren image data. Optical flow [20] based methods have been applied for schlieren image analysis *e.g.* [21, 22], but are outside the scope of this research. Jonassen, Settles, and Tronosky [23] used commercial PIV (Particle Image Velocimetry) software and hardware to measure velocities in schlieren images where turbulent structures in the



(a) Schlieren image of scramjet model.



(b) Gradient image.



(c) Detected oblique shock.

Figure 1.5: Oblique Shock detection method from reference [17].

images were used in place of flow seeding particles. Hargather *et al.* [24] improved this method using custom software to perform correlations, accounting for size of turbulent structures for matching and an the addition of an LED light source. Velocity measurements were performed in a boundary-layer using focused and standard schlieren in addition to shadowgraphy. Average velocity profiles compared well to pitot surveys.

Laurence and Hornung [25] measured the displacement of models in images allowing force and moment data to be determined. A semi-automated routine selected edges detected by the Sobel operator which corresponded to the model in each frame. Simple geometries were fit using least squares; more complex geometries required a non-linear fitting with a cross-sectional geometry assumed. This gave trajectory information of the center of mass and pitching moment from which accelerations and forces were inferred. The technique was improved [26] by employing the Canny edge detector and sub-pixel localization. A more stringent fitting criterion was used for more complex geometries. Laurence [27] modified the technique to detect visible portions of edges in instances where model bodies are obscured or unknown using a method termed “edge-tracking”. These methods appear limited to solid bodies under rigid motion.

1.5 Concepts from Vision

The computer vision field focuses on the interpretation of images and video. Vision, or image understanding, deviates from image processing in that its goal is to extract features from an image or sequence to understand the phenomena contained in the image. Image understanding involves analysis, description and abstraction of information from images or sequences that are meaningful descriptions relevant to the scene, and is specific to analytical goals and task-dependent. Items crucial to the successful application of vision include: the means by which images are interpreted, the understanding process, the role of noise, and the concept of scale.

A final concern is the validity of computer vision for the interpretation of schlieren images. No matter how complex a technique may seem, it is largely based on two mathematical principles: gradient similarity, and spatial proximity. These are among the cues that assist experts in making inferences from schlieren data.

1.5.1 Image Interpretation and Understanding

Humans interpret images using visual evidence (*i.e.* image quality, intensity values, and spatial relations) and expert or domain knowledge, the latter playing a much more important role than the former. For example, schlieren, MRI, ultrasound, and many medical images have several common factors. These include low dynamic range or low SNR (Signal to Noise Ratio) images, illusory contours (those which are perceived rather than actual, as shown in Figure 1.7), weak or ambiguous structure boundaries, and required training or knowledge is necessary for their analyses. In such imagery, interpretations are based much less on image evidence and rely more heavily on scene knowledge; they are best left to interpretation by their respective experts.

Knowledge of how images are interpreted by the human vision system is paramount to implementation of any successful image understanding task. Image content has no meaning without knowledge of the scene. Perceptual grouping mechanisms of the human vision system seem to be based on factors which include proximity, similarity, symmetry, continuity, closure, and common fate, among other factors. Through use of these grouping mechanisms, known as Gestalt principles

(*e.g.* reference [28]), humans are able to establish and interpret scenes from images.

An MRI scan is shown in Figure 1.6 from reference [29]. Background is largely indiscernible from foreground. To the untrained eye, little more than noise is visible. However, trained technicians were able to make decisions to delineate a physical boundary based on expert knowledge. Two separate interpretations of the image are shown in Figures 1.6b and 1.6c. Information about pertinent features can be retrieved, although, such information is not always based on image information, but rely more on expert knowledge. This figure also illustrates the need for expert consensus in certain instances [29]. Engineers, without appropriate training, for example, would not be capable of this interpretation. Likewise, only with knowledge of fluid mechanics can a schlieren sequence be analyzed. This is the role knowledge and domain specific information play in bounding the problem.

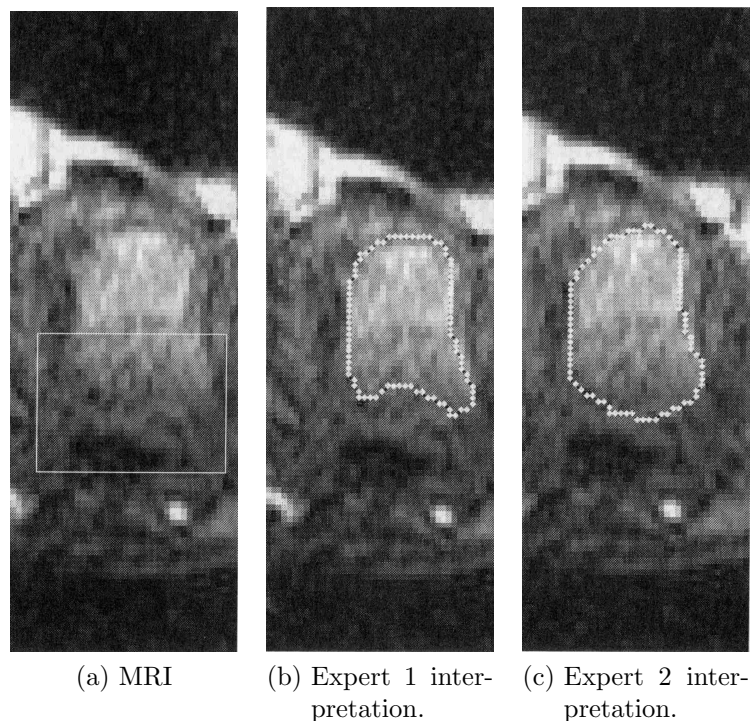


Figure 1.6: Expert knowledge: Medical image interpretation [29].

The role of prior knowledge is illustrated in the optical illusions in Figure 1.7. Both the triangle and the circle are observed and appear brighter than their surroundings, yet neither shape is present: the perceived contours are illusory. The contour completion, due to many factors, can be explained due to strong priors associated with knowledge of geometry. Shapes can be perceived due to evidence that suggests them. Similar phenomena can occur when interpreting schlieren images. The observer may see a completed shock wave, yet there is often not enough evidence in the image to support this. Here, the engineer’s prior knowledge consists of conservation laws. For example, a bow shock in front of a cylinder must be continuous, and is completed when viewing, although it may be broken upon closer inspection. This effect can limit what can be interpreted in an image sequence.

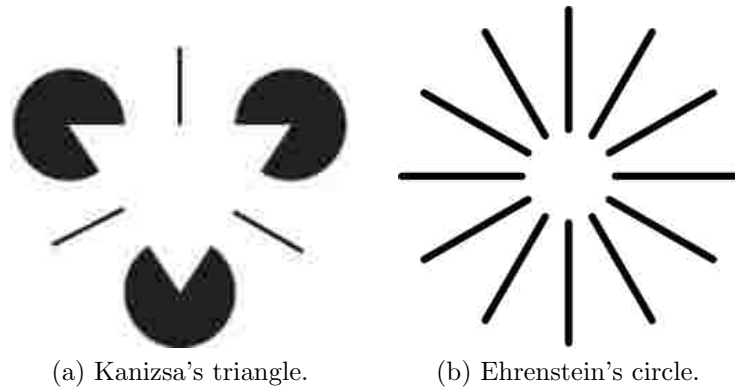


Figure 1.7: Examples of illusory contours believed to be completed with prior knowledge of geometry.

The incorporation and representation of knowledge is required for machine interpretation of images, the “image understanding” portion of this work. Knowledge is incorporated whenever possible in this work to assist the machine in making inferences. In this research, a “bottom-up” image understanding control strategy is

generally applied (*e.g.* [29]). Generally, the strategy controls the representations of the image data. The data is first pre-processed and labeled, or segmented. Further data partitioning is achieved by the application of descriptors. Finally, descriptors can sometimes be checked against a known definition, for example, that supplied by a knowledge base.

1.5.2 The Role of Noise

From an analysis perspective, signal refers to the regions in the image from which information is desired and can be extracted; noise may be defined for all intents and purposes as structure or areas in the image or sequence which are either undesirable, need not, or cannot be examined (*e.g.* Forsythe and Ponce [30]). Qualitative analysis of “noisy” or low SNR images by an observer suggests the utility of image enhancement techniques such as histogram equalization, contrast stretching, and unsharp masking [31]. These techniques may yield results which can be more pleasing visually, yet two problems can be created. First, enhancement can introduce false structures by emphasizing high frequency components, amplifying noise, while denoising can remove critical information on the same scale as important features in the image. Secondly, the resulting images may complicate algorithmic grouping and analysis due to the same false artifact creation and true artifact removal mentioned. When noise is indiscernible from structure at intensity and scale, with large temporal variance, denoising should be kept to a minimum if done at all.

Slight noise can interfere with or make impossible tasks relating to edge and object detection. Blurring or denoising can be a trade off with signal preservation and loss of structure. Figure 1.8 from reference [32] shows row-wise a ramp image followed by its first and second derivative images. Below each image is a horizontal intensity trace. From the first column, random Gaussian noise of increasing standard deviation corrupts the original image. Note that the addition of noise imperceptible to the human renders the second derivative indistinguishable from white noise. Once the level of noise or degradation is detectable in the image by the naked eye, as in the bottom row of Figure 1.8, differentiation becomes unreliable as a means to determine salient structure of the image. The second and third columns in the bottom row are directly equivalent to edge detection using the Sobel filter, and enhancement using unsharp masking, respectively.





Figure 1.8: Corruption of a step image with Gaussian noise and its derivatives [32].


1.5.3 The Role of Scale

Just as flow structures can exist at different scales, so can their representation in images. Existence at scale, and persistence through scale space play a crucial role in many vision algorithms. When information is ambiguous, consistency through scale can provide a means of structure or measurement confirmation. Therefore, where possible, a consistent scale representation is sought for each measurement objective in this research. Scale space representations used in image understanding include wavelets, image pyramids, and the use of Gaussian functions and their derivatives with scale parameter σ . Detailed discussions on scale space are given in references [33] and [34]. Burt and Adelson [35] introduced the image pyramid structure wherein an image is convolved with a Gaussian function and decimated. The repeated process creates a coarse to fine series of images decimated in a manner which preserves structure. Pyramid structures are used in numerous vision applications where a multiscale view is required.

Kalias and Narasimha [36] used image scales to infer physical scales. Wavelet transforms at different scales were performed on images of laminar and turbulent mixing layers to infer and clarify structure and organization, some of which was not evident in the original images. Figure 1.9 shows an image of the turbulent mixing layer. Wavelet transforms thresholded at a small coefficient illustrate the small scale homogeneous structure, while a larger threshold reveals the roller structures.



(a) Original image.



(b) Small scale wavelet thresholding.



(c) Large scale wavelet thresholding.

Figure 1.9: Mixing layer scales visualized with the wavelet transform [36].

Chapter 2: Double-Cone Flows and Schlieren Data

2.1 Introduction

The double-cone schlieren data sequence used in this research was provided by Arnold Engineering Development Center (AEDC) Hypervelocity Wind Tunnel 9. This chapter gives a brief overview of double-cone flows and significant findings of numerical studies and the issues that they have identified. An overview of the Tunnel 9 facility and the experiments performed by Coblish *et al.* [37] that generated the schlieren sequence follow. The schlieren sequence is discussed in order to establish the flow structures whose extraction and measurement is the scope of the current study. The focus of this research is on the application and development of computer vision techniques for schlieren analysis. The sequence used here is largely illustrative in purpose as discussed in section 1.1 and selected due to uniqueness in feature complexity and size.

Due to the complex and coupled viscous interactions that occur in the double-cone flows, this geometry was originally selected by the aerospace community to assist in the validation of high fidelity simulations of hypersonic flows with separation and shock interactions [38]. Due to its role in validation of simulations, many studies have focused on the numerical aspects, and simulations have been performed

at low Reynolds numbers to ensure steady laminar flow. Under certain flow conditions, difficulties are still encountered in predicting separation zone lengths and peak heating [39]. A recent review of previous experimental and computational studies on this geometry under various flow conditions is provided in reference [40].

2.2 The Double-Cone Flowfield and Previous Findings

The hypervelocity flow about the 25-55° double-cone model is shown in Figure 2.1 from reference [41]. The angle of the first cone is chosen such that an attached shock forms at the fore-cone tip. The angle of the second is large enough to ensure a detached bow shock. Flow behind the bow shock can be wholly or partially subsonic. These two shocks intersect to establish a complex flowfield in which shock-shock interaction and viscous interactions are coupled. The attached shock intersects with the detached shock, forming a shock-shock interaction. A transmitted shock emits from the intersection point and impinges on the cone surface causing a large adverse pressure gradient. The transmitted shock causes severe aerothermal loading at the point of impingement. A supersonic jet propagates from the point of impact along the length of the aft-cone. The adverse pressure gradient in the region of the cone junction causes flow separation upstream on the fore-cone surface. A shear-layer and contact surface originate from the point of separation. The separation zone in turn affects the interaction, developing a complex feedback mechanism. This feedback mechanism changes the pattern of the interaction.





Figure 2.1: Double-cone flowfield schematic from [41].


Druguet, Candler, and Nompelis [42] showed that due to the complexity of the double-cone flowfield, accurate predictions are extremely sensitive to the numeric schemes used and require careful grid studies. Of note is that even for low Reynolds number flows (*e.g.* 130×10^3 1/m), due to the sensitivity of the interaction feedback mechanisms, simulations must be carried out to at least 100 characteristic time scales to reach steady state, “if a steady state exists.” The characteristic time scale was defined as the time a particle moving at the freestream velocity takes to travel the length of the double-cone model. The location of the shock interaction is crucial to resolve correctly during simulations, as it establishes the flowfield interaction pattern [41,42]. Additionally, it was shown that hypervelocity double-cone flow can serve as an important benchmark case to evaluate the suitability of computational methods for hypersonic flows with separation and shock interactions.

Pressure and heat transfer rates on the double-cone surface in Figure 2.1 from Nompelis' work [41] compare experimental results with those from a blind simulation. Freestream conditions for the experiment were $M = 9.59$ with a unit Reynolds number of $130.9 \times 10^3 \text{ 1/m}$ ($39.9 \times 10^3 \text{ 1/ft}$). The low Reynolds number ensures laminar flow on the cone surface. These surface quantities from Nompelis' blind comparisons reveal further information on the flowfield physics, and indicate some of the numerical difficulties associated with prediction. The surface pressure in Figure 2.2a remains low where the fore-cone flow is attached. Pressure rises at the point of separation, and remains constant in the separated region. The sharp peak indicates the location of the transmitted shock impingement. The supersonic jet on the aft-cone surface in Figure 2.1 is evidenced by the compressions and expansions downstream of the impingement point. Nompelis states that the disagreement with experiment in the supersonic jet region is due to the inexact location of the intersection (triple) point, which alters the flow structure. Heat transfer rates are shown in Figure 2.2b, showing a significant over-prediction of heating rate on the fore-cone, and a peak heating rate at the point of shock impingement. The sharp drop and region of constant heat transfer rate indicate the point of separation and recirculation zone respectively. The embedded jet causes oscillations in heat transfer on the aft-cone surface.

Nompelis, Candler, and Holden [43] discovered the reason for the disagreement between the experimental and predicted heat transfer rate on the fore-cone. They postulated that vibrational freezing in the facility nozzle throat, if present, would lower the heat transfer rate to the model surface, as this quantity scales approx-



(a) Surface pressure distribution.



(b) Surface heat transfer distribution.

Figure 2.2: Numerical and experimental pressure and heat transfer distributions on the double-cone surface from Nompelis [41].

imately with $\rho_\infty u_\infty^3$. Heat transfer rate would therefore be more sensitive to the reduced freestream total enthalpy than surface pressure. Simulations verified the assumption of freezing at the facility nozzle throat conditions, reducing the heat transfer rate to the model surface. They also observed that, when, accommodations for slip and allowances for slight freestream non-uniformity were made, the computed heat transfer rate at the model surface was brought into agreement to within the experimental measurement uncertainty. It was also discovered that freestream non-uniformity reduced the size of the separation zone, bringing it to closer agreement with the experiment.

2.3 Hypervelocity Wind Tunnel 9

AEDC Tunnel 9 is a unique hypervelocity blow down facility capable of producing run times on the order of seconds, much greater than traditional impulse type facilities [44]. Tunnel 9's operational envelope is shown in Figure 2.3; it is capable of simulation at Mach numbers of 7, 8, 10 and 14. Flight duplication is achieved at Mach 7. Pure nitrogen is used as the working fluid. Plenum pressures and temperatures are capable of up to 26 kpsi (186 MPa) and 3040° F (1888 K), respectively. Shown in Figure 2.4, the test section has a 5 ft. (1.52 m) diameter and is 12 ft. (3.66 m) long; capable of testing full-scale models. The test article mounting system is capable of pitching a 200 lb (91 kg) model through angles of attack between -5° and 45°, and at rates up to 80° per second.

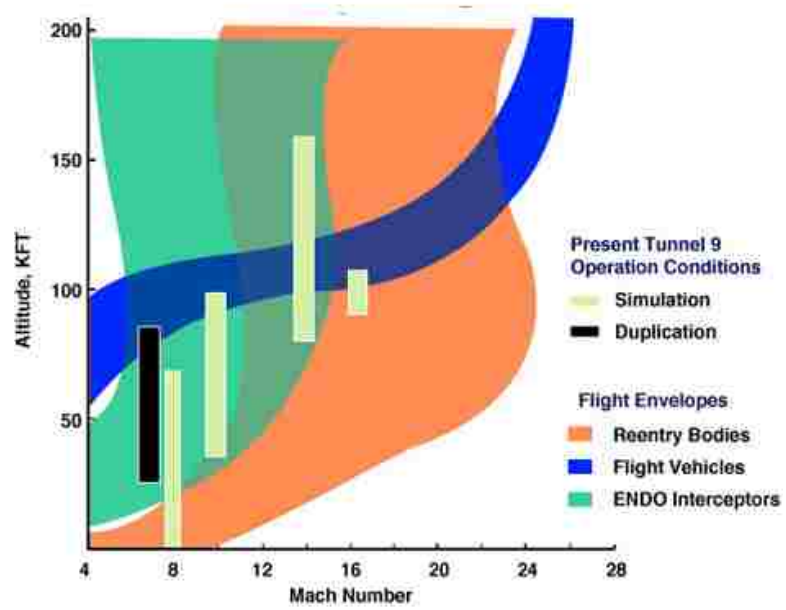


Figure 2.3: Tunnel 9 operational envelope.



Figure 2.4: Tunnel 9 test section.

To produce hypervelocity conditions, a fixed volume of nitrogen test gas is pressurized and heated in a vertical heating element, separated from a driver gas section and the nozzle by a pair of diaphragms. The nozzle and test section are evacuated to pressure values near 1 mm Hg. Diaphragms are burst once the test gas reaches desired pressure and temperature conditions. Cold nitrogen from the driver section acts as a piston, pushing the hot test gas into the test section in a manner that maintains constant test conditions. The period of usable test time occurs when the flow maintains a constant freestream Reynolds number.

2.4 Tunnel 9 Double-Cone Experiment

This section details the findings of Coblish *et al.* [37], whose experiment provided the optical data set used in this research. The double-cone configuration was tested at AEDC Hypervelocity Wind Tunnel 9 at a nominal Mach number of 14 at varying Reynolds numbers. Figure 2.5 illustrates the double-cone test article dimensions, and shows the model installed in the tunnel test section. The test article has a fore-cone angle of 25° and an aft-cone angle of 55° . The total length of the model is 6.051 in. (15.37 cm). The largest diameters of the fore and aft-cone sections measure 3.381 in. (8.588 cm) and 10.309 in. (26.185 cm), respectively. The model was instrumented with Kulite pressure transducers and Type-E thermocouples. All data was sampled at 500 Hz and passed through a 30 Hz low-pass filter prior to the data acquisition system. Unfortunately, this did not allow for high-speed measurement of the surface quantities. Schlieren video was also taken for each run (discussed in

the following section). Unit Reynolds number test conditions were 1.31×10^6 , 0.479×10^6 , 0.0983×10^6 , and 0.0682×10^6 1/ft (4.30 , 1.57 , 0.323 , and 0.234×10^6 1/m). At the two highest Reynolds numbers tested, CFD was not able to obtain a converged solution. Researchers therefore concluded that the flow would be unsteady in these cases. However, all pre-test CFD at the lower Reynolds converged without difficulty, suggesting that as in all previous validation studies, the flow would be steady at these Reynolds numbers.




Figure 2.5: The double-cone schematic and model installed in the test section from [37].


Motivated by Nompelis *et al.* [43], who showed that vibrational freezing in the nozzle was responsible for previous discrepancies in the fore-cone heat transfer rate, experiments were conducted to measure the degree of vibrational non-equilibrium in the tunnel core freestream. M. Smith and Coblisch [45] performed non-intrusive freestream measurements of vibrational temperature upstream of the double-cone model. Results showed that the Tunnel 9 Mach 14 nozzle operating conditions are at equilibrium within measurement uncertainty. Prior to testing, calculations were performed to assess the effects of non-equilibrium, slip effects, and freestream non-

uniformity. Computations showed that these effects were negligible. All simulations were therefore conducted with a uniform freestream at equilibrium conditions.

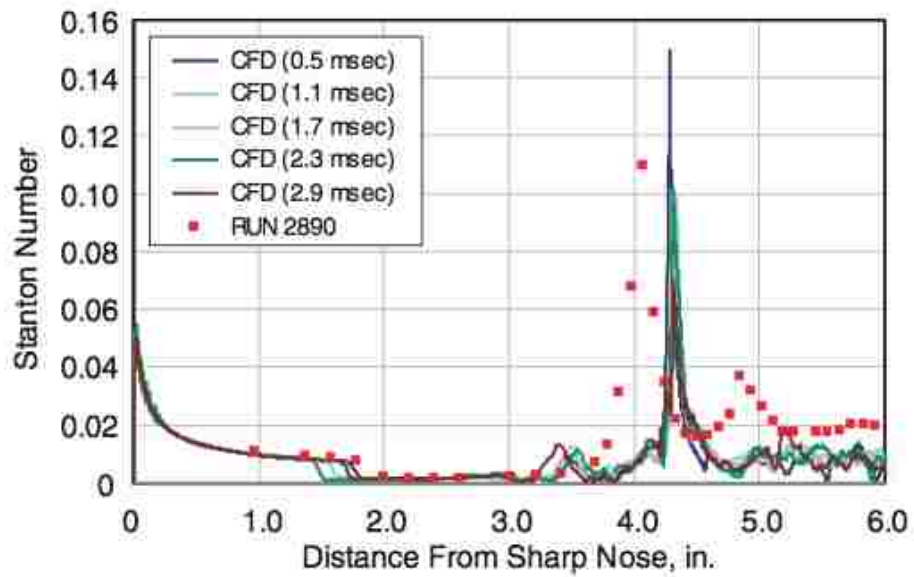
When compared with experimental results, the fore-cone heat flux was not over-predicted as it had been in all previous computational efforts. Additionally, computations appeared to have severely over-predicted the peak heating rate on the aft-cone surface associated with impingement of the transmitted shock from the interaction point. In stark contrast to previous findings, Coblish *et al.* observed significant flowfield unsteadiness, even at low Reynolds numbers. Temporal changes were evident from both pressure coefficient and heat flux data from the surface instrumentation. Frequency analysis, however, could not be performed due to lack of high-speed surface instrumentation and due to filtering prior to the data recording system. Significant unsteadiness was also evidenced from viewing the schlieren data at all Reynolds numbers. Coblish *et al.* [37] report in more detail the experiment and analysis of traditional surface instrumentation and comparison with numerical simulations.

In light of the unsteadiness observed during the test time, Coblish *et al.* conducted time-accurate CFD studies at the two highest unit Reynolds numbers tested, as initial CFD results were unable to obtain a converged solution for these conditions. Results from the unsteady CFD analysis from [37] are shown in Figure 2.6, and experimental measurements are also shown. These results, from Run 2890, correspond to a freestream Mach number of 13.603, and a unit Reynolds number of 1.31×10^6 1/ft. (4.30×10^6 1/m), the highest Reynolds number used in the experiments. Time accurate pressure coefficient and Stanton number computational results on

the double-cone surface are shown in Figures 2.6a and 2.6b, respectively. Both plots show that the separation shock foot experiences considerable motion, as evidenced by the pressure rise and heat flux decrease at approximately 1.5 in. from the model tip in Figures 2.6a and 2.6b, respectively. This unsteady shock foot motion is indicative of the “breathing” motion found in turbulent shock-wave/boundary-layer interactions. The supersonic jet on the aft-cone surface appears to be unsteady as well. Figure 2.6 also shows that the location and value of maximum aerothermal loading changes, which Coblish *et al.* noted was random. This indicates that the point of impingement of the transmitted shock on the aft-cone surface undergoes considerable motion, suggesting that the point of shock interaction is moving during the test time. The effect of the unsteadiness on heat transfer rates is summarized in Figure 2.7 from [37]. Time-averaged unsteady CFD results are shown with experimental measurement. The effect of the unsteadiness seems to broaden the region of peak heating, and over time, reduces its magnitude.



(a) Surface pressure distribution.



(b) Surface heat transfer distribution.

Figure 2.6: Unsteady CFD results from Coblish *et al.* [37].

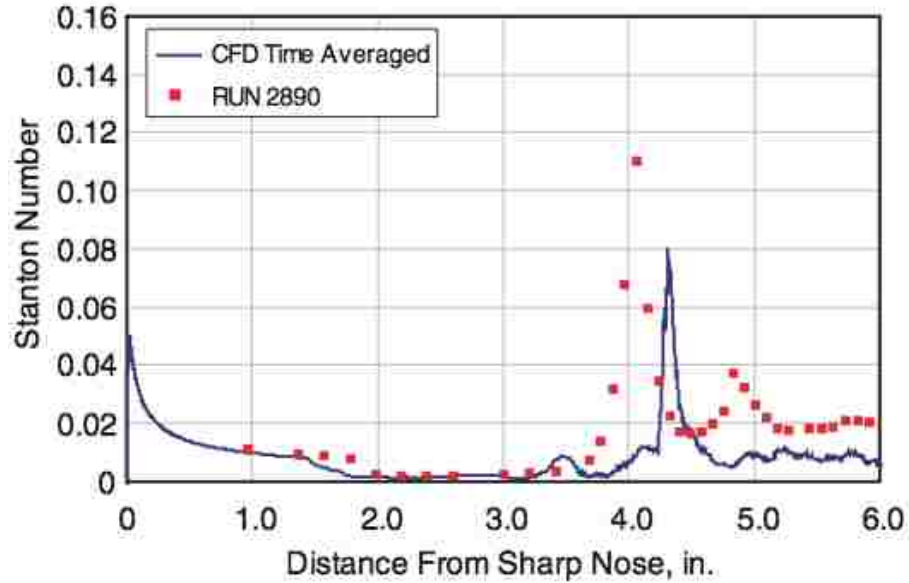


Figure 2.7: Time averaged computational results from Coblish *et al.* [37].

2.5 The Double-Cone High-Speed Schlieren Sequence

High-speed schlieren was acquired at 10 kHz for the two large Reynolds numbers tested. For these cases, the 312 x 416 resolution was used image the entire flow-field, capturing the double-cone flow about both sides of the model’s line of symmetry. This resulted in insufficient spatial resolution of the fluid structures. The laser illumination source malfunctioned prior to the low Reynolds number test conditions. Therefore, the low Reynolds number (and hence low density) test schlieren videos were acquired at 500 Hz. This rendered only a single optical data set, originally scheduled as merely a data repeatability run, with temporal and spatial resolution sufficient for further analysis using computer vision methods. This data set consisted of only the upper plane of symmetry, effectively doubling the spatial resolution of the initial image sequences.

The freestream conditions for the data set from the experiment performed in [37] had a nominal Mach number of 14 with a Reynolds number of 1.31×10^6 1/ft (4.30×10^6 1/m). Images were collected at the maximum framing rate of the Redlake model HG-100K camera of 10 kHz at an 8-bit resolution. Illumination was provided by a copper-vapor laser model LS-50 from Oxford Lasers. The laser has low coherence with a 25 ns pulse time and was triggered by the camera configured to operate at a continuous pulse rate of 10 kHz. This provided nearly instantaneous flowfield imaging, giving a Nyquist frequency of 5 kHz.

Although the nominal test time for this run was on the order of one second, the laser power source failed midway through the run. Viewing both the sequence, and the average intensity of the image during the test time allowed for clear identification of the failure point. From the sequence of 10,000 images of the potentially usable test time, only 5,000 were acquired before the laser failure.

Two consecutive images from the sequence of 5,000 and their absolute difference are shown in Figure 2.8, starting from the nominally steady-state run time t_o . To achieve the 10 kHz frame rate of the Redlake camera, images were acquired at a lower spatial resolution than the maximum allowed by the camera. Each image is 312 x 416 pixels, with flow structures of interest existing on a much smaller scale. Each pixel is 2 mm x 2 mm in the physical plane. The 10 kHz frame rate also necessitated a short exposure time, requiring the schlieren system to have high sensitivity. These factors resulted in images with low SNR, exhibiting a high degree of laser speckle, with temporally non-uniform illumination. In addition, freestream turbulence can be indistinguishable from speckle, and both exist on similar scales

as flow features of interest. The difference image (inverted for clarity) in Figure 2.8c shows the motion difference between frames. This difference is due both to the laser pulse and structure motion. White represents constancy; black represents the regions of most change between frames. The freestream area is the region of most change, misleading due to the fact that it does not correspond to an accurately captured timescale. Figure 2.8c also shows that the most textured regions (feature rich) of the images do not correspond to structures of interest in this study. These regions are the most rapidly changing in the data set, and therefore the strongest candidates for “features”. The largest challenge to any computer vision-based analysis method is sensitivity to illumination. The laser light source used in the double-cone data sequence has up to a 20% illumination variation between frames, typical of pulsed laser light sources. Also, the inclusion of freestream schlieres may occlude flow structures of interest in the three-dimensional optical path of integration.

From inspection of the image in Figures 2.8a and 2.8b, most of the structures shown in the flowfield sketch in Figure 2.1 are evident. The point of shock intersection and the separated flow region are clearly discernible. Above the separation zone, the oblique shock and contact surfaces are well imaged and distinct. The point of shock impingement appears obscured, although the diamond pattern associated with the supersonic jet is visible. The separation point on the fore-cone is seen, although the structure is blurred and difficult to make out, as is the attached conical shock. The bow shock on the aft-cone is sharp and distinct from the freestream and shocked flow regions, with evidence of curvature. The fore-cone attached shock is perhaps the most poorly resolved artifact of the flowfield, and unfortunately the

only structure for which an analytical solution is attainable. The “missing” attached conical shock and the point at which it intersects the separation shock may be obscured for several reasons. These may include factors such as a potential optical misalignment or structure occlusion by the optical path of integration through the boundary-layer at the optical window. It may also be partially due to the fact that at such a large freestream Mach number, oblique structures and shock layers are sufficiently steepened that they are close to the surface angle.




Figure 2.8: Two consecutive images and their absolute difference.

Chapter 3: Flow Isolation Using Segmentation

3.1 Motivation

A general bottom-up image understanding strategy requires the implementation of multiple sequential processing steps. The first steps in the understanding process involve any pre-processing of the image and its segmentation, or labeling into useful parts that correspond to information desired for the understanding task. Segmentations are later used to describe objects or for feature recognition, or in some cases can be results whose properties describe important information in and of themselves in the analysis tasks. The results of the segmentation in this chapter will be used to identify specific features in subsequent chapters through parallel and hierarchical understanding implementations, and be used as a measurement to provide information about the unsteady flowfield directly.

Two images from the test sequence are shown in Figure 3.1. Inspection of Figure 3.1 suggests that the most useful segmentation scheme for this particular problem would be to classify the freestream, the shocked flowfield region, and the model respectively. Only a small portion of the images shown in Figure 3.1 is of interest to the present study. Here we are concerned with the shocked flow region contained between the freestream and the model surface, and all else can be

considered noise. A bottom-up schematic of the image interpretation useful to the expert is shown in Figure 3.2, corresponding to the desired segmentation scheme.




Figure 3.1: Consecutive images from the sequence.

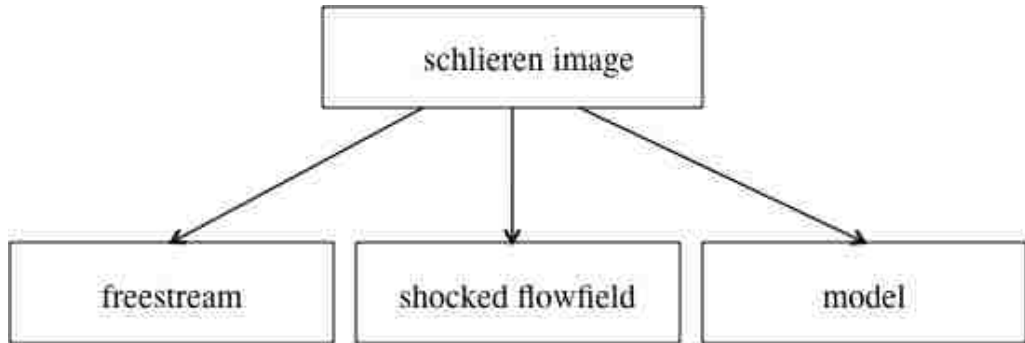


Figure 3.2: Schematic of image interpretation.

The labeling of each schlieren image, or segmentation of the images into meaningful groups representing the freestream, flowfield, and body allows for the following: 1) we are interested only in the flowfield region itself; it can be isolated, 2) the flowfield itself can serve as a measurement of the unsteadiness; the area of the three-dimensional flowfield projected onto the schlieren image is directly proportional to the global unsteadiness, and finally, 3) a labeled representation of physical structure

allows for rule-based searching and classification of other desired flow structures, such as the shock wave, shear layer, and intersection points.

This chapter encompasses pre-processing of the sequence followed by a manual segmentation of the double-cone model. Segmentation methods relevant to schlieren images are discussed. Heuristic measures are incorporated into an algorithm that results in a physically meaningful segmentation of the double-cone flowfield. The segmentation results are then evaluated.

3.2 Segmentation

Segmentation is the assignment of label to regions of interest. The world is perceived in terms of objects, or parts of a whole, not pixels; the goal of segmentation is to convert the latter to the former. This allows the conversion of images in Figure 3.1 into the format shown in Figure 3.2. Once grouped into meaningful objects, these partitions isolate regions for analysis, and allow for Cartesian descriptors of the segments. Segmentation of images into partitions that are meaningful to humans is difficult, as discussed in the introduction. Ideally, segmentation allows for a meaningful description of the information contained in the image. The concepts of foreground and background have clear desirable partition schemes in most natural images and scenes. Here the segmentation schemes must also be meaningful to the expert, that is, they must be useful from an analysis perspective. In this case, the physical world is governed by fluid mechanics structures, *i.e.* shear layer, shock wave, separation region, shock structure and body model.

Decisions on a desired segmentation scheme must be made on a case-by-case basis, depending on schlieren image content, and desired measurement outcomes. The object partitioning here is an example based on the test data sequence properties, although the segmentation methodology is largely general. Methods of segmentation or partitioning are numerous and can involve techniques based on thresholding, statistical data-clustering and graph theory. The focus here is limited to graph theoretic partitioning since these methods tend to incorporate a global approach to segmentation which is in line with the Gestalt theory of perception and human vision [28] (discussed in section 1.5.1), by incorporation of both feature similarity and spatial closeness. This is advantageous over methods such as thresholding which only incorporate intensity similarity, and therefore do not correspond to models of human vision.

3.3 Region Preserving Diffusion-Based Filtering

All image understanding tasks in this research involve the extraction of information from images dominated by white noise spatially and temporally, from which unsteady structures of interest must be identified, isolated and described. Blur is needed to reduce noise which can make tasks such as object detection difficult if not impossible. However, in many cases structures in the image can be on the order of noise or less. In the schlieren images shown in Figure 3.1, image evidence is not sufficient to provide strong boundaries between the freestream, flowfield, and body: neither in the individual images nor from frame to frame. For the analysis of im-

ages with low SNR to detect structures that exist at small scale and resolution, the best denoising scheme can be no denoising at all. Denoising, as in image enhancement, can remove true image artifacts and potentially create false structure. This does not mean that certain edge-preserving schemes cannot lead to intermediate representations of the image that can be useful for referencing as a marker image.

To encourage the desired segmentation scheme shown in Figure 3.2, the images must be pre-grouped so that affinities between and dissimilarities across structures are more pronounced for assistance in higher-level grouping. A global visual impression of a noisy image can be used as a pre-processing step for higher-level tasks such segmentation, which is performed in the following sections. Smoothing with a Gaussian kernel is a common technique for pre-processing images to remove small detail. The kernel is edge preserving in that no new maxima are produced, but it blurs across boundaries by spreading the influence of an edge region. Statistical rank filtering processes can be undesirable for similar reasons.

A filtering scheme is desired which blurs within boundaries but not across them, encouraging similarity between like regions and enhancing dissimilarities across them. A method that provides this effect is the bilateral filter [46], which denoises while strongly preserving edges, and may be considered a member of the class of diffusion-based methods. Diffusion-based schemes are preferable for edge-preserving denoising and grouping for low SNR images over non region preserving methods, such as smoothing or rank filtering. Non-linear diffusion for image segmentation and denoising based on Fick's law of diffusion was first proposed by Perona and Malik [47]. An overview of diffusion-based techniques is given in Weickert [48].

The bilateral filter introduced by Tomasi and Manduchi [46] allows a single pass non-iterative scheme while providing comparable results to PDE diffusion-based methods. The bilateral filter employs a Gaussian kernel both on the domain and range of an image. The effect is a strong edge preserving diffusion effect. The bilateral filter response $\mathbf{I}_{BF}[\mathbf{I}(\mathbf{p})]$ at an image pixel \mathbf{p} is given below as a function of spatial distance between points \mathbf{p} and \mathbf{q} , and their difference in intensity values

$$\mathbf{I}_{BF}[\mathbf{I}(\mathbf{p})] = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{N}_p} \exp\left(-\frac{\|\mathbf{q} - \mathbf{p}\|^2}{2\sigma_D^2}\right) \exp\left(-\frac{|\mathbf{I}(\mathbf{q}) - \mathbf{I}(\mathbf{p})|^2}{2\sigma_R^2}\right) \mathbf{I}(\mathbf{q}). \quad (3.1)$$

The response is normalized by a weight $W_{\mathbf{p}}$, which is the product of the exponentials unscaled by the intensity value at pixel \mathbf{q} ,

$$W_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{N}_p} \exp\left(-\frac{\|\mathbf{q} - \mathbf{p}\|^2}{2\sigma_D^2}\right) \exp\left(-\frac{|\mathbf{I}(\mathbf{q}) - \mathbf{I}(\mathbf{p})|^2}{2\sigma_R^2}\right). \quad (3.2)$$

The first grouped term in Equation 3.1 represents the spatial filter response, while the second term represents the intensity response, over a neighborhood \mathcal{N}_p about pixel \mathbf{p} . Domain and range scale parameters are given by σ_D and σ_R respectively. The range scale σ_R acts a weighting scale to the traditional domain filter σ_D .

Figure 3.3 presented below from reference [46] shows the effects of edge preservation on a step corrupted with Gaussian noise (first column). The second column shows the response of the filter placed at the bright side (right) of discontinuity, the bright values of the images are therefore smoothed together, whereas the dark region

to the left of the discontinuity only weakly contributes to the filtering of the bright values on the right of the edge. The converse holds when the filter is centered in the dark portion of the image. The filtered image is shown in the rightmost portion of the figure. Note the edge preservation quality, whereas a typical Gaussian blur would have introduced a ramping effect, the diffusion filter has preserved the edge almost completely with the added intensity weighting parameter.





Figure 3.3: Bilateral filter response on a step edge [46].

A sample result of applying the bilateral filter to two randomly selected images in the sequence shown in Figure 3.4 with grouping scales σ_D and σ_R of 4 and 0.1 respectively with a kernel support width of $(3\sigma_D + 1)$ is shown in Figures 3.4c and 3.4d. The filter has accomplished the following: the regions of “flow” and “freestream” are much more distinguishable to the human eye with a stronger visible partition across the main shock in a manner which strongly preserves and enhances the edge; regions now have more meaningful boundaries in an individual image and are more similar in frame-to-frame comparisons of the sequence, as compared with their unfiltered counterparts in Figures 3.4a and 3.4b. This is in stark contrast to the non-region preserving Gaussian filtered in Figure 3.5 where the region delineating boundaries have been obscured, which is equivalent to a bilateral filter with range


scale $\sigma_R = 0$; at the increased spatial scale alone provided by σ_D , much information has been lost. The freestream region in Figure 3.5 is much darker than in Figure 3.4d; as there is no range exclusion from the neighboring dark regions closer to the model. Although any of the diffusion-based schemes discussed in this section are viable techniques for edge-preserving denoising and grouping for low SNR schlieren images, the bilateral filtered images were visually preferable to those grouped using traditional diffusion-based methods.

3.4 Manual Segmentation of the Model

As the boundary between the model and the flowfield is largely obscured due to density gradients on most of the upper surface, the contour is ambiguous. The boundary of the model is visible in all images, but it is blurred with the flow structures near the surface walls. This makes the perception of the delineating contour illusory rather than actual. For this reason, the model must be segmented manually. Segmentation schemes cannot delineate a missing boundary. The model must be segmented manually, under the assumption that it is stationary. Model motion is sometimes observed in Tunnel 9, however this occurs at low frequencies (on the order of 10-15 Hz or less), and was not observed by the human eye in this data set (more robust motion estimations are considered later). In addition, the symmetry of the model should ensure stability. Model and camera motion can never be ruled out entirely. It is assumed that the camera is level with the model, and that any relative motion between the two is negligible in comparison to gross flow motion.



(a) Raw image $\mathbf{I}(t_o - 1832)$



(b) Raw image $\mathbf{I}(t_o + 4960)$



(c) Diffused image $\mathbf{I}_{BF}(t_o - 1832)$



(d) Diffused image $\mathbf{I}_{BF}(t_o + 4960)$

Figure 3.4: Effect of the bilateral filter with parameters $\sigma_D = 4$ and $\sigma_R = .10$.




Figure 3.5: Non-region preserving Gaussian filtering for image $\mathbf{I}(t_o + 4960)$, $\sigma_D = 4$

In order to manually segment the test article, an approximate estimate of its motion relative to the flowfield is required. Many methods exist to separate foreground and background regions in image sequence, the simplest of which is the threshold based. The model motion estimate was performed by approximating a motion energy image \mathbf{I}_{diff} through the sequence using a semi-thresholded cumulative difference image given by

$$\mathbf{I}_{diff} \approx \sum_{t=t_o+1}^{t_{max}} |\Delta \mathbf{I}| \geq T, \quad (3.3)$$

where $|\Delta \mathbf{I}|$ is the absolute difference of the current image t referenced from the first frame in the sequence, summed over the number of images in the usable test time t_{max} . The value T was experimentally selected such that the main flowfield structures would be discernible from freestream region sensitivities caused by turbulent eddies and non-uniform illumination due to the laser. The cumulative difference

image of 5,000 frames is shown in Figure 3.6. From this image, it was determined that the model is still with regard to gross shock structure. Based on this evidence, a model segment was manually constructed that was added to the labeled image resulting from the higher-level segmentation performed in the next section. This approach for separating the test article may be applied to any stationary physical structure contained in a schlieren sequence, such as a tunnel wall or optical window boundary.



Figure 3.6: Motion energy approximation image.

3.5 The Image as a Graph: Graph Terminology

Image partitioning using graph-theoretic techniques have been established as powerful segmentation methods [49–51], and have seen wide spread use in. A graph $\mathbf{G}(\mathbf{N}, \mathbf{E})$ is defined as a set of vertices or nodes $n \in \mathbf{N}$ connected by a set of edges $e \in \mathbf{E}$. A weight ω is assigned to each edge connecting two arbitrary nodes,

representing a cost function. To construct a graph from an image, each pixel becomes a node, and edge weights are constructed from one or more affinities such as: the spatial distance, a function of the intensity difference, or a relationship based on textures between pixels. References [49,50] provided a detailed description of graph theory as it relates to image analysis. An $N \times M$ image will have $NM = |\mathbf{N}|$ nodes. An image graph has the advantageous property of being undirected, *i.e.* there exists no preferred direction between nodes in which the edge cost function would change. The image graph is also symmetric; that is, if node n_i is adjacent to node n_j , with associated edge weight $e_{n_i n_j}$, then node n_j is adjacent to node n_i with the same edge or weight value. Finally, image graphs are often locally connected, which lead to block diagonal representations that are often sparse. Figure 3.7 illustrates a simple undirected weighted graph with nine nodes, showing node, edge and weight labels. Nodes 1 and 2 are labeled n_i and n_j , with the connecting edge weight denoted by ω_{ij} . Numbers next to the edges represent the cost between adjacent nodes.




Figure 3.7: Example of a weighted undirected graph.

The partitioning of a graph into two segments S and \bar{S} assigns a binary labeling \mathbf{x} (characteristic vector) for each pixel (graph vertex n_i) in an image such that its value is unity if it belongs to a segment $S \subset \mathbf{N}$, and zero otherwise (it belongs to the set complement),

$$\mathbf{x} = \begin{cases} 1 & \text{if } n_i \in S. \\ 0 & \text{if } n_i \in \bar{S}. \end{cases}, \quad (3.4)$$

where $S \cup \bar{S} = |\mathbf{N}|$ and $S \cap \bar{S} = \{\emptyset\}$. Each graph edge connecting any two nodes has an associated weight, which is generally an appropriate measure of affinity between nodes n_i and n_j . This is a generalized adjacency list between all nodes n_i and n_j ,

$$\omega_{n_i n_j} = \exp\left(-\alpha_I (I_{n_i} - I_{n_j})^2\right), \quad (3.5)$$

where α_I is a weighting parameter of the selected feature affinity, here (node indexed) intensity. The degree d_{n_i} of a node n_i for a weighted undirected graph is defined in Equation 3.6 as the sum of all edge weights or total connectedness leaving node n_i ,

$$d_{n_i} = \sum_{e_{n_i n_j}} \omega(e_{n_i n_j}) \forall e_{n_i n_j} \in \mathbf{E}. \quad (3.6)$$

The degree matrix and the weight matrix are used to define the graph Laplacian. The Laplacian of a graph is defined as the connectedness d_{n_i} of node n_i minus the weighting function between locally connected edges $\omega_{n_i n_j}$, is defined below in Equation 3.7

$$L_{n_i n_j} = \begin{cases} d_{n_i} & \text{if } i = j, \\ -\omega_{n_i n_j} & \text{if } e_{n_i n_j} \in \mathbf{E}, \\ 0 & \text{otherwise.} \end{cases} \quad (3.7)$$

The Laplacian therefore characterizes the graph with a feature affinity measure through $\omega_{n_i n_j}$ and a spatial measure of Euclidean distance between features through d_{n_i} . The matrix is $\mathbf{N} \times \mathbf{N}$ and positive semi-definite, symmetric, and sparse for locally connected graphs.

3.6 Graph Cut Segmentation

Standard graph partitioning algorithms used for image segmentation include the min-cut/max-flow algorithms [49], and spectral methods, most notably the normalized cut of Shi and Malik [50]. The minimum cut partitions the graph along the weakest connection between groups, which can lead to small isolated segments that may not correspond to a desirable grouping. The normalized cut algorithm forms partitions by maximizing an affinity within groups while simultaneously minimizing associations across groups, using the Fiedler value for the partition, requiring long computational times and can become unstable. An alternative to these methods which favors large partitions and offers an increase in speed by reducing the labeling problem to a linear system is the isoperimetric graph cut developed by Grady and Schwartz [51].

Figures 3.8 and 3.9 from reference [50] help to visualize the graph partitioning problem. The data cluster in Figure 3.8 show the minimum cut, which can favor small isolated segments through partitioning along the weakest edges. The vertical line in the figure labeled “better cut” is the normalized graph cut developed by Shi and Malik, which in this case yields an improved segmentation. Shi and Malik demonstrate the normalized cut algorithm in Figure 3.9 on a synthetic image. Three segments are seen in the noisy image in Figure 3.9a; the recursive normalized cut has isolated the segments shown in Figures 3.9b, 3.9c, and 3.9d.



Figure 3.8: Graph partitioning illustrated on a cluster from [50].

Figure 3.10 shows a partition of the graph in Figure 3.7. The graph has been partitioned (cut) into two disjoint regions S_1 and S_2 that comprise \mathbf{G} by the minimization of a function; here the minimum cost. The minimum cut is represented by the dashed line in Figure 3.10 and given by the minimum possible sum of broken edges to segment the graph.

The isoperimetric graph partitioning method minimizes a perimeter of a segment over a maximal area on a graph. This algorithm is well suited for the seg-

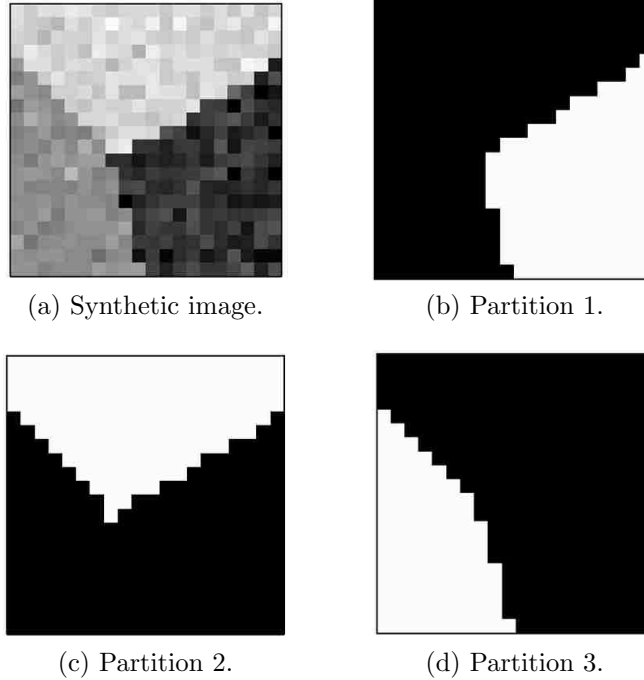


Figure 3.9: The normalized cut performed on a synthetic image. [50].

mentation scheme in Figure 3.2, in that large segments (freestream and shocked flow regions) are typically produced. The desired segmentation of images could be considered an under-segmentation (which is desired here), and the isoperimetric problem is appropriate for this by definition. A summary of the isoperimetric graph partitioning algorithm from reference Grady and Schwartz [51] is presented here.

The isoperimetric constant of a region $S \subset \mathbf{G}$ (segment) of a graph is based on the Cheeger constant of a manifold, and given by the minimum of the region boundary divided by the total volume of the region S

$$h_G = \min_S \frac{|\partial S|}{Vol_S}. \quad (3.8)$$

The area of a region boundary $|\partial S|$ of a set S is defined the sum of the edge weights on the boundary ∂S which is the set of edges in the graph that have a single graph




Figure 3.10: Illustration of the minimum graph cut.

node in the set. Defined in below in Equation 3.9 for a weighted graph

$$|\partial S| = \sum_{e_{n_i n_j} \in |\partial S|} \omega(e_{n_i n_j}). \quad (3.9)$$

For a weighted undirected graph, the above can be rewritten in terms of the graph Laplacian in Equation 3.7 as

$$|\partial S| = \mathbf{x}^T L \mathbf{x}. \quad (3.10)$$

The volume Vol_S of a graph region S is defined as the sum of all node degrees in the segment, which is a measurement of the total connectedness of a region S on the graph, which is a measure of total connectivity of the set S .

$$Vol_S = \sum_{n_i \in S} d_{n_i}. \quad (3.11)$$

In terms of the degree matrix and the indicator vector, this can be written as

$$Vol_s = \mathbf{x}^T d. \quad (3.12)$$

Using the definitions given by Equations 3.5 and 3.6, the isoperimetric constant in Equation 3.8 is written in terms of the label vector in Equation 3.4 and the graph Laplacian given by Equation 3.7 as

$$h_G = \min_{\mathbf{x}} \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T d}. \quad (3.13)$$

A solution of h_G in Equation 3.13 is accomplished by minimizing a cost function which is reduced to the linear system

$$L \mathbf{x} = d. \quad (3.14)$$

The Laplacian matrix in 3.14 is singular; to overcome this problem, Grady and Schwartz remove the row and column of L corresponding to the node with the degree of largest connectivity, or the “ground node”, along with the corresponding entries in \mathbf{x} and d allowing the system to be solved. Grady and Schwartz [51] present in detail the conversion of Equation 3.14 into the nonsingular system below in terms of the reduced Laplacian

$$L_o \mathbf{x}_o = d_o. \quad (3.15)$$

Grady [52] shows that the solution to the isoperimetric problem using Equation 3.15 is the “combinatorial equivalent” to solving the equivalent circuit and minimizing its power with specified current sources and a ground node needed to uniquely solve for potentials.

This system is relaxed to solve \mathbf{x}_o in the modified Equation 3.4 for real values. To convert this into a graph partition, the real values must be thresholded back to binary values to create the graph partitions. A threshold, or cut value α_c is chosen such that S contains the nodes corresponding to $x_i > \alpha_c$ and the complement \bar{S} contains the nodes that meet the criterion $x_i \leq \alpha_c$. The cut value α_c is selected from \mathbf{x} such that the minimum value of 3.13 is obtained. This yields a two class segmentation. The partitioning is continued recursively on each segment until the isoperimetric value of each segment has exceeded a minimum specified value corresponding to the largest allowable isoperimetric ratio that defines a quality segmentation.

3.6.1 Determining a Physically Meaningful Flow Partition

An arbitrary partitioning which is desirable to an aerodynamicist is not easily obtained; segmentations that are physically meaningful are not often based on image evidence, but on the interpretation of the particular image by the expert. This is true for images in general and for this particular data sequence. In this case, the desired segmentation scheme is to assign labels to the freestream, the flowfield, and the test article as shown in Figure 3.2.

Prior to the diffusion filtering of the images, there was little affinity between aerodynamically meaningful regions in either textures or intensities, which inhibits grouping algorithmically. The use of the bilateral filter for a pre-segmentation step partially assures that the graph cut will segment along the outer-shock structure by visually enhancing the boundary along the shock, increasing the measure of similarity within these regions and decreasing similarity across them as evident from Figure 3.4. This does not, however, ensure that the desired partition will occur. In certain images, the region downstream of the aft-cone (the expansion region) is visibly the brightest and most dissimilar region, and therefore the most likely candidate for a single partition. In most images this is not the case; the desired partition along the shock is the most visible region delimiter. Due to the expansion fan region discussed previously, the recursive method of partitioning is used as multiple partitions may be required to isolate the freestream from the shocked flowfield.

To evaluate whether the completed segmentation has isolated the freestream as desired, a form of “prior knowledge” was incorporated. The exact shape of the shocked flow between the freestream and the model is unknown; therefore there are no rules or observations that can be used to formally construct a knowledge base for an understanding model in this particular case. However, heuristics can be incorporated to ensure the desired partitioning has occurred. With an unknown number of segments, we require that the freestream region is a single segment, while any other pixels are allowed to belong to any number of segments, subject to the constraint that they are not part of the freestream.

Heuristic measures may be employed based on knowledge of certain static regions in the image which must be inclusive or exclusive of certain regions. The image was segmented until at least 2 partitions were obtained. After the partition was performed, the upper left corner p_o was labeled freestream. Several additional points \mathbf{P} were randomly selected closer to the shock surface; these pixels in must also be labeled freestream. The lower right hand corner of the image will always contain the model; this point was labeled q_o . Several randomly selected points moving toward the outer-shock were placed in \mathbf{Q} . The points in \mathbf{Q} are allowed to take on any label other than the freestream. In a sense this is a prescription of foreground and background priors required to initialize the segmentation scheme in many interactive algorithms, after the fact. Inclusive and exclusives points \mathbf{P} and \mathbf{Q} overlaid on two images and their respective segmentation boundaries are shown in Figure 3.11 at time instances $(t_o + 107)$, $(t_o + 1682)$, and $(t_o + 4590)$, which were selected due to the number of segments produced. Figure 3.11 shows the diffused images used for segmentation, below are the segmented images with 4, 5 and 6 segmented regions respectively. Points \mathbf{P} in all cases are located in the freestream segment. The points \mathbf{Q} which must be exclusive of the freestream each take on multiple labels, and have an arbitrary label other than the freestream. Note that sets of points \mathbf{P} and \mathbf{Q} are static and do not change during the sequence; these were experimentally selected and visual evaluation of the results indicate that this heuristic demonstrated in Figure 3.11 works well through the sequence irrespective of the number of segments returned from the partitioning scheme.




Figure 3.11: Illustration the heuristic to determine a good segmentation from N partitions. The points \mathbf{P} indicated by \circ must be labeled freestream. Points \mathbf{Q} marked $+$ may be assigned any label other than freestream.

3.6.2 Flow Segmentation Algorithm

The algorithm for determining a physically meaningful image scheme is given in Algorithm 1. Inputs are the diffused image \mathbf{I}_{BF} and the raw image \mathbf{I} , and labeled body image \mathbf{I}_{body} . The parameters required for the isoperimetric segmentation algorithm are the intensity affinity weighting parameter $\alpha_I = 100$ and the recursion stop parameter $\mathbf{stop} = 1e^{-5}$. The points \mathbf{P} and \mathbf{Q} are selected manually and remain the same for the sequence. Image graphs are constructed using 8-connectivity of each pixel (small world topological view). Segmentations are performed directly on the diffused image using the Graph Analysis Toolbox developed by Grady [52, 53], partitioning the input images into $N \geq 2$ segments. If a segmentation is not achieved (trivial solution), the image is set aside. To evaluate a meaningful segmentation, points \mathbf{P} are checked that should be contained in the freestream segment and labeled fs . Points \mathbf{Q} must have any label l other than fs . If these criterion are not met, a segmentation failure is identified. Following this, all points in the labeled region containing \mathbf{P} are labeled fs , a point having any other label is labeled flowfield, or ff , creating a two-class image irrespective of the number of segments originally found. A modified version of the unfiltered image \mathbf{I}_M is also calculated, with the freestream region mapped to white; partitions were performed directly on the diffused image \mathbf{I}_{BF} , although segments were mapped back to the unfiltered images \mathbf{I} . Finally, the manually segmented body image is mapped onto the two-class image creating the three-class labeled image \mathbf{I}_L consisting of the freestream, flowfield, and model.

Algorithm 1: Determining a physically meaningful segmentation.

Input: \mathbf{I} Raw Image
 \mathbf{I}_{BF} diffused image
 \mathbf{P} points for freestream inclusion
 \mathbf{Q} points for freestream exclusion
 \mathbf{L}_{body} labeled body image
 α_I intensity affinity weight
 h_G isoperimetric, or **stop** bound

Output: \mathbf{I}_L the three-class image labeled fs , ff , and $body$
 \mathbf{I}_M modified raw image with freestream region removed

for $i = [t_o, \dots t_f]$ **do**
 Perform isoperimetric cut on \mathbf{I}_{BF} giving N segments;
 if $N < 2$ **then** // segmentation failure
 | $fail[i] = \mathbf{true}$;
 | **continue** to next i ;
 end
 // Ensure the segmentation is physically meaningful
 Label upper left pixel as freestream, $\mathbf{I}_L[p_o] = fs$;
 Ensure all points $\in \mathbf{P}$ also have label $\mathbf{I}_L[\mathbf{P}] = fs$;
 if $\mathbf{I}_L[\mathbf{P}] \neq fs \forall p \in \mathbf{P}$ **then** // segmentation failure
 | $fail[i] = \mathbf{true}$;
 | **continue** to next i ;
 end
 Label a known *body* point arbitrarily, $\mathbf{I}_L[q_o] = l$;
 Check that the points in \mathbf{Q} have any label other than fs , $\mathbf{I}_L[\mathbf{Q}] \neq fs$;
 if $\mathbf{I}_L[\mathbf{Q}] = fs \exists q \in \mathbf{Q}$ **then** // segmentation failure
 | $fail[i] = \mathbf{true}$;
 | **continue** to next i ;
 end
 // Create the Labeled Image
 Relabel the image such that all pixels in the segment containing points \mathbf{P} are labeled fs ;
 Map all other segments to the flowfield, $\mathbf{I}_L[x_i \neq fs] = ff$;
 Map the manually segmented model image \mathbf{L}_{body} to the partition, giving the three-class labeled image \mathbf{I}_L ;
 // Create the Image with freestream removed
 Map all points in the raw image \mathbf{I} labeled freestream in the labeled image to white $\mathbf{I}_L[x_i = fs] = 1$ giving the modified image with freestream removed \mathbf{I}_M ;
 return $\mathbf{I}_L, \mathbf{I}_M$;
end

3.7 Segmentation Results

Each successfully segmented image from the sequence consisted of between 2 and 6 partitions, with 2 and 3 segments being the most common. All but three of the 5,000 images were segmented successfully using the Algorithm outlined in the previous section. Raw images at selected time steps in the sequence and their labeled counterparts are shown in Figure 3.12. These results are indicative of segmentations obtained throughout the sequence. Visual inspection of the labeled images \mathbf{I}_L with respect to their raw counterparts \mathbf{I} indicates a performance of the partitioning algorithm as desired. Figure 3.13 displays two random images from the sequence in their original form and with the freestream segment removed and mapped to white in the corresponding images \mathbf{I}_M . These enlarged images visually illustrate the accuracy of the cut along the outer-shock structure, and show that segmentation and removal can be done in place of background subtraction and regional noise removal.

Three segmentation failures from a sequence of 5,000 were detected automatically using the method of known points to be included in and excluded from the freestream region in Algorithm 1. These failures are shown against their raw image counterparts in Figure 3.14. These Figures illustrate that in all cases the freestream segment had merged with the flowfield, demonstrating that a simple segmentation scheme does not always produce meaningful results; and that failures may be detected reliably and in an automated fashion. To correct the improperly labeled images, the isoperimetric cut parameters were fine-tuned and successful results were evaluated manually.






Figure 3.12: Evaluation of labeled Images.



(a) Raw image $\mathbf{I}(t_o - 1832)$



(b) Raw image $\mathbf{I}(t_o + 4960)$



(c) Modified image $\mathbf{I}_M(t_o - 1832)$



(d) Modified image $\mathbf{I}_M(t_o + 4960)$

Figure 3.13: Two images freestream segmentation removed.




Figure 3.14: Automatically detected segmentation failures.

3.7.1 Segmentation Accuracy

It is infeasible to individually examine all segmented images to judge their accuracy for large data sets. The inspection of images at different time steps through the sequence displayed in Figure 3.12, and the fact that the segmentation scheme devised in Algorithm 1 was able to detect any obviously poor labeling schemes lend a measure of trust to the methodology. There are additional statistical measures and graph properties that can be employed to measure the accuracy of the segmentation scheme. Treating the area of the segment as a variable with unknown behavior, the observation of that variable over time should exhibit some distribution over time. If the variable is binned into unique values, outliers from the distribution can be defined and selected for manual inspection. Small world 8-connected graphs were used to perform the image partitioning in the previous section. For a sensitivity study, Algorithm 1 was run again with a larger connectivity of radius-3 using the same intensity affinity α_I and stop criterion h_G . This significantly increases the number of entries in the graph Laplacian and the computational time, segmentation results were evaluated visually and remain similar to those obtained using the smaller connectivity.

A histogram of flowfield segment areas is given in Figure 3.15 showing an approximately Gaussian distribution of the flow partition for each graph-connectivity of the flowfield segment. Nearly 1,000 discrete area values were obtained for both the 8-connected and radius-three connected graphs; although the histogram displayed below has fewer divisions for clarity. Distribution properties of the segments are

given in Table 3.1. The near normality of the distributions as evidenced by values of skewness and kurtosis in Table 3.1 allows detection of outliers using the 3σ event rule. Of the 5,000 observations, only 15 images gave flowfield partitions that could be considered statistical outliers. These three-class image partitions of these statistical outliers were visually evaluated against their raw counterparts. Each outlier was visually judged to be an acceptable partitioning scheme and these labeled images were retained. Figure 3.16 shows the % difference between the flowfield segment area with graph- connectivity which is nearly negligible. From human observation of selected results, the fact that little difference exists in the flowfield segment results with graph-connectivity along with a successful means of identifying poor segmentation results, it can be concluded that the segmentation algorithm performs accurately and as desired. Partitioning using the 8- connected topology were visually preferable over the radially connected results, and will be used for all further results.

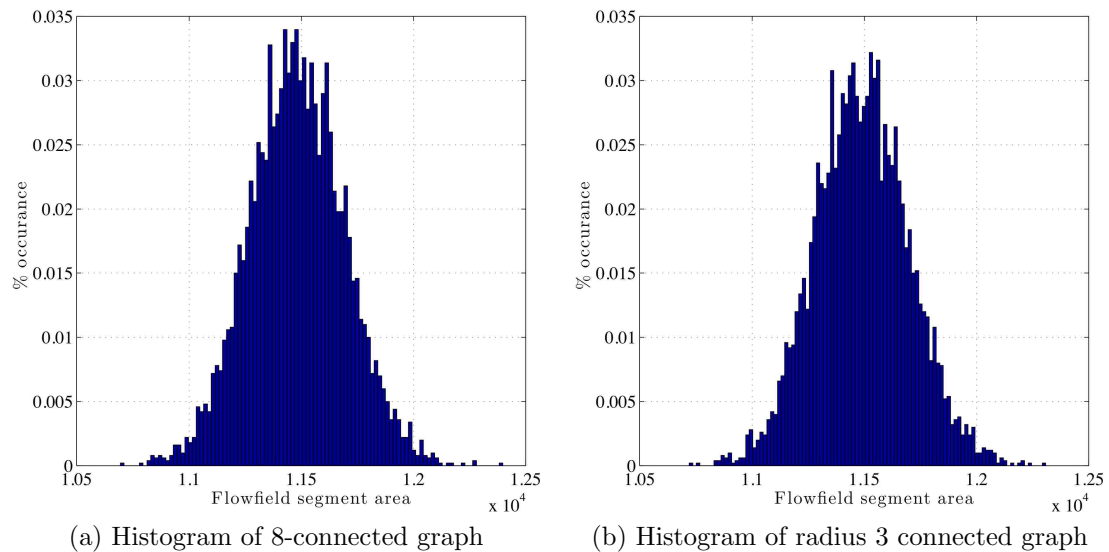


Figure 3.15: Histogram distribution of flowfield segment area with graph connectivity.

Table 3.1: Flowfield segment distributions for each graph-connectivity [pixles].

connectivity	μ	σ	$\frac{\mu_3}{\sigma^3}$	$\frac{\mu_4}{\sigma^4}$	failures	outliers	bins
8-connected	11482	211	.0483	3.110	3	15	984
radius 3	11490	209	.0731	3.081	-	15	970




Figure 3.16: Convergence of flowfield segment area with graph connectivity.

3.7.2 The Importance of the Labeled Image

The three-class labeling scheme is shown in Figure 3.17. The labeled image \mathbf{I}_L is given in Figure 3.17b, with the freestream segment shown in black, the flowfield in white and the body in gray. Figures 3.17c and 3.17d show the segmented regions of the freestream and flowfield respectively. The flowfield segment in Figure 3.17d can serve as a measurement of the flow unsteadiness and motion patterns, directly, as its projection on the image is directly proportional to the unsteadiness of the

flowfield. This has been partially illustrated in the previous section and will be performed to a greater extent in a later chapter. Also, structures of interest such as the separation shock and the contact surface are embedded within the flowfield labeled region. Use of image labels will provide a region to search for these structures and a means of identifying them in the following chapters.




Figure 3.17: Image labeling scheme and segmentations for image $\mathbf{I}(t_o)$.

Chapter 4: Extracting Oblique Structures and Outer Contours

4.1 Motivation

The segmentation results from the previous chapter gives region labels to gross structure, which allow for Cartesian and statistical descriptors of the regions themselves; and a means to navigate through the image and enforce physical laws during search for smaller scale structures such as the separation shock. Edge detection is a technique better suited to identify smaller level details of boundaries and gradient changes in the image and sequence. In conjunction with the labeled images, edge images and the structures which can be generated from them can be used to represent the outer-shock structure and fluid structures within region labels.

This chapter focuses on representations of shock and flow structure using edge maps. A motion history map of the isolated complete outer-shock structure will be constructed using edge maps. The separation shock and the contact surface will be isolated and measured using linear segment detection. A knowledge base from physical laws and observation of the sequence will be employed to assist the measurement and classification process. An employment of a bottom-up image understanding approach is required; each task is illustrated in Figure 4.1.

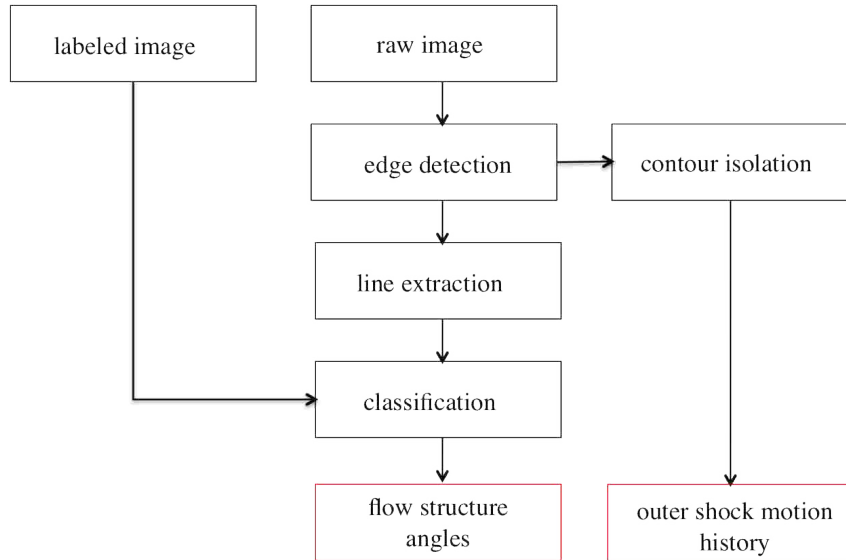


Figure 4.1: Shock structure image understanding schematic.

4.2 Edge Detection

Edge detection is the identification of object boundaries in images based on image gradients, which ideally correspond to features in the image. An edge image is a binary representation of an image which conveys information relevant to objects and their boundaries in the image. An edge map represents a simplification of an image into its most primitive content and can be considered a segmentation. Edges are classified as predominant changes in intensity gradient values; as such, their role in schlieren image analysis where flow structure is represented by changes in the density gradient is inherently applicable and valuable. It is rarely the case where boundaries or objects of interest (shock waves and flow features in the present case) are represented clearly by gradient similarity, or even discernible from image noise. This is particularly true in low SNR images such as flowfield and medical imaging. A previous section of this work discusses the ambiguity between the image noise

versus flow structure, and the role that differentiation plays in noise enhancement (see Figure 1.8). Edge detection in and of itself therefore remains a non-trivial task.

Classical edge finding techniques involve a masking approach to determine image gradients through approximation of first derivatives include the Roberts, Prewitt and Sobel edge detection methods [12], or through location of the zero-crossings of second derivatives [54]. These methods may suffer from poor localization and accuracy, produce thick edge responses, and can enhance noise and result in streaking and contour loss. In addition, some methods prefer certain angle orientations.

The Canny edge detection algorithm [15, 55] (or optimal edge detector) has largely replaced the above approaches as the method of choice for boundary detection, and is used at numerous stages for the identification of flow structures in this work. The Canny edge detector is based on the formalization of the following requirements of an ideal edge detection: a low error rate (high probability of detection with low false positives), proper localization (proximity of edge points to the true edge), and a single response criterion [15]. Through optimizing a detector for a step edge corrupted with Gaussian noise, Canny shows that the simplest function with which to detect step that most closely meets the above requirements is the derivative of the Gaussian function.

Canny’s algorithm typically employs a Gaussian smoothing step on the image \mathbf{I} on the image to suppress noise and remove high frequency content, given by

$$\mathbf{I}_\sigma = g(\sigma_I, x, y) \otimes \mathbf{I}, \quad (4.1)$$

where the separable Gaussian kernel is defined

$$g(\sigma_I, x, y) = \frac{1}{2\pi\sigma_I^2} \exp\left[-\frac{x^2 + y^2}{2\sigma_I^2}\right]. \quad (4.2)$$

Smoothing operations are typically done in each direction using the separated filters from Equation 4.2. Equation 4.1 accounts for a scale representation of the image at integration scale parameter σ_I ; the smoothed image is now a function of scale. Marr and Hildreth [54] state why the Gaussian smoothing kernel is optimal. It guarantees that no new maxima are created during smoothing, and it is the only filter that simultaneously localizes in both the spatial and frequency domains.

Gradients are calculated convolving directional derivatives with the image

$$\frac{\partial \mathbf{I}}{\partial x} = \mathbf{I}_\sigma \circledast \frac{\partial g}{\partial x} \quad (4.3)$$

$$\frac{\partial \mathbf{I}}{\partial y} = \mathbf{I}_\sigma \circledast \frac{\partial g}{\partial y} \quad (4.4)$$

where the Gaussian derivatives are computed using separability. The function used to calculate the term \mathbf{I}_x is given below in equation 4.5 (similarly for \mathbf{I}_y), where an additional differentiation scale parameter has been introduced. Typically, integration and differentiation scales are the same and will be referred to as σ_c to represent the scale of the Canny edge detector.

$$\frac{\partial g}{\partial x} = \frac{-x}{\sqrt{2\pi}\sigma_c^3} \exp\left(-\frac{x^2}{2\sigma_c^2}\right). \quad (4.5)$$

The magnitude of the image gradient and its direction of change are calculated from

$$\|\nabla\mathbf{I}\| = \sqrt{\left(\frac{\partial\mathbf{I}}{\partial x}\right)^2 + \left(\frac{\partial\mathbf{I}}{\partial y}\right)^2} \quad (4.6)$$

and

$$\theta_{\nabla\mathbf{I}} = \arctan\left(\frac{\mathbf{I}_y}{\mathbf{I}_x}\right). \quad (4.7)$$

Many classical edge detectors threshold the gradient magnitude image and employ morphological thinning to produce an edge image. Canny introduced non-maximal suppression of the gradient magnitude image to apply the single response criterion. Equations 4.6 and 4.7 allow a search for the largest gradient in the direction of change. Values of the gradient magnitude image in Equation 4.6 are retained for potential edge assignment if they are maximal along the direction of the gradient using Equations 4.6 and 4.7. The value is followed by searching in the direction perpendicular to the gradient, giving the maximum edge response in the direction perpendicular to the edge. All other pixels are set to zero. Wide ridges around the local extrema, typical of gradient magnitude images, are therefore removed.

Many edge detectors employ a single threshold on the gradient magnitude of the image, which can lead to both thick and spurious responses which are more likely to be streaked or broken. This is somewhat alleviated with Canny's use of dual thresholding. The remaining elements of the gradient magnitude contains thin lines whose values are maximal within the direction of the gradient, all of which are edge candidates. The edge map is completed using dual thresholding on the

remaining ridges in the gradient map. Hysteresis thresholding declares values in the gradient image greater than a high threshold T_H as true edges. Weak edges above a low threshold T_L and below T_H are labeled true edges if they are 8-connected to a strong edge. This process is recursive and completes the edge map, retaining completed contours from the non-maximally suppressed gradient image. Many weak or spurious responses that can be considered nuisance edges, are eliminated and weak responses that are likely to correspond to a true edge are retained.

4.2.1 Limitations and Noise

The ambiguity of the gradient direction at a corner for any edge detection method may result in the representation of sharp corners as curved corners, and the detector may fail here. Edge maps will nearly always display the stair-casing effect, which results in some degree of discontinuity and aliasing. In addition the only true lines in a discrete representation are determined by 8-connectivity, oriented at 0° , 45° and 90° . This is due to the finite number of pixels available to represent boundaries at a small scale, and may be alleviated somewhat with sub-pixel edge localization. Sub-pixel representation is not considered here, as an edge image is not the final result, rather a means to assist in feature description and detection.

Also, the use of the Gaussian kernel in Equations 4.2 and 4.5 adds scale space to the detection method. Not only do spurious responses need to be eliminated by dual thresholding, but the proper Canny scale-space parameter σ_c at which desired structure is represented must be selected. The scale at which the object of interest

occurs can change based on the size and shape of the structure (*i.e.* coarse to fine scale). Due to these considerations, experiments must be conducted and evaluated by the expert to determine meaningful edges which correspond to shocks.


4.2.2 Representation of Shock Structure with Edge Maps

Ascertaining the edge images which best represents the underlying shock structure is subject to the visual interpretation of the data by the aerodynamicist. The uncertainty principle identified by Canny in determining the ideal edge detector states an inversely proportional trade-off between localization and accuracy (detection of true edges) of the detector as a function of scale σ_c . Both a study of scale and hysteresis thresholding must be conducted to form the best shock structure in the image sequence; these factors are crucial to ensure proper representation. Edge detection was performed directly on the raw images. Image enhancement or denoising was avoided, as this can lead to removal of true image structure and the creation of false artifacts, as per the discussion in the introduction.


Experiments with edge images for three increasing detection scales σ_c of 1, 1.5, and 3 with naive hysteresis thresholding for image $\mathbf{I}(t_o)$ are shown in Figure 4.2. The top row of Figures illustrate the edge maps \mathbf{I}_e (inverted for clarity). The bottom row is the overlay of the edge image on the raw image \mathbf{I} , used to judge the accuracy of the edge response for flow structure representation. As scale parameter σ_c is increased, certain flow features become more visually apparent (for example the transmitted roughly oblique shock) as fewer false responses are detected. With

increased scale, some features become aliased or blurred (the triple point near the bow shock). The model tip becomes rounded at increased scale. At fine scale, the edge maps contain spurious responses to information in the freestream. In this case, freestream structure is not of interest and may therefore be considered noise. All images in figure 4.2 are acceptable edge images; however some contain both physically irrelevant details and poor representation of actual shock features.


Hysteresis thresholding was varied at fixed scales in order to eliminate physically uninteresting and spurious results while retaining the most salient shock features. Figure 4.3 shows the best thresholding results obtained (evaluated visually) at the same scales used in Figure 4.2. The same trend exists as in the previous figure: with increasing scale from left to right, spurious responses are reduced as localization becomes poorer. Weak edges have been removed, clarifying global structure, but along with them went actual structure. Note the complete lack of body information. It may be argued that this is a poor edge image, but an *edge image* is not of interest here; a *shock image* is desired. The triple point is best represented at a scale of $\sigma_c = 1$, becoming rounded and aliased at larger scale space representations. The same may be said of the transmitted shock. The contact surface however, although shorter at a larger scale, appears less jagged. Due to the uncertainty principal, the best localization is achieved at the smallest scale of $\sigma_c = 1$, this is also the scale most be susceptible to noise. Since this scale appears to give the best structural representation of the fluid mechanics, it is selected for further analysis.




(c) Edge map \mathbf{I}_e , $\sigma_c = 3$




(b) Edge map \mathbf{I}_e , $\sigma_c = 1.5$




(a) Edge map \mathbf{I}_e , $\sigma_c = 1$



(f) Overlay of $\sigma_c = 3$




(e) Overlay of $\sigma_c = 1.5$




(d) Overlay of $\sigma_c = 1$


Figure 4.2: Edge images \mathbf{I}_e at multiple scales σ_c with naive thresholding for image $\mathbf{I}(t_o)$.




(c) Edge map \mathbf{I}_e , $\sigma_c = 3$




(b) Edge map \mathbf{I}_e , $\sigma_c = 1.5$




(a) Edge map \mathbf{I}_e , $\sigma_c = 1$



(f) Overlay of $\sigma_c = 3$



(e) Overlay of $\sigma_c = 1.5$



(d) Overlay of $\sigma_c = 1$

Figure 4.3: Edge images \mathbf{I}_e at multiple scales σ_c with thresholding $T_L = 0.2, T_H = 0.5$ for image $\mathbf{I}(t_o)$.

4.2.3 Variation Through the Sequence

Figures 4.2 and 4.3 illustrated the existence of edges at mathematically differing levels of noise and scale. The small scale detection results ($\sigma_c = 1$) yield the most appropriate representations of shock structure; although they contain the most false responses, as expected from the uncertainty principle. However, due to non-uniform illumination from the laser, and the changing orientations of the shocks through the test time; selection of hysteresis threshold parameters must be evaluated at several images to ensure that shocks can be well defined through the sequence.

Experiments on hysteresis at several time instances are shown in Figure 4.4. Edge maps \mathbf{I}_e were obtained using a scale parameter of $\sigma_c = 1$. Each row in Figure 4.4 shows the same image, columns contain edge maps at different hysteresis threshold parameters $\mathbf{T} = [T_L, T_H]$. The leftmost column has the lowest values of high and low threshold, therefore retains the most information. Both threshold values increase from left to right. Noise is difficult to discern from structure, particularly in the region between the model and the outer-shock. Noticeable is a small structure in the upper left corner of the edge maps; also evident in the raw images \mathbf{I} and is likely a marking on the optical access window or mirror. In the middle column, considerable flow structure information has been lost between the aft-cone bow shock and the body surface. In the final column in Figure 4.4, a significant portion of the bow shock is missing in some images, and the contact surface is incomplete. These figures demonstrate that consistent flow representation using edge detection is difficult through the sequence; parameters that yield near complete information

retainment for many images may produce poor edge maps for certain images. This may be due to a number of factors, including non-uniform illumination or the partial occlusion of flow structure with a freestream schlieren. In all of the edge maps the body surface is not evident; this is due to its contour being illusory and not image gradient based. This issue was addressed with the creation of a labeled body image in the previous chapter. From the experiments conducted in the section, in order to best represent the shock structures while considering the effects of the uncertainty principle, all edge images are constructed using a scale value $\sigma_c = 1$ and hysteresis thresholding values of $T_L = 0.0$ and $T_H = 0.35$ unless otherwise stated.

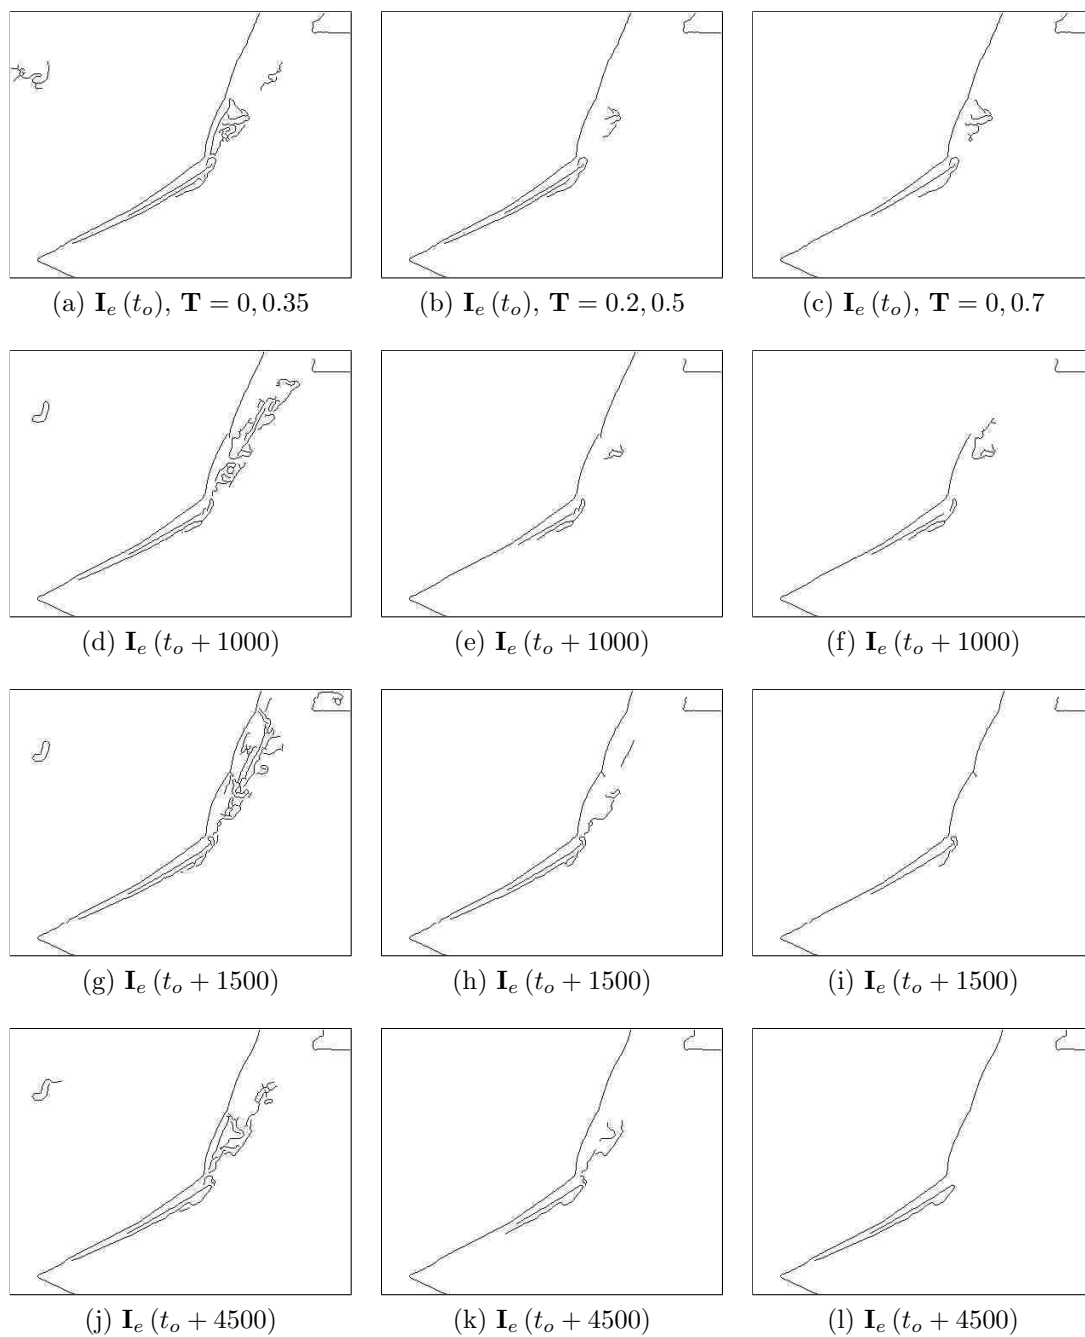


Figure 4.4: Representations edge maps at constant scale $\sigma_c = 1$ with hysteresis through the run sequence.

4.3 Isolation of the Outer-Shock

Edge maps of the global images (those shown in Figure 4.4) contain not only global and local flow structure information, but also spurious responses to strong gradients. To capture the isolated path history of the outer-shock, much of the information in the edge maps must be removed. Edge images were converted into contour images in order to remove small isolated fragments. Contours were extracted from the edge image and labeled using a portion of the corner detection algorithm developed by He and Yung [56,57]. Contour extraction from edge images preserves a more global impression of continuous edge structures. A kernel with scale parameter $\sigma_c = 1$ and aperture width of 9 was used to construct all edge maps. Contours were retained subject to a minimum arc length requirement s_{min} scaled with the image perimeter. Short contours that may not correspond to true structure were removed (in this case 40 pixels). This can remove small and disconnected edges such as those responding to strong schlieres in the freestream, and turbulent structures within the bow shock region.

To isolate the outer-shock structure, contour images were scanned with three horizontal rays. If all rays intersected a single contour, this contour was retained and a successful isolation is declared. If a single unique contour remained, it was assumed that the outer-shock was represented as an unbroken structure in the edge space. A successful isolation is shown in Figure 4.5. Spurious responses and fine detail that correspond to features other than the outer-shock in the edge map in Figure 4.5a, have been removed; these structures have labels other than the outer-shock

structure. Figure 4.6 illustrates a detected outer contour that is not representative of the shock structure. In this case, three contours were uniquely identified during row scanning process. Broken contours are due to reasons discussed earlier and may include nonuniform illumination or a freestream schlieren in the integration path.

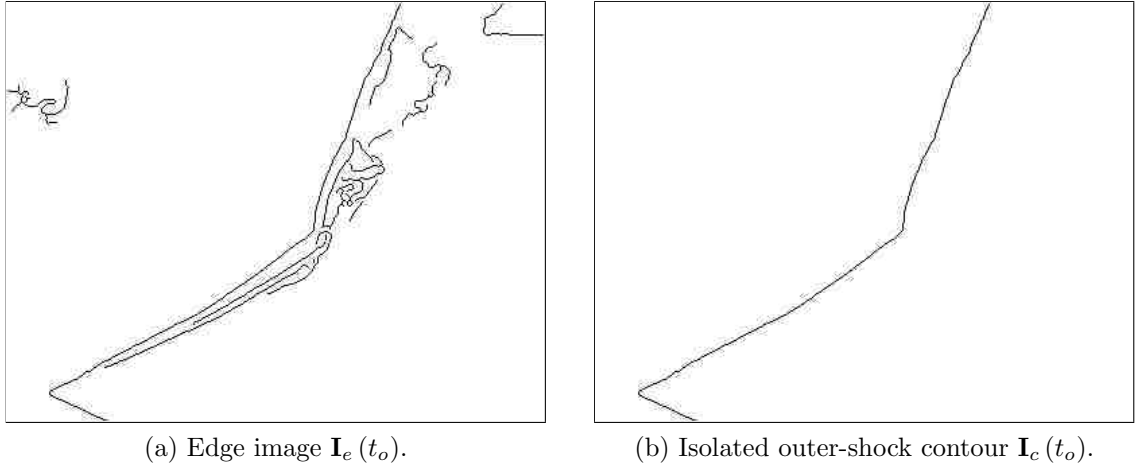


Figure 4.5: Determining a usable contour from an edge map. Fine detail and features other than the outer-shock have been removed from the edge map.

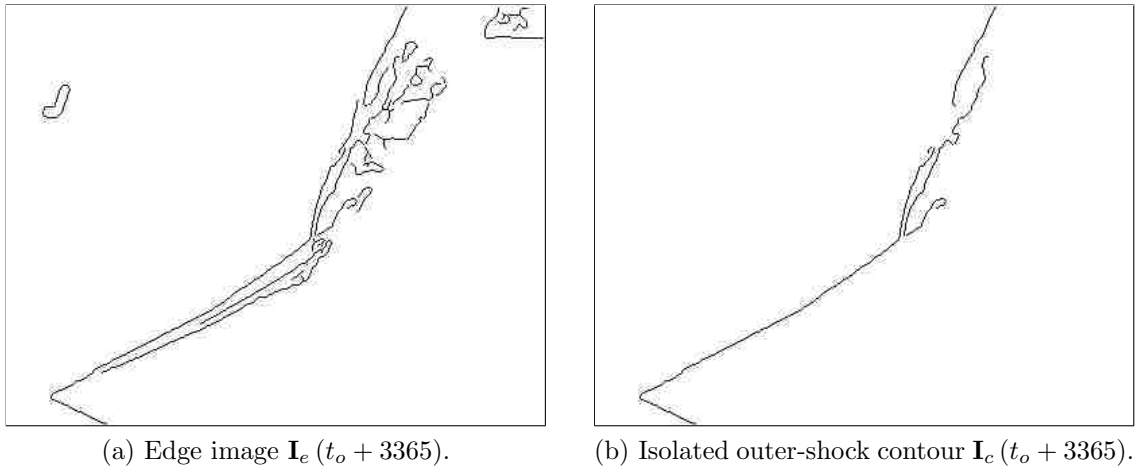


Figure 4.6: An unsuccessful contour isolation. Outer-shock is broken and consists of three contours.

4.3.1 Outer-Shock Motion History

Out of the sequence of 5,001 images, 4,487 (89.7%) were successfully isolated with a scale of $\sigma_c = 1$. For the remaining 511 images, multiple contours were returned from the ray tracing step. The process was restarted at a larger scale of $\sigma_c = 1.5$. This brought the total number of images where the outer-shock structure could be isolated to 4,912. Beyond the edge detection scale of 1.5, edges drifted from the outer-shock in the raw image. The remaining 89 images were isolated for manual measurement. Images which met this criterion were summed over the sequence to create a motion history of the global outer-shock structure.

The resulting summation of unbroken isolated outer contours is shown in Figure 4.7, a binary image that illustrates the path history of the complete outer-shock structure. A pixel is labeled black if the outer-shock has occupied that location during the sequence, and white otherwise. Significant motion and shape change of the upper bow shock is evident from the thickness of the region. The region containing the separation shock seems to indicate small pitching motion of the shock; however the motion history also includes any translation and length change if present. This illustration of unsteady behavior was evident from viewing the raw image sequence; and a computational solution did not converge to a steady state for the test conditions [37]. The knot structure near the fore-cone tip in Figure 4.7 is not representative of shock motion. The edge maps were unable to detect the full attached conical shock on the fore-cone tip, due to insufficient image evidence (gradients) in this region. This can be seen by the human eye on close examination of Figure 3.1




Figure 4.7: Motion history of isolated outer-shock structure

4.3.2 Instances of False Shock Structure

During the summation of the isolated contours, several false representations of shock structure were evident from visual inspection of Figure 4.7. False contours accounted for 15 out of the 4,912 contours which were successfully isolated, and were removed from the motion history shown in Figure 4.7. Examples of false contours are shown in Figure 4.8. The shock intersection structure is represented by a loop in Figures 4.8a and 4.8b. The upper portion of the bow shock turns back toward itself in Figure 4.8c. In Figure 4.8d, a sharp corner protrudes from the middle portion of the shock structure. In all instances in Figure 4.8, the outer-shock representation violates conservation of mass. These representations may be due to a freestream schlieren which partially occludes the true shock shape, or image gradients of flow structures downstream of the shock that are larger than the outer-shock.

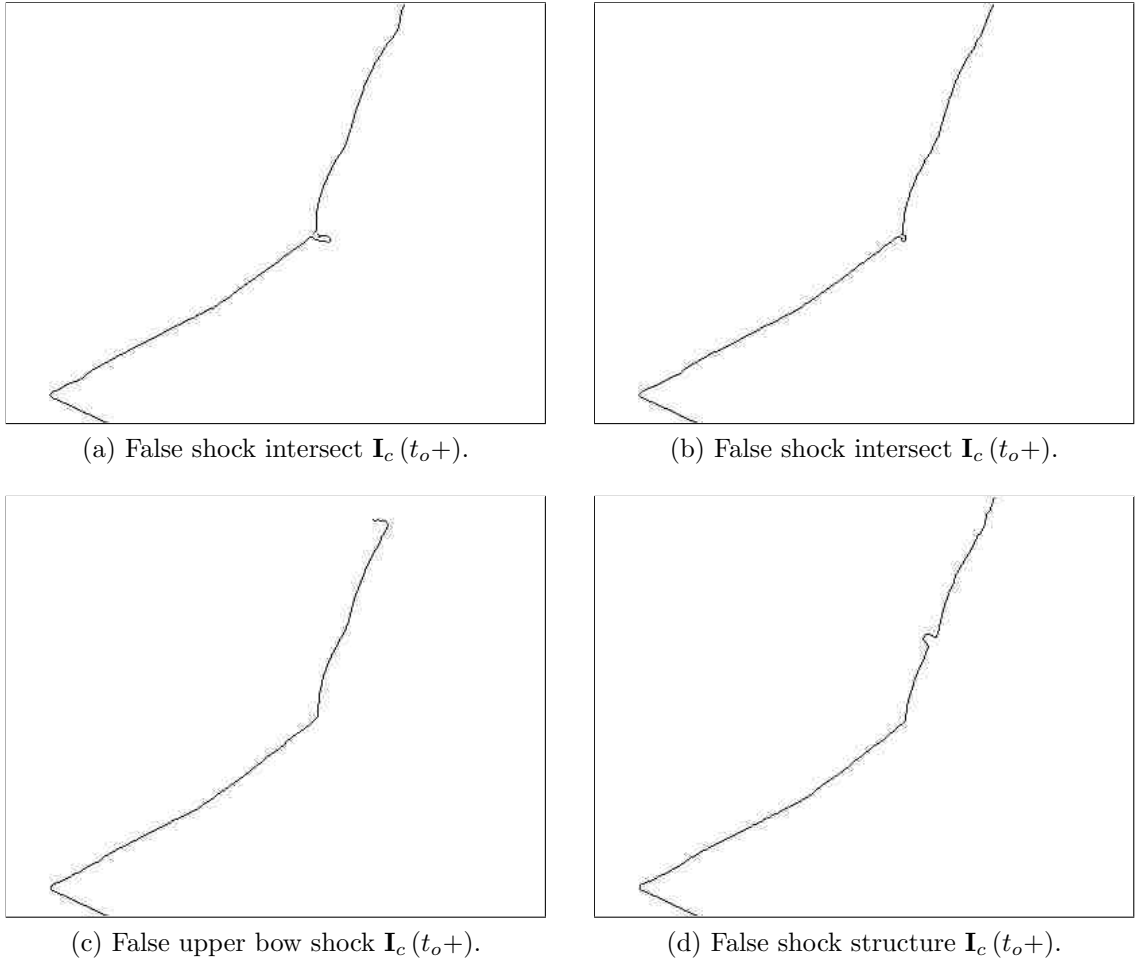


Figure 4.8: Contour images \mathbf{I}_c indicating false structure.

4.4 Line Fitting

Once an edge map of an image has been constructed, the edge points must be assembled into more useful representations of aerodynamic interest, such as lines representing shock waves, from which desirable physical properties can be extracted. Cues from computational studies of the double-cone [41], as well as visual inspection of the schlieren images themselves suggest that both the transmitted separation shock and the contact surface may be well characterized by linear segments. A slight deviation from linearity is observable at close scrutiny in certain individual

frames; however, any apparent curvature is at such a small scale in the image that it appears indiscernible from noise associated with discrete representation of a digital image. It is assumed that these structures can be appropriately modeled as linear.

Techniques for line detection and fitting edge pixels include least squares fitting, M-estimators, and random sample consensus (RANSAC) [58], among others. These methods are best suited when the line is well defined, or when the points the line contains are known *a priori*. When searching for a line that may undergo translation or pitching, or at a scale where line endpoints are difficult to discern, the Hough transform [18, 19] provides a better estimate for the impression of a line.

4.4.1 The Hough Transform

The Hough transform is a statistical binning method used to cluster collinear or near collinear nearby points into the most likely linear fit. Binary edge maps, for reasons discussed in section 4.2.2 rarely characterize even the most simple linear segment. The Hough transform is a means to link broken or even mis-characterized pixel segments into meaningful lines. The transform fits binary image points to the normal representation of a line. A point (x, y) in the image axis plane is transformed to a parameter space (r, θ) given by

$$x \cos(\theta) + y \sin(\theta) = r, \tag{4.8}$$

where the angle is bound by $\theta \in [-\pi/2, \pi/2]$ and the normal distance from the

origin is bound by $r \in [-D, D]$, where D is the diagonal of the image. Equation 4.4.1 allows for the representation of vertical and horizontal segments. Collinear points in the image plane will fit a line given by 4.4.1, with constant r and θ .

The Hough transform technique is illustrated in Figure 4.9 from reference [32]. Two points (x_i, y_i) and (x_j, y_j) are collinear in the the image plane in Figure 4.9a. The two points share the same perpendicular distance from the line to the image origin r , and are both inclined at the same angle θ , and described by the same normal representation in Equation 4.4.1. In the $r\theta$ parameter space plane in Figure 4.9b, all possible lines passing through a given non-background point are represented as sinusoids. Each sinusoid in the parameter space represents all possible lines through a fixed point in the image plane, given by Equation 4.4.1. The two sinusoids shown are all lines going through points (x_i, y_i) and (x_j, y_j) . Each point in the parameter space plane represents a line in the image plane. The point where the two sinusoids intersect at (r', θ') represents the straight line in Figure 4.9a. Figure 4.9c shows the discretization of the $r\theta$ space into a two-dimensional accumulator bin. For each intersection point in Figure 4.9b, the corresponding accumulator cell in the transform plane is incremented. Each vote in the accumulator represents a corresponding line in the image. The accumulator cells with the highest counts (peaks in the transform space) represent the line with the most number of collinear points, and therefore the most statistically likely and dominant lines in the image.

After the transform space has been discretized, non-maximal suppression is performed on the accumulator array $\mathbb{H}(\Delta r, \Delta \theta)$. Peaks in the resulting operation represent potential line candidates in the image.

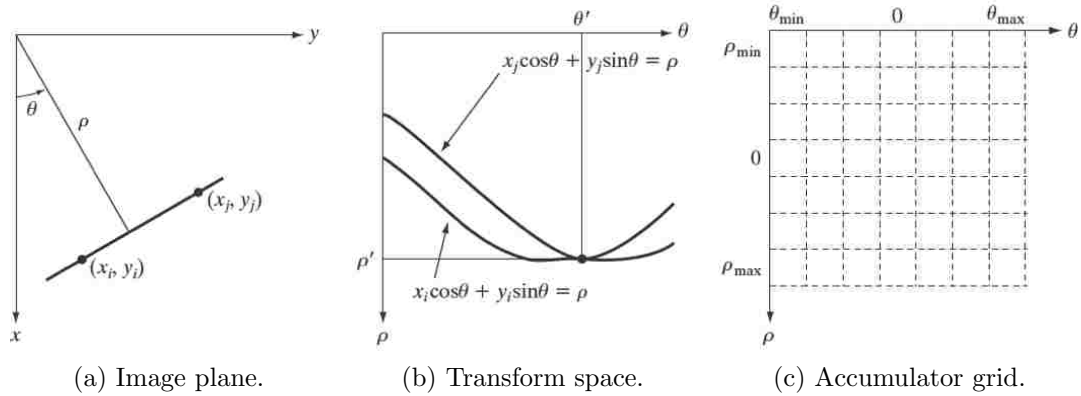


Figure 4.9: The Hough transform parameter space [32].

4.4.2 Grid Sensitivity Studies

Difficulties associated with Hough transform peak detection include the role of noise in the image, and the proper grid size selection of the binning parameters $[\Delta r, \Delta \theta]$. Transform parameters must be determined through experimentation. If the transform grid of the $r\theta$ plane is too coarse, the line resulting line is largely useless, as points well outside of the desired line will be included in the voting plane. If the plane is too fine, the collinearity of points will be artificially decreased. Proper discretization is non-trivial: lines in the image plane may be clustered together on a coarse grid, and too fine a grid will be unable to capture a jagged line (a likely representation for many line segments), restricting the definition of near collinear points. A fine grid will also collect votes for meaningless edge fragments. A study of the $r\theta$ accumulator space is necessary to see how fine a scale the lines can be reliably detected. Accuracy of collinearity is directly determined by grid sizing. For computational simplicity and the incorporation of domain knowledge, only the corner Region Of Interest (ROI) will be searched for linear segments that represent

the shock wave and the contact surface. Figure 4.10 shows the corner region of the image at instance $(t_o + 1500)$ without pre-processing, and its edge image (inverted for clarity). The ROI images in Figure 4.10 are 88 x 128 pixels.





Figure 4.10: Corner ROI of image $\mathbf{I}(t_o + 1500)$.

Grid sensitivity studies are performed on the binary image in Figure 4.10b. Three transform resolution spaces are tested with the following $[\Delta r, \Delta\theta]$ resolutions: $[2, 2^\circ]$, $[2, .125^\circ]$ and $[1, .125^\circ]$; the transform bins are plotted in Figures 4.11a, 4.11b and 4.11c respectively. Each sinusoid in these Figures represent all linear fits through a given point in the edge map in Figure 4.10b. Dark values indicate large vote counts for a line fit. High accumulator bin counts indicate a larger the number of collinear points on the line corresponding to Equation 4.4.1 at the (r', θ') locations in Figure 4.11. Note that θ is plotted with respect to the transform axis. Scale is decreasing from left to right in Figure 4.11; the coarsest accumulation grid is in Figure 4.11a where the peaks in the transform are more closely grouped. At the finest scale in Figure 4.11c, the intersection of sinusoids are the clearest, although the fine division shows a looser grouping of the transform peaks. This suggests that single lines

corresponding to the peaks in Figure 4.11a are subdivided into finer segments in Figure 4.11c.

The four largest peaks in the transform planes in Figure 4.11 were found after non-maximal suppression of the grid, and are listed in Table 4.1; listing the vote, inclination θ' (with respect to lab reference) and perpendicular distance to origin r' for their respective transform grid plots above. The separation shock and contact surface angles in Figure 4.10b measure approximately 36° and 31° respectively; similar values are seen in some of the tables. The strongest line persists through cases I and II in Tables 4.1a and 4.1b, within the $\Delta\theta$ discretization differences, and has the same r' and number of votes. However, at the finest scale in Table 4.1c, this line does not appear to have been captured. Instead, the strongest peak in Table 4.1c seems to correspond to the second strongest peak in Tables 4.1a and 4.1b, although it has received considerably less votes. The variation in detected peaks with grid sensitivity suggest that a better method to assure feature capture than a single grid would be to compare multiple grids.



(a) $\Delta r = 2, \Delta \theta = 2^\circ$

(b) $\Delta r = 2, \Delta \theta = .125^\circ$

(c) $\Delta r = 1, \Delta \theta = .125^\circ$

Figure 4.11: Hough transform accumulator grid sensitivities for corner ROI $\mathbf{I}_c(t_o + 1500)$.

Table 4.1: Hough transform peak votes for corner ROI $\mathbf{I}_c(t_o + 1500)$.

(a) $\Delta r = 2, \Delta \theta = 2^\circ$				(b) $\Delta r = 2, \Delta \theta = .125^\circ$				(c) $\Delta r = 1, \Delta \theta = .125^\circ$			
Peak	Vote	θ' [$^\circ$]	r' [pixels]	Peak	Vote	θ' [$^\circ$]	r' [pixels]	Peak	Vote	θ' [$^\circ$]	r' [pixels]
1	87	34	70	1	87	34.125	70	1	66	30.750	77
2	77	32	78	2	78	32.375	78	2	62	35.500	71
3	55	28	72	3	75	29.625	76	3	55	28.625	81
4	44	36	82	4	71	30.000	82	4	49	32.875	70

4.5 Detecting Shock Wave and Contact Surface Angles

To ensure a shock wave or shear layer is extracted as opposed to a random edge segment, insight into why a particular representation is chosen must be built on image evidence and physical knowledge. For example, the oblique shock is best represented as a straight line segment, of some specific length, it is unlikely or physically impossible to take on certain angle values, and must lie in a certain region of the flow. It is known in this case of this specific sequence that the shock wave must be above the contact surface, which in turn must be above the body. This would be an acceptable classification if we knew that only these lines could be detected in each image in the sequence, or that multiple line segments would not be returned to represent each line.

To differentiate linear segments from nuisance lines, and shock waves from contact surfaces, a classifier must be built to make sure the extracted line is the flow structure of interest. This can be done by extracting all linear segments which are potential candidates for these flow features, and applying physical rules which these segments must obey before they can be classified as a shock wave or a random segment. A scale-space representation must be selected that accurately constructs edge maps which correspond to the flow structure in the images. Linear segments are then extracted from these edge maps using the Hough Transform technique. Finally, a rule-based classification scheme is developed based on observations from the particular sequence. An enforcement of physical rules using the segmentation scheme is employed.

4.5.1 Scale Selection for Shock Representation

Arbitrary edge detection on a schlieren image will not lead to a useful representation of shock structure. Ensuring that the detected edges correspond to flow structure requires both a proper scale selection, knowledge of fluid mechanics, and evaluation by the researcher. Scale selection is a trade-off between localization and accuracy, and proper selection is non-trivial. To find the appropriate scale for the edge map, a raw primal sketch, similar to that employed by Marr and Hildreth [54] was used. Multiple representations of the flow structure are potentially available; built from the raw image or the diffused image, each available at a different scale.

Two raw primal sketches are shown in Figure 4.12: Figure 4.12c presents a combined representation of scale $\sigma_c = 1$ of the raw image shown in Figure 4.12a and its diffused counterpart in Figure 4.12b; Figure 4.12d shows edge maps built from a raw image using the Canny scale representation at two scales $\sigma_c = 1$, and 2 (a single octave). Black pixels represent the logical **AND** operation between representations; gray pixels denote the logical **OR** of detected edges. The separation shock at first glance seems to overlap well. However, the ambiguity of the endpoints of the shock has grouped pixels to shocks on either side of the structure and aliased them together at a larger scale. Considerable difference is evident between representations, the shock becoming curved and aliased at larger scale, while the contact surface is entirely different. Although the diffused image has better grouped the image globally, the composite sketches do not exactly coincide. The trade-off between accuracy and localization is a factor in many vision implementations. In this case, it is best

to choose the smallest scale at which the desired structure is represented. Although noise is more prevalent at the smaller scale, better localization is attained. The raw image ROI is therefore used for all edge maps. For the edge images shown in this research, a Gaussian kernel of unit scale with support width of 9 was used.

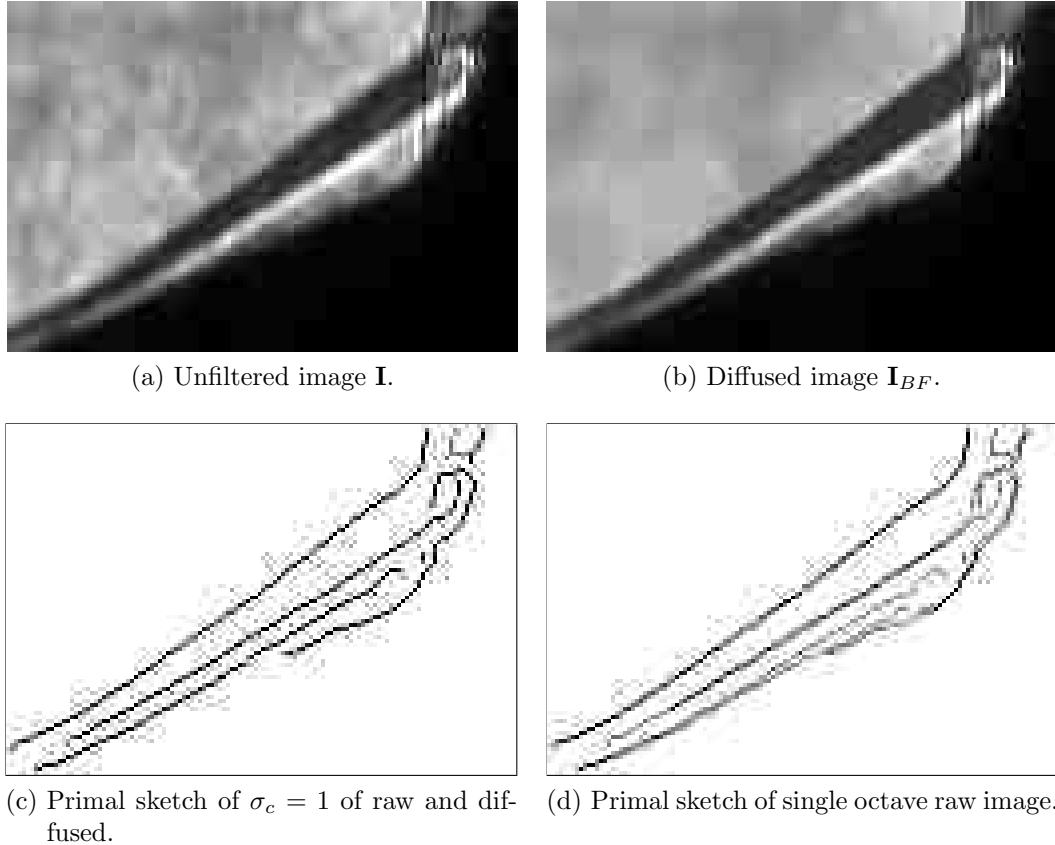


Figure 4.12: Scale selection and raw primal sketches for Corner ROI $\mathbf{I}(t_o)$.

4.5.2 Knowledge Representation for Structure Definitions

Methods for edge detection and extraction of lines in an image have been shown. Expert knowledge must be explicitly incorporated into the a detection algorithm to ensure that a separation shock can be distinguished from the contact

surface, and each from fragments or nuisance lines. Physical knowledge can be represented as machine instructions, allowing the algorithm to differentiate between features and be more likely to return the desired line segment. The separation shock and the contact surface are drawn in Figure 4.13a, its corresponding labeled region, which can be used to enforce a knowledge base is shown in Figure 4.13b.




Figure 4.13: Illustration of knowledge base incorporation.

To develop the classification schemes, both observations of the sequence and expert knowledge need to be incorporated into the decision-making process. Both the separation shock and contact surface are visibly dominant lines in the edge images. Throughout the sequence, it is physically impossible for the angle of the separation shock to go above or below certain upper or lower bounds (θ_u and θ_l); likewise for the contact surface. For a line to be a shock wave, it must be immediately below the freestream region and immediately above the flowfield, a rule which can be enforced by the segmentation shown in Figure 4.13b. Similarly, the contact surface must be below the shock and therefore must be labeled both above and below by the flowfield. The image coordinate system is used here, a right hand coordinate

system with the origin in the upper left hand corner of the image. The following statements can therefore form the knowledge base, which is later incorporated into the detection algorithm:

1. Both structures may be approximated by a line.
2. Both structures are among the most dominant and longest lines in the ROI.
3. Angles are constrained; they are unlikely to take on values outside $\theta \in [\theta_l, \theta_u]$.
4. Known location: a point several pixels above the shock wave will be labeled freestream; pixels immediately above and below the contact surface must be labeled flowfield.
5. Structures should be the longest remaining line after constraints applied.

4.5.3 Algorithm for Measuring Oblique Features

The algorithm used to extract and measure the separation shock β_s and contact surface θ_{cs} angles is listed in Algorithm 2; this is a bottom-up image understanding model as depicted in Figure 4.1. When measuring the angle of the line, it is assumed that the camera array is aligned with the model axis, providing the datum. Each feature is considered separately. Potential failures of the algorithm are listed in the order in which they may be encountered. If at any point in the sequence, one of the failure modes above has been encountered, the algorithm records the mechanism and continues to the next image. The first step of the extraction of oblique features is the identification of linear segments. Given the raw image ROI **I**

shown in Figure 4.13a and its corresponding label image \mathbf{I}_L in Figure 4.13b, a fixed Canny scale $\sigma_c = 1$, hysteresis threshold parameters T_L and T_H , and a fixed Hough grid resolution $[\Delta r, \Delta\theta]$, the following steps are performed:

- 1 Canny edge detection on the raw image \mathbf{I} at scale σ_c to create edge image \mathbf{I}_e .
- 2 Hough transformation of \mathbf{I}_e to form the transformation matrix $\mathbb{H}(\Delta r, \Delta\theta)$.
- 3 Extract k_1 linear segments from \mathbb{H} corresponding to large vote counts.

These first steps provide anywhere from two to six linear segments for the test sequence used here, some or none of which are viable candidates to represent the angle of both the shock wave and contact surface. Considering line segments which correspond to large vote counts in the Hough transform accumulator bin encourage that the more visually evident lines in the edge image are returned. No steps have been taken yet that specify the detection algorithm to the double-cone geometry.

To eliminate nuisance linear segments and test for candidate viability for a physical representation, the remaining portion of the algorithm is the implementation of the knowledge base that allows to determine which of the remaining line segments (if any) is a reasonable representation of the structure of interest, and is a rule-based classification. Since both the separation shock and contact surface are visibly dominant in the edge images, elimination of some false positives was accomplished by selecting large vote counts in the Hough transform accumulator bin in step 3 of the line finding algorithm. Considering only line segments greater than a certain length l_{min} for candidates can further restrict the selection. Potential failure

methods and locations are listed in the algorithm below, which continues where the line extraction algorithm left off:

4 Remove lines whose length $l[k_1] < l_{min}$ giving k_2 segments.

- return line failure if $k_2 = \emptyset$.

failure mode 1; line segments are too short to represent either structure.

5 Remove lines whose angles $\theta_l \leq \theta[k_2] \leq \theta_u$ giving k_3 segments.

- return angle failure if $k_3 = \emptyset$.

failure mode 2; no segment is within the expected angle bounds, meaning the model surface or a nuisance fragment has been identified.

6 Calculate the midpoint locations (x_m, y_m) of the remaining k_3 lines.

a To extract the separation shock angle, check the label immediately above the midpoint, removing lines that do not satisfy the label check $\mathbf{I}_L[x_m - \Delta x, y_m] = fs$, returning k_4 segments.

b To extract the contact surface angle, check the label above and below the midpoint, removing lines that do not satisfy the label check $\mathbf{I}_L[x_m - \Delta x, y_m] = ff$ and $\mathbf{I}_L[x_m + \Delta x, y_m] = ff$, returning k_4 segments.

- return region failure if $k_4 = \emptyset$.

failure mode 3; no lines exist which satisfy the location requirements.

7 If more than a single line remains, extract the longest remaining. This is likely the best fit as it incorporates a larger number of points in the edge map, while maintaining a high accumulator vote.

The only information specific to the application of the algorithm to the double-cone sequence is the particular segmentation scheme and knowledge base, which gives relative structure location, as well as angle bounds for line segments. Identification of failure modes is critical to the successful employment of the of the knowledge base used to extract the angle; this is an implementation of decision making.

Algorithm 2: Determination of separation shock and contact surface angles.

Input: \mathbf{I} raw image ROI
 \mathbf{I}_L labeled image ROI
 $\sigma_c, [T_L, T_H]$ Canny scale and hysteresis thresholds
 $\Delta r, \Delta \theta$ fixed Hough accumulator resolution
 l_{min} minimum allowed line length
 θ_u, θ_l upper and lower angle bounds

Output: separation shock $\beta_s [i]$ or contact surface angle $\theta_{cs} [i]$ for each image

```

for  $i = [t_o, \dots t_f]$  do
  // Line extraction
  Canny Edge Detection of  $\mathbf{I}$  to form edge image  $\mathbf{I}_e$  ;
  Hough transformation of  $\mathbf{I}_e$  to form  $\mathbb{H}(\Delta\theta, \Delta r)$  ;
  Extract  $k_1$  linear segments  $l$  from transform peaks ;
  // Structure identification
  remove lines whose length  $l[k_1] < l_{min}$ , giving  $k_2$  segments;
  FailCheck( $k_2, 1$ ) ;
  remove lines whose angles  $\theta_l \leq \theta[k] \leq \theta_u$ , giving  $k_3$  segments;
  FailCheck( $k_3, 2$ ) ;
  Enforce location constraints;
  LocationCheck( $k_3, \mathbf{I}_L$ ) ;
  FailCheck( $k_4, 3$ ) ;
  return angle of the longest remaining line  $\beta_s [i]$  or  $\theta_{cs} [i]$  ;
end
// Function to find and identify failure mode
FailCheck(line set  $k$ , mode  $j$ ) ;
if  $k = \{\emptyset\}$  then
  fail[ $i$ ] = true;
  failmode =  $j$  ;
  continue to next  $i$ ;
end
// Function to evaluate labels
LocationCheck(line set  $k_3, \mathbf{I}_L$ ) ;
calculate midpoints of all linear segments  $(x_m, y_m)$  ;
if measuring  $\beta_s$  then
  retain only lines whose midpoints are immediately above the freestream,
   $\mathbf{I}_L[x_m - \Delta x, y_m] = fs$ ;
  return  $k_4$  segments ;
else
  measuring  $\theta_{cs}$  ;
  retain only lines whose midpoints are immersed in the flowfield
   $\mathbf{I}_L[x_m - \Delta x, y_m] = ff \wedge \mathbf{I}_L[x_m + \Delta x, y_m] = ff$ ;
  return  $k_4$  segments
end

```

4.6 Results for the Shock and Contact Surface Angle

A partial indication of the usefulness of the extraction technique to capture the line is to study how finely the angle can be resolved using the Hough transformation. The transform grid studies in Figure 4.11 and Table 4.1 illustrated sensitivities to line measurements of the structures in the ROI. It is therefore desired to measure as fine an angle possible subject to minimizing the failure rate of the detection algorithm.

Due to the small size of the ROI, the pixel bin spacing Δr in the Hough transform space was held constant at 2. The angular resolution $\Delta\theta$ was varied between 0.125, 0.25 and 0.5°. These are listed as cases I, II and III respectively in Tables 4.2 and 4.3. Edge detection was performed at scale $\sigma_c = 1$ for all grid resolutions. The detection algorithm developed in the previous sections was applied to a sequence of 5,000 images for each grid resolution case, with success rates between 98% and 99%. From these tables, the effects of grid resolution have little effect on the mean μ and standard deviation σ_{std} of the data spread which remain reasonably constant. The shock angle distribution has a mean μ of approximately 36° and standard deviation σ_{std} close to 1°, suggestive of a fairly tight grouping over of the angle value through the test time. The contact surface angle has a mean value of approximately 31° with nearly the same standard deviation of the shock wave angle. However, the number of discrete bins is too small for any analysis over time at lower grid resolutions. Case I in Tables 4.2 and 4.3 illustrate that the finest resolution offers the largest number of discrete angle values, and has the lowest failure rate

for each angle measurement. Increasing failure rate suggests that over too coarse a grid, the angles possess a smaller statistical accumulator in the Hough transform space and the line is missed. Statistical outliers were defined using the $3\sigma_{std}$ rule, and are few.

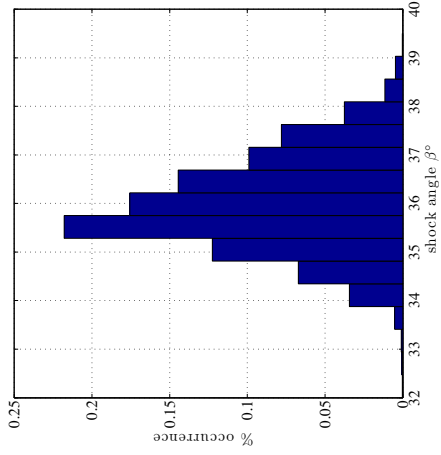
Table 4.2: Test cases for shock angle β_s .

case	σ_c	$\mathbb{H}[\Delta\theta^\circ, \Delta r]$	failures	$\mu(\beta_s)^\circ$	$\sigma(\beta_s)^\circ$	outliers	$\Delta\theta$ bins
I	1.0	[0.125, 2]	42	36.026	1.033	6	54
II	1.0	[0.25, 2]	53	36.007	1.031	5	30
III	1.0	[0.5, 2]	69	35.987	1.047	5	16

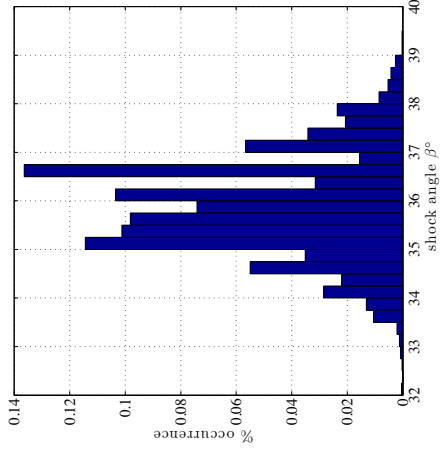
Table 4.3: Test cases for contact surface angle θ_{cs} .

case	σ_c	$\mathbb{H}[\Delta\theta^\circ, \Delta r]$	failures	$\mu(\theta_{cs})^\circ$	$\sigma(\theta_{cs})^\circ$	outliers	$\Delta\theta$ bins
I	1.0	[0.125, 2]	53	30.977	1.010	12	53
II	1.0	[0.25, 2]	64	30.955	1.024	22	30
III	1.0	[0.5, 2]	87	30.919	1.024	35	15


The histograms for the shock angles at the different cases of $\Delta\theta$ divisions of the probabilistic transform bins tested are shown in Figures 4.14 and 4.15 for the shock and contact surface angles respectively. All distributions shown are relatively normal for both flow structures across all grid resolutions. This preservation of the distribution over grid refinement demonstrates consistent performance of the extraction algorithm irrespective of Hough transform resolution. However, the finest angle resolution is desired for a quantitative analysis. The persistent distribution of angle values is suggestive of motion patterns grouped tightly about a mean value of the flow structures through the sequence.



(a) case I: $\Delta\theta = 0.125^\circ$




(b) case II: $\Delta\theta = 0.25^\circ$

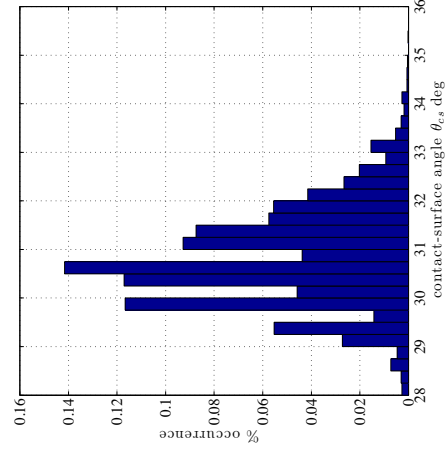


(c) case III: $\Delta\theta = 0.5^\circ$


Figure 4.14: Histograms of detected shock angle β_s across grid resolution.



(a) case I: $\Delta\theta = 0.125^\circ$



(b) case II: $\Delta\theta = 0.25^\circ$



(c) case III: $\Delta\theta = 0.5^\circ$

Figure 4.15: Histograms of detected contact surface angle θ_{cs} across grid resolution.

4.6.1 Analysis of Detected Failure Modes

The total number of failures for shock and contact surface angle measurements are given for their respective cases show success rates ranging from 98.3% (contact surface angle coarse grid) to 99.16% (shock wave angle fine grid) over the sequence of 5,000 frames. A break down of the failures listed in Tables 4.2 and 4.3 are given in Tables 4.4 and 4.5, respectively. Failures were identified during the classifier portion of the detection algorithm, and were defined when either lines extracted were not sufficiently long to represent a flow structure, no lines had angle values within the expected bounds of the flow structure, or a line segment was not located in the proper physical place in order to be a shock wave or contact surface. The most common cause of failure was incorrect location for proper classification.

Failures can always be measured manually when they are few. Detection failures may be caused by a multitude of factors including the tight solution bounds imposed by the algorithm, and the potential alignment with a freestream disturbance (turbulent eddy) in a manner which partially occludes the shock structure. This can inhibit the edge map from attaining a reasonable representation of the features of interest.

Table 4.4: Shock failure modes.

case	Lines	Angles	Region	Total
I	3	4	35	42
II	6	4	43	53
III	9	6	54	69

Table 4.5: Contact surface failure modes.

case	Lines	Angles	Region	Total
I	1	0	52	53
II	2	0	62	64
III	2	0	85	87

4.6.2 Combining Evidence Across Grids and Convergence

The high success rates for the grid studies do not guarantee a converged or accurate solution. Due to the size of the structures in the image sequence, a search across scale space for measurement verification was not possible. The lines quickly became curves before the smallest edge detection scale σ_c was doubled, as illustrated by the raw-primal sketches shown in Figure 4.12. Likewise, a coarse-to-fine pyramidal downsampling using the methods of Burt and Adelson [35] provided images which were too small for non-aliased representation. Either of these methods would be a typical approach in the computer vision field to verify an observation given higher resolution images, or if the structure existed in a larger ROI than in the present case of 88 x 128 pixels.

Without persistence across scale in the traditional sense to verify measurements, an approach varying the Hough transform grid resolution was applied using the $\Delta\theta$ bin spacing. Convergence was examined across the two finest grids, cases I and II for both the shock and contact surface angle. A measure of bin distance in the histograms in each case was used as a convergence criterion as twice the accumulator width of $\Delta\theta$ for case II minus $\Delta\theta$ for case I,

$$\epsilon \leq (2\Delta\theta_{II} - \Delta\theta_I). \quad (4.9)$$

This convergence criterion is based on the inherent uncertainty of the accumulator bin divisions in the Hough transform itself, and by the grid spacing defined in Tables 4.2 and 4.3, $\epsilon = \pm 0.375^\circ$.

Results illustrating convergence across Hough grid resolutions for cases I and II are shown in Figures 4.16 and 4.17, for the shock and contact surface angles respectively. In addition to structure detection failures noted in the previous section, there are an additional 109 shock angle values which lie outside of the convergence criterion ϵ , and an additional 229 for the contact surface angle value. Over the sequence, these non-converged values represent less than 2.2% and 4.6% of the images for the shock and shear layer respectively.

To obtain a final measurement, an indirect Bayesian approach was employed to verify observations at the fine grid resolution. The observation from case I is accepted, subject to the constraint that the angle is within $\pm\epsilon$ of of the coarser case II observation at the same time instant. This means that any failures in either case I or II cannot be considered a valid observation if a value is not observed at both scales. Given these measurements, the remaining values from the fine scale observation are removed if they are statistical outliers, defined by the $3\sigma_{std}$ rule. After removal of failure modes, statistical outliers and non-convergent results, the successful detection rate for the shock and contact surface angles are nearly 97% and 94% respectively.

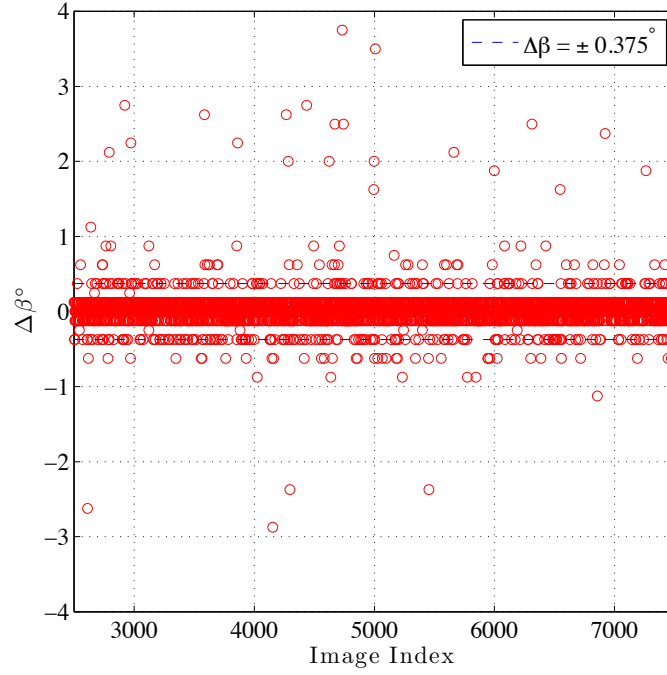


Figure 4.16: Shock angle convergence across $\Delta\theta$ spacing.




Figure 4.17: Contact surface angle convergence across $\Delta\theta$ spacing.

Chapter 5: Feature Tracking of Large Scale Interest Points

5.1 Motivation

Segmentation, the extraction of oblique structures, and outer-shock structure motion of the double-cone sequence have been performed in previous chapters. This portion of the research focuses on the identification and tracking of the model tip and the shock intersection point in the sequence using feature identification and tracking concepts from vision. From inspection of the sequence, the shock intersection point motion appears to be one of the dominant unsteady features of the flowfield. Tracking through feature identification can provide a trajectory and motion history of the structure throughout the test duration. In addition, the location of the triple point can provide a location at which the segmented flowfield can be split and further analyzed for spectral and stochastic relationships of and between the full, fore, and aft-cone flow regions.

Tracking the model cone tip can provide an estimate of model motion, if present. Model motion during wind tunnel test time can be crucial in extracting meaningful conclusions from image or sensor data. Nominally still models in supersonic wind tunnel are prone to experiencing large forces due to shock pressure rise, and vibration may occur. Estimation of vibration or illustrating model stillness,

in this particular data set, can help to isolate reasons for the unsteadiness in the double-cone flowfield. As the model is axisymmetric and mounted on a stiff sting, any motion caused by vibration will be largest at the tip.

5.2 Feature Concepts

A feature is defined as a point or region in an image which is perceptually interesting. Also referred to as corners or interest points, these features should be unique, and account for some definition of saliency so that they may be identified and potentially tracked through an image sequence. From a vision perspective, a feature or corner should be unique, well defined, tractable and meaningful. These properties are well suited for identification and tracking these features between frames. The concept of saliency, and the interpretation of “meaningful” is largely mathematical in this sense. False structures can also be detected along with true features. Even in the case of edge labeling, spurious and unphysical artifacts can be returned in addition to actual edges in simple image scenes.

The definition of a point of interest or desirable feature can change from application to application. For example, in the present work, interest is in the model tip and the shock intersection point. In other data sets such as boundary layer images, structures of interest may include a transition region, a separation point, or the movement of a turbulent structure as it is convected downstream. There is no unifying theme, physical or mathematical, of these potentially interesting points. Their only common factor is that they are physically salient, and of interest to their

respective studies. From an aerodynamics perspective, a feature, in addition to the above considerations, must be useful, or interesting from a physical standpoint. Saliency must also account for aerodynamic definitions. The concepts of feature-ness or cornerness in vision are not immediately reconcilable with those from a fluid dynamics perspective. Even the most advanced feature identification algorithms in vision, such as the Scale Invariant Feature Transform, or SIFT [59] can provide structures that are undesirable for any analysis; may not correspond to anything of interest or what is expected from a fluid dynamics perspective; are matched by chance, or at worst, correspond to false structure.

Schlieren images, like ultrasound images, contain structures that appear as “white noise” or “sand”, particularly when system sensitivity is high. Images can be dominated by textures that are physically uninteresting, for example the large freestream regions. Although these correspond to true structure, from a human standpoint, they are considered “noise”. Many feature or corner detectors rely implicitly on texture and regions of rapid change in the image through the use of differential operators. These regions can be the most rapidly changing features in the data. Most corner or feature detection algorithms may return a physically interesting point, but they are guaranteed to return most responses from the texture rich regions which can dominate the images. Responses are highly sensitive to freestream texture patterns and other “noise”, few of which (if any) are meaningful. Those that are returned often have no physical basis or objective metric by which they can describe a flow feature, other than human evaluation. Uniqueness is difficult to define, as mathematical scores of the same feature can vary drastically within the sequence.

5.2.1 Scale in Features

Feature or interest point detection schemes in vision are heavily reliant on the concept of scale. Persistence through, and existence at, scale as discussed in previous chapters, is of paramount importance in vision. This is particularly true in feature detection algorithms where search across scales is common. For example, an integration scale results from the convolution of an image with a Gaussian kernel with scale parameter σ . The convolution of an image with a Gaussian derivative provides a differentiation scale.

Scale space filtering of a signal was first proposed by Witkin [60] as a means of examining and inferring structure at multiple scales. Through successive levels of signal smoothing, the number of signal features, for example local extrema or zero-crossings, will gradually diminish. At a large level of smoothing (large scale), only true features of the signal will remain. Those due to noise will not persist through increased smoothing. This can be applied to employ a coarse to fine search of such features by identifying structure at a coarse scale (large amount of smoothing), and searching through a finer scale (less smoothing) to update location. This can avoid the marking of noise and spurious signal properties as salient points, which can be indistinguishable from true features at fine or unsmoothed scale observations. Mathematically, this is represented by

$$\mathbf{feature}_{\sigma_k} \in [\mathbf{feature}_{\sigma_{k+1}} - \Delta : \mathbf{feature}_{\sigma_{k+1}} + \Delta], \quad (5.1)$$

which states that a feature at a fine level with scale parameter σ_k is located within some distance Δ of the feature at the next coarsest level of the scale parameter σ_{k+1} .

Yuille and Poggio [61] showed that of all the smoothing filters, only the Gaussian avoids the creations of new zero crossings with increasing scale parameter σ in one and two dimensions. The Gaussian kernel and its family of derivatives have been singled out as the only possible functions that satisfy the requirements of a scale space representation of an image or signal [62]. The Gaussian functions have the ideal properties of non causality, semi grouping, separability and ensures that the number of zero-crossings and values of extrema are strictly decreasing with increasing scale parameter σ , properties which make them ideal and unique for scale space implementation of vision algorithms [34].

A one dimensional signal and its scale space representation are show in Figure 5.1 from reference [62]. The original signal is show below. Scale is calculated by convolution of the signal with a Gaussian. The scale space description of the signal is computed by finding the zero-crossings of the second derivative of the signal at each scale. The ordinate is indexed to the original signal length, the abscissa is scale parameter σ . The scale space is binary; a black mark represents a zero crossing at the corresponding scale. The tree structure provides a hierarchical description of the inflection points of the signal at all scales simultaneously. Through successive smoothing operations, fine features such as those due to noise gradually disappear, while the more dominant points are retained.




Figure 5.1: A one-dimensional signal and its scale space representation [62].

5.2.2 Tracking

Tracking is closely coupled to feature definition. The tracking problem involves the definition of objects or features through modeling in one or more frames in a sequence. Objects are then matched by the application of an appropriate tracking algorithm. Matching can be achieved through an optimal criterion, taking advantage of the sequence; or in cases where this is not possible, features can be tracked by searching for a unique definition in each frame to circumvent the correspondence problem. The latter approach is used in this research. Tracking may be complicated by many factors including feature ambiguity, noise, spatial or temporal nonuniform illumination, object or feature occlusion and the entry and exit of objects from the imaging plane during the sequence. Constraints are often employed to simplify the problems, which may be probabilistic or deterministic. Knowledge of the scene is crucial and should be implemented if available. For example, a known direction of object motion or a constant velocity should be included in the search algorithm. A recent review of tracking methods and strategies is presented in [63].

5.2.3 Specific Challenges in Schlieren Data

A challenge for feature identification in this particular data sequence is the nonuniform laser pulse illumination. The laser pulsed at 25 ns was desired from a sensitivity standpoint and provided nearly instantaneous image acquisition at each frame. Intensity can vary up to 20% between pulses, typical of pulsed laser light sources. This can make the structural representation of objects or regions in the sequence inconsistent in an intensity-based representation. Laser speckle can also compete with true structure and noise. Another difficulty is the framing rate of the camera itself. The camera was operated at the maximum framing rate of 10 kHz, however there are multiple time scales existing simultaneously in the flowfield, most of which are unknown. This means that once a general interest point is identified (presuming it can be), if the structure is moving at a time scale greater than the Nyquist frequency, it will have exited the image plane long before the next frame is acquired, and a new set of features will have entered the frame. Therefore any feature match between frames within the freestream region moving at the freestream velocity would be rejected on a physical basis.

The problem is complicated by the three-dimensional optical path of integration. The wall boundary layers at the imaging station windows can cause structure occlusion, and allow for the possibility of the inclusion of numerous additional time scales to the freestream region. These are reasons why tracking cannot be performed on the finer flow structures in the sequence. The focus remains on the shock intersection point and the model tip.

5.3 Grayscale Features

Many feature detectors operate directly on the image, *i.e.* they are intensity-based. Harris and Stephens [64] proposed a corner detector that accounts for rapid spatial changes within the image through the identification of localized gradients. With tracking as a motivation, they stated that any features should be unique and invariant to translation and rotation. The Harris matrix is defined by the convolution of a two-dimensional Gaussian kernel with the structure tensor (Hessian) of the image in a windowed region, given in Equation 5.2. The Harris matrix incorporates scale through Gaussian convolution and allows for inclusion of scale with the differential operator. The Harris matrix \mathbf{H} is a function of both position and scale(s):

$$\mathbf{H} = g(\sigma_I, x, y) \otimes \begin{bmatrix} \mathbf{I}_x^2 & \mathbf{I}_x \mathbf{I}_y \\ \mathbf{I}_y \mathbf{I}_x & \mathbf{I}_y^2 \end{bmatrix}, \quad (5.2)$$

where subscripts denote partial differentiations with respect to the image \mathbf{I} . The Gaussian kernel g is a function of both x, y and scale parameter σ_I .

Corner responses are computed from an approximation to the eigenvalues λ of the Harris matrix in Equation 5.2 through the determinant and trace

$$\mathbf{R} = \det(\mathbf{H}) - k \text{trace}(\mathbf{H})^2 \approx \mathbf{R}(\lambda), \quad (5.3)$$

where k is an empirical constant, typically between 0.04 and 0.15. The response \mathbf{R} is positive at a corner or a region of rapid change in two directions, a negative value indicates an edge or a region of change in one direction, and a flat region of near constant intensity gives a small value, with little change in either direction. These changes correspond to the hypothetical eigenvalues of \mathbf{H} as two large eigenvalues, one small and one large, and two small eigenvalues, respectively. The approximation of the eigenvalues in Equation 5.3 ensures a degree of rotational and translational invariance under constant illumination through a sequence. Once the corner responses from Equation 5.3 have been determined, scores are retained only if they are a local maxima within an 8-neighbored region using non-maximal suppression. The resulting responses may then be thresholded or statistically ranked by discarding scores smaller than a specified value.

Shi and Tomasi [65] proposed a modification to corner detection, where the eigenvalues in Equation 5.2 are solved for directly. The model was based on building a tracking algorithm that by definition, detected “good features.” The rotation and translation of the structure tensor is monitored during tracking. Affine feature deformation is tracked in consecutive frames; the feature is discarded if it becomes too degraded. A threshold is employed on the minimum allowed eigenvalues directly. The Features from Accelerated Segment Testing (FAST) [66, 67] was developed for speed in real time tracking applications. This method constructs a Bresenham circle around each pixel and tests for dissimilarity at several points. Scale is implicit based on the circle radius, and hard thresholding is employed for each image based on the dissimilarity metric from the radius test.

5.3.1 Corner Detection Evaluation

Harris corner detection was performed on the double-cone sequence using an integration scale parameter of $\sigma_I = 1$. All corner detection was performed directly on the raw images. Reasons for not denoising the data have been discussed previously in Chapter 1. The Hessian was calculated using 3-tap central difference operators which eliminated the use of a differentiation scale, and provided the finest level of localization. The trace in equation 5.3 was weighted with $k = 0.04$.

Thresholding or ranking was not employed, due to the temporally nonuniform illumination through the sequence. Only 8-neighbored non-maximal suppression is used, responses above zero were kept. As response signatures and their distributions may change throughout the sequence, it is unwise to employ a threshold at this point. This is the reason why the Shi-Tomasi method is not shown, as a hard value of $\min(\lambda_1, \lambda_2)$ is required in the operation. A similar argument holds for FAST. Both methods were attempted and gave similar results, which are not shown.

Figure 5.2 shows the Harris corners in red overlaid on image $\mathbf{I}(t_o)$. 2,502 corners were found, largely in the texture dominated freestream region. The corners are clustered due to the small interrogation window used in the Hessian. As a consequence, features are well localized but numerous. Features identified on the model body are of little use, and could be discarded by employing the labeling scheme from the previous chapters. This result is typical of images in the sequence. Several thousand features are returned, few of which seem particularly meaningful, and are highly responsive to freestream texture patterns and other “noise”.




Figure 5.2: Image $\mathbf{I}(t_o)$ showing all Harris corner responses.

5.3.2 Performance in The Regions of Interest

Figure 5.3 shows typical Harris responses (black \circ surrounding white $*$) in the cone tip ROI. Images are shown at t_o and $(t_o + 1055)$, times selected to illustrate characteristic differences in detection results. The cone tip ROI is 41 x 41 pixels, with 44 and 38 interest points found in Figures 5.3a and 5.3b respectively. The detector is sensitive to noise, as expected. The model tip gives a corner response, and can therefore be considered a good detection method. However, there are far too many responses, and there is no physical reasoning in the detection algorithm on why the cone tip has been detected, and more importantly, a means of isolating the response point of the actual cone tip. This last item is the problem of localization.

Harris responses in the shock intersection ROI shown in Figure 5.4 corresponds to the time instances in Figure 5.3. The shock intersection ROI is 51 x 51 pixels.




Figure 5.3: Harris corner responses in the model tip ROI shown in \otimes . ROI is 41 x 41 pixels. Responses are spurious, although the cone tip is identified.

Sensitivity to “noise” or freestream structure is the most obvious attribute of the detector performance in the vicinity of the shock intersection. It can be argued that the shock intersection point has been reasonably identified in image t_o in Figure 5.4a by the Harris detector among the 79 feature points found. However, in Figure 5.4b, with 74 features, the intersection structure is difficult to determine even by the human eye, and does not appear to have been identified. Note that aside from the fact of considerable motion of the intersection structure, the triple point itself seems to take on a different representation. The separation shock angle seems to have steepened in Figure 5.4b. The Harris operator fails to identify the intersect point. Furthermore, it appears that tracking of the intersect point may require a larger region than a small windowed operator can account for. A larger scale must be accounted for when defining the intersection point.

The following can be concluded from Figure 5.3: the model tip is a true physical corner and responds as such using the Harris operator, and its identification




Figure 5.4: Harris corner responses in the shock intersection ROI shown in \otimes . ROI is 51 x 51 pixels. Responses are spurious and the intersect point is not always captured.

is largely a question of both a hard definition and localization. The cone tip should be identified at a larger scale by an alternate algorithm, that should account for some physical measure of saliency, and localized by a Harris corner point as a final step to provide the finest scale assignment. This larger scale algorithm should more importantly include some form of logic as to why the returned point can be defined as the model tip. Regarding the performance of the Harris feature detector for the shock intersection point, as shown in Figure 5.4 the desired interest point is not found, or well defined. More importantly, a method which accounts for the fact the the shock intersection is only meaningful when examining a large region. A larger scale must be accounted for when defining this feature, and the Harris method cannot be applied for a final localization step. Finally, the Harris detector incorporates no measure of physical saliency that an ideal detection algorithm for these features to account for.

5.4 Curvature as a Feature

As intensity-based feature detectors have proven problematic for application to schlieren data, at least as a general method, attention is turned to saliency measures based on boundary descriptors. Direct comparison of edges are non-unique; there is no one-to-one correspondence from any edge point to another in monocular vision. This is due to an ambiguous match from a point on any one line to any on the next in from frame to frame caused by the aperture problem. Therefore, edge point matching is by definition intractable and under-constrained. However, boundary descriptors built from edge representations can provide unique and isolated features. Given that a contour has been extracted from the image, for example from segmentation or edge detection, salient features describing the shape can be defined in terms of contour properties. Changes in curvature on a silhouette or contour provide a set of salient points or knot points, that can be used to describe shape. These salient points are localized boundary descriptors.

Outer-shock contours have been isolated in the previous chapters using edge maps and boundary shot detection. The use of a contour property for a feature detection scheme in this case is beneficial due to localization and saliency. Given that the outer contour is a reasonable representation of the complete outer-shock structure, both the model tip and the triple point must be defined in terms of, and exist on this contour. The curvature of this contour can be used to uniquely define the structures of interest on a physical basis, namely the signed extrema.

Given a parameterized curve in terms of $(x(t), y(t))$, the analytical expression of curvature is defined as the rate at which the tangent vector slope which traverses the curve changes. Analytical curvature κ is given by

$$\kappa = \frac{x\ddot{y} - y\ddot{x}}{(\dot{x}^2 + \dot{y}^2)^{\frac{3}{2}}}. \quad (5.4)$$

Positive curvature is defined when the tangent vector traversing the outside of a curve moves clockwise. The curvature is negative when the tangent vector sweeps in the counter-clockwise direction. Extrema in curvature, its changes, and inflection points measured from the second derivative can all be used to describe the features of a contour. Therefore, In addition to sharp corners, features can also be identified by curvature properties.

Asada and Brady [68] proposed that curvature changes consists of two fundamental primitives, a corner and a smooth join. Additional primitives were defined as compound representations of the first two, consisting of an end, a crank and a bump (or dent). Analytical forms of convolutions of curvature with a Gaussian and its derivatives were performed. These representations were computed at multiple scales to form the curvature primal sketch, consisting of curvature profiles and their derivatives. The resulting tree structure was searched, and instances of primitives in the primal sketch were assigned a descriptor based on their behavior through the scale space, or tree.

5.4.1 Curvature Evolution and Scale Space

Mokhtarian and Mackworth [69,70] proposed a method of curve representation using the zero-crossings of contour curvature at multiple scales for descriptions of, and matching between, curves. Gaussian convolution was used to compute curvature and its derivatives of the discrete parameterized coordinates. This gave representations of a contour and its curvature through multiple scales. Successive convolutions are termed evolutions of the contour. Evolutions were used to construct the Curvature Scale Space (CSS) image, given in Equation 5.5

$$\kappa(s, \sigma) = 0. \tag{5.5}$$

this a binary image tree structure which constitutes the number of zero-crossing of curvature along a contour at multiple scales (vector of σ), providing a view of contour inflection points at all scales simultaneously, a direct analogy to scale space filtering shown in Figure 5.1. A zero-crossing (feature) can be found at a large scale where noise has been eliminated. Localization through a coarse to fine search is accomplished by searching the neighborhood near the zero-crossing at successively smaller scales using Equation 5.1.

5.4.2 The Curvature Scale Space Corner Detector

Mokhtarian and Suomela [71] developed a CSS-based corner detector. Corner points were identified at a large scale, where noise is assured to be removed. Since

locations of the corner points may not be well localized at large scales, they were updated by searching from a coarse to fine scale in the neighborhood of the corner on the previous scale, using Equation 5.1. The true corners are determined as local extrema in absolute curvature at the largest scale. In addition to representing a true corner, extrema in curvature can be due to noise and rounded peaks caused by rounded corners. Mokhtarian and Suomela state that noise and rounded corners should not be considered; they are removed through the introduction of a threshold ensuring that only sharp corners remain. Local extrema are defined as having values twice as large as their neighboring local minimum. A rounded corner will have a rounded peak in local curvature, and therefore be discarded. The detector was later refined [72] to account for relative scales for the smoothing of long and short contours, and for smoothing of long contour curvature. Both of these steps increased the detection rate of true corners while decreasing the false positive detections. This highlights the fact that it is difficult to select a proper scale for a general class of applications.

He and Yung [56, 57] considered curvature at a single small scale and removed corners based on an adaptive threshold. Additional corners were removed subject to an angle threshold based on a dynamic region of support surrounding corner candidates. By design, rounded corners were removed. A true scale space view was not taken advantage of due to the use of a single assumed scale. Zhang *et al.* [73] proposed a CSS-based corner detector that examined extrema in the curvature product at several scales. The need from coarse to fine tracking is claimed to be eliminated, although the assumption is made that true features, exclusive of noise,

are included in the small scale calculation. A global threshold was used. Zhang *et al.* [74] used contour evolution directly to calculate locations of corners. The difference in evolution was defined as the difference between two evolved contours computed with convolutions with Gaussians of different scale, or the Difference of Gaussians filter (DoG). Evolution distance was noted to be a local maximum at true corner locations. The use of two scales in the DoG filter superseded the coarse to fine tracking step. It was shown that extrema in the norm of the DoG evolution distance were equivalent to those in curvature. The definition of curvature used did not take into account the contour shrinking during the evolution process.

Problems in reconciling these measures of cornerness include: only a well defined (sharp) corner will have a distinct peak in local curvature, and a thresholding criterion can arbitrarily discard weak or rounded corners which may correspond to true physical events. Too large a threshold (sensitive to smoothing scale) can remove true events on a shock contour, while one that is too small may return false positives. Most importantly, there is again the lack of an objective metric for which to define a physically salient point. These problems can be overcome with the introduction of a knowledge base while taking advantage of the fact that significant events on shock contours have curvature responses that may be weak.

5.5 Contour Extraction and Processing

Bottom-up image understanding requires image segmentation to reduce image content. Here, image content is reduced to edge maps, from which outer-shock

structures are isolated. Given that a reasonable representation of the outer-shock structure can be isolated, salient features describing the shape can be defined as properties of the contour. The use of a contour property for a feature detection scheme is beneficial due to inherent localization and saliency. Both the model tip and the triple point must be defined by, and exist on, this contour; these features are automatically localized. The contour curvature can be used to uniquely define the structures of interest on a physical basis, namely the signed extrema.

5.5.1 Contour and Curvature Evolution

Application of a CSS-based method requires that the outer-shock be represented as a single unbroken contour. Contour evolution and curvature calculation is performed as in Mokhtarian's work [69–71] discussed in the previous section. Scale selection is crucial, and can often only be determined by trial and error through inspection. The method for the extraction and isolation of the continuous outer-shock contour has been discussed in the previous chapter. The Canny edge detection algorithm [15] was used. A portion of the corner detection algorithm code developed by He and Yung [56, 57] was employed to extract and label contours given the edge image. The outer-shock contour was then isolated. This is by no means the only method by which image contours can be retrieved. Alternative methods such as level sets or snakes [75] may also be used for contour extraction.

Once an image has been reduced to a single contour, we assumed that this contour represents the outer-shock. This is justified since the structure is represented

by the strongest gradients in the image and therefore the most likely to be detected. This was confirmed by inspecting sample results. The isolated outer-shock contours are reasonable but “noisy” representations of true structure. This is largely due to the scale of the original images (resolution), but partially due to the low dynamic range and the high sensitivity at which they were acquired. The discrete curvature will be noisy at a small scale, dominated by contour noise rather than true features. Noise is inherent on discrete curves. At the level of single pixel inspection, lines may only be formed by 8-neighbored movements (chessboard distance) from one pixel to the next. The only pure discrete lines in an image are oriented at 0, 45 and 90° *etc.* For these reasons, it is assumed that the dominant features of the shock contour are retained at large scale.

The discrete curve or contour in an image is represented by Γ_o in Equation 5.6, where x and y are the parameterized image coordinates along the discrete arclength s of length N .

$$\Gamma_o(s) = \{x(s), y(s)\}, s \in [1, \dots, N]. \quad (5.6)$$

The contour Γ_o from Equation 5.6 is successively evolved at multiple scales σ through Gaussian convolutions of increasing scale with the parameterized functions. During contour evolution, the curve is simplified with increasing scale (as in scale space filtering) during this evolution and the shape is simplified. The evolution is given in Equation 5.7 below, where g is the Gaussian. Although the curve remains integer index, the evolved curve is mapped from $\mathbb{Z} \mapsto \mathbb{R}$:

$$\begin{aligned}
\mathbf{\Gamma}(s, \sigma) &= g(s, \sigma) \otimes \mathbf{\Gamma}_o(s) \\
&= \{g(s, \sigma) \otimes x(s), g(s, \sigma) \otimes y(s)\}.
\end{aligned} \tag{5.7}$$

Scale parameters σ from coarse to fine of 6, 4, 2 and 1.5 were used for evolution. During this process, the contour shrinks although s remains discrete and the total integer arclength N remains constant.

The discrete curvature is computed directly from its continuous definition in Equation 5.4 using the evolved contours,

$$\kappa(s, \sigma) = \frac{\dot{X}(s, \sigma) \ddot{Y}(s, \sigma) - \dot{Y}(s, \sigma) \ddot{X}(s, \sigma)}{\left(\dot{X}(s, \sigma)^2 + \dot{Y}(s, \sigma)^2\right)^{\frac{3}{2}}}. \tag{5.8}$$

Where the variables X denotes a Gaussian smoothing at scale σ has been applied to the image indexed coordinates x, y . As differentiation and convolution are commutative, the smoothed parameters and their derivatives are calculated as follows

$$X(s, \sigma) = x(s) \otimes g(s, \sigma) \tag{5.9}$$

$$\dot{X}(s, \sigma) = x(s) \otimes g_s(s, \sigma) \tag{5.10}$$

$$\ddot{X}(s, \sigma) = x(s) \otimes g_{ss}(s, \sigma), \tag{5.11}$$

and similarly for y and Y . All Gaussian convolution kernels are constructed directly from their analytical definitions given by

$$g(s, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{s^2}{2\sigma^2}} \quad (5.12)$$

$$g_s(s, \sigma) = \frac{\partial g}{\partial s} = \frac{-s}{\sqrt{2\pi}\sigma^3} e^{-\frac{s^2}{2\sigma^2}} \quad (5.13)$$

$$g_{ss}(s, \sigma) = \frac{\partial^2 g}{\partial s^2} = \frac{1}{\sqrt{2\pi}\sigma^3} \left[\frac{s^2}{\sigma^2} - 1 \right] e^{-\frac{s^2}{2\sigma^2}}. \quad (5.14)$$

5.5.2 Kernel Construction and Convolution

For accurate Gaussian kernel construction, all functions in Equation 5.14 were sampled at $\pm 6\sigma$ and renormalized. This was done to minimize truncation errors in the construction of a finite representation of a function with infinite support and to maintain scale space behavior in the discrete analog case. Due to the use of large kernel convolutions, dictated by large scale parameters σ , potential edge effects from padding the contours prior to convolution was a concern. All contours are open in the cases of the present data set and will be for most schlieren data. Circular convolution is unwise in this case. Zero-padding and reflected boundary conditions introduced large artifacts in the curvature scale space. Boundary conditions were therefore copied on either end of the X and Y coordinates to match the extent of the kernel. This method yielded results with minimum effects of false artifacts near the edges relative to the order of the extrema.

5.6 Algorithm Design

Construction of a feature detector for a specific application requires implementation of expert knowledge to construct a formal knowledge base. This incorporates descriptions of the feature into rules to construct a definition that ensures proper detection and localization; providing machine instructions that implement visual observations and physical laws. This should also negate or minimize false detections and maximize the probability of true detections. Given the computed curvature at multiple scales, the data must be searched for instances which correspond to relevant features. Whereas most CSS-based corner detectors consider absolute curvature, the retention of signed curvature can be advantageous in building a feature definition to search for salient regions on a shock contour.

Figure 5.5 displays a representative isolated outer contour from the sequence showing locations of the model tip and shock intersection along with curvature notations and the locations of contour endpoints. Although inspection of the contours extracted from the sequence shows that they approximate the outer-shock structure well in most instances, this may not always be the case. Contour noise may be significant, or a freestream schlieren in the optical path may be aligned such that false structures are generated in the contour image. This can manifest itself as protrusions in the aft-cone bow shock, or a shock wave which appears to turn into itself, due to strong, but unphysical gradients in the image, causing sharp corners on the contour other than the model tip. These factors must be considered when building the detector. It is assumed that the extracted contours are reasonable rep-

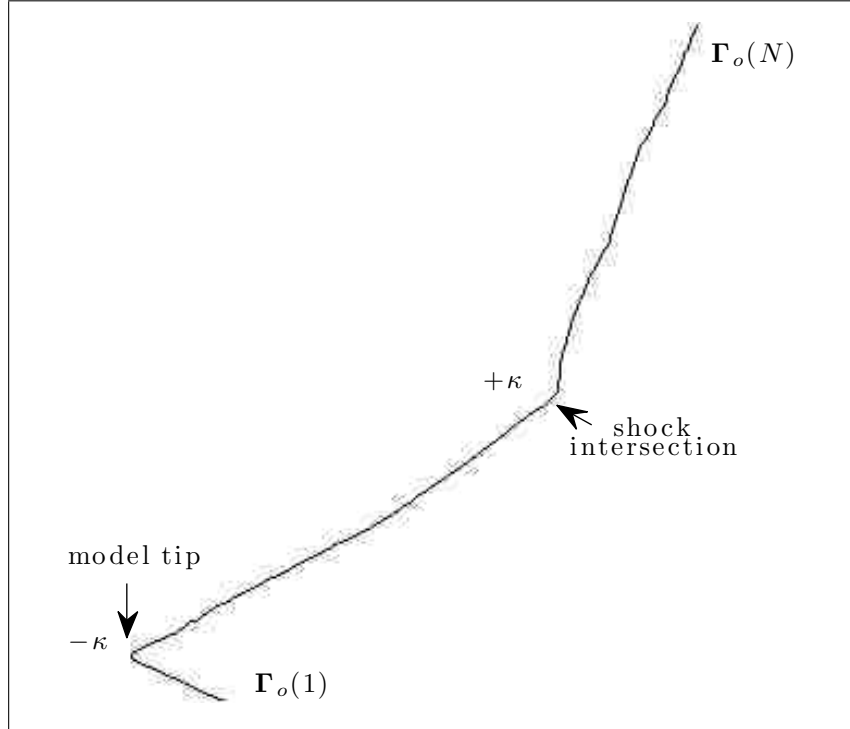


Figure 5.5: Isolated outer-shock contour showing feature locations and sign conventions.

representations of the outer-shock structure, and that significant departures from the ideal can be recognized with a carefully constructed knowledge base.

From a tracking perspective, it is assumed that both the model tip and the shock intersection region occur at most once in each frame, and can be well defined. Additionally, the assumption is made that the motion of both the model tip and shock intersection point is planar and translational. The small motion assumption is imposed; the motion of both features are smooth and the distance moved between frames is small. Finally, observations of an event should be similar. This means that observations that seem like outliers in a final motion distribution have a physical basis for being re-examined.

5.6.1 Large Scale Curvature as a Canonical Descriptor

Figures 5.6 and 5.7 show curvature of the continuous outer-shock contours computed at the largest and smallest scales, σ of 6 and 1.5 for images $\mathbf{I}(t_o)$ and $\mathbf{I}(t_o + 2158)$, illustrating typical variations of the shock structure through the sequence. Curvature at small scale ($\sigma = 1.5$) is represented by the dashed line in Figures 5.6 and 5.7; both instances show significant noise. In Figure 5.6 (corresponding to the contour shown in Figure 5.5) at small scale, two significant positive peaks and a negative peak are discernible among the salient structures. At time $(t_o + 2158)$ in Figure 5.7, the dominant structure at small scale is the negative peak. Significant zero-crossings are evident for both time instances at small scale. At the larger scale of $\sigma = 6$ (the solid line), the number of zero-crossings has been reduced significantly and much of the contour “noise” and fine structure has been removed for both time instances. Most importantly, at both time instances for increased scale, there are only two dominant features remaining in curvature: a negative peak at $s \approx 45$, and a small positive rounded structure near $s \approx 240$. The large scale view allows for a consistent and general description of the contours that is not possible at the small scale. This provides a canonical view of all contours in terms of two events: a large negative curvature structure, and a local (possibly weak) maximum. These large scale events correspond directly to the features of interest the shock: the model tip and the shock intersection.




Figure 5.6: Curvature κ for $\mathbf{I}_c(t_o)$ at two scales.




Figure 5.7: Curvature κ for $\mathbf{I}_c(t_o + 2158)$ at two scales.

5.6.2 Model Tip Identification

From Figure 5.5, the cone model tip can be represented by the most negative value of curvature. Furthermore, the value of curvature at the model tip is the largest

curvature in the absolute sense. Figures 5.6 and 5.7 illustrate a sharp negative peak at the large and small scale curvatures at an arclength position $s \approx 45$. This is the curvature primitive for a sharp corner (cusp) made by an acute physical angle (obtuse angle swept by the tangent vector). The primitive persists through scale, maintaining its general appearance, and both of values of curvature at this point are similar at both scales at each time instance. The cone tip is a solid corner, its representation and curvature primitive will not change during the sequence aside from potential contour noise. Since the aft-cone can be oriented at most vertically, the cone tip forms the most acute angle in the image; its curvature must be the largest absolute curvature on the shock contour, and negative. Any additional sharp cusps protruding from the shock contour would be due to poor contour representation and indicate false structure. From observation of Figure 5.5, the cone tip must occur within, say the first N_1 points of the contour. This provides further localization and minimizes computational time for tracking. Although a sharp corner has an easily identifiable signature in curvature and would survive arbitrary thresholding selection, this method lacks a meaningful and searchable definition of the model tip. The following knowledge is therefore incorporated to construct a physical definition

- 1 The model tip occurs at most once in each image; $\exists! \mathbf{p}_{tip}$
- 2 The approximate cone tip location is known; $\mathbf{p}_{tip} \in [\mathbf{\Gamma}(1), \mathbf{\Gamma}(N_1)]$
- 3 The model tip is physically acute; κ_{tip} must be negative.
- 4 The tip must be the only acute angle on the contour at large scale: $\max(|\kappa|) = |\kappa_{tip}|$.

The above statements form the knowledge base; this allows the location of model tip s_{tip} to be defined at large scale in Equation 5.15, which uniquely defines the feature.

$$\mathbf{p}_{tip}(s) = \underset{s}{\operatorname{argmin}} (\kappa(s, \sigma_{\text{MAX}})) : s \in [\mathbf{\Gamma}(1) \mathbf{\Gamma}(N_1)] \quad (5.15)$$

5.6.2.1 Localization

Search from coarse to fine scale may arrive at final incorrect locations due to local peaks created by contour noise. This is a trade-off between accuracy and localization; the more accurate the scale of final assignment, the more susceptible that assignment is to local peaks created by noise. This uncertainty principal is common in many vision algorithms. Locations of curvature extrema are inherently based on some region of support on either side of the defined corner, this is true from the mathematical definition and explicit through the definition of the Gaussian functions used in convolution. Therefore some error, which may be difficult to quantify, will always be inherent at the scale of pixel assignment to a feature. In order to minimize the effect, an “attractor” feature is compared to that returned by the CSS method.

Figures 5.3a and 5.3b illustrate that the Harris response detected the cone tip well. Although it did not account for a definition and feature values were numerous, it can be used as a final localization step. This is because the true feature has been uniquely defined with a knowledge base, narrowing down the potential candidates. The model tip is a true physical corner persisting through illumination changes and

is well suited for a final localization step as it detected well against the many and spurious responses. The final CSS model tip point is therefore updated to the closest Harris response in the vicinity of the model tip in the raw image using Equation 5.16

$$\mathbf{p}_{tip} = \underset{j}{\operatorname{argmin}} (\|\mathbf{R}(x_j, y_j) - \mathbf{\Gamma}_o(s_{min})\|) : (x_j, y_j) \in ROI, \quad (5.16)$$

which gives the final location of the model tip \mathbf{p}_{tip} at time t .

5.6.2.2 Detection Algorithm

The procedure to identify and localize the cone tip for the successfully isolated shock contours is given in Algorithm 3. Both the raw image and the isolated contour are required. Curvature is calculated at four scales using σ values of 6, 4, 2 and 1.5 for Equation 5.8. The minimum value of curvature within the first $N_1 = 100$ arclength points along the contour, and the location and curvature value of the model tip s_{min} and κ_{min} are identified at large scale using Equation 5.15. Localization is accomplished by a coarse to fine search across scale using Equation 5.1, searching within a distance $\Delta = \pm 5$ arclength points on either side of the previous minimum. The Harris response closest to the minimum identified curvature is then returned as the final model tip. In the event that Equation 5.16 returns multiple solutions, this instance is recorded and marked for visual inspection.

Algorithm 3: Identifying the cone tip for known contours.

Input: \mathbf{I} raw image
 \mathbf{I}_c isolated outer contour image
 $IMAX$ descending scale parameters σ


Output: $\mathbf{p}(t)_{tip}$: model tip coordinates at time t
calculate Harris response \mathbf{R} for \mathbf{I} in ROI;
// calculate curvature at all scales
for $i = 1 : IMAX$ **do**
| smooth contours;
| calculate curvature $\kappa(s, \sigma[i])$;
end
// initialize the minimum
 $s_{min} = \underset{s}{\operatorname{argmin}}(\kappa(s, \sigma_1)) : s \in [1, \dots, N_1]$;
// localize the result from coarse to fine search
for $i = 2 : IMAX$ **do**
| $s_{min} = \underset{s}{\operatorname{argmin}}(\kappa(s, \sigma[i])) : s \in [s_{min}[i-1] - \Delta, s_{min}[i-1] + \Delta]$;
end
 $\kappa_{min} = \kappa(s_{min}, \sigma[IMAX])$;
// localize to the closest Harris response
 $\mathbf{p}_{tip} = \underset{j}{\operatorname{argmin}}(\|\mathbf{R}(x_j, y_j) - s_{min}\|) : (x_j, y_j) \in \text{ROI}$;
// is argmin multivalued?
if $j > 1$ **then**
| mark the image for manual evaluation;
end
return \mathbf{p}_{tip}

Three of the four contour evolution stages are illustrated for images $\mathbf{I}(t_o)$ and $\mathbf{I}(t_o + 1678)$ in Figure 5.8 at scales 6, 4 and 1.5 with decreasing scale from left to right. Contours are evolving from right to left. The initial contour Γ_o extracted from the edge map is labeled with a \square , each evolved contour is labeled \circ . The respective minimums detected at each stage in the algorithm are denoted with a solid marker, corresponding to the coordinate on the edge and its contour in \mathbb{R} , illustrating the coarse to fine search through CSS in the detection algorithm. In Figures 5.8a, 5.8b, and 5.8c, for $\mathbf{I}(t_o)$, the location of the minimum curvature indexed to the


initial contour does not change through evolution. At time $(t_o + 1678)$ in Figures 5.8d, 5.8e, and 5.8f, the initial contour is noisy, but the curve simplification through contour evolution is able to capture the structure well at large scale, and the location of the minimum changes. This illustrates both the benefits of a large scale canonical view, and the importance of the final localization step.

5.6.3 Shock Intersection Identification


In Figures 5.6 and 5.7, the shock intersection is represented by a small rounded peak near $s \approx 240$ at the large scale; although the peak at time $(t_o + 2158)$ is nearly half the value at both scales as it was at t_o . In Figure 5.6, a positive peak is evident at small scale, surrounded by smaller structure or noise. At a larger scale, the peak has become clarified, seeming to take on a global maximum. To contrast, at the large scale in Figure 5.7, the peak positive value is not evident at the small scale. The intersection angle is more physically obtuse, and is dominated by fine detail and noise at small scale. At a larger scale, the fine detail and noise has been removed, revealing the dominant structure. The increased aperture width of the Gaussian at larger scale allows the primitive of the more obtuse shock angle to be captured accurately. Although it persists through scale, the intersection point is difficult to spot at small scale. In both instances, the intersect curvature has the opposite sign and is much smaller in absolute value than that of the model tip.




(a) $\mathbf{I}(t_o) \Gamma(\sigma = 6)$




(b) $\mathbf{I}(t_o) \Gamma(\sigma = 4)$




(c) $\mathbf{I}(t_o) \Gamma(\sigma = 1.5)$



(d) $\mathbf{I}(t_o + 1678) \Gamma(\sigma = 6)$



(e) $\mathbf{I}(t_o + 1678) \Gamma(\sigma = 4)$



(f) $\mathbf{I}(t_o + 1678) \Gamma(\sigma = 1.5)$

Figure 5.8: Contour evolution near the cone tip for images $\mathbf{I}(t_o)$ and $\mathbf{I}(t_o + 1678)$.

5.6.3.1 Intersection Curvature Primitives

The region of the shock intersection, unlike the model tip, changes during the sequence and may take on multiple appearances, many of which have different curvature primitives. The intersection may appear as a sharp corner, a rounded corner, or a smooth join. During the sequence, the bow shock angle is observed to pitch backward. This is due to the angle made by the intersecting shocks changing due to changes in the orientation of the aft-cone bow shock, which has an at most vertical orientation near the intersection (further inclination of the bow shock is not physically possible). Variants of the shock intersection representation are illustrated in Figure 5.9; a “typical” representation is shown in addition to a representation which occurs when the bow shock angle is pitched back, and a nearly vertical inclination of the bow shock is shown. These are referred to here as a typical, a smooth join and a sharp curvature primitives due to their corresponding contour shapes in Figure 5.10. Curvature primitives for the three types are shown in Figure 5.11; large and small scales curvatures κ are plotted. A value for a typical shock intersect primitive is shown in Figure 5.11a (corresponding to the contour in Figure 5.10a), and has values at large and small scales of κ 0.12 and 0.4 respectively. The smooth join curvature in Figure 5.11b (corresponding to the contour in Figure 5.10b) is less than half that of the typical value, and on the level of contour noise at large scale; such a “feature” would typically be discarded by a CSS-based corner detector. The sharp corner contour in Figure 5.10c has curvature scales shown in Figure 5.11c. The value of κ is nearly three times the typical value, and in this case would be

detected through CSS-based methods. However, all potential representations must be accounted for when searching for the shock intersection point.

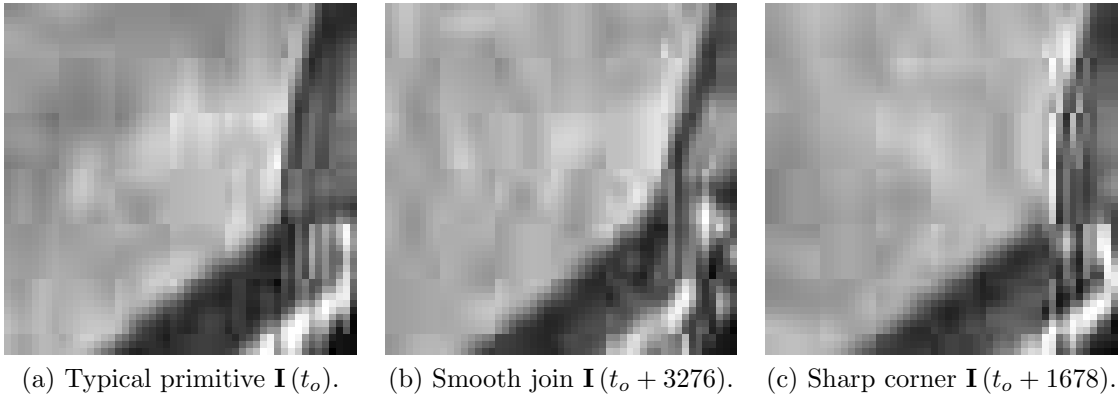


Figure 5.9: Shock intersection primitive structures in the raw image.

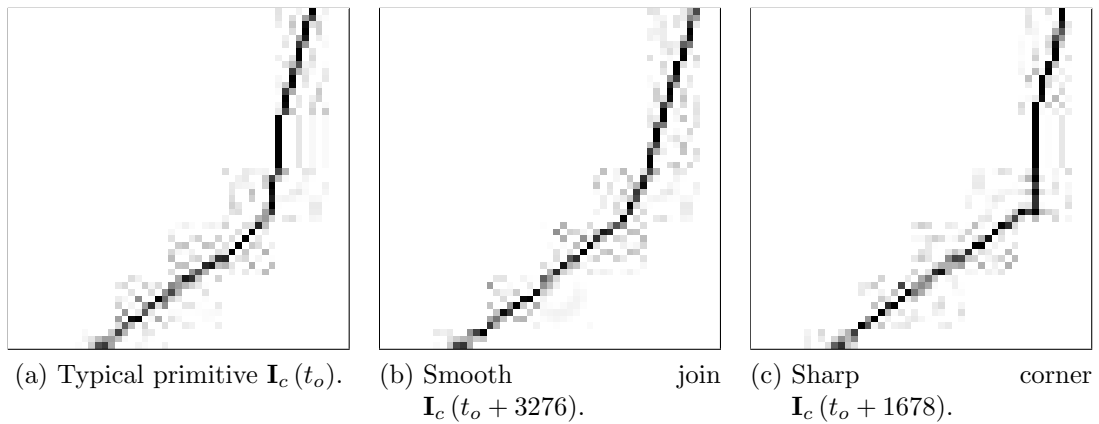


Figure 5.10: Shock intersection primitive structures in the contour image.

Many CSS-based corner detectors would likely reject all but the model tip during the detection phase through an arbitrary threshold. The sharp corner is the only feature with a strong curvature signature. The typical intersect primitive and the smooth join ($\kappa < 0.02$) would be classified as weak features and likely be purged from the candidate list. While not strong corners, they correspond to true physical events. Lowering the threshold would cause false positive detections. The CSS algorithm must be modified to allow for these primitives. Thresholding is therefore




Figure 5.11: Shock intersection curvature primitives.

avoided, as a rounded corner or smooth join may be weak. All primitives in Figure 5.11, while potentially weak, appear to take on at least a local maximum in curvature at large scale. Viewing the intersect in this manner as opposed to a threshold allows rounded corners and smooth joins a definition insensitive to a threshold based on a physical event.

At a large scale, the shock intersection region can be represented by the largest positive curvature on the contour, since the angle must be physically obtuse (acute with respect to motion of the contour tangent). The intersect point occurs within the first $N_2 > N_1$ points from inspection of Figure 5.5. The curvature must also be

less than the absolute value of the model tip to account for potential false structure.

The following rules and observations can therefore constitute the knowledge base:

- 1 The intersection occurs at most once in each image; $\exists! \mathbf{p}_{int}$
- 2 The approximate intersection location is known; $\mathbf{p}_{int} \in [\mathbf{\Gamma}(1), \mathbf{\Gamma}(N_2 > N_1)]$
- 3 The intersection may take on multiple primitive curvatures, all of which must be physically obtuse and therefore have positive curvature.
- 4 The intersection must be the only obtuse angle at large scale; $\kappa_{int} = \max(\kappa)$
- 5 As the bow shock is inclined at most vertically, it is physically impossible for the intersection curvature to be larger than that of the absolute value of tip curvature, which provides an upper bound indicating false structure; $\kappa_{int} < |\kappa_{tip}|$.

The above knowledge base allows the shock intersection to be uniquely defined at large scale by Equation 5.17.

$$\mathbf{p}_{int}(s) = \underset{s}{\operatorname{argmax}}(\kappa(s, \sigma_{\text{MAX}})) : s \in [\mathbf{\Gamma}(1) \mathbf{\Gamma}(N_2)] \quad (5.17)$$

The shock intersection is defined and localized using the CSS method to determine the value and location of the maximum curvature and location κ_{max} and s_{max} , and localized using the same methods as identifying the model tip. At the finest scale, the value of curvature must be checked against a physically impossible representation of the contour image. This is asserted by evaluating

$$|\kappa_{max}| < |\kappa_{min}|. \quad (5.18)$$

If the above expression is true, a failure is noted and the time instance is recorded for manual evaluation. An example of an impossible contour representation is shown in Figure 5.12a where the intersection is represented by a loop structure. This image was identified by the algorithm due to a curvature which violated Equation 5.18, and shown in Figure 5.12b.




Figure 5.12: Unphysical intersection representation violating predicate logic.

5.6.3.2 Localization

Instances of significant feature drift for the shock intersection were observed viewing the algorithm progress, due to contour noise and complex primitives. Due to nonuniform illumination and more importantly the changing primitive of the intersect structure, a Harris response (shown in Figure 5.4) can not be used as a final attractor to correct the intersection for feature drift through scale, as it is not a

true physical corner. Contour evolution distance, which is maximal at corners, was therefore used as a localization step. Evolution is calculated between two large scales $\sigma_2 > \sigma_1$, using a similar method to that in [74] to ensure that any distance moved by a point on the arclength is not due to local noise, this can be used to guarantee that structure motion as opposed to noise evolution is measured. Evolution distance is given by Equation 5.19

$$\mathbf{D}(s, \sigma_2, \sigma_1) = \|\Delta\mathbf{\Gamma}\| = \sqrt{(X(s, \sigma_2) - X(s, \sigma_1))^2 + (Y(s, \sigma_2) - Y(s, \sigma_1))^2}, \quad (5.19)$$

which is the absolute distanced traveled by a point on the arclength during evolution. This method was not used as an initial search criterion as it does not allow for incorporation of a physics-based knowledge base, and can therefore only be used as a final localization step. The shock intersection point returned from the CSS-based method in Equation 5.17 is updated to the nearest maximum of Equation 5.19, using

$$\mathbf{p}_{int} = \underset{s}{\operatorname{argmin}}\{\|\mathbf{\Gamma}_o(s_{max}) - \underset{s}{\operatorname{argmax}}(\mathbf{D}(s))\|\}: s \in [s_{max} - \Delta s : s_{max} + \Delta s]. \quad (5.20)$$

It was observed during the algorithm development that this method consistently corrected for feature drift.

5.6.3.3 Detection Algorithm

Algorithm 4 shows the procedure for identification and localization of the shock intersection point for images with known contours. Unlike the model tip, this algorithm requires only the continuous isolated contour. Contours and curvatures at **IMAX** decreasing scales are calculated as in the cone tip identification, in this instance the maximum value of curvature is used. The search is initiated within the first $N_2 = N - 100$ arcelngth points using Equation 5.17. Coarse to fine search for localization is again performed using Equation 5.1. The final value of κ_{max} is compared to the value of κ_{min} to identify false structure in the contour, any such image is recored and set aside for manual evaluation. The final localization is defined as the closest maximum in the evolution distance to the final CSS location using Equation 5.20.

Algorithm 4: Identifying the shock intersection for known contours.

Input: \mathbf{I}_c isolated outer contour image
IMAX descending scale parameters σ

Output: $\mathbf{p}(t)_{int}$ shock intersection coordinates at time t

// calculate curvature at all scales

for $i = 1 : IMAX$ **do**

 | smooth contours;

 | calculate curvature $\kappa(s, \sigma[i])$;

end

// calculate κ_{min} as in previous algorithm

// initialize the maximum

$s_{max} = \underset{s}{\operatorname{argmax}}(\kappa(s, \sigma_1)) : s \in [1, \dots, N_2]$;

// localize the result from coarse to fine search

for $i = 2 : IMAX$ **do**

 | $s_{max} = \underset{s}{\operatorname{argmax}}(\kappa(s, \sigma[i])) : s \in [s_{max}[i-1] - \Delta, s_{max}[i-1] + \Delta]$;

end

$\kappa_{max} = \kappa(s_{max}, \sigma[IMAX])$;

// asset physical impossibility

if $\kappa_{max} > |\kappa_{min}|$ **then**

 | mark the image for manual evaluation;

end

// calculate evolution distance at two large scales


$\mathbf{D}(s) = \|\Delta\mathbf{\Gamma}_o(\sigma[1], \sigma[2])\|$;

// localize to the nearest maximum in evolution distance


$\mathbf{p}_{int} = \underset{s}{\operatorname{argmin}}\{\|\mathbf{\Gamma}_o(s_{max}) - \underset{s}{\operatorname{argmax}}(\mathbf{D}(s))\|\} : s \in [s_{max} - \Delta s : s_{max} + \Delta s]$;

return \mathbf{p}_{int}


Figure 5.13 displays contour evolution near the vicinity of the shock intersection for instances $\mathbf{I}(t_o)$ and $\mathbf{I}(t_o + 1678)$ (corresponding to the sharp corner primitive). Scales and labeling schemes correspond to those in Figure 5.8. Figures 5.13a, 5.13b, and 5.13c for (t_o) show the typical primitive which has a single localization step from coarse to fine search as the evolution scale decreases from left to right. A similar trend occurs in Figures 5.13d, 5.13e, and 5.13f; in both cases, with decreasing evolution and reduction of shape simplification, the evolved contour closely approaches the initial curve $\mathbf{\Gamma}_o$.




(a) $\mathbf{I}(t_o) \Gamma(\sigma = 6)$




(b) $\mathbf{I}(t_o) \Gamma(\sigma = 4)$




(c) $\mathbf{I}(t_o) \Gamma(\sigma = 1.5)$



(d) $\mathbf{I}(t_o + 1678) \Gamma(\sigma = 6)$



(e) $\mathbf{I}(t_o + 1678) \Gamma(\sigma = 4)$



(f) $\mathbf{I}(t_o + 1678) \Gamma(\sigma = 1.5)$

Figure 5.13: Contour evolution near the shock intersection for images $\mathbf{I}(t_o)$ and $\mathbf{I}(t_o + 1678)$.

5.6.4 Handling of Missing Contours

When a contour representation in the schlieren sequence is physically not representative of the outer-shock structure, neither continuous nor single valued, the relative extrema of curvature are not accessible to describe the physical features as they have been defined. For the model tip, the CSS method was used for a unique definition and localization criterion. The final localization was assigned by updating the final value to the Harris response subject to the minimization of the euclidean distance between the two features. When a contour at time t is not available, given the previous model point at instant $(t - 1)$, the tracking problem can be constrained deterministically using the small motion assumption. The location of the model tip at time t can be determined by minimizing the euclidean distance between the known and unknown feature between time instances, or

$$\mathbf{p}_{tip}(t) = \underset{j}{\operatorname{argmin}} (\|\mathbf{p}_{tip}(t - 1) - \mathbf{R}(x_j, y_j, t)\|) : (x_j, y_j) \in ROI, \quad (5.21)$$

which is simply the assignment of the model tip closet to the previous known location. Note that this approach allows for consecutive missing contours.

For the shock intersection point, Harris point estimators are not available for final candidate localization, a separate contour evolution based criterion was used, as this feature is not consistently represented by a true corner. To extract the missing data points, a graphical user input was constructed. The point was selected

manually from the raw image for each missing frame which was processed with a diffusion like bilateral filter [46] allowing a visual grouping, as the instances are by definition difficult to evaluate by the human eye. A total of 89 frames, less than 2% of the data, needed to be measured manually. Although a human evaluation is not repeatable, the small percentage of data for which this method is used will render any potential manually generated anomalies statistically insignificant, and will have no effect on the final motion distributions. Also, outliers arising from this method will be visible through inspection of the distributions. This program later served to visually verify and evaluate, and correct if necessary, any potential problem images, or those identified as outliers.

5.7 Results

The schlieren sequence was processed using MATLAB. Results are plotted in pixel coordinates (row, column) to best illustrate feature motion history. There are several reasons for not localizing the results beyond pixel coordinates. These include the resolution (outer scale) of the features, spurious gradients caused by schlieren system sensitivity, and potential artifacts in the optical integration path. Combined, these reasons create structures which are difficult to discern by the human. As there is no ground truth to compare against in this application, further localization would provide a false impression that the methods used on this data allow for resolution capabilities beyond a discrete (binned distribution) assignment. Sub-pixel localization is of course possible, and potentially desirable given other data.

5.7.1 Model Detection and Motion

Three instances of detection and localization of the cone model tip are shown in Figure 5.14, which correspond to the model tip ROI of the primitives in Figure 5.3. The smallest scale CSS result is shown by \triangle , and the final tip location selected by the Harris attractor using Equation 5.16 is labeled with a \square . At t_o in Figure 5.14a, all CSS scales and the final Harris assignment are coincident, there was no feature drift. For Figure 5.14b, a single localization through scale occurred followed by a final localization change by the Harris update. The same number of updates occurred at $(t_o + 1678)$ in Figure 5.14c, although the finest CSS point drifted in the opposite direction, and is located closer to the final assignment as opposed to Figure 5.14b. In all illustrated cases, the model tip has been localized well by the final assignment.




Figure 5.14: Model tip localization results. Fine scale CSS point labeled \triangle , final assignment in \square .

A total of 4,912 images were solved using the combined modified CSS-method fused with the Harris detector. The small motion assumption was employed for missing contours, and the remaining 89 images were solved by minimizing distance

between the Harris response operators alone in Equation 5.21. A total of 43 images had multiple solutions to the closest Harris response \mathbf{R} in Equation 5.16 for the final localization step, and had to be evaluated manually. In some cases, the correct response had been returned. In others, the model cone tip solution had drifted from the true response, which was equidistant to a spurious response. Initial distributions of model motion histories were inspected for anomalies and violations of the small motion assumption (deterministic constraint) and suspect points were re-examined. This does not imply that the reported values are incorrect, simply an indication that manual evaluation is called for. A total of 9 responses fit this criterion. Several were corrected and determined visually to be a result of feature drift, due to contour noise at small scale which caused the closest Harris response \mathbf{R} to be located away from the true model tip. Although handling of the small number of outliers and broken curves may appear insignificant, these numbers will increase with larger data sets.

After visual verification and manual corrections, the final distributions of the cone tip motion obtained from the sequence obtained are shown in Figures 5.15a and 5.15b, in terms of row and column position, respectively. Figures are oriented with respect to the image viewing axes for clarity. Each 2 mm division represents a single pixel. The row and column position of the model tip each occupy three discrete bins over the entire sequence of 5,000 images, indicating a degree of model motion small enough to conclude that the double-cone model remains steady throughout the test time. The two-dimensional histogram (point distribution) in Figure 5.16 shows that the cone tip is most often located in the center pixel of the 3 x 3 region. It can

be concluded that model motion is negligible, As all locations of the model tip are 8-connected to the most likely location, within the inherent localization errors associated with the identification methods. This confirms the expected results from viewing the sequence, yet the analysis method developed here provides a means for a quantitative assessment and confirmation.




Figure 5.15: Row and column distributions of the model tip.

5.7.2 Shock Intersection Detection and Motion

Figure 5.17 shows three instances of the shock intersect detection and localization procedure corresponding to the primitive intersection types in Figures 5.9 and 5.10. The finest scale CSS point is shown in \triangle , the final assignment corrected by distance evolution from Equation 5.19 is labeled \square . For t_o in Figure 5.17a, one jump from coarse to fine scale and a final jump due to the evolution distance metric occurred. Figure 5.17b (smooth join) is the most difficult primitive structure to identify and localize. At $(t_o + 3276)$, two position changes occurred during the




Figure 5.16: Two-dimensional histogram of the model tip.

coarse to fine search, followed by a change in the final update step. The algorithm has captured the subtle structure well. For the sharp corner primitive in Figure 5.17c ($t_o + 1678$), a single location update was performed in the scale search; the fine scale and final assignment points are coincident. An example of the coarse to fine search and localization process of the smooth join is shown in Figure 5.18b; the maximum κ at each scale is identified with a solid marker on its plot of curvature *vs.* arclength s . Moving upward from the zero-datum, the maximum drifts leftward toward one of two extrema in noise. The vertical dashed line represents the maximum distance evolution point on the arc closest to the smallest scale curvature. Overall, the shock intersection point has been modeled well. Potential strong violations of the structure definition were defined in the predicate logic, an example of which was shown in Figure 5.12.




Figure 5.17: Shock intersection localization results. Fine scale CSS point labeled \triangle , final assignment in \square .




Figure 5.18: CSS localization of the contour primitives. Local extrema shown in solid.

As with the model tip, 4,912 images were solved using the CSS-based method. The final location was assigned to the closest maximum evolution distance. The remaining 89 shock intersection points were manually located using a graphical user interface. A total of 33 images were marked during the identification algorithm as violating a constraint on the maximum allowed value of the intersection point curvature defined by Equation 5.18. These images were revisited and evaluated visually. Several results were found to be from an edge representation of the image which was physically impossible, for example, a sharp inward directed turn on the outer-shock contour. Other images had been properly localized, but were marked for manual inspection due to large values of curvature at the smallest scale.

An example of a failure which was not detected by the algorithm is shown in Figure 5.19. Several hundred images selected at random were visually viewed at all stages of the detection algorithm, and run multiple times to find this instance. It can be concluded that such instances are few. The method failed due to a poor edge image arising from the outer-shock being defined having a much weaker gradient than a neighboring artifact, as illustrated by the contour image in Figure 5.19b.




Figure 5.19: Failure instance undetected by the algorithm.

In the initial shock intersection motion distributions, statistical outliers were identified using the 3σ rule for the row and column histories. These images were set aside for visual verification. The combined 19 outliers from the row and column distributions corresponded to 14 unique images. Many needed correction due to feature drift caused by contour noise in the localization process. In other images, the algorithm had selected the proper location, which simply appeared as outliers from viewing the distributions.

Final motion distributions for the row and column locations of the shock intersection point are shown in Figures 5.20a and 5.20b, respectively. Data were converted to fluctuating components by subtracting the median. Distribution moments corresponding to these plots are given in Table 5.1. The discrete number of bins in each distribution denotes the pixel-wise extent of the structure motion and gives the unique locations through the sequence. The small number of unique bins in the distribution precludes the use of spectral or temporal analysis. This is a function of the limiting scales of the data set, which are the image resolution of the structure and its absolute motion. This could be remedied by the use of a lens with a larger magnification factor (at the cost viewing the entire flowfield), or by an increased resolution (at the cost of framing rate). The limiting scale would not be a factor when the methodology is applied to data sets acquired with modern cameras with similar framing rates.

From an inspection of the distributions and their moments in Figure 5.20 and Table 5.1, the motion individual motion histories of the intersect coordinates appear Gaussian. This is confirmed by their skewness μ_3/σ_{std}^3 and kurtosis μ_4/σ_{std}^4 , which




Figure 5.20: Row and column distributions of the shock intersection point.

Table 5.1: Distribution properties for the triple point location.

coordinate	σ_{std} [mm]	$\frac{\mu_3}{\sigma^3}$	$\frac{\mu_4}{\sigma^4}$	bins
row	5.291	0.224	3.054	19
col.	3.672	-0.247	2.965	14

are near the nominal values of a normal distribution. Although independent random variables produce normal distributions by definition, the converse is not true. The observed distributions neither confirm nor discount any underlying random motion of the intersection point. Row motion in the image plane has a larger extent than the column location, indicating a greater extent of motion in the vertical than the streamwise direction. This is evident from the values of σ_{std} and the number of discrete bins in Table 5.1, as well as confirmed by viewing Figure 5.20.

The two-dimensional histogram of the shock intersection is shown in Figure 5.21. The motion history as plotted is similar to a discrete joint distribution function where bar height indicates probability of feature location. An alternate view of the

motion is presented in a joint probability distribution style in Figure 5.22. Motion of the shock intersection structure was apparent from viewing the sequence. However, Figures 5.21 and 5.22 show a definite orientation of the motion, with vertical (column) motion of the structure as dominant. The algorithms developed here offer a quantitative evaluation of the motion, and illustrate a pattern and motion extent of the shock intersection motion that could not be inferred by a human observer. From Figure 5.21 and Table 5.1, the dominant motion is in the vertical (column) direction. Comparing values of σ_{std} , the motion spread is 44 % greater in the column than in the streamwise direction. This shows that the inherent breathing motion associated with the shock-wave/boundary-layer interaction on the fore-cone surface dominates the streamwise motion of lower portion of the bow shock in the vicinity of the shock intersection. The motion of the shock intersection observed here explains the cause of the oscillations in peak heating rate on the aft-cone surface reported by Coblish *et al.* [37] The peak heating rate on the model surface is caused by the impinging transmitted shock, which emanates from the shock intersection point. The motion shown in Figure 5.21 is consistent with the description by Druguet, Candler, and Nompelis [42] who state that the separation zone affects the shock impingement, which can alter the entire interaction pattern. The unsteady behavior observed here is consistent with previous findings of the interaction sensitivities.

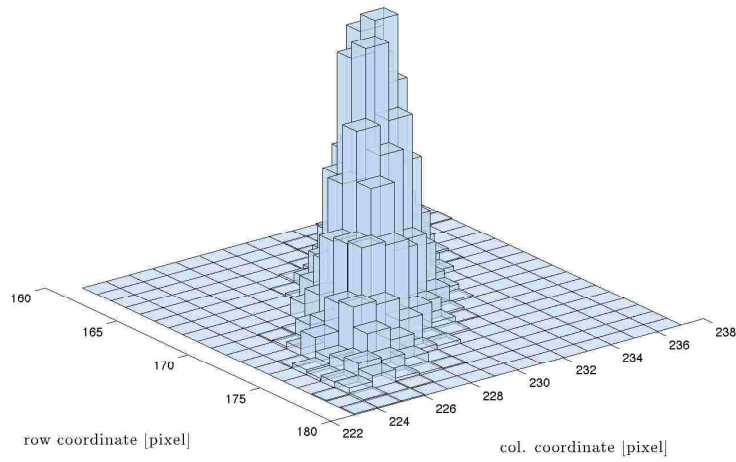


Figure 5.21: Two-dimensional histogram of shock intersection motion.

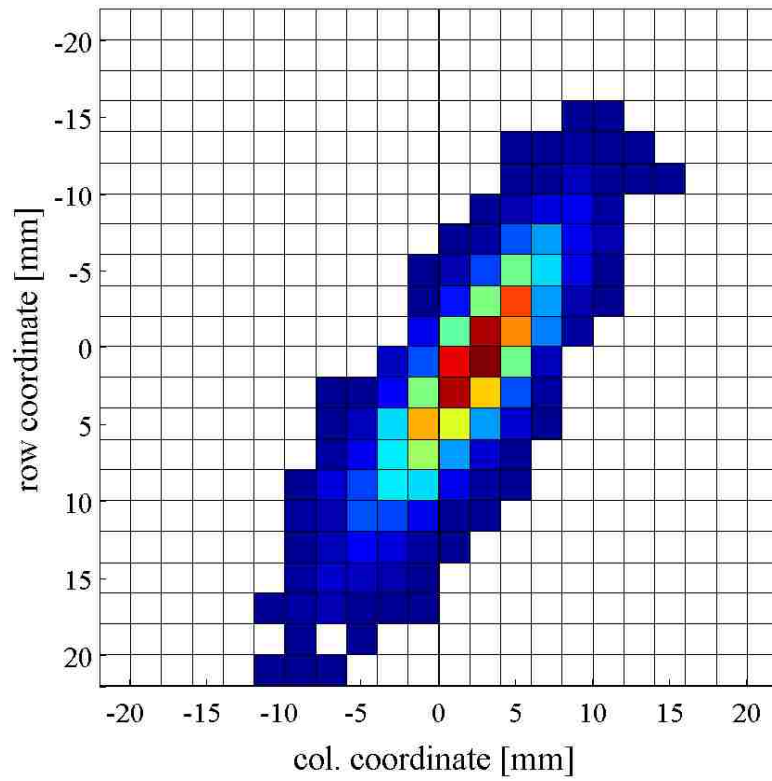


Figure 5.22: Joint probability distribution of shock intersection motion.

Chapter 6: Flowfield Motion Analysis

6.1 Introduction

The preceding chapters of this work have focused on the development and application of computer vision-based techniques for interpreting large schlieren data sets. The methods of measurement and extraction of flow structures is largely general, and applicable to a broad class of data sets which contain these common phenomena. Here, focus is on the analysis and understanding of the double-cone flowfield. Of all the techniques applied, the only data for which there are both enough unique data points to approximate a continuous distribution, and for which further statistical and spectral analysis can provide additional physical understanding of the flowfield, are the segmentation data. The goals of this analysis are twofold: first, to demonstrate that segmentation can permit measurement of gross flow motion, and second, to see what this analysis can reveal about the double-cone flow.

Further analysis of the outer-shock motion identified in Chapter 4 is first performed. The flowfield segment isolated in Chapter 3 is further partitioned by the value of the shock intersection point determined in Chapter 5 into fore and aft-cone segments. Flowfield areas and centroids of all three segments are calculated to perform spatial and temporal analysis on the flowfield as a whole, and its sub-partitions.

6.2 Outer-Shock Motion Analysis

The motion history of the isolated outer-shock structure from Figure 4.7 and discussed in section 4.3.1 is repeated in Figure 6.1. Further analysis of the shock motion can assist in understanding the segmentation data. Black indicates the outer-shock has occupied the image location during the test time. This Figure indicates significant motion of the aft-cone bow shock. The knot which appears just downstream of the model tip is false structure and is not indicative of shock motion. As discussed in section 4.3.1, this is due to insufficient image evidence to capture the faint structure of the fore-cone shock. This issue was addressed at length in section 2.5, and may be due to potential misalignment in the schlieren system.




Figure 6.1: Motion history of isolated outer-shock structure.

The small number of unique bins in the outer-shock motion history precludes temporal or correlation analyses. This does not, however, prevent the estimation of distribution moments which can provide insight into the outer-shock contour mo-

tion. The standard deviation of the outer-shock contour motion history in Figure 6.1 is plotted in Figure 6.2 in units of mm. Locations of the model tip and shock intersection bounds identified in Chapter 5 are shown. The largest motion spreads occur toward the upper portion of the bow shock, increasing toward the frame edge. A significant change in the pattern of σ_{std} occurs near the intersection point. This region, as with the region containing the separation shock identified in Chapter 4, contains both rotational and translational motion. The standard deviation significantly decreases near where the attached conical shock should be on the fore-cone. As discussed in sections 2.5 and 4.3.1, the entirety of this structure could not be consistently captured and was difficult to discern with the human eye. The sharp increase in σ_{std} reflects this. The motion spread is smallest at the model tip, as expected from results presented in the previous chapter.

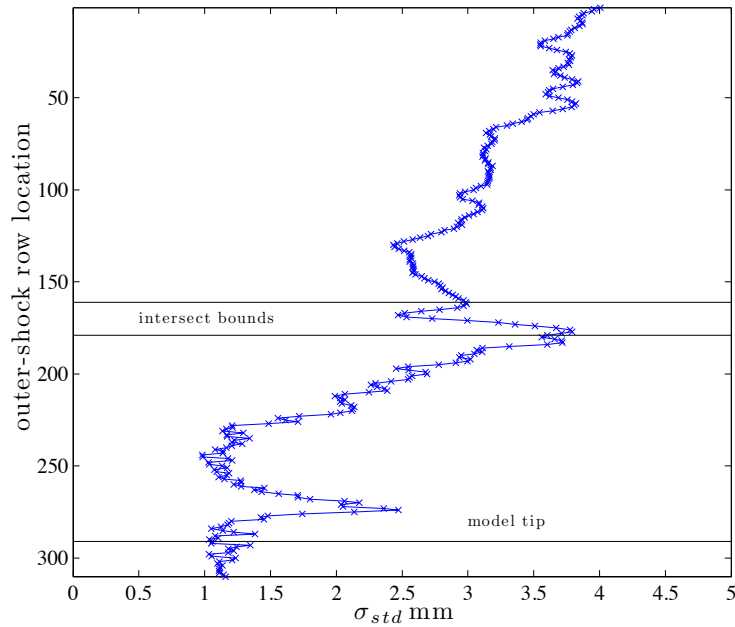


Figure 6.2: Outer-shock motion standard deviation.

Figure 6.3 shows the kurtosis of the outer-shock motion history from Figure 6.1. The model tip and shock intersection motion bounds are indicated. For a normal distribution, the kurtosis should be near 3, which is approximated at most of the outer-shock contour locations. The two-locations where kurtosis is significantly non-normal are the model tip and the region near the “knot structure” shown in Figure 6.1. The low kurtosis at the model tip location is due to its lack of a true distribution caused by its stationarity. The large kurtosis indicates that the knot structure in Figure 6.1 has a large peak with heavy tails; this is confirmation that this portion of the outer-shock was only captured intermittently.




Figure 6.3: Outer-shock motion kurtosis.

6.3 Full Flowfield Motion Analysis

To calculate the properties of the full flowfield segment, images segmented in Chapter 3 (Figure 3.17) were tightly cropped, as illustrated in Figure 6.4. Figure 6.4 shows the image at t_o , with the geometric centroid indicated by \otimes . Spatial and temporal analysis of the full flowfield area A_{FF} and its centroid $\mathbf{x}_{c_{FF}}$ can provide information that can describe the gross motion of the shocked flow region in the double-cone flow. The change in the segment shape, and the motion of its centroid through the sequence are directly related to the fluctuations in, and motion of, the shocked flowfield. Unless otherwise stated, all subsequent analysis is performed on the fluctuating data components.




Figure 6.4: Cropped flowfield segment showing centroid location for $\mathbf{I}(t_o)$. Flowfield centroid indicated by \otimes .

The motion history of the full segment centroid $\mathbf{x}'_{c_{FF}}$ fluctuation is plotted in Figure 6.5 in units of mm. The centroid is resolved in the image coordinate system. The distribution is resolved into real numbers \mathbb{R} as opposed to integers; this is from

the definition of the centroid. The centroid spans approximately 8 mm and 12 mm in the streamwise and normal direction, respectively, through the test time. Pearson's correlation coefficient r_p , defined in Equation 6.1 below, was calculated to examine the linearity of the shocked flowfield motion

$$r_p = \frac{\sum_{i=1}^N (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^N (X_i - \bar{X})^2} \sqrt{\sum_{i=1}^N (Y_i - \bar{Y})^2}}. \quad (6.1)$$

A value of $r_p = 0.7221$ does not indicate a strong linear trend. The motion is bounded, with a slight Gaussian appearance from inspection of Figure 6.5. The centroid motion is overlaid on the outer-shock motion history in Figure 6.6, and shown in red. The motion bounds of the full flowfield centroid appear small. However, this centroid encompasses the entire flowfield area change throughout the test time. A large deviation in the upper portion of the bow shock may be offset by a decreasing area in the fore-cone region, or a smaller motion in the lower portion of the bow shock. The greater distance covered by the vertical (column) motion of the centroid suggests a larger fluctuation in the upper bow shock region. This is substantiated by Figure 6.2, which shows that the upper portion of the bow shock has a larger motion than the lower segment.

The distribution moments of the flowfield centroid corresponding to Figure 6.5 are listed in Table 6.1. Inspection of σ_{std} of the row and column distributions indicate a nearly 40% greater motion spread in the row (vertical) direction of motion through the sequence as opposed to the column (streamwise) direction. This is consistent with the orientation of the motion in Figure 6.5. The number of unique




Figure 6.5: Full flowfield segment centroid $\mathbf{x}'_{c_{FF}}$ fluctuation motion history. Axes oriented to image.




Figure 6.6: Full segment centroid motion overlaid on outer-shock motion history.

values distributions in Table 6.1 are sufficient for spectral and statistical analysis. The skewness and kurtosis of the centroid $\mathbf{x}'_{c_{FF}}$ in Table 6.1 are close to the normal distribution. Temporal behavior cannot be deduced from a spatial distribution, and is considered in a later section.

Table 6.1: Full segment centroid fluctuation \mathbf{x}'_c distributions.

\mathbf{x}'_c	σ_{std} [mm]	$\frac{\mu_3}{\sigma^3}$	$\frac{\mu_4}{\sigma^4}$	bins
row	1.8845	-0.0020	2.9687	494
col.	1.3670	0.1235	3.1217	381

6.4 Fore and Aft-Cone Flowfield Relations

The flowfield segmentation results in Figure 6.4 were further partitioned by the shock intersection point determined in Chapter 5. This operation parsed the isolated flowfield into fore and aft-cone flowfield segments, which roughly isolates the shock-wave/boundary-layer interaction region in the fore-cone from the aft-cone flow phenomena behind the detached bow shock. Figure 6.7 shows the segment at time t_o from Figure 6.4 split by the column (vertical) coordinate of the detected shock intersection. Figure 6.7 illustrates a typical flowfield partitioning, determined from the visual inspection of multiple sample results. The fore-cone segment is identified in black; the isolated aft-cone segment is gray. The centroid locations of the fore and aft-cone centroids are indicated by \otimes .

Distribution properties of the full flowfield A_{FF} , the fore-cone A_{FC} , and aft-cone A_{AC} segment areas are given in Table 6.2 in pixels. All segments are based on 8-connected graph partitions. Properties of A_{FF} are repeated from Table 3.1. Distribution moments for each segment indicate near-normal distributions, evidenced by skewness $\frac{\mu_3}{\sigma^3}$ and kurtosis $\frac{\mu_4}{\sigma^4}$. The number of bins in each distribution justify spectral and stochastic estimates for the image sequence. From inspection of the three flowfield standard deviations σ_{std} in Table 6.2, the full segment and the aft-cone




Figure 6.7: Cropped flowfield segment split into fore and aft-cone regions by the detected shock intersect for $\mathbf{I}(t_o)$. Flowfield centroids indicated by \otimes .

segment appear to have the same data spread in terms of area size. The standard deviation of the fore-cone flowfield is approximately 50% less than its counterparts; this is not surprising, as from viewing the sequence, the aft-cone segment seems to constitute the bulk of the flow motion. This was also evident in the outer-shock motion history shown in Figures 6.1 and 6.2.

Table 6.2: Flowfield segment area distributions for the full, fore-cone, and aft-cone [pixels].

Area	μ	σ_{std}	$\frac{\mu_3}{\sigma^3}$	$\frac{\mu_4}{\sigma^4}$	bins
A_{FF}	11482	212	0.0483	3.110	984
A_{FC}	2425	118	0.0158	3.009	604
A_{AC}	9056	217	0.1019	3.095	1003

6.4.1 Fore-Cone Region Motion Analysis

The motion history of the fore-cone segment centroid $\mathbf{x}'_{c_{FC}}$ is shown in Figure 6.8. The fluctuating component, given in mm, is plotted in the image coordi-

nate system. The motion history of the fore-cone flow corresponds to the shock-wave/boundary-layer interaction region of the flow shown in black in Figure 6.4. Pearson's correlation coefficient was calculated to examine the linearity of the segment motion. A coefficient of $r_p = 0.9807$ indicates that the movement is dominantly linear, and along the direction of the model surface. This is indicative of the expansion motion of shock separated flows. Figure 6.8 shows that the largest motion of the separated region is in the streamwise direction. The centroid motion is overlaid on the motion history of the outer-shock in Figure 6.9.




Figure 6.8: Fore-cone flowfield segment centroid $\mathbf{x}'_{c_{FC}}$ fluctuation motion history. Axes oriented to image.

The distribution moments of the fore-cone centroid are given in Table 6.3. The column (streamwise) motion distribution has a value of σ_{std} nearly 50% greater than that of the row (vertical) motion. This indicates that the streamwise motion of the fore-cone flow region spans a much larger spatial area than the vertical component.




Figure 6.9: Fore-cone segment centroid motion overlaid on outer-shock motion history.

The span of streamwise motion of the separated flow region is approximately 30 mm, significantly larger than the 8 mm span of the full flowfield segment shown in Figure 6.5. Vertical extent of the centroid motion is approximately 18 mm, compared to 12 mm for the full flowfield segment centroid $\mathbf{x}'_{c_{FF}}$.

Table 6.3: Fore-cone segment centroid fluctuation \mathbf{x}'_c distributions.

\mathbf{x}'_c	σ_{std} [mm]	$\frac{\mu_3}{\sigma^3}$	$\frac{\mu_4}{\sigma^4}$	bins
row	3.0346	0.1561	3.0699	749
col.	4.4985	-0.2532	3.0988	1040

The large scale streamwise motion evident in Figure 6.8 and Table 6.3 is likely due to the motion of the separation shock foot at the point of separation onset. This shock can be a significant source of unsteady motion in typical shock-wave/boundary-layer interactions. The shock foot could not be captured by any of the vision-based methodologies, since it was too faint to visualize clearly, as discussed in section 2.5.

6.4.2 Aft-Cone Region Motion Analysis

The motion history of the aft-cone segment centroid $\mathbf{x}'_{c_{AC}}$ is shown in Figure 6.10, oriented to the image axis and scaled to mm. Pearson's correlation for the isolated aft-cone centroid is 0.7276, which is on the order of the value for the full flowfield centroid motion. The streamwise and vertical span of the centroid motion of 9 and 14 mm, respectively, are slightly larger than that those of the full flowfield. This is consistent with the observation that the apparently small motion of the full flowfield may be due to mitigation caused by events in the fore-cone segment. This suggests large scale bulk motions in the aft-cone flow region, which is consistent with motion patterns of the aft-cone bow shock in Figures 6.1 and 6.2. The centroid motion history plotted along with the outer-shock motion history in Figure 6.11 shows a different location and larger spread of the aft-cone centroid motion as compared to that of the full flowfield shown in Figure 6.6.

Distribution moments of the aft-cone centroid motion are listed in Table 6.4. The row (vertical) σ_{std} is approximately 33% larger than the column (streamwise) centroid motion, consistent with the motion history in Figure 6.10. Data spread in the motion of the aft-cone centroid is larger than that of the entire flowfield by approximately 24% and 31% in the vertical and streamwise directions, respectively.

Table 6.4: Aft-cone segment centroid fluctuation \mathbf{x}'_c distributions.

\mathbf{x}'_c	σ_{std} [mm]	$\frac{\mu_3}{\sigma^3}$	$\frac{\mu_4}{\sigma^4}$	bins
row	2.3434	0.2169	3.0438	600
col.	1.7497	-0.1877	3.0505	471




Figure 6.10: Aft-cone flowfield segment centroid \mathbf{x}'_{CAC} fluctuation motion history. Axes oriented to image.




Figure 6.11: Aft-cone segment centroid motion overlaid on outer-shock motion history.

6.4.3 Flowfield Area Correlations

Pearson's correlation coefficients are used here to test for relationships between the partitioned flowfield regions, as opposed to motion linearity in the preceding

sections. All segment areas are normalized; the fluctuating components are scaled by their respective standard deviations σ_{std} . Correlations between the segmented areas are shown in Figure 6.12. Figure 6.12a plots the aft-cone area fluctuations A'_{AC} against the full area fluctuations A'_{FF} . From inspection, these events are correlated strongly, and a correlation coefficient of $r_p = 0.8491$ confirms this. An important conclusion can be drawn from Figure 6.12a, namely, that the bulk of the unsteady motion in entire the double-cone flowfield is due to fluctuations in the aft-cone area.

Figure 6.12b shows the fore-cone area fluctuations A'_{FC} against those of the full flowfield A'_{FF} . A coefficient of $r_p = 0.2341$ suggests a weak positive correlation between the fore-cone segment and the full flowfield, indicating a slight correlation with sign change. This is not surprising, as from viewing the sequence it appeared that the motions of the aft-cone bow shock region occur on a much larger scale than those of the fore-cone region associated with the shock separated flow. From Figures 6.12a and 6.12b and their correlation coefficients, it can be concluded that that the large scale motions seen in the sequence are primarily due to motion of the largely subsonic region behind the bow shock, and occur downstream of the shock intersection. This is further supported by Figure 6.2 which showed that the motion spread in the upper portion of the bow shock is much larger than that of the fore-cone area.

The aft-cone area A'_{AC} fluctuations are plotted against the fore-cone flowfield A'_{FC} fluctuations in Figure 6.12c. The correlation coefficient between the fore and aft-cone area is weak; $r_p = -0.3148$ indicates only a slight inverse relation between the two partitions. A stronger relationship between the two fluctuations might have

been expected. However, the aft-cone region may be either wholly or partially subsonic, so a direct proportionality between this region and the shock separated portion of the flow may not exist, due to the lack of a dominant feedback mechanism between these regions as a whole. Also, the aft-cone region is treated as a single entity, whereas in actuality the outer-shock structure may be more sensitive to changes in the supersonic jet along the model surface than to any upstream disturbances from the shock separated region of the flow.

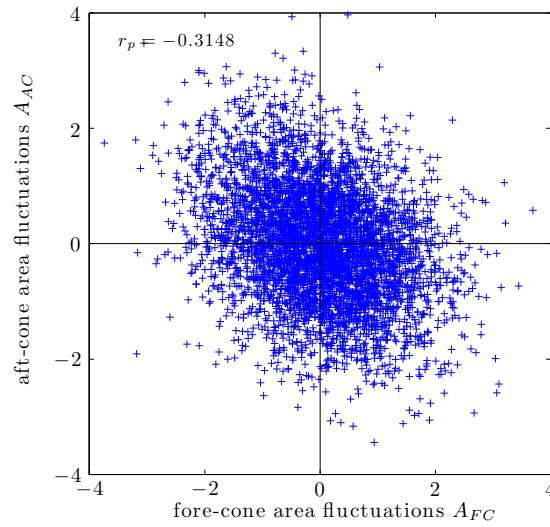



Figure 6.12: Standardized flowfield segment area change correlations.

6.5 Flowfield Spectral Content

The changes in flowfield segment shape and its sub-partitions into fore and aft-cone regions encompass all unsteady events in the shocked flowfield. All frequencies and length scales associated with the flow will affect segmented areas. In addition to frequency content within the camera's Nyquist frequency of 5 kHz, aliased phenomena may be present at harmonics much lower than their true frequency. In an unsteady Mach 14 flowfield, multiple phenomena can be expected in a frequency range much greater than 5 kHz. The flowfield segments will therefore contain a summation of an unknown number of events, and are likely to be noisy for two reasons: there may be random motion(s) present in the flowfield, and the segmentations may contain small errors due to segment resolution scales and 8-connectivity. This can cause the appearance, or a potential source, of noise. All processes are assumed Wide Sense Stationary (WSS). The assumption is made that frequency content is present in the segmented data fluctuations, and that it may be obscured by significant noise. All data are therefore tested heuristically for whiteness by examining content of the autocorrelation functions, and flatness of power spectral densities.

Without high-speed surface instrumentation data, there are no expected values for the number of motion components in the flowfield, or their associated frequency content. The lack of high-speed instrumentation on the double-cone model significantly hinders power spectral estimation of the computer vision results. The absence of known signal information precludes the use of more advanced signal recovery and parametric spectral estimation methods than those employed here, and make fre-

quency content difficult to recover. Initial studies to ascertain tunnel noise have been conducted at Tunnel 9 [76]. Results indicate that noise in the freestream is broadband in nature, at least up to the 20 kHz range of the flush mounted pitot acoustic probes used in the study.

6.5.1 The Full Flowfield

Plots of the fluctuating components of full segment flowfield area A'_{FF} and its centroid motion $\mathbf{x}'_{c_{FF}}$ against test time revealed no obvious sinusoidal patterns. As the full flowfield segment consists of all length and time scales in the entire shocked flowfield, autocorrelations were examined to test for any underlying frequency content. The autocorrelation coefficient in Equation 6.2 is measure of signal similarity with itself at lag τ .

$$R_{xx}(\tau) = \frac{1}{N} \sum_{t=0}^{N-\tau} x'(t) x'(t + \tau) \quad (6.2)$$

If the data are truly white, no patterns whatsoever should appear in the autocorrelation data, even in the initial lags. The presence of noise can drive all coefficients at $\tau \neq 0$ to zero exponentially. When SNR is sufficiently low, data trends can be obscured. If the data are truly random, there will be no significantly non-zero coefficients aside from $\mathbf{R}_{xx}(\tau = 0)$. Significantly non-zero values are defined as those outside the bounds $\pm 1.96/\sqrt{N}$ (*e.g.* [77]). In the presence of noise or randomness, inspection of the correlogram can assist in the determination of underlying periodicity.

Normalized autocorrelations are plotted in Figure 6.13. Rows in Figure 6.13 from top to bottom give autocorrelations for the area, row, and column centroid motion, respectively. The first column plots the autocorrelations for the first 500 lags (0.05s). The second column is a closeup on the first 100 lags (0.01s). From initial inspection, any periodicity in the area and centroid row autocorrelations over the 1st 500 lags in Figure 6.13 is difficult to determine. The column (streamwise) centroid motion, however, appears to have a definite “chirp” pattern over the first 500 lags, indicating the presence of multiple signals. Close inspection of the first 100 lags in the second column show approximately 7.5 cycles in each trend. This pattern is most clear in the streamwise autocorrelation, but evident in the area and centroid row functions correlograms as well, although it appears partially obscured by noise, and possibly additional frequency content. Periodicity is more clear in the centroid components; this is likely due to the averaging operation of the area from the calculation of the centroid. The 7.5 cycles visible through the first 100 lags should correspond to a frequency peak at 750 Hz.




Figure 6.13: Normalized autocorrelations for area and \mathbf{x}_c fluctuations.

Distributions for the full flowfield segment area A'_{FF} in Table 6.2 appeared normal. Normality of the distribution does not imply randomness, when significant noise and multiple frequency content exist in the signal. Welch's method is used to estimate spectral density. Lack of known signal structure that could have been provided by surface instrumentation precludes the use of parametric methods which are better suited to recover partially obscured signals. Periodogram averaging provides a good estimation in the presence of noise. A heuristic estimate of noise was gathered for each periodgram presented here. As the distributions of all measurements appear normal in the previous section, a Gaussian noise estimate was created using Equation 6.3 below

$$\hat{\eta} = \sigma_{std} \mathcal{N}(\mu = 0, \sigma = 1). \quad (6.3)$$

where the normal white noise distribution function \mathcal{N} is shaped with the σ_{std} of the variable of the measured variable. If the data are truly random, then the spectral density estimate of the data should be bound by the maximum frequency in the spectrum of Equation 6.3, as they are both approximately normal, with the same σ_{std} . For all power spectral estimates presented here, the calculated power spectrum was compared against that of the test white noise sequence shaped by σ_{std} . This provided an upper estimate of noise, above which, a legitimate frequency peak could be declared.

The power spectral density of the full flowfield area fluctuations A'_{FF} is shown in Figure 6.14. All power spectral densities are plotted in terms of arbitrary units

(a.u.); and only relative levels in the peaks are of concern. A 1024 point Hamming window with 8 overlapping segments yielded the best trade-off between spectrum estimation and leakage in the periodogram estimate, providing a frequency resolution of $\Delta f = 4.883$ Hz. Spectral density of the noise and the upper value of the noise estimation $\max(\hat{\eta})$ are shown. Any peak above the noise estimate $\hat{\eta} \approx 23$ can be considered true frequency as opposed to artifact. Significant frequency peaks are evident at 740 and 835 Hz. A lesser peak exists at approximately 884 Hz. Potential peaks at approximately 5 and 54 Hz are just below the upper noise estimate. The 740 Hz frequency is consistent with the conclusion from the first 100 lags of the autocorrelations in Figure 6.13.




Figure 6.14: Power spectral density of full segment area fluctuation A'_{FF}

Figures 6.15 and 6.16 show the power spectral densities of the row and column centroid motions, respectively. Noise estimates generated using Equation 6.3 are also shown. The frequency peaks of 740, 835 and 884 Hz from the area fluctuation periodogram in Figure 6.14 are also evident in the row and column centroid motions.

Additionally, peaks have become more clarified at approximately 54 and 5 Hz; an additional peak is observed at 10 Hz. The peaks at 5 and 10 Hz occur within the first two bins of the Δf periodogram resolution. The reason the lesser peaks are evident in Figures 6.15 and 6.16 as opposed to the corresponding area spectral estimation is likely due to the effect of averaging across the large segment to calculate the centroid, as was observed in the correlogram discussion.




Figure 6.15: Power spectral density of full segment row centroid fluctuation x'_{cFF} .




Figure 6.16: Power spectral density of full segment col. centroid fluctuation x'_{cFF} .

6.5.2 The Fore-Cone Flow

Normalized autocorrelations of the fore-cone isolated segment A'_{FC} and its centroid \mathbf{x}'_{cFC} fluctuations are shown in Figure 6.17. The first column of Figure 6.17 shows the autocorrelations through the first 500 lags (0.05s); the second column is a closeup of the first 100 lags (0.01s). The first, second, and third rows of Figure 6.17 show the area, row, and column centroid autocorrelations, respectively. The first 500 lags of the autocorrelations show no evidence of periodic motion in the signals. From close inspection of the first 100 lags of the area fluctuations, and with assistance from the correlograms of the full segment in Figure 6.13 to interpret the data, there is faint evidence of 7.5 cycles in the data. There is no evidence of periodicity in either component of the centroid data. The fact that the autocorrelations of the row and column motion are above the confidence bounds may indicate that the fore-cone centroid motion is non-WSS through the test time. In section 2.5, it was noted that

the laser began to fail approximately half way through the test. As a result, the images became darker; this may have an effect on the visualization of the attached flow near the cone tip. The visualization, could in turn affect the segment size of the fore-cone. The weak appearance of the 7.5 cycles in the area data could be due to either actual signal noise, or the presence of random motion. Here, actual signal noise may arise from segmentation errors. The small fore-cone segment would be more sensitive to slight errors in the segmentation, since this area is much smaller than the full flowfield area. A segmentation error of several pixels may therefore result in the appearance of significant noise which may dominate the signal. This is simply a function of the small size of the fore-cone segment compared to the limiting scale (resolution) at which the images were acquired.

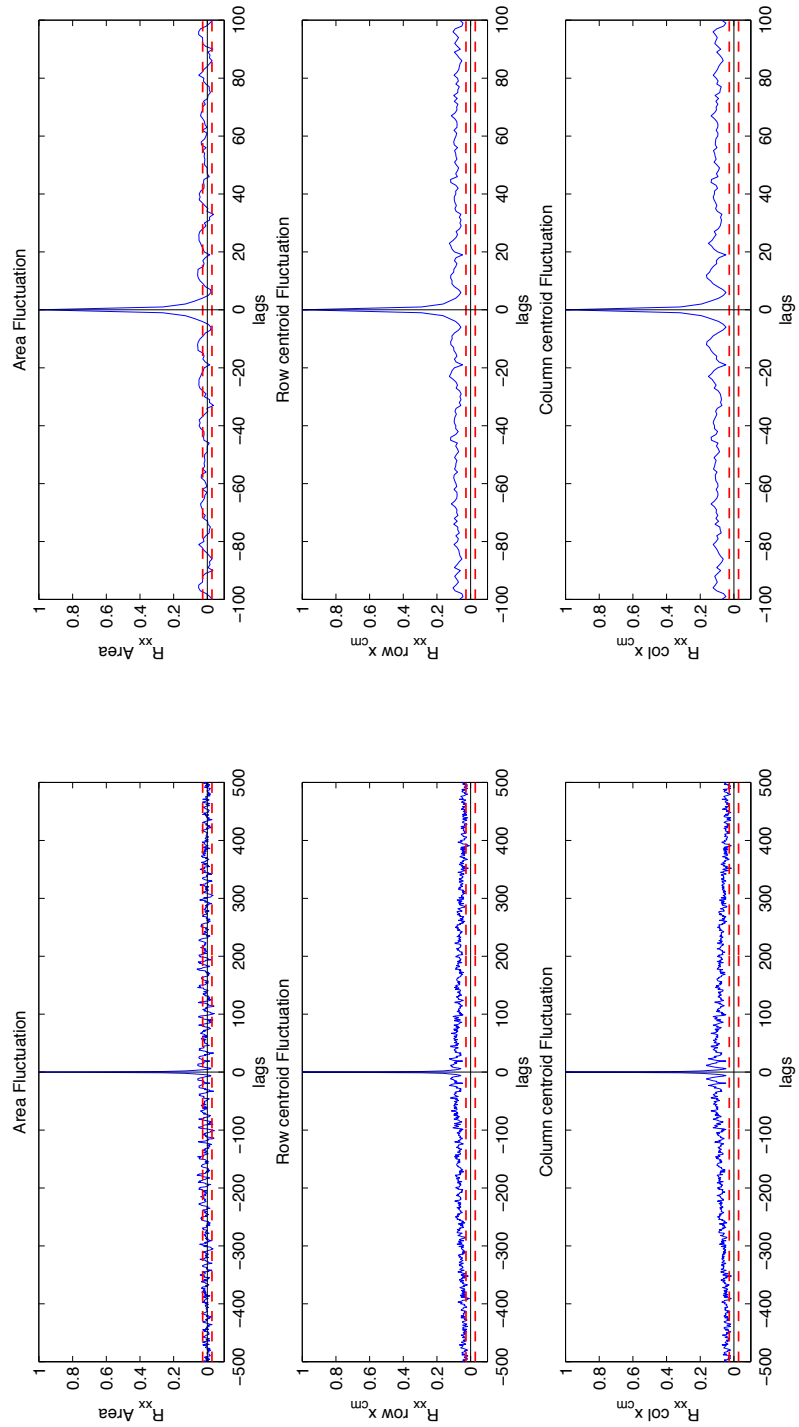


Figure 6.17: Normalized autocorrelations for fore-cone area and x_c fluctuations.

Power spectral density estimations of the fore-cone area fluctuation A'_{FC} are presented in Figure 6.18. Peaks above the estimated noise-band in the power spectrum of A'_{FC} appear at 740 and 835 Hz, and approximately 50 and 5 Hz. A frequency cluster that may not be properly resolved appears at approximately 434 Hz. A potential peak may exist near 1 kHz, but this is very close to the noise limit. The power spectrum in Figure 6.18 contains more noise than that for the full segment A'_{FF} in Figure 6.14, due to reasons discussed during the correlogram interpretation, yet the peaks are still distinct. The higher noise level relative to the peaks is likely caused by small segmentation errors. Such an error may be caused, for example, by the lack of consistent image evidence to visually identify the attached conical shock on the model tip, as discussed in section 2.5. This is consistent with the assumption that the segmentations contain frequency content which may be partially obscured by noise.




Figure 6.18: Power spectral density of fore-cone segment area fluctuation A'_{FC}

Spectral estimates of the fore-cone centroid $x'_{c_{FC}}$ row and column motion are shown in Figures 6.19 and 6.20, respectively. There is no significant spectral content in the row centroid motion in Figure 6.19. Peaks observed in the full area periodogram in Figure 6.14 are evident but weak in the column centroid motion at 742 and 884 Hz. These frequencies are dominated by peaks at 0 and approximately 5 Hz, the first two frequency bins in the spectrum. Whereas the averaging inherent in the centroid calculation served to highlight features in the full segment; the same effect is not seen here. This is likely due to the averaging process in the centroid calculation occurring at a reduced scale over the smaller fore-cone segment, where small pixel errors can constitute a larger percentage of the segment.




Figure 6.19: Power spectral density of fore-cone row centroid fluctuation $x'_{c_{FC}}$.




Figure 6.20: Power spectral density of fore-cone col. centroid fluctuation x'_{cFC} .

6.5.3 The Aft-Cone Flow

Figure 6.21 shows the autocorrelations of the aft-cone area fluctuations A'_{AC} and centroid row and column fluctuations \mathbf{x}'_{cAC} . The first and second columns of Figure 6.21 are again the first 500 and 100 lags corresponding to the first 0.05s and 0.01s of the sequence, respectively. There is no evidence of periodicity whatsoever in any of the autocorrelations in either the first 500, or 100 lags. The correlograms are indicative of data whiteness, and correspond physically to large scale broadband motion of the aft-cone bow shock region of the double-cone flow. This seemed evident from viewing the sequence, and the large scale motion of the bow shock was shown in Figure 6.2. From Figure 6.21, it appears that the aft-cone motion is random, at least within the Nyquist frequency of 5 kHz.

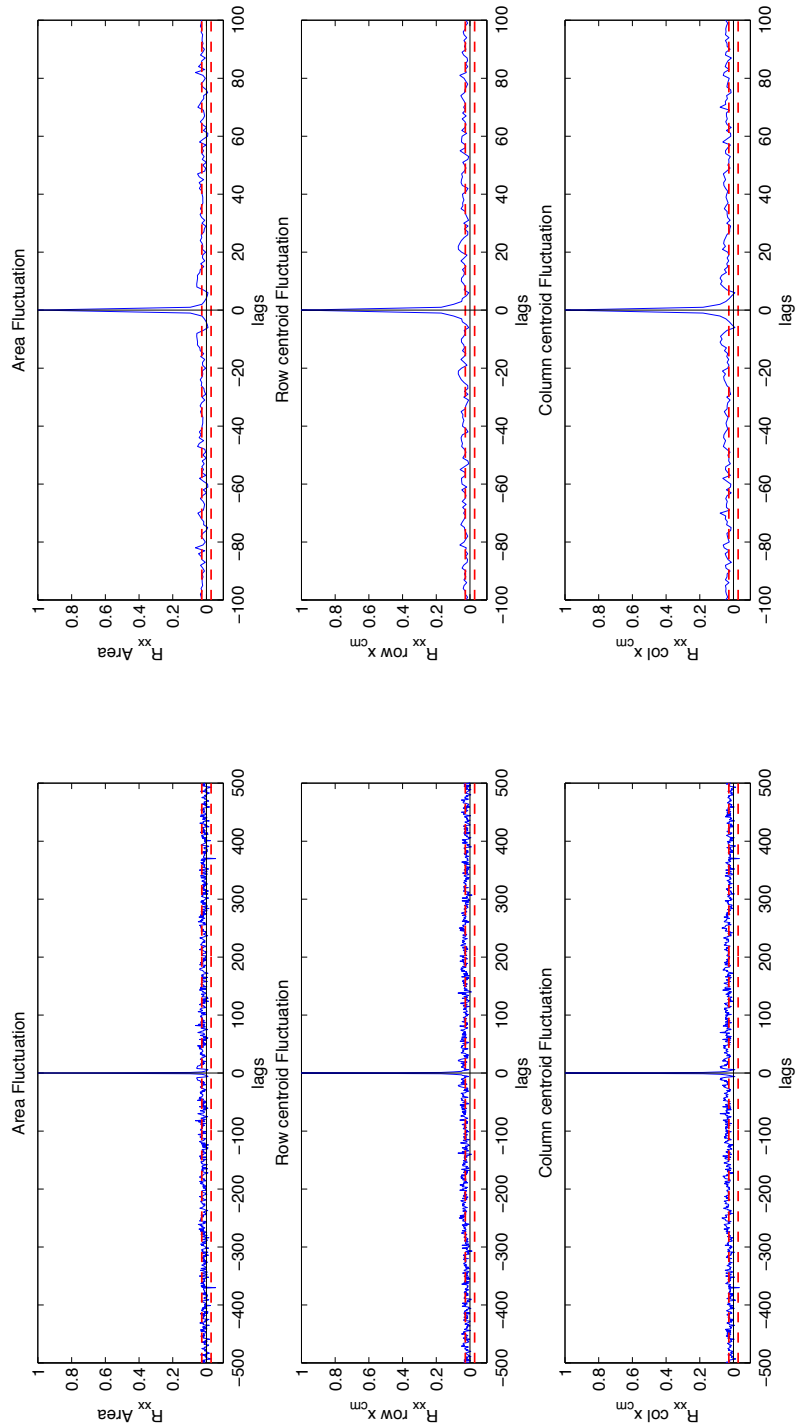


Figure 6.21: Normalized autocorrelations for aft-cone area and x_c fluctuations.

The power spectral estimates of the aft-cone flowfield fluctuations are shown in Figure 6.22. Compared against the test white noise sequence, the spectrum of A'_{AC} has no significant frequency content above the estimated noise level; the apparent peak at 884 Hz is too close to $\max\hat{\eta}$ to be considered significant. The aft-cone fluctuations can therefore be considered broadband and random, within the 5 kHz bandwidth. The same observation holds for the column centroid periodogram in Figure 6.24, and the row centroid component in Figure 6.23 appears to be noise. From the power spectrum estimates of the full A'_{FF} and fore-cone A'_{FC} areas in Figures and 6.14 6.18, it can be concluded that the distinct frequency peaks of 740, 835, and 50 Hz are associated with the shock separated flow in the fore-cone region, while the flow downstream of the shock intersection is large scale and random within the 5 kHz bandwidth.




Figure 6.22: Power spectral density of aft-cone segment area fluctuation A'_{AC}




Figure 6.23: Power spectral density of aft-cone row centroid fluctuation x'_{cAC} .




Figure 6.24: Power spectral density of aft-cone col. centroid fluctuation x'_{cAC} .

6.6 Flowfield Motion Summary

Motion analysis of the double-cone isolated outer-shock contour across the entire sequence revealed that the upper portion of the aft-cone bow shock undergoes the most spatial change during the test time. The lower portion of the bow shock moves to a lesser extent. Considerable motion was evident near where the separation shock interacts with the bow shock; this motion includes both rotation and translation. The motion analysis confirmed that the attached conical shock and the point of separation were difficult to consistently capture, most likely due to a misalignment of the schlieren system discussed in section 2.5.

The segmentation results from Chapter 3 were further partitioned into fore and aft-cone regions using the shock intersection point detected in Chapter 5. Analysis of spatial motion showed that the motions of the flowfield as a whole, and the aft-cone segment appear Gaussian. The motion of the smaller scale fore-cone segment is nearly linear, and has the largest extent of centroid motion in the three flow partitions, consistent with observations in typical shock-wave/boundary-layer interactions. Correlations showed that the dominant unsteadiness observed in the sequence are due to fluctuations in the largely subsonic aft-cone region associated with the detached bow shock. No significant correlations were observed in comparing the fore-cone fluctuations to both the full flowfield, and the aft-cone region. This suggests that fore-cone motion, although more organized, occurs at a much smaller scale.

It was demonstrated that true frequency content, distinct from artifacts of signal processing or noise, can be recovered from segmentation data. The aft-cone motion was shown to be broadband in frequency content from spectral estimation, and large scale from motion analysis of the outer-shock structure. Definite frequency content at approximately 50, 740, and 835 Hz was shown to exist in the fore-cone region from peridogram estimations, and this motion was shown to be small scale and organized from centroid motion analysis. Researchers have stated that one of the most difficult aspects of simulating double-cone flows is accurate modeling of the multiple length scales and the coupling of the shock- separation and shock- interaction [43]. Due to the inherent complexity of the flowfield, and lacking high-speed surface instrumentation, we are not presently able to draw definite conclusions about the source of these frequencies, other than their association with the shock-separated flow on the fore-cone.

Chapter 7: Conclusions and Future Work

7.1 Conclusions

The analysis of large schlieren data sets has been recast as a computer vision problem, using a unique and interdisciplinary approach. The complex double-cone sequence was used as a testbed for the development and application of vision-based data extraction and measurement of features and flow structures of interest. Success rates of the techniques presented here when applied to a noisy sequence, with features at small spatial resolution, implies the methodology can be applied to classic or degraded data sets, and promises to produce better quantitative results given sequences with higher resolution and greater temporal uniformity of illumination. The application of computer vision techniques has yielded physical information from the schlieren sequence that was previously difficult, or not possible, to obtain. Emphasis has been given to vision techniques and approaches which apply not only to a broad spectrum of schlieren sequences, but also to shadowgraph visualizations and other flowfield imaging results. Computer vision analysis provides a natural and meaningful extension to existing data analysis capabilities.

Vision techniques were used to describe linear structures in schlieren sequences. Using physical rules as prior knowledge, a classifier was developed to extract aero-

dynamic features of interest. Image representations were used which allowed for incorporation of rules which the flow structure needed to obey to be a reasonable candidate for a shock or contact surface, eliminating a random line fragment from consideration. Over a noisy sequence of 5,000 images, separation shock wave and contact surface angles in a small region of interest were measured at each frame with nearly 97% and 94% success rates, respectively. Results were substantiated by extracting measurements across accumulator grids at two scales. Final measurements were taken from the finest resolution distribution, subject to a similar observation at a coarser scale. A concise description of this technique is presented in reference [78].

Feature detection algorithms from computer vision have been used to define, localize, and track structures that occur on shock contours. Grayscale feature detection methods can be spurious, have difficulty measuring large scale structures, and may not allow for incorporation of expert knowledge in a manner that can be used to uniquely define shock structures of interest from a fluid mechanics perspective. Features based on curvature offer an inherent localization to structures of interest on shock contours, and the curvature properties may define points of interest directly. The CSS method was modified; signed curvature was retained and no thresholding was employed to provide a low-level canonical description of shock contours. Features were defined by constructing their definitions through implementation of a knowledge base rather than by employing an arbitrary threshold, which lacks physical meaning. This allowed the hard definition (as opposed to a potential probable correlation) of a sharp model tip and a shock intersection region, which changes appearance through the test time, to be localized and tracked.

Results were localized through scale, and a final update was performed using a separate metric. Feature uniqueness circumvents the correspondence problem, which justifies its use as a tracking method. Multiple failure instances were accounted for when designing the detection algorithm to handle potential false representations of the shock contour based on fluid dynamic rules. Strong violations of the knowledge base rejected physically impossible image representations, and an example of a weak violation that was missed by the knowledge base was shown. This resulted in a robust semi-supervised method necessary for the evaluation of large data sets which is summarized in [79].

The double-cone model appeared still from examining the image sequence; the implementation of the algorithm provided a quantitative confirmation of model stillness. Motion of the shock intersection point was evident from human inspection of the sequence. The application of a feature detection and tracking technique from computer vision methods provided a quantitative measurement of the movement and revealed a definite pattern of motion. This motion correlates directly to oscillations in peak heating observed on the aft-cone surface in previous computational studies of the double-cone. Due to the techniques developed here and the analysis performed, model motion can be ruled out as a cause of flow unsteadiness.

The shocked flowfield was isolated using image segmentation. This segment was further partitioned by the detected shock intersection point. This grouped the flowfield into the fore- cone region, which contains the shock-wave/boundary-layer interaction and the aft-cone region, which contains the region downstream of the shock intersection associated with the detached bow shock and the supersonic

jet along the model surface. Motion analysis and power spectral estimates were performed on all three segments. It was shown that the bulk of the unsteadiness observed in the double-cone sequence is due to large scale random motions of the aft-cone flow. The shock-separated flow region occurs on a smaller spatial scale than the either the bulk or the aft-cone flow, but is strongly organized and appears to contain frequency content at approximately 50, 740, and 840 Hz. The analysis holds within the 5 kHz bandwidth. Lacking high-speed surface instrumentation, it is difficult to attribute these frequencies to anything other than the gross motion of the shock-separated flow, or to apply more advanced spectral estimates and signal recovery methods to better isolate the spectral peaks from potential noise sources.

The motion analysis and power spectral density estimates performed on the segmentation results and its sub-partitions show that image partitioning of a schlieren sequence into physically meaningful events can permit a measurement of flow unsteadiness. It can also provide for motion analysis and correlations which are not possible to achieve through traditional surface instrumentation. Small segmented regions can be more sensitive to noise than larger areas. This is simply a function of image resolution. Future benchmark data sets, coupled with high-speed surface instrumentation that was lacking in this study, can be used to calibrate segmentation methods as a means of measurement.

7.2 Original Contributions

- Computer vision-based tools have been developed that can be used to analyze supersonic schlieren data. To the author's knowledge, this is the first application of computer vision to large schlieren data sets.
- A method and means for recognizing and classifying oblique structures was developed. A separate methodology was designed to define, localize, and track key features on shock contours. The methods enable automation of the interpretation of large image sequences, adding quantitative statements to qualitative human inference.
- It has been demonstrated that image segmentation can be used to partition and label flowfields providing information that was not previously possible to obtain, such as flowfield motion correlations and spectral content.
- General methodologies for automatically identifying and classifying common flow structures and features have been developed so that the approaches used here may be applied to other schlieren and shadowgraphy data.
- The integration of a computer vision approach with compressible flow visualization data provides a future direction for analyzing and understanding large data, as such tools are necessary.

7.3 Directions for Future Work

Researchers are capable of acquiring thousand of frames of data for a given test sequence. This research has shown that computer vision-based analysis is viable, and should be used to better understand unsteady flows. Optical data can provide insight into mechanisms which surface instrumentation may not. Vision is a rapidly evolving field, and not yet ready to be applied without a significant understanding of the discipline. Computer vision has had significant impacts on medical imaging and understanding through cooperation between vision and medical researchers; vision has the potential for a similar impact on schlieren and shadowgraph understanding. For example, Piponnier *et al.* [80] proposed a mechanism for large scale motion of the separation bubble in shock-wave/boundary-layer interactions through mass drain across the discriminating streamline. Given a sequence with sufficient contrast, segmentation can be used to measure not only the time scale, but also the length scales of separation length and bubble height directly from the optical data.

Application of these techniques to sequences obtained with modern cameras can yield significantly improved results in one of two ways: a nearly three-fold resolution increase is possible at the same sampling frequency of 10 kHz used in this study, or similar spatial resolutions to the sequence used here can be acquired at over 70 kHz. Larger framing rates are possible still (1 MHz), but at the cost of decreased resolution (128 x 16) and should be used to study smaller scale phenomena with simpler flow structures. Larger resolution data will increase the limiting scale at which flow phenomena can be observed. This increased scale can provide more

accurate segmentation results, and a better approximation to contours, edges, and lines in images. Larger scale will also increase the motion of detected features, which would provide a larger domain over which spectral analysis of feature motion would be appropriate. Physical scale may also be increased by a lens with a larger magnification factor, but at the cost of viewing the entire flowfield.

A problem that is often encountered in all areas of vision (including this research), is the lack of a ground truth that can be used to benchmark results. Multiscale analysis can mitigate this effect by providing confirmation of measurement across scale, and should be used whenever possible. Modern segmentation methods and edge detection schemes are typically benchmarked against the “Berkeley Segmentation Database” [81]. This database was created to provide human marked ground truths on a set of test images, as it is largely human interpretation that vision aims to reproduce. The test images were marked by hand by several hundred participants, and results were combined for each image. Thicker lines represent what more humans marked as a line or region, as interpretation of scene can be subjective. Corner and feature detectors are similarly ground truthed against specific sets of images, where locations of corners are known.

Ground truth data sets should be created for schlieren and shadowgraph analysis. Ideal experiments would be simple two-dimensional geometries with analytical solutions, or computational solutions that can be easily obtained. These could include a wedge pitching during test time, and a circular cylinder during tunnel start. Mach number variation during these tests would also be helpful. High-speed imaging with moderate resolution could benchmark the linear feature, shock extraction, and

feature detection methods used here. Initially, high resolution would be preferable over framing rate to study algorithm performance over a wide range of spatial scales. Synthetic schlieren generated from numerical studies could be useful in benchmarking performance against “noise.” For more complex phenomena, it is better to begin with single images that are hand-marked by multiple researchers, so as to gather a consensus in expert knowledge against which to benchmark.

The effects of optical integration across the test section on image quality can be mitigated when two-dimensional test articles are used, or more appropriately, focusing schlieren techniques can limit the effects of boundary layer sidewalls. Focusing schlieren has been successfully applied at AEDC Tunnel 9 [82], and should benefit from vision-based analysis. High frequency surface instrumentation should also be employed. With knowledge of some of the frequency content, the use of modern signal recovery and parametric spectral estimation methods can be used to further process vision-based results and better interpret their frequency content.

Appendix A: Extensions of Salient Point Identification Methodology

The generality of the modified curvature scale space (CSS) method developed in Chapter 5 for identification of salient points on isolated shock contours is illustrated here. The techniques are applied to a still schlieren image of a “waverider” (a body designed for hypersonic flight that generates lift by capturing the shock wave created underneath it) in order to identify the foremost point on a bow shock. Figure A.1 shows a shock forming on a three-dimensional waverider installed in Arnold Engineering Development Center Hyper Velocity Wind Tunnel 9 with a nominal freestream Mach number of 8. Details of the test are given by Norris [83]. The foremost region of the waverider has been cropped resulting in an 376 x 226 pixel image in Figure A.1.




Figure A.1: Waverider test image.

Figure A.2 shows the shock contour extracted after applying a diffusion-type

bilateral filter [46]. The shock contour was extracted using the methods discussed in Chapter 4. The foremost point on the shock is represented by the minimum in curvature. As the contour is simple, the knowledge base simply comprises the global minimum in contour curvature, given by Equation A.1

$$\mathbf{p}_{tip}(s) = \underset{s}{\operatorname{argmin}} (\kappa(s, \sigma_{\text{MAX}})) : s \in \Gamma. \quad (\text{A.1})$$

Curvature was calculated using four scales of $\sigma = 12, 8, 6$ and 4 due to long contours in the image and the width of the primitive structure.




Figure A.2: Waverider contour image.

Curvatures at the largest and smallest scale are shown in Figure A.3. The minimum in curvature, corresponding to the foremost shock position, is plotted on the grayscale image in Figure A.4. The final solution was updated to the nearest maximum in contour evolution distance large scales as in Chapter 5, due to the large aperture needed to define the nearly circular contour segment. Figure A.4 shows that the foremost point on the bow shock has been identified well, thus demonstrating the general applicability of the detection method.




Figure A.3: Waverider contour curvature $\kappa(\sigma, s)$.




Figure A.4: Bow shock identification with modified CSS method. Fine scale CSS point labeled \triangle , final assignment in \square .

Appendix B: Performance Estimate of Salient Point Detection

In general, a traditional uncertainty analysis (*e.g.* the method of Coleman and Steele [84]) of computer vision-based algorithms designed to perform on data where expert knowledge is needed for interpretation is challenging. There is no objective metric, only judgment by an expert, with which to judge how well schlieren, ultrasound, or magnetic resonance images represent the true scene. Feature detection methods discussed in Chapter 5, excluding our modified curvature scale space (CSS) method, have been benchmarked against known ground truths in the cited references. The issue of detection accuracy here is related to contour distortion; that is, how well the contour represents the underlying structure. While direct comparisons cannot be made without a ground truth, the effect of noise on contour distortions of a known shape, and how this affects the feature detection methodology developed in Chapter 5 can be studied.

A 256 x 256 grayscale image with a square in the lower quadrant was constructed as a test image. The square was chosen since the Canny edge detector has difficulty at 90° corners due to ambiguity of the gradient direction, which may exclude the true corner from the edge map. The image was corrupted with zero-mean Gaussian noise with a variance of 0.025, and multiplicative speckle noise with

a variance of 0.04. Closeups of the original and noise corrupted shapes are shown in Figure B.1, each image is 32 x 32 pixels.

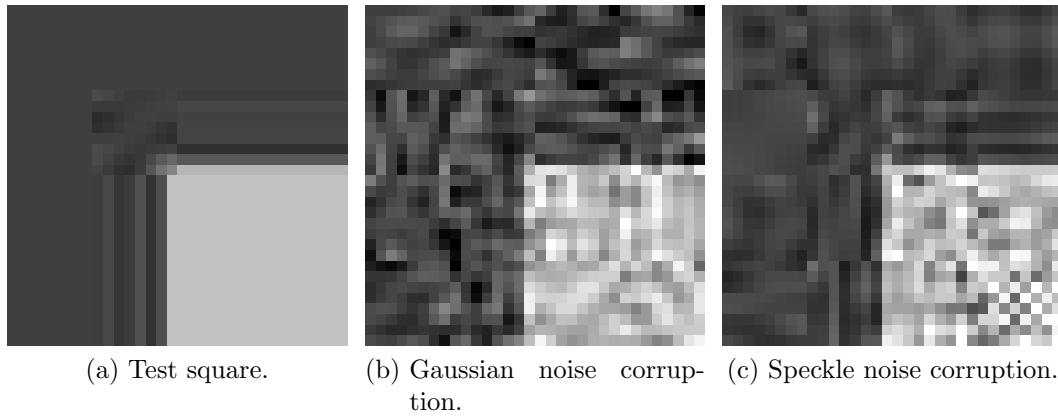


Figure B.1: Test shape and effects of noise corruption. Each image is a 32 x 32 pixel closeup.

Edges and contours were extracted in the same manner used to isolate the outer-shock contour discussed in Chapter 4. The contours extracted from the images in Figure B.1 are shown in Figure B.2, where significant distortions are evident under both Gaussian and speckle noise.

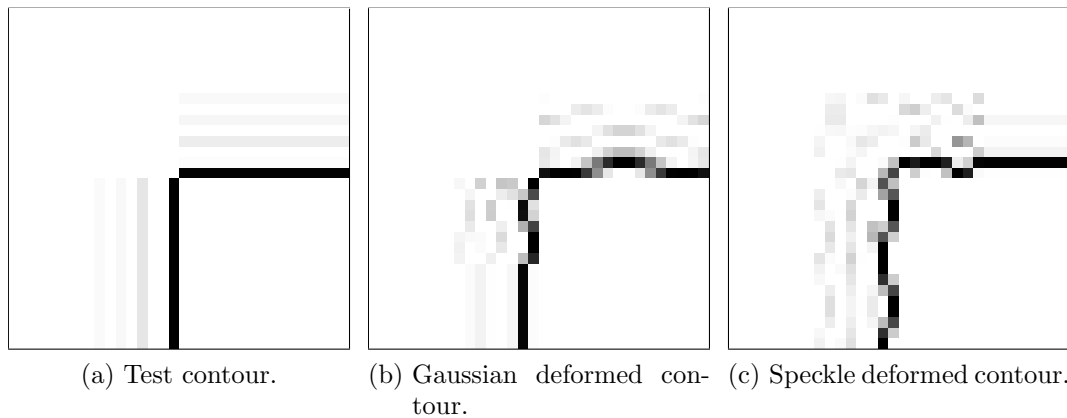


Figure B.2: Test shape and deformed isolated contours. Each image is a 32 x 32 pixel closeup.

Following contour extraction, the corner locations s_{min} in the images shown in Figure B.1 were found using Equation B.1 to search for the global minimum in

curvature that defines the corner in the image

$$\mathbf{p}_{tip}(s) = \underset{s}{\operatorname{argmin}} (\kappa(s, \sigma_{\text{MAX}})) : s \in \Gamma. \quad (\text{B.1})$$

Scales σ of 6, 4, 2 and 1.5 were used to calculate curvature. As the true corner may be missed by the contour, the final location was updated to the closest Harris response \mathbf{R} using Equation B.2 below

$$\mathbf{p}_{tip} = \underset{j}{\operatorname{argmin}} (\|\mathbf{R}(x_j, y_j) - \Gamma_o(s_{min})\|) : (x_j, y_j) \in \mathbf{I}. \quad (\text{B.2})$$

Sub-pixel localization was not performed, as a pixel-level comparison to the shock contour features is desired.

The error was defined as the distance between the known and detected corner locations. The algorithm applied to the distorted contours in Figure B.2 matched the true corner exactly. As noise generation is not repeatable, the test shape was corrupted with Gaussian and speckle noise 500 times each. Outliers were removed from the resulting error distributions, which give an average error of 0.6157 and 0.8483 pixels for the additive and speckle noise, respectively. From this analysis, it can be concluded that the salient point identification method developed in Chapter 5 is detected within an accuracy of less than one pixel at the level of pixel-scale localization. It should be noted that this error estimation is based on the pixel-level assignment of a feature, and may be further reduced in cases where sub-pixel localization is appropriate.

Bibliography

- [1] P. Wu, W. L. Lempert, and R. B. Miles. Megahertz pulse-burst laser and visualization of shock-wave/boundary-layer interaction. *AIAA Journal*, 38(4):672–679, 2000.
- [2] P. Wu and R. B. Miles. Megahertz visualization of compression-corner shock structures. *AIAA Journal*, 39(8):1542–1546, 2001.
- [3] A. Ben-Yakar and R. K. Hanson. Ultra-fast-framing schlieren system for studies of the time evolution of jets in supersonic crossflows. *Experiments in Fluids*, 32(6):652–666, 2002.
- [4] G. S. Settles. *Schlieren and Shadowgraph Techniques: Visualizing Phenomena in Transparent Media*. Springer Verlag, Berlin, Heidelberg, New York, 2001.
- [5] M. J. Hargather and G. S. Settles. Recent developments in schlieren and shadowgraphy. *AIAA Paper 2010-4206*, 2010.
- [6] M. J. Hargather and G. S. Settles. A comparison of three quantitative schlieren techniques. *Optics and Lasers in Engineering*, 50(1):8 – 17, 2012.
- [7] A. L. Kastengren, J. C. Dutton, and G. S. Elliott. A method for measuring recompression shock unsteadiness applied to two supersonic wakes. *Experiments in Fluids*, 39(1):140–151, 2005.
- [8] P. J. K. Bruce and H. Babinsky. Unsteady shock wave dynamics. *Journal of Fluid Mechanics*, 603:463–473, 2008.
- [9] B. Timmerman, A. J. Skeen, P. J. Bryanston-Cross, P. G. Tucker, R. J. Jefferson-Loveday, J. Paduano, and G. R. Guenette Jr. High-speed digital visualization and high-frequency automated shock tracking in supersonic flows. *Journal of Optical Engineering*, 47(10), 2008.
- [10] J. Beloki Perurena, C.O. Asma, R. Theunissen, and O. Chazot. Experimental investigation of liquid jet injection into mach 6 hypersonic crossflow. *Experiments in Fluids*, 46(3):403–417, 2009.

- [11] T. Kouchi, T. Hoshino, K. Sasaya, and G. Masuya. Time-space trajectory of unsteady jet into supersonic crossflow using high-speed framing schlieren images. *AIAA Paper 2009-7316*, 2009.
- [12] A. Rosenfeld and A. C. Kak. *Digital Picture Processing*, volume 2, pages 96–97. Academic Press, San Diego, New York, Berkley, 1982.
- [13] T. Kouchi, K. Sasaya, J. Watanabe, H. Shibayama, and G. Masuya. Penetration characteristics of pulsed injection into supersonic crossflow. *AIAA Paper 2010-6645*, 2010.
- [14] D. Estruch, N. J. Lawson, D. G. MacManus, K. P. Garry, and J. L. Stollery. Measurement of shock wave unsteadiness using a high-speed schlieren system and digital image processing. *Review of Scientific Instruments*, 79(12):1261081–1261083, 2008.
- [15] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.
- [16] D. Estruch, D.G. MacManus, J.L. Stollery, N.J. Lawson, and K.P. Garry. Hypersonic interference heating in the vicinity of surface protuberances. *Experiments in Fluids*, 49(3):683–699, 2010.
- [17] J. Wolfram and J. Martinez Schramm. Pattern recognition in high speed schlieren visualization at the high enthalpy shock tunnel göttingen (heg). In A. Dillmann, G. Heller, M. Klaas, H.-P. Kreplin, W. Nitsche, and W. Schröder, editors, *New Results in Numerical and Experimental Fluid Mechanics VII*, volume 112 of *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, pages 399–406. Springer Berlin / Heidelberg, 2010.
- [18] P.V.C. Hough. Method and means for recognizing complex patterns. U.S. Patent no. 3069654, filed 18 Dec. 1962.
- [19] R.O. Duda and P.E. Hart. Use of the hough transform to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1972.
- [20] B.K.P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1-3):185 – 203, 1981.
- [21] S. Fu and Y. Wu. Detection of velocity distribution of a flow field using sequences of schlieren images. *Journal of Optical Engineering*, 40(8):1661–1666, 2001.
- [22] B. Atcheson, W. Heidrich, and I. Ihrke. An evaluation of optical flow algorithms for background oriented schlieren imaging. *Experiments in Fluids*, 46(8):467–476, 2009.
- [23] D. R. Jonassen, G. S. Settles, and M. D. Tronosky. Schlieren “piv” for turbulent flows. *Optics and Lasers in Engineering*, 44(3-4):190–207, 2006.

- [24] M. J. Hargather, M. J. Lawson, G. S. Settles, and L. M. Weinstein. Seedless velocimetry measurements by schlieren image velocimetry. *AIAA Journal*, 49(3):611–620, 2011.
- [25] S. Laurence and H. Hornung. Image-based force and moment measurement in hypersonic facilities. *Experiments in Fluids*, 46(2):343–353, 2009.
- [26] S. J. Laurence and S. Karl. An improved visualization-based force-measurement technique for short-duration hypersonic facilities. *Experiments in Fluids*, 48(6):949–965, 2010.
- [27] S. Laurence. On tracking the motion of rigid bodies through edge detection and least-squares fitting. *Experiments in Fluids*, 52(2):387–401, 2012.
- [28] S. Palmer. *Vision Science: Photons to Phenomenology*, chapter Organizing Objects and Scenes, pages 254–310. MIT Press, 1999.
- [29] R. Boyle, V. Hlavac, and M. Sonka. *Image Processing, Analysis, and Machine Vision*. Thomson Learning, Toronto, 3rd edition, 2007.
- [30] D.A. Forsyth and J. Ponce. *Computer Vision: A Modern Approach*. Prentice Hall, Upper Saddle River, NJ, 2003.
- [31] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*, chapter Intensity Transformations and Spatial Filtering, pages 104–198. Prentice Hall, Upper Saddle River, NJ, 3rd edition, 2008.
- [32] R. C. Gonzalez and R. E. Woods. *Digital Image Processing, 3rd Edition*. Prentice Hall, Upper Saddle River, NJ, USA, 2008.
- [33] T. Lindeberg. Scale-space. In *Encyclopedia of Computer Science and Engineering*, volume IV, pages 2495–2504. John Wiley and Sons, 2008.
- [34] T. Lindeberg. Scale-space theory: A basic tool for analysing structures at different scales. *Journal of Applied Statistics*, 21(1-2):224–270, 1994.
- [35] P.J. Burt and E.H. Adelson. The laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31(4):532–540, 1983.
- [36] S. V. Kailas and R. Narasimha. The eduction of structures from flow imagery using wavelets part i. the mixing layer. *Experiments in Fluids*, 27(2):167–174, 1999.
- [37] J. Coblisch, M. Smith, T. Hand, G. Candler, and I. Nompelis. Double-cone experiment and numerical analysis at aedc hypervelocity wind tunnel no. 9. *AIAA Paper 2005-902*, 2005.

- [38] J. K. Harvey, M. S. Holden, and T. P. Wadhams. Code validation study of laminar shock boundary-layer and shock/shock interactions in hypersonic flow part b: Comparison with navier-stokes and dsmc solutions. *AIAA Paper 2001-1031*, 2001.
- [39] I. Nompelis and G. V. Candler. Numerical investigation of double-cone flow experiments with high-enthalpy effects. *AIAA Paper 2010-1283*, 2010.
- [40] M. S. Holden, T. P. Wadhams, and M. MacLean. A review of experimental studies with the double-cone and hollow cylinder/flare configurations in the lens hypervelocity tunnels and comparisons with navier-stokes and dsmc computations. *AIAA Paper 2010-1281*, 2010.
- [41] I. Nompelis. *Computational Study of Hypersonic Double-Cone Experiments for Code Validation*. PhD thesis, University of Minnesota, 2004.
- [42] M.-C. Druguet, G. V. Candler, and I. Nompelis. Effect of numerics on navier-stokes computations of hypersonic double-cone flows. *AIAA Journal*, 43(3):616–623, 2005.
- [43] I. Nompelis, G. V. Candler, and M. S. Holden. Effect of vibrational nonequilibrium on hypersonic double-cone experiments. *AIAA Journal*, 41(11):2162–2169, 2003.
- [44] D. Marren and J. Lafferty. *Advanced Hypersonic Test Facilities*, volume 198 of *Progress in Aeronautics and Astronautics*, chapter The AEDC Hypervelocity Wind Tunnel 9, pages 467–477. AIAA, Reston, Virginia, 2002.
- [45] M. Smith and J. Coblish. Measurements to assess the degree of thermal nonequilibrium at aedc hypervelocity wind tunnel no. 9. *AIAA Paper 2004-2399*, 2004.
- [46] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *IEEE Sixth International Conference on Computer Vision.*, pages 839–846, 1998.
- [47] P. Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- [48] J. Weickert. A review of nonlinear diffusion filtering. In B. Haar Romeny, L. Florack, J. Koenderink, and M. Viergever, editors, *Scale-Space Theory in Computer Vision*, volume 1252 of *Lecture Notes in Computer Science*, pages 1–28. Springer Berlin Heidelberg, 1997.
- [49] Z. Wu and R. Leahy. An optimal graph theoretic approach to data clustering: Theory and its application to image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1101–1113, 1993.

- [50] J.B. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- [51] L. Grady and E. L. Schwartz. Isoperimetric graph partitioning for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(3):469–475, 2006.
- [52] L. Grady. *Space-Variant Computer Vision: A Graph-Theoretic Approach*. PhD thesis, Boston University, Boston, MA, 2004.
- [53] L. Grady and E. L. Schwartz. The Graph Analysis Toolbox: Image processing on arbitrary graphs. Technical Report CAS/CNS-TR-03-021, Department of Cognitive and Neural Systems, Boston University, Boston, MA, 2003.
- [54] D. Marr and E.C. Hildreth. Theory of edge detection. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, B-207(1167):187–217, 1980.
- [55] J. Canny. Finding edges and lines in images. Master’s thesis, Massachusetts Institute of Technology, 1983.
- [56] X. C. He and N. H. C. Yung. Corner detector based on global and local curvature properties. *Optical Engineering*, 47(5):057008, 2008.
- [57] X.C. He and N.H.C. Yung. Curvature scale space corner detector with adaptive threshold and dynamic region of support. In *IEEE 17th International Conference on Pattern Recognition*, volume 2, pages 791–794 Vol.2, 2004.
- [58] D.A. Forsyth and J. Ponce. *Computer Vision: A Modern Approach*, chapter Segmentation by Fitting a Model, pages 329–353. Prentice Hall, Upper Saddle River, NJ, 2003.
- [59] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, 2004.
- [60] A. Witkin. Scale-space filtering: A new approach to multi-scale description. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 9, pages 150 – 153, 1984.
- [61] A. L. Yuille and T. A. Poggio. Scaling theorems for zero crossings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(1):15 – 25, 1986.
- [62] J. Babaud, A. P. Witkin, M. Baudin, and R. O. Duda. Uniqueness of the gaussian kernel for scale-space filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(1):26 –33, 1986.
- [63] A. Yilmaz, O. Javed, and M. Shah. Object tracking: A survey. *ACM Comput. Surv.*, 38(4), 2006.

- [64] C. Harris and M. Stephens. A combined corner and edge detector. In *Proceedings of the 4th Alvey Vision Conference*, pages 147–151, 1988.
- [65] J. Shi and C. Tomasi. Good features to track. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 593 – 600, 1994.
- [66] E. Rosten and T. Drummond. Fusing points and lines for high performance tracking. In *IEEE International Conference on Computer Vision*, volume 2, pages 1508–1511, 2005.
- [67] E. Rosten and T. Drummond. Machine learning for high-speed corner detection. In *European Conference on Computer Vision*, volume 1, pages 430–443, 2006.
- [68] H. Asada and M. Brady. The curvature primal sketch. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(1):2 –14, 1986.
- [69] F. Mokhtarian and A.K. Mackworth. Scale based description and recognition of planar curves and two-dimensional shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):34–43, 1986.
- [70] F. Mokhtarian and A.K. Mackworth. A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(8):789 –805, 1992.
- [71] F. Mokhtarian and R. Suomela. Robust image corner detection through curvature scale space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1376 –1381, 1998.
- [72] F. Mokhtarian and M. Bober. Curvature scale space representation: Theory, applications, and mpeg-7 standardization. volume 25 of *Computational Imaging and Vision*, pages 215–242. Kluwer Academic Publishers, 2003.
- [73] X. Zhang, M. Lei, D. Yang, Y. Wang, and L. Ma. Multi-scale curvature product for robust image corner detection in curvature scale space. *Pattern Recogn. Lett.*, 28(5):545–554, 2007.
- [74] X. Zhang, H. Wang, M. Hong, L. Xu, D. Yang, and B. C. Lovell. Robust image corner detection based on scale evolution difference of planar curves. *Pattern Recognition Letters*, 30(4):449 – 455, 2009.
- [75] R. Szeliski. *Computer Vision: Algorithms and Applications*, chapter Segmentation, pages 225–250. Texts in Computer Science. Springer, 2011.
- [76] J. Lafferty and J. Norris. Measurements of fluctuating pitot press, “tunnel noise,” in the aedc hypervelocity wind tunnel no. 9. *AIAA Paper 2007-1678*, 2007.
- [77] C. Chatfield. *The Analysis of Time Series: An Introduction*, pages 56–57. Chapman & Hall/CRC, Boca Raton, London, New York, sixth edition, 2004.

- [78] N. T. Smith, M. J. Lewis, and R. Chellappa. Extraction of oblique structures in noisy schlieren sequences using computer vision techniques. *AIAA Journal*, 50(5):1145–1155, May 2012.
- [79] N. T. Smith, M. J. Lewis, and R. Chellappa. Detection, localization and tracking of shock contour salient points in schlieren sequences. Revisions submitted to *AIAA Journal*, May 2013.
- [80] S. Piponniau, Dussauge J. P., J. F. Dediève, and P. Dupont. A simple model for low-frequency unsteadiness in shock-induced separation. *Journal of Fluid Mechanics*, 629:87–108, 2009.
- [81] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV 2001. Proceedings of the Eighth IEEE International Conference on Computer Vision*, volume 2, pages 416–423, 2001.
- [82] C. P. Vandercreek, M. S. Smith, and K. H. Yu. Focused schlieren and deflectometry at aedc hypervelocity wind tunnel no. 9. *AIAA Paper 2010-4209*, 2010.
- [83] J. D. Norris. Mach 8 high reynolds number static stability capability extension using a hypersonic waverider at aedc tunnel 9. *AIAA Paper 2006-2815*, 2006.
- [84] H. W. Coleman and G. W. Steele. *Experimentation and Uncertainty Analysis for Engineers*. John Wiley & Sons, Hoboken, New Jersey, 2 edition, 1999.