

**Improved Pressure Sensitive Paint Measurement  
Using Natural Feature Tracking and Piecewise Linear Resection**

by

**Jeremy Kuzub, B.Sc. E.E.**

A thesis submitted to the  
Faculty of Graduate Studies and Postdoctoral Affairs  
in partial fulfillment of the requirements for the degree of

**Master of Applied Science in Electrical and Computer Engineering**

Ottawa-Carleton Institute for Electrical and Computer Engineering (OCIECE)  
Department of Systems and Computer Engineering  
Carleton University  
Ottawa, Ontario, Canada, K1S 5B6  
December 2010

Copyright © 2010 Jeremy Kuzub



Library and Archives  
Canada

Published Heritage  
Branch

395 Wellington Street  
Ottawa ON K1A 0N4  
Canada

Bibliothèque et  
Archives Canada

Direction du  
Patrimoine de l'édition

395, rue Wellington  
Ottawa ON K1A 0N4  
Canada

*Your file* *Votre référence*  
ISBN: 978-0-494-79538-5  
*Our file* *Notre référence*  
ISBN: 978-0-494-79538-5

**NOTICE:**

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

---

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

**AVIS:**

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

---

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

  
**Canada**

---

## Abstract

When a physical model is placed in a wind tunnel, aerodynamic forces on the model can deform and displace it, making Pressure Sensitive Paint (PSP) ratio techniques less accurate, since they rely on pixel-by-pixel comparison of a "wind-on" image to a "wind-off" reference image. While some of this mis-registration can be corrected by whole-image re-projection of the displaced model, local bending and twisting cannot be accounted for, resulting in noise in final pressure image. One strategy for noise reduction is to track the distortion and displacement of the model at many points in the image and resect (i.e. reconstruct) the distorted image based on a 2D mesh created by these local motion vectors. This Piecewise Linear Resection (PLR) method becomes more accurate as the mesh density and the number of accurate motion vectors increases, so it is desirable to maximize the number of motion vectors by tracking as many points in the image as is practical. This work presents a method to increase the accuracy of PLR by tracking natural features in addition to standard fiducial markers, significantly increasing the mesh density. The software system presented uses the SURF algorithm to track natural features of a model coated in PSP between wind-off to wind-on images, generating an initial field of local motion vectors. Mismatched vectors are then filtered using a modified disparity gradient filtering technique that integrates weighted fiducial marker motion data to improve accuracy. The resulting motion gradient field is used as a guide to subdivide the wind-on image into a 2D mesh using Delaunay triangulation. Each triangle in this mesh is then independently distorted according to motion vectors at its vertices, completing the image resection process. The ratio between this resected wind-on image and the reference wind-off is used to obtain pressure data. This work shows that this method of automatic PLR is feasible on wind tunnel imagery and that the resulting pressure data has reduced mis-registration noise without the need to perform 3D resection using virtual models and pre-defined deformation equations. It is demonstrated that a denser resection mesh can be generated using a combination of natural features and fiducial markers. It is also shown that the disparity gradient filter using prioritization, subdivision, and vector weighting is an effective way of removing erroneous motion vectors from motion gradient fields. Finally, it is demonstrated that the work flow can be developed using Open Source Software (OSS) libraries, including OpenCV, tifLib, freeGLUT. Overall, this work flow for automatically creating a dense 2D resection mesh and using it to remove model deformation and displacement is demonstrated to be effective in decreasing error in PSP applications where model deformation and displacement are present.

---

# Acknowledgments

It is a pleasure to thank the many people who made this thesis possible.

I would like to extend my gratitude and thanks to my supervisor, Dr. Anthony Whitehead, who provided inspiration and guidance throughout my time as a research assistant. His vision, intuition, and enthusiasm were invaluable.

I would like to extend my sincere thanks and appreciation to Dr. Youssef Mebarki of the Institute of Aeronautical Research for his guidance and support in developing the software solutions in this thesis and for his enthusiasm for students and researchers who aspire to contribute to aeronautics.

I would like to gratefully acknowledge Matthew Brown of the Carleton University Faculty of Computer Science for his outstanding contribution to user interface implementation for this project.

Many thanks and gratitude to the faculty and staff at Carleton University including: Dr. Treavor Pearce for his guidance and confidence, Dr. Gabriel Wainer for his enthusiasm and invaluable assistance, Erenia Oliver and Jennifer Poll for their support in all things logistic.

I cannot overstate my appreciation for the opportunities given me by Dr. David Routledge of the University of Alberta, who inspired me to post graduate research by offering me the opportunity while still in undergraduate studies. Also from the University of Alberta, Dr. Peter Flynn, who encouraged me to pursue entrepreneurial paths throughout my education and professional life. Thank you!

Fellow students and alumni Keith Holman, Matthew Hagman, and Dr. Daniel LaRussa, who provided incalculable support and guidance: my thanks and gratitude.

To Burt Rutan, who has been my inspiration in aeronautics and engineering throughout my life and provided a visionary example of what creativity and engineering can do together, I extend my gratitude.

Finally, I would like to thank my family, who have provided uplifting support, inexhaustible encouragement and inspired example that are an art and science in themselves - I cannot thank you enough, except by pursuing what I love and sharing it with you, even the not-so-simple things!

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Wind Tunnel Testing . . . . .	1
1.2	Pressure and the Registration Problem . . . . .	3
1.3	Compensating for Model Displacement and Deformation . . . . .	5
1.4	Motivation . . . . .	6
1.5	Problem Statement . . . . .	7
1.6	Thesis Structure . . . . .	7
1.7	Contributions to the Field . . . . .	8
<b>2</b>	<b>Background</b>	<b>10</b>
2.1	PSP in Wind Tunnel Testing . . . . .	10
2.1.1	Physical Arrangement . . . . .	11
2.1.2	Imaging System . . . . .	13
2.1.3	Processing PSP data . . . . .	14
2.2	Problem of Model Deformation and Displacement . . . . .	17
2.3	Overview of Registration and Resection . . . . .	19
2.4	Survey of Registration and Resection Methods . . . . .	21
2.4.1	Global Image Transforms . . . . .	21
2.4.2	3D Reprojection Transforms . . . . .	22
2.4.3	Piecewise Linear Resection . . . . .	24
2.5	Introduction to Natural Features . . . . .	30

---

2.5.1	Overview of SURF . . . . .	30
2.5.2	Reducing Feature Mismatches . . . . .	33
2.6	Applying Constraints: Disparity Gradient Filtering . . . . .	35
2.6.1	Background . . . . .	35
2.6.2	Applying Disparity Gradients . . . . .	37
2.7	Fiducial Marker Tracking . . . . .	39
2.8	Conclusion . . . . .	43
<b>3</b>	<b>Solution</b>	<b>44</b>
3.1	Overview of Work Flow . . . . .	45
3.2	Fiducial Marker Detection and Localization . . . . .	47
3.3	SURF Feature Matching Between Wind-On and Wind-Off Images . . . . .	47
3.4	Motion Vector Field Generation . . . . .	48
3.4.1	Fiducial Marker Matching . . . . .	48
3.4.2	SURF Matching . . . . .	48
3.4.3	Disparity Gradient Filtering . . . . .	49
3.4.4	Extending DGF With Vector Weighting and Prioritization . . . . .	50
3.5	Resection . . . . .	53
3.5.1	Creating the Mesh . . . . .	53
3.5.2	Undistortion and Reprojection . . . . .	56
3.6	Generating Ratio Images . . . . .	57
3.7	Interface Design . . . . .	58
<b>4</b>	<b>Tests and Results</b>	<b>62</b>
4.1	Establishing Ground Truth in Photogrammetry . . . . .	62
4.2	Overview of Testing Process . . . . .	63
4.2.1	Software Implementation . . . . .	63
4.2.2	Testing Order . . . . .	64
4.3	Fiducial Marker Detection and Localization . . . . .	66
4.3.1	Synthetic Target Tracking . . . . .	66

---

4.4	Natural Feature Detection and Localization . . . . .	72
4.4.1	SURF Natural Feature Matching with PSP . . . . .	73
4.4.2	Disparity Gradient Filter Testing on Synthetic Imagery . . . . .	82
4.4.3	Testing SURF with DGF on Wind Tunnel Imagery . . . . .	87
4.4.4	Testing of Weighted, Prioritized DGF With Fiducial Marker Data . . . . .	94
4.4.5	Testing the Effect of Image Segmentation on DGF . . . . .	100
4.5	Image Subdivision and Resection . . . . .	104
4.5.1	Subdivision Using Delaunay Triangulation . . . . .	104
4.5.2	Resection . . . . .	106
4.6	Resection Performance Using PLR . . . . .	108
4.6.1	Visualization and Measurement of Misregistration . . . . .	108
4.6.2	PLR Using Natural Feature Tracking Only . . . . .	110
4.6.3	Validating Workflow Using Pressure Tap Data . . . . .	113
4.6.4	Analyzing Spanwise Ratio Image Slice Data . . . . .	116
4.6.5	Comparing Local and Global Resection on Highly Deformed Imagery . . . . .	122
4.6.6	Comparison of SURF-PLR to 3D Global Resection . . . . .	126
4.6.7	Visual Analysis of SURF-PLR Work Flow Results on a Range of Models . . . . .	129
4.7	Interface Design . . . . .	132
4.7.1	Project Level Parameters and Extensibility . . . . .	132
4.7.2	Slice Tool . . . . .	134
4.7.3	Findings From User Testing . . . . .	134
<b>5</b>	<b>Conclusion</b> . . . . .	<b>137</b>
5.1	Summary of Findings . . . . .	137
5.2	Contributions to the Field . . . . .	141
5.2.1	Used Natural Features to Automatically Improve Resection Performance . . . . .	141
5.2.2	Demonstrated a Method for Improving Mesh-Edge Resection Accuracy . . . . .	142

---

5.2.3	Developed a Weighted Extension of the Disparity Gradient Filter Algorithm . . . . .	142
5.2.4	Developed a Modular, Open Source Workflow for 2D Resection and Ratio Image Generation . . . . .	143
5.2.5	Designed Graphical User Interface to Improve Workflow Efficiency . . . . .	143
5.2.6	Findings Applicable to Methods of PSP Coating . . . . .	144
5.3	Criteria for Choosing a Resection Method . . . . .	144
5.4	Future Research Topics . . . . .	146
5.4.1	Confidence Map for Resection Result . . . . .	146
5.4.2	PSP Coating Methods to Improve Natural Feature Tracking . . . . .	147
5.4.3	Arbitrary Disparity Gradient Filter Segmentation . . . . .	148
5.4.4	Additional Matching and Filtering Constraints on Motion Vector Field . . . . .	148
5.4.5	Extension to Multi-frame Tracking . . . . .	149
5.4.6	Integration of Calibration Data from Binary PSP . . . . .	150
5.4.7	Automated Optimization of Work-Flow Parameters . . . . .	150
	<b>References</b>	<b>151</b>

# List of Figures

1.1	Scale model positioned in high speed wind tunnel. . . . .	2
1.2	Cross section of airfoil model with integrated pressure taps . . . . .	3
1.3	Scale model with PSP applied . . . . .	4
2.1	Electromagnetic absorption and re-emission characteristic of a PSP sample . . . . .	11
2.2	Physical arrangement for PSP experiments . . . . .	12
2.3	Camera for experiments . . . . .	13
2.4	Sources of Error in PSP Measurement . . . . .	16
2.5	PSP intensity variations with partial pressure of oxygen . . . . .	18
2.6	Effects of mis-registration due to displacement, rotation and warping. . . . .	19
2.7	PLR subsectioning of a wing . . . . .	25
2.8	Illustration of PLR resection error on a curved surface. . . . .	27
2.9	Relation of PSP area as a fraction of model area to the maximum distance to the nearest trackable landmark. . . . .	29
2.10	Generating a SURF descriptor . . . . .	34
2.11	Calculating Disparity and filtering Disparity Gradient Fields . . . . .	38
2.12	Wind tunnel model with fiducial markers . . . . .	40
2.13	Typical fiducial marker with falloff at edges . . . . .	42
3.1	Software architecture with YML configuration variables and output data . . . . .	46
3.2	Effect of weighting on local disparity gradient filter . . . . .	51

---

3.3	Hierarchy of vector subsets within a field . . . . .	52
3.4	Resection of a polygon within a mesh . . . . .	54
3.5	Delaunay subdivision to include image edges . . . . .	55
3.6	Graphical User Interface design . . . . .	60
4.1	ONERA test target tracking results . . . . .	67
4.2	ONERA target template matching results . . . . .	68
4.3	Gradient background removal before fiducial marker sub-pixel center localization . . . . .	68
4.4	Matching fiducial markers by displacement. . . . .	70
4.5	Hessian Score Distribution of PSP Texture . . . . .	74
4.6	Non-maximal suppression on SURF features . . . . .	75
4.7	Number of filtered outliers as minimum descriptor error in increased	77
4.8	Number of matched SURF features as maximum allowable displacement is increased . . . . .	78
4.9	SURF feature tracking on synthetic PSP texture . . . . .	83
4.10	Motion tracking of the Yosemite sequence . . . . .	85
4.11	Effect of increasing the disparity gradient filter's maximum disparity threshold . . . . .	88
4.12	SURF feature tracking of wind tunnel model with rotation . . . . .	90
4.13	Natural feature tracking of wing bend and motion. . . . .	93
4.14	Local decrease in outlier tolerance near weighted vectors . . . . .	95
4.15	Effect of increasing fiducial marker weight relative to natural features.	96
4.16	Comparing the effect of Dmax to weighting for a given filtering effect.	97
4.17	Effect of segmentation on gradient vector field disparity . . . . .	101
4.18	Delaunay Triangulation to create resection meshes . . . . .	105
4.19	Bicubic resampling introduces artificial haloing artifacts at contrast boundaries . . . . .	106
4.20	Grid reference image to show bilinear resampling blur effect . . . . .	107
4.21	The effect of misregistration on pressure data . . . . .	109
4.22	Resection using natural feature only . . . . .	111

---

4.23 Alignment error after natural feature-based resection . . . . .	112
4.24 Calibration of PSP using pressure tap data . . . . .	114
4.25 Calibration of PSP using pressure tap data . . . . .	115
4.26 Resection effect on ratios cross sections . . . . .	117
4.27 Resection on sting-mounted delta wing . . . . .	119
4.28 Natural feature tracking error on delta wing . . . . .	121
4.29 Comparison of global versus local resection using a non-linearly distorted wind-on image. . . . .	123
4.30 Comparison of global versus local resection using a non-linearly distorted wind-on image. . . . .	124
4.31 Comparison of SURF-PLR to 3D Global Resection . . . . .	127
4.32 Comparison of SURF-PLR to 3D Global Resection at Station 4 . . . .	128
4.33 Sample results of PLR resection . . . . .	130
4.34 PLR user interface . . . . .	133
4.35 Sample Configuration and Output YML Files . . . . .	135

---

# Glossary of Acronyms

CCD - Charge-Coupled Device

DGF - Disparity Gradient Filter

DT - Delaunay Triangulation

HIL - Human-in-the-Loop

FDM - Fiducial Marker

PLR - Piecewise Linear Resection

PSP - Pressure sensitive paint

SURF - Speeded up Robust Features.

TIF - Tagged Image Format

# Chapter 1

## Introduction

### 1.1 Wind Tunnel Testing

Wind tunnel testing is an integral part of aircraft design and development. Although computational fluid dynamic techniques have informed preliminary designs, these are ultimately verified and tuned using wind tunnel data.<sup>1</sup> Measuring forces on a wind tunnel model as a whole can be done through its mounting platform,<sup>2</sup> point pressure can be monitored using static pressure taps on the model skin, and air flow surrounding the model can be visualized by methods such as particle flow velocimetry<sup>3</sup> or airflow density visualization.<sup>4</sup>

Monitoring pressure across an model's surface is fundamental to design, as it indicates areas of flow compression and separation which directly affect real world performance. Pressure data would ideally be continuous across the whole model, and by extension pressure sensor size should be as small a possible to maximize resolution. However surface pressure measurements had been limited to static pressure transducers (taps), as shown in Figure 1.2 on page 3 which provide

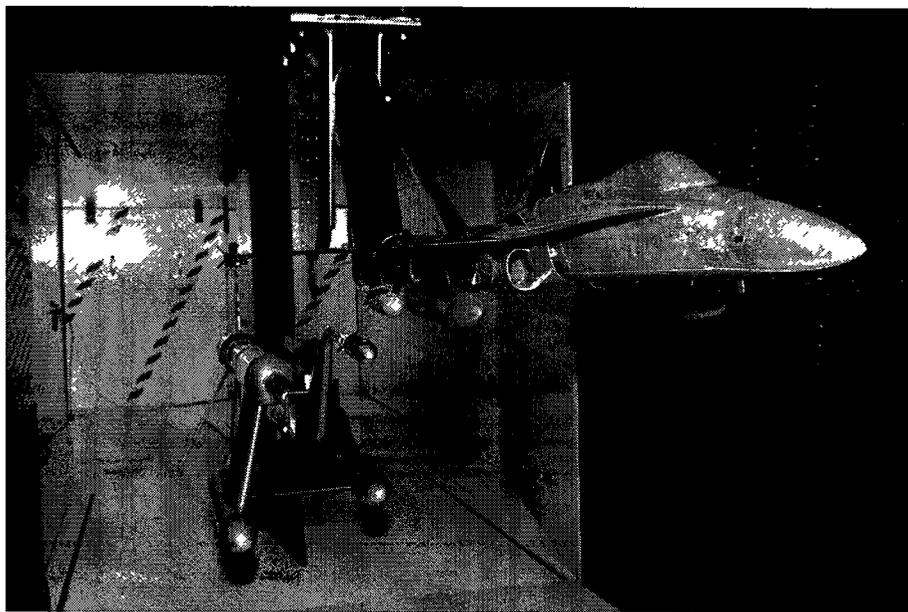


Figure 1.1: Scale model of a CF-18 fighter mounted in the 1.9m Trisonic wind tunnel at Uplands Facility. Wind tunnel testing validates computational fluid dynamic methods, so accurate data collection is vital.

(Courtesy, Institute for Aeronautical Research)

point-data but are expensive to integrate into the wind tunnel model.<sup>5</sup> In addition, these taps only cover a specific set of points across the model, numbering at most in the hundreds, so pressure data across the model surface must be interpolated. Discontinuous pressure features such as transonic shock fronts, which are especially critical to understand in aircraft design, may appear between pressure taps, making the feature's precise position ambiguous. Localized pressure features may be entirely missed, and pressure tap placement on articulated or thin components such as control surfaces can be difficult or impossible. One solution to address this problem of sparse surface pressure data is the use of pressure sensitive paint (PSP) which can provide continuous, high resolution pressure measurement across arbitrary sections of the model. A model coated in PSP is shown in Figure 1.3 on page 4.

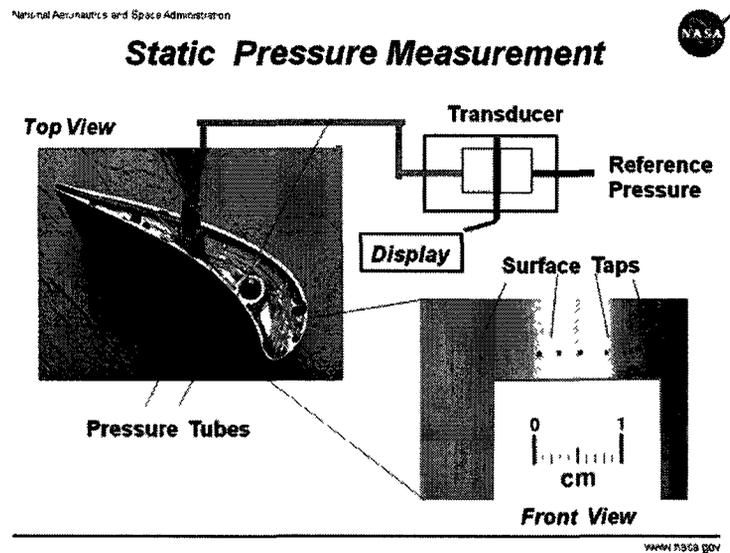


Figure 1.2: Cross section of airfoil model with integrated pressure taps. Each tap must have an individual hole in the model surface connected to a static pressure line running to a transducer.

(<http://www.grc.nasa.gov/WWW/K-12/airplane/tunpsm.html>)

## 1.2 Pressure and the Registration Problem

Recently, PSP has been shown to provide pressure information across the entire model surface with accuracy approaching static pressure taps.<sup>6</sup> PSP contains compounds whose fluorescence can be correlated to incident air pressure in a highly localized and continuous way.<sup>7</sup> PSP has the advantage of being comparatively easy to apply to the model in comparison to static pressure taps and can provide information about air pressure continuously across the entire painted. This greatly reduces cost of testing while increasing data density, however a system of capturing and processing this data is required. Specialized illumination, digital cameras, and image processing are necessary to extract pressure data from PSP coatings. Furthermore, the pressure data can be extracted only as a ratio between a reference or "wind-off" image, and a "wind-on" image, a technique that removes variation

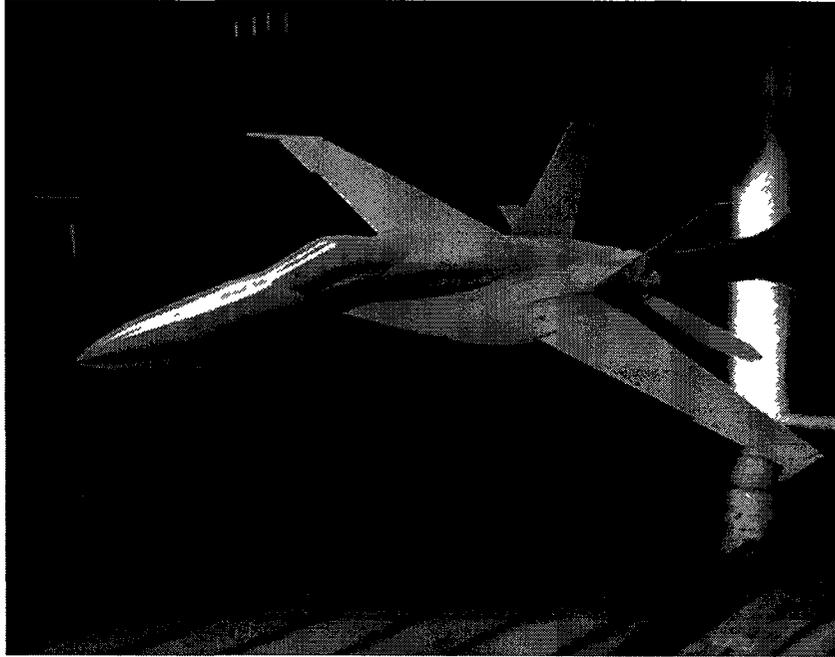


Figure 1.3: Scale model of an F-18 fighter with PSP applied.  
(<http://www.aerospaceweb.org/design/psp/>)

in the absolute brightness of the PSP coating due to factors other than fluorescence. An example of a wind-off and wind-on image pair and the resulting ratio image is shown in Figure 2.5 on page 18.

Because two images separated in time must be used to get a final pressure distribution, it is critical that these images are identical with the exception of the changes in the PSP reflectivity. Any misregistration or misalignment of the images, the camera and illumination system, or the model will introduce errors into the pressure measurement. While camera and illumination systems can be locked down outside of the wind tunnel airstream, the model itself is exposed to these forces and can undergo displacement and deformation. Displacement due to the force of the airflow acting on the model's support structure and distortion due to local bending and deformation of the model, which is not completely rigid. These distortions of the image from wind-off to wind-on directly affect the pres-

sure measurements obtained from the ratio of these two images since even small misalignments can result in significant measurement errors. Despite this technical challenge, the promise of PSP to provide full pressure data across a model makes it desirable to find methods of removing image misalignments post-process in order to obtain the most accurate data possible.

### **1.3 Compensating for Model Displacement and Deformation**

There are several existing methods for obtaining approximate realignment of a wind-on image to its associated reference wind-off image based on tracking the positions of visual marker points (fiducial markers) on the wind tunnel model. These markers are typically of minimal size and number so as not to obscure the PSP coating. The movement of these markers in bulk and relative to one another can provide information about model displacement and deformation, but cannot readily differentiate between the two and across the entire model surface; motion in the areas between the markers must be interpolated.

To address this problem, two approaches can be taken. First, a virtual 3D model of the wind tunnel model can be created, including predefined deformation equations. A system of linear equations for motion and deformation can be solved to minimize error between the virtual model and the wind tunnel model imagery, providing information on displacement and deformation.<sup>8</sup> A second method that does not rely on a virtual model is piecewise linear resection (PLR).<sup>9</sup> In PLR the image is subdivided based on trackable features points such as fiducial markers, and each subdivision within the image is independently realigned using the mo-

tion of these tracking points. This method can operate in 2D and does not require a 3D virtual model, however it can exhibit decreased accuracy in areas of the image that are distant from any tracking points. It becomes desirable to minimize the distance of any given portion of the image from a tracking point in order to have the best fidelity to actual model displacement and deformation.

## 1.4 Motivation

The PLR method can be automated; no prior knowledge of model deformation characteristics is assumed and no virtual 3D model is needed. It can reduce registration errors greatly in the local image areas around tracking points, making it desirable to track as many points across the image as possible.

The motivation behind the work presented in this thesis is to increase the accuracy of PLR by increasing the number of tracking points in the image without adding additional tracking markers that obscure the PSP coating. It follows that some method of tracking pre-existing features within the images of the wind tunnel model is needed. Methods of tracking these natural features in an image have been established for other image processing applications,<sup>10</sup> however they have not previously been applied to the wind tunnel resection problem. It is desirable to determine their suitability as an accurate method of increasing tracking points for the PSP application.

## 1.5 Problem Statement

In this thesis we set out to investigate if natural feature tracking can improve PLR, and by extension the PSP process, in the presence of model displacement and deformation. In order to accomplish this, we determine if a specific natural feature identification algorithm (SURF)<sup>11</sup> is suitable for tracking motion in PSP images. We seek to determine if combining accuracy-weighted motion information from natural feature and fiducial marker tracking improves overall tracking accuracy in capturing local and global model movement. We also attempt to determine if an extended disparity gradient filtering algorithm can remove erroneous motion vectors and improve the fidelity of the overall motion gradient field.

Natural feature tracking has the potential to improve resection accuracy because the number of tracking points in the image is increased, resulting in a smaller distance between any two tracked points and therefore a smaller interpolation error for pixels resampled between these points. The overall resection error should decrease because the number of accurately tracked points has increased, resulting in more accurate pressure data across the PSP image.

## 1.6 Thesis Structure

Chapter 2 describes the characteristics of PSP and its application to wind tunnel testing in addition to a survey of existing methods of removing model displacement and deformation. Characteristics and methods of tracking natural features are described, focusing on the SURF algorithm. An overview of the disparity gradient filtering system is presented in the context of improving the accuracy of the

motion gradient field for a wind tunnel model. Finally, a general overview of resection work flow is given in the context of improving PSP measurement accuracy. The chapter concludes with a statement of the questions investigated in this thesis.

Chapter 3 describes methods and algorithms to be used in the software tool chain. Each functional module is detailed in the order of application in the full work flow. This chapter concludes with a discussion of the advantages of modular software architecture and a graphical user interface for efficient user feedback.

Chapter 4 includes specific validation and performance tests using the software work flow, and discusses specific details of implementation. Fiducial marker and natural feature tracking accuracies are determined, and the performance of the extended disparity gradient filtering algorithm is evaluated. Finally, the resection algorithm functionality and overall system performance is evaluated in the context of PSP measurement error. The results of a full resection work flow are compared to unresected images and to ground-truth data from pressure tap sensors. This Chapter concludes with a discussion of the final graphical user interface architecture and the benefit of human-in-the-loop (HIL) work-flow.

Chapter 5 discusses findings and contributions to the field, and states possible extensions of this work and new problems that can be explored. A discussion of best practices for wind tunnel image capture is also presented in the context of selecting a resection method based on available data.

## **1.7 Contributions to the Field**

- **Use of Natural Features to Automatically Improve PLR Performance**

The work described here demonstrates that it is feasible to track natural fea-

ture using the SURF algorithm for PSP applications and that the resulting increase in mesh density results in a higher fidelity resected wind-on image. This in turn increased accuracy of pressure measurement from the ratio image.

- **Developed a Weighted Extension of the Disparity Gradient Filter Algorithm**

In order to combine the highly accurate data from fiducial marker tracking with the denser data from natural feature tracking, a weighted, prioritized extension of disparity gradient filter was developed as part of this work, allowing more accurate vectors to have more influence on the final motion gradient field, resulting in a higher quality resection mesh.

- **Developed a Modular, Open Source Workflow for Improving PSP Ratio Images using 2D PLR**

This software tool chain for this project was developed using open source libraries and demonstrates that relatively complex software for this application can be developed efficiently without licensing costly analysis tools.

- **Designed Graphical User Interface to Improve Work Flow Efficiency**

An intuitive user interface was developed to manage work flow and focus the user on tuning key parameters using a visual feedback system of intermediate processed images. The user interface also included measurement tools that provide virtual pressure taps and pressure cross sections.

- **Developed Best Practices Guidelines for use of PSP with Natural Feature-Based Resection**

Test findings were used to assist in the choice of a resection procedure, whether it be planar reprojection, 3D model reprojection, or PLR.

# Chapter 2

## Background

### 2.1 PSP in Wind Tunnel Testing

PSPs are partly composed of one or more lumiphor compounds whose reflectivity varies with the partial pressure of oxygen at their immediate boundary.<sup>12</sup> In wind tunnel applications, this oxygen partial pressure varies in proportion with air pressure across the surface of the model. The PSP is applied to the model or a section of the model using simple spray aerosol method to a minimal thickness, since the surface of the PSP is the reaction area to oxygen. PSP coatings are typically made as homogenous as possible, however measurement methods described below remove most variance due to uneven application.<sup>13</sup> A review of PSP formulations and surface treatments can be found in Gregory et al.<sup>6</sup>

The change in PSP reflectivity is primarily in specific wavelengths, and is a result of reemission of energy absorbed from other wavelengths. In practice, a high signal-to-noise ratio is desirable, so the wind tunnel model is illuminated specifically in the excitation band and the CCD camera incident on the model is equipped

with a optical filters transparent to the re-emission band. A sample PSP spectrum for absorption and re-emission can be seen in 2.1 on page 11. It is necessary to arrange the illumination source, model and camera to best use this characteristic of the PSP.

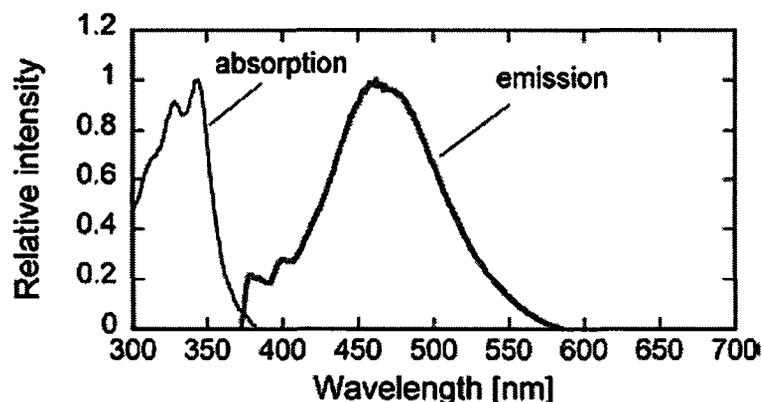


Figure 2.1: Electromagnetic absorption and re-emission characteristic of a PSP sample formulation containing Pyrene-1-butyric acid (PBA). The PSP coating is excited by an illumination source centered at 325nm which causes fluorescence centered at 460nm. Spectra reproduced from Gregory et al. [2008].<sup>6</sup>

### 2.1.1 Physical Arrangement

The experimental setup for PSP measurement typically requires that only the model itself be in the wind tunnel. The camera and illumination source are situated outside the tunnel with access via optical glass ports or windows. This arrangement allows the camera and illumination source to remain static with only the model moving due to aerodynamic forces. A diagram of a typical set-up is shown in 2.2 on page 12. Illumination sources are mounted where practical to provide an even illumination of the model. The camera(s) are placed such that the region of interest on the model is approximately perpendicular to the imaging

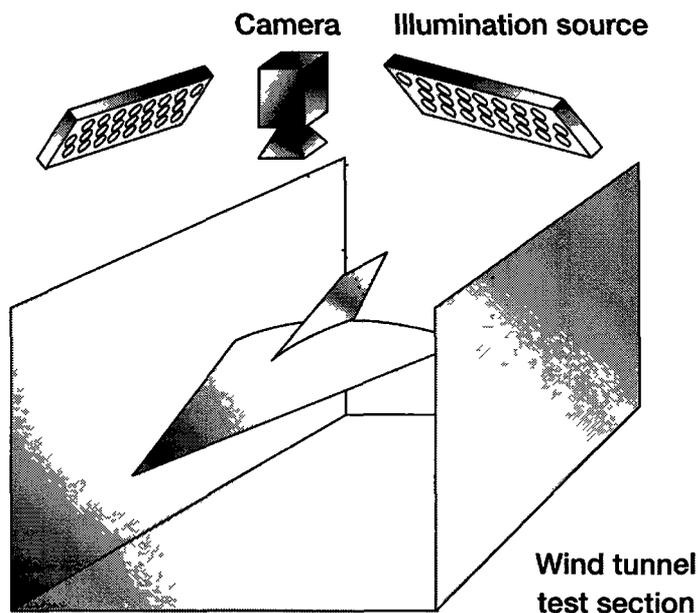


Figure 2.2: Physical arrangement for PSP experiments. Illumination source(s) and camera system are located outside the wind tunnel test section with visibility provided either by optical ports or windows.

plane, and camera lens and focal length selection is based on projecting this area of interest on as much of the image sensor as possible while minimizing depth-of-field blurring. In cases where optical view ports are relatively small, compromises may have to be made in camera and illumination arrangement, resulting in limited field-of view or illumination fall-off, both of which can effect post processing. For example, model displacement and deformation may move fiducial markers out of frame, preventing their use, and uneven illumination can reduce the signal to noise ratio in areas of lower illumination. Data is typically captured on-site; wind-off reference images are captured with no airflow in the tunnel and full illumination while wind-on images are captured when the desired air velocity is reached and steady state airflow is established. The images are then processed off-line.

## 2.1.2 Imaging System

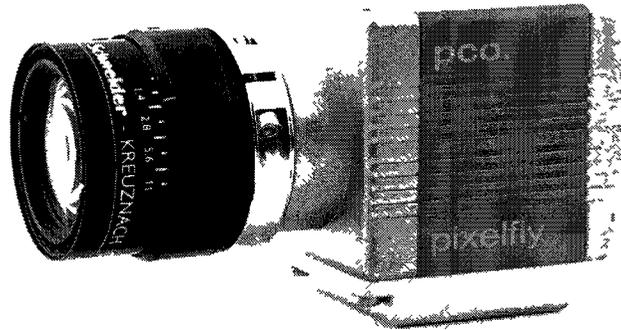


Figure 2.3: CCD camera for PSP experiments. A typical instrumentation-grade camera with a dynamic range of approximately 3000:1. Output is 12 bit linear and noise levels intrinsic to the camera system are approximately 1 count rms. Lens system is replaceable for wider angles of view, and aperture can be set to increase depth of field.

In order to make the workflow in PSP data acquisition efficient and robust, the same cameras are used for model motion tracking and PSP imagery.<sup>13</sup> Typical cameras used in the PSP application are monochrome charged-coupled device (CCD) digital cameras with a standard 2D grid pattern sensor. Each pixel on the CCD is digitized to 12 or more bits of linear intensity and stored as a standard 16-bit per channel, single channel Tagged Image Format (TIF). The camera system can introduce artefacts into the image, such as:

1. Non-linear optical distortion from the lens and CCD system
2. Blur in near and distant parts of the model due to finite depth-of-field
3. Per-pixel random variation due to thermal and shot noise in the CCD

Item (1) is invariant in this application and since the camera image is not being matched to a 3D model mesh or other virtual data, does not have to be removed. Item (2) directly effects per-pixel data since higher frequency data components are effectively lost in areas that are out of the focal plane. This is mitigated by reducing the aperture of the camera lens to increase depth-of-field. This reduces image brightness on the sensor, requiring an increase in exposure time. If the model and air flow patterns are static, this is not an issue. Item (3) is an electronic property of the CCD and introduces random per-pixel variation in the output images. This can be partially mitigated by subtracting a dark-frame image from any output image of the camera. This dark-frame image is the output of the camera with no light incident on the sensor for an exposure time equal to that used in model imaging. Increasing model illumination in turn increases the signal-to-noise ratio, reducing the effect of this noise source.

The illumination system can be LED, laser, or UV-lamp based, since the excitation band of PSP is typically in the 300nm range. Illumination must be constant between wind-on and wind-off images as any change will be indistinguishable from a change in global air pressure. The illumination is ideally uniform across the model to maximize the dynamic range of the pressure signal that can be recovered, however this may not be the case in high-speed wind tunnels where smaller optical port make point sources necessary.

### **2.1.3 Processing PSP data**

Once the wind-off/wind-on image pairs have been captured they require pre-processing and merging. The brightness (intensity) of the PSP coating is recorded by CCD cameras as gray scale values. The brightness values vary locally with air

---

pressure changes, however the absolute intensity values for a given pixel location are not representative of the absolute pressure values for several reasons. For example, a non uniform illumination field is a practical reality because of the point sources used and the 3D geometry of the wind tunnel model. Additionally, the PSP coating may be slightly heterogeneous when applied to the model, which can result in local differences baseline reflectivity.

To separate these factors from the pressure data, the ratio of the wind-off image intensity ( $I_0$ ) with no airflow is made with the wind-on image with airflow ( $I$ ). By taking a pixel-by-pixel ratio, local variations in illumination field and PSP coating are removed from the final image, and variation in the ratio-image intensity can be related to pressure change, as described by in<sup>14</sup> using the Stern-Volmer equation, where  $p$  and  $p_0$  are the wind-on and wind-off pressures, respectively, and  $A$  and  $B$  are coefficients specific to the PSP chemistry and the linear contrast characteristics of the imaging pair.

$$\frac{I_0}{I} = A + B \frac{p}{p_0}$$

This approach is used to make PSP measurements a feasible tool; however there are practical considerations that can influence this result, such as self illumination and PSP response to temperature. Sources of PSP measurement error are discussed and modeled by Liu et. al.,<sup>7</sup> showing that a critical variable that affects measurement accuracy is model displacement and distortion which modifies the derived pressure ratio, introducing error where there is displacement or deformation.

A summary of error sources is shown in Figure 2.4 on page 2.4. Image mis-registration between wind-off and wind-on images plays a role in either increasing

## Error Sources in PSP Measurement

Error Source	Also Affected by	Method of Compensation
<b>Ratio image intensity calibration to pressure taps</b>	<b>All error and uncertainty sources</b>	<b>Removing all variables that affect PSP intensity other than pressure</b>
Illumination variance over time	PSP degradation over time	Consistent illumination source
Illumination variance over camera field of view	Image misregistration	Ratio image method, reference illumination image
Responsivity of PSP coating over model	Image misregistration	Ratio image method
Temperature variance over model	Image misregistration	Binary PSP additional reference lumphor with temperature dependent flourescence
Motion blur due to model vibration and oscillation	Image misregistration, camera position, aperture and focal-length	Reduction of camera exposure time
Camera system shallow Depth-of-field	Camera position, aperture and focal-length	Reduction in camera aperture and corresponding increase in exposure time
Misregistration between wind-on and wind-off images	Model deformation, displacement	Image tracking and resection

Figure 2.4: Sources of Error in PSP Measurement. The Ratio method compensates for errors due to illumination and PSP responsivity, and Binary PSP formulations can compensate for temperature variations, however these corrective methods are less effective in the presence of misregistration; it is therefore desirable to remove image registration error to increase the effectiveness of other corrective methods.

---

error and uncertainty or in reducing the effectiveness of correction strategies; it is therefore of primary interest to remove this source of measurement error. All error sources ultimately impact the calibration of ratio image intensity to pressure values. While Liu's treatment of error sources is analytical, practical approaches to reducing the effect of these error sources are described by Crafton et al.<sup>13</sup>

## **2.2 Problem of Model Deformation and Displacement**

The aerodynamic loads on a wind tunnel model can be significant enough to displace and deform the model. This causes local and global misregistration of the wind-off and wind-on images, reducing the accuracy of measured data from PSP. LeSant et al. show that model deformation and displacement is a key source of pressure measurement error.<sup>15</sup>

Model displacement is primarily due to the movement of the mounting armature or "sting" that holds the model stationary within airflow. In the case of wall-mounted half-models, there can be slight changes in rotation about pivot points used to change the model's angle relative to air flow. These changes are global, affecting the model as a rigid body. In the case of a 3D object such as a wind-tunnel model, rigid motion typically appears as nonlinear motion when projected on a 2D image sensor.

In contrast, model deformation is local rather than global, since a model is not a fully rigid body under load; for example, bending and twisting of wings is nonlinear. This introduces local changes in model perspective and geometry from the point of view of the camera system which are typically nonlinear.

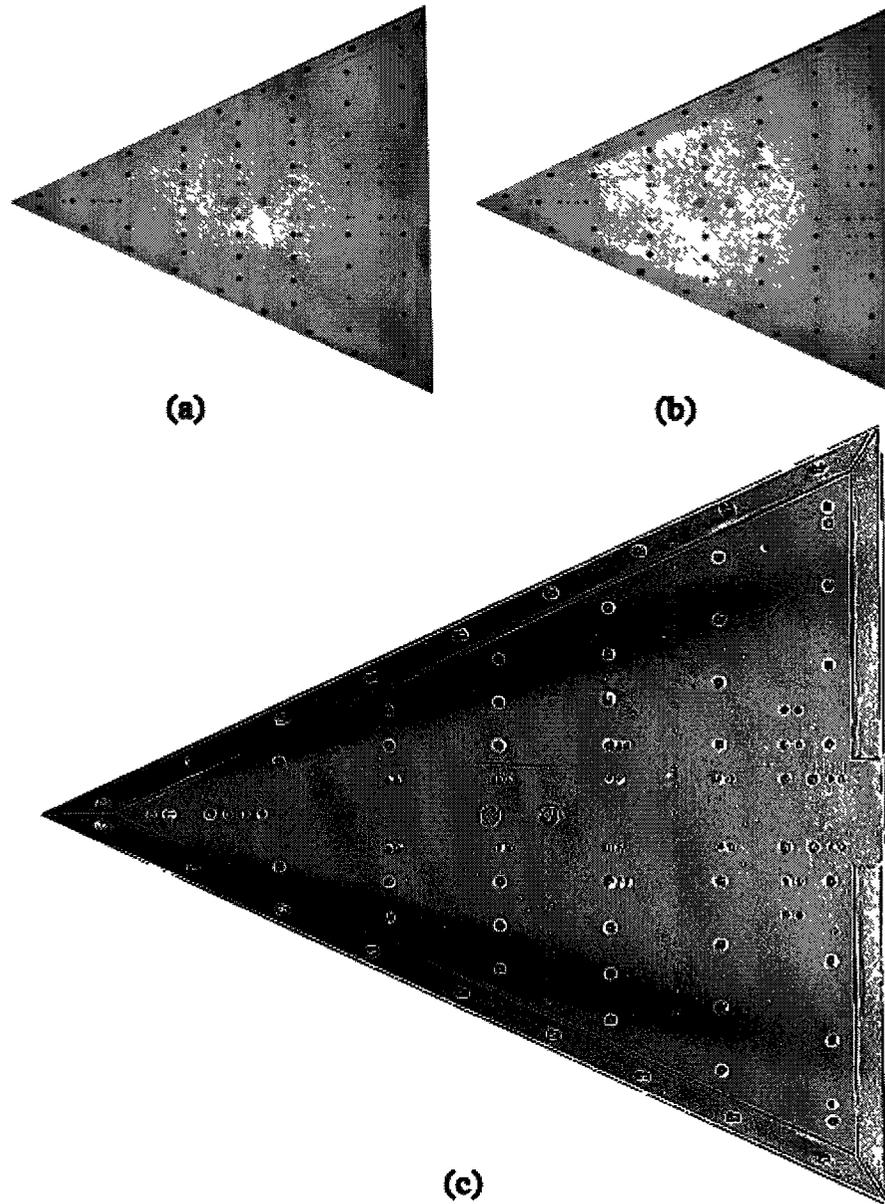


Figure 2.5: Deriving pressure change from PSP imagery. The change of pressure from static to steady airflow is of interest, and is represented by the intensity ratio, taken per-pixel across the image. Images (a) and (b) are the wind-off and wind-on images, respectively, which are combined to create a ratio image shown in (c). The image intensity levels have a linear relationship to pressure changes across the model.

These two types of motion have the net effect of reducing the pixel-pixel alignment accuracy of the wind-off and wind-on images, which in turn reduces the accuracy of the ratio images generated from these image pairs. A misregistration of even a fraction of a pixel is visually apparent and increasing misalignment reduces the accuracy further as shown in Figure 2.6 on page 19.



Figure 2.6: Effects of misregistration due to displacement, rotation and warping. In 2.6(a) misregistration is clearly seen at fiducial markers and as noise in the PSP coating. In 2.6(b) the wind-off and wind-on images are registered to less than 1 pixel, reducing noise in the PSP data.

## 2.3 Overview of Registration and Resection

Different techniques have been developed to detect and correct model displacement deformation. These typically rely on a process of image registration followed by image resection. Registration is the process of determining the relative positions of the model and camera for both the wind-off and wind-on images, and determining any difference between the two. Resection is the process of modifying the wind-on image to reduce or remove this difference and thereby reducing misreg-

istration artefacts in the final ratio image generated from the wind-off/wind-on pair.

The registration process uses known landmarks in the images to determine model movement and deformation. A typical landmark is a fiducial marker. These markers are visible to the cameras monitoring the model during testing and are designed to balance ease of application to the model with accuracy of tracking.<sup>16</sup> Changes in the position of these landmarks between wind-off and wind-on images allow partial or complete information about model position and deformation changes to be solved for.<sup>17</sup>

In this application, the resection process involves distorting the wind-on image to remove as much geometric distortion as possible when compared to the paired wind-off image; in effect removing model deformation due to aerodynamic loads. This deformation process aims to minimize the displacement of landmarks such as fiducial markers between the wind-on and wind-off images. Removing as much distortion as possible improves the accuracy of the pressure values determined in subsequent ratios of the two images.

The distortion information recovered from landmarks such as fiducial markers is sparse, so the resection process requires assumptions about how the wind tunnel model deformed in order to interpolate the intensity values at positions located away from the markers. Various methods have been developed to make this interpolation as accurate as possible in the presence of model displacement and deformation.

## 2.4 Survey of Registration and Resection Methods

Several methods have been developed to resect the wind-on image based on the motion of a sparse set of landmarks such as fiducial markers. The goal of all these methods is to reduce the displacement error of known landmarks in comparison to their positions in the wind-off image. This is done by distorting the wind-on image according to specific constraints. Throughout this section, "registration error", or "error" refers to the displacement between a given landmark in the wind-off and wind-on images.

### 2.4.1 Global Image Transforms

This class of resection methods attempts to reduce error by applying an affine or perspective transform to the entire wind-on image. If the wind tunnel model was a flat plane tilted arbitrarily to the camera, this method can reduce error significantly. In addition, only two to four landmarks need to be tracked by a single camera. However, most models are 3D objects projected onto the 2D image plane, and residual error is present in any part of the image which is not on the same plane as the tracked landmarks. This does not correct local or nonlinear deformation of the model and any camera system distortion such as barrel distortion can reduce accuracy and should be removed before resection.

This method has been extended to include other deformations of the image plane, including polynomial-based nonlinear bending using curve-fitting. These methods can provide improvement where nonlinear deformation of the model is known and is the dominant source of error, however the number of landmarks tracked may have to be increased to provide data for the curve-fitting solver. Sam-

---

ple testing and errors for this 2D distortion methods, including affine, and perspective are discussed in by Bell.<sup>12</sup> Specifically, the required number of tracked landmarks is shown to increase as the distortion model becomes more complex (i.e. higher polynomial order).

## 2.4.2 3D Reprojection Transforms

Since most wind tunnel models are 3-dimensional, reprojection methods have been developed to improve on global image transformations. These methods attempt to reproduce the camera-model system in 3D using the 2D image information and require that a virtual 3D model of the actual wind-tunnel model be available. This model is typically represented as a polygon mesh, and a key requirement for accurate reprojection is that landmarks on the actual wind tunnel model are accurately located on the virtual model. This method required several more landmarks to be tracked, although they do not have to be on the same plane. Solving for model pose relative to the camera typically requires 7 landmarks to be tracked.<sup>8</sup> The camera and model system pose for the wind-off and wind-on images are solved, and the 2D wind-on image is back projected through the camera onto the model geometry, forming a texture map that is linked to the model geometry in 3D space. The model is then transformed to its wind-off pose and a virtual 2D wind-on image is generated, which ideally will have all model displacement and rotation relative to the wind-off image removed. LeSant [2004] discussed this method with examples of registration and resection using 3D virtual models, showing that registration error can be reduced if the system is well calibrated.<sup>15</sup>

The advantage of this system is the potential to remove all error due to model displacement and rotation. There is flexibility in landmark placement since

---

these are not limited to a 2D plane. In terms of visualization, the final image ratio can also be projected onto the 3D model to aide in analysis.

A disadvantage of this method is that camera system distortion must be determined and removed for an accurate reprojection. Camera internal and external parameters can be determined simultaneously with enough landmarks on the wind tunnel and virtual model, or a calibration of the camera system can be made before testing by imaging a known target.

Another limitation is the accuracy of landmark registration between the real and virtual wind tunnel models. Because the solver attempts to minimize the difference between these two point sets, any error in specifying landmark location on the either model will limit the accuracy of the resection process.

Most critically, this system assumes the wind tunnel model is non-deformable. The solver can reduce error to zero only if markers are accurately placed, the camera parameters are known, and markers have not moved relative to one another in 3D space between wind-off and wind-on images. In the case model deformation, the markers do move relative to one another, and global rotation and displacement transforms of the virtual model will not remove all error.

In order to account for model deformation, systems have been developed which add additional parameters that specify model deformation. For example, wing bend along a single axis can be modeled as a second-order polynomial equation describing landmark motion at the wing tip versus the wing root. The additional terms in the polynomial must be solved for, and are added to the RANSAC solver system of equations.<sup>8</sup> These additional terms require additional data points either from another camera, or in the form of more landmarks on the model.

If deformation is accurately modeled, this can be a viable solution, however

the number of additional equations to be solved increases with the number of deformations described. These deformation models must be determined ahead of time, and resection accuracy is limited to the fidelity of these deformation models.

This system is effective, however requires significant preparation of a virtual model, landmark placement and measurement, camera calibration and preparation of deformation equations which balance fidelity with number of unknowns introduced into the solver.

### **2.4.3 Piecewise Linear Resection**

Using the 3D resection method, the problem of reducing localized, nonlinear distortion of the wind tunnel model requires significant preparatory work. As stated by LeSant et. al in,<sup>8</sup> the simplest solution is the best in terms of experimental setup. A method which uses local solutions to model displacement and deformation is piecewise linear resection (PLR), which involves segmenting the image into subsections and resecting each of these individually based on local landmarks. PLR has a relationship to piecewise linear approximations to continuous functions, where line segments are fit to known points of a function that will not easily be fit to a polynomial. In practice, the wind-on image is segmented 2D planes segments which are distorted to match the wind-off to wind-on displacement of local landmarks.

This system operates in 2D and has the advantage of becoming no more complex to solve as model deformation becomes non-linear. It requires no knowledge of the camera system parameters, and no virtual 3D model is required for reprojection. Landmarks on the wind tunnel model do not have to have their positions determined in absolute terms. All of the information used for resection is

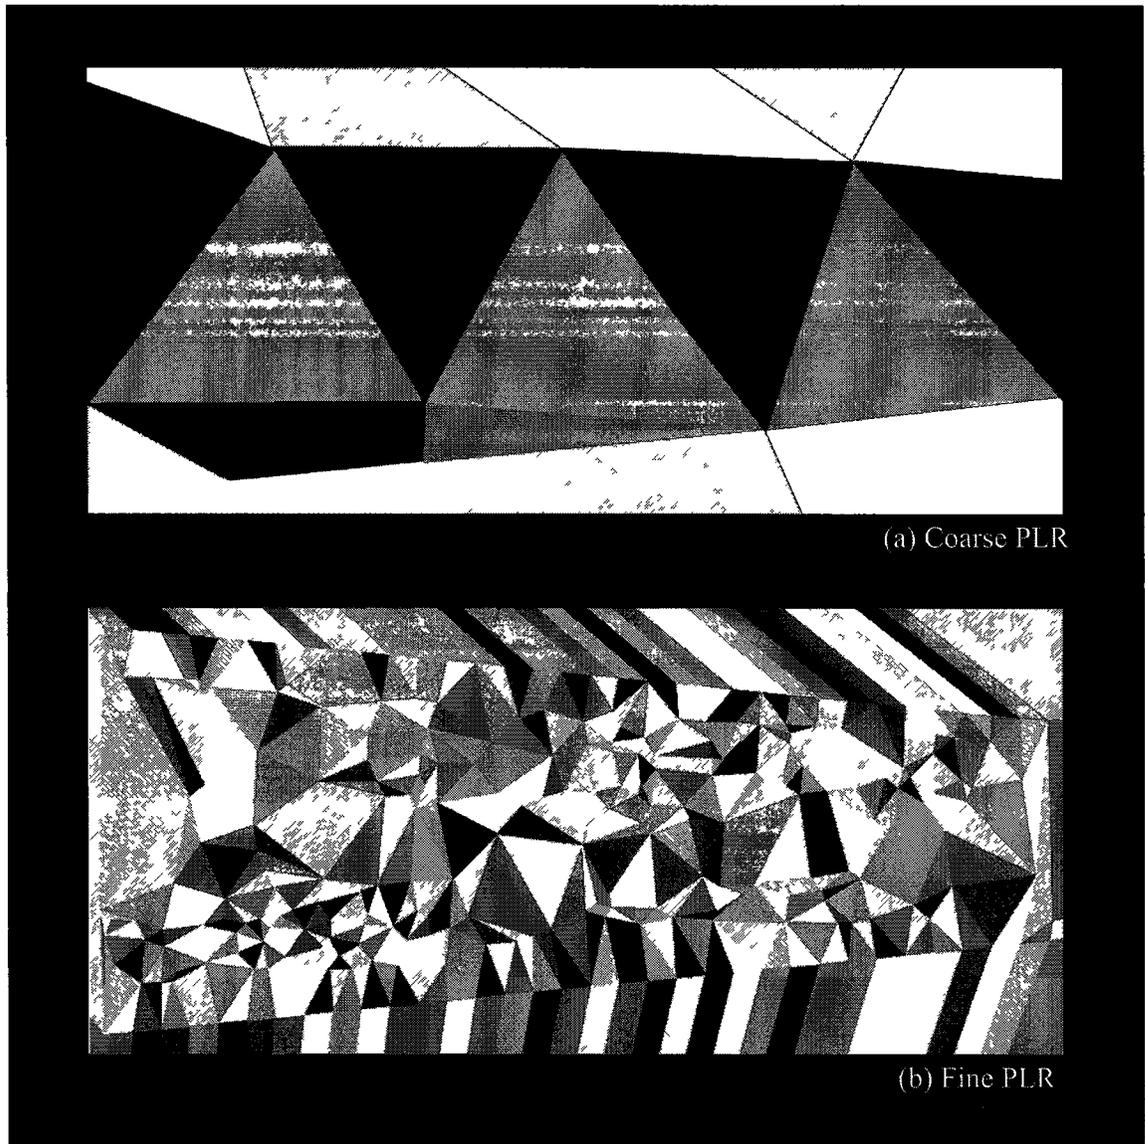


Figure 2.7: PLR applied to an wing image. Each shaded polygon is resected independently based on motion vectors at its vertices. Pixels on the interior of the polygons are linearly interpolated. (b) uses more landmarks and therefore generates more polygons than (a), resulting in a smaller amount of interpolation throughout the image when each polygon is resected.

---

derived from the displacement of landmarks in 2D from the wind-off image to the wind-on image. This has the potential of reducing preparation and calibration time while removing the requirement that model deformation modes be estimated.

### **Delaunay Triangulation**

PLR relies on subdivision to produce a mesh of the wind-on image before resection. This mesh must be produced using an automated method because of the high density of the segmentation. This is done using Delaunay triangulation, which can create a triangular mesh using set of landmark locations as vertices. The algorithm works to create a mesh where each mesh has triangles has similar interior angles. This method will produce the same mesh given the landmark locations, assuring repeatability.<sup>18</sup> Delaunay triangulation has been shown an effective way of subdividing an image for PLR. Cazals et. al. discuss algorithms for creating a mesh from a set of points<sup>19</sup> demonstrating that Delaunay triangulation is an efficient method for subdividing a surface. Devilliers shows that points can be added sequentially to a mesh, and that recomputing a mesh as points are added is efficient for the densities that would appear in the PSP application.<sup>20</sup>

### **Improving PLR Performance**

The flexibility of PLR is offset by the fact that motion between landmarks is, as in other methods, subject to interpolation. Methods described previously interpolate image data based on either a 2D model of image distortion or a 3D virtual model mesh. These use global constraints on model deformation and motion; however PLR has no such constraints and can reduce the displacement error of any landmark by resecting the local image area. However the confidence in this resection

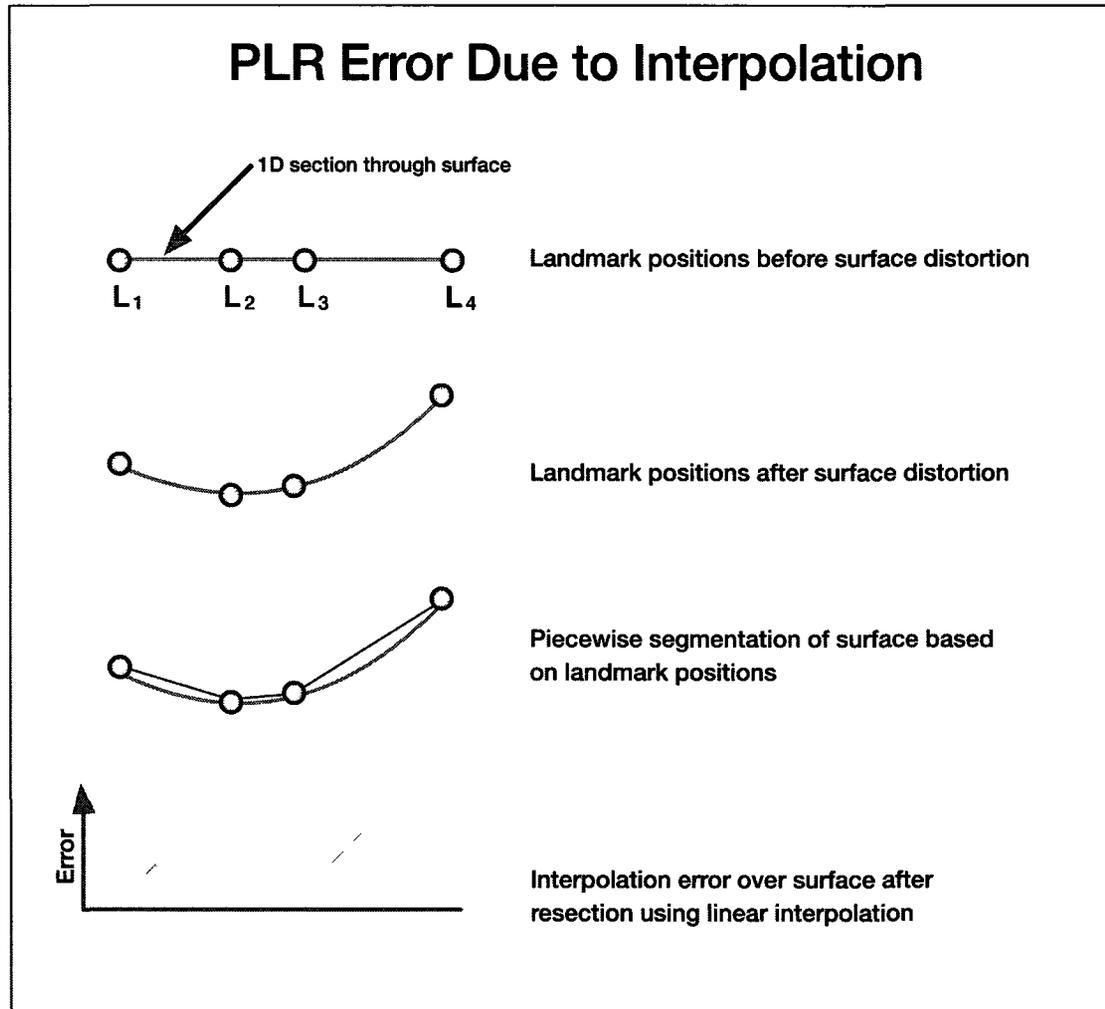


Figure 2.8: Illustration of PLR resection error on a curved surface. Four landmarks are used to track the distortion of a surface, shown here as a 1-D cross section. The surface undergoes a continuous nonlinear distortion, and linear interpolation between landmarks diverges from actual curvature. Larger distances between landmarks are more prone to this interpolation error.

---

accuracy is reduced as distance from landmarks increases. Performance of PLR in sparse landmark testing is discussed by Bell and McLachlan, where scenarios used between 4 and 28 landmarks.<sup>12</sup> They noted the high accuracy of this method near landmarks and its lower accuracy away from them. Venkatakrishnan [2004], in a comparative study of 2D image resection techniques, found that PLR was the most accurate method near markers, but less accurate away from these landmarks.<sup>21</sup> The error in resection of a continuously distorted surface due to linear interpolation is illustrated in 2.8 on page 27 as a one-dimensional simplification. Interpolation error is proportional to the distortion of the surface and the distance between landmarks. The goal then is to reduce interpolation error by reducing the distance between landmarks in the image.

To understand the effect of increasing landmark count on interpolation distance in a two-dimensional image, a simulation was created. Figure 2.9 on page 29 shows that increased landmark count decreases the required interpolation distance, with the largest gains in accuracy in the first 100 new landmarks. This indicates that increasing the number of landmarks beyond the minimal number available in standard wind tunnel testing arrangements will provide significant initial decreases in the interpolation distance and a corresponding increase in PLR resection accuracy.

Adding more landmarks typically requires more fiducial markers. These fiducial markers must be of finite radius and therefore obscure a portion of the PSP coating. The fiducial marker itself can interfere with airflow if it is not match to the surface height and texture of the surrounding PSP. A practical balance must be struck between resection accuracy and density (coverage) of trackable features.

Ideally, each pixel could be identified and mapped between wind-off

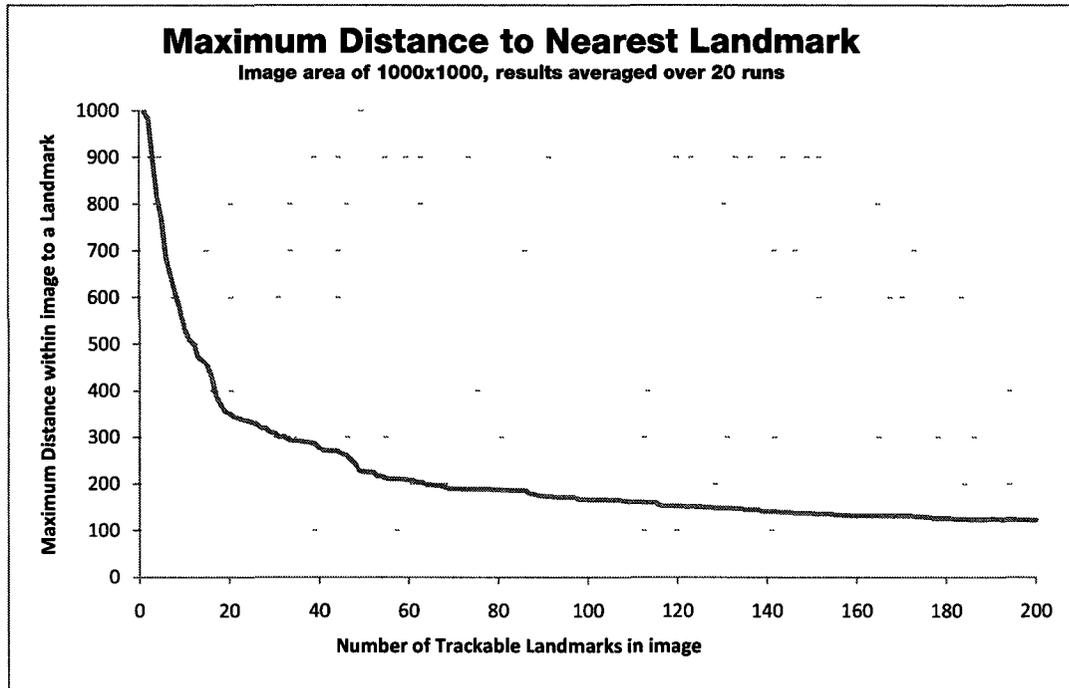


Figure 2.9: Maximum distance to the nearest trackable landmark. As the number of trackable landmarks in an image increases, distance of any given pixel to the nearest landmark decreases, reducing error due to interpolation. Largest decreases in interpolation distance are seen in the initial increase in landmarks count, with an asymptotic approach to zero distance.

and wind-on image pairs, allowing a perfectly accurate resection of the image. uniquely identifying pixel motion is difficult because the intensity of each pixel containing PSP data can change between wind-on and wind-off due to pressure changes as well as model deformation - the two cannot be decoupled at a pixel level.

## 2.5 Introduction to Natural Features

A compromise must be made between the fraction of wind tunnel model image area dedicated to fiducial markers and area dedicated to the PSP. This compromise exists as long as the fiducial marker has a finite area. The more accurate the desired tracking, the more fiducial markers must be tracked across the model, reducing the area available to PSP.

Natural Feature tracking can improve this situation by using image features that are already present in the image and not explicitly placed for tracking; this can include non-uniformities in the PSP itself. Natural feature identification includes a class of algorithms that identify and describe areas of intensity patterns in an image such that these areas can be identified in a subsequent image frame, even if that subsequent frame includes changes in rotation, scale, skew, brightness, and contrast, all of which can occur to some degree between the wind-on and wind-off images in the PSP application.

Natural feature tracking depends on identifying and tracking subsections of an image, referred to generally as "features". A feature is typically composed of a descriptor and a set of parameters with the aim of making the descriptor minimally variant across a wide range of image geometry and intensity transforms while maintaining discrimination from other features. Natural feature tracking can be divided into three stages: identification, description, and matching.

### 2.5.1 Overview of SURF

"Speeded Up Robust Features" (SURF) is an algorithm for natural feature identification and description within image sequences. SURF is a multistage process:

an image is first scanned for suitable features, and responses to a series of convolution filters are then used to generate a descriptor for each feature. To find matching features in sequential images, an element-by-element comparison can be made of candidate features' descriptors. These matches are tolerant to scale, rotation, brightness, and contrast changes from image to image, and can be localized to sub-pixel accuracy.<sup>22</sup> Developed by Bay et. al., SURF is a refinement of other natural feature identification algorithms and optimized for image processing speed while maintaining robust identification characteristics.<sup>11</sup>

This robustness makes SURF suitable for the PSP application, where a feature can undergo both geometric distortion and changes in intensity due to the PSP coating's reaction to local air pressure. SURF was chosen as the natural feature identification and characterization algorithm primarily because of its invariance to brightness and contrast changes in the image. Juan et. al. showed in testing that SURF has better tolerance to illumination variance than related methods.<sup>23</sup> Additionally, it is not covered by non-open source licensing terms like some other methods. The speed of SURF is also a desirable characteristic, since the PLR method has improved accuracy with more trackable features.

### **Identifying Features to Track**

SURF uses the determinate of the Hessian to determine localized areas that are suitable for tracking. The Hessian is a 2nd order square matrix in which each element is the response of the subregion to a specific convolution filter. This method of feature identification has been shown to be as robust under affine transforms of the image as other feature identification systems.<sup>10</sup> These filters are designed such that their maximum response is to image areas which are localizable and

rotation-variant. SURF has optimized these filters for speed with little or no loss in performance. This is used as a measure of an image subregion's suitability for tracking.<sup>11</sup> This detected is carried out across multiple image scales by varying the size of convolution kernels used to determine the Hessian, providing a final location of the feature in  $x$ ,  $y$  and scale axis. The scale space value is used in descriptor generation to decouple the descriptor from the feature's pixel dimensions, making it scale-invariant.

### **Generating a Feature Descriptor**

A descriptor is generated using a square subregion around an identified feature. A reproducible orientation of the descriptor is first calculated, and all descriptor array elements are generated with respect to the orientation, increasing rotational tolerance. The descriptor itself is a sequence of 64 real numbers. This is generated by subdividing a square region surrounding the feature center into 16 square subregions, as shown in Figure 2.10 on page 34. Each of these subregions' responses to four Haar wavelet kernels is then added to the descriptor.

### **Matching Features Using the Descriptor**

Features can be matched from image to image by comparing their descriptors. The descriptor remains similar for a feature under changes in rotation, scale, brightness and contrast, so image features can be compared simply as the sum of squared errors (SSE) between each element of their descriptor sequences. To match features between two images, their descriptors are matched as pairs with minimum SSE. A feature pair is only matched once, then removed from the matching pool. Any unmatched descriptors constitute unmatched features, which are discarded. Be-

---

cause feature similarity is proportional to the Euclidean distance of two features' descriptor arrays, matching is a straightforward operation, equivalent to finding closest point-pairs in 64-dimensional state space.

## 2.5.2 Reducing Feature Mismatches

Highly localizable, rotation-variant image features yield the best descriptors for unique matching. In PSP applications the quality of the descriptor, as measured by its Hessian determinant, tends to be less than other in imaging applications. This is attributable to the high self-similarity of the PSP coating over the model, and the long, straight edges of the model itself, characteristic of aerodynamic bodies. Features such as pressure taps and circular fiducial markers also are not ideally suited as they have radial symmetry, resulting again in a lower Hessian determinant. Prior work in the field of medical imaging processing has shown that these types of low-contrast, self-similar images can be challenging to register and match by natural features alone without further constraints.<sup>24</sup>

If sufficient natural features are to be extracted from a PSP wind tunnel model image, it is necessary to accept a lower quality of feature as measured by the Hessian determinant  $h$ . Typical values may be as low as  $h = 10$ , whereas standard test images can get sufficient feature counts with a lower Hessian threshold of  $h = 100$  or greater. This lower quality implies more similarity between feature descriptors, which in turn implies that more mismatches may occur when tracking features from image to image. In order to improve the fraction of valid matches, additional constraints must be applied during and after descriptor matching to yield a satisfactory motion vector field that closely describes the motion of the wind tunnel model.

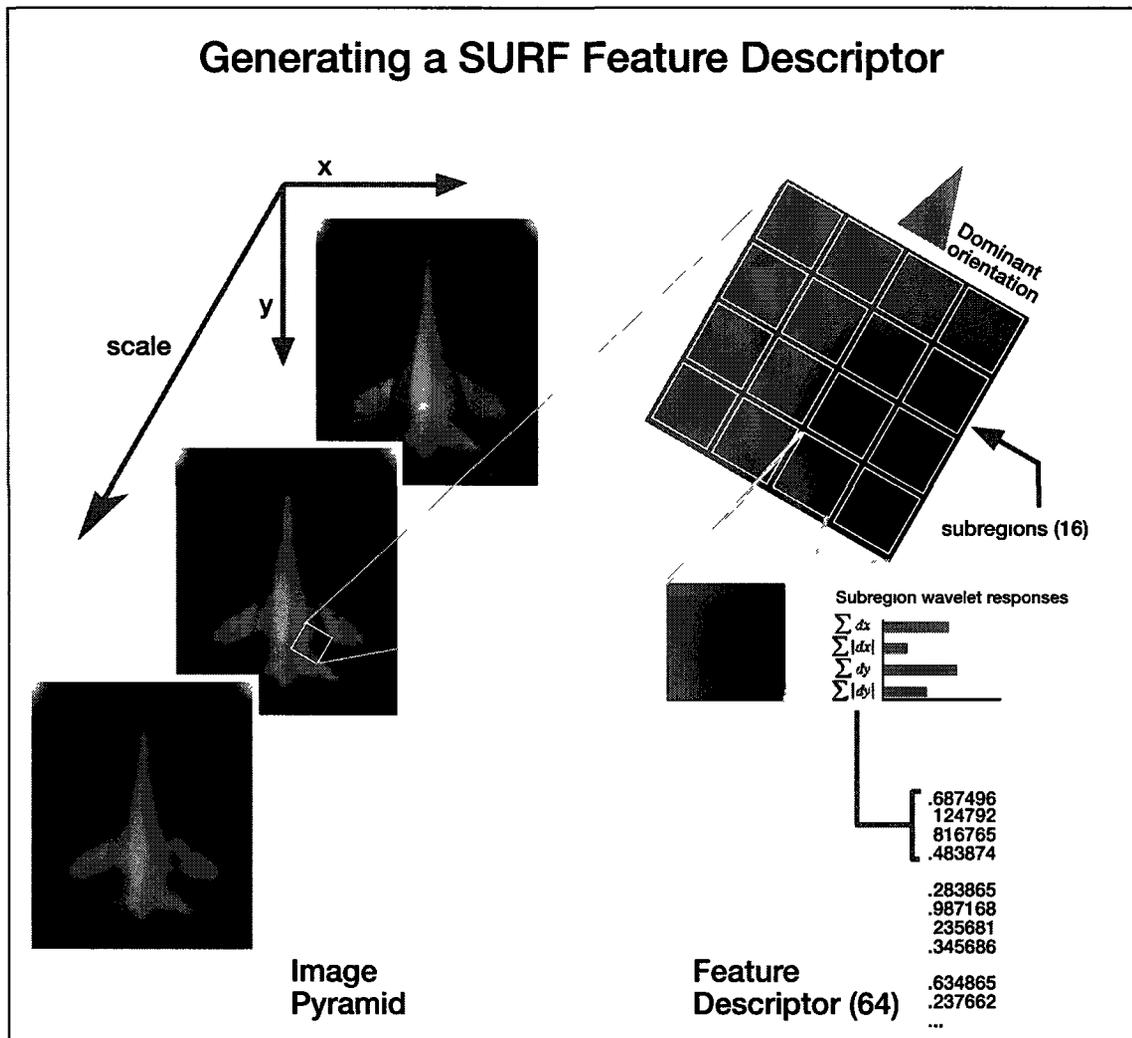


Figure 2.10: Generating a SURF descriptor. A visual work flow for generating a single SURF descriptor from a specified point in an image. The resulting feature descriptor will be invariant under limited changes in rotation, scaling, brightness, contrast, and distortion, allowing SURF features to be re-identified in multiple images.

One example of a constraint that can be applied is that of limiting maximum feature displacement between images. Since the wind tunnel model system is designed to have limited motion between wind-on and wind-off images, a constraint can be applied that no matches are permitted between features with a large displacement between images. In application, this method is typically not sufficient to eliminate all mismatches and additional constraints must be applied using characteristics of the entire vector field, as described in the following section.

## **2.6 Applying Constraints:**

### **Disparity Gradient Filtering**

#### **2.6.1 Background**

Constraining SURF feature matches based on descriptor similarity and physical displacement between images is not sufficient to remove all mismatches. These methods do not take advantage of global characteristic of the total motion vector field produced between images and a typical result of natural feature matching between images can include significant fraction of mismatches.

In the wind tunnel application images are of physical systems which include only small displacements and deformations; large disparities in direction or magnitude between physically adjacent motion vectors should not occur. The goal is to determine which vectors represent true model deformation and motion, and which are erroneous. Feature motion vectors which do not represent true motion between images can be classified as outliers. It is desirable to remove these outliers to recover a vector field that represents true motion, and methods of applying con-

---

straints to the whole global motion vector field provide a tool for accomplishing this.

Global methods of removing mismatches typically rely on geometric constraints: assumptions about how features may move collectively between images. In the case of a non-deformable object moving between frames, a known set of perspective changes can occur. Techniques such as RANSAC combined with epipolar geometry can be used to recover this motion by attempting to find the motion which makes the largest fraction of all features inliers.<sup>8</sup> This method is effective for motions that are global - the whole scene changes due to a camera motion, for example.

Local changes due to object deformation produce motion fields which may not be effectively handled by global methods, however. A class of constraints derived from the concept of optical flow<sup>25</sup> can be more effective in these cases. Optical flow algorithms can constrain feature motion locally within the image - flow vectors are compared with their neighbours as opposed to all vectors in the image.<sup>26</sup>

Dense optical flow techniques rely on the changes in local pixel intensity as indicators of frame-to-frame motion.<sup>27</sup> These methods of registration use area-based correlation, which has been shown to be preferable in image matching applications where data is low-contrast greyscale features. Although these are the types of textures seen in the PSP application, these methods were found by Aitova et. al. to not be suitable in applications where illumination levels change, which is also the case in PSP<sup>18</sup>. An example of one of these methods, Pyramidal L-K<sup>28</sup> is shown in 4.10 on page 85. Although this method works well for perspective motion, as shown in the common 'Yosemite Sequence',<sup>29</sup> the method is not tol-

---

erant to the local intensity changes due to the wind tunnel model PSP coating, resulting in significant prediction error. Another method that is tolerant of PSP intensity changes is required, and basing this method on the brightness and contrast invariant SURF feature displacement vectors combined with a disparity gradient filtering can serve as a foundation.

## 2.6.2 Applying Disparity Gradients

A disparity gradient filter works on a set of motion vectors, filtering outliers based on their deviation from the median characteristics of the whole set. The disparity between any two vectors describes, as a single scalar value, their differences in angle and magnitude in relation to their mutual distance. Two dissimilar vectors located a large distance apart (compared with their magnitude) can have a similar disparity to two similar vectors that are separated by a smaller distance. Figure 2.11 on page 38 shows a physical interpretation of disparity calculation between two vectors. In application, this means that a filtered vector field will have local agreement in local magnitude and direction while allowing smooth changes in global magnitude and direction. Furthermore, the algorithm allows the amount of local-to-global variance to be adjusted based on a global scalar parameter.

Disparity gradient is a measure of a set of vectors' mutual disparity. It is calculated as a median value from the mutual disparities of all combinations of vector pairs in the field. Any one vector in the field can then be classified in terms of its variation from this scalar median. A vector with a disparity  $d$  times greater than the median disparity of the field is classified here as having a disparity of  $d$ . By filtering out vectors from the field whose disparities exceed a threshold value of  $d_t$ , the field can be effectively smoothed to varying tolerances.

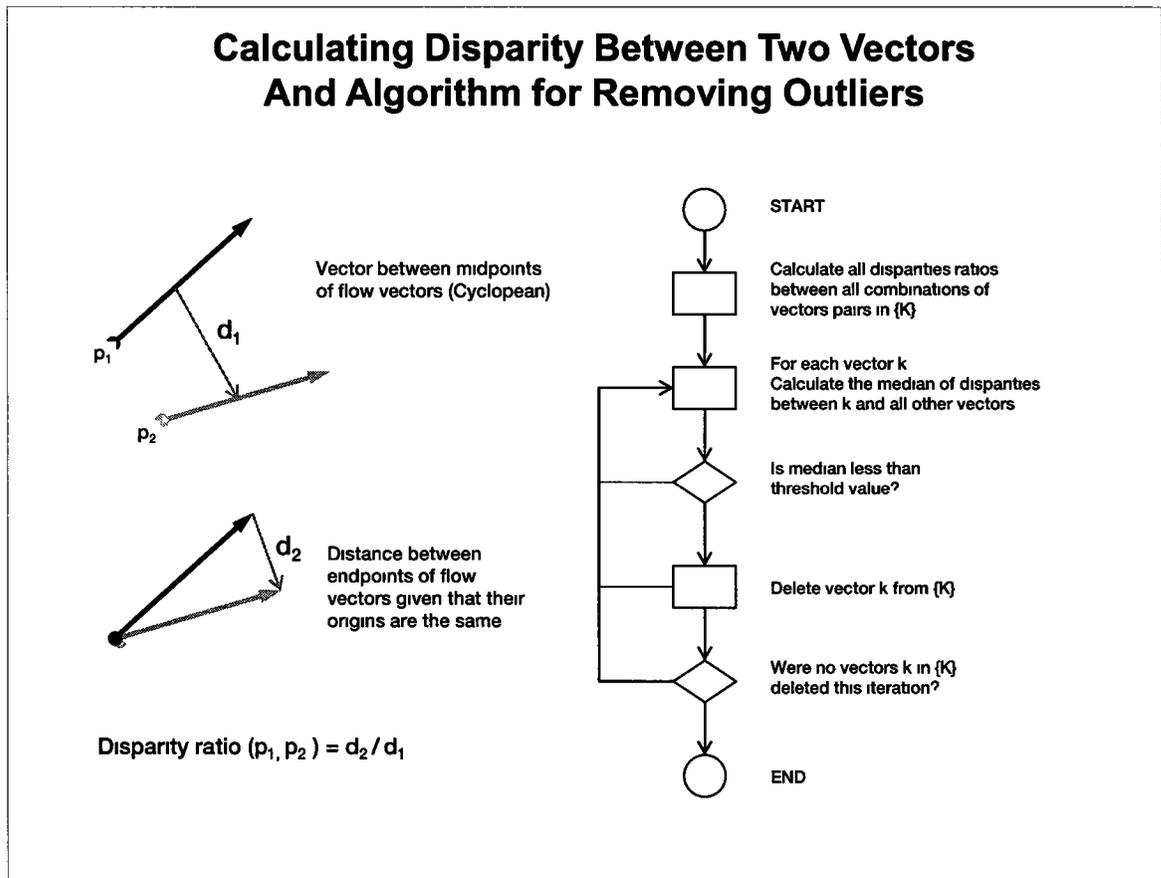


Figure 2.11: Calculating Disparity and filtering Disparity Gradient Fields. Physical interpretation of disparity between two motion vectors for generating their mutual disparity. The algorithm for filtering a disparity gradient field is shown on the right.

---

Using this measure, vector fields can be filtered to include a subset of vectors that represent desired levels of local variance. An iterative process is shown in Figure 2.11 on page 38 which can filter a vector field to a desired level of consistency, essentially a self-similar local vector flow.

Disparity gradient was developed originally as a method of quantifying feature correlation between images in stereo vision systems.<sup>30</sup> It is demonstrated by Triveti et. al. that a vector field with a  $d_t=1.0$  represents a stereo vision field of view that has no occlusion or overlap in the target scene.<sup>31</sup> A value of 2.0 or greater indicates occlusion or other discontinuities. The concept is applicable when the two images are separated in time rather than in space, and can therefore be applied to a PSP wind-on/wind-off image pair. Experimentally it was determined that a disparity threshold  $1.0 < d_t < 2.0$  is also suitable for the wind tunnel model deformation applications where vector field flow varies because of model movement and deformation.

## 2.7 Fiducial Marker Tracking

Fiducial markers are designed to appear as distinct, easily identifiable features in an image which can be localized to sub-pixel accuracy. The primary purpose of fiducial markers is to allow tracking of model movement and deformation in the wind tunnel environment using camera systems. In wind tunnel applications in general and especially in PSP applications, they are designed to not interfere with aerodynamic data gathering; they cannot change the local airflow, nor can they cover a significant area of the wind tunnel model that could otherwise be covered in PSP. A wind tunnel model with fiducial markers is shown in figure 2.12 on page 40.

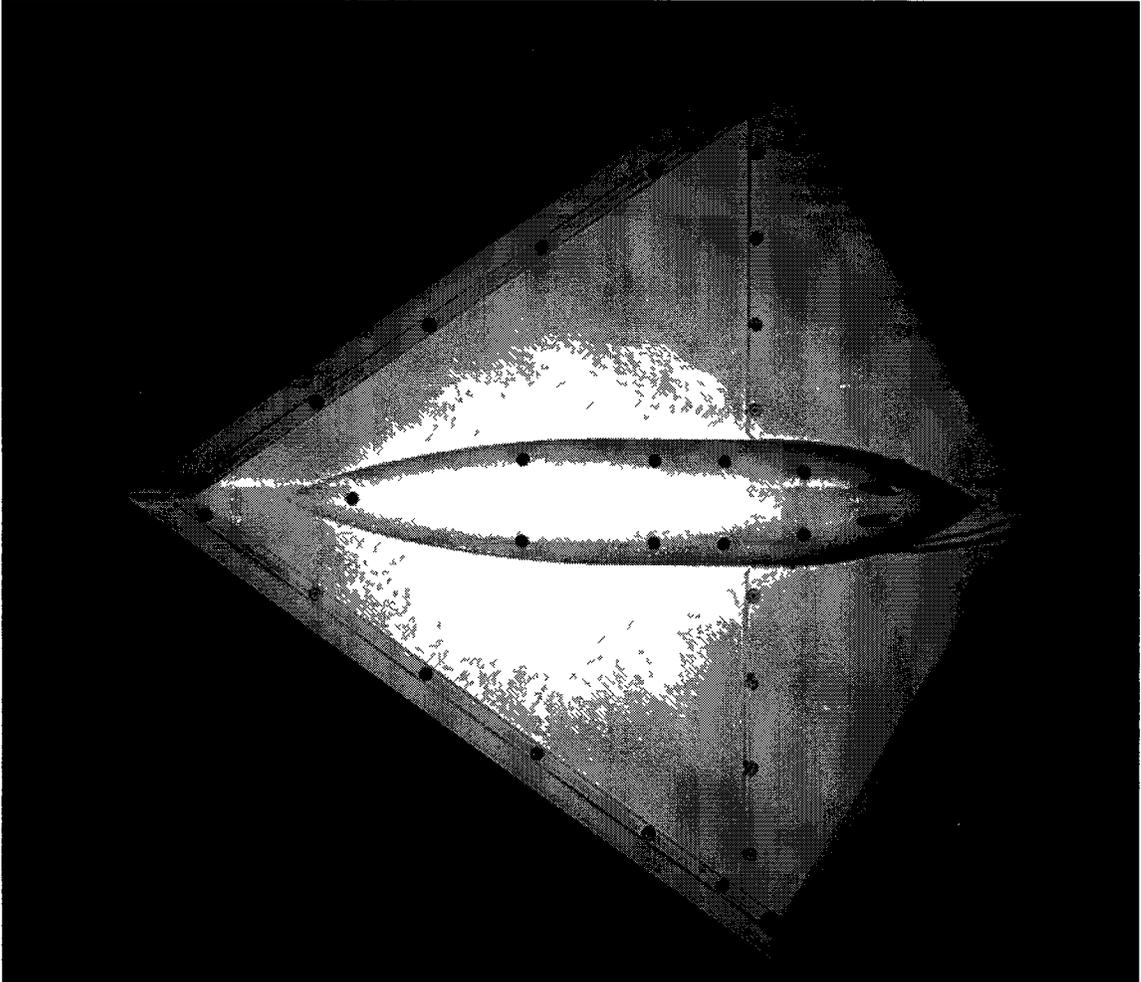


Figure 2.12: Wind tunnel model with fiducial markers. Markers are between 5 and 15 pixels in diameter to allow best sub-pixel registration without obscuring too large an area. In this model static pressure tapes are located at the center of some of the markers, creating two types of marker images to track.

As a result of these requirements, fiducial markers are typically circular marker of a minimum radius such that they appear to the camera system approximately 5-10 pixels in diameter. A circular mark is desirable for ease of localization - the center of a circular marker is readily determined to sub-pixel accuracy even when distorted due to perspective changes, which can be approximated as affine transformations for small motions. LeSant [2004] describes optimal fiducial marker size as proportional to the desired accuracy and camera system resolution, with minimal benefit in accuracy beyond 10 pixels diameter.<sup>15</sup>

Fiducial markers can be located approximately within the image in a number of ways. Shortis et. al. compare methods of detection and localization, concluding that a combination of methods chosen to suit the specific application is preferable in terms of balancing speed and accuracy.<sup>32</sup> Initial location can be done using template matching, in which the wind tunnel image is convoluted with the smaller template image containing a single marker, creating a correlation image whose maxima or minima represent best match locations. In the case of circular fiducial markers, the template match is rotationally independent, reducing the search time. Contrast and brightness-tolerant methods such as normalized cross-correlation increase the matching success further.<sup>33</sup>

Establishing the sub-pixel location of a fiducial marker is desirable in wind tunnel applications since model motion and deformations of less than one pixel can still degrade the accuracy of techniques such as those used in PSP. A minimum radius of circular fiducial marker provides visual information that allows localization to sub-pixel accuracy, since the perimeter of the marker tends to have a gradual fall-off in intensity that can be fit to a Gaussian model of the marker using a method of least squared error (LSE). The maximum of the Gaussian at the point of lowest LSE is taken as the center of the marker.

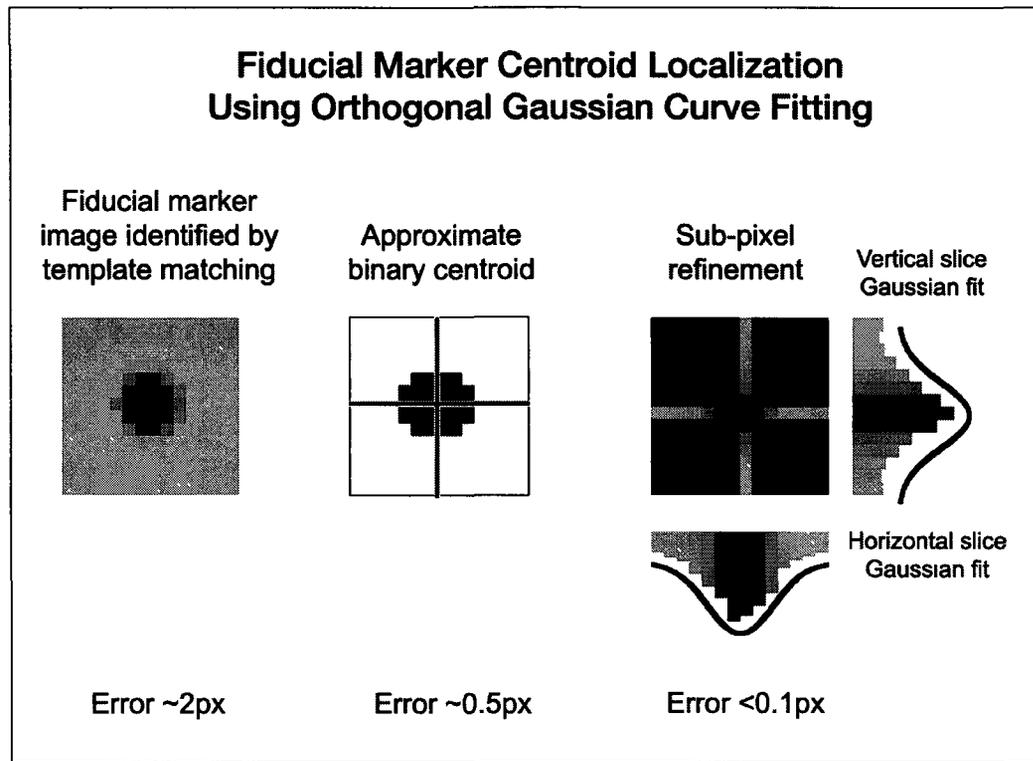


Figure 2.13: Typical fiducial marker with falloff at edges, approximate binary centroid localization, and sub-pixel refinement using Gaussian fit to orthogonal intensity cross sections.

This method can provide accuracies of approximately 0.1 pixels for markers with a 5-10 pixel diameter,<sup>16</sup> however fitting a 2-D Gaussian surface can be computationally inefficient. To address this, it is hypothesized that the marker image can be fit to a pair of 1-D Gaussian curves orthogonal to one-another, crossing at the approximate center of the marker. The concept of template interpolation from Frischholz et. al.<sup>34</sup> can be used to construct a series of sub-pixel shifted 1-D Gaussian curve templates, improving localization speed. An example of marker fitting using this technique can be seen in Figure 2.13 on page 42.

## 2.8 Conclusion

PSP's advantages and disadvantages as a pressure measurement technique in wind tunnel applications was described. The problem of misregistration was introduced and methods to reduce this through image resection were surveyed. The specific method of PLR was introduced as an alternative that can remove misregistration due to non-linear model deformation without the need for virtual models, camera system calibration, or physical marker-position measurement. The accuracy of PLR was shown to increase as the number of trackable landmarks within the image was increased, and natural feature tracking in the form of the SURF algorithm was reviewed as a method of increasing the landmark count. Filtering of mismatches was introduced as an issue affecting natural feature tracking and a method to remove mismatches based on the disparity gradient filter were described. The issue of feature tracking accuracy was introduced and sub-pixel fiducial marker tracking accuracy described in literature was established as baseline. This provides a background to define the problem and select tools and methods to address it.

# Chapter 3

## Solution

This chapter describes the methods used to implement a piecewise linear resection workflow based on the combination of motion data from natural features and fiducial markers. While natural feature tracking can increase the total number of points tracked and the accuracy of resection, it was determined that combining the motion vectors from these features with those from potentially more reliable sources, specifically fiducial marker tracking, was desirable. First, the SURF natural feature identification system was chosen from those previously developed. Second, a method of accurately tracking fiducial markers to sub-pixel accuracies was developed and implemented. Third, a robust method of identifying each feature uniquely between the wind-off and wind-on image was developed to produce a motion vector field. Finally, a system was devised to combine motion data from both sources and filter out any erroneous or inaccurate vectors to create a final motion gradient across the image.

This system allowed the weighting of vector sets for relative reliability and accuracy, and provided various degrees of thresholding, allowing more or less dis-

parity in the final vector field. This allowed a user to balance overall coverage of the image with accuracy of tracking.

### **3.1 Overview of Work Flow**

In order to test the hypothesis that natural feature tracking could improve PLR, a software work flow was developed using open source libraries. This work flow was designed to be modular so that performance at each stage of the process could be evaluated. An overview of the architecture can be seen in Figure 3.1 on page 46. Each module is a discrete executable that has specific input parameters defined via standard YML, and output results in the same format. Modules also accept images for processing to show graphical results; for example locations of found landmarks.

Once a module is configured and executed, its output data is used as input for the next module. Because this data is simply YML formatted files, any other software platform that can work with YML can use a subset or all of this work flow. The proposed architecture can run automatically as a batch process; however the user can set multiple parameters that change the performance of each module depending on the specific characteristics of the images being processed. The input images are the 16-bit gray scale wind-on and wind-off images, and the output is a 16-bit gray scale pressure ratio image that has distortion reduced using PLR.

Each module is described in detail in the following sections. Since the modules operate sequentially, they are described in the order of execution.

## Software Module Architecture and Work Flow

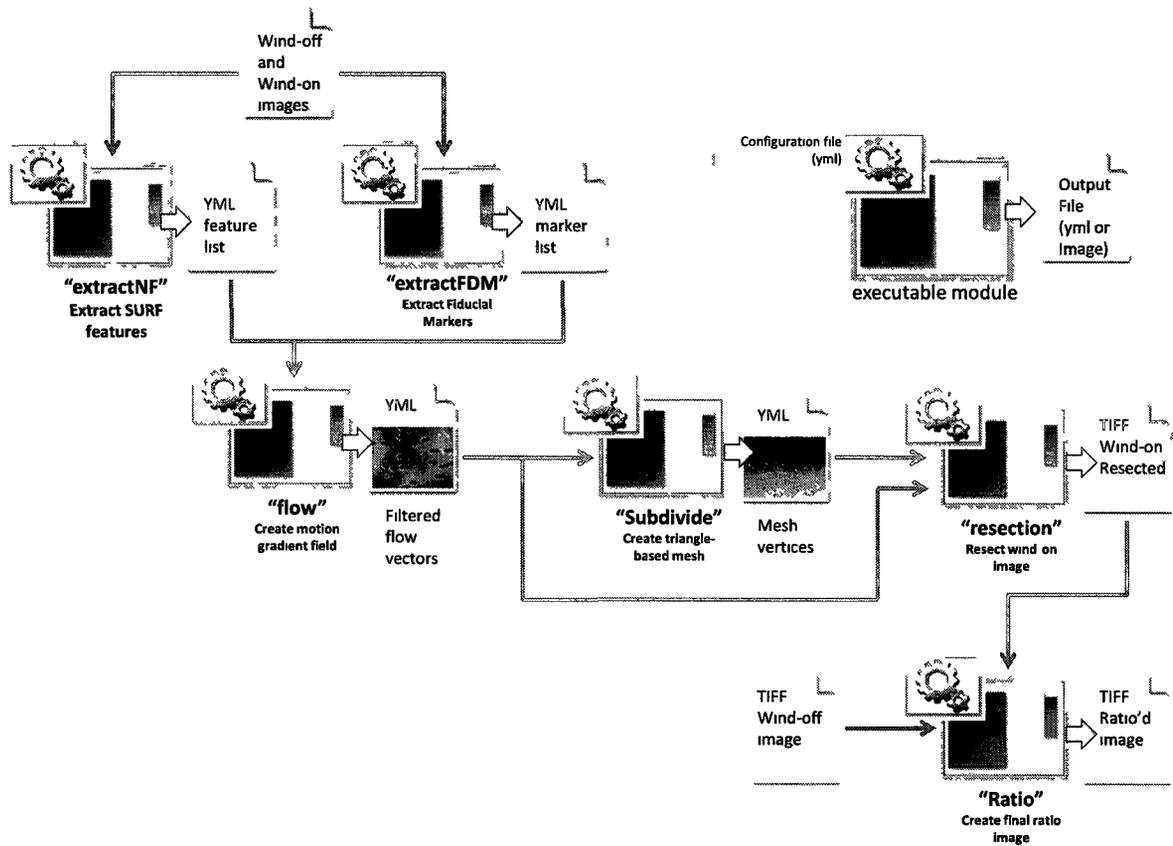


Figure 3.1: Software architecture and work flow. The work flow is handled by several discrete executable modules that use YML configuration files and TIF images as input. The output of each module is a set of processed images and YML data. Work flow begins here in the top left and finishes in the bottom right.

## 3.2 Fiducial Marker Detection and Localization

The "extractFDM" Module locates the fiducial marker position to an accuracy better than 0.1pixel using the a multi-step process of template matching, binary centroid localization, and sub-pixel refinement, as shown in Chapter 2, Figure 2.13 on page 42. Template matching is done using normalized cross-correlation with a synthetic marker template that is generated from user parameters, and sub-pixel localization accuracy can also be adjusted to trade speed for accuracy by setting the range and step-size for Gaussian fitting. The module applies gradient background removal to improve accuracy in areas of high model curvature, and applies pre-blur before determining the sub-pixel center of the marker in order to reduce the effect of image noise. The module will output the 2D sub-pixel coordinates of all markers found in a single YML file. One file is generated for each of the wind-on or wind-off images.

## 3.3 SURF Feature Matching Between Wind-On and Wind-Off Images

SURF feature detection and descriptor generation is performed by the "extractNF" module and uses the method described in by Bay et. al. [2006]<sup>11</sup> on an 8-bit gray scale version of the wind-on and wind-off images. The user is able to apply non-maximal suppression (NMS) to the resulting feature set, filtering based on the maximum Hessian score in a local area of the image. The module outputs the 2D sub-pixel coordinates and descriptors of all features retained after NMS filtering in a single YML file. One file is generated for each of wind-on or wind-off image.

## 3.4 Motion Vector Field Generation

The "flow" module performs two related tasks. The first is to generate motion vectors from wind-off to wind-on. This is done by matching SURF features and fiducial markers using the output of "extractNF" and "extractFDM". Fiducial marker and SURF feature motion vectors are determined separately and then combined. The second task is filtering of this motion vector field to remove outliers by implementing a modified disparity gradient filter (DGF). An overview of each of these tasks is provided below.

### 3.4.1 Fiducial Marker Matching

There is no unique descriptor generated for each fiducial marker position extracted by "extractFDM" so the constraint of minimum displacement is used to match markers between images. This is suitable for the wind tunnel application since the displacement of the model typically far less than the distance between any two markers. Unmatched markers are discarded from the motion vector set. The user is able to adjust the maximum allowable marker displacement in the matching algorithm.

### 3.4.2 SURF Matching

Natural features are matched on a combination of constraints. The first constraint is that a match must not exceed a specified maximum displacement in the image. The second uses the feature descriptor as described in Chapter 2. The difference between any two SURF features' descriptors in the entire set of SURF features can

be used to determine matches. An algorithm is used to sort pairs SURF features from the wind-off and wind-on images based on closest descriptor match. Since the descriptor is a vector of 64 float values, a match score is the Euclidean distance between these descriptors with a lower distance giving a better score. Since the number of natural feature found in wind-off and wind-on images is typically unequal, unmatched features are discarded. The user is able to adjust the maximum allowable displacement in the image and the maximum Euclidean distance between feature descriptors in order to minimize mismatches.

Other methods typically used as constraints in SURF matching cannot be used effectively in this application since the Hessian scores of the features tend to be lower than in other applications. For example, pre-filtering by the sign of the Laplacian was found in initial testing to filter out some potentially good matches. The constraint that a feature match must have no more than half the error of the next best match also was found to discard too many potential good matches. Constraints based on similar Hessian score were considered, however it was found that the Hessian of a given feature could vary from wind-off to wind-on images.

### **3.4.3 Disparity Gradient Filtering**

The two sources of motion vectors - fiducial marker matching and SURF matching - yield a set of displacement vectors representing the combined motion and distortion detected in the wind tunnel image. Although each motion vector is subject to matching constraints as described in the previous section, no constraint has been applied to the entire vector field; this global constraint is valuable in reducing the number of false matches. Given that model motion and deformation yield flow vectors which are at least locally consistent due to the relative rigidity of the

model, constraints can be applied to remove flow vectors which are significantly different from their local neighbours in direction or magnitude. Disparity gradient filtering was proposed in Chapter 2 as this global/local operator and in this module is extended in order to account for the two different sources of motion vectors.

### **3.4.4 Extending DGF With Vector Weighting and Prioritization**

As described in Chapter 2, the disparity gradient filter algorithm can remove outliers from the vector field in such a way that the remaining vectors are those more likely to be part of the actual model motion. However this algorithm assumes that all vectors in the field have equal weighting - specifically, all vectors have an equal chance of being valid or invalid and no vector is defined to be more reliable or more accurate than any other.

In this application, however, the fiducial markers can be localized to greater sub-pixel accuracy and with less mismatch error than SURF features. It follows that their motion vectors can be treated as more accurate and reliable than motion vectors from SURF features. This can potentially increase the accuracy of the filtering process in the local neighbourhood around heavily weighted motion vectors.

To test the potential benefits of this system, a flexible solution was developed to implement vector weighting and prioritization. The system increases the weighting factor of a vector by duplicating it multiple times in the vector field, effectively stacking the vector in the same location. This skews the disparity averages within the disparity gradient filter algorithm in favour of the heavily weighted vector, achieving the same effect as having many parallel local vectors of the same magnitude, as shown in Figure 3.2 on page 51 A disparate vector in that

local area would be considered an outlier with a lower disparity than would otherwise be the case, resulting in a local vector field that has a lower overall disparity from the heavily weighted vector.

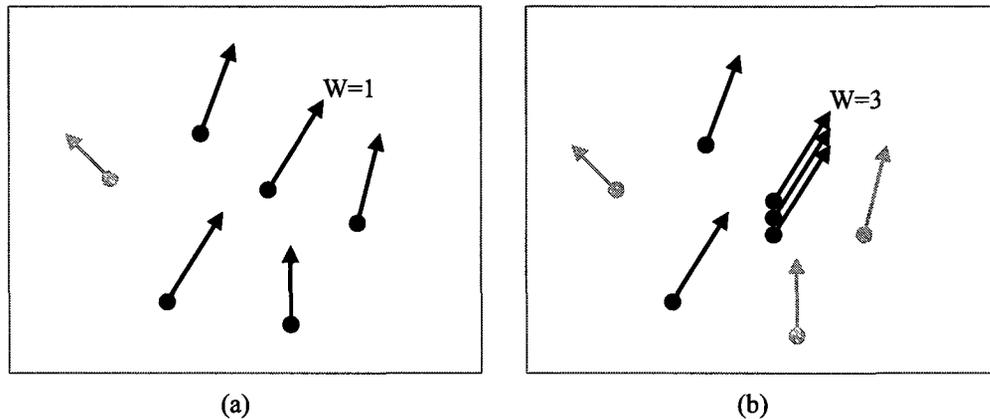


Figure 3.2: Effect of weighting on local disparity gradient filter. The filter removes vectors which exceed a local disparity threshold, here shown in as lighter-shaded vectors. Weighting a given vector by an integral factor of  $W$  gives it influence within the disparity gradient filter equivalent to  $W$  superimposed vectors in the field, effectively given the vector more local influence in its locality. In this example, a vector is assigned a weighting factor of  $W=3$ , resulting in an additional local vector being classified as an outlier.

Because some vectors, such as those created from the motions of fiducial marker, are considered very reliable, and some, such as those from natural feature tracking, are considered less reliable, a method of prioritizing vector subsets within the disparity gradient filter was developed. Within the overall gradient field, subsets of vectors could be assigned a priority number. Those sets with the higher priority number are included in the disparity calculations for all vectors with lower priority number, but not vice versa. Effectively this prevents vectors with a lower priority number from filtering out vectors with a higher priority number, creating a hierarchy of vector subsets. A visual representation of this can be seen in Fig-

ure 3.3 on page 52 Vectors having a maximum priority are a special case in that they cannot be filtered out of the gradient field at all. In this specific application, fiducial marker motion vectors were assigned this maximum priority number and vectors from natural feature motion were assigned the same lower priority.

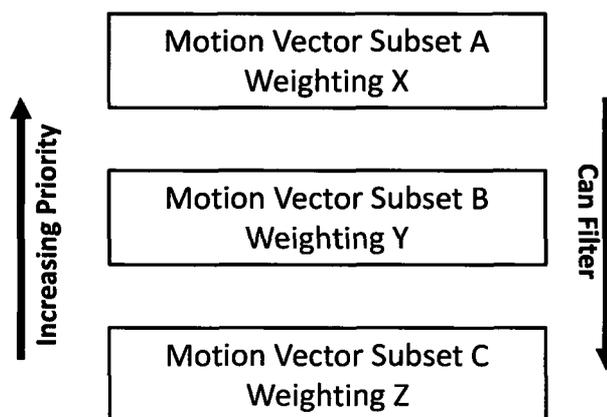


Figure 3.3: Hierarchy of vector subsets within a field. Priority levels can be used to explicitly prevent subsets of vectors from removing higher priority subsets. Vector removal can only operate from highest to lowest priority, which is useful when multiple sources of motion vectors are present and have different reliabilities or accuracies. Each subset can be independently assigned a weighting factor.

The "Flow" module also allows segmentation of the image into a regular grid in order to test the effect of applying DGF to the image in separate segments. Motion vector filtering in each segment is decoupled from all other segments, potentially removing more outliers in the overall image without discarding valid local motions. The user can specify the number of segments vertically and horizontally.

This modified DGF implementation can be tuned by the user based on the specific characteristics of the wind tunnel imagery under consideration. The overall maximum disparity can be specified, and the relative weight of fiducial marker-

based motion vectors can be specified. The goal is to extract a motion vector field that is locally consistent yet flexible enough to allow for distortions in the model that are due to bending and warping as well as simple movement. Once this motion vector field is saved as a YML file, it is used to generate an image subdivision mesh and to un-distort that mesh to resect the wind-on image, as described in the following section.

## **3.5 Resection**

As described in Chapter 2 resection is the process of un-distorting an image by reprojecting it onto a polygonal mesh. In PLR, the mesh is generated automatically from the motion flow vectors extracted by the "flow" module. The end point of each motion vector can be used as the vertex of the mesh, and the mesh can then be undistorted by moving each vertex along its displacement vector, as shown in Figure 3.4 on page 54.

Because each vertex is displaced individually, the whole image is effectively undistorted as many local areas, each based on motion detected in that specific region. Because wind tunnel model deformation is a local distortion of the image, this method of using many areas allows a more accurate resection of the wind-on image when compared to fewer, larger areas. The resection process is carried out in sequence by the "subdivide" and "resect" modules, described below.

### **3.5.1 Creating the Mesh**

Creating a polygonal mesh from a set of points must be done such that all areas are covered by a polygon and no polygons overlap. For this application, the

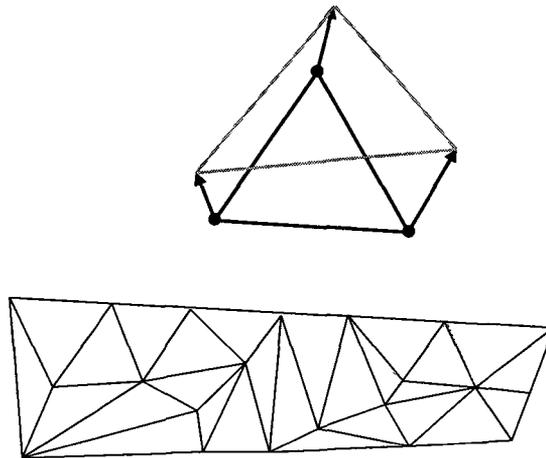


Figure 3.4: Resection of a polygon within a mesh. A resection mesh is shown with a single triangle detailed above it. Each triangle in the mesh is individually undistorted by moving each vertex along its respective motion vector. Because triangles share vertices across the mesh, the whole image undergoes a seamless resection.

mesh is composed of non-overlapping triangles, which are easily handled during un-distortion and re-projection operations described in the following section. As discussed in Chapter 2, the Delaunay triangulation algorithm for creating a non-overlapping mesh from a set of points is an optimal solution. Delaunay triangulation maximizes the minimum internal angles for all triangles in the mesh, minimizing the number of triangles where any given interior angle is significantly less than the others. The subdivision functions within OpenCV<sup>35</sup> provide a solution for implementing this algorithm on the point set made from the origins of all vectors in the image motion field. Each point is added to the subdivision and the Delaunay triangulation output as a sequence of vertices, most of which are repeated since they can belong to multiple triangles.

Areas outside this mesh are not assigned any explicit motion by the motion vector set and can therefore have significant error in comparison to their true mo-

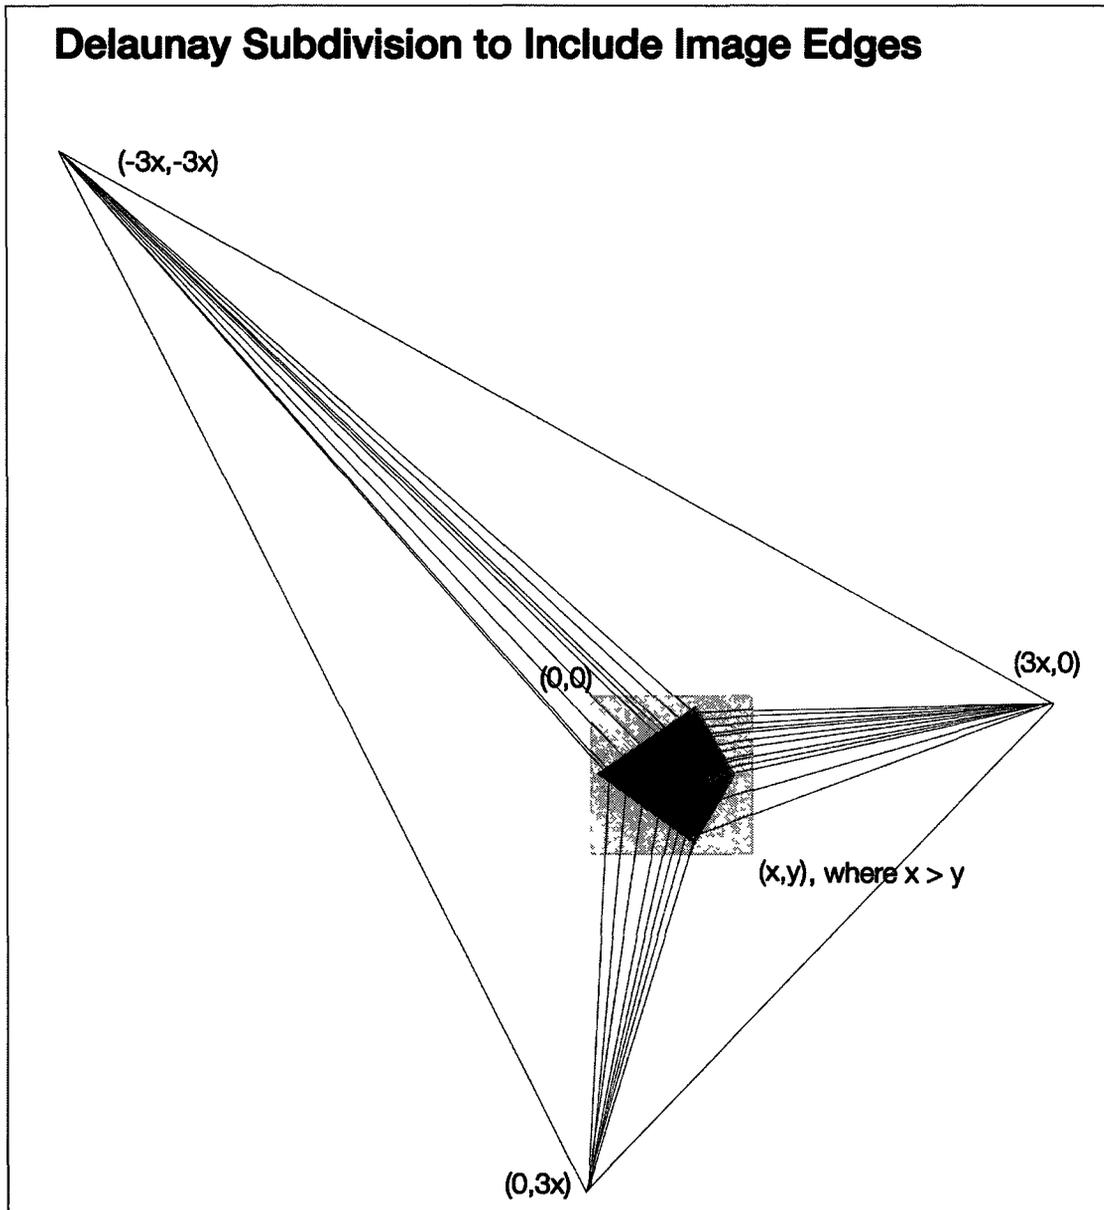


Figure 3.5: Delaunay subdivision to include image edges. Original wind-on image is seen in the center of a much larger triangle. This large triangle's vertices are also vertices in for all edge triangles which are interior to the resection mesh. In this way image edges can be included in the resection process.

tion.<sup>36</sup> In this implementation, an attempt is made to address this by extending the mesh beyond the visible portion of the image so that the triangulation covers the whole image. Subdivisions of the image which are not contained within the defined points of the mesh are instead contained by triangles which have a vertex at one of three points off the image. These points are defined as  $(-3x, -3x)$ ,  $(0, 3x)$ ,  $(3x, 0)$  where  $x$  is the maximum pixel dimension of the image. Any edge triangle will have one or two vertices attached to one or two of these points, as shown in Figure 3.5 on page 55. Because the points are located well beyond the edge of the image, any resection that affects an edge triangle will degrade from a full undistortion of all three vertices to an affine transform using information from one or two vertices. This is an approximation of motion that is automatically based on the motion of the remaining vectors composing the triangle, and in this application can be more accurate than simply "pinning" edge triangles to the edges of the image.

### 3.5.2 Undistortion and Reprojection

The Delaunay triangulation subdivides the wind-on and wind-off images into a triangle-based mesh which can then be undistorted per-triangle, allowing for local correction of wind tunnel model deformation and bending. The goal is to undistort the wind-on image to remove model motion and deformation, aligning it feature-for-feature with the wind-off image. This alignment process is carried out by moving the mesh vertices (each of which represents a tracked feature) backwards along the flow vectors describing those vertices' motions between the wind-off and wind-on images. The image information within each triangle is linearly interpolated based on this undistortion.

In practice, the FreeGLUT implementation of the OpenGL graphics standard was chosen to implement this process since it provides a facility for importing the triangular mesh and projecting a texture onto the mesh at a per-triangle level. Using OpenGL, triangles composing the mesh are defined separately by setting their vertex positions in 2D space. Each triangle is textured with a subsection of a bitmap image - in this case a triangular subsection of the wind-on image. Because OpenGL allows any arbitrary triangular piece of the bitmap to be selected to texture a triangle, the texture can be selected such that it uses the endpoints of the flow vectors, while the triangle it is projected on uses the start points of those same motion vectors. OpenGL includes a bilinear resampling algorithm that smoothly distorts the texture triangle to fit the mesh triangle on which it is projected. Bilinear interpolation has been shown to have comparable accuracy to more computationally intensive methods by Park and Sung [2004], specifically in PSP applications.<sup>37</sup> The net effect is that the wind-on image is distorted at a per-triangle level to fit the mesh geometry of the wind-off image, effectively undistorting the wind-on image based on motion vectors. This resected wind-on image is hypothesized to have reduced distortion because model motion and deformation have been tracked and corrected at a per-triangle level, and can therefore provide more accurate local PSP pressure data when used to create the final ratio image, as described in the next section.

### **3.6 Generating Ratio Images**

The procedure described to this point has acted to resect the wind-on image to remove detected wind tunnel model distortion and deformation. As described in Chapter 2, this directly affects the accuracy of PSP pressure measurements. If the

images are not aligned due to wind tunnel model deformation, the resulting pressure data is no longer accurate in the areas of the image where this misalignment takes place.

To extract the pressure data the final step in the work flow takes the per-pixel ratio between the resected wind-on image and the original wind-off image using the "Ratio" module. This is a per-pixel division of the wind-off image by the wind-on image. Values where the wind-on image is zero are set to zero. This procedure is carried out in floating point, single channel image space to maximize dynamic range. The resulting image is stored as 16-bit gray scale and the module has provision for the user to select the black point and white-point of the image for visualization. The pixel intensities of the 16-bit image are correlated with the pressure at that point on the wind tunnel model in a process which is carried out during final analysis outside this workflow.

### **3.7 Interface Design**

This process is fully automated and the software modules described previously are designed to run as a batch process in sequence, with an estimated run time of approximately 30 seconds for 1 megapixel camera images.

Each functional module described in the previous section is configured via a YML text file and can provide intermediate output as YML data files and imagery. This allows integration of some or all of the work flow into other applications, such as MATLAB.

However, the user does have a critical role to play in the work-flow. Since fiducial marker detection, natural feature detection, feature tracking and gradient

field filtering have several adjustable parameters that affect their performance and outcome, a human-in-the-loop iterative process is desirable to tune these parameters when new image batches are being processed. In order to make this intuitive and efficient for the user, a graphical user interface (GUI) is implemented to allow tuning of each software module in the work-flow. Output images and data from each module are read back into the interface for display to the user, allowing for fast iteration of parameters. A sample interface concept can be seen in Figure 3.6, page 60.

The interface was built on the cross platform GUI toolkit Qt<sup>38</sup> with additional interface elements written using OpenCV,<sup>39</sup> such as YML file handling and 16-bit image handling. The GUI works at a project level, where a project is defined as a workflow and parameters applied to a specific wind-off/wind-on image pair. All working and configuration files are stored under a project folder allowing simple relocation, and all configuration files are human readable and editable as text. All feedback images are stored as 8 bit colour or gray-scale TIF files.

In addition to the work-flow management tools, the GUI provides mask and slice tools at the start and end of the work-flow, respectively. The mask tool allows the user to select a polygonal region of interest in the images to be processed, allowing explicit exclusion of parts of the image which are not part of the wind tunnel model. This can improve the performance of the modules when identifying features and motion gradients and when applying automatic contrast enhancement.

The slice tool is available at the end of the work flow and operates on the ratio image. It allows the user to specify a linear section through the image and to view a cross sectional representation of pixel intensities along this section. This

allows quantitative analysis of shock fronts or correlation of PSP pressure values with pressure tap data. Multiple slices can be stored or recalled for any image set, allowing comparison of different wind-off/wind-on image pairs, and effectively creating virtual instrumentation on the ratio image. The slice data can also be exported in standard text format for further analysis.

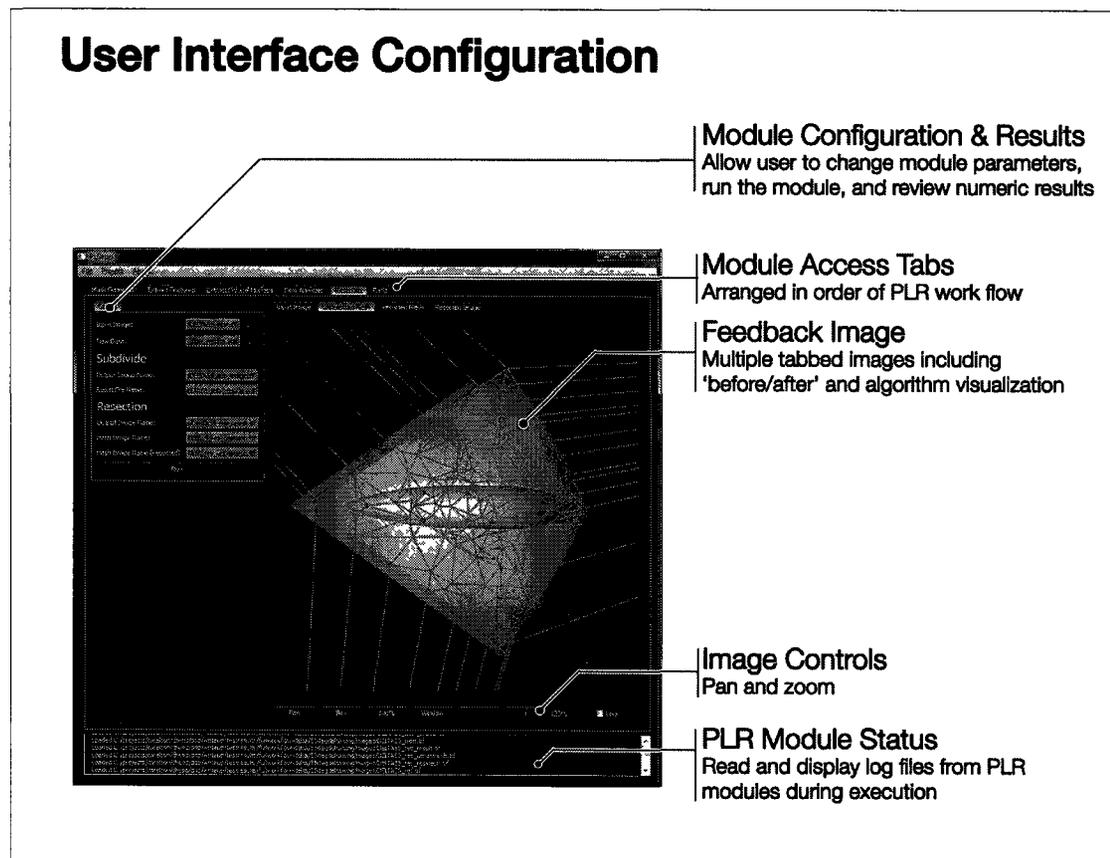


Figure 3.6: Graphical User Interface design. The user can adjust parameters and view visual feedback for each module in the PLR work flow. All configuration and file management is automatic, reducing processing time and potential errors.

This interface is designed to focus the user on the task iteratively extracting an acceptable motion field and measuring the final pressure data, while minimizing exposure to the underlying modular nature of the work-flow. Because no text

files need be directly manipulated, the interface is primarily visual, taking advantage of the user's strengths in image recognition and evaluation.

# Chapter 4

## Tests and Results

### 4.1 Establishing Ground Truth in Photogrammetry

In photogrammetry, ground truth can be difficult to ascertain because of the large set of parameters affecting transformation from a real 3D scene to a 2D image produced by a digital camera system. In order to establish the performance of any image manipulation process in terms of absolute accuracy, a ground truth must be established. In this application, there are several software modules that can be tested individually for performance metrics; however the metrics themselves will be unique for each module.

The ground truths used for module testing in this application can be divided into two categories: measurement of performance on synthetic imagery, and measurement of performance on real imagery against other sensor data. In the case of natural feature tracking, synthetic imagery cannot ultimately be used to characterize performance since the nature of the simulated imagery (i.e. synthetic PSP paint patterns) will have a dominant effect of test results; actual wind tunnel im-

agery must be used. However in the case of fiducial marker tracking, accuracy and robustness can be determined using synthetic imagery. A synthetic marker can have realistic parameters which can be accurately generated, and the accuracy needed in tracking them is high enough to require an absolute ground truth generated with the synthetic image. By first establishing fiducial marker tracking accuracy against synthetic images, that system can then be used as a reference against which to measure natural feature tracking accuracy.

The goal of PSP is to have pressure ratio measurements across the surface of the model. Static pressure taps on the model are used to calibrate PSP image intensities. This relation to pressure ratio is not required in testing and validation of the system described here when comparing the quality of the PLR-resected ratio image to others generated with no correction or with global correction. Improvements should be seen as visible reduction in PSP noise across the image and measurable in any sections of these images.

## **4.2 Overview of Testing Process**

The testing process consisted of a sequential series of experiments, such that the findings of the prior experiment could provide a basis for the subsequent one. In this was ground truths could be established and those findings used for further performance measurements.

### **4.2.1 Software Implementation**

Testing was conducted using the modular software system described in Chapter 2. The system was successfully implemented using C++ in the Microsoft Visual

---

C++ Express 2008 development environment<sup>40</sup> using open source libraries including OpenCV 2,<sup>41</sup> TIFlib,<sup>42</sup> freeGLUT,<sup>43</sup> Qt 4, and Qwt graphing extension to Qt.<sup>44</sup> Testing was conducted on consumer-grade PC hardware running Windows 7 SP1, typically Core2 Duo 1.8 Ghz with 4GB of RAM. Software was compiled for optimized Intel instruction set and multi-threading with standard optimization settings. A full work flow from raw wind-on and wind-off images to final ratio image could be produced in less than 30 seconds without user parameter tuning.

### 4.2.2 Testing Order

Testing order was important so that accuracy of modules could be verified by other modules whose performance is measured against absolute metrics. Testing began by determining the accuracy of automatic fiducial marker identification and tracking using synthetic data. Fiducial marker tracking was then tested on wind tunnel imagery with deflections in 3D space in order to validate image-to-image tracking. Once the accuracy of fiducial marker tracking was established, it was used to validate natural feature tracking accuracy in the next tests on wind tunnel imagery.

Natural feature identification and tracking were tested in combination with motion vector field filtering throughout. This combined algorithm is referred to as "SURF and DGF". The constraints used to match and filter natural features were characterized first on synthetic imagery to determine basic accuracy and image coverage and to provide absolute error data on image textures similar to those seen in wind tunnel images. Testing was then extended to wind tunnel imagery to determine the effect of model motion and changes in the PSP local intensity. SURF and DGF accuracy on wind tunnel imagery was then measured against the fiducial marker tracking module tested previously.

In testing the weighted, prioritized extension of the disparity gradient filter, performance was characterized on wind tunnel imagery, using changes in overall motion gradient field disparity as a metric. The image set chosen for this testing had localized nonlinear deformation, so the disparity of the motion gradient would decrease as filtering improved. A lower bound on this decrease was indicated by the loss of full image coverage by motion vectors, indicating that constraints were too strict to allow for nonlinear deformation. Understanding this general characterization of the image provided a suitable test data set.

In order to verify that the workflow was improving the quality of PSP data, comparisons were made between uncorrected and PLR-resected ratio images using calibrated pressure tap data. Performance improvements could then be measured as reduction in error with respect to pressure ratio readings from an independent sensor system.

Once resection data quality was validated using tap data, image resection performance away from taps was tested. Ratio images were measured and compared to uncorrected and globally resected images. Performance of natural feature tracking using residual fiducial marker displacement was first established. This was a form of regression testing since the images of the fiducial markers were resected with the rest of the image, allowing comparison of their final positions with those in reference wind-off images. Image intensity data was also analyzed for noise improvement versus uncorrected ratio images. Finally, a test against global resection was conducted on wind tunnel imagery. The wind-on image in this test had additional nonlinear deformation introduced in order to test more extreme conditions than were present in available images sets.

Further details on testing and results are presented in this chapter.

## 4.3 Fiducial Marker Detection and Localization

Fiducial Marker detection and localization tests were first carried out on synthetic target images to validate accuracy against a known software tool chain. Once accuracy was validated, testing on wind tunnel imagery was carried out to verify detection and tracking of targets between wind-off and wind-on images.

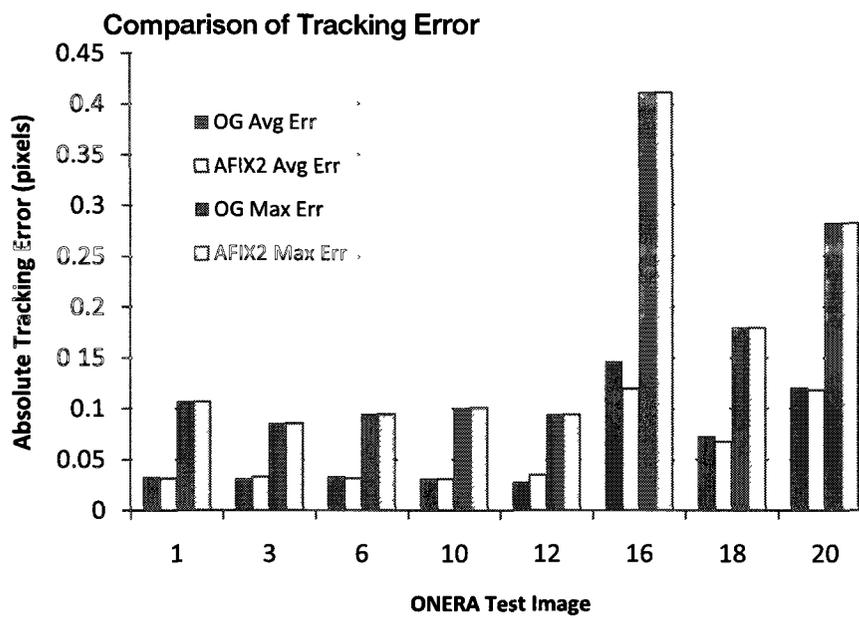
### 4.3.1 Synthetic Target Tracking

Synthetic test images with arrays of 400 fiducial marker targets were automatically generated by the SPOT tool developed at ONERA. Its method of synthetic marker generation was described by LeSant et. al. [2005]<sup>8</sup> Each image had a specified contrast level, random noise level, target size variation, and target aspect variation (affine warping), as well as a known sub-pixel center location for each of the 400 targets. This provided an absolute baseline for measuring the fiducial marker location and localization procedure.

#### Testing and Results

The first step in testing the fiducial marker tracking software was to determine if template matching would detect the relatively dense target arrangement without multiple detections or false detections. The results on the ONERA test targets are shown in 4.2 on page 68, and indicated that template matching resulted in a 100 percent detection rate with no false matches or duplication detections over a the full range of noise, contrast, blur and distortion parameters. No change in detection parameters was needed over the image set.

ONERA Targets Parameters			
ONERA Synthetic Image	Gray Levels (contrast level)	Noise (%)	Blur (pixels)
1	50	0	0
3	50	0	2
6	100	0	1
10	200	0	1
12	200	0	2
16	200	50	0
18	200	25	2
20	50	50	2



#### Template Matching and Sub-pixel Refinement Settings

Synthetic marker radius = 4px  
 Synthetic Template Width = 20px  
 Synthetic Template Pre-blur Radius = 3px  
 Repeat-Detection Exclusion Radius = 8px  
 Template Correlation detection threshold = 0.50  
 Orthogonal Gaussian fit pre-blur radius = 23  
 Orthogonal Gaussian search range  $\pm 2.0$ px at step size = 0.002

Figure 4.1: ONERA test target tracking results. The orthogonal Gaussian matching system's accuracy was compared to that of AFIX2 for the ONERA test target patterns. AFIX2 is a purpose-designed PSP resection software system and the test patterns were generated to validate its performance. Accuracy of the two software systems was comparable across a wide range of noise, contrast, blur, and distortion levels.

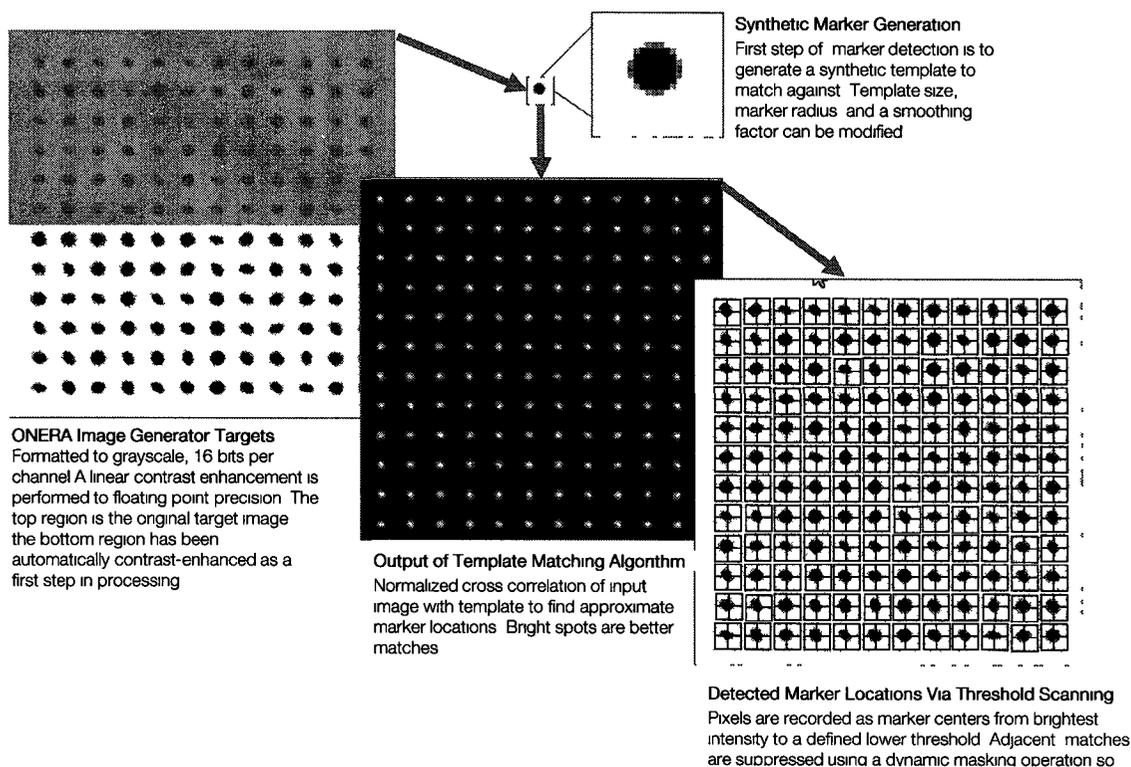


Figure 4.2: ONERA target template matching results, The template matching algorithm successfully detected all targets with no false detections or duplicate detections over the full range of noise, contrast, blur, and distortion parameters with no change in detection parameters.

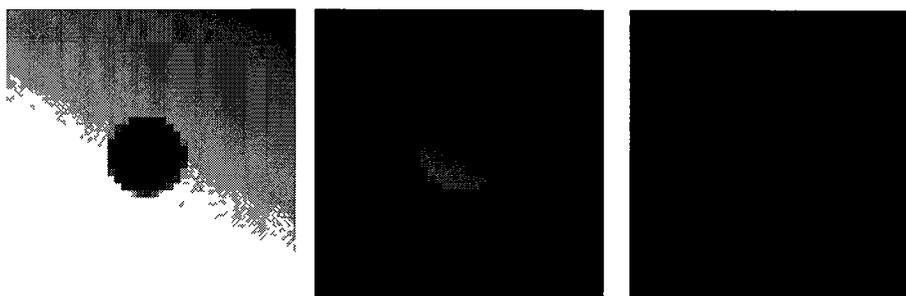


Figure 4.3: Gradient background removal before fiducial marker sub-pixel center localization. From left to right: marker before background removal, after background subtraction (inverted by the algorithm), and after the pre-blur step which precedes Orthogonal Gaussian sub-pixel refinement. In this extreme test case the gray scale edges of the marker were preserved and refinement error was reduced from 0.3px to 0.2px compared to the same marker with no background noise.

Testing then moved to the orthogonal Gaussian sub-pixel refinement algorithm. Results are shown in Figure 4.1 on page 67. The tests were conducted with the same ONERA synthetic marker set since an absolute sub-pixel center for each marker was known. When configured with the correct parameter range, the software successfully detected all markers and localized their centers to sub-pixel accuracies comparable to AFIX2, a software package designed in part to detect and track fiducial markers in wind tunnel PSP applications. Performance was comparable over a range of noise, contrast, blur and distortion levels indicating that the software met accuracy requirements. During testing it was noted that accuracy values were robust for a given set of similar images. In practical terms this indicates that extensive fine-tuning of parameters is not required to achieve sufficient accuracy for this application. Localization accuracy was affected most at high noise levels and very low contrast levels. Both of these conditions affected the quality of the information available at the fiducial marker perimeter, which in turn had an impact on the Gaussian fitting process.

The gradient background removal process developed for this application was also tested on synthetic imagery, as seen in 4.3, page 68. Steep background gradients like these can occur in wind tunnel imagery near areas of high model curvature, such as wing leading edges where fiducial markers tend to be placed, so this performance metric is of practical importance. Accuracy was improved significantly when compared to displacement versus the same marker image with a solid background. In this extreme test case the gray scale edges of the marker were preserved and refinement error was reduced from 0.3px to 0.2px compared to the same marker with no background noise. In practical applications, the gradient would typically not vary from to this degree between wind-on and wind-off images, so the low absolute accuracy found in the boundary test is acceptable, since

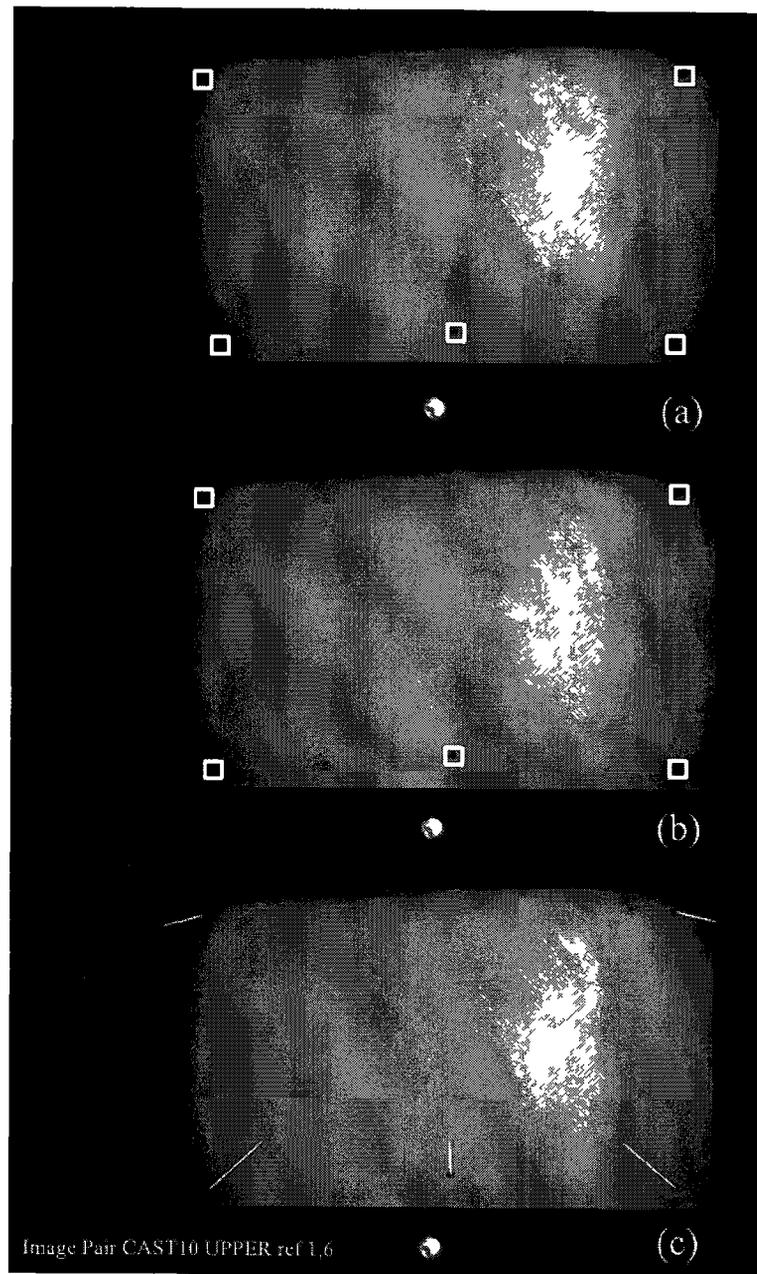


Figure 4.4: Matching fiducial markers by displacement. Images (a) and (b) show a wing section rotated about its center horizontal axis 10 degrees such that the lower edge is approaching the camera and upper edge is receding. Five fiducial markers of a radius of approximately 2 pixels are present and all are identified in both images without false matches. The resultant motion tracking vectors are shown in (c). In this application the fiducial marker displacement is small enough that there is no possible transposition or rotation ambiguity, making minimum-displacement a robust matching method.

it demonstrates a significant improvement.

## **Discussion**

On actual wind tunnel imagery, the fiducial marker detection was robust within a dataset once the appropriate detection parameters had been manually set to approximate marker characteristics. Marker detection success rate was most directly affected by the synthetic template marker radius and the detection threshold parameters. Rejection of false matches was controlled by increasing the size of the template around the synthetic marker; however this technique is limited near the edges of the wind-tunnel model where markers tend to be placed. Enough tuning parameters were available to get acceptable detection results on a range of synthetic and real images.

Marker localization to sub-pixel accuracy was most directly affected image pre-blur and the step-size used when fitting a Gaussian curve to the orthogonal fiducial marker image slices. The accuracy of this localization did not vary appreciably with the shape or magnitude of the Gaussian, but was affected by the amount of pre-blur applied to the image before fitting. Up to an image-specific threshold, increasing the pre-blur was found to increase accuracy.

## **Results Summary**

Testing of the fiducial marker tracking module on synthetic targets showed that detection and localization was robust over several simulated image distortions and as accurate as AFIX2, a purpose-built software tool.<sup>15</sup> Tests on wind tunnel imagery demonstrated successful localization of markers over a broad range of image sets. The three-stage approach to fiducial marker location and localization was shown

to be efficient, taking less than 5 seconds to localize approximately 20 markers in a single 1-megapixel image. The system was shown to be flexible enough to locate both fiducial markers and pressure taps, which have similar visual characteristics. Gradient background removal was shown to be effective in enhancing localization accuracy in synthetic test cases where the fiducial marker was modeled on a surface of high curvature.

In summary, this module met the fiducial marker tracking requirements of the PLR application work flow. In practical application it was found that setting synthetic marker template radius to the same value as the markers in the PSP image was most effective in locating markers, and that setting a Gaussian step size equivalent to half the center accuracy desired yielded satisfactory results. Increasing Gaussian pre-blur of the marker image before curve fitting was found to increase center accuracy up to the point where the marker image was clipped by the edge of the curve fitting subregion.

## **4.4 Natural Feature Detection and Localization**

Natural feature extraction and tracking is used to increase the number of triangles used in the resection mesh, thereby increasing the local accuracy of resection. This will only be valid if natural features can be tracked to a sufficient accuracy. Since natural feature detection and tracking can result in mismatches which yield false motion vectors, testing was conducted with the disparity gradient filter as part of the filtering and matching process.

Testing addressed three aspects of natural feature detection using the SURF algorithm. Synthetic imagery was used to determine tracking error against ground

---

truth values. Feature identification and matching on wind tunnel data determined how natural feature tracking worked with actual wind tunnel imagery. Finally, fiducial marker tracking data was used as a ground truth to determine natural feature tracking accuracy in wind tunnel imagery.

#### **4.4.1 SURF Natural Feature Matching with PSP**

This set of tests used wind tunnel imagery in order to characterize SURF feature detection, tracking and filtering over a range of parameters. The matching algorithm between SURF features creates motion vectors from pairs of features with minimum error between their descriptors arrays. Methods suggested by Bay et al. [2006] provide additional constraints on this method initially comparing the Laplacian of the descriptors and by discarded matches that are not significantly better than the next best match.<sup>11</sup> In initial testing it was found that these constraints resulted in few feature matches on wind tunnel imagery, primarily because of the low-contrast, self-similar PSP texture. Since additional constraints would be introduced by disparity gradient filtering after feature matching, the constraints during the matching stage were adapted for the PSP application. The following tests measure the performance of these modified constraints on wind tunnel imagery.

##### **Generating SURF Features from PSP Imagery**

High quality features (i.e. those with a high Hessian score) have distinct descriptors which in turn provide improved discrimination between matches and mismatches of features between images. Figure 4.5 on 74 shows the distribution of SURF Hessian scores in a sample wind tunnel image. The majority of features have

a lower Hessian score, attributable to the low contrast and lack of hard edges in the PSP texture. This implies that modified approaches to identifying and matching SURF features are required.

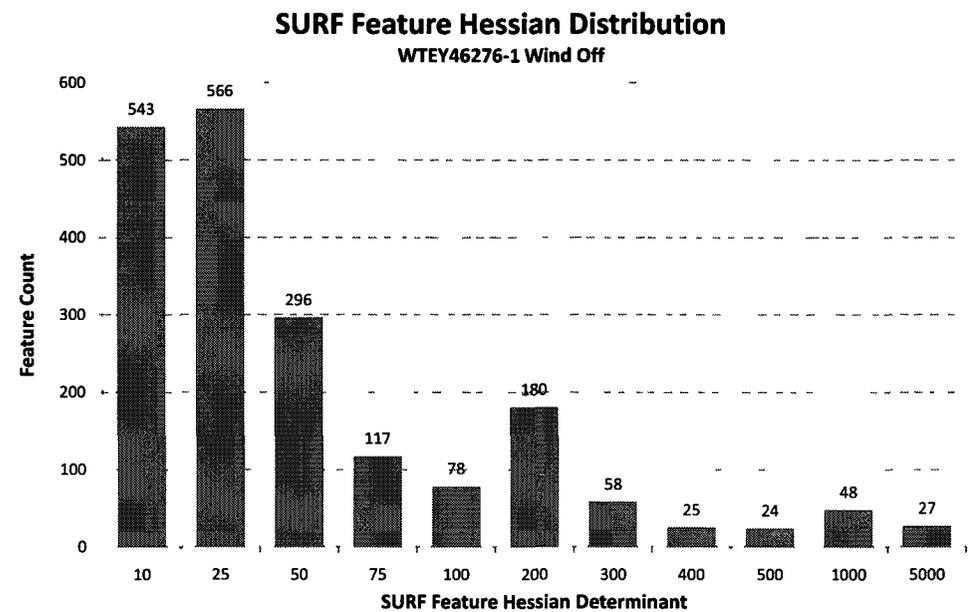


Figure 4.5: Hessian Score Distribution of PSP Texture. Hessian scores for a wind tunnel image, primarily based on PSP texture. Extremely large values are considered outliers. The distribution is skewed towards lower scores due to lack of high contrast local features.

In order to provide a set of high quality features across the whole image and remove lower quality features, a non-maximal suppression (NMS) filtering process was applied to the full set of SURF features. The software allowed the user to specify a radius of exclusion in the image which defines the minimum distance between any two filtered SURF features. These features were filtered based on maximum Hessian score (excluding extreme values) within their neighbourhood, defined by the NMS radius. The resulting feature distribution was consistent but

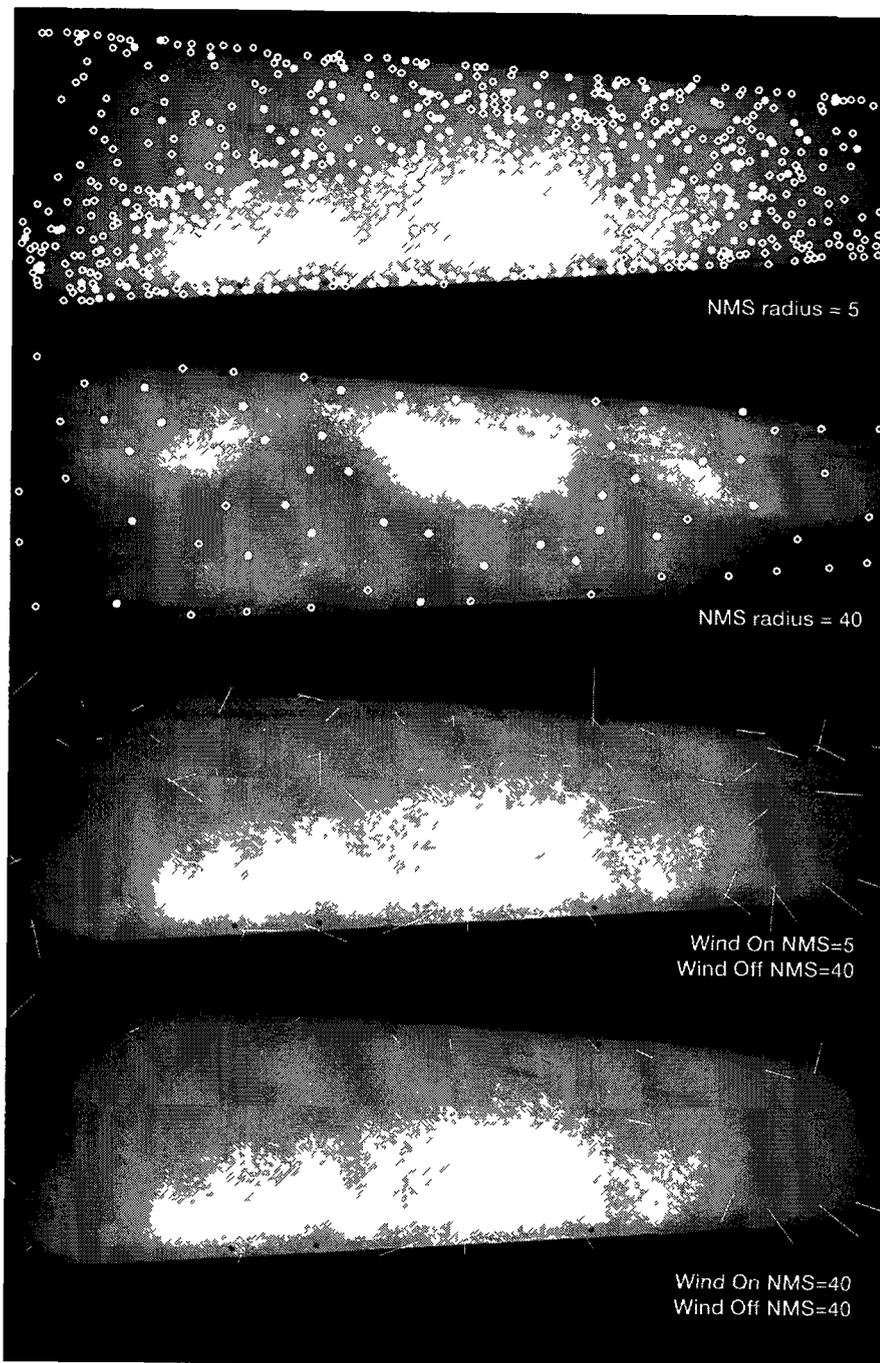


Figure 4.6: Non-maximal suppression on SURF features. The algorithm is shown to select the SURF feature with the highest Hessian determinant in its local area, as defined by a radius value. The resulting distribution is uniform without being regimented or grid-based. Images (3) and (4) show the results of matching between wind-on and wind-off images. The fraction of erroneous vectors is not decreased by increasing NMS for either image.

not constrained to an artificial grid, as shown in Figure 4.6 on page 75.

Images (3) and (4) of Figure 4.6 show two test cases. Various combinations of NMS radii in both the wind-on and wind-off images were tested, and the results were consistently in favour of a minimal NMS radius of approximately 5 pixels or less. This yielded the best combination of vector density and lower gradient field disparity.

### **Results Summary for SURF Feature Detection and Quality Measurement**

These tests determined that filtering out lower quality SURF features did not improve the quality of the resulting motion field density or accuracy, and in fact had the opposite effect. The results indicate that the Hessian score of a given matched feature does not strongly correlate with the accuracy of the resulting motion vector, and that generating the maximum number of features for matching and subsequent filtering is desirable in the PSP application.

### **Effect of Varying Descriptor Error Threshold**

Once testing established the subset of SURF features that were best used for matching, the matching constraints developed specifically for the PSP application were tested. The first constraint parameter that was tested was the maximum allowable Euclidean distance between valid matches. This value set a limit on how much error there could be between any two SURF feature descriptors before they were discarded as a potential match. Figure 4.7 on page 77 plots the fraction of features matched as the maximum allowable error is increased for a wind tunnel image pair.

The matched fraction of features increases with the allowable error until all

features are matched. As allowable error approached 1.0 many of the new matches were erroneous, however there was no indication in the trend that further increase in the matched fraction was due to false matching. Reducing the maximum allowable error did not reduce the average disparity of the vector set, which could have been used to establish a threshold beyond which most matches were false. Instead, the average disparity remained relatively constant, and in fact increased at lower error thresholds primarily because the number of vectors was very low, resulting in a small sample population for disparity average calculation. As the number of threshold was increased, match count increased, however there was not a significant change in average gradient field disparity. Coverage of the image by motion vectors did become more homogeneous as the threshold was raised towards 0.1, especially in areas of high PSP reflectivity change.

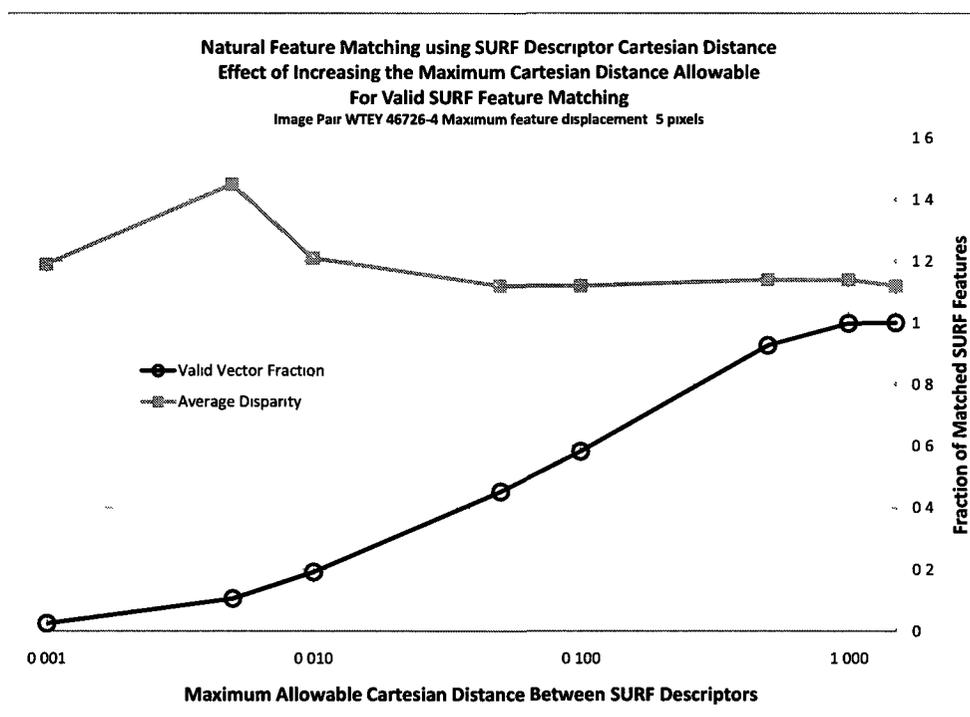


Figure 4.7: Number of valid SURF feature matches between wind-off and wind-on images as maximum descriptor error threshold is increased.

In summary, this lack of a clear transition point between primarily valid and primarily false matches indicated that additional filtering of SURF matches was required and that additional constraints were desirable in feature matching. A typical matching threshold of approximately 0.1 yielded the best balance between matched feature coverage of the image and fraction of false matches.

### Effect of Varying Feature Displacement Threshold

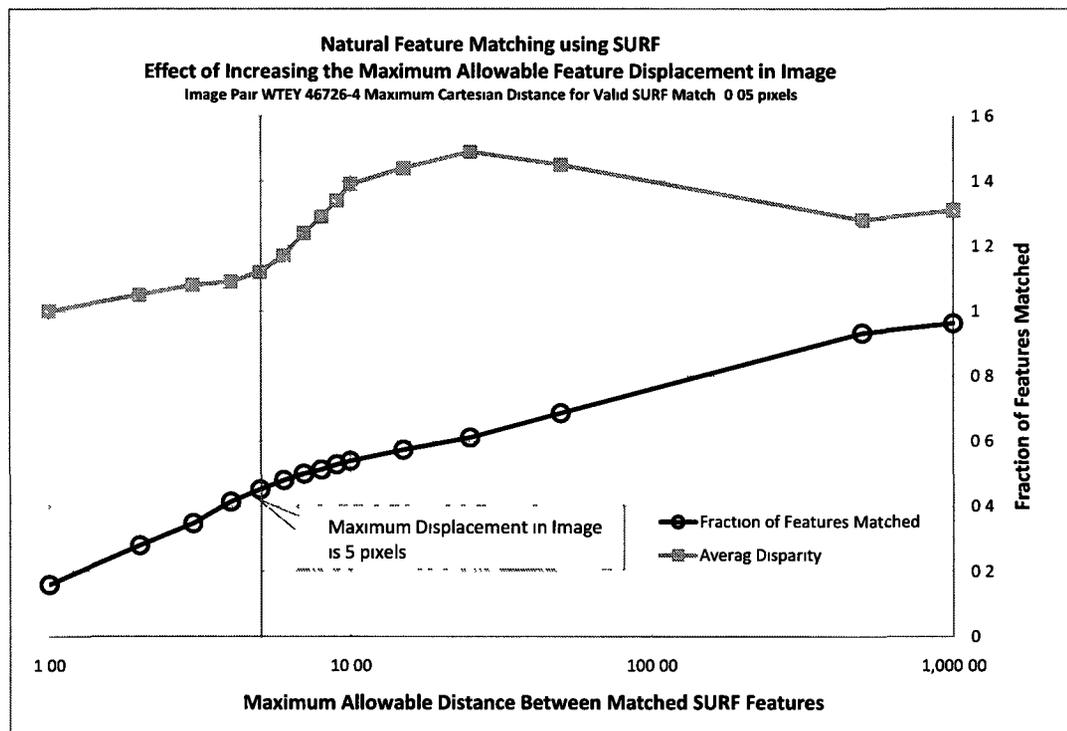


Figure 4.8: Number of matched SURF features as maximum allowable displacement is increased. The number of matches continues to increase even past the maximum actual displacement in the image, indicating additional invalid matches beyond 5 pixels. Selecting a radius that is slightly greater than the maximum expected displacement or distortion is critical, however an overly large value can result in mismatches due to self similarity of the PSP coating.

---

An additional matching constraint was developed and tested which limited the maximum displacement of any SURF feature pair between the wind-off and wind-on images. This constraint is viable in the PSP application since all model motion and deformation is expected to be relatively small in comparison with the overall image dimensions. Figure 4.8 shows the resulting fraction of features matched as the maximum allowable feature displacement is increased. It was determined before the test that the maximum displacement of any part of the model was no more than 5 pixels, and this point is labeled in the plot. Any increase in the matched fraction of features beyond this displacement were erroneous matches due to self-similarity of the PSP texture. An exponential relationship was noted, and was expected since the search area for a matching feature increased with the square of allowable maximum search radius.

Unlike the relatively constant average disparity in the previous plot, a clear increase is seen immediately after the actual maximum feature displacement is exceeded. This is confirmation that new matches are increasing the average disparity and are most likely false matches. This result indicates that placing an upper limit on the displacement allowed between matched features is a viable constraint that can reduce false matches. This maximum displacement can however vary across the image, for example from wing root to wing tip, and setting a single maximum value across the whole image may be too lenient in some areas resulting in increased false matches, indicating that additional filtering of matches is required.

In summary, the results showed that assuring a high proportion of valid matches required setting the maximum displacement parameter just above the maximum observed displacement between the wind-on and wind-off images. The maximum-displacement constraint was found to be an effective method of reducing the fraction of mismatches, however additional filtering for false matches is

still required as a subsequent step.

### **Discussion and Summary of SURF Feature Tracking**

Natural feature tracking using SURF with modified matching constraints was found to meet the requirement of increasing the number of valid motion vectors across the wind tunnel model when applied to several real data sets. Tracking was independent of shifts in contrast and brightness due to PSP reflectivity, and was tolerant of the limited changes in feature distortion due to wind tunnel model bending, distortion, and motion.

As expected, SURF tracking was not as consistent or accurate as fiducial marker tracking. Specifically, relocation of the same feature from wind-off to wind-on could be in error; however these errors tended to be large and readily filtered as outliers. Tuning the algorithm parameters to remove sufficient outliers resulted in a motion field that matched closely to that of the fiducial marker tracking. Synthetic image testing showed that absolute accuracy of feature tracking is approximately 0.2px which is approximately 10 times less accurate than typical fiducial marker tracking. This accuracy is sufficient for this application, therefore it was critical to strike a balance between sufficient image coverage and sufficiently stringent outlier removal.

SURF tracking vector density was reduced in areas of high local contrast change such as shock fronts. This is not necessarily of concern since SURF features on either side of a shock front are still tracked successfully, implying that a resection mesh will include triangles which will simply cross over the shock front, allowing resection and un-distortion to proceed.

It was found that wind tunnel images where illumination decreased to-

---

wards the edges of the frame, tracking vector density decreased. This is attributed to the reduced contrast in these areas. Additional tests conducted to flatten the illumination field in post processing showed an increase in number of tracking features near the edges of the frame; however the most effective method would be assuring that the illumination field during testing, since post processing can increase noise levels and operates on a reduced image intensity dynamic range.

A non-homogeneous PSP paint application may be of benefit for natural feature tracking. Since the image ratio process at the end of the PSP work flow removes artefacts due to inconsistent reflectivity of PSP, it follows that inconsistencies could be deliberately introduced into the PSP application to facilitate natural feature tracking. A method similar to this was explored by Park [2005]<sup>45</sup> and Sung [2005],<sup>46</sup> where a speckle pattern was applied on top of the PSP to facilitate dense optical flow tracking. The disadvantage of that solution was the necessity to filter out the speckle signal from the final PSP data and the requirement to apply the coating as an additional step. Several data sets had PSP coating that varied enough in local intensity to allow SURF feature tracking, demonstrating that methods could be developed to intentionally vary the coating reflectivity without resorting to applying a specific pattern.

SURF tracking was initially hypothesized to work best with high quality features, as measured by the Hessian determinant. This metric is suitable for applications with high contrast, distinct features, however the PSP application deals primarily with lower contrast, non-distinct features; a wind tunnel model has significant self-similarity when considered as sub-images. Based on this hypothesis a non-maximal-suppression algorithm was implemented to pre-filter the natural features detected by their Hessian score before matching wind-off and wind-on images. It was found that a given feature's Hessian score can vary significantly be-

tween the images, potentially reducing the number of available features to match between images. Better coverage of the image with motion vectors was achieved by allowing the Hessian to be very low, and instead using a combination of maximum feature displacement and disparity gradient filter constraints to maximize the number of valid motion vectors.

In summary, SURF feature identification and tracking met the requirement for a natural feature tracking system. Tracking was found to be robust under PSP imaging conditions and more features were tracked than with fiducial markers alone. However, additional filtering to remove outliers was found to be required as a subsequent step to feature matching.

#### **4.4.2 Disparity Gradient Filter Testing on Synthetic Imagery**

In the previous section matching accuracy and constraints to limit mismatches were examined. It was found that erroneous motion vectors tended to be a significant fraction of matches. It is therefore necessary to apply a constraint that acts on the motion vector set as a whole. As described in Chapter 3, disparity gradient filter was chosen as a potential solution because of its characteristics in removing local outliers while allowing the overall motion vector field (i.e. "gradient field") to vary. This closely models the expected behaviour of a rigid wind tunnel model that exhibits some continuous, non-linear deformation and perspective change.

##### **Motion Tracking with Synthetic PSP Texture**

Synthetic imagery provided a tool for testing feature tracking accuracy against exact known displacements and motions. Initial testing on synthetic images was used to verify sub-pixel accuracy of SURF feature tracking with imagery having

no distinct edges or corners.

### SURF Tracking on Synthetic PSP Pattern

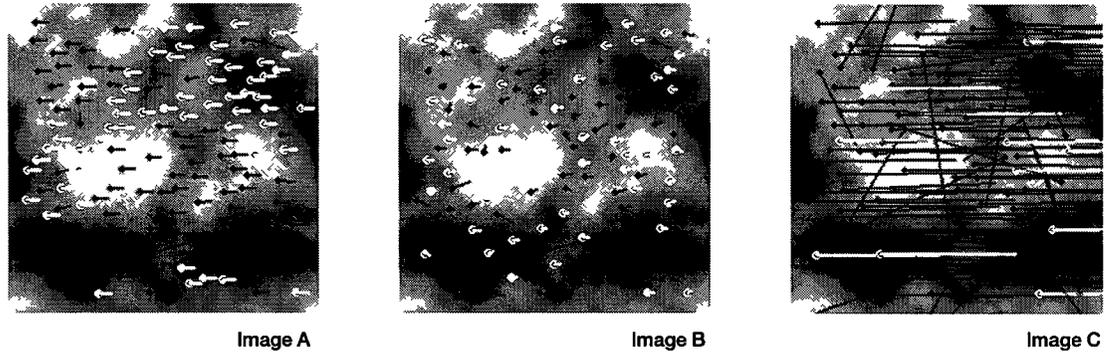


	Image A	Image B	Image C
Actual Image Displacement (pixels)	+1.000	+0.500	+10.500
Measured Average Displacement (pixels)	1.000	+0.490	+10.486
Error (pixels)	0.000	0.010	0.014
Standard Deviation (pixels)	0.000	0.056	0.021

**Synthetic Image Details.**  
 Grayscale Perlin noise,  
 Supersampled at 200%  
 Scaled using bilinear interpolation

**Tracking and Filtering Parameters**  
 0.5 maximum error between SURF descriptors  
 5.0 pixel maximum feature displacement  
 10 pixel NMS radius (wind on and wind-off image)  
 1.25 maximum vector deviation in disparity gradient

Figure 4.9: SURF feature tracking on synthetic PSP texture generated using super-sampling and bilinear interpolation to minimize synthetic artefacts. Tracking accuracy is independent of feature displacement, although mismatches increase due to the self-similar pattern of the Perlin noise.

Images of Perlin noise were used to simulate PSP patterns seen on wind tunnel test models. These Perlin noise images were generated at double resolution and translated by odd numbers of pixels before being downsized using bilinear interpolation. The resulting images were then tracked using SURF feature detection followed by disparity gradient filter as seen in Figure 4.9 on page 83.

Image displacements of 0.5 pixels and 10.5 pixels were tracked using SURF with DGF, resulting in an average tracking error of 0.02 pixel or better, independent of the displacement distance of the features. The larger translation distance at 10.5

---

pixels resulted in more feature mismatches; however these were effectively filtered by the disparity gradient filter. This increase in matching errors can be attributed to the larger image area over which a feature could be matched, increasing the probability of a false match due to the self-similarity of the Perlin noise pattern. This indicated that motion vector gradient filtering is a necessary constraint in determining accurate model motion.

The reduction in matches as displacement increases indicates that a large set of initial motion tracking vectors is desirable. Vector density across the image was increased by reducing the NMS radius to approximately 10 pixels. Although lowering the NMS radius potentially results in less distinct SURF feature descriptors and an increased chance of mismatches, it was found that gradient field had improved image coverage after tracking and filtering.

In summary, this test of SURF tracking accuracy on simulated PSP texture demonstrated high tracking accuracy is possible despite the lack of high quality features. Larger feature displacements resulted in more mismatches, which were effectively handled by subsequent disparity gradient filtering. This indicates that SURF in conjunction with DGF can provide accurate motion data without explicit tracking markers.

### **Comparison to Pyramidal L-K Optical Flow Tracking**

Tracking of motion in a flat translational plane is a subset of the types of motion expected in wind tunnel model deformation and perspective changes. In order to extend testing of SURF feature tracking and disparity gradient filter performance to a more complex motion environment, the standard Yosemite image sequence was used.<sup>29</sup> This synthetic image sequence of a fly-over of terrain is typically used to

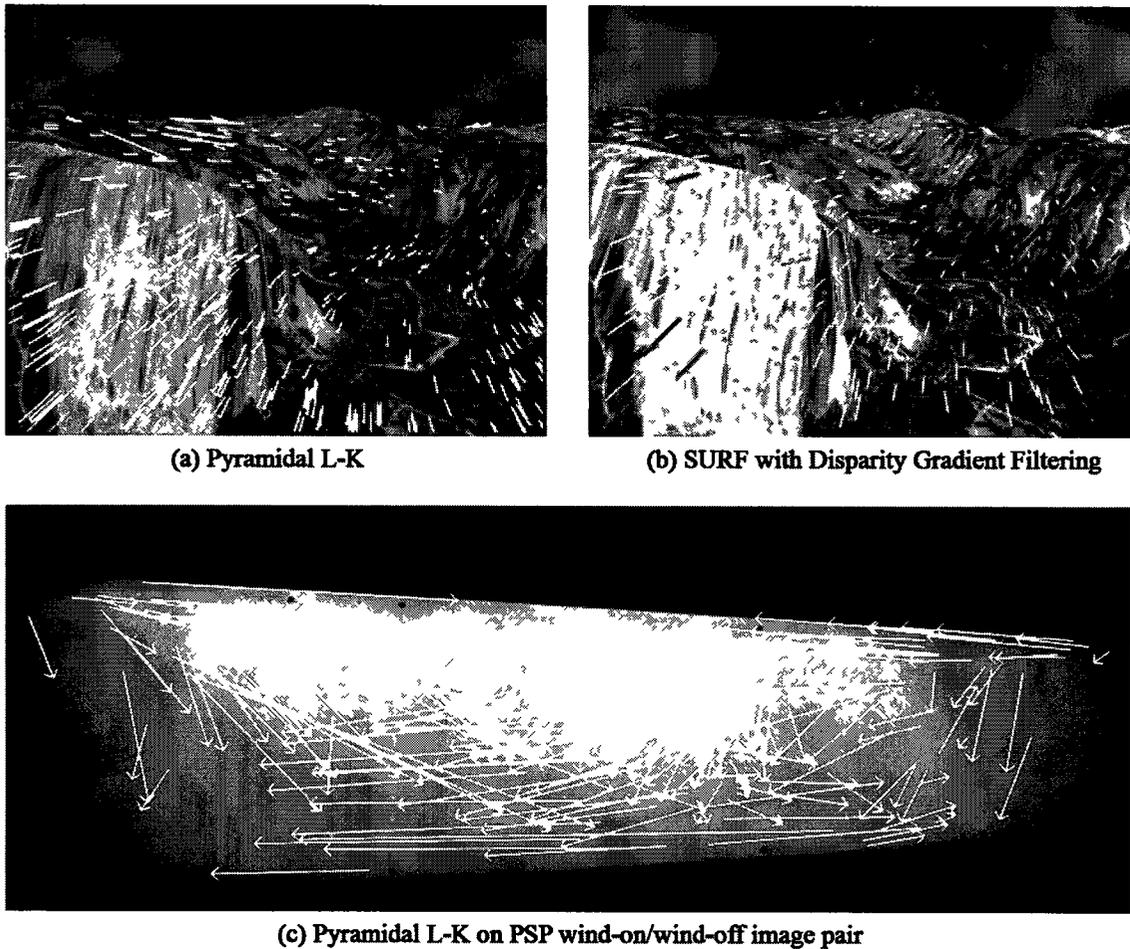


Figure 4.10: Motion tracking of the Yosemite sequence, frames 0 and 3. Pyramidal L-K performance is shown in (a) with good tracking of perspective motion and rejection of the random sky noise pattern. SURF feature tracking combined with disparity gradient filter is shown in (b) with filtered outlier vectors darkened. Tracking and rejection performance is similar to Pyramidal L-K. Image (c) shows Pyramidal L-K applied to a PSP wind tunnel image pair. Variations in PSP image intensity cause significant tracking error.

---

compare the performance of dense optical flow algorithms, as described in Chapter 3. Motions include perspective shift and parallax as well as extraneous random noise in the sky portion of the image. A large image displacement was chosen, specifically between images 0 and 3 which would best reflect the displacements seen in wind tunnel imagery. The results can be seen in Figure 4.10 on page 85. This also offered the opportunity to compare SURF and DGF performance to that of an established motion tracking algorithm, specifically Pyramidal L-K.

Results showed that both algorithms identified and tracked features across the entire landscape portion of the image, and both rejected extraneous motion vectors due to the random cloud region. The SURF and DGF successfully filtered erroneous motion vectors due to mismatches, shown as darker vectors in image (b). Identification and tracking of valid motion were comparable for both algorithms. Total time for calculating motion was under 0.5 seconds for SURF and DGF. This demonstrated that the combination of SURF feature tracking and disparity gradient filter filtering, when combined, could track nonlinear motions with discontinuities, such as those due to perspective changes combined with a noise region with similar performance to a combined global/local dense optical flow algorithm.

Because the performance of synthetic data was comparable, a test on PSP data was carried out to note any differences that might appear in this application. Figure (c) shows as demonstration the performance of the same Pyramidal L-K algorithm on wind tunnel data. When tested on a wind-on/wind-off image pair, this motion tracking algorithm was unable to distinguish between the motion of the model and variations in local intensity due changes in the PSP reflectivity. The calculated motion vector field was mainly random or erroneous, indicating that conventional dense field motion tracking is likely not a suitable candidate for the

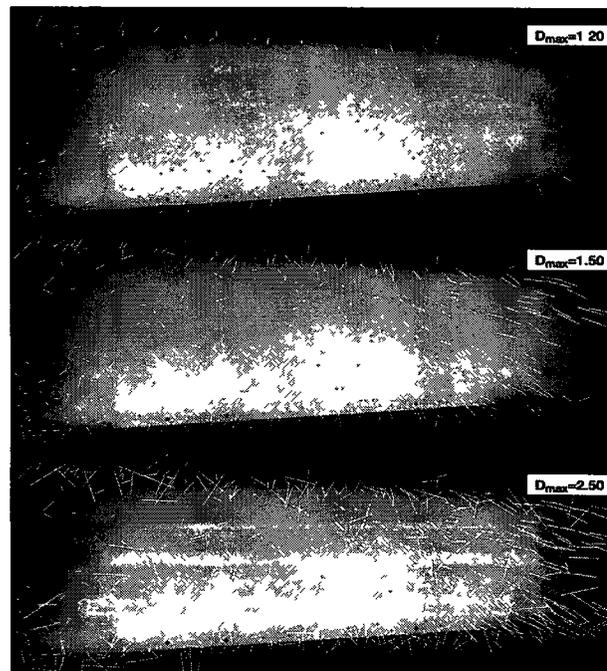
PSP application.

In summary, this test demonstrated that SURF and DGF had similar motion tracking performance to an established dense optical flow motion tracking algorithm, successfully discriminating perspective motions from noise. The test also showed that the same dense optical flow technique is not suitable for PSP image motion tracking, indicating that SURF and DGF is a preferable algorithm for this application.

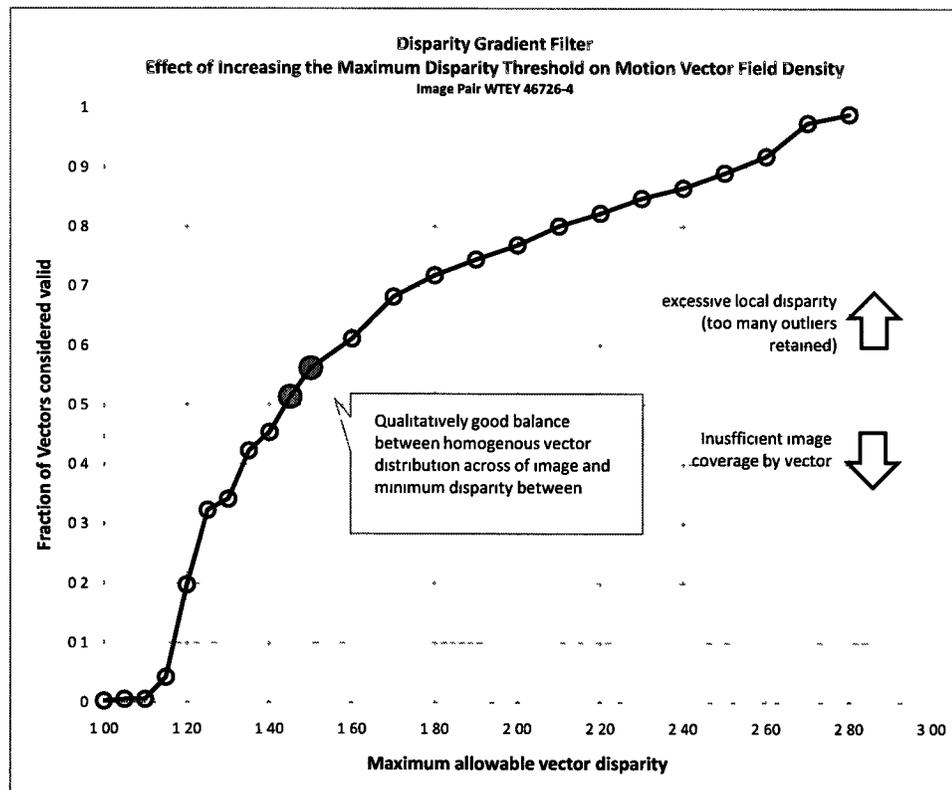
#### **4.4.3 Testing SURF with DGF on Wind Tunnel Imagery**

The goal of SURF with DGF is to produce an accurate motion gradient field over the surface of the model with sufficient density and coverage. What is 'sufficient' and 'accurate' is partially subjective when only using a single source of motion data, such as SURF feature tracking. For this reason the effect of adjusting the maximum allowable disparity threshold ( $D_{max}$ ) of the DGF was tested, as it directly affects the sub-set of motion vectors retained for subsequent PLR workflow. Figure 4.11 on page 88 shows the effect of adjusting this threshold parameter on the final motion gradient field. The number of vectors filtered as outliers drops as the threshold was increased, however there was no clear indication in this trend that would indicate at which setting the highest fraction of erroneous vectors was filtered out. This indicates that the filter setting was partially dependent on requirements of the user and the specific image set. This result was the motivation for developing a weighting and prioritization extension of the disparity gradient filter algorithm, as described later in this chapter.

In summary, this test showed that the optimal filtering strength of the DGF is at least partially dependent on the users requirements and on the specific wind



(a)



(b)

Figure 4.11: Effect of increasing the disparity gradient filter's maximum disparity threshold ( $D_{max}$ ). Image series (a) shows an increase in the number of retained vectors at three  $D_{max}$  values. The fraction of the initial gradient vector field that is retained as  $D_{max}$  is increased shows a nonlinear characteristic, as plotted in (b).

tunnel being analyzed. In practical application it is straightforward to quickly select a  $D_{max}$  which balances motion vector density with erroneous outliers, but because model motion and deformation are not known, the definition of 'erroneous' is based on the user's experience and expertise. This tool provides a useful tuning parameter to adjust filtering based on user experience.

### **Effect of Model Motion and PSP Intensity Changes**

A test was conducted in order to determine model motion tracking performance independently from PSP intensity changes. This test was conducted using by tracking motion between two wind-off images. These images showed different model positions relative to the camera, specifically a different rotational position along the horizontal axis. Projected onto the 2D image plane, model motion was nonlinear, and included perspective changes across the 3D surface, distortion from the camera system and changes in illumination due to rotation of the wing with respect to the illumination source. The results are shown in Figure 4.12(a) on page 90.

With SURF and DGF parameters coarsely adjusted for the image set, features were tracked across the whole model image and the filtered motion gradient field agreed qualitatively with the fiducial marker motions. Even with large displacements (i.e. large changes in aspect resulting in occlusion of some portion of the wing), the number of filtered features still remains large enough for this application. The SURF motion vectors clearly showed additional motion information in areas not covered by fiducial markers and showed perspective changes due to the curvature of the wing surface. This demonstrated that SURF and DGF can track 3D model motion projected into 2D images, in this case a type of motion which

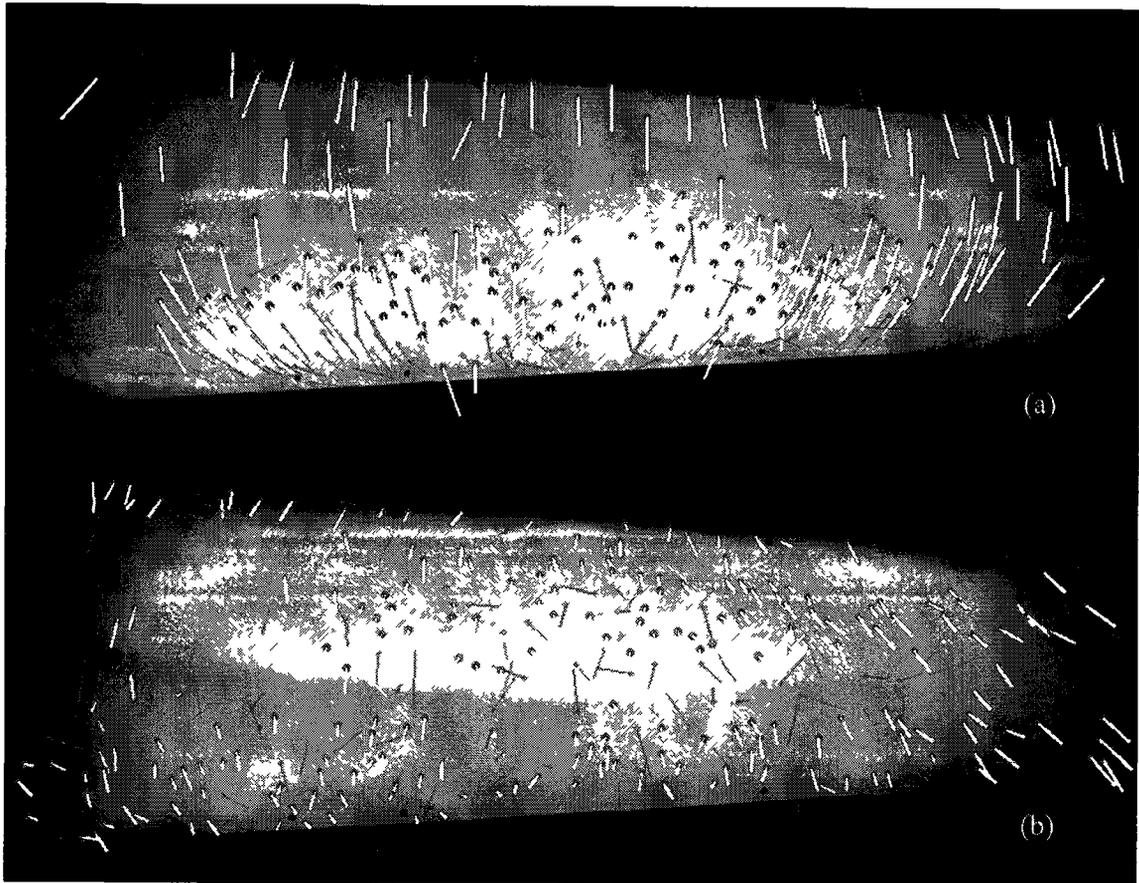


Figure 4.12: SURF feature tracking of wind tunnel under rotation. Darker vectors have been filtered out of the final motion gradient. Figure (a) shows feature tracking between two wind-off images rotated with respect to one-another, showing motions of the 3D wing about its horizontal axis. The PSP texture is readily tracked. Figure (b) uses a mismatched wind-on and wind-off image pair to show tracking with a combination of a smaller 3D motion as in (a) as well as tracking with local PSP intensity change and local deformation of the right wingtip. The resulting motion gradient shows tracking of both the smaller rotation about the horizontal axis as well as the wing tip deformation towards the camera.

---

was locally consistent but globally variant.

A test using the same data set with wind-off and wind-on image pairs was conducted to determine performance in the presence of local changes in PSP reflectivity. The image pair included wing rotation, non-linear wingtip deformation on the right side of the image, and PSP intensity changes, including a significant transonic discontinuity. The results are shown in Figure 4.12(b) on page 90. In this case, the motion of the wing rotation was significantly less than image (a) and the resulting motion vectors across the wing have a lower magnitude. Because of the filtering characteristic of the DGF, these vectors appear to have a higher directional variance than in (a), however their disparity with respect to the motion field median was below the specified threshold. At the right of the image wing tip bending was detected and filtered successfully by SURF and DGF. The PSP coating exhibited significant intensity changes across the model surface; however a sufficient number of motion vectors were retained, with the exception of a local discontinuity running horizontally along the wing span. This discontinuity was due to a transonic shock front that was highly localized, reducing the feature matching rate in this local area. However features on either side of this discontinuity were successfully matched and tracked, indicating that PLR work flow will still be able to create a dense mesh that effectively bridges this discontinuity.

In summary, this test shows that SURF and DGF can track model motions due to rotations towards and away from the camera, as well as nonlinear model deformations. Tracking coverage was shown to remain acceptable even when significant, localized PSP intensity changes are present, such as those found near shock fronts.

---

## Tracking Performance on Different Wind Tunnel Models

Several image sets were compared with varying amounts displacement distortion, as shown in Figure 4.13 on page 93. In all these tests only SURF features were tracked, not fiducial markers. These tests were used to evaluate tracking density in a range of conditions. In all tests, the wind-off image is shown under the motion vectors for texture reference.

The "WTEY" model deformation is primarily a wing tip bend towards the camera on the right of the image. This test set images also showed changes in the PSP pattern due to transonic flow shocks and stall patterns, both of which result in significant changes in local and global PSP reflectivity. These changes in PSP pattern resulted in standard optical flow algorithms failing, as shown in the previous section. There was no prior knowledge of the type of motion, illumination changes, or model distortion that would be present in these images sets; the software package was tasked with automatically identifying these issues and resecting the wind-on image to match the wind-off image with only minimal parameter tuning for feature detection and motion vector filtering tolerance.

Two delta-wing models with less local distortion and more fiducial markers were also tested to show the flexibility of the tracking work flow. There was less deformation because of the model geometry, but significant global motion due to changes in angle of attack of the models. SURF tended to track the outer edges of the model and areas which combined features from model seams and pressure taps. The filtered motion field shows the change in angle of attack towards the camera, with the nose of the model having more motion. As before, there was no a priori knowledge of the model motion and distortion.

Finally, tracking performance using a 2D wing section was evaluated. This

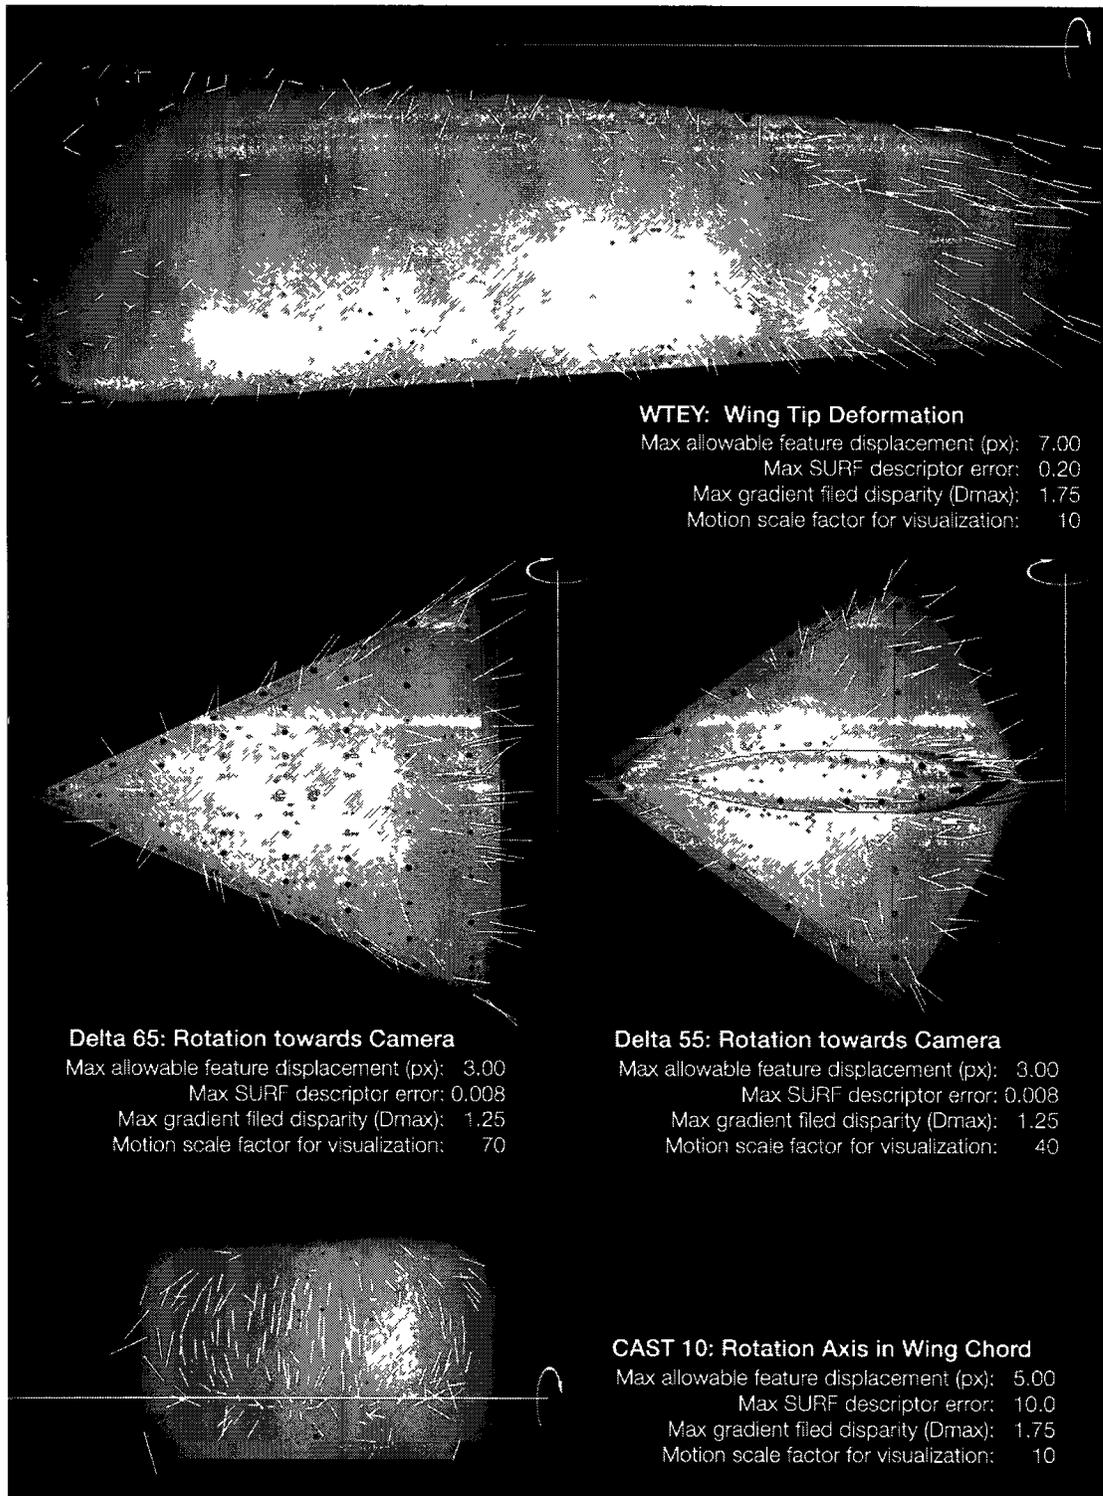


Figure 4.13: Tracking of wing bend and motion for a range of PSP-coated wind tunnel models with SURF and DGF. Approximate rotation axis is indicated with the types of motions. Motion gradient field vectors show detected motion and are scaled by varying amounts for visibility; motions range from sub-pixel to more than 7 pixels.

section underwent rotation about an axis located on the section, creating a local motion discontinuity through the rotation axis. SURF and DGF tracked the motion successfully, removing outliers while maintaining tracking coverage across the entire section. This demonstrated the flexibility of the tracking algorithm near areas of discontinuous motion.

In summary, these tests demonstrated that SURF and DGF could track a range of model motions and deformations using wind tunnel imagery from several different test scenarios. Motions including global rotation and local bending were tracked, and motion vector coverage of models was shown to exceed the number of fiducial markers. This result demonstrates that in comparison to fiducial marker tracking, additional tracking information can be extracted using SURF and DGF.

#### **4.4.4 Testing of Weighted, Prioritized DGF With Fiducial Marker Data**

The disparity gradient filter was extended to allow assignment of weighting factors and priorities to subsets of the gradient field. This provided a tool for combining flow vectors from multiple sources. The system of weighting and prioritization of vector sub-sets described in Chapter 3 was tested using wind tunnel model imagery.

The goal of the test was to determine how vector weighting works in conjunction with the disparity threshold. It was hypothesized that increasing the weighting of more accurate flow vectors would smooth the gradient field locally around the weighted vectors, acting as a localized decrease in disparity gradient threshold. This would effectively improve resection accuracy in the local area surrounding more accurate motion vectors.

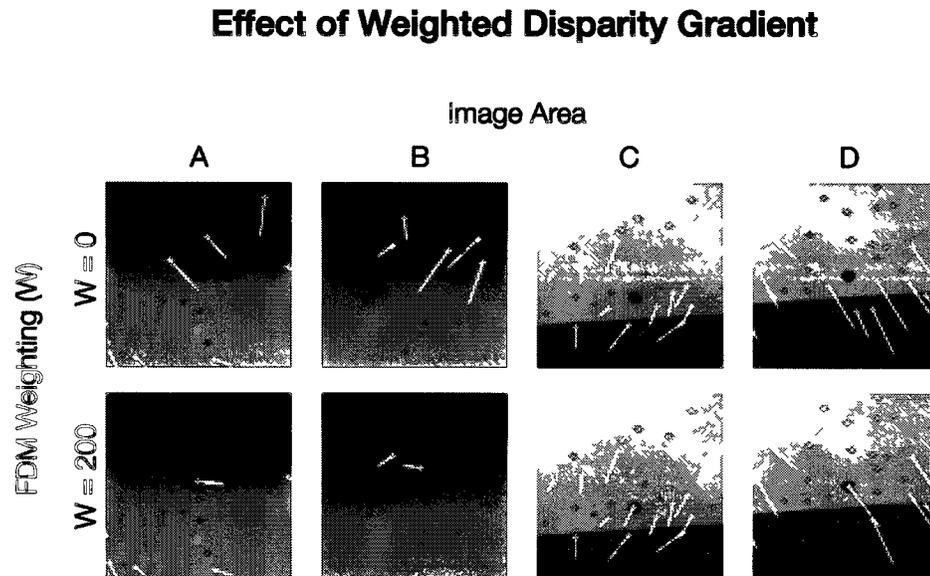


Figure 4.14: Local decrease in outlier tolerance near weighted vectors. Four sample regions of an image centered on fiducial markers vectors weighted relative to surrounding natural feature vectors. The top row shows a weighting of 0 (no fiducial marker vectors) for 4 regions and the bottom row shows a weighting of 200 for those same regions. Global  $D_{max}$  is adjusted to maintain the same gradient field vector count for the entire image. The local area around the fiducial marker motion vector has significantly reduced disparity tolerance, and as a result a higher proportion of surrounding vectors classified as outliers and removed.

In these tests, the priority value for the fiducial marker motion vectors was set to zero, meaning they could not be filtered out of the vector set by the disparity gradient algorithm. Their weight relative to SURF-based motion vectors was then varied across a range larger than what would be expected in practical application in order to test extremes. Figure 4.15 on page 96 shows the result of increasing fiducial marker vector weighting while holding maximum disparity constant at three different values.

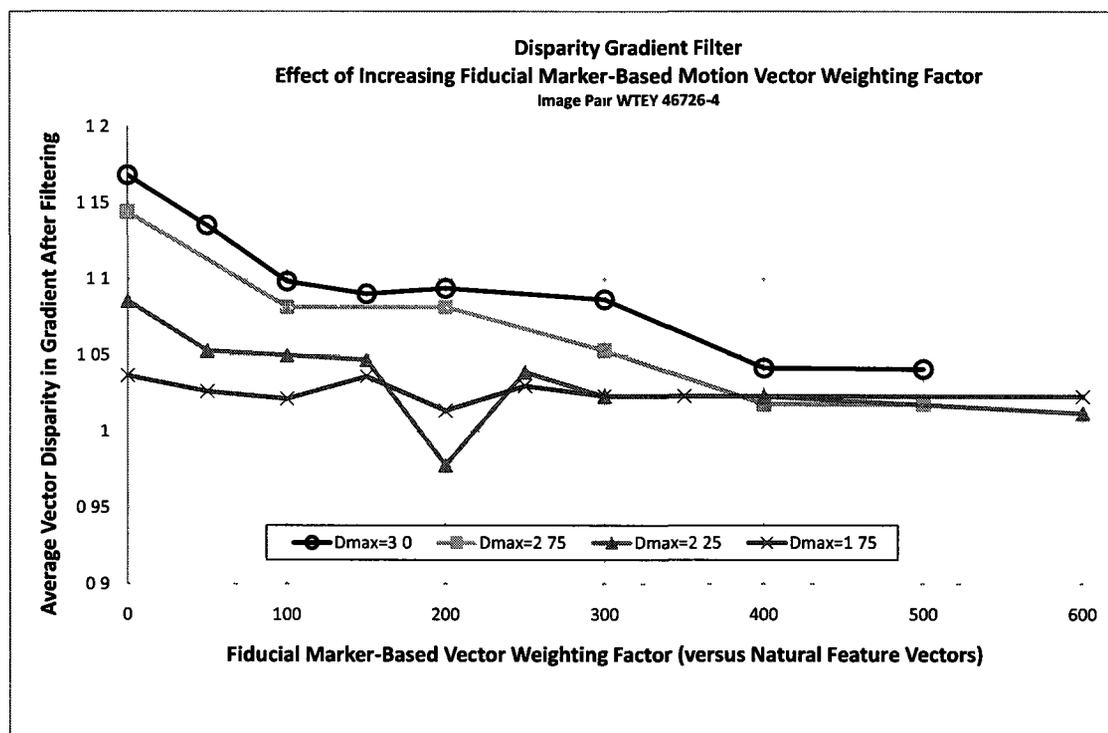


Figure 4.15: Effect of increasing fiducial marker-based vector weight relative to natural features. Increasing vector weight in turn increases the number of natural feature vectors classified as outliers in the vicinity of fiducial markers, reducing overall field disparity.

The effect of changing weighting was determined by observing the max-

imum disparity in the filtered gradient field. Tests were conducted at different Dmax settings to determine how Dmax and Weighting were coupled. Increasing the weighting factor was found to decrease the disparity of the overall field, indicating that at even small weighting factors, the effect of filtering was increased across the image. The similarity of the plots show that increasing Dmax or the weighting factor have similar effect if the fiducial marker distribution was homogeneous in the image.

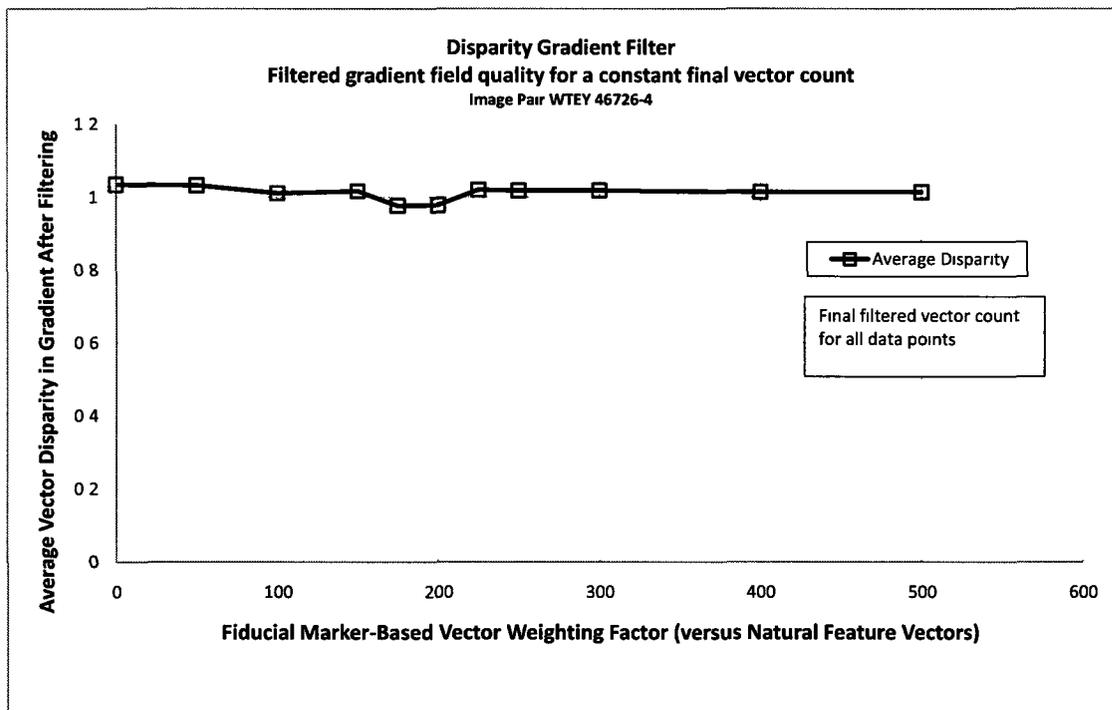


Figure 4.16: Comparing the effect of Dmax to weighting for a given filtering effect. Vector weighting is increased while in coordination with Dmax such that the total number of vectors in the filtered field is constant, effectively shifting the filtering role from weighting to global disparity threshold. The average disparity of the gradient field remains constant for this image set, indicating that the two methods have similar effect when the fiducial marker distribution is homogeneous.

In all cases a minimum disparity value was reached after which increasing weighting has little effect. This minimum can be interpreted as the disparity between the fiducial marker vectors since these vectors will not be in agreement and cannot remove one another from the field; the disparity of the overall field cannot be improved beyond the mutual disparity of these vectors. The overall disparity of the field was based on this and  $D_{max}$ , so as  $D_{max}$  was decreased, this minimum was also decreased.

These results indicate that a balanced approach to local fiducial marker weighting and global maximum disparity should be taken to get the best filtering performance.

Further testing was done to determine which of these filtering methods was more effective at creating a high density gradient field with low disparity. Figure 4.16 shows the average disparity of the gradient field for combinations of maximum disparity threshold and weighting factor which filtered the same number of vectors.

If one method were preferable to the other throughout the image there would be a variation in average disparity. This was not shown to be the case, with average disparity staying relatively constant, indicating that at the global image level, the two methods yield the same quality of gradient field, given a relatively homogeneous fiducial marker distribution.

On visual inspection of the gradient fields, it was noted that more vectors were classified as outliers in the local areas around the heavily-weighted vectors even though the final vector count in the field was held constant. The local area around several fiducial markers is shown in Figure 4.14 on page 95, where the overall vector count between the two tests was kept the same. This demonstrated

that weighted vectors have the effect of locally decreasing the disparity threshold.

### **Discussion and Summary of Vector Weighting Test**

These tests demonstrate that weighted-vector filtering was an effective local-region complement to global disparity gradient filtering. With a distribution of fiducial markers across the whole image, as was typically the case in practical testing, the effect of increasing their relative weight can effectively act across the whole image. It was visually determined that the filtering effect was greater near the region of the weighted vectors.

Prioritization of vector subsets was found to have the effect of presenting a floor or minimum disparity that was dependent on the disparity between these prioritized vectors. This applied in the case these vectors are assigned a maximum priority, preventing them from filtering one-another. In practical applications this indicates that a balance must be struck between weighting factor and global  $D_{max}$  to obtain the best filtering of the gradient field.

In terms of practical use, it was of benefit to more heavily weight vectors with a better accuracy, since this did decrease the tolerance for erroneous vectors in their local vicinity. This had the practical effect of allowing more accurate motion vectors to have a more dominant role in the resection process in their local vicinity. Using a weighting allowed the global  $D_{max}$  to be raised, which allowed more local disparity in areas not dominated by the effect of the weighted vectors. For example, a weighting of 200 and a  $D_{max}$  2.0 had a similar filtering effect to a weighting of 1 and a  $D_{max}$  of 1.7, but with less disparity in the local areas around weighted vectors. This indicates that weighting is a practical local-effect complement to setting  $D_{max}$  globally.

In summary these tests show that weighted, prioritized DGF can improve filtering performance near tracked fiducial markers without significantly affecting the overall disparity of the gradient field. The user of the system can reduce the DGF filtering threshold,  $D_{max}$ , if the fiducial marker weighting factor is increased. In PSP applications, a starting point for tuning of these parameters is a fiducial marker weighting factor of 100 and  $D_{max}$  of 1.7.

#### **4.4.5 Testing the Effect of Image Segmentation on DGF**

The method of dividing the image into sections and applying filtering to each section independently was described in Chapter 3. The goal was to maintain a high vector count in areas with have significantly different motion magnitude or direction from adjacent areas. Without a method of separating these areas, filtering could remove valid motion vectors in areas with significantly different motions. Examples include areas of high deformation or articulation, such as wing tips and control surfaces.

Testing involved separating the wind tunnel imagery into a regular grid with varying grid density vertically and horizontally. Each grid sector was filtered independently by the weighted disparity gradient method. The hypothesis tested was that better overall gradient field coverage of the image can be maintained if areas of different movement can be separated. This would allow the local disparity threshold  $D_{max}$  in each sector to be decreased, removing more outliers.

Test results are shown in Figure 4.17 on page 101. The same image pair was tested increasing segmentation. Horizontal subsections were chosen since it was known the the right are of the image exhibited bending towards the camera, and had significantly different motions than the left side of the wing. When the

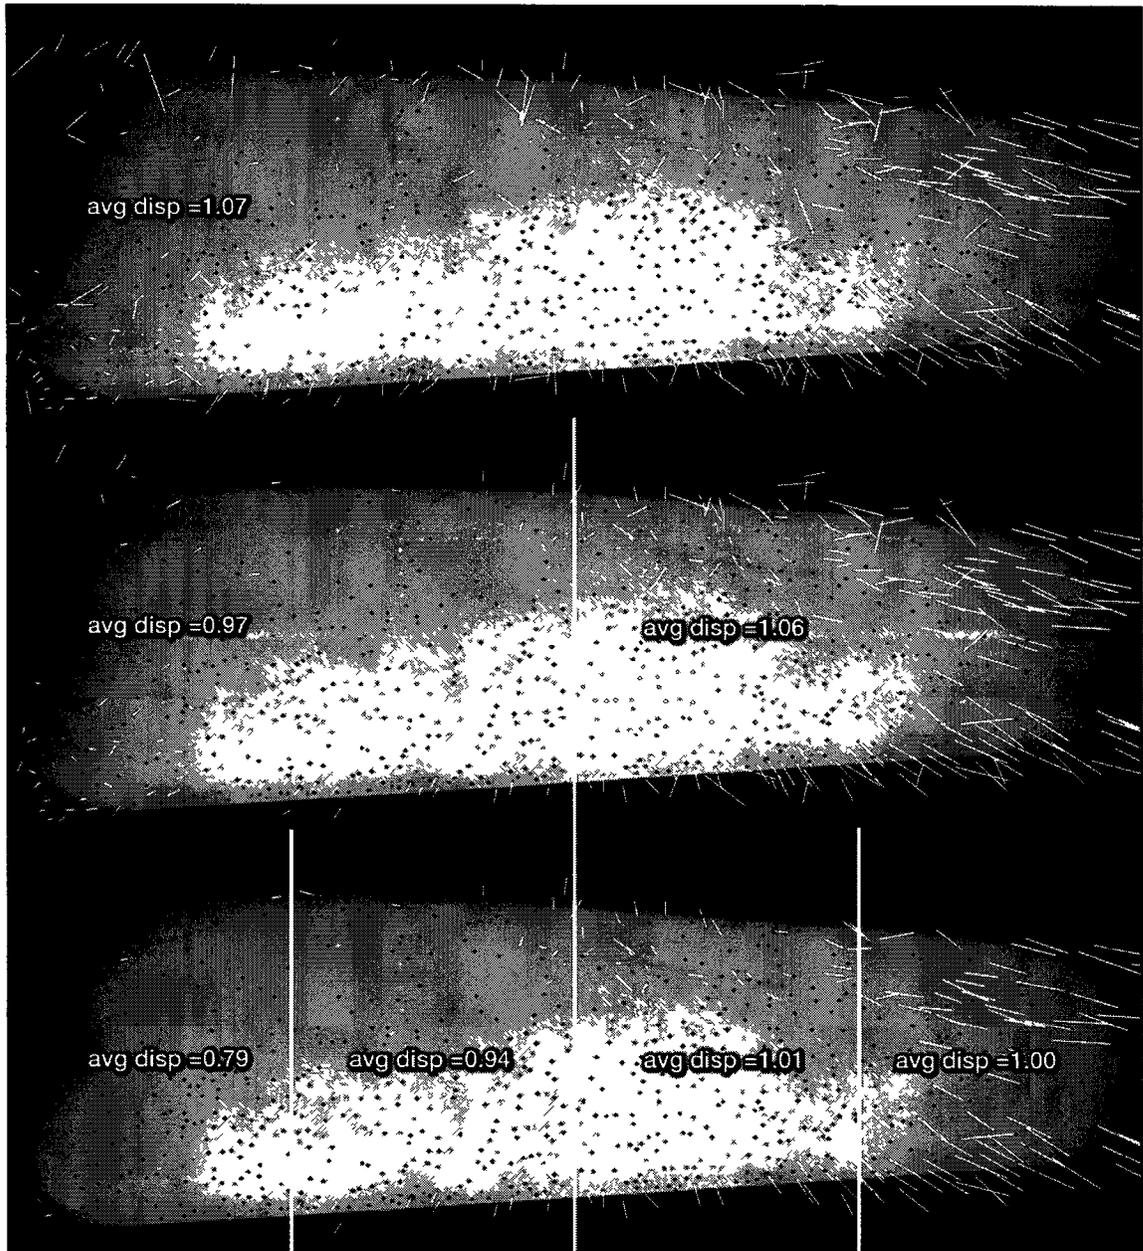


Figure 4.17: Effect of segmentation on gradient vector field disparity. For all images  $D_{max}$  and weighting are held constant at 2.5 and 100 respectively. The same image is tested with 1, 2 and 4 segments seen here from top to bottom with a visualization of the slice lines. The average disparity of each segment is labelled.

image was segmented, it was found that the average disparity within each segment decreased relative to the same area when it was unsegmented. Areas of minimal motion, where vector direction was determined primarily by tracking error, had the largest decrease in disparity. Areas with significant motion, such as the wing tip, also improved in terms of disparity, though not as significantly.

In terms of the global image, more erroneous vectors were removed while preserving vectors reflecting model deformation, improving the quality of the overall gradient field when compared to no segmentation. Segments with minimal motion had the greatest number of vectors classified as outliers; although this decreased the mesh density in these areas, a dense mesh was not necessarily desirable in areas with minimal motion, especially when the motion vectors were primarily based on tracking noise. In areas of large motion and deformation, more vectors were preserved while decreasing the local average disparity, indicating an increase in the local gradient field quality. Together, these two effects were effective in simultaneously removing more erroneous errors in stable areas of the image while improving vector field quality in areas undergoing more deformation.

### **Discussion and Summary Segmentation Test Results**

In these tests the image was segmented according to an approximation of the primary deformation motions. Segmentation performance can likely be improved further through better understanding of motions in the image, possibly through an iterative process.

It was shown that segmenting does have a practical limit beyond which the number of motion vectors within a segment is not large enough to create a significant consensus in the disparity gradient filter. Segmentation was found to be

more effective when each segment included one or two weighted fiducial marker-based vectors. Since each segment is independently filtered, a segment with only one weighted vector can be primarily influenced by that vector, whereas multiple weighted vectors in the same segment create a lower limit in segment disparity as seen in previously in Figure 4.15. Together with the results in from weighted and prioritized vectors, this indicates that a combined approach of segmentation and weighting would be effective. If the image is segmented in such a way that each segment is centered on a weighted vector, it could yield an improved overall gradient field with better discrimination of erroneous vectors and improved tolerance to large local deformations.

Wind tunnel model imagery can have regions of intentional vector field discontinuity, specifically where articulated section of the model move relative to one another. This case is an extension of what was measured in Figure /refsegmentationsample, where wing tip deformation was significant compared to the wing root. In this case, the test results indicate that segmenting the articulated area from the main area of the image can yield improved gradient field coverage with reduced disparity.

In summary, performing image segmentation before applying DGF improved local filtering performance compared to global filtering in images with model deformation. In practical application, it was determined that too many segments should be avoided, as this can prevent the filter from forming a meaningful median disparity, resulting in filtering of good vectors. If fiducial markers are heavily weighted, no more than one should be included in any segment.

## 4.5 Image Subdivision and Resection

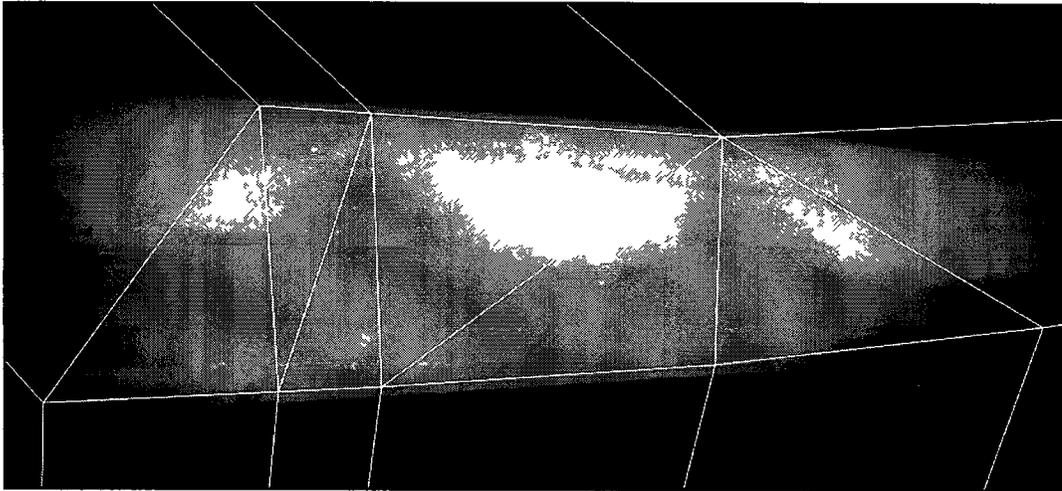
After a gradient field has been generated, the PLR work-flow must resect the wind-on image. As discussed in Chapter 3, resection involved the per-triangle un-distortion of this image to realign tracked features to their wind-off locations, resulting in reduced ratio image misregistration error due to model displacement and distortion. This section tests the two software modules which work together to use motion vector data to undistort the wind-on image: subdivision and resection.

### 4.5.1 Subdivision Using Delaunay Triangulation

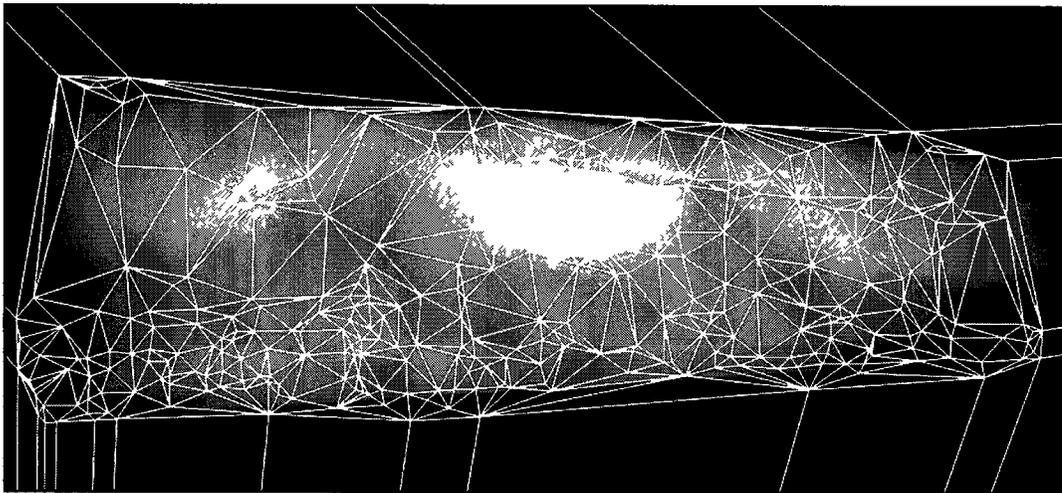
As described in Chapter 3, the motion vector field was used as a guide for un-distortion of the wind-on image. Each vector was mapped to a vertex in a triangular mesh covering the wind-on image, and these triangles are used to subdivide the image in the Resection module. The triangles are created from the vector field by constructing a Delaunay Triangulation using the endpoints of each vector. Figure 4.18 on page 105 shows the resulting triangulation on wind tunnel imagery using flow vectors from fiducial markers.

The mesh based on fiducial markers only was composed of relatively large triangles. As described in Chapter 3 the edge of the mesh between the fiducial marker and the edge of the image was subdivided based on points well outside the image, effectively appearing as parallel line segments.

Including motion vectors from natural feature tracking, the mesh became significantly denser, as seen in 4.18. Triangulation time is less than 1 second, even with several hundred vertices. In summary, this module met requirements for subdivision of the image bases on motion vector origins.



(a)



(b)

Figure 4.18: Delaunay Triangulation to create resection meshes. Figure (a) shows a coarse mesh generated from fiducial markers only and (b) shows a mesh generated from natural features. Polygon counts are 18 and 650 respectively.

## 4.5.2 Resection

As described in Chapter 3, resection subdivides the wind-on image into triangular texture patches and their vertices are shifted along their associated motion vectors, effectively un-distorting each triangle in the mesh.

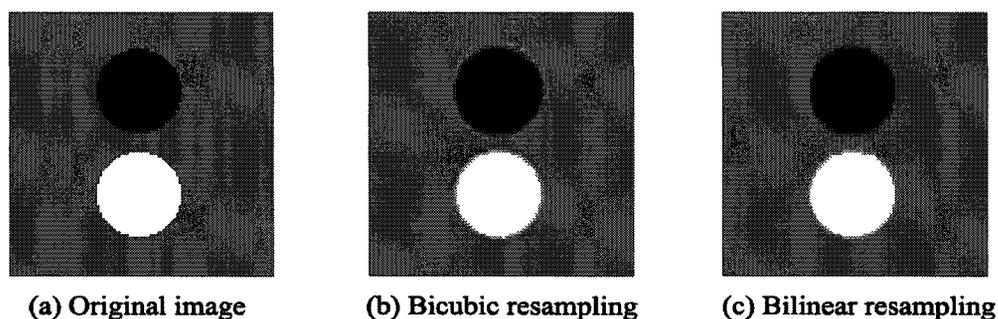


Figure 4.19: Effect of resampling algorithms. The original image in (a) is reduced then enlarged using two resampling methods. Bicubic resampling, seen in (b), introduces artificial brightness halo artefacts at contrast boundaries which could be misinterpreted as pressure data. Bilinear resampling, seen in (c) does not exhibit this characteristic.

Initial tests verified sub-pixel performance of the OpenGL routines in distorting triangles and displaying the resected image. As shown in Figure 4.20 on page 107, some blurring of the texture was evident due to bilinear resampling of the texture.

Although bilinear resampling was considered inferior to bicubic resampling for applications where the final result was evaluated by perceived image quality, it did not introduce any halo effects which would appear as additional pressure data in this application, as demonstrated in Figure 4.19.

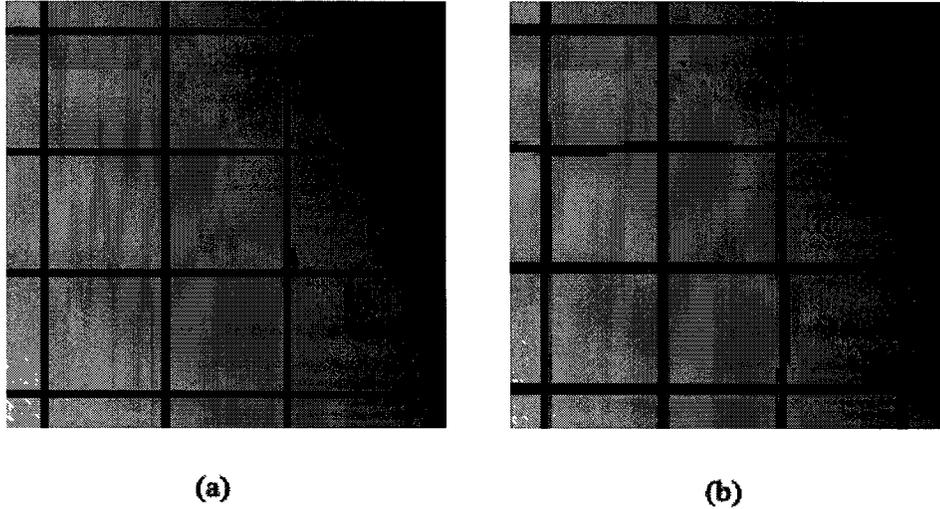


Figure 4.20: Resampling blur due to bilinear interpolation. A grid pattern is resected in the PLR work-flow using a motion gradient field generated from wind tunnel imagery. Image (a) is uncorrected, and image (b) has been resected with several polygon boundaries in this image section. Although blurring is visible due to non-integer pixel motions in parts of the image, no halo artefacts are created.

### Discussion and Summary of Resection Test

Resection module testing demonstrated that sub-pixel resampling of the undistorted mesh is feasible using OpenGL. The bilinear pixel interpolation introduced some sub-pixel blur to the image as expected, but did not introduce any significant image artefacts that could be mistaken for pressure differences, and no seams between triangles were visible. The smoothing of the image due to resampling was at the sub-pixel level, and because the image features of interest are much larger, this will not have an impact on the final ratio image. In practical application, the resampling artefacts from un-distortion were an acceptable trade-off in order to produce the more accurate resected ratio image.

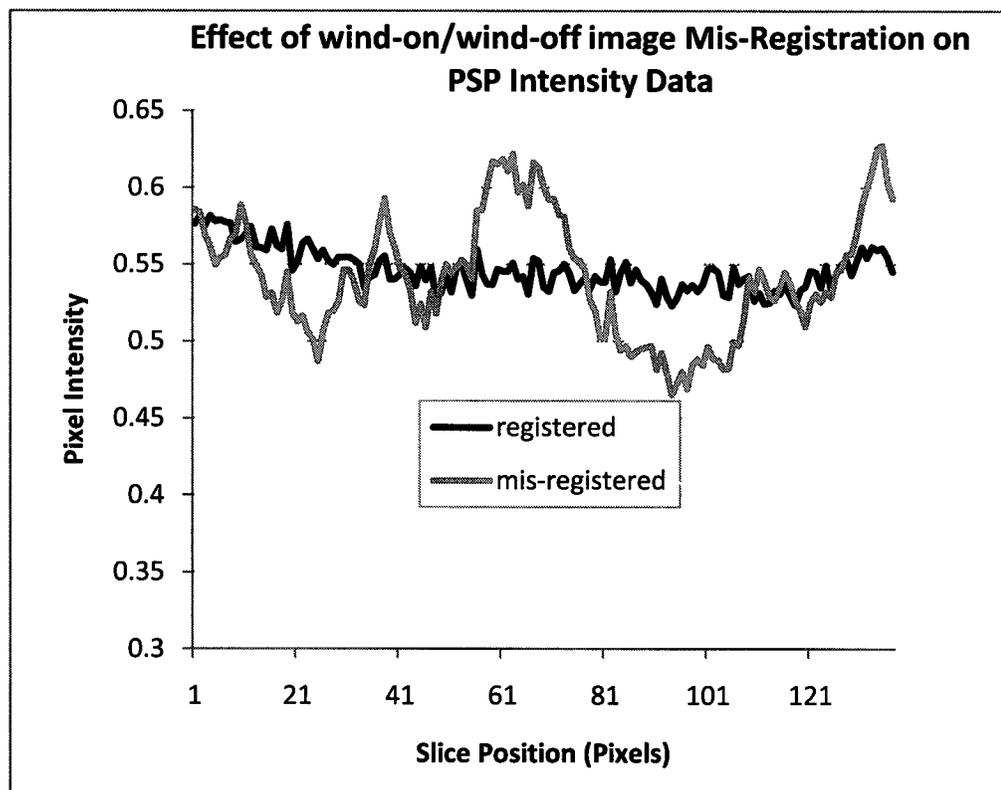
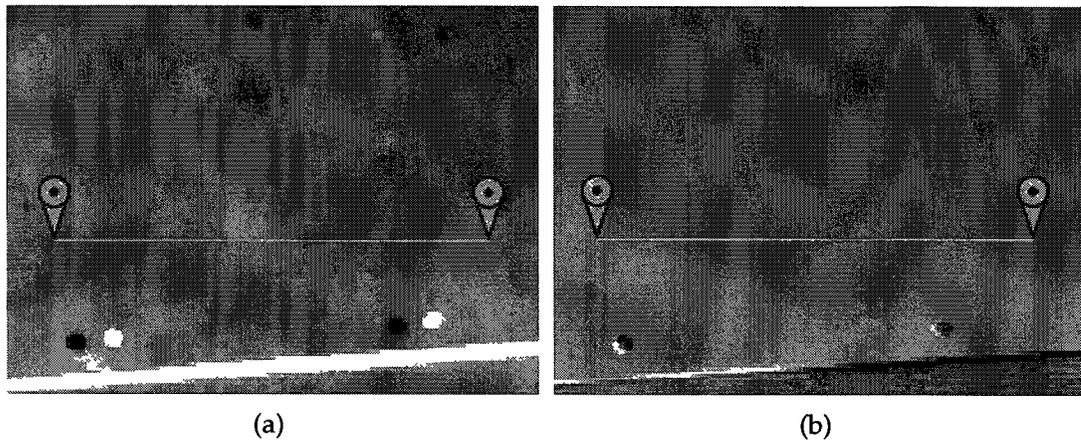
## 4.6 Resection Performance Using PLR

As described in Chapter 3, taking the ratio between the wind-off and wind-on images produces an intensity change map across the PSP-covered area of the wind tunnel model. This procedure effectively removes image intensity variation due to illumination and PSP coating non-homogeneity, producing a cleaner data. The intensity of the ratio image can then be directly correlated with a pressure value at any point in the image.

### 4.6.1 Visualization and Measurement of Misregistration

Without high-accuracy alignment between the wind-on and wind-off images, the ratio image does not represent the true pressures across the model. To visually observe the effect on wind tunnel imagery, a 10 pixel translation misalignment between a wind-on/wind-off image pair was generated and the ratio image produced. The results can be seen in Figure 4.21 on page 109. The effect was qualitatively visible image (a) as noise in the PSP coating that was not present in the correctly registered image seen in (b). The offset was also visible in the black and white circular fiducial marker pairs. The black marker image was from the wind-off image fiducial marker location, and the white image from the same marker's wind-on location. When these were correctly aligned, as in image (b), the images overlapped and the contrast was reduced. The same effect was seen in a vertical row of pressure taps at the center of the image. Finally, linear features such as a wing edge, shown in the bottom of these images, became very dark or very bright depending on the misregistration direction.

In summary, this test demonstrated that local intensity of the image devi-



(c)

Figure 4.21: The effect of misregistration on pressure data. Image (a) shows the ratio image of a PSP coated wind tunnel model with a wind-on image offset of approximately 10 pixels to the right of the wind-off image. Noise is lower in an aligned image pair in (b). A comparison of an image intensity slice shows the increase in noise due to misregistration.

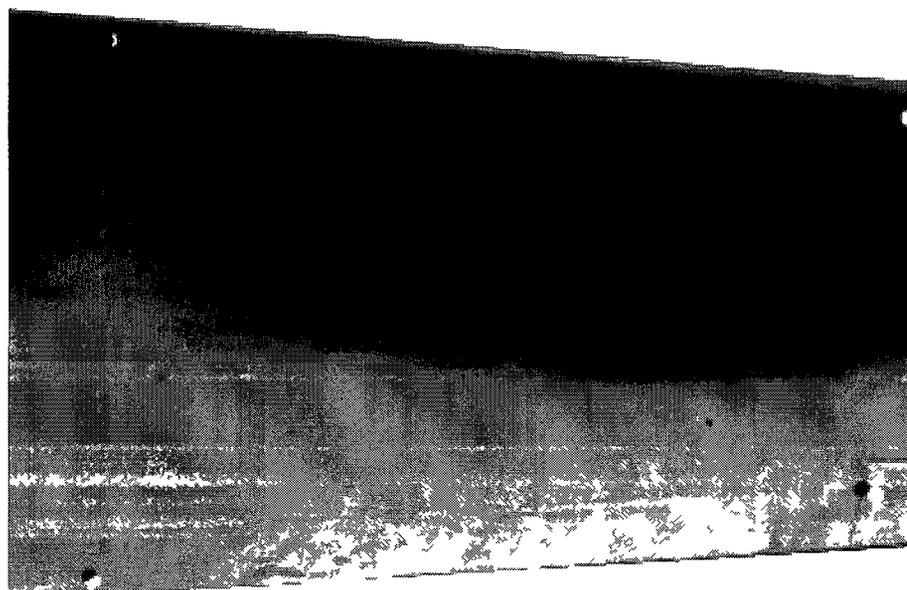
ated significantly in the unaligned image pair, directly affecting measured pressure data via PSP intensity. This demonstrated the importance of correct registration of wind-on and wind-off images.

#### 4.6.2 PLR Using Natural Feature Tracking Only

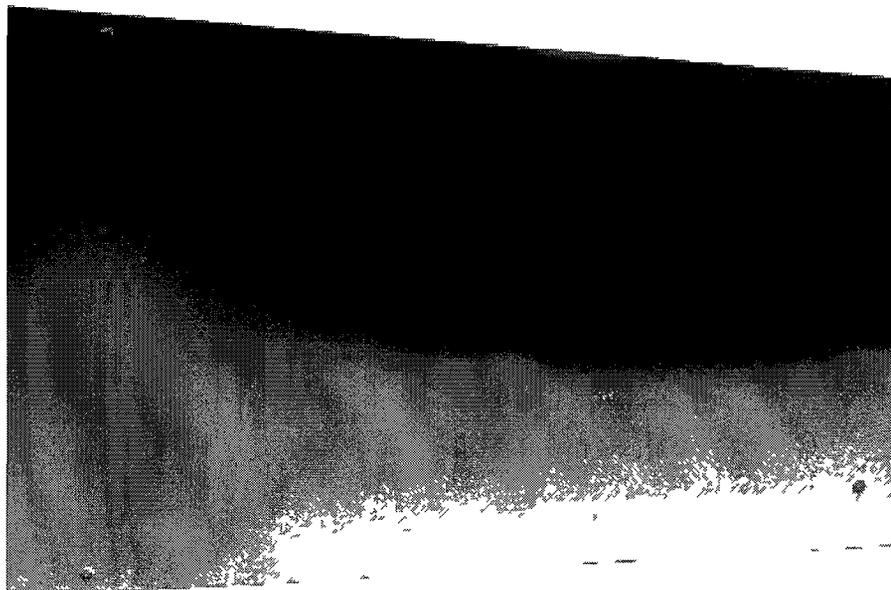
To test the accuracy of resection based on natural feature tracking only, a test was conducted using fiducial marker tracking information to determine natural feature-based motion tracking error. This test does not use fiducial marker tracking data for resection, only natural feature data, allowing the error to be evaluated using fiducial marker positions.

Figure 4.22 on page 111 shows the results. The uncorrected image (a) exhibits registration errors in fiducial markers position towards the wingtip (on right) which are visible as partially overlapping black and white marker images. The same image based on SURF-based resection is shown in (b) with reduced registration error and the alignment is improved compared to no resection. As shown in Figure 4.23 the error in fiducial marker centers averages 0.36 pixels with a maximum of 0.79 pixels and a standard deviation of 0.26 pixels. Comparing tracking error against feature displacement indicates there was a weak or minimal correlation between, indicating that accuracy was not affected by feature displacement.

In summary, this test showed that SURF tracking alone provided good motion tracking data for realignment, successfully reducing mis-registration error in PSP images without the use of fiducial marker tracking. As expected, tracking accuracy for any given SURF feature was not as accurate as tracking of any given fiducial marker. In practical application, this result indicates that a combination of sparse, accurate fiducial marker data and dense, less accurate SURF data would be



(a)



(b)

Figure 4.22: Resection using natural feature only, crops of full images. Image (a) shows a the ratio image without resection. The registration error in fiducial markers position towards the wingtip (on right) is visible as overlapping black and white marker images. The same image based on natural feature-based resection is shown in (b) with reduced registration error.

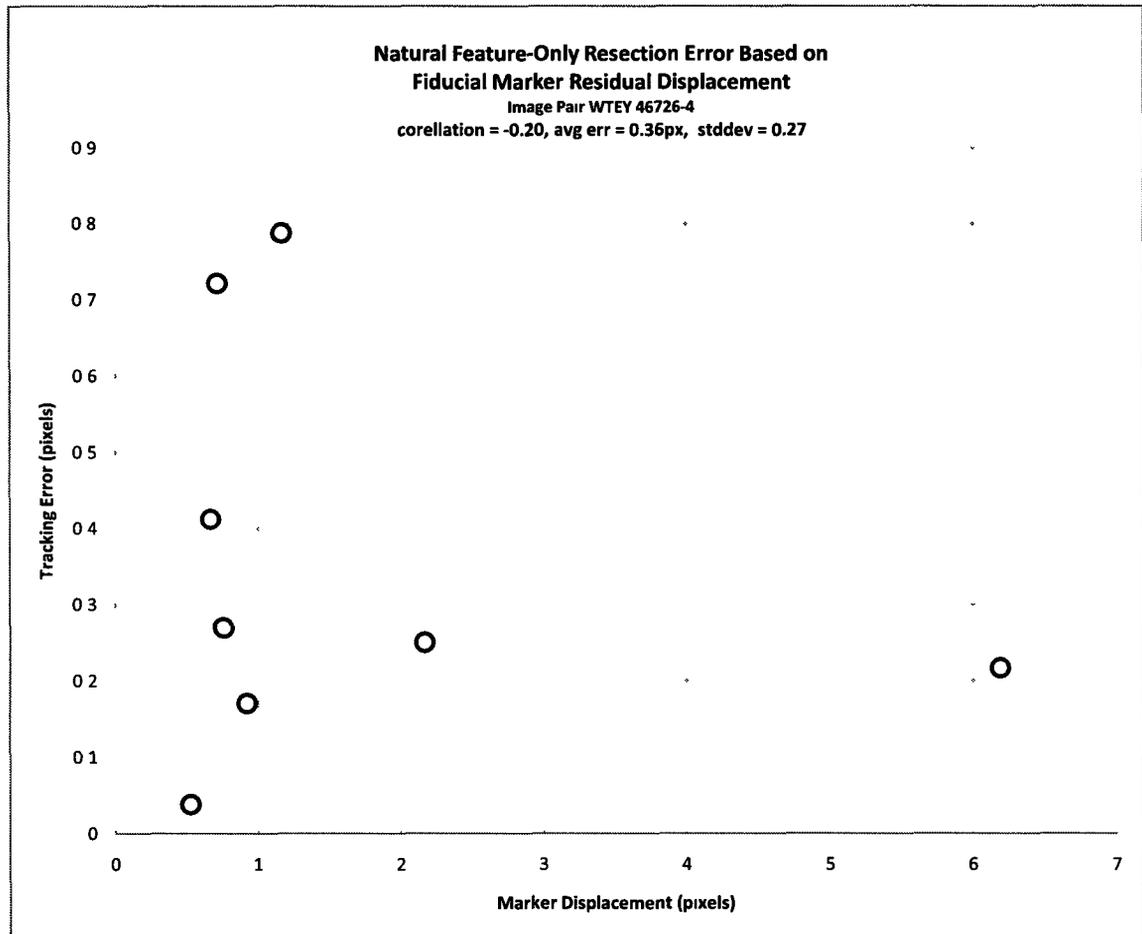


Figure 4.23: Alignment error after natural feature-based resection. Fiducial marker residual error is used as a measure of natural feature tracking accuracy. Tracking accuracy is better than 1 pixel, and there is a weak (-0.2) correlation between image deformation and tracking accuracy.

---

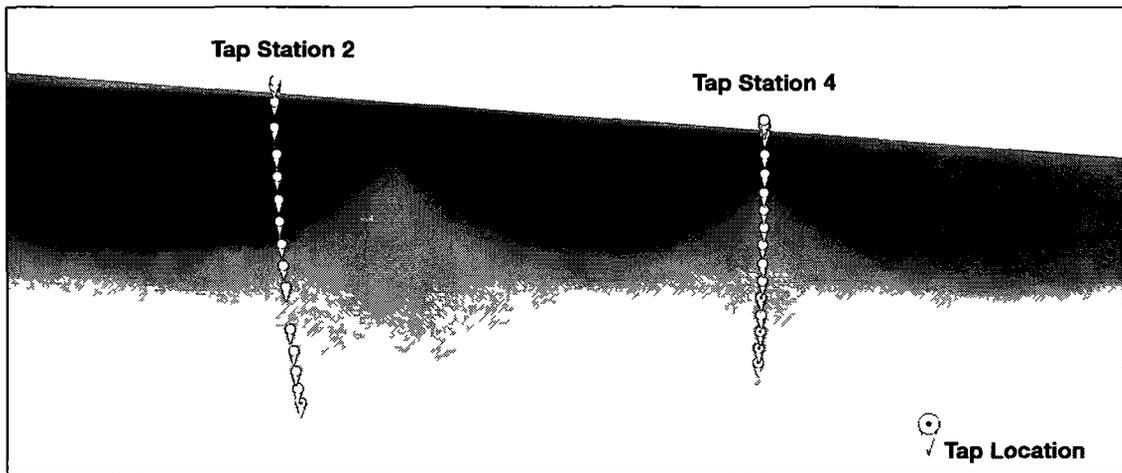
more accurate than either alone.

### 4.6.3 Validating Workflow Using Pressure Tap Data

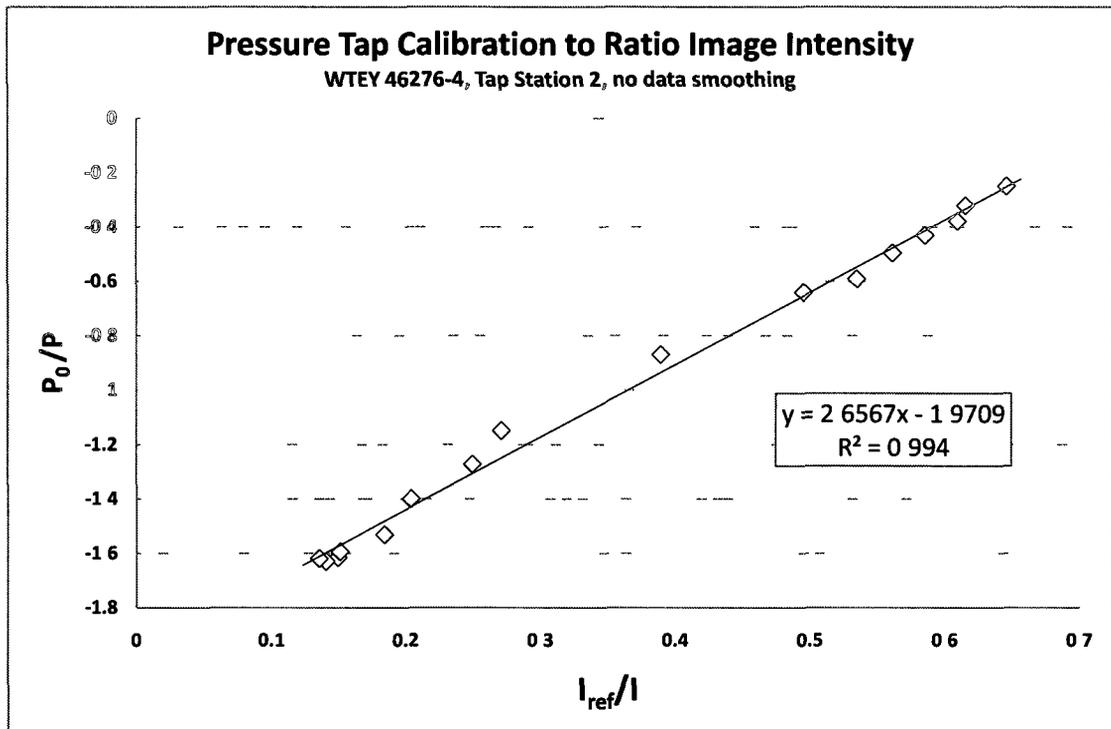
Pressure taps are distributed across the model, typically linearly across wing chords or other areas of interest. The tap data is reported as a ratio between pressure between wind-off and wind-on conditions, and this ratio has a linear relationship to the PSP response ratio. In order to calibrate pressure tap readings the PSP ratio image is used. As shown in Figure 4.24 on page 114, the image intensity values in the immediate vicinity of each pressure tap is correlated with the tap pressure ratio, and the resulting data is used to determine the linear transform between image intensity and pressure ratio. This allows image intensity throughout the image to be converted into pressure values.

To validate the SURF-PLR workflow against the ground-truth data from pressure taps, this calibration was used to measure the accuracy of SURF-PLR resection in the vicinity of a different pressure tap station in the same ratio image. This test station was chosen in an area of the model that shows moderate deformation. The results are shown in Figure 4.25 on page 115. The reduction in image noise due to SURF-PLR resection is visible in figure (a), notably where taps have become less visible due to improved image alignment. Figure (b) compared the resected and unresected PSP data to the calibrated pressure tap data. The sum of squared error of image ratio at the pressure taps is reduced from 0.104 to 0.043, a 59% improvement. In this test process, no data smoothing or interpolation of data was done in either the calibration or measurement steps.

The results of this test show that the SURF-PLR work-flow is valid when compared to ground-truth data. The image slices used for calibration and mea-

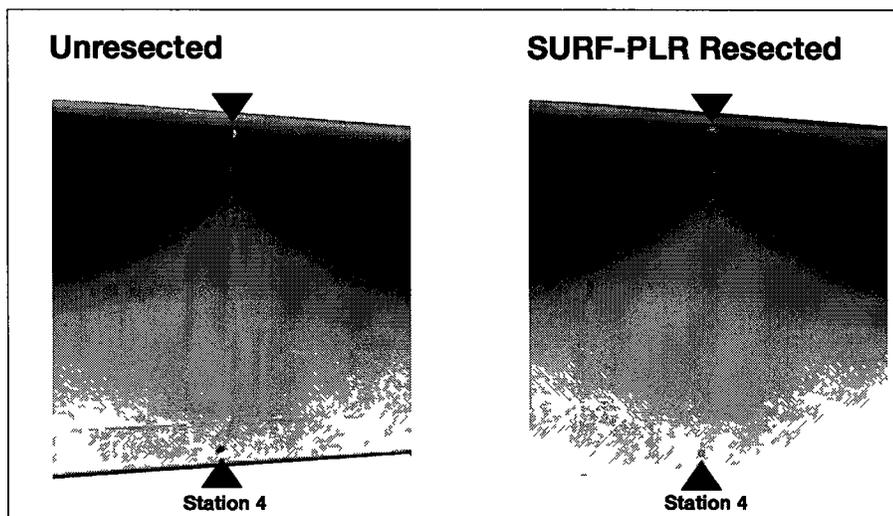


(a)

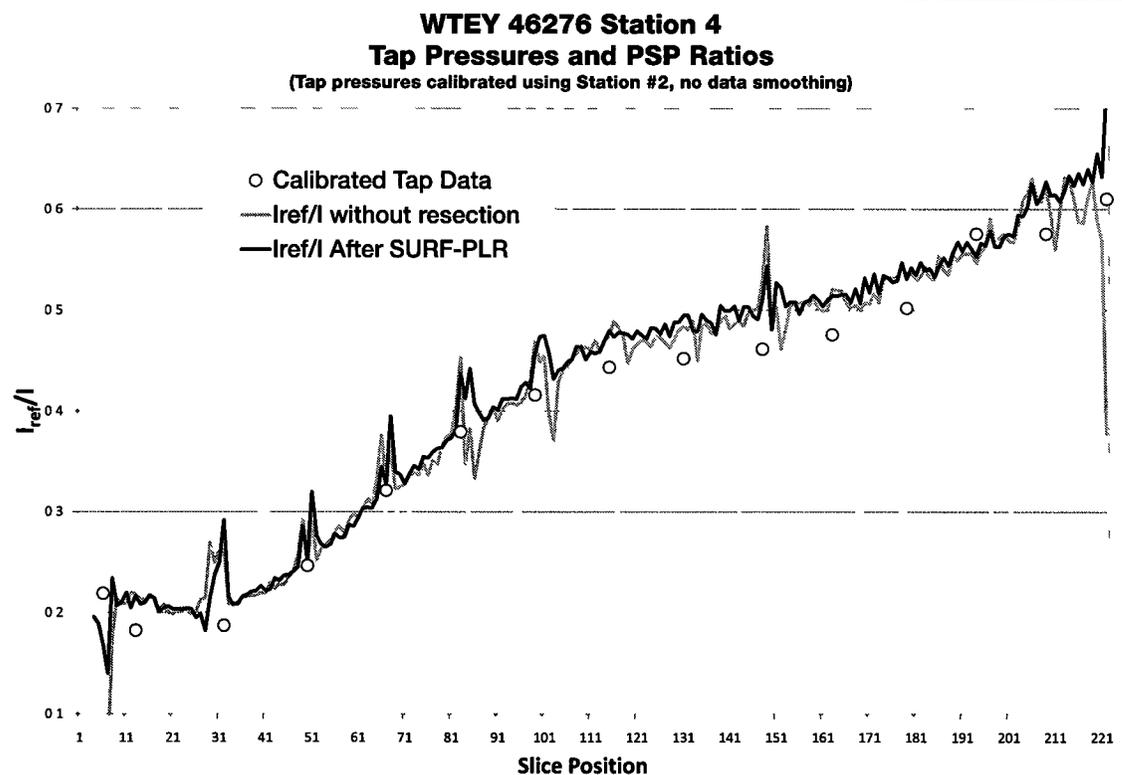


(b)

Figure 4.24: Calibration of PSP using pressure tap data. Figure (a) shows tap station locations along wing chord. The slices intersect several pressure taps (referred to collectively as a "station"). Station 2 intensity data was used to calibrate tap pressure ratios to image intensity ratios. The resulting linear transform is used to relate pressure tap readings to image intensity readings throughout the ratio image. No data smoothing or manual correction is used for the calibration. Figure (a) image is gamma-corrected (1.5) for clarity.



(a)



(b)

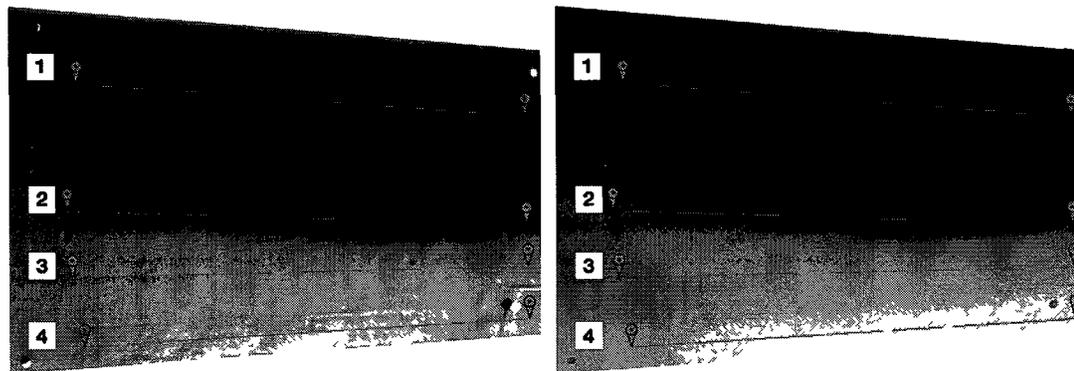
Figure 4.25: Calibration of PSP using pressure tap data. Data is shown for an image slice through Station 4 of 4.24(a). Both SURF-PLR-resected and unresected intensity data are compared to tap data. The sum of squared error of image ratio at the pressure taps is reduced from 0.104 to 0.043, a 59% improvement. No data smoothing is used, so tap locations are visible in PSP data. Figure (a) images are gamma-corrected (1.25) for clarity.

surement each pass through at least 10 resection triangles without noticeable discontinuity or deviation from the unresected image data, demonstrating that resection is not introducing significant distortion. The improvement in measurement accuracy, though small in absolute terms, demonstrates that resection does improve measurement accuracy. This calibration and measurement was done using unsmoothed PSP data, showing that resection can potentially reduce the data smoothing steps required in the pressure measurement workflow.

In summary, this test shows that SURF-PLR provides an improvement to ratio image fidelity when compared to pressure tap data. It validates the SURF-PLR workflow with ground-truth data from pressure taps and provides confidence in further tests on image areas away from pressure taps.

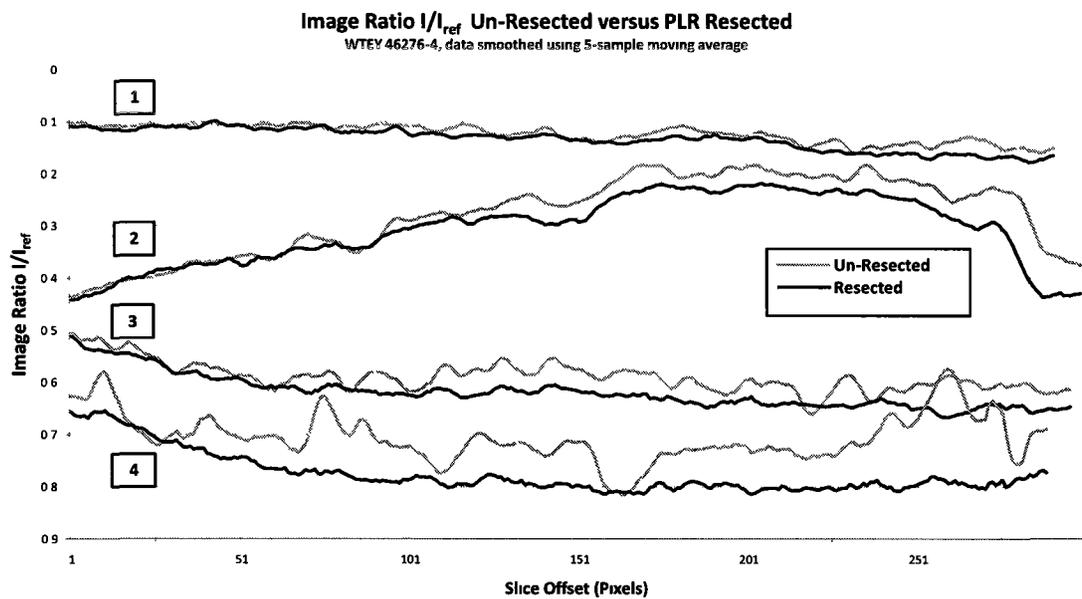
#### **4.6.4 Analyzing Spanwise Ratio Image Slice Data**

The effect of PLR on cross-sections through ratio image data are shown in 4.26. Spanwise slices were chosen to visualize the increasing wing deformation from wing root to wing tip. Several span-wise slices of intensity data were sampled near the wingtip area, where model deformation ranged from 1.2 to 6.0 pixels, with maximum deformation near the outboard trailing edge of the wing. These slices are compared before and after SURF-PLR resection, showing a reduction in noise, especially near the trailing edge of the wing. As expected, in areas of lower deformation, such as the leading edge of the wing at Slice [1], there was closer agreement between the two ratio images. Slice [2] and [3] are relatively close on the wing chord, and display approximately the same level of disparity. Slice [3] cuts through a nonlinearity in air flow, specifically a flow separation, and the comparison of uncorrected and corrected intensity data shows that model distortion can



(a)

(b)



(c)

Figure 4.26: Resection effect on ratios cross sections (a) shows the ratio image without resection, (b) has been resected. Comparison between these images using specific slices of intensity data are shown in (c)

change the apparent location and magnitude of this discontinuity. Slice [4] shows an overall increase in intensity disparity between uncorrected and corrected ratio images, notably near the wing tip on the right of the images. This was expected since the ratio is being determined in these areas using a wind-on image which was offset towards the left of the image due to wing bend towards the camera; this moves higher-intensity pixel values left of their actual locations, effectively increasing the local intensity of wind-on image.

### **Discussion and Summary of Spanwise Slice Test**

It was noted that the fiducial marker in the top-right of these images was partially out of frame in the wind-on image, however natural feature tracking provided data to resect this area of the image nonetheless. This demonstrates that in practical applications where fiducial markers cannot be used, natural feature tracking can provide good tracking data.

In summary, this test showed improvement in fidelity of PSP pressure data when SURF-PLR was compared to unresected ratio images, notably in areas of higher misregistration. Both higher frequency noise and global drift in pressure readings towards the wingtip were reduced. Areas with little misregistration were effectively left as found, indicating that local, accurate resection was taking place throughout the image. This demonstrates that SURF-PLR improved measurement accuracy across large areas of the image which had undergone nonlinear distortion with no prior knowledge of the deformation modes or 3D model geometry.

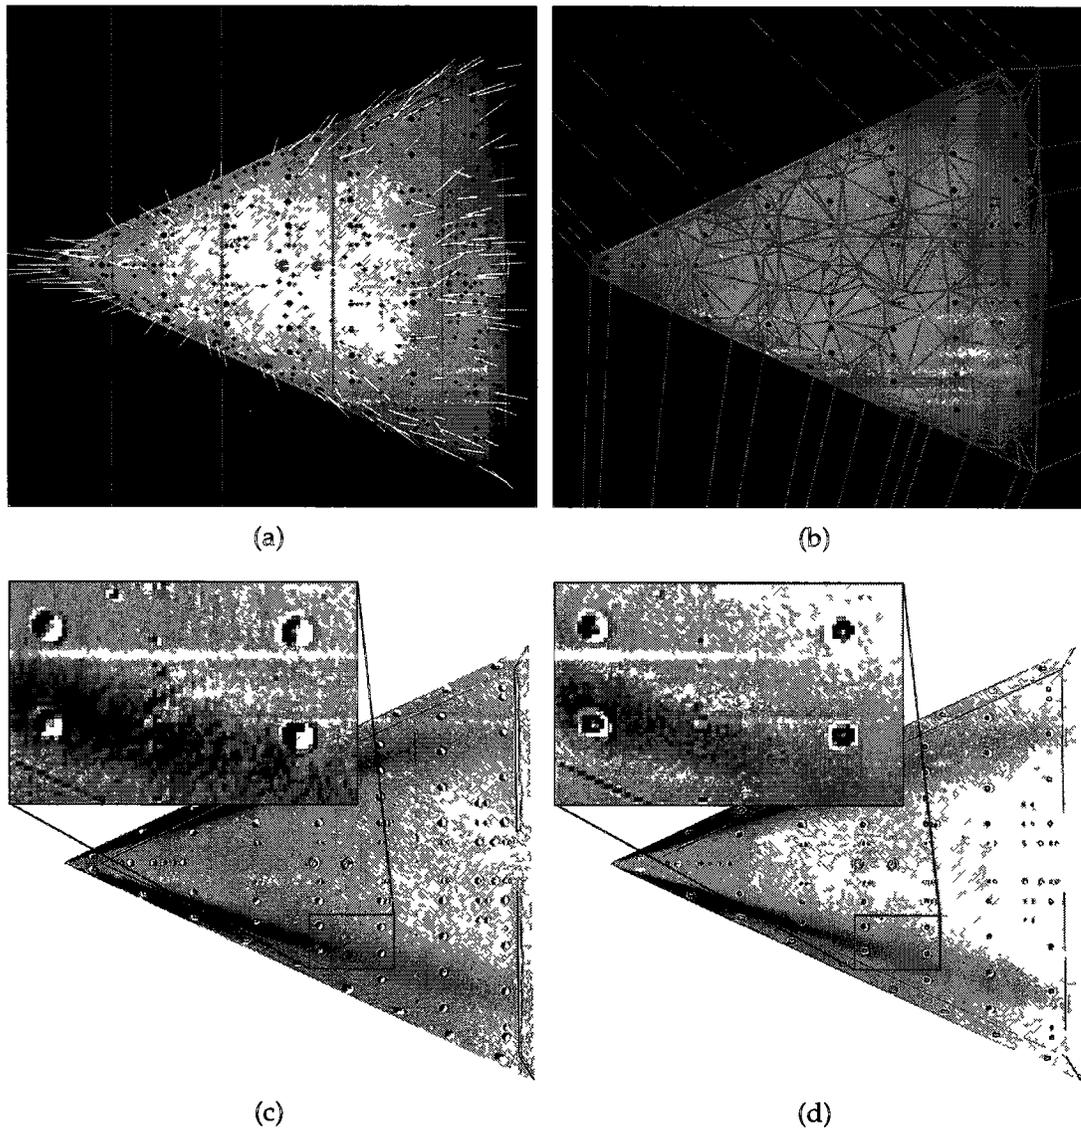


Figure 4.27: Natural feature-only resection of delta wing. The model is sting-mounted and relatively non-deformable, and motion is primarily a rotation of the nose towards the camera, with a center of rotation at approximately one-third of the model length. (a) shows motion vectors as determined from natural feature tracking, lengthened by a factor of 5 for visibility. Four subregions are used for outlier filtering. (b) shows the resulting resection mesh. (c) shows the ratio image without resection, (d) has been resected, showing reduced registration error around the fiducial markers and pressure taps. Image gamma in (c) and (d) increased to 1.5 for clarity.

---

## Analyzing Ratio Images for a Delta Wing Model

A similar test was completed on a supersonic delta wing model, as shown in Figure 4.27. This model differs from the transonic wing in the previous test in that it was sting-mounted and relatively stiff, so the majority of registration error from wind-off to wind-on was due to displacement of the model in the air stream. The PSP paint pattern was more homogeneous, providing less local texture contrast for natural feature tracking. It also had a higher density of fiducial markers allowing a distribution chart of natural feature tracking error magnitudes to be generated, as seen in Figure 4.28.

Initial natural feature tracking indicated this motion was a rotation towards the camera, with the greatest displacement at the nose of the model. Resection using only natural features resulted in reduced error as measured by residual displacement of fiducial markers. Average residual displacement error was reduced from 0.51 pixels in the uncorrected image to 0.16 pixels, and maximum error was reduced from 0.71 to 0.47 pixels.

## Discussion and Summary of Delta Wing Test

This improvement as a percentage of initial misregistration was found to be relatively small, however the initial displacement of the model was less than 1 pixel, demonstrating that natural feature tracking can improve reduce even this sub-pixel displacement. Because error was measured across the entire surface of the model due to the large number of fiducial markers, this test demonstrated that resection with the dense natural feature-based mesh was viable for small global perspective changes of the model as well as local deformations. It also demonstrated that with a different, more homogenous PSP coating, natural features were

still be found and tracked with a high density.

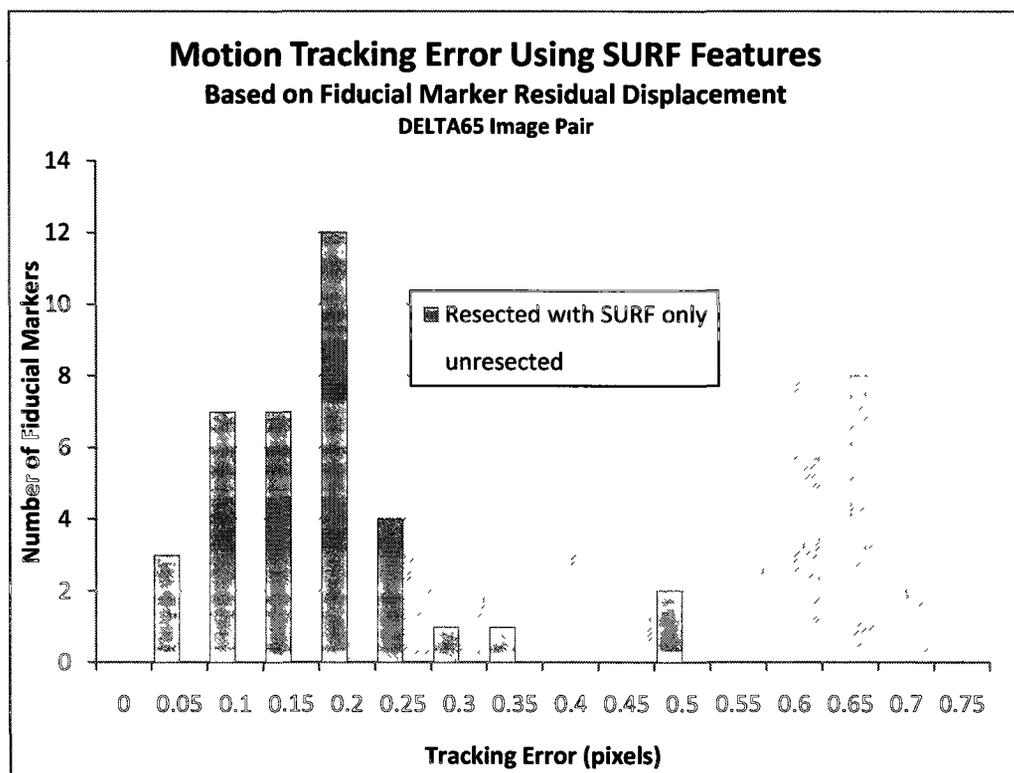


Figure 4.28: Natural feature tracking error on supersonic delta wing. Tracking error is the residual displacement error on fiducial markers in the image.

Some of the SURF features at least partly incorporated fiducial marker imagery, however a fiducial marker alone would not be considered a good feature to track by the SURF algorithm, since it was rotationally invariant. Many tracked natural features included a combination of the linear seams near the perimeter of the model and the adjacent pressure taps, demonstrating that SURF can track model features other than PSP texture. In comparing the tracking accuracy results between the delta wing and the straight wing it was observed that the tracking of surface features such as pressure taps and seams has similar accuracy to tracking

PSP paint texture, with a median error of approximately 0.2 pixels.

In summary, this test demonstrated that SURF-PLR was effective at reducing misregistration due to global model rotation towards the camera even without fiducial marker tracking data. SURF successfully tracked features which were not part of the PSP coating demonstrating the flexibility of this tracking method. In conjunction with earlier testing on models with significant PSP texture and minimal fiducial markers, this indicates that SURF-PLR can reduce misregistration under a wide range of conditions.

#### **4.6.5 Comparing Local and Global Resection on Highly Deformed Imagery**

PLR is a local resection method, since each triangle of the mesh created from natural features is undistorted individually. Global resection methods act on the whole image, distorting it using perspective and affine transforms in order to minimize a metric such as fiducial marker residual displacement between the wind-on and wind-off images. In order to compare these two methods in a visually apparent way, an existing wind tunnel data pair was modified to introduce additional deformation. The wind-on image of the pair was distorted using a Photoshop displacement tool to introduce localized, non-linear displacement into the image such that a global correction method would be unable to remove much of the deformation error. Displacements of approximately 10 pixels were made at select control points which varied image distortion continuously in direction and magnitude both vertically and horizontally. This extreme deformation was chosen to test boundary performance of the two methods such that any advantage of PLR would be immediately visible.

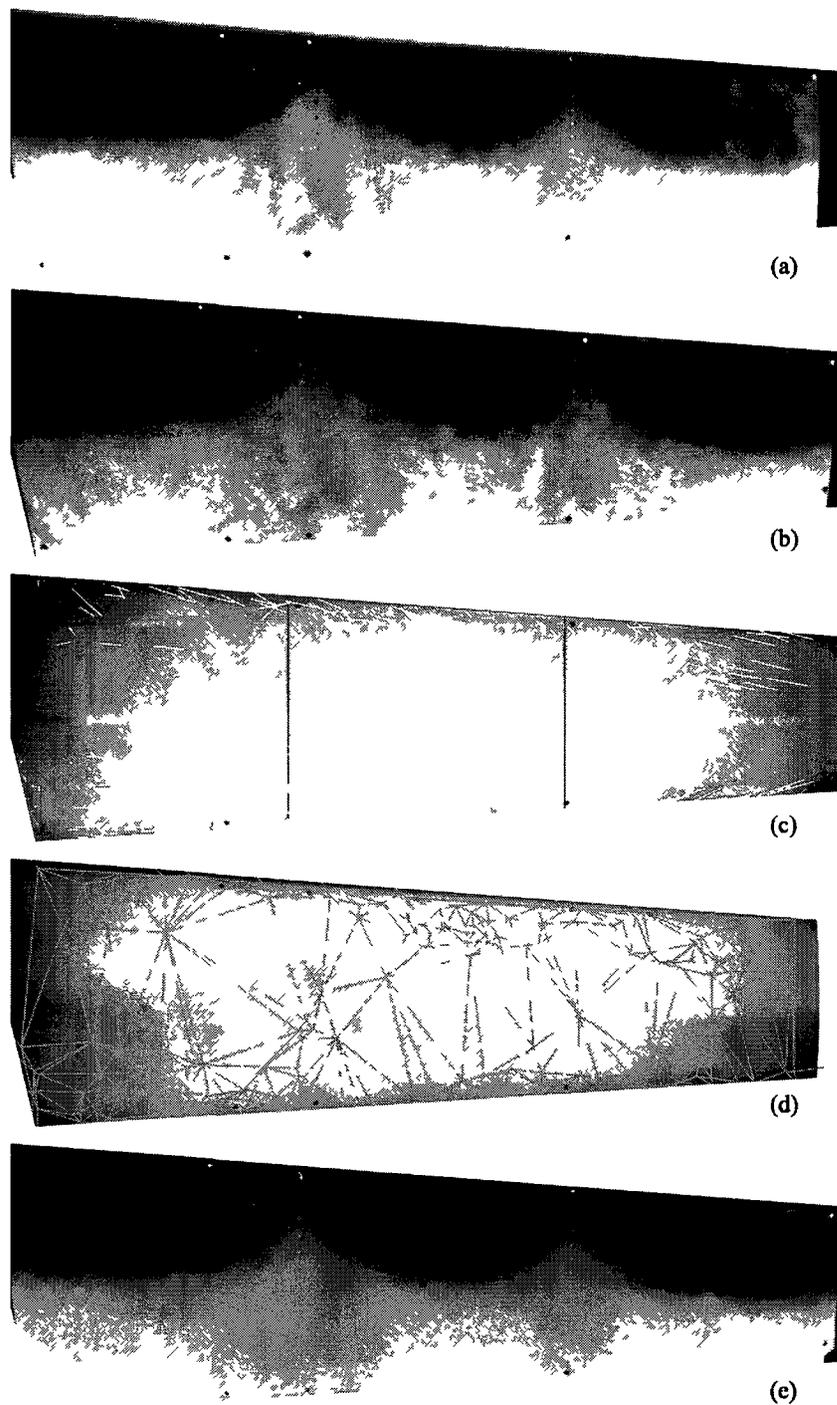


Figure 4.29: Comparison of global versus local resection using a non-linearly distorted wind-on image. Image (a) shows the ratio image with no resectioning (uncorrected). The wind-on image is corrected using a single global correction in (b). Misregistration is visible in the residual fiducial marker displacements and the PSP noise. The wind-on image is processed using the PLR system, with (c) and (d) showing the filtered gradient field and resection mesh, respectively. The PLR-corrected ratio is shown in (e) with reduced PSP noise and residual marker displacement compared to (b). Image gamma increased to 1.75 for clarity.

## Resection Error Comparison

Global Resection versus Natural Feature-Based PLR

WTEY46276-4 with additional non-linear distortion added  
Measured as residual displacement at fiducial markers

<b>Residual Distortion (pixels)</b>	<b>Average</b>	<b>Maximum</b>	<b>Std. Dev.</b>
Uncorrected	9.6	23.9	7.9
Global Resection	8.3	16.3	3.8
Peicewise Linear Resection	1.1	3.5	1.0

Figure 4.30: Comparison of global versus local resection using a non-linearly distorted wind-on image. Based on the image series in Figure 4.29, performance is measured as residual displacement error for fiducial markers in the resected images. Global resectioning provides some improvement in both the average and maximum residual errors, however PLR reduces errors further in this nonlinear distortion test.

Figure 4.29 on page 123 shows the results of this test. Figure (a) shows the ratio using the uncorrected wind-on image with synthetic distortion. Noise is immediately apparent due to mis-registration of the PSP coating, appearing as local brightness variations. Correction in (b) is done through global distortion of the image as a single plane, using fiducial markers as a reference. The displacement of marker errors is reduced by the global method, but because distortion is local and nonlinear, realignment of some areas comes at the expense of misaligning others. Figures (c) through (e) show the process of PLR for comparison. In this instance, fiducial marker tracking is not used to resect the image and is used instead as a reference for measuring performance. The filtered motion vector field is shown in (c) with the image segmented into 3 sections, and (d) shows the resulting resection mesh. The resulting ratio with the reference image is shown in (e) with reduced misregistration noise across the majority of the image. A summary of resection accuracy is given in Table 4.30 on page 124. These measurements confirm what is

visible in the images; specifically that global resection does not perform as well as PLR when distortion is nonlinear.

### **Discussion and Test Summary**

This method of introducing additional distortion into wind tunnel imagery required resampling of the original image in an intermediate software package and potentially reduced the natural feature matching performance. Original camera image data is preferable in cases where natural feature tracking is used, however artificial distortions of these types were not practical to replicate with a PSP coated model. Despite this, enough features were tracked to demonstrate the strength of PLR in handling nonlinear distortion; error was significantly reduced at measured control points (fiducial markers).

It was noted in this test that it was critical to filter the disparity gradient field of outliers, since the distortions in the image were large. Any mesh triangle resected based on erroneous tracking data introduced increased local error and misregistration. A method of providing an accuracy confidence value in any given area of the resected image would be desirable, since the accuracy of PLR can vary across the image. This could be based a combination on the local disparity in the gradient field and characteristics of the local natural features used in tracking, such as their Hessian determinant or quality of match between wind on and wind-off. For example, areas where tracking vectors have a lower local disparity and higher natural feature Hessian determinant could be assigned a relatively higher confidence rating.

In summary, this test demonstrated that SURF-PLR outperforms planar global resection when model distortion is highly localized and nonlinear. In prac-

tical application, this indicates that SURF-PLR can be used in cases where global resection methods are insufficient, specifically where deformation is significant, localized, and nonlinear.

#### 4.6.6 Comparison of SURF-PLR to 3D Global Resection

As discussed in Chapter 2, 3D global resection projects the wind-on image onto a virtual 3D mesh of the wind tunnel model, then applies a global rotation and translation to this mesh in order to minimize displacement error between fiducial markers in the wind-on and wind-off images. A test was conducted to compare the performance of SURF-PLR with that of 3D global resection in order to determine relative performance in a case where model bending and twisting were present.<sup>5</sup>

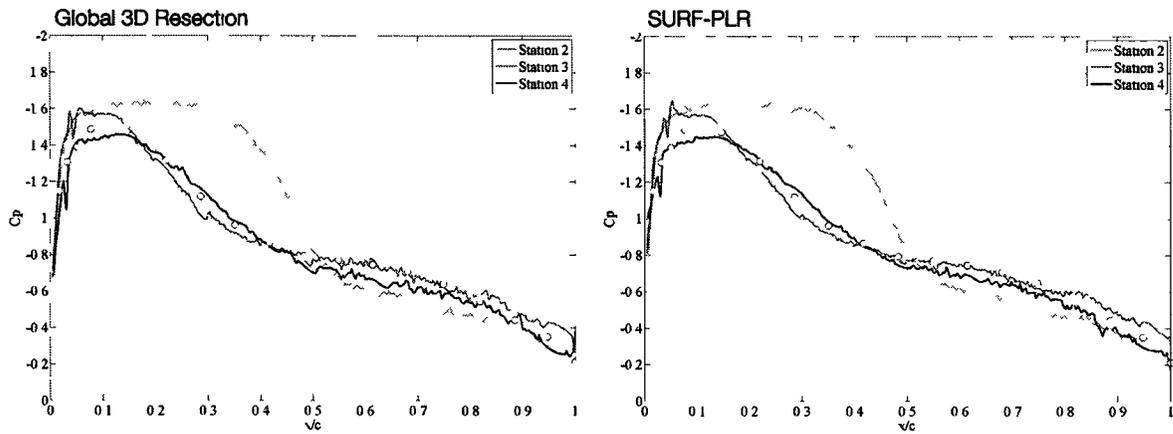
Figure 4.31 on page 127 compares a span-wise and chord-wise slices of the ratio images from SURF-PLR and 3D global resection. The wing model had undergone nonlinear wingtip bending and twisting towards the right hand side of the image, both of which were local deformations. The global resection workflow acted to minimize the total residual displacement error of all tracked fiducial markers; however residual marker displacements were noted in the resected wind-on image from the workflow, with an average residual error of 0.7 pixels and a standard deviation of 0.5 pixels. This can be attributed to the constraint that global 3D resection was acting on a rigid model, whereas local deformation is a non-rigid model motion. SURF-PLR was able to make nonlinear, local corrections to the image, and no meaningful fiducial marker error was noted, due to the algorithm aligning the local marker-areas explicitly. Both data sets showed similar improvement over uncorrected image ratio data.

Pressure tap data was next compared to these resected ratio images at Sta-

## Comparison of SURF-PLR to Global 3D Resection

(WTEY-46276-4, 1px smoothing, tap discontinuities removed)

Chord-wise Slices, Leading to Trailing Edge



Span-wise Slices, Wing Root to Wing Tip

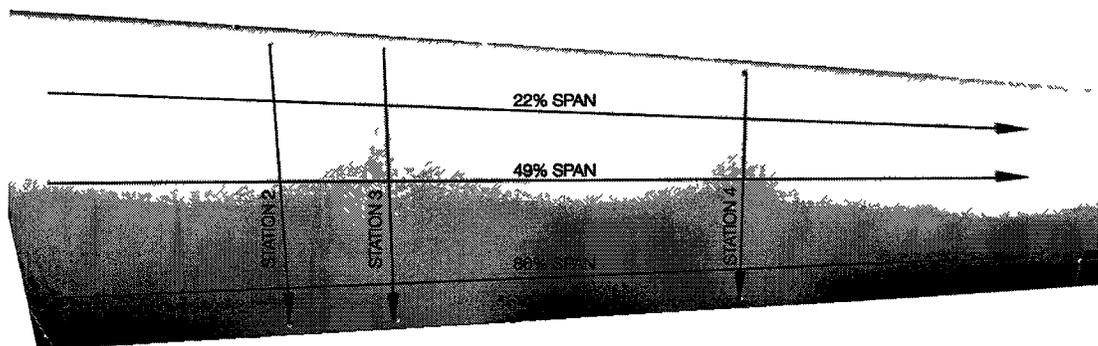
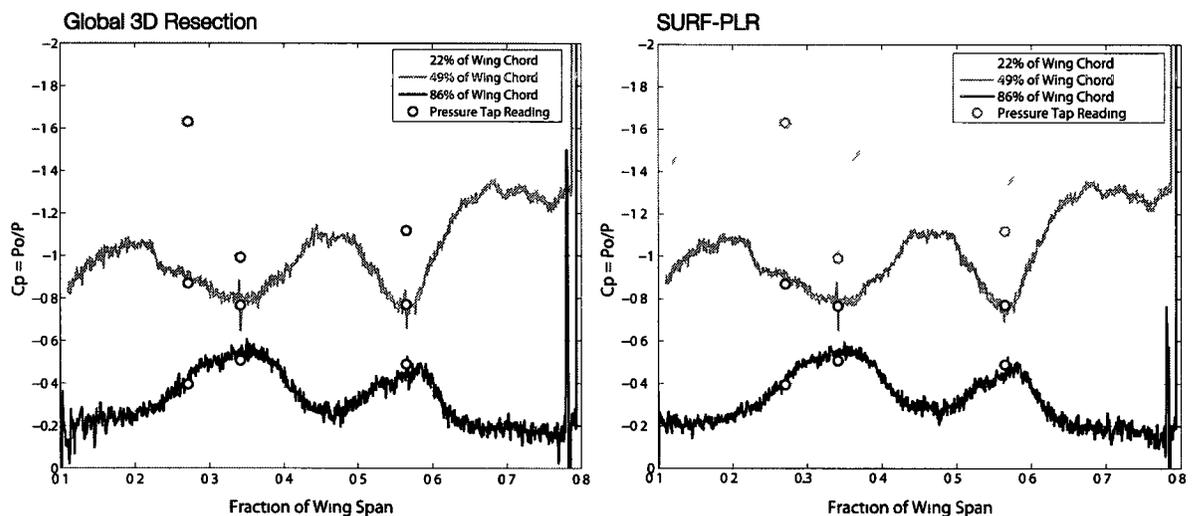


Figure 4.31: Comparison of SURF-PLR to 3D Global Resection. Model deformation took place as a spanwise bend towards the right wing tip, and as a chord-wise twist, again mainly at the right wing tip. Chord-wise slices at Stations 2,3, and 4 indicate that SURF-PLR reduced registration noise in comparison to global resection span-wise slices show the same result, with reduced misregistration noise especially towards the trailing edge of the wing.

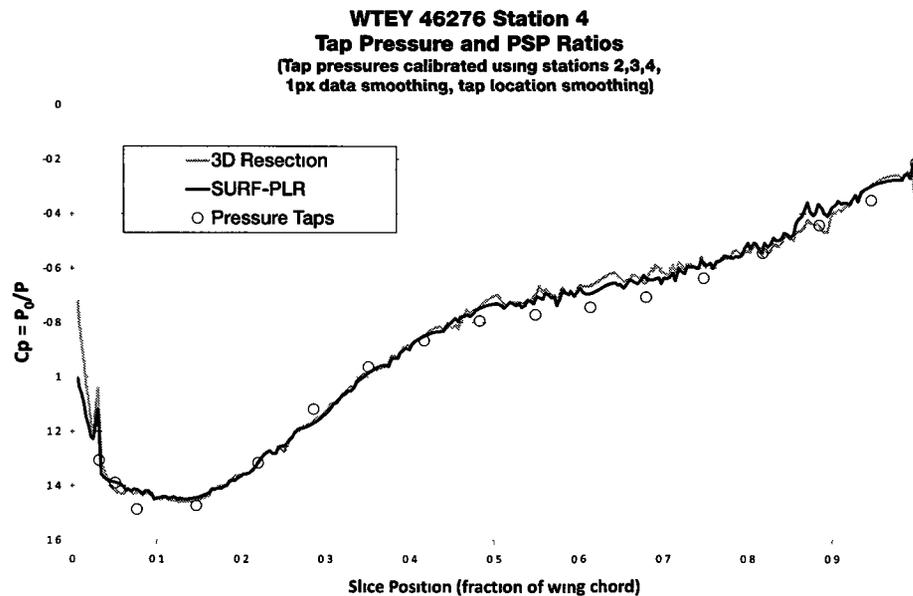


Figure 4.32: Comparison of SURF-PLR to 3D Global Resection. Sum of squared error adjacent to the pressure taps was determined for both methods. SURF-PLR showed an improvement over 3D global resectioning of SSE=0.07 versus 0.13.

tion 4 near the wing tip. Figure 4.32 on page 128 shows the results. The sum of squared errors of PSP intensity ratios adjacent to the pressure taps was determined for both methods. SURF-PLR showed an improvement over 3D global resectioning of SSE=0.07 versus 0.13.

In summary, these tests demonstrate that in a situation where model deformation was nonlinear the global 3D resection method produced a ratio image with more misregistration noise than SURF-PLR. The noise improvement using the SURF-PLR was noted in areas of higher local deformation, confirming that local correction for deformation can improve on global methods which use virtual 3D models. The SURF-PLR was simpler to set up and calibrate as it required no 3D model mesh, camera system calibration, or fiducial marker registration between the physical and virtual models.

#### **4.6.7 Visual Analysis of SURF-PLR Work Flow Results on a Range of Models**

Because the improvement in wind-off/wind-on image registration is visually apparent, the full SURF-PLR workflow was run on wind tunnel models to visually assess the work flow performance. The results of running PLR on three models are shown in Figure 4.33 on page 130. In each case a different type of motion and different set of variables affects performance, providing a wide range of test conditions. The first model exhibits nonlinear wingtip deformation, which was successfully tracked and reduced by PLR despite large changes in local PSP reflectivity due to discontinuous pressure changes. The second model exhibited displacement due to rotation towards the camera, resulting in sub-pixel mis-registration which was successfully tracked and reduced. This model's PSP coating exhibited very uniform texture, demonstrating tracking of model features such as combinations of seams and pressure taps. The third model was an airfoil section with a horizontal rotation axis through the model representing a complex nonlinear motion with a discontinuity. This image pair was composed of two wind-off images, so a perfectly resected ratio image would have had a uniform intensity with minimal noise. Again it was shown that noise was reduced after PLR resection. In all cases a combination of all techniques in the work flow are used, including natural feature tracking, fiducial marker tracking, and vector weighting.

##### **Discussion: Approximating Motion at Mesh Outer Edge**

The motion gradient vector data in all tests did not extend to the edges of the models due to masks applied in the work-flow. This forced the PLR algorithm to

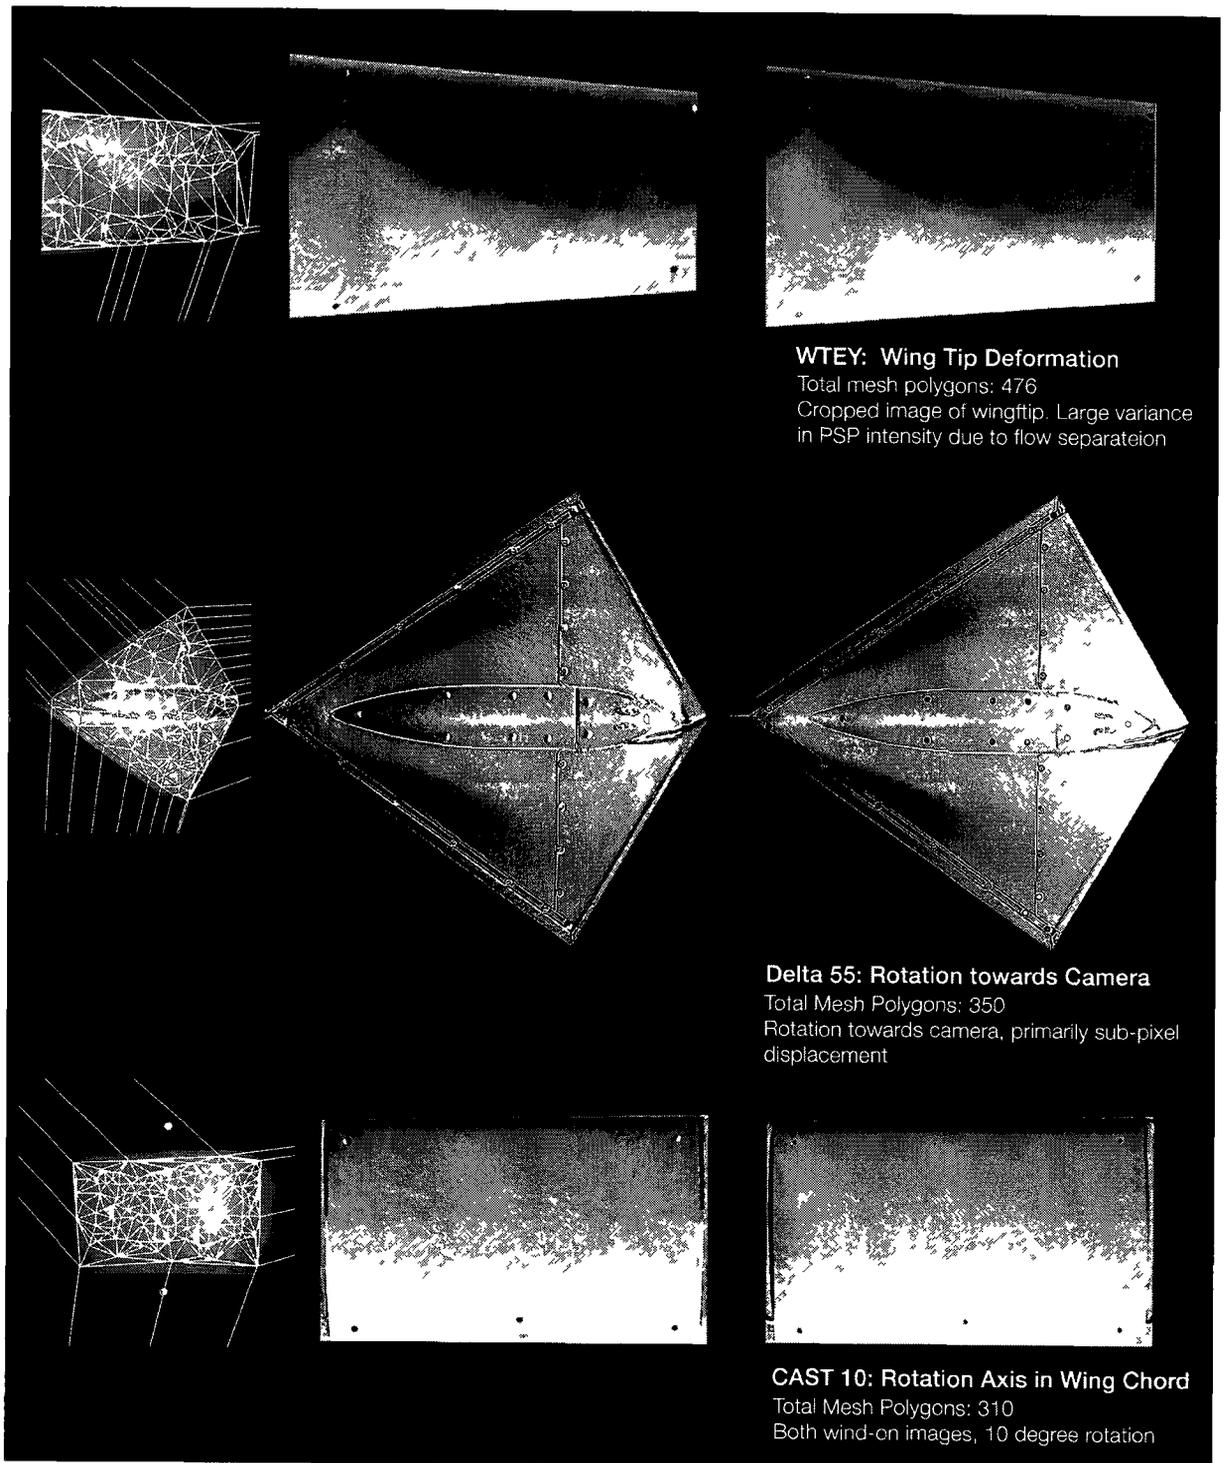


Figure 4.33: Sample results of PLR resection. Three model resections are shown from top to bottom with each model exhibiting a different type of motion and deformation. Three steps of the resection are shown from left to right: the resection mesh, the uncorrected ratio image, and the PLR resected image. Image contrast curve for each image pair modified for clarity.

use an image resection approximation at the edges of the model, specifically an affine, or skew, transformation. The resection accuracy in these areas was reduced compared to areas inside the explicitly defined mesh; however it was still visibly improved over the uncorrected ratio in terms of noise level. The 'mesh edge problem' is a limitation of PLR methods that has been noted in previous work,<sup>21</sup> and was dealt with here by the method of affine approximation. This approximation used known motion vectors for one or two vertices of the boundary triangles, and extended the remaining vertices a large distance from the center of the image to pin it at a fixed coordinate. This approximated an affine transform that used as much motion data as was available at the boundary of the mesh. The result was an improvement over using no motion data at all on the outer boundary of the motion gradient field area.

### **Summary of Test Results**

In summary, these tests showed that SURF-PLR was effective at visibly reducing image noise due to misregistration on several different wind tunnel models when compared to unresected ratio images. This method increased the polygon count considerably over what was available using fiducial markers alone, and successfully combined SURF and fiducial marker tracking information. SURF-PLR was also shown to reduce registration errors outside the resection mesh area by using an affine motion approximation, an improvement over unresected imagery. In practical application this demonstrated that SURF-PLR can be used to reduce misregistration error in a range of model geometries and motions including half-models with wind-bend, sting mounted models with global rotation, and 2D sections with changing angle of attack.

## 4.7 Interface Design

Although this software work-flow can operate automatically from feature identification through resection and ratio image generation, there are several parameters within each module that can be tuned with user intervention (or by system-level closed loop tuning software). Specifically, fiducial marker identification and flow vector filtering affect the quality of the final pressure data significantly, tuning these processes quickly using visual feedback is desirable. Human-in-the-loop systems are efficient in image processing applications, since a level of intuition and experience can be used to produce a more accurate final result.

The work-flow was designed to operate on a single pair of images at a time: a wind-off reference image and a wind-on image, however once parameters are tuned for a specific data gathering run, the parameters can typically be used for all image pairs in the run. For this reason, the user interface was designed to quickly save and reapply parameter settings to multiple image pairs. Screen shots of the user interface can be seen in Figure 4.34 on page 133

### 4.7.1 Project Level Parameters and Extensibility

The software tool chain was tested successfully both with the user interface front end and as standalone PLR executables managed by batch and YML configuration files. Any scripting language or external scripting system can conceivably use the modules. For example, any or all of the modules can be used via MATLAB, which can read and write YML configuration and data files. Since the natural feature and fiducial marker data are generated by their own modules and used as input to the motion vector 'Flow' module, it is possible to use other fiducial marker localiza-

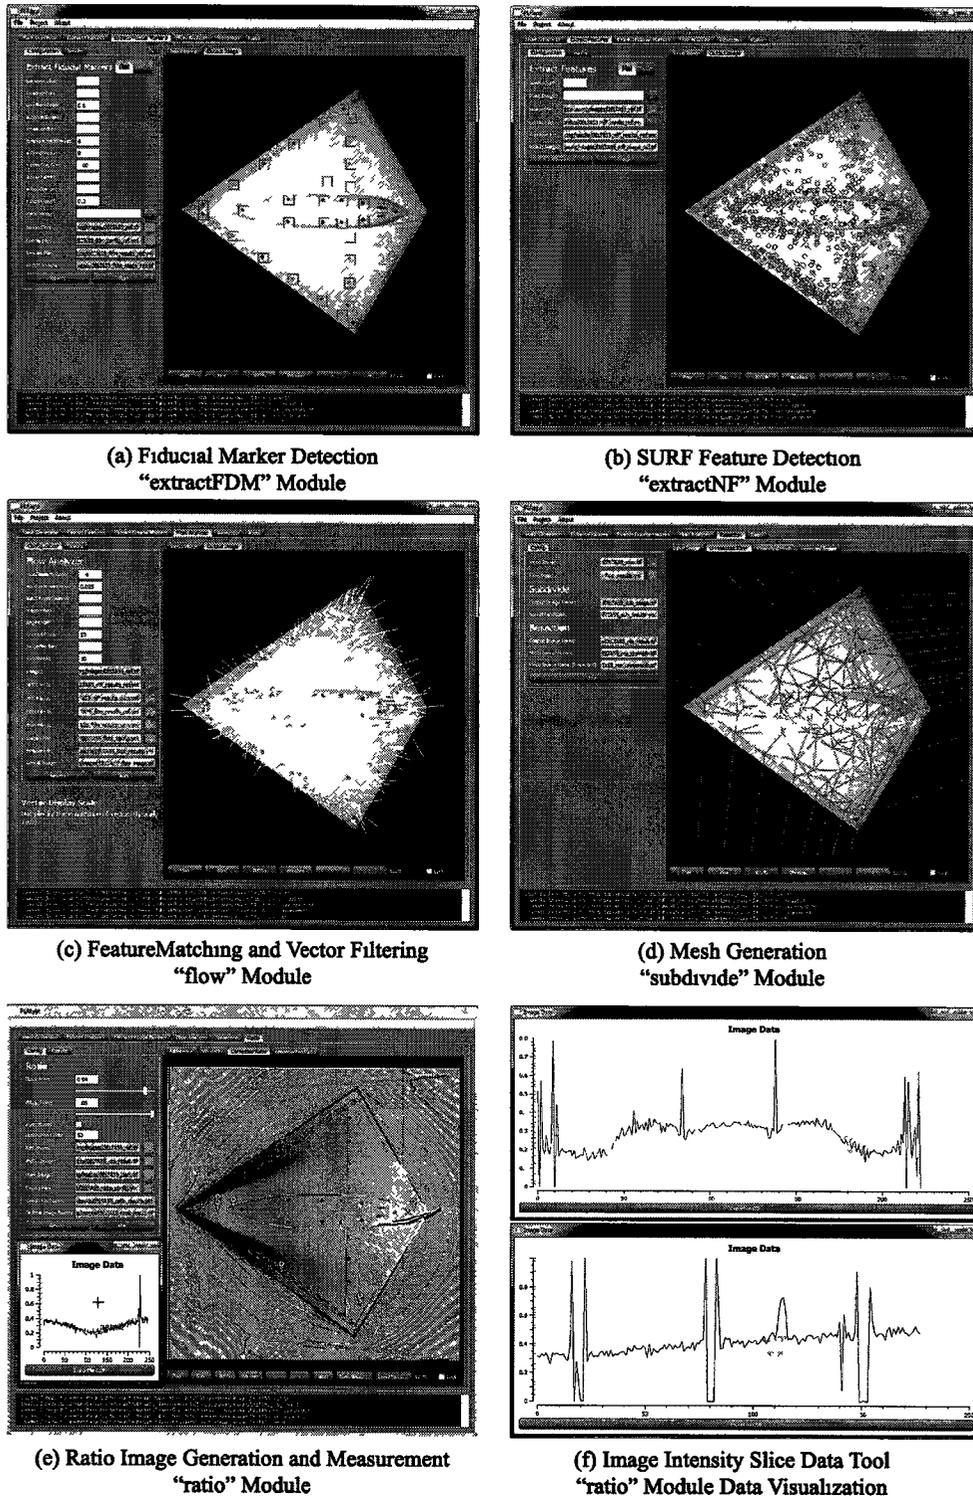


Figure 4.34: PLR user interface screen shots. Each stage in image pair analysis and resection was managed through the interface. Parameters and numeric results were presented on the left and visualizations are presented on the right. A tab system allowed parallel access to every stage of the work-flow and the ability to quickly switch between wind-on and wind-off images at any stage. A dedicated 'slice' tool provided direct visualization of specified ratio image sections.

tion or natural feature extraction methods, allowing the tool chain to be selectively integrated with other software. Examples of the YML configuration and output files are shown in Figure yml on page 135.

### **4.7.2 Slice Tool**

PSP imagery provides quantitative 2D surface pressure data, however it is often necessary to extract a specific pressure cross section set from an image, as shown in the previous section tests. Slices effectively act as virtual point or line segment pressure transducers that are repeatable in position across multiple ratio images. For example, mapping intensity over a single slice that intersects several pressure taps allows absolute PSP pressure data verification.

A slice tool was created which plots intensity across an arbitrary line or line set on an image. The slices positions could be saved and applied to other images, allowing comparisons at the same pixel locations for multiple ratio images. Slice data was exported as YML data that could be plotted in other applications.

### **4.7.3 Findings From User Testing**

The software was designed to leverage the human capability for pattern recognition which provides the outermost feedback loop in the resection process. For this reason visualization of each step was found to be critical in minimizing the total processing time for a given image pair project. This informed some additions to the interface during development, including a before/after visualization comparing the PLR-resected ratio image with an unresected ratio image. Overall, an image pair could be processed within 30 seconds, including basic parameter tuning for

### Sample Output of SURF Feature Finder “ExtractNF”

```
%YAML 1.0
module extractNF
version 0 2000000000000000
description locates natural features in
an image using SURF algorithm
# SURF feature extraction on file
file "wtey46726-4/images/WTEY46726-4_ref.tif"
# Features extracted
total 1073
# extraction parameters used
nmsRadius 2
# features -(x, y, -hessian, -descriptor[real 64]>)
features
-
  x 41 0835762023925780
  y 78 9590682983398440
  size 1118325044
  hessian 86 1891098022460940
  descriptor
    - 2 8036765288561583e-003
    - 2 7403724379837513e-004
    - 2 8036765288561583e-003
    - 3 9891185588203371e-004
    - 0 0695258080959320
-
  x 137 6676330566406200
  y 84 1429748535156250
  size 1118593235
  hessian 86 7850875854492190
  descriptor
    - 5 0178513629361987e-004
    - 4 2759733332786709e-005
    - 5 0178513629361987e-004
    - 3 2600099802948534e-004
```

### Sample Configuration for Motion Tracker “Flow”

```
%YAML 1.0
# Flow extraction parameters
parameters
  maxDisparityRatio 1 500000000000000e+000
  maxCartesianDistance 2 000000000000001e-001
  maxCenterDistance 7
  segmentsH 3
  segmentsV 1
  displayScale 10
  targetMedian 0
  fdmWeight 10
```

### Sample Input/Output File for Ratio Image Slice Tool

```
%YAML 1.0
totalSlices 3
imageWidth 1024
imageHeight 471
slices
-
  x1 735
  y1 333
  x2 835
  y2 336
  data
    - 735
    - 333
    - 6 3024902343750000e-001
-
  x1 735
  y1 333
  x2 835
  y2 336
  data
    - 275
    - 199
    - 6 6638487000000000e-001
```

Figure 4.35: Sample configuration and output YML files. YML files are used to configure each functional module in the workflow and to store output results. The files start with optional global information for test tracking and are editable as text and human-readable. These files can be read and written by external software tools, providing options for integration of all or part of the SURF-PLR workflow.

fiducial marker detection and disparity gradient filtering. The ability to duplicate an existing project was also found to increase overall efficiency when dealing with multiple image pairs from the same wind tunnel experimental run.

In summary, the user interface was found SURF-PLR workflow efficiency by automatically managing image files, parameter input, and intermediate imagery and used for visual feedback. In practical application, this user interface successfully leveraged the user's strength in pattern recognition to accelerate parameter tuning and image processing, and was found to be a valuable tool for rapid resection and measurement of PSP wind tunnel imagery.

# Chapter 5

## Conclusion

### 5.1 Summary of Findings

PLR using a combination of natural feature and fiducial marker tracking has been shown to reduce wind tunnel model motion and deformation errors in the PSP application. In cases of local and nonlinear deformation and motion, it was shown to remove these artefacts better than global resection methods. It was shown to operate with no prior data about the 3D model geometry, deformations, or camera calibration parameters. The software tool-chain written to perform this task successfully and automatically determined model motion and deformation to sub-pixel accuracies using a combination of SURF natural feature and fiducial marker tracking.

It was shown that SURF natural feature were tolerant of the types of small 3D transformations that occur between a typical wind-on/wind-off image pair. SURF natural feature were also tolerant of illumination changes due to model aspect change with respect to the illumination source and due to changes in reflec-

---

tivity of the PSP itself. Good coverage of the wind tunnel model by motion vectors could be achieved by allowing the Hessian score of the SURF features to approach a lower bound of 5-10.

SURF matching used a three stage process. Matches were discarded if their displacement between wind-on and wind-off images exceeded a specified threshold or if the total error between their descriptors exceeded a separate threshold. This was found to remove more false matches than other matching methods, without removing potential inliers. The resulting motion gradient field was found to require further global filtering by a disparity gradient filter.

Tracking of fiducial markers using an efficient 3-step method was demonstrated. This algorithm used a configurable synthetic fiducial marker template to detect markers and a two-stage center localization algorithm that was shown to be as accurate as currently used systems on synthetic marker test targets. The sub-pixel localization method of orthogonal one-dimensional Gaussian fitting was shown to be as accurate as existing algorithms in the presence of noise and low contrast. A novel background gradient detection and subtraction method to improve localization accuracy was developed specifically for this application since many fiducial markers are located near areas of high curvature, such as wing leading edges.

Matching of SURF features and fiducial markers was shown to produce a motion gradient field with good coverage that required additional filtering of mismatches. A disparity gradient filter was implemented to provide a combination of local and global motion gradient filtering that was shown to effectively reduce the fraction of erroneous motion vectors. It was demonstrated that these motion vectors could be split into subsets which could be weighted and prioritized based on

---

the accuracy of their tracking method, which could in turn provide an improved distribution and quality of motion gradient field throughout the wind tunnel image. It was further demonstrated that this algorithm supported segmentation of the motion gradient field into geometrically discrete sections, which allowed for parts of the wind tunnel model to move independently of others while still being filtered for outlying motion vectors.

This motion gradient field was used to subdivide wind-on and wind-off images into a mesh by Delaunay triangulation to provide a basis for performing piecewise linear resection of the wind-on image, reducing model motion and deformation. The resected image was shown to be sub-pixel accurate when using the OpenGL framework. Bilinear interpolation of the individual triangles of the mesh was not found to contribute significant image artefacts that would adversely affect pressure measurements. The resection portion of the tool-chain was shown to function with a 16-bit dynamic range, assuring significant headroom for cameras with dynamic ranges.

Ratio images generated from the PLR-resected images showed a significant decrease in misregistration error due to misalignment and deformation when compared to uncorrected ratio images. A combination of fiducial marker and natural feature tracking was found to provide better overall results than either alone, since the more accurate fiducial marker tracking data was successfully combined with the denser natural feature data to produce a final resection mesh. The tool-chain was shown to function in several different wind tunnel test image sets with different qualities of illumination, PSP coating, fiducial marker distributions, and illumination fields, as well as different model motions and deformations. The system successfully tracked motions such as model rotations towards the camera and wing bending and twisting motions from both natural feature and fiducial marker

sources.

The system was measured to run in less than 30 seconds from start to finish on pairs of images in the 1 megapixel, 16-bit resolution class with approximately 20 fiducial marker features and 500 natural features tracked. The modular architecture of the tool-chain and reconfigurability using text-based files was demonstrated using both a graphical user interface and batch files, demonstrating a range of user interaction or automation options.

This software tool chain was developed with open source software including OpenCV, freeGLUT, TIFFlib, Qt, and Qwt. These open source libraries were found to streamline development by providing functions such as template matching, SURF extraction, image and file management, and user interface frameworks.

In conclusion, natural feature-based tracking combined with an extension of the disparity gradient filter algorithm was found to improve 2D piecewise linear resection in pressure sensitive paint applications where model deformation and motion would otherwise degrade measurement accuracy. This system was shown to have equivalent performance to 3D global resection methods without the need for 3D models and associated fiducial marker registration and camera calibration steps with improved performance where local deformation was present. Although natural feature tracking was less accurate than fiducial marker tracking, the two motion vector sources could be combine, weighted, and segmented using an extension of the basic disparity gradient filter algorithm to produce a higher density motion gradient field. This field in turn provided a basis for a denser resection mesh for the PLR algorithm, resulting in a reduction in ratio image measurement error. It was demonstrated that a modular, open source framework could be leveraged to create an automated resection tool-chain that allowed various levels of user

interaction and integration with other software.

## **5.2 Contributions to the Field**

### **5.2.1 Used Natural Features to Automatically Improve Resection Performance**

PSP as a wind tunnel tool is dependent on the accuracy of the ratio image. This image is highly sensitive to errors caused by misregistration between wind-off and wind-on imagery due to model deformation. Eliminating this deformation using image resection of the wind-on image has been developed using various methods. The most flexible of these methods is PLR, which requires no prior knowledge of model geometry or expected deformation and is as accurate as the model motion tracking method used. A common method is tracking of dedicated fiducial markers which provide high accuracy tracking. The more markers used, the denser the PLR mesh, and the higher the fidelity of the wind-on image resection. However these fiducial marker markers cover the PSP and can produce aerodynamic disturbances. An alternative it to track natural features in the image to increase the density of the mesh without increasing the marker density. The work described here demonstrates that it is feasible to track natural feature using the SURF algorithm for PSP applications and that the resulting increase in mesh density results in a higher fidelity resected wind-on image which in turn increases accuracy of pressure measurement from the ratio image. This method was shown to track model motion and deformation with sufficient fidelity when combine with disparity gradient filter as described below.

### **5.2.2 Demonstrated a Method for Improving Mesh-Edge Resection Accuracy**

PLR is most accurate when the vertices of each triangle in an image mesh can be tracked; however image areas outside the mesh lack full motion information. These areas typically consist of the perimeter of the image including a portion of the wind tunnel model. It was demonstrated that extending the Delaunay triangulation to include stationary vertices well outside the image area allowed the image edge to be resected based on adjacent mesh vertices motion. This resection was an affine approximation of the motion at the mesh edge, however was more accurate than leaving the image perimeter stationary. This addressed an issue stated in prior literature in which PLR accuracy suffered outside of the mesh area.,<sup>1236</sup>

### **5.2.3 Developed a Weighted Extension of the Disparity Gradient Filter Algorithm**

Tracking additional features via SURF can work independently of fiducial marker tracking or in conjunction with it. In order to combine the highly accurate data from fiducial marker tracking with natural feature data, a weighted, prioritized extension of disparity gradient filter was developed. The results show that this extension allows motion vector sets from different sources to be weighted for accuracy and prioritized in terms of their effect on one another, allowing more accurate vectors to have more influence on the final motion gradient field, resulting in higher resection fidelity

## **5.2.4 Developed a Modular, Open Source Workflow for 2D Resection and Ratio Image Generation**

This software tool chain for this project was developed using open source libraries, including OpenCV, freeGLUT, tiffLib, and Qt. This demonstrates that relatively complex software for this application can be developed efficiently without proprietary licensing and analysis tools. The software tool chain was developed as a series of stand-alone modules configurable using YML test files, providing output in standard TIF images and YML data files. This modular architecture allows for batch processing and integration with other tools and software such as MATLAB. The software is developed entirely in C and C++ and is available to the research community for download and extension in the future.

## **5.2.5 Designed Graphical User Interface to Improve Workflow Efficiency**

The software Tool chain developed as part of this work can be run fully automatically; however additional tuning parameters are exposed to the User via YML configuration files, providing a HIL capability. This leverages the pattern and feature recognition strengths of the User to verify motion gradient fields, for example. As part of this work, an intuitive user interface was developed to manage work flow and focus the User on tuning key parameters using a visual feedback system of intermediate processed images. This increased the overall efficiency of resection. The GUI also included measurement tools that provide virtual pressure taps and pressure cross sections that can be saved, exported, and repeated for multiple tests.

### **5.2.6 Findings Applicable to Methods of PSP Coating**

The PSP process relies on image ratios to remove inconsistencies in illumination field and PSP coating response. This can be used to the advantage of the SURF tracker, which relies on distinct, non-uniform features in an image. These inconsistencies in PSP coating response do not have to be structured or patterned, but can be effectively random, as shown in Chapter 4 where Perlin noise was used as an approximation of an inconsistent PSP application. It follows that optimizing a wind tunnel model for SURF feature tracking benefits from non-uniformity of the PSP coating, and that this non-uniformity should be maximized. A further refinement may be the addition of a random speckle pattern in the PSP coating, as initially demonstrated in image correlation work by Park and Sung.<sup>46</sup> The speckle pattern could be a second coat of PSP, removing the distinction between tracking landmarks and sensor. The effort to implement this is minimal, as the pattern can and should be random. Potential methods for testing are described in the Future Research Topics section.

## **5.3 Criteria for Choosing a Resection Method**

The PLR method described in this work addresses specific shortcomings in existing PSP applications. Existing methods can be sufficient or superior in other circumstances. These other methods can also be informed by the results of this tool-chain, making it a complementary tool rather than a replacement. As a practical guideline, the circumstances under which each method is preferred are listed here.

**Where SURF-PLR may be preferred:**

- No 3D virtual model is available
- No camera calibration information is present
- Model deformation or displacement is present, especially localized, nonlinear, undefined distortion
- Insufficient fiducial markers are present
- Model motion and deformation need to be estimated for use in other methods
- Model structure is highly flexible or nonrigid (ornithopters, helicopter blades, parachutes, airships, inflatable structures)

**Where 3D reprojection methods may be preferred:**

- A virtual 3D model is available
- Camera parameters are known
- Fiducial markers are present in sufficient quantity and distribution to solve unknown parameters
- Fiducial markers are accurately registered between virtual model and real model
- Model deformation can be sufficiently described by a small set of deformation equations, or distortion is minimal

**When 2D global resectioning using polynomials may be preferred:**

- Area of interest on wind tunnel model is relatively planar and fiducial markers lie on that plane.

- No 3D virtual model is available.
- Camera parameters are known
- Distortion is can be adequately described as a first-, second- or third-order polynomial.
- Sufficient fiducial markers are present to solve for the polynomial equation above.

### **Summary of Method Selection**

SURF-PLR is a rapid method of reducing model displacement and distortion that requires little input information other than the image pair to be resected. It is this flexibility under a wider variety of practical circumstances that makes it an effective complement to existing systems.

## **5.4 Future Research Topics**

In the process of developing the solutions and tool-chain described here, several new avenues for exploration were noted. Those with the most potential to improve performance are listed here in order of interest.

### **5.4.1 Confidence Map for Resection Result**

Of key interest is not just the final output of the work-flow, but the confidence in this result. Because a ground truth is not available for comparison, a confidence measure is desirable. In the case of PLR, confidence in the resection accuracy can be measured and displayed per-triangle. It is hypothesized that a confidence measure

---

can be generated based on the Disparity Gradient Filter algorithm. The disparity calculation would be made using the set of three motion vectors at the vertices of each triangle in the resection. A larger disparity between these vectors implies more interpolation (i.e. potentially motion error), reducing confidence in that triangle's validity. A small disparity indicates general agreement about the motion of the triangle and implies less interpolation (i.e. a translation as opposed to a rotation or perspective transformation). Because of the structure of the DGF algorithm, this confidence rating implicitly includes vector angles, magnitudes, and the triangle area. This would be a straightforward implementation and addition to the "Resection" or "Subdivide" modules. In conjunction with this, improved methods of 'pinning' the exterior of the resection mesh could be explored, including use of more than three pin points so as to improve resection performance outside the tracked mesh area.

#### **5.4.2 PSP Coating Methods to Improve Natural Feature Tracking**

In the work described here, the pattern of the PSP coating (i.e. its visible non-homogeneity) was a variable present in the data, but not explicitly controlled. In future applications it is feasible to tailor the PSP coating to have a pattern that provides better natural feature tracking. Methods of tracking such as SURF can remove the distinction between tracking landmarks and the PSP sensor, since the trackable features can be created using PSP instead of non-varying fiducial markers. This implies that PSP coatings can be tailored for denser, more accurate natural feature tracking. Furthermore, these features can be random, reducing labor and planning time. Further research could focus on optimal PSP patterns and application techniques for improved SURF feature tracking.

### 5.4.3 Arbitrary Disparity Gradient Filter Segmentation

Based on the results of grid-based image segmentation before the DGF is applied, it is hypothesized that an improved automatic segmentation can be achieved by dividing the image such that each segment centered on a single fiducial marker. This could be achieved using the Voronoi graph complement of the Delaunay triangulation of these markers. This would provide a 'cell' structure in which a single highly accurate fiducial marker could influence the SURF-based vectors in its local area, improving the effectiveness of weighted DGF.

### 5.4.4 Additional Matching and Filtering Constraints on Motion Vector Field

SURF feature matching was found to create a significant amount of outliers which were subsequently filtered using DGF. Reducing these mismatches using additional constraints can potentially lead to more valid matches, since more features would be freed up for additional matching attempts. Constraints which compare candidate matches with respect to their scalar space, difference in Hessian score or difference in dominant orientation could be explored. Additional constraints using calibrated stereo cameras and the epipolar geometry could also be explored as a method of removing outliers.

The DGF algorithm could potentially be further modified for this application. In the current implementation, the disparity between any two vectors is a scalar function of their mutual distance, magnitude and difference in angle. It may be beneficial to provide a method of tuning the DGF disparity score to be more or less sensitive to each of these parameters. This would effectively change the single

“maximum disparity threshold” parameter into three parameters: angle weighting, magnitude weighting, and distance weighting. This parameter space could be explored for optimal weighting combinations for PSP applications.

Finally, the weighted DGF algorithm, as a proof-of-concept, uses an algorithm that increases computation time by the square of the weighting factor. A modified method which does not have this penalty should be tested against the current method in order to supersede it. Such a method has been developed and requires validation.

#### **5.4.5 Extension to Multi-frame Tracking**

Transient model deformations are of increasing importance. For example: helicopter rotor active damping, elastic composite structures, deformation-based control surfaces, and active vortex control. As aerodynamic understanding moves into the realm of transient control, it becomes necessary to track transient events. Multi-frame tracking at high speed is one solution; however this requires either specialized equipment and techniques, and still can require a wind-off reference image. PSP formulations and application techniques for high speed imaging have been studied Gregory et al.<sup>6</sup> and response times in the tens of microseconds are possible. An appropriate high speed camera system is required, for example a line transfer CCD camera with suitable high intensity illumination sources. With issues of PSP responsivity and imaging time addressed, it is possible to minimize error from motion blur and PSP response time constants. The methods described here have been shown to resect images based on landmarks in a reference image, and this can be extended to use a single reference image and multiple wind-on images, separated in time. This would allow visualization of airflow as a movie sequence,

enhancing understanding of transient events. The software tool chain developed for this work can be automated with simple script files to add this functionality, since the tool chain, once tuned for the specific image set, can automatically resect any specified pair of wind-on/wind-off images.. In the case of a constant deformation and displacement in the wind-on condition but varying PSP response, the same resection mesh can be used for all images in the wind-on series with a similar script file configuration.

#### **5.4.6 Integration of Calibration Data from Binary PSP**

PSP with a combination of two lumiphors, one of which has no response to pressure changes, is be used to calibrate pressure ratio readings which may be coupled with other variables such as local temperature or illumination changes. Integrating these additional images into the SURF-PLR workflow could produce improvements in pressure measurement accuracy.

#### **5.4.7 Automated Optimization of Work-Flow Parameters**

The current tool-chain allows user adjustment of parameters to find and match SURF features and fiducial markers. This iterative process is amenable to automation. Fiducial marker location and Disparity Gradient filtering are good examples where the work-flow efficiency could be optimized by a system that adjusts specific parameters and re-runs the modules. Pre-caching a set of results would allow the user to very quickly evaluate the effects of parameter adjustment and choose the best settings.

# References

- <sup>1</sup> J. Jouhaud, P. Sagaut, and B. Labeyrie, A Kriging Approach for CFD/Wind-Tunnel Data Comparison, *Journal of Fluids Engineering* **128**, 847 (2006).
- <sup>2</sup> W. Woods, S. Holland, and M. DiFulvio, Hyper-X stage separation wind-tunnel test program, *Journal of Spacecraft and Rockets* **38**, 811–819 (2001).
- <sup>3</sup> C. Willert and M. Gharib, Digital particle image velocimetry, *Experiments in fluids* **10**, 181–193 (1991).
- <sup>4</sup> L. Venkatakrisnan and G. Meier, Density measurements using the Background Oriented Schlieren technique, *Experiments in Fluids* **37**, 237–247 (2004).
- <sup>5</sup> youssef mebarki, Personal Communications, IAR, 2010, 2010.
- <sup>6</sup> J. Gregory, K. Asai, M. Kameda, T. Liu, and J. Sullivan, A review of pressure-sensitive paint for high-speed and unsteady aerodynamics, *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* **222**, 249–290 (2008).
- <sup>7</sup> T. Liu, M. Guille, and J. Sullivan, Accuracy of pressure-sensitive paint, *AIAA journal* **39**, 103–112 (2001).

- 
- <sup>8</sup> Y. Le Sant, A. Durand, and M. Mérienne, Image Processing Tools Used for PSP and Model Deformation Measurements, AIAA Paper 5007 (2005).
  - <sup>9</sup> R. Shanmugasundaram and J. Samareh-Abolhassani, Modified scatter data interpolation used to correct pressure sensitive paint images, in *AIAA, Thermophysics Conference, 30 th, San Diego, CA, 1995*.
  - <sup>10</sup> K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Gool, A comparison of affine region detectors, *International journal of computer vision* **65**, 43–72 (2005).
  - <sup>11</sup> H. Bay, T. Tuytelaars, and L. Van Gool, Surf: Speeded up robust features, *Computer Vision–ECCV 2006* , 404–417 (2006).
  - <sup>12</sup> J. Bell and B. McLachlan, Image registration for pressure-sensitive paint applications, *Experiments in Fluids* **22**, 78–86 (1996).
  - <sup>13</sup> J. Crafton, S. Fonov, E. Jones, L. Goss, and C. Tyler, Simultaneous Measurements of Pressure and Deformation on a UCAV in the SARL, in *11th Annual Flow Visualization Congress, Notre Dame University, 2004*.
  - <sup>14</sup> B. McLachlan, J. Kavandi, J. Callis, M. Gouterman, E. Green, G. Khalil, and D. Burns, Surface pressure field mapping using luminescent coatings, *Experiments in Fluids* **14**, 33–41 (1993).
  - <sup>15</sup> Y. Le Sant, A model deformation method applied to PSP measurements, ONERA, TP no. 2004-46 **2004** (2004).
  - <sup>16</sup> W. Ruyten, A. I. O. AERONAUTICS, and A. N. YORK, *Toward an Integrated Optical Data System for Wind Tunnel Testing*, (1999).

- 
- 17 Y. Le Sant and M. Merienne, An image resection method applied to mapping techniques, in *Instrumentation in Aerospace Simulation Facilities, 1995. ICIASF'95 Record., International Congress on*, page 46, IEEE, 2002.
  - 18 B. Zitova and J. Flusser, Image registration methods: a survey, *Image and vision computing* **21**, 977–1000 (2003).
  - 19 F. Cazals and J. Giesen, Delaunay triangulation based surface reconstruction, *Effective Computational Geometry for Curves and Surfaces* , 231–276 (2006).
  - 20 O. Devillers, Filtering Relocations on a Delaunay Triangulation, (2009).
  - 21 L. Venkatakrisnan, Comparative Study of Different Pressure-Sensitive-Paint Image Registration Techniques, *AIAA journal* **42** (2004).
  - 22 C. Evans, Notes on the OpenSURF Library, (2009).
  - 23 O. Juan, Luo Gwon, A comparison of SIFT, PCA-SIFT and SURF, (2007).
  - 24 P. Koshevoy, T. Tasdizen, and R. Whitaker, Automatic assembly of TEM mosaics and mosaic stacks using phase correlation, Technical report, SCI Institute Technical Report, No. UUSCI-2007-004, University of Utah, 2007.
  - 25 B. Horn and B. Schunck, Determining optical flow, *Artificial intelligence* **17**, 185–203 (1981).
  - 26 S. Beauchemin and J. Barron, The computation of optical flow, *ACM Computing Surveys (CSUR)* **27**, 466 (1995).
  - 27 J. Barron, D. Fleet, and S. Beauchemin, Performance of optical flow techniques, *International journal of computer vision* **12**, 43–77 (1994).

- 
- <sup>28</sup> A. Bruhn, J. Weickert, and C. Schnorr, Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods, *International Journal of Computer Vision* **61**, 211–231 (2005).
- <sup>29</sup> M. J. Black, *Image Sequences for Optical Flow*, 2010.
- <sup>30</sup> P. Burt and B. Julesz, A disparity gradient limit for binocular fusion, *Science* **208**, 615–617 (1980).
- <sup>31</sup> H. Trivedi and S. Lloyd, The role of disparity gradient in stereo vision, *Perception* **14**, 685–690 (1985).
- <sup>32</sup> M. Shortis, T. Clarke, and T. Short, A comparison of some techniques for the subpixel location of discrete target images, *Videometrics III* **2350** (1994).
- <sup>33</sup> J. Lewis, Fast normalized cross-correlation, in *Vision Interface*, volume 10, pages 120–123, Citeseer, 1995.
- <sup>34</sup> R. Frischholz and K. Spinnler, A class of algorithms for real-time subpixel registration, *Proceedings of SPIE, Computer Vision for Industry* **1989**, 50–59 (1993).
- <sup>35</sup> G. Bradski, The OpenCV library, *Doctor Dobbs Journal* **25**, 120–126 (2000).
- <sup>36</sup> H. Ochis and E. Russell, Comparison of a piecewise transformation to polynomial-based geometric correction algorithms, Master's thesis, Colorado State University, 1997.
- <sup>37</sup> S. Park and H. Sung, Assessment Of Image Registration Interpolation Methods For Pressure-sensitive Paint Measurements, in *11th International Symposium on Flow Visualization, August*, volume 9, page 12, 2004.

- 38 Nokia, Qt Corss-Platform User Interface Toolkit, 2010.
- 39 G. Bradski and A. Kaehler, *Learning OpenCV: Computer vision with the OpenCV library*, O'Reilly Media, 2008.
- 40 Microsoft, Visual C++ Express Integrated Development Environment, 2009.
- 41 willowgarage.com, OpenCV Computer Vision Library, 2010.
- 42 F. Warmerdam, A. Kiselev, B. Friesenhahn, J. V. Damme, and L. Howard, LibTiff tagged image format processing library, 2010.
- 43 P. W. Olszta, A. Umbach, and S. Baker, FreeGLUT open source implementation of OpenGL, 2010.
- 44 Nokia, Qwt Graphing and visualization toolkit for Qt, 2010.
- 45 H. Sung, S. Park, and M. Kim, Accuracy of correlation-based image registration for pressure-sensitive paint, *Experiments in Fluids* **39**, 630–635 (2005).
- 46 S. Park and H. Sung, Correlation-based image registration for applications using pressure-sensitive paint, *Aiaa Journal* **43** (2005).