

University of Southampton Research Repository
ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON

**Statistical Modelling and Variability of
the Subtropical Front, New Zealand**

by

Joanne E. Hopkins

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the
Faculty of Engineering, Science and Mathematics
School of Ocean and Earth Science

August 2008

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF OCEAN AND EARTH SCIENCE

Doctor of Philosophy

Statistical Modelling and Variability of the Subtropical Front, New Zealand

by Joanne E. Hopkins

Ocean fronts are narrow zones of intense dynamic activity that play an important role in global ocean-atmosphere interactions. Of particular significance is the circumglobal frontal system of the Southern Ocean where intermediate water masses are formed, heat, salt, nutrients and momentum are redistributed and carbon dioxide is absorbed. The northern limit of this frontal band is marked by the Subtropical Front, where subtropical gyre water convergences with colder subantarctic water. Owing to their highly variable nature, both in space and time, ocean fronts are notoriously difficult features to adequately sample using traditional *in-situ* techniques. We therefore propose a new and innovative statistical modelling approach to detecting and monitoring ocean fronts from AVHRR SST images. Weighted local likelihood is used to provide a non-parametric description of spatial variations in the position and strength of individual fronts within an image. Although we use the new algorithm on AVHRR data it is suitable for other satellite data or model output.

The algorithm is used to study the spatial and temporal variability of a localized section of the Subtropical Front past New Zealand, known locally as the Southland Front. Twenty-one years (January 1985 to December 2005) of estimates of the front's position, temperature and strength are examined using cross correlation and wavelet analysis to investigate the role that remote atmospheric and oceanic forcing relating to the El Niño-Southern Oscillation may play in interannual frontal variability.

Cold (warm) anomalies are observed at the Southland Front three to four months after peak El Niño (La Niña) events. The gradient of the front changes one to two seasons in advance of extreme ENSO events suggesting that it may be used as a precursor to changes in the Southern Oscillation. There are strong seasonal dependencies to the correlation between ENSO indices and frontal characteristics. In addition, the frequency and phase relationships are inconsistent indicating that no one physical mechanism or mode of climate variability is responsible for the teleconnection.

Acknowledgements

Thanks to my supervisors Peter Challenor and Andrew Shaw for their help and advise throughout this thesis. Thanks also to Alex Mustard, Anna Hickman, Becky Bell, Chris Jeffery, Chris Nicolle, Dave Poulter, Doug McNeill, Kirsty Edgar, Mark Vardy, Phil Wallhead and Sarah Bennett all of whom have contributed significantly to making the last three years so memorable. I appreciate all the advise and support from those within the Ocean Observing and Climate, and Ocean Modelling and Forecasting Groups at the National Oceanography Centre, Southampton. Lastly, thanks to my parents, Liz and Dave for all their encouragement.

Contents

Acronyms and Abbreviations	xi
Symbols	xiii
1 Introduction	1
1.1 Overview	1
1.2 Significance of Ocean Fronts	2
1.3 Research Objectives	5
1.4 Thesis Outline	5
2 Oceanography and Hydrodynamics of Ocean Fronts	8
2.1 Introduction	8
2.2 Continental Slope and Shelf Break Systems	10
2.2.1 Taylor-Proudman Theorem	10
2.2.2 Boundary Upwelling Fronts	11
2.2.3 Flow Nonlinearity	12
2.2.4 Tidal Mixing Fronts	13
2.2.5 Shelf Break Fronts	14
2.3 Open Ocean Convergences in the Southern Ocean	15
2.3.1 Topographic Steering and Vorticity	15
2.3.2 Enhanced Productivity at the Subtropical Front	18
2.3.3 Southern Ocean Subtropical Front	18
2.4 Physical Oceanography of New Zealand	20
2.4.1 Introduction	20
2.4.2 Frontal Systems and Circulation	22
2.4.2.1 North East Shelf	22
2.4.2.2 South West Shelf and Cook Strait	22
2.4.2.3 South East Shelf: The Subtropical Front	23
2.4.2.4 The Southland Front and Current	25
3 Ocean Front Detection: An Overview	30
3.1 Derivative Based Edge Detection	31
3.2 Automated Classification and Edge Discrimination	34

3.3	Gradient Magnitude Criteria	41
3.4	Surface Fitting	42
3.5	Supervised Learning Techniques	44
3.6	Ocean Color and Altimetry in Front Detection	44
4	Statistical Modelling of Ocean Fronts	46
4.1	Algorithm Design Objectives	46
4.2	A Statistical Ocean Front Model	48
4.2.1	Maximum Likelihood Theory	48
4.2.2	A Model Ocean Front	51
4.2.3	Maximum Likelihood Properties and Confidence Bounds	51
4.3	Modelling Spatial Variability	56
4.3.1	Regression Analysis	56
4.3.1.1	Parametric Regression	57
4.3.1.2	Weighted Non-Parametric Regression	57
4.3.2	Local Likelihood Estimation	60
4.4	Bandwidth Selection	63
4.4.1	Bias and Variance Considerations	63
4.4.2	Cross Validation	65
4.4.3	Likelihood Cross Validation	66
4.4.4	Bayesian Information Criterion	66
4.5	Summary	68
5	Practical Aspects of Front Detection	70
5.1	Data Acquisition and Preparation	70
5.1.1	AVHRR-SST Pixel Quality	72
5.1.2	Rotation and Extraction of SST Data	75
5.2	Practical Aspects of Front Detection	76
5.2.1	Optimization	76
5.2.1.1	Newton-Raphson Optimization	77
5.2.1.2	Transformation of Variables and Parameter Bounds	78
5.2.1.3	Convergence Criteria	79
5.2.2	Bandwidth Selection	81
5.2.3	Quality Control	85
5.2.3.1	Quality Control Criteria	87
5.3	Implementation: Stability, Sensitivity and Limitations	91
5.3.1	Sensitivity to Initialization Conditions	92
5.3.2	Noise Limitation	92
5.3.3	Curvature Limitation	95
5.3.4	Sparsity of Observations	97
5.4	Summary	102
6	Spatial and Seasonal Variability of The Southland Front	103
6.1	Introduction	103
6.2	Methods	103
6.2.1	Weighted Statistics	104
6.2.2	Gradient and Error Propagation	105

6.3	Results	106
6.3.1	Southland Front Characteristics and Position	106
6.3.2	Seasonal SST Cycle	114
6.4	Discussion	115
6.4.1	Comparison to Alternative Techniques	116
6.4.2	Data Resolution Considerations	120
6.4.3	Temperature	124
6.4.4	Temperature Range	125
6.4.5	Width	126
6.4.6	Gradient	127
6.4.7	Frontal Position and Stability	128
6.5	Summary	132
7	Southland Front-ENSO Teleconnections	133
7.1	Introduction	133
7.1.1	The El Niño-Southern Oscillation	133
7.1.2	ENSO: a New Zealand Perspective	135
7.1.3	The Southland Front and ENSO: The Questions	136
7.2	Data Sets	136
7.3	Methods	138
7.3.1	Cross Correlation	138
7.3.1.1	Significance Levels	140
7.3.2	Wavelet Analysis	140
7.3.2.1	Local Wavelet Power Spectrum	142
7.3.2.2	Global Wavelet Spectrum	144
7.3.2.3	Scale Averaged Wavelet Power	145
7.3.2.4	Cross Wavelet Power	145
7.3.2.5	Wavelet Coherence	146
7.3.2.6	Phase Relationship	146
7.4	Results and Discussion	147
7.4.1	Temperature of the Southland Front and ENSO	147
7.4.1.1	Seasonal ENSO Correlations	148
7.4.1.2	Oceanic Advection of SST Anomalies	151
7.4.1.3	A Variable Phase Relationship	153
7.4.1.4	PDO Modulation of ENSO Teleconnections	156
7.4.2	Strength of the Southland Front and ENSO	159
7.4.2.1	Impact of Mernoo Saddle Currents	159
7.4.2.2	An Inconsistent Teleconnection	160
7.4.2.3	Upper Layer Heat Content, Sea Level and Current Speeds	164
7.4.2.4	Seasonal Correlations	166
7.4.2.5	A Precursor to ENSO events?	169
7.4.2.6	The Antarctic Oscillation Teleconnection	171
7.4.3	Location of the Southland Front and ENSO	173
7.5	Summary	174
8	Conclusions and Further Work	177
8.1	Statistical Modelling of Ocean Fronts	177

8.2 Southland Front Variability	181
A Appendix	185
A.1 Hessian matrix for local likelihood estimates	185
A.2 AVHRR Pixel Quality	187
A.3 Practical Aspects of Front Detection	188
A.4 Southland Front Variability	189
A.5 ARIMA Model Fitting	191
A.6 Southland Front and ENSO	193
Bibliography	194

List of Figures

2.1	Boundary upwelling front dynamics	11
2.2	Residual current driven upwelling	12
2.3	Tidal mixing front dynamics	14
2.4	Southern Ocean SeaWiFS Chlorophyll concentrations	16
2.5	Open ocean frontal dynamics	17
2.6	Bathymetry and circulation around New Zealand	21
2.7	Bathymetry, topography and circulation around North and South Islands, New Zealand	24
2.8	Mean March SST around New Zealand	26
2.9	Salinity cross-section of the Southland Front	27
3.1	Sobel, Prewitt, LoG and Canny edge detection	33
3.2	Segmentation of data using Fishers Linear Discriminant	35
3.3	Construction of the gray-level occurrence matrix	36
3.4	Flow chart of the SIED algorithm	37
3.5	Temperature histogram analysis during SIED	38
3.6	Calculation of cohesion coefficients during SIED	39
3.7	Fronts detected in the Gulf Stream using a multi-image algorithm	40
3.8	Multi-spectral front map of the Western Approaches, English Channel	41
3.9	The ‘front following’ algorithm	43
4.1	Schematic illustrating normal distribution assumption of SST observations	50
4.2	Physical representation of model parameters	52
4.3	Standard maximum likelihood model function fitting	55
4.4	Kernel smooth estimate versus a linear regression	59
4.5	Graphic representation of constructing local likelihood estimates	62
4.6	Bias-variance control of the bandwidth	65
5.1	AVHRR-SST pixel quality	74
5.2	Regridding, rotation and extraction of data	75
5.3	Comparison of bandwidth selection methods	82
5.4	Local likelihood estimates for varying bandwidths	83
5.5	Distribution of optimal bandwidths	84
5.6	Data quality flag assignment flowchart	88

5.7	Analysis of standard maximum likelihood fit and convergence criteria I	89
5.8	Analysis of standard maximum likelihood fit and convergence criteria II	90
5.9	Sensitivity to optimization initialization point	93
5.10	Robustness of algorithm to noise	94
5.11	Orientation of an artificial front	95
5.12	Effect of frontal orientation on parameter estimates	96
5.13	Percentage of missing data	97
5.14	Robustness to sparse data sets I	100
5.15	Robustness to sparse data sets II	101
6.1	Mean twenty-one year characteristics along the Southland Front	107
6.2	Mean twenty-one year position of the Southland Front	108
6.3	Seasonal variability along the Southland Front	109
6.4	Seasonal positions of the Southland Front	111
6.5	Location of the Southland Front with respect to depth	112
6.6	Seasonal meandering intensity along the Southland Front	113
6.7	Regional surface water mass locations	114
6.8	Location of Southland Front on 28th March 1990 as estimated by Shaw and Vennell (2000a), the local likelihood algorithm and the SIED.	117
6.9	Parameter estimates along the Southland Front on 28th March 1990 made by Shaw and Vennell (2000a) and the local likelihood algorithm.	119
6.10	Effects of temporal blurring on algorithm estimates	121
6.11	Location of Southland Front on 28th March 1990 from 1km and 4km resolution data sets	122
6.12	Characteristics of the Southland Front on 28th March 1990 from 1km and 4km resolution data sets	123
6.13	Histograms of absolute differences in estimated month-to-month positions	129
6.14	Potential temperature, salinity and potential density structure off Oamaru	131
7.1	Map of Niño SST Regions	137
7.2	Cross correlation between the SOI and mean frontal temperature	140
7.3	The Morlet Wavelet	141
7.4	Local wavelet power spectrum for Niño3 SST	142
7.5	Local wavelet power spectrum for the SOI	143
7.6	Local wavelet power for frontal temperature anomalies	144
7.7	Latitude-time plot of frontal temperature anomalies	148
7.8	Seasonal correlation patterns between MSLP and SOI	150
7.9	Spatially averaged seasonal SOI-SST correlations	152
7.10	Wavelet coherence between mean frontal temperature and ENSO indices .	154
7.11	Scale averages wavelet power for the SOI and SFT	155
7.12	Warm and cold phases of the PDO and ENSO	157
7.13	Smoothed time series of temperature anomalies, the SOI and PDO	158
7.14	Latitude-time plot of frontal gradient	160
7.15	Wavelet coherence between SFG and ENSO	162
7.16	Smoothed time series of the gradient, temperature difference, width, SOI and Niño3.4 SST	163
7.17	Seasonally averaged cross correlations between the temperature, gradient, temperature difference and ENSO indices	167

7.18	Seasonal correlation matrices	168
7.19	Schematic of the ‘Van Loon Hypothesis’	170
7.20	Wavelet coherence between the AAO and frontal characteristics	172
7.21	Interannual variability in frontal position and stability	174
A.1	Flowchart to derive overall AVHRR pixel quality	187
A.2	Orientation of the Southland Front as estimated by Shaw (1998)	188
A.3	Percentage of missing data in each processed AVHRR SST image	188
A.4	Mean three year characteristics along the Southland Front	189
A.5	Three year seasonal variability along the Southland Front	190
A.6	Wavelet coherence between Southland Front width and SOI	193
A.7	Interannual variability in frontal position and RMSSD	193

List of Tables

2.1	New Zealand watermass characteristics	20
3.1	First derivative convolution masks	32
5.1	Scaling bounds and transformation of variables	78
5.2	Interpretation of IFAIL values	80
5.3	Data quality flag weightings	86
5.4	Percentage of estimates on upper or lower bounds	91
5.5	Impact of sparse data sets on standard errors	98
6.1	Weighted mean twenty-one year and seasonal characteristics of the Southland Front	110
6.2	Seasonal deviation and meandering intensity at the Southland Front	113
6.3	Annual SST range and phase of surface water masses	115
6.4	Parameter estimates over sections A-C on March 28th 1990	118
6.5	Three year mean characteristics from Shaw and Vennell (2001) and the local likelihood algorithm	125
7.1	Seasonal correlations between ENSO and frontal temperature	149
7.2	Phase relationship between Niño SST's and frontal temperature anomalies	156
7.3	Seasonal lag correlations of gradient and ENSO indices	169
A.1	AVHRR pixel quality tests	187
A.2	Mean three year seasonal characteristics from Shaw and Vennell (2001) and the local likelihood algorithm	189

Acronyms and Abbreviations

AAO	Antarctic Oscillation
ACC	Antarctic Circumpolar Current
AMSR	Advanced Microwave Scanning Radiometer
AR	Autoregressive
ATSR	Along Track Scanning Radiometer
AVHRR	Advanced Very High Resolution Radiometer
BIC	Baysian Information Criterion
COI	Cone of Influence
CVSS	Cross Validation Sum of Squares
CV	Cross Validation
CWT	Wavelet Coherence
DJF	December-January-February (Summer)
EAC	East Australia Current
ENSO	El Niño-Southern Oscillation
EOF	Empirical Orthogonal Function
ERS	European Remote Sensing
FASINEX	Frontal Air-Sea Interaction Experiment
GHRSSST	Global High-Resolution Sea Surface Temperature
IR	Infrared
JJA	June-July-August (Winter)
LCV	Likelihood Cross Validation
MAM	March-April-May (Autumn)
MJO	Madden-Julian Oscillation
MLE	Maximum Likelihood Estimate
MODIS	Moderate Resolution Imaging Spectroradiometer
MSE	Mean Squared Error

MSLP	Mean Sea Level Pressure
NIWA	National Institute of Water and Atmospheric Research
NOAA	National Oceanic and Atmospheric Administration
OCCAM	Ocean Circulation and Climate Advanced Modelling
PDO	Pacific Decadal Oscillation
PSD	Power Spectral Density
RMSD	Root-Mean-Square Spatial Displacement
RSS	Residual Sum of Squares
SAW	Subantarctic Water
SeaWiFS	Sea-viewing Wide Field-of-view Sensor
SF	Southland Front
SFG	Southland Front Gradient
SFT	Southland Front Temperature
SFTD	Southland Front Temperature Difference
SFW	Southland Front Width
SIED	Single Image Edge Detector
SL	Sea Level
SLP	Sea Level Pressure
SMMR	Scanning Multichannel Microwave Radiometer
SOI	Southern Oscillation Index
SON	September-October-November (Spring)
SSS	Sea Surface Salinity
STCZ	Subtropical Convergence Zone
STF	Subtropical Front
STW	Subtropical Water
SST	Sea Surface Temperature
SURVOSTRAL	Surveillance de l'Océan Austral
TOPEX	Topography Experiment
TRMM	Tropical Rainfall Measuring Mission
WOCE	World Ocean Circulation Experiment
XWT	Cross Wavelet Spectrum

Symbols

Z	Vector of temperature observations $\mathbf{Z} = \{z^{(1)}, z^{(2)} \dots, z^{(n)}\}$. [°C]
Y	Vector of distances $\mathbf{Y} = \{y^{(1)}, y^{(2)} \dots, y^{(n)}\}$. [arc degrees]
H	Fisher information matrix (second order partial derivatives)
g	Gradient vector. Partial derivatives of function $F(\cdot)$
d	Vector of search directions
S	Matrix of weightings
QW	Vector of quality weightings
$p(\cdot)$	Probability density function
$l(\cdot)$	Likelihood function
$L(\cdot)$	Log likelihood function
$\phi(\cdot)$	Normal distribution
$K(\cdot)$	Gaussian kernel (smoothing) function
$m(\cdot)$	Parametric model function
$\Phi(x)$	Smooth non-parametric function describing along-front trend w.r.t x
$F(\cdot)$	Negative log-likelihood function. Objective function to optimize
θ	Vector of model parameters $\theta = \{\theta_1, \theta_2, \theta_3, \theta_4, \sigma\}$
$\hat{\theta}$	Vector of maximum likelihood model parameter estimates $\hat{\theta} = \{\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\sigma}\}$
θ_1	Parameter. Mean model frontal temperature. [°C]
θ_2	Parameter. Half the temperature differential across the model front. [°C]
θ_3	Parameter. Half the width of the model front. [arc degrees]
θ_4	Parameter. Position of model front. [arc degrees]
σ (σ^2)	Parameter. Standard deviation (variance) of noise about model function $m(\cdot)$
ϵ	Normally distributed random noise with mean zero and standard deviation σ

v	Degrees of freedom
x	Along-front distance. [arc degrees]
h	Constant smoothing parameter (bandwidth). [arc degrees]
h_{op}	Optimal smoothing parameter. [arc degrees]
$[a_q, b_q]$	Lower and upper optimization bounds for q^{th} unknown model parameter
α	Distance stepped along search direction \mathbf{d}
θ_{tol}	Accuracy of estimated parameters
i, k, q	Sub- and superscripts referencing matrix elements
j	Spatial index
\bar{x}_w	Weighted mean
$\sigma_{\bar{x}_w}^2$	Variance of the weighted mean
$\bar{\sigma}_w^2$	Weighted variance
\prod	Product of
\sum	Sum of
\mathbb{E}	Expectation
\mathbb{V}	Variance

CHAPTER 1

Introduction

1.1 Overview

The ocean and atmosphere are dynamic, complex and intricately linked systems that help maintain the present global climate. Within the climate system large-scale ocean and atmosphere anomalies may be related to one another over vast geographical distances (typically thousands of kilometres). These teleconnection patterns occur naturally and may persist for several weeks, months or even years. They therefore form an important part of both interannual and interdecadal variability of the ocean and atmosphere systems. Perhaps the best known teleconnection is the ‘seesaw’ in sea-level pressure between Tahiti and Darwin, Australia, which defines the Southern Oscillation, the atmospheric component of the coupled ocean-atmosphere phenomenon known as the El Niño-Southern Oscillation (ENSO).

Ocean fronts, the narrow transition zones between different water masses, play a significant role in global ocean-atmosphere interactions, they are zones of enhanced primary production and sites of intense water mass formation, advection and mixing.

This thesis is concerned with the variability of the southern Subtropical Front (STF) as it passes New Zealand, one of the most interconnected regions of the world. The STF is the near continuous hydrographic boundary between warm, saline subtropical gyre waters and cooler, fresher subantarctic water found at approximately 40°S. As it passes New Zealand, where it is known locally as the Southland Front, it experiences extreme climatic and oceanic conditions owing to the islands’ links to both the Southern Ocean and south-west Pacific. Atmospheric circulation changes, sea level variability and sea

surface temperature anomalies around New Zealand, resulting from ENSO variability, are well documented. The response of the Southland Front however has not yet been fully explored. It has been suggested that the strength of the front decreased during the 1991 El Niño event (Shaw and Vennell, 2001), but with only three years worth of analysis no firm conclusions could be drawn. This thesis is motivated to further investigate possible teleconnections between ENSO and the Southland Front, the nature of which may influence the structure and strength of the global STF as it enters the open Pacific Ocean.

In order to conduct this research an objective means of monitoring the front's position and characteristics over extensive length and time scales is required. In a broader context, accurate knowledge of frontal zones and how they change in space and time is of growing importance to many organizations worldwide with such diverse tasks as climate variability, monitoring, seasonal forecasting, operational weather and ocean forecasting, validating ocean and atmospheric models, ecosystem assessment, and tourism and fisheries research.

There are two main components to this work; firstly the development of a front detection tool for use with remotely sensed sea surface temperature (SST) data; and secondly application of the new technique to investigate the variability and ENSO teleconnections of the Southland Front.

The following sections afford a more detailed look at the significance of ocean fronts within the climate system, and why development of front detection techniques is so valuable. The chapter concludes by setting out the key scientific questions to be addressed and an outline of the thesis structure.

1.2 Significance of Ocean Fronts

Ocean fronts are climatologically important features. They are responsible for the re-distribution of heat, salt, nutrients and momentum throughout the worlds oceans, particularly so in the Southern Ocean where frontal boundaries are associated with strong circumpolar flows that connect all three major ocean basins.

Fronts are sites of water mass convergence and are therefore potentially zones of active subduction and downwelling; important components of global ocean-atmosphere interaction. The surface layer of the ocean may exchange heat, fresh water and gases with the atmosphere. This exchange stops once the surface water is subducted into the ocean interior and properties set by the atmosphere are carried away by intermediate water masses that participate in global ocean circulation. For example, the heat content of South Atlantic Mode Water formed at the Subantarctic Front may affect the strength

of the South Atlantic subtropical gyre (Morris et al., 2001). Greenhouse gases may also be sequestered into the ocean via subduction at fronts.

Instabilities within frontal zones resulting in cross frontal mixing, upwelling and the exchange of different limiting nutrients between water masses give rise to some of the most unique and highly productive ecosystems on earth (Mann and Lazier, 1996). Enhanced biological activity at fronts supports some of the worlds largest fisheries; valuable catches of hoki, orange roughy and oreos come from the Subtropical Front around New Zealand (McClatchie et al., 2001). Furthermore, increased primary production driven by unstable meanders and eddies stimulates carbon export via the ‘biological pump’ (Allen et al., 2005).

Frontal zones are recognized to be a major sink for the increasing levels of carbon dioxide in the atmosphere. This is particularly true of the frontal zone between 40° and 60° in the northern and southern hemispheres where warm subtropical waters meet colder, higher latitude water masses. In these regions the partial pressure of dissolved carbon dioxide ($p\text{CO}_2$) is decreased by the cooling of poleward flowing warm waters, and by the biological drawdown of $p\text{CO}_2$ in subpolar water. High wind speeds over these waters further increases the CO_2 drawdown (Takahashi et al., 2002). The Subtropical Convergence over the Chatham Rise in the south-west Pacific has been identified as a carbon dioxide sink during austral spring (Murphy et al., 1991; Currie and Hunter, 1998, 1999). Greater primary productivity (increased Chlorophyll-*a* concentrations) in the frontal zone together with an increase in biological carbon assimilation is thought to be responsible for the increased carbon dioxide uptake rate.

High levels of eddy kinetic energy are coincident with the frontal systems of the Antarctic Circumpolar Current. Determination of the correct frontal positions is important for studies of eddy-mean flow interaction (Hughes and Ash, 2001) and poleward heat fluxes (Phillips and Rintoul, 2000). Slight shifts in location can reverse the eddy-mean flow relationship.

Variability of the strength and location of sea surface temperature fronts is important to the coupling of winds and upper-ocean processes (O’Neill et al., 2003; Chelton et al., 2004). Surface winds in the Marine Atmospheric Boundary Layer (MABL) are modified by sharp gradients in both the sea surface temperature and current speed. There is a decrease (increase) in wind speed when moving from warm (cold) to cold (warm) water (Song et al., 2006). Major ocean currents and frontal zones such as the Gulf Stream and Kuroshio Extension where there are prominent gradients in temperature and current speed can induce geostrophic jets in the MABL, the speeds of which are dictated by the intensity of the front. These oceanographic fronts change in strength and position on time scales of weeks to years which can force similar changes in the atmospheric fronts and flows, and hence in the free atmosphere above (Feliks et al., 2004).

The response of winds to sea surface temperature may also feedback into the ocean itself and modify the ocean current system. The near-surface wind stress curl and divergence in the Southern Ocean are linearly related to the crosswind and downwind components of the sea surface temperature gradient respectively (O'Neill et al., 2003). Wind stress curl is an important driving force behind Ekman vertical velocity implying that winds and sea surface temperature may have an important feedback effect on upper ocean processes near fronts.

For the many reasons outlined above a considerable amount of time and expense has been invested in ocean front focused research. Ocean temperature, salinity and currents are all readily measured by a vast array of *in-situ* sensors deployed from ships. Towed and moored CTD (Conductivity Temperature Depth) sensors, XBT's (Expendable Bathythermographs), free drifting ARGO floats, ADCP's (Acoustic Doppler Current Profilers) and water bottle samples all provide information about the structure, dynamics, biology and chemistry of frontal zones. Many long-term field programmes and individual cruises have been and continue to be conducted aimed specifically at determining ocean front variability, structure and atmospheric interaction, e.g. FASINEX, SURVOSTRAL, Drake Passage Repeat Hydrography, and the Subantarctic Zone Project to name but a few. As such, the position of fronts has traditionally been ascertained from ship transects. The climatological mean surface and subsurface expressions of the Polar Front in the Southern Ocean have been examined in a number of studies using historical *in-situ* measurements (Orsi et al., 1995; Belkin and Gordon, 1996). More recently, high resolution repeat hydrographic sections have enabled the time varying component of frontal structures to be better resolved (Sprintall, 2003; Cunningham et al., 2003). However the large, highly variable nature of these synoptic features is difficult to capture using relatively sparse shipboard measurements. For example, the position of the fronts will vary according to which ship track is taken, and even change during the survey, making the temporal and spatial components of their variability difficult to partition.

Sensors onboard orbiting satellites on the other hand are able to capture snapshots of the surface expression of fronts over extensive areas at regular time intervals providing oceanographers with the most complete view to date of ocean front dynamics. The Southern Ocean frontal system has been the focus of many studies using all flavors of remote sensing data. Infrared and microwave sensors (e.g. AVHRR, ATSR, AMSR) pick out the sharp gradient in SST that marks the transition between water masses (Moore et al., 1997, 1999; Dong et al., 2006). SeaWiFS measures the optical properties of the ocean surface and is used to identify zones of increased chlorophyll concentration that are associated with frontal structures (Moore and Abbott, 2000). Lastly, altimeters detect changes in sea surface height that drive localized frontal currents (Gille, 1994; Gille and Kelly, 1996).

The techniques with which fronts may be automatically detected from remotely sensed SST images are wide ranging. Few however are well suited to a study of the relationship

between ocean fronts and modes of climate variability as proposed here. Firstly, the majority of front detection techniques are ‘non-specific’ i.e. they do not target any one particular structure. Furthermore, extracting information about the strength of identified features is not always possible. The ‘front following’ algorithm developed by Shaw and Vennell (2000a) was the first attempt at targeting and then tracking individual fronts. Additionally, the algorithm made measurements of the front’s mean temperature, the step change in temperature from one side to the other, and the distance over which this transition occurred (i.e. frontal width). A statistical model fitting approach such as this has the potential to generate a detailed data set of frontal characteristics ideally suited to teleconnection studies.

1.3 Research Objectives

The first half of this thesis aims to develop a front detection algorithm based upon statistical modelling techniques for use with satellite images of the sea surface temperature. The new approach should be robust to poor quality (noisy) data sets and missing temperature measurements, two problems that often hamper current front detection approaches. Additionally, the ability to quantify the uncertainty around estimates made by the algorithm is desirable. As yet, no other method is able to do this.

The following scientific questions concerning the Southland Front will then be addressed:

1. How and why do the characteristics of the Southland Front change spatially and seasonally?
2. Is there a persistent teleconnection between the temperature of the Southland Front and ENSO?
3. Is the strength of the Southland Front modified by ENSO variability?
4. Is there any interannual variability in the position and stability of the Southland Front?

1.4 Thesis Outline

Work here is organized into a further six chapters:

Chapter 2

In the first of two background chapters an overview of the oceanography and hydrodynamics of continental slope, shelf break and open ocean convergence fronts is given. This highlights the wide range of length and time scales over which ocean fronts may persist

and the dynamical processes governing their formation. Particular emphasis is placed on the Subtropical Front and the frontal systems and circulation around New Zealand.

Chapter 3

This chapter reviews the existing techniques used to automatically detect ocean fronts from satellite imagery. Six different approaches are identified and discussed: derivative based edge detection, automated classification and edge discrimination techniques, gradient magnitude criteria, surface fitting, supervised learning and the use of ocean color and altimetry.

Chapter 4

The statistical theory of weighted local likelihood as a means of providing a non-parametric description of spatial variation in the position and strength of ocean fronts from SST images is outlined. The reader is firstly introduced to the concept of maximum likelihood and how a model function may be fitted to a set of SST observations. Regression analysis and kernel smoothing techniques are then extended into the idea of local likelihood which forms the basis of a new innovative front detection algorithm. Lastly, the role and importance of selecting an appropriate smoothing function is discussed.

Chapter 5

Following on from the theoretical introduction to local likelihood, Chapter 5 deals with the practicalities of applying this technique to the Southland Front, New Zealand. The acquisition and preparation of the AVHRR SST data set used in subsequent chapters is described as is the Newton-Raphson optimization scheme used to obtain local likelihood estimates. A series of tests on simulated data sets demonstrates the algorithms abilities and assess its limitations.

Chapter 6

The front detection model is used to present a spatial and seasonal variability study of the Southland Front. Results are discussed in relation to the findings of other authors and alternative techniques.

Chapter 7

Chapter 7 examines the interannual variability of the Southland Front with respect to the El Niño-Southern Oscillation. Cross correlation and wavelet analysis are used to identify the nature and stability of the relationship between extreme ENSO variability and the structure of the Southland Front. Findings are discussed within the context of previous investigations into the regional ocean-atmosphere response to El Niño and La Niña events. Potential mechanisms explaining the results are proposed.

Chapter 8

The final chapter of this thesis summarizes results and findings and discusses possible avenues for future research.

Computing

All data preparation and analysis of the model results is undertaken in MATLAB. Wavelet software was provided by A.Grinsted and is available at URL: <http://www.pol.ac.uk/home/research/waveletcoherence/>. The algorithm itself is coded in FORTRAN 90 and the Newton-Raphson routine (E04JYF) taken from the NAG Fortran Library, Mark 21 is used for optimization. All maps were prepared in GMT.

CHAPTER 2

Oceanography and Hydrodynamics of Ocean Fronts

2.1 Introduction

Throughout the world's oceans variables such as temperature, salinity and density do not vary smoothly with horizontal distance. Instead we observe narrow high gradient regions where water masses of different temperatures, salinities, chemical or biological properties meet and interact separated by regions where horizontal variations are small. In addition to a marked change in horizontal property gradients these narrow frontal zones are also characterized by a change in vertical structure.

Fronts are formed from a wide variety of processes and are therefore observed across a wide range of scales. Small transient fronts are generated and destroyed everyday in the upper layers of coastal waters. They are typically only a few kilometres long and separate waters that differ by only 1-2°C. At the other extreme, large open ocean convergences in the Southern Ocean persist for thousands of kilometres around the Antarctic continent and exhibit much greater temperature and salinity differences. Over the Crozet Plateau region between 35°E and 60°E in the Southern Ocean the Crozet Front exhibits temperature and salinity ranges of 11°C and 1.8psu (Belkin and Gordon, 1996). The property values within these fronts change both temporally and spatially as a result of the gradual modification of the adjacent water masses by across frontal mixing, vertical transport and non-frontal processes, such as air-sea interaction.

Fronts may be divided into two distinct categories. Dynamically *active density fronts* are defined as regions where strong density gradients are capable of inducing an along-front geostrophic current - a flow balancing the across-front pressure gradient that results from

inclined density surfaces or sea surface elevation gradients, with the coriolis acceleration. At *passive fronts* there are changes in chemical and/or biological tracer properties rather than density gradients. At density compensated fronts, where there is no density gradient and therefore no pressure gradient forces there is no dynamic locally driven frontal development and destruction implied by the momentum equations. In these situations a ‘biological front’ exists in the absence of a physical one (Olson, 2002).

The dynamics of persistent upper ocean density fronts are governed by the ratio between turbulent transport processes (e.g. entrainment) and the Rossby Radius of deformation - the length scale at which the earth’s rotational effects become as important as buoyancy in the evolution of a flow (Garvine, 1979). On the smallest scale (order of 100m) the dynamics of *riverine plume* (e.g. Mississippi Plume) and *tidal intrusion fronts* (e.g. James River Estuary, Virginia) are dominated by turbulent dissipative processes. Plume fronts form where fresh river or estuarine outflow meets more saline coastal waters. The less saline and therefore less dense river outflow is carried seaward over the top of the coastal waters forming a steeply sloping pycnocline and vigorous convergence and downwelling. Tidal intrusion fronts occur within estuaries where local changes in depth give rise to small scale variations in the degree of tidal mixing. The fronts that form between mixed and stratified waters are created and destroyed with every tide.

At the other extreme, large scale *eastern* (e.g. California and Benguela Current Fronts) and *western* (e.g. Gulf Stream, Kuroshio, and Agulhas Current Fronts) *boundary upwelling fronts*, driven by large scale circulation and offshore Ekman flow, are governed by the effects of rotation. They have length scales of the order of 100km. Also found within coastal waters are *tidal mixing* and *shelf break fronts* (e.g. Georges Bank, Ushant Front, Celtic Sea Front) where sharp density gradients separate well mixed and highly stratified waters over the continental slope and shelf. The largest frontal systems are the *open ocean convergences* separating two different water masses such as the the Polar, Subantarctic and Subtropical Fronts found within the Southern Ocean. The mean widths of these fronts are correlated with the size of the baroclinic Rossby Radius (Gille, 1994), the horizontal length scale for stable mesoscale features.

This background chapter is divided into three main parts. Firstly we examine the biophysical dynamics of shelf seas and the three main frontal systems found in them: boundary upwelling fronts, tidal mixing fronts and shelf break fronts (Section 2.2). Secondly we examine open ocean convergence fronts and provide details of the Southern Ocean Subtropical Front (Section 2.3). Lastly we focus on a localized section of the Subtropical Front as it flows along the continental shelf break of South Island, New Zealand (Section 2.4). The Subtropical Front in this area is known locally as the Southland Front and is the example used in this thesis.

2.2 Continental Slope and Shelf Break Systems

We begin our discussion of continental slope and shelf break systems by considering how the equations of motion in the open ocean are modified at the continental slope. Shelf-ocean connectivity is expressed by the Taylor-Proudman Theorem.

2.2.1 Taylor-Proudman Theorem

Consider a steady, linear, frictionless flow. In this situation the equations of motion reduce to a state of geostrophic balance:

$$\begin{aligned}\frac{\partial \psi}{\partial t} - fv &= -\frac{1}{\rho} \frac{\partial P}{\partial x} + F_x \\ \frac{\partial \psi}{\partial t} + fu &= -\frac{1}{\rho} \frac{\partial P}{\partial y} + F_y\end{aligned}\tag{2.1}$$

There are no time varying components to the flow ($\frac{\partial u}{\partial t}, \frac{\partial v}{\partial t}$), no frictional or tidal forces (F_x, F_y), and no second order or squared terms. ρ is the sea water density. At geostrophic equilibrium, the horizontal pressure gradient forces ($\frac{\partial P}{\partial x}, \frac{\partial P}{\partial y}$) are balanced by the coriolis force (f) acting at right angles to any moving particle as a result of the earth's rotation ($-fv, fu$) (Pond and Pickard, 1983). The resulting geostrophic current always follows isopycnals (lines of constant density).

Considering also the hydrostatic and advection equations and the continuity of volume, the Taylor-Proudman theorem states that at the shelf edge where the bathymetry is rapidly changing, *the current speed can vary with depth but the direction of the current is always along isobaths* (Brink, 1998). This means that at the shelf break where steep density gradients separate well mixed and stratified waters geostrophic flow can not move across the shelf edge or slope, only along it.

If any of the assumptions involved in the Taylor-Proudman theorem are relaxed (linear, frictionless, steady) then mass or momentum may be exchanged across isobaths - cross-shelf edge exchange. We will now examine those cases where the theorem is violated and cross shelf edge flows are permitted. Firstly, by wind driven flow along a coastal boundary that leads to an upwelling front. Secondly, where departures from linearity caused by tidal currents and topographic interactions, for example lead to a break down in geostrophy, and the growth of nonlinear meanders, eddies and vertical transports. In this way, heat, salt, nutrients and momentum are transferred across the interface between stratified and well mixed waters at tidal mixing and shelf break fronts.

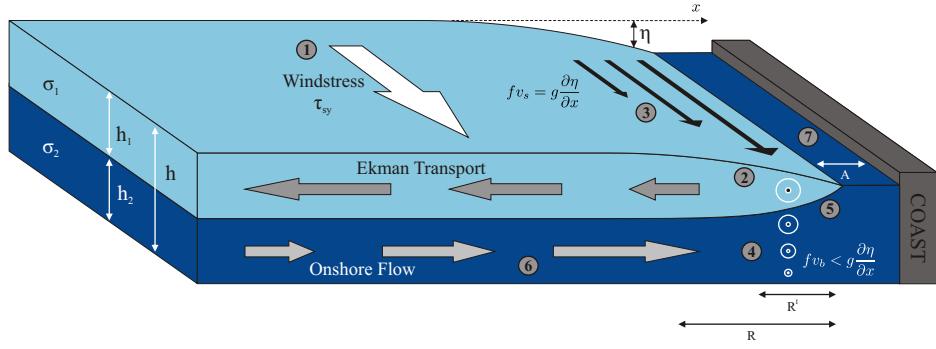


FIGURE 2.1: Wind driven boundary upwelling front dynamics (N.Hemisphere). Numbers relate to stages described in the main text. h_1 , ρ_1 and h_2 , ρ_2 are the upper and lower layer depths (in m and kg.m^{-3} respectively). h is the total depth (m) and ρ the mean density (kg.m^{-3}) of the water column. f is the coriolis frequency (s^{-1}) and g is gravity (m.s^{-2}). v_s and v_b are the surface and near bed along shore velocities (m.s^{-1}). $R = \sqrt{\frac{gh}{f}}$ is the Rossby Radius of Deformation (m). $R^I = [g \frac{\Delta \rho}{\rho_2} \frac{h_1 h_2}{h_1 + h_2}]^{\frac{1}{2}} / f$ is the Baroclinic (internal) Rossby Radius (m). A is the distance (m) between the frontal zone and the coast. The offshore Ekman Flux of water = $\frac{\tau_{sy}}{\rho f} (\text{m}^2 \cdot \text{s}^{-1})$.

2.2.2 Boundary Upwelling Fronts

Boundary upwelling fronts provide the first example where cross-shelf edge exchange develops. The sequence of events leading to an upwelling front, summarized in Figure 2.1, is as follows:

1. Alongshore wind stress (τ_{sy}) drives an alongshore current and an offshore Ekman transport of water.
2. The sea surface slopes down toward the coast as Ekman Transport removes more water.
3. An alongshore geostrophic current increases in response to the sloping sea surface.
4. Bed friction reduces the near bed alongshore current leading to weak onshore flow at the bottom of the boundary layer.
5. The thermocline rises as water is upwelled along the coast.
6. As the slope of the sea bed increases the alongshore current increases and the onshore near bed flow must also increase.
7. Eventually the sea surface slope must reach a steady state with offshore Ekman Transport balanced by onshore flow in the bottom boundary layer. For a sufficiently strong wind the pycnocline may break the surface and the whole system moves offshore. Cool bottom water appears as a band against the coast in contrast to the stratified warmer offshore surface waters.

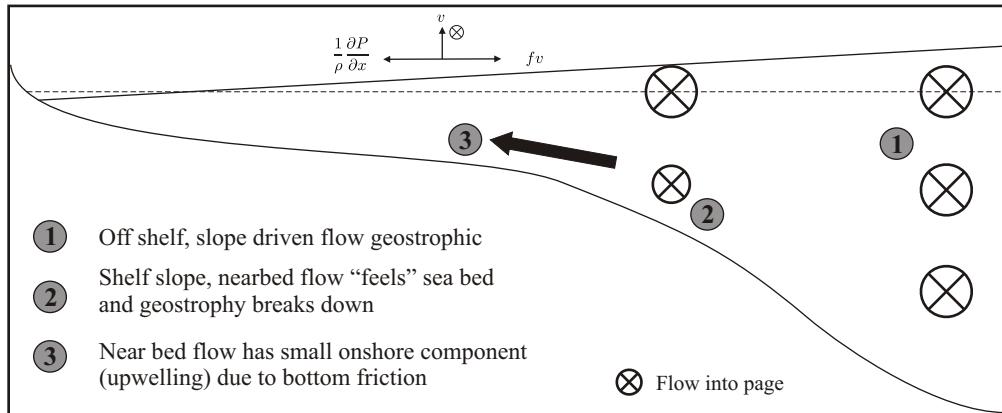


FIGURE 2.2: Residual current driven upwelling (N.Hemisphere).

Strong wind driven upwelling fronts are found against the Peruvian, south-west and north-west African coasts.

The cross-shelf sea surface slope does not necessarily have to be generated by wind stress and Ekman Transport. Current driven cross-shelf flow as depicted in Figure 2.2 is a possible upwelling mechanism in Western Boundary Currents such as the Gulf Stream.

We now consider a second mechanism for cross-shelf exchange: flow nonlinearity.

2.2.3 Flow Nonlinearity

The second means by which the Taylor-Proudman theorem breaks down is when instabilities in the flow lead to a break down in geostrophy at the shelf edge causing meanders, eddies and upwelling events.

The nonlinearity of a flow is described by the nondimensional Rossby Number (R_o).

$$R_o = \frac{V^2/L}{fV} = \frac{V}{Lf},$$

where V is the flow speed (m.s^{-1}), L is the length scale (m) associated with the flow, and f is the coriolis frequency (s^{-1}). When $R_o > \approx 0.1$ the nonlinear terms are large enough to break down geostrophy (Pond and Pickard, 1983). An example of flow perturbation is seen where the Gulf Stream encounters the Charleston Bump as it flows along the north-east American seabord. Meanders that develop downstream of the bump shed warm and cold core eddies that transport heat and momentum across the shelf break (Miller, 1994). In the northern hemisphere cyclonic (anticlockwise) eddies are associated with a sea level depression causing the thermocline to dome upwards bringing nutrient rich waters into the photic zone and increasing the primary productivity at the shelf edge and outer shelf.

This leads us into a discussion of two different frontal regimes where sharp density gradients separate well mixed and stratified waters over the shelf slope and edge: tidal fronts and shelf break fronts, each of which are characterized by meanders and eddies.

2.2.4 Tidal Mixing Fronts

Tidal fronts are a common feature of continental shelf waters during the spring and summer (e.g. Ushant Front, English Channel and Georges Bank, north-east US coast). They mark the transition between thermally stratified deep water and vertically well-mixed waters of shallower areas. Enhanced mixing on the unstratified side is maintained by high levels of turbulence generated at the sea bed by tidal currents. In deeper waters, where currents are not strong enough to break down stratification by solar heating a thermocline is maintained (Figure 2.3). Through consideration of the energy balance between tidal mixing and solar insolation Simpson and Hunter (1974) imply that tidal fronts should follow a critical contour of $\frac{h}{u^3}$, where h is water depth and u is the tidal current amplitude. This situation is shown to be complicated by the control of surface wind mixing and boundary layer thickness on the frontal position (Simpson et al., 1978; Soulsby, 1983).

Tidal fronts are frequently found to be sites of enhanced planktonic biomass and productivity. The mechanisms that cause the transfer of nutrients to the surface layer on the stratified side of the front fall into two categories: Cross frontal exchange and vertical transport. The former may be further subdivided into spring-neap frontal adjustment, baroclinic eddies and residual currents. Instabilities leading to small departures from geostrophy, eddy formation and upwelling is discussed in Section 2.2.3. Weak residual cross-shelf flow resulting from internal friction has also been identified as a potential upwelling mechanism in Section 2.2.2.

The spring-neap frontal adjustment is responsible for a considerable flux of nutrients into the upper stratified layer and for an oscillation in frontal position. As current speeds strengthen toward spring tides the energy available for mixing increases and bottom turbulence pushes the front towards deeper water (Figure 2.3). As neap tides approach turbulence weakens and the stratified waters return shoreward. Newly stratified water behind the advancing front will contain nutrient levels characteristic of the previously mixed waters (Pingree et al., 1975). Vertical eddy diffusion also has a key role in enhancing primary production at tidal mixing fronts. Increased vertical diffusion at the front, along the pycnocline and near the sea bed results in an upward flux of nutrients (Figure 2.3).

The Ushant Front in the English Channel has been the focus of numerous studies (Pingree et al., 1975; Morin et al., 1985; Loder and Platt, 1985; Pedersen, 1994). Estimates of nitrogen transport into the mixed layer of the stratified side of the Ushant Front by

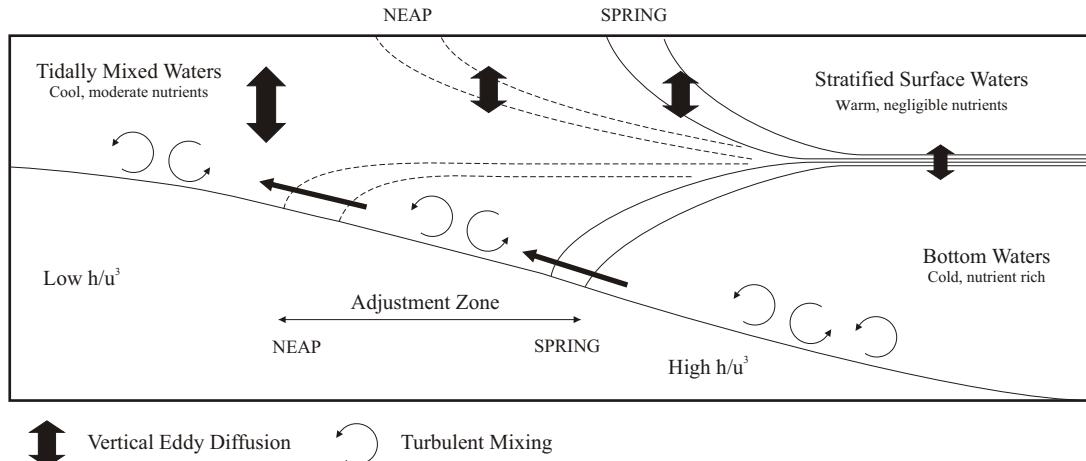


FIGURE 2.3: Density section through a tidally mixed front. Vertical arrows indicate an increase in the rate of vertical eddy diffusion from the stratified side to the mixed side. At the bottom higher levels of turbulence help to transport nutrients into the fully mixed region. The dotted and solid isopycnals indicate the position of the front during neap and spring tides respectively.

spring-neap adjustment, eddies, residual currents and vertical eddy diffusion are 0.12, 0.08, 0.01 and $0.20 \text{ mg.N.m}^{-1}.\text{s}^{-1}$ respectively (Loder and Platt, 1985). This nitrogen supply equates approximately to a carbon flux of $0.28 \text{ g.C.m}^{-2}.\text{d}^{-1}$ over a 10km wide frontal zone (Mann and Lazier, 1996).

2.2.5 Shelf Break Fronts

Shelf break fronts usually straddle the shelf break itself and mark the interface between fresher shelf water and warmer, more saline offshore water (e.g. Celtic Sea and north-east New Zealand shelf breaks). The position of the front may change under the influence of wind forcing. In these regions enhanced primary production is often driven by an internal baroclinic tidal wave that propagates both seaward and shoreward decreasing in amplitude with distance from its origin.

At the shelf break the barotropic tidal current interacts with the topography and on an ebb tide draws the thermocline down off the shelf. This disturbance propagates along the thermocline as an internal tide supported by the density structure. It creates zones of divergence and convergence and therefore zones of up and downwelling respectively and pulses of enhanced primary production along the thermocline. This twice daily mechanism is in contrast to that of tidally mixed fronts where nutrient supply occurs in pulses related to the fortnightly spring-neap adjustment.

The dissipation of internal waves is an important area of shelf edge research. Sharples et al. (2001) estimate a contribution of 100 g.C.m^{-2} to the annual new production at the north-east New Zealand shelf edge driven by a summer flux of $\sim 12 \text{ mmol.N.m}^{-2}.\text{d}^{-1}$ into the photic zone due to increased vertical eddy diffusivity.

We now move away from continental shelf and slope dynamics and consider the large convergence systems where open ocean water masses of differing properties meet and interact. Our focus will be on the Southern Ocean and in particular the Subtropical Front (STF) which will be used as a test bed for model development. Its variability and relationship with ENSO will be investigated in Chapters 6 and 7.

2.3 Open Ocean Convergences in the Southern Ocean

The Antarctic Circumpolar Current (ACC) dominates the circulation in the Southern Ocean. At a latitude of 40°-50°S the predominant winds are westerly, and as Ekman transport in the Southern Hemisphere is to the left, the sea surface slopes down toward the Antarctic continent. An easterly flowing geostrophic current is thus set up extending to the sea floor. Circulation in the Southern Ocean consists of frontal bands separating different water masses of the flow. Three zones of convergence may be identified: 1. The Subtropical Zone, which is to be the example used in this thesis; 2. The Antarctic Polar Frontal Zone delimited by the Subantarctic Polar Front to the north, and the Antarctic Polar Front to the south; and, 3. The Southern Antarctic Zone between the Southern Antarctic Circumpolar Current Front and the Antarctic Continent itself. The position of these fronts is shown on the SeaWiFS chlorophyll image of Figure 2.4.

Away from the continental boundaries and regions of large current shear topography places vorticity constraints on the dynamics of open ocean fronts that strongly controls their temporal and spatial variability. This control is known as *topographic steering*.

2.3.1 Topographic Steering and Vorticity

Put simply, vorticity is the tendency for a fluid parcel to rotate. The potential vorticity (PV) of a body of water in the ocean is defined as:

$$PV = \frac{\zeta + f}{h}.$$

ζ is the *relative vorticity* caused by wind stress and/or current shear, and is measured relative to the earth. f is the *planetary vorticity* (coriolis) of the fluid parcel due to it being on a rotating earth. The planetary vorticity changes with latitude increasing from zero at the equator to a maximum at the poles. It is defined as $2\Omega \sin \Phi$ where Φ is the latitude and Ω the angular velocity of the earth's surface about a vertical axis ($= 7.29 \times 10^{-5} s^{-1}$). h is the depth of the water column. In the open ocean, for large scale processes, $f \gg \zeta$ and the potential vorticity may be approximated as $\frac{f}{h}$ - the *potential planetary vorticity*, PPV. Barotropic flow over smoothly varying topography tends to conserve angular momentum and therefore follows lines of constant

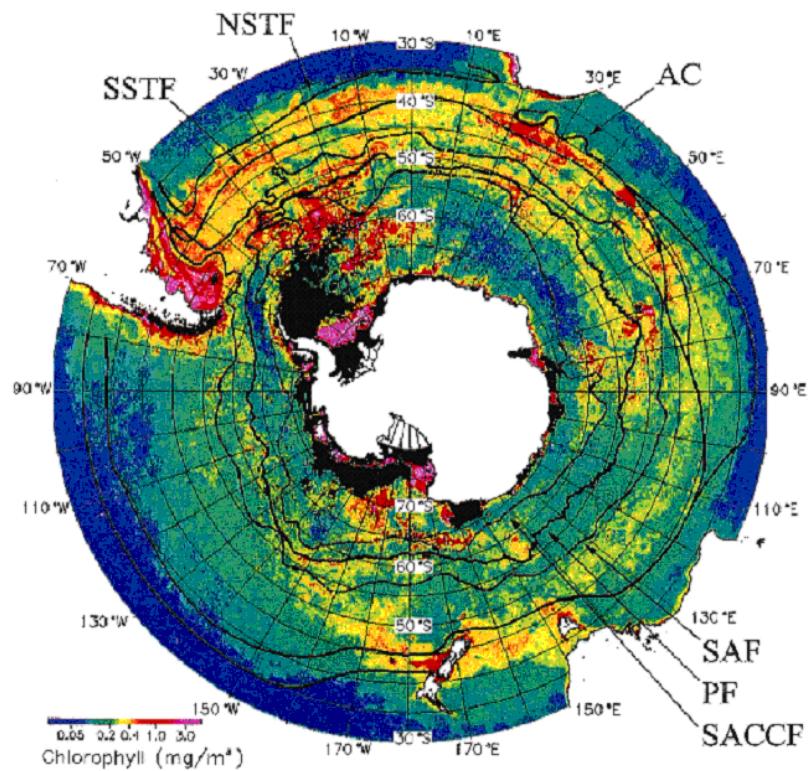


FIGURE 2.4: Mean SeaWiFS chlorophyll concentration during Austral summer (December 1997 to February 1998) and the mean location of the major Southern Ocean Fronts. The Polar Front (PF), the Southern Antarctic Circumpolar Current Front (SACCF), the Subantarctic Front (SAF), the Northern Subtropical Front (NSTF) and Southern Subtropical Front (SSTF) and the Aghulas Current (AC). From Moore and Abbott (2000)

PPV (Pond and Pickard, 1983) - a requirement that leads to topographic steering. The water column is stretched and compressed as it flows over troughs and ridges respectively and f must either decrease or increase in order to maintain a constant value of $\frac{f}{h}$. Since f is a function of latitude only the current must either swing equatorward over ridges or poleward over troughs (in both hemispheres).

In the Southern Ocean sharp mid ocean ridges and troughs mean that frontal jets such as the Polar Front can not follow lines of constant PPV circumglobally (Moore et al., 1999). In regions where the flow is forced suddenly across isolines of PPV (e.g. Drake Passage, Scotia Ridge, Falklands Plateau, Kerguelen Plateau) a large amount of relative vorticity is input into the water column through shrinking/stretching of vortex lines. This energy is dissipated downstream through nonlinear meanders and eddies (Moore et al., 1997, 1999) that cause localized areas of up and downwelling and increased chlorophyll concentrations (Moore et al., 1999; Moore and Abbott, 2000). The horizontal scale of eddies ranges between 10 and 100km, close to the internal Rossby Radius. These departures from geostrophy are depicted in Figure 2.5.

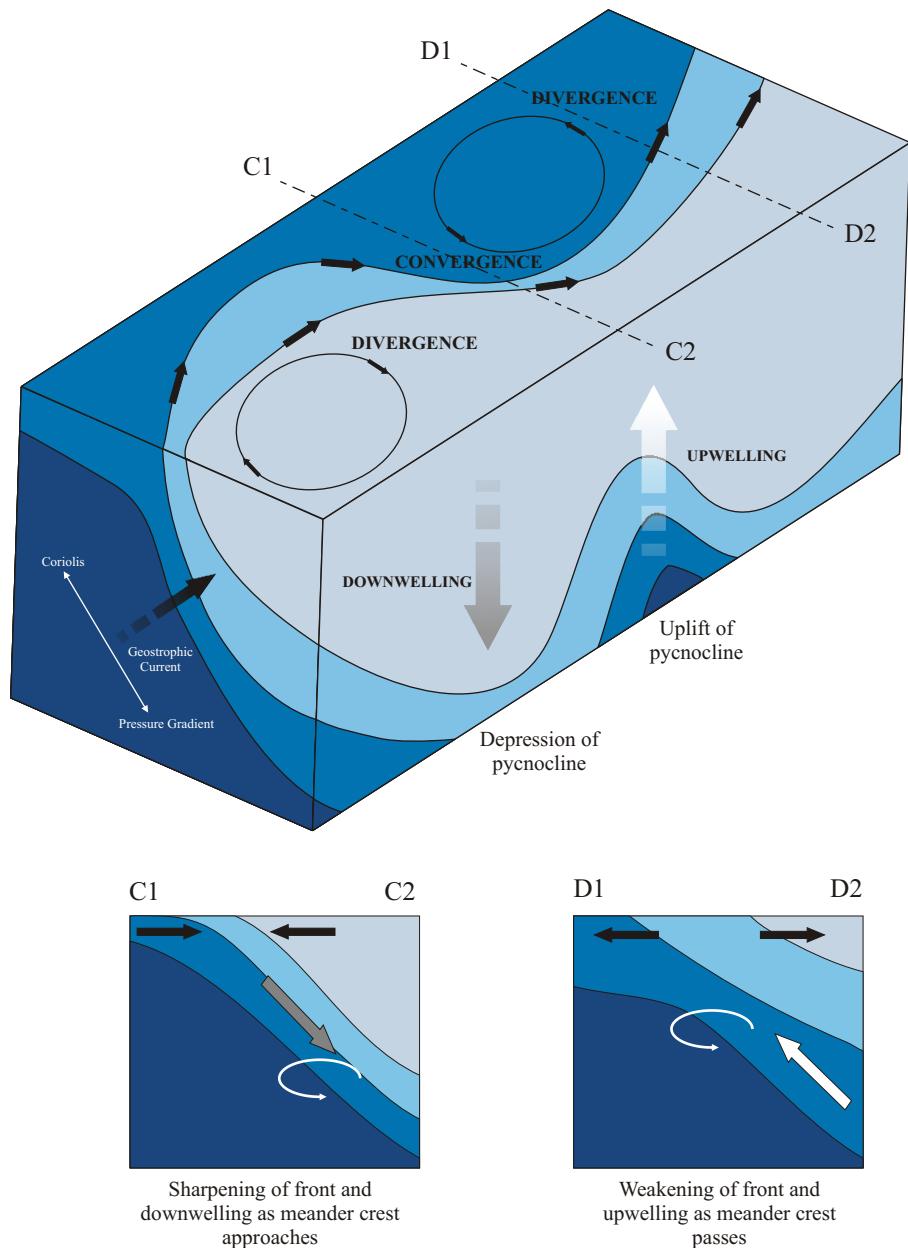


FIGURE 2.5: Open ocean frontal dynamics. The density structure, geostrophic flow, meanders and vertical motions typical of open ocean convergences.

The trough of a meander translating eastwards consists of a cyclonically rotating dome of upwelled water which is continually exchanged due to divergence and convergence at its leading and trailing edges respectively. As the crest of a meander approaches the front steepens and water must move out of the way of the tilted pycnocline (downwelling). As the crest passes the front weakens and waters containing micronutrients (e.g. iron) upwell behind it into the photic zone (Olson, 2002). In the Southern Ocean where micronutrient limitation keeps chlorophyll concentrations low, despite high levels of macronutrients (e.g. nitrate), these biophysical dynamics stimulate localized regions of increased productivity (Moore et al., 1999).

2.3.2 Enhanced Productivity at the Subtropical Front

In the subtropical frontal zone, nitrate availability in the gyre regions likely limits phytoplankton growth, while iron limitation is more probable to the south. Meandering and eddy formation injects iron southwards and nitrate equatorward, resulting in increased productivity (Lutjeharms et al., 1985). Figure 2.4 illustrates elevated chlorophyll concentrations observed along Southern Ocean Fronts. Moore and Abbott (2000) observe phytoplankton blooms with chlorophyll concentrations exceeding 1 mg.m^{-3} at the northern and southern Subtropical Front in the Indian and Atlantic Oceans where there is increased eddy kinetic energy. Mean chlorophyll surface water concentrations in the Southern Ocean are otherwise $<0.3\text{--}0.4 \text{ mg.m}^{-3}$.

We conclude Section 2.3 with a more detailed look at the Southern Ocean Subtropical Front (STF).

2.3.3 Southern Ocean Subtropical Front

The hydrographic boundary between warm, saline subtropical gyre waters to the north and cooler, fresher subantarctic waters to the south is known as the **Subtropical Convergence Zone** (STCZ), (Deacon, 1933). It is a region characterized by negative wind stress curl in the atmosphere and associated Ekman transport and convergence in the ocean, between the core of the trade winds and the maximum westerlies, at an approximate latitude of 40°S . There is a decrease in both the sea surface temperature and salinity from the tropics toward the temperate zone. Since the increase in surface density produced by the decrease in temperature is greater than the poleward decrease in density caused by lower salinities, the surface density tends to increase with distance from the equator (James et al., 2002). According to Deacon (1933), the convergence '*is marked by a sudden change of surface temperature of at least 4°C , and a change of salinity of at least 0.5*'. Within the STCZ and on its poleward side the enhanced meridional temperature and salinity gradients are known as the **Subtropical Front** (STF) and the sloping density surfaces maintain geostrophic equilibrium below the wind driven layer.

The STF is a nearly continuous feature in the southern hemisphere. From the Brazil-Malvinas confluence in the South Atlantic ($\sim 30^\circ\text{S}$) it flows eastward across both the South Atlantic and Southern Indian Oceans at a latitude of approximately $35\text{--}45^\circ\text{S}$. It passes south of Australia, Tasmania and New Zealand and continues across the South Pacific at a latitude of approximately 40°S , before turning north off the west coast of South America to diffuse in lower latitudes near Chile (25°S). Latitudinal fluctuations of 6° or more are common due to seasonal variations (Deacon, 1937). The Antarctic Polar Front and Subantarctic Front on the other hand are fairly stable, meandering about 75km either side of their mean locations (Gille, 1994).

There is confusion in the literature as to whether the STF has a double (Belkin, 1993; Belkin and Gordon, 1996; Kostianoy et al., 2004; Lutjeharms et al., 1993) or a single structure (Stramma and Peterson, 1990; Deacon, 1937). It's temperature and salinity gradients do not remain constant across the three major ocean basins and it's strength, position and volume transport vary (Stramma and Peterson, 1990; Stramma, 1992; Stramma et al., 1995) as it bifurcates and merges with the Subantarctic, Agulhas and Polar Fronts along it's path (Belkin and Gordon, 1996). In the Indian Ocean, where density gradients and geostrophic currents are weak, it has been suggested that density compensation is maintaining the structure of the front (Rudnick and Ferrari, 1999).

A formal criteria for the identification of the STF does not exist. Deacon (1937) uses the 34.9 isohaline as the boundary between subtropical and subantarctic waters, whereas Lutjeharms et al. (1993) use the 18°C isotherm in the summer and the 15°C isotherm during the winter months. Other authors identify the STF on the basis of temperature and salinity gradients at depth, as opposed to near surface maximum, the locations of which do not always coincide (Lutjeharms and McQuaid, 1986). In the Southern Atlantic sector Park et al. (1993) specify temperature and salinity ranges of 8-12°C and 34.6-35.0 respectively at 200m depth, with axial values of 10°C and 34.8. Orsi et al. (1995) on the other hand use a reference depth of 100m. Ambiguity in the identification of the STF, due partly to discrepancies between data collection methods (bottle vs CTD vs remote sensing), natural seasonal/interannual variability, and frontal interactions, introduces uncertainty into monitoring it's strength and position.

This thesis is concerned with the variability of the STF as it passes New Zealand. This localized segment was chosen on a number of accounts. Firstly, it has been suggested that the strength of the front was modified by the 1991 El Niño (Shaw and Vennell, 2001) which motivates an investigation into possible ENSO teleconnections. Secondly, the STF in this area is known to be strong, persistent and stable (Shaw and Vennell, 2001; Uddstrom and Oien, 1999; Chiswell, 1994) providing an ideal test bed for algorithm development. Additionally, the STF off South Island, New Zealand has already been used to develop a previous front detection algorithm (Shaw and Vennell, 2000a). This will allow us to compare the performance of our statistical model and the findings of our variability study. The following is a discussion of the oceanography of the waters surrounding New Zealand and of the fronts and circulation in the region.

2.4 Physical Oceanography of New Zealand

2.4.1 Introduction

New Zealand consists of two main islands that are situated on an isolated continental shelf extending from approximately 33°S to 49°S (Figure 2.6). The shelf crosses a range of climates from subtropical in the north to temperate in the south. There are four main water masses around New Zealand: Subtropical Surface Water (STW), Subantarctic Surface Water (SAW), Antarctic Intermediate Water (AIW), and Pacific Deep Water (PDW). Neritic water (close inshore water) along the coastal zone is also identified. Table 2.1 summarises the characteristics and origins of these water masses.

TABLE 2.1: Temperature (**T**), salinity (**S**) and origins of the main water masses around New Zealand (Shaw, 1998)

Water Mass	T (°C)	S (psu)	Origin
Neritic Water (NW)	9-14	33.8-34.6	Coastal waters diluted by fresh riverine input
Subtropical Surface Water (STW)	10-13	34.6-34.9	Originates central S.Pacific. Flows south in East Australia and Tasman Currents
Subantarctic Surface Water (SAW)	7-12	34.3-34.5	Subantarctic water driven north in west wind drift within Circumpolar Current
Antarctic Intermediate Water (AAIW)	4-6	34.3-34.4 at 700-1200m	Low salinity surface water south of Polar Front sinks at convergence and travels north
Pacific Deep Water (PDW)	1-2	34.6-34.8 at core of 3500m	Originates at surface in N.Atlantic. Travels south and then eastward in Circumpolar Current then north in south-west Pacific

New Zealand forms the western boundary to the South Pacific Gyre and occupies the unique position as being the only major region of continental shelf to interrupt a Western Boundary Current in mid-ocean (Sharples, 1998). Indeed New Zealand may play a determining role in the separation point of the Western Boundary East Australian Current from the continent (Godfrey et al., 1980). It also lies across the path of the Antarctic Circumpolar Current system and presents a significant barrier to the Subtropical and Subantarctic Fronts.

Tides around New Zealand are characterized by a coastally trapped Kelvin wave that travels anticlockwise around the shelf at around $10-20 \text{ cm.s}^{-1}$. The width of the shelf varies considerably around the islands from less than 1km off Fiordland to over 250km on the Central Western Shelf (Figure 2.7) (Sharples, 1998).

New Zealand's unique, isolated location and range of climatic/oceanic conditions results in a diverse physical regime over the continental shelf. Barotropic and internal tides, ocean driven shelf edge flows, buoyancy and wind driven coastal currents, and coastally

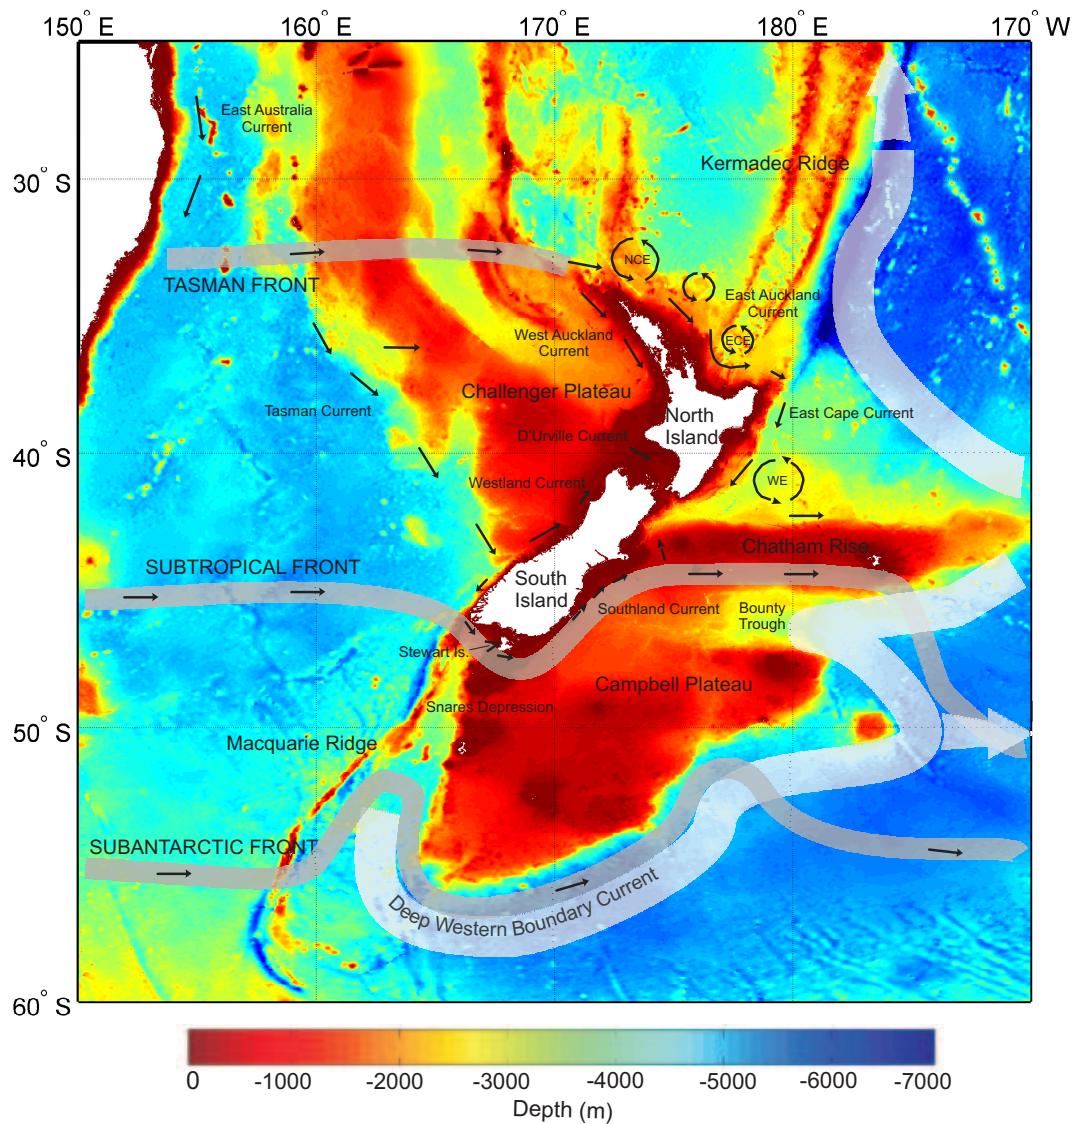


FIGURE 2.6: A bathymetric map showing the main frontal systems and circulation around New Zealand. There are three main fronts: The Tasman Front, The Subtropical Front and the Subantarctic Front. WE = Wairarapa Eddy, ECE = East Cape Eddy, NCE = North Cape Eddy.

trapped waves all contribute to the shelf dynamics and help drive a highly productive ecosystem.

New Zealand has the fourth largest marine Exclusive Economic Zone in the world which has an average chlorophyll-*a* concentration of $0.26\text{--}0.43 \text{ mg.m}^{-3}$ (Murphy et al., 2001). The highest concentrations are found at the Subtropical Front (Vincent and Howard-Williams, 1991; Murphy et al., 2001) where warm macronutrient-poor, and relatively iron rich subtropical waters mix with cold, macronutrient-rich and iron poor subantarctic waters (Boyd et al., 1999; Butler et al., 1992). Spring and autumn pigment concentrations may be $>3 \text{ mg.m}^{-3}$ and up to 1.5 mg.m^{-3} respectively over the Chatham Rise (Bradford-Grieve et al., 1997). Increased productivity over the Chatham Rise supports some of New Zealand's most valuable deep water fisheries including hoki, orange roughy,

oreos and ling (McClatchie et al., 2001). Coastal and pelagic species (e.g. snapper, rock lobster, tuna, mackerel) are also of significant economic importance.

We continue with a description of the frontal systems and circulation around New Zealand recognizing the two features that dominate the oceanographic environment: the Western Boundary Current of the South Pacific Subtropical Gyre and the Subtropical Convergence. Specific examples of boundary upwelling fronts, tidal mixing and shelf break fronts are highlighted.

2.4.2 Frontal Systems and Circulation

2.4.2.1 North East Shelf

New Zealand is surrounded by warm subtropical waters originating from the Coral Sea off the northeast coast of Australia. A warm Western Boundary Current - the East Australian Current - flows southwards along the Australian shelf and then east across the Tasman Sea toward North Island, New Zealand forming the Tasman Front (Figure 2.6).

A major portion of the Tasman Current impinges on the north of North Island to form the East Auckland Current. This topographically steered ocean flow has a mean transport of approximately 9Sv (Stanton and Sutton, 2003; Roemmich and Sutton, 1998) and maintains a shelf break front and near bed onshore flows at the shelf edge. Wind forcing and strong summer thermal stratification decouple the surface water from the topographic constraint of the shelf leading to an onshore frontal movement and intrusion of subtropical surface waters over the shelf edge (Sharples, 1997). Strong seasonal wind induced up- and downwelling (Zeldis et al., 2004) and internal tidal dissipation during the summer (Sharples et al., 2001) make the northeastern shelf one of New Zealand's most productive shelf seas.

At East Cape most of the East Auckland Current deflects south forming the East Cape Current and eventually the northern side of the Subtropical Front (Tilburg et al., 2001; Heath, 1985a). Offshore of this North Island boundary current three permanent, anticyclonic eddies are identified (Figure 2.6): North Cape Eddy, East Cape Eddy and Wairarapa Eddy (Roemmich and Sutton, 1998). The eddy field that is set up as currents pass North Island may be responsible for the formation of Subtropical Mode Water (Roemmich and Cornuelle, 1992).

2.4.2.2 South West Shelf and Cook Strait

Approaching New Zealand a smaller portion of the Tasman Current is diverted southwards by the Lord Howe Rise/Challenger Plateau (Figure 2.6). It reaches the west coast

of South Island and forms the beginnings of the Southland Current past Fiordland (Figure 2.7). As is the case for most of the shelf waters around New Zealand this region is not dominated by any one physical process. The offshore Tasman Current aligns isotherms and isohalines parallel to the shore and warm, saline offshore waters meet colder, fresher inshore waters at the shelf break. Highly variable coastal currents are driven by wind forced coastally trapped waves generated over the Cook Strait (Stanton, 1995). These waves may be involved in shelf to offshore exchange through coastal squirts (Moore and Murdoch, 1993).

North of Jackson Head, a weak northward flow known as the Westland Current (Figures 2.6 and 2.7) is driven by upwelling favorable winds (Brodie, 1960). These prevailing winds induce a fall in sea level near Cape Farewell and the resulting sea surface slope accelerates flows of deep water over the bathymetric rise (see Section 2.2.2). The uplift of the thermocline coupled with the convergence of the bottom Ekman flow results in a strong upwelling source near Kahurangi Point and an upwelling front - The Farewell Front - that extends northwards into Cook Strait (Figure 2.7). Within Cook Strait itself tidal mixing fronts have been identified in the Cook Strait Narrows and off D'Urville Island and Marlborough Sounds (Bradford-Grieve et al., 2006).

2.4.2.3 South East Shelf: The Subtropical Front

The landmass of New Zealand and its submarine platform are a natural barrier to the circumpolar flow of the Subtropical Front (SFT) around the Southern Ocean. There is much debate over the path and continuity of the STF as it passes across the Tasman Sea on its way to the Pacific Ocean. Unlike Belkin and Gordon (1996), James et al. (2002) do not find the STF to be continuous in this region. During winter, between 38°S and 39°S, the STF encounters strong westward flow south of Tasmania, and ends at the continental shelf west of the Bass Strait, near 40°S. In the summer, the temperature expression of the front is covered by a 50-70m thick mixed layer, but continues to exist at depth in the permanent thermocline (James et al., 2002). Deacon (1933) describes the STF as passing south of Tasmania and then turning north-east toward the southern end of North Island, terminating on the eastern side of New Zealand. Jeffrey (1986)¹ tracks two bands of enhanced salinity gradient. He follows the greater 34.7 isohaline zonally from Tasmania to the south of South Island, and the 35.1 isohaline along a northerly track toward the south of North Island. The extreme variation in position is a result of the area being one of the most diverse and interconnected regions of the world's oceans, influenced by the Antarctic Circumpolar Current, the Indo-Pacific Throughflow and global thermohaline circulation. Furthermore, the Tasman and western South Pacific experience some of the largest seasonal and interannual atmospheric conditions in the

¹cited within Stramma et al. (1995)

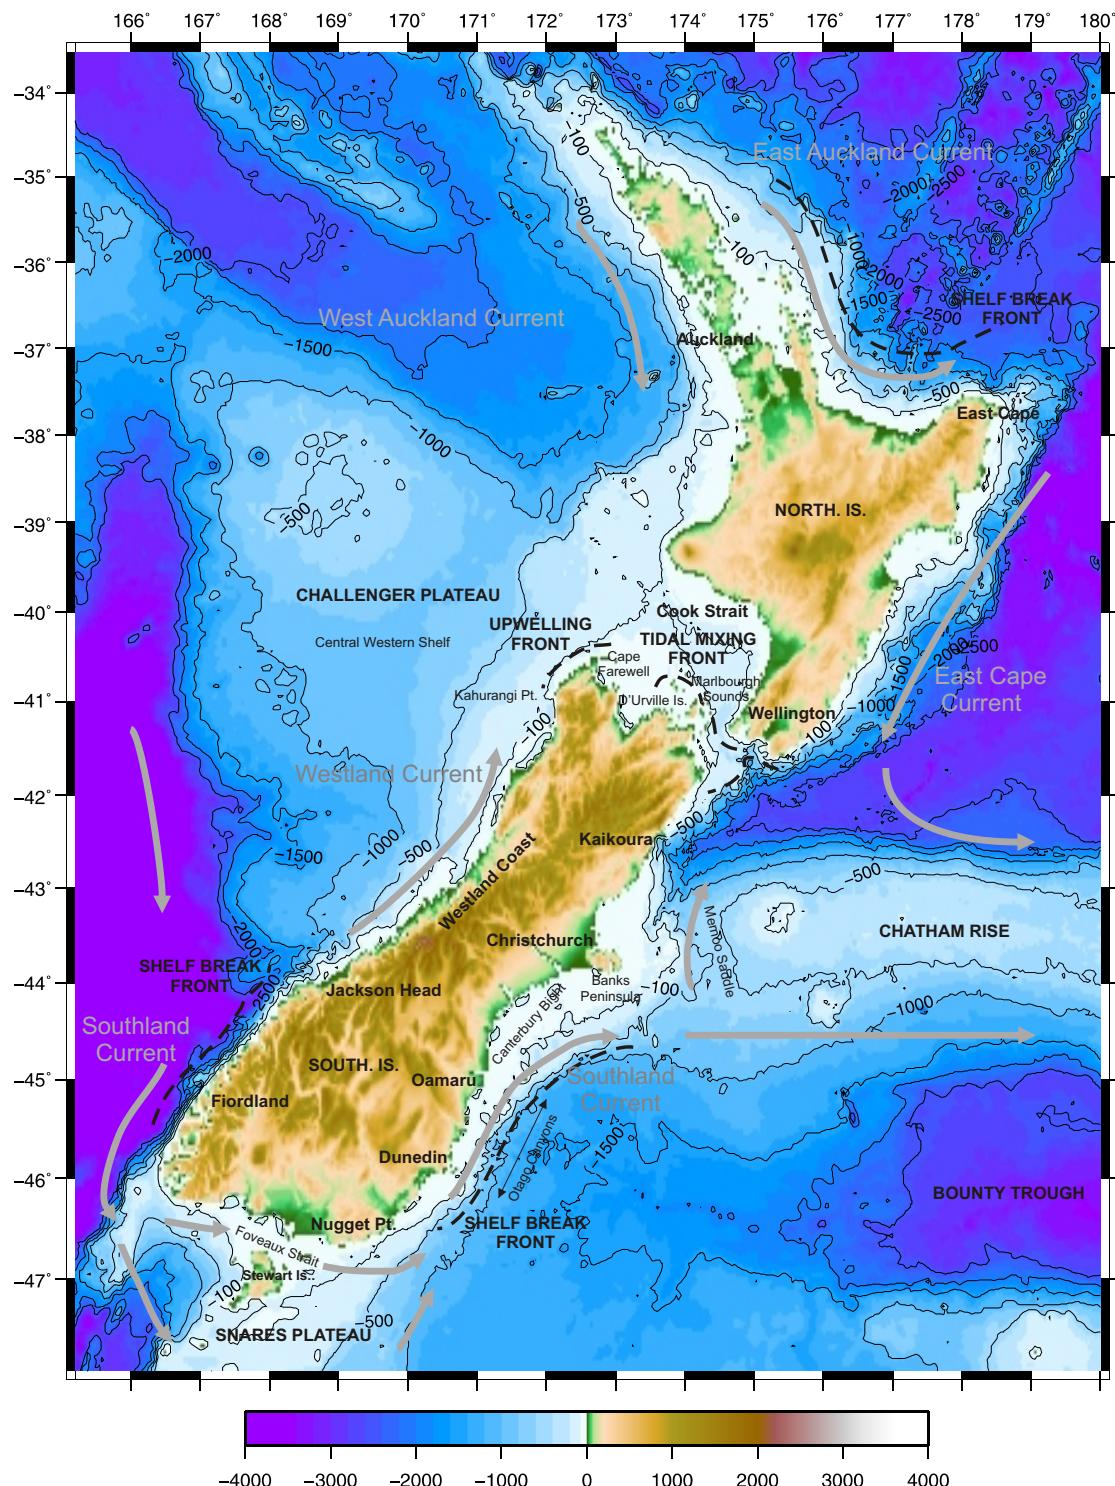


FIGURE 2.7: Bathymetry and topography (metres) of New Zealand. Frontal structures discussed in the main text are marked with thick black dashed lines.

world, causing variability in wind stress and thus Ekman Pumping (Stramma et al., 1995).

West of New Zealand the STF is directed southward toward Snares Shelf (Figure 2.6) (Jillett, 1969; Heath, 1985a; Brodie, 1960; Houtman, 1966). It then turns northwards just east of South Island bounded on its western side by an induced northward geostrophic flow that forms part of the **Southland Current** (Heath, 1972, 1985a). It is subsequently steered eastwards along the Chatham Rise out into the open Pacific Ocean. The front is bounded on its eastern side by subantarctic waters (SAW). Figure 2.6 shows the path of the deep current associated with the Subantarctic Front (SAF). The strong flows are deflected abruptly northwards as they encounter the Macquarie Ridge and Campbell Plateau in order to conserve potential planetary vorticity. Subantarctic waters follow the eastern edge of the Campbell Plateau and flow into Bounty Trough hugging the western bathymetry contours before exiting into the south-east Pacific Basin just south of Chatham Rise. This northwards excursion of the SAF is responsible for the presence of cool SAW immediately to the east of New Zealand, further north than in the Tasman Sea and south-west Pacific (Morris et al., 2001), that forms the eastern boundary of the STF along the shelf. The presence of the Southland Current on the eastern coast of South Island is thought to be due to remote forcing from the strong topographically constrained subantarctic abyssal currents located approximately 12° east of New Zealand (Tilburg et al., 2002).

The surface expression of the STF (Figure 2.8) follows approximately the 15°C isotherm in the summer (10°C in winter), and the 34.7-34.8 isohaline (Heath, 1985a). This thesis is concerned with the structure and variability of this surface expression as it passes the east coast of South Island, New Zealand where is known locally as **The Southland Front**.

2.4.2.4 The Southland Front and Current

Traditionally, the Southland Current, the western boundary of the Southland Front is described as consisting mainly of warm, saline subtropical waters (Jillett, 1969; Burling, 1961). This water originates from the southward flow of the Tasman Current that is deflected poleward, passing Stewart Island to the south through the Snares Depression, and to the north through Foveaux Strait (Jillett, 1969; Brodie, 1960; Houtman, 1966). More recently however the current has been shown to advect a much higher percentage of SAW than previously thought. Sutton (2003) finds the mean transport to comprise 90% SAW and 10% STW owing to the current extending further offshore than was captured in former surveys. The core of its geostrophic flow was found to be seaward of the Southland Front.

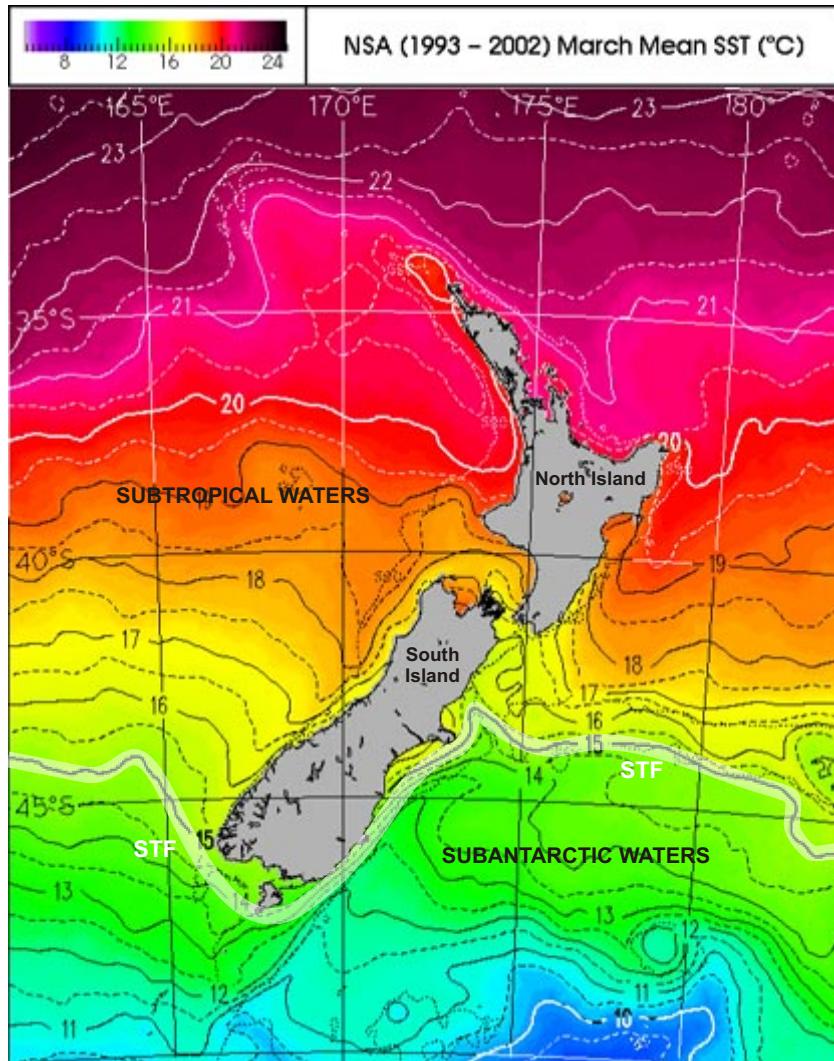


FIGURE 2.8: Mean March SST between 1993 and 2002. 15°C isotherm marks the approximate location of the Subtropical Front (STF). Source: NIWA Satellite Data Service SST Climatology.

The Southland Current is topographically steered as it flows equatorward (Heath, 1972; Chiswell, 1994) closely following the 500m isobath (Shaw and Vennell, 2001). It is characterized along the Westland and Otago coastlines by a high subsurface salinity core (Jillett, 1969; Heath, 1972; Chiswell, 1996). A number of estimates of current speed and transport for the Southland Current have been made. Heath (1972) estimates the Southland Current to transport in total 3.5Sv south of Peninsula Banks. Based on geostrophic calculations with a reference level of no motion at 1000dbar, Chiswell (1996) calculate current magnitudes of $<10\text{cm.sec}^{-1}$ off Nugget Point, and $30\text{-}50\text{cm.sec}^{-1}$ off Oamaru. The total volume transports were estimated to be 2.2Sv and 10.4Sv respectively. Heath (1975) calculates a geostrophic velocity of 20cm.sec^{-1} using a higher reference depth (500dbar). Much of the variability in velocity of the Southland Current is a delayed response to the winds and currents over the Snares Plateau and in Foveaux Strait (Chiswell, 1996).

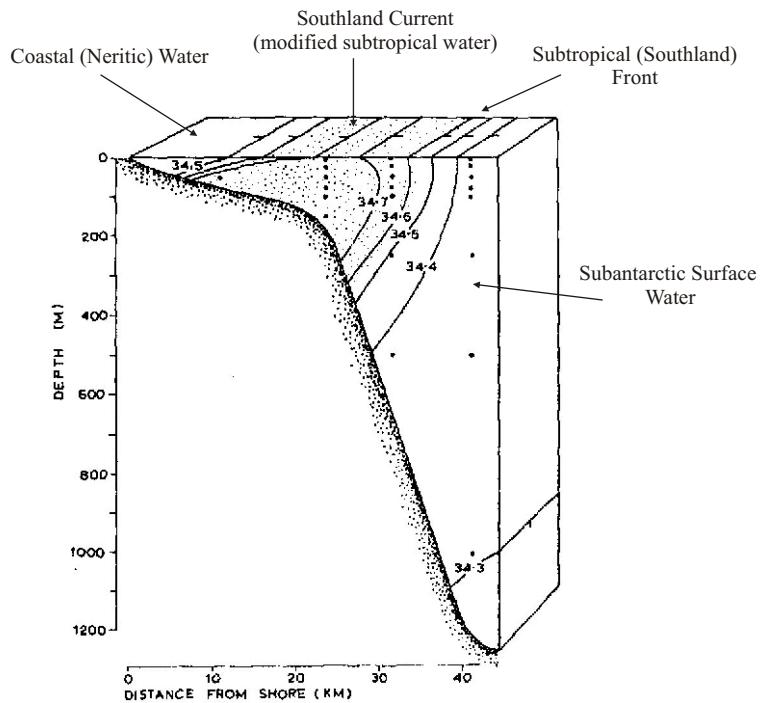


FIGURE 2.9: Vertical salinity section across the Southland Front off the Otago coast in October 1966. Adapted from Jillett (1969).

Figure 2.9 is a vertical salinity section across the Otago shelf. To the east it shows the Southland Front separating modified subtropical shelf waters flowing northwards in the Southland Current from offshore subantarctic waters. Over a three year period, the thermal surface expression of the Southland Front was found to have a mean width of 8.36km, a mean temperature of 11.01°C , a mean temperature difference between its two bounding water masses of 1.76°C , and a mean thermal gradient (strength) of $0.28^{\circ}\text{C}.\text{km}^{-1}$ (Shaw and Vennell, 2001). Exact definitions of the mean temperature, width and temperature difference adopted in this study are presented in Section 3.4. Essentially, they are the parameters controlling the shape of a hyperbolic tangent function fitted to a set of SST observations made across a front. Annual modulation in the strength of the Southland Front accounts for about 20% of its variance (Chiswell, 1996) suggesting that remote forcings may play a significant role in its variability (e.g. ENSO). The front is strongest and narrowest in the winter (Chiswell, 1996; Shaw and Vennell, 2001). In contrast, Uddstrom and Oien (1999) find the Southland Front to be strongest in the spring and autumn, and weakest in the winter. This study hopes to shed light on these conflicting findings.

On its passage northwards the thermal subtropical characteristics of the Southland Current become diluted as a result of mixing at its boundary with cool SAW being brought to the surface (Butler et al., 1992) and strong fresh outflows from fiords and rivers. The main coastal input of fresh water is from the Clutha River (Jillett, 1969). Waters cool by approximately 4°C and the core of maximum salinity is freshened by approximately

0.5 from the currents upstream source (Morris et al., 2001). Warm subtropical plumes protrude seaward from the Southland Front north of the Otago Canyons (Shaw, 1998), providing evidence of instabilities and mixing between STW and SAW. The front's salinity signature however remains unaffected and is therefore a more appropriate indicator of the subtropical water (Houtman, 1966).

Circulation north of Peninsula Banks is complex. An extension of the Southland Current splits northwards through the Mernoo Saddle (Heath, 1972, 1985a; Shaw and Vennell, 2000b), and the remainder turns eastwards and continues offshore. North of 44°S, the Southland Front is identified by a cool, low salinity tongue of SAW forced upwards through the western side of the Mernoo Saddle from the continental shelf off Dunedin. Water deeper than 800m does not pass through the gap and flows eastwards, south of the Chatham Rise (Heath, 1985a). On the eastern side of the channel a periodic southwards flow of STW has been recorded by Greig and Gilmour (1992) and Shaw and Vennell (2000b). North of the Mernoo Gap the Southland Current extension encounters a highly dynamic eddy system. The large, anticyclonic, permanent Wairarapa Eddy, centered at approximately 42°S, 176°E off the Wellington Coastline, interacts with the meandering subtropical East Cape Current and sheds smaller eddies south-eastwards toward the Hikurangi Trench (Heath, 1968; Barnes, 1985; Heath, 1975; Vincent and Howard-Williams, 1991). In this region, the Southland Current extension splits into three branches: eastward to combine with the East Cape Current, north-east across the southern end of Cook Strait, and northwards to enter the Cook Strait on its south-western side (Heath, 1972).

The exact location and structure of the STF as it flows eastwards along Chatham Rise is unclear. Uddstrom and Oien (1999) find the STF to be bathymetrically locked to the southern flank of the rise, looping southwards beyond 177°W to the northern edge of Bounty Plateau where it meets the SAF (Figure 2.6). In contrast, Heath (1985b) and Chiswell (1994) describe a north-south oscillation in the position of STF over the Chatham Rise. It lies along the southern edge of the rise in spring, and moves to the northern flank during late summer. Sutton (2001) describes a double structure of northern and southern branches of the STF either side of the Chatham Rise marking a Subtropical Frontal Zone similar to that described by Belkin and Gordon (1996) elsewhere.

The Southland Front and Current are important on a number of different levels. Regionally the Southland Current system determines the local oceanographic conditions off the south-east coast of South Island. The narrow band of STW results in approximately a 2°C difference in ocean temperature between inshore and offshore waters during the winter. It also helps maintain a cooler climate than on North Island (Heath, 1972). On the larger scale the Southland Current transports heat and salt into the convergence region over the Chatham Rise. The variability in strength and position of the current may result in fluctuations in the stability, gradients and position of the STF as it enters

the Southern Pacific. The convergence over the rise is important from a carbon ocean-atmosphere exchange perspective. It is known to be a carbon dioxide sink during the austral spring (Murphy et al., 1991; Currie and Hunter, 1998, 1999) because of its high productivity rates and increased biological carbon assimilation.

This chapter has provided a comprehensive review of the biophysical dynamics of different frontal regimes. A thorough description of the fronts and circulation around New Zealand provides the background material to the discussion of results in Chapters 6 and 7. In the following chapter the different techniques used to detect ocean fronts from SST satellite images are introduced.

CHAPTER 3

Ocean Front Detection: An Overview

The advent of environmental satellites in the 1970's has revolutionized the oceanographers ability to study rich and highly energetic oceanic structures such as upwellings, eddies, fronts and jet currents that are important not only to the study of oceanic circulation, but also to the productivity and ecology of life around them (Moore and Abbott, 2000; Lutjeharms et al., 1993; Mann and Lazier, 1996). The availability of such a large and rapidly expanding data set of remotely sensed sea surface temperature (SST), altimetry and ocean color has fueled a growing interest and demand for objective automatic techniques to detect and monitor these mesoscale features and processes. Applications for such algorithms can be found in the world of scientific research, where they are used, for example, in dynamical and variability studies and for model validation purposes. However, applications also extend into the more commercial domain of fisheries management and exploitation and coastal monitoring.

Successfully detecting and monitoring ocean fronts is a nontrivial problem. Fronts exhibit nonlinear flows and processes on a range of different temporal and spatial scales (Chapter 2). Capturing these features as they grow, merge, split, shrink and disappear is a considerable challenge. Complicating the task are the resolution limitations imposed by the instruments being used to image these phenomena. Noise and images obscured by cloud further add to the challenge of accurately detecting fronts. In this chapter we present an overview of techniques that are used to detect ocean fronts and eddies, ranging from the earliest attempts at feature extraction by Gerson et al. (1979) in the Gulf Stream, to the more modern and sophisticated approaches of Cayula and Cornillon (1992, 1995), Marcello et al. (2005) and Shaw and Vennell (2000a).

Finding ocean fronts in remotely sensed imagery is a problem that has been approached in five different ways:

1. Derivative based edge detection
2. Statistical/probabilistic edge discrimination and classification
3. Gradient magnitude criterion
4. Surface fitting
5. Supervised learning methods

Each of these areas will now be addressed in turn.

3.1 Derivative Based Edge Detection

Traditionally, derivative based edge detectors that are quick and easy to apply are used to compute gradients, detect edges and enhance features in satellite data from their background (Simpson, 1990). These classical gradient edge detection techniques rely on locating brightness intensity gradient discontinuities in an image. For example in a sea surface temperature image where there is a rapid transition from cold to warm water at a frontal boundary. Locating these gradient discontinuities is achieved by finding the maxima/minima and zero crossings in the first and second order derivatives of an image respectively.

The first order derivatives of an image are approximated by convolving the image with what are known as convolution masks or kernels. Convolution is the process of sub-sampling an image, $M(I, J)$, into tiles, $m(i, j)$, upon which a mask, $k(i, j)$, is applied. The sum of all the tile elements multiplied by the corresponding mask elements yields the convolution product (CP) for the pixel about which the tile is centred. Mathematically, for a 3×3 pixel tile size this is expressed as:

$$CP = \sum_{i=1}^3 \sum_{j=1}^3 m(i, j) \cdot k(i, j).$$

The first order Prewitt, Sobel and Kirsch convolution masks are shown in Table 3.1. Two masks (k_x and k_y) are applied calculating the convolution product (intensity change) in both the east-west (Δx) and north-south (Δy) directions. The mean (or normalized) intensity change ($|\Delta x|, |\Delta y|$) is calculated by dividing the convolution product from each mask by the number of differences computed. For the Prewitt masks three differences are calculated and added together therefore the normalization factor is 3 (see Table

TABLE 3.1: First derivative convolution masks. Each of the masks place different weightings and bias toward different pixels. From Simpson (1990).

Mask	k_x			k_y			Normalization Factor
Prewitt	-1	0	-1	-1	-1	-1	3
	-1	0	1	0	0	0	
	-1	0	1	1	1	1	
Sobel	-1	0	1	-1	-2	1	4
	-2	0	2	0	0	0	
	-1	0	1	1	2	1	
Kirsh	-3	-3	5	5	5	5	15
	-3	0	5	-3	0	-3	
	-3	-3	5	-3	-3	-3	

3.1). The total gradient magnitude $G = |\Delta x| + |\Delta y|$. A pixel whose gradient magnitude exceeds some threshold value is identified as an edge.

Other gradient operators falling into this category are the Roberts and Isotropic kernels. These first order derivative techniques are used when the actual values of the brightness gradient are required. Dividing the mean gradient intensity ($|\Delta x|, |\Delta y|$) by a spatial factor returns a gradient that may be used in the initialization and updating of numerical models for example. For a 3×3 convolution mask applied to a 1km resolution AVHRR image this spatial factor is twice the pixel dimension, i.e. 2km. Janowitz (1985)¹ uses the Kirsch edge detector to help locate Gulf Stream eddies, and Sauter and Parson (1994) employ the Sobel filter to aid feature identification in GLORIA (Geostationary/Low-Earth Orbiting Radar Image Acquisition) images.

A more sophisticated gradient edge detector is the Laplacian of the Gaussian (LoG), a second order derivative approach (Nixon and Aguado, 2002). Horizontal and vertical Laplacian convolution masks are designed to approximate the second derivative of a brightness intensity image, just as those masks in Table 3.1 act as first order difference operators. To counter the exaggeration in high frequency noise caused by the second order masks the image is first smoothed by convolving it with a Gaussian kernel. The standard deviation, σ , or width of the Gaussian smoothing mask determines the scale of detectable feature. Increasing σ biases the detection of large scale edges. A smaller σ allows much finer scale fronts to be detected.

Another important edge detector is the Canny operator (Nixon and Aguado, 2002). It works as a multi-stage process. Firstly the image is smoothed by Gaussian convolution. A simple two-dimensional first derivative operator is then applied to highlight regions with high first order derivatives. Lastly the algorithm tracks along ridges in the gradient

¹cited in Holyer and Peckinpaugh (1989)

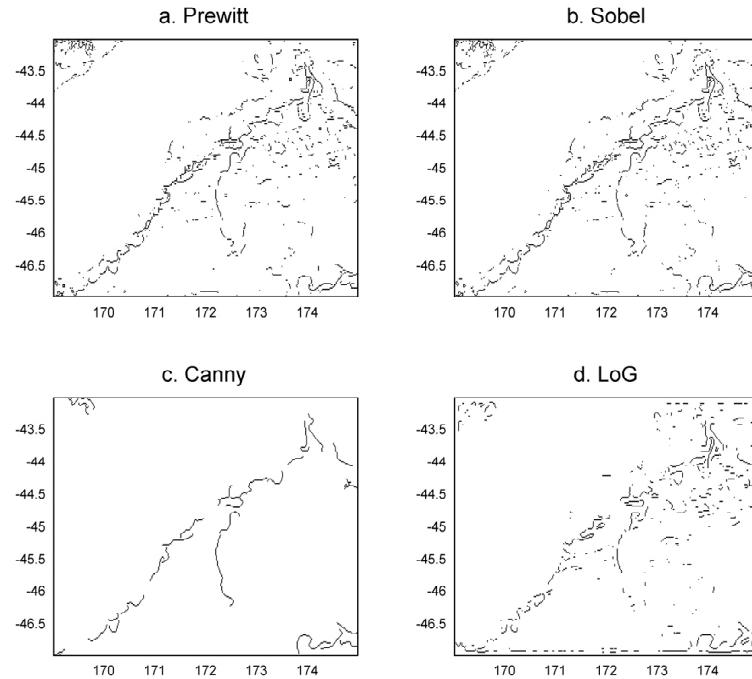


FIGURE 3.1: Examples of gradient edge detectors applied to a 1km resolution image of the east coast of South Island, New Zealand. a. Prewitt, b. Sobel, c. Canny, and d. LoG. The Prewitt and Sobel operators were set at a threshold of $0.2^{\circ}\text{C}.\text{km}^{-1}$. The lower and upper thresholds of the Canny detector are 0.1 and $0.2^{\circ}\text{C}.\text{km}^{-1}$ respectively. The gaussian smoothing mask used for the LoG operator has a standard deviation of 2km and a threshold of 0.04.

magnitude image setting to zero all pixels that are not actually on the top of the ridge. This means that a thin line is output as the edge location as opposed to a wide band of pixels. The tracking process is controlled by two threshold values which helps ensure that noisy edges are not broken up into segments.

Figure 3.1 shows examples of the Sobel, Prewitt, LoG and Canny edge detectors applied to a 1km resolution AVHRR image of the south-east coast of South Island, New Zealand, from 28th March 1990.

Many of these gradient edge detection techniques have been developed in the world of machine vision to detect, for example, fish from underwater video images (Savage et al., 1994). However, these traditional edge detectors are not well suited to oceanographic remote sensing applications (Hoyer and Peckinpaugh, 1989). They struggle to discriminate between weak, small scale features and noise due to atmospheric interference (e.g. cloud), and sensor inaccuracies. Often a pre-filtering stage is required, but this smoothing blurs features and sharp gradients and makes subsequent edge detection more difficult.

3.2 Automated Classification and Edge Discrimination

Using an automated segmentation technique, combined with a specific edge detection or contour connection component, is an alternative to front detection and feature identification in satellite imagery. Broadly speaking, such methods involve the classification of pixels or windows of data and some form of statistical/probabilistic analysis to determine the presence of an edge. Algorithms that fall into this category may be further subdivided: distribution diversity (entropy) based methods (Vazquez et al., 1999; Shimada et al., 2005); histogram analysis (Cayula and Cornillon, 1992, 1995; Marcello et al., 2005); a clustering based approach (Hoyer and Peckinpaugh, 1989); and an examination of the moments combined with *a-priori* knowledge of the region (Gerson et al., 1979; Coulter, 1983). Other edge detectors tested in an oceanographic frontal context include the wavelet based approach of Simhadri et al. (1998), mathematical morphology algorithms (Lea and Lybanon, 1993; Krishnamurthy et al., 1994), the Ordered Structural Edge Detector of Holland and Yan (1992), and a structural relationship based identification scheme (Nichol, 1987). The following text provides a brief overview of the above techniques. Particular attention is given to the Single Image Edge Detector of Cayula and Cornillon (1992) as this technique is compared to the front detection model developed in Chapter 4.

Gerson et al. (1979) were among the first to demonstrate that an automated pattern recognition and frontal identification algorithm was possible. They used Fishers Linear Discriminant, a measure of the separation between two classes of data and of the dispersion within them, to evaluate first and second order statistics on a frame by frame basis to separate different water masses in the Gulf Stream and to identify its northern wall. Given two classes of data, C_1 and C_2 , a line may be drawn (Figure 3.2) that gives good separation between them. Fishers Linear Discriminant is a measure of the separation between the two classes and of the dispersion within them. The optimal projected line will maximize the distance between classes while minimizing the variance within them. It is determined by maximizing the Fisher Distance (FD):

$$FD_{1,2} = \frac{|M_1 - M_2|^2}{S_1^2 + S_2^2},$$

where M_1 and M_2 represent a mean, and S_1 and S_2 a standard deviation (Gerson et al., 1979). The statistics used to separate Gulf Stream, slope and Sargasso Sea water within 16×16 frames were the mean temperature and standard deviation, and the difference histogram maximum entropy. Difference histograms measure the frequency with which temperature differences occur adjacent to or at a fixed distance between pixels. The entropy measures the diversity of events within that distribution. Exactly pinpointing the Gulf Stream northern front was subsequently achieved by scanning each 16×16 frame with a smaller 5×5 pixel window within which the skewness of the temperature

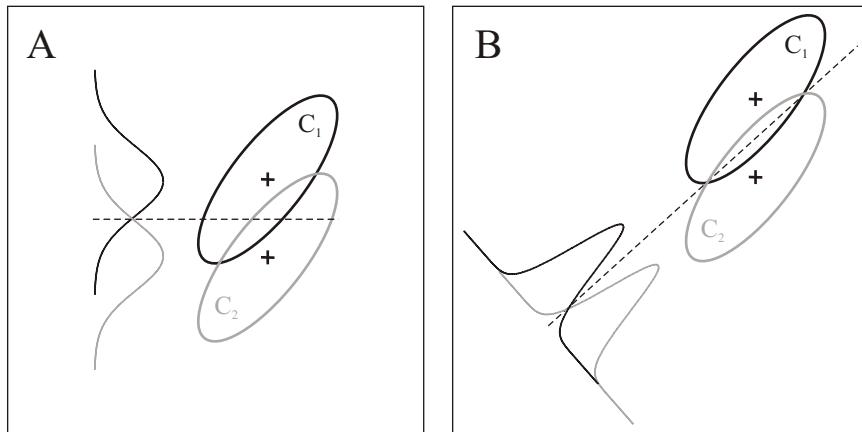


FIGURE 3.2: Segmentation of data classes C_1 and C_2 using Fisher's Linear Discriminant. Although the separation between the two class means is the same for lines A and B, the direction of the line B minimizes the overlap between the two classes and reduces the within class variance. From Hastie et al. (2001).

distribution was calculated. The Gulf Stream was identified where the skewness changed from positive to negative as it was crossed from east to west. Following this work Coulter (1983) performed an automatic detection of Gulf Stream Rings by combining the mean, standard deviation and gradient within windows with *a-priori* frontal probabilities (based on extensive historical data sets) to classify pixels according to Bayes' decision theory.

A further attempt at extracting eddy and frontal structure from images was made by Nichol (1987) using a novel structural relationship approach. Eddy like structures were identified by searching for isolated regions of high temperature enclosed within areas of cold temperature (and vice versa) on a 'region adjacency graph' showing the connections between regions of constant gray-level.

Holland and Yan (1992) made use of a simple Ordered Statistical Edge Detector as part of a tracking scheme for thermal features off the Delaware and New Jersey Coast and in the California Current System. Incrementally moving a window across a image, ordered sets of vectors representing the temperature values from highest to lowest in each window were created. An edge was selected when the range of these values, or subset thereof, exceeded a certain threshold.

The edge detection algorithm proposed by Holyer and Peckinpaugh (1989) is the first of three more modern and successful attempts at frontal detection in this category. The gray-level occurrence (GLC) matrix, whose elements $P(T1, T2 | \Delta x, \Delta y)$, are a measure of the frequency with which two intensities (temperatures), $T1$ and $T2$, occur adjacent to, or at some fixed distance ($\Delta x, \Delta y$) from each other in an image is calculated (see Figure 3.3). The GLC matrix for overlapping sub-windows is then used to calculate the Cluster Shade ($C(\Delta x, \Delta y)$), a measure of 'edginess', at the centre point of each window.

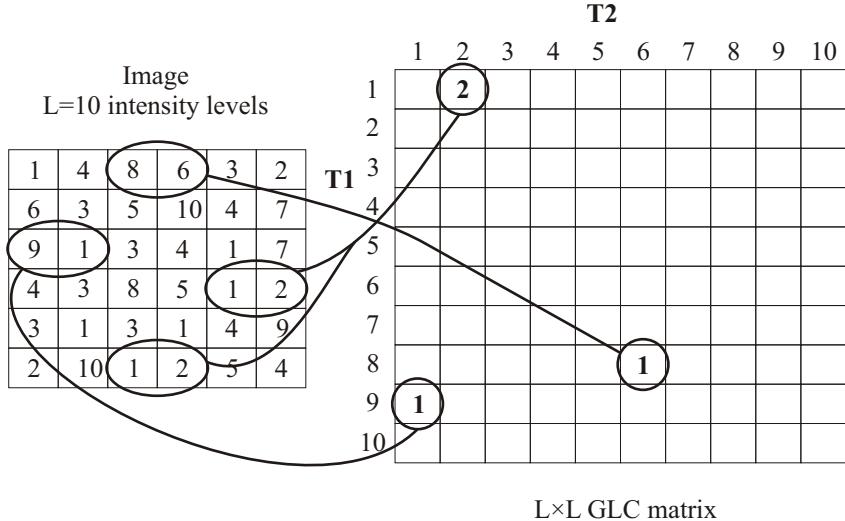


FIGURE 3.3: Construction of the gray-level occurrence matrix (GLC matrix). The $(T_1, T_2)^{th}$ element of the GLC matrix, $P(T_1, T_2 | \Delta x, \Delta y)$, is the relative frequency with which two image intensity levels T_1 and T_2 occur in the image separated by distance $\Delta x, \Delta y$.

For an image with L intensity levels ranging from 0 to $(L - 1)$ the cluster shade is as follows:

$$C(\Delta x, \Delta y) = \sum_{T_1=0}^{L-1} \sum_{T_2=0}^{L-1} (T_1 + T_2 - \mu_{T_1} - \mu_{T_2})^3 \cdot P(T_1, T_2 | \Delta x, \Delta y),$$

where μ_{T_1} and μ_{T_2} are estimates of mean intensity based on weighted summations of rows and columns within the GLC matrix (see Holyer and Peckinpaugh (1989) for full details). When the length displacement vector $(\Delta x, \Delta y)$ is set to zero $C(\Delta x, \Delta y)$ simplifies to being a measure of the asymmetry in the intensity histogram of the image:

$$C(\Delta x, \Delta y) = 2 \sum_{T_1=0}^{L-1} (T_1 - \mu)^3 H(T_1),$$

where $H(T_1)$ is the histogram of intensity values and μ is the mean intensity within each window. Looking for the transition between large positive and negative values of the cluster shade proved an effective way to identify the edge of the Gulf Stream. A threshold is used to determine which zero crossings are significant and may be adjusted to eliminate the detection of weak edges. The algorithm gives a much cleaner delineation of mesoscale features than the Sobel filter seen in Section 3.1. It is interesting to note that Gerson et al. (1979) also report that finding zero crossings in local histogram skew is the best way to locate the north wall of the Gulf Stream.

The histogram based Single Image Edge Detection (SIED) algorithm designed by Cayula and Cornillon (1992) forms the basis for the most widely applied and adapted frontal

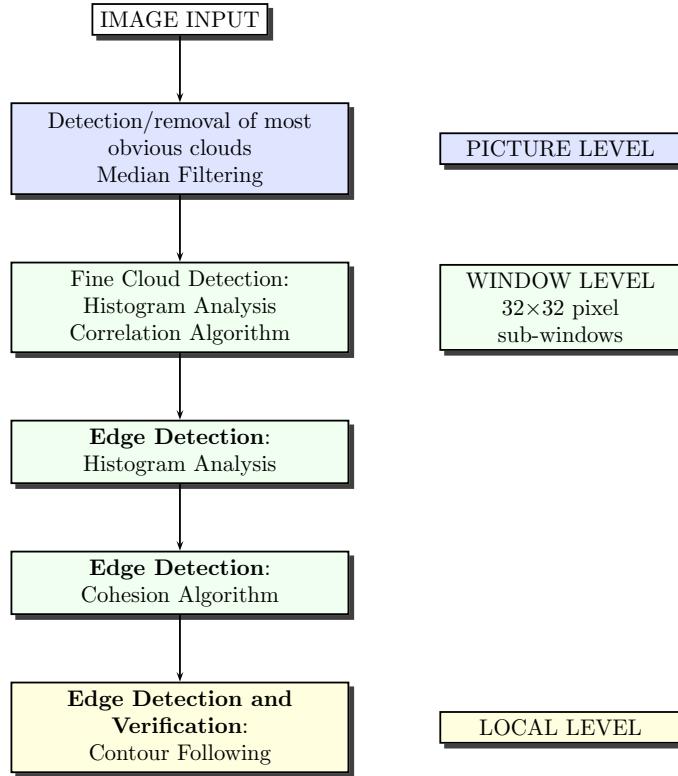


FIGURE 3.4: Flow chart outlining the multistage process of the Single Image Edge Detection (SIED) Algorithm. Adapted from Cayula and Cornillon (1992)

detection scheme to date. It will later be compared to the front detection method proposed in this thesis. The algorithm is implemented in a number of stages. It operates at the picture, window and local level, combining histogram and cohesion analysis with contour-following techniques. Figure 3.4 outlines the various stages of processing involved. The image is firstly cleared of all the most obvious cloud contamination and median filtered. It is subsequently segmented into overlapping 32×32 pixel windows within which small scale cloud detection procedures are performed (details in Appendix A of Cayula and Cornillon (1992)). The first stage of ocean front detection is implemented in the form of histogram analysis.

Figure 3.5 shows that the temperature distribution histogram ($h(t)$) of each window can be used to determine the presence of an edge. If the distribution is unimodal (A), then only one population (watermass) is present. On the other hand if the histogram is bimodal (B and C) then two populations and an edge (front) is present. The statistical significance of an assumed two class segmentation is assessed using a criteria based on the variances of the two populations. The total variance S_{tot} within each window is expressed as the sum of two terms: $J_e(\tau)$ and $J_b(\tau)$. $J_e(\tau)$ is the sum of the variances within each of the two populations resulting from segmenting the region with respect to the threshold τ . This is referred to as the within cluster variance. $J_b(\tau)$ is the variance resulting from the separation of the two clusters. Assuming the histograms exhibit two well defined peaks, if τ is chosen optimally then most of the variance in

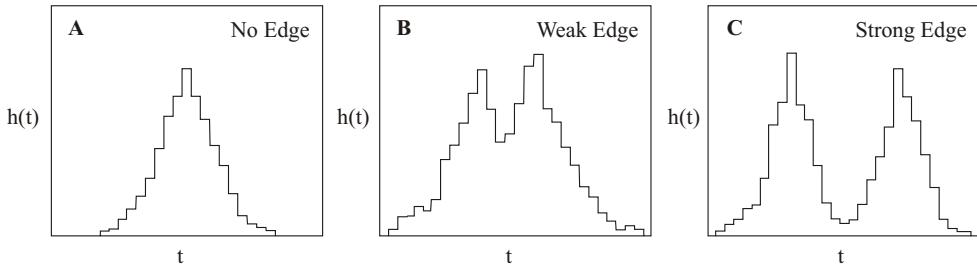


FIGURE 3.5: Example temperature distribution histograms showing A. no edge, B. a weak edge, and C. a strong edge

the window will be due to the difference in temperature between the two populations - $J_b(\tau)$ is large compared to $J_e(\tau)$. If pixels from one population are mistakenly included in the other, some of the terms contributing to the between cluster variance will be shifted to terms contributing to the within cluster variance. The optimal threshold τ_{opt} is therefore the value that maximizes the between cluster variance $J_b(\tau)$. The ratio $\theta(\tau) = \frac{J_b(\tau)}{S_{tot}}$ measures the proportion of the total variance due to the separation between clusters. $\theta(\tau_{opt})$ therefore gives an indication of how good the best segmentation is and is used to decide whether two populations are present or not. A value of 0.7 for $\theta(\tau_{opt})$ is determined to be a reliable discriminant between bimodal and unimodal distributions (Cayula and Cornillon, 1992).

The second stage of the front detection takes into account that clouds and water masses do not always form spatially distinct populations. A bimodal distribution may result from scattered clouds over a sea surface of uniform temperature where no frontal structures are present. The cohesion and smoothness of temperatures within each window is therefore assessed to confirm or reject the presence of any fronts identified by the histogram analysis. Cohesion coefficients for each population (C_1, C_2) and for all data within each window (C) are calculated. Calculation of the cohesion coefficients is explained in Figure 3.6 using an example 10×10 pixel window. High cohesion confirms the presence of an edge. Low cohesion means that the two populations are not spatially distinct and the presence of an edge is unlikely. A threshold of 0.92 for C and 0.90 for C_1 and C_2 is used to eliminate edges resulting from noisy distributions.

The final stage of the algorithm operates at the local level and detects the location of edge pixels creating an image of zeros and ones. The edge image is then subject to a contour following algorithm that links independent edge pixels to form continuous contours.

Improving on the performance of the SIED a multi-image procedure capable of producing maps of persistent fronts was subsequently developed (Cayula and Cornillon, 1995). After initial application of the SIED on an image by image basis, fronts showing similarities in shape and location, and having gradient vectors of similar magnitude within a preset window of time, are identified as persistent. Contour thinning is performed

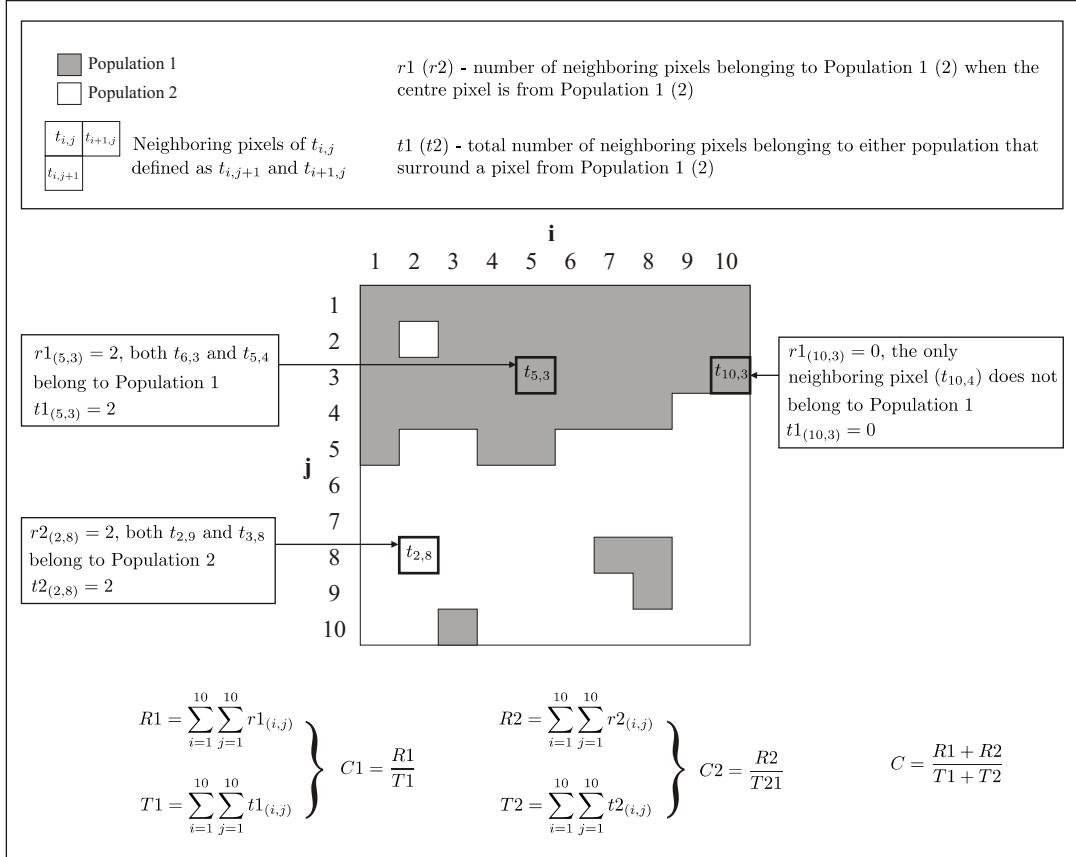


FIGURE 3.6: Calculation of cohesion coefficients C_1 , C_2 and C for an example 10×10 pixel window of data. $R1$ is defined as the total number of comparisons between centre pixels and neighboring pixels that also belong to population one (*vice versa* for $R2$). $T1$ is the total number of comparisons between centre pixels belonging to population one, and neighbor pixels belonging to either population (*vice versa* for $T2$).

and then the SIED reapplied, taking as input both the original series of images and the associated map of persistence. This allows the detection of weak or cloud obscured contours. A second map of persistence is then created using the contours from the second application of the SIED. Finally the SIED is applied a third time using the revised map of persistent contours, the original images and a multi-image cloud detection algorithm.

Kahru et al. (1995) use the SIED to study fronts in the Baltic Sea. The multi-image technique has been applied to timeseries of satellite images by many authors to study, for example, the continental shelf fronts off the northeast US coast (Ullman and Cornillon, 1999, 2001; Mavor and Bisagni, 2001) and fronts in the East China Seas (Hickox et al., 2000). Figure 3.7 shows fronts detected in the Gulf Stream using this multi-image approach (Cayula and Cornillon, 1995).

Miller (2004) uses the SIED as the basis for constructing five day composite front maps with successful application in the English Channel Western Approaches and north-west Iberian Peninsula. The SIED algorithm is adapted for use with SeaWiFS chlorophyll-*a* scenes and the normalized water leaving radiance at 555nm to reveal sediment fronts.

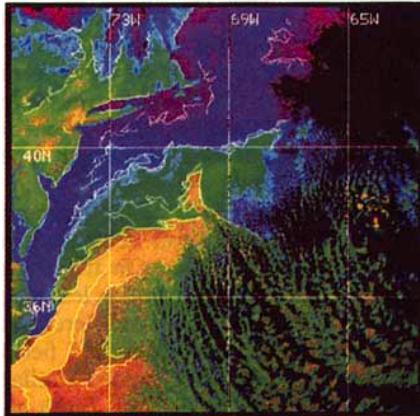


FIGURE 3.7: Fronts detected by the multi-image algorithm of Cayula and Cornillon (1995) in the Gulf Stream. All images within 2.5 days of the base image of 26th March 1985 were used. From Cayula and Cornillon (1995)

The significance of each front detected using the SIED is determined by its gradient and persistence. The gradient magnitude at each frontal pixel is weighted by the probability of observing a front at that location during the time sequence and by its spatial proximity to features detected at different times. More significant fronts are indicated by darker lines in the composite maps. Composite sediment, chlorophyll and thermal front maps are then combined into a single multi-spectral map, an example of which is given in Figure 3.8. These images are designed to help understand the relationships between physical and biological properties in the ocean.

Chin and Mariano (1997) use both spatial and temporal information to estimate the position of ocean fronts through cloud occluded regions of IR AVHRR images. A front's position is reconstructed by smoothing in both space and time fragmented position data (a single contour) obtained from satellite observations with the added dimension of an adaptive feature detection/matching scheme. This motion compensation technique is able to account for the advection and deformation of frontal features (e.g. meanders) between time frames and therefore reduces the chances of introducing spurious features. The algorithm has been shown to reliably interpolate position measurements of the surface temperature fronts associated with the highly dynamic Gulf Stream and Kuroshio currents.

Based on work by Barranco-López et al. (1995), Vazquez et al. (1999) and Shimada et al. (2005) both use the Jensen-Shannon Divergence (JSD) as a means of segmenting SST images and identifying frontal features. The JSD is a measure of the cohesion between two probability distributions and is calculated for each pixel by sliding a moving window over the image. As the JSD increases so too does the difference between the two distributions, and thus the probability of a front. By introducing a ridge-line extraction filter to delineate frontal features more sharply, and producing composite frontal probability, mean SST gradient magnitude, and warm/cold front probability maps, Shimada et al. (2005) provide a detailed study of the finer scale Japanese coastal fronts.

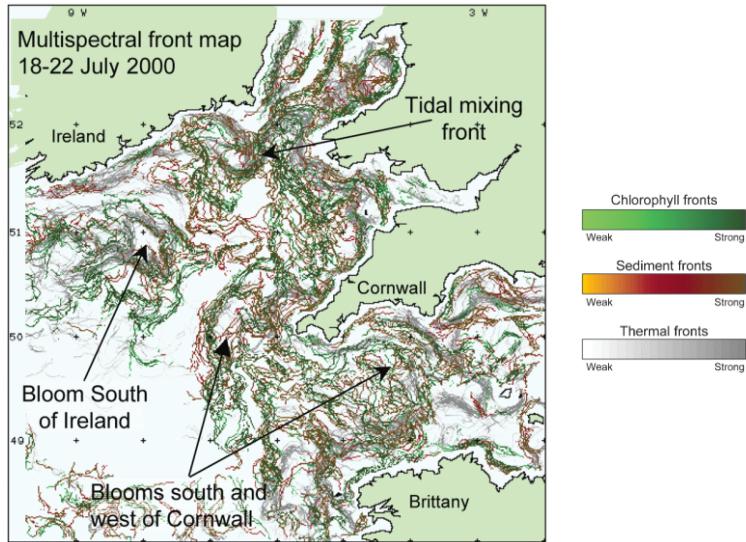


FIGURE 3.8: Multi-spectral front map of the Western Approaches between 18-22 July 2000. The darker the shading the more significant the front. From Miller (2004)

3.3 Gradient Magnitude Criteria

Perhaps the simplest way in which large basin scale fronts may be mapped and studied from thermal satellite images is by setting a gradient magnitude threshold. Southern Ocean Fronts in particular have been mapped in this way by many authors who look for strong temperature gradients within windows of data (Belkin and Gordon, 1996; Kostianoy et al., 2004) - the Polar Front being particularly well studied (Moore et al., 1997, 1999; Dong et al., 2006). Using infrared AVHRR SST images Moore et al. (1997, 1999) define a strong gradient at the Polar Front, calculated in four different directions, as a change in SST greater than or equal to 1.35°C over a distance of 45-65km. The width and gradient of any identified feature is estimated by measuring the distance equatorward from the southern side of the front until the SST no longer increases over 20-30km. Dong et al. (2006) use AMSR-E microwave data and calculate the absolute SST gradient per window. They then define the Polar Front as the southern most location at which the absolute gradient exceeds $1.5 \times 10^{-2} \text{ }^{\circ}\text{C}.\text{km}^{-1}$. In case of weak gradients this threshold is relaxed, and consideration given to the spatial and thermal continuity in the region. Park et al. (2004) calculate the gradient magnitude from AVHRR SST images to study fronts in the East Japan Sea, and Wang et al. (2001) use the average of the three highest gradients from calculations over eight different directions to map frontal features in the northern South China Sea. To study the seasonal variability, Wang et al. (2001) and Hickox et al. (2000) use the Pathfinder Monthly SST Climatology to calculate mean monthly temperature differences across the mapped fronts.

3.4 Surface Fitting

The surface fitting approach to be discussed is the ‘front following’ algorithm of Shaw and Vennell (2000a), which has been used to study the variability of The Southland Front, New Zealand (Shaw and Vennell, 2001). The approach was later modified by Lou et al. (2005) to describe the Zhejiang-Fujian Coastal Front in the East China Sea. An S-shaped function - the hyperbolic tangent (Equation 3.1) - is fitted using least-squares to data extracted within a 20×30 km window centered at the front. The window is rotated to the approximate direction of the front to give a more homogeneous distribution of extracted temperatures. The approach is unique in the sense that it provides estimates not only of the frontal location, but also of key frontal parameters: the mean temperature at the front, its width, and the temperature difference between the water masses on either side. The function used to describe an ocean front is shown in Figure 3.9 and is expressed as:

$$\text{Temperature} = T0 + b \tanh \left[\frac{-X' \cos(\theta) - Y' \sin(\theta) + C}{a} \right]. \quad (3.1)$$

The parameters are defined with respect to the line of inflection (see Figure 3.9) where $T0$ ($^{\circ}\text{C}$) is the mean temperature of the two water masses, b ($^{\circ}\text{C}$) is half the temperature difference across the front, and a (km) is the horizontal distance over which the temperature change is 76.2% of b . $2a$ (km) is taken to be the width of the front. θ ($^{\circ}$) is the direction of the front relative to the extraction window, and C (km) is the distance from the centre of the extraction window to the line of inflection. Calculation of the orientation allows the extraction window to be stepped 2km further along the projected line of inflection resulting in a tracking routine across the image. Parameters estimated in one window are used as initialization conditions for the nonlinear least squares optimization in the next. In the first window of each image parameter values estimated to be typical of the Southland Front were used to start the optimization. Limits are placed on the parameter values and on the number of interations taken for the optimization to converge on a solution. This helps to stabilize the algorithm in the presence of cloud, background noise and where the frontal position has a high level of curvature.

In comparison to the traditional derivative based Prewitt edge detector, which outperforms most of the other gradient operators when there is a significant level of noise present in an image (Simpson, 1990), the ‘front following’ algorithm performs well. It provides the most satisfactory estimates of gradient from an artificially generated set of fronts, and deviates less than the Prewitt operator from the true value when the level of noise increases (Shaw and Vennell, 2000a). No comparisons however have been made with the edge detection techniques of Cayula and Cornillon (1992, 1995) or Vazquez et al. (1999) for example. The ‘front following’ algorithm is the standard against which results of this study are to be compared. A statistical model fitting approach such as this is able to generate a data set ideally suited to teleconnection and variability studies.

One of the disadvantages of this technique is the bias introduced by user interaction. Since the algorithm is unable to track through cloud, the front must be broken up into a series of segments, the processing of each requiring initialization parameters provided by the user. In an attempt to eliminate this bias Lou et al. (2005) apply a Prewitt gradient operator to automatically locate the front and initiate the ‘front following’ algorithm at each segment.

The fixed window size used by Shaw and Vennell (2000a) is also a limitation of the technique since it restricts the smallest resolvable feature to 20km in an along-front direction. This is evident when the algorithm is compared to a 3×3 Prewitt edge detector. The ‘front following’ technique fails to capture the smaller scale features resolved by the gradient operator. The fixed size may be optimal for one segment, providing a sufficient amount of smoothing while not blurring oceanic features of interest, but be sub-optimal for another where the length scale of features or number of available observations has changed. In regions where long narrow plumes protrude from the front Shaw and Vennell (2000a) reduce the window width from 30km to 20km but this is not a global solution to the problem of navigating small scale, narrow features. Lou et al. (2005) propose an adaptable window size based on the semivariance - a statistical measure of the spatial correlation between points of data separated by specific lags. The semivariance γ at lag h is calculated by summing the squared differences of all data pairs $(X_i$ and $X_{i+h})$ along a line of length n :

$$\gamma(h) = \sum_i^{n-h} \frac{[X_i - X_{i+h}]^2}{2n}.$$

The window size is determined by the range (distance) at which data are no longer correlated in the x and y directions.

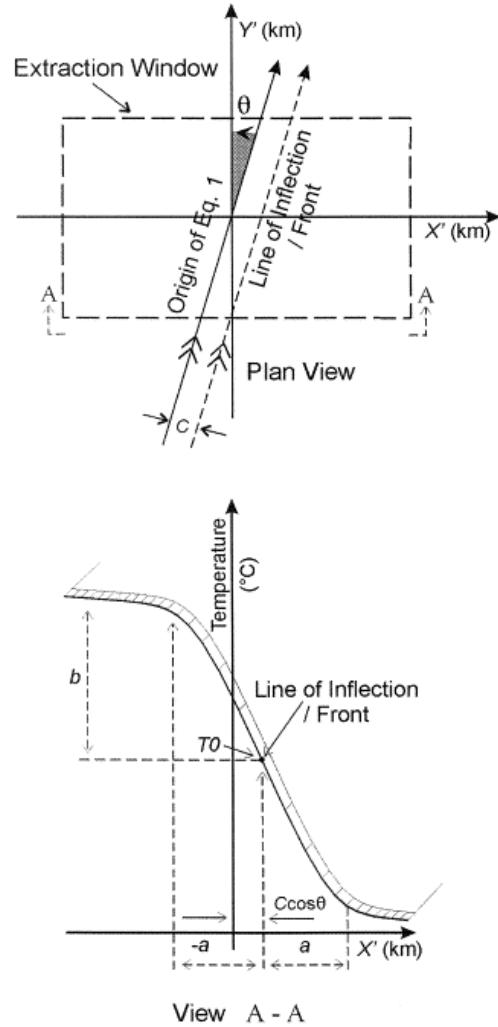


FIGURE 3.9: A plan and side view of the function (Equation 3.1) fitted to temperature values within each extraction window. From Shaw and Vennell (2000a). Note the reference to Equation 1 in this figure refers to Equation 3.1 in the main text.

3.5 Supervised Learning Techniques

An alternative to the statistical classification techniques discussed thus far are neural networks, a supervised machine learning approach. In comparison to the unsupervised learning techniques already presented supervised learning relies on a background of previous input-output training examples in order to classify different data sets.

Askari and Zerr (1998) propose a network trained on idealized shapes, patterns, textures and intensity changes as a means of automatically detecting and classifying different types of oceanographic fronts from multi-source satellite imagery. Relationships between intensity changes in the roughness, colour and temperature fields represent different physical processes in the ocean. For example, distinct step like changes in surface roughness seen in SAR (Synthetic Aperture Radar) images tend to be due to two physical mechanisms: wind stress changes induced by thermal stability variations near SST fronts, or wind stress changes caused by impulsive gusts of wind. In the temperature field distinct boundaries are associated with large scale fronts, currents and eddies. Narrow pulses of intensity change above or below a background temperature or colour level suggest river and estuarine plumes, current filaments and jets. In SAR imagery this pattern is related to velocity fronts or slicks. By combining SAR and thermal AVHRR imagery, a methodology for discriminating between wind induced and sea surface temperature induced roughness fronts is described. A strategy for automatically detecting salinity fronts is also proposed based on fusing the signatures from the roughness and color fields.

Benediktsson et al. (1990) explore the performance of neural networks as a means of classifying multi-source remote sensing data versus the more classical statistical approach. Neural networks have the advantage of being distribution free, thus no prior knowledge about the statistical distribution of classes is required. However, their ability to successfully classify is dependent upon the size and variety of the training data set. Statistical approaches need only an appropriate model of each class.

An alternative technique, that of support vector machines, has been more recently applied to satellite data (Pal and Mather, 2005) and represents a promising development in supervised machine learning not yet widely used by remote sensing scientists.

3.6 Ocean Color and Altimetry in Front Detection

As intimated above, ocean surface temperature observations from infrared (AVHRR, ATSR, MODIS), and microwave (SMMR, TRMM, AMSR) sensors are not the only data sets available from which ocean fronts can be detected. Observations of ocean color from the visible bands of sensors such as SeaWiFS, and ocean surface topography fields

from the microwave bands of altimeters onboard the ERS and TOPEX/POSEIDON satellites, are increasingly being used as complimentary or alternative sources of data. Gille (1994) maps the paths of the Subantarctic and Polar Fronts using Geosat altimeter data. Miller (2004), Takahashi and Kawamura (2005) and Valavanis et al. (2005) all use a combination of SeaWiFS chlorophyll-*a* images and thermal data to map ocean fronts. Legeckis et al. (2002) also make use of SeaWiFS images to establish the exact path of ocean fronts during the animation of GOES (Geostationary Operational Environmental Satellite) SST composites to visualize the daily motions of thermal fronts.

There are a number of situations in which using ocean colour rather than IR SST images is advantageous. Firstly, riverine plume fronts where water masses with distinctly different optical properties but similar temperatures meet may be more visible in ocean colour images. Likewise, during the summer months when phytoplankton blooms occur along ocean fronts, colour images may be more effective in distinguishing between productive and non-productive bodies of water. Additionally, seasonal variability of the upper mixed layer owing to different levels of surface heating and wind forcing may mask and/or decouple the thermal signature of a front from its structure below. During the summer months increased solar insolation makes the SST uniform and the surface thermal signature of a front may disappear (Ullman and Cornillon, 1999; Hickox et al., 2000). Diurnal warming can also create a thin isothermal surface layer that masks the separation between water masses. In these circumstances ocean colour images that are not affected by surface heating may give a clearer and more accurate delineation between water masses bordering a front.

In the following chapter a new front detection technique is developed based around the idea of fitting a model function to a set of SST observations taken at right angles across a front.

CHAPTER 4

Statistical Modelling of Ocean Fronts

Chapter 3 discusses the current techniques employed to detect and monitor ocean fronts, from the early classical gradient edge detection techniques (Simpson, 1990; Janowitz, 1985), through to the more sophisticated histogram based approach of Cayula and Cornillon (1992, 1995), and the surface fitting technique of Shaw and Vennell (2000a). This chapter details the design and theory behind the new front monitoring algorithm proposed in this thesis. The method relies upon a weighted local likelihood technique in which estimates of frontal parameters are based around weighted contributions from surrounding data points. The local likelihood expression for the fit of the model function is maximized with respect to unknown model parameters, representative of the front's location, temperature and strength. Such an approach allows us to estimate a smooth non-parametric spatial trend in frontal characteristics. The following sections introduce the theory of maximum likelihood and detail the parametric function used to model the surface temperature change across a front. Local likelihood regression is then discussed together with the role and choice of an optimal smoothing parameter.

4.1 Algorithm Design Objectives

An overview of the variety of different techniques used to detect ocean fronts from remote sensing data has already been given. We wish to combine the desirable features of some of these techniques in the design of our front detection tool, while avoiding as many of their pitfalls and limitations as possible. Our design objectives are as follows:

I Target Specific

With the exception of the ‘front following’ algorithm (Shaw, 1998; Shaw and Vennell, 2000a), all other front detection methods are ‘non specific’, i.e. no particular frontal structure is targeted. The classical derivative based edge detectors, the statistical feature segmentation and contour following techniques of Cayula and Cornillon (1992), Holyer and Peckinpaugh (1989), and Vazquez et al. (1999), and the multi-image algorithms of Miller (2004) and Shimada et al. (2005), all produce maps with multiple fronts. This approach is ideal if a broad, perhaps exploratory, variability study of the surface structure and dynamics of a region is required. Indeed such techniques may reveal previously unknown features and characteristics of an area that are persistent only on very short time scales. In this thesis we aim to develop a technique best suited to a detailed study of a known, persistent front, such as the Southland Front. The multiple, smaller and weaker features picked out by many current algorithms are often oceanographically less important than the stronger fronts associated with geostrophically driven currents and the transport of heat, nutrients, chemicals and momentum.

II Measurement and uncertainty estimation of frontal characteristics

The front detection scheme developed by Shaw and Vennell (2000a) is unique in the sense that it has the ability to monitor along-front changes in key frontal characteristics; temperature, width, temperature difference and position on an image-by-image basis. We aim to build on this idea with the added dimension of being able to assess the uncertainty of each estimate. As yet, no other technique is able to do this.

III Robust to missing, noisy and poor quality data

Radiometric SST images are rarely cloud free, particularly around coastal and frontal zones where cloud formation is common. Even when a cloud free over-pass is achieved, quality control checks lead to suspect quality data being removed, sometimes mistakenly in areas of strong gradients (see Section 5.1.1). SST images are also noisy owing to sensor inaccuracies, atmospheric interference with the signal, and resolution limitations.

Derivative techniques are highly sensitive to impulsive noise (Simpson, 1990), and often require a pre-filtering stage which blurs features of interest and makes detection more difficult. Even some of the more advanced techniques, such as the SIED of Cayula and Cornillon (1992) require a median filtering stage. The Jensen-Shannon segmentation and contour connection schemes of Vazquez et al. (1999) and Shimada et al. (2005) on the other hand, are robust against noise, avoid pre-filtering and retain edge sharpness.

The multi-spectral front mapping approach of Miller (2004) solves the problem of partial cloud cover and avoids a time averaged composite and the subsequent blurring of dynamic mesoscale features. The multi-spectral map however is still a somewhat confusing mixture of frontal structures shaded according to a measure

of their persistence or gradient. We aim to design an algorithm robust to partially cloudy regions and noisy data that does not require a pre-filtering stage.

IV Automation

With the vast quantities of remotely sensed data now available, having a front detection routine that requires minimal user interaction is essential in order to quickly and objectively process large data sets. One of the main draw backs of the ‘front following’ algorithm (Shaw, 1998), is the segmentation of the front due to missing data and the subsequent user interaction required to manually restart processing.

V Adaptability

Attempts to automate image segmentation using thresholding techniques is complicated by the fact that each technique will only be appropriate for certain applications - a definitive and globally applicable solution does not exist (Marcello et al., 2005). We therefore aim to design an algorithm that may be applied in any region, for any frontal length scale, from smaller estuarine fronts (e.g. Rio de La Plata, South America), to large open ocean convergences (e.g Polar Front, Southern Ocean), and to any image resolution.

4.2 A Statistical Ocean Front Model

4.2.1 Maximum Likelihood Theory

An ocean front occurs when water masses of different thermohaline properties meet and interact. There is a marked change in the vertical structure and a sharp horizontal gradient in temperature, salinity and/or chemical properties as the interface between two water masses is crossed (see Chapter 2). Remote thermal infrared and passive microwave sensors allow us to measure the rapid horizontal change in surface temperature in these regions. It is this surface expression that we aim to model.

Let us suppose that $\mathbf{Z} = \{z^{(1)}, z^{(2)}, \dots, z^{(n)}\}$ is a vector of independently observed temperatures at right angles across an ocean front. The probability of these observations being drawn from a given model front may be expressed in terms of a likelihood function. This may be thought of as the formula for the joint probability distribution of the sample \mathbf{Z} . If $p(\mathbf{Z}; \boldsymbol{\theta})$ represents the probability density function of \mathbf{Z} with a vector of unknown model parameters $\boldsymbol{\theta}$, then:

$$\text{Likelihood Function} \equiv l(\boldsymbol{\theta}; \mathbf{Z}) \equiv p(\mathbf{Z}; \boldsymbol{\theta}) = \prod_{i=1}^n p(z^{(i)}; \boldsymbol{\theta}), \quad (4.1)$$

where n is the number of observations across the front. The aim of maximum likelihood estimation is to find the set of values of the unknown parameters $\boldsymbol{\theta}$, that make the likelihood $l(\cdot)$, a maximum.

Now assume that each SST observation $(z^{(i)})$ is drawn from a normal distribution, $\phi(\cdot)$, with mean $\mathbb{E}(z^{(i)})$ and variance $\mathbb{V}(z^{(i)}) = \sigma^2$, which we shall assume remains constant for all i . A normal distribution is chosen in the absence of any other information about the observations and could be changed when necessary to, for example, a Poisson distribution. If the expectation \mathbb{E} is determined by the function $m(\mathbf{Y}; \boldsymbol{\theta})$, where $\mathbf{Y} = \{y^{(1)}, y^{(2)}, \dots, y^{(n)}\}$ is a vector of known distances across the front corresponding to observations $\mathbf{Z} = \{z^{(1)}, z^{(2)}, \dots, z^{(n)}\}$, and the parametric model $m(\cdot)$ has a vector of unknown parameters, $\boldsymbol{\theta} = \{\theta_1, \theta_2, \dots, \theta_q\}$, we may express the likelihood of observed temperatures \mathbf{Z} by:

$$\begin{aligned} l(\boldsymbol{\theta}, \sigma; \mathbf{Z}) &= \prod_{i=1}^n \phi\left(\mathbb{E}(z^{(i)}), \mathbb{V}(z^{(i)})\right) \\ &= \prod_{i=1}^n \phi\left(m(y^{(i)}; \boldsymbol{\theta}), \sigma^2\right) \\ &= \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2}} e^{\left[-\frac{[z^{(i)} - m(y^{(i)}; \boldsymbol{\theta})]^2}{2\sigma^2}\right]}. \end{aligned}$$

Sigma, σ , may be thought of as the standard deviation of noise about the model function. These concepts and assumptions are illustrated graphically in Figure 4.1.

It is often more convenient to maximize the log of the likelihood. This transformation converts products (\prod) into sums (\sum), a much easier mathematical expression to program. Letting $L(\cdot)$ denote the log of the likelihood, $\ln[l(\cdot)]$, we derive:

$$\begin{aligned} L(\boldsymbol{\theta}, \sigma; \mathbf{Z}) &= \sum_{i=1}^n \ln\left(\frac{1}{\sqrt{2\pi\sigma^2}} e^{\left[-\frac{[z^{(i)} - m(y^{(i)}; \boldsymbol{\theta})]^2}{2\sigma^2}\right]}\right) \\ &= \sum_{i=1}^n \ln\left(\frac{1}{\sqrt{2\pi\sigma^2}}\right) + \sum_{i=1}^n \ln\left(e^{\left[-\frac{[z^{(i)} - m(y^{(i)}; \boldsymbol{\theta})]^2}{2\sigma^2}\right]}\right) \\ &= \sum_{i=1}^n \left(\ln(1) - (\ln \sqrt{2\pi} + \ln \sigma)\right) - \sum_{i=1}^n \frac{[z^{(i)} - m(y^{(i)}; \boldsymbol{\theta})]^2}{2\sigma^2}, \end{aligned}$$

Since $\ln(1) = 0$, and ignoring $-\ln \sqrt{2\pi}$ as an irrelevant constant,

$$L(\boldsymbol{\theta}, \sigma; \mathbf{Z}) = -n \ln \sigma - \frac{1}{2\sigma^2} \sum_{i=1}^n (z^{(i)} - m(y^{(i)}; \boldsymbol{\theta}))^2. \quad (4.2)$$

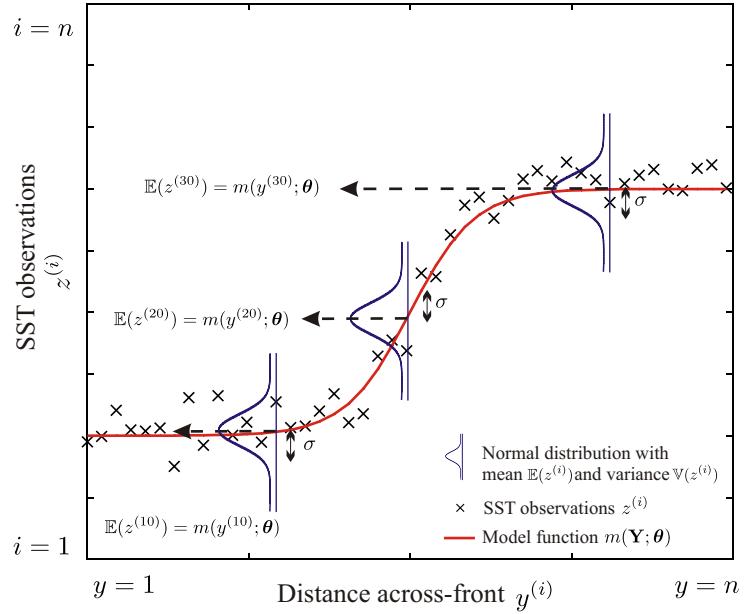


FIGURE 4.1: Schematic illustrating the assumption that each SST observation $z^{(i)}$ at distance $y^{(i)}$ across the front is drawn from a normal distribution with mean $\mathbb{E}(z^{(i)})$ and standard deviation σ . The expectation \mathbb{E} is determined by the model function $m(y^{(i)}; \boldsymbol{\theta})$.

The maximum likelihood estimates $\hat{\boldsymbol{\theta}}$ and $\hat{\sigma}$ are attained when the rate of change of $L(\cdot)$, with respect to the unknown parameters $\boldsymbol{\theta}$ and σ , equals zero. The first order partial derivatives are:

$$\frac{\partial L}{\partial \theta_q} = -\frac{1}{\sigma^2} \sum_{i=1}^n -\left(z^{(i)} - m(y^{(i)}; \boldsymbol{\theta})\right) \frac{\partial m}{\partial \theta_q} = 0 \quad (4.3a)$$

$$\frac{\partial L}{\partial \sigma} = -\frac{n}{\sigma} + \frac{1}{\sigma^3} \sum_{i=1}^n \left(z^{(i)} - m(y^{(i)}; \boldsymbol{\theta})\right)^2 = 0 \quad (4.3b)$$

To ensure that $L(\cdot)$ is maximized, it must be the case that the slope of $L(\cdot)$ is decreasing near the maximum likelihood estimate (MLE). This is the second order condition, given by the expression for the second partial derivative of $L(\cdot)$ with respect to the unknown parameters, $\frac{\partial^2 L}{\partial \theta^2} < 0$ and $\frac{\partial^2 L}{\partial \sigma^2} < 0$.

The maximum likelihood estimates are obtained via an iterative optimization routine known as Newton-Raphson. Details of this algorithm together with the practicalities of its implementation are given in Section 5.2.1 of Chapter 5.

Note that in the following sections we will refer collectively to all unknowns using the vector $\boldsymbol{\theta}$, the last element of which being σ .

4.2.2 A Model Ocean Front

We now turn our attention to the specific form of the function $m(\cdot)$ used to model the change in surface temperature across a front. We require a sigmoid (S-shaped) function that is able to emulate the steep thermal gradient at the interface between two water masses with disparate surface temperatures. Previous work by Shaw (1998) and Lou et al. (2005) uses the hyperbolic tangent function, and we adopt the same function here to represent a cross section of sea surface temperature observations.

$$\mathbf{Z} = m(\mathbf{Y}; \boldsymbol{\theta}) + \epsilon = \theta_1 + \theta_2 \tanh \left[\frac{\mathbf{Y} + \theta_4}{\theta_3} \right] + \epsilon, \quad (4.4)$$

where θ_1 is the front's mean temperature. $2\theta_2$ and $2\theta_3$ define the temperature difference and width respectively. θ_4 is a translation parameter determining the position of the front. The noise, ϵ , is assumed to be normally distributed with zero mean and standard deviation σ . Figure 4.2 is a graphical representation of the how the model parameters may be interpreted in a more physical sense.

Using this model, and setting $\boldsymbol{\theta} = \{\theta_1, \theta_2, \theta_3, \theta_4, \sigma\}$, the log-likelihood may now be expressed in full as follows:

$$L(\boldsymbol{\theta}; \mathbf{Z}) = -n \ln \sigma - \frac{1}{2\sigma^2} \sum_{i=1}^n \left(z^{(i)} - \left(\theta_1 + \theta_2 \tanh \left[\frac{y^{(i)} + \theta_4}{\theta_3} \right] \right) \right)^2 \quad (4.5)$$

This likelihood expression forms the basis of the statistical ocean front model.

4.2.3 Maximum Likelihood Properties and Confidence Bounds

The advantage of using maximum likelihood over other parameter estimation techniques such as least-squares are its statistical properties that allow the construction of confidence intervals around $\hat{\boldsymbol{\theta}} = \{\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\sigma}\}$. Note that in this particular case, where the errors in predicting the observations are assumed to follow a normal distribution, maximum likelihood estimation is equivalent to least-squares. Asymptotically (as $n \rightarrow \infty$), the maximum likelihood estimator is unbiased, has the smallest possible variance and is consistent. Thus as $n \rightarrow \infty$, it holds that

$$\begin{aligned} \mathbb{E}(\hat{\boldsymbol{\theta}}) &= \boldsymbol{\theta} \\ \text{Var}(\hat{\boldsymbol{\theta}}) &= -\mathbf{H}^{-1}(\boldsymbol{\theta}) \\ \hat{\boldsymbol{\theta}} &\approx \phi(\boldsymbol{\theta}, \text{Var}(\hat{\boldsymbol{\theta}})), \end{aligned}$$

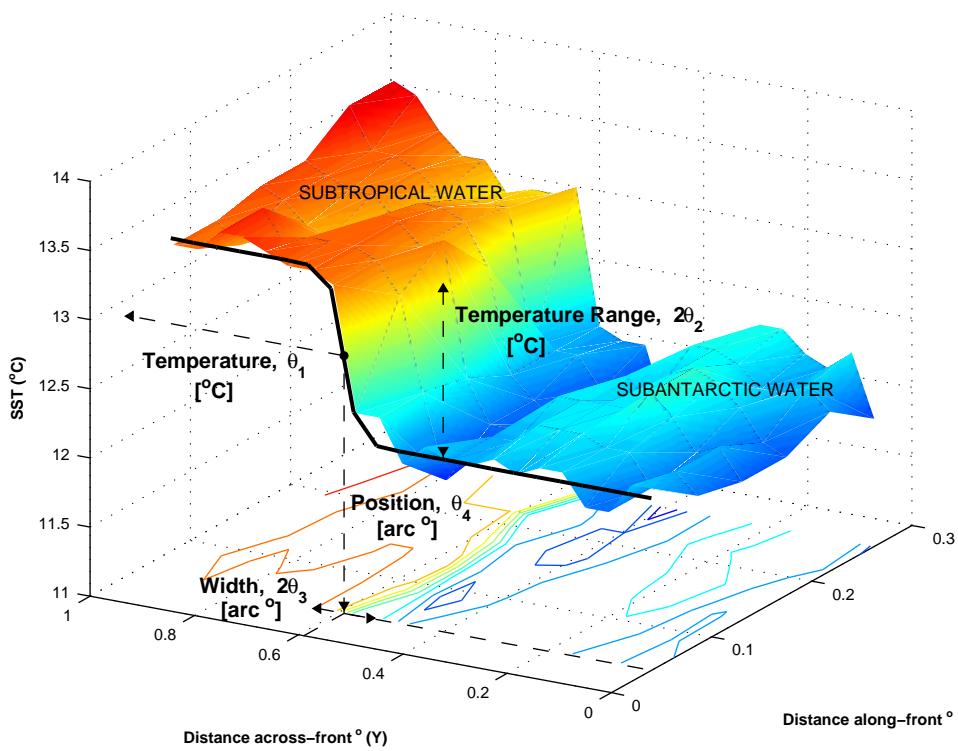


FIGURE 4.2: Physical interpretation of parameters in model

$$m(\mathbf{Y}; \boldsymbol{\theta}) + \epsilon = \theta_1 + \theta_2 \tanh \left[\frac{\mathbf{Y} + \theta_4}{\theta_3} \right] + \epsilon$$

where \mathbf{H} , the Fisher Information Matrix, is a matrix of second order partial derivatives with respect to the unknown parameters. That is, the distribution of estimates is asymptotically normal with the true (but unknown) parameters $\boldsymbol{\theta}$ as expectation. For small n this is an approximation (Davison, 2003). So the estimated variances of the maximum likelihood estimates are equal to the diagonal elements of, $-\mathbf{H}^{-1}(\hat{\boldsymbol{\theta}})$, the asymptotic variance-covariance matrix. The off-diagonal elements are the co-variances between different estimated parameters. The square root of the diagonal elements yields the standard errors. Given the asymptotically normal distribution of the maximum likelihood estimates, confidence intervals for $\hat{\boldsymbol{\theta}}$ may be constructed,

$$\hat{\boldsymbol{\theta}} - t_{\alpha, v} \cdot \text{Var}(\hat{\boldsymbol{\theta}})^{\frac{1}{2}} ; \quad \hat{\boldsymbol{\theta}} + t_{\alpha, v} \cdot \text{Var}(\hat{\boldsymbol{\theta}})^{\frac{1}{2}},$$

where $t_{\alpha, v}$ is the t -value for a t -distribution with v degrees of freedom for a $100(1 - \alpha)\%$ confidence interval. v is defined as $n - 5$, the number of observations minus the number of unknowns being estimated. If $t_{0.025, v}$ is the 97.5th percentile of the t -distribution with v degrees of freedom, for repeated estimates there is a 2.5% chance that $\hat{\boldsymbol{\theta}} > \hat{\boldsymbol{\theta}} + t_{0.025, v} \cdot \text{Var}(\hat{\boldsymbol{\theta}})^{\frac{1}{2}}$, and a 2.5% chance that $\hat{\boldsymbol{\theta}} < \hat{\boldsymbol{\theta}} - t_{0.025, v} \cdot \text{Var}(\hat{\boldsymbol{\theta}})^{\frac{1}{2}}$. The probability that $\hat{\boldsymbol{\theta}}$ lies

between $\hat{\boldsymbol{\theta}} \pm t_{0.025, n} \cdot \text{Var}(\hat{\boldsymbol{\theta}})^{\frac{1}{2}}$ is thus 95%. Throughout this thesis we will always use $\alpha = 0.025$ so that our confidence intervals are for a 95% probability.

For small n , as is the case here, the distribution of estimates is only approximately normal, therefore a t -distribution ¹, rather than a z -distribution (normal distribution) is used. As the number of observations and therefore the degrees of freedom increase the t -distribution approaches a normal distribution with a mean of zero and unit variance.

The relationship between confidence and the matrix of second derivatives (\mathbf{H}) is not immediately obvious. Consider the likelihood, $L(\boldsymbol{\theta}; \mathbf{Z})$, as defining a multidimensional surface, its rate of curvature being described by \mathbf{H} . If at its maximum the likelihood surface is reasonably flat, then there is considerable uncertainty in the estimate, and the low rate of curvature leads to a small second derivative, which once inverted produces a large standard error. Conversely a sharp peak at the maximum implies a considerable degree of certainty in the estimate. A high rate of curvature makes the second derivative large and the standard error small.

From Equation 4.5 we derive the first and second order partial derivatives necessary for the construction of confidence bounds. Setting

$$T = \tanh \left[\frac{y^{(i)} + \theta_4}{\theta_3} \right] \quad \text{and} \quad R = z^{(i)} - \left(\theta_1 + \theta_2 \tanh \left[\frac{y^{(i)} + \theta_4}{\theta_3} \right] \right),$$

$$S = \operatorname{sech} \left[\frac{y^{(i)} + \theta_4}{\theta_3} \right]^2$$

the set of first partial derivatives are:

$$\begin{aligned} \frac{\partial L}{\partial \theta_1} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n -2R \\ \frac{\partial L}{\partial \theta_2} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n -2 \tanh \left[\frac{y^{(i)} + \theta_4}{\theta_3} \right] R \\ \frac{\partial L}{\partial \theta_3} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n \frac{2\theta_2(\theta_4 + y^{(i)}) \operatorname{sech} \left[\frac{y^{(i)} + \theta_4}{\theta_3} \right]^2 R}{\theta_3^2} \\ \frac{\partial L}{\partial \theta_4} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n \frac{-2\theta_2 \operatorname{sech} \left[\frac{y^{(i)} + \theta_4}{\theta_3} \right]^2 R}{\theta_3} \\ \frac{\partial L}{\partial \sigma} &= -\frac{n}{\sigma} + \frac{1}{\sigma^3} \sum_{i=1}^n R^2. \end{aligned}$$

¹the probability distribution that arises in the problem of estimating the mean of a normally distributed population when the sample size is small

The second partial derivatives are as follows:

$$\begin{aligned}
\frac{\partial^2 L}{\partial \theta_1^2} &= -\frac{n}{\sigma^2} \\
\frac{\partial^2 L}{\partial \theta_1 \partial \theta_2} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n 2T \\
\frac{\partial^2 L}{\partial \theta_1 \partial \theta_3} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n -\frac{\theta_2}{\theta_3^2} (\theta_4 + y^{(i)}) S \\
\frac{\partial^2 L}{\partial \theta_1 \partial \theta_4} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n \frac{2\theta_2}{\theta_3} S \\
\frac{\partial^2 L}{\partial \theta_1 \partial \sigma} &= \frac{1}{\sigma^3} \sum_{i=1}^n -2R \\
\frac{\partial^2 L}{\partial \theta_2^2} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n 2T^2 \\
\frac{\partial^2 L}{\partial \theta_2 \partial \theta_3} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n \frac{2(\theta_4 + y^{(i)})}{\theta_3^2} SR - \frac{2\theta_2(\theta_4 + y^{(i)})}{\theta_3^2} ST \\
\frac{\partial^2 L}{\partial \theta_2 \partial \theta_4} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n \frac{2\theta_2}{\theta_3} ST - \frac{2}{\theta_3} SR \\
\frac{\partial^2 L}{\partial \theta_2 \partial \sigma} &= \frac{1}{\sigma^3} \sum_{i=1}^n -2TR \\
\frac{\partial^2 L}{\partial \theta_3^2} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n \frac{2\theta_2^2}{\theta_3^4} S^2 - \frac{4\theta_2(\theta_4 + y^{(i)})}{\theta_3^3} SR + \frac{4\theta_2(\theta_4 + y^{(i)})^2}{\theta_3^4} STR \\
\frac{\partial^2 L}{\partial \theta_3 \partial \theta_4} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n -\frac{2\theta_2^2(\theta_4 + y^{(i)})}{\theta_3^3} S^2 + \frac{2\theta_2}{\theta_3^2} SR - \frac{4\theta_2(\theta_4 + y^{(i)})}{\theta_3^3} STR \\
\frac{\partial^2 L}{\partial \theta_3 \partial \sigma} &= \frac{1}{\sigma^3} \sum_{i=1}^n \frac{2\theta_2(\theta_4 + y^{(i)})}{\theta_3^2} SR \\
\frac{\partial^2 L}{\partial \theta_4^2} &= -\frac{1}{2\sigma^2} \sum_{i=1}^n \frac{2\theta_2^2}{\theta_3^2} S^2 + \frac{4\theta_2}{\theta_3^2} STR \\
\frac{\partial^2 L}{\partial \theta_4 \partial \sigma} &= \frac{1}{\sigma^3} \sum_{i=1}^n -\frac{2\theta_2}{\theta_3} SR \\
\frac{\partial^2 L}{\partial \sigma^2} &= \frac{n}{\sigma^2} - \frac{3}{\sigma^4} \sum_{i=1}^n R^2.
\end{aligned}$$

Note that the model function (Equation 4.4) is nonlinear and that the partial derivatives involve the parameters themselves. An iterative procedure must therefore be used to obtain the maximum likelihood estimate. Further details are given in Section 5.2.1.

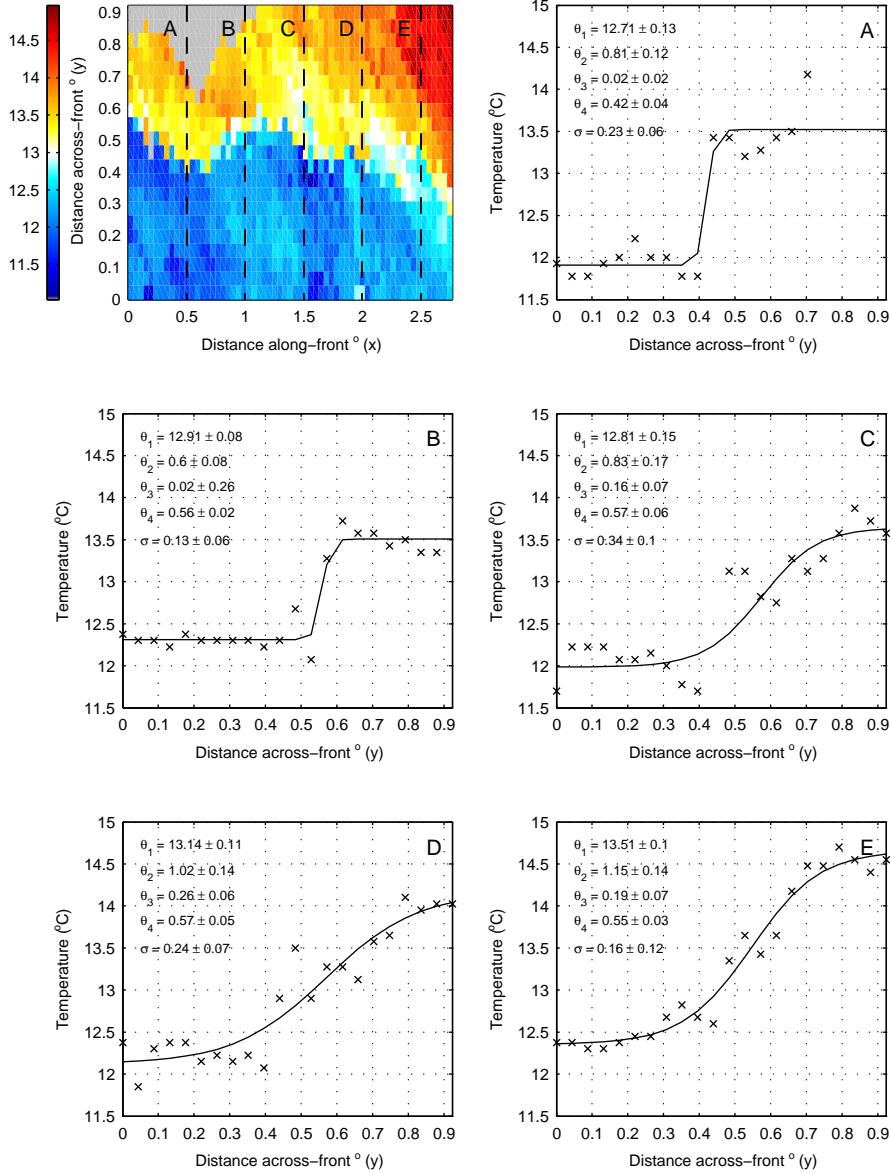


FIGURE 4.3: Examples of the model function fitted to five different sets of SST observations (A-E) across the Southland Front during April 1985 (top left). For each cross section the observed temperature values (Z) are represented by black crosses and the fitted model (Equation 4.4) by a solid black line. Also shown are the specific parameter estimates $\pm 95\%$ confidence intervals. Color scale in $^{\circ}\text{C}$.

Figure 4.3 is an example of the model function fitted to five different sets of SST observations taken along the Southland Front. 95% confidence intervals are given for each estimated parameter. This will henceforth be referred to as standard maximum likelihood estimation. It is clear that the temperature, width, temperature range and position do not remain the same at different locations along the front.

4.3 Modelling Spatial Variability

We are not interested solely in the position and strength of a front at one specific location. We wish to monitor how these characteristics change across ocean basins in response to localized oceanic and atmospheric conditions and interactions such as; stratification, wind stress and mixing, sea level pressure, bathymetry, vorticity constraints, regional eddy activity, circulation dynamics and remote ocean-atmosphere forcing. Ocean fronts are recognized as important carbon dioxide sinks. They are capable of inducing localized perturbations in the wind field which then force changes in the atmosphere above. They are zones of enhanced productivity relied upon by fishing industries worldwide, and are responsible for the redistribution and formation of water masses as part of the thermohaline conveyor. For these reasons it is important that we understand not only the temporal changes but also the spatial variability of ocean fronts.

In this section we begin with a discussion of basic non-parametric regression and how it may be used to smooth a series of observations of frontal characteristics, assumed to be independent and unbiased, in order to reveal the underlying spatial variability (along-front). However, the parameter estimates that we are able to make through standard maximum likelihood (Section 4.2.1) are not without bias. We expand therefore the concept of non-parametric regression into the context of local likelihood in order to provide a smooth description of spatial variations in the position, temperature and strength of an ocean front. Along-front estimates are made by considering a weighted sum of likelihood contributions from surrounding data points. This is a new approach to monitoring ocean fronts from remotely sensed data.

4.3.1 Regression Analysis

Let us begin by considering the general concept of regression analysis. Take, for example, observations of the mean temperature (θ_1) at positions x_j along a front, which for the purpose of this current discussion we shall take to be independent and unbiased. The subscript $j = \{1, 2, \dots, m\}$ is used here as a spatial index rather than as a reference to a matrix element. Smoothing by means of regression analysis may be used to crudely reveal the underlying structure (or trend). The mathematical model of the relationship between the dependent (response) variable that we shall denote θ (the temperature), and the independent variable x_j is known as the regression equation. We drop the subscript from θ_1 , using instead θ to represent the temperature for the purpose of clarity. The regression equation is as follows:

$$\theta_j = \Phi(x_j) + \epsilon_j.$$

The error term ϵ_j picks up the unpredictable part of the response variable not captured by the function $\Phi(x_j)$ (Davison, 2003).

4.3.1.1 Parametric Regression

If the form of the underlying function $\Phi(x_j)$ was known and could be parameterized in terms of a basis function, then a parametric regression could be performed on the data pairs $\{\theta_j, x_j\}$. In its simplest form this is a linear regression.

$$\theta_j = \Phi(x_j) = \alpha + \beta x_j + \epsilon_j,$$

where α and β are coefficients to be estimated from the data using the methods of least-squares or maximum likelihood. Note that for a model with normally distributed errors ϵ_j , the methods of least-squares and maximum likelihood coincide. Parametric models may also take the form of a polynomial, $\theta_j = \alpha x_j^2 + \beta x_j + \gamma$, or a nonlinear exponential function, $\theta_j = \alpha e^{\beta x_j}$.

4.3.1.2 Weighted Non-Parametric Regression

In the context of revealing along-front changes in frontal characteristics, where there may be a considerable nonlinear component to the data, and no obvious regression function to be applied, non-parametric regression provides a means of smoothing in order to identify the underlying structure. The predictor does not take on any specific form but is constructed according to information derived from the data. Important non-parametric or smoothing methods include kernel smoothers, local regression and smoothing splines. These all fall into the class of linear smoothers. The simplest of the smoothing methods is the kernel smoother. It involves making smooth composites by applying a weighted filter to the data. A non-parametric regression often captures the trend of the data better than a parametric model fit since the smoother does not make as rigid an assumption as the specified model about the relationship between variables (Figure 4.4).

Once again, consider that the data θ , available at x , may be described by $\theta = \Phi(x) + \epsilon$, where $\Phi(x)$ is the unknown underlying smooth ² function (or trend), and ϵ denotes an independent normally distributed error term with mean zero and unknown variance. As before, a normal distribution is assumed in the absence of any other information to the contrary. A smooth curve estimate of $\Phi(x)$ may be obtained by constructing a weighted local mean estimator, sometimes called the Nadaraya-Watson estimate (Bowman and Azzalini, 1997).

²after Tibshirani and Hastie (1987) we loosely define *smooth* in a *non-mathematical sense* as a function less smooth than a straight line, but smoother than an interpolating polynomial (as seen in Figure 4.4)

$$\hat{\Phi}(x) = \frac{\sum_{j=1}^m K(x_j - x; h) \cdot \theta_j}{\sum_{j=1}^m K(x_j - x; h)}, \quad (4.6)$$

θ_j being the discretized version of θ , and h a constant smoothing parameter. Weightings are determined by a smooth symmetric and positive kernel function $K(\cdot)$ that peaks at zero to ensure that most weight is given to data closest to the point of interest. The smoothing parameter or bandwidth h controls the width of the kernel and hence the degree of smoothing applied to the data (see Section 4.4 for further discussion). In matrix notation we may write $\hat{\theta} = \mathbf{S}_h \theta$, where $\hat{\theta}$ and θ are $1 \times m$ vectors of estimates and observations respectively, and \mathbf{S}_h is an $m \times m$ matrix of weightings determined by the kernel function $K(\cdot)$.

The normal density function is commonly used as the kernel smoother such that

$$K(x_j - x; h) = \frac{1}{\sqrt{2\pi h^2}} e^{\left[\frac{-(x_j - x)^2}{2h^2} \right]}. \quad (4.7)$$

In this case the smoothing parameter or bandwidth h is the standard deviation of the distribution, and observations over an effective range of $4h$ will contribute to the local mean estimate (Bowman and Azzalini, 1997).

The Nadaraya-Watson estimate is the solution to a weighted least-squares problem where the estimate $\hat{\beta}$ is the minimizer of

$$\min_{\beta} \sum_{j=1}^m (\theta_j - \beta)^2 \cdot K(x_j - x; h). \quad (4.8)$$

$\hat{\beta}$ corresponds to locally approximating $\Phi(x)$ with a constant.

Figure 4.4 shows a series of standard maximum likelihood temperature estimates made along the Southland Front in April 1985 at positions x_j (examples of which are given in Figure 4.3), to which both a kernel smooth and a linear regression have been fitted. The linear model fails to capture the non-linear component of data and misses important features. The kernel smooth gives a more realistic representation of how the mean temperature changes with increasing distance (measured in arc degrees).

An alternative to the construction of a local mean estimate is to fit a local linear regression. The underlying principle is that a smooth function can be well approximated by a linear function in the neighborhood of any point x .

$$\Phi(x) \approx \alpha + \beta(x_j - x)$$

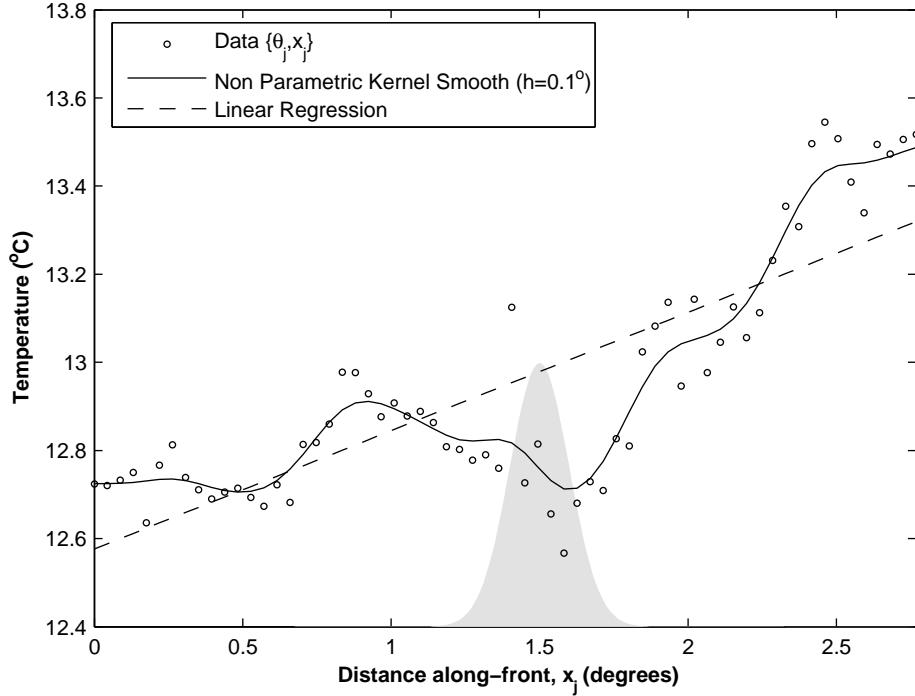


FIGURE 4.4: A local Nadaraya-Watson smooth with a bandwidth of 0.1° (solid line), and a linear regression (dashed line) on the standard maximum likelihood temperature estimates (open circles) along the Southland Front for April 1985 (see Figure 4.3). The width of the kernel smoother is shown by the solid gray gaussian.

The local approximation is fitted by locally weighted least-squares. Coefficient estimates $\hat{\alpha}$ and $\hat{\beta}$ are chosen to minimize

$$\min_{\alpha, \beta} \sum_{j=1}^m \{\theta_j - \alpha - \beta(x_j - x)\}^2 \cdot K(x_j - x; h).$$

The regression estimate for each local least-squares problem is defined as $\hat{\theta} = \hat{\Phi}(x) = \hat{\alpha}$ as this defines the position of the local regression line at point x . As x changes so to do the smoothing weights $K(x_j - x; h)$ and estimates $\hat{\alpha}$ and $\hat{\beta}$. The local mean estimator in Equation 4.8 can be derived by removing the $\beta(x_j - x)$ term from the problem.

Local linear regression as a means of visualizing along-front trends in frontal characteristics assumes that independent and unbiased observations of temperature, gradient and position are available. Using the standard maximum likelihood technique discussed in Section 4.2.1 however provides us with a set of biased parameter estimates $\hat{\theta}$ at each point x . We cannot therefore simply perform a non-parametric regression on a series of standard MLE's. The statistical properties of maximum likelihood (Section 4.2.3) that allow the construction of confidence intervals around each estimate would be lost once the data has been smoothed. Furthermore, poor or missing maximum likelihood estimates in regions where there is a lack of SST measurements could introduce outliers and therefore spurious features into the resulting smoothed profile. Instead we adopt the

method of *local likelihood*, a more robust and statistically correct means of estimating along-front trends.

4.3.2 Local Likelihood Estimation

First introduced by Tibshirani and Hastie (1987), local likelihood is an extension of the local kernel based fitting techniques, seen above, to likelihood based regression models. We recall the maximum likelihood method seen in Section 4.2.1 and expand it into the local likelihood context to create a smooth non-parametric description of spatial variations in frontal position, temperature and strength.

Assume that temperature observations $\mathbf{Z}_j = \{z_j^{(1)}, \dots, z_j^{(n)}\}$ are available at fixed and known positions x_j along the front, where $j = \{1, \dots, m\}$. The observations \mathbf{Z}_j are realizations from the parametric model previously discussed (Equation 4.4):

$$\mathbf{Z}_j = m(x_j, \mathbf{Y}_j; \boldsymbol{\Phi}(x_j)) = \theta_{1,j} + \theta_{2,j} \tanh \left[\frac{\mathbf{Y}_j + \theta_{4,j}}{\theta_{3,j}} \right] + \epsilon_j,$$

where ϵ_j is normally distributed random noise with a mean of zero and unknown variance σ_j . $\mathbf{Y}_j = \{y_j^{(1)}, \dots, y_j^{(n)}\}$ is a vector of known distances across the front indexed by the superscript $i = \{1, \dots, n\}$. We now assume that the unknown parameters $\boldsymbol{\theta}$ are themselves a smoothly varying function of x denoted by $\boldsymbol{\Phi}(x_j)$. $\boldsymbol{\theta}_j = \{\theta_{1,j}, \theta_{2,j}, \theta_{3,j}, \theta_{4,j}, \sigma_j\}$ are as before the front's mean temperature, the across front temperature change, the width, position and standard deviation of noise about the model fit.

We are interested in estimating the smooth function of parameters $\boldsymbol{\theta}$ as they vary along the front, i.e. the function $\boldsymbol{\Phi}(x)$.

A standard approach would be to assume a parametric model for the form of $\boldsymbol{\Phi}(x)$, such as a linear regression where $\boldsymbol{\theta}_j = \boldsymbol{\alpha} + \boldsymbol{\beta}x_j$. The likelihood equation $l(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \prod_{j=1}^m m(x_j, \mathbf{Y}_j; \boldsymbol{\Phi}(x_j; \boldsymbol{\alpha}, \boldsymbol{\beta}))$ is then solved to obtain parameter estimates $\hat{\boldsymbol{\alpha}}$ and $\hat{\boldsymbol{\beta}}$ and a fitted across-front temperature profile $\hat{\mathbf{Z}}_j = m(x_j, \mathbf{Y}_j; \hat{\boldsymbol{\Phi}}(x_j))$. However, as before, such a parametric approach is not justified in the case of frontal modelling. We can not presume to know how a parameter may vary in such a complex system.

In contrast, the **local likelihood** method assumes only that $\boldsymbol{\theta}$ is a smooth function of x and is an ideal alternative approach. We estimate the coefficients of the function $\boldsymbol{\Phi}(x)$ locally at each discrete point x_j . Firstly consider the most basic case where $\boldsymbol{\Phi}(x_j)$ is assumed approximately constant at points close to x_j , i.e. no particular model for the behavior of the parameters near x_j . Denoting the log-likelihood associated with $m(\cdot)$ from the j^{th} set of temperature observations as:

$$L(\boldsymbol{\theta}_j; \mathbf{Z}_j) = -n \ln \sigma_j - \frac{1}{2\sigma_j^2} \sum_{i=1}^n \left[z_j^{(i)} - \left(\theta_{1,j} + \theta_{2,j} \tanh \left[\frac{y^{(i)} + \theta_{4,j}}{\theta_{3,j}} \right] \right) \right]^2, \quad (4.9)$$

a simple local likelihood estimator for $\boldsymbol{\theta}_j$ is of the form:

$$\hat{\boldsymbol{\theta}}_j = \max_{\boldsymbol{\theta}_j} \sum_{k=1}^m K(x_k - x_j; h) \cdot L(\boldsymbol{\theta}_j; \mathbf{Z}_k), \quad (4.10)$$

where $K(x_k - x_j; h)$ is a normal smoothing function (Equation 4.7) such that $\sum_{k=1}^m K(x_k - x_j; h) = 1$, with bandwidth $h > 0$. Note that j is fixed with k varying over the points $k = \{1, \dots, m\}$. The estimator $\hat{\boldsymbol{\theta}}_j$ is the value of $\boldsymbol{\theta}_j$ which maximizes the weighted sum of likelihood contributions $w(x_k, x_j) L(\boldsymbol{\theta}_j; \mathbf{Z}_k)$ in which the weights $w(x_k, x_j) = K(x_k - x_j; h)$ are dependent upon the separation of x_k and x_j . By solving each of these weighted local likelihood problems at each position x_j we obtain a series of smooth parameter estimates $(x_j, \hat{\boldsymbol{\theta}}_j)$, and a fitted set of temperature profiles $\hat{\mathbf{Z}}_j = m(x_j, \mathbf{Y}_j; \hat{\boldsymbol{\theta}}_j)$. This concept is laid out more visually in Figure 4.5.

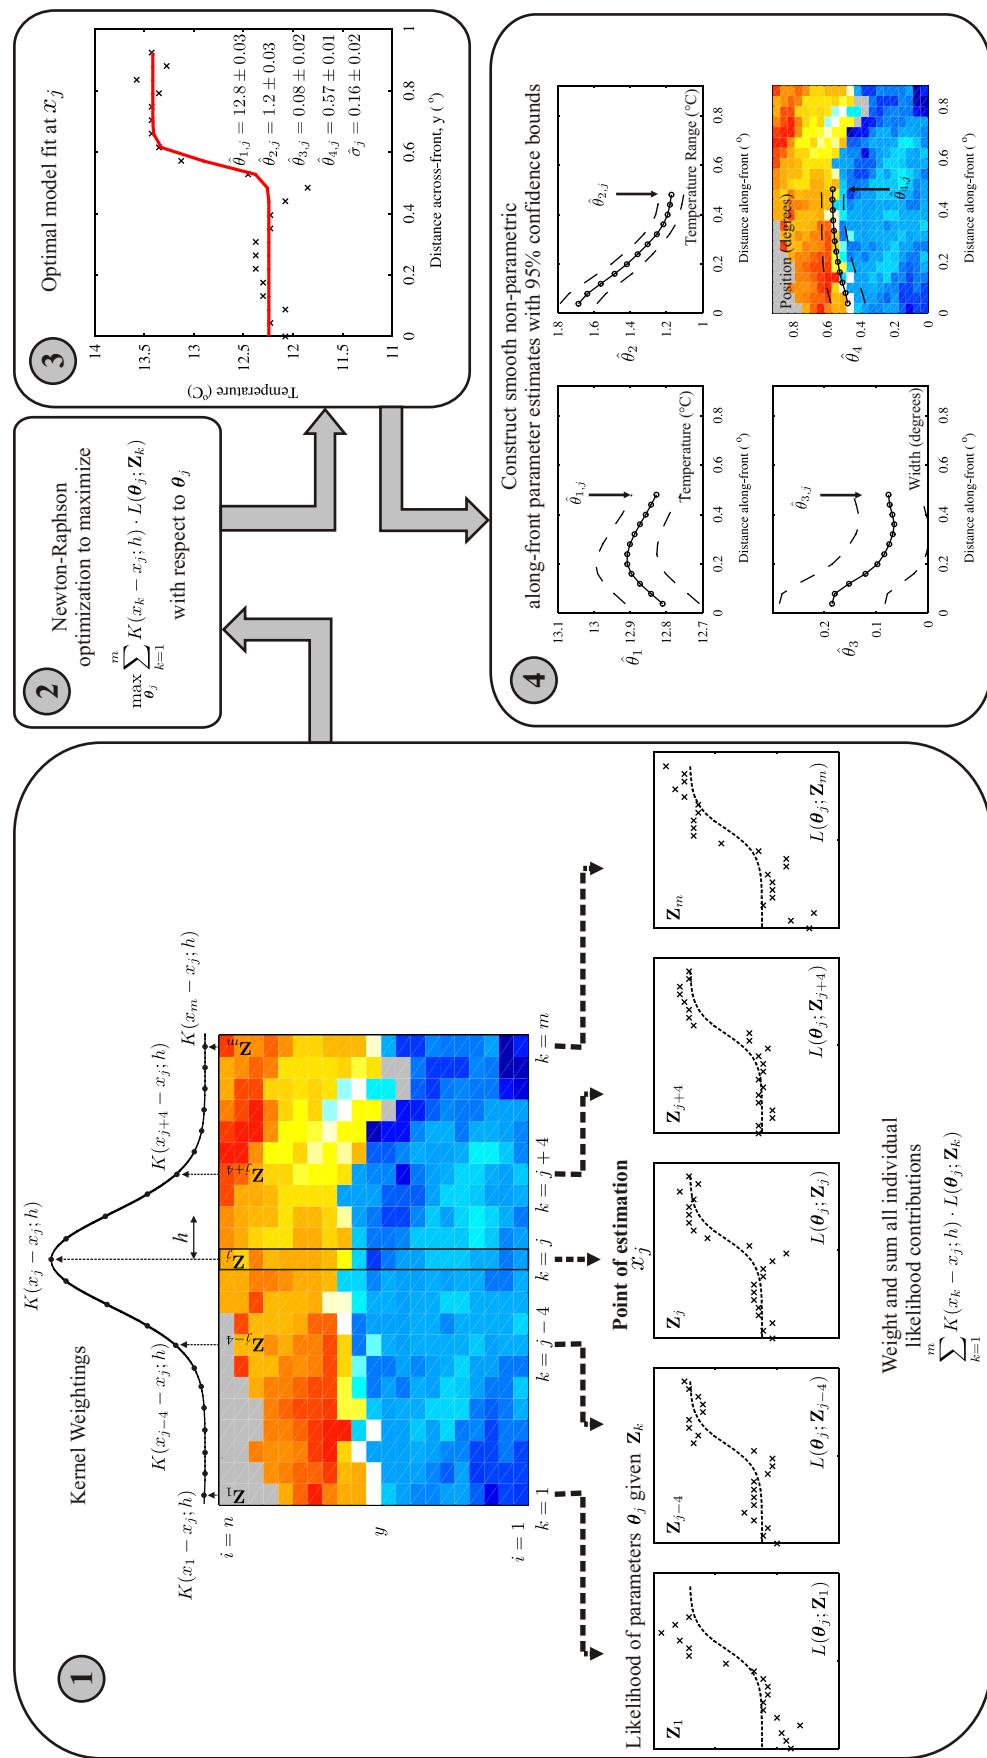
These ideas can be expanded into the local regression context by assuming a particular model for the behavior of $\boldsymbol{\theta}$ near x . Taking again a local linear model, $\boldsymbol{\theta}_j = \boldsymbol{\alpha}_j + \boldsymbol{\beta}_j x_j$, the local likelihood estimate $\hat{\boldsymbol{\theta}}_j$ is then $\hat{\boldsymbol{\theta}}_j = \hat{\boldsymbol{\alpha}}_j + \hat{\boldsymbol{\beta}}_j x_j$, where $\hat{\boldsymbol{\alpha}}_j$ and $\hat{\boldsymbol{\beta}}_j$ maximize the local log-likelihood:

$$L(\boldsymbol{\alpha}_j, \boldsymbol{\beta}_j) = \max_{\boldsymbol{\alpha}_j, \boldsymbol{\beta}_j} \sum_{k=1}^m K(x_k - x_j; h) \cdot L(\boldsymbol{\alpha}_j, \boldsymbol{\beta}_j; \mathbf{Z}_k).$$

This *local likelihood* approach is a significant improvement over the standard maximum likelihood estimates. The local likelihood estimates are constructed from a much larger set of observations and therefore have reduced confidence intervals. In addition, estimates in those regions with a sparsity or complete absence of temperature observations are made possible by drawing on surrounding information.

As already discussed in Section 4.2.3, for any given position x_j , the likelihood calculations generate an estimated covariance matrix of the parameters. From this, a confidence interval for each estimate can be calculated. The diagonal elements of the matrix of second derivatives for the kernel weighted local likelihood are given in Appendix A.

In the final section of this theoretical chapter we discuss how the width of the gaussian kernel ($K(\cdot)$) is chosen. This is the determining factor in the amount of smoothing applied and the level of spatial variability captured by the algorithm.

FIGURE 4.5: Graphic representation of constructing weighted local likelihood estimates at position x_j

4.4 Bandwidth Selection

We are concerned with the estimation of the smooth variation in parameters $\boldsymbol{\theta} = \{\theta_1, \theta_2, \theta_3, \theta_4, \sigma\}$ along a front. The bandwidth h controls the width of the kernel smoother used to create the local likelihood estimates $\hat{\boldsymbol{\theta}}$, and determines the form of the resulting non-parametric function $\hat{\Phi}(x)$. Choosing an appropriate smoothing parameter therefore deserves careful consideration. The choice will always be influenced by the purpose of the model. In some situations it is sensible to choose h subjectively. This may be done by visually assessing the fitted non-parametric trends obtained using different values of h , and choosing the smoothing parameter that generates estimates most in accordance with previous expectations of the trend. A subjective choice may also be based upon *a priori* information about the scale of features expected in the region of interest. A natural choice for our purposes would be the Rossby Radius of Deformation, the length scale at which rotational effects become as important as buoyancy effects in the evolution of flow about some disturbance. The Baroclinic Rossy Radius off the east coast of South Island, New Zealand is approximately 20km (Chelton et al., 1998).

Alternative to subjectively choosing a smoothing parameter are a variety of automatic methods, useful in determining the degree of smoothing best supported by the data themselves. Automatic methods are essential if smoothing is being performed on a large number of data sets or when different results are being compared and reference to a standard method is needed. In this next section we explore the consequences of increasing and decreasing the smoothing parameter, and compare a variety of automated methods designed to select the bandwidth most appropriate to the length scale of features and density of data in each image.

4.4.1 Bias and Variance Considerations

Setting the bandwidth equal to zero ($h = 0$), the local likelihood estimator $\hat{\boldsymbol{\theta}}_j$ (Equation 4.10) reduces to the standard maximum likelihood estimate, and the parameters $\boldsymbol{\theta}_j$ are estimated using only data at point x_j . Taking $h \rightarrow \infty$ on the other hand sets the parameters constant globally with distance. Somewhere inbetween there is an optimal value of h , which may be considered as a measure of model order or complexity.

We are concerned with trying to minimize the mean squared error (MSE) between the estimate $\hat{\theta}_q^{[h]}$, and the true function value θ_q at each point x_j with respect to the smoothing parameter h . $\hat{\boldsymbol{\theta}}_q$ is the q 'th element in the vector of parameter estimates $\hat{\boldsymbol{\theta}}$.

$$\begin{aligned}
\text{MSE}(h) &= \frac{1}{m} \sum_{j=1}^m \mathbb{E}\{\hat{\Phi}_q^{[h]}(x_j) - \Phi_q(x_j)\} \\
&= \frac{1}{m} \sum_{j=1}^m \{\mathbb{E}(\hat{\Phi}_q^{[h]}(x_j)) - \Phi_q(x_j)\}^2 + \frac{1}{m} \sum_{j=1}^m \mathbb{E}\{\hat{\Phi}_q^{[h]}(x_j) - \mathbb{E}(\hat{\Phi}_q^{[h]}(x_j))\}^2 \\
&= \frac{1}{m} \sum_{j=1}^m \text{Bias}^2(\hat{\Phi}_q^{[h]}(x_j)) + \frac{1}{m} \sum_{j=1}^m \text{Var}(\hat{\Phi}_q^{[h]}(x_j)).
\end{aligned}$$

The MSE is therefore comprised of both the variance and squared bias of the estimate $\hat{\theta}_q^{[h]}$. The bias is the amount by which the expectation of the estimate, $\mathbb{E}(\hat{\theta}_q^{[h]})$, differs from the true value θ_q . The variance is the expected squared deviation of $\hat{\theta}_q^{[h]}$ around its mean $\mathbb{E}(\hat{\theta}_q^{[h]})$ (Hastie et al., 2001).

Considering a simple linear smoother, such as the Nadaraya-Watson estimate, approximate expressions for the mean and variance of the estimate $\hat{\Phi}_q(x)$ may be derived. These give insight into the response of the bias and variance to the bandwidth h . Parameterizing the kernel function $K(x_j - x; h)$ with $\frac{1}{h}k(\frac{x}{h})$,

$$\begin{aligned}
\mathbb{E}\{\hat{\Phi}_q(x)\} &\approx \Phi_q(x) + \frac{h^2}{2} \sigma_k^2 \Phi_q''(x) \\
\text{Var}\{\hat{\Phi}_q(x)\} &\approx \frac{\sigma^2}{nh} \frac{\alpha(k)}{f(x)},
\end{aligned}$$

(Bowman and Azzalini, 1997)

where σ_k^2 denotes the variance of the kernel function, and $\alpha(k) = \int k(x)^2 dx$. $f(x)$ is a function expressing the local pattern or density of data points with respect to x . An attractive property of local linear estimators is that the bias does not depend on this pattern (Hastie et al., 2001). Since $\Phi_q''(x)$ measures the curvature of the function $\Phi_q(x)$, there is a tendency for $\hat{\Phi}_q(x)$ to smooth out peaks and troughs. This bias increases with larger bandwidths (h), and with functions showing higher degrees of curvature. The variance is inversely proportional to the local sample size $nhf(x)$ and therefore increases with smaller smoothing parameters and where there is a sparsity of data.

Figure 4.6 demonstrates the balance between bias and variance using the estimated temperatures $\hat{\theta}_1^{[h]}$ along the Southland Front from April 1985 as an example. As h increases weights attached to the underlying temperature observations become increasingly similar and the estimator $\hat{\theta}_1^{[0.5]}$ approaches a fitted least squares constant. Its variance diminishes while its bias increases. Too large a bandwidth will result in interesting oceanic features being blurred. At the other extreme, as h decreases ($\hat{\theta}_1^{[0.025]}$), and approaches

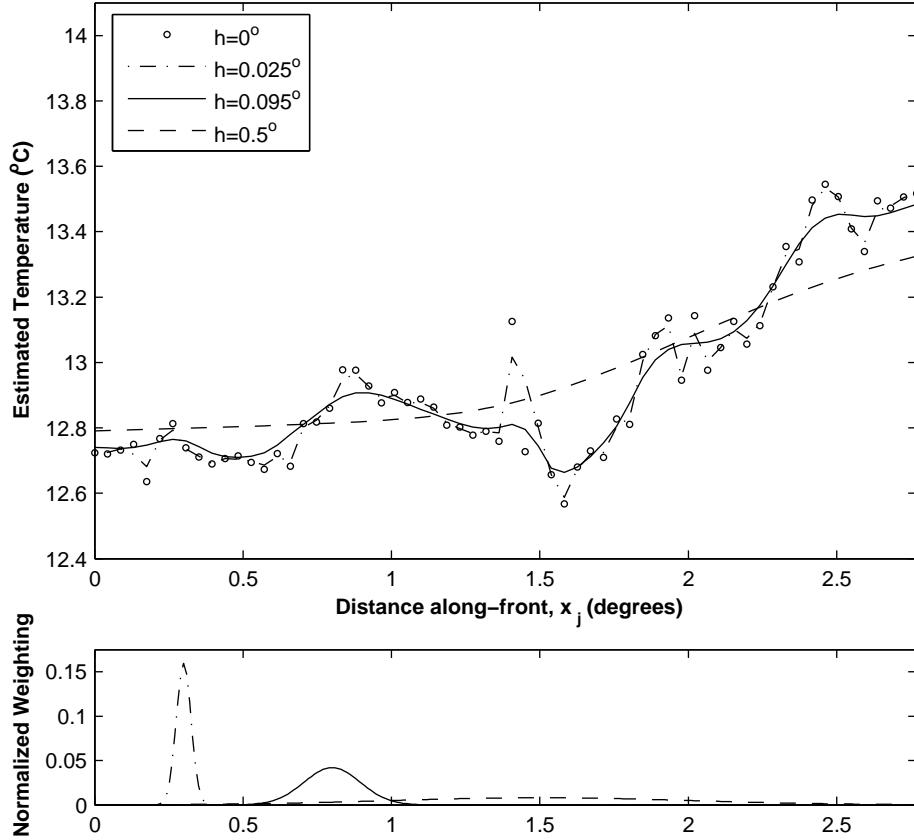


FIGURE 4.6: Bias-variance control of the bandwidth. Local likelihood estimates for the mean temperature of the Southland Front in April 1985 are plotted for bandwidths of 0° , 0.025° , 0.095° and 0.5° . Note the increasing tendency for overestimation at troughs and underestimation at peaks as the bandwidth gets larger. The width and weightings of each kernel function are shown in the bottom panel.

the standard MLE $\hat{\theta}_1^{[0]}$, the estimator tracks the underlying temperature profiles too closely resulting in overfitting, decreasing the bias but increasing the variance. It is this concept of *bias-variance tradeoff* that requires consideration when choosing a smoothing parameter since it will ultimately effect the structure of the estimated non-parametric function. However since we do not know the true function $\Phi_q(x)$ we are unable to calculate the MSE.

There are a number of techniques that may be used to calculate an appropriate smoothing parameter that do not require calculation of the MSE. The techniques of Cross Validation, Bayesian Information Criterion and Likelihood Cross Validation are detailed below and a comparison between them carried out in Section 5.2.2.

4.4.2 Cross Validation

An alternative approach to calculating the MSE is to perform some form of cross validation and then minimize a cross validation prediction error. Let

$$\hat{\boldsymbol{\theta}}_{jCV}^{[h]} = \max_{\boldsymbol{\theta}_j} \sum_{k=1 \setminus j}^m K(x_k - x_j; h) \cdot L(\boldsymbol{\theta}_j; \mathbf{Z}_k) \quad (4.11)$$

denote the local likelihood estimate for parameters $\boldsymbol{\theta}_j = \{\theta_{1,j}, \theta_{2,j}, \theta_{3,j}, \theta_{4,j}, \sigma_j\}$ computed by leaving out the j^{th} set of temperature observations \mathbf{Z}_j .

For each parameter estimate $\hat{\boldsymbol{\theta}}_{q,j}^{[h]}$, where $q = \{1, 2, 3, 4, 5\}$, the cross validation sum of squares is then

$$CV(h) = \frac{1}{m} \sum_{j=1}^m \{\hat{\boldsymbol{\theta}}_{q,j}^{[h]} - \hat{\boldsymbol{\theta}}_{q,jCV}^{[h]}\}^2, \quad (4.12)$$

and is minimized to find an optimal smoothing parameter h_{op} .

h_{op} will be different for each of the five parameters because of differences in their natural structure and variability.

Cross validation is based on minimizing prediction error in parameter space. Likelihood cross validation is an alternative approach whereby the prediction error is instead minimized in likelihood space.

4.4.3 Likelihood Cross Validation

The method of likelihood cross validation uses likelihood to judge the adequacy of fit for a statistical model (Silverman, 1986).

$$LCV(h) = \sum_{j=1}^m L(\hat{\boldsymbol{\theta}}_{jCV}^{[h]}; \mathbf{Z}_j), \quad (4.13)$$

where $L(\hat{\boldsymbol{\theta}}_{jCV}^{[h]}; \mathbf{Z}_j)$ is the likelihood associated with the leave-one-out cross validation estimate $\hat{\boldsymbol{\theta}}_{jCV}^{[h]}$. The bandwidth which then maximizes the function $LCV(h)$ is chosen as the smoothing parameter. The distance between $\hat{\boldsymbol{\theta}}_{jCV}^{[h]}$ and $\hat{\boldsymbol{\theta}}_j^{[h]}$ is being measured in likelihood space rather than in parameter space.

An indication of the optimal bandwidth may also be given by the Bayesian Information Criterion.

4.4.4 Bayesian Information Criterion

Like the Akaike Information Criterion (Hastie et al., 2001) the Bayesian Information Criterion (BIC) is a technique used for model selection applicable in settings where

fitting is carried out by maximization of a log-likelihood. The BIC is a measure of statistical model fit (Hastie et al., 2001). It attempts to find the minimum model fit that correctly explains the data by examining the complexity of the model i.e. the size of the bandwidth, together with its goodness-of-fit to the observed data. The BIC is calculated as follows:

$$\text{BIC}(h) = -2 \ln(\mathbb{L}) + df \ln(N) \quad (4.14)$$

- N denotes the total number of observations
- df , the number of degrees of freedom is given by $df = \sum_{j=1}^m w(x_j, x_j)$. The weightings $w(x_j, x_j)$ correspond to the diagonal elements of the smoother matrix \mathbf{S}_h , an $m \times m$ matrix composed of weightings from a kernel function with bandwidth h . Each weighting $w(x_j, x_j)$ is a measure of the potential for the observations at x_j to exert influence on the estimates $\hat{\theta}_j$ made at x_j . The sum of these weightings is therefore a global measure of the amount of smoothing provided. This is the measure of model complexity.
- $\mathbb{L} = \sum_{j=1}^m L(\hat{\theta}_j^{[h]})$ is the sum of likelihoods associated with each set of estimates $\hat{\theta}_j^{[h]}$.

The bandwidth which minimizes the function $\text{BIC}(h)$ is the optimal smoothing parameter.

4.5 Summary

This chapter has outlined the statistical theory behind a new ocean front detection scheme. The key points and equations are summarized as follows:

- Measurements of SST across an ocean front are modelled by the function:

$$m(\mathbf{Y}; \boldsymbol{\theta}) + \epsilon = \theta_1 + \theta_2 \tanh \left[\frac{\mathbf{Y} + \theta_4}{\theta_3} \right] + \epsilon,$$

where θ_1 , $2\theta_2$ and $2\theta_3$ represent the front's mean temperature, temperature range and width respectively (Figure 4.2). The position of the front is determined by θ_4 and ϵ is normally distributed noise with zero mean and standard deviation σ . \mathbf{Y} is a vector of known distances at which SST observations are available.

- The log-likelihood associated with the model $m(\cdot)$ from a set of SST observations \mathbf{Z}_j is given by:

$$L(\boldsymbol{\theta}_j; \mathbf{Z}_j) = -n \ln \sigma_j - \frac{1}{2\sigma_j^2} \sum_{i=1}^n \left[z_j^{(i)} - \left(\theta_{1,j} + \theta_{2,j} \tanh \left[\frac{y^{(i)} + \theta_{4,j}}{\theta_{3,j}} \right] \right) \right]^2.$$

- Model parameter estimates at position x_j are calculated by maximizing the weighted sum of likelihood contributions from all available sets of SST observations in the image. This local likelihood estimator for $\boldsymbol{\theta}_j$ is of the form:

$$\hat{\boldsymbol{\theta}}_j = \max_{\boldsymbol{\theta}_j} \sum_{k=1}^m K(x_k - x_j; h) \cdot L(\boldsymbol{\theta}_j; \mathbf{Z}_k).$$

- Weightings are provided by the normal density function such that:

$$K(x_j - x; h) = \frac{1}{\sqrt{2\pi h^2}} e^{\left[\frac{-(x_j - x)^2}{2h^2} \right]}.$$

The bandwidth h determines the degree of spatial smoothing.

- Cross Validation, Bayesian Information Criterion or Likelihood Cross Validation may be used to automatically select the optimal bandwidth of the smoothing function $K(\cdot)$ (i.e. the standard deviation).
- Confidence bounds about each parameter estimate may be constructed via the estimated variance-covariance matrix.

To obtain the most realistic estimate of frontal parameters at each point k along the front data \mathbf{Z}_k should be extracted at perfect right angles across the front. The algorithm presented here fits the model function based on the assumption that the front

is orientated east-west across the image and does not take into account changes in the front's orientation that would lead to observations \mathbf{Z}_k not always being normal to the front's direction. As will be seen in Section 5.3.3 of Chapter 5 deviations from normal gradually lead to overestimation of the width. An orientation parameter however adds a considerable amount of additional complexity to the likelihood solution presented here and has therefore not been taken into account. Firstly, the orientation of the front is, like the mean temperature and position etc., an unknown variable, but one that is not easily incorporated into the optimization. Its value would need to be allowed to vary spatially between points k during the optimization such that it resulted in the smallest possible estimate of the width at each position. As such, the vectors of observations \mathbf{Z}_k would change with each iteration introducing programming complexities. Allowing \mathbf{Z}_k to continually vary would result in a certain number of SST observations being used more than once. The affect that repeated use of the same data would have on the underlying statistics of the algorithm is somewhat unknown and difficult to quantify. The assumption that vectors of data \mathbf{Z}_k are independent would be overturned since for example, \mathbf{Z}_k and \mathbf{Z}_{k+1} might both include some of the same measurements. Solutions to the omission of a rotational parameter are discussed in Chapter 8.

CHAPTER 5

Practical Aspects of Front Detection

Chapter 4 details the theory of local likelihood that is used as the basis of the front detection scheme proposed in this thesis. In this chapter a full account of how this model is used to build up a detailed picture of the Southland Front is given. Firstly, the acquisition and preparation of the AVHRR SST data is described. This is followed by a description of the Newton-Raphson optimization routine used to obtain the local likelihood estimates, together with an explanation of the scaling parameters specific to the Southland Front that are used. The choice of smoothing parameter is revisited and justification for selecting likelihood cross validation is given. A weighting criteria, additional to the smoothing kernel is also introduced as a means of downweighting the influence of poor quality or sparse regions of data. A series of tests on simulated data sets demonstrates the algorithms abilities and assesses its limitations.

5.1 Data Acquisition and Preparation

There are a range of thermal infrared (IR) and passive microwave SST products available, each with their associated advantages and weaknesses. Thermal IR sensors (e.g. AVHRR, ATSR, GOES and MODIS etc.) provide good resolution and accuracy, but require atmospheric correction and are obscured by clouds. In areas where there is persistent cloud cover such as in the Southern Ocean this can prevent the structure and position of ocean fronts from being observed for considerable periods of time. Passive microwave sensors (e.g. SMMR, TRMM and AMSR etc.) on the other hand, are mostly transparent to cloud and are relatively insensitive to atmospheric effects making them

'all-weather' devices. However, thermal emission in the microwave channel is weaker than at the shorter IR wavelengths and the signal received by the sensor is therefore weaker. To maintain signal strength and overcome noise levels the field-of-viewing is wider than for IR radiometers resulting in a loss of spatial resolution and accuracy which may limit the performance of the front detection algorithm. Microwave sensors are also sensitive to surface roughness and precipitation.

For this study we take advantage of the 4km twice-daily global Advanced Very High Resolution Radiometer (AVHRR) IR Pathfinder SST data set (Level 3 processed ¹), available from the mid 1980's. Although higher (1km) resolution images and more accurate SST measurements are available from the Along Track Scanning Radiometer (ATSR) (Robinson, 2004), data is only available from 1991. Spatial resolution was sacrificed in order to gain a longer time series more suitable for long term variability, teleconnection and climatological studies. Additionally, owing to the scanning geometry of the ATSR sensor global coverage cannot be achieved in less than four days. In comparison, with AVHRR the whole earth can be viewed at least twice a day.

Monthly climatologies of 4km Pathfinder V5 AVHRR SST data were downloaded through NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC) POET data server (<http://poet.jpl.nasa.gov/>) in netcdf format. Observations are globally gridded into equiangle $0.0439453 \times 0.0439453^\circ$ pixels. Data were obtained for the period between January 1985 though to December 2005 (inclusive), providing a twenty-one year time series of 252 images from which both the seasonal and interannual variability of the Southland Front are studied. Only night-time overpasses were used so as to avoid any surface skin created by diurnal warming masking the true surface frontal structure. In addition, four daily images (28nd, 29th, 31st March 1990 and 2nd April 1990), and a weekly (8 day) composite (28th March 1990 - 4th Apr 1990), were obtained in order to assess the level of detail lost through blurring in a monthly composite (see Section 6.4.2 of Chapter 6). Monthly data are adequate for climate studies but investigations into meso-scale variability would require weekly and daily data in order to capture the more rapid changes in frontal position and strength. The choice of data set is therefore partly dependent upon the intended use of the algorithm results.

The Pathfinder V5 product is a reanalysis of the V4 dataset that is based on the NOAA/-National Environment Satellite Data and Information Service (NESDIS) nonlinear SST algorithm (Kilpatrick et al., 2001; Walton et al., 1998). The reprocessing has improved spatial resolution, ice mask identification and uses a more accurate, consistent land mask. The improved spatial resolution greatly enhances the utility of the Pathfinder V5 dataset in coastal zones and areas of strong SST gradients, both improvements beneficial to this study. Note that the AVHRR sensor itself measures only a skin brightness temperature at a depth of $\sim 10\text{-}20\mu\text{m}$. The SST algorithm coefficients are estimated using a match-up data base of *in-situ* buoys and produces an SST product that is more

¹Geophysical parameters derived and data resampled spatially and/or temporally

closely tuned to a bulk temperature (Kilpatrick et al., 2001) representative of the mixed layer temperature (\sim 1-20m). This is a more useful measurement to oceanographers and climatologists.

The absolute accuracy of the AVHRR-derived SST's can not be exactly determined. The accuracy is regionally and seasonally dependent and is influenced by the water vapor content in the atmosphere. Globally, the Pathfinder SST (PFSST) is found to be within $0.02 \pm 0.53^\circ\text{C}$ of the match-up data base of buoys (Kilpatrick et al., 2001). However, this aggregate global statistic varies across latitudinal bands as the state of the atmosphere (water vapor, temperature profile, absorbing aerosols) and ocean stratification varies. In addition, the *in-situ* buoys used in this comparison are themselves not calibrated to an accuracy of 0.02°C , introducing another factor of uncertainty. Examination of the PFSST and match-up buoy measurements by latitudinal bands suggests that the median value of the residual is probably closer to $0.1 \pm 0.5^\circ\text{C}$. These values will be weighted by the relative distribution of match-ups which is continually changing through time.

Different intended uses for the Pathfinder data may involve a trade-off between data coverage and SST pixel quality. Each SST value is therefore examined to assess the likelihood that it is of suspect quality.

5.1.1 AVHRR-SST Pixel Quality

Nine individual pixel-by-pixel tests are performed and combined to define eight overall quality flags for each pixel. These tests involve:

- the brightness temperature (BT) of AVHRR channels,
- the relative difference in BT within 3×3 pixel boxes,
- the satellite zenith angle (angle at earth's surface between a line pointing straight up and the direction to the satellite),
- the difference between the PFSST and a reference SST field,
- the Earth-Sun-satellite configuration,
- sun glint, and
- cloud tests specific to each satellite and year.

Details of these tests and how they are combined to derive an overall pixel quality level may be found in Figure A.1, Table A.1 and in Kilpatrick et al. (2001).

To give an appropriate balance between geographic coverage, i.e. minimizing the percentage of missing data in an image, and SST pixel quality the monthly composites were

downloaded with a quality flag of 4 (Figure A.1). Under this flag a uniformity test of brightness temperatures is failed (Table A.1). The test is designed to detect contamination by small clouds, but in doing so may erroneously identify sharp, small scale frontal features instead (Kilpatrick et al., 2001). Under cloud free conditions the test assumes that the range of brightness temperatures within each 3×3 pixel window should be less than 0.7°C (i.e. relatively uniform temperatures). Sharp frontal features on scales less than 12km may exceed this threshold and be flagged as small clouds. Given that we are interested in frontal temperature gradients over relatively small spatial scales, SST pixel quality was sacrificed so as to not lose possible smaller scale features in the image.

Figure 5.1 shows, on a daily level, the increase in suspect data as the SST pixel quality flag improves. It also highlights the bias toward flagging pixels in high gradient areas. A quality flag of 1 starts to result in data loss along the Southland Front in this relatively cloud free example from 28th March 1990. The pixels identified here have failed not only the uniformity test described above but also a test designed to identify pixels where the absolute difference between the Pathfinder SST and Reynolds Reference SST field is greater than 2°C . This process is once again prone to erroneous flagging in frontal regions. As the quality flag improves eventually all SST measurements along the coastal zone and at the front are lost.

The daily and weekly images were downloaded with the lowest possible quality flag of zero. This makes the maximum sacrifice on SST pixel quality but ensures that no frontal features are lost. As seen in Figure 5.1, on a daily basis erroneous flagging of pixels over a front results in significant data loss over important areas. Over a month, multiple overpasses and changes in the weather allow many of these gaps to be filled in and a higher quality flag of 4 was set. Given that the purpose of the daily images was to quantitatively assess the level of detail lost through temporal blurring it was important that all SST structures were visible. In addition, for so few images it was possible to visually check for and take into account any large areas of cloud contamination. Furthermore, the detection of spurious fronts due to clouds is minimized through bounded optimization and data quality weightings (see Sections 5.2.1.2 and 5.2.3).

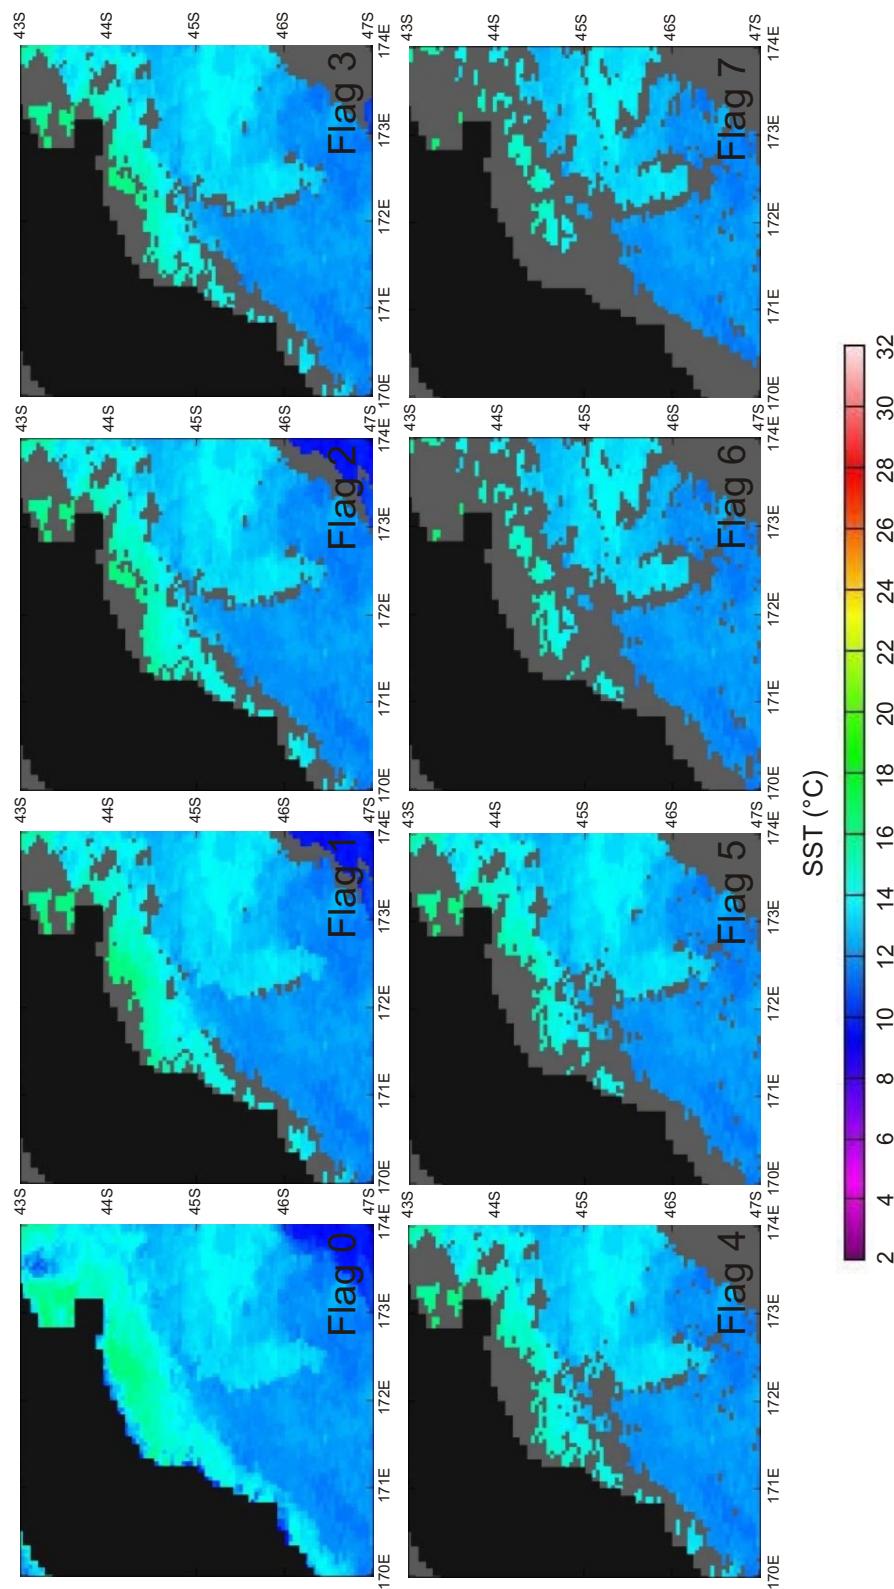


FIGURE 5.1: 1km resolution SST image of the east coast of South Island on 28th March 1990 with quality flags 0 to 7 applied.

5.1.2 Rotation and Extraction of SST Data

Local likelihood parameter estimates described in Chapter 4 are made based on the assumption that each set of SST observations (\mathbf{Z}) are taken at approximately right angles across the front i.e. the front is oriented east-west across the image being analyzed. The Southland Front approximately follows the 500m isobath (Shaw and Vennell, 2001), that runs south-west to north-east along the east coast of South Island, New Zealand. It is therefore necessary to extract from the original AVHRR image observations within a rotated window before parameter estimates are made. In this way more accurate estimates of frontal characteristics are obtained (discussed in Section 5.3.3). The original equiangle latitude-longitude grid ($0.0439453 \times 0.0439453^\circ$) is converted to a complex coordinate system centred about the point $45^\circ\text{S}, 172^\circ\text{E}$ (Figure 5.2). From this new origin a 1° by 2.8° window is defined and rotated by 45° aligning the window approximately along the Southland Front. The rotated complex coordinates are converted back into the original latitude-longitude system and pixels of data within the AVHRR SST image corresponding most closely with those locations are extracted.

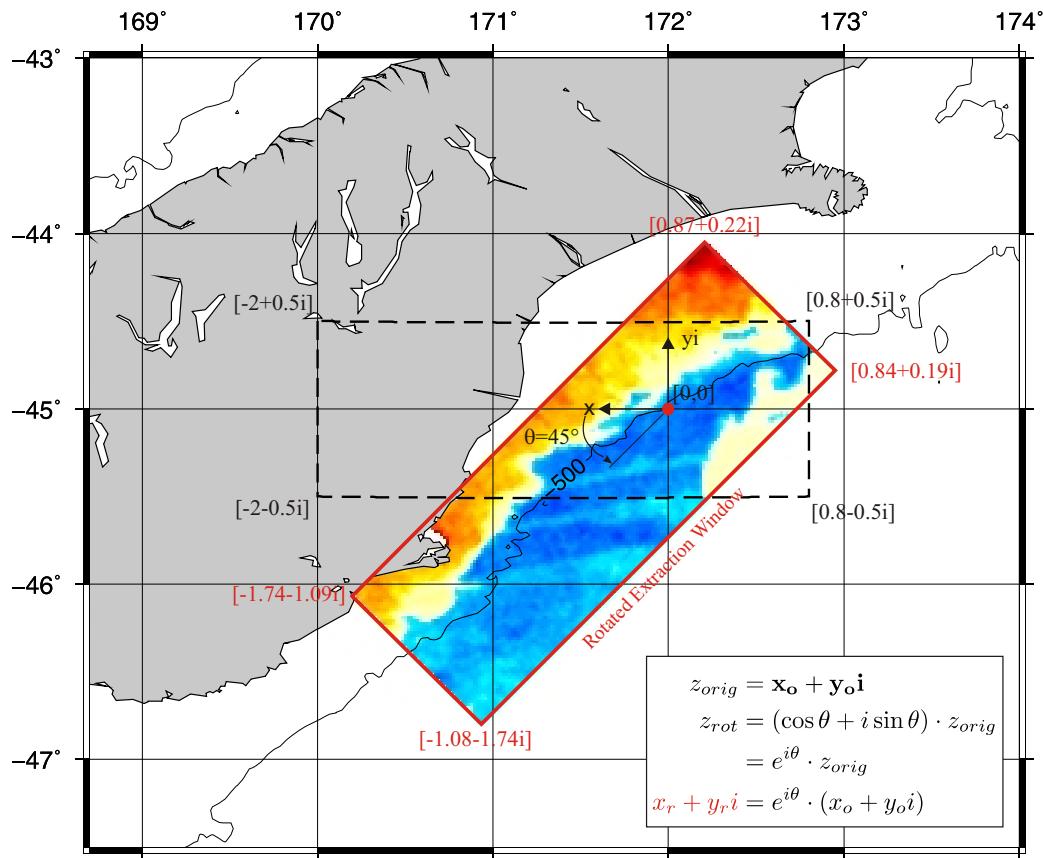


FIGURE 5.2: Rotation and extraction of a data window in the complex plane. The equiangle latitude-longitude coordinate system is converted into the complex plane with the origin at $45^\circ\text{S}, 172^\circ\text{E}$ (red dot) and then rotated by 45° . The 500m isobath is marked.

5.2 Practical Aspects of Front Detection

5.2.1 Optimization

Chapter 4 describes the statistical theory of weighted local likelihood that forms the basis of the proposed front detection algorithm. We derive a local likelihood expression (Equation 4.10) that must be maximized with respect to the unknown parameters θ , at each position x_j along the Southland Front:

$$\hat{\theta}_j = \max_{\theta_j} \sum_{k=1}^m K(x_k - x_j; h) \cdot L(\theta_j; \mathbf{Z}_k).$$

Solving this equation is an optimization problem involving a nonlinear function and multiple unknowns. Since the problem of *maximizing the likelihood function L* , is equivalent to *minimizing $-L$* , the following text will henceforth refer to *minimization* of an objective function F , where $F = -L$, the negative local log-likelihood.

All optimization techniques are based on an iterative procedure requiring a set of initial parameter estimates, a first best guess at the solution. On the basis of either estimates of the function gradient and curvature (gradient based methods), or on multiple evaluations of the objective function F (direct search methods), a sequence of estimates intended to converge to a local minimum is generated (Box et al., 1969).

We use the gradient based method of **Newton-Raphson** in order to make estimates $\hat{\theta}_j$. Newton-Raphson benefits from rapid quadratic convergence when near a minimum, but can be sensitive to the initialization parameters (Gill et al., 1995). In order to stabilize the optimization and create a good first approximation to the solution of Equation 4.10, the bandwidth is initially set to zero ($h = 0$). In this way Equation 4.10 simplifies to:

$$\hat{\theta}_j = \max_{\theta_j} L(\theta_j; \mathbf{Z}_j), \quad (5.1)$$

the *standard maximum likelihood estimate*. Estimates from this simplified problem (i.e. no smoothing function) are then used as a starting point for the full local likelihood estimate when $h > 0$. The following sections describe the Newton-Raphson algorithm, the transformation of variables, parameter bounds and scaling performed to facilitate rapid and successful convergence.

5.2.1.1 Newton-Raphson Optimization

We use a gradient based Quasi-Newton scheme subject to fixed upper and lower bounds taken from the NAG (Numerical Algorithms Group, 2006) Fortran Library, Mark 21 to make standard (Equation 5.1) and local likelihood parameter estimates (Equation 4.10).

Newton-Raphson approximates the objective function (F) at each iteration with a quadratic, given by the second order Taylor expansion of F , and then moves to the minimum of this quadratic. At each iteration (t) the routine calculates the projected gradient vector $\mathbf{g}^{(t)}$, whose elements are finite difference approximations to the derivatives of the objective function (F), and a positive definite approximation to the matrix of second derivatives, $\mathbf{H}^{(t)}$, both with respect to the current free parameter estimates, $\boldsymbol{\theta}^{(t)}$. The equations

$$\mathbf{H}^{(t)} \mathbf{d}^{(t)} = -\mathbf{g}^{(t)}, \quad (5.2)$$

are then solved to give a search direction $\mathbf{d}^{(t)}$. Then α , the distance stepped along the direction $\mathbf{d}^{(t)}$, is found such that $F(\boldsymbol{\theta}^{(t)} + \alpha \mathbf{d}^{(t)})$ is a minimum with respect to α , subject to the fixed bounds of Section 5.2.1.2. For the next iteration $t + 1$, $\boldsymbol{\theta}^{(t)}$ is replaced by $\boldsymbol{\theta}^{(t)} + \alpha \mathbf{d}^{(t)}$ and the procedure is repeated until convergence criteria (Section 5.2.1.3) are satisfied. If a variable reaches a bound during the search along \mathbf{d} it is fixed and the number of free variables is reduced for the next iteration (Numerical Algorithms Group, 2006). Knowing which variables have reached these bounds forms an important part of the subsequent quality control of estimates (see Section 5.2.3). A detailed review of Newton-Raphson is given by Gill et al. (1995).

For a general nonlinear function, Newtons method benefits from rapid quadratic convergence if the initial parameter inputs are close to the estimated optimal values. We calculate an initial estimate and a set of starting values by first solving the local likelihood case where $h = 0$, i.e. the standard maximum likelihood equation (Equation 5.1). The mean temperature (11.01°C), width (8.36km), and temperature range (1.76°C) of the Southland Front estimated by Shaw and Vennell (2001) are used as a starting point for these iterations. The front is assumed to be orientated south-west, north-east and therefore initially positioned centrally across the rotated extraction window. A value of 0.15 is used as a starting point for σ , the standard deviation about the mean model, chosen based on user experience. These standard maximum likelihood estimates are then used as a starting point for the full local likelihood solution with an optimal bandwidth determined by likelihood cross validation (see Section 5.2.2). We also use these results to apply a quality flag and an additional weighting to the data (see Section 5.2.3).

TABLE 5.1: Upper (b_q) and lower (a_q) scaling bounds for linear transformation of variables using a $1^\circ \times 2.8^\circ$ window of data with a 4km pixel resolution.

Parameter	Geophysical Lower Bound	Geophysical Upper Bound	Additional transformation	a_q	b_q
Mean Temp θ_1	5°C	23°C	none	5°C	23°C
Temp Range $2\theta_2$	0.1°C	6°C	none	0.05°C	3°C
Width $2\theta_3$	0.044° (~4km)	0.6° (~60km)	logarithmic	log(0.022)	log(0.3)
Position θ_4	-1	0	none	-1	0
s.d σ	1e-06	6	logarithmic	log(1e - 06)	log(3)

The Newton-Raphson technique is most efficient when the unknown variables being estimated are of a similar magnitude and of order unity near the minimum. The following section describes how this is achieved.

5.2.1.2 Transformation of Variables and Parameter Bounds

To facilitate rapid and successful convergence we apply a linear transformation to all variables by imposing upper and lower bound constraints. In this way all variables are of order unity when in the region of likely solutions. Transformed variables $\boldsymbol{\theta}_{new}$ are obtained as follows:

$$\boldsymbol{\theta}_{new} = \mathbf{D}^{-1}(\boldsymbol{\theta}_{old} - \mathbf{C}), \quad (5.3)$$

where \mathbf{D} is a diagonal matrix with q^{th} diagonal element $\frac{1}{2}(b_q - a_q)$ and \mathbf{C} is a vector with q^{th} element $\frac{1}{2}(a_q + b_q)$. a_q and b_q are the lower and upper bounds respectively imposed on the $1 \times q$ matrix of variables $\boldsymbol{\theta}_{old}$. As before the subscript q refers to the specific elements $\theta_1, \theta_2, \theta_3, \theta_4$ and σ of the vector $\boldsymbol{\theta}$. This transformation guarantees that $-1 \leq \theta_{new} \leq +1$ for all q , within the interval $[a_q, b_q]$ (Gill et al., 1995). The limits are determined by the resolution and accuracy of the data set, the size of the image and reasonable geophysical values expected for SST's and frontal scales in the region. In addition, unknown parameters θ_3 and σ are log transformed to avoid any arithmetic errors in the evaluation of the objective function.

Table 5.1 summarizes the bounds set for the Newton-Raphson routine. Based on work by Heath (1985a); Chiswell (1994); Uddstrom and Oien (1999) and Shaw and Vennell (2001), the upper and lower bounds set for the mean temperature values are 23°C and 5°C respectively. Sea surface temperatures around New Zealand are unlikely to exceed 23°C during either the summer or winter. As in Shaw and Vennell (2000a), a lower bound of 5°C is set so as to exclude high level cloud. Taking into account the range of the seasonal SST cycle around New Zealand (Heath, 1985a; Chiswell, 1994; Uddstrom

and Oien, 1999), an upper limit of 6°C was set for the temperature difference across the front. Any greater difference seems unreasonable and likely to include cloud. This bound is particularly important during the processing of daily images where data were download with a lower quality flag (Section 5.1.1), and the chances of detecting spurious fronts related to cloud contamination is higher. The lower bound for the temperature range was set at 0.1°C, the approximate relative accuracy of the Pathfinder Data Set (K. Casey personal comment). If the relative accuracy of the temperature between two bodies of water is no more than 0.1°C in the original data, then a resolution greater than this can not be expected from the algorithm. The maximum and minimum bounds for the frontal width are determined by the size ($1^\circ \times 2.8^\circ$) and resolution (4km) of the images. Note that only one front is allowed to exist. Each image is approximately 100km wide (across-front scale), therefore an upper width range was set at 60km. This allows for roughly 20km of data from each plateau region on either side of the front. Having a sufficient amount of data to identify these regions was found to be important by Shaw and Vennell (2000a). The lower limit is set by the 4km data resolution. Features smaller than this can not be resolved. The frontal position is bounded by the limits of the extracted image.

Optimization is most efficient when the objective function, F , is of order no larger than unity near the minimum (Gill et al., 1995). The magnitude of the objective function is therefore estimated before each minimization and used as a scaling constant during subsequent iterations.

5.2.1.3 Convergence Criteria

There are a set of convergence criteria that the optimization routine must meet in order to report successful convergence, and estimates that are good approximations to the true value. With each set of estimates $\hat{\theta}$, a quality flag between zero and ten (IFAIL value), is reported indicating the relative success or failure of the optimization (Numerical Algorithms Group, 2006). The interpretation of IFAIL values are summarized in Table 5.2. An IFAIL value of zero indicates that convergence criteria designed to test whether the sequence of parameter estimates and function values are converging, and whether the gradient vector is zero at the solution, have been met. More formally, the solution $\hat{\theta}_{sol}$ is most likely to be an estimate of the minimum θ_{true} , to the accuracy specified in θ_{tol} , when the following criteria B1, B2 and B3, or B4 hold (Numerical Algorithms Group, 2006):

$$\begin{aligned}
 B1 &\equiv \alpha^{(t)} \times \|\mathbf{d}^{(t)}\| < (\theta_{tol} + \sqrt{\epsilon}) \times (1.0 + \|\theta^{(t)}\|) \\
 B2 &\equiv |F^{(t)} - F^{(t+1)}| < (\theta_{tol}^2 + \epsilon) \times (1.0 + |F^{(t)}|) \\
 B3 &\equiv \|\mathbf{g}^{(t)}\| < (\epsilon^{\frac{1}{3}} + \theta_{tol}) \times (1.0 + |F^{(t)}|) \\
 B4 &\equiv \|\mathbf{g}^{(t)}\| < 0.01 \times \sqrt{\epsilon},
 \end{aligned} \tag{5.4}$$

TABLE 5.2: Interpretation of IFAIL values (Numerical Algorithms Group, 2006). Given that there were no instances of IFAIL equaling 1, 4, 9 or 10, we may have confidence that the program is correctly setup, the problem is well scaled and that there are no mistakes in the programming of the user supplied objective function F .

IFAIL	% Standard MLE	Error
0	96.76	Succesfull convergence (see main text)
1	0	Problem not correctly set up. Invalid or incorrect parameters input
2	0.94	Algorithm not converging after $400 \times q$ function evaluations or F has no minimum
3	0.29	Conditions for a minimum not all met and a lower point not found
4	0	Computational Overflow
5	1.91	
6	0.02	Doubt as to whether a minimum has been found. Degree of confidence decreases as IFAIL increases
7	0.04	
8	0.02	
9	0	Mistake in the user supplied function to evaluate F , the problem has no finite solution or the problem needs rescaling
10	0	Error calculating the set of forward difference intervals

where $\alpha^{(t)}$ is the distance stepped along the search direction $\mathbf{d}^{(t)}$ at the t^{th} iteration. $F^{(t)}$ and $\boldsymbol{\theta}^{(t)}$ are the objective function value and a vector of parameter estimates at the t^{th} iteration respectively. $\mathbf{g}^{(t)}$ is a finite difference approximation to the projected gradient vector at the t^{th} iteration. $\theta_{tol} = 100\sqrt{\epsilon}$ is the desired accuracy of the estimated parameter, where ϵ is the machine precision ($\epsilon = 1.11e-20$). Conditions B1 and B2 are designed to test whether the sequences of parameter estimates and function values are converging. B3 is based on the condition that the gradient vector should be zero at the solution. If the initial position is so close to the solution that no further progress can be made, then the algorithm will also terminate successfully if B4 is satisfied (Gill et al., 1995).

IFAIL values between 1 and 10 result from errors or warnings detected by the routine. Table 5.2 details these warnings together with the percentage of standard maximum likelihood estimates ($h=0$) resulting in each flag.

5.2.2 Bandwidth Selection

Section 4.4 highlights the importance of selecting an appropriate smoothing parameter so as to achieve an acceptable balance between the bias and variance of parameter estimates. A number of methods available to automatically identify the optimal bandwidth h_{op} are presented; cross validation (CV), likelihood cross validation (LCV) and the Bayesian Information Criterion (BIC). Here we compare the performance of these techniques on an image of the Southland Front and justify why likelihood cross validation was chosen as a means of bandwidth selection.

Figure 5.3 shows the value of the LCV, the BIC, and the leave-one-out cross validation sum of squares (CVSS) for bandwidths between 0.01° ($\approx 1\text{km}$), and 1° ($\approx 110\text{km}$) for the Southland Front in April 1985. The patterns seen here are representative of the majority of the Southland Front data set.

There is a well defined maximum at 0.095° in the LCV. This corresponds to the optimal range of values suggested by the BIC (0.05 - 0.15°). The BIC curve is dominated by the relationship between the bandwidth and the degrees of freedom of the smoother; as the bandwidth increases so the degrees of freedom are reduced. There is a rapid decrease in the BIC followed by a slow leveling out of values, and as such no clear cut optimal value stands out.

The CVSS is based in parameter space and demonstrates that each frontal characteristic has its own optimal bandwidth. There is a distinct change in gradient of the CVSS curve for θ_1, θ_4 and σ at $h = 0.05^\circ$ (bottom panel inset of Figure 5.3). This suggests that a smoothing parameter smaller than 0.095° (as identified by the LCV) may be more appropriate for these variables. Figure 5.4 shows that for the temperature (θ_1), and position (θ_4), a bandwidth of 0.05° (solid black line), produces an acceptably smooth curve that is comparable to using the larger LCV optimal bandwidth (solid red line). Mesoscale variability of the Southland Front's position is well defined. The model is robust to the noise of the data set, while not oversmoothing and missing important meanders and curvature, as is the case for $h = 0.3$.

The change in gradient of the CVSS for θ_2 and θ_3 is more gradual, occurring over a range of approximately 0.05 - 0.1° . Figure 5.4 shows that a smoothing parameter toward the upper end of this range is most suitable for the temperature difference (θ_2) and width (θ_3). When $h = 0.05^\circ$ too much of the higher frequency variability introduced by the noisy data set is modelled. These smaller scale fluctuations disappear when h increases to 0.095° . Similar to the BIC, there is no well defined CVSS minimum making automatic selection of an optimal bandwidth using the CVSS difficult.

It is no surprise that different frontal parameters require different amounts of smoothing. The optimal bandwidth h_{op} will vary depending on both the natural along-front stability and variability of each parameter, and the algorithms ability to make accurate estimates

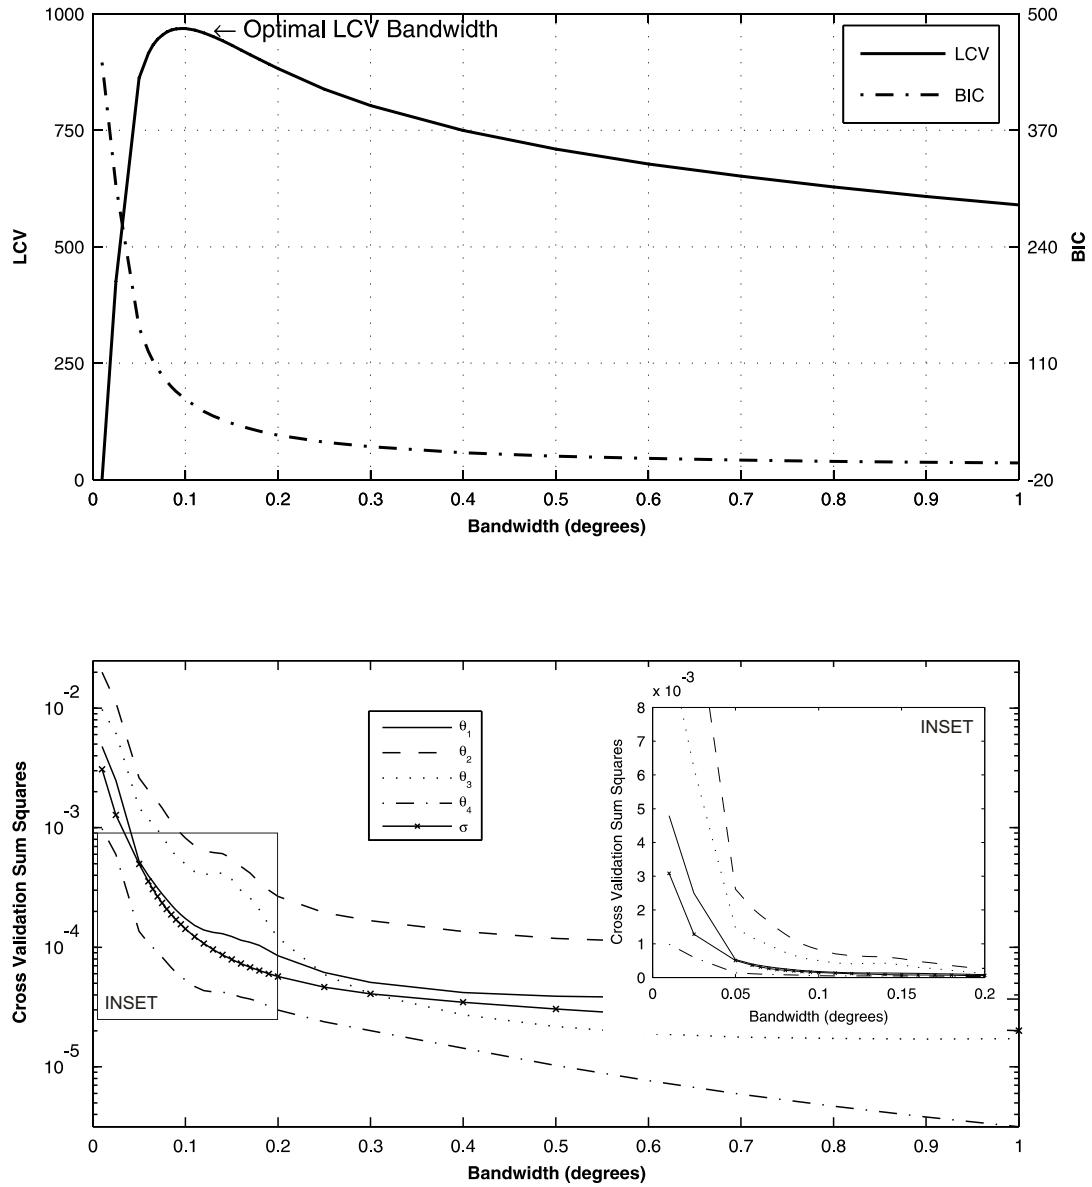


FIGURE 5.3: Comparison of bandwidth selection methods using the April 1985 image of the Southland Front. Top: likelihood cross validation (LCV) and Bayesian Information Criterion (BIC). Bottom: cross validation sum-of-squares for all five parameters ($\theta_1, \theta_2, \theta_3, \theta_4$ and σ).

given the level of noise and missing data. As will be seen in Section 5.3 the algorithms ability to make a realistic estimate of frontal width (θ_3) is degraded as the level of noise and amount of missing data increases. This helps explain why a larger bandwidth is required to smooth along-front estimates of this parameter.

The LCV and BIC both work in likelihood rather than parameter space and select larger optimal bandwidths than the CVSS. This is not an unexpected result since it is already known that cross validation has a tendency to undersmooth (Simonoff, 1996; Hastie and

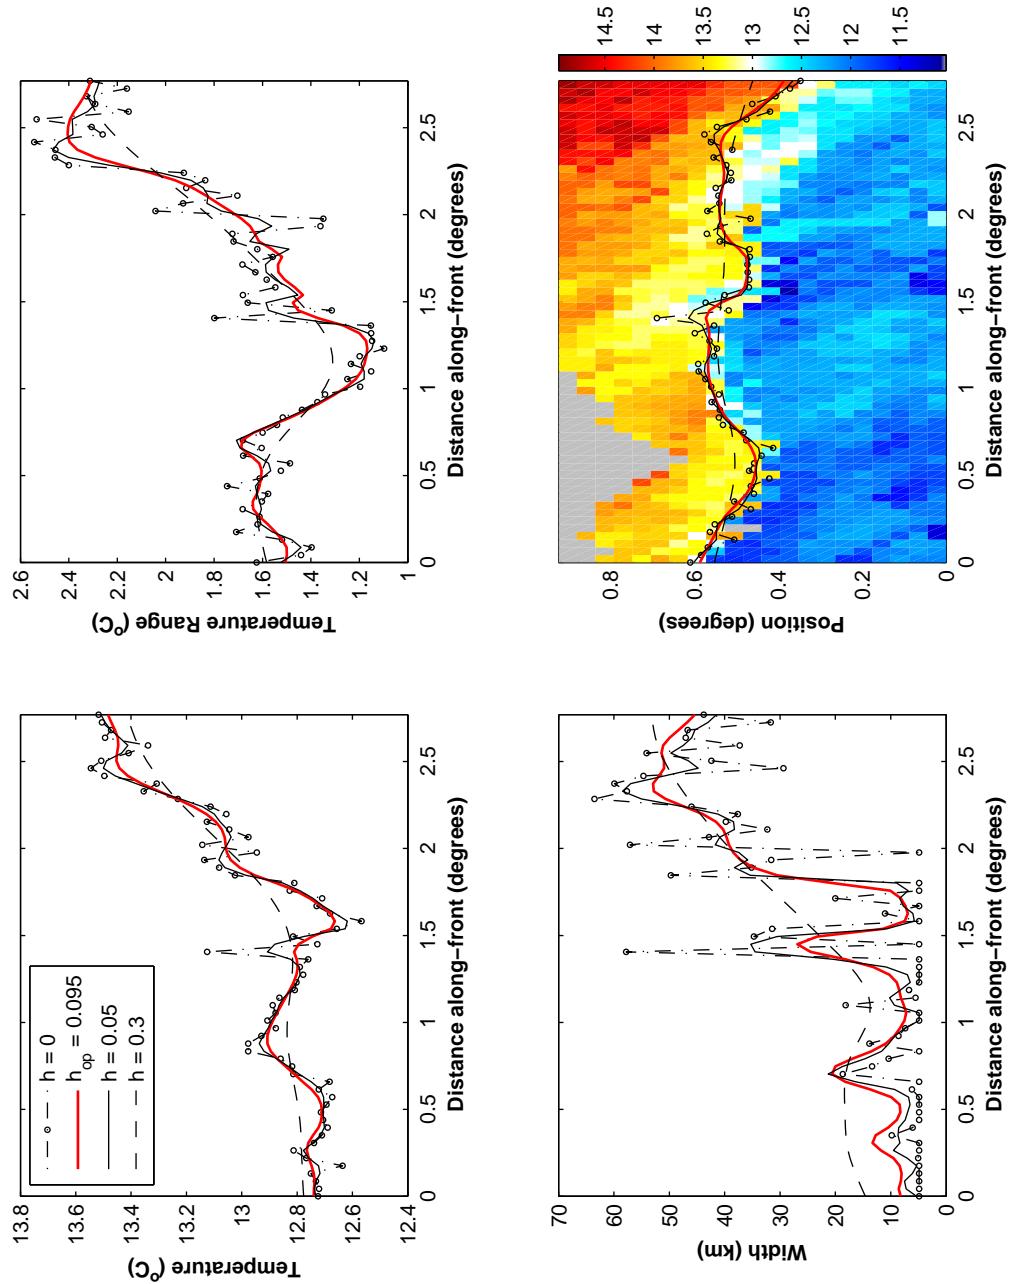


FIGURE 5.4: Local likelihood parameter estimates for bandwidths $h = 0$, $h = 0.05^{\circ}$, $h = 0.3^{\circ}$ and $h_{op} = 0.095^{\circ}$ (as identified by LCV), for the April 1985 image of the Southland Front.

Tibshirani, 1990), whereas LCV is prone to oversmoothing (Silverman, 1986). Short trends and outliers in small data sets may be interpreted as high frequency structure during a standard cross validation. For high weightings, the fitted non-parametric regression will be drawn to these values irrespective of the general trend implied by the remainder of the data set, leading to undersmoothing. An attempt to dampen this effect through additional downweighting is discussed in Section 5.2.3.

Given that different bandwidths are optimal for different parameters, and that no clear minimum is defined for either the CVSS or BIC, the only fully automated and objective way in which to choose a smoothing parameter is to use LCV. Simonoff (1996) considers the effect of autocorrelation. A cyclical pattern in positively autocorrelated errors may be viewed as a high frequency regression relationship, and the bandwidth set small enough to track the cycles, leading to undersmoothing. Autocorrelation analysis within each image is not performed here, but a certain amount of autocorrelation is to be expected. This would further favor using an automated bandwidth selection procedure such as LCV that is prone to oversmoothing in order to compensate.

Out of all 252 monthly images of the Southland Front the median value of h_{op} based on LCV is 0.095° (~ 10.6 km). The median instead of the mean was chosen as it is less biased by outliers. The distribution of the optimal bandwidths selected for each image is given in Figure 5.5. For a gaussian kernel with a standard deviation of 0.095° the highest 68% of weightings are assigned to data within 10.6km (1 s.d) of the point of interest, and 95% within 21.6km (2 s.d's). The baroclinic Rossby Radius around South Island, New Zealand is approximately 20km (Chelton et al., 1998). We can be confident therefore that LCV is selecting a physically appropriate length scale for the smoothing parameter as well as one that balances the bias and variance of the estimates.

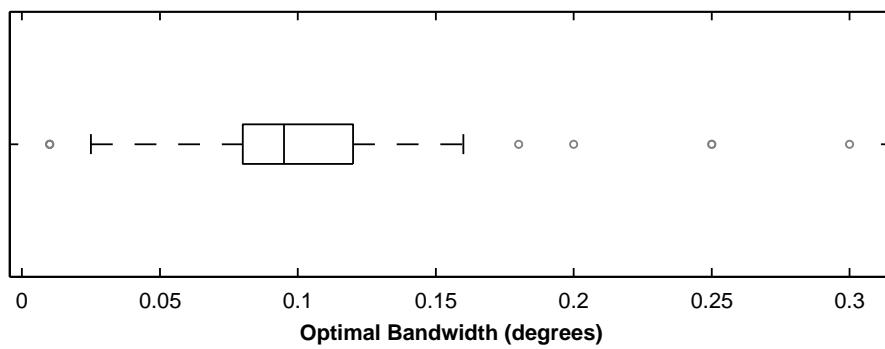


FIGURE 5.5: Distribution of optimal bandwidths calculated for the Southland Front data set (4km resolution)

The remaining practical aspect of making local likelihood model parameter estimates is an assessment of the quality of each likelihood contribution to the weighted local likelihood sum.

5.2.3 Quality Control

The local likelihood estimator (Equation 4.10) developed in Section 4.3.2 is as follows:

$$\hat{\theta}_j = \max_{\theta_j} \sum_{k=1}^m K(x_k - x_j; h) \cdot L(\theta_j; \mathbf{Z}_k).$$

The Gaussian kernel function $K(\cdot)$ distributes weightings to likelihood contributions $L(\cdot)$ at positions k based upon the distance between x_k and x_j . This is designed to ensure that data closest to the point of estimation j have most influence on the local likelihood parameter estimates. As yet however no consideration has been given to the quality or reliability of these likelihood contributions. This is heavily dependent upon the data \mathbf{Z} that they are calculated from.

There are a number of situations where downweighting of likelihood contributions $L(\cdot)$, *in addition to the weightings of the kernel $K(\cdot)$* is desirable. Firstly, in regions where there is a sparsity of across-front observations, the ability to make reliable estimates of frontal characteristics is reduced. The quality of these estimates is somewhat dependent upon the amount of data missing and on how this percentage is distributed across the front. Considering standard MLE's where $h = 0$. If a set of observations \mathbf{Z} is reduced by 50%, where every other data point across the profile is absent, then the overall horizontal temperature structure is maintained and reliable estimates of the temperature and position etc. may be made. At the other extreme, if all the missing values are on one particular side of the front then vital information concerning one water mass is lost and estimates may not be a true reflection of the structure of the front at that point. Of course when there are no observations available then no estimate can be made. If when $h > 0$ the data \mathbf{Z}_k are located close to the point of estimation j and thus assigned a high weighting by the kernel $K(\cdot)$, the localized sparsity and uneven distribution of observations may heavily bias the final local likelihood estimate away from the more realistic values supported by other neighbouring observations.

The second situation in which we might wish to downweight data is where there is no discernible change in gradient across the front, and the best fit to observations \mathbf{Z}_k would be a straight line. This of course may be an accurate reflection of the state of the ocean, a strong subsurface front can exist without a marked surface thermal signature. Strong winds and mixing may break down the surface structure and increased solar insolation in the summer stratifies the water column resulting in an isothermal top layer that can mask the sub-surface structure. Alternatively it may be the case that thermal gradients have been blurred through temporal smoothing. This is particularly true of areas where the front has a high degree of spatial variability and an increased number of plumes, as is the case toward the north of the Canterbury Bight, New Zealand (Shaw, 1998). A noisy set of observations may also make it difficult to pick out any rapid change in

TABLE 5.3: Percentage of data assigned to each quality flag (QF), and the additional weighing (QW) applied during calculation of the local likelihood. The coloured boxes relate to flag assignment detailed in Figure 5.6

Quality Flag (QF)	Weighting (QW)	% of data
Flag 1	1	59.59%
Flag 2	0.8	26.25%
Flag 3	0.6	7.89%
Flag 4	0.4	5.04%
Flag 5	0.2	0.15%
Flag 6	0.1	0.43%
Flag 7	0	0.65%

temperature between the two water masses. In all of these cases, the standard maximum likelihood estimate ($h = 0$) of the width has a tendency to reach the upper bound of 60km. As part of a local likelihood estimate this data would bias the width toward larger values.

Thirdly, we must be cautious where we estimate a very sharp decrease in temperature over spatial scales that we are unable to resolve (<4km). A lack of data directly at the front where SST's are changing most rapidly often results in a very narrow estimate of the width. Unfortunately cloud cover over frontal regions is common and may not always be remedied by monthly composites.

To identify data that could bias local likelihood estimates toward values uncharacteristic of the surrounding area, a set of criteria are used to allocate a quality flag (between 1 and 7) to each set of observations \mathbf{Z} in an image. The criteria are based upon the results from the initial standard maximum likelihood estimate ($h = 0$). Flag 1 is the highest quality and 7 the poorest. The exact criteria used to assign each quality flag are detailed in Figure 5.6 and are discussed in Section 5.2.3.1 below. They are categorized as follows:

- The termination IFAIL value of the Newton-Raphson algorithm
- The value of the gradient (\mathbf{g}) and tolerated accuracy of the estimates (θ_{tol})
- Whether an estimate reached one of the bounds set in Section 5.2.1.2
- The percentage of missing data
- The residual sum of squares (RSS) of the model fit.

Each flag is associated with a quality weighting (QW) between 1 and 0 shown in Table 5.3. This is combined with the weightings assigned by the kernel $K(\cdot)$ to produce a new weighting that takes into account both the location and reliability of each likelihood

contribution $L(\cdot)$. If the quality weightings for each \mathbf{Z} in an image are combined into a $1 \times m$ vector \mathbf{QW} , the j^{th} row of weightings of the final smoother matrix \mathbf{S}_j is given by:

$$\mathbf{S}_j = \frac{K(x_k - x_j; h) \cdot \mathbf{QW}}{\sum_{k=1}^m [K(x_k - x_j; h) \cdot \mathbf{QW}_k]}. \quad (5.5)$$

Normalization ensures that $\sum_{k=1}^m \mathbf{S}_{j,k} = 1$. Using weightings from the matrix \mathbf{S} when calculating the local likelihood results in a smooth non-parametric along-front trend in temperature and position etc. less likely to be interrupted by outliers.

The value of the maximum likelihood, the residual sum of squares of the model fit, the estimated noise $\hat{\sigma}$, the percentage of missing data and the optimization IFAIL value are all used to help select quality flags (**QF**). The specific criteria are discussed in the following section.

5.2.3.1 Quality Control Criteria

Figures 5.7 and 5.8 show the relationships between the value of the standard maximum likelihood (Equation 4.5), the percentage of missing data from vector \mathbf{Z} , the residual sum of squares (RSS) of each model fit (normalized by the number of available observations), the IFAIL value of the optimization, and the estimated standard deviation of noise about each fit ($\hat{\sigma}$).

The first point to note is that the quality of the solution found by the optimization algorithm should not be judged solely on whether a non-zero IFAIL has been reported or not. The stringent accuracy demanded by the NAG optimization routine may result in an otherwise acceptable solution being reported as a failure. An IFAIL value of 5 warns that there is doubt as to whether a minimum has been found (Table 5.2). Estimates carrying this exit flag can not be selected or rejected based on the percentage of missing data, $\hat{\sigma}$, or the RSS (Figures 5.7 and 5.8b.). Therefore, when an exit flag of 5 is reported, the following two checks are made to determine whether $\hat{\theta}_{sol}$ is still a good estimate of θ_{true} :

$$\|\mathbf{g}(\hat{\theta}_{sol})\|^2 < 10 \times \epsilon \quad (5.6a)$$

$$cond < \frac{1.0}{\|\mathbf{g}(\hat{\theta}_{sol})\|}, \quad (5.6b)$$

(Numerical Algorithms Group, 2006)

where $cond$ is an estimate of the condition number of the projected Hessian at $\hat{\theta}_{sol}$. The condition number of the Hessian is a measure of the susceptibility of the estimates

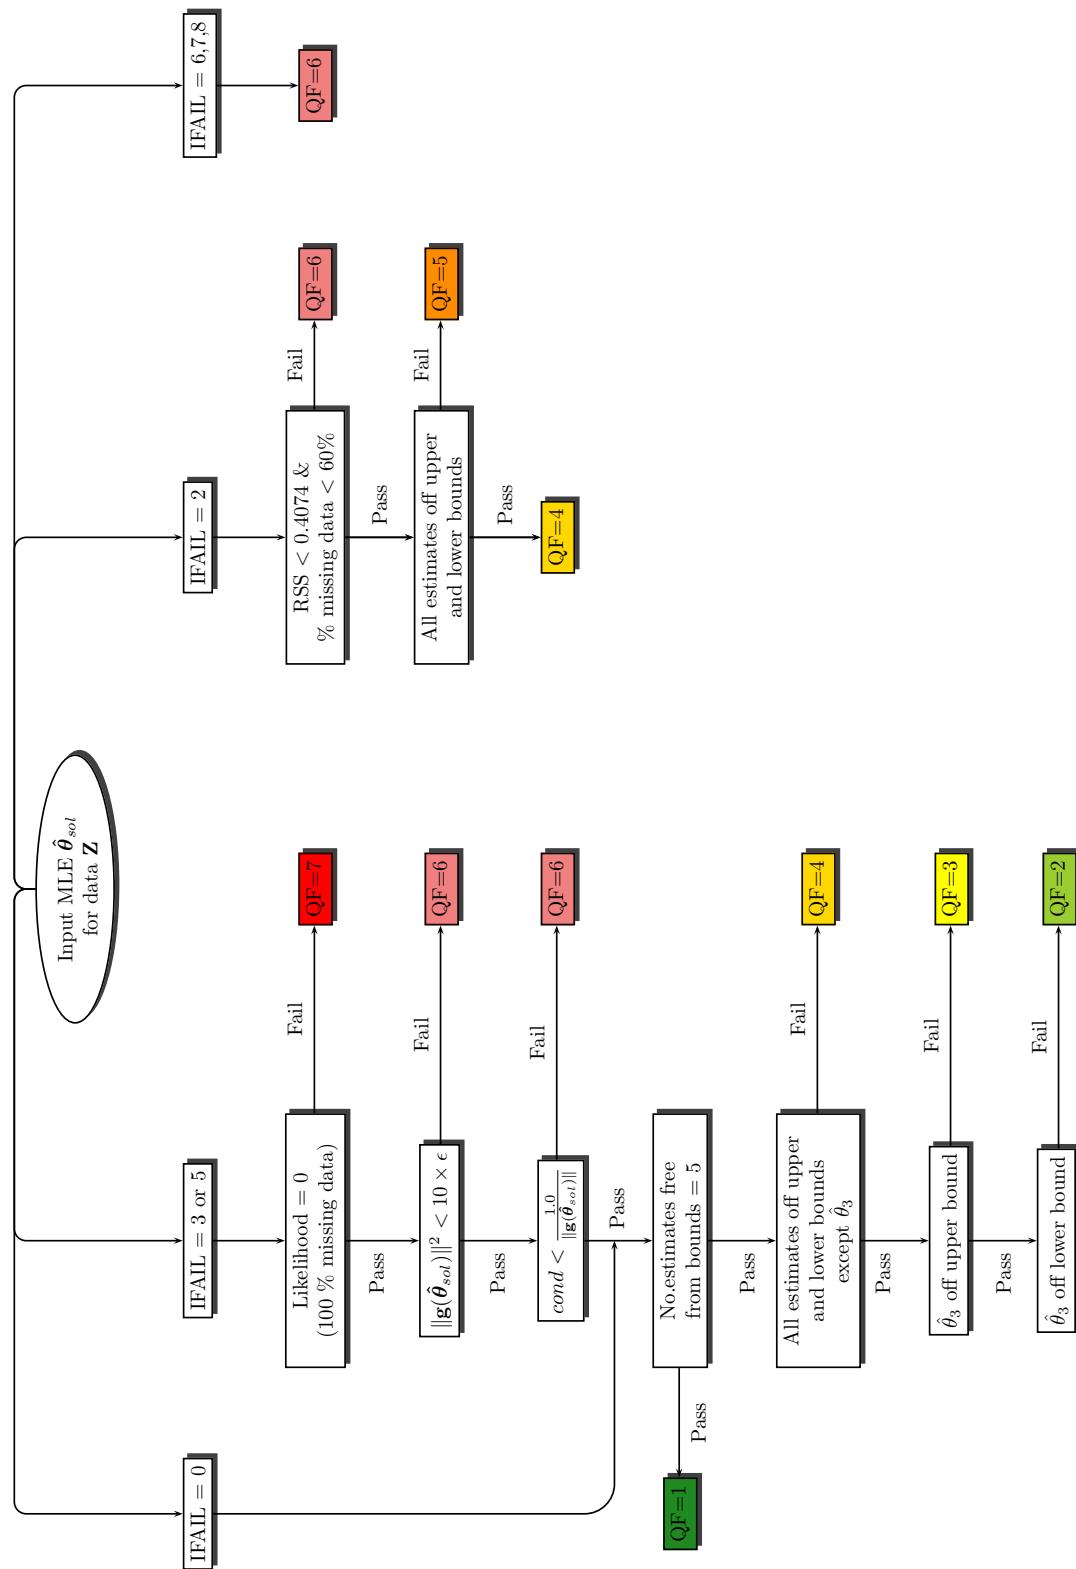


FIGURE 5.6: Flowchart used to assign quality flags (QF) to data. Refer to Table 5.3 for the associated weightings.

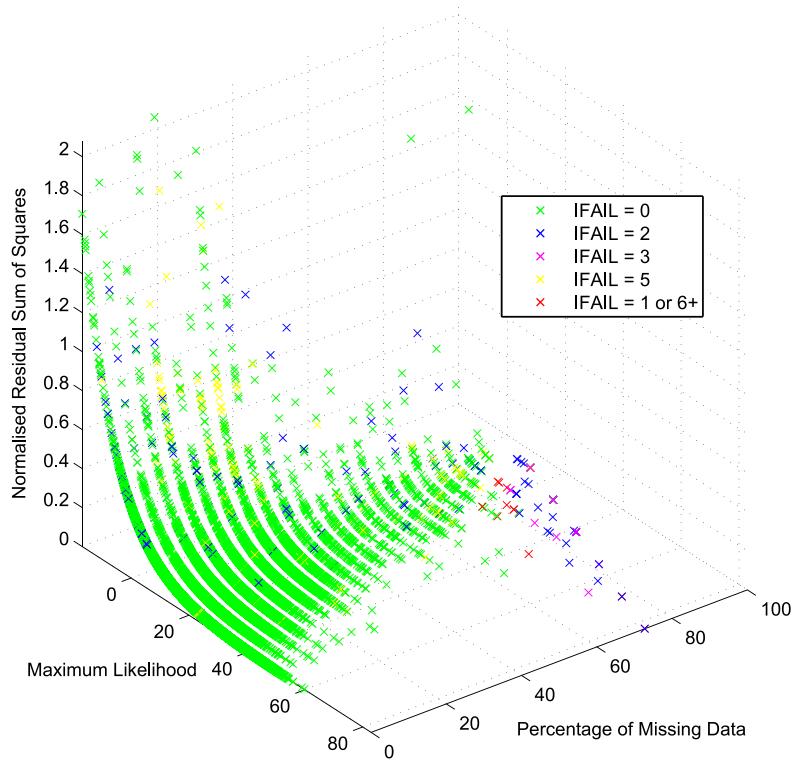


FIGURE 5.7: The final function value, the percentage of missing data and the residual sum of squares of each standard maximum log likelihood model fit ($h = 0$).

to errors associated with the quadratic approximation of the likelihood (Gill et al., 1995). A large condition number indicates that the solution is sensitive to errors in the approximation. It implies inaccurate estimates and slow convergence. A small condition number (close to one), indicates a high rate of convergence and reassurance of the accuracy of the gradient vector ($\mathbf{g}(\hat{\theta}_{sol})$). When both Equations 5.6a and 5.6b are true we may be more confident that $\hat{\theta}_{sol}$ is a close approximation to the position of the minimum. This test is also performed when IFAIL = 3.

Similarly, the number of observations, the value of the likelihood, the RSS nor the estimated noise alone can be used to identify when the algorithm will struggle to converge from its initialization point, or where there is no minimum (IFAIL = 2). Data that results in estimates falling into this category are assigned a weighting based on the RSS, the percentage of missing data, and whether any estimates have reached their bounds during optimization. Figure 5.8a shows how the value of the standard maximum likelihood, as defined in Equation 4.5, is controlled by the number of observations and the level of noise. The pivotal point, above which the first term of the likelihood expression ($-n \ln \sigma$), is larger than the second ($-\frac{1}{2\sigma^2} \sum (z^{(i)} - m(\cdot))^2$) i.e. the likelihood is negative, is where the $RSS = 0.4074$ (Figure 5.8c). This corresponds to a $\hat{\sigma}$ of 0.6063. Since the accuracy of the estimates deteriorates with increasing amounts of noise (Section 5.3.2), this seems a natural cutoff point with which to help assign weightings. Estimates not on their bounds with an RSS less than 0.4074 that are based on data (\mathbf{Z}) with less than 60%

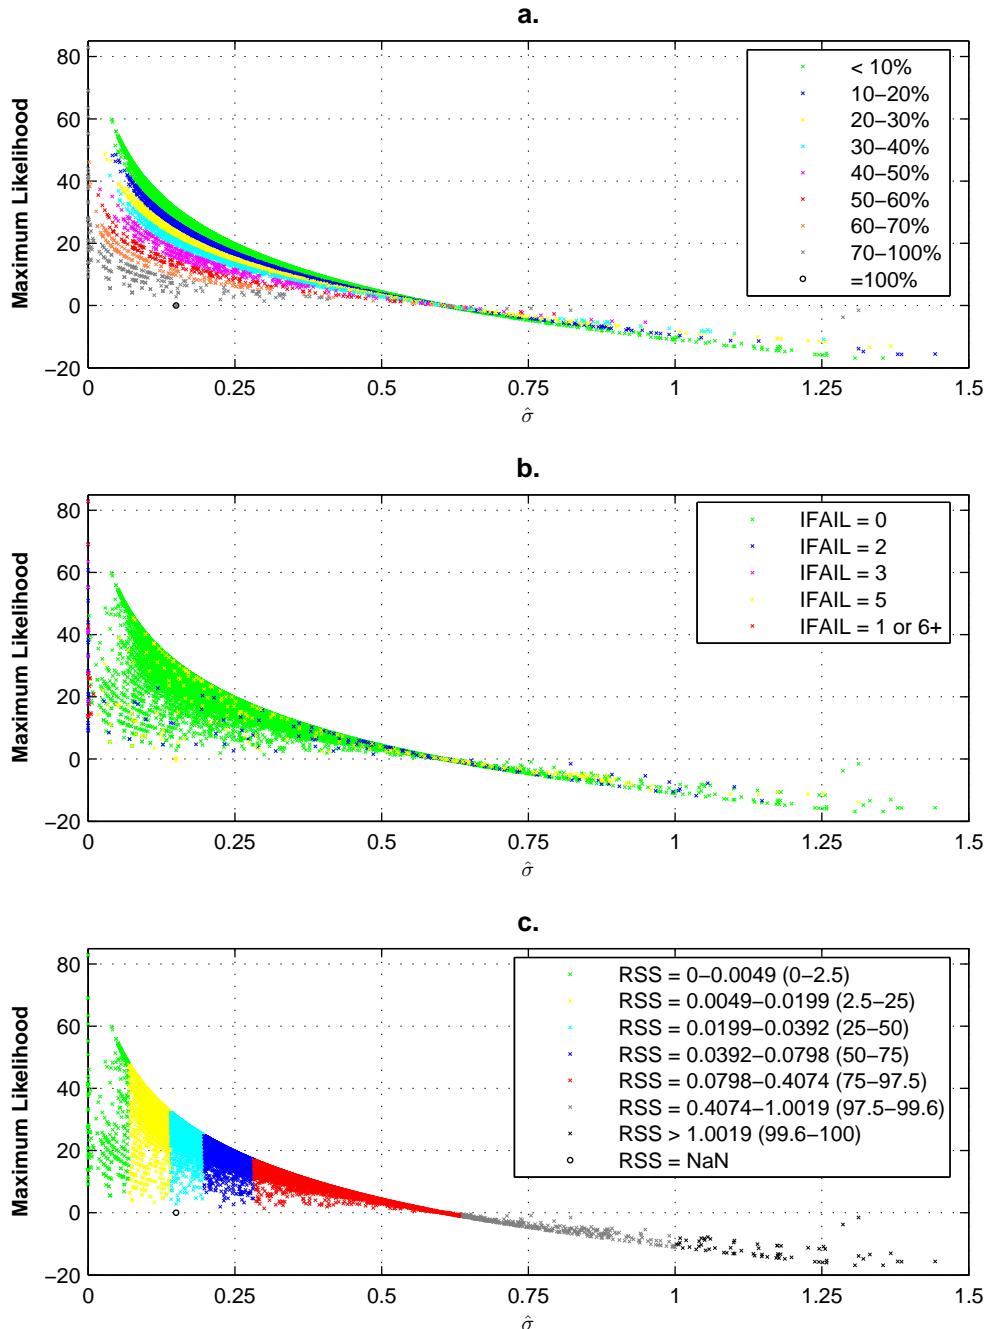


FIGURE 5.8: **a.** Estimated Noise ($\hat{\sigma}$) vs Maximum log likelihood (L) vs percentage of missing data. **b.** $\hat{\sigma}$ vs L vs optimization exit flag (IFAIL). **c.** $\hat{\sigma}$ vs L vs normalised residual sum squares (RSS). Bracketed figures in the key indicate the percentile range corresponding to the RSS.

of observations missing are not too heavily penalised. Convergence may be achieved by starting the optimisation from a different place.

Alternative to the scenario where acceptable estimates have been flagged with a warning, is the situation where we wish to downweight noisy or sparse data that result in estimates with an IFAIL value of 0, but which do not contain any useful oceanographic information (the model has not been fitted to a frontal structure). This is achieved by monitoring which estimates have been pushed onto their upper or lower bounds. For example, an estimated temperature difference of 6°C is likely the result of cloud contamination. With the exception of high latitudes clouds are generally radiometrically colder than the sea surface. A contrast in temperature of the order of $5\text{-}6^{\circ}\text{C}$ between two neighboring bodies of water is generally accepted as the threshold above which cloud contamination is suspected (Shaw and Vennell, 2001). Table 5.4 shows the percentage of standard maximum likelihood estimates ($h = 0$) that reached upper or lower bounds during optimization. Extreme value estimates of the width are responsible for most of the downweighting due to boundary conditions.

TABLE 5.4: Percentage of standard maximum likelihood estimates that reached an upper or lower bound (for all values of IFAIL)

Parameter	Percentage
Mean Temperature $\hat{\theta}_1$	0.30
Temperature Difference $2\hat{\theta}_2$	3.21
Width $2\hat{\theta}_3$	37.20
Position $\hat{\theta}_4$	3.14
Standard Deviation $\hat{\sigma}$	0.70

5.3 Implementation: Stability, Sensitivity and Limitations

In this next section we examine; a. the sensitivity and stability of the Newton-Raphson optimization to its initialization conditions, b. the errors associated with a noisy data set, c. the implications of a meandering front and d. the robustness of the algorithm to missing data. We investigate a, b and c using a simulated data set with known parameters. An artificial image of the same size ($1^{\circ}\times 2.8^{\circ}$) and resolution (4km) as the Southland Front data set was constructed. A straight front with a mean temperature of 12°C , a temperature range of 2°C and a width of 20km (10km for test c.) was built centrally across the image and orientated east-west. Normally distributed random noise with a standard deviation of 0.1 was added to this base image. This value seemed appropriate based on experience with the Southland Front data set and development of the algorithm. 0.1°C is also the approximate relative accuracy of the AVHRR Pathfinder data (K. Casey personal comment).

5.3.1 Sensitivity to Initialization Conditions (a)

The Newton-Raphson algorithm is known to be sensitive to its initialization conditions (Gill et al., 1995). In an attempt to approximately quantify this instability with respect to the specific nonlinear sigmoid function used here, we perform a series of ‘sensitivity’ tests. Standard maximum likelihood estimates ($h = 0$), were made repeatedly at all positions x_j along the simulated front starting from different sets of initial estimates. Only one parameter was varied at a time. All other initial estimates were set at their correct value (not taking into account the effect of additional random noise). This test was performed without any parameter bounds being set during the optimization.

In general, we find that for a percentage error of $\pm 100\%$ in the initialization point of one unknown, correct estimates of all five parameters are made. Figure 5.9 shows the exception to this rule. The optimization is sensitive to the initialization value of the width (θ_3). Once $\hat{\theta}_3$ falls below 2km the routine is unable to recover the true value of ~ 20 km. This cutoff is related to the resolution of the data set and provides further justification for setting a lower bound of 4km on the width. In an attempt to compensate, under and over estimates of θ_2 and σ respectively are made. θ_4 defaults to its starting value of 0.5 beyond a -90% error. θ_1 fails in the same way beyond a -99.5% error suggesting that it is the more robust parameter. This is not unexpected since θ_1 is a ‘translation’ parameter of the model, whereas the width (θ_3), and temperature range (θ_2) are nonlinear components of the model function, and control the shape of the sigmoid.

An indicator of a successful model fit is given by the normalized residual sum of squares (RSS). This is seen to increase by an order of magnitude from 0.0078 to over 0.05 once the optimization is unable to converge on the correct solution.

5.3.2 Noise Limitation (b)

An examination of the errors expected to be introduced to estimates by measurement errors in the AVHRR SST, and noise introduced by resolution limitations and processing inaccuracies is assessed in Figure 5.10. The loss in accuracy and precision of estimates as the level of noise increases in an image is shown for bandwidths of 0° and 0.15° . It demonstrates the improvement made by introducing the smoothing function. Normally distributed random noise with standard deviations between 0.05 and 0.55 was added to the simulated data set. For $\sigma = 0.05$ the front is exceptionally well defined. When $\sigma > 0.55$ no frontal structure can be distinguished through the noise. Standard maximum likelihood estimates ($h = 0$), and local likelihood estimates with $h = 0.15^\circ$ were made at each point.

For a bandwidth of zero, the range of estimates increases rapidly as the amount of noise becomes more significant. This is brought under control by increasing the smoothing

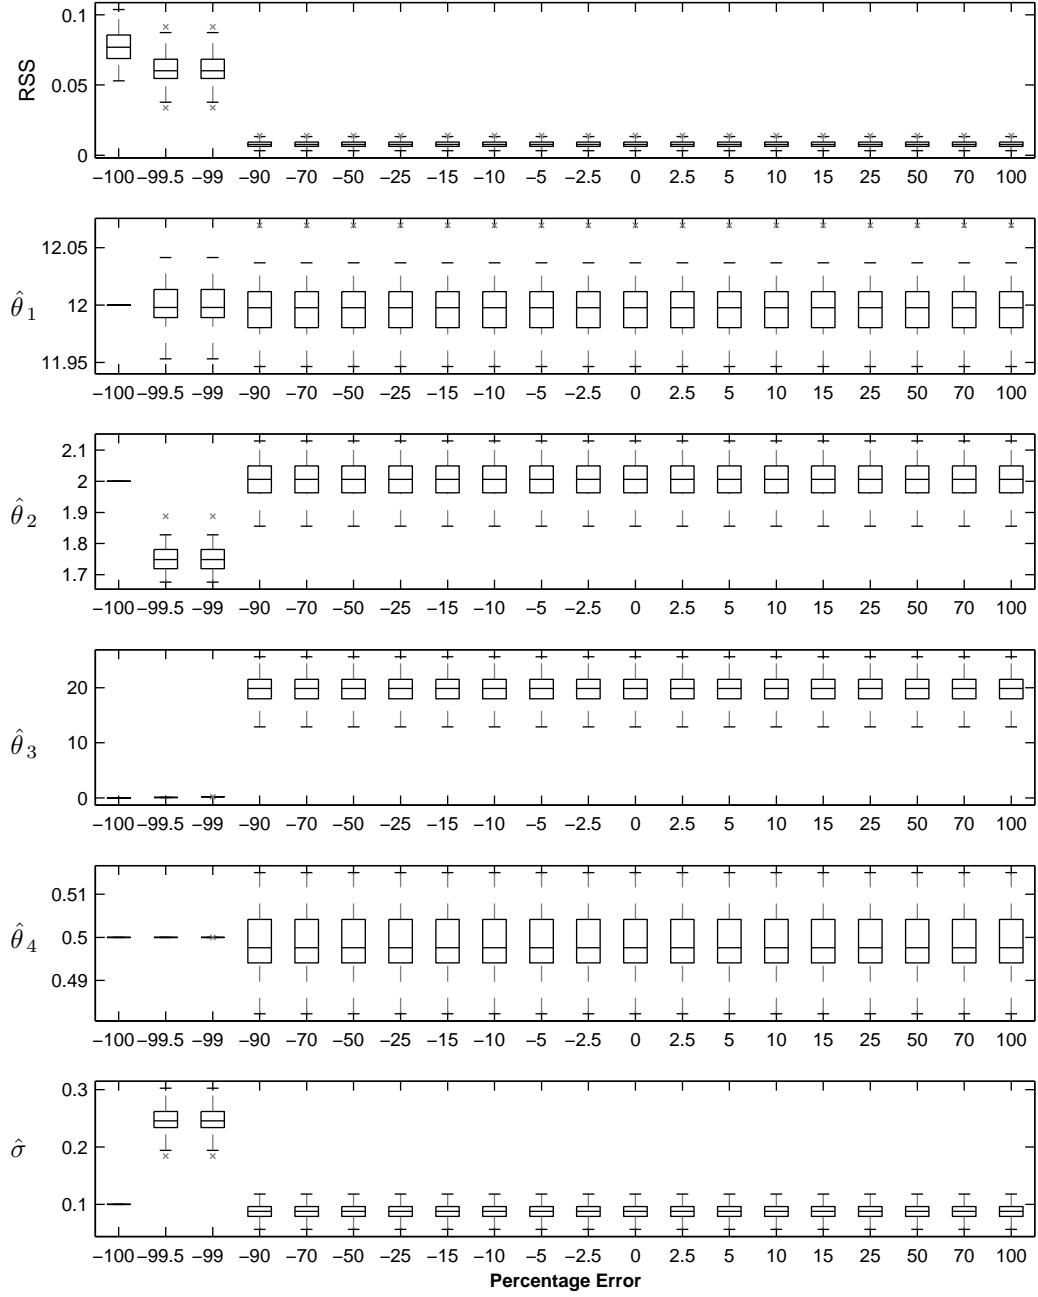


FIGURE 5.9: Sensitivity of Newton-Raphson optimization to percentage error in the starting point of the estimated width. All other parameters are initialized with a zero percent error. A -90% error corresponds to a width of 2km.

parameter to $h = 0.15^\circ$. Note that for $h = 0$, the level of noise is increasingly underestimated as the true standard deviation of normally distributed noise increases. This is corrected by using a bandwidth greater than zero.

If the relative accuracy of AVHRR measurements between pixels is of the order 0.1°C (K. Casey personal comment) we conclude that the local likelihood estimates are not overly sensitive to noise and errors in the AVHRR SST measurements.

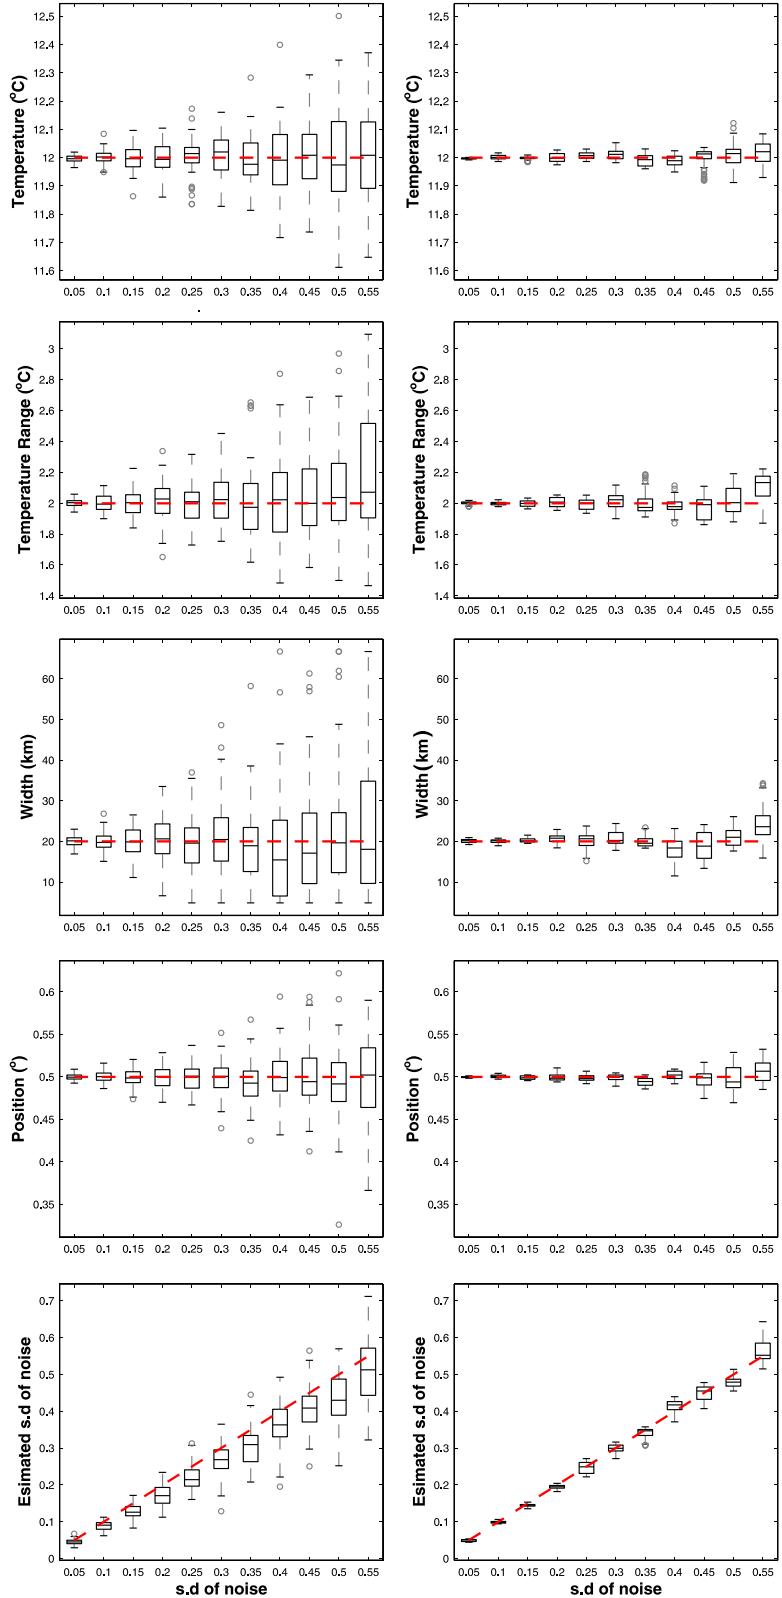


FIGURE 5.10: Distribution of estimated temperature, temperature range, width, position and s.d. of noise from a simulated data set with a smoothing parameter $h = 0$ (left column), and $h = 0.15$ (right column). Normally distributed random noise with standard deviations between 0.05 and 0.55 was added to a base image of known parameters (dotted red line).

5.3.3 Curvature Limitation (c)

With the exception of the initial extraction of data within a rotated window (Section 5.1.2), the current version of the algorithm is unable to adapt to slight changes in the orientation of the front. We fit our model based on the assumption that the front is orientated east-west across the image (i.e. each cross section (\mathbf{Z}) is at right angles to the front). In reality fronts meander and the angle at which each cross section bisects the front may change. It is important to be aware of the implications that such an assumption has on results.

The artificial front (with a width of 10km) was rotated between 0° and 90° from the horizontal (Figure 5.11). A front with an angle of 0° is bisected at right angles by each cross section. An angle of 90° represents a front orientated north-south across the image.

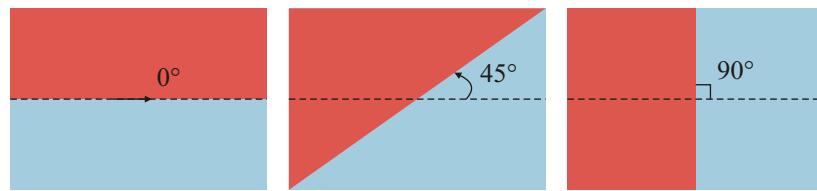


FIGURE 5.11: From left to right, a front oriented east-west (0°), at 45° and north-south (90°) across an image.

Figure 5.12 shows how each parameter estimate deviates from its true value as the angle of rotation increases. Estimates of the mean temperature, temperature range and position remain unaffected by the orientation of the front. Estimates of the width and therefore the gradient however are more sensitive. The estimated width exponentially increases as the angle steepens (Figure 5.12c). Beyond 25 - 30° the width is noticeably overestimated resulting in a serious underestimate of the gradient. The overestimate in width increases from 14% of the true value at an angle of 25° to over 400% at 80° . The corresponding percentage underestimate of the gradient ranges from 13% to 80%. The standard deviation of noise about the model fit (Figure 5.12f) is slightly underestimated by 0-5% of the true value between 0° and 40° . As the angle increases beyond 40° this parameter starts to become overestimated, reaching a maximum of 18% of the true value when the front is rotated by 80° . This experiment highlights that estimates of the width and gradient must be treated with caution in regions where the front meanders at angles greater than 30° .

The point at which the algorithm ‘collapses’ is shown by the gray shaded box. Beyond an angle of 80° no useful output may be obtained. Eddies and meanders that loop back on themselves are therefore not resolved.

Using estimates of the front’s orientation made by Shaw (1998) between April 1989 and March 1992 the Southland Front is calculated to have a mean (± 1 s.d) orientation of $46.34 \pm 29.88^\circ$ relative to due north. Only points falling within the extraction window

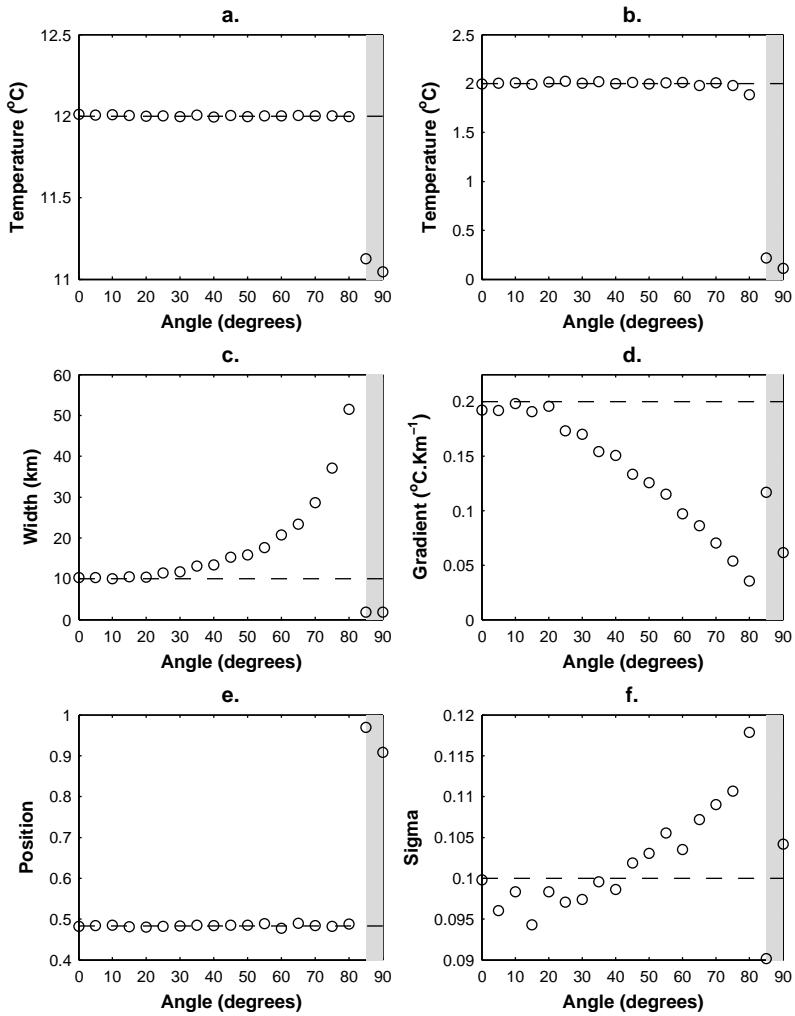


FIGURE 5.12: Effect of frontal orientation on parameter estimates. Angle vs a. Mean Temperature, b. Temperature Range, c. Width, d. Gradient, e. Position and f. s.d of noise about mean. The dashed line indicates the true value. The circles are the algorithm estimates. The shaded gray box represents the angle past which the algorithm ‘collapses’

defined in Figure 5.2 were used in this calculation. Firstly, this justifies the 45° rotation of the extraction window. Secondly, we may be confident that the majority of estimates made using the local likelihood algorithm should not be too severely effected by local changes in the orientation of the front as it flows north-east; a large percentage ($\sim 68\%$) of observations (\mathbf{Z}) should be taken at angles $\leq 30^\circ$ across the front. Infact this may even be considered a conservative estimate since the data set used by Shaw (1998) was of higher temporal and spatial resolution than the images used in this study and therefore small scales meanders and plumes would be better defined.

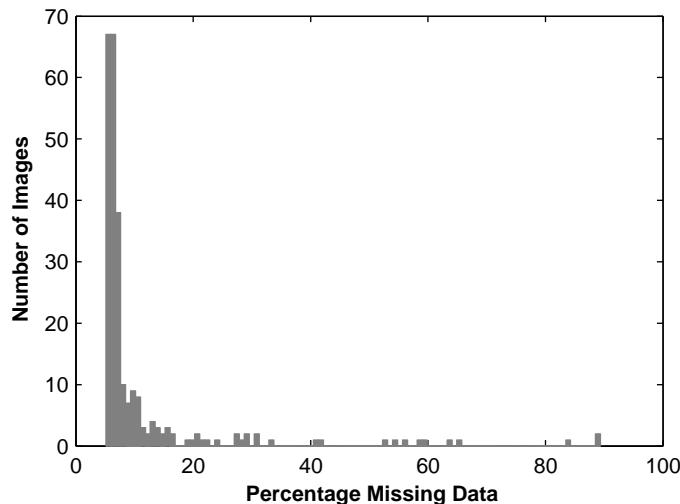


FIGURE 5.13: Histogram of the percentages of missing data in the 252 monthly images of the Southland Front

5.3.4 Sparsity of Observations (d)

The Southern Ocean is a region where IR sensors struggle to make measurements because of the high levels of cloud cover. Studies of Southern Ocean frontal dynamics have in the past been hampered by a high percentage of missing data. A study of the Antarctic Polar Front by Moore et al. (1999) was hindered by more than 45% of their AVHRR SST data set being missing. One of the objectives in the development of a new front detection algorithm was to minimize the loss of information owing to missing data in an image. Figure 5.13 is a histogram of the percentage of missing data in each monthly image of the Southland Front data set. Figure A.3 in the Appendices gives a more detailed breakdown of this distribution over time. Between February 1995 and December 1995 (inclusive) anomalously high percentages of observations are missing in each image. February 1995 coincides with a change over between the NOAA-9 and NOAA-14 polar orbiting satellites used in the Pathfinder Version 5.0 processing for SST (NODC, 2008). The reduced number of observations at this time is believed to be due to errors in the brightness temperatures for channel 4 in the AVHRR instrument onboard the NOAA-14 satellite (P.Miller personal comment). The missing and erroneous values, most apparent at low brightness temperatures, appear to be related to problems with the instruments analog-to-digital conversion of radiances (Podestá et al., 2003). Pixels are assigned incorrect SST values, and when processed with the usual quality tests (Figure A.1 and Table A.1), a greater number than usual are assigned low quality flags (the total number of observations remains constant). Note that Podestá et al. (2003) describe the effect of digitizer error on the AVHRR-14 data in October 2000. The description of the error however also applies to the problems seen in 1995 in the Pathfinder Version 5 data.

Injection of large amounts of infrared absorbing aerosols into the atmosphere from significant volcanic eruptions such as Mt Pinatubo in 1991 can also reduce the number of

TABLE 5.5: Mean standard errors of estimates made with increasing amounts of missing data

Parameter	Cloud Free	16% missing	41% missing	42% missing
	Fig 5.14(a)	Fig 5.14(b)	Fig 5.15(a)	Fig 5.15(b)
$\hat{\theta}_1$	0.035°C	0.036°C	0.047°C	0.042°C
$\hat{\theta}_2$	0.043°C	0.042°C	0.042°C	0.040°C
$\hat{\theta}_3$	2.59km	2.67km	3.07km	3.12km
$\hat{\theta}_4$	0.015°	0.017°	0.017°	0.020°

SST retrievals from the AVHRR sensor and negatively bias SST's (Reynolds, 1993). Mt Ruapehu on North Island, New Zealand became active during 1995 which may account for some of the poor SST retrievals.

The majority of images are missing less than 20% of SST observations. The mean percentage of missing data in each image is 11% (median = 6.6%).

To test the robustness of the local likelihood model to data loss, SST observations are removed from a cloud free image of the Southland Front (September 1997). Instead of trying to simulate the real spatial patterns of cloud cover and contaminated SST retrievals, the patterns of missing data from a selection of other months provides us with a means of removing observations.

Figure 5.14a. shows along-front estimates of frontal parameters for the cloud free image of September 1997. Missing data in the top-left corner is the headland off Dunedin and accounts for 5% of all the percentages that will be quoted in the following discussion. In Figure 5.14(b), a small (16%), portion of the observations have been removed, mostly over the front itself. Estimates of the temperature, temperature range, width and position remain, on the whole, unaffected. In Figures 5.15a. and 5.15b. over 40% of observations have been removed. Estimates of the front's position are the most robust to a sparse data set, deviating very little from the original estimate made with a complete set of observations (red line). Where SST measurements across the front have been removed the width is underestimated, although the estimate made with a complete data set falls, for the most part, within the 95% confidence interval. There are two peaks ($>50\text{km}$) in the estimated width (Figure 5.14a.). These points correspond to increases in the angle of the front relative to the horizontal and subsequent overestimation of the width (Section 5.3.3).

The distribution of missing values is important in determining the quality of estimates. In Figures 5.15a., and 5.15b. 41% and 42% of observations have been removed respectively. Although there is little difference between the amounts of missing data in these two examples the algorithm returns better estimates of the temperature, width and temperature range in Figure 5.15b.

Table 5.5 shows the mean standard error for each estimated parameter as a result of a reduction in the number of available observations. As expected, the mean standard error increases with higher percentages of missing data for estimates of the temperature, width and position. Confidence in the estimates of temperature range however increases - the standard error decreases from 0.043 to 0.040°C.

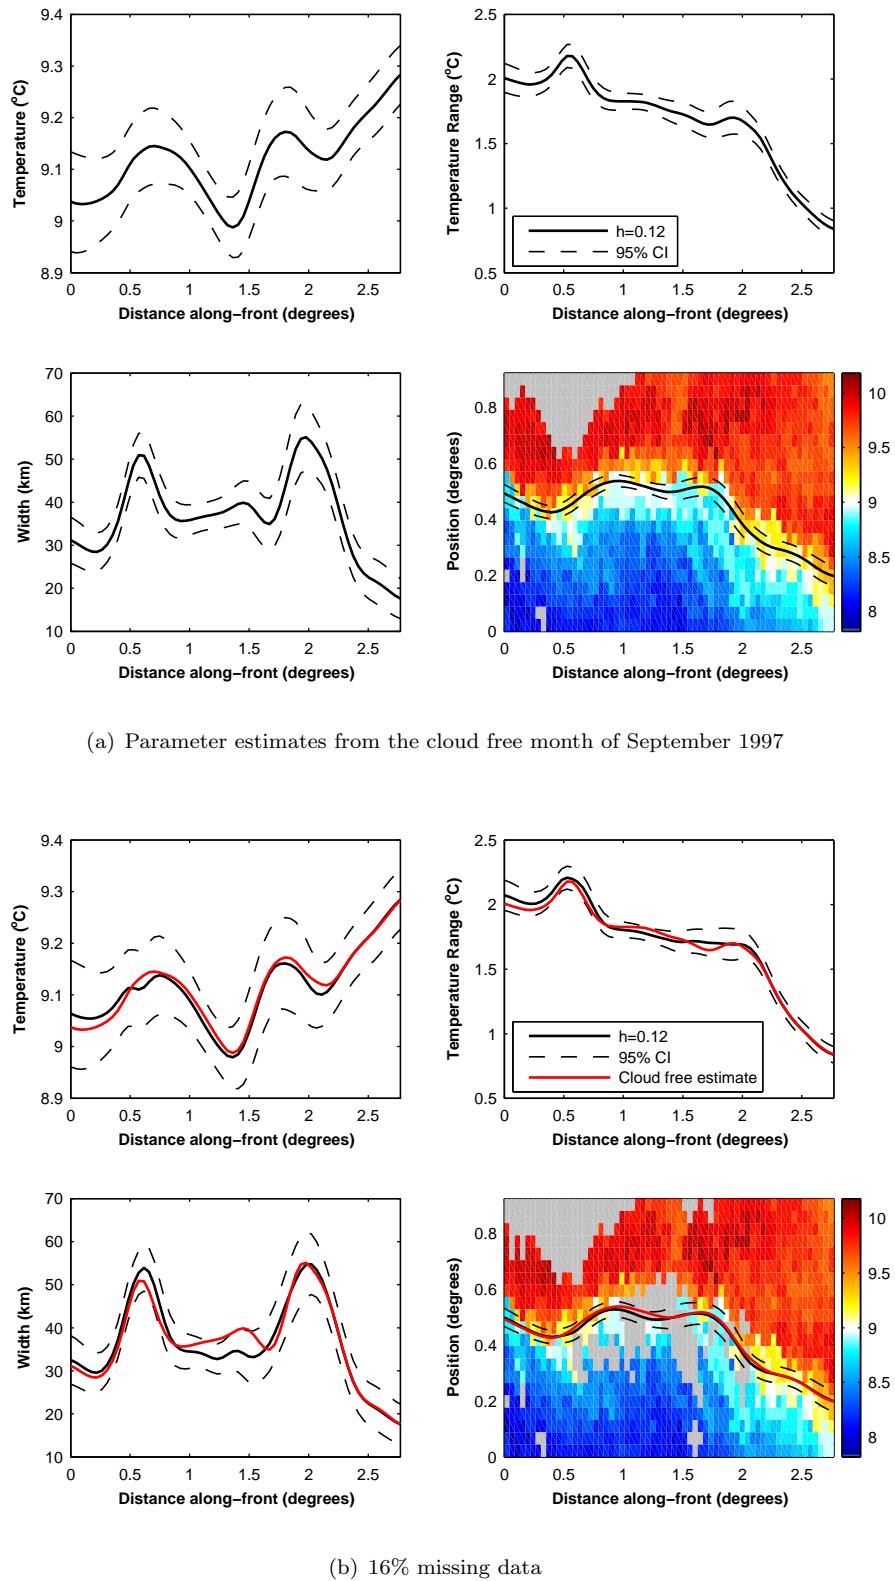
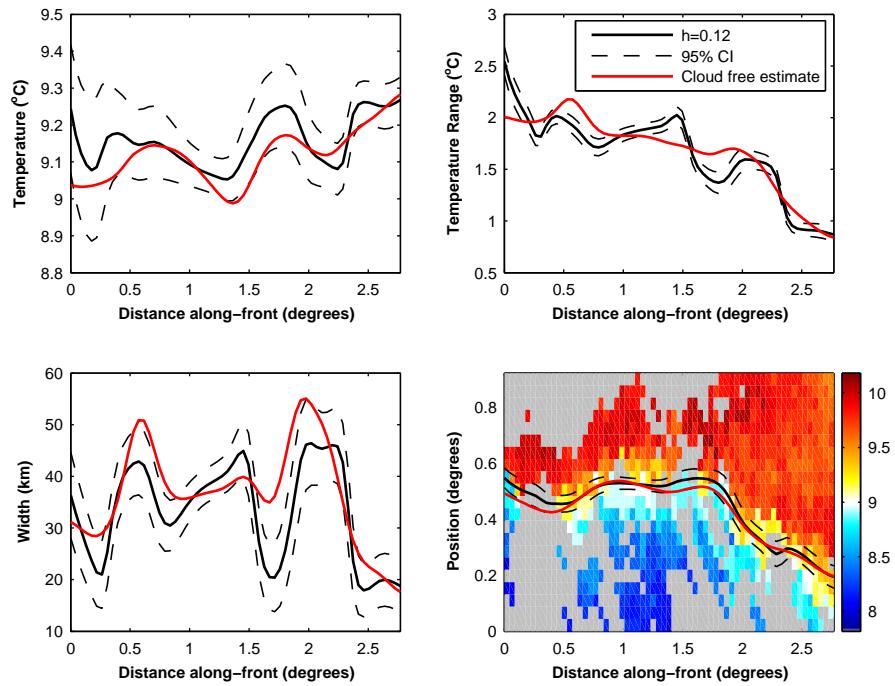
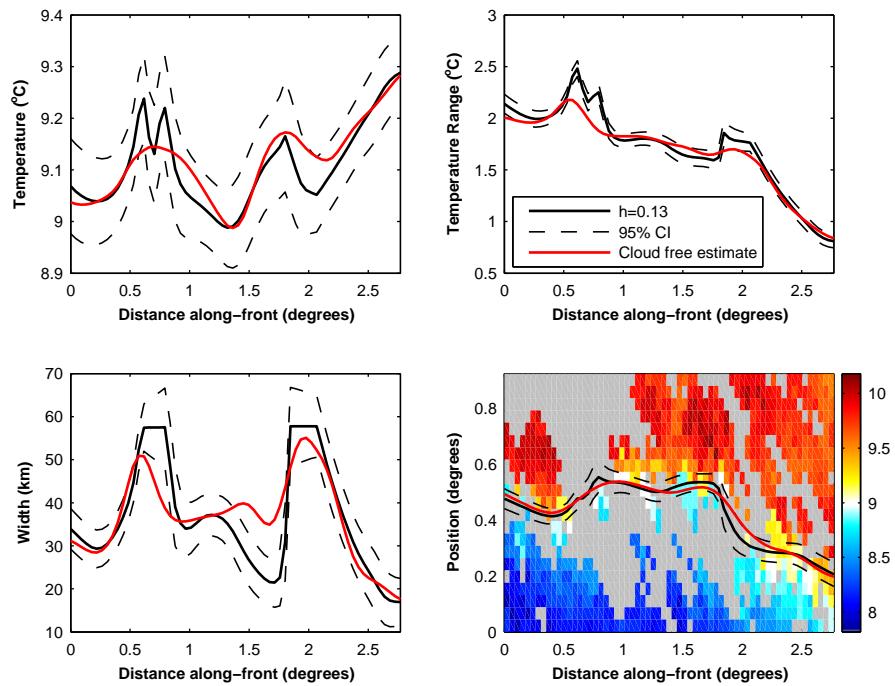


FIGURE 5.14: Robustness to sparse data sets (1)



(a) 41% missing data



(b) 42% missing data

FIGURE 5.15: Robustness to sparse data sets (2)

5.4 Summary

The main content of this chapter may be summarized as follows:

- 4km resolution monthly AVHRR SST composite images of the Southland Front between January 1985 and December 2005 downloaded for analysis.
- SST measurements are extracted from within a $1^\circ \times 2.8^\circ$ window orientated normal to the approximate direction of the Southland Front (rotation angle = 45°).
- Newton-Raphson optimization scheme used to solve the likelihood equations. A linear transformation is applied to all unknown variables.
- Initial parameter estimates made with a bandwidth of 0° (standard maximum likelihood). Optimization initialized using mean frontal characteristics estimated by Shaw and Vennell (2001).
- Optimization of the local likelihood with a bandwidth $>0^\circ$ initialized using the standard maximum likelihood estimates.
- An optimal bandwidth calculated using likelihood cross validation.
- Gaussian kernel weightings ($K(\cdot)$) are combined with an additional quality weighting (QW) between 1 and 0 designed to downweight the influence of those areas with a sparsity of data.
- Optimization is most sensitive to the starting value of the width. Optimization is unsuccessful if the estimated width drops below 2km during any iteration.
- Local likelihood estimates are not overly sensitive to noise and errors in the AVHRR SST measurements.
- Where the front meanders at angles greater than 30° from the horizontal the width is overestimated by approximately 14%. This leads to an underestimate of the gradient. Estimation of the mean temperature, temperature difference and position is more robust to changes in the front's orientation.
- Local likelihood estimates are robust to large percentages of missing observations in an image ($\sim 40\%$). The errors associated with parameter estimates increase as the number of available measurements decreases.

CHAPTER 6

Spatial and Seasonal Variability of The Southland Front

6.1 Introduction

In this Chapter we present a detailed study of the seasonal and spatial variability of the Southland Front using the front detection algorithm introduced in Chapters 4 and 5.

High biological activity at the Southland Front means that it is a commercially valuable zone (Jackson et al., 2000). Fisheries management schemes, put in place to ensure that fishing within the frontal zone is a sustainable activity, require an understanding of the seasonal movements of the front and the timing and extent of fluctuations in its strength. Furthermore, variability in the front's position reflects its stability and fluctuations in the gradient provide an indication of the associated geostrophic currents. These are both characteristics that will determine the structure, strength and stability of the global Subtropical Front as it enters the open Pacific basin.

6.2 Methods

A twenty-one year series of monthly AVHRR SST images of the Southland Front between 46.4°S and 44.5°S form the base data set for this study. Further details are given in Section 5.1 of Chapter 5. The position and characteristics (temperature, width and temperature range) of the front in each image were determined using the front detection algorithm outlined in Chapters 4 and 5.

Data were analyzed spatially and over two different timescales; a twenty-one year period and seasonally. Southern hemisphere seasons are used throughout, so summer is taken as December, January, February (DJF); autumn as March, April, May (MAM); winter as June, July, August (JJA), and spring as September, October, November (SON). We also separate the three year period (April 1989-March 1992) studied by Shaw and Vennell (2001) as a means of independently verifying our results. Only those estimates receiving a Quality Flag of 1 as derived in Figure 5.6 were accepted for the study. The standard errors of each estimate are used to calculate weighted mean statistics for each of the time periods specified.

6.2.1 Weighted Statistics

Attached to each estimate of position and temperature etc. is a standard error representing the estimated standard deviation of the error in our algorithm (see Section 4.2.3). The level of confidence in an estimate increases as its standard error gets smaller. This measure of uncertainty is used to calculate weighted mean statistics for each frontal parameter.

The *weighted mean* \bar{x}_w of estimates x_i , where $i = \{1, 2, \dots, n\}$, assumed to be made from the same parent distribution but each having its own standard error σ_i representing the confidence with which that particular point was estimated is expressed as:

$$\bar{x}_w = \frac{\sum_{i=1}^n x_i w_i}{\sum_{i=1}^n w_i} \quad \text{where} \quad w_i = \frac{1}{\sigma_i^2}. \quad (6.1)$$

Each estimate x_i in the sum is weighted by the inverse of its own variance σ_i^2 . In this way the more reliable estimates contribute more to the overall mean for a given time/space period than do estimates with large variances (standard errors). The weighted mean is the maximum likelihood estimator of the mean of the probability distributions assuming that they are independent and normally distributed with the same mean (Bevington and Robinson, 2003).

Each x_i contributes some uncertainty to the determination of \bar{x}_w . This uncertainty is expressed as the *variance of the weighted mean*, $\sigma_{\bar{x}_w}^2$:

$$\sigma_{\bar{x}_w}^2 = \frac{\sum_{i=1}^n w_i^2 \sigma_i^2}{(\sum_{i=1}^n w_i)^2}.$$

When $w_i = \frac{1}{\sigma_i^2}$ Graybill and Deal (1959) show that \bar{x}_w is the unbiased estimator of the mean with a minimum variance among all the weighted estimates. In this situation the variance of the weighted mean simplifies to:

$$\sigma_{\bar{x}_w}^2 = \frac{1}{\sum_{i=1}^n \frac{1}{\sigma_i^2}}, \quad (6.2)$$

and $\sigma_{\bar{x}_w}$ is known as the *standard error of the mean*.

The spread of estimates x_i is expressed in terms of the unbiased *variance* of the data, $\bar{\sigma}_w^2$:

$$\bar{\sigma}_w^2 = \frac{\sum_{i=1}^n w_i}{(\sum_{i=1}^n w_i)^2 - \sum_{i=1}^n w_i^2} \sum_{i=1}^n w_i (x_i - \bar{x}_w)^2. \quad (6.3)$$

This equation corrects for the fact that \bar{x}_w was itself estimated from the data (Bevington and Robinson, 2003).

$\sigma_{\bar{x}_w}^2$ and $\bar{\sigma}_w^2$ may be viewed as measures of the internal and external variances respectively. $\sigma_{\bar{x}_w}^2$ is an indicator of the models performance and the amount of confidence to be placed in an estimate. $\bar{\sigma}_w^2$ is a quantity more interpretable oceanographically. It gives an indication of the variance and stability of a parameter over time or space.

6.2.2 Gradient and Error Propagation

The thermal gradient (in $^{\circ}\text{C}.\text{km}^{-1}$), a measure of the strength of the Southland Front is calculated by dividing the weighted mean of the temperature change (in $^{\circ}\text{C}$) between subtropical and subantarctic waters for a given area and/or timescale by the corresponding weighted mean estimate of the width (in km's). Error propagation may be used to calculate the variance associated with the resulting gradient. For a general function $x = f(u, v)$ where u , v and x represent the temperature difference, width and gradient respectively, the error propagation equation is:

$$\sigma_x = \sqrt{\sigma_u^2 \left(\frac{\partial x}{\partial u}\right)^2 + \sigma_v^2 \left(\frac{\partial x}{\partial v}\right)^2 + 2\sigma_{uv}^2 \left(\frac{\partial x}{\partial u}\right) \left(\frac{\partial x}{\partial v}\right)}, \quad (6.4)$$

(Bevington and Robinson, 2003)

where

- σ_x , σ_u and σ_v are the standard deviations of the x , u and v measurements,
- $\frac{\partial x}{\partial u}$ and $\frac{\partial x}{\partial v}$ are the partial derivative of x with respect to u and v and,
- σ_{uv}^2 is the covariance between u and v .

For a division of u by v ($x = \frac{u}{v}$), the partial derivatives and their squares are:

$$\begin{aligned}\frac{\partial x}{\partial u} &= \frac{1}{v} & \left(\frac{\partial x}{\partial u}\right)^2 &= \frac{1}{v^2} \\ \frac{\partial x}{\partial v} &= -\frac{u}{v^2} & \left(\frac{\partial x}{\partial v}\right)^2 &= \frac{u^2}{v^4}\end{aligned}$$

Substituting into Equation 6.4 above the variance of x is given by:

$$\sigma_x^2 = \frac{\sigma_u^2}{u^2} + \frac{\sigma_v^2}{v^2} - \frac{2\sigma_{uv}^2}{uv}. \quad (6.5)$$

Assuming that the measurements of u and v are independent, then the associated covariance term (σ_{uv}^2) equals zero and Equation 6.5 reduces to:

$$\sigma_x^2 = \frac{\sigma_u^2}{u^2} + \frac{\sigma_v^2}{v^2}.$$

We use this theory to calculate the error associated with the estimated gradient.

Note that the assumption of independence between u and v (the temperature difference $2\theta_2$, and width $2\theta_3$) is not strictly correct. The two variables are dependent within the model algorithm and also likely to have some dependence physically as a result of the complex dynamics that govern the behaviour of ocean fronts. The latter of these two is difficult to assess. Information on the former however is available from the off-diagonal elements of the estimated variance-covariance matrix as discussed in Section 4.2.3.

6.3 Results

6.3.1 Southland Front Characteristics and Position

The mean position and characteristics of the Southland Front between January 1985 and December 2005 (21 years) are shown in Figures 6.1 and 6.2. From south to north the mean temperature of the front decreases from 10.67°C at 46.4°S to 10.2°C at 45.25°S near Oamaru. It then increases again to 10.67°C as it continues northwards across the Canterbury Bight toward Timaru at 44.58°S. Along its path the difference in temperature between subantarctic (SAW) and subtropical waters (STW) on either side of the front gradually decreases. South of the Dunedin Peninsula the two water masses differ in temperature by just over 1.9°C. Once past Oamaru at 45°S this difference has been reduced to less than 1.7°C, a change of 0.2°C in approximately 200km. The mean width

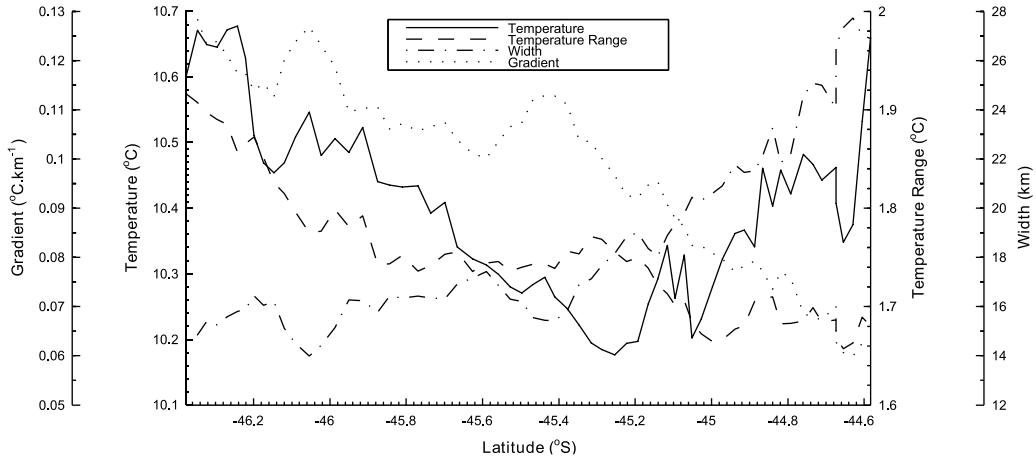


FIGURE 6.1: Along-front changes in temperature, temperature range, width and gradient between January 1985 and December 2005. The latitude refers to the mean position of the front shown in Figure 6.2

of the Southland Front over the twenty-one years increases gradually northwards from 14km to 28km. Combining the weighted mean profiles of temperature range and width reveals a decrease in thermal gradient along the front as it moves equatorward. The front is strongest south of Dunedin reaching $0.128^{\circ}\text{C}.\text{km}^{-1}$ at 46.4°S . It is weakest at 44.65°S where its gradient is $0.06^{\circ}\text{C}.\text{km}^{-1}$.

The mean twenty-one year position of the Southland Front follows the contours of the continental slope (between the 250m and 1000m isobaths) that gradually swing eastward north of 45°S (Figure 6.2). The front sits over a median water depth of 560m. Between Dunedin and Oamaru it runs along the steepest part of the continental slope.

The top panel in Table 6.1 shows the weighted mean twenty-one year estimates of the temperature, width, temperature range and gradient over the length of the Southland Front measured here plus or minus two standard errors ($\bar{x}_w \pm 2\sigma_{\bar{x}_w}$). The weighted standard deviation ($\bar{\sigma}_w$) is also given. The front has a mean temperature of 10.38°C , a temperature range of 1.75°C , a width of 18.0km and a gradient of $0.097^{\circ}\text{C}.\text{km}^{-1}$. The standard errors are all significantly below the spatial resolution of the data set ($0.044^{\circ} \sim 4\text{km}$) and the approximate accuracy of the AVHRR temperature measurements ($0.1 \pm 0.5^{\circ}\text{C}$). This indicates that it is the quality and resolution of data processed that is the main source of uncertainty rather than estimation errors introduced by the algorithm itself.

The seasonal differences in the structure and characteristics of the front can be seen in Figure 6.3 and Table 6.1. The Southland Front is warmest during the summer and coldest in the winter. It may be approximately defined by the 9.36°C , 13.27°C , 11.87°C and 9.12°C contours in the spring, summer, autumn and winter respectively (Table 6.1). The largest standard deviations associated with these weighted mean estimates are for the summer and autumn months. This may partly be attributed to the greater increase

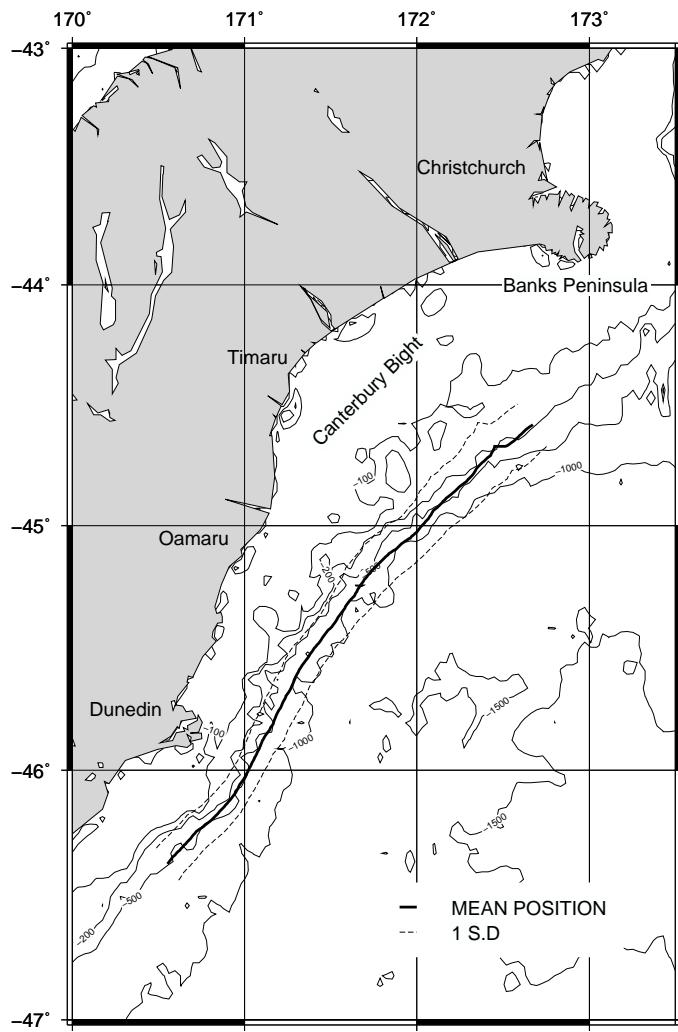


FIGURE 6.2: Mean position of the Southland Front between January 1985 and December 2005

in temperature as the front flows north during the summer. Throughout the winter (JJA) the front warms by 0.45°C between 46.4°S and 44.5°S whereas in the summer (DJF) a more pronounced increase of 1.3°C is observed (Figure 6.3).

Variability of the temperature change across the front for individual seasons is more complex. Overall, the greatest difference between STW and SAW is in the summer (1.99°C) and spring (1.9°C). The smallest temperature change is seen during the winter (1.55°C). The along-front trends of this parameter are not the same for each season. During the spring and winter, the temperature difference across the front gradually decreases as it flows northwards. This decrease is achieved mainly during the first and last 80km. Between approximately 45.57°S and 45°S the temperature difference is reasonably stable at 1.94°C in the spring and 1.53°C in the winter. During the autumn it remains between 1.63°C and 1.76°C until a latitude of 45°S is reached. North of this point it increases to 2.0°C . This sudden change also happens in the summer, although it starts from a local minima in temperature range at a slightly lower latitude of 44.9°S .

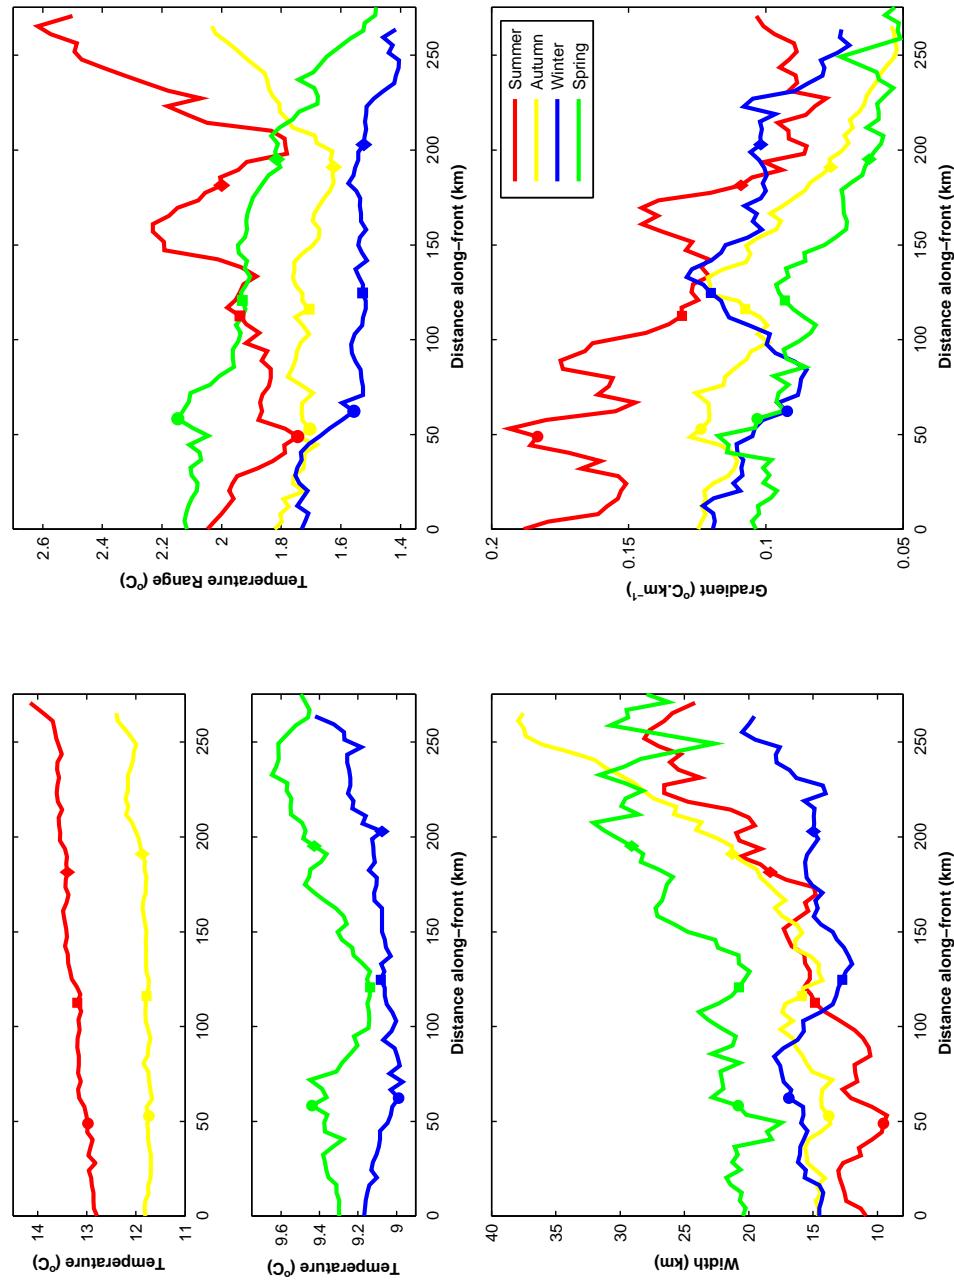


FIGURE 6.3: Seasonal and spatial variability in temperature, temperature range, width and gradient (January 1985 to December 2005). Filled circles, squares and diamonds mark 46°S , -45.5°S and -45°S (Oamaru) respectively along the mean seasonal positions (Figure 6.4). Oamaru marks the northern limit of the submarine canyons and is where isobaths start to diverge.

TABLE 6.1: Characteristics of the Southland Front (January 1985 to December 2005). Weighted mean (\bar{x}_w), $2\times$ standard error of the mean ($2\sigma_{\bar{x}_w}$) and weighted standard deviation ($\bar{\sigma}_w$) for the temperature, temperature difference, width and gradient.

	\bar{x}_w	$\pm 2\sigma_{\bar{x}_w}$	$\bar{\sigma}_w$
21 YEAR			
Temperature (°C)	10.38	0.001	1.80
Temperature Range (°C)	1.75	0.001	0.51
Width (km)	17.99	0.04	10.7
Gradient (°C.km $^{-1}$)	0.097	0.003	
SPRING			
Temperature (°C)	9.36	0.002	0.89
Temperature Range (°C)	1.90	0.002	0.51
Width (km)	23.88	0.09	11.79
Gradient (°C.km $^{-1}$)	0.079	0.004	
SUMMER			
Temperature (°C)	13.27	0.003	1.08
Temperature Range (°C)	1.99	0.003	0.54
Width (km)	15.02	0.1	10.22
Gradient (°C.km $^{-1}$)	0.13	0.007	
AUTUMN			
Temperature (°C)	11.87	0.002	1.14
Temperature Range (°C)	1.74	0.002	0.54
Width (km)	18.02	0.09	10.89
Gradient (°C.km $^{-1}$)	0.097	0.005	
WINTER			
Temperature (°C)	9.12	0.002	0.68
Temperature Range (°C)	1.55	0.002	0.39
Width (km)	15.10	0.08	7.51
Gradient (°C.km $^{-1}$)	0.10	0.005	

The summer increase in temperature difference is also more pronounced; it increases by over 0.8°C in 67km. At the most northern latitude (44.6°S) there is a change in temperature of over 2.6°C across the Southland Front. During the summer there is also a local maxima in the temperature range at 45.2°S (160km), and a second local minima at 46°S .

On average the Southland Front is widest during the spring (23.9km) and narrowest in the winter (15.1km). Locally however there are exceptions to this generalization (Figure 6.3). South of 45.5°S the Southland Front is narrowest during the summer and north of 44.7°S (240km along the profile) the front is widest during the autumn. As for the temperature range different spatial trends in the width are seen between seasons. Generally during the spring, summer and autumn months a northward widening of the front dominates. During the summer this increase occurs gradually from 10.9km in the south to 28.1km in the north. In the autumn the majority of the increase takes place

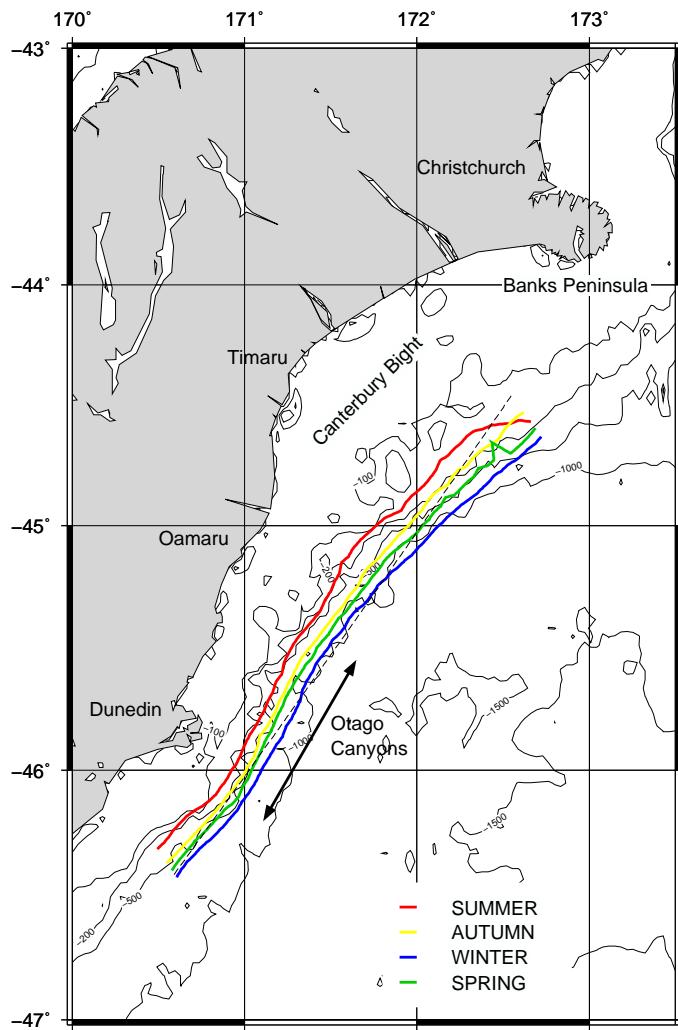


FIGURE 6.4: Mean seasonal positions of the Southland Front between January 1985 and December 2005. The dashed line marks the reference point from which the mean seasonal deviations shown in Table 6.2 were calculated.

very rapidly north of 45.5°S . There is an exponential increase from 14.26km to 38km over a distance of approximately 140km. During the spring most of the northward increase in width takes place between 45.5°S and 45°S (south of Oamaru). In the winter the general widening trend is much weaker. It is interrupted by a definite decrease in width between 45.8°S and 45.4°S .

In general the thermal gradient of the Southland Front decreases equatorward. Taking the weighted mean value for each season over the full twenty-one years and the entire length of the front, the Southland Front is found to be strongest during the summer ($0.13^{\circ}\text{C}.\text{km}^{-1}$) and weakest during the spring ($0.08^{\circ}\text{C}.\text{km}^{-1}$). As before this generalization is modified by focusing on localized areas. South of 45.5°S the front is significantly stronger during the summer ($0.13\text{--}0.19^{\circ}\text{C}.\text{km}^{-1}$). In the spring, autumn and winter its strength over this southern section ranges between $0.08\text{--}0.12^{\circ}\text{C}.\text{km}^{-1}$. North of 45.5°S the front is not always strongest in the summer. The winter gradient, having suddenly

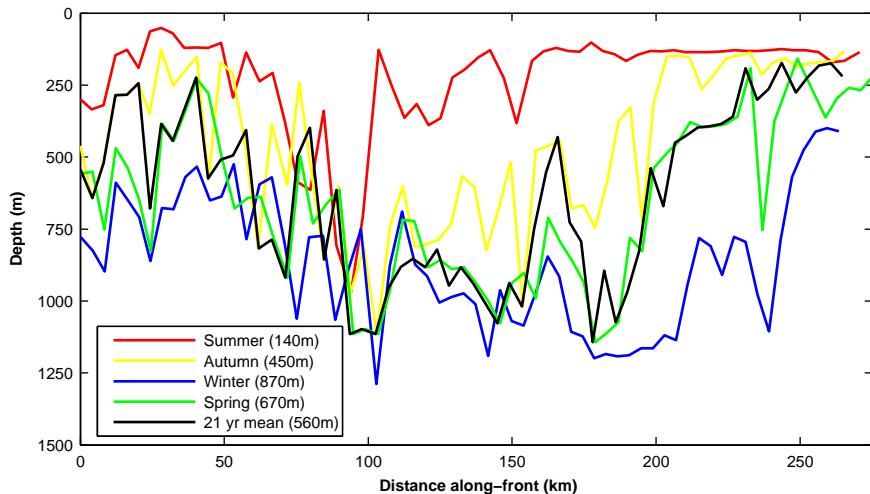


FIGURE 6.5: Mean twenty-one year and seasonal positions of the Southland Front with respect to depth. Bracketed distances in the legend refer to the median depth.

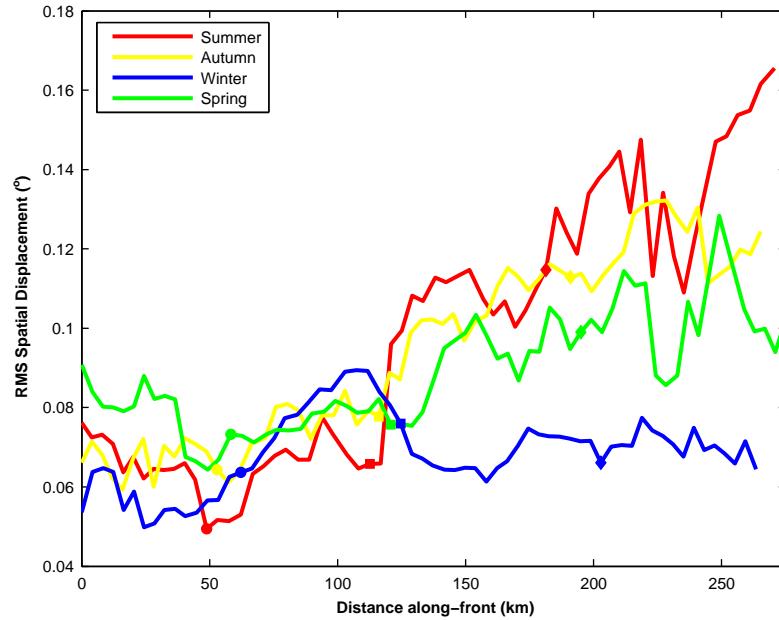
increased from 0.086 to $0.13^{\circ}\text{C}.\text{km}^{-1}$ between 45.82°S and 45.45°S , is greater than the weakening summer gradient over a distance of approximately 50km north of Oamaru. The summer decrease in gradient toward the Chatham Rise is punctuated by three local maxima at 46°S , 45.72°S and 45.1°S .

The seasonal locations of the Southland Front between January 1985 and December 2005 are shown in Figure 6.4. The water depths corresponding to these positions are shown in Figure 6.5. In all four seasons the front follows the continental shelf break. It is located furthest inshore during the summer and moves furthest offshore in the winter. More specifically, it follows the 140m , 450m , 670m and 870m contours in the summer, autumn, spring and winter respectively (median depths).

Variability of the front's position is further quantified in Figure 6.6 and Table 6.2. Firstly, the deviation in location of the front during all four seasons is assessed with respect to one common position. The reference path, a subjective estimate of the mean position was taken to be a line oriented north-east between 46.42°S , 170.59°E and 44.46°S , 172.55°E along the continental shelf break (shown in Figure 6.4). The mean deviation of all seasonal estimates of the front's location is then calculated at right angles to this line. Secondly, in order to quantify variability in the overall mean and seasonal positions the root-mean-squared spatial displacement (RMSSD) of all estimates from the mean paths shown in Figures 6.2 and 6.4 was calculated. This may be interpreted as a measure of meandering intensity (Lee and Cornillon, 1995).

Considering the mean RMSSD shown in Table 6.2 there is little difference in the meandering intensity between the spring, summer and autumn. The lowest meandering intensity is observed during the winter. Figure 6.6 shows how the RMSSD changes along the path of the front during each of the four seasons. During the spring, summer and autumn there is an overall increase in meandering intensity as the Southland Front

FIGURE 6.6: Seasonal root-mean-squared spatial displacement (RMSSD) about the mean seasonal positions of the Southland Front (Figure 6.4). Filled circles, squares and diamonds mark -46°S , -45.5°S and -45°S (Oamaru) respectively.



flows north along the east coast of South Island. In the winter the increasing RMSSD trend does not persist beyond 45.6°S . There is a sudden decrease in meandering intensity at this latitude after which a relatively constant value of 0.07° is maintained.

South of 46°S the highest RMSSD is seen in the spring and there is a slight decrease in RMSSD in both the spring and summer as the front flows toward Dunedin Peninsula. North of 46°S the RMSSD during the spring, summer and autumn begins to increase. Once the front has passed 45.5°S the RMSSD is strongest in the summer and weakest (ignoring the winter values), in the spring.

TABLE 6.2: Variability in the position of the Southland Front (January 1985 to December 2005). Overall mean offset in position of the Southland front (weighted mean deviation) relative to a straight reference path (\bar{x}_w), $2\times$ standard error of the mean ($2\sigma_{\bar{x}_w}$), weighted standard deviation ($\bar{\sigma}_w$), and root-mean-squared spatial displacement (RMSSD) about the mean twenty-one year and seasonal positions. Units are in degrees. Positive \bar{x}_w indicates a more shoreward location. Negative \bar{x}_w are positions seaward of the reference path.

	\bar{x}_w	$\pm 2\sigma_{\bar{x}_w}$	$\bar{\sigma}_w$	RMSSD
21 YEAR	0.0094	2.66E-04	0.11	0.11
SPRING	-0.0047	5.21E-04	0.10	0.090
SUMMER	0.11	6.31E-04	0.10	0.10
AUTUMN	0.036	5.22E-04	0.094	0.097
WINTER	-0.060	4.86E-04	0.090	0.070

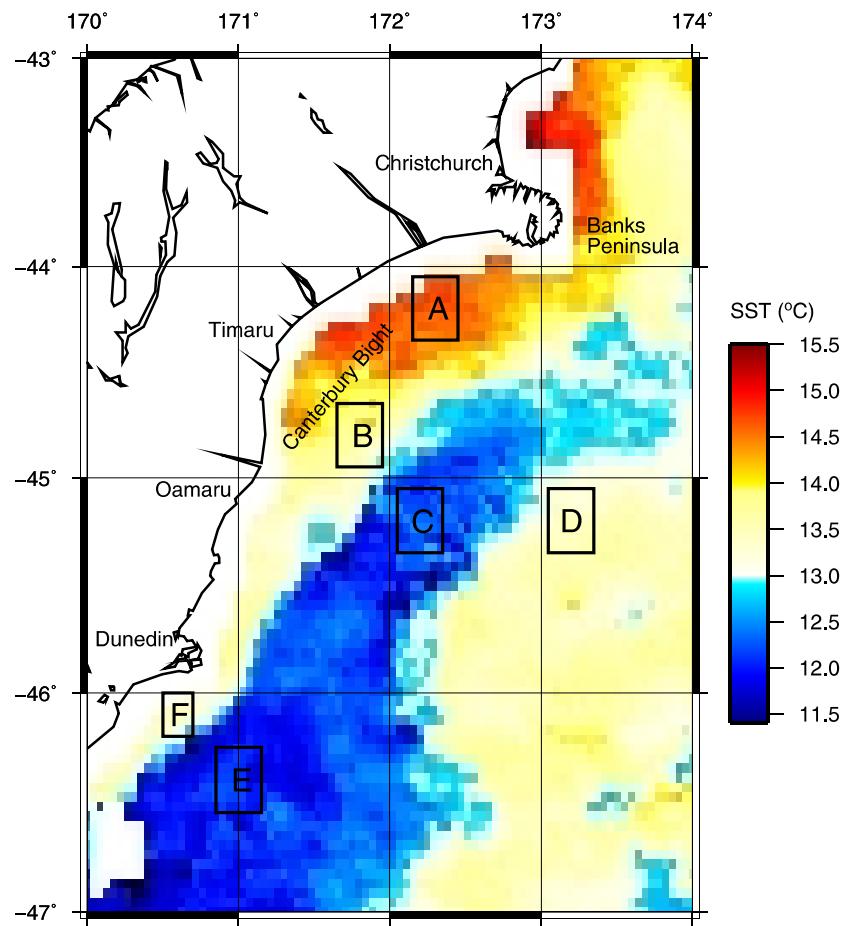


FIGURE 6.7: SST ($^{\circ}\text{C}$) during April 1985 showing; A. Neritic Waters, B. Northern STW, C. Northern SAW, D. Subantarctic Warm Pool, E. Southern SAW, F. Southern STW. The cool subantarctic tongue (E and C) and offshore subantarctic warm pool (D) are clearly visible.

6.3.2 Seasonal SST Cycle

In an attempt to better understand why the temperature difference between SAW and STW varies both spatially and temporally, monthly AVHRR SST data were used to observe the seasonal SST cycles of different water masses in the region between January 1985 and December 2005. These regions are identified in Figure 6.7 and the seasonal SST range in each of them is shown in Table 6.3. The amplitude and phases of each time series were calculated from Fourier power spectra. The SAW's of regions C and E form part of the cold subantarctic tongue described by Shaw (1998). Region D is located in the offshore pool of warm SAW. Its low salinity has shown that it is not of subtropical origin. Both the SAW tongue and warm water pool are more than surface features, they extend down to approximately 100db (Shaw, 1998).

The Neritic (inshore, A) and offshore SAW's (D) have the greatest range in SST over the year in the region. SAW and STW have a similar annual variability but the amplitude of their cycles varies with latitude. In the north the seasonal SST range of SAW and

STW is 5.2°C and 5.1°C respectively while further south this variability has decreased to 4.68°C and 4.48°C . This latitudinal dependence was also noted by Chiswell (1994).

Seasonal temperature changes appear first in the neritic waters and last (8 days later) in the offshore warm pool. This is in agreement with Shaw et al. (1999) who find inshore waters to reach maximum temperatures ten days earlier than offshore SAW. STW temperatures peak 2-3 days earlier than SAW at a similar latitude. For each water mass, the annual signal in the north leads changes in the south by 1-2 days. Within the annual cycle, temperatures peak in February and are at a minimum during August, agreeing with Chiswell (1994) and Heath (1985a). The maximum along-front increase occurs during February whereas in August a constant temperature is maintained south to north.

TABLE 6.3: Range and phase (relative to neritic waters) of the seasonal SST cycle for regions A to F identified in Figure 6.7

Region	Range $^{\circ}\text{C}$	Phase (Days behind A)
A: Northern Neritic (inshore)	6.96	0
B: Northern Subtropical	5.10	3.4
C: Northern Subantarctic	5.20	5.14
D: Offshore Subantarctic Warm Pool	6.00	8.23
E: Southern Subantarctic	4.68	7.16
F: Southern Subtropical	4.48	4.5

6.4 Discussion

An important part of the following discussion will be a comparison between results presented here and the findings of other authors, particularly Shaw and Vennell (2001) whose methods are most similar to our own. Previous studies however have used different spatial and temporal resolution data sets. They have analyzed shorter length time series and arrived at their conclusions using a variety of different techniques such as; Empirical Orthogonal Function (EOF), gradient and harmonic analysis of AVHRR images (Chiswell, 1994; Uddstrom and Oien, 1999; Shaw et al., 1999), *in-situ* measurements (Jillett, 1969; Heath, 1985a; Sutton, 2001, 2003), and specifically designed front detection techniques (Shaw and Vennell, 2001). In the following sections we firstly directly compare the performance of the ‘front following’ algorithm (Shaw and Vennell, 2000a) and likelihood based model using a 1km daily image of the Southland Front. Secondly we consider the effects of using 4km rather than 1km data and examine the amount of structural detail lost to temporal smoothing in a monthly composite. The results presented in Section 6.3 will then be fully discussed.

6.4.1 Comparison to Alternative Techniques

One of the key reasons for focusing on the Southland Front during this thesis was that Shaw and Vennell (2000a) developed their ‘front following’ algorithm (Section 3.4, Chapter 3) for use in this area. This provides an important independent set of estimates with which results from the local likelihood approach adopted here can be compared. In Figure 6.8 the location of the Southland Front on 28th March 1990 as estimated by Shaw and Vennell (2000a) is compared to results obtained from the local likelihood technique using the same 1km data set. Estimates having a quality flag greater than 1 (as defined in Figure 5.6) are highlighted in green. Also plotted are the fronts identified by the Single Image Edge detector (SIED) of Cayula and Cornillon (1992) described in Section 3.2 of Chapter 3. Only those features with a 0.2°C across-front difference are shown. Edge detection using the SIED was kindly performed by Dr. Peter Miller from the Plymouth Marine Laboratory.

The local likelihood and ‘front following’ algorithms produce very similar estimates of the front’s position and both agree with the western most structure picked out by the SIED. The local likelihood estimates highlight slightly more of the mesoscale variability than the ‘front following’ algorithm. This is likely due to the limited resolvable along-front length scale imposed by the 20km wide moving window used by Shaw and Vennell (2000a). There are gaps around 44.5°S and 45.8°S in estimates made by the ‘front following’ algorithm where the routine was unable to identify the front. There are no gaps in the local likelihood estimates although the location of the front over the Dunedin Headland is clearly incorrect. This is identified by the quality flagging system and all estimates made at this location would be eliminated from further analysis.

The SIED helps identify frontal structures further offshore and reveals a possible double structure to the Southland Front. North of Dunedin both the ‘front following’ algorithm and local likelihood estimates most closely follow the more shoreward of the two SIED fronts. This would suggest that the strongest of the two structures is found further west.

Figure 6.9 compares estimates of the temperature, temperature range and width. Slight differences are expected since the estimates of position do not exactly coincide. A series of markers have been placed along the profiles at latitudes of 46°S , 45.5°S , 45°S and 44.4°S to aid comparison. South of 44.8°S where the estimated positions are most closely aligned the two techniques return similar results. Corresponding increases and decreases in frontal characteristics are highlighted by both techniques. For example, a drop in temperature, temperature range and width is recorded south of Dunedin ($<50\text{km}$ distance). Note that no other front detection techniques are capable of providing these details.

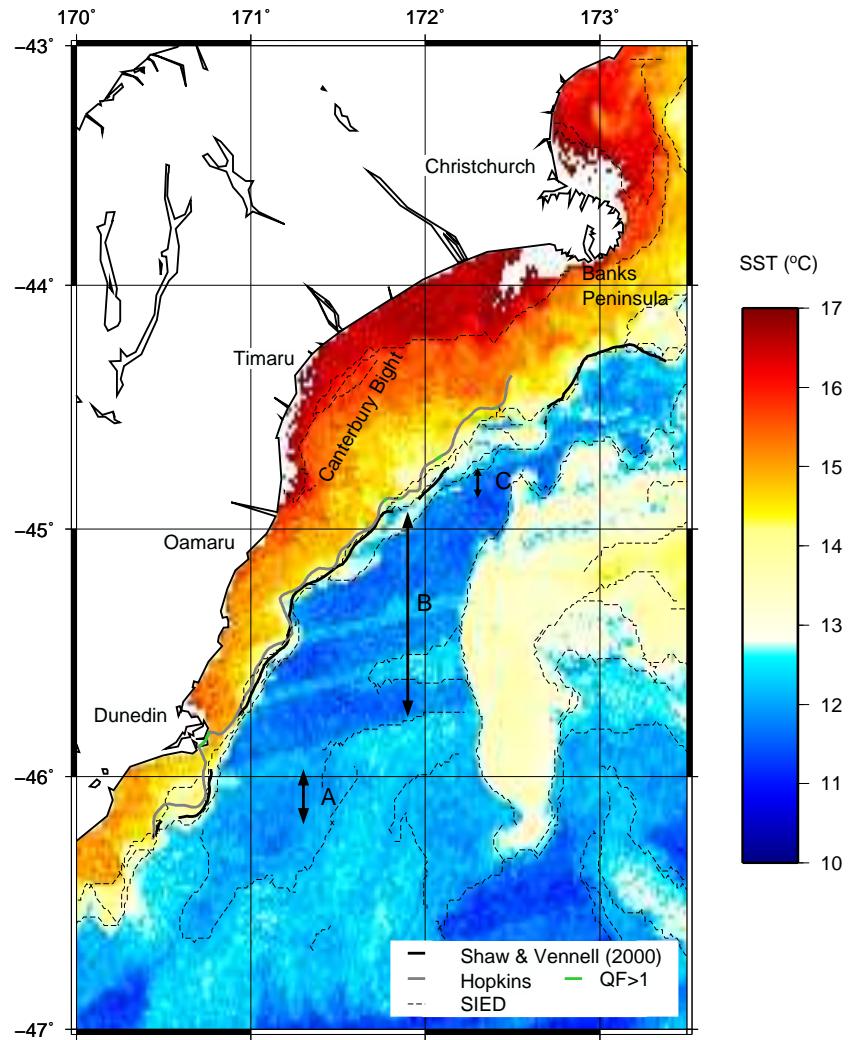


FIGURE 6.8: Location of the Southland Front on 28th March 1990 from Shaw and Vennell (2000a) compared with the estimated position using the local likelihood algorithm where $h_{op} = 0.025^\circ$ (Hopkins). Gaps in the estimates made by Shaw and Vennell (2000a) are where the algorithm failed to converge. Estimates marked in green have quality flags (QF) greater than 1 (see Figure 5.6). Fronts identified using the single image edge detection (SIED) of Cayula and Cornillon (1992) are also shown (minimum cross-front difference of 0.2°C). Color scale in $^\circ\text{C}$. Sections A, B and C refer to estimates in Table 6.4.

Erroneous estimates made around the Dunedin Headland (50-70km along the profile) from the local likelihood algorithm are identified by estimates of the temperature difference being on the upper bound of 6°C set during optimization, and by wider confidence intervals.

Columns a. and b. of Table 6.4 compare the mean parameter estimates made by the 'front following' and local likelihood algorithms from the same 1km image over three discrete sections A-C (shown in Figure 6.8). Identifying latitudinal sections over which *both* algorithms performed successfully is the best way to quantitatively compare the performance of the two techniques. Using the same image estimates of the temperature and temperature range are highly comparable. Estimates of the width however are not in

TABLE 6.4: Mean parameter estimates made over sections A-C on 28th March 1990 (as marked in Figure 6.8) by a. Shaw and Vennell (2000a) from a 1km image, b. the local likelihood algorithm (Hopkins) from the same 1km image ($\pm 2\sigma_{\bar{x}_w}$), and c. the local likelihood algorithm from the equivalent daily 4km image.

Latitude Section	Parameter	a. Shaw (2000)	b. Hopkins	c. Hopkins
		1km image	1km image	4km image
A 46.16 to 45.96°S	Temperature (°C)	13.07	13.47±0.02	12.19±0.10
	Temperature Range (°C)	1.95	2.53±0.02	0.61±0.10
	Width (km)	8.10	16.09±0.60	5.67±6.00
B 45.76 to 44.93°S	Temperature (°C)	13.26	13.41±0.01	13.03±0.03
	Temperature Range (°C)	3.24	3.00±0.01	2.11±0.03
	Width (km)	12.85	9.30±0.21	15.14±1.00
C 44.88 to 44.75°S	Temperature (°C)	13.09	13.39±0.03	13.38±0.06
	Temperature Range (°C)	2.63	2.44±0.03	1.60±0.07
	Width (km)	14.81	9.43±0.80	25.10±3.00

such close agreement. The disparity is most pronounced in section A, south of Dunedin, where the local likelihood algorithm picks out a tight meander in the position of the front (Figure 6.8). Using the ‘front following’ algorithm the width of the Southland Front over this section is estimated to be 8.10km. The local likelihood approach returns a much greater estimate of 16.09 ± 0.6 km. This is likely the result of overestimation due to the front’s orientation (see Section 5.3.3). The moving extraction window used in the ‘front following’ algorithm is aligned normal to the front and therefore this technique does not suffer the same problem. Over sections B and C the local likelihood estimates are 3.5km and 5.4km respectively less than those made by Shaw and Vennell (2000a).

Column c. of Table 6.4 shows the mean characteristics of each section estimated from the equivalent 4km image on 28th March 1990. There are pronounced differences between estimates of all three parameters made from different resolution data sets. For example, a mean temperature difference of 2.53°C is estimated over section A using 1km resolution data. For a 4km resolution image this estimate reduces to 0.61°C . These differences are due to the loss of smaller scale structures and sharp gradients and will be discussed further in Section 6.4.2.

In light of the above comparisons, we can be confident that the local likelihood algorithm is capable of making realistic and reliable estimates of the location, temperature, and temperature difference across ocean fronts. Estimates of the width are the correct order of magnitude but should be interpreted alongside consideration of the front’s orientation. In the next section we explore the effects of using different resolution data sets. In this way we hope to account for any differences between the results of our variability study and those of other authors, particularly Shaw and Vennell (2001) whose technique when compared directly with our own produces very similar estimates.

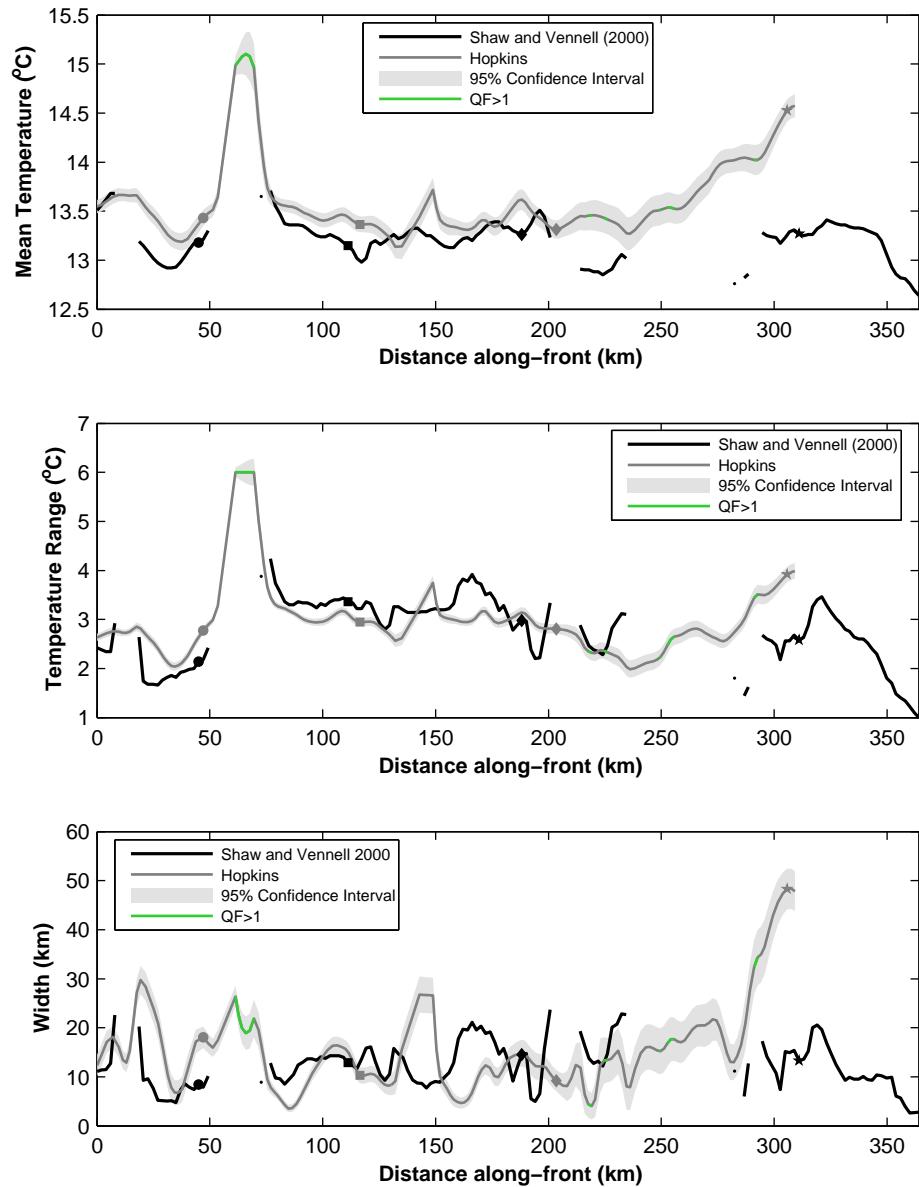


FIGURE 6.9: Estimates of temperature, temperature range and width made by Shaw and Vennell (2000a) on 28th March 1990 compared with estimates made using the local likelihood algorithm ($h_{op} = 0.025^\circ$). Error bars are for the 95% confidence interval. 46°S , 45.5°S , 45°S and 44.4°S latitude are marked by filled circles, squares, diamonds and stars respectively. Estimates with a quality flag (QF) greater than 1 (see Figure 5.6) are highlighted in green.

6.4.2 Data Resolution Considerations

Fronts are dynamic features, continually changing in strength and location in response to the modification of adjacent water masses by non frontal processes, such as air-sea interaction, and across frontal mixing (Belkin and Gordon, 1996). A single cloud free overpass reveals a synoptic snap shot of the structure and intensity of a front, its meanders, eddies and rings. Averaged over time, many of these transient features become blurred or are lost altogether. A composite image can not hope to retain the level of detail and clarity seen in a single overpass. To consider the effects of temporal smoothing in the monthly SST fields used in this study Figure 6.10 compares the location, temperature and strength of estimates made from 4km resolution daily, weekly and monthly SST maps.

The three daily images, taken only a couple of days apart, reveal how quickly the surface temperature field changes and the difference that this makes to the front's estimated location. For example, on 28th March 1990 there is a large warm pool of water in the top right hand corner of the image and the surface expression of the Southland Front is pushed further offshore. On the 2nd April 1990 this warm pool is significantly smaller and the algorithm tracks the front further shoreward. These day to day differences result from the complex vertical structure of the upper ocean. Ocean turbulence and air-sea fluxes of heat, momentum and moisture mean that diurnal temperature variability of surface waters is high. Strong solar heating leads to daytime warming and stratification within a thin surface layer that may mask the true watermass distribution and frontal structure below. The depth to which this layer extends is dependent upon the balance between positive buoyancy provided by incoming solar radiation, and wind driven turbulent energy that mixes warming down through the water column (Stuart-Menteth et al., 2003). Although the effects of diurnal variability have been minimized in our study by using only measurements collected during nighttime overpasses, the data set used by Shaw and Vennell (2001) includes daytime observations.

Much of the smaller scale structures and sharp gradients seen in the daily and weekly images are lost in the monthly climatology through temporal blurring. Consequently, along-front parameter estimates from the monthly SST field are more stable. Additionally, better estimates, particularly of the width, are made around the Dunedin Headland. In daily images the delineation between surface water masses around the headland is weak and noisy, due possibly to a localized increase in turbulence and mixing. This leads to wider confidence intervals around all estimates made in this area. Over a month there is a much clearer distinction between subtropical and subantarctic waters and the front is more confidently located.

The final consideration when comparing findings between authors should be the spatial resolution of observations. Figure 6.11 compares the location of the Southland Front on 28th March 1990 as estimated from 1km and 4km resolution data sets using the

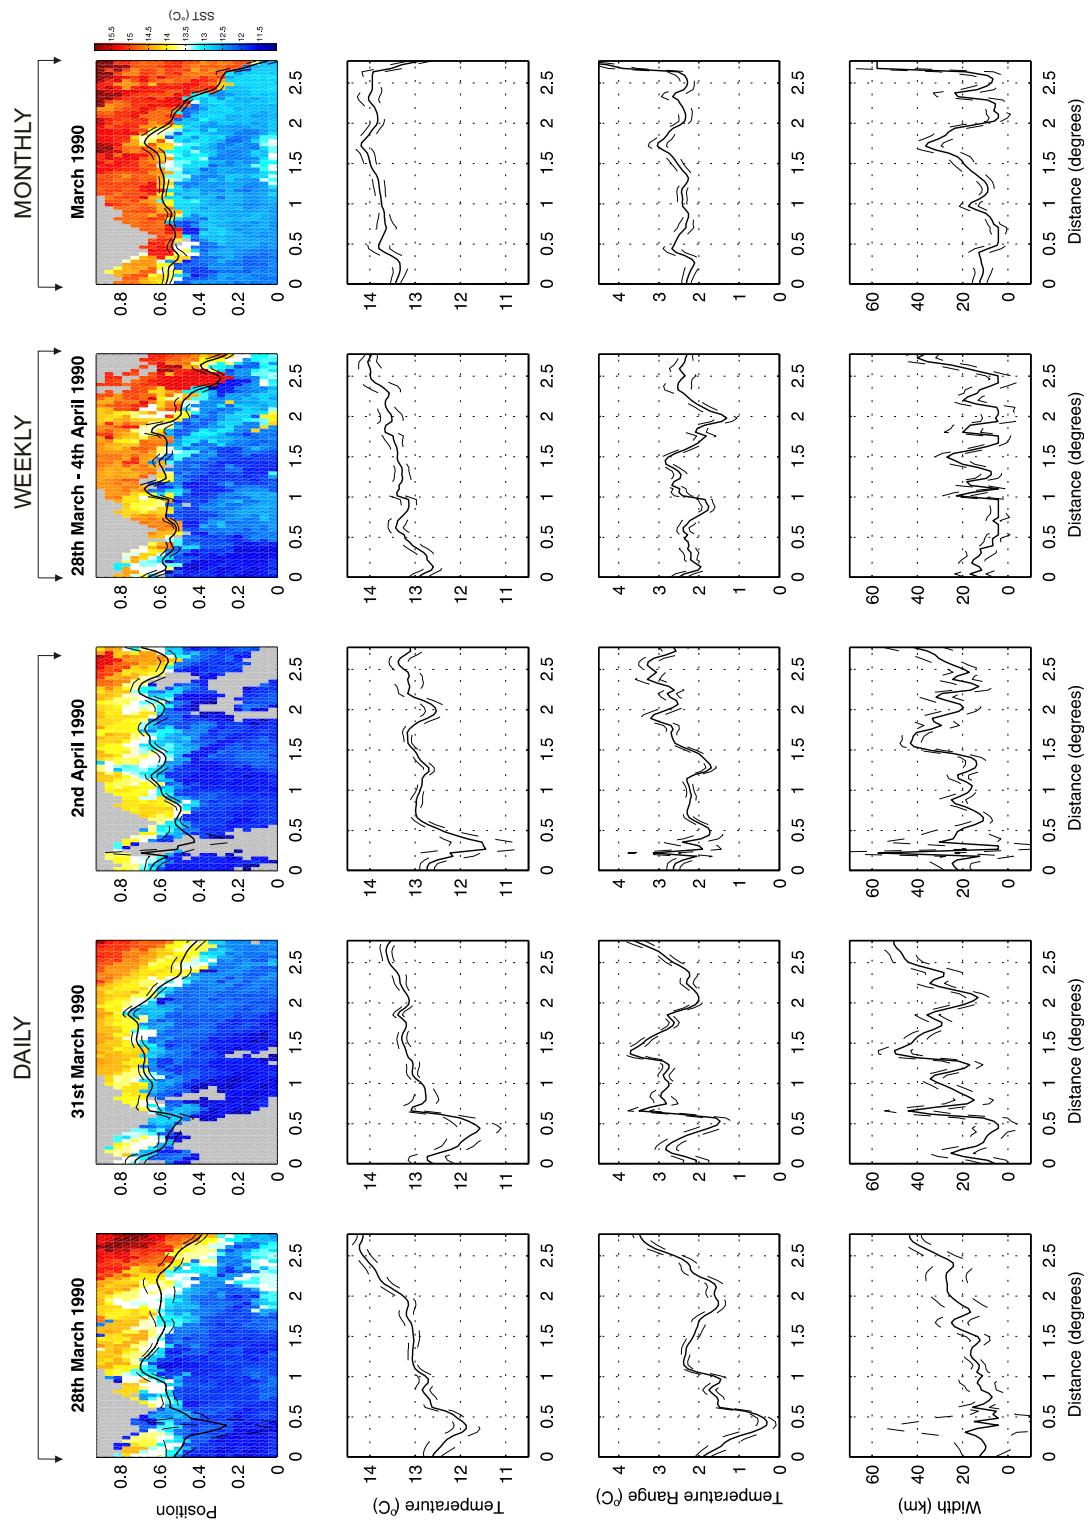


FIGURE 6.10: Estimates made using daily, weekly and monthly images of the Southland Front. 95% confidence intervals shown. The likelihood cross validation optimal bandwidths used as smoothing parameters for the images were $0.065^\circ, 0.07^\circ, 0.06^\circ, 0.05^\circ$ and 0.07° from left to right respectively.

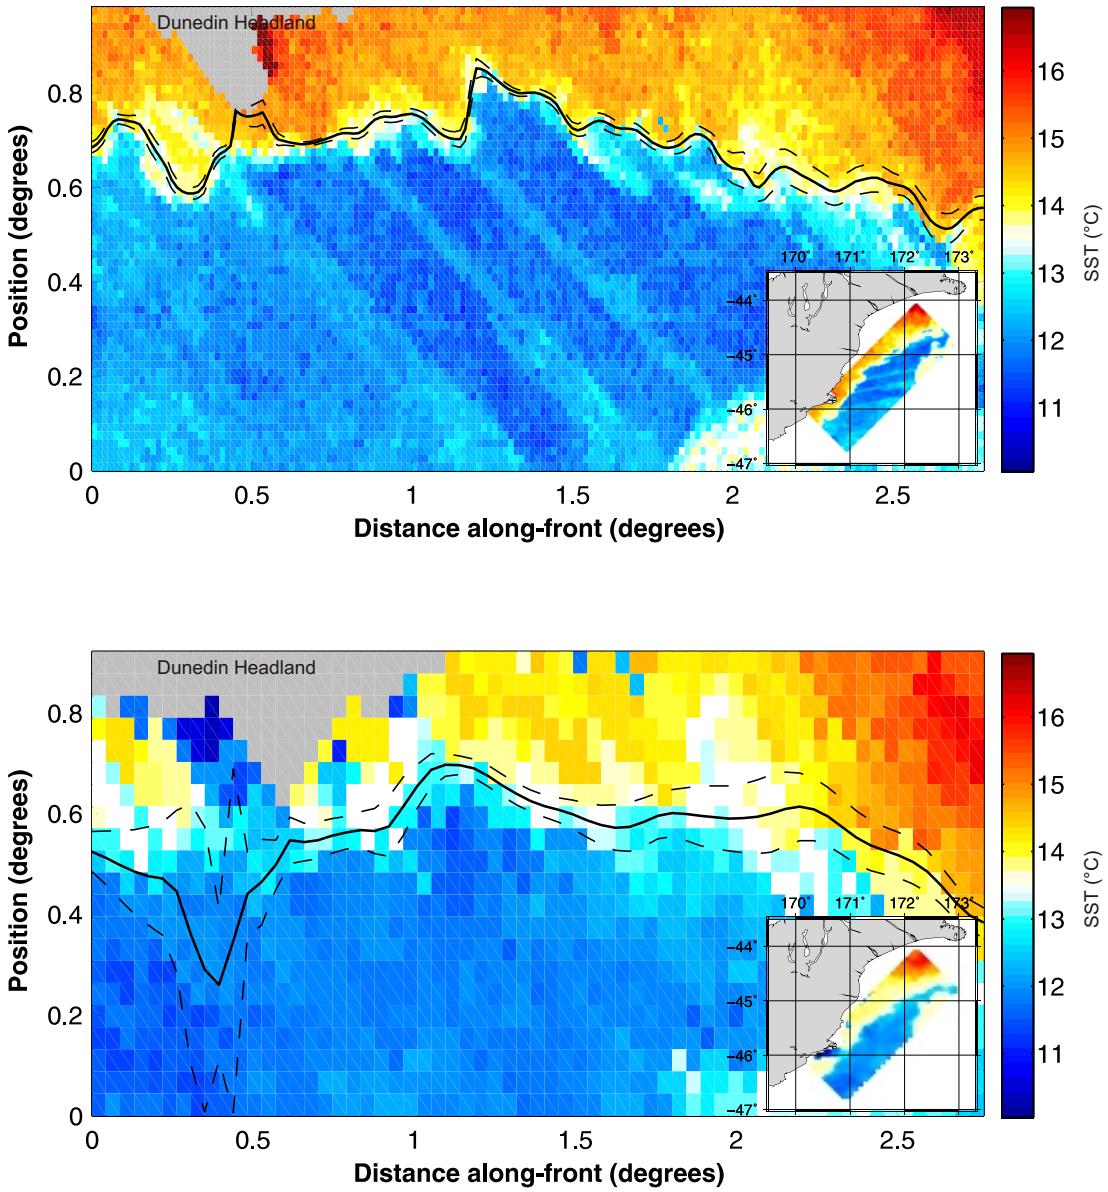


FIGURE 6.11: Local likelihood estimates of the location of the Southland Front on March 28th 1990 using 1km (top) and 4km (bottom) resolution data and optimal bandwidths of 0.025° and 0.065° respectively. Dashed lines mark the 95% confidence bounds.

local likelihood algorithm. There is a marked difference in the level of detail resolved between the two images and consequently a difference in the estimated location. Using the 1km resolution image, the front is located further inshore and a much higher degree of curvature is resolved. Where the thermal surface signature of the front is weak at approximately 0.25 - 0.5° along the profile (south of Dunedin), there is not enough structure in the 4km image to confidently estimate the front's location. This is also reflected in the wider confidence intervals around all other parameter estimates made in this area (Figure 6.12).

Figure 6.12 and column c. of Table 6.4 compare the effect that a coarser resolution data set has on estimates of the temperature, width and temperature range. In the example

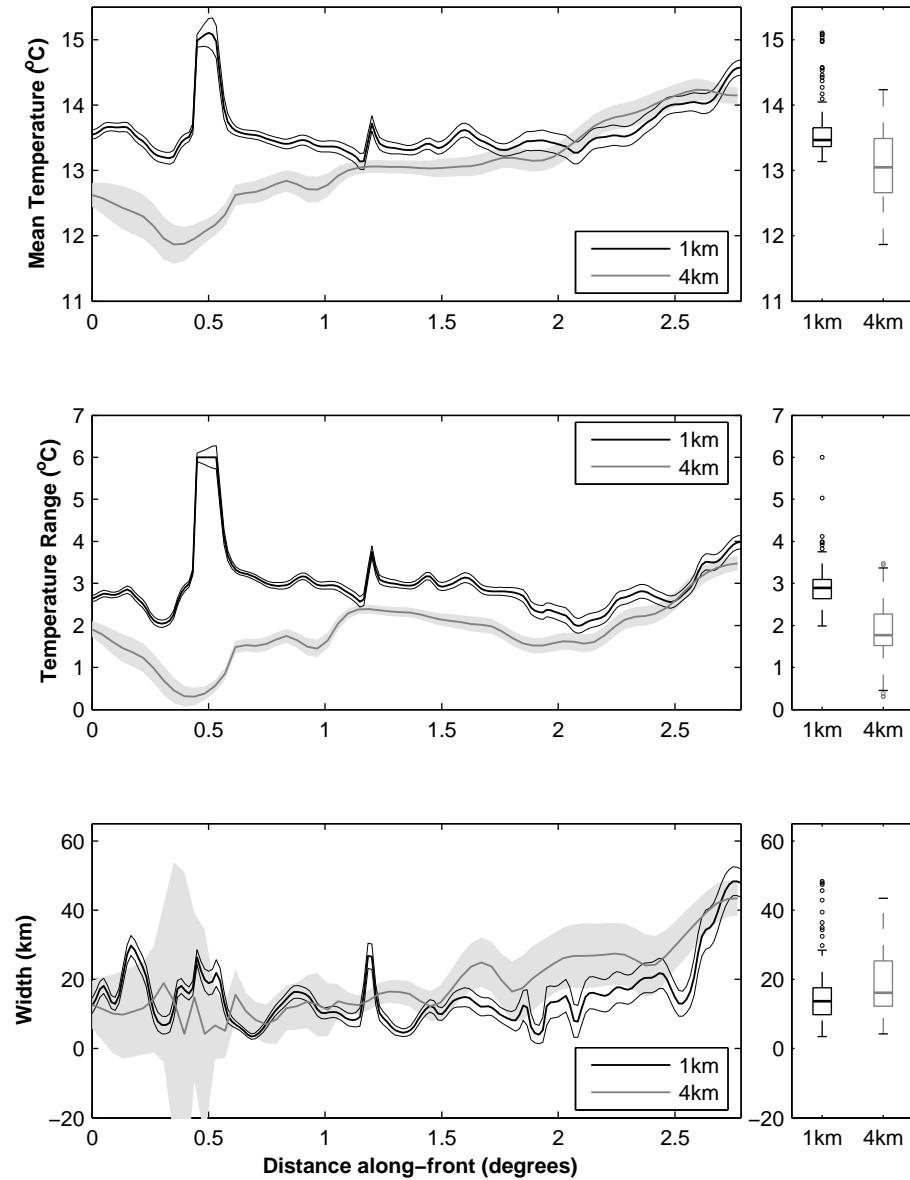


FIGURE 6.12: Local likelihood estimates of the temperature, temperature range and width of the Southland Front on 28th March 1990 using 1km and 4km resolution images. Confidence intervals are for the 95% level. Boxplots on the right indicate the range of estimates made.

image used here, the mean temperature of the front is estimated to be warmer, it has a narrower width and a greater temperature range when SST observations are available every 1km. The boxplots on the right of Figure 6.12 show that a much wider range of estimates is made using a sparser data set.

In each of the three sections A-C defined in Table 6.4 error estimates for the width are an order of magnitude greater when using 4km resolution data. In lower resolution images fewer observations are available to define the transition between water masses. In a 1km resolution image an 8km wide front is defined by eight pixels, enough data to confidently fit the unknown model width. For a feature this size in a 4km resolution

image there would only be two SST observations defining the transition between water masses. This ultimately leads to a less reliable estimate of the width and is reflected in the wider confidence intervals. All other parameters are less sensitive to a sparser data set.

We have demonstrated that the spatial and temporal resolution of an image plays a determining role in the quality and detail of results obtained using the likelihood based front detection algorithm. The intended use of parameter estimates should therefore be considered before a data set is chosen. In light of the above discussion the results presented in Section 6.3 will now be examined.

6.4.3 Temperature

Estimates of the mean and seasonal frontal temperature (Figures 6.1, 6.3 and Table 6.1) are well supported by previous work. Heath (1985a) uses the 15°C isotherm in the summer and the 10°C isotherm in the winter to separate subtropical Southland Current waters from offshore subantarctic waters. Jillett (1969) suggests summer and winter temperatures of 12°C and 9.5°C. Our summer and winter estimates of 13.27°C and 9.12°C fit well with these values. Uddstrom and Oien (1999) use the 10-11°C isotherm to delineate the Southland Current which agrees with the twenty-one year mean estimate of 10.38°C made here.

Comparison with estimates made by the ‘front following’ algorithm (Shaw and Vennell, 2000a) suggests that our results veer toward the lower range of expected temperatures. Taking only those estimates made over the three year period (April 1989–March 1992) studied by Shaw and Vennell (2001), we calculate a mean frontal temperature of 9.93°C (see Table 6.5). This is more than 1°C cooler than the mean temperature calculated by Shaw and Vennell (2001) over the same period of 11.01°C. The same approximate 1°C underestimation is maintained throughout the seasons (Table A.2). This disparity is likely a result of the front being tracked approximately 80km further north-east by Shaw and Vennell (2001) into the warmer waters surrounding Peninsula Banks. The standard deviation of estimates made over the three year period (an indicator of parameter variability) is comparable for both techniques. This is to be expected since the same feature is being examined. The 95% confidence intervals, a measure of how well the model functions are fitted to the AVHRR SST observations (i.e. algorithm performance) are an order of magnitude greater for the ‘front following’ technique. However, the estimation accuracies for both techniques ($\leq 0.01^\circ\text{C}$) are much smaller than the errors introduced by sensor inaccuracies, atmospheric corrections and resolution limitations and are therefore insignificant.

Concurrent with work by Chiswell (1994), the annual seasonal cycle in SST dominates variance in SST along the Southland Front. Using EOF analysis of AVHRR images,

TABLE 6.5: Mean three year (April 1989 - March 1992) characteristics of the Southland Front as calculated by Shaw and Vennell (2001) and by the local likelihood algorithm (Hopkins). The 95% CI's represent the performance of the algorithm used to make the estimate. The standard deviations (s.d) are a measure of the distribution of all estimates made over the three years. The less than sign (<) represents a confidence interval less than 0.005.

	Shaw and Vennell (2001)		Hopkins	
	Mean \pm 95% CI	s.d	Weighted Mean \pm 95% CI	Weighted s.d
Temperature (°C)	11.01 \pm 0.01	2.29	9.93 \pm 0.002	1.78
Temperature Range (°C)	1.76 \pm <0.01	0.68	1.53 \pm 0.003	0.47
Width (km)	8.36 \pm 0.02	4.8	19.20 \pm 0.12	11.06
Gradient (°C.km⁻¹)	0.28 \pm <0.01	0.25	0.08 \pm 0.007	

Chiswell (1994) estimates that the annual cycle accounts for 92% of the variance in SST, but for less of the variance in SST gradients. They find a temperature range of 3.72-6.82°C over the entire New Zealand region and observe a lobe of low amplitude extending into the Southland current of 4.96°C, comparable to our estimates of 5.10°C and 4.48°C for northern and southern STW respectively (Table 6.3).

6.4.4 Temperature Range

A mean temperature difference of 1.75°C across the Southland Front between January 1985 and December 2005 (Table 6.1) is in line with Jillett (1969) who report a jump of 2°C between SAW and STW. Shaw and Vennell (2001) calculate a mean three year temperature range of 1.76°C (Table 6.5). Over the same period we estimate a slightly lower differential of 1.53°C. Once again the standard deviation of estimates (referring to the distribution over the three years) is comparable, but the confidence placed in estimates made by the local likelihood algorithm is an order of magnitude greater than for the ‘front following’ technique.

Contrary to Shaw and Vennell (2001) this study revels clear seasonal differences in the temperature range across the Southland Front on both the twenty-one (Figure 6.3, Table 6.1) and three (Figure A.5, Table A.2) year timescales. Variable along-front trends are also observed between the seasons. Differences in the range and phase of the regional water masses off the east coast of South Island account for these changes.

Inshore waters have a substantial annual SST range of 6.95°C (Table 6.3), and in the summer heat up much more than the already cool SAW on the opposite side of the front. This large fluctuation in temperature over the year may be explained by the limited water depth (\sim 50m south of Banks Peninsula). During the summer, turbulence vertically mixes the heat provided by increased insolation down through the mixed layer, depressing the depth of the thermocline. In limited water depths, where the mixed layer may occupy the entire water column, higher temperatures are reached compared to deep

water offshore sites where the additional heating may be taken up by further expansion of the mixed layer. In the winter, when the surface waters experience intense cooling, convective overturning cells replace the colder and therefore denser sinking surface waters with warmer water from depth. In shallower regions this overturning is interrupted resulting in cooler surface waters than those further offshore.

North of 45°S, particularly during the summer, the front moves into shallower waters of the Canterbury Bight (Figure 6.4) where intense heating south of Banks Peninsula results in maximal temperature differences between the water masses bordering the front. Figure 6.3 confirms that during the summer and autumn months there is a distinct increase in the temperature difference across the front north of 45°S. During the winter more dramatic cooling of inshore waters brings the temperature of SAW and inshore STW's closer together. Through the Canterbury Bight there is a gradual decrease in an already minimal temperature range (Figure 6.3). Additionally, warming (cooling) of STW leading SAW temperature changes by 2-3 days would exaggerate (dampen) the temperature difference during the summer (winter) months.

6.4.5 Width

Estimation of the width has already been shown to be highly sensitive. It is of no surprise therefore that estimates of the width made in this study are contrasting both in magnitude and trend to those made by Shaw and Vennell (2001) from daily 1km resolution images. Three and twenty-one year mean widths of 19.20km and 18.00km respectively are considerably greater than the 8.36km estimated by Shaw and Vennell (2001). This disparity is a function of temporal and spatial blurring that limits the length scale and clarity of features that can be resolved (see Section 6.4.2). Fewer available observations also decreases the confidence of each estimate. Furthermore, overestimation of the width where the front is orientated more than 30° away from the assumed east-west path helps account for the much greater estimate made here. The 'front following' algorithm adjusts the orientation of its extraction window and therefore does not suffer the same problem.

Figures 6.1 and 6.3 show a northward increase in width. This was expected by Shaw and Vennell (2001) but for unknown reasons was not observed. The most noticeable widening occurs during the autumn. Plumes extending seaward from the Canterbury Bight that are most prevalent during the autumn and winter months are likely responsible. They have a mean length and width of 48 ± 23 km and 18 ± 10 km respectively (Shaw, 1998) and are therefore not always well defined in 4km resolution monthly composites leading to overestimation of the width.

6.4.6 Gradient

The mean gradient of the Southland Front over this twenty-one year study was $0.1^{\circ}\text{C}.\text{km}^{-1}$. This is an order of magnitude greater than the estimates of Chiswell (1994) who finds the Southland Front and Subtropical Convergence south of the Chatham Rise to have gradients exceeding $0.002^{\circ}\text{C}.\text{km}^{-1}$, and amplitudes within the Southland Front to reach $0.04^{\circ}\text{C}.\text{km}^{-1}$. On the other hand, Shaw and Vennell (2001) estimate the front to be 2-3 times stronger ($0.28^{\circ}\text{C}.\text{km}^{-1}$). The different resolution data sets used in all three studies is once again responsible for these very different results. Chiswell (1994) use 25-min of arc resolution images ($\approx 46\text{km}$) whereas Shaw and Vennell (2001) analyze a 1km data set. This study, based on $4 \times 4\text{km}$ sized pixels, produces estimates that fall between these two extremes. As explained above, the image resolution places a lower bound on the minimum frontal width detectable and consequently an upper bound on the maximum resolvable gradient.

Overall, the strongest gradients occur during the summer ($0.13^{\circ}\text{C}.\text{km}^{-1}$) and winter ($0.10^{\circ}\text{C}.\text{km}^{-1}$) months, although poleward of 45.5°S the front is much stronger during the summer (Table 6.1 and Figure 6.3). This completely contradicts the strong spring to autumn gradients, and weak winter structure reported by Uddstrom and Oien (1999). The seasonal dependence seen here is most in agreement with Chiswell (1996, 1994) and Shaw and Vennell (2001) who find the Southland Front to be strongest during the winter months.

The majority of previous studies that comment on the gradient of the Southland Front (Chiswell, 1994; Uddstrom and Oien, 1999) have interpreted their results based only on the variability of STW and SAW temperatures. They have not examined spatial and temporal changes in the width of the front that we find here to play an important role. Chiswell (1994, 1996) attribute higher winter gradients to greater annual SST variability of SAW compared to STW within the Southland Current, and the subsequent enhanced cooling offshore during the winter. Although we find STW's to have a marginally smaller SST range (Table 6.3), which is a possible reflection of their non local source (Chiswell, 1996), the 2-3 day phase lag between the two cycles results in the greatest temperature differential occurring in the summer. This is especially true in the Canterbury Bight where intensely warmed inshore waters may extend seaward over the STW of the Southland Current (Jillett, 1969) and exaggerate further the surface temperature difference. The distribution of surface water masses is discussed more thoroughly in the next section. The strong winter gradients across the Canterbury Bight are therefore due to the front being narrowest in this region during the winter rather than disparate water mass temperatures (Figure 6.3). The high southerly summer gradients are not maintained beyond 45.5°S , despite a sharp increase in temperature difference, because of the overall increase in width northward.

We conclude that although the gradient of the Southland Front is modulated by the seasonal SST cycle of surrounding water masses, the stability and therefore the width of the front introduces variability on a more localized scale.

6.4.7 Frontal Position and Stability

The ability of the algorithm to successfully locate the Southland Front was assessed by visual inspection of the front's position in each image and through an analysis of absolute month-to-month differences in estimated position.

Under visual inspection estimates at all locations in 189 out of 252 images (75%) were accepted. In 23 images between two and eight of the most northern points reached the upper bound. This however is easily identified using the quality control checks and these specific points removed from further analysis. The remainder of the estimated positions during these months were all acceptable. A small number of suspect estimates over the Dunedin Headland were made in 14 images although once again flagged by the control criteria. There were 18 images in which, over short distances (\sim 10-40km), the algorithm picked out strong inshore frontal structures rather than the Southland Front. This led to sharp jumps in the estimated position. Owing to the exceptionally high percentage of missing data in 1995 (see Section 5.3.4) estimates of the front's position in February, April, June and August-December during this year are not reliable.

Figure 6.13 shows histograms of the absolute month-to-month differences in estimated position of the Southland Front using: a. all estimates, and b. only those estimates with a quality flag of 1. Both histograms have a narrow distribution suggesting that no major changes in position occur between months. This gives further confidence in the ability of the algorithm to successfully locate the front. Figure 6.13b. demonstrates that the quality flagging scheme identifies those estimates with high (\sim >60km) month-to-month variations in position that are more likely to be the result of estimation error rather than real fluctuations in frontal location.

The results presented in Section 6.3 show that the Southland Front is strongly topographically steered following the 560m isobath along the south-east coast of South Island. The position agrees with Jillett (1969), Heath (1972) and Shaw and Vennell (2001) who all observe the front over the continental slope. Meandering intensity increases as the front flows northward suggesting that topographic control weakens across the Canterbury Bight where isobaths start to diverge. Shaw and Vennell (2001) report a minimum variation in position off Dunedin which is coincident with the minimum meandering intensity observed here just south of the headland (Figure 6.6). A decrease in stability is supported by an observed increase in plume formation (Shaw, 1998), divergence of flow (Shaw and Vennell, 2001) and intermittent throughflow from the Mernoo Saddle (Shaw and Vennell, 2000b). A northward increase in width and decrease in gradient is

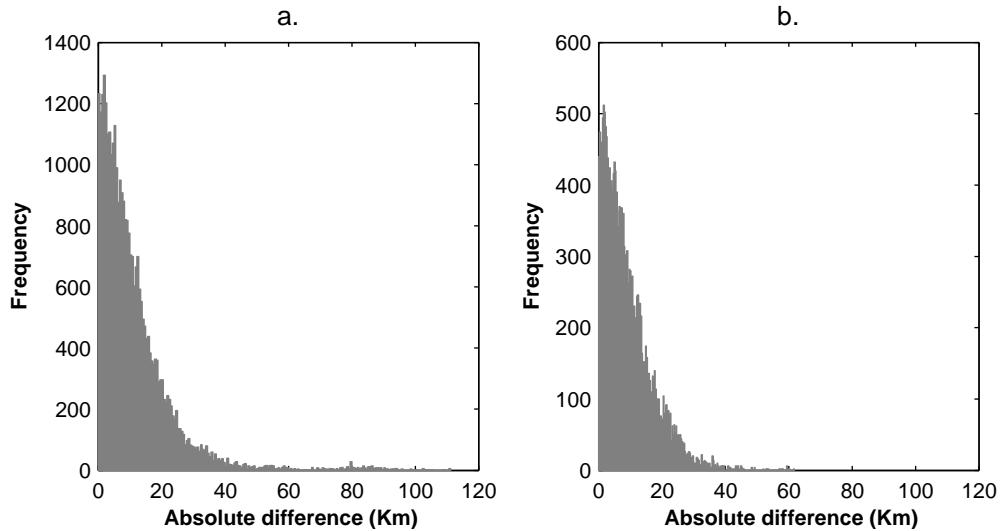


FIGURE 6.13: Histograms of absolute differences in estimated month-to-month positions for a. all estimates, and b. estimates with a quality flag of 1. The median values for a. and b. are 8.3km and 7.3km respectively.

consistent with greater instability and mixing. Rough topography of the Otago Canyons may create instabilities in the flow that propagate and grow downstream contributing to increased meandering intensity as the isobaths diverge through the Canterbury Bight.

The Southland Front is located closest inshore during the summer and farthest offshore in the winter (Figure 6.4). This result is supported by both Shaw and Vennell (2001) and Jillett (1969). The seasonal variation in water masses present in the Mernoo Saddle coincides with the annual cycle in frontal position. A wisp of SAW extending northward through the Mernoo Saddle from the cool subantarctic tongue (identified in Figure 6.7) pushes furthest west (shoreward) during the summer months (Shaw and Vennell, 2000b). In the winter and early spring a southward extension of STW through the saddle prevents the northward intrusion of SAW and directs the Southland Front further offshore.

Seasonal variability in the distribution of surface water masses may be an important factor to consider when interpreting the front's position and seasonal characteristics. Neritic water sometimes extends seawards as a superficial layer over the whole continental shelf (Jillett, 1969). At other times it is barely present. During the spring the Southland Current may often be hidden. Warming and dilution of inshore waters lowers its density so that it spreads seawards above the denser subtropical Southland Current water mass from which it is separated by a sharp salinity difference. At the same time, surface warming of SAW leads to a decrease in its density and movement shoreward over the top of STW's. The Southland Current extends inshore beneath the neritic water, and offshore under the SAW but can not be seen at the surface. A similar situation occurs in the autumn; the salinity of inshore waters is lowered by increased land water runoff and a layer of neritic water spreads seawards over the Southland Current and may reach the shelf edge. At the same time, warm SAW extends inshore to cover the STW's. The

potential masking of Southland Current waters during these seasons introduces some uncertainty into how well the estimated positions represent the subsurface signature of the front. It may also contribute to the large temperature difference during the spring south of 45.5°S (Figure 6.3) and could potentially result in a misrepresentation of the strength of the front.

According to Jillett (1969) the subtropical waters of the Southland Current are more visible during the summer. They cover the continental slope and extend seawards beneath the warm less saline SAW. In the winter, convective overturning breaks down stratification in the upper 100m of the water column. The resulting isothermal and isohaline conditions mean that all surface water masses are visible and more confidence can be placed in estimates made between July and August. Increased vertical mixing over the continental shelf during the winter may help push the front further offshore.

Using the meandering intensity as a measure of stability this study finds the front to be most stable during the winter, especially north of 45.5°S. There is a pronounced decrease in RMSSD (Figure 6.6) at this latitude during the winter months that is coincident with a decrease in width and an increase in gradient (Figure 6.3). These trends contradict the expected response to seasonal prevailing winds. North-easterlies during the summer have been shown to decrease the speed of the Southland Current and Shaw (1998) suggests that this might increase the stability. Instead, we find greater stability during the winter when south-westerly winds dominate and current speeds increase.

The surface thermal expression of a front does not necessarily coincide with the subsurface expression. Previous investigations using *in-situ* CTD observations confirm that the subsurface structure of the Southland Front does follows the continental shelf break (Jillett, 1969; Heath, 1972, 1985a; Shaw, 1998; Morris et al., 2001; Sutton, 2003). Figure 6.14 shows the subsurface structure of the water column across the Southland Front at Oamaru from CTD data collected in April 1993. Also plotted (dashed line) is the average position of the front during this time as estimated by the ‘front following’ algorithm (Shaw, 1998). The position of the front is estimated at the mid-point of the continental slope (~500m) and coincides with changes in the subsurface structure. Isotherms and isopycnals peel steeply upwards from 50-100m in the thermocline and break the surface 50-80km across the profile. Shaw (1998) estimates the front to lie along the shoreward side of this structure (~55km across) i.e. the western edge of the SAW tongue. Further inshore subtropical waters are well mixed all the way to the bottom. During this study *in-situ* observations were not available with which to compare the front’s estimated position (from IR SST images) to the subsurface expression. It is left for future research to validate the algorithm results against CTD measurements.

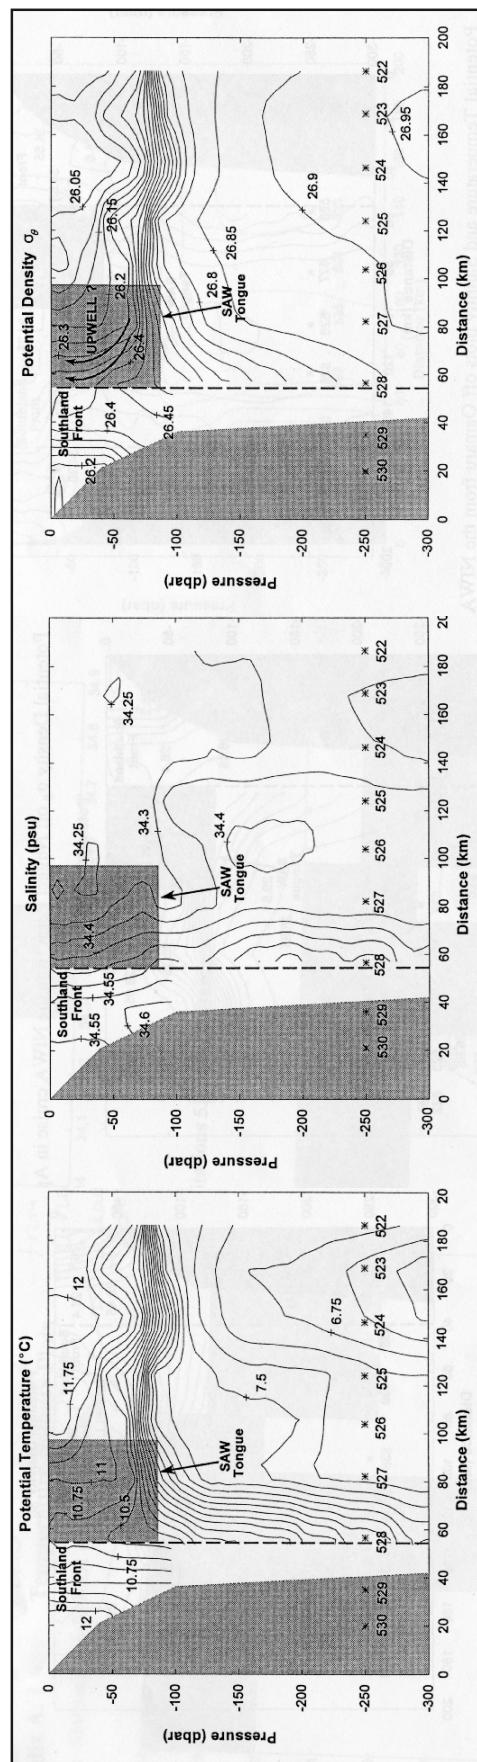


FIGURE 6.14: Potential temperature, salinity and potential density structure off Oamaru from CTD data collected on the NIWA (NZOI) cruise 3006 in April 1993 from the ship *RV Akademik Lавентьев*. The average position of the Southland Front in April 1993 as estimated using the 'front following' algorithm is shown by the vertical dashed line. From Shaw (1998).

6.5 Summary

The new front detection algorithm has allowed us to examine in detail seasonal changes in the position and strength of the Southland Front between January 1985 and December 2005. During this time the front had a mean temperature of 10.4°C and was locked to the bathymetry of the continental slope. The difference in temperature between SAW and STW on either side of the front was 1.75°C and the transition between these two water masses took place on average over 18km. The mean gradient of the front was approximately $0.1^{\circ}\text{C}.\text{km}^{-1}$. Overall, the temperature and width of the front increase northward and the temperature differential and gradient both decrease. This is consistent with an increase in meandering intensity through the Canterbury Bight.

All characteristics of the Southland Front are seasonally and spatially variable. The front is narrowest in the winter and widest in the spring. The temperature difference is maximal during the summer and at a minimum in the winter. The gradient is strongest in the summer south of 45.5°C . To the north the front is equally as strong during the winter. This suggests an increase in induced velocity at the front during summer and winter. The strength of the front is controlled by a combination of the magnitude and phase of the seasonal SST cycles of STW and SAW, and by more local fluctuations in the stability and width of the front, particularly during the winter.

Estimates of the width and gradient are strongly dictated by the sampling resolution. This makes a direct quantitative comparison of these variables between authors and methods difficult. The temperature, position and temperature range however do not appear to be affected by the image resolution and our results are supported by previous research.

The front is located furthest inshore during the summer and toward the seaward edge of the continental slope in the winter. The distribution of water masses and flow reversals in the Mernoo Saddle display a similar seasonal pattern. The stratification of the water column during periods of increased insolation may result in the STW's of the Southland Current that form the western edge of the Southland Front being masked by overlying neritic and SAW's. This highlights the limitations of studying fronts from remote sensing data and stresses the importance of *in-situ* measurements that are able to locate the subsurface expression of these features.

CHAPTER 7

Southland Front-ENSO Teleconnections

7.1 Introduction

In Chapter 6 a detailed description of the seasonal and spatial variability of the Southland Front was presented. In this chapter interannual trends and anomalies in frontal characteristics are investigated in relation to variability of the El Niño-Southern Oscillation (ENSO).

A brief overview of ENSO is given together with a summary of how the New Zealand climate is affected by extreme El Niño and La Niña events. The main data analysis techniques (cross correlation and wavelet transform) are discussed and results relating to each of the main scientific objectives are presented. Owing to the short time series the physical mechanisms proposed to explain the teleconnections observed here should be regarded as speculative.

7.1.1 The El Niño-Southern Oscillation

The El Niño-Southern Oscillation is the dominant mode of coupled interannual ocean-atmosphere variability on the planet. Although it is most apparent in the tropics its signal appears in climatic and oceanic records around the globe. El Niño refers to the warm oceanic component of the phenomena (La Niña to the cold phase). The atmospheric counterpart is the Southern Oscillation measuring the ‘seesaw’ in sea-level pressure (SLP) between the western and south-eastern tropical Pacific. During normal conditions in the tropical Pacific, westward blowing trade winds cause warm surface

water to accumulate west of the date line forming the western Pacific warm pool. The thermocline across the Pacific therefore tilts upward toward the eastern side of the basin allowing cold bottom waters to upwell along the coasts of Peru and Ecuador. Associated with the warm pool is intensive atmospheric convection, whereby air rises and flows aloft toward the eastern Pacific where it sinks over the cold water to close the zonal circulation cell known as the Walker Circulation.

During an El Niño event, the warm phase of ENSO, an unusually high SLP develops in the western tropical Pacific and anomalously low pressure is observed in the south-east. This results in a weakening and reversal of the easterly trade winds. Since the higher sea-level in the western Pacific can no longer be maintained against weaker trade winds, the warm pool migrates eastward across the Equator and upwelling in the east is reduced. At the same time atmospheric convection shifts from the western to the central Pacific. The diminished east-west contrast in SST enhances the anomalous pressure gradient which in turn feeds back into a large scale weakening of the trade winds within which westerly wind bursts are embedded. This drives additional surface waters eastward amplifying the effect. This process is known as positive Bjerknes feedback (Bjerknes, 1969) and can keep El Niño episodes alive for the better part of a year. A La Niña is the opposite of an El Niño. During this cold ENSO phase the easterly trade winds strengthen and intensify the upwelling of cold water in the east and further depress the thermocline in the west.

The evolution and strength of an ENSO event is quantified both in terms of SST anomalies across the equatorial Pacific and by the Southern Oscillation Index (SOI), the difference in SLP between Tahiti and Darwin, Australia, two stations located near the centres of action of the Southern Oscillation (see Section 7.2 for further details). El Niño episodes differ from one another not only in their relative strengths, but also their season of onset, maturity, and demise, as well as the location of their maximum SST anomaly within the tropical Pacific (Wang and Weisberg, 2000; McPhaden, 1999; Trenberth and Tepaniak, 2001).

The mechanisms triggering the onset of an El Niño are not yet fully understood (Philander and Fedorov, 2003). According to the delayed oscillator theory (Bjerknes, 1969) the evolution of the climate system in the tropical Pacific is governed by the interplay between large-scale equatorial wave processes and atmospheric feedbacks. A build up of heat content in the western Pacific, mediated by trade wind-forced downwelling equatorial Rossby waves, is a precursor to El Niño. Reflection of Rossby waves off the western boundary can initiate El Niño events by generating downwelling equatorial Kelvin waves that propagate eastward to depress the thermocline along the coast of South America and cause warming in the eastern cold tongue. Westerly wind bursts associated with the Madden-Julian Oscillation (MJO) may also have a role to play in the initiation of an El Niño event (McPhaden, 1999; Edwards et al., 2006). Before the 1997 El Niño stronger than normal trades built up the heat content in the western Pacific warm pool.

The event did not start however until the intensification of the MJO in late 1996 which was the catalyst for eastward propagating downwelling Kelvin waves; reflecting Rossby waves were not so apparent (McPhaden, 1999).

7.1.2 ENSO: a New Zealand Perspective

The local climatic response to the ENSO around New Zealand is well documented. Spatially averaged air temperatures over New Zealand are positively correlated with the SOI (Gordon, 1986; Mullan, 1998). During an El Niño event when the SOI is negative, changing SLP patterns around New Zealand result in anomalous south-westerly flow that introduces colder air from the Antarctic. The positive phase of the SOI (La Niña) enhances north-easterly flows against a mean background westerly circulation bringing warmer subtropical air southward toward New Zealand (Gordon, 1986). In general agreement with these anomalous flows, annual rainfall in the east of the North Island is positively correlated with the SOI i.e. drier conditions during an El Niño. Rainfall in the south of South Island is negatively correlated i.e. wetter conditions during El Niño (Salinger and Mullan, 1999). There is a marked seasonal variation in the pattern of correlation between pressure in the Australasian-New Zealand region and the SOI (Gordon, 1986).

New Zealand SST anomalies are also found to be positively correlated with the SOI (Mullan, 1998). Colder than usual SST's are observed, particularly around North Island, during El Niño episodes and warm SST anomalies occur during La Niña (Gordon, 1986; Greig et al., 1988; Mullan, 1998; Shaw et al., 1999). New Zealand air temperatures lead SST's by approximately 0.5-1 month (Basher and Thompson, 1996) suggesting that SST anomalies can be explained in terms of the atmospheric circulation forcing oceanographic anomalies through a combination of advection and heat fluxes. More northerly sector airflow during positive SOI events advects warmer surface water and air southward. Heat fluxes from the ocean to the atmosphere are reduced because of the warmer overlying air and generally lighter winds resulting in higher SST's. Conversely, during El Niño stronger and colder southerly winds blow cold water into the region and ocean to atmosphere heat fluxes increase. The air temperature anomalies and fluctuations are generally greater than those seen in the SST (Greig et al., 1988). There is both a seasonal and regional dependence of the significance of correlations between SST anomalies around New Zealand's North and South Islands and the atmospheric circulation patterns associated with ENSO (Mullan, 1998).

There is also a positive correlation between the SOI and sea level around New Zealand. El Niño events result in a drop in sea level and *vice versa* for La Niña periods (Goring and Bell, 1999). For northern New Zealand the fall in sea level during an El Niño is derived mainly from a decrease in regional SST above the thermocline; similarly, the rise in sea level during La Niña is primarily the result of a rise in regional heat content.

7.1.3 The Southland Front and ENSO: The Questions

To date there has been very little investigation into the interannual variability of the Southland Front (SF) and to how it may respond to changes in atmospheric circulation. Using their ‘front following’ algorithm Shaw and Vennell (2001) record a significant decrease in the mean temperature at the interface of the SF between April 1989 and March 1992. This cooling agrees with a 10-year fitted AVHRR time series in the same region showing similar effects due to the 1991 El Niño event (Shaw et al., 1999). In addition, a decrease in gradient was noted over the three years, coincident with the decrease in SOI. Given the direct relationship between SST gradients and the velocity of induced geostrophic currents at large scale fronts, Shaw and Vennell (2001) suggest that there was a decrease in velocity of the Southland Front as a consequence of the El Niño event. This conclusion has never been substantiated with a longer time series.

In this chapter we propose to investigate further the response of the Southland Front to ENSO variability. More specifically we ask:

- **How does the temperature of the Southland Front respond to oceanic and climatic changes induced by ENSO?**
- **Is the strength of the Southland Front modulated by ENSO?**
- **Does the position of the front change during ENSO events?**

Each of these key areas is addressed in turn using a variety of spectral analysis techniques from cross correlation to wavelet coherence and power.

7.2 Data Sets

A combination of atmospheric and oceanic indices are used as a measure of ENSO variability: the Southern Oscillation Index (SOI), Niño3 SST, Niño3.4 SST and Niño4 SST. Additional data sets used in the discussion are described where appropriate.

a *Southern Oscillation Index (SOI)*

The Southern Oscillation (or atmospheric) component of ENSO is measured by the SOI. It is defined as the anomalous SLP in the eastern Pacific at Tahiti (17.6°S,149.6°W) minus the SLP in the western Pacific at Darwin, Australia (12.4°S,130.9°E). These two centres of action are marked in Figure 7.1. The anomalies are departures from the 1951-1980 base period. Monthly values of the SOI were downloaded from the NOAA Climate Prediction Center between January 1985 and December 2005 inclusive (<http://www.cpc.ncep.noaa.gov/data/indices/>).

b Regional Niño SST Anomalies

The El Niño (or oceanic) component of ENSO is measured by the Niño3, Niño3.4 and Niño4 SST indices. These consist of the area averaged SST over the eastern (5°N - 5°S , 150°W - 90°W), east central (5°N - 5°S , 170°W - 120°W) and central (5°N - 5°S , 160°W - 150°W) tropical Pacific. Each Niño SST Region is defined in Figure 7.1. Monthly values were downloaded from the NOAA Climate Prediction Center.

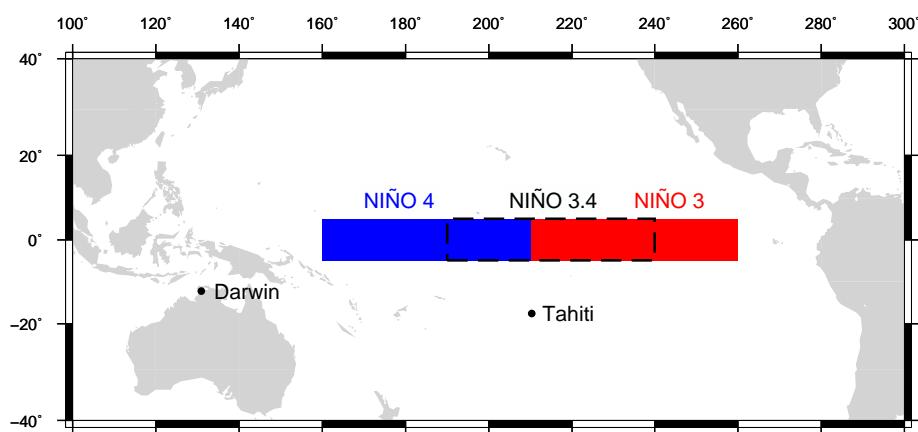


FIGURE 7.1: Niño SST Regions 3, 4 and 3.4 and the stations measuring SLP used to calculate the SOI.

A twenty-one year monthly time series of the temperature, gradient and position of the Southland Front was created by calculating weighted mean local likelihood estimates for each month. As in Chapter 6 the inverse of the standard error was used to weight each estimate. Only estimates receiving the highest quality flag of 1 (see Figure 5.6) were used. The annual cycle was removed by subtraction of the 1985-2005 monthly mean for each calendar month. Any linear trend was also removed from each time series in order to make them stationary, an assumption made by the techniques to be applied. Weighted spatially averaged monthly anomalies of each parameter were left behind. Only very small linear trends were found suggesting that no global warming signal was present.

In order to reduce dependency between data points, an autoregressive (AR) model is fitted to each anomaly time series leaving behind only the white noise residuals (Priestley, 1992). This technique, known as prewhitening, removes any autocorrelation between samples making each point statistically independent from all others and allowing time series to be repeatedly cross correlated. Autocorrelation in individual time series (e.g. SST) where there is likely to be considerable persistence between months can introduce non-zero cross correlations at non-zero lags, even if there is no real interrelationship present. Prewhitening results in correlations that are generally smaller in magnitude than those produced from the original series but, since the residual series have near zero-persistence, the full number of data points is used to estimate significance levels from a Student's *t*-test. Details of autoregressive modelling may be found in Appendix A.5.

7.3 Methods

The relationship between ENSO and variability of the Southland Front is investigated using cross correlation and wavelet analysis techniques. Unlike traditional Fourier analysis (and cross correlation) wavelets assume that the time series in question are stationary i.e. the amplitudes and phases of the different harmonic constituents of the signal do not change in space or time. When dealing with atmospheric, oceanic or climatic signals this is an invalid assumption to make. Such phenomena vary as the local conditions change (e.g. wind stress, density profiles, sea level pressure and bathymetry etc.). Wavelet analysis is a technique that expands the time series into time-frequency space and therefore yields localized estimates for the amplitude and phase of each spectral component of the data set. Where a Fourier transform would smear out any detailed information on the changing processes, wavelet analysis allows local variations in power of the dominant modes of variability to be studied. We use the wavelet power spectrum to find localized intermittent variations in power of ENSO events, as characterized by the SOI and regional Niño SST indices, and in the strength, temperature and position of the Southland Front. More importantly, we calculate the cross wavelet power and coherence between pairs of time series to examine any links by physical atmospheric or oceanic processes (teleconnections). Regions in time-frequency space with large common power or coherence and with a consistent phase relationship suggest causality between two series. The significance of any relationship is tested against a background of red or white noise. A white noise signal is that of a purely random process that consists of a sequence of uncorrelated random variables. It has a flat power spectrum i.e. it has the same power at each frequency. On the other hand, geophysical time series are often referred to as having a red noise spectrum where the energy in the series grows with increasing period.

In the following section we provide some of the mathematical details behind cross correlation and wavelet transform.

7.3.1 Cross Correlation

Cross correlation is used to investigate any persistent relationships (teleconnections) between the position, strength and temperature of the Southland Front, and modes of ENSO variability.

The *cross-covariance* of two signals g and h at a lag of time τ is defined as:

$$C_{gh}(\tau) = g(-\tau) \otimes h(\tau) = \int_{-\infty}^{\infty} g(t)h(\tau + t) dt, \quad (7.1)$$

where \otimes is the convolution, an integral which expresses the amount of overlap of one function h as it is shifted over another function g - as a function of the shift τ . The

cross-covariance therefore is a measure of how similar the two signals are as they are translated past each other. A high covariance is suggestive of a possible relationship between the two series. Convolving a function g with itself is known as *autocovariance*:

$$C_{gg}(\tau) = \int_{-\infty}^{\infty} g(t)g(\tau + t) dt. \quad (7.2)$$

Dividing the cross-covariance function (Equation 7.1) by the geometric mean of the zero lag autocovariances gives the *cross-correlation coefficient function* (or *normalized cross-covariance function*), $\hat{C}_{gh}(\tau)$:

$$\hat{C}_{gh}(\tau) = \frac{C_{gh}(\tau)}{\sqrt{C_{gg}(0)C_{hh}(0)}}, \quad (7.3)$$

where $C_{gg}(0) = \sigma_g^2$ and $C_{hh}(0) = \sigma_h^2$ (i.e. the variances of the two signals). Correlation is thus a scaled version of the covariance with maximum and minimum values of 1 and -1 respectively.

Figure 7.2 shows the cross correlation between the SOI and the frontal temperature for $\tau = \pm 150$ months. The dashed lines mark the 95% confidence levels for a two tailed Student's t -test. The effective number of degrees of freedom (N^*) for the t -test is calculated using information on the autocorrelation of each time series.

$$N^* = \frac{N}{\sum_{-\infty}^{\infty} \hat{C}_{gg}(\tau) \hat{C}_{hh}(\tau)}, \quad (7.4)$$

(Emery and Thomson, 1998)

where N are the number of data points being used. $\hat{C}_{gg}(\tau)$ and $\hat{C}_{hh}(\tau)$ are the normalized autocovariances (or autocorrelations) where $\hat{C}_{gg}(\tau) = \frac{C_{gg}(\tau)}{C_{gg}(0)}$. When all points in the series are neither cross -nor serially correlated the effective number of degrees of freedom is N . Where the series are correlated $N^* \ll N$. Greater correlation leads to a decrease in the number of independent values in the series and a smaller N^* . Prewhitening the times series removes any large autocorrelations and therefore increases the number of effective degrees of freedom for each cross correlation. This increases the accuracy of the correlation coefficients. The confidence levels in Figure 7.2 are therefore much smaller than they would be for a cross correlation between the original (non prewhitened) data sets. Note how the confidence level increases toward higher lags as the sample size decreases.

Interpretation of the lags identified as significant in Figure 7.2 should be approached with caution. An understanding of what a 95% (or 99% etc.) significance level means is important.

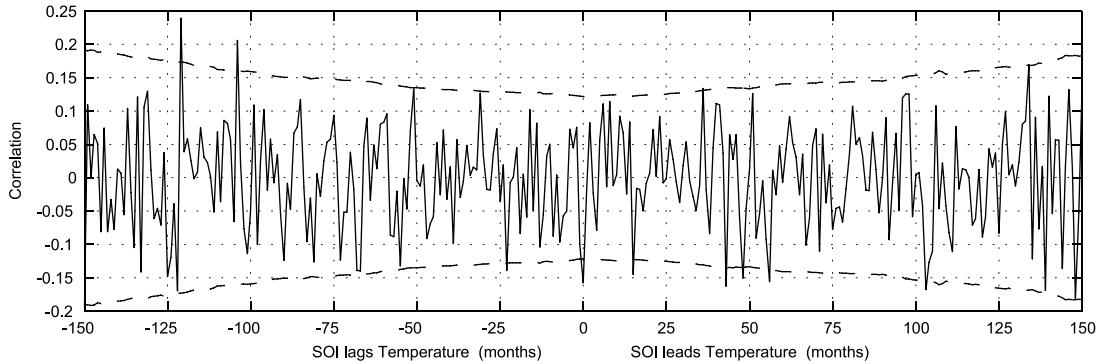


FIGURE 7.2: Cross correlation between the prewhitened SOI and mean temperature of the Southland Front for a lag of ± 150 months. Dashed lines mark the 95% confidence interval.

7.3.1.1 Significance Levels

The significance level is a fixed probability of wrongly rejecting the null hypothesis of ‘*no correlation between two time series*’, if it is in fact true. The aim is to make the significance level as small as possible to protect the null hypothesis and to prevent, as far as possible, false claims being inadvertently made. Testing for significance at the 95% level we would expect to erroneously identify a significant correlation for every 5 in 100 tests. Considering the cross correlation between the estimated temperature of the Southland Front and the SOI in Figure 7.2 where 301 different tests are performed (-150 to 150 months), we would expect to find, even if the two series were in no way correlated, approximately fifteen significant correlations. Therefore even though twelve lags are identified as significant in Figure 7.2 it is still a possibility that no real teleconnection between the two series exists.

Adjustment of the alpha value ($\alpha = 0.05$ for a 95% significance level) to a stricter threshold can be made to safeguard against multiple tests of statistical significance on the same data set falsely giving the appearance of significance. This is known as the Bonferroni Correction but is seldom used in oceanography and will not be applied here. Instead, to guard against over interpretation of cross correlation results times series are split up into shorter sub-series. If the null hypothesis is repeatedly rejected at the same lags for each of the sub-series then confidence in a real and persistent teleconnection having been identified increases.

7.3.2 Wavelet Analysis

Wavelet analysis involves the convolution of a real time series x_t ($t = 1, \dots, N$), with uniform time steps δt , with a set of wavelet functions that are derived through translation, scaling and normalization of a ‘mother wavelet’. The Morlet Wavelet, $\psi_0(t)$, a function with zero mean that is localized in both frequency and time, is the standard

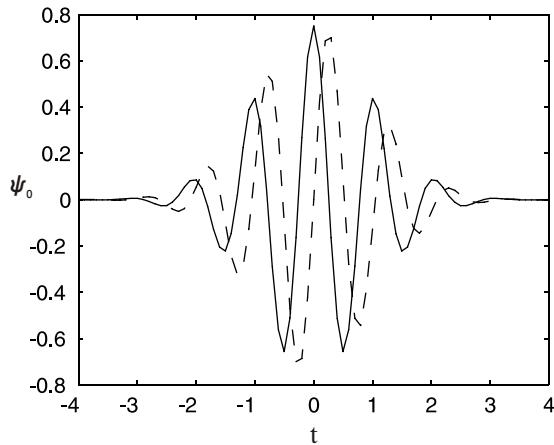


FIGURE 7.3: Real (solid) and imaginary (dashed) parts of the Morlet Wavelet plotted in the time domain. $\omega_0 = 6$.

‘mother wavelet’ used in the analysis of geophysical time series and is the base function used here. The Morlet Wavelet (Figure 7.3) consists of a plane wave of dimensionless frequency ω_0 modulated by a Gaussian of unit width (Emery and Thomson, 1998):

$$\psi_0(t) = e^{i\omega_0 t} e^{-\frac{1}{2}t^2}. \quad (7.5)$$

The continuous wavelet transform of the series x_t with respect to the wavelet $\psi_0(t)$ is defined as:

$$W_\tau^X(s) = \sqrt{\frac{\delta t}{s}} \sum_{t=1}^N x_t \psi_0^* \left[(t - \tau) \frac{\delta t}{s} \right], \quad (7.6)$$

where $*$ indicates complex conjugation. The mother wavelet is translated by τ along the time axis. The scale dilation parameter s corresponds to the width of the wavelet. The constant $\sqrt{\frac{\delta t}{s}}$ is a normalization applied such that the wavelet function at each scale has unit energy. This ensures that the transforms at each scale are directly comparable to each other and to the transforms of other series (Torrence and Compo, 1998). By varying the scale and translating by τ a picture showing both the amplitude of any features versus scale, and how this amplitude varies with time is constructed. In practice the continuous wavelet transform (Equation 7.6) is performed much faster in Fourier space, the details of which are given in Torrence and Compo (1998).

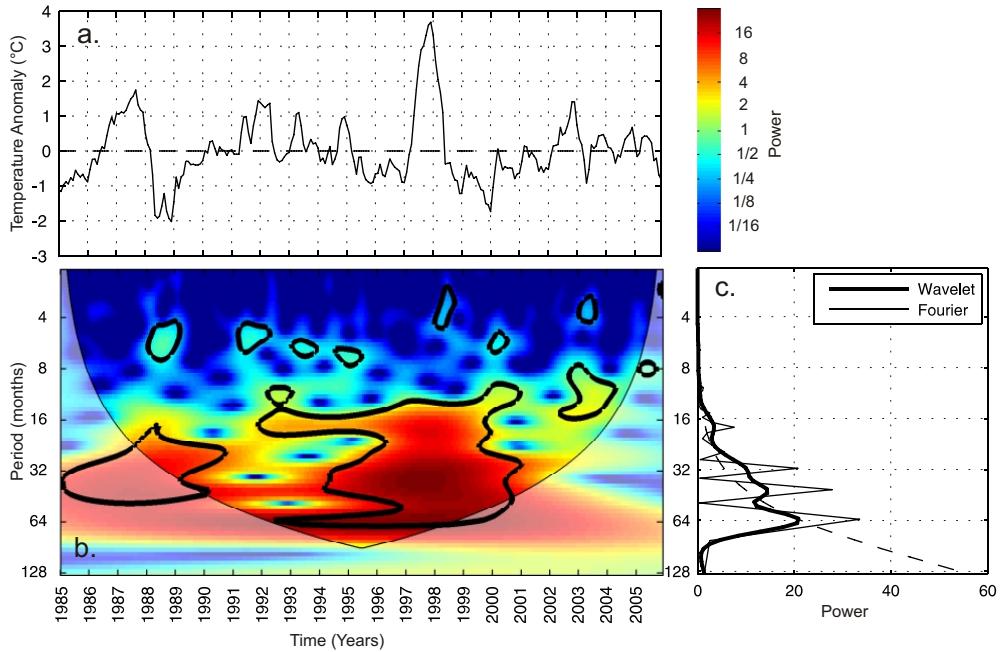


FIGURE 7.4: a. The time series of Niño3 SST anomalies. b. Local wavelet power spectrum for Niño3 SST normalized by $\frac{1}{\sigma^2}$ ($\sigma^2 = 0.94^{\circ}\text{C}^2$). The thick black contour designates the 95% significance level against a red noise background ($\alpha = 0.94$). The cone of influence (COI) where edge effects might distort the picture is shown as a lighter shade. c. Global wavelet power spectral density (PSD) and Fourier PSD both normalized by $\frac{1}{\sigma^2}$. Dashed line is the 95% confidence level for the global wavelet spectra based on a mean background spectrum of red noise ($\alpha = 0.94$).

7.3.2.1 Local Wavelet Power Spectrum

Analogous to Fourier analysis, the *local wavelet power spectrum* is defined as $|W_{\tau}^X(s)|^2$. Figures 7.4 and 7.5 show the normalized wavelet power spectra, $\frac{|W_{\tau}^X(s)|^2}{\sigma^2}$, for the Niño3 SST and SOI indices, where σ^2 is the variance of each time series (0.94 and 1.78°C^2 respectively). For a white noise process the expectation value for the wavelet transform is $|W_{\tau}^X(s)|^2 = \sigma^2$, for all scales and times. The normalization by $\frac{1}{\sigma^2}$ therefore gives a measure of the power relative to white noise.

The Niño3 SST index (Figure 7.4b) shows dominant power at periods of approximately 12-70 months (\sim 1-6 years). Note that the frequency content changes with time; something that the conventional Fourier analysis does not detect. Between 1985 and 1991 significantly high variance is observed at periods of 18-45 months. This corresponds to the 1987 El Niño and the following 1988/89 La Niña seen clearly in the plot of Niño3 SST anomalies above (Figure 7.4a). The following few years then display a decrease in power at these periods when there are no strong ENSO events characterized by the Niño3 SST record. The La Niña of 1996 followed by the strong El Niño of 1997/98 and subsequent La Niña period in the early 2000's introduce a wide band of high power across the full

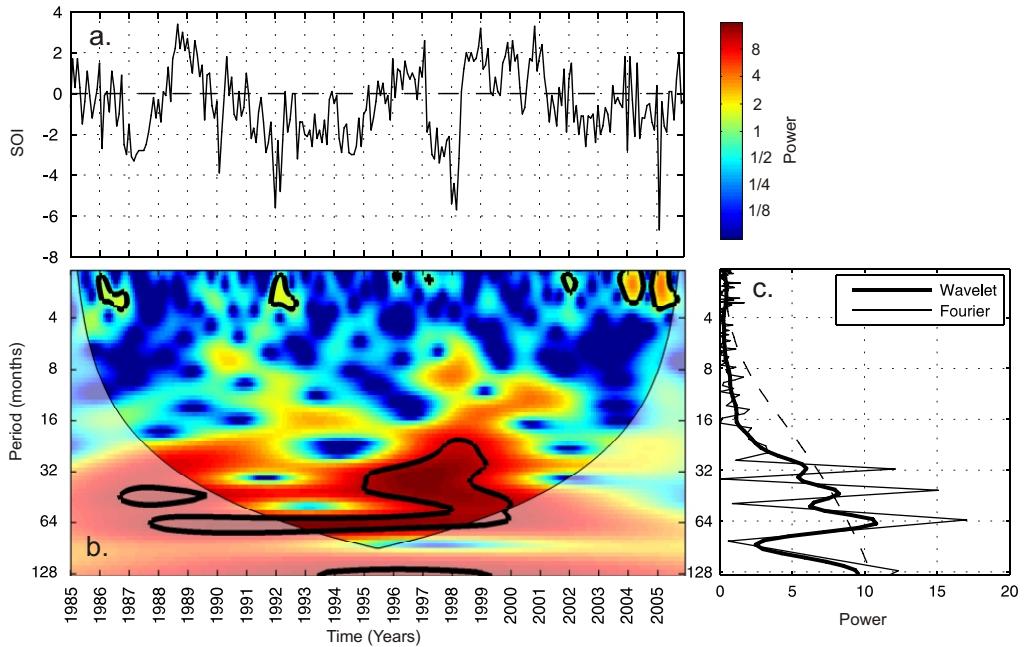


FIGURE 7.5: a. The time series of SOI. b. Local wavelet power spectrum for SOI normalized by $\frac{1}{\sigma^2}$ ($\sigma^2 = 1.78^{\circ}\text{C}^2$). The thick black contour designates the 95% significance level against a red noise background ($\alpha = 0.59$). The cone of influence (COI) where edge effects might distort the picture is shown as a lighter shade. c. Global wavelet power spectral density (PSD) and Fourier PSD both normalized by $\frac{1}{\sigma^2}$. Dashed line is the 95% confidence level for the global wavelet spectra based on a mean background spectrum of red noise ($\alpha = 0.59$).

range of periods, 12-70 months (1-6 years). This injection of high frequency power associated with the strong 1997/98 El Niño is also seen in the local wavelet power spectrum of the SOI (Figure 7.5b).

The lighter shaded area of the local power spectrum marks what is known as the Cone of Influence (COI). Each finite length series is zero padded before wavelet analysis to reduce errors introduced at the beginning and end of the times series because of the Fourier transform assumption that the data is cyclic. At larger scales, introducing these discontinuities decreases the amplitude near the edges as more zeros enter the analysis. The COI is the region of the wavelet spectrum in which edge effects become important. The peaks within this region have presumably been reduced in magnitude due to the zero padding. Additionally, a lack of data near the end of time series introduces some uncertainty into the calculation of power at each frequency. In Figure 7.4b for example it is unclear whether the decrease in power at periods of approximately 32 months near the end of the time series is a real loss or the result of edge effects and lack of data.

The 95% significance level of the wavelet power spectrum is calculated by assuming a background spectrum of red noise. The background red noise is modelled by the lag-1 autoregressive process (AR(1)) of the data set: $x_t = \alpha x_{t-1} + Z_t$ where α is the lag-1 autoregressive coefficient, $x_0 = 0$, and Z_t is taken from Gaussian white noise. Unlike for

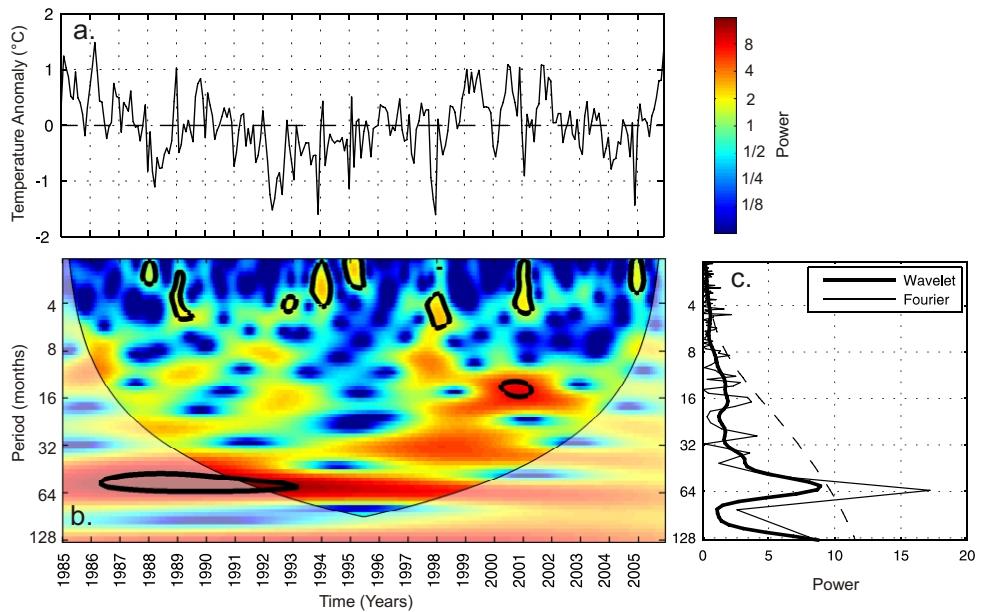


FIGURE 7.6: a. The time series of mean frontal temperature anomalies. b. Local wavelet power spectrum for frontal temperature anomalies normalized by $\frac{1}{\sigma^2}$ ($\sigma^2 = 0.56^{\circ}\text{C}^2$). The thick black contour designates the 95% significance level against a red noise background ($\alpha = 0.62$). COI shown as a lighter shade. c. Global wavelet power spectral density (PSD) and Fourier PSD both normalized by $\frac{1}{\sigma^2}$. Dashed line is the 95% confidence level for the global wavelet spectra based on a mean background spectrum of red noise ($\alpha = 0.62$).

white noise processes, the value at time t depends, at least in part, on the value at $(t-1)$. The appropriate coefficient α is estimated for each time series. It has been shown that the local wavelet power spectrum (a vertical slice through Figure 7.4b) follows the mean Fourier power spectrum of the time series (Torrence and Compo, 1998). Assuming that the time series has a mean red noise spectrum, any peak in wavelet power significantly above this background can be assumed to be a true feature with a certain degree of confidence. The significance levels for the wavelet power spectra are therefore relative to the null hypothesis that the signal is generated by a stationary process with a red background spectrum.

7.3.2.2 Global Wavelet Spectrum

Averaging all the local wavelet spectra over the entire time period yields the *global wavelet spectrum* $\bar{X}^2(s)$, equivalent to the Fourier Spectrum of the series (Torrence and Compo, 1998):

$$\bar{X}^2(s) = \frac{1}{N} \sum_{t=0}^{N-1} |W_t^X(s)|^2. \quad (7.7)$$

The global wavelet spectra for Niño3 SST anomalies, the SOI and the temperature of the Southland Front are shown in Figures 7.4c, 7.5c and 7.6c respectively, together with the 95% confidence levels. Note that for the frontal temperature, no particular scale is *globally* significant. The wavelet power spectrum however reveals that *locally* there are times when power at periods of around 55 months (~ 4.5 years) are significantly above the background red noise.

7.3.2.3 Scale Averaged Wavelet Power

Averaging can also be performed in terms of the scale. The *scale averaged wavelet power* can be used to examine fluctuations in power over a range of scales. It is defined as the weighted sum of the wavelet power spectrum over scales s_1 to s_2 (Torrence and Compo, 1998):

$$\bar{W}_t^2 = \frac{\delta j \delta t}{C_\delta} \sum_{j=j_1}^{j_2} \frac{|W_t(s_j)|^2}{s_j}, \quad (7.8)$$

where $\delta t = 1$ month is the sampling interval of the time series, $\delta j = \frac{1}{12}$ is the sampling interval between discrete scales, and $C_\delta = 0.776$ is a normalization constant specific to the Morlet Wavelet (Table 2 in Torrence and Compo (1998)). Scales $s_j = s_0 2^{j \delta j}$, $j_1 = \delta j^{-1} \log_2(\frac{s_1}{s_0})$, $j_2 = \delta j^{-1} \log_2(\frac{s_2}{s_0})$ where $s_0 = 2$ months is the smallest resolvable scale.

7.3.2.4 Cross Wavelet Power

Wavelet analysis may be used to compare two time series thought to be connected by some physical ocean/atmosphere process. The *cross wavelet spectrum* (XWT) of two series x_t and y_t with wavelet transforms $W_\tau^X(s)$ and $W_\tau^Y(s)$ is defined as $W_\tau^{XY}(s) = W_\tau^X(s) W_\tau^{Y*}(s)$, where $*$ denotes complex conjugation. The cross wavelet power is $|W_\tau^{XY}(s)|$ and will reveal regions of high power common to both series. When two data sets are compared the prewhitened versions of the two series are used to reduce leakage from the most intense spectral components, and low frequency components of the series that are poorly resolved. This greatly improves the statistical reliability of the spectral estimates made. Note that prewhitening does not remove the correlation between the time series.

Since the time series have been prewhitened the AR(1) coefficients (α) used to determine a background spectrum and confidence level are zero (i.e. white noise). The significance levels for the cross spectral power are thus relative to the null hypothesis that the signal is generated from a white rather than a red background spectrum.

For there to be a simple cause and effect relationship between the two series we would expect the oscillations to be phase locked. The complex argument $\arg(W_{\tau}^{XY}(s))$ is the local relative phase between x_t and y_t in time frequency space.

7.3.2.5 Wavelet Coherence

Perhaps a more useful measure is the *wavelet coherence* (CWT). The XWT reveals only regions of *high* common power in *both* series, whereas the CWT finds locally phase locked behavior even where the power of the signal at that frequency is low. It is defined as the square of the cross spectra normalized by the individual power spectrum of each series:

$$CWT_{\tau}^2(s) = \frac{|S(s^{-1}W_{\tau}^{XY}(s))|^2}{S(s^{-1}|W_{\tau}^X(s)|^2) \cdot S(s^{-1}|W_{\tau}^Y(s)|^2)}, \quad (7.9)$$

where S is a smoothing operator designed so that it has a similar foot print to the wavelet being used (Torrence and Compo, 1998). The wavelet coherence can be thought of as a localized cross correlation coefficient in time and frequency space. The statistical significance of the coherence is estimated using Monte Carlo methods. The coherence of 10,000 artificially generated data pairs of white noise is used to calculate the significance level for each scale (Grinsted et al., 2004). As for the XWT the phase relationship between the two series can be determined from the wavelet coherence. Note however that they will be slightly different owing to the smoothing operator applied during calculation of the coherence. All angles presented in the results section are calculated from the CWT.

7.3.2.6 Phase Relationship

A constant phase angle across all scales argues for a constant time lag due to a physical mechanism of signal propagation between the two series. The circular mean of the phase over regions with higher than 95% statistical significance that are outside the cone of influence is used to quantify the phase relationship. The circular mean of a set of angles a_i ($i = 1, \dots, n$) is used:

$$a_m = \arg(X, Y) \text{ with } X = \sum_{i=1}^n \cos(a_i) \text{ and } Y = \sum_{i=1}^n \sin(a_i). \quad (7.10)$$

It is difficult to calculate the confidence interval of the mean angle reliably because the phase angles are not independent (Grinsted et al., 2004). Instead, the circular standard deviation *csd* (scatter) of the angles around the mean is calculated as:

$$csd = \sqrt{-2 \ln \left(\frac{R}{n} \right)}, \quad (7.11)$$

where $R = \sqrt{(X^2 + Y^2)}$. If the mean phase angle is calculated for a particular scale then it can be quantified as a number of months.

7.4 Results and Discussion

7.4.1 Temperature of the Southland Front and ENSO

We begin our investigation into the relationship between ENSO and the Southland Front by examining the interannual variability of the front's mean temperature in relation to the SOI and regional Niño SST indices. Figure 7.7 shows a time series of frontal temperature anomalies estimated using the local likelihood based algorithm. Also displayed is a latitude-time plot of the temperature anomalies at the front. The seasonal cycle and linear trend have been removed at each latitude. All estimates with a quality flag greater than 1 (Figure 5.6) have been removed.

Anomalies reaching $\pm 1.5^{\circ}\text{C}$ persist along the entire length of the Southland Front (as measured here between 46.4°S and 44.5°S). There does not appear to be any spatial preference for this variability; the northern and southern limits of the front experience the same magnitude of fluctuations. Notable periods of warmer than usual conditions occurred between 1985 and mid 1986, briefly toward the end of 1989, and between 1999 and 2002. Colder than usual conditions were experienced between 1988 and 1989, and during the early to mid 1990's. The far two right plots in Figure 7.7 display the Niño3.4 SST and SOI indices. Broadly speaking, anomalously low temperatures at the Southland Front are associated with periods of strong negative SOI values and positive SST anomalies in Niño3.4. Conversely, warm anomalies at the front coincide with positive SOI values and colder than usual temperatures in the eastern central equatorial Pacific Ocean. A clear exception to this pattern occurs in 1988 when a positive SOI coincides with a negative temperature anomaly. This is a result of inconsistent phase lags between ENSO activity and changes at the front (see Sections 7.4.1.3 and 7.4.2.2).

Based on this initial evidence we may conclude that in general the temperature of the Southland Front decreases during El Niño events and increases during La Niña episodes. This is consistent with wider studies of the SST field around New Zealand (Greig et al., 1988; Sutton and Roemmich, 2001; Mullan, 1998). Sutton and Roemmich (2001) report the surface layer to be 0.7°C cooler than average between 1991 and 1994, and 0.7°C warmer off north-east New Zealand in 1999. Off the east coast of South Island inshore waters at Portobello cooled by 1.4°C between 1990 and mid-1992, while offshore waters

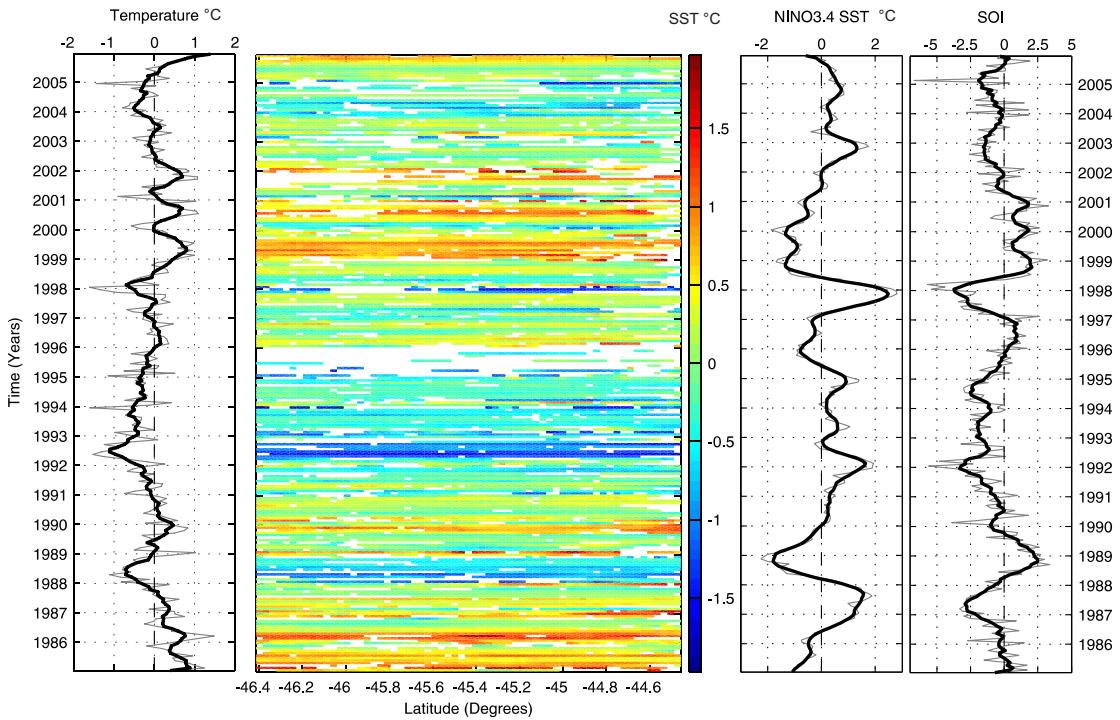


FIGURE 7.7: From left to right, plots of the weighted mean frontal temperature through time, a time-latitude plot of temperature anomalies at the front, the Niño3.4 SST index and the SOI. Thick black lines are an eight month running mean to help identify the main features. Gaps in the latitude-time plot are where estimates have been assigned a quality flag greater than 1.

showed even greater variation, cooling by 2.2-3.4°C over this same period (Shaw et al., 1999). In direct comparison to our measurements Shaw and Vennell (2001) report a significant decrease ($>1^{\circ}\text{C}$) in the temperature at the interface of the Southland Front between April 1989 and March 1992.

7.4.1.1 Seasonal ENSO Correlations

From Figure 7.7 the correlation between the SOI (Niño3.4 SST) and the Southland Front temperature (SFT) is expected to be positive (negative). Table 7.1 examines correlation between the time series where all monthly values and three monthly seasonal averages are considered. Correlations where the data have been stratified by season are also shown. Southern hemisphere seasons are used throughout, so summer is taken as December, January, February; autumn as March, April, May and so on. The one season lag results are also considered where ENSO indices lead the frontal temperature by one season. The seasonal lag for summer in Table 7.1 is therefore the summer SFT correlated with the spring SOI (or Niño SST) indices.

TABLE 7.1: Correlations between the temperature of the Southland Front, the SOI and regional Niño SST's over 21 years (01/85-12/05) for all monthly data combined, all seasonal data combined, and seasonally stratified data. * and † indicate significance at the 95% and 99% levels respectively. Seasonal leads are for the ENSO indices leading temperature at the front.

	SOI		Niño4 SST		Niño3.4 SST		Niño3 SST	
	0 Lag	1 Season Lead						
All Monthly	-0.13*		-0.07		-0.01		-0.01	
All Seasonal	0.03	0.09	-0.06	-0.14	-0.1	-0.12	0.02	-0.05
Spring	0.41	0.38	-0.17	-0.26	-0.35	-0.23	-0.21	-0.29
Summer	0.48*	0.63†	-0.50*	-0.50*	-0.67†	-0.61†	-0.72†	-0.48†
Autumn	0.31	0.37	-0.35	-0.43	-0.5*	-0.60†	-0.37	-0.57†
Winter	-0.01	0.066	-0.16	-0.35	-0.04	-0.33	0.23	-0.25

When all months are considered simultaneously (zero lag) the SOI and SFT are significantly negatively correlated (-0.13) at the 95% level. A very weak positive correlation (0.03) is revealed once the data has been seasonally averaged. Stratifying the data by season shows that the SOI-SFT correlation is seasonally dependent. There is a positive correlation during the spring, summer and autumn months. This relationship is strongest during the summer where a correlation of 0.48 is significant at the 95% level. During the winter there is a very weak negative correlation from which we conclude that there is no relationship between the SOI and temperature at the Southland Front between June and August. With the exception of spring, all SOI-SFT correlations are strengthened when the SOI leads by one season. This is most striking during the summer where the correlation increases from 0.48 (zero lag) to 0.63 when the spring SOI values are correlated with summer SFT's. The significance of this relationship also increases from the 95% to 99% level.

The seasonal dependence of the significance of SOI-SFT correlations may be attributed to both the water column structure, and the seasonal variation in correlation patterns between the mean sea level pressure (MSLP) in the Australasian region and the SOI. SST's are likely to be more sensitive to atmospheric variations in the summer half of the year when oceanic thermocline depths are shallower and vertical mixing is reduced. Also, during the summer the MSLP around South Island is significantly positively correlated with the SOI (Figure 7.8). For a negative SOI indice during the summer months a low pressure anomaly with clockwise rotating winds dominates over South Island (Gordon, 1986). Under this cyclonic anomaly skies will be cloudier and less radiation will reach the sea surface; furthermore, divergence of surface waters will occur in the centre of the low pressure anomaly, leading to upwelling and cooler SST's. The reverse is true for positive summer SOI values; under an anticyclonic anomaly more radiation reaches the sea surface and convergence of surface waters leads to downwelling and increased temperatures. In the autumn and spring months this region of significant positive correlation is displaced to the east and south of New Zealand respectively (Figure 7.8) and the radiative and up/downwelling effects of pressure anomalies are not so strongly felt

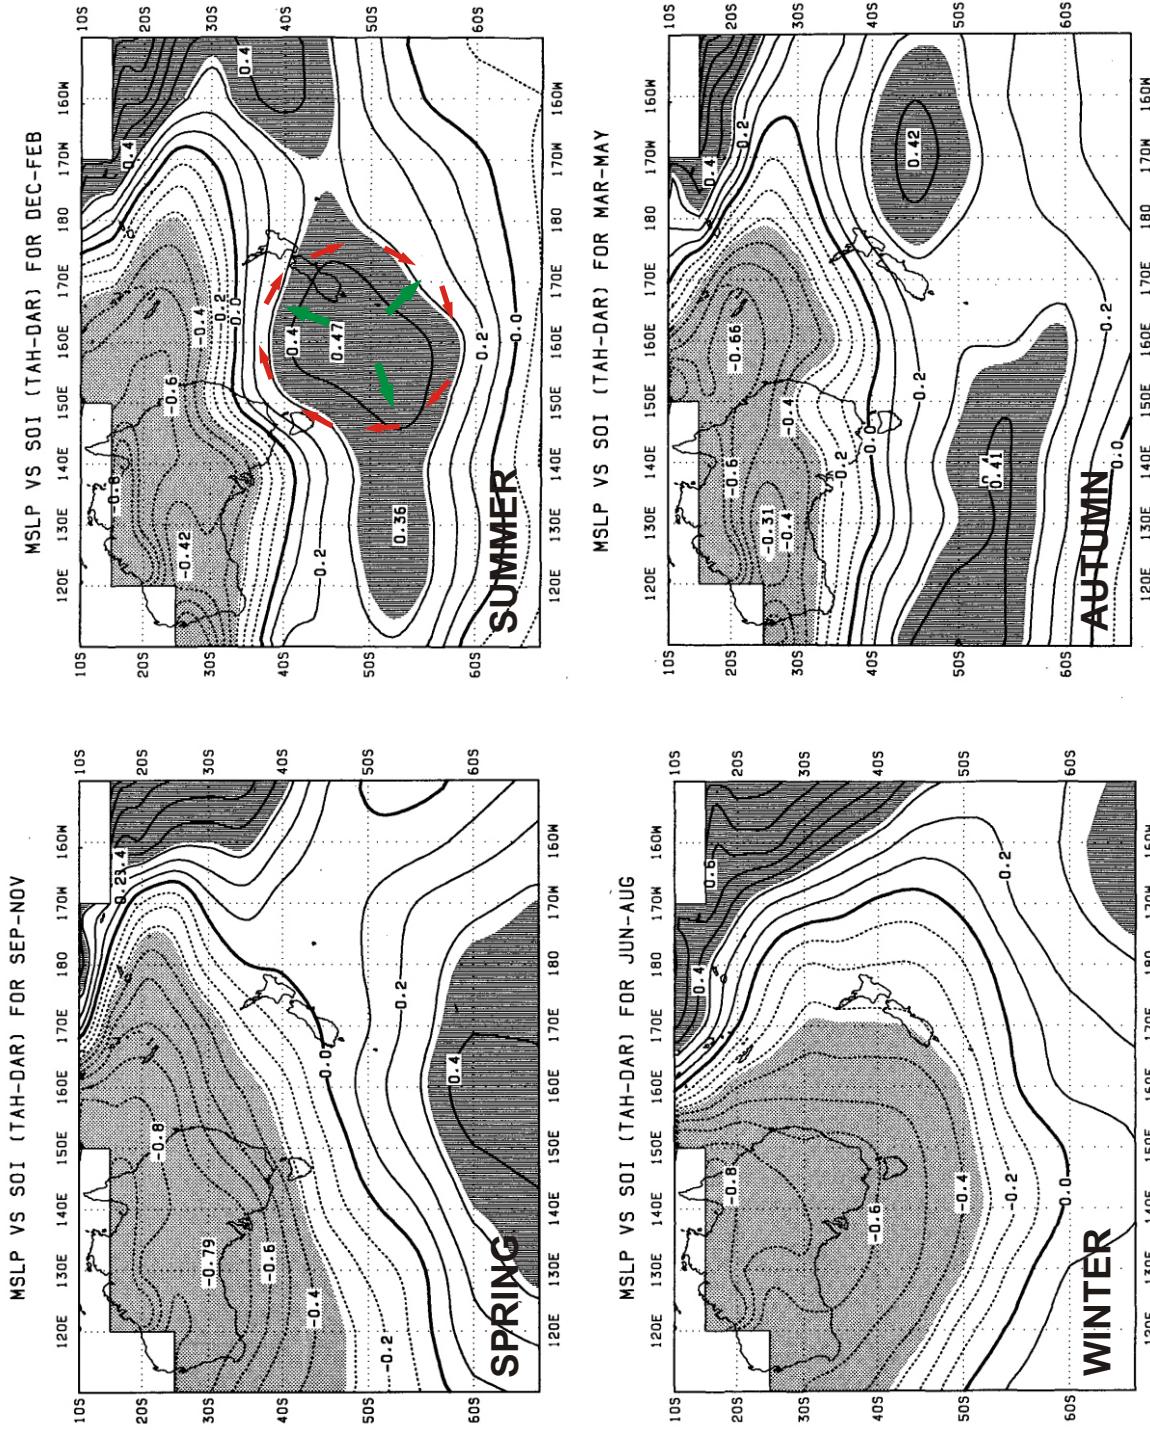


FIGURE 7.8: Correlation coefficients between three month average MSLP and the SOI for all four seasons. Coefficients in the shaded regions exceed the 90% significance value of 0.32. Red arrows indicate the summer wind anomaly during El Niño (negative SOI) events and the green arrows show the divergence of surface waters. Modified from Gordon (1986).

around South Island. Weaker winter correlations are consistent with Mullan (1998). Using the leading rotated empirical orthogonal functions of monthly mean SST anomalies around New Zealand, Mullan (1998) found that in general simultaneous correlations with the atmospheric circulation were markedly weaker during the winter.

A similar seasonal dependence applies to the Niño SST regional anomalies and SFT correlations (Table 7.1). Only a very weak negative correlation is revealed when all months and seasons are considered together. This correlation is strengthened by seasonally stratifying the data and is strongest, for a zero lag correlation, during the summer. Unlike for the SOI-SFT relationship a one season lag (spring-summer) does not strengthen the summer Niño SST-SFT correlation, although the relationship remains significant at the 95% level or above. Instead, the lagged summer-autumn correlation (Niño SST's leading) coefficients are strengthened. This is especially true for Niño3-SFT; summer anomalies in Niño Region 3 are significantly negatively correlated (-0.57) at the 99% confidence level with the following autumn temperatures at the Southland Front. The contemporary autumn correlation is -0.37. So a warming of the central equatorial Pacific in the summer during an El Niño corresponds to a decrease in STF the following autumn.

Monthly lags within seasonally stratified data representing month to month relationships within a season were also examined (not shown). Note that these results represent correlations ending during the specified season. For example a one month lag during the summer is November-December, December-January and January-February. During the summer a negative Niño4-SFT correlation (-0.28), significant at the 95% level, is maintained for a one and two month lag. It is generally more difficult to predict variations within a season than variations of the season as a whole (Mullan, 1998) therefore these one month lag correlations between summer months serves to reinforce the stronger summer teleconnection suggested in Table 7.1.

7.4.1.2 Oceanic Advection of SST Anomalies

The significant seasonal lag responses are suggestive of an oceanic connection between the SFT and ENSO, additional to the fast (order 1 month) heat fluxes driven by local atmospheric circulation processes. Mullan (1998) calculates the pattern of SOI and SST correlations over the southwest Pacific for the period 1949-1991. At zero lag, for all seasons combined, South Island lies outside the region of significant positive correlation that stretches from Indonesia, east-southeast into the central South Pacific (Figure 7.9). This is consistent with the weak (0.03) positive SOI-SFT correlation in Table 7.1 for three monthly averaged data.

Figure 7.9 shows that the Coral Sea and Northern Tasman Sea experience the most significant SST anomalies during ENSO events. Section 2.4 and Figure 2.6 describe

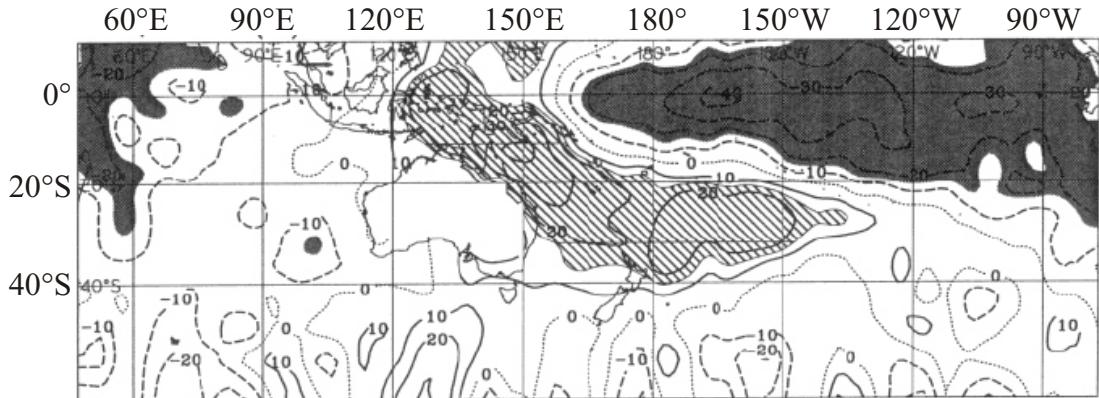


FIGURE 7.9: Seasonal correlations at zero lag ($\times 100$) between the SOI and SST, all seasons combined. Crossed (black) shading denotes positive (negative) correlations significant at the 95% level. From Mullan (1998).

how flow in the Tasman Sea is either directed eastward toward North Island, or south-southeast toward South Island forming the Southland Current. Current speeds across the Tasman Sea are not well documented. Chiswell et al. (1997) report maximum southerly current speeds of 2m.s^{-1} and 0.7m.s^{-1} at 30°S and 47°S respectively. These values were recorded between 1990 and 1994 along WOCE repeat hydrographic sections at the mentioned latitudes. More conservative estimates are likely to be of the order 0.2m.s^{-1} (Mullan, 1998). At 20km.day^{-1} (0.2m.s^{-1}) it would take approximately 3-4 months for an ENSO related temperature anomaly in the Coral Sea to be advected to the southern tip of South Island, and into the northward flowing Southland Current. Assuming a faster current of 0.7m.s^{-1} this transport time would be of the order of one month. A 3-4 month advection time across the Tasman Sea fits well with stronger one season lead correlations.

From rotated EOF SST patterns Mullan (1998) identifies two regions of significant SOI-SST correlation; the East Australia Current (EAC) region off the east coast of Australia, and waters south of the South Pacific Convergence Zone. The South Pacific Convergence Zone is located approximately along the diagonal line in Figure 7.9 where the correlation reverses sign. Both regions have correlations significant at the 99% level in at least one season and are overall strongest in the winter and spring (Mullan, 1998). Goring and Bell (1999) find no significant lag between the SOI and SST to the north of New Zealand. This all fits well with the strong SOI-SFT seasonal lead correlations (winter-spring and spring-summer) seen in Table 7.1, and a 3-4 month advection time across the Tasman Sea. Significant ENSO related SST anomalies in the northern Tasman Sea, EAC and south-west Pacific during winter and spring take 3-4 months to be advected southward by Tasman Sea currents reaching the STF in spring and summer respectively, approximately one season later.

7.4.1.3 A Variable Phase Relationship

So far we have analyzed the seasonal dependence of the SFT-ENSO relationship and suggested a possible mechanism by which ENSO related SST anomalies may be advected into the Southland Current approximately one season after peak events. In this section we take a closer look at the response of the mean temperature of the front to individual ENSO events, investigate the frequencies at which correlation exists and the consistency of the phase relationship between the two signals.

There is a substantial amount of evidence to indicate that the timing and magnitude of the temperature changes at the Southland Front in response to ENSO events is neither consistent nor easily predictable. Cross correlation analysis is expected to reveal the lead or lag time of any regular teleconnection between two time series and was performed for prewhitened monthly data. As explained in Section 7.3.1 the reliability of significant lead or lag correlations should be treated with caution. A correlation that is significant at the 99% level does not necessarily indicate the existence of any real connection between two time series. The stability of possible teleconnections identified from cross correlations using the entire twenty-one year data set were therefore tested by calculating lag correlations with respect to the two halves of the time series (01/85-06/94 and 07/94-12/05). Any real persistent connections should be identified in all three sets of data. On no occasion was the same lead or lag correlation significant in all three tests. This provides the first indication that the atmospheric and oceanic processes connecting ENSO variability to the SFT are inconsistent.

Figure 7.10 displays the coherence between the prewhitened SFT and ENSO indices and reveals not only the specific frequencies at which teleconnections occur but further evidence that the cause and effect relationship between ENSO and SFT is temporally variable. There is significant coherence in an 8-32 month period band between 1989 and 2003. Within this band a number of features are noted. From 1991 to 1994 a wide band of high and significant coherence (95% level) between 8 and 23 months coincides with a series of weaker El Niño events that took place in rapid succession during the early 1990's. From 1995 to 2001 there are two separate bands of high SOI-SFT coherence; the first spanning periods of 23-32 months persists throughout (1995-2002), and the second, at higher frequencies (\sim 10 months) occurs between 1996 and mid 1998. This coincides with the strong 1997-1998 El Niño event.

The oscillations within the zones identified above are not phase locked indicating that there is no one simple physical mechanism linking the SOI and SFT. The arrows on Figure 7.10 indicate the relative phase difference between the two series. Arrows pointing to the right show the two series to be in-phase (0°). An anti-phase relationship is represented by arrows pointing to the left ($\pm 180^\circ$). Arrows pointing straight down (90°) show ENSO to lead the temperature at the front. Arrows pointing straight up result from the temperature leading the ENSO signal (-90°).

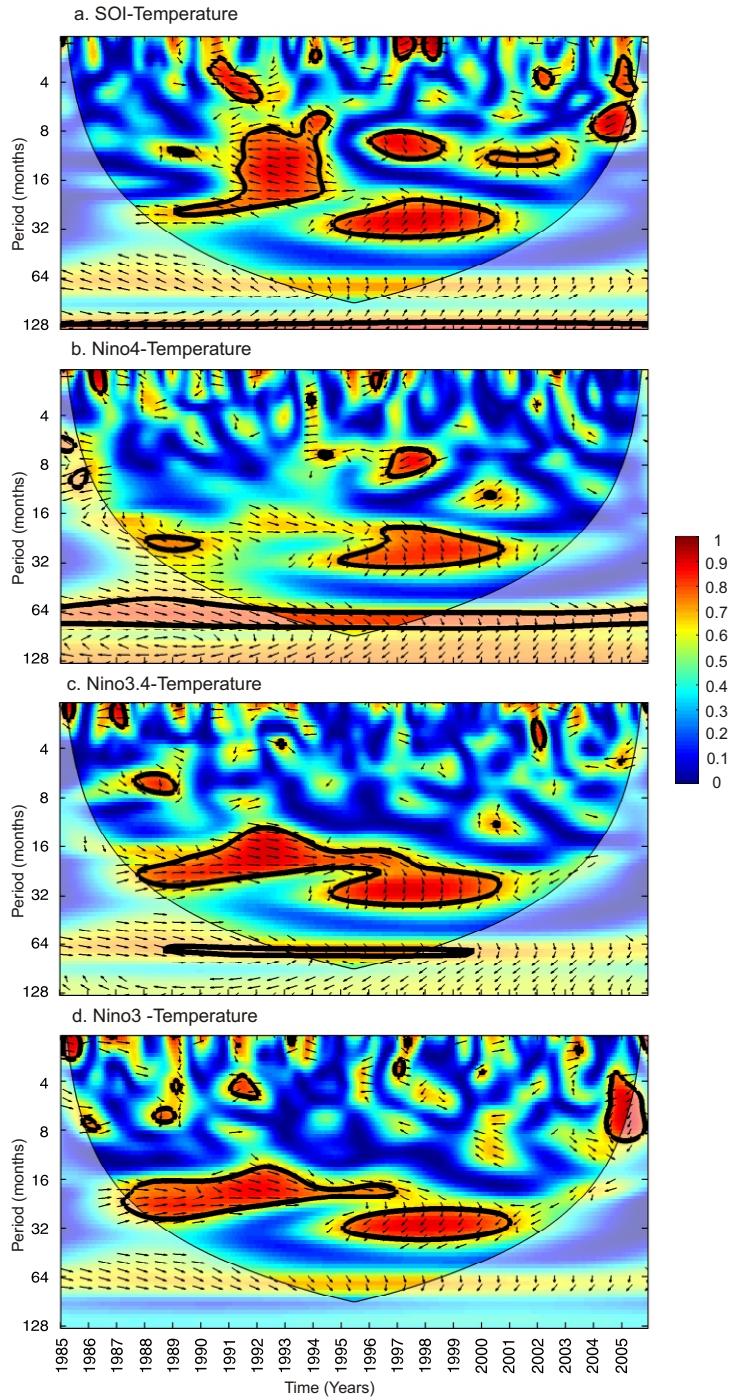


FIGURE 7.10: Wavelet squared coherence between the temperature of the Southland Front and a. SOI, b. Niño4 SST, c. Niño3.4 SST, d. Niño3 SST. Thick black contours are the 95% confidence level above a white noise background spectrum. Arrows indicate the relative phase relationship. In phase pointing right (0°), antiphase pointing left ($\pm 180^\circ$), ENSO indice leading temperature by 90° straight down.

Prior to 1994 the phase relationship (for SOI-SFT) in regions with higher than 95% significance that are outside the COI is $-156 \pm 22^\circ$ (circular mean \pm circular s.d), suggesting that the temperature at the front leads the SOI. For oscillations with a period of 8 months this translates to an approximate 3-4 month lead. Alternatively this may be

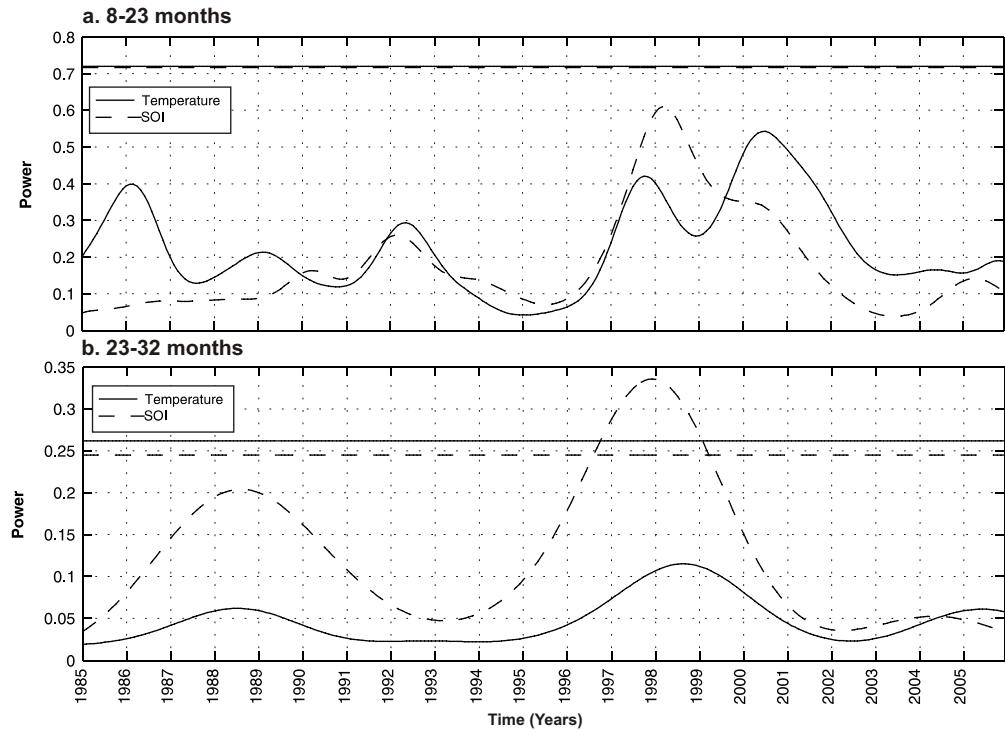


FIGURE 7.11: Scale averaged wavelet power over the a. 8-23 month and b. 23-32 month period bands for the SOI and temperature of the Southland Front. Power has been normalized by the variance of each time series. The horizontal lines are the 95% confidence levels for the SOI (assuming red noise $\alpha = 0.59$) and for the temperature (assuming red noise $\alpha = 0.62$).

interpreted as the SOI leading the temperature by 4-5 months (for an 8 month frequency component). The phase relationship within the 23-32 month band between 1995 and 2001 is $-72 \pm 10^\circ$. For signals with a period of 23 months this translates to the temperature at the front leading the SOI by 4-5 months. There appears therefore to be a shift in the phase lag between the SOI and temperature relationship around 1994. This is coincident with a brief change in the mode of the Pacific Decadal Oscillation (see Section 7.4.1.4).

The scale averaged wavelet power over the 8-23 and 23-32 month bands confirms the correlation at these frequencies (Figure 7.11). The common phase locked behavior in power between periods of 8-23 months is clearly seen for the El Niño episodes in the early 1990's and for the strong 1997-1998 event. Power within the 23-32 month band is significantly positively correlated (0.90) at the 95% level across the entirety of the record.

Phase locked behavior between the temperature of the Southland Front and Niño SST regions 3, 3.4 and 4 is also highlighted in Figure 7.10b-d. A wide band of significant coherence (12-36 months) is identified between 1987 and 2001 with variations in temperature across the equatorial Pacific leading temperature changes at the Southland Front. Before 1994 oscillations are more or less in-phase. The circular mean angle between 12

TABLE 7.2: Mean phase relationship (\pm circular s.d) for regions of significant coherence between 12 and 36 months prior to- and post 1994. Positive phase indicates that Niño SST's lead the temperature at the Southland Front.

	Pre 1994	Post 1994
Niño4	-	73 \pm 12°
Niño3.4	9 \pm 20°	95 \pm 5°
Niño3	10 \pm 25°	113 \pm 5°

and 36 months for Niño 3.4 and Niño 3 is 9 \pm 20° and 10 \pm 25° respectively. This is a lead of less than one month but should be treated with caution owing to the large circular standard deviations. After 1994 coherence between the SFT and the Niño SST regions shifts to slightly lower frequencies and the relative phase increases to approximately 6 \pm 3 months (depending on the frequency). Table 7.2 summaries the mean circular phase for all significantly coherent periods between 12 and 36 months prior to- and post 1994. Note how the lag response increases as the Niño region in question is located further eastward across the Pacific. The response time is greater for Niño 3 than for Niño 4 which is further west and the closest Niño region to Indonesia.

In Figure 7.10 there is a band of high coherence persistent across the entire data record at longer periods of approximately 64 months (5 years) for all ENSO indices. This band however is not always significant at the 95% level and stretches across the COI. It is strongest for the Niño 4-SFT relationship between 60 and 90 months (5-7.5 years). This suggests a correlation at typical ENSO timescales but should be treated with caution since it is at the limit of resolvable frequencies given the length of the time series.

The variable phase relationship evident in Figure 7.10 is likely the result of a number of factors. Firstly, each ENSO events is different; La Niña and El Niño events develop at different speeds, are initiated by different triggers and decay at different rates. The magnitude and distribution of SST anomalies across the equatorial Pacific also vary between events and do not always coincide with peak SOI values. As seen in Section 7.4.1.1 the season within which an ENSO event takes place can also play a determining role in the strength and timing of ocean-atmosphere interactions.

Consideration should be given to other modes of climate variability that have a signal in the southwest Pacific and/or are able to modulate ENSO teleconnections. An explanation for the variable response of the SFT to changes in the SOI may in part be sought in the Pacific Decadal Oscillation (PDO).

7.4.1.4 Pacific Decadal Oscillation Modulation of ENSO Teleconnections

The PDO is a mode of climate variability that influences the SST's, sea level pressure and surface winds in the Pacific Ocean in a similar way to ENSO (Gershunov and

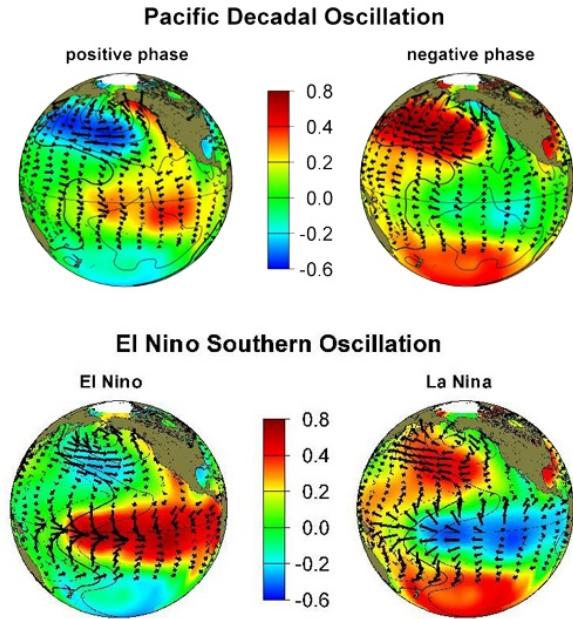


FIGURE 7.12: Sea surface temperature and winds during positive and negative phases of the PDO (top) and El Niño and La Niña phases of ENSO (bottom). Dark red and dark blue indicate temperature anomalies of $+0.8^{\circ}\text{C}$ and -0.6°C , respectively. Wind stress directions and magnitudes are indicated by vectors. From Mantua (2000).

Barnett, 1998). Their differences lie in the time scale and primary region of influence. Whereas ENSO events tend to persist for approximately one year, the PDO signature can persist for up to 30 years and is most prominent in the North Pacific (Mantua and Hare, 2002). When the PDO is positive waters in the north central Pacific tend to be cool, and SST's along the west coast of North America are warm (Figure 7.12). The opposite is true for a negative PDO. ENSO may be thought of as lying ontop of this large scale PDO temperature distribution.

The PDO can have a modulating effect on ENSO teleconnections (Gershunov and Barnett, 1998; Power et al., 1999; Folland et al., 2002). The time series of the PDO has sign reversals that match some of the observed change points in the New Zealand climate over the last 70 years (Salinger and Mullan, 1999). Since 1978 the PDO has been predominantly in a positive phase where associated SST anomalies in the eastern tropical Pacific have enhanced the pattern of El Niño warming. The PDO is quantified by the leading principal component of the monthly SST anomalies in the North Pacific Ocean poleward of 20°N . The monthly mean global average SST anomalies are removed to separate this pattern of variability from any ‘global warming’ signal that may be present in the data. The PDO index was obtained from the NOAA Climate Data Centre ([Figure 7.13 compares the SOI, PDO and SFT between 1985 and 2005. Also shown are the zero lag correlation coefficients between the SOI and PDO \(non smoothed data\),](http://www.cdc.noaa.gov/ClimateIndices>List/).</p>
</div>
<div data-bbox=)

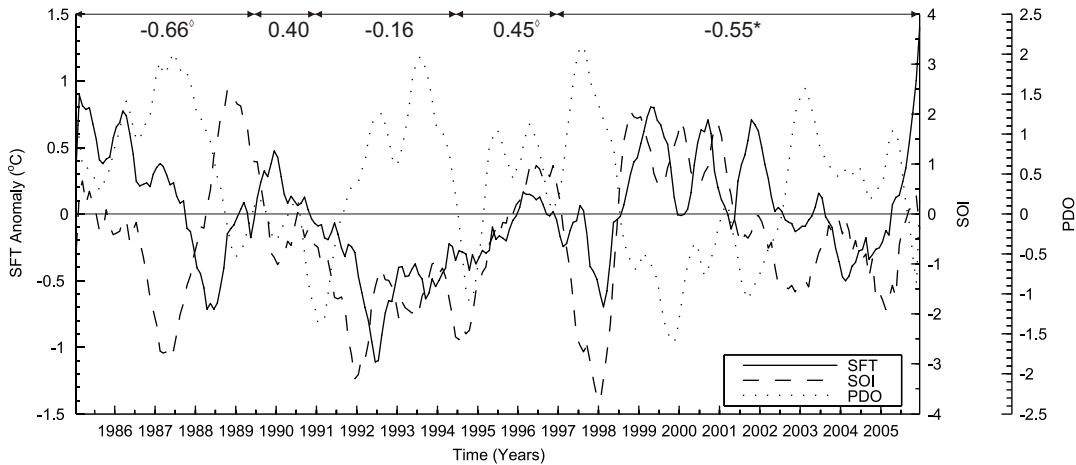


FIGURE 7.13: Eight month running mean of temperature anomalies at the Southland Front (SFT), the Southern Oscillation Index (SOI) and the Pacific Decadal Oscillation (PDO). Zero lag correlation between the SOI and PDO (non smoothed data) for the time periods indicated by the horizontal arrows are also shown. $^{\circ}$ and $*$ indicate significance at the 90% and 95% levels respectively. The correlation coefficient for the entire twenty-one year series is -0.35^* .

for five separate sub-periods of time. For the most part the SOI and PDO are in-phase; negative SOI values (El Niño) are matched with a positive PDO, the correlation coefficient between them is negative and ENSO related SST anomalies around New Zealand are enhanced. However there are two short periods when the warm phase of the PDO is interrupted by a brief shift into the cold (negative) regime. This happens between 1989 and 1991, and again in 1994. These episodes correspond to positive SOI-PDO correlation and a dampening of the ENSO signal at the Southland Front and may help explain the variable response of the mean frontal temperature to ENSO events.

ENSO shifts into a warm phase of El Niño episodes (negative SOI) at the end of 1989. A corresponding decrease in the temperature of the Southland Front however is delayed by several months. The minimum SOI at the end of 1991 is matched by a temperature minimum at the front five months later in 1992. During 1990 when the SOI starts to fall the PDO has shifted into a negative regime resulting in anomalously warm waters in the southwest Pacific Ocean around New Zealand (Figure 7.12). This background distribution of temperatures may be acting to dampen the full impact of cooling that is expected in this region during an El Niño. In early 1991 the PDO shifts back into a positive phase where ENSO SST's are enhanced. It is at this point that a rapid decrease in the mean frontal temperature occurs and the minimum anomaly is observed much later than the minimum SOI value. Throughout the early 1990's the PDO remains positive but is interrupted once again in 1994 by a temporary shift to a negative regime. This results in a weaker SST anomaly than expected at the Southland Front in response to the 1994 El Niño. It may also be a contributing factor to the end of the prolonged warm phase of ENSO. In 1995 the PDO reverts back to a positive state and is accompanied by a weak La Niña in 1996 which sustains the positive SOI-PDO correlation of 0.45° . The

warm anomaly at the front may be restrained by the background cooling of the PDO in the southeast Pacific.

The PDO clearly has the potential to disrupt the ENSO-SFT teleconnection and is a determining factor in the strength and timing of temperature anomalies at the front following El Niño and La Niña events. The brief change in 1994 to a negative PDO regime coincides with the shift in phase relationship between the SFT and ENSO indices highlighted by the wavelet coherence in Figure 7.10 and Table 7.2.

7.4.2 Strength of the Southland Front and ENSO

In this next section we explore the relationship between the gradient of the Southland Front, which may be used as a proxy for the strength of geostrophic flow, and ENSO events. Shaw and Vennell (2001) report a decrease in the gradient of the Southland Front between 1989 and 1992 but with only three years worth of data were unable to provide conclusive evidence that this was linked to the coincident decrease in SOI. With a much longer time series we are able to independently verify or otherwise a weakening of the front during the early 1990's and to investigate whether similar anomalies have occurred in response to other El Niño and La Niña events.

7.4.2.1 Impact of Mernoo Saddle Currents

The thermal gradient (strength) of the Southland Front is determined by a combination of the temperature difference between Subtropical (STW) and Subantarctic waters (SAW), and the distance over which the transition between these two water masses occurs (i.e. the width of the front). We wish firstly to establish whether interannual variability of these parameters is spatially dependent and if so determine why.

Figure 7.14 examines the interannual variability of the temperature range, width and gradient of the Southland Front between 46.4°S and 44.4°S. There are distinct periods of positive and negative anomalies in the temperature range that show no strong spatial dependence i.e. the entire length of the front is affected. There is a less well defined structure to the interannual variability of the width and a tendency for larger and more noisy anomalies further north within the Canterbury Bight. This is likely the result of a combination of factors. Firstly, the northern section of the front as it passes the Canterbury Bight is influenced by flows through the Mernoo Saddle (Shaw and Vennell, 2000b). SAW extends north through the western edge of the Mernoo Saddle throughout most of the year, however in winter and early spring a southward extension of STW may occur. This flow reversal is thought to be linked to the pressure force between the STW of an anticyclonic eddy at Kaikoura (possibly the Hikurangi Eddy), and the cooler, less saline water of the Southland Current (Greig and Gilmour, 1992; Heath, 1975). Heath

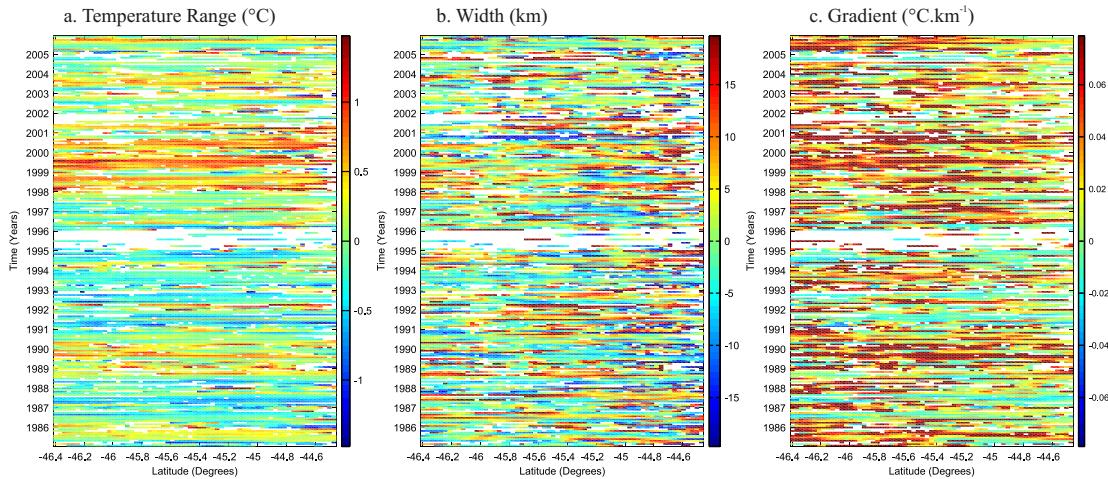


FIGURE 7.14: Latitude-time plots of a. temperature range, b. width and c. gradient anomalies. The annual cycle has been removed by subtraction of the 1985-2005 monthly means for each month at each latitude. Only those estimates with a quality flag of 1 are shown.

(1975) noted that should this eddy exist just north of the Saddle then the Southland Current will be directed further offshore, east along the Chatham Rise. This increased angle of rotation would introduce a tendency for the algorithm to overestimate the width of the front as explained in Section 5.3.3. Lastly, the algorithms ability to make accurate estimates of frontal characteristics is weakest for the width (Section 5.3) which explains the more noisy results in the middle panel of Figure 7.14. A bias toward overestimation of the width in the north is reflected in the latitude-time plot of the gradient (right panel); stronger positive anomalies are preferentially observed further south.

We conclude that there are processes north of the Chatham Rise that may periodically influence the location of the Southland Front. We are unable to confidently differentiate between increases (decreases) in the width that are due to over (under) estimation by the algorithm and increases that reflect real oceanographic changes. Therefore only data poleward of 45.5°S is used to investigate any ENSO teleconnections related to the strength of the front. Owing to its more robust estimation the temperature differential between the SAW and STW is perhaps a better indicator of changes in thermal gradient (assuming that the width remains stable).

7.4.2.2 An Inconsistent Teleconnection

Wavelet coherence analysis confirms the existence of a teleconnection between the Southland Front gradient and ENSO. Figure 7.15 reveals bands of significant coherence at periods between 8 and 32 months for all ENSO indices used in this study. The scale averaged wavelet power spectra (not shown) over the 8-32 month period band for the

SOI, Niño4, Niño3.4 and Niño3 are all significantly positively correlated (95% confidence) with the corresponding power spectra for the gradient (0.70, 0.67, 0.73 and 0.74 respectively). This correlation increases (>0.8) when the frequency band is narrowed to 16-32 months.

The negative relative phase relationship in Figures 7.15b-d suggests that the gradient and Niño SST anomalies are in anti-phase, the gradient leading changes in the equatorial SST by 6-12 months. The SOI-gradient phase relationship however is positive ($77.6 \pm 1.4^\circ$) indicating the SOI leading gradient variability. As for the mean temperature of the front the phase relationship does not remain consistent across the entirety of the record suggesting that the mechanism connecting the strength of the Southland Front to ENSO is neither consistent nor straight forward.

To help better understand the response of the Southland Front to individual ENSO events Figure 7.16 compares changes in the gradient, temperature range and width to the SOI and Niño3.4 SST anomalies. The raw time series have been smoothed with a twelve month running mean to dampen any signals from the seasonal cycle. Estimates during 1995 are considered to be highly unreliable owing to the anomalously high and unexplained lack of data during this period (Figure A.3). As can be seen in Figure 7.14 very few estimates along the length of the front passed the quality control checks over this period leaving behind only a handful of numbers with which to calculate any overall mean and trend. The increase in gradient during 1995 is to be treated with caution.

From late 1989 to mid 1991 there is a well defined and steady decrease in the gradient from 0.13 to $0.07^\circ\text{C}.\text{km}^{-1}$, coincident with a drop in the SOI. This is comparable to the decrease of $0.1^\circ\text{C}.\text{km}^{-1}$ reported by Shaw and Vennell (2001). Note that the minimum gradient is observed during mid 1991, approximately six months ahead of the minimum SOI index and maximum Niño3.4 SST anomaly. A further two local gradient minima occur during the following period of weak but persistent El Niño events; the first in 1992, a few months ahead of the second peak in equatorial SST's, and the second in 1994 coincident with the 1994 El Niño. A similar weakening of the gradient occurs during the 1987 El Niño. In this instance the gradient and SOI increase and decrease in phase. The pattern of decreasing gradient during negative SOI episodes is broken in 1997. Contrary to the 1987 and 1991 events the strong 1997/1998 El Niño is accompanied by a sharp increase in gradient.

Figure 7.16 also provides evidence for increased gradients during La Niña events. There is a distinct strengthening of the Southland Front from mid 1998 to 2001 during a period of persistently positive SOI values.

The middle and bottom panels of Figure 7.16 show interannual changes in the temperature difference and width respectively, the two variables controlling the gradient of the Southland Front. This allows us to determine which characteristic is responsible for driving changes in the gradient at each individual event. On the whole, variability of

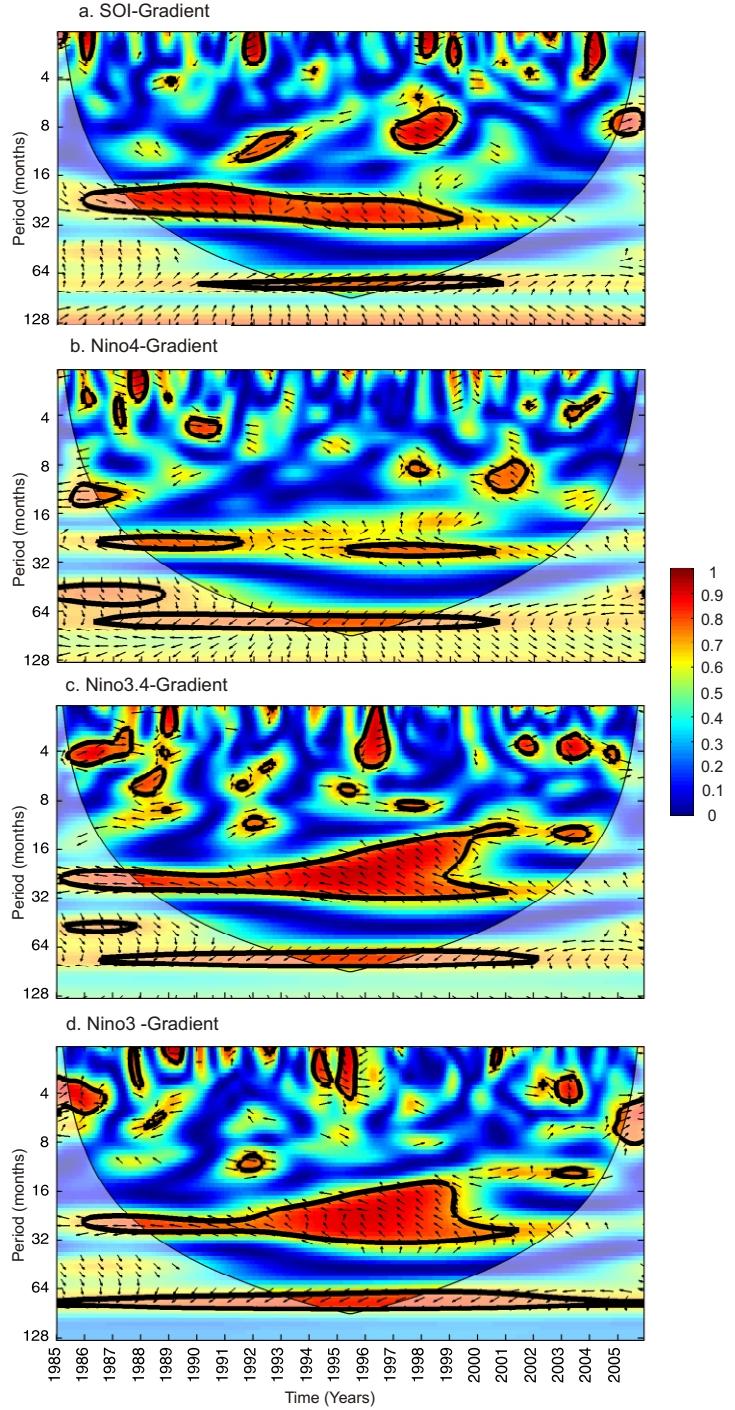


FIGURE 7.15: Wavelet squared coherence between the gradient of the Southland Front and a. SOI ($77.6 \pm 1.4^\circ$), b. Niño4 SST ($-139 \pm 1.5^\circ$), c. Niño3.4 SST ($-149.7 \pm 1.2^\circ$), d. Niño3 SST ($-134.8 \pm 1.1^\circ$). The mean relative phase (± 1 s.d) for all correlations outside the COI and above the 95% confidence level are shown in brackets. Thick black contours are the 95% confidence level above a white noise background spectrum.

Arrows as for Figure 7.10.

the temperature difference exhibits a much stronger relationship with the SOI than does the width. It is unclear however whether this is in part due to the more unreliable local likelihood estimates of the width.

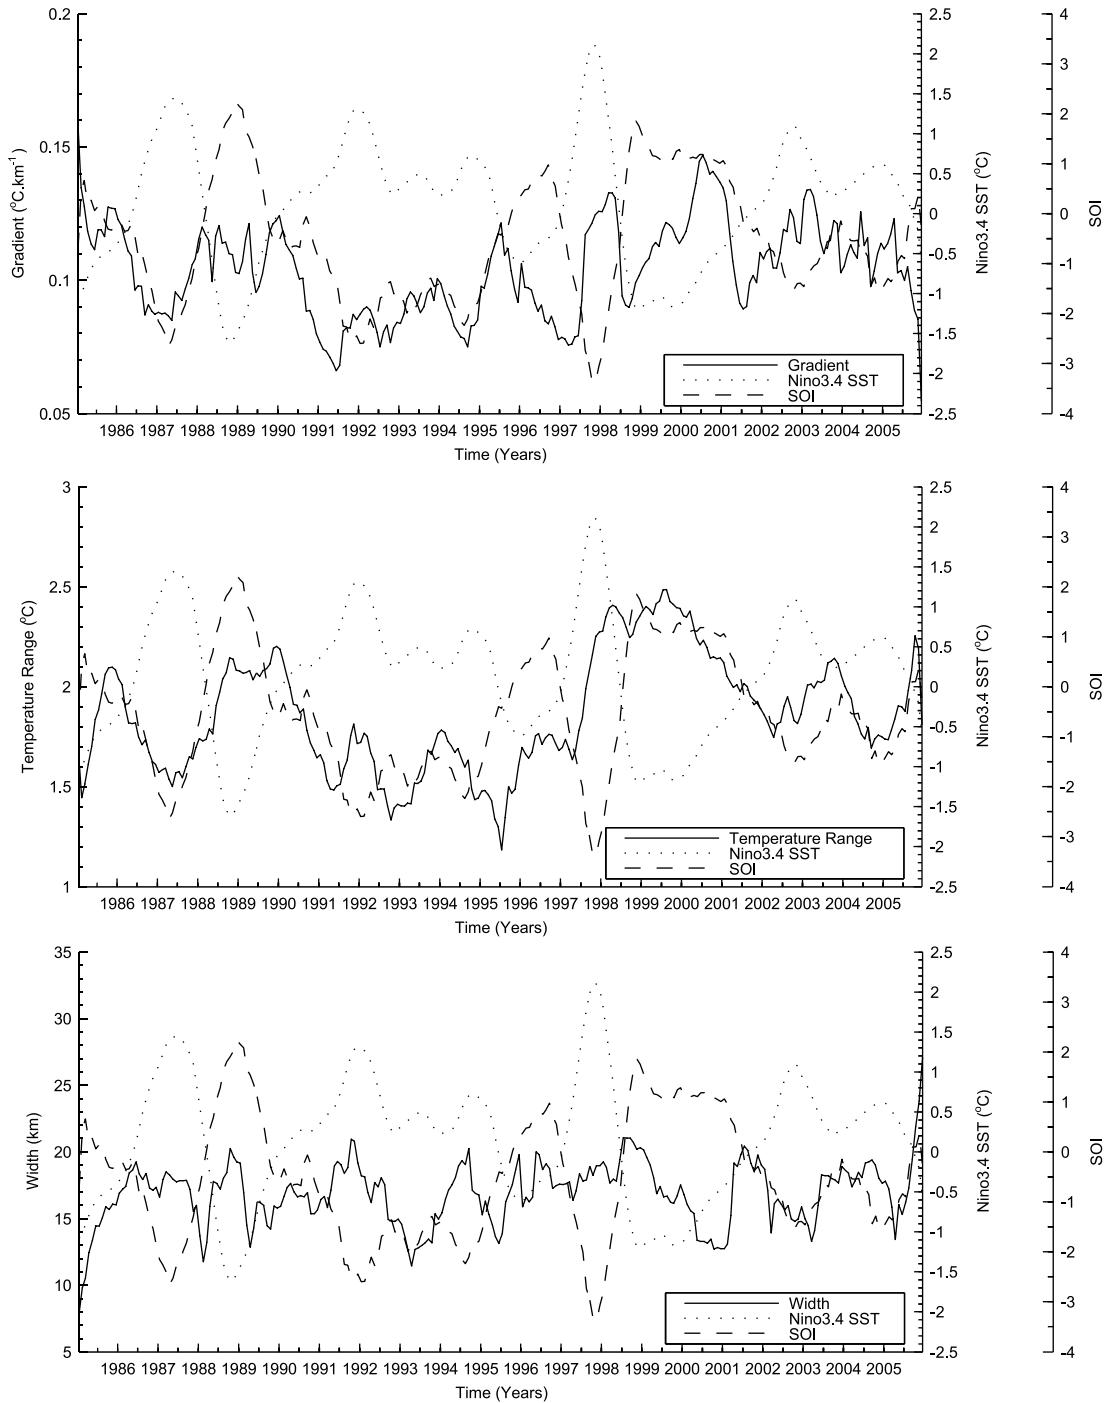


FIGURE 7.16: Twelve month running mean time series of the gradient (top), temperature range (middle) and width (bottom) of the Southland Front against the SOI and Niño3.4 SST.

Generally, where there are rapid changes in the SOI it is the temperature range that controls the response of the gradient. For example, between 1990 and mid 1991 the drop in gradient is driven by a decrease in the temperature difference and not by a sudden increase in frontal width. The 1986/87 decrease in gradient is also the result of a reduced temperature differential. The same control is seen in 1997; the width remains

stable whereas a sharp increase in temperature difference between STW and SAW drives an increase in gradient at the front.

Between 1999 and 2002 when the SOI is positive the pronounced increase in gradient is driven by a decrease in the width. The temperature range at this time remains high but relatively stable following a rapid increase during the previous El Niño. During the 1989 La Niña such a pronounced increase in gradient is not observed. The well defined increase in temperature range however suggests that over estimation of the width masks an increase in gradient.

Post 2002 no definitive trend is observed in the gradient. This is due to the opposing influences of the temperature difference and width. While the temperatures of SAW and STW diverge and converge in parallel with increases and decreases in the SOI, the width of the front also increases and decreases respectively. These trends have countering influences on the gradient which therefore does not appear to respond to ENSO variability.

7.4.2.3 Upper Layer Heat Content, Sea Level and Current Speeds

Consideration of changes in the upper layer heat content of Subtropical Waters, regional sea level and current speeds during ENSO events provides a plausible mechanism by which the gradient of the Southland Front may vary.

Firstly, the more pronounced interannual signal in the temperature range suggests that variability of the thermal gradient at the Southland Front is dominated by the relative temperature difference between the two bordering water masses. If SST anomalies are advected across the Tasman Sea and enter the Southland Current, as suggested in Section 7.4.1.2, in addition to the overall cooling/warming experienced across the region through atmospheric circulation changes, then variation in the temperature difference between the STW flowing equatorward within the Southland Current and the SAW further offshore is expected to predominantly come from changes in the STW. During an El Niño when there is a reduction in SST above the thermocline in the northern Tasman Sea and Australian Bight, cold anomalies may be advected southward. These anomalies reach the Subtropical Front and subsequently flow northward along the east coast of South Island as part of the subtropical Southland Current. The temperature differential and hence the gradient between the SAW and STW at the front therefore decrease. The advection and entrainment of warm ENSO related anomalies during La Niña events would result in the opposite - increasing gradients.

Using our algorithm results we are unable to determine which water mass causes the observed variations in temperature range. Shaw et al. (1999) provide evidence that the STW within the Southland Current is likely responsible. Between 1989 and 1992 an overall decrease of 2-3°C was observed in time series of AVHRR SST's and manually

collected water samples from Otago Harbour, coincident with a decrease in the SOI. The STW sites however cooled by $\sim 0.5^{\circ}\text{C}$ more than the SAW at both 45°S and 46°S , particularly when well within the Southland Current itself. The disparate cooling rates of STW and SAW appear to account for increases and decreases in the gradient and support the advection of ENSO related SST anomalies across the Tasman Sea.

ENSO related changes in ocean heat storage has implications for regional sea level (SL); the SOI is positively correlated with SL variability around New Zealand (Goring and Bell, 1999). In general there is a rise in SL during La Niña derived from increased heat storage in the upper layers of the water column, and a fall in SL during El Niño events resulting from a reduction in regional SST's above the thermocline. The north-south sea surface slope and pressure gradient between the Southern Ocean and subtropics set up the geostrophic easterly flowing circumpolar currents. A localized reduction (increase) in SL would decrease (increase) this pressure gradient and decelerate (accelerate) flows within the global Subtropical Front as it crosses the Tasman Sea. This could potentially widen (narrow) the front and weaken (strengthen) its gradient as it approaches New Zealand.

Finally, changes in upper layer heat content and fluctuations in the gradient may be associated with transport variability of the East Australia Current (EAC). The EAC arises in the Coral Sea and flows down the southeast coast of Australia to Tasmania, ultimately feeding into the Subtropical Front. A portion of the current also flows across the Tasman Sea along the Challenger Plateau where it joins the south coast of South Island and forms the Southland Current. There is evidence to suggest that El Niño events result in reduced flow and transport in the East Australia Current (Sprintall et al., 1995).

In late 1991 and early 1992, strong eastward flow was observed in the equatorial band in the western and central Pacific (Roemmich et al., 1994). As the warm water flowed eastward the surface elevation of the western tropical Pacific dropped, reducing the downward slope to the south along the east coast of Australia. Godfrey (1973) suggests that the EAC is driven by the along shore slope of the sea surface. Thus in 1992 there was a reduced source of warm water for the EAC and a reduced alongshore pressure gradient to drive the southward current. This is confirmed by reduced southward transport between Brisbane and Fiji (Sprintall et al., 1995). Divergence of mass in the upper water column of the southwest Pacific in early 1992 led to upwelling and cold SST anomalies. Reduced current velocities in the region coupled with cooler water being advected into the Southland Current would lead to a widening of the front and a reduction in the temperature difference between STW and SAW and consequently a decrease in gradient. Maximum and minimum geostrophic transports representative of the Indonesian Throughflow were also recorded during the La Niña of 1988-1989 and the El Niños of 1986-1987 and 1991-1994 respectively (Meyers, 1996).

The anomalous increase in gradient during the 1997/1998 El Niño (Figure 7.16) at first glance appears to be an exception to the mechanisms of reduced flows and disparate water mass cooling proposed above. Instead of STW and SAW temperatures becoming increasingly similar they rapidly diverge and the strength of the front is enhanced. This implies that, contrary to the expected cooling, STW's of the Southland Current were anomalously warm. The progression of the 1997/1998 El Niño around New Zealand was in fact unexpected. Instead of the anticipated negative summer SST anomaly of the order of -0.5°C , warmer than usual ocean conditions off the north-east coast were recorded (Goring and Bell, 1999; Basher and Thompson, 1996; Sutton and Roemmich, 2001). There is evidence that since 1997 the northern Tasman Sea has been filling with warm water (Uddstrom and Oien, 1999) because of the combination of a relatively strong East Australian Current advecting warm water into the region and a relatively weak East Auckland Current not removing it (Roemmich et al., 2005). These changes are opposite to the conditions hypothesized by Sprintall et al. (1995) suggesting that this signal may not be related to the ENSO event.

7.4.2.4 Seasonal Correlations

Figure 7.17 shows the cross correlation results for seasonally averaged estimates of the Southland Front's mean temperature, gradient and temperature difference with the corresponding SOI, Niño3.4 and Niño3 SST indices. The gradient and temperature difference are both significantly positively correlated at the 95% level with the SOI, and negatively correlated with Niño SST anomalies when the ENSO indices lag changes at the front by 1-2 seasons (3-6 months). No significant lagged correlations for seasonally averaged estimates of the width are observed (and therefore not shown). A weakening (strengthening) of the gradient is therefore followed 3-6 months later by a decrease (increase) in the SOI and warming (cooling) across the equatorial Pacific. The opposite lag response is true for the mean frontal temperature; ENSO variability and the temperature are most highly correlated when the SOI and Niño SST indices lead by two seasons. The same lag being significant for all examples reinforces confidence that these are real physical connections.

Figure 7.18 reveals that both the strength and sign of correlation between the gradient of the Southland Front and ENSO is seasonally dependent. During the spring, autumn and winter the SOI (Niño SST's) is positively (negatively) correlated with the gradient. During the summer months however this correlation is negative (positive). So outside of the summer, warming across the equatorial Pacific and a decrease in the SOI as an El Niño event develops is associated with a decrease in strength of the Southland Front. Between December and February (summer) the reverse situation occurs - an El Niño episode corresponds to an increase in the frontal gradient.

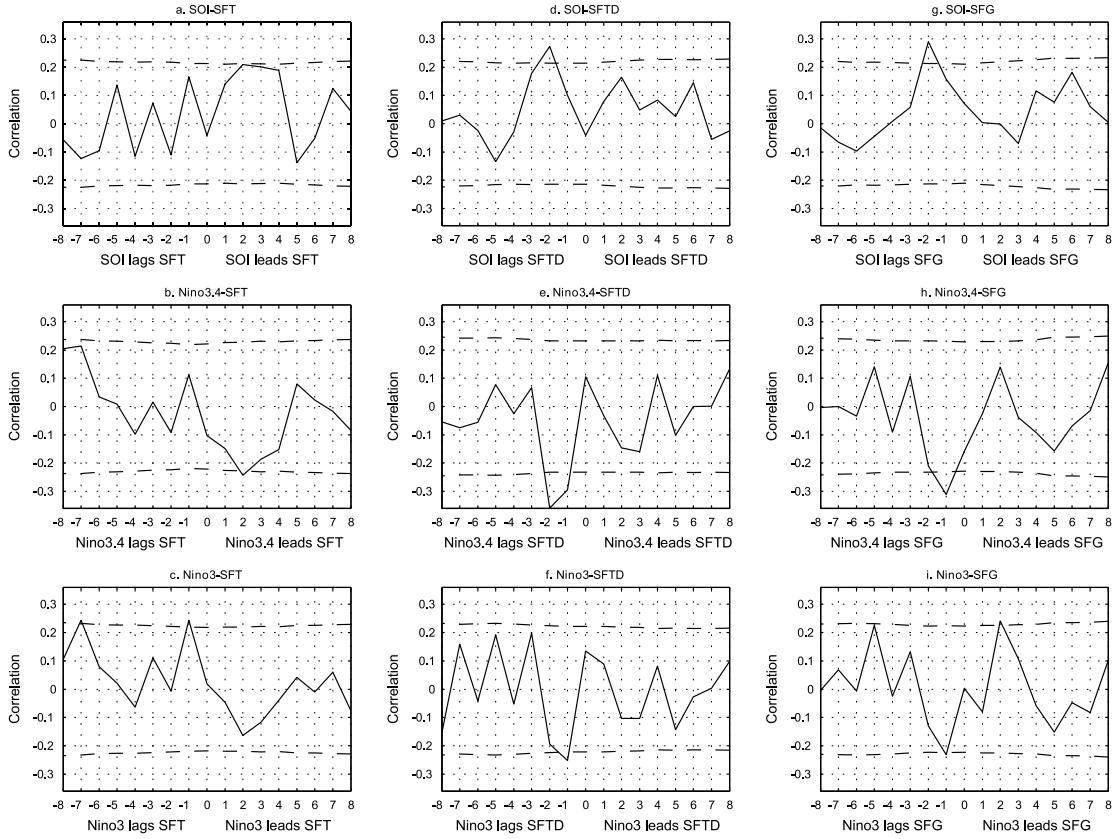


FIGURE 7.17: Cross correlation analysis of seasonally averaged prewhitened time series. Lead and lag times on the x-axis are for seasons. Dashed lines mark the 95% confidence levels. SFT = Southland Front Temperature, SFTD = Southland Front Temperature Difference, SFG = Southland Front Gradient.

The positive relationship between the SOI and gradient is strongest during the spring and winter (0.51 and 0.44 respectively). The negative regional Niño SST-gradient correlations are strongest in the winter. Seasonally stratified one month lag correlations (not shown) are significant at the 95% level during spring and winter for Niño4-gradient correlations, and in the winter months for Niño3.4 and frontal gradient.

Figure 7.18 also lets us explore which component of the gradient - the temperature difference or the width - determines the seasonal correlations. Looking initially at the SOI (Figure 7.18a.), the strong positive spring and winter correlations are matched by significant (95%) positive correlations between the temperature difference and SOI (0.47 and 0.55 respectively). Correlation with the width during these seasons is weak in comparison to summer and autumn. During the summer when the SOI and gradient are negatively correlated there is no correlation between the SOI and the temperature difference across the front. Instead, the width of the Southland Front is weakly positively correlated with the SOI (0.26). So during the summer a decrease in the SOI coincides with a narrowing of the front and a subsequent increase in gradient.

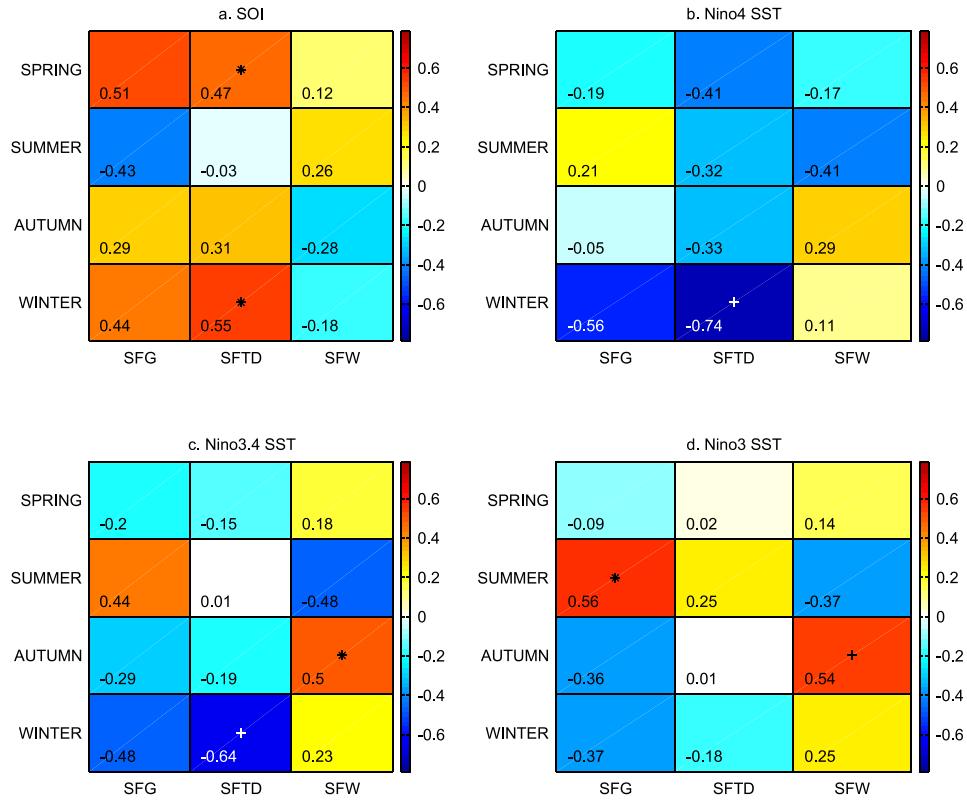


FIGURE 7.18: Zero lag seasonally stratified correlations between the ENSO indices a. SOI, b. Niño4 SST, c. Niño3.4 SST, d. Niño3 SST, and the gradient (SFG), temperature difference (SFTD) and width (SFW) of the Southland Front. * and + indicate significance at the 95% and 99% levels respectively. Only prewhitened data south of 45.5°S has been used in the correlations.

A similar pattern is evident for the Niño SST correlations. In Figures 7.18b and 7.18c the strongest negative Niño SST correlation with the gradient is during the winter. At this time there is a significantly strong (99%) negative correlation between the equatorial SST anomalies and the temperature difference (-0.74 and -0.64), and a much weaker correlation with the width. This suggests that it is changes in the temperature difference between the two Southland Front water masses that dominates ENSO related gradient variability in the winter. During the summer when the correlations with the temperature difference are weakest, especially with Niño3.4, it is the negative correlation with the width that determines the gradient.

The connection between SST anomalies in Niño3 and the gradient of the Southland Front is dominated by changes in the width (Figure 7.18d). There is no correlation with the temperature difference in spring and autumn and a significant (99%) positive Niño3-width relationship (0.54) in the latter of these two seasons. There is also a significantly positive correlation between the width and Niño3.4 during the autumn.

The seasonally stratified lead and lag correlations for the gradient of the Southland Front are shown in Table 7.3. Unlike the temperature, in general, stronger correlations

TABLE 7.3: Seasonally stratified correlations between the gradient of the Southland Front, the SOI and regional Niño SST's. * indicates significance at the 95% level. One season leads (*1 Lead*) are for the ENSO indices leading the gradient at the front. One and two season lags (*1 Lag* and *2 Lag*) are for ENSO indices lagging the gradient at the front. e.g. 2 season winter lag correlations refer to the winter gradient correlated with the following summers ENSO indice.

Season	SOI				Niño4 SST			
	<i>0 Lag</i>	<i>1 Lead</i>	<i>1 Lag</i>	<i>2 Lag</i>	<i>0 Lag</i>	<i>1 Lead</i>	<i>1 Lag</i>	<i>2 Lag</i>
Spring	0.51	-0.02	0.27	0.41	-0.19	-0.30	-0.32	-0.27
Summer	-0.43	-0.34	-0.01	0.32	0.21	0.24	-0.11	-0.35
Autumn	0.29	0.03	0.25	0.17	-0.05	0	-0.13	-0.10
Winter	0.44	0.49	0.58*	0.50	-0.56	-0.44	-0.43	-0.42

Season	Niño3.4 SST				Niño3 SST			
	<i>0 Lag</i>	<i>1 Lead</i>	<i>1 Lag</i>	<i>2 Lag</i>	<i>0 Lag</i>	<i>1 Lead</i>	<i>1 Lag</i>	<i>2 Lag</i>
Spring	-0.20	-0.16	-0.33	-0.41	-0.09	-0.12	-0.30	-0.31
Summer	0.44	0.48*	0.21	-0.35	0.56*	0.46*	0.37	-0.12
Autumn	-0.29	-0.16	-0.20	-0.16	-0.36	-0.19	-0.30	-0.18
Winter	-0.48	-0.37	-0.57*	-0.60*	-0.37	-0.28	-0.51*	-0.60*

with the gradient are not observed when ENSO indices lead by one season. The notable exception to this rule are the summer Niño3.4-gradient and Niño3-gradient correlations. Spring anomalies in the equatorial Pacific are significantly positively correlated with the gradient the following summer.

Winter correlations, with the exception of Niño4, increase and become significant when a one season lag is introduced. For example, the contemporary zero lag correlation between Niño3 and the gradient during the winter is -0.37. When the winter gradient is compared to the following spring Niño3 anomalies this correlation increases to -0.51 and is significant at the 95% level.

7.4.2.5 A Precursor to ENSO events?

Wavelet coherence analysis, cross correlation and plots of the raw time series all suggest that changes in the gradient of the Southland Front preceded ENSO events from anywhere between 3-12 months. The reason behind this is unclear. A possible explanation may lie in the mid-latitude pressure changes in the southwest Pacific that lead extremes of the Southern Oscillation and equatorial Pacific warming (Van Loon and Shea, 1985; Stephens et al., 2007) - the so called 'Van Loon Hypothesis'. These changes are summarized in Figure 7.19.

The position and intensity of the South Pacific trough in the surface westerlies centered over Australia reaches its farthest northward extent and maximum intensity between the

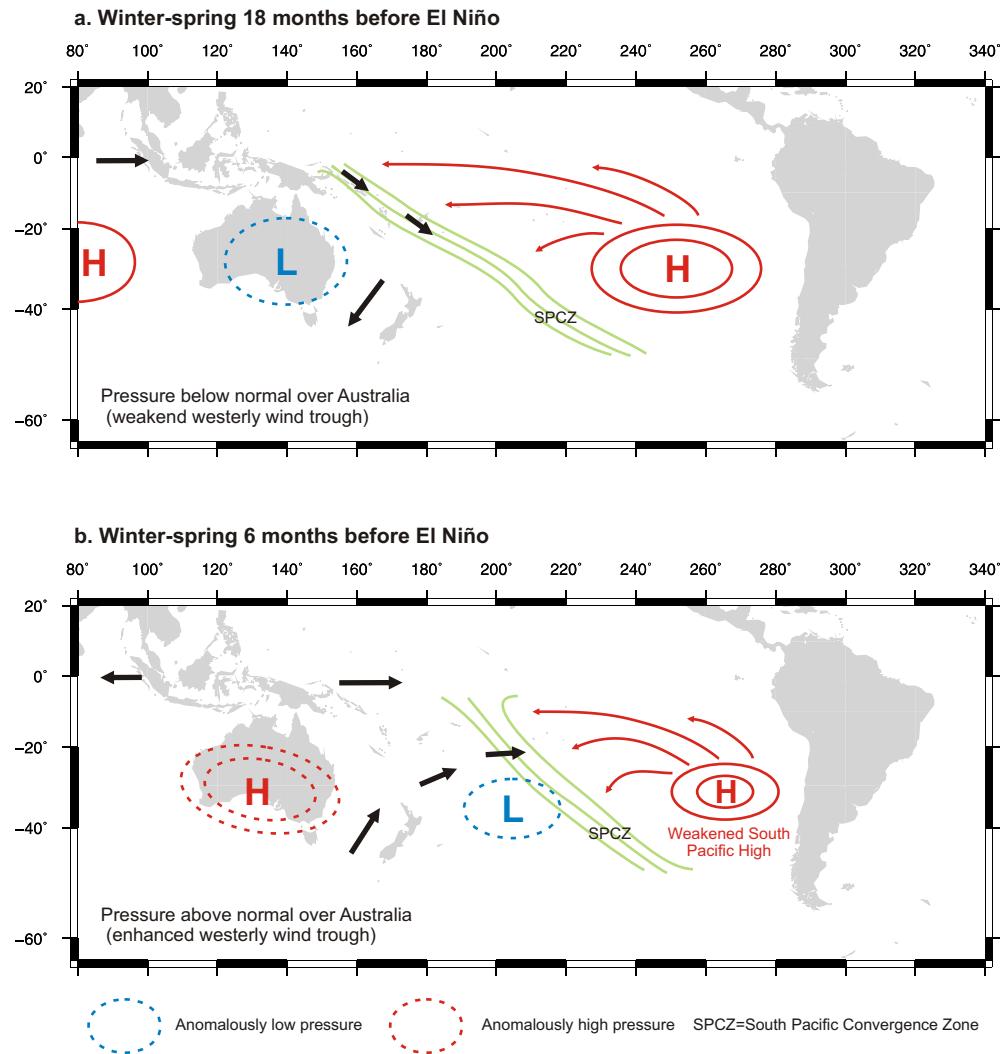


FIGURE 7.19: Schematic showing the pressure and wind anomalies in austral winter-spring in a. the year before an El Niño and b. the year of the El Niño event. High and low SLP are indicated by 'H' and 'L' respectively. Wind anomalies represented by black arrows.

months of May, June and July i.e. austral winter (Van Loon, 1984). During this period of the troughs semi-annual oscillation the South Pacific high and associated trade winds are weakened. Leading into an El Niño, during the two winter seasons (i.e. 18 and 6 months) preceding an extreme in the SOI, that typically occurs in austral summer, the trough changes from a weakened to an enhanced state respectively (Van Loon, 1984; Stephens et al., 2007). Associated with this is an west-to-east pressure reversal in the southern mid-latitudes across the Pacific and a change from anomalous north-easterly to south-westerly winds between 15°S and 45°S, west of 140°W (Van Loon and Shea, 1985). The westerly wind changes are a vital part of a gradual progression of negative SLP anomalies into the wider Pacific that weaken the South Pacific high (Figure 7.19). Synonymous with the different wind anomalies, the SST is higher two winter seasons (18 months) before an El Niño when the trough fails to amplify, than in the winter six months before the event between 15°S and 45°S, from Australia to 140°W, when higher

sea level pressures and southerly wind stress drive colder water toward the western Equatorial Pacific. The opposite is true preceding an extreme La Niña event with an enhanced trough two winters before the cold event, and a depressed trough in the year of the event.

The timing of these circulation anomalies fits well with the significant 1-2 season lagged correlations seen in Figure 7.17 and in particular the strong winter-summer Niño3.4-gradient and Niño3-gradient correlations in Table 7.3. There are peaks at a lag of 5-6 seasons (15-18 months) opposite in sign to the strong two season lag coefficients for all gradient and temperature difference correlations in Figure 7.17. This is consistent with changes in regional SLP 18 months before an El Niño or La Niña event, but since the correlations are not significant a teleconnection at this longer timescale between the Southland Front and ENSO can not be concluded.

As yet, it is unclear how these large anomalies of wintertime SLP and cyclonic activity over the Australian-New Zealand region may affect the gradient of the Southland Front. Nonetheless there is evidence to support teleconnections to an El Niño event that become evident up to two winter seasons preceding the extreme in the SOI. It is left for future research to better understand this mechanism.

7.4.2.6 The Antarctic Oscillation Teleconnection

Interannual variability in the width of the Southland Front does not appear to be strongly coupled to ENSO variability (Figures 7.16 and A.6). Nonetheless fluctuations in the width still influence the gradient of the front and interact with ENSO induced changes in the temperature difference. A second important mode of climate variability that may effect the New Zealand region not yet considered is the Antarctic Oscillation (AAO). Just as the Pacific Decadal Oscillation can modulate ENSO teleconnections so the AAO and ENSO have the potential to partially off-set or enhance their influence on each other and the Southern Hemisphere as a whole.

The AAO, also known as the Southern Annular Mode is the leading mode of climate variability on time scales from intraseasonal to interannual in the Southern Hemisphere. It is characterized by a large scale alternation in surface pressure anomalies between the Indian-Australasian region and the southeastern tropical Pacific, and by a meridional shift in the atmospheric westerly winds. These winds are the driving force behind the Antarctic Circumpolar Current. The AAO index is defined by the leading principal component of the 700mb geopotential height south of 20°S in the atmosphere. Positive phases of the AAO are associated with negative pressure anomalies over the mid-latitudes, and a strengthening of the zonal atmospheric pressure gradient. The enhanced pressure gradient strengthens westerly winds at $\sim 55^{\circ}\text{S}$ (over the Antarctic and Polar Fronts), but weakens the westerlies (easterly anomalies) at $\sim 35^{\circ}\text{S}$ (Lovenduski

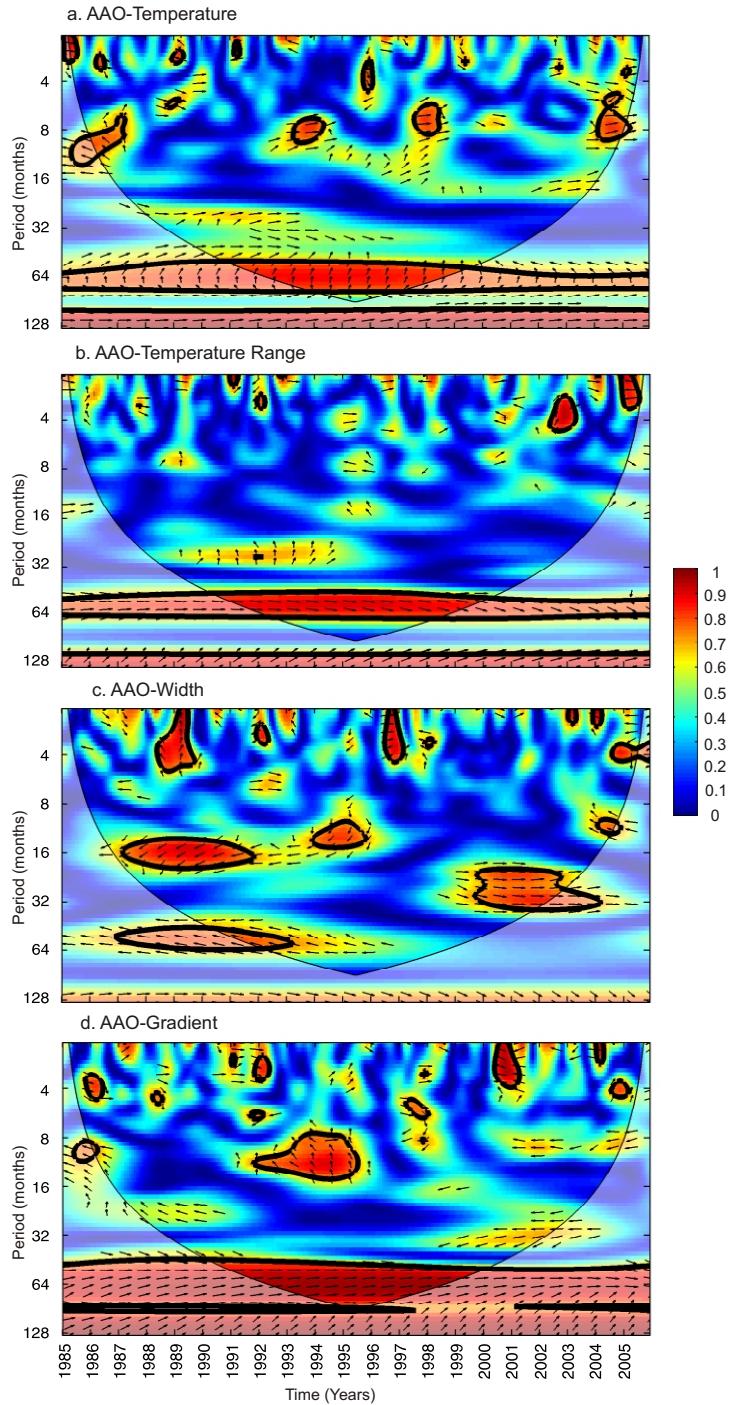


FIGURE 7.20: Wavelet squared coherence between the Antarctic Oscillation (AAO) and a. Temperature , b. Temperature Range, c. Width, d. Gradient of the Southland Front.

and Gruber, 2005). Surface easterlies result in southward Ekman transport anomalies in the Subtropical Zone and therefore warm SST anomalies. This SST response is most pronounced during austral summer. AAO related wind stress variability influences ACC transport and frontal positions (Meredith et al., 2004; Fyfe and Saenko, 2006; Hall and Visbeck, 2002). We should therefore also consider the influence that changes in the AAO

may have on the temperature and strength of the Southland Front and how these may interact with ENSO variability.

Figure 7.20 identifies the significant periods of coherence between the AAO and the temperature, width, temperature range and gradient of the Southland Front. There are bands of significantly high coherence that persist across the entire data set for the temperature, temperature range and gradient centered around an oscillation period of five years. This is suggestive of a lower frequency teleconnection on timescales greater than individual ENSO events. Coherence between the AAO and the width of the front at periods of approximately five years is only evident between 1987 and 1992. However, the AAO and width are at times significantly correlated at higher frequencies, a connection that is not observed in the SOI-width wavelet coherence (Figure A.6). Between 1987 and 1992 the AAO and width are coherent at periods of 16 months. Later, between 1999 and 2004 oscillations with periods of 16-32 months are correlated. It is possible that the AAO-width relationship at these times complicates and modifies the response of the Southland Front gradient to ENSO events.

7.4.3 Location of the Southland Front and ENSO

Figure 7.21 examines interannual variability in the position of the Southland Front, relative to its twenty-one year mean (as seen in Figure 6.2, Chapter 6). Also plotted is a measure of the front's average meandering intensity; the root-mean squared spatial displacement (RMSSD).

There is much greater interannual variability in the RMSSD than in the position of the front. Neither however appear to be related to changes in the SOI. The most pronounced deviation in position takes place during 2001 when the front is displaced offshore by more than 10km from its twenty-one year mean. This coincides with an increase in meandering intensity suggesting a decrease in stability. The movement does not coincide with a strong ENSO event but does however match a peak in the AAO index (Figure A.7).

Wavelet coherence confirms that no persistent teleconnection exists between the position and stability of the Southland Front and ENSO, or the AAO. There are small patches of significant coherence but nothing to suggest a consistent physical relationship. Modelling studies suggest that the current structure of the ACC should move poleward in response to a strengthening of the zonal winds during positive AAO periods (Fyfe and Saenko, 2006). It is likely that the strong bathymetric control of the continental slope over the position of the front inhibits such a response at this particular location.

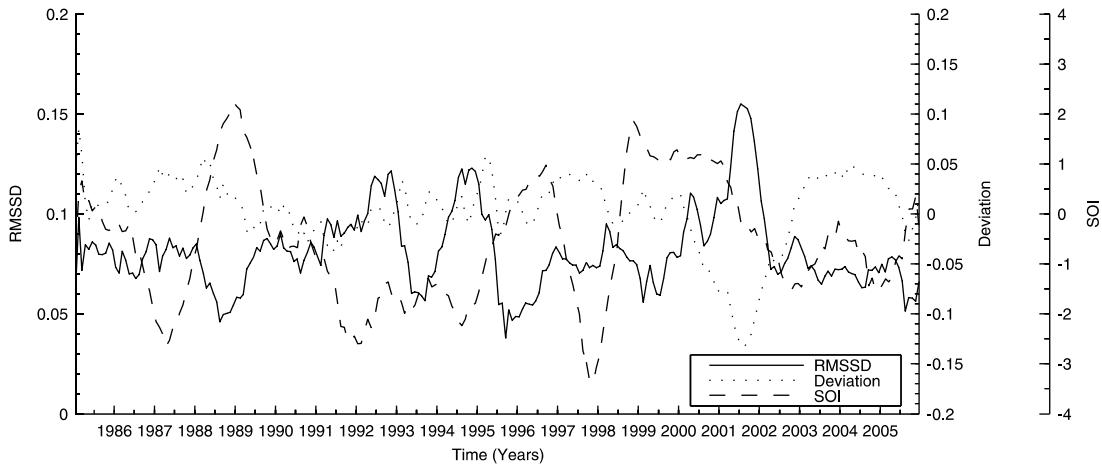


FIGURE 7.21: Twelve month running mean time series of the front's deviation from its mean twenty-one year position, its root-mean squared spatial displacement (RMSSD), a measure of stability, and the SOI. Positive (negative) deviation is shoreward (seaward) of the mean position.

7.5 Summary

ENSO related atmospheric circulation and climate changes (air temperature and precipitation) around New Zealand are already well documented, as are the SST anomalies that dominate interannual variability of the surrounding waters, particularly in the north. In this study we have presented the first in-depth investigation of the connection between the Southland Front, a major oceanographic feature off the east coast of South Island, and ENSO activity.

The El Niño-Southern Oscillation dominates interannual changes in the temperature of the Southland Front in the 8-36 month band. El Niño episodes are associated with cooling at the Southland Front and La Niña events result in warm anomalies. The teleconnection between ENSO and the mean temperature of the front has a strong seasonal dependence. We find the temperature to be significantly positively correlated with the SOI, and negatively correlated with the Niño SST regional anomalies during the summer. In the winter no strong relationship is detected. Summer stratification results in surface waters being more sensitive to anomalous winds and heat fluxes. Additionally, the summer pattern of SOI-MSLP correlation favors enhanced radiative and up/downwelling effects from localized cyclonic/anticyclonic circulation features centered over South Island.

Maximum and minimum SST anomalies at the front lag the peak El Niño and La Niña events by approximately 3-4 months (one season). The spring SOI is significantly positively correlated with summer temperatures at the front. Likewise, there is a significant negative spring-summer and summer-autumn correlation between the Niño SST's and the mean temperature of the front. A time lag of this magnitude can not be attributed to atmospheric teleconnections. Instead, we propose that SST anomalies in the Coral and

Northern Tasman Seas, where the greatest ENSO related SST anomalies are recorded, are advected across the Tasman Sea and reach the Southland Front approximately one season later.

The phase lag and magnitude of frontal temperature changes is inconsistent. There appears to be a change in the phase relationship around 1994 that is coincident with a brief period of negative PDO values that reverses the correlation (negative to positive) between the SOI and PDO. We suggest that the PDO is capable of modulating the magnitude of ENSO related SST anomalies at the Southland Front and disrupting the timing and mechanics of the teleconnection. More evidence however is needed to conclusively determine whether SST anomalies are enhanced or dampened.

The gradient (strength) of the Southland Front is also modified by ENSO variability. Confirmation that the gradient of the front decreased during the 1991 El Niño event as suggested by Shaw and Vennell (2001) is confirmed. This is suggestive of a decrease in velocity of the northward flowing Southland Current that ultimately sets up the global Subtropical Front across the open Pacific Ocean. A well defined decrease in gradient is also observed in response to the 1987 El Niño. Conversely during the 1997/1998 event the gradient of the front increases, indicative of increased geostrophic velocities along the continental slope. There is also evidence for increased gradients during La Niña events. The temperature difference appears to be the main driving force behind changes in the gradient. The width of the front is more coherent with the AAO and may modify ENSO related gradient fluctuations.

ENSO and gradient of the Southland Front are significantly coherent in the 8-36 month band. Just as correlations between ENSO indices and the mean temperature of the front are seasonally dependent, the relationship between ENSO and the gradient varies throughout the year. The strongest positive (negative) correlations between the SOI and the gradient (Niño SST's and the gradient) occur in the winter. This means that during an El Niño period when the SOI is negative and warm SST anomalies are observed across the central equatorial Pacific the gradient at the Southland Front is weaker than normal (*vice versa* for a La Niña). Interestingly, during the summer, these correlations are reversed suggesting that an El Niño event should result in a strengthening of the front. The winter correlations are the result of a strong connection between ENSO and the temperature difference across the front, whereas during the summer, the width has a much stronger relationship to ENSO variability.

Changes in the upper ocean heat content and advection of SST anomalies across the Tasman Sea and into the STW's of the Southland Current is a possible mechanism for ENSO related gradient variability. There is some evidence to suggest that changes in the temperature difference between STW and SAW is driven mainly by changes in STW temperatures. Further data is needed to confirm this. Positive correlation between the SOI and mean sea level around New Zealand may also influence the strength of the

Southland Front. A drop in sea level in the northern Tasman Sea and Australian waters may reduce the north-south pressure difference driving circumpolar flows sufficiently enough to weaken the Subtropical Front as it crosses the southern Tasman Sea. Reduced transport of the East Australian Current and Indonesian Throughflow have previously been linked to movement of the warm water pool over Indonesia during El Niño events and subsequent reduction in the sea level and associated current transport.

Wavelet analysis, cross correlation and an examination of the raw time series all show that the gradient leads peak ENSO events by 3-12 months. Why the gradient of the front should change so far in advance of El Niño and La Niña events, especially given the delayed response (3-4 months) of the frontal temperature to SOI and Niño SST anomalies, remains unclear. Further investigation into the atmospheric changes around New Zealand associated with the westerly wind trough 1-2 winter seasons in advance of extreme summer SOI indices is necessary. Seasonally stratified lagged correlations reveal a strong winter-summer correlation between the temperature of the front (winter) and equatorial SST's (summer).

Interannual variability in the position of the Southland Front is small owing to the strong topographic control of the continental slope. Fluctuations in stability are more pronounced but do not appear to be correlated with El Niño or La Niña events. ENSO's influence on the position of the Subtropical Front would be better investigated in areas of open ocean away from topographic features such as the Crozet, Kerguelen and Amsterdam Plateaus that limit latitudinal variability.

In conclusion, teleconnections between the Southland Front and the ENSO cycle are by no means simplistic or necessarily predictable. Different correlation patterns and air-sea interactions operate at different times of the year and the front's response to these changes is also seasonally dependent. Each ENSO event is different; the order and strength in which the ocean and atmospheric components of ENSO develop and peak do not remain constant. The speed and timing of anomalies at the Southland Front are therefore also expected to vary. Furthermore, other modes of climate variability such as the PDO and AAO may modify the timing and magnitude of teleconnections. Nonetheless, we have provided in this study the first substantial evidence that given the correct conditions the strength of the Subtropical Front as it passes New Zealand may be modified by ENSO activity.

CHAPTER 8

Conclusions and Further Work

The main objectives of this research were two fold; firstly to develop a technique capable of mapping the position, strength and temperature of major ocean fronts from remotely sensed SST data; secondly to use the front detection algorithm to investigate possible teleconnections between the Southland Front, New Zealand and ENSO activity. In this final chapter we return to the scientific problems which motivated this work, summarize the main achievements and findings, interpret our results within the broader scientific context and suggest some possible extensions.

8.1 Statistical Modelling of Ocean Fronts

The first half of this thesis is dedicated to the development of a new front detection tool able to provide a sufficiently detailed data set with which to study the interaction between ocean fronts and modes of atmospheric variability. We extend the idea of statistical model fitting by using a weighted local likelihood approach to provide a non-parametric description of spatial variations in the position and strength of ocean fronts from remotely sensed SST images.

The rapid change in SST between two different water masses on either side of a front is modelled using an S-shaped (tanh) function. The unknown parameters of this model are determined by maximizing a weighed sum of likelihood contributions from all available cross sections of SST observations in an image. A Gaussian smoothing function assigns weightings based on the distance of observations from the point of estimation.

The bandwidth of the kernel function determines the smoothness of the fitted non-parametric regression, with larger values corresponding to stronger levels of smoothing. Likelihood cross validation is used to determine the optimal level of smoothing best supported by the available data. A bandwidth of 0.095° ($\sim 10.6\text{km}$), the median value identified in this study, assigns 95% of weight to observations within a 21.6km window centred about the point of interest. The weightings assigned by the kernel smoother are modified by an additional weighting based on an assessment of the quality of the likelihood contributions from each set of SST measurements. The maximum local likelihood estimates are obtained via a Newton-Raphson optimization scheme. The optimization is scaled and bounded based on *a-priori* knowledge about the front being studied. Estimation uncertainty is quantified by standard errors calculated from the variance-covariance matrix of each local likelihood solution.

The limitations and sensitivity of the algorithm were investigated using artificially generated images and by direct comparison to alternative front detection techniques; primarily the ‘front following’ algorithm of Shaw and Vennell (2000a). Estimates of frontal position, temperature and temperature range made by the local likelihood algorithm are both realistic and robust to sparse and noisy data sets, and to the orientation of the front. This solves a problem that has hampered previous front detection schemes. The width of the front is the most difficult and sensitive parameter to resolve during optimization. The quality of an estimate is heavily dependent upon not only the number of available SST measurements but also on the way in which these observations are distributed within the image.

Applications of the algorithm developed here are numerous. Climate variability and monitoring studies, ocean forecasting, validating ocean models, ecosystem and fisheries research all require an accurate understanding of the spatio-temporal behavior of ocean frontal systems. Specifically, the algorithm may be used by fisheries agencies to help correlate oceanic fronts with fish catches and thereby put in place suitable fishing regulations. The detailed results from such a front detection scheme would compliment *in-situ* data sets where variability of a front is difficult to resolve. It may also help put into a wider context findings from research cruises only able to sample a limited area.

The algorithm introduced here is adaptable to a wide variety of remotely sensed data sets (e.g. altimetry and ocean colour), model output (e.g. OCCAM) and data assimilation products (e.g. GHRSSST). Ocean colour and altimetry data may be used simply by adopting a model function that describes the change in colour/height across a front rather than the temperature. In situations where: a. the surface thermal signature of the front is masked by diurnal/seasonal heating or, b. no change in SST gradient is present (e.g. estuarine plume fronts) then measurements of chlorophyll/sediment concentration or changes in relative sea surface height may be more reliable sources of data with which to locate ocean fronts.

Seasonal variability in the surface layer of the water column is a significant disadvantage of using SST images to monitor the position and strength of ocean fronts. For example, in the summer the thermal expression of the Subtropical Front as it crosses the Tasman Sea is covered by a 50-70m thick mixed layer, although it continues to exist in the permanent thermocline below (James et al., 2002). Additionally, periodic wind forcing moving the Ekman layer back and forth creates eddies and inversions that shear the mixed layer off from the underlying structure. In these conditions estimates of a front's position become unreliable.

Traditionally, both temperature and salinity are used to classify different water masses and together determine the density of sea water. In turn, density is an important factor driving ocean circulation and frontal dynamics. The salinity signature of water masses is in fact often a more appropriate indicator of a front's position. The thermal expression of the Subtropical Front in the mid-South Atlantic and South Pacific may be more easily distinguishable through its salinity distribution, which is more stable than the temperature (Lutjeharms et al., 1993; Stramma et al., 1995). Sea surface salinity (SSS) is to be measured by sensors flown on two forthcoming missions: ESA's Soil Moisture and Ocean Salinity (SMOS) mission and the joint NASA/CONAE mission called Aquarius, both scheduled for launch in 2009. The availability of salinity measurements in the future will allow the distributions of both key properties (SST and SSS) to contribute to the estimation of frontal position.

There are a number of improvements that could be made in future versions of the algorithm. The bandwidth of the smoothing function used to process each image is determined automatically using likelihood cross validation. An improvement on this scheme would be to introduce a variable bandwidth within each image. This would allow the smoother to adjust to localized variations in the density, distribution and quality of SST observations. For well defined structures that are not obscured by small scale clouds a smaller bandwidth would be favored, minimizing bias in parameter estimates. Where the frontal structure is poorly defined or masked by cloud then a larger bandwidth would increase the weightings assigned to data further afield and decrease variance in the parameter estimates. The along-front structure and variability of each parameter is different. Adopting different smoothing functions and bandwidths for each parameter may therefore be advantageous. The local likelihood framework however does not naturally allow the degree of smoothing applied to each parameter to be controlled because smoothing takes place in likelihood rather than in parameter space. Further investigation is needed to assess the possibilities of incorporating such flexibility.

The local likelihood algorithm only uses spatial information to estimate the position and characteristics of fronts in regions occluded by cloud cover. Incorporating temporal information about the structure and local motion of the front from images taken before and after the time of interest could improve estimates made in regions of very sparse data. If the position of a front is well defined at times $t - 1$ and $t + 1$, but obscured

by clouds at time t , then estimates made at $t - 1$ and $t + 1$ could be used to constrain optimization at time t . Accounting for temporal evolution would be most effective when using daily or weekly images. Natural meso-scale variability (meanders, eddies) and the temporal smoothing inherent in monthly composite images means that the position and structure of a front in one month will not necessarily bear any resemblance to the location of features in the previous or following months. The advection and deformation of features over daily and weekly time scales is likely to be significantly less which would allow specific structures to be detected and matched between time frames.

Under the current design, changes in the orientation of the front can lead to overestimation of the width and subsequent underestimation of the gradient. Incorporating some form of angular dependency into the algorithm will be a high priority during future developments and will help to stabilize and correct estimates of the width. For angles less than 80° we have shown that the location of the front can be successfully estimated regardless of orientation. Knowing the relationship between orientation and the error in the estimated width, overestimation could be corrected for post-processing. A better solution would be to extract data \mathbf{Z}_k at variable angles along the front during optimization. A standard edge detection technique could be used to locate the front and estimate its orientation. Alternatively, the standard maximum likelihood model fit (no smoothing kernel) at each position k could be optimized with respect to the angle at which SST observations \mathbf{Z}_k are extracted across the front to obtain the minimum possible estimate of the width. This information could then be used to extract vectors \mathbf{Z}_k at optimal angles during the full local likelihood calculation. Both of these solutions however would result in a certain number of SST observations being used more than once and assumptions of independence violated. Even so, it may be more appropriate to sacrifice precise statistical correctness in order to gain more realistic estimates of the front's strength.

Currently the model function allows for only one front to be resolved. In reality of course fronts may bifurcate and then merge back together, temporarily creating a double frontal structure. Adding a second S-shaped function to the model in theory allows two fronts in an image to exist. The possibility of allowing multiple fronts was investigated during testing and development of an early version of the algorithm where the optimal solutions at each position were all interdependent. This condition resulted in a very large, complex and unstable optimization problem for which no satisfactory solution could be found. The final local likelihood solution presented here is less complex and more effective than the initial approach and the re-introduction of a double frontal structure more feasible.

8.2 Southland Front Variability

Information collected using the algorithm enabled the spatial and temporal characteristics of the Southland Front between January 1985 and December 2005 to be determined. More importantly, work here has contributed to better understanding the variability of the front in response to remote ENSO forcing. Four scientific questions posed in Chapter 1 were addressed and we summarize here the main findings relating to each of them:

1. *How and why do the characteristics of the Southland Front change spatially and seasonally?*

Over the twenty-one year period studied the front had a mean temperature of $10.4 \pm 1.8^\circ\text{C}$ ($\pm 1\text{s.d}$) and was found to be locked to the 560m contour of the continental slope. The difference in temperature between SAW and STW on either side of the front was $1.75 \pm 0.51^\circ\text{C}$, a transition that took place on average over a distance of $18 \pm 10.7\text{km}$. The mean thermal gradient of the front over this time period was therefore estimated to be $0.1^\circ\text{C}.\text{km}^{-1}$.

In general, the width of the front increases northward and the temperature difference between the STW and SAW on either side decreases. This leads to an overall decrease in gradient as the front approaches the open Pacific basin and suggests a decrease in the velocity of the Southland Current, a partially geostrophically driven flow. The meandering intensity of the front increases as it flows northward along the east coast of South Island indicating that the bathymetric control of the continental shelf weakens where isobaths start to diverge. This decrease in stability is supported by the northward increase in width.

All frontal characteristics are seasonally variable. Overall, the front is widest during the spring and narrowest in the summer and winter. The greatest change in temperature across the front occurs during the summer and may reach 2.6°C in the Canterbury Bight. The minimum differential is observed during the winter months. The slightly larger annual SST range of SAW and the phase lag of 2-3 days behind STW further inshore accounts for these seasonal temperature differences. The front is overall strongest in the summer peaking locally at 46°S . North of 45.5°S the front is also strong during the winter. The weakest gradients occur during the spring.

Seasonally the front is located furthest inshore during the summer and furthest offshore in the winter. The potential masking of STW within the Southland Current during the spring and autumn months however introduces some uncertainty into the reliability of estimates made during these particular seasons.

2. *Is there a persistent teleconnection between the temperature of the Southland Front and ENSO?*

El Niño (La Niña) events are associated with cool (warm) SST anomalies at the front. On interannual timescales ENSO dominates the variability at periods of 8-36 months. The mean temperature of the front is positively correlated with the SOI and negatively correlated with Niño SST anomalies during the summer when the stratified surface waters are most sensitive to localized anomalous winds and heat fluxes. This teleconnection does not appear to exist during the winter months and is weaker in the spring and autumn.

Maximum and minimum SST anomalies lag peak El Niño and La Niña events by 3-4 months (one season). A possible explanation may be the advection of SST anomalies across the Tasman Sea and into the Subtropical Front and Southland Current. This provides an oceanic as well as an atmospheric coupling mechanism.

The phase relationship between ENSO indices and the frontal temperature is inconsistent. This indicates that there is no one simple physical mechanism governing interannual temperature fluctuations. The PDO potentially has a modulating effect on the timing and magnitude of ENSO related SST anomalies at the Southland Front. There is a distinct phase shift in the relationship between the frontal temperature, the SOI and Niño SST's around 1994 coincident with a brief negative PDO period.

3. Is the strength of the Southland Front modified by ENSO variability?

ENSO and the strength of the Southland Front are significantly coherent in the 8-36 month period band. During the 1987 and 1991 El Niño events the gradient of the front decreased implying a reduction in geostrophic velocity. In contrast, the strong 1997 El Niño is associated with an increase in gradient, although this may be due to non ENSO related mechanisms.

There is a strong seasonal dependence to the correlation between the gradient and both atmospheric and oceanic ENSO indices. El Niño (La Niña) events during the winter are associated with a decreasing (increasing) gradient that is dominated by a decrease (increase) in the temperature difference between the STW and SAW. During the summer, El Niño (La Niña) episodes correspond to an increase (decrease) in frontal strength driven by a decrease (increase) in the width of the Southland Front.

Possible mechanisms for ENSO related gradient variability include: changes in upper ocean heat content in the Northern Tasman Sea and advection of subtropical SST anomalies into the Southland Current, and sea-level pressure changes modifying transport of the East Australian Current and flows within the global Subtropical Front.

As for the frontal temperature, the phase relationship between ENSO and the front's strength is variable. In general, peak gradient anomalies are found to lead ENSO events by 3-12 months. The temperature range and gradient are both significantly positively (negatively) correlated with the SOI (Niño SST's) at a 1-2 season lead. There is a strong negative winter-summer correlation between the gradient and SST's across the

Pacific. This fits well with observed regional atmospheric circulation changes six months proceeding a extreme summer SOI event.

ENSO related variability may be masked or modified by the influence of other modes of climate variability such as the Antarctic Oscillation.

4. *Is there any interannual variability in the position and stability of the Southland Front?*

The location of the Southland Front shows minimal interannual variability. The most pronounced deviation in position occurs during 2001 but this does not coincide with a strong ENSO event. The stability of the front displays a higher degree of variability but this is not well correlated with ENSO activity. It is possible that in this region the strong bathymetric control of the continental slope limits any substantial fluctuations in position.

The unique nature of the timing, strength, and duration of each ENSO event, coupled with the complex interactions of other modes of variability, and the strong seasonal dependence of the oceanic response to atmospheric anomalies means that the exact response of the Subtropical Front to El Niño and La Niña episodes is difficult to predict. Nonetheless, this study has conclusively shown for the first time that a real connection exists between the gradient of the Southland Front and ENSO variability. We suggest that the strength of the Southland Front may potentially be used as a precursor to extremes in the Southern Oscillation. The physical mechanism by which circulation changes around New Zealand preceding ENSO events modify the strength of the front remains unclear and demands further investigation using a wider range of oceanographic (e.g. altimetry, *in-situ* observations) and meteorological data (e.g. winds, pressure). Additionally, interannual changes in the relative geostrophic velocity of the Southland Current need to be verified.

The wider implications of interannual fluctuations in the strength, stability and location of the Southland Front with respect to: the structure and strength of the Subtropical Front as it enters the open Pacific Ocean; the consequences for CO₂ drawdown, particularly over the Chatham Rise; and feedbacks into the wind field, which may in turn modify ocean current structure, is left for future research.

As the Subtropical Front passes New Zealand it is tightly constrained by the bathymetry. Consequently, any variations in its location in response to ENSO are less likely to be detected. Investigations of this nature are much better suited to open ocean sections of the front, away from pronounced topographic features such as the Kerguelen and Crozet Plateaus, where the front meanders more freely over the deep abyssal plains.

The focus of Chapter 7 has been primarily on links to the ENSO cycle. Although the Antarctic Oscillation was briefly considered, less attention has been paid to the rich variety of atmospheric modes over the Southern Ocean and the complex interaction

between them. Further consideration into how the Antarctic Oscillation, the Antarctic Circumpolar Wave, and the semiannual oscillation may drive changes at the Subtropical Front is necessary. Furthermore, the potential for ENSO related Rossby and Kelvin waves (Jacobs et al., 1994; Galanti and Tziperman, 2003) to introduce variability at the front is an area worth investigation.

Numerous studies using the algorithm developed here are possible. Of particular personal interest is the variability of Southern Ocean fronts in relation to long term wind shifts driven by modes of climate variability in the Southern Hemisphere such as the Antarctic Oscillation (AAO). Modelling studies have demonstrated that latitudinal shifts and modulation in the strength of the wind stress are capable of generating ocean circulation variations and temperature fluctuations at intermediate water depths that are tightly in phase with the AAO (Hall and Visbeck, 2002; Oke and England, 2004). Antarctic circumpolar transport should significantly increase and the current structure move poleward in response to a strengthening of zonal winds during positive phases of this climatic mode (Fyfe and Saenko, 2006). Using two decades worth of tide gauge measurements Meredith et al. (2004) provide the first observational evidence that interannual changes in transport variability through Drake Passage are forced by the Antarctic Oscillation. The front detection technique developed here is an ideal tool with which to investigate AAO-frontal relationships further. High resolution temporal and spatial variability in the front's location may be monitored over vast distances and long periods of time. In addition, following Dong et al. (2006) the SST gradient may potentially be used as a proxy for zonal velocity and transport at the fronts.

APPENDIX A

Appendix

A.1 Hessian matrix for local likelihood estimates

The diagonal elements of the hessian matrix from each local likelihood calculation at position x_j are given by:

$$\begin{aligned}
 \frac{\partial^2 L}{\partial \theta_{1,j}^2} &= \sum_{k=1}^m K \cdot -\frac{n_k}{\sigma_j^2} \\
 \frac{\partial^2 L}{\partial \theta_{2,j}^2} &= \sum_{k=1}^m K \cdot -\frac{1}{2\sigma_j^2} \sum_{i=1}^{n_k} 2T^2 \\
 \frac{\partial^2 L}{\partial \theta_{3,j}^2} &= \sum_{k=1}^m K \cdot -\frac{1}{2\sigma_j^2} \sum_{i=1}^{n_k} \frac{2\theta_{2,j}^2}{\theta_{3,j}^4} S^2 - \frac{4\theta_{2,j}(\theta_{4,j} + y^{(i)})}{\theta_{3,j}^3} SR + \frac{4\theta_{2,j}(\theta_{4,j} + y^{(i)})^2}{\theta_{3,j}^4} STR \\
 \frac{\partial^2 L}{\partial \theta_{4,j}^2} &= \sum_{k=1}^m K \cdot -\frac{1}{2\sigma_j^2} \sum_{i=1}^{n_k} \frac{2\theta_{2,j}^2}{\theta_{3,j}^2} S^2 + \frac{4\theta_{2,j}}{\theta_{3,j}^2} STR \\
 \frac{\partial^2 L}{\partial \sigma_j^2} &= \sum_{k=1}^m K \cdot \frac{n_k}{\sigma_j^2} - \frac{3}{\sigma_j^4} \sum_{i=1}^{n_k} R^2
 \end{aligned}$$

where

$$T = \tanh \left[\frac{y^{(i)} + \theta_{4,j}}{\theta_{3,j}} \right] \quad S = \operatorname{sech} \left[\frac{y^{(i)} + \theta_{4,j}}{\theta_{3,j}} \right]^2 \quad K = K(x_k - x_j; h)$$

and

$$R = z_k^{(i)} - \left(\theta_{1,j} + \theta_{2,j} \tanh \left[\frac{y^{(i)} + \theta_{4,j}}{\theta_{3,j}} \right] \right)$$

A.2 AVHRR Pixel Quality

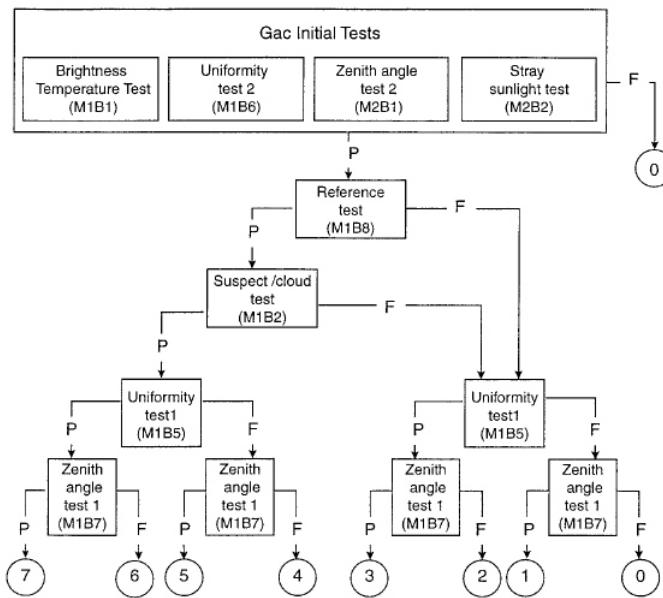


FIGURE A.1: Combination of individual quality tests to derive an overall pixel quality level. From Kilpatrick et al. (2001). Tests are described in Table A.1

TABLE A.1: Description of individual quality tests used to derive AVHRR pixel qualities. Full details are given in Kilpatrick et al. (2001).

Quality Test	Description
Brightness Temperatures	BT's for AVHRR channels 3,4, and 5 must be $\geq -10^{\circ}\text{C}$ and $\leq 35^{\circ}\text{C}$. Identifies sensor errors or cold pixels associated with high cloud.
Cloud Test	Initial cloud-flagging tests include absolute thresholds for brightness temperature, radiance values and albedo, and thresholds for differences between max and min values within a 3×3 pixel extraction box for each channel. Further screening is refined for each satellite and year.
Uniformity Test 1	Difference between max and min BT's for channels 4 and 5 in a 3×3 pixel window must be $< 0.7^{\circ}\text{C}$. Identifies contamination by small scale clouds.
Uniformity Test 2	Same as Uniformity Test 1 but for a threshold of 1.2°C .
Zenith Angle Test 1	Zenith angle must be $< 45^{\circ}$. At higher zenith angles radiation emitted by the ocean travels further through the atmosphere before reaching the AVHRR sensor. This increases attenuation of the ocean signal and a greater proportion of radiance reemitted from the atmosphere makes up the received radiance.
Zenith Angle Test 2	Same as Zenith Angle Test 1 but for an angle $< 55^{\circ}$.
Reference Test	The absolute difference between the Pathfinder SST and the 3 week Reynolds reference SST must be $\leq 2^{\circ}\text{C}$.
Stray Sunlight Test	This flag identifies Earth-Sun-satellite configurations under which stray sunlight may contaminate pixels.

A.3 Practical Aspects of Front Detection

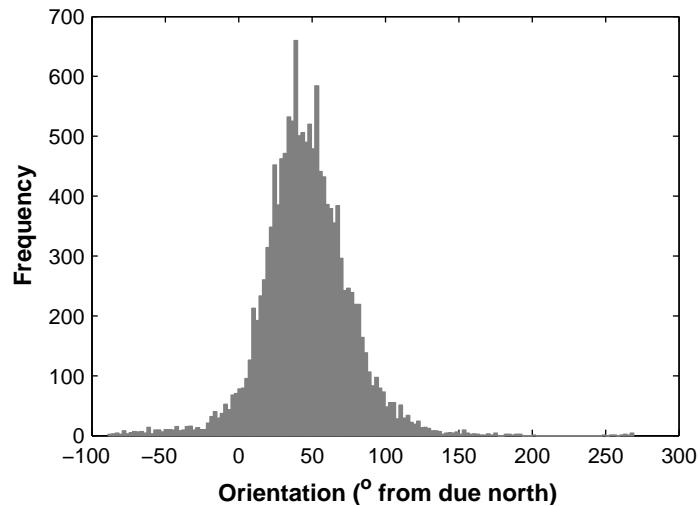


FIGURE A.2: Histogram showing the frequency of the Southland Fronts estimated orientation (with respect to due north) calculated from the data set generated by Shaw (1998) using the ‘front following algorithm’. Positive (negative) values relate to an angle measured clockwise (anticlockwise) from due north.

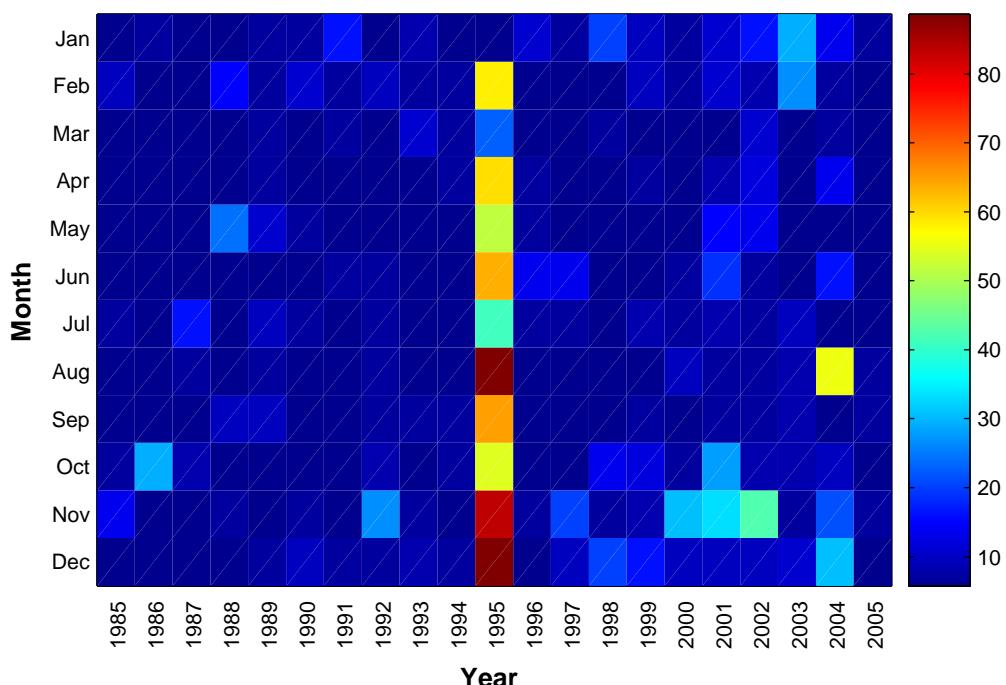


FIGURE A.3: Percentage of missing data in each 4km resolution AVHRR SST image of the Southland Front between January 1985 and December 2005

A.4 Southland Front Variability

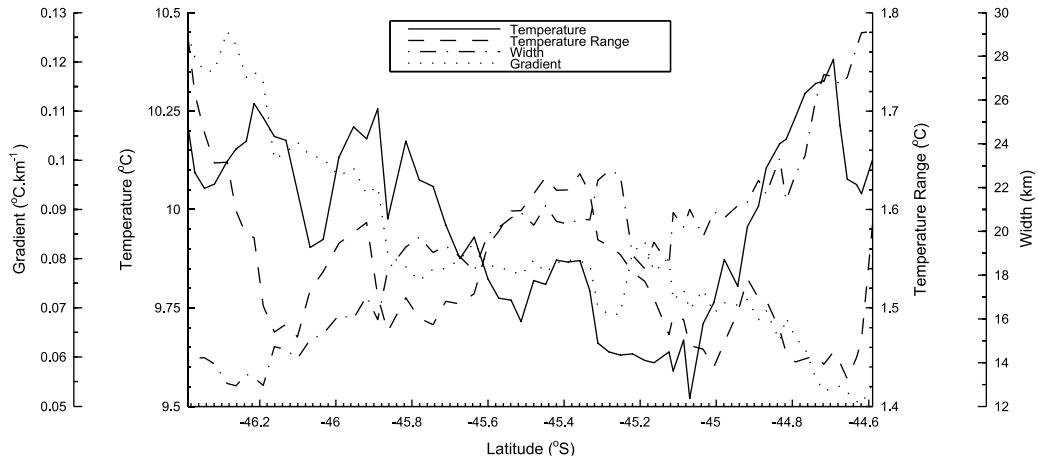


FIGURE A.4: Along-front changes in temperature, temperature range, width and gradient between April 1989 and March 1992.

TABLE A.2: Mean three year (April 1989 - March 1992) seasonal characteristics of the Southland Front as calculated by Shaw and Vennell (2001) from 1km data and by the current algorithm (Hopkins) using 4km resolution images. The less than sign ($<$) represents a confidence interval less than 0.005.

	Shaw and Vennell (2001)		Hopkins	
	Mean \pm 95% CI	s.d	Weighted Mean \pm 95% CI	Weighted s.d
SPRING				
Temperature (°C)	10.02 \pm 0.02	1.35	8.90 \pm 0.004	1.03
Temperature Range (°C)	1.63 \pm <0.01	0.63	1.54 \pm 0.005	0.49
Width (km)	8.69 \pm 0.03	4.81	24.40 \pm 0.25	12.00
Gradient (°C.km⁻¹)	0.25 \pm <0.01	0.19	0.063 \pm 0.01	
SUMMER				
Temperature (°C)	13.37 \pm 0.02	1.36	12.88 \pm 0.007	0.84
Temperature Range (°C)	1.83 \pm 0.01	0.70	1.96 \pm 0.008	0.47
Width (km)	9.37 \pm 0.04	4.77	16.71 \pm 0.28	11.51
Gradient (°C.km⁻¹)	0.24 \pm <0.01	0.15	0.12 \pm 0.02	
AUTUMN				
Temperature (°C)	12.27 \pm 0.02	1.24	11.46 \pm 0.004	1.13
Temperature Range (°C)	1.73 \pm <0.01	0.63	1.54 \pm 0.005	0.51
Width (km)	8.19 \pm 0.03	4.73	19.07 \pm 0.24	10.92
Gradient (°C.km⁻¹)	0.28 \pm <0.01	0.32	0.08 \pm 0.01	
WINTER				
Temperature (°C)	9.17 \pm 0.01	0.76	8.71 \pm 0.004	0.55
Temperature Range (°C)	1.86 \pm <0.01	0.72	1.37 \pm 0.005	0.30
Width (km)	7.63 \pm 0.03	4.74	16.41 \pm 0.24	7.98
Gradient (°C.km⁻¹)	0.34 \pm <0.01	0.28	0.08 \pm 0.02	

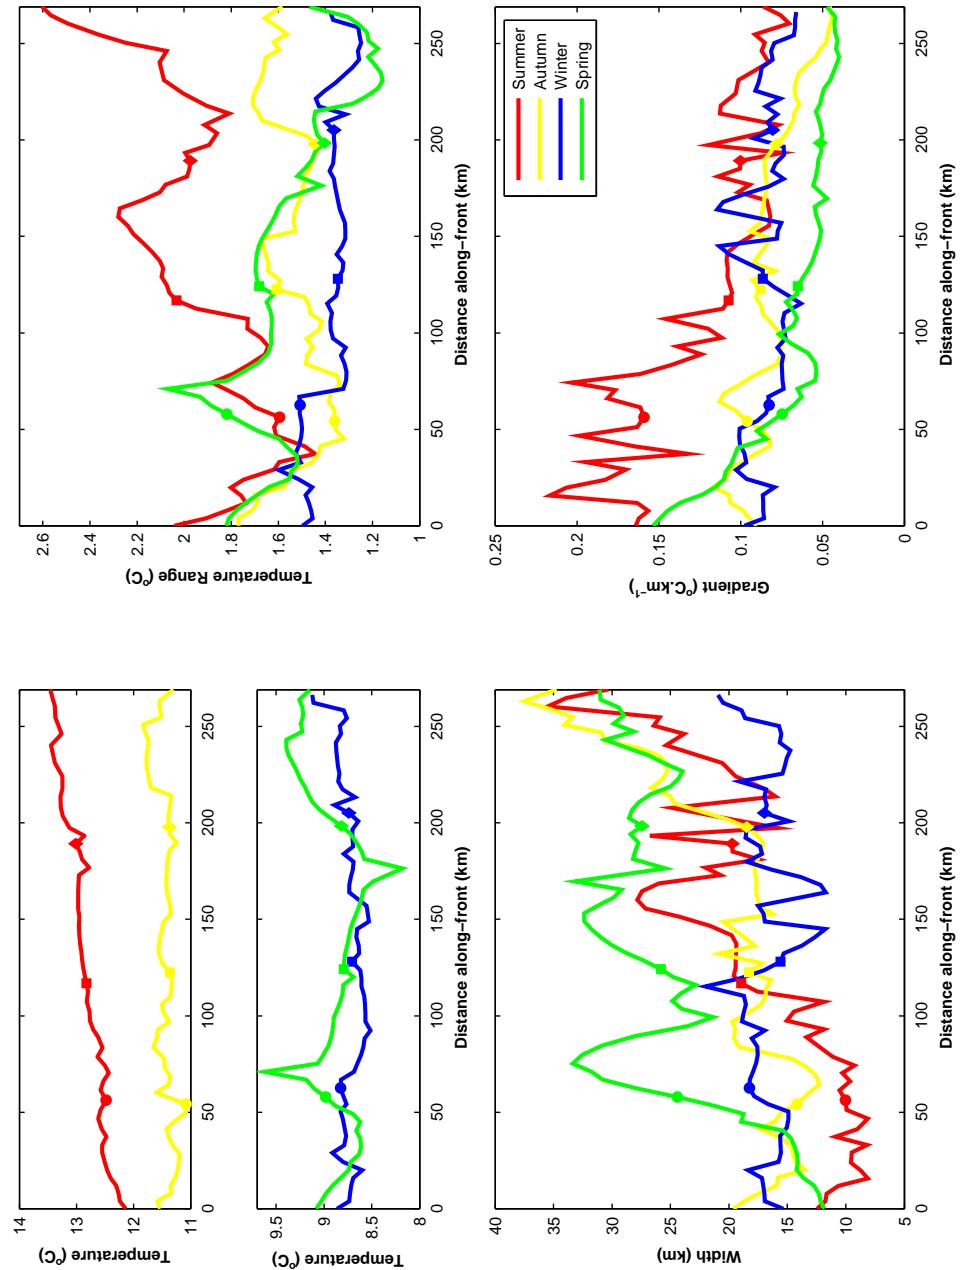


FIGURE A.5: Seasonal variability in along-front temperature, temperature range, width and gradient between April 1989 and March 1992. Filled circles, squares and diamonds mark -46°S , -45.5°S and -45°S (Oamaru) along the mean seasonal positions respectively.

A.5 ARIMA Model Fitting

In preparation for analysis all timeseries are detrended and the monthly anomalies calculated. An autoregressive integrated moving average (ARIMA) model is then fitted, following the methodology of Box and Jenkins (1970), to remove any autocorrelation between samples in each time series, thus making each point statistically independent from all others. This process, known as prewhitening, preserves the structure between time series and allows them to be repeatedly cross correlated.

ARIMA(p,d,q) is a generalization of an ARMA(p,q) model that may be applied to a stationary time series. A process is said to be stationary if its statistical properties do not change over time, that is, it should have a constant mean, variance and autocorrelation. The autoregressive part of the model, AR(p), expresses the time series as a linear function of its prior p observations. The moving average process, MA(q), which is independent from the autoregressive model, regards the time series as an unevenly weighted moving average of its past q error terms. Each observation is made up of a random error and a linear combination of prior random ‘shocks’. An ARMA(p,q) model therefore includes both types of lagged term, lags on both the series itself and on the residuals. The introduction of a differencing term, d, in order to make stationary a non-stationary time series leads to an Autoregressive Integrated Moving Average Model, ARIMA(p,d,q), where d represents the order of differences required to achieve stationarity.

Fitting each ARIMA model comprises three stages of model identification, estimation and checking. Identification of the order of differencing required, if any, relies heavily on examination of the autocorrelation (ACF) and partial autocorrelation (PACF) functions. Having already detrended and removed any seasonal cycle from the time series stationarity was expected in most cases. The ACF of a stationary process should converge rapidly to zero as the values of the lag increases. An ACF that exhibits a smooth pattern of significant values at high lags indicates a non-stationary series. A model with one order of differencing was therefore deemed most appropriate for these series and makes the assumption that the original series has a constant average trend. Note that differencing tends to introduce negative autocorrelations. If the ACF at lag 1 is zero or negative then no further differencing is necessary. Over-differencing leads to an increase in the standard deviation of the residuals.

After the series has been made stationary AR and/or MA terms were identified to correct for any autocorrelations that remained. This was done using both the Akaike Information Criterion, a measure of the goodness-of-fit of the model and by examining the ACF and PACF of the differenced series. The theoretical ACF of a pure MA process of order q, has zero autocorrelation at all lags greater than q. The corresponding PACF is liable to decay gradually towards zero. For a purely AR process, it is the PACF that serves most clearly to identify the appropriate model order; the PACF displays a sharp

cutoff at lag p , while the ACF decays more slowly. To judge whether the sample ACF or PACF showed significant evidence of truncation approximate 95% confidence bounds for autocorrelations of a white noise sequence were constructed.

Once a model has been chosen its coefficients are estimated using the method of maximum likelihood. If the chosen model is doing an effective job at describing the persistence then the residuals should be random and uncorrelated in time. The ACF of the residuals should be zero at all lags except zero. The model order was adjusted until this requirement was met. Box and Jenkins (1970) cover the topic of model checking in all the necessary detail.

A.6 Southland Front and ENSO

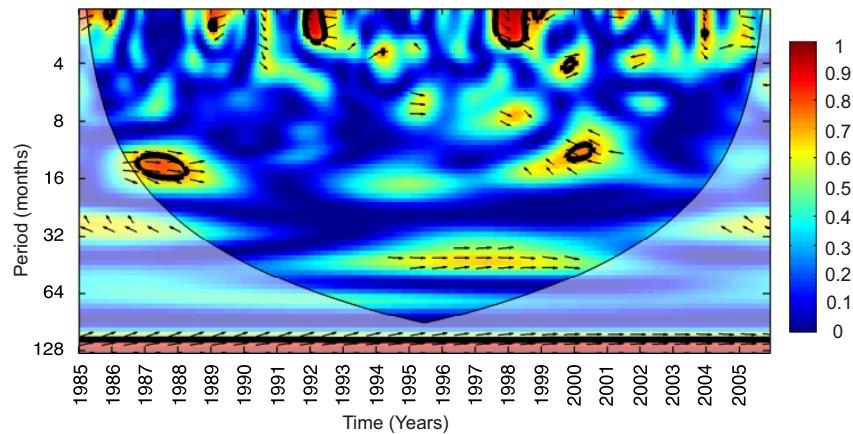


FIGURE A.6: Wavelet squared coherence between the width of the Southland Front and the SOI. Thick black contours mark the 95% significance level above a white noise background spectrum. Arrows indicate the relative phase relationship.

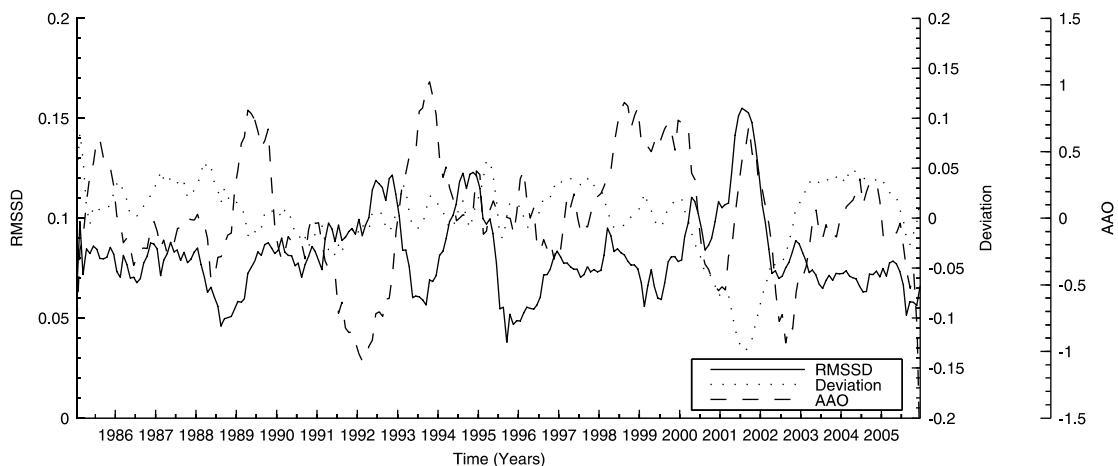


FIGURE A.7: Twelve month running mean time series of the fronts deviation from its mean twenty-one year position, its root-mean squared spatial displacement (RMSSD), a measure of stability, and the AAO. Positive (negative) deviation is shoreward (seaward) of the mean position.

Bibliography

Allen, J. T., L. Brown, R. Sanders, C. M. Moore, A. Mustard, S. Fielding, M. Lucas, M. Rixen, G. Savidge, S. Henson, and D. Mayor (2005). Diatom carbon export enhanced by silicate upwelling in the northeast Atlantic. *Nature* 437(7059), 728–732.

Askari, F. and B. Zerr (1998). A neural network architecture for automatic extraction of oceanographic features in satellite remote sensing. In *Proc.IGARSS 1998*, pp. 1017–1021.

Barnes, E. J. (1985). Eastern Cook Strait region circulation inferred from satellite-derived, sea-surface, temperature data. *New Zealand Journal Of Marine And Freshwater Research* 19(3), 405–411.

Barranco-López, V., P. Luque-Escamilla, J. Martínez-Aroza, and R. Román-Roldán (1995). Entropic Texture-Edge Detection For Image Segmentation. *Electronics Letters* 31(11), 867–869.

Basher, R. E. and C. S. Thompson (1996). Relationship of air temperatures in New Zealand to regional anomalies in sea-surface temperature and atmospheric circulation. *International Journal of Climatology* 16(4), 405–425.

Belkin, I. M. (1993). Frontal structure of the South Atlantic (in Russian). In N. Voronina (Ed.), *Pelagic Ecosystems of the Southern Ocean*, pp. 40–53. Moscow: Nauka.

Belkin, I. M. and A. L. Gordon (1996). Southern Ocean fronts from the Greenwich meridian to Tasmania. *Journal of Geophysical Research-Oceans* 101(C2), 3675–3696.

Benediktsson, J. A., P. H. Swain, and O. K. Ersoy (1990). Neural network approaches versus statistical-methods in classification of multisource remote-sensing data. *IEEE Transactions on Geoscience and Remote Sensing* 28(4), 540–552.

Bevington, P. and D. Robinson (2003). *Data Reduction and Error Analysis for the Physical Sciences* (3rd ed.). New York: McGraw Hill.

Bjerknes, J. (1969). Atmospheric teleconnections from the equatorial Pacific. *Monthly Weather Review* 97, 163–172.

Bowman, A. and A. Azzalini (1997). *Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations*. Oxford Statistical Science Series. Oxford: Oxford University Press.

Box, G. and G. Jenkins (1970). *Time Series Analysis: forecasting and control*. San Francisco: Holden-Day.

Box, M., D. Davies, and W. Swann (1969). *Non-Linear Optimization Techniques*. Mathematical and Statistical Techniques for Industry. Edinburgh: Oliver and Boyd.

Boyd, P., J. La Roche, M. Gall, R. Frew, and R. M. L. McKay (1999). Role of iron, light, and silicate in controlling algal biomass in subantarctic waters SE of New Zealand. *Journal of Geophysical Research-Oceans* 104 (C6), 13395–13408.

Bradford-Grieve, J. M., F. H. Chang, M. Gall, S. Pickmere, and F. Richards (1997). Size-fractionated phytoplankton standing stocks and primary production during austral winter and spring 1993 in the Subtropical Convergence region near New Zealand. *New Zealand Journal of Marine and Freshwater Research* 31(2), 201–224.

Bradford-Grieve, J. M., K. Probert, K. Lewis, P. Sutton, J. Zeldis, and A. Orphin (2006). New Zealand shelf region. In A. Robinson and K. Brink (Eds.), *The global coastal ocean: interdisciplinary regional studies and syntheses. Part B: The coasts of Africa, Europe, Middle East, Oceania and Polar Regions*, The Sea (14B), pp. 1451–1492. Cambridge, MA: Harvard University Press.

Brink, K. (1998). Deep-sea forcing and exchange processes. In K. Brink and A. Robinson (Eds.), *The global coastal ocean: processes and methods*, The Sea (10), pp. 151–167. New York: John Wiley and Sons.

Brodie, J. (1960). Coastal surface currents around New Zealand. *New Zealand Journal of Geology and Geophysics* 3, 235–252.

Burling, R. (1961). Hydrology of Circumpolar Waters South of New Zealand. Bulletin 143, New Zealand Department of Scientific and Industrial Research.

Butler, E. C. V., J. A. Butt, E. J. Lindstrom, P. C. Tildesley, S. Pickmere, and W. F. Vincent (1992). Oceanography of the Subtropical Convergence Zone around Southern New-Zealand. *New Zealand Journal of Marine and Freshwater Research* 26(2), 131–154.

Cayula, J. F. and P. Cornillon (1992). Edge-detection algorithm for SST images. *Journal of Atmospheric and Oceanic Technology* 9(1), 67–80.

Cayula, J. F. and P. Cornillon (1995). Multi-image edge-detection for SST images. *Journal of Atmospheric and Oceanic Technology* 12(4), 821–829.

Chelton, D. B., R. A. De Szoeke, M. G. Schlax, K. El Naggar, and N. Siwertz (1998). Geographical variability of the first baroclinic Rossby radius of deformation. *Journal of Physical Oceanography* 28(3), 433–460.

Chelton, D. B., M. G. Schlax, M. H. Freilich, and R. F. Milliff (2004). Satellite measurements reveal persistent small-scale features in ocean winds. *Science* 303(5660), 978–983.

Chin, T. M. and A. J. Mariano (1997). Space-time interpolation of oceanic fronts. *IEEE Transactions On Geoscience And Remote Sensing* 35(3), 734–746.

Chiswell, S. M. (1994). Variability in sea-surface temperature around New-Zealand from AVHRR images. *New Zealand Journal of Marine and Freshwater Research* 28(2), 179–192.

Chiswell, S. M. (1996). Variability in the Southland Current, New Zealand. *New Zealand Journal of Marine and Freshwater Research* 30(1), 1–17.

Chiswell, S. M., J. Toole, and J. Church (1997). Transports across the Tasman Sea from WOCE repeat sections: the East Australian Current 1990-94. *New Zealand Journal of Marine and Freshwater Research* 31(4), 469–475.

Coulter, R. (1983). Application of the Bayes decision rule for automatic water mass classification from satellite infrared images. *Proc. 17th Int. Symp. Remote Sensing Environment* 2, 589–597.

Cunningham, S. A., S. G. Alderson, B. A. King, and M. A. Brandon (2003). Transport and variability of the Antarctic Circumpolar Current in Drake Passage. *Journal Of Geophysical Research-Oceans* 108(C5), doi:10.1029/2001JC001147.

Currie, K. I. and K. A. Hunter (1998). Surface water carbon dioxide in the waters associated with the subtropical convergence, east of New Zealand. *Deep-Sea Research Part I-Oceanographic Research Papers* 45(10), 1765–1777.

Currie, K. I. and K. A. Hunter (1999). Seasonal variation of surface water CO₂ partial pressure in the Southland Current, east of New Zealand. *Marine and Freshwater Research* 50(5), 375–382.

Davison, A. (2003). *Statistical Models* (1st ed.). Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge: Cambridge University Press.

Deacon, G. (1933). A general account of the hydrology of the South Atlantic Ocean. *Discovery Reports* 7, 171–238.

Deacon, G. (1937). The hydrology of the Southern Ocean. *Discovery Reports* 15, 1–124.

Dong, S., J. Sprintall, and S. T. Gille (2006). Location of the Antarctic Polar Front from AMSR-E satellite surface temperature measurements. *Journal of Physical Oceanography* 36, 2075–2089.

Edwards, L. A., R. E. Houseago-Stokes, and P. Cipollini (2006). Altimeter observations of the MJO/ENSO connection through Kelvin waves. *International Journal of Remote Sensing* 27(5-6), 1193–1203.

Emery, W. and R. Thomson (1998). *Data Analysis Methods in Physical Oceanography*. Oxford: Permagon Press.

Feliks, Y., M. Ghil, and E. Simonnet (2004). Low-frequency variability in the midlatitude atmosphere induced by an oceanic thermal front. *Journal of the Atmospheric Sciences* 61(9), 961–981.

Folland, C. K., J. A. Renwick, M. J. Salinger, and A. B. Mullan (2002). Relative influences of the Interdecadal Pacific Oscillation and ENSO on the South Pacific Convergence Zone. *Geophysical Research Letters* 29(13), doi:10.1029/2001GL014201.

Fyfe, J. C. and O. A. Saenko (2006). Simulated changes in the extratropical Southern Hemisphere winds and currents. *Geophysical Research Letters* 33(6), doi:10.1029/2005GL025332.

Galanti, E. and E. Tziperman (2003). A Midlatitude-ENSO teleconnection mechanism via baroclinically unstable long Rossby waves. *Journal of Physical Oceanography* 33(9), 1877–1888.

Garvine, R. W. (1979). Integral hydrodynamic model of upper ocean frontal dynamics. 1. Development and analysis. *Journal of Physical Oceanography* 9(1), 1–18.

Gershunov, A. and T. P. Barnett (1998). Interdecadal modulation of ENSO teleconnections. *Bulletin of the American Meteorological Society* 79(12), 2715–2725.

Gerson, D., E. Khedouri, and P. Gaborski (1979). Location of oceanic fronts from digital satellite temperature data by automated pattern analysis. In *Intergovernmental Oceanographic Commission Workshop Report No.17, Supplement 143-163*.

Gill, P., W. Murray, and M. Wright (1995). *Practical Optimizaion*. London: Academic Press.

Gille, S. T. (1994). Mean sea-surface height of the Antarctic Circumpolar Current from Geosat data - Method and application. *Journal of Geophysical Research-Oceans* 99(C9), 18255–18273.

Gille, S. T. and K. A. Kelly (1996). Scales of spatial and temporal variability in the Southern Ocean. *Journal of Geophysical Research-Oceans* 101(C4), 8759–8773.

Godfrey, J. S. (1973). Comparison of East Australian Current with Western Boundary flow in Bryan And Cox's (1968) numerical model ocean. *Deep-Sea Research* 20(12), 1059–1076.

Godfrey, J. S., G. R. Cresswell, T. J. Golding, and A. F. Pearce (1980). The separation of the East Australian Current. *Journal of Physical Oceanography* 10(3), 430–440.

Gordon, N. D. (1986). The Southern Oscillation and New-Zealand weather. *Monthly Weather Review* 114(2), 371–387.

Goring, D. G. and R. G. Bell (1999). El Niño and decadal effects on sea-level variability in northern New Zealand: a wavelet analysis. *New Zealand Journal of Marine and Freshwater Research* 33(4), 587–598.

Graybill, F. and R. Deal (1959). Combining Unbiased Estimators. *Biometrics* 15(4), 543–550.

Greig, M. J. and A. E. Gilmour (1992). Flow through the Mernoo Saddle, New-Zealand. *New Zealand Journal of Marine and Freshwater Research* 26(2), 155–165.

Greig, M. J., N. M. Ridgway, and B. S. Shakespeare (1988). Sea-surface temperature-variations at coastal sites around New-Zealand. *New Zealand Journal of Marine and Freshwater Research* 22(3), 391–400.

Grinsted, A., J. C. Moore, and S. Jevrejeva (2004). Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Processes In Geophysics* 11(5-6), 561–566.

Hall, A. and M. Visbeck (2002). Synchronous variability in the southern hemisphere atmosphere, sea ice, and ocean resulting from the annular mode. *Journal Of Climate* 15(21), 3043–3057.

Hastie, T. and R. Tibshirani (1990). *Generalized Additive Models*. London: Chapman and Hall.

Hastie, T., R. Tibshirani, and J. Friedman (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series In Statistics. New York: Springer-Verlag.

Heath, R. (1968). Geostrophic currents derived from oceanic density measurements north and south of the Subtropical Convergence east of New Zealand. *New Zealand Journal of Marine and Freshwater Research* 2, 659–677.

Heath, R. (1975). Oceanic Circulation off the East Coast of New Zealand. Memoir 55, New Zealand Oceanographic Institute.

Heath, R. A. (1972). The Southland Current. *New Zealand Journal of Marine and Freshwater Research* 6, 497–533.

Heath, R. A. (1985a). A review of the physical oceanography of the seas around New Zealand - 1982. *New Zealand Journal of Marine and Freshwater Research* 19(1), 79–124.

Heath, R. A. (1985b). Large-scale influence of the New-Zealand seafloor topography on Western Boundary Currents of the South-Pacific Ocean. *Australian Journal of Marine and Freshwater Research* 36(1), 1–14.

Hickox, R., I. Belkin, P. Cornillon, and Z. Shan (2000). Climatology and seasonal variability of ocean fronts in the east China, Yellow and Bohai Seas from satellite SST data. *Geophysical Research Letters* 27(18), 2945–2948.

Holland, J. A. and X. H. Yan (1992). Ocean thermal feature recognition, discrimination, and tracking using infrared satellite imagery. *IEEE Transactions on Geoscience and Remote Sensing* 30(5), 1046–1053.

Holyer, R. J. and S. H. Peckinpaugh (1989). Edge-detection applied to satellite imagery of the oceans. *IEEE Transactions on Geoscience and Remote Sensing* 27(1), 46–56.

Houtman, J. (1966). A note on the hydrological regime in Foveaux Strait. *New Zealand Journal of Science* 9, 472–483.

Hughes, C. W. and E. R. Ash (2001). Eddy forcing of the mean flow in the Southern Ocean. *Journal of Geophysical Research-Oceans* 106(C2), 2713–2722.

Jackson, G. D., A. G. P. Shaw, and C. Lalas (2000). Distribution and biomass of two squid species off southern New Zealand: *Nototodarussloani* and *Moroteuthisingsens*. *Polar Biology* 23(10), 699–705.

Jacobs, G. A., H. E. Hurlburt, J. C. Kindle, E. J. Metzger, J. L. Mitchell, W. J. Teague, and A. J. Wallcraft (1994). Decade-Scale Trans-Pacific Propagation and Warming Effects of an El-Nino Anomaly. *Nature* 370(6488), 360–363.

James, C., M. Tomczak, I. Helmond, and L. Pender (2002). Summer and winter surveys of the Subtropical Front of the southeastern Indian Ocean 1997-1998. *Journal of Marine Systems* 37(1-3), 129–149.

Janowitz, M. (1985, June). Automatic detection of Gulf Stream Rings. Technical report, Office of Naval Research TR-J8501.

Jeffrey, M. (1986). Climatological features of the Subtropical Convergence in Australian and New Zealand Waters. University of Sydney Ocean Sciences Institute Report 17, pp 96.

Jillett, J. (1969). Seasonal hydrology of waters off the Otago Peninsula, south-eastern New Zealand. *New Zealand Journal of Marine and Freshwater Research* 3(3), 349–375.

Kahru, M., B. Hakansson, and O. Rud (1995). Distributions of the sea-surface temperature fronts in the Baltic Sea as derived from satellite imagery. *Continental Shelf Research* 15(6), 663–679.

Kilpatrick, K. A., G. P. Podesta, and R. Evans (2001). Overview of the NOAA/-NASA advanced very high resolution radiometer Pathfinder algorithm for sea surface temperature and associated matchup database. *Journal of Geophysical Research-Oceans* 106(C5), 9179–9197.

Kostianoy, A. G., A. I. Ginzburg, M. Frankignoulle, and B. Delille (2004). Fronts in the Southern Indian Ocean as inferred from satellite sea surface temperature data. *Journal of Marine Systems* 45(1-2), 55–73.

Krishnamurthy, S., S. S. Iyengar, R. J. Holyer, and M. Lybanon (1994). Histogram-based morphological edge detector. *IEEE Transactions on Geoscience and Remote Sensing* 32(4), 759–767.

Lea, S. M. and M. Lybanon (1993). Automated boundary delineation in infrared ocean images. *IEEE Transactions on Geoscience and Remote Sensing* 31(6), 1256–1260.

Lee, T. and P. Cornillon (1995). Temporal variation of meandering intensity and domain-wide lateral oscillations of the Gulf-Stream. *Journal of Geophysical Research-Oceans* 100(C7), 13603–13613.

Legeckis, R., C. W. Brown, and P. S. Chang (2002). Geostationary satellites reveal motions of ocean surface fronts. *Journal of Marine Systems* 37(1-3), 3–15.

Loder, J. and T. Platt (1985). Physical controls on phytoplankton production at tidal fronts. In P. Gibbs (Ed.), *19th European Marine Biology Symposium*, pp. 3–21. Cambridge University Press.

Lou, X., W. Huang, A. Shi, and B. Fu (2005). Satellite measurements of the Zhejiang-Fujian coastal front in the East China Sea with an improved front-following algorithm. In D. Lu and G. Matvienko (Eds.), *The International Society for Optical Engineering*, Volume 5832- Optical Technologies for Atmospheric, Ocean and Environmental Studies, Bellingham, WA.

Lovenduski, N. S. and N. Gruber (2005). Impact of the Southern Annular Mode on Southern Ocean circulation and biology. *Geophysical Research Letters* 32(11), doi:10.1029/2005GL022727.

Lutjeharms, J. R. E. and L. H. McQuaid (1986). Changes in the structure of thermal ocean fronts south of Africa over a 3-month period. *South African Journal of Science* 82(9), 470–476.

Lutjeharms, J. R. E., H. R. Valentine, and R. C. Vanballegooien (1993). On the Sub-tropical Convergence in the South-Atlantic Ocean. *South African Journal of Science* 89(11-12), 552–552.

Lutjeharms, J. R. E., M. Walters, and B. Allanson (1985). Oceanic Frontal Systems and Biological Enhancement. In W. Siegfried, P. Condy, and R. Law (Eds.), *Antarctic Nutrient Cycles and Food Webs*, pp. 11–21. Berlin: Springer-Verlag.

Mann, K. and J. Lazier (1996). *The dynamics of marine ecosystems: biological-physical interactions in the oceans* (2nd ed.). Oxford: Blackwell Science.

Mantua, N. (2000, 31st March 2008). The Pacific Decadal Oscillation. <http://jisao.washington.edu/pdo/>.

Mantua, N. J. and S. R. Hare (2002). The Pacific Decadal Oscillation. *Journal of Oceanography* 58(1), 35–44.

Marcello, J., F. Marques, and F. Eugenio (2005). Automatic tool for the precise detection of upwelling and filaments in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing* 43(7), 1605–1616.

Mavor, T. P. and J. J. Bisagni (2001). Seasonal variability of sea-surface temperature fronts on Georges Bank. *Deep-Sea Research Part II-Topical Studies in Oceanography* 48(1-3), 215–243.

McClatchie, S., R. Coombs, and G. Macaulay (2001). Bugs, physics and production in the Subtropical Front. *Water and Atmosphere* 9(3), 19–21.

McPhaden, M. J. (1999). Genesis and evolution of the 1997-98 El Niño. *Science* 283(5404), 950–954.

Meredith, M. P., P. L. Woodworth, C. W. Hughes, and V. Stepanov (2004). Changes in the ocean transport through Drake Passage during the 1980s and 1990s, forced by changes in the Southern Annular Mode. *Geophysical Research Letters* 31(21), doi:10.1029/2004GL021169.

Meyers, G. (1996). Variation of Indonesian throughflow and the El Niño Southern Oscillation. *Journal of Geophysical Research-Oceans* 101(C5), 12255–12263.

Miller, J. L. (1994). Fluctuations of gulf-stream frontal position between cape-hatteras and the straits of florida. *Journal of Geophysical Research-Oceans* 99(C3), 5057–5064.

Miller, P. (2004). Multi-spectral front maps for automatic detection of ocean colour features from SeaWiFS. *International Journal of Remote Sensing* 25(7-8), 1437–1442.

Moore, J., M. Abbot, J. Richman, W. Smith, T. Cowles, K. Coale, W. Gardner, and R. Barber (1999). SeaWiFS satellite ocean colour data from the Southern Ocean. *Geophysical Research Letters* 26, 1465–1468.

Moore, J. and M. Abbott (2000). Phytoplankton chlorophyll distributions and primary production in the Southern Ocean. *Journal of Geophysical Research-Oceans* 105, 28 709–28 722.

Moore, J., M. Abbott, and J. Richman (1997). Variability in the location of the Antarctic Polar Front (90-20° W) from satellite sea surface temperature data. *Journal of Geophysical Research-Oceans* 102, 27 825–27 833.

Moore, J. K., M. R. Abbott, and J. G. Richman (1999). Location and dynamics of the Antarctic Polar Front from satellite sea surface temperature data. *Journal Of Geophysical Research-Oceans* 104(C2), 3059–3073.

Moore, M. I. and R. C. Murdoch (1993). Physical and biological observations of coastal squirts under non-upwelling conditions. *Journal of Geophysical Research-Oceans* 98(C11), 20043–20061.

Morin, P., P. Lecorre, and J. Lefevre (1985). Assimilation and regeneration of nutrients off the west-coast of Brittany. *Journal of the Marine Biological Association of the United Kingdom* 65(3), 677–695.

Morris, M., B. Stanton, and H. Neil (2001). Subantarctic oceanography around New Zealand: preliminary results from an ongoing survey. *New Zealand Journal of Marine and Freshwater Research* 35(3), 499–519.

Mullan, A. B. (1998). Southern hemisphere sea-surface temperatures and their contemporary and lag association with New Zealand temperature and precipitation. *International Journal of Climatology* 18(8), 817–840.

Murphy, P. P., R. A. Feely, R. H. Gammon, D. E. Harrison, K. C. Kelly, and L. S. Waterman (1991). Assessment of the air-sea exchange of CO₂ in the South-Pacific during austral autumn. *Journal of Geophysical Research-Oceans* 96(C11), 20455–20465.

Murphy, R. J., M. H. Pinkerton, K. M. Richardson, J. M. Bradford-Grieve, and P. W. Boyd (2001). Phytoplankton distributions around New Zealand derived from SeaWiFS remotely-sensed ocean colour data. *New Zealand Journal of Marine and Freshwater Research* 35(2), 343–362.

Nichol, D. G. (1987). Autonomous extraction of an eddy-like structure from infrared images of the ocean. *IEEE Transactions on Geoscience and Remote Sensing* 25(1), 28–34.

Nixon, M. and A. Aguado (2002). *Feature Extraction and Image Processing* (First Edition ed.). Oxford: Newnes.

NODC (2008, 24th March 2008). 4km Pathfinder Version 5.0 User Guide. <http://www.nodc.noaa.gov/sog/pathfinder4km/userguide.html>.

Numerical Algorithms Group (2006). NAG Library Manual, Mark 21. <http://www.nag.co.uk/numeric/fl/manual/html/FLLibraryManual.asp>.

Oke, P. R. and M. H. England (2004). Oceanic response to changes in the latitude of the Southern Hemisphere subpolar westerly winds. *Journal of Climate* 17(5), 1040–1054.

Olson, D. (2002). Biophysical Dynamics of Ocean Fronts. In A. Robinson, J. McCarthy, and B. Rothschild (Eds.), *Biological-Physical Interactions in the Ocean*, The Sea (12), pp. 187–218. New York: John Wiley and Sons.

O'Neill, L. W., D. B. Chelton, and S. K. Esbensen (2003). Observations of SST-induced perturbations of the wind stress field over the Southern Ocean on seasonal timescales. *Journal of Climate* 16(14), 2340–2354.

Orsi, A., T. Whitworth, and W. Nowlin (1995). On the meridional extent and fronts of the Antarctic Circumpolar Current. *Deep-Sea Research Part A-Oceanographic Research Papers* 42(5), 641–673.

Pal, M. and P. M. Mather (2005). Support vector machines for classification in remote sensing. *International Journal of Remote Sensing* 26(5), 1007–1011.

Park, K. A., J. Y. Chung, and K. Kim (2004). Sea surface temperature fronts in the East (Japan) Sea and temporal variations. *Geophysical Research Letters* 31(7), doi:10.1029/2004GL019424.

Park, Y., L. Gamberoni, and E. Charriaud (1993). Frontal structure, water masses and circulation in the Crozet Basin. *Journal of Geophysical Research-Oceans* 98, 12,361–12,385.

Pedersen, F. (1994). The oceanographic and biological tidal cycle succession in shallow sea fronts in the north-sea and the english-channel. *Estuarine Coastal And Shelf Science* 38(3), 249–269.

Philander, S. G. and A. Fedorov (2003). Is El Niño sporadic or cyclic? *Annual Review Of Earth And Planetary Sciences* 31, 579–594.

Phillips, H. E. and S. R. Rintoul (2000). Eddy variability and energetics from direct current measurements in the Antarctic Circumpolar Current south of Australia. *Journal of Physical Oceanography* 30(12), 3050–3076.

Pingree, R. D., P. R. Pugh, P. M. Holligan, and G. R. Forster (1975). Summer phytoplankton blooms and Red Tides along tidal fronts in approaches to English-Channel. *Nature* 258(5537), 672–677.

Podestá, G. P., M. Arbelo, R. Evans, K. Kilpatrick, V. Halliwell, and J. Brown (2003). Errors in high-latitude SSTs and other geophysical products linked to NOAA-14 AVHRR channel 4 problems. *Geophysical Research Letters* 30(11), doi:10.1029/2003GL017178.

Pond, S. and G. Pickard (1983). *Introductory Dynamical Oceanography* (2nd ed.). Oxford: Butterworth-Heinemann.

Power, S., T. Casey, C. Folland, A. Colman, and V. Mehta (1999). Inter-decadal modulation of the impact of ENSO on Australia. *Climate Dynamics* 15(5), 319–324.

Priestley, M. (1992). *Spectral Analysis and Time Series*. Probability and Mathematical Statistics. London: Academic Press.

Robinson, I. (2004). *Measuring the oceans from space: The principles and methods of satellite oceanography*. Berlin: Springer/Praxis.

Roemmich, D. and B. Cornuelle (1992). The Subtropical Mode Waters of the South-Pacific Ocean. *Journal of Physical Oceanography* 22(10), 1178–1187.

Roemmich, D., J. Gilson, J. Willis, P. Sutton, and K. Ridgway (2005). Closing the time-varying mass and heat budgets for large ocean areas: The Tasman Box. *Journal of Climate* 18(13), 2330–2343.

Roemmich, D., M. Morris, W. R. Young, and J. R. Donguy (1994). Fresh equatorial jets. *Journal of Physical Oceanography* 24(3), 540–558.

Roemmich, D. and P. Sutton (1998). The mean and variability of ocean circulation past northern New Zealand: Determining the representativeness of hydrographic climatologies. *Journal of Geophysical Research-Oceans* 103(C6), 13041–13054.

Rudnick, D. L. and R. Ferrari (1999). Compensation of horizontal temperature and salinity gradients in the ocean mixed layer. *Science* 283(5401), 526–529.

Salinger, M. J. and A. B. Mullan (1999). New Zealand climate: Temperature and precipitation variations and their links with atmospheric circulation 1930–1994. *International Journal of Climatology* 19(10), 1049–1071.

Sauter, D. and L. Parson (1994). Spatial-filtering for speckle reduction, contrast enhancement, and texture analysis of GLORIA images. *IEEE Journal of Oceanic Engineering* 19(4), 563–576.

Savage, C., R. Petrell, and T. Neufeld (1994). Underwater fish video images: Image quality and edge detection techniques. *Canadian Agricultural Engineering* 36, 175–183.

Sharples, J. (1997). Cross-shelf intrusion of subtropical water into the coastal zone of northeast New Zealand. *Continental Shelf Research* 17(7), 835–857.

Sharples, J. (1998). Physical processes on the New Zealand shelf and the rest of the world's islands. In A. Robinson and K. Brink (Eds.), *The global coastal ocean: regional studies and syntheses*, The Sea (11), pp. 965–996. New York: John Wiley and Sons.

Sharples, J., C. M. Moore, and E. R. Abraham (2001). Internal tide dissipation, mixing, and vertical nitrate flux at the shelf edge of NE New Zealand. *Journal of Geophysical Research-Oceans* 106(C7), 14069–14081.

Shaw, A. G. P. (1998). *The temporal and spatial variability of the Southland Front, New Zealand using AVHRR SST imagery*. Ph.D, University of Otago, Dunedin, New Zealand.

Shaw, A. G. P., L. Kavalieris, and R. Vennell (1999). Seasonal and inter-annual variability of SST off the east coast of the South Island, New Zealand. *Geocarto International* 14(3), 27–32.

Shaw, A. G. P. and R. Vennell (2000a). A front-following algorithm for AVHRR SST imagery. *Remote Sensing of Environment* 72(3), 317–327.

Shaw, A. G. P. and R. Vennell (2000b). Variability of water masses through the Mernoo Saddle, South Island, New Zealand. *New Zealand Journal of Marine and Freshwater Research* 34(1), 103–116.

Shaw, A. G. P. and R. Vennell (2001). Measurements of an oceanic front using a front-following algorithm for AVHRR SST imagery. *Remote Sensing of Environment* 75(1), 47–62.

Shimada, T., F. Sakaida, H. Kawamura, and T. Okumura (2005). Application of an edge detection method to satellite images for distinguishing sea surface temperature fronts near the Japanese coast. *Remote Sensing of Environment* 98(1), 21–34.

Silverman, B. (1986). *Density Estimation for Statistics and Data Analysis*. Monographs on Statistics and Applied Probability. London: Chapman and Hall.

Simhadri, K. K., S. S. Iyengar, R. J. Holyer, M. Lybanon, and J. M. Zachary (1998). Wavelet-based feature extraction from oceanographic images. *IEEE Transactions on Geoscience and Remote Sensing* 36(3), 767–778.

Simonoff, J. (1996). *Smoothing Methods in Statistics*. Springer Series in Statistics. New York: Springer Verlag.

Simpson, J. H., C. M. Allen, and N. C. G. Morris (1978). Fronts on continental-shelf. *Journal of Geophysical Research-Oceans and Atmospheres* 83(NC9), 4607–4614.

Simpson, J. H. and J. R. Hunter (1974). Fronts in Irish Sea. *Nature* 250(5465), 404–406.

Simpson, J. J. (1990). On the accurate detection and enhancement of oceanic features observed in satellite data. *Remote Sensing of Environment* 33(1), 17–33.

Song, Q. T., P. Cornillon, and T. Hara (2006). Surface wind response to oceanic fronts. *Journal of Geophysical Research-Oceans* 111(C12), doi:10.1029/2006JC003680.

Soulsby, R. (1983). The bottom boundary layer of shelf seas. In B. Johns (Ed.), *Physical oceanography of coastal and shelf seas*, pp. 189–266. Amsterdam: Elsevier.

Sprintall, J. (2003). Seasonal to interannual upper-ocean variability in the Drake Passage. *Journal of Marine Research* 61(1), 27–57.

Sprintall, J., D. Roemmich, B. Stanton, and R. Bailey (1995). Regional climate variability and ocean heat-transport in the southwest Pacific-Ocean. *Journal Of Geophysical Research-Oceans* 100(C8), 15865–15871.

Stanton, B. and P. Sutton (2003). Velocity measurements in the East Auckland Current north-east of North Cape, New Zealand. *New Zealand Journal of Marine and Freshwater Research* 37(1), 195–204.

Stanton, B. R. (1995). Intraseasonal sea-level variability on the west-coast of New Zealand. *New Zealand Journal of Marine and Freshwater Research* 29(2), 213–222.

Stephens, D. J., M. J. Meuleners, H. Van Loon, M. H. Lamond, and N. P. Telcik (2007). Differences in atmospheric circulation between the development of weak and strong warm events in the Southern Oscillation. *Journal of Climate* 20(10), 2191–2209.

Stramma, L. (1992). The South Indian-Ocean Current. *Journal of Physical Oceanography* 22(4), 421–430.

Stramma, L. and R. G. Peterson (1990). The South-Atlantic Current. *Journal of Physical Oceanography* 20(6), 846–859.

Stramma, L., R. G. Peterson, and M. Tomczak (1995). The South-Pacific Current. *Journal of Physical Oceanography* 25(1), 77–91.

Stuart-Menteth, A., I. Robinson, and P. Challenor (2003). A global study of diurnal warming using satellite-derived sea surface temperature. *Journal of Geophysical Research-Oceans* 108(C5), doi:10.1029/2002JC001534.

Sutton, P. (2001). Detailed structure of the Subtropical Front over Chatham Rise, east of New Zealand. *Journal of Geophysical Research-Oceans* 106(C12), 31045–31056.

Sutton, P. J. H. (2003). The Southland Current: a subantarctic current. *New Zealand Journal of Marine and Freshwater Research* 37(3), 645–652.

Sutton, P. J. H. and D. Roemmich (2001). Ocean temperature climate off north-east New Zealand. *New Zealand Journal of Marine and Freshwater Research* 35(3), 553–565.

Takahashi, T., S. C. Sutherland, C. Sweeney, A. Poisson, N. Metzl, B. Tilbrook, N. Bates, R. Wanninkhof, R. A. Feely, C. Sabine, J. Olafsson, and Y. Nojiri (2002). Global sea-air CO₂ flux based on climatological surface ocean pCO₂, and seasonal biological and temperature effects. *Deep-Sea Research Part II-Topical Studies in Oceanography* 49(9–10), 1601–1622.

Takahashi, W. and H. Kawamura (2005). Detection method of the Kuroshio front using the satellite-derived chlorophyll-*a* images. *Remote Sensing of Environment* 97(1), 83–91.

Tibshirani, R. and T. Hastie (1987). Local Likelihood Estimation. *Journal of the American Statistical Association* 82(398), 559–567.

Tilburg, C. E., H. E. Hurlbert, J. J. O'Brien, and J. F. Shriver (2001). The dynamics of the East Australian Current system: The Tasman Front, the East Auckland Current, and the East Cape Current. *Journal of Physical Oceanography* 31(10), 2917–2943.

Tilburg, C. E., H. E. Hurlbert, J. J. O'Brien, and J. F. Shriver (2002). Remote topographic forcing of a baroclinic western boundary current: An explanation for the Southland Current and the pathway of the subtropical front east of New Zealand. *Journal of Physical Oceanography* 32(11), 3216–3232.

Torrence, C. and G. P. Compo (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society* 79(1), 61–78.

Trenberth, K. E. and D. P. Tepaniak (2001). Indices of El Niño evolution. *Journal of Climate* 14(8), 1697–1701.

Uddstrom, M. J. and N. A. Oien (1999). On the use of high-resolution satellite data to describe the spatial and temporal variability of sea surface temperatures in the New Zealand region. *Journal of Geophysical Research-Oceans* 104(C9), 20729–20751.

Ullman, D. S. and P. Cornillon (1999). Satellite derived sea-surface temperature fronts on the continental shelf off the northeast U.S. Coast. *Journal of Geophysical Research-Oceans* 104, 23 459–23 478.

Ullman, D. S. and P. C. Cornillon (2001). Continental shelf surface thermal fronts in winter off the northeast US coast. *Continental Shelf Research* 21(11-12), 1139–1156.

Valavanis, V. D., I. Katara, and A. Palialexis (2005). Marine GIS: Identification of mesoscale oceanic thermal fronts. *International Journal of Geographical Information Science* 19(10), 1131–1147.

Van Loon, H. (1984). The Southern Oscillation.3. Associations with the trades and with the trough in the westerlies of the South-Pacific Ocean. *Monthly Weather Review* 112(5), 947–954.

Van Loon, H. and D. J. Shea (1985). The Southern Oscillation.4. The precursors south of 15°S to the extremes of the oscillation. *Monthly Weather Review* 113(12), 2063–2074.

Vazquez, D. P., C. Atae-Allah, and P. L. L. Escamilla (1999). Entropic approach to edge detection for SST images. *Journal of Atmospheric and Oceanic Technology* 16(7), 970–979.

Vincent, W. F. and C. Howard-Williams (1991). Distribution and biological properties of oceanic water masses around the South Island, New-Zealand. *New Zealand Journal of Marine and Freshwater Research* 25(1), 21–42.

Walton, C. C., W. G. Pichel, J. F. Sapper, and D. A. May (1998). The development and operational application of nonlinear algorithms for the measurement of sea surface temperatures with the NOAA polar-orbiting environmental satellites. *Journal of Geophysical Research-Oceans* 103(C12), 27999–28012.

Wang, C. Z. and R. H. Weisberg (2000). The 1997-98 El Niño evolution relative to previous El Niño events. *Journal of Climate* 13(2), 488–501.

Wang, D. X., Y. Liu, Y. Q. Qi, and P. Shi (2001). Seasonal variability of thermal fronts in the northern South China Sea from satellite data. *Geophysical Research Letters* 28(20), 3963–3966.

Zeldis, J., R. Walters, M. Greig, and K. Image (2004). Circulation over the northeastern New Zealand continental slope, shelf and adjacent Hauraki Gulf, during spring and summer. *Continental Shelf Research* 24, 543–561.