A FIELD-WISE WIND RETRIEVAL ALGORITHM FOR THE NASA SCATTEROMETER

by

Charles G. Brown

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

Department of Electrical and Computer Engineering

Brigham Young University

December 1998

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Charles G. Brown

This thesis has been read by each member of the following graduate committee and by majority vote has been found to be satisfactory.

Date	David G. Long, Chair
Date	Randal Beard
Date	Michael Jensen

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Charles G. Brown in its final form and have found that (1) its format, citations, and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

Date

David G. Long
Chair, Graduate Committee

Accepted for the Department

Michael D. Rice
Graduate Coordinator

Accepted for the College

Douglas M. Chabries
Dean, College of Engineering and Technology

ABSTRACT

A FIELD-WISE WIND RETRIEVAL ALGORITHM FOR THE NASA SCATTEROMETER

Charles G. Brown

Electrical and Computer Engineering

Master of Science

An alternative to traditional wind retrieval techniques that does not rely on numerical weather prediction winds is developed and applied in this thesis. Traditional satellite scatterometer wind retrieval algorithms consist of point-wise wind estimation and ambiguity removal. Point-wise estimation yields multiple estimates, or ambiguities, for the average near-surface ocean wind in each scatterometer resolution element. The presence of ambiguities requires subsequent ambiguity removal to select unique wind estimates. Even though some point-wise ambiguity removal techniques perform well, they are still subject to error. Further, the most successful ones, such as the "nudging" method employed by the Jet Propulsion Laboratory (JPL), incorporate outside information. One alternative to the traditional methods is field-wise wind retrieval.

Like point-wise retrieval, the field-wise method consists of estimation and ambiguity removal stages. However, field-wise estimation and ambiguity removal employ a wind field model that exploits the spatial correlation in wind fields to process large numbers of resolution elements simultaneously. A field-wise retrieval algorithm

that uses only scatterometer measurements is designed for the NASA Scatterometer (NSCAT) and is applied to data from that instrument.

Test results indicate that the field-wise method compares favorably with the standard NSCAT data product obtained from the Jet Propulsion Laboratory (JPL); however, the field-wise method has the advantage of not relying on outside information. It is further concluded that the field-wise method can be used to repair severe ambiguity removal failures in the output of the median filter, a traditional pointwise ambiguity removal scheme. An additional application of the field-wise method is validation of the JPL product. The field-wise method and a related model-based technique are used to assess the reliability of the JPL product, and both methods suggest over 96% accuracy, which is consistent with other recent validation results.

ACKNOWLEDGMENTS

Life often turns on small pivots. Such is the case with my initiation into remote sensing research. A brief visit with Dr. David Long about writing an honor's thesis culminated in this work on scatterometer wind retrieval. Without Dr. Long's patient guidance, this thesis would never have been possible. I especially appreciate the independence he granted me in my research. I have made many make mistakes, but I know I have grown in the process. Finally, I thank my wife, Mary Kaye, for her continual encouragement and for hours of proofreading. I am grateful to her and my children, Stephen and Kristina, for their love, and I dedicate this work to them.

Contents

Chapter 1

Introduction

Satellite wind scatterometers employ an indirect measurement technique to observe near-surface ocean wind. The scatterometer transmits microwave pulses to the ocean's surface and measures the backscattered power. Since changes in surface wind velocity alter surface roughness, the wind velocity can be inferred from the scatterometer measurements.

Satellite scatterometers are the only instruments that can provide accurate and frequent measurements of winds over all the oceans. Other methods, such as meteorological buoy and ship observations, are deficient in coverage, accuracy, or both. Buoy observations provide accurate wind measurements, but those data primarily cover coastal areas of the Northern Hemisphere [1]. Ship observations lack accuracy and coverage. They are inaccurate due to human error and ship motion. Their measurements are limited to coastal areas and shipping lanes, and avoidance of storms for safety and economic reasons further limits their coverage [1].

1.1 Brief History of Scatterometers

Scatterometers are relatively new radar systems. While radar was being developed during World War II, the effect of ocean waves on radar backscatter was first observed. These purely accidental observations were regarded as an annoyance. In World War II applications of radar over the ocean, operators noticed noise on their screens that blocked the signatures of boats and aircraft. They labeled the noise "sea clutter" and did not realize that it was radar returns from ocean waves. Scientists in the 1950's and 60's realized that they could optimize radar instruments specifically for

the purpose of exploiting the sea clutter phenomenon. They discovered a link between ocean wind, waves, and radar backscatter, and the modern science of scatterometry was born [2].

Much of the early development of scatterometers was a by-product of the Apollo program. In 1964, a system originally designed to make measurements predicting the behavior of the Apollo landing radar was used to measure scattering characteristics of the land and sea. After the Apollo program was cut short, a group of scientists and instrument designers met to decide what to do with the remaining Apollo instruments and spacecraft. Out of these meetings grew the Skylab program. The first satellite scatterometer, S-193, orbited on Skylab in 1973 and 1974. S-193 proved that ocean wind measurements could be made from space [3, 4].

In 1978, the SEASAT Scatterometer (SASS) was orbited on SEASAT, a remote sensing satellite. However, the SEASAT power supply failed after the first 90 days. Even though the SEASAT mission was short-lived, it was a major landmark in satellite scatterometry, providing the first ocean winds measured from space by a scatterometer [3, 4, 5].

Because of the success of SASS, several other scatterometers have been orbited. The European Space Agency deployed European Remote Sensing Satellites one and two (ERS-1 and ERS-2), and the United States orbited the NASA Scatterometer (NSCAT) on August 16, 1996. Follow-on missions, such as QuikScat and Seawinds, have already been planned by the United States.

1.2 NSCAT Mission

NSCAT was launched in a near-polar, sun-synchronous orbit, aboard the Advanced Earth Observing Satellite (ADEOS), which was designed by Japan's space agency, NASDA. NSCAT began collecting data during September 1996, but ADEOS lost power in June 1997. NSCAT used 14 GHz microwave pulses to observe ocean winds and measured 286000 globally distributed wind vectors per day [6].

Although the mission terminated prematurely as a result of a spacecraft power failure, its potential impact on meteorology, oceanography, and climatology is

Wind Velocity	2 m/s (rms) for 3-20 m/s 10% for 20-30 m/s
Wind Direction	20° (rms) 3-30 m/s
Spatial Resolution	25 km or 50 km
Coverage	90% of oceans every 2 days

Table 1.1: NSCAT mission requirements.

still limitless. NSCAT provided frequent and global coverage (Table 1.1) under all-weather conditions—coverage that is necessary for understanding complex worldwide weather and climate change. NSCAT data are expected to lead to improved methods of global weather forecasting and modeling, and possibly to a better understanding of large scale environmental phenomena, such as El Niño [6]. However, all of these applications require accurate retrieval of ocean winds from NSCAT data.

1.3 Introduction to Wind Retrieval

The traditional method of retrieving ocean wind is termed point-wise wind retrieval. In the point-wise method, estimation of ocean wind is performed for single resolution elements individually. Due to the nature of the relationship between wind and radar backscatter, point-wise estimation yields several wind vector solutions, or ambiguities. An additional processing step, known as point-wise ambiguity removal, is necessary in order to select a unique wind vector for each resolution element. Although some point-wise ambiguity removal techniques achieve high accuracy, none are perfect. Further, the most successful ones, such as the "nudging" method employed by the Jet Propulsion Laboratory (JPL), rely on outside information like numerical weather prediction winds. This thesis presents field-wise retrieval, which is an alternative to the point-wise methods that retrieves wind from scatterometer measurements alone.

Like its point-wise counterpart, field-wise retrieval consists of estimation and ambiguity removal stages. However, the field-wise method operates on many resolution elements simultaneously using a wind field model.

Field-wise estimation is a form of model-based retrieval [7]. Model-based processing starts with an initial wind field and seeks to improve that initial field through *local* optimization constrained by a wind field model. Model-based retrieval yields a single solution for each initial field. In contrast to model-based retrieval [7], field-wise estimation does not require initialization. Field-wise estimation developed in this thesis generates its own initial values and refines them via *global* optimization techniques constrained by a wind field model.

As with point-wise estimation, field-wise estimation yields multiple solutions, necessitating an ambiguity removal stage. Field-wise ambiguity removal has the advantage that adjacent wind field solutions can be made to overlap, and continuity considerations can be invoked to greatly simplify the process of selecting a unique solution [7, 8].

This thesis focuses on the design and application of a full field-wise wind retrieval algorithm. Each of the component stages, field-wise estimation and ambiguity removal, are developed separately, and the complete algorithm is tested on NSCAT data. Results are used to analyze the accuracy of NSCAT derived winds.

1.4 Summary of Contributions

The first major contribution of this work is the development and implementation of field-wise estimation. A global optimization algorithm is adapted to field-wise estimation and is tested on actual NSCAT data. The development and implementation of a field-wise ambiguity removal method is the second major contribution. The algorithm selects and assembles multiple solutions from the field-wise estimation algorithm into unique estimates for large areas of the ocean's surface. The final contribution is the testing of the full field-wise retrieval algorithm on NSCAT data. The algorithm output compares favorably with the standard "nudged" wind product provided by JPL; however, the field-wise method has the advantage of not using outside information. Finally, evidence is gathered to support recent validation results of the JPL product.

Although the field-wise methods presented in this thesis are developed for NSCAT, they can be extended for use with ERS-1, ERS-2, and the upcoming QuikScat and Seawinds missions.

1.5 Thesis Organization

Chapter 2 provides an overview of wind scatterometry, which is the science of inferring wind speed and direction from scatterometer backscatter measurements. Point-wise estimation and ambiguity removal are elaborated, and the chapter culminates in the introduction of field-wise wind retrieval. Definitions of both field-wise estimation and ambiguity removal are provided, which are expanded in the following chapters.

In Chapter 3 the field-wise estimation problem is characterized. A survey of several possible algorithms is provided, and one algorithm is selected and tested on NSCAT data. A field-wise ambiguity removal algorithm is developed in Chapter 4. The full field-wise wind retrieval method is evaluated in Chapter 5 using NSCAT data. The thesis concludes with a summary of contributions and recommendations for future work in Chapter 6.

Chapter 2

Background

2.1 Introduction to Wind Scatterometry

Wind scatterometry is the science of inferring the speed and direction of near-surface ocean wind from scatterometer backscatter measurements. The theory of scatterometry relies on a simple transitive relationship:

- Near-surface wind drives ocean waves
- Ocean waves affect radar backscatter
- Radar backscatter is related to the near-surface ocean wind
- Thus from observations of radar backscatter, the near-surface wind can be determined.

Indirect measurement of the wind by observing sea state is facilitated by the coupling between wind and waves. This section explores the relationship between ocean wind and waves from both historical and scientific points of view. The functional relationship between wind and radar backscatter is introduced in the following section.

2.1.1 History of Indirect Measurement of Ocean Wind

Aristotle (384-322 B.C.) was one of the first scientists to record that wind is important in wave generation [9]. Benjamin Franklin and his co-workers observed in 1774 that, "Air in motion, which is wind, in passing over the smooth surface of the

water, may rub, as it were, upon the surface and raise it into wrinkles which, if wind continues, are the elements of future waves [9]."

Early sailors also noticed that ocean waves are related to wind. They developed a visual scale by which they could estimate wind speed from observing the sea state [9]. In 1805, British Rear Admiral Sir Francis Beaufort took the sailors' visual scale and applied numbers to it to construct the Beaufort Scale (Fig. 2.1). The British Navy adopted the Beaufort Scale for general use in 1834. This scale is now an international means of reporting sea-state [9]. The principles that underlie the Beaufort Scale, that the wind determines sea state and that the wind speed can be inferred from ocean waves, are also key relationships in the theory of scatterometry. In the next section, scientific evidence supporting the first principle of scatterometry, that ocean winds drive ocean waves, is discussed.

2.1.2 Ocean Winds Drive Ocean Waves

The ocean's surface is a complex mixture of waves. When Lord Rayleigh, a 17th century mathematician and physicist, was confronted with the complexity of the ocean's surface, he said, "The basic law of the seaway is the apparent lack of any law [9]." However, the complexity of the ocean's surface can be reduced by examining the basic types of ocean waves.

The vast spectrum of ocean waves can be divided into five basic types: sound, capillary, gravity, inertial, and planetary, based on wavelength and dominant restoring force. This discussion of ocean waves is limited to capillary and gravity waves, as these are the most important ocean waves in satellite scatterometry [11].

Capillary waves are the waves that ripple across the ocean's surface when a light breeze blows. They have wavelengths on the order of a centimeter and are only a few millimeters in height. The surface tension of water is their dominant restoring force. Capillary waves dominate the ocean's surface for wind speeds up to 1 m/s. For higher wind speeds, capillary waves increase in size and wavelength to become gravity waves [9]. However, capillary waves are still present.

Force	Speed		Conditions
	knots	km/h	
0	<1	<1	Calm, sea like a mirror.
1	1–3	1–5	Light air, ripples only.
2	4–6	6–11	Light breeze, small wavelets (0.2m). Crests have a glassy appearance.
3	7–10	12–19	Gentle breeze, large wavelets (0.6m), crests begin to break.
4	11–16	20–29	Moderate breeze, small waves (1m), some white horses.
5	17–21	30–39	Fresh breeze, moderate waves (1.8m), many white horses.
6	22–27	40–50	Strong breeze, large waves (3m), probably some spray.
7	28–33	51–61	Near gale, mounting sea (4m) with foam blown in streaks downwind.
8	34–40	62–74	Gale, moderately high waves (5.5m), crests break into spindrift.
9	41–47	76–87	Strong gale, high waves (7m), dense foam, visibility affected.
10	48–55	88–102	Storm, very high waves (9m), heavy sea roll, visibility impaired. Surface generally white.
11	56–63	103–118	Violent storm, exceptionally high waves (11m), visibility poor.
12	64+	119	Hurricane, 14m waves, air filled with foam and spray, visibility bad.

Figure 2.1: The Beaufort scale was devised as a method of visually inferring wind speed from sea state. The coupling between sea state and wind is also a key relationship in modern satellite scatterometry. Figure from [10].

Gravity waves are long, undulating ocean waves. They are the most familiar type of wave. They are longer than capillary waves, having periods ranging from 1 to 10 seconds under normal conditions, but sustained ocean winds can form gravity waves with periods greater than 25 seconds. Gravity is the dominant restoring force in gravity waves [9].

One of the reasons why capillary and gravity waves are important in wind scatterometry is that the density of capillary and short-gravity waves changes in

response to wind speed. Reference [11] reports results from a wavetank study which relates the near-surface wind speed to the density of the resulting waves. This study shows that capillary and short-gravity wave density increases as wind speed increases. Even though wavetank studies cannot completely replicate open-ocean conditions, the study in [11] is supported by measurements over the sea.

2.1.3 Ocean Waves Affect Radar Backscatter

In order to discuss how ocean waves affect radar backscatter, it is necessary to define the normalized radar cross section, which quantifies radar backscatter. But first "backscatter" is qualitatively introduced by explaining the difference between reflection and scattering.

Reflection Versus Scattering

The difference between reflection and scattering may be elucidated by examining how radiation interacts with three different types of surfaces: smooth, slightly rough, and very rough (refer to Fig. 2.2). Reflection occurs when radiation strikes a smooth surface, which is also called a specular surface. Radiation reflects off a specular surface at the same incidence angle at which the incident radiation struck the surface. No radiation returns along the angle of incidence after striking a perfectly specular surface [11]. An example of a specular surface is a mirror or the ocean's surface on a calm day.

Radiation incident on a slightly rough surface is both reflected and scattered. Figure 2.2 represents the reflected portion of the incident radiation as a single large arrow and the scattered portions as smaller arrows. The reflected portion dominates the scattered portion. This type of mixed reflection and scattering, where the reflected portion dominates, is known as quasi-specular reflection [11]. In quasi-specular reflection, most of the incident radiation is reflected. However, some is scattered in various directions, including back along the angle of incidence. This scattered radiation is called "backscatter".

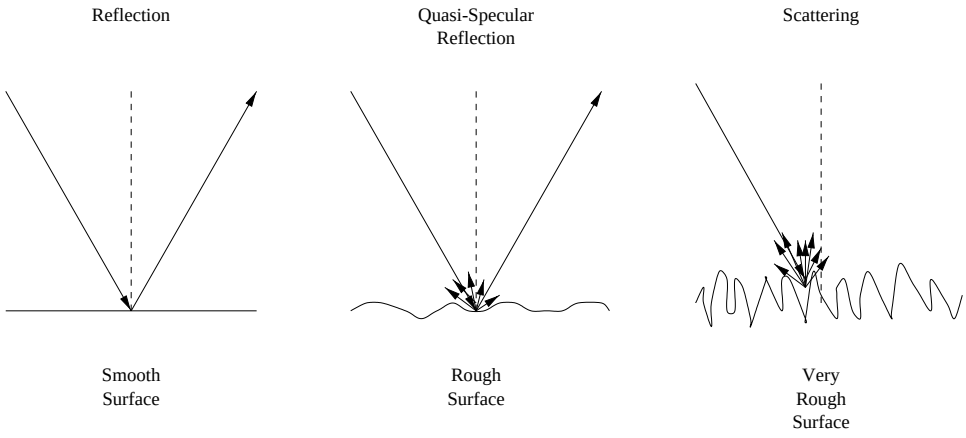

Figure 2.2: Difference between reflection and scattering. Adapted from [11].

The interaction between incident radiation and a very rough surface results mostly in backscatter. A perfectly rough surface reflects absolutely no incident radiation but scatters all of it [11].

Radar Cross Section

The normalized radar cross section σ_0 quantifies backscatter and varies with the roughness of the targeted surface. It is defined by the basic radar equation:

$$\sigma_0 = \frac{P_s}{P_t}\left(\frac{(4\pi)^3 R^4 L}{G^2 \lambda^2 A}\right),$$

where P_t is the power of the incident (transmitted) radiation, and P_s is the power of the received backscatter. R is the slant range to the surface, L accounts for known system losses, G is the antenna gain, λ is the wavelength of the incident radiation, and A is the effective area of the ocean's surface illuminated by the radiation [1]. This equation assumes a monostatic radar system, which transmits and receives with the same antenna.

In general for large incidence angles, rougher surfaces are associated with larger σ_0 values [11]. This is because more of the radiation is scattered back to the receiver than is reflected away from it. On the other hand, for very small incidence

angles (looking almost straight down), rougher surfaces are associated with smaller σ_0 values.

Bragg Scattering

Satellite scatterometer measurements of the roughness of the ocean's surface are primarily influenced by a phenomenon known as "Bragg scattering". Bragg scattering is the dominant backscattering mechanism for scatterometer observations at incidence angles between 20° and 65°, which range is known as the "plateau" region [1, 11]. Bragg scattering is a constructive resonance effect. The Bragg effect is so pronounced that it can dominate reflection from relatively large reflectors in the same area, such as the larger waves on which capillary or short-gravity waves are superimposed [11]. The main condition for Bragg scattering is that the incident radiation be about the same wavelength as the undulations of the surface. Capillary and short gravity waves cause Bragg scattering at the wavelengths used by scatterometers. For example, NSCAT transmits microwaves at a wavelength of about 2 cm, which is on the order of the wavelength of capillary and short gravity waves.

2.2 Geophysical Model Function

Functions that relate wind speed and direction to σ_0 are called geophysical model functions. The particular geophysical model function used in this thesis is NSCAT-1, which was derived from 3 months of NSCAT data and relates the wind 10 m above the ocean's surface to σ_0 [12]. NSCAT-1 may be expressed as follows:

$$\sigma_0 = \mathcal{M}(U, \chi, \theta, f, p)$$

where U is the true wind speed and χ is the relative azimuth angle, i.e. $\chi = \psi - \phi$, where ψ is the radar azimuth angle and ϕ is the true wind direction (refer to Fig. 2.3). The parameter θ is radar incidence angle, and f and p are radar frequency and polarization, respectively. Polarization is either horizontal or vertical. The geophysical model function may also be written as $\sigma_0 = \mathcal{M}(U, \psi)$ or $\sigma_0 = \mathcal{M}(\mathbf{w})$, with the dependence on χ, θ, f, and p implied. The vector $\mathbf{w}$ is the (U, ψ) pair in component

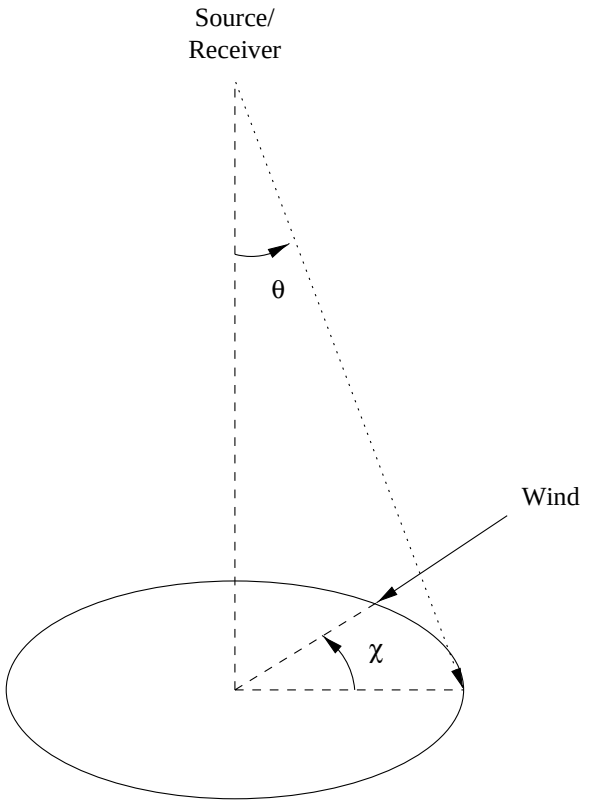

Figure 2.3: Definition of relative azimuth angle χ and incidence angle θ. Adapted from [1].

form, where $\mathbf{w} = (u, v)^T$, and $u = Usin(\phi)$ and $v = Ucos(\phi)$ by convention [13]. The omission of factors such as salinity, sea temperature, and long gravity waves introduces geophysical modeling errors. These errors are not considered in this work but are explored in [13, 14].

Figure 2.4 offers one perspective on the NSCAT-1 geophysical model function. It plots the radar cross section σ_0 versus relative azimuth angle χ for several different wind speeds U and incidence angles θ. The radar cross section generally increases with wind speed. However, σ_0 also varies with relative azimuth angle. The relationship between σ_0 is nearly $cos(2\chi)$, where $\chi = 0°$ is upwind and $\chi = 180°$ is downwind. The slight difference between σ_0 for $\chi = 0°$ and $\chi = 180°$ is called the upwind-downwind asymmetry [1].

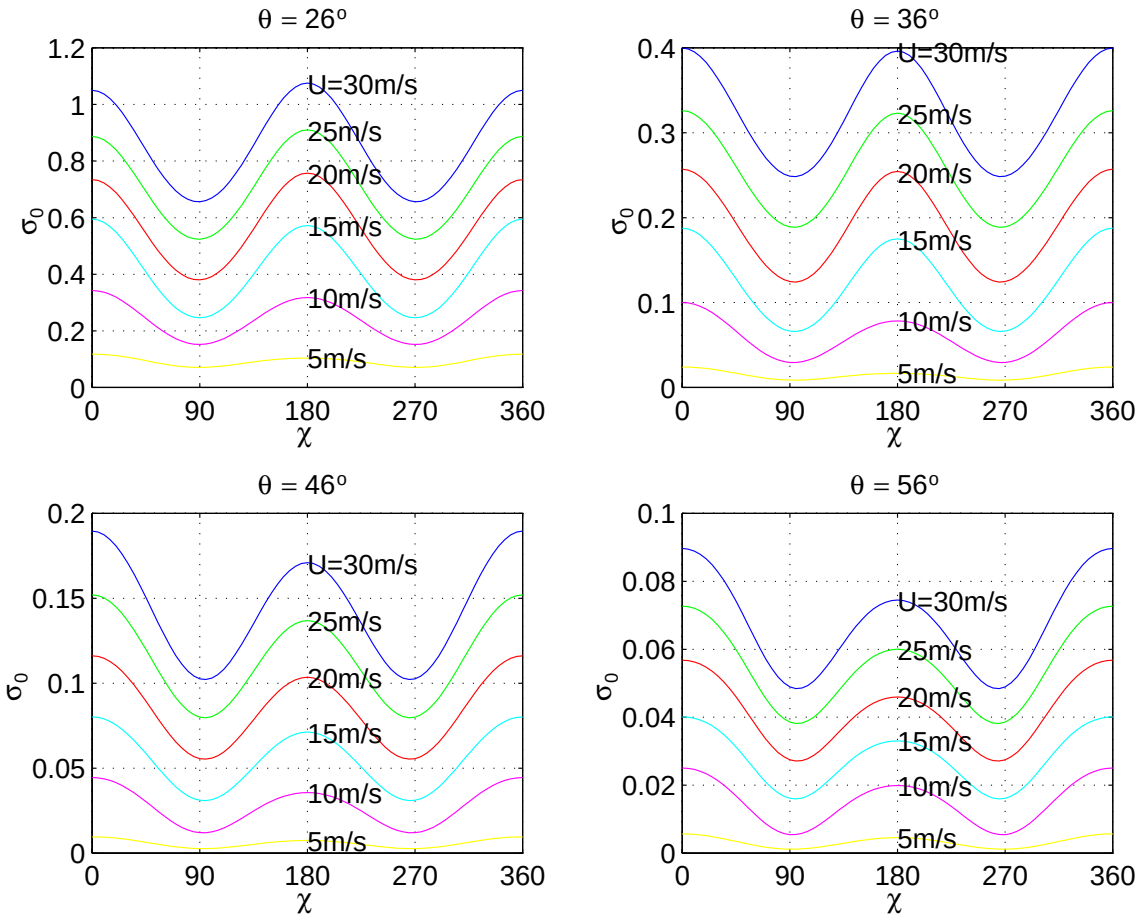

Figure 2.4: NSCAT-1 radar cross section σ_0 versus relative azimuth angle χ for several different wind speeds U and incidence angles θ. Adapted from [13].

2.3 Inverting the Geophysical Model Function

To retrieve wind from radar measurements it is necessary to invert the geophysical model function. The geophysical model function describes σ_0 as a function of wind velocity but is not one-to-one, i.e. $\mathcal{M}(U_1, \chi_1, \theta, f, p) = \mathcal{M}(U_2, \chi_2, \theta, f, p)$ does not guarantee that $(U_1, \chi_1) = (U_2, \chi_2)$. The many-to-one nature of NSCAT-1 is evident in Fig. 2.4. Observe that a value of σ_0 may be linked to many different (U, χ) pairs. Since the geophysical model function is not one-to-one, it does not have a unique inverse, which greatly complicates wind estimation.

The inverse relationship between a single measurement and wind velocity yields an infinite number of possible (U, ψ) pairs. The upper left subfigure in Fig. 2.5 illustrates this situation for noiseless measurements. A measurement from a single azimuth angle ϕ does not provide enough information to yield a unique wind velocity. However, two collocated measurements, each taken from different azimuth angles, narrow the number of solutions to a finite set, as depicted in the upper right subfigure of Fig. 2.5. Intersections of the curves identify the solutions. That is, a (U, ψ) pair is a solution to the inverse problem if $\sigma_1 = \mathcal{M}(U, \chi_1 = \psi - \phi_1, \theta_1, f_1, p_1)$ and $\sigma_2 = \mathcal{M}(U, \chi_2 = \psi - \phi_2, \theta_2, f_2, p_2)$, where the subscripts denote the measurement number.

Taking more measurements further refines the set of possible solutions. The bottom two subfigures in Fig. 2.5 illustrate the case for one and two more measurements from different azimuth angles. Note that the number of solutions is reduced to one with three measurements.

Although the series of plots in Fig. 2.5 illustrates the need for multiple collocated measurements from different azimuth angles, it is an idealized situation with no measurement noise. Inversion with measurements from real scatterometers rarely results in single solutions. Noise corrupts the measurements, shifting the curves in Fig. 2.5 up or down. Exact intersections are changed to near intersections, and the inversion of the geophysical model function becomes a statistical estimation problem.

Before continuing with the statistical estimation of ocean wind, several pertinent aspects of the NSCAT Scatterometer are discussed, particularly how NSCAT provides collocated measurements from various azimuth angles.

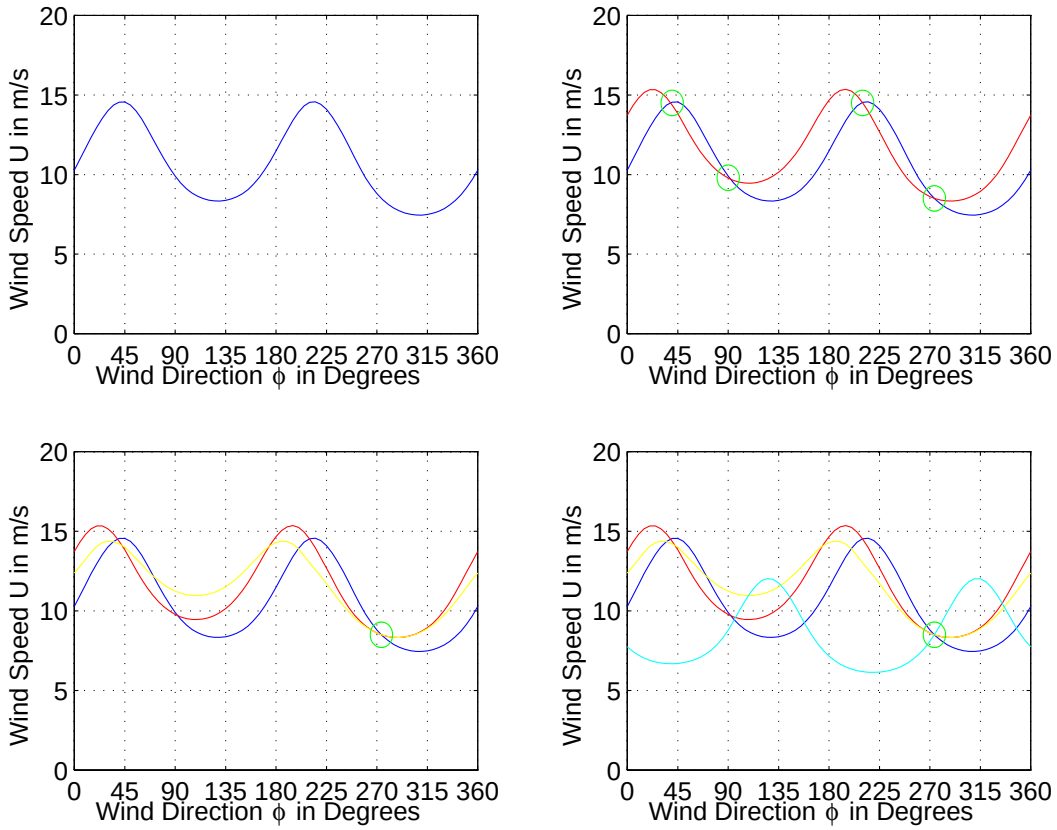

Figure 2.5: Reduction of multiple wind estimates through additional collocated measurements. One measurement yields an infinite number of estimates (upper left subfigure). Two collocated measurements, each taken from different azimuth angles, narrow the number of solutions to a finite set (upper right subfigure). Solutions occur at intersection points and are circled. Taking more measurements further refines the set of possible solutions (lower subfigures). The bottom two subfigures illustrate the case for one and two more measurements from different azimuth angles. Note that the number of solutions is reduced to one with three measurements (lower left subfigure).

2.4 An Overview of NSCAT

NSCAT is a fan beam scatterometer with six long, narrow antennas protruding from the instrument, three on each side [1]. The antennas cast illumination patterns several hundred kilometers long but only a few kilometers wide. Figures 2.6 and 2.7 illustrate the antenna illumination pattern on the ocean's surface. The multiple antennas produce measurements from different azimuth angles. For example, as NSCAT passes over an area of the ocean to its left (relative to the subsatellite track), it is first observed by antenna 6 from an azimuth angle of 45°. The footprint of antenna 5 then passes over it, and the antenna observes from an angle of 65°. Finally, antenna 4 looks behind the instrument to view the area from the 135° vantage point.

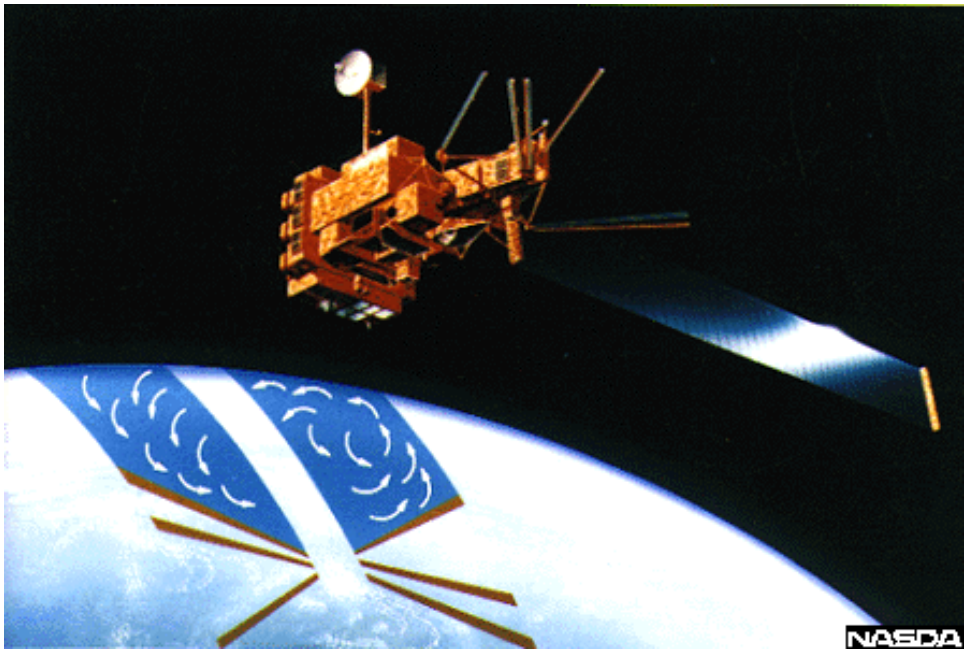

Figure 2.6: An artist's depiction of NSCAT. Obtained from [15].

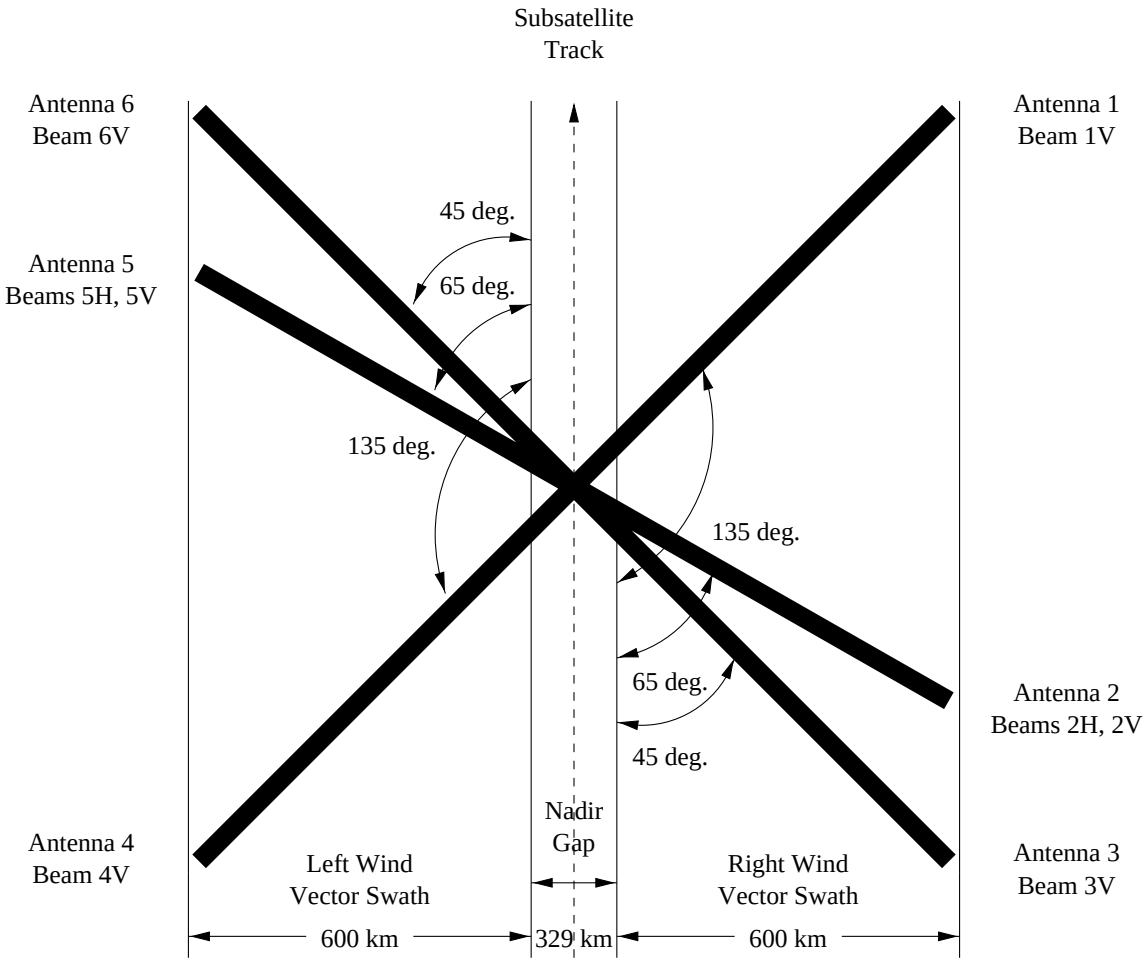

Figure 2.7: NSCAT illumination pattern. Adapted from [1].

NSCAT takes radar backscatter measurements from two 600 km swaths on either side of the subsatellite track. These two wind vector swaths are separated by a nadir gap of 329 km. The extent of the nadir gap and the wind vector swaths is constrained so that the wind vector swaths are observed from incidence angles within the plateau region, 20° to 65°.

Since the individual antenna footprints are much longer than the desired resolution, the illumination patterns are resolved into smaller elements using signal processing techniques. Doppler filtering and measurement timing resolve the NSCAT antenna footprints into resolution elements, or "cells" [1]. The resolution of each cell is 25 km square. However, the high resolution cells are usually grouped into 50 km cells in order to increase the number of measurements used in wind estimation. All of the wind retrieval methods used in this thesis operate on the 50 km scale.

Figure 2.8 graphs a portion of an NSCAT measurement swath for the 50 km scale. Note the two wind vector swaths on either side of the nadir gap. The cells of the wind vector swaths are located by their along-track and cross-track indices. The along-track numbers increase in the subsatellite track direction, and they range from 1 to 401. The cross-track numbers range from 1 at the leftmost cell of the left wind vector swath to 24 at the rightmost cell of the right wind vector swath, where left and right are relative to the subsatellite track.

This measurement geometry permits measurement from multiple azimuth angles to facilitate wind retrieval.

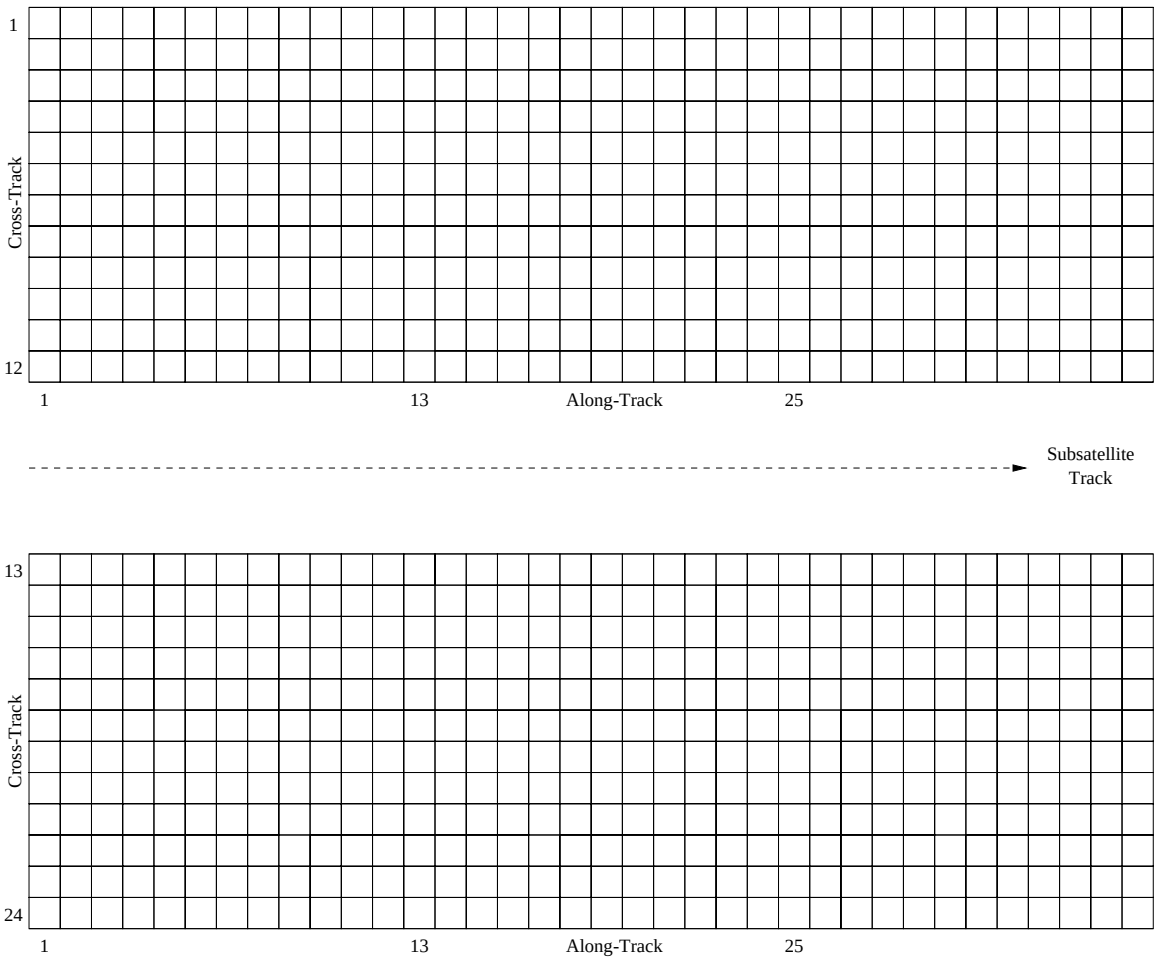

Figure 2.8: Portion of an NSCAT 50 km resolution measurement swath.

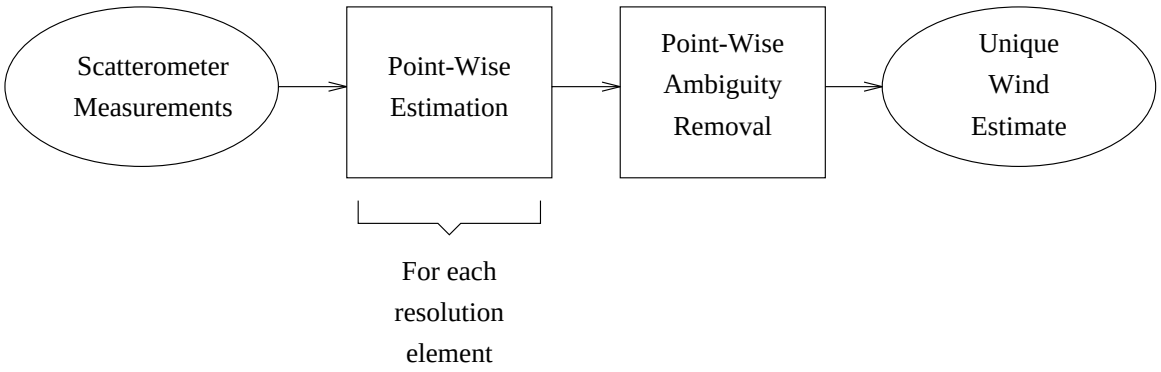

Figure 2.9: Point-wise wind retrieval.

2.5 Point-Wise Wind Retrieval

Point-wise wind retrieval is the traditional method of recovering wind velocity from backscatter measurements. Figure 2.9 is a flow diagram of point-wise retrieval, which consists of point-wise estimation and ambiguity removal.

2.5.1 Point-Wise Estimation

Point-wise estimation inverts the geophysical model function for each resolution element separately. An actual inversion scenario for a 50 km resolution cell in NSCAT ascending revolution 2454 is depicted in Fig. 2.10. Observe that there is no point where all 17 measurement curves intersect. (The 17 measurements are taken at 25km resolution and then grouped into a single 50km cell.) However, there are several areas of near intersection, and it is reasonable to assume that the wind that affected the measurements lies somewhere in these areas. In order to narrow the estimate down to a set of points, it is necessary to define an estimator to locate an "optimal" set of solutions. Multiple solutions arise because of the harmonic nature of the geophysical model function [16]. The noisy measurements make it difficult to discern between valid near intersections and spurious near intersections. Although there are numerous possible estimators, the one used in this thesis is the maximum likelihood approach, since it is believed to work best for the 50 km case [17].

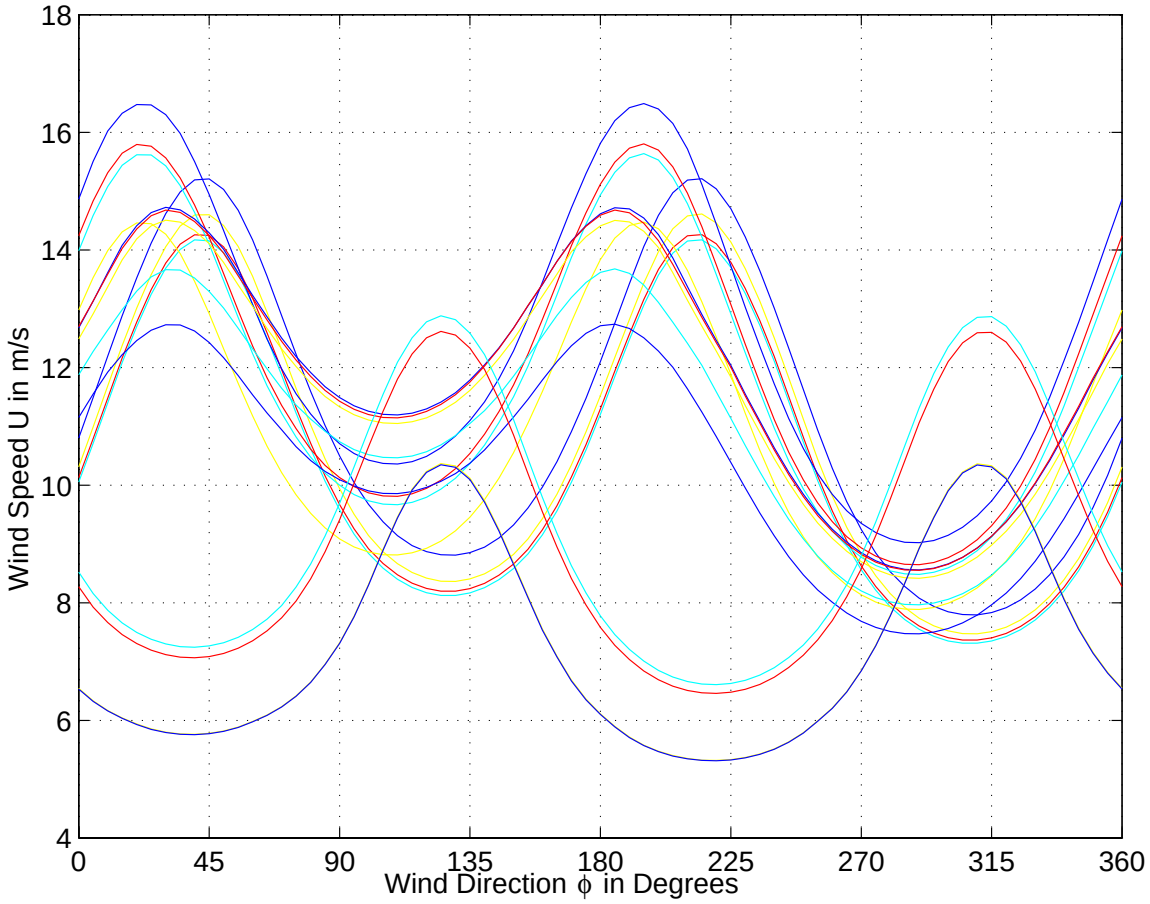

Figure 2.10: An actual point-wise estimation scenario for a cell in NSCAT ascending revolution 2454. Observe that there is no point where all 17 measurement curves intersect. (The 17 measurements are taken at 25km resolution and then grouped into a single 50km cell.) However, there are several areas of near intersection, and it is reasonable to assume that the wind that affected the measurements lies somewhere in these areas.

The formulation of the point-wise estimator begins with characterization of the measurement noise. Once the statistical nature of the noise is described, the probability that NSCAT would observe a given measurement given that the true wind has a certain velocity is described. The probability is extended for the case of multiple collocated measurements, and the maximum likelihood estimator is introduced.

Let σ_{0t} denote the true radar cross section for the wind velocity (U_t, ψ_t), i.e. $\sigma_{0t} = \mathcal{M}(U_t, \psi_t)$, where the dependence on χ, θ, f, and p is implied. The observed σ_0, denoted z, is modeled as a gaussian random variable with mean σ_{0t} and variance ς^2 [13]. Hence, the probability that NSCAT would observe the measurement z given that the true radar cross section is σ_{0t} is

$$p(z|\sigma_{0t}) = \frac{1}{\sqrt{2\pi\varsigma^2}} \exp\left[-\frac{(z-\sigma_{0t})^2}{2\varsigma^2}\right].$$

Since the equality $\sigma_{0t} = \mathcal{M}(U_t, \psi_t)$ is assumed to be true, the probability may be written as

$$p(z|(U_t, \psi_t)) = \frac{1}{\sqrt{2\pi\varsigma^2}} \exp\left[-\frac{(z-\mathcal{M}(U_t, \psi_t))^2}{2\varsigma^2}\right].$$

In order to simplify notation in the case of several measurements, denote the k^{th} measurement as z_k and define the true radar cross section σ_{0tk} as $\mathcal{M}_k$ [13]. Write the variance of the random variable z_k as

$$\varsigma_k^2 = \alpha_k \mathcal{M}_k^2 + \beta_k \mathcal{M}_k + \gamma_k.$$

The noise model parameters α_k, β_k, and γ_k are known for each measurement. Denote the vector wind in component form, and represent the measurements as the vector $\mathbf{z}$. Then the probability that NSCAT would observe the measurements $\mathbf{z}$ given that the true vector wind is $\mathbf{w}$ is [13]

$$p(\mathbf{z}|\mathbf{w}) = \prod_{k=1}^{K} \frac{1}{\sqrt{2\pi\varsigma_k^2}} \exp\left[-\frac{(z-\mathcal{M}(\mathbf{w}))^2}{2\varsigma_k^2}\right].$$

The maximum likelihood approach to wind estimation selects the set of vector winds that locally maximizes $p(\mathbf{z}|\mathbf{w})$, which is called an objective function. Since the dynamic range of the probability can be very large, the natural logarithm

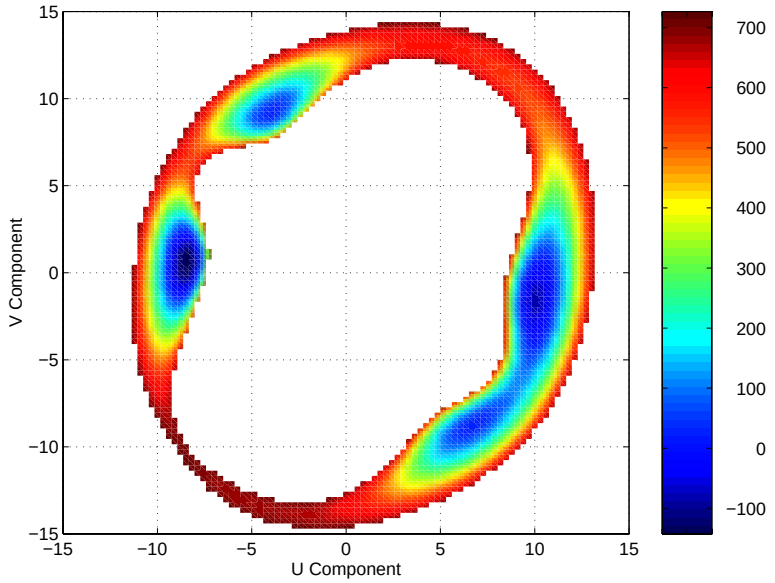

Figure 2.11: Point-wise objective function for NSCAT ascending revolution 2454, along-track 81 and cross-track 3.

is optimized. Most local optimizers minimize functions, so the point-wise objective function used in practice is [1]

$$J_{PW}(\mathbf{w}) = -\sum_{k=1}^{K}\left[\frac{(z_k - \mathcal{M}(\mathbf{w}))^2}{\varsigma_k^2} + \ln(\varsigma_k^2)\right].$$

Figure 2.11 is an image of the objective function for NSCAT ascending revolution 2454, along-track 81 and cross-track 3. Each point in the image represents the objective function value for a particular vector wind in (u, v) form. The high objective function values have been omitted, leaving a ring-shaped area where the most likely solutions lie. Solutions may be ranked according to their likelihoods, with solutions having lower objective function values being more likely and more highly ranked. There are 4 local minima in this example. They are the point-wise estimates, and are, in order from most likely to least likely, $(-8.5, 0.6)$, $(10.0, -1.6)$, $(6.6, -8.8)$, and $(-4.0, 9.2)$. The objective function values at the minima are -141.4, -77.6, 27.0, and 28.7, respectively. Figure 2.12 plots the point-wise wind estimates for a portion of NSCAT ascending revolution 2454 that includes the cell in the example above.

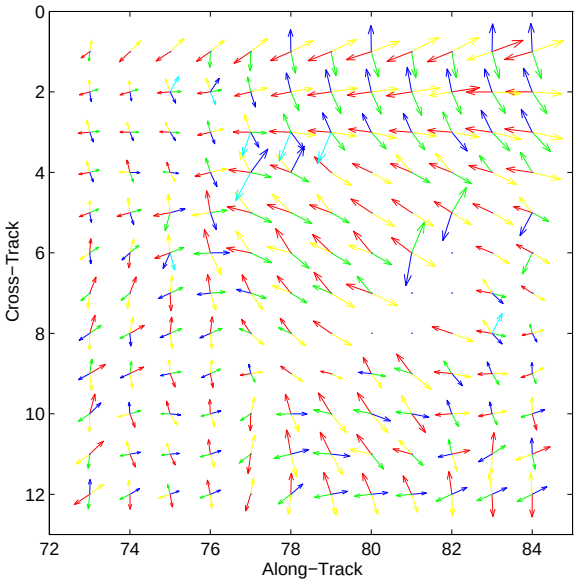

Figure 2.12: Point-wise wind estimates for a portion of NSCAT ascending revolution 2454 that includes the cell in Fig. 2.11. The rankings of the ambiguities are indicated by color. Red is the highest rank, yellow the second, then green, blue, cyan, and magenta.

2.5.2 Point-Wise Ambiguity Removal

The presence of multiple solutions, or ambiguities, in point-wise estimates requires an additional processing step, point-wise ambiguity removal, to select a unique wind vector for each cell. Many different ambiguity removal schemes exist, and a brief overview of some of the techniques is provided.

Several techniques are reported in [16] for use with three-beam scatterometers like NSCAT. The three techniques are the University of Wisconsin, Remote Sensing Systems (RSS), and University of Kansas methods. The University of Wisconsin method attempts to recognize patterns in streamline of the highest ranked ambiguities, while the RSS method uses wind field continuity considerations. The University of Kansas method incorporates meteorological procedures into ambiguity removal. All of the three techniques use only scatterometer data.

Reference [18] presents a technique that incorporates surface pressure analysis fields. The pressure fields are used in conjunction with a boundary layer model to compute unique wind velocity from wind speeds provided by the scatterometer. The author asserts that the method could be used as an ambiguity removal algorithm.

One of the more recent ambiguity removal algorithms is based on an extension of the median filter to vector processing and was developed for NSCAT wind retrieval [1, 19]. The algorithm begins with an initial field, which is usually constructed using the most likely ambiguities. The vector median filter attempts to smooth the initial field without removing abrupt changes in the field due to fronts, cyclones, etc. Processing using the median filter proceeds as follows:

1. Center a 7x7 element window on the cell at location (i, j)

2. For each of the k ambiguities $\mathbf{A}_{ij}^k$ at (i, j), calculate

$$E_{ij}^k = (L_{ij}^k)^{-2} \sum_{m=i-3}^{i+3} \sum_{n=j-3}^{j+3} \left\| \mathbf{A}_{ij}^k - \mathbf{U}_{mn} \right\|,$$

   where L_{ij}^k is the likelihood of the k^{th} ambiguity in the center of the window. $\mathbf{U}_{mn}$ denotes the vector at the location (m, n) in the window

3. Substitute the ambiguity $\mathbf{A}_{ij}^k$ that minimizes E_{ij}^k into location (i, j) for the next iteration

4. Move the filter center (i, j) to the next cell, and repeat process from step 1; continue until desired convergence is reached.

Reference [1] reports ambiguity removal skills in excess of 96% on simulated NSCAT data.

A particularly successful application of the median filter is the "nudged" median filter algorithm for NSCAT data processing. National Center for Environmental Protection (NCEP) 2.5° wind fields determine the initial value field for the median filter [12]. The initial field is constructed from the closest first or second ranked ambiguity to the NCEP wind [20]. The median filter is then run on this field.

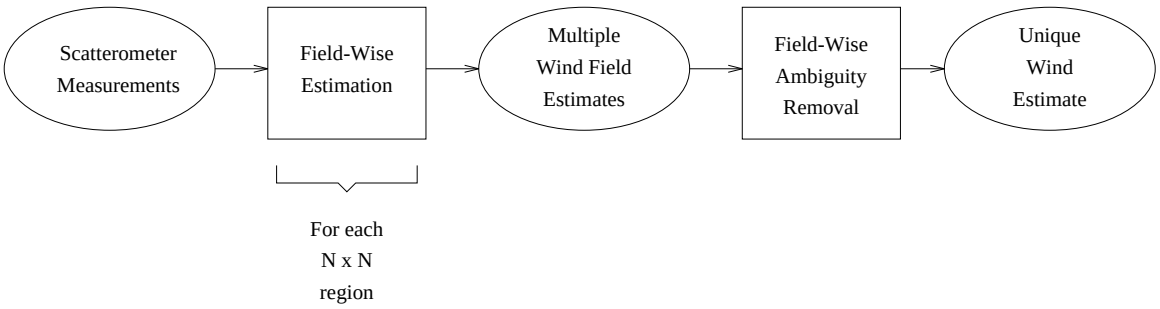

Figure 2.13: Field-wise wind retrieval.

This technique is remarkably successful, due to NSCAT's high first and second ambiguity skill. About 90% of the time one of the two most likely ambiguities is the closest ambiguity to the true wind, and the first ambiguity is the closest about 60% of the time [19]. The output of the "nudged" median filter is the standard NSCAT JPL product.

2.6 Introduction to Field-Wise Wind Retrieval

One alternative to point-wise wind retrieval that relies only on scatterometer data is field-wise retrieval [7, 8, 21]. Although field-wise estimation is the topic of this thesis, an overview of the entire process of field-wise retrieval proves useful, as the estimation stage is an integral part of the retrieval process. An introduction to the topic is presented and field-wise estimation is expanded in the following chapter.

Figure 2.13 displays an overview of field-wise retrieval. The major processing steps are field-wise estimation and field-wise ambiguity removal, which are analogous to their point-wise counterparts, except that they process multiple cells simultaneously using a wind field model. Before proceeding with field-wise estimation and ambiguity removal, the Karhunen-Loeve wind field model is introduced.

2.6.1 Karhunen-Loeve Wind Field Model

The Karhunen-Loeve (KL) model [13] is linear and data-driven and models MxN wind fields. It is expressed by the model matrix F, which describes the functional relationship between the parameter vector $\mathbf{X}$ and the model-based wind field $\mathbf{W}$.

By convention, $\mathbf{W}$ and $\mathbf{X}$ are column vectors. $\mathbf{W}$ contains the vector components of the wind for a MxN wind field:

$$\mathbf{W} = \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ u_{MN} \\ v_1 \\ v_2 \\ \vdots \\ v_{MN} \end{pmatrix}.$$

The u and v components are mapped from the MxN wind field so that along-track varies more rapidly than cross-track. For example, if the upper left corner of the modeled region is at along-track 73 and cross-track 1, or $(73, 1)$, then (u_1, v_1) is the wind in the cell at $(73, 1)$, (u_2, v_2) is the wind at $(74, 1)$, etc.

The model-based wind field $\mathbf{W}$ is expressed as a linear combination of the columns of F, which are called basis vectors, and the number of basis vectors is the order of the model. The linear combination is written compactly in matrix notation:

$$\mathbf{W} = F\mathbf{X}.$$

Several techniques for calculating KL models exist. The KL model used in this work is obtained from [22]. A biased estimate of the autocorrelation matrix R from over 50000 MxN regions of NSCAT wind fields is generated, where $M = N = 12$:

$$R \approx \frac{1}{P} \sum_{i=1}^{N} W_i W_i^T.$$

The matrix F is obtained by singular value decomposition of R. Since R is Hermitian and positive semidefinite, the singular value decomposition is $R = F\Lambda F$, where

$$\Lambda = \begin{pmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_{2MN} \end{pmatrix}$$

and $\lambda_1, \lambda_2, ..., \lambda_{2MN}$ are the eigenvalues of R, sorted so that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{2MN}$ [23]. The columns of F are orthonormal eigenvectors of R, and they appear in the same order as their corresponding eigenvalues. Thus, the first column of F is the eigenvector corresponding to λ_1, the second column corresponds to λ_2, and so on. The associated eigenvalues determine the importance of the columns of F in modeling wind, with the larger eigenvalues corresponding to more important basis vectors [13]. This fact becomes important when using models with fewer than 2MN columns, or truncated models, to represent a MxN wind field.

Selecting the order of a truncated model is a trade-off between modeling error and number of model parameters. It is desirable to model wind with as few parameters as possible because the computational complexity of field-wise retrieval increases dramatically with the number of model parameters. However, the lower the order, the higher the modeling error. See [24, 25] for further development of model order selection. For the purposes of this thesis, a 22 parameter model is adequate [22]. It is used throughout this work and is denoted F_0.

The least-squares model fit [7] obtains a model-based wind field $\hat{\mathbf{W}}$ that optimally models the wind field $\mathbf{W}$ in a least-squares sense. Since F_0 is not invertible (it is not full rank), there may be no $\mathbf{X}$ that exactly solves $\mathbf{W} = F\mathbf{X}$. However, the left pseudoinverse of F_0 can be used to find a solution that minimizes the mean squared difference from $\mathbf{W}$. The left pseudoinverse of F_0 is $F_0^\dagger = (F_0^T F_0)^{-1} F_0^T$, and the least-squares solution of $\mathbf{W} = F_0\mathbf{X}$ is $\hat{\mathbf{X}} = F^\dagger\mathbf{W}$[26, 27]. $\hat{\mathbf{W}}$ is the model-based wind field.

2.6.2 Field-Wise Estimation

As mentioned in the introduction, model-based retrieval [7] is the predecessor of field-wise estimation. Model-based retrieval requires an initial wind field $\mathbf{W}$. A least-squares model fit is performed: $\hat{\mathbf{X}} = F^{\dagger}\mathbf{W}$ to yield the model-based wind field $\hat{\mathbf{W}}$. $\hat{\mathbf{X}}$ is refined by locally optimizing the field-wise objective function with $\hat{\mathbf{X}}$ as the initial value. The field-wise objective function is described below.

The formulation of the field-wise estimator is similar to that of the point-wise estimator. However, the field-wise approach deals with many resolution elements at once using the wind field model F_0.

Using the noise model described for the point-wise case, the probability that NSCAT would observe a given *region* of measurements given that the true wind field is a certain *vector field* can be described. In order to simplify the expression of this probability, it is necessary to introduce notation from [13].

Let the wind field vector $\mathbf{W}$ be indexed by l, which runs from 1 to MN, where the wind field mapped into $\mathbf{W}$ is MxN. Group the radar cross section measurements for the whole region into a single block vector:

$$\mathbf{Z} = \begin{pmatrix} \mathbf{z_1} \\ \mathbf{z_2} \\ \vdots \\ \mathbf{z_l} \\ \vdots \\ \mathbf{z_{MN}} \end{pmatrix},$$

where the vectors $\mathbf{z_l}$ composing the block vector are identical to the single measurement vectors defined for point-wise estimation. The vectors $\mathbf{z_l}$ are variable length. Let the length of $\mathbf{z_l}$ be denoted as $K(l)$. The k^{th} measurement in the l^{th} cell is written $Z_{k,l}$. Write the true radar cross section for the k^{th} measurement in the l^{th} cell as $\mathcal{M}_{k,l}(\mathbf{w_l})$. The subscript k denotes the implied dependence on the incidence angle, azimuth angle, frequency, and polarization for the k^{th} measurement. The noise model parameters for the k^{th} measurement in the l^{th} cell are denoted $\alpha_{k,l}$, $\beta_{k,l}$, and $\gamma_{k,l}$.

The desired probability may be expressed in terms of these notational conventions as

$$\begin{aligned} p(\mathbf{Z}|\mathbf{W}) &= \prod_{l=1}^{MN} \prod_{k=1}^{K(l)} \frac{1}{\sqrt{2\pi\varsigma_{k,l}^2}} \exp\left[\frac{-(Z_{k,l} - M_{k,l}(\mathbf{w_l}))^2}{2\varsigma_{k,l}^2}\right], \\ \varsigma_{k,l}^2 &= \alpha_{k,l}\mathcal{M}_{k,l}^2 + \beta_{k,l}\mathcal{M}_{k,l} + \gamma_{k,l}. \end{aligned}$$

The maximum likelihood approach to field-wise wind estimation selects the set of wind fields that locally maximize the objective function $p(\mathbf{Z}|\mathbf{W})$. Since it is much less computationally intensive to search the space of model parameters $\mathbf{X}$ than to explore all possible wind fields $\mathbf{W}$, the relation $\mathbf{W} = F_0\mathbf{X}$ is invoked to write the objective function equivalently as $p(\mathbf{Z}|\mathbf{X})$.

As in point-wise estimation, the objective function used in practice is the negative of the natural logarithm of $p(\mathbf{Z}|\mathbf{X})$:

$$J_{FW}(\mathbf{X}) = -\sum_{l=1}^{MN} \sum_{k=1}^{K(l)} \left[\frac{(Z_{k,l} - M_{k,l}(\mathbf{w_l}))^2}{\vartheta_{k,l}^2}\right].$$

Since $J_{FW}(\mathbf{X})$ has multiple local minima, the model parameter sets that locally optimize $J_{FW}(\mathbf{X})$ are denoted $\hat{\mathbf{X}}_\mathbf{i}$, and they represent wind field solutions via the modeling relationship $\hat{\mathbf{W}} = F_0\hat{\mathbf{X}}$.

2.6.3 Field-Wise Ambiguity Removal

As in the point-wise scheme, field-wise estimation produces multiple ambiguities because of the multiple local minima of the objective function $J_{FW}(\mathbf{X})$. However, since the ambiguities are fields of vectors rather than individual vectors, adjacent solutions can be made to overlap, and continuity considerations can be invoked to greatly simplify ambiguity removal [7, 8]. Model-based estimation is used to further refine the solutions selected in field-wise ambiguity removal.

Chapter 3

Field-Wise Estimation Techniques

3.1 Introduction

Field-wise estimation is essentially global optimization of the objective function $J_{FW}(\mathbf{X})$. However, due to the nature of the geophysical model function, the global minimum is not the only desired solution. Field-wise estimation must locate the *multiple* near-global local minima of the field-wise objective function. Since the process of locating the minima must be repeated for each region of data, it is essential that the global optimization routine be reasonably fast. In addition, the nature of the field-wise objective function places further constraints on the choice of global optimization methods.

In order to clarify the choice of field-wise estimation routines, the general features of the field-wise objective function are explored, and several global optimization methods are considered in light of these features. It is determined that since multiple local minima must be located in a reasonable amount of time, a simple multistart algorithm is the most reasonable choice for field-wise estimation.

A variant of the traditional multistart algorithm is developed for field-wise estimation. A local optimization scheme using median filtered data is used to augment the multistart solutions. The multistart algorithm is analyzed and both of the algorithms are tested on NSCAT data.

3.2 Characterization of the Field-wise Objective Function

The characteristics of the objective function constrain the viability of global optimization algorithms in field-wise estimation. Some of the salient features of the objective function are discussed in this section. These characteristics include capture region size, number of local minima, and solution symmetry. Before probing these aspects of the objective function, a brief overview of the general features of the objective function is provided.

3.2.1 General Features

If the time required to run field-wise estimation were not an issue, it would be a relatively straightforward problem, because the objective function is relatively tame. Even though the objective function is nonlinear and exhibits multiple local minima, it is moderately low-order, since the wind field model F_0 has only 22 basis vectors. Further, the time to evaluate the objective function is approximately 0.1 CPU seconds per evaluation, and the objective function has an analytic gradient. However, since the field-wise global optimization problem must be solved repeatedly in a reasonable amount of time, the difficulty is compounded, and the field-wise estimation problem is intractable for many optimization algorithms.

3.2.2 Capture Region Size

A gradient-descent algorithm started from a given point in the model parameter space defines a path that leads to a local minimum. The capture region of a local minimum is defined as the set of all points that locally optimize to the local minimum using a gradient-descent algorithm. Capture region size describes how difficult it is for the global optimization routine to locate a given local minimum. Capture region size is usually represented by the fraction of the domain of interest occupied by the capture region.

Consider, for example, a simple objective function with two parameters. The camel hump function [28],

$$f(x, y) = ax^2 + bx^4 + cx^6 - xy + dy^2 + ey^4,$$

where $a = 4, b = -2.1, c = 1/3, d = -4$, and $e = 4$, has two local minima in the domain $x \in [-0.5, 0.5], y \in [-1, 1]$. Figure 3.1 is a plot of the camel hump function. The leftmost plot in Fig. 3.2 shows the normalized negative gradient of the camel hump function. The negative gradient points in the "downhill" direction of the objective function, which is the direction that a gradient-descent algorithm follows. For reference, all graphs in Fig. 3.2 also display contours of the camel hump function.

In order to graph the capture regions of the camel hump function, 5000 (x, y) pairs are randomly selected from the domain $x \in [-0.5, 0.5], y \in [-1, 1]$. Each of the pairs is used as an initial value in a gradient-descent local optimization. The center and rightmost plots of Fig. 3.2 depict the initial value points that optimize to the two different local minima. The number of points that optimize to the local minimum at $(0.09, 0.7)$ is 2453, and the number that optimize to the local minimum at $(-0.09, -0.7)$ is 2547. Reasonable estimates of the capture region sizes are $2453/5000 = 0.49$ and $2547/5000 = 0.51$. These numbers represent the fraction of the domain area, or solution space, that is occupied by each capture region. In higher dimensional objective functions, the extent of the capture region and the size of the domain are hypervolumes, but the capture region size is a dimensionless quantity.

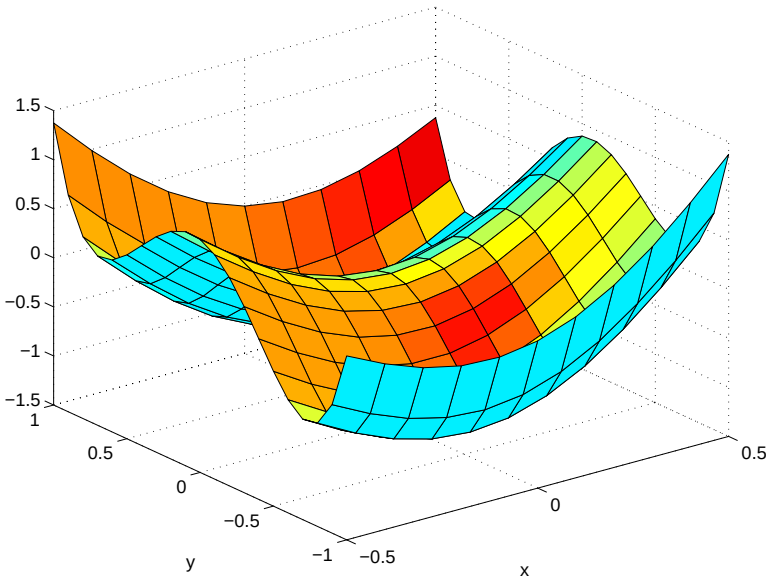

Figure 3.1: Camel hump function. The camel hump function is an example of a function of 2 parameters that exhibits multiple minima. The domain of the graph is restricted to display only 2 of the 6 minima of the function.

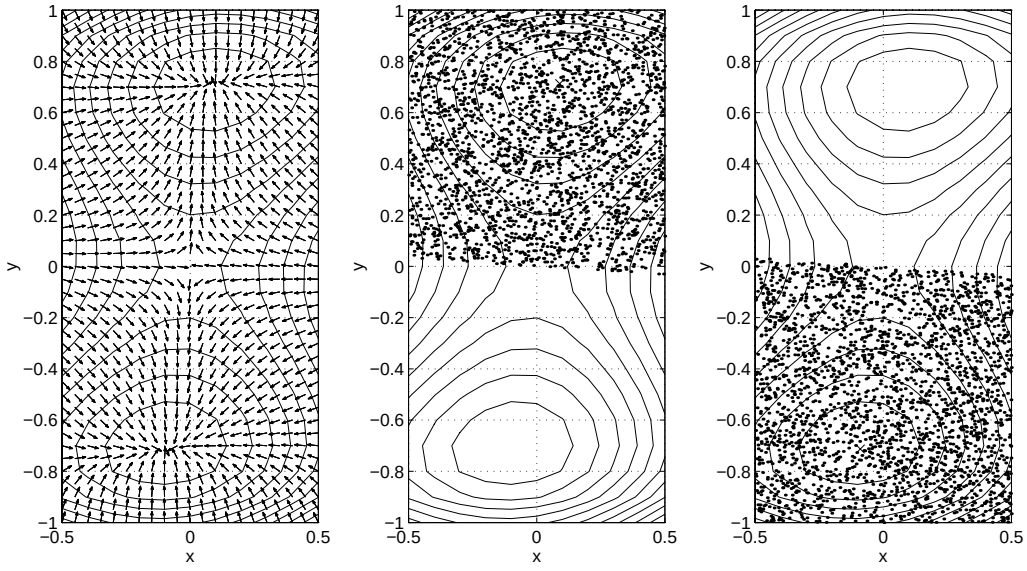

Figure 3.2: Capture regions of the camel hump function. The capture regions of the two minima in Fig. 3.1 are illustrated in the left plot by a vector field of the normalized negative gradient, which indicates the "downhill" direction of the function. The center and rightmost plots depict the initial value points that descend to each minimum as a result of minimization. Contour plots in each figure indicate the structure of the function. The local minima are at $(0.09, 0.7)$ and $(-0.09, -0.7)$.

Capture region sizes for the field-wise objective function are estimated using a method similar to that employed in the camel hump example. Unlike the camel hump function, the field-wise objective function displays different characteristics for different wind field phenomena. Accordingly, three representative wind fields with various features are selected from NSCAT revolution 2454 ascending. Figures 3.3, 3.4, and 3.5 show the JPL product for the three wind fields. Figure 3.3 is a smooth wind field, while Figs. 3.4 and 3.5 are more interesting features. Figure 3.4 is a front and Fig. 3.5 is a cyclone. Since the JPL product for these wind fields appears reasonable, the optimized least-squares model fit to the JPL product is assumed to be within the capture region of, or identical to, the optimized least-squares model fit to the true wind; thus, the capture region properties of the desired solution can be investigated by probing the capture region of the optimized JPL product. Figures 3.6, 3.7, and 3.8 display the optimized model fits to the JPL product.

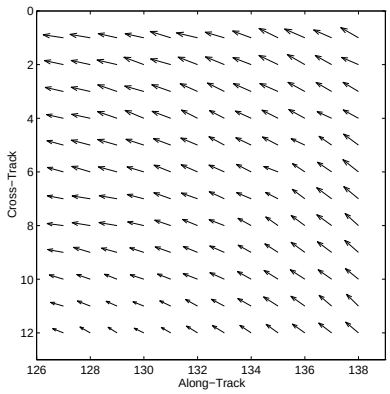

Figure 3.3: Smooth wind field taken from JPL product from NSCAT ascending revolution 2454.

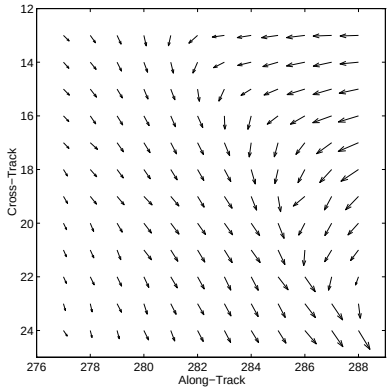

Figure 3.4: Frontal feature taken from JPL product from NSCAT ascending revolution 2454.

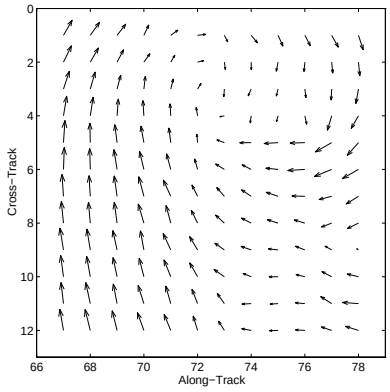

Figure 3.5: Cyclone taken from JPL product from NSCAT ascending revolution 2454.

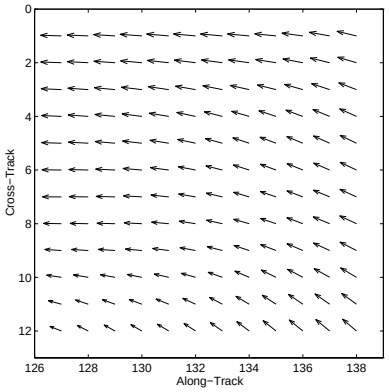

Figure 3.6: Smooth wind field. The wind field in Fig. 3.3 is optimized to yield this plot.

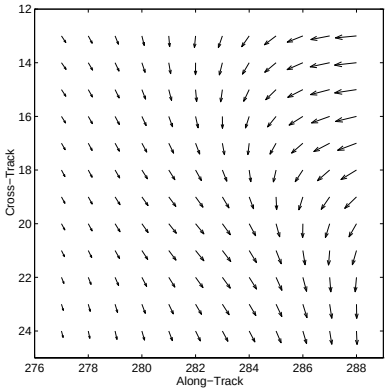

Figure 3.7: Frontal feature. The wind field in Fig. 3.4 is optimized to yield this plot.

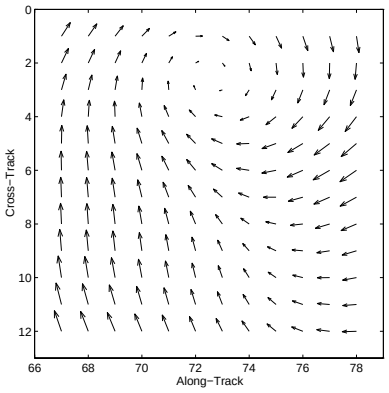

Figure 3.8: Cyclone. The wind field in Fig. 3.5 is optimized to yield this plot.

The multistart global optimization method is used to gauge the capture region size of the three representative regions. Multistart consists of three basic steps [29]:

1. Randomly generate an initial value (uniform distribution)
2. Locally optimize the initial value
3. If stopping criteria have not been met, continue with (1).

The ratio of the the number of times a local minimum is reached to the total number of initial values is an estimate of the local minimum's capture region.

In order to determine capture region size, the multistart algorithm is run on the three representative regions using 1000 initial values. The initial values are randomly chosen according to a uniform product distribution defined on a limited portion of the model parameter space. Justification of the limits on the model parameter space are deferred until the description of the field-wise multistart algorithm.

The product distribution is formed by examining the variances of the KL model parameters, which are simply the eigenvalues of the autocorrelation matrix R used in calculating the KL model. Let $\sigma_1^2 = \lambda_1$ and $\sigma_2^2 = \lambda_2$, where λ_1 and λ_2 are the eigenvalues corresponding to the first 2 KL model basis vectors. Let the model parameter vector $\mathbf{X}_i = (x_1, x_2, ..., x_{12}, 0_{1\times 10})^T$ denote the i^{th} initial value, where $0_{1\times 10}$ is a 10 element row vector with all zero entries. Each nonzero element of $\mathbf{X}_i$ is drawn independently from different uniform populations:

$$\begin{aligned} x_1 &\in [-2\sigma_1, 2\sigma_1] \\ x_2 &\in [-2\sigma_2, 2\sigma_2] \\ x_3 &\in [-120, 120] \\ &\vdots \\ x_{12} &\in [-120, 120]. \end{aligned}$$

Each initial value is locally optimized, and the number of initial values that converge to the optimized JPL solution is recorded. An initial value is said to converge to the

Revolution	Location	Feature	Capture Region Size
2454 ascending	Along-Track 127 Cross-Track 1	Smooth Field	0.094
2454 ascending	Along-Track 277 Cross-Track 13	Front	0.028
2454 ascending	Along-Track 67 Cross-Track 1	Cyclone	0.015

Table 3.1: Estimates of capture region size. The multistart algorithm is run with 1000 initial values on the three representative regions: smooth wind field, cyclone, and front. The fraction of the initial values that converge to the optimized JPL solution (within 0.75 m/s vrms) is an estimate of capture region size.

optimized JPL solution if the vrms difference between the optimized initial value and the optimized JPL solution is less than 0.75 m/s. The vrms difference between two model-based wind fields $\mathbf{W_1}$ and $\mathbf{W_2}$ is calculated as follows:

$$\text{Vrms Difference} = \sqrt{\frac{(\mathbf{W}-\mathbf{W})^T(\mathbf{W}-\mathbf{W})}{MN}},$$

where M and N are the dimensions of the wind fields from which $\mathbf{W_1}$ and $\mathbf{W_2}$ are scanned.

Table 3.1 summarizes the capture region size for the three representative regions. Again, the size estimates are for the capture region of the optimized least-squares model fit to the JPL product. For the smooth wind field region, 94 out of the 1000 initial values converge to the optimized JPL solution. An estimate of the capture region size of the optimized JPL solution for the smooth region is 94/1000=0.094. For the region with the front, 28 out of the 1000 initial values converge to the optimized JPL solution. For the cyclone region, only 15 out of the 1000 initial values converge to the optimized JPL solution. Thus, the capture region sizes of the front and cyclone regions are 28/1000=0.028 and 15/1000=0.015, respectively.

Capture region size is probably determined by several factors. Certainly if the data are sufficiently ambiguous to allow a large number of point-wise ambiguities per resolution element, then the field-wise objective function will also have a large number of field-wise ambiguities. Thus, the capture region size will tend to

be smaller. Figures 3.9, 3.10, and 3.11 show the ambiguous point-wise winds for the three representative regions. Observe that there are more point-wise ambiguities in Fig. 3.11 than in Figs. 3.9 and 3.10. Accordingly, the capture region for Fig. 3.11 is larger than those of the other regions.

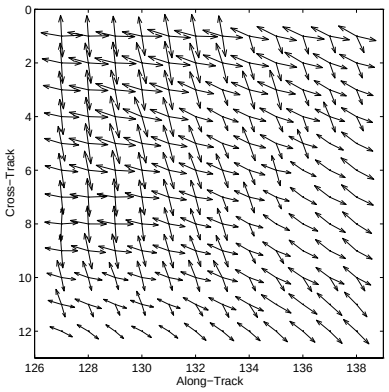

Figure 3.9: All ambiguity plot—smooth wind field region

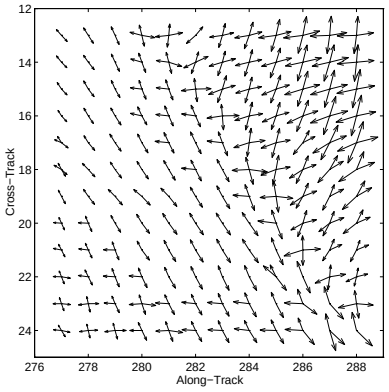

Figure 3.10: All ambiguity plot—frontal feature region

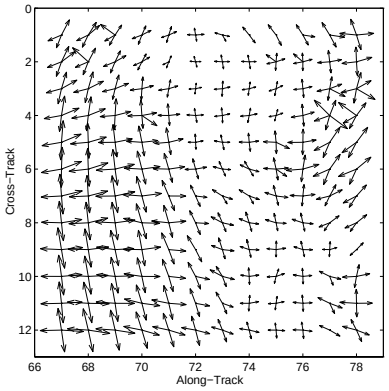

Figure 3.11: All ambiguity plot—cyclone region

However, the number of point-wise ambiguities does not completely explain the capture region sizes in Table 3.1. Even though the region with the front has fewer point-wise ambiguities than the smooth region, it has a smaller capture region. Another factor in determining capture region size is the complexity of the wind field, which would explain the discrepancy between the number of point-wise ambiguities for the frontal region and its capture region size.

3.2.3 Number of Local Minima

The multistart algorithm can also estimate the number of local minima of the field-wise objective function. An obvious method of bounding the number of local minima is to run the multistart algorithm and count the number of unique local minima that it identifies. This method is explored in this subsection.

The multistart algorithm is run on the three representative regions (Figs. 3.3, 3.4, and 3.5) with 200 initial values. The initial values are chosen according to the uniform product distribution described in the previous subsection. The unique local minima are enumerated by retaining only those that differ by at least 0.75 m/s, in a vrms sense. This process is repeated 5 times for each region, and the results are summarized in Table 3.2.

Further results varying the vrms threshold justify the 0.75 m/s value as a reasonable choice for the minimum distance between distinct local minima. The data from the first of the 5 runs are examined for unique local minima using the thresholds of 0.75, 1.5, 2.25, 3, 3.75, 4.5, 5.25, and 6 m/s. The apparent number of unique local minima for each region and each threshold is tabulated in Table 3.3. The rapid decrease in number of apparent local minima, most noticeably observed for the cyclone region, brings the threshold value of 0.75 m/s into question. Since optimization routines locate local minima only approximately, it is possible that the threshold separates solutions that approximate the same local minimum. However, subjective examination of the solutions that the 0.75 m/s value declares unique indicates that they should indeed be treated as separate solutions. Thus, the anomaly observed in

Run Number	Region		
	Smooth	Front	Cyclone
1	147	153	170
2	157	145	174
3	152	150	160
4	144	140	159
5	157	141	164
Ave	151.4	145.8	165.4
Std	5.9	5.6	6.5
Total Number of Unique Minima for All Five Runs	672	668	643

Table 3.2: Number of local minima for three different regions. The data are obtained by running the multistart algorithm 5 times with 200 initial values for each run. The vrms threshold value of 0.75 m/s is used to define distinct local minima.

Table 3.3 for the cyclone region is probably due to very small capture regions, not to an inordinately restrictive threshold.

Although there are a large number of possible solutions, it is acceptable to miss some of the local minima because only a fraction of these are reasonable. In a geophysical sense, the field-wise objective function provides an effective ranking criterion that tends to rank more reasonable solutions better than less reasonable solutions. However, due to the nature of the geophysical model function and to the presence of noise, the most likely solution is not always the true solution. This is the motivation for locating the multiple near-global local minima.

Figures 3.12, 3.13, and 3.14 graph the objective function value versus solution number, sorted according to ascending objective function value, for the smooth, front, and cyclone regions, respectively. The jump in objective function values in Figs. 3.12 and 3.13 around solutions 4 and 15, respectively, marks the transition between the best solutions and less reasonable ones. Figures 3.15, 3.16, 3.17, and 3.18 plot solutions on either side of the transitions for the smooth and front regions. Figure 3.15 is the first solution for the smooth wind field and is very reasonable. It is the

Threshold (m/s)	Region		
	Smooth	Front	Cyclone
0.75	147	153	170
1.5	147	153	170
2.25	147	152	161
3	147	151	152
3.75	147	149	137
4.5	145	148	134
5.25	144	147	120
6	142	145	109

Table 3.3: Apparent number of local minima as a function of threshold.

closest to the optimized JPL wind. However, the 14^{th} solution (plotted in Fig. 3.16) is somewhat less reasonable. Figures 3.17 and 3.18 plot the first and 21^{st} solutions for the front region. Figure 3.17 is a very reasonable wind field, although it is not the closest to the optimized JPL wind. Figure 3.18, however, is an unreasonable solution.

Figure 3.14 does not display such a sharp transition. Subjective inspection places the transition from the best solutions to the less reliable ones in the neighborhood of the 40^{th} ranked solution. Figures 3.19 and 3.20 plot solutions on either side of the transition. Figure 3.19 is the first solution, and it is reasonable contrasted with the unreasonable 45^{th} solution in Fig. 3.20.

Even though the objective function tends to rank more reasonable solutions higher than less reasonable ones, the global minimum is not always the desired solution. For the smooth wind field, the global minimum is the desired solution (i.e. the closest to the optimized JPL product), but for the region with the front, the desired solution is ranked 5^{th}. The best solution for the cyclone region is ranked 16^{th}. Since the desired solution is not always the global minimum, a field-wise estimation algorithm must locate a *set* of local minima. This required set is difficult to define. But it does not consist of all possible minima—a field-wise estimation algorithm can afford to miss some of the solutions, provided that they are less reasonable ones.

Thus, the required set includes the global minimum and all of the near-global minima, which occur before the transition from the more desirable solutions to the less desirable solutions.

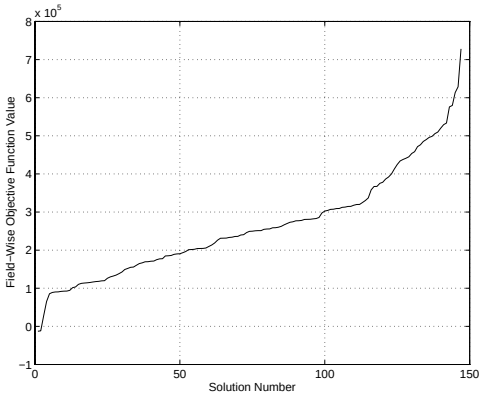

Figure 3.12: Objective function values versus sorted solution number—smooth wind field region. The transition between reasonable and less reasonable solutions is around the 4^{th} ranked solution.

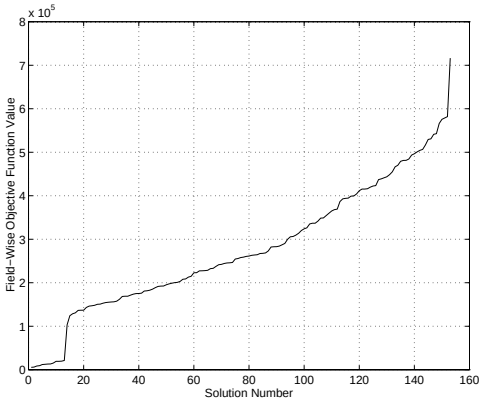

Figure 3.13: Objective function values versus sorted solution number—frontal feature region. The transition between reasonable and unreasonable solutions is around the 15^{th} ranked solution.

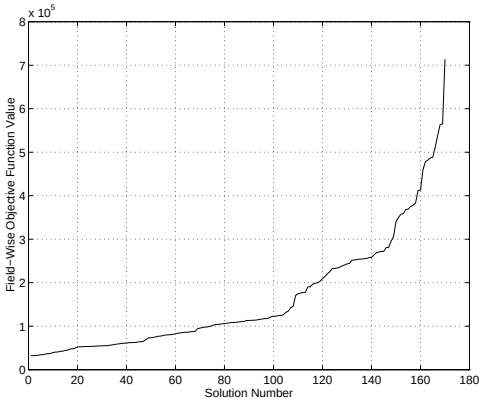

Figure 3.14: Objective function values versus sorted solution number—cyclone region. Subjective inspection of the solutions places the transition from the best ones to the less ones around the 40^{th} ranked solution.

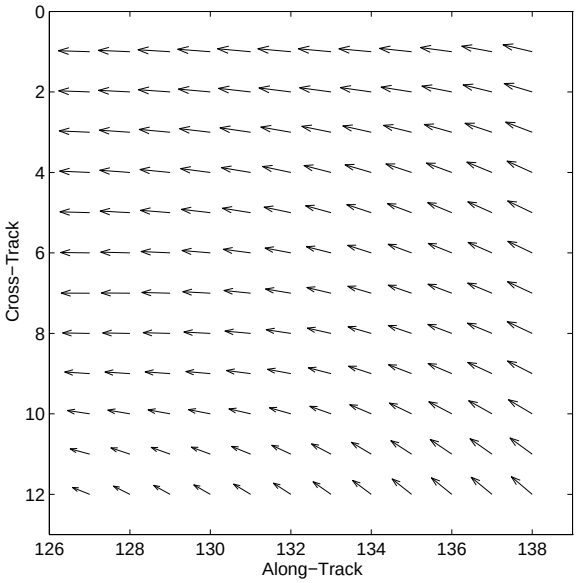

Figure 3.15: First ranked solution—smooth region. This solution occurs before the transition in Fig. 3.12 and is a reasonable solution.

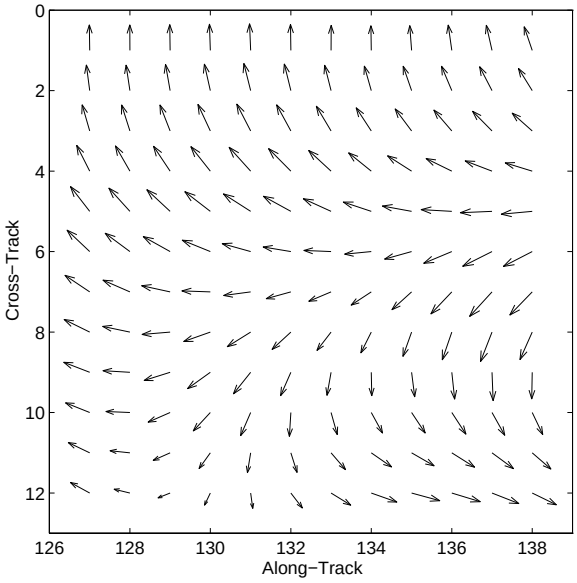

Figure 3.16: 16^{th} ranked solution—smooth region. This solution occurs after the transition in Fig. 3.12 and is a less reasonable solution.

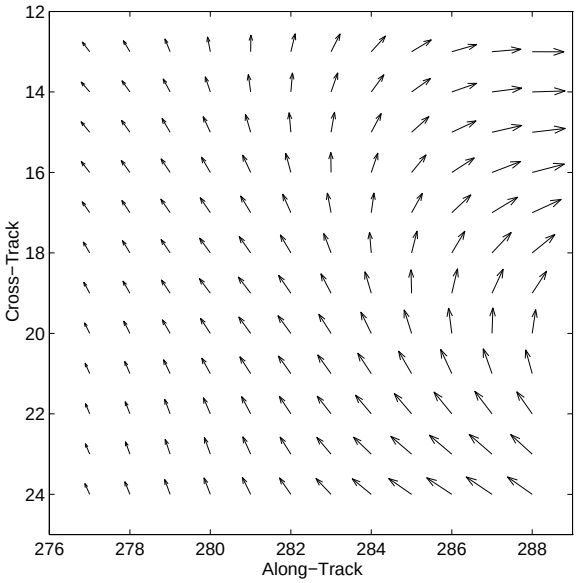

Figure 3.17: First ranked solution—front region. This solution occurs before the transition in Fig. 3.13 and is a reasonable solution.

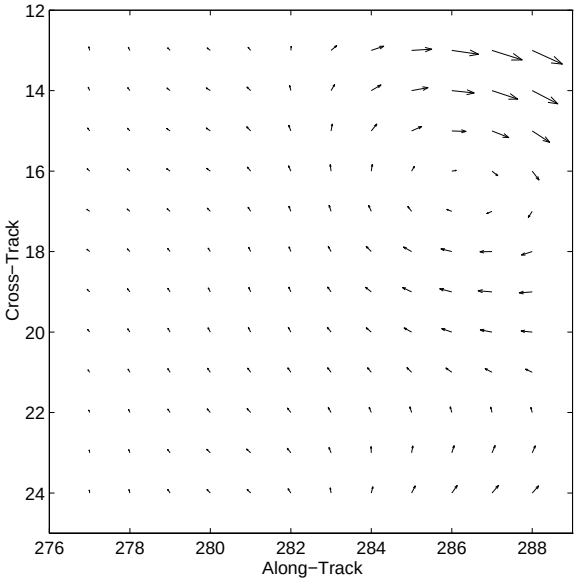

Figure 3.18: 21^{st} ranked solution—front region. This solution occurs after the transition in Fig. 3.13 and is an unreasonable solution.

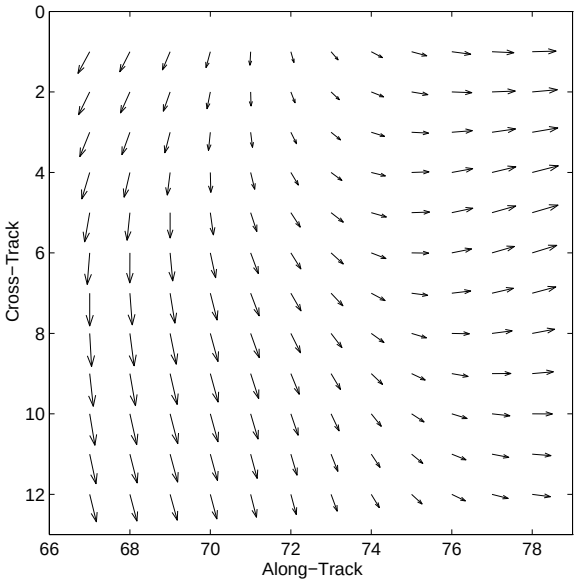

Figure 3.19: First ranked solution—cyclone region. This solution occurs before the transition in Fig. 3.14 and is a reasonable solution.

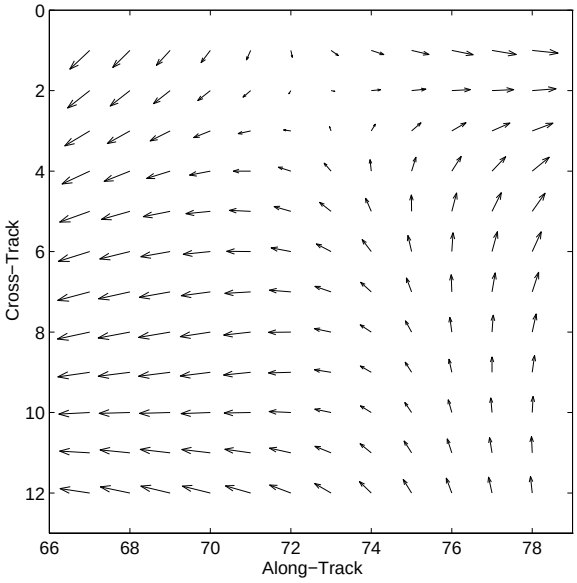

Figure 3.20: 45^{th} ranked solution—cyclone region. This solution occurs after the transition in Fig. 3.14 and is classified as an unreasonable solution.

3.2.4 Symmetry

The field-wise objective function displays a useful symmetry: if there is a local minimum at $\hat{\mathbf{X}}$, there will tend to be a local minimum close to $-\hat{\mathbf{X}}$, where $\hat{\mathbf{X}}$ is the solution model parameter vector. The symmetry in the objective function arises from the ambiguity in the geophysical model function.

Evidence of this symmetry is observed in the representative regions used in the previous subsections. In order to demonstrate the objective function symmetry, the multistart algorithm, as described previously, was run with 5000 initial values to yield 5000 solutions. This set is referred to as initial solutions. Each of the initial solutions was negated, and the resulting parameter vectors were optimized.

If the objective function were perfectly symmetric, all of the negated initial solutions would be identical to the optimized negated initial solutions. A simple two-dimensional example of perfect solution symmetry is plotted in the leftmost graph of Fig. 3.21. However, the field-wise objective function is not perfectly symmetric, so the situation is more like that depicted in the rightmost graph of Fig. 3.21. In the case of imperfect symmetry, the optimized negated initial solutions are close to the negated initial solutions, but not coincident.

Figs. 3.22, 3.23, and 3.24 plot the results of the multistart runs for the smooth, front, and cyclone regions, respectively. The figures plot the 5000 optimized negated initial solutions ($-\hat{\mathbf{X}}$) versus the 5000 negated initial solutions ($-\hat{\mathbf{X}}$) by graphing one model parameter per subplot (for the first 9 model parameters). Thus, the first subplot graphs the first parameter of the optimized negated initial solutions (optimized $-\hat{x}_1$) versus the first parameter of the negated initial solutions ($-\hat{x}_1$), and so on for the first 9 parameters. The strong linear trend (the line with a slope of one through the origin) in all of the subplots indicates that if there is a local minimum at $\hat{\mathbf{X}}$, the optimized value of $-\hat{\mathbf{X}}$ is usually close to $-\hat{\mathbf{X}}$. The horizontal lines are probably due to the upper and lower speed limits of the NSCAT-1 geophysical model function.

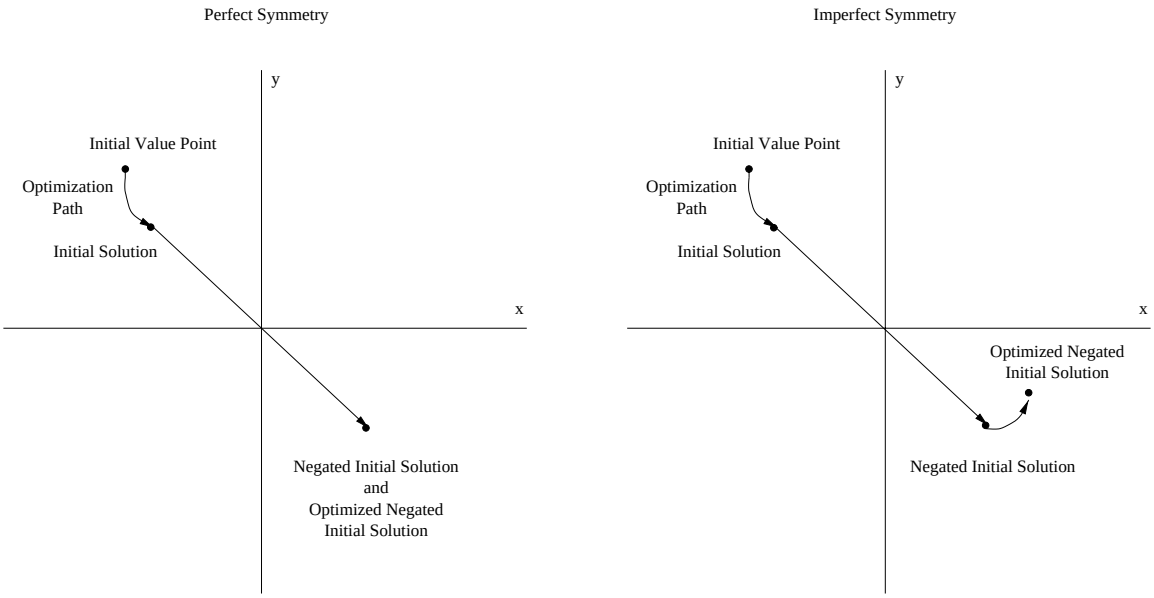

Figure 3.21: Two-dimensional example of perfect and imperfect solution symmetry.

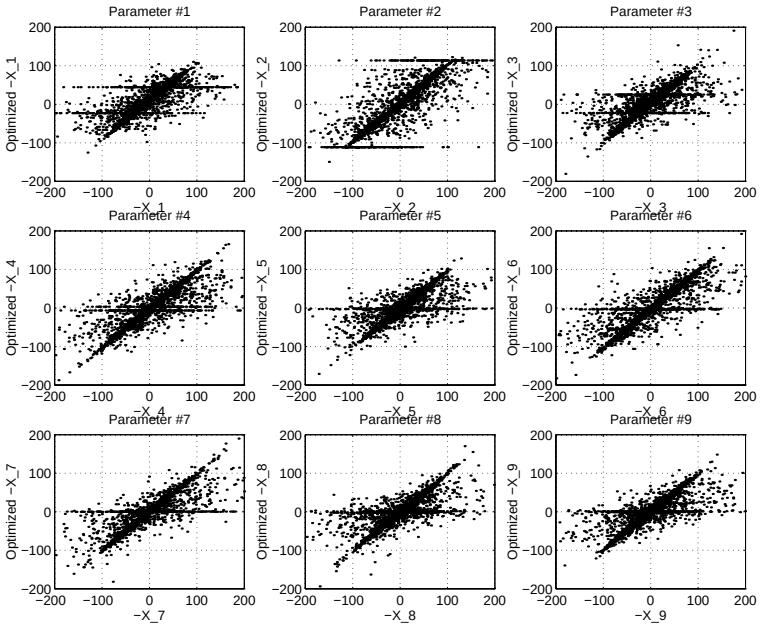

Figure 3.22: Solution symmetry of smooth wind field region. The strong linear trend (the line with a slope of one through the origin) in all of the subplots indicates a pronounced solution symmetry. The horizontal lines are probably due to the upper and lower speed limits of the NSCAT-1 geophysical model function. (X in the legend is equivalent to $\hat{x}$ in the text.)

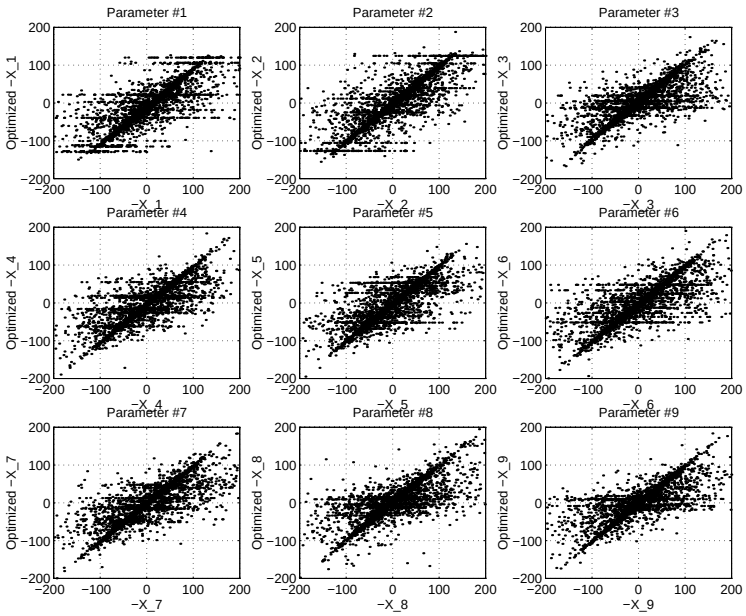

Figure 3.23: Solution symmetry of frontal feature region. The strong linear trend (the line with a slope of one through the origin) in all of the subplots indicates a pronounced solution symmetry. The horizontal lines are probably due to the upper and lower speed limits of the NSCAT-1 geophysical model function. (X in the legend is equivalent to $\hat{x}$ in the text.)

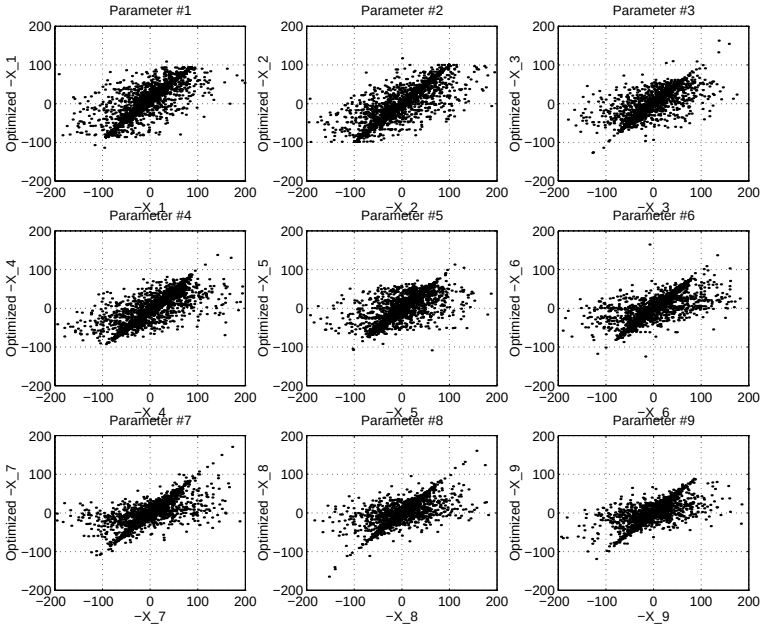

Figure 3.24: Solution symmetry of cyclone region. The strong linear trend (the line with a slope of one through the origin) in all of the subplots indicates a pronounced solution symmetry. The horizontal lines are probably due to the upper and lower speed limits of the NSCAT-1 geophysical model function. (X in the legend is equivalent to $\hat{x}$ in the text.)

3.3 Field-Wise Estimation Algorithms

The choice of global optimization algorithms for field-wise estimation is based primarily on the demand for speed in field-wise estimation and on the nature of the field-wise objective function. The following is a survey of several global optimization algorithms considered for field-wise estimation. Each of the methods is introduced briefly, and strengths and weaknesses relative to the field-wise estimation problem are discussed. Two other novel estimation algorithms are presented in [30, 31]. The multistart method is selected for field-wise estimation because of its simplicity, speed, and ability to locate multiple local minima.

3.3.1 Survey of Global Optimization Methods Considered

Global optimization techniques can be organized into two broad categories: deterministic and stochastic. Several deterministic methods considered for field-wise estimation are interval branch and bound, the Kronecker-Picard (KP) integral method, and a novel systematic search technique. Each of these methods is discussed in turn, and they are evaluated in light of the demands of the field-wise estimation problem.

Interval Branch and Bound

A simple interval branch and bound approach to global minimization divides the parameter space into large hyperintervals, and then seeks to determine lower and upper bounds for the objective function on each of the hyperintervals. If the lower bound on a given hyperinterval exceeds the upper bound on another hyperinterval, then it may safely be assumed that the first does not contain the global minimum of the objective function. The hyperintervals that can possibly contain the global minimum are divided, and the process is repeated until only one hyperinterval remains. The final hyperinterval contains the global minimum of the objective function. A description of a more general branch and bound algorithm and expositions of bounding methods is found in [32]. It is a difficult matter to bound the objective function tightly enough so that the convergence to the hyperinterval containing the

global minimum is rapid. One commonly used technique of bounding the objective function is the interval arithmetic method [32, 33, 34]. The objective function is translated into interval arithmetic functions, which are analogous to point functions, except that they operate on intervals and yield interval bounds as results. Interval arithmetic methods can also incorporate information from the partials of the objective function to produce even tighter bounds.

Several attractions of interval branch and bound is that it can be altered to locate multiple near-global local minima and that it can guarantee convergence to the global minimum [35]. However, the need for fast convergence limits its utility in field-wise estimation.

KP Integral Method

The KP integral method [36, 37] is similar to the branch and bound technique in that the parameter space is divided into increasingly smaller hyperintervals, and the hyperintervals that cannot contain the global minimum are discarded. Instead of bounding the objective function in order to eliminate hyperintervals, the KP integral method evaluates the KP integral, which yields the number of stationary pints within the hyperinterval. Hyperintervals with zero stationary points are discarded, and the remaining hyperintervals are divided. The process is continued until the only remaining hyperintervals contain only one stationary point each. A local optimization method then can be used to find the local minima, from which the global minimum is selected.

The major weakness of the KP integral method is that the KP integral is prohibitively expensive to evaluate for objective functions much higher than two dimensions [37].

Systematic Search

The final deterministic method considered is a novel systematic search [28]. The basic idea is to pass from one local minimum to another via points on the objective function that are like mountain passes. First, a single initial value is locally

optimized. Then, the gradient field is altered in such a way that the "mountain passes", actually called decomposition points, are local minima. Each one of the decomposition points is located by a series of local optimizations from the single local minimum. Finally, points around the decomposition points are locally optimized, using the original gradient field, to find new local minima. These local minima are then used to find additional decomposition points, which are used to find additional local minima. The process is repeated until all of the local minima have been visited.

In spite of its esthetic appeal, this method is also inappropriate for field-wise estimation. The numerous local minima and the dimensionality of the problem make its application to field-wise estimation computationally intractable.

Stochastic Algorithms

In addition to the deterministic methods, two stochastic techniques were considered: genetic algorithms and multistart.

Genetic algorithms treat solutions like chromosomes. A genetic algorithm generates an initial population of chromosomes, and then it evaluates the fitness of each chromosome by computing the objective function value for each corresponding solution. Chromosomes with high fitness are selected to reproduce. In the process of replication, crossover takes place. Crossover produces offspring with solution characteristics inherited from both parent chromosomes. Mutation, the random alteration of the chromosomes, provides further solution variability. As the processes of natural selection and reproduction progress, the population tends to approach the global minimum [38]. An encoding of a simple genetic algorithm from [39] is provided by [40].

Variants of the basic genetic algorithm have been successful at locating multiple local minima [41]. However, genetic algorithms converge too slowly to the minima to be used without extensive modification in field-wise estimation.

The final stochastic algorithm considered was the multistart algorithm, which is outlined in the previous section. Unlike the other methods, which are useful for lower dimensional problems or for high dimensional problems where execution

time of the algorithm is not an issue, the multistart method is designed to locate multiple local minima very rapidly.

Further, multistart is appealing because it is easily described as a simple statistical process, which leads to estimates of capture region size and sequential stopping rules [42, 43], and because it is easily parallelized. Further, more efficient variants of the multistart algorithm, such as clustering [44] and repulsion [29], also exist.

Clustering algorithms seek to speed up the multistart method by intelligently eliminating initial value points. First, as in the multistart method, initial value points are chosen. In the clustering algorithm, however, the initial value points are locally optimized only for a few iterations. These partially optimized initial value points are clustered. For each cluster, only a specified fraction of the points in each cluster are retained. The process of partially optimizing and clustering continues until the desired stopping criteria are achieved.

The repulsion algorithm is a more recent variant of multistart. The repulsion algorithm proceeds exactly as multistart, except that past initial value points and all of the intermediate points of the local optimizations are retained. When a new initial value is selected, the stored points are used to repel the initial value away from regions which have already been explored. In this way, the repulsion algorithm seeks to locally optimize only one initial value per capture region.

3.3.2 Field-Wise Multistart Algorithm

I have developed a variant of the multistart algorithm suited for field-wise estimation. It was introduced in a different context in the subsection that explores objective function symmetry. However, for the sake of completeness and clarity, it is described again here.

In the field-wise multistart method, the initial value points are chosen in two groups, where one is the negative of the other. This approach uses the symmetric structure of the objective function to make the multistart method more efficient. First, 50 points in the 22-parameter model space are chosen according to a uniform

product distribution defined on that space. The product distribution is the same as was described previously. Recall that $\sigma_1^2 = \lambda_1$ and $\sigma_2^2 = \lambda_2$, where λ_1 and λ_2 are the eigenvalues corresponding to the first 2 KL model basis vectors. Let the model parameter vector $\mathbf{X}_i = (x_1, x_2, ..., x_{12}, 0_{1x10})^T$ denote the i^{th} initial value. Each nonzero element of $\mathbf{X}_i$ is drawn independently from different uniform populations:

$$\begin{aligned} x_1 &\in [-2\sigma_1, 2\sigma_1] \\ x_2 &\in [-2\sigma_2, 2\sigma_2] \\ x_3 &\in [-120, 120] \\ &\vdots \\ x_{12} &\in [-120, 120]. \end{aligned}$$

The bounds on all of the parameters were subjectively chosen after some experimentation as a balance between maximizing algorithm reliability and minimizing domain size. Smaller domains require fewer initial value points. However, algorithm reliability depends on the domain including points from the capture region of the local minimum of interest. It was found that specifying nonzero values only for the first 12 initial value parameters was a reasonable balance between these two conflicting requirements. Further, the generous interval bounds on the 12 nonzero parameters ensures that the domain allows significant freedom in wind features.

The 50 initial values $\mathbf{X}_i$ are then locally optimized using the IMSL implementation of the quasi-Newton method with analytic gradient. Although the constraints within the objective function introduce discontinuities, the quasi-Newton method optimizes adequately, even though it is traditionally used only with continuous objective functions. If difficulties due to the discontinuities arise in future tests and applications of this method, local optimization routines for discontinuous objective functions can be employed, but with a possible increase in computation time.

After the 50 initial values are optimized, the negative of the resulting solutions are also optimized, yielding a total of 100 optimized solutions. Optimizing the negative of the first 50 optima increases the chance that both the upwind and

downwind versions of the solutions are present in the solution set, which improves the success rate of the algorithm. Finally, the 100 optima are examined to remove redundant solutions. If any two solutions have a vrms difference within 0.75 m/s, only the solution with the lower field-wise objective function value is retained. The following is a summary of the algorithm:

1. Randomly select 50 initial values
2. Locally optimize initial values
3. Locally optimize negatives of the resulting solutions
4. Remove redundant solutions with a vrms threshold of 0.75 m/s, retaining the solutions with the lowest objective function values.

The performance of the field-wise multistart algorithm is considered in the next section. However, it is necessary to note that even though the algorithm is reliable, it can fail and that the failures are more likely to occur in regions with interesting features, such as fronts and cyclones, since those regions have smaller capture regions. In order to minimize the failure rate of the algorithm, its solution set is augmented using a local optimization algorithm initialized with median filtered wind.

First, the median filter is run using the first and second ambiguity fields as initial values. The median filter outputs are used as initial values in model-based retrieval. Accordingly, least-squares model fits are performed, and the model parameters are locally optimized using the same gradient-descent methods employed in the multistart optimizations. The resulting solutions are added to those from the multistart algorithm.

3.4 Test Results

The augmented multistart algorithm is tested on 10 revolutions from week 23 of the NSCAT mission. The NSCAT swaths are divided into 12x12 regions with 50% overlap, and the algorithm is run on each region. In addition, the JPL product

for the overlapping regions is fit to the wind field model (in a least-squares sense) and optimized, using the same optimization routine as in the multistart algorithm. The output of the multistart algorithm is compared with the optimized JPL product to determine the success rate of the algorithm.

Note that the JPL product is only about 96% accurate due to ambiguity removal errors [22]. In lieu of a better choice, only the reliable portions of it are used for comparison as "ground truth". The "poor regions", those regions with more than 20% of the cells having possible ambiguity removal errors, are identified using the quality assessment algorithm developed in [22] and are not used in the comparison.

Figure 3.25 is a histogram of the vrms difference between the optimized JPL product and the closest field-wise estimate for the over 700 reliable JPL regions in the test set. In a high percentage of the regions, the augmented multistart algorithm is very successful at locating the JPL solution. The actual percentage of times that the algorithm locates the optimized JPL solution is 97% (0.75 m/s vrms threshold). However, the success rate of finding a solution subjectively close to the JPL solution is higher. Even if the algorithm fails to locate the true solution exactly, a solution close to true one is admissible. In the field-wise ambiguity removal stage the overlapping solutions are joined by an averaging process which can repair incorrect solutions if the other overlapping solutions are closer to correct.

A variation of the usual vrms difference can be used to quantify how much deviation from the desired wind is admissible. Let the vrms difference between two model-based wind fields $\mathbf{W_1}$ and $\mathbf{W_2}$ be calculated as follows:

1. Fill the *valid* vector: $valid(i) = \begin{cases} 0 & : \text{ no data in resolution element } i \\ 1 & : \text{ otherwise} \end{cases}$

2. Vrms Difference $= \sqrt{\frac{(\mathbf{W}-\mathbf{W})^T D_v^2 (\mathbf{W}-\mathbf{W})}{\sum_{i=1}^{144} valid(i)}}$. D_v is the diagonal matrix with the elements of *valid* on the main diagonal.

Subjective inspection of solutions indicates that vrms differences between 2 and 3 m/s are sufficiently close. Since approximately 99% of the regions have

solutions this close to the optimized JPL product, the subjective skill of the multistart algorithm is probably closer to 99%.

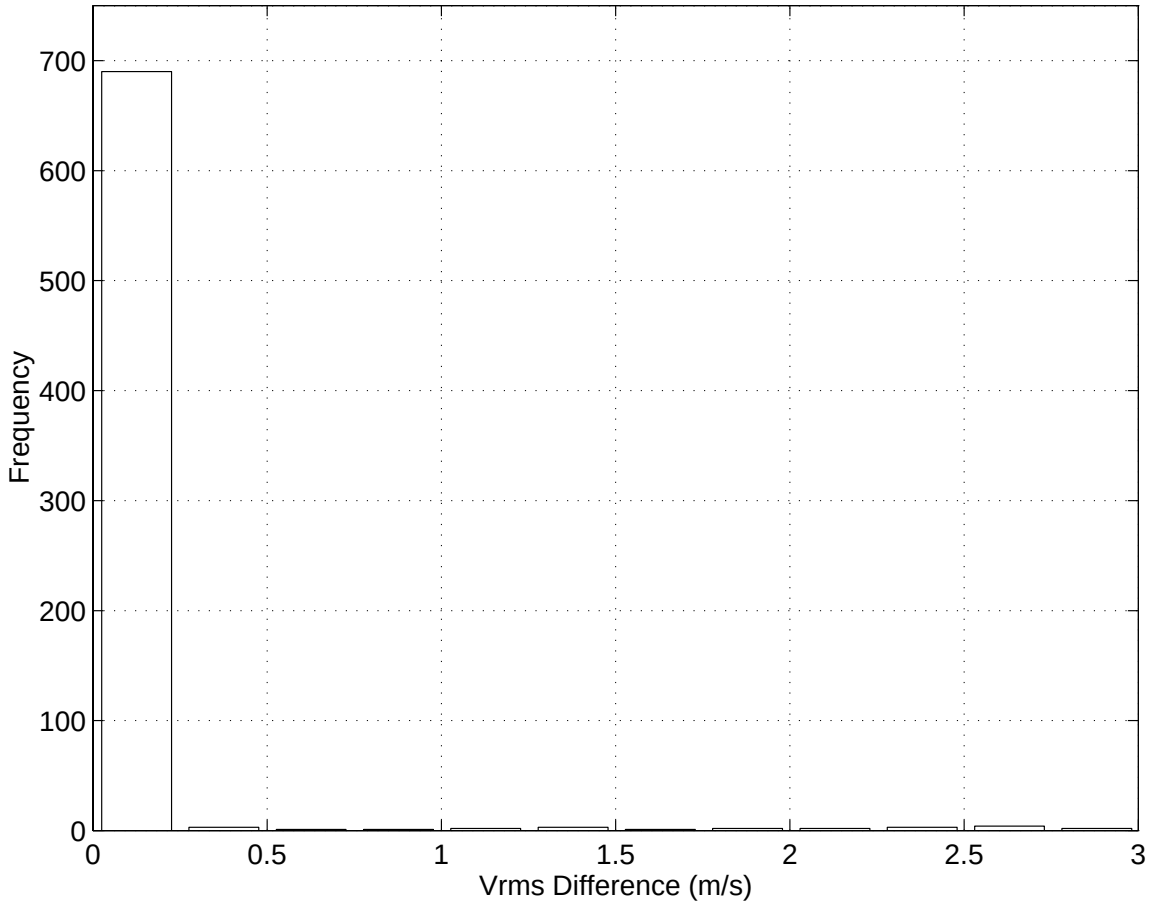

Figure 3.25: Histogram of skill of field-wise multistart algorithm measured in vrms difference from optimized JPL product. Approximately 99% of the regions have solutions subjectively close to the optimized JPL product.

3.5 Summary

In this chapter, the characteristics of the field-wise objective function are explored. It is found that the objective function exhibits numerous local minima. However, the field-wise objective function tends to rank more reasonable solutions higher than less reasonable ones, in a geophysical sense. It is also determined that the field-wise objective function displays solution symmetry, which can assist in the estimation procedure.

Since it is observed that the global minimum of the objective function is not always the desired solution, it is necessary to find a field-wise estimation algorithm that could locate multiple near-global local minima. A further constraint on the field-wise algorithm is that it be reasonably fast, since it would need to be executed repeatedly, once for each region. Several global optimization algorithms are considered, and the multistart method is chosen as the most favorable candidate.

A variant of the multistart method is developed to exploit the solution symmetry of the field-wise objective function, and the algorithm is tested on NSCAT data. The algorithm displays about 99% success at locating local minima acceptably close to the desired solution.

Chapter 4

A Field-Wise Ambiguity Removal Algorithm for the NASA Scatterometer

4.1 Introduction

The field-wise estimation technique developed for NSCAT in the previous chapter generates multiple solutions (also called ambiguities or estimates) for each region of the scatterometer swath. A field-wise ambiguity removal (FWAR) algorithm is necessary to determine a unique wind swath estimate. This report discusses the structure of a FWAR algorithm that does not rely on numerical analysis winds. First, preliminary definitions are offered for several terms used in describing the FWAR algorithm. Then, an overview of the algorithm is provided and each of the major stages in the algorithm is discussed.

4.2 Preliminary Definitions

In order to simplify the description of the FWAR algorithm, it is necessary to introduce several useful terms. Recall that the output of a field-wise estimation algorithm is a set of possible solutions, which are also referred to as ambiguities or estimates.

The FWAR algorithm introduced in the following section deals with multiple regions overlapping by 50%. A set of solutions, one solution for each overlapping region, is called an extended solution. The model-based wind on a multi-region scale is also referred to as a wind swath solution or estimate.

An additional objective function that is useful in ambiguity removal is the F_I objective function, which uses the wind field model F_I instead of F_0, where

$$F_I = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}.$$

The F_I objective function is then

$$J_I(\mathbf{X_I}) = -\sum_{l=1}^{MN}\sum_{k=1}^{K(l)}\left[\frac{(Z_{k,l} - M_{k,l}(\mathbf{w_l}))^2}{\vartheta_{k,l}^2}\right],$$

and the model parameter sets that locally optimize $J_I(\mathbf{X}_I)$ represent wind field solutions via the modeling relationship $\hat{\mathbf{W}}_I = F_I\hat{\mathbf{X}}_I$. Observe that $\hat{\mathbf{X}}_I$ is a vector of 288 elements. Those elements are identical to the (u, v) components of a scanned vector wind field. Inspection of $J_I(\mathbf{X}_I)$ reveals that each of the wind vector cells can be optimized independently, just as in point-wise estimation. Thus, the F_I objective function can be calculated by selecting the closest point-wise ambiguity to each wind vector represented by $J_I(\mathbf{X}_I)$ and totaling their point-wise objective function values.

Optimization of the F_I objective function can be expedited by choosing the closest point-wise wind to each of the vectors of the initial value wind field. For example, in order to perform an F_I optimization of $\hat{\mathbf{X}}$, a field-wise estimate, $\hat{\mathbf{X}}$ must be transformed into its corresponding model-based wind field: $\hat{\mathbf{W}} = F_0\hat{\mathbf{X}}$. Then, the initial value of the F_I optimization is $\mathbf{X}_I = \hat{\mathbf{W}}$. The optimized parameter vector is $\hat{\mathbf{X}}_I$, which contains the closest point-wise winds to the vectors represented by the initial value $\mathbf{X}_I = \hat{\mathbf{W}}$.

4.3 Field-Wise Ambiguity Removal Algorithm

The algorithm (refer to Fig. 4.1) begins with an initial extended solution, assembled from the most likely field-wise estimates. The algorithm detects discontinuities in the initial extended solution, and then it repairs them by trying combinations of lower-ranked solutions. Finally, the most likely extended solution is refined through

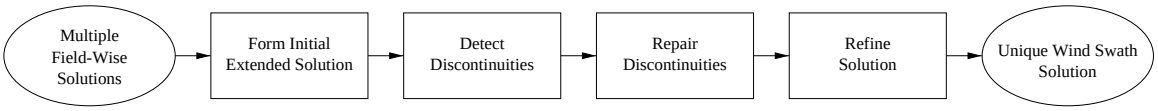

Figure 4.1: The FWAR algorithm. The algorithm begins with an initial extended solution. It detects discontinuities in the initial extended solution, and then it attempts to repair them. The repaired extended solution is further refined, and a unique wind swath solution is output.

model-based estimation, and a unique wind swath estimate is output. Each of the major divisions of the algorithm is discussed in detail in a separate subsection.

4.3.1 Form Initial Extended Solution

The FWAR algorithm acts on sets of overlapping solutions to piece together the solutions into extended solutions. The first step in determining the extended solution for a given set of regions is to form an initial extended solution. In the FWAR algorithm, the initial extended solution is assembled from the most likely field-wise estimates. First, the field-wise estimates are ranked according to their field-wise objective function values, and only the 20 highest ranked solutions are retained. The field-wise objective function tends to rank more reasonable solutions higher than less reasonable ones, so retaining only the top solutions eliminates unlikely solutions. The 20 retained solutions are ranked according to their F_I objective function values. The highest ranked ambiguities from each region according to this scheme are assembled into the initial extended solution.

4.3.2 Detect Discontinuities

Although the initial extended solution is generally realistic, there may be some discontinuities due to improper choice of initial solutions from the field-wise ambiguities. However, these discontinuities can be detected and repaired. In order to detect the discontinuities, the vrms difference between overlapping portions of

solutions are computed. If a pair of solutions exhibits an overlap difference in excess of a fixed threshold value, that pair is marked as a discontinuity.

A threshold value of 4.5 m/s is selected by subjectively locating discontinuities in 10 revolutions of NSCAT data from week 23 and determining which threshold level locates a high percentage of the discontinuities. Initial extended solutions are constructed, and the discontinuities in each extended solution are identified. A threshold of 4.5 m/s locates approximately 86% of the discontinuities. The undetected discontinuities are successfully identified only with a threshold of about 2 m/s, which generates an excessive number of discontinuities. Even though false alarms are not as critical as missed detections, the scheme performs sufficiently well. It produces good results in the next stage of the algorithm, which involves identifying the extent of the initial value errors around the discontinuities. It is possible that many of the undetected discontinuities are sufficiently close to the detected discontinuities that they are included in the clusters of possibly erroneous regions. Further support of the 4.5 m/s threshold value is provided below.

4.3.3 Repair Discontinuities

All of the discontinuous pairs are grouped to define continuous sets of possibly erroneous initial solutions. The groups are called "clusters". The clusters are centered on the discontinuous pairs and include two pairs of solutions to either side of the discontinuous pair. This conservative approach tends to exclude initial solutions of high confidence from the clusters, and these solutions can be used to assist in repairing discontinuities within the clusters, as they help establish the general wind flow. However, it has been observed that the initial solutions adjacent to the clusters can still be incorrect, so an additional test is necessary to determine which solutions not included in the clusters help establish the correct wind flow.

Identification of Anchors

Initial solutions that help establish the correct wind flow are called "anchors". Model-based thresholding [13] is employed to identify the anchors.

Model-based thresholding takes the form of a hypothesis test, but not one in the traditional Neyman-Pearson sense. The decision rule employed in model-based thresholding eliminates improbable solutions while keeping the probability of eliminating the correct solution below a fixed threshold α_0.

In order to construct the decision rule, the log-likelihood-ratio statistics between the highest ranked field-wise ambiguity and each of the lower ranked solutions is formed:

$$\begin{aligned} \lambda_2(z) &= \ln\left(\frac{p_z(z|\hat{w}^2)}{p_z(z|\hat{w}^1)}\right) && (4.1)\\ \lambda_3(z) &= \ln\left(\frac{p_z(z|\hat{w}^3)}{p_z(z|\hat{w}^1)}\right) && (4.2)\\ &\vdots \\ \lambda_n(z) &= \ln\left(\frac{p_z(z|\hat{w}^n)}{p_z(z|\hat{w}^1)}\right), && (4.3) \end{aligned}$$

where $\hat{w}^n$ represents the n^{th} ranked field-wise estimate for the measurement vector z, according to field-wise objective function value. Since the measurement vector z is a random variable, $\lambda_2(z), \lambda_3(z), ..., \lambda_n(z)$ are also random variables. Denote a given realization of z by z_0. Then, using the techniques developed in [13], calculate the following probabilities:

$$\begin{aligned} \alpha_2 &= Prob[\lambda_2(z) \leq \lambda_2(z_0)|w = \hat{w}^2] && (4.4)\\ \alpha_3 &= Prob[\lambda_3(z) \leq \lambda_3(z_0)|w = \hat{w}^3] && (4.5)\\ &\vdots \\ \alpha_n &= Prob[\lambda_n(z) \leq \lambda_n(z_0)|w = \hat{w}^n], && (4.6) \end{aligned}$$

with the estimates $\hat{w}^2, \hat{w}^3, ..., \hat{w}^n$ obtained from the deterministic measurement vector z_0. As stated in [13], these probabilities are approximations since the true wind should be used on the given side of the argument. With these definitions in place, the decision rule is stated as

$$\text{Retain the solutions } \hat{w}^1, \hat{w}^2, ..., \hat{w}^i \text{ if } \alpha_i \leq \alpha_0.$$

The first ranked field-wise estimate (by the same criterion as used to define the initial extended solution) is compared to the solution that is closest to its negative.

This test is roughly analogous to performing point-wise thresholding on the first and second point-wise ambiguities. The test, in effect, determines to what degree the highest ranked field-wise ambiguity can be trusted above its negative. If there is a high degree of separation in F_I objective function value, then there is a good chance that the first field-wise ambiguity correctly defines the true wind. In terms of the test defined in the above paragraph, the first field-wise ambiguity is assumed to define the general wind flow correctly if the test determines it is reasonable to retain the solution $\hat{w}^i$, where $\hat{w}^i$ is the closest solution to the negative of the first field-wise ambiguity. In this test, the F_I objective function values are used for ranking and to calculate $\lambda_i(z_0)$.

Solutions not initially included in clusters are anchors if the size of the above test is locally low. That is, the size of the hypothesis test involved in the field-wise thresholding is calculated for each of the solutions in the initial extended solution, and local minima of the graph of the size versus location are identified. The local minima mark the anchors.

The anchor location algorithm is run on 10 revolutions of NSCAT data. First, initial extended solutions are constructed and clusters are formed. Anchors are chosen from regions not included in the clusters. The number of anchors used in later stages of the field-wise ambiguity removal algorithm is 82, and out of those 82, only two are subjectively identified as poor anchors.

In order to determine which anchors (if any) are associated with each cluster, the endpoints of the clusters are expanded until they arrive at either the extended solution boundaries or an anchor. The anchors adjacent to the extended clusters are associated with the clusters and are used in later processing.

Retention Threshold

Each of the clusters is processed in turn. Within the cluster, the compatibility of pairs of solutions is tested by examining their vrms overlap differences. Those exceeding an overlap threshold of 7.5 m/s, termed the "retention threshold", are considered discontinuous and are discarded. Initial support for the retention threshold

is in Fig. 4.2. This figure is generated by parsing the JPL product for 10 revolutions in week 23 into 7468 pairs of regions that overlap by 50%. For each overlapping pair, model fits are performed using a pseudo-inverse: $\hat{\mathbf{X}} = F_0^{\dagger}\mathbf{W}$, where $\mathbf{W}$ denotes the JPL vector wind. The resulting model parameter vectors $\hat{\mathbf{X}}$ are locally optimized using the IMSL quasi-Newton algorithm with analytic gradient. Finally the vrms overlap difference is calculated according to the following procedure:

1. Fill the *valid* vector: $valid(i) = \begin{cases} 0 & : \text{ no data in resolution element } i \\ 1 & : \text{ otherwise} \end{cases}$

2. VRMS Overlap Difference $= \sqrt{\frac{(\hat{w}_l - \hat{w}_r)^T D_v^2 (\hat{w}_l - \hat{w}_r)}{\sum_{i=1}^{72} valid(i)}}$, where $\hat{w}_l$ and $\hat{w}_r$ are the column vectors containing the wind in the left and right overlapping regions, respectively. D_v is the diagonal matrix with the elements of *valid* on the main diagonal

The vrms overlap differences for overlap areas with more than 57 resolution elements of data are used to generate the histogram in Fig. 4.2. Note that 7.5 m/s is clearly in the tail of the histogram. For a threshold of 7.5 m/s, only roughly 1.7% of the JPL overlapping pairs are declared discontinuous. The JPL pairs declared as discontinuous probably suffer from ambiguity removal errors or the model fit to the JPL solutions is poor.

The threshold for this portion of the algorithm is selected as low as possible in order to reduce the number of pairs that survive. This significantly reduces the computational time of the algorithm. A higher threshold, in addition to increasing the processing time, would also allow the undesirable introduction of discontinuous pairs of solutions into the solution set. Of course, a lower threshold does increase the chance of pruning desirable solution pairs. However, the field-wise ambiguity removal algorithm is quite robust with respect to the choice of threshold.

In order to quantify the effects of varying the retention threshold, 5 revolutions from week 23 are processed by the FWAR algorithm with different thresholds: 5, 7.5, and 9 m/s. The discontinuity detection threshold is kept constant at 4.5 m/s. The unrefined product of the FWAR algorithm is compared with the JPL product,

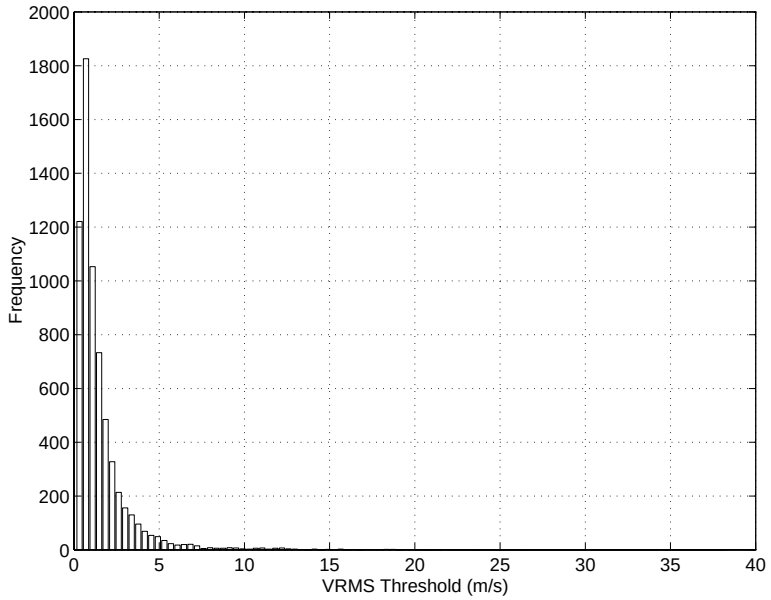

Figure 4.2: Histogram used to select the vrms threshold that defines what vrms overlap difference is allowed between two overlapping solutions. This figure is generated by parsing the JPL product for 10 NSCAT revolutions into regions that overlap by 50%. The vrms overlap differences for overlap areas with more than 57 resolution elements of data are used to generate the histogram.

and wind vectors having a vector difference within 2 m/s are declared as matches. Only wind vector cells where the JPL wind is above 4 m/s are considered. Table 4.1 displays the results of the tests. It is apparent that the exact value of the threshold is not critical to the overall performance of the FWAR algorithm. However, examination of the individual swaths indicates that the local skill can vary with the choice of threshold. When desirable pairs of solutions are pruned because the threshold is too restrictive, the resulting errors can be significant. For instance, it is observed that with a threshold of 7 m/s the FWAR algorithm reverses the direction of the wind for a section of revolution 2457 descending. The correct solution is pruned because the threshold is too restrictive. Such aberrations in the algorithm affect the local skill, but occur infrequently enough not to significantly affect the overall agreement with JPL. Varying the discontinuity threshold from 4.5 to 8 m/s produces similar results.

Retention Threshold (m/s)	FWAR Algorithm Skill
5	91%
7	89%
7.5	91%
9	91%

Table 4.1: Effect of threshold variation on FWAR algorithm skill. This table quantifies the overall effects of varying the FWAR algorithm retention threshold. The discontinuity detection threshold is kept constant at 4.5 m/s. The unrefined product of the FWAR algorithm is compared with the JPL product, and wind vectors having a vector difference within 2 m/s are declared as matches. Even though the local skill can vary, the overall skill seems robust with respect to the threshold.

Extended Solution Assembly

Multiple extended solutions spanning each cluster are formed by joining solutions from adjacent pairs of regions, moving from left to right (Fig. 4.3), although the assembly direction has no effect on the final product. A multi-stage filtering process is used to select which solutions, out of the many that field-wise estimation provides, participate in the pair-wise assembly.

Only the top 20 solutions (in order of field-wise objective function value) from each region are used to form pairs. This filtering step tends to eliminate unrealistic solutions from the solution set. The remaining solutions are sorted according to their F_I objective function values, and only the top 6 are considered in the remaining portion of the algorithm. In order to increase the chances that the set of 6 solutions contains the correct one, the set is searched for the closest to the negative of the most likely solution. The set is also searched for the solution that is the closest to its closest point-wise field. Further, it is searched for the closest to the negative of that solution. If any of these solutions are missing from the set of 6, they are added to the set in place of the least likely solutions.

All of the possible pairs of the remaining solutions are enumerated. Those pairs that meet the 7.5 m/s vrms continuity criterion described above are retained,

and the rightmost solution in each pair is compared with the solutions in the overlapping region immediately to its right (Fig. 4.3). The surviving pairs from this comparison are appended to the growing solution sets, and this process is repeated across the cluster. Clusters with anchors have the further constraint that all of the extended solutions in the cluster must be continuous with the anchors (7.5 m/s vrms overlap difference) or they are discarded. After all of the admissible extended solutions have been enumerated, the model-based wind corresponding to each extended solution is constructed via an overlap-and-average process. Each of the model parameter sets comprising the extended solution generates its corresponding model-based wind field. The resulting wind fields overlap by 50%. The model-based wind field represented by the extended solution is calculated by adding the overlapping wind fields and averaging them with trapezoidal along-track weighting, with the highest weighting in the center of the wind fields. Trapezoidal weighting is used because it is known that field-wise estimates are less reliable farther away from their centers [13].

The overlap-and-average method can be expressed in mathematical terms. The following derivation is only for the 12x12 50% overlap case with trapezoidal weighting, although the derivation could be made more general. First form the 12x12 weighting matrices

$$W_1 = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & .75 & .75 & .5 & .5 & .25 & .25 \\ 1 & 1 & 1 & 1 & 1 & 1 & .75 & .75 & .5 & .5 & .25 & .25 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ 1 & 1 & 1 & 1 & 1 & 1 & .75 & .75 & .5 & .5 & .25 & .25 \end{pmatrix},$$

$$W = \begin{pmatrix} .25 & .25 & .5 & .5 & .75 & .75 & .75 & .75 & .5 & .5 & .25 & .25 \\ .25 & .25 & .5 & .5 & .75 & .75 & .75 & .75 & .5 & .5 & .25 & .25 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ .25 & .25 & .5 & .5 & .75 & .75 & .75 & .75 & .5 & .5 & .25 & .25 \end{pmatrix},$$

and

$$W_N = \begin{pmatrix} .25 & .25 & .5 & .5 & .75 & .75 & 1 & 1 & 1 & 1 & 1 & 1 \\ .25 & .25 & .5 & .5 & .75 & .75 & 1 & 1 & 1 & 1 & 1 & 1 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ .25 & .25 & .5 & .5 & .75 & .75 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix}.$$

Let $D_1 = Diag(Vec(W_1^T))$, $D = Diag(Vec(W^T))$, and $D_N = Diag(Vec(W_N^T))$, where the $Vec(\cdot)$ represents the column scan operation and $Diag(\cdot)$ forms a matrix from a vector with the vector on the main diagonal and zeros elsewhere. Then

$$\hat{\mathbf{W}} = G_1 F_0 \hat{\mathbf{X}}_1 + G_1 F_0 \hat{\mathbf{X}}_2 + \cdots + G_1 F_0 \hat{\mathbf{X}}_N \tag{4.7}$$

where

$$G_i = I_{2\times 2} \otimes \left\{ \left[I_{12\times 12} \otimes \begin{pmatrix} 0_{6(i-1)\times 12} \\ I_{12\times 12} \\ 0_{6(N-i)\times 12} \end{pmatrix} \right] D_i \right\},$$

with

$$D_i = \begin{cases} D_1 & , \quad i = 1 \\ D & , \quad 1 < i < N \\ D_N & , \quad i = N \end{cases}.$$

Equation 4.7 can be further simplified to yield

$$\hat{\mathbf{W}} \;=\; (G_1 F_0 \vdots G_2 F_0 \vdots \cdots \vdots G_N F_0) \begin{pmatrix} \hat{\mathbf{X}}_1 \\ \cdots \\ \hat{\mathbf{X}}_2 \\ \cdots \\ \vdots \\ \cdots \\ \hat{\mathbf{X}}_N \end{pmatrix} \tag{4.8}$$

$$= (G_1 \vdots G_2 \vdots \cdots \vdots G_N) F_0 \begin{pmatrix} \hat{\mathbf{X}}_1 \\ \cdots \\ \hat{\mathbf{X}}_2 \\ \cdots \\ \vdots \\ \cdots \\ \hat{\mathbf{X}}_N \end{pmatrix} \tag{4.9}$$

$$= F_\Phi \hat{\mathbf{X}}_\Phi, \tag{4.10}$$

with

$$F_\Phi = (G_1 \vdots G_2 \vdots \cdots \vdots G_N) F_0$$

and

$$\hat{\mathbf{X}}_\Phi = \begin{pmatrix} \hat{\mathbf{X}}_1 \\ \cdots \\ \hat{\mathbf{X}}_2 \\ \cdots \\ \vdots \\ \cdots \\ \hat{\mathbf{X}}_N \end{pmatrix}.$$

These manipulations reveal the linear nature of the overlap-and-average output. It is simply a linear transformation of the model parameter set $\hat{\mathbf{X}}_\Phi$ using the model F_Φ.

The extended solution that meets the vrms overlap difference constraints and has the lowest F_I objective function value is selected as the optimum extended solution for that cluster. It is then grafted into the initial extended solution in place of the original cluster.

The method of selecting solutions that repair the discontinuities in the initial extended solution is an optimization process. With

$$\hat{\mathbf{X}}_\Phi = \begin{pmatrix} \hat{\mathbf{X}}_1^{k_1} \\ \cdots \\ \hat{\mathbf{X}}_2^{k_2} \\ \cdots \\ \vdots \\ \cdots \\ \hat{\mathbf{X}}_N^{k_N} \end{pmatrix},$$

where $\hat{\mathbf{X}}_1^{k_1}$ is the k_1^{th} solution for region 1, $\hat{x}_2^{k_2}$ is the k_2^{th} solution for region 2, etc., form the extended solution objective function

$$J_S(\hat{\mathbf{x}}_\Phi) = J_I(F_\Phi \hat{\mathbf{X}}_\Phi),$$

where J_I denotes the F_I objective function. Since solution indices $k_j, j = 1, 2, ..., N$ in $\hat{\mathbf{X}}_\Phi$ take on only natural number values, the optimization problem is discrete. It is stated as

$$(k_1, k_2, ..., k_N) = arg \left\{ \min_{s.t.\ vrms(\hat{\mathbf{W}}_j^{k_j}, \hat{\mathbf{W}}_{j+1}^{k_{j+1}}) \leq 7.5} \left[J_I(F_\Phi \hat{\mathbf{X}}_\Phi) \right] \right\},$$

where the function $vrms(\hat{\mathbf{W}}_j^{k_j}, \hat{\mathbf{W}}_{j+1}^{k_{j+1}})$ calculates the vrms overlap difference for the model-based wind fields $\hat{\mathbf{W}}_j^{k_j}$ and $\hat{\mathbf{W}}_{j+1}^{k_{j+1}}$. Note that $\hat{\mathbf{W}}_j^{k_j} = F_0 \hat{\mathbf{X}}_j^{k_j}$, and $\hat{\mathbf{W}}_{j+1}^{k_{j+1}} = F_0 \hat{\mathbf{X}}_{j+1}^{k_{j+1}}$. The extended solution that satisfies the optimization criteria is represented by the index numbers of the individual model-based solutions $(k_1, k_2, ..., k_N)$. The wind swath corresponding to the extended solution is $F_\Phi \hat{\mathbf{X}}_\Phi$.

This process is repeated for each cluster. Occasionally, the vrms constraints listed above make it impossible to find a continuous extended solution for a given cluster. If this is the case, the cluster is split at the point where no valid solution pairs are available, and processing resumes at the first portion of the new cluster. All of the regions included in the original cluster are marked with a warning flag. Warnings also occur if 100000 extended solutions have been accumulated. Again,

each region of the entire cluster is flagged with warnings, and the cluster is split at the point where the 100000 extended solutions accumulated. Processing resumes at the first portion of the new cluster. Another special case is that of clusters with single regions. Such clusters are unable to utilize the continuity information from overlapping regions and are also marked as potential errors. In such cases, the initial extended solution ambiguity is used.

4.3.4 Refine Solution

Finally, the closest point-wise ambiguity to each vector in the resulting swath, after overlap-and-average, is selected and the closest-ambiguity swath is used as an initial value for model-based estimation. The swath is segmented into 12x12 regions overlapping by 50%, and the F_0 model is fit to each of the point-wise winds for each region. The model fit is locally optimized using the IMSL finite-difference quasi-Newton algorithm, and the solutions for overlapping regions are averaged using trapezoidal weighting. The final output of the algorithm is the model-based wind or the closest point-wise field to the model-based wind.

4.3.5 Example of Field-Wise Ambiguity Removal Algorithm

As an example of the FWAR algorithm, consider the portion of NSCAT revolution 2454 ascending, from along-track 43 to 139. For reference, Fig. 4.4 shows the point-wise all-ambiguity plot for this section, and Fig. 4.5 plots the JPL product.

Form Initial Extended Solution

The initial extended solution is shown in Fig. 4.6, where the model-based wind fields comprising the extended solution are plotted with 50% overlap.

Detect Discontinuities

Observe the minor overlap differences around along-track 50 and between along-track 70 and 100. These minor differences are admissible. However, the overlap differences in the along-track ranges from 120 to 130 and 130 to 140 are clearly errors

in the initial extended solution. Table 4.2 shows vrms overlap differences for the solutions in the extended solution in Fig. 4.6. Note that the vrms threshold of 4.5 m/s detects the major and minor incongruities between the overlapping solutions beginning at along-track 115 and 121; beginning at 127 and 133; and at 67 and 73. In Table 4.2, these overlapping solutions are marked with "1" in the "Discontinuity?" column. Since it is not always obvious which regions are at fault, the clustering algorithm takes a very conservative approach and flags the two overlap regions to either side of each discontinuity. For example, observe that in Table 4.2 the clustering algorithm marks all overlapping pairs of solutions from 55 to 85 and 103 to 139 as being possible errors.

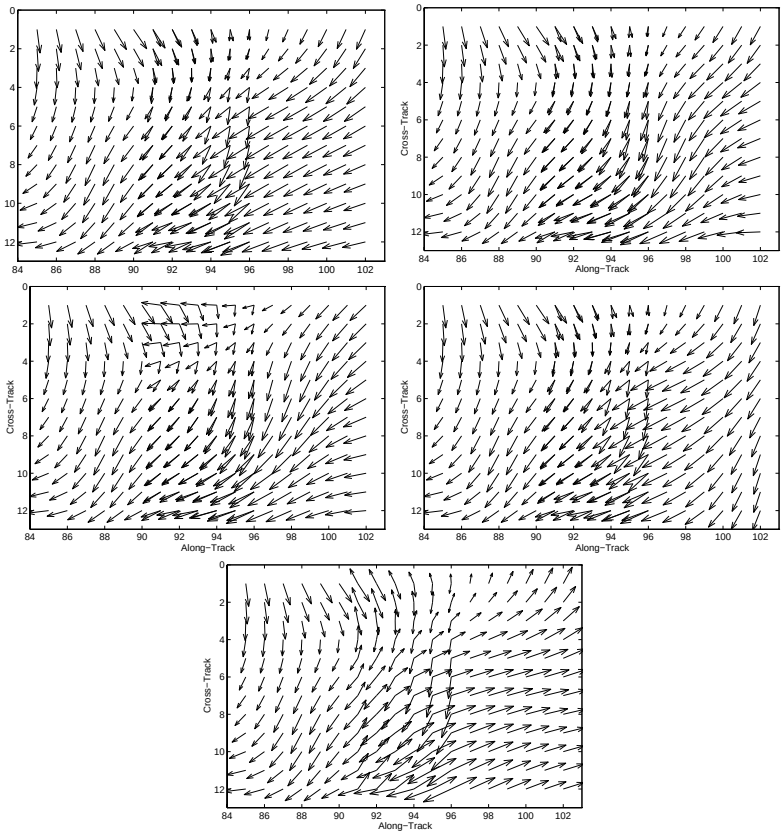

Figure 4.3: Extended solution assembly. Extended solutions are formed by joining pairs of overlapping solutions from adjacent regions, moving from left to right. This figure plots 5 overlapping solution pairs from 2 adjacent regions of NSCAT revolution 2454 ascending. Only the first 4 overlapping solution pairs are retained for further consideration because they are reasonably continuous and the last is not. The rightmost solutions of the retained pairs are paired with all solutions in the next region to the right, and discontinuous solution pairs are discarded. The process of pairing solutions and discarding discontinuous pairs is repeated for the length of the extended solution.

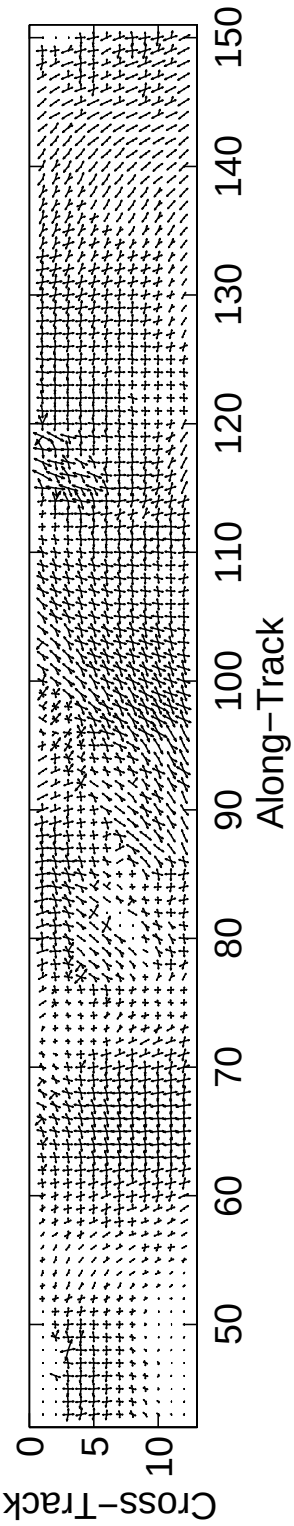

Figure 4.4: Plot of all point-wise ambiguities for a portion of NSCAT revolution 2454 ascending

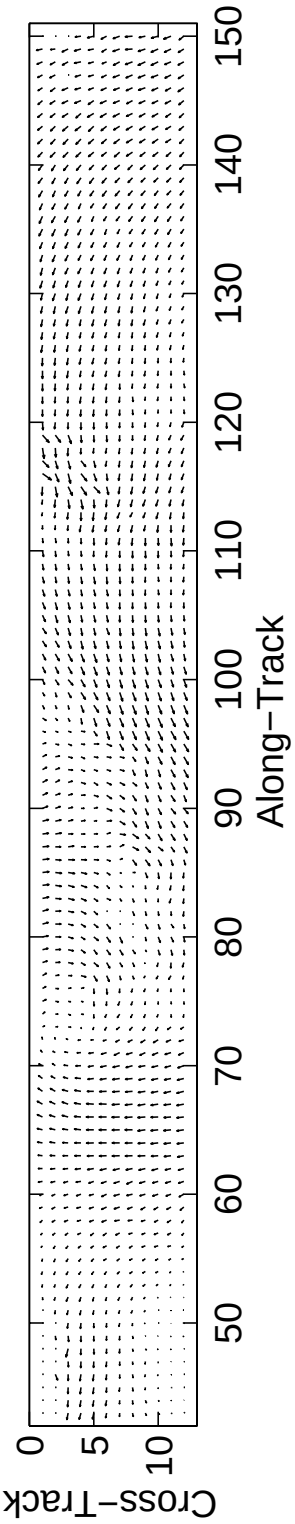

Figure 4.5: JPL product for a portion of NSCAT revolution 2454 ascending

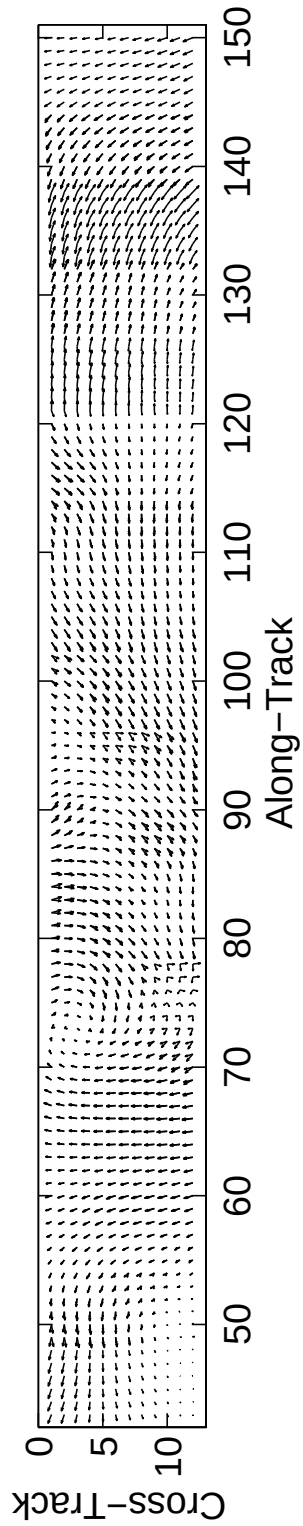

Figure 4.6: Initial extended solution for a portion of NSCAT revolution 2454 ascending

Region 1 Along-Track	Region 2 Along-Track	VRMS Overlap Difference (m/s)	Discontinuity? (0–No, 1–Yes)	After Clustering Algorithm
43	49	1.7	0	0
49	55	0.5	0	0
55	61	0.4	0	1
61	67	1.9	0	1
67	73	4.9	1	1
73	79	2.3	0	1
79	85	2.9	0	1
85	91	4.1	0	0
91	97	2.0	0	0
97	103	1.0	0	0
103	109	1.5	0	1
109	115	1.5	0	1
115	121	19.4	1	1
121	127	0.7	0	1
127	133	20.0	1	1
133	139	1.0	0	1

Table 4.2: Vrms overlap differences and resulting clusters. This figure details the processing of a sample swath for NSCAT revolution 2454A. The vrms overlap differences for the solutions in the extended solution in Fig. 4.6 are listed in the "VRMS Overlap Difference" column. Note that the vrms threshold of 4.5 m/s detects the major and minor incongruities between the overlapping solutions beginning at along-track 115 and 121; beginning at 127 and 133; and at 67 and 73. These overlapping solutions are marked with "1" in the "Discontinuity?" column. The next stage of the processing, the clustering algorithm, marks all overlapping pairs of solutions from 55 to 85 and from 103 to 139 as being possible errors.

Repair Discontinuities

The next step in the algorithm identifies anchors. First, the clusters from Table 4.2 are extended to form the clusters in Table 4.3 in the "Extended Clusters" column. In order to extend the clusters, the regions that are a part of the original clusters are marked with a "1" in the "Extended Clusters" column of Table 4.3. The extended clusters are important in determining anchors, as anchors cannot be selected

within the extended clusters. This serves to prevent selection of anchors from initial solutions that could be erroneous.

Figure 4.7 shows a plot of the log bound of the hypothesis test size ($log(\alpha_i)$) as a function of region. Again, the hypothesis test for each region is between the solution in the initial extended solution and the solution that is closest to its negative. The local minima in Fig. 4.7 outside of the extended clusters represent regions for which the confidence of the initial extended solution is high, at least relative to the surrounding regions. Since the region beginning at along-track 43 is the only local minima outside of the extended clusters, its initial solution is the anchor.

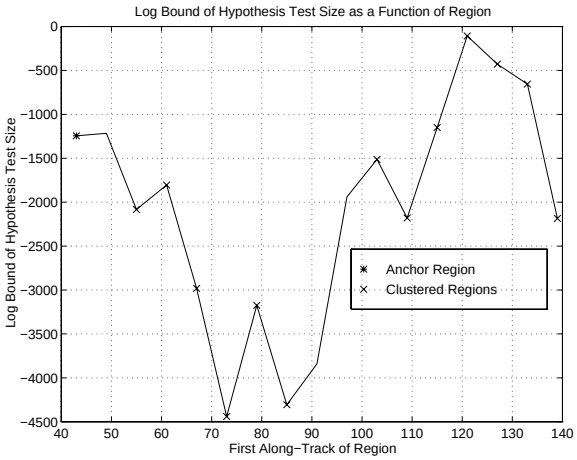

Figure 4.7: Log bound of the hypothesis test size ($log(\alpha_i)$) as a function of region. The local minima of this plot outside of the extended cluster are possible anchors. Since the region beginning at along-track 43 is the only local minimum outside of the extended cluster, its initial solution is the anchor.

Once the anchors are identified, the clusters are expanded until both edges of the extended clusters touch either an anchor or the beginning or ending of the initial extended solution. Table 4.3 shows the result of the expansion operation on the extended cluster. The cluster now includes the regions from along-track 43 to

139. The initial extended solution for the region beginning at along-track 43 is the anchor for the cluster.

The solutions for each region in the cluster, except for the anchor, are filtered after they are ranked in order of their field-wise objective function values. Only the top 20 are retained (after removing solutions with a vrms difference within 0.75 m/s). The remaining solutions are sorted according to their F_I objective function values, and only the top 6 are considered in the remaining portion of the algorithm.

Multiple extended solutions spanning each cluster are formed by assembling solutions from adjacent pairs of regions, moving from left to right. Those pairs that meet the 7.5 m/s vrms continuity criterion described above are retained, and the rightmost solution in each pair is compared with the solutions in the overlapping region immediately to its right. The surviving pairs from this comparison are appended to the growing solution sets, and this process is repeated across the cluster.

After the process of growing possible extended solutions is complete, the model-based wind for each extended solution is constructed, and overlap-and-average is used to produce wind swaths corresponding to the extended solutions. The wind swath with the lowest F_I objective function value is selected as the optimum extended solution for the cluster. It is then grafted into the initial extended solution in place of the original cluster. Figure 4.8 shows the repaired initial extended solution after overlap-and-average. In this case, since the cluster includes the entire range of the extended solution, the optimal extended solution replaces the initial extended solution. Figure 4.9 shows the closest point-wise ambiguity field.

Refine Solution

Finally, Fig. 4.9 is divided into overlapping regions (50% overlap), and least-squares model fits are performed on each of the overlapping wind fields. Each model fit is optimized, and the resulting model-based winds, after overlap-and-average, are displayed in Fig. 4.10. Figure 4.11 shows the closest point-wise ambiguity field to the final model-based product.

First Along-Track of Region	Extended Clusters (Marked with 1)	Extended Clusters with Anchors (Clusters Marked with 1 and Anchors Marked with "A")
43	0	A
49	0	1
55	1	1
61	1	1
67	1	1
73	1	1
79	1	1
85	1	1
91	0	1
97	0	1
103	1	1
109	1	1
115	1	1
121	1	1
127	1	1
133	1	1
139	1	1

Table 4.3: Extended clusters and anchor regions. The example from Table 4.2 is continued in this table. The clusters from Table 4.2 are extended to form the clusters in the "Extended Clusters" column. In order to extend the clusters, each region that is a part of the original clusters are marked with a "1" in the "Extended Clusters" column of this table. After anchors are chosen, the extended clusters are further expanded until they border anchors or the boundaries of the initial extended solution.

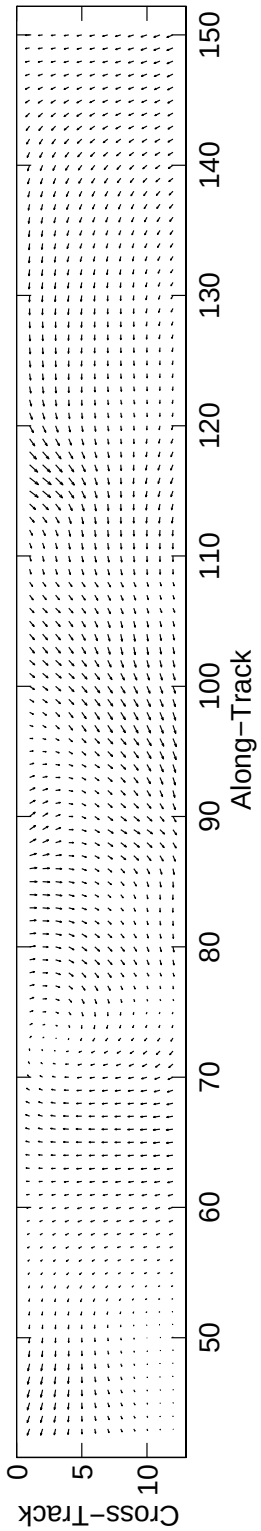

Figure 4.8: Repaired initial extended solution after overlap-and-average

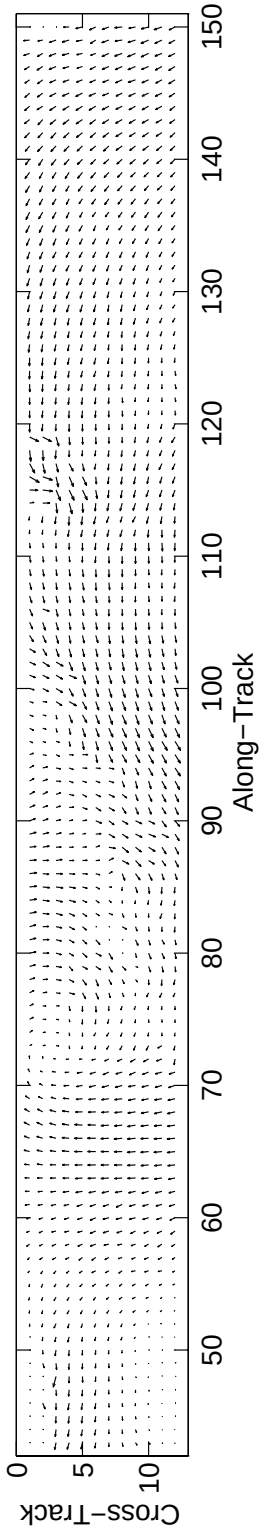

Figure 4.9: Closest point-wise ambiguity field to Fig. 4.8

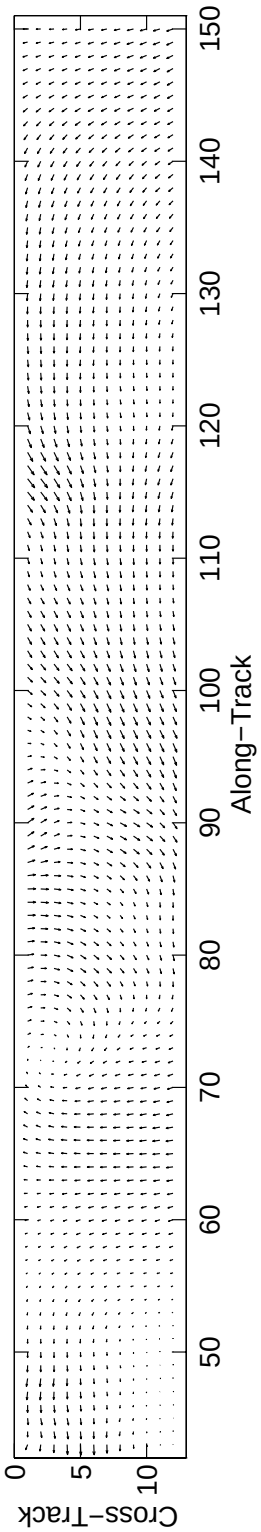

Figure 4.10: Model-based winds after overlap-and-average

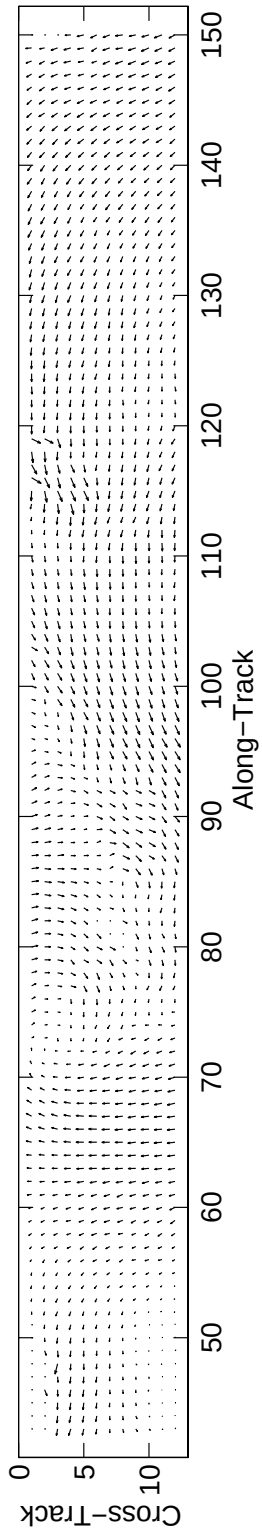

Figure 4.11: Closest point-wise ambiguity field to Fig. 4.10

4.4 Summary

A FWAR algorithm is developed and tested in this chapter. The FWAR algorithm is designed to use scatterometer measurements only.

The first stage in the algorithm is to form an initial extended solution from the individual highest ranked field-wise estimates. Discontinuities in the initial extended solution are then detected, and the vrms threshold value used in the detection algorithm is supported by tests on NSCAT data.

Once the discontinuities are detected, they are repaired by patching in lower ranked solutions that "fit" according to another vrms threshold value, called the retention threshold. Choice of this threshold value is also supported by tests on NSCAT data. It is observed that the algorithm is not sensitive to the particular choice of discontinuity and retention thresholds.

Finally, the repaired initial solution is refined using model-based estimation. The chapter concludes with a detailed example of the FWAR algorithm.

Chapter 5

Results

To demonstrate and test the field-wise wind retrieval algorithm, the algorithm is run on a withheld set of 45 revolutions from week 27. The set covers approximately half of the week. In addition, the median filter is run on the data set, using the first point-wise ambiguities as the initial field. Further, the refined JPL product is obtained by performing model-based estimation on the raw JPL product.

The data listed above are used to evaluate the field-wise method. It is determined that the method compares favorably with the JPL product and that it can be used to repair severe ambiguity removal errors in the median filter output. Another application of the field-wise method is validation of the JPL product. Since the algorithm does not employ the median filter, it can be considered an independent validation of the JPL method. Validation using the field-wise technique is discussed, and a faster method using only model-based retrieval is introduced.

5.1 Preliminary Definitions

Several comparison tools and metrics are used in this chapter to evaluate the field-wise, median filter, and JPL ambiguity removal algorithms. These are described here.

5.1.1 Ambiguity Removal Skill

Ambiguity removal skill is perhaps the most straightforward evaluation tool. Ambiguity removal skill is defined as the percentage of wind vectors selected by an ambiguity removal algorithm that are the closest vectors to the true wind [19].

Since the true wind is not available, the skill comparison is performed relative to the JPL product, which has been shown to be only about 96% correct [22].

5.1.2 Greater-Than-90° Error Percentage

The greater-than-90° error percentage is the percentage of wind vectors that are more than 90° in error with respect to reference winds [12]. In the case of comparison with the JPL product, this measure is almost certain to reflect retrieval errors, since the JPL product is assumed usually to display the correct general wind flow. The reliability of the JPL product in this respect arises from the NCEP initial value field that correctly orients the streamline before the median filtering.

5.1.3 Vector Correlation

Vector correlation is a method of determining the similarity of two vector sequences and is defined in Ref. [45]. Let the individual vectors of the first vector sequence be denoted $\mathbf{W_{1i}}$, $i = 1, 2, ..., N$. Denote the corresponding vectors in the second sequence as $\mathbf{W_{2i}}$, $i = 1, 2, ..., N$. Form the vectors

$$\mathbf{X_i} = \begin{pmatrix} \mathbf{W_{1i}} \\ \mathbf{W_{2i}} \end{pmatrix},$$

where $i = 1, 2, ..., N$. The vector correlation ρ_v^2 is estimated from the sample covariance matrix of the $\mathbf{X_i}$ sequence:

$$\begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix} = \frac{1}{N-1} \sum_{i=1}^{N} (\mathbf{X_i} - \bar{\mathbf{X}}_i)(\mathbf{X_i} - \bar{\mathbf{X}}_i)^T.$$

The estimate of vector correlation is then

$$\hat{\rho}_v^2 = Tr(S_{11}^{-1} S_{12} S_{22}^{-1} S_{21}).$$

The result is a scalar value which describes the similarity of the vector sequences $\mathbf{W_{1i}}$, $i = 1, 2, ..., N$ and $\mathbf{W_{2i}}$, $i = 1, 2, ..., N$. $\hat{\rho}_v^2$ varies from 0 to 2, with higher values indicating greater correlation. Two vector sequences with a $\hat{\rho}_v^2$ of 1.5 or greater are considered highly correlated [45].

5.2 Evaluation of Field-Wise Wind Retrieval Algorithm

Evaluation of the the field-wise algorithm is accomplished by comparing its output with the JPL product, using the closest point-wise ambiguities to the field-wise output. The median filter output also is compared with the JPL product, in order to provide a benchmark for the field-wise performance.

In the evaluation process, the wind swath solutions from the field-wise, median filter, and JPL methods are segmented into 12x12 regions with 0% overlap. Recall from Chapter 4 that the ambiguity removal stage of the field-wise algorithm sets warning flags for regions with possible ambiguity removal errors. Accordingly, vectors in regions flagged with warnings are not considered in the evaluations. In addition, only wind vectors where the speeds of the JPL product are above 4 m/s are included, as it has been shown that ambiguity removal performance degrades for low wind speeds [22]. Since the performance of field-wise techniques can be poor in regions with sparse data or large numbers of low wind speed vectors, only vectors in regions where more than 85% of the resolution elements have valid data *and* the corresponding JPL vectors have speeds greater than 4 m/s are considered. About 75% of the data satisfy these constraints.

Table 5.1 displays the ambiguity removal skills of the field-wise and median filter methods, compared with the JPL product, for the constraints listed above. The marginally better statistics for the median filter are misleading. Investigation of 90° error percentages and vector correlation illuminates the advantages of the field-wise method.

Table 5.2 lists the 90° error percentages for the field-wise and median filter products, relative to the JPL product. Note that the field-wise method has fewer significant deviations from the JPL product than the median filter. That is, the field-wise method tends to predict the correct general wind flow better than the median filter. Table 5.3 tabulates the vector correlation between the JPL product and the field-wise and median filter outputs. The higher correlation value for field-wise retrieval output indicates that it is in better general agreement with the JPL product than the median filter is.

Algorithm	Skill
Field-Wise	95.0%
Median Filter	95.4%

Table 5.1: Ambiguity removal skills, relative to the JPL product. Vectors from regions with field-wise ambiguity removal warnings are not included in the skill calculations. Vectors corresponding to JPL winds with speeds 4 m/s or lower are not included in the skill calculations. Further constraints are listed in the text.

Algorithm	Greater-Than-90° Error Percentage
Field-Wise	1.8%
Median Filter	2.1%

Table 5.2: Greater-Than-90° error percentages of field-wise and median filter methods. Vectors from regions with field-wise ambiguity removal warnings are not included in the skill calculations. Vectors corresponding to JPL winds with speeds 4 m/s or lower are not included in the skill calculations. Further constraints are listed in the text.

Algorithm	Vector Correlation Value
Field-Wise	1.83
Median Filter	1.78

Table 5.3: Vector correlation values of field-wise and median filter methods with the JPL product. Vectors from regions with field-wise ambiguity removal warnings are not included in the skill calculations. Vectors corresponding to JPL winds with speeds 4 m/s or lower are not included in the skill calculations. Further constraints are listed in the text.

Algorithm	Skill
Field-Wise	90.9%
Median Filter	92.6%

Table 5.4: Skill values of field-wise and median filter methods for JPL wind vectors with speeds greater than 4 m/s.

Algorithm	Greater-Than-90° Error Percentage
Field-Wise	5.0%
Median Filter	4.0%

Table 5.5: Greater-than-90° error percentage for field-wise and median filter methods for JPL wind vectors with speeds greater than 4 m/s.

When the constraints on the vectors used in the evaluations are relaxed, the performance of the field-wise method degrades. Table 5.4 lists skill values and Table 5.5 presents the greater-than-90° error percentages using *all* vectors with corresponding JPL wind speeds above 4 m/s. Observe that the field-wise algorithm differs more than 90° from the JPL product 5% of the time, while the median filter differs to that extent only 4% of the time. The 1% higher greater-than-90° error percentage of the field-wise method accounts for most of the skill difference from the median filter, and the other 0.7% is probably due to the similarity between the median filter and the nudging algorithm. At least 1% of the difference is not a result of modeling error problems, since the greater-than-90° error percentages are about the same when the two algorithms are compared with the refined JPL product.

Several plots of the outputs of the field-wise, median filter, and JPL algorithms are included to illustrate typical performance of the field-wise method. The first two sets of plots demonstrate the effectiveness of the algorithm on difficult features, and the third set of plots depicts a region of greater-than-90° errors.

Figures 5.1, 5.2, and 5.3 show the field-wise, median filter, and JPL products for a portion of revolution 2856 descending, and all three algorithms perform

reasonably well in this example. However, the field-wise and JPL algorithms turn the cyclone located at along-track 305 better than the median filter.

The outputs of the three algorithms are displayed for revolution 2881 descending in Figs. 5.4, 5.5, and 5.6. All three algorithms resolve the cyclone at along-track 337, but the JPL algorithm differs from the field-wise and median filter methods in the area around along-track 343 and cross-track 3. It is not obvious which interpretation of this area is correct.

Figures 5.7, 5.8, and 5.9 illustrate a case where the field-wise algorithm fails by greater than 90°. Figures 5.8 and 5.9 are the median filter and JPL products for a segment of revolution 2945 ascending, and Fig. 5.7 is the field-wise product. Observe that the field-wise product reverses the wind direction from along-track 169 to about 205; however, this set regions is flagged with warnings by the field-wise ambiguity removal stage.

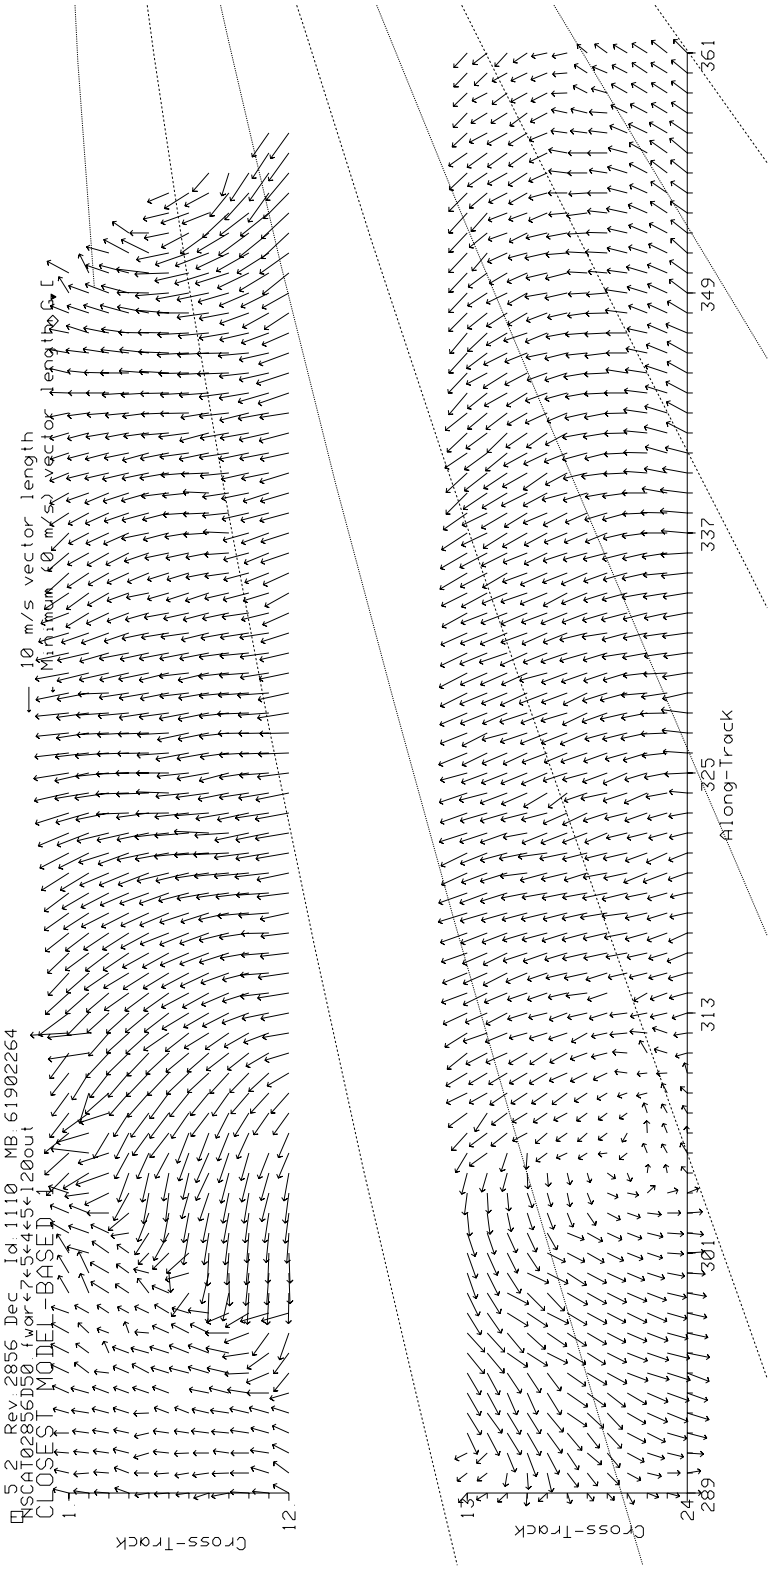

Figure 5.1: Closest ambiguity to field-wise product for a portion of revolution 2856 descending.

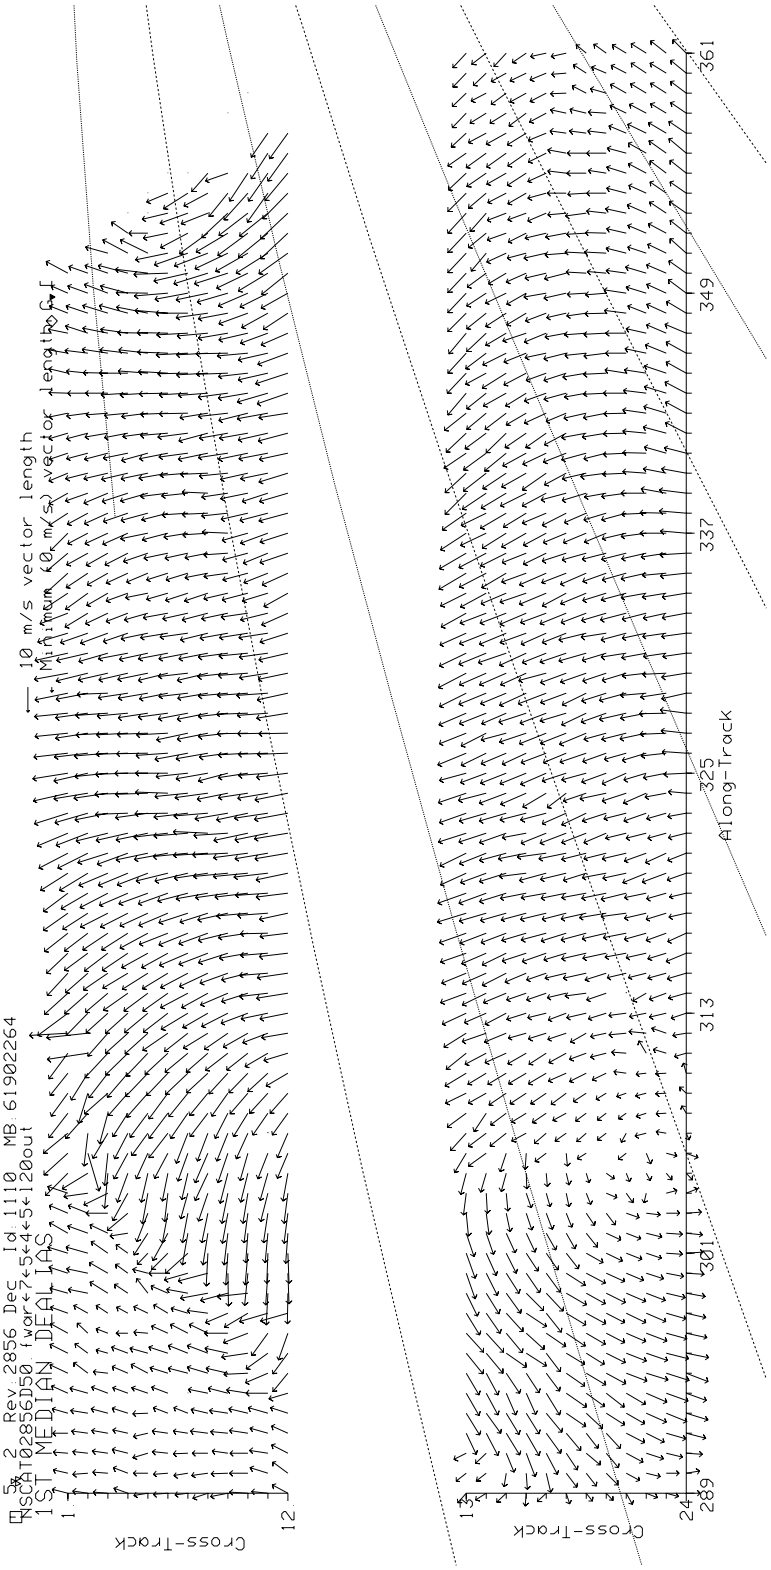

Figure 5.2: Median filter product for a portion of revolution 2856 descending.

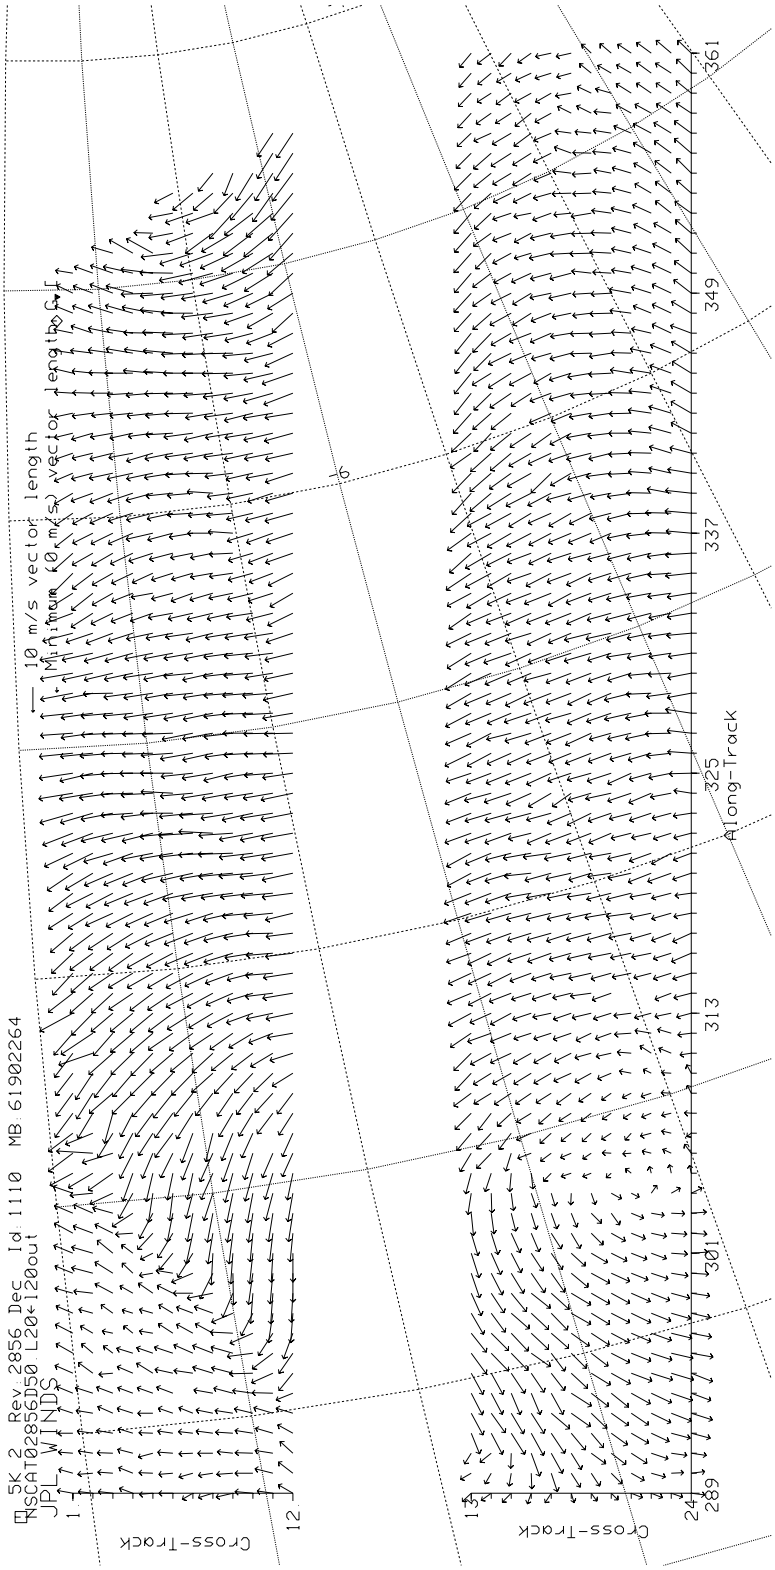

Figure 5.3: JPL product for a portion of revolution 2856 descending.

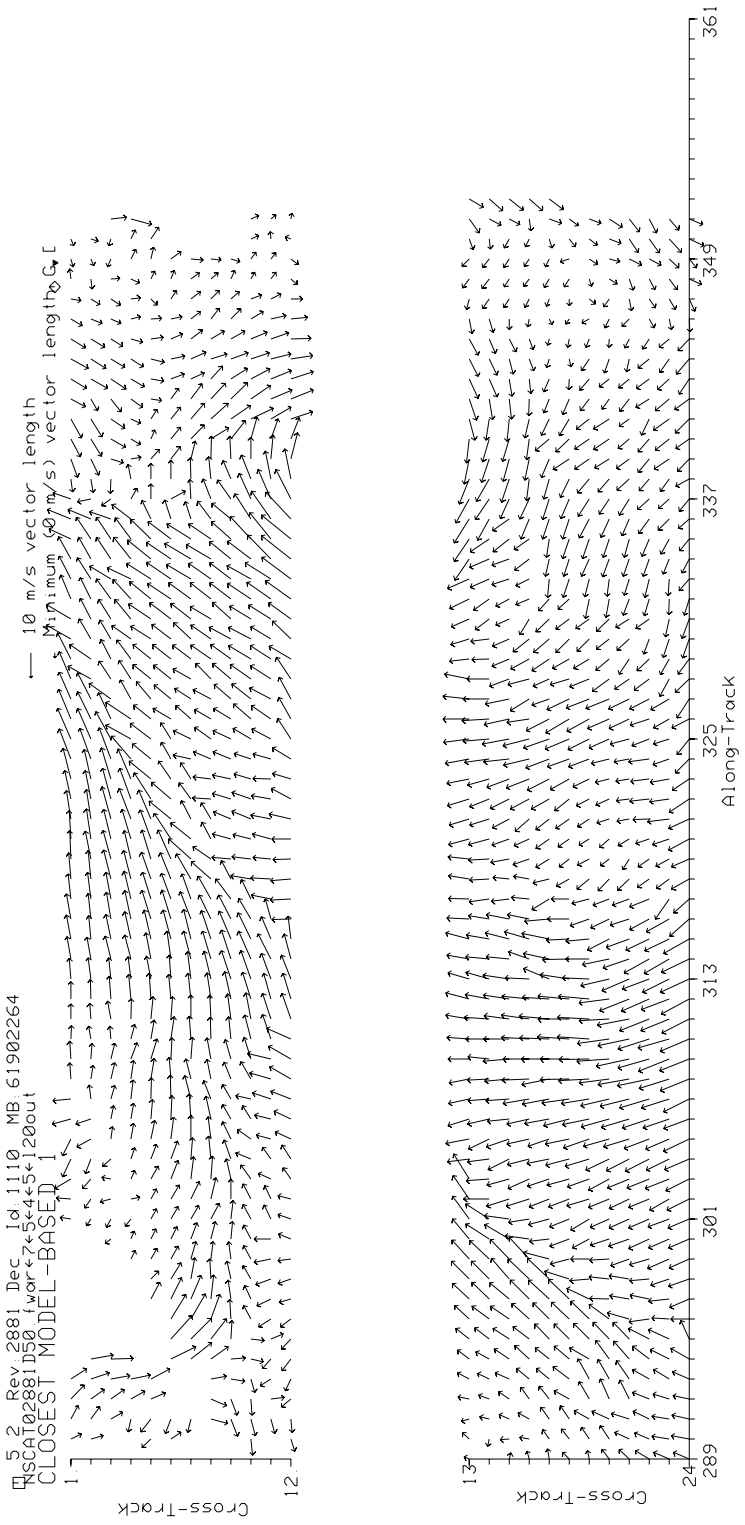

Figure 5.4: Closest ambiguity to field-wise product for a portion of revolution 2881 descending.

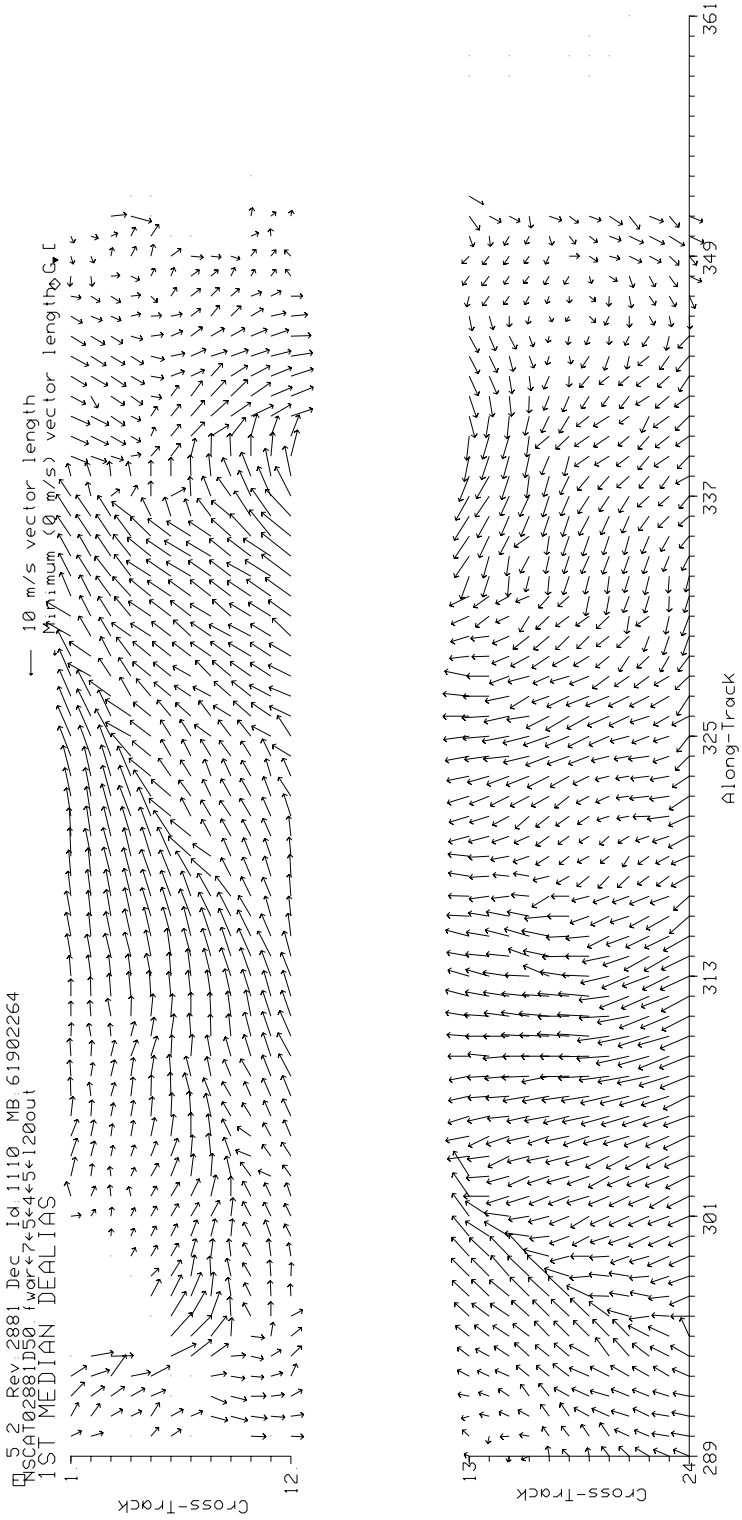

Figure 5.5: Median filter product for a portion of revolution 2881 descending.

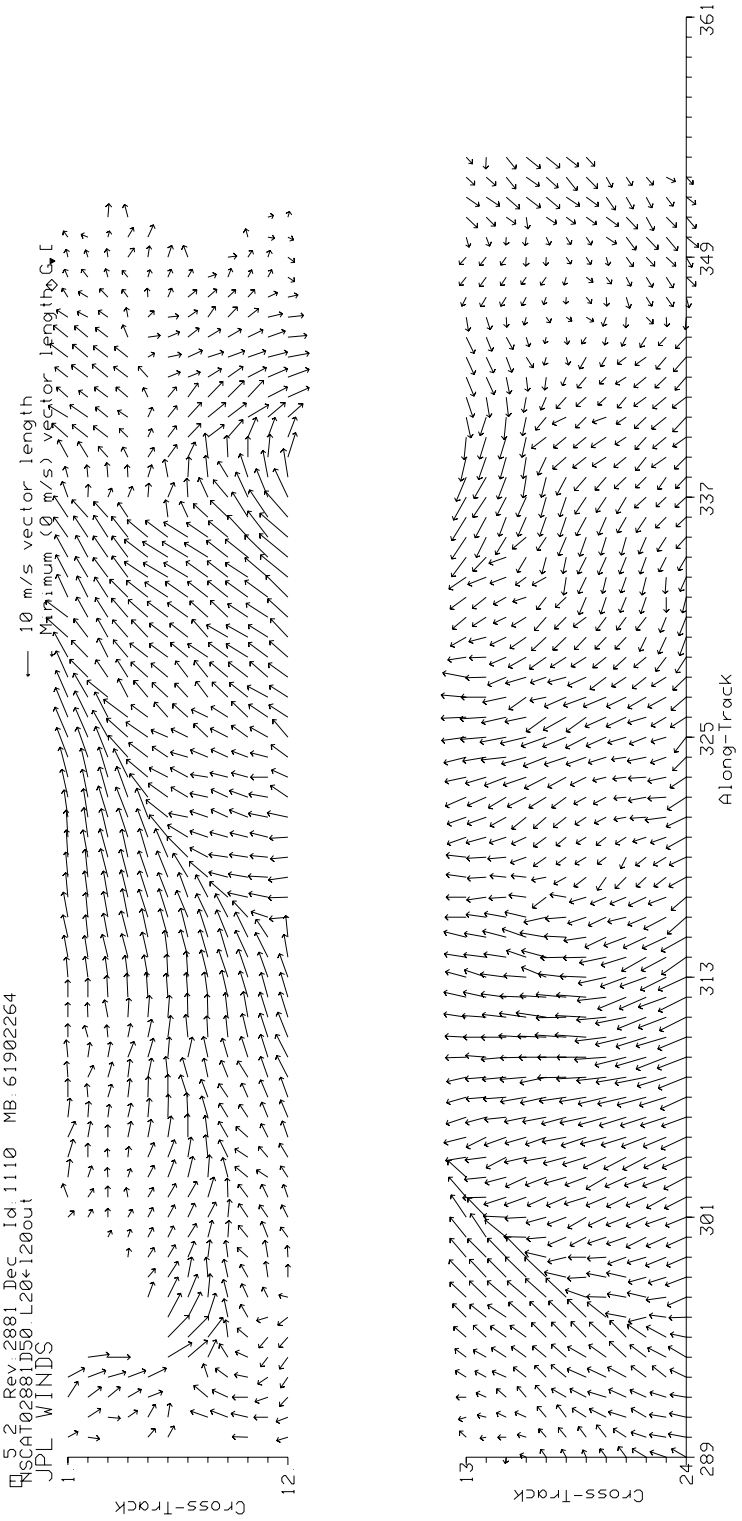

Figure 5.6: JPL product for a portion of revolution 2881 descending.

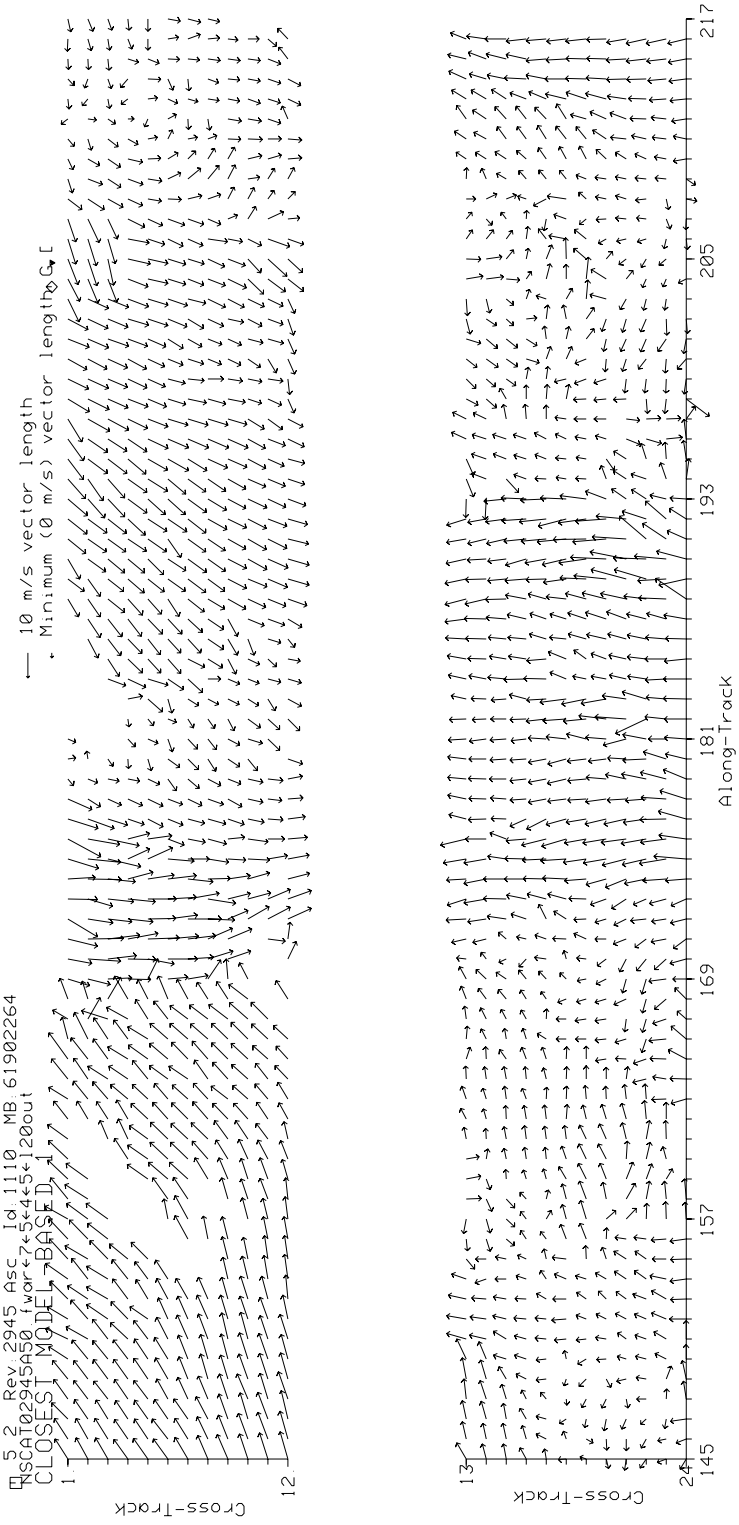

Figure 5.7: Closest ambiguity to field-wise product for a portion of revolution 2945 ascending.

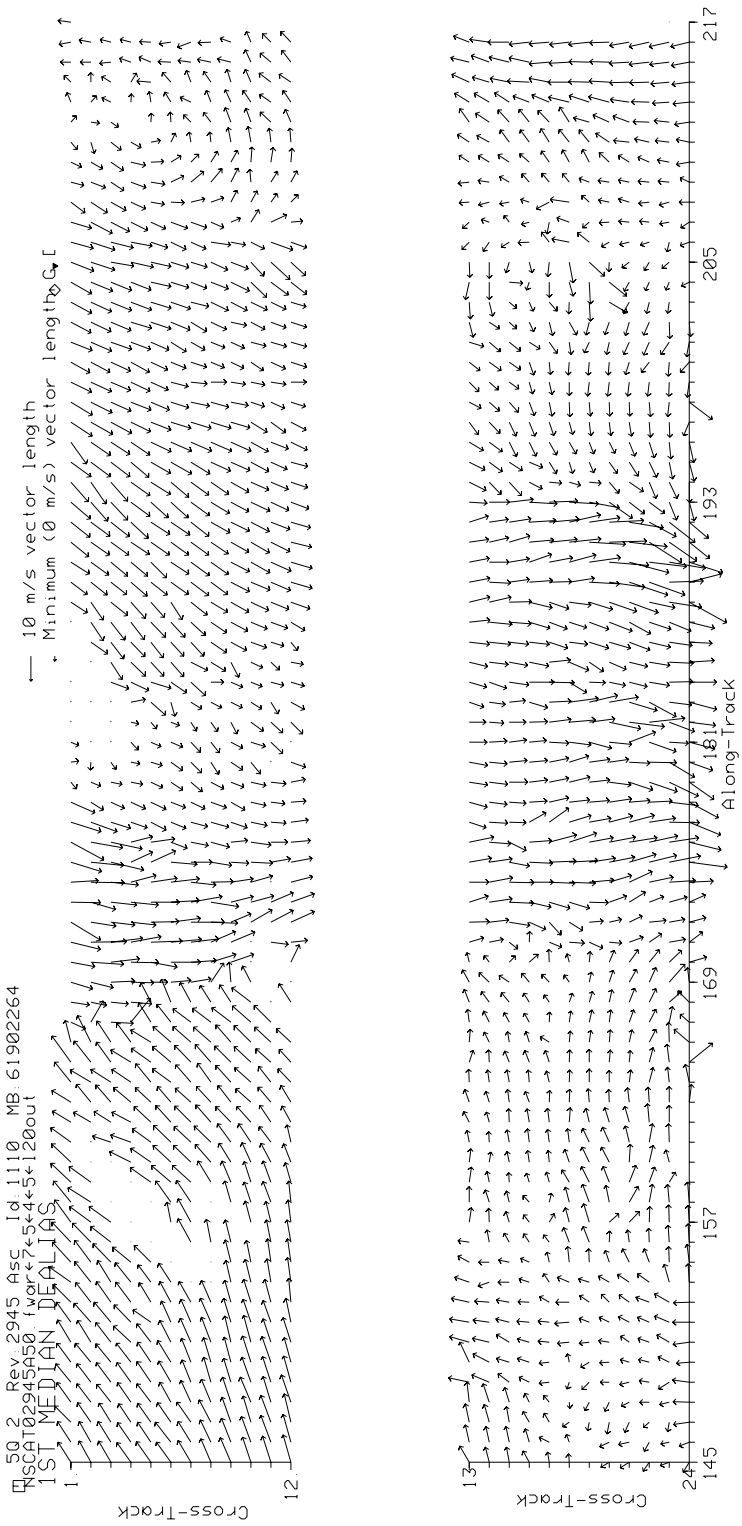

Figure 5.8: Median filter product for a portion of revolution 2945 ascending.

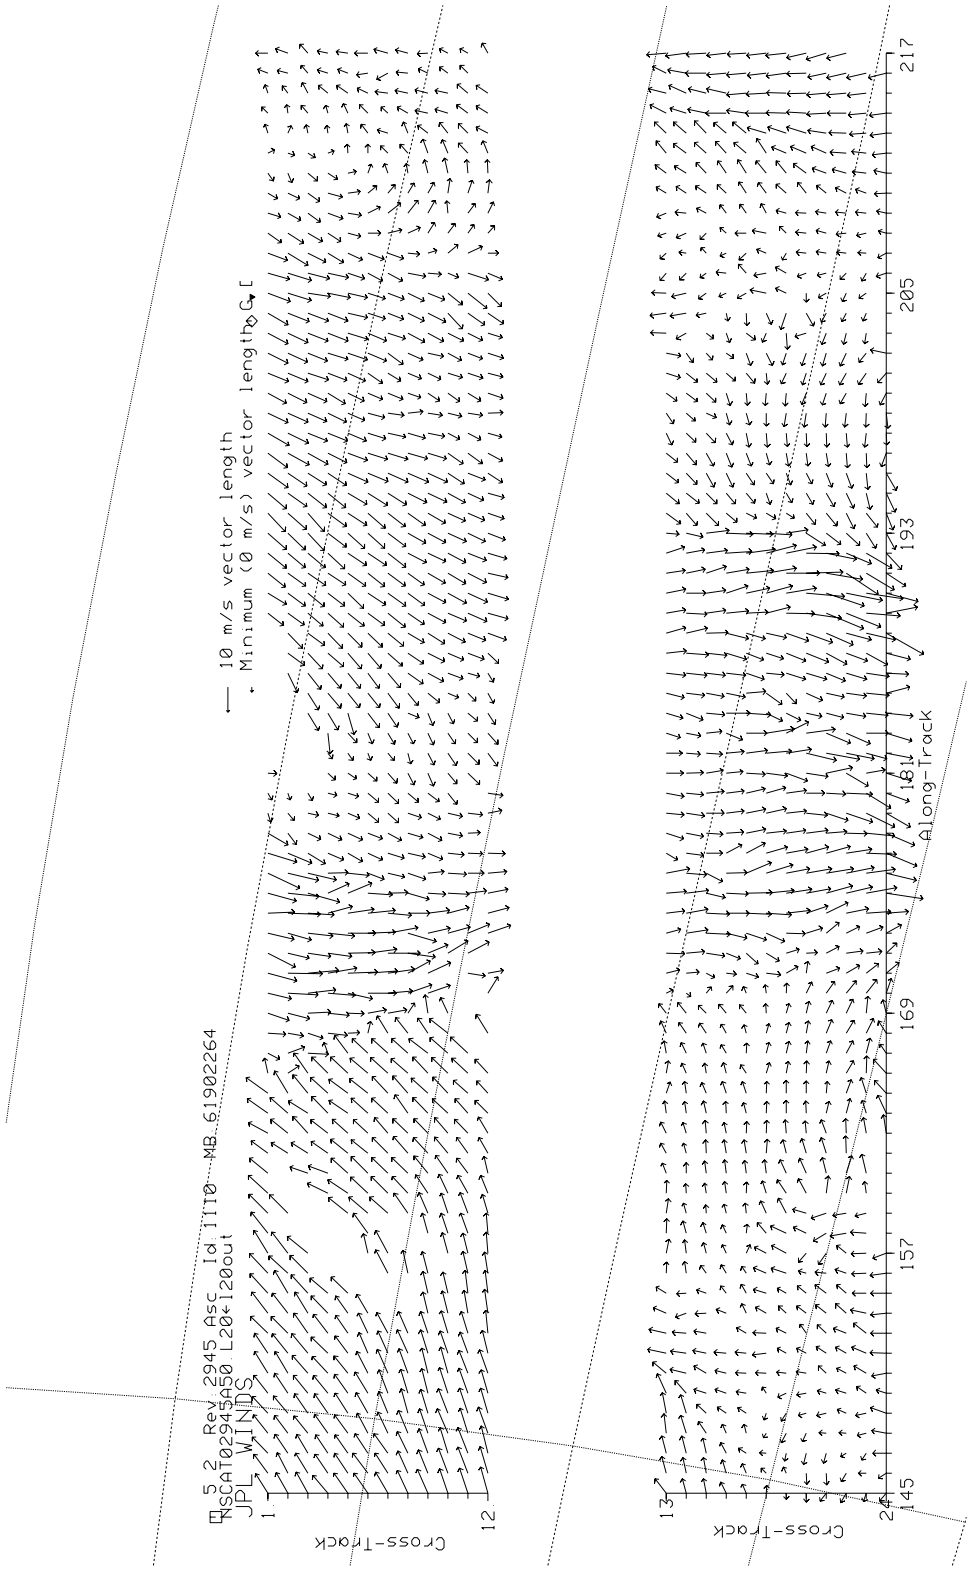

Figure 5.9: JPL product for a portion of revolution 2945 ascending.

5.3 Correction of Median Filter Errors

Evidence suggests that the field-wise method could be a powerful correction technique for the median filter. Table 5.2 indicates that the field-wise method exhibits fewer errors in excess of 90° than the median filter for the constraints defined above. Computation of collocated greater than 90° errors shows that only 1% of the errors in excess of 90° are in common to both algorithms. If all of the excessive errors in the median filter output could be identified, then the field-wise method could correct more than half of them. Perfect detection is probably impossible. However, it is reasonable to expect a detection scheme that could locate regions with extensive blocks of errors, such as occur where the scatterometer skill is low. These regions could be bounded, and the field-wise method, with regions flanking the errors as anchors, could be implemented to correct the errors.

Although such a method improves the median filter product by only about 1%, it corrects the major errors that are otherwise difficult to repair automatically and that are probably the most disruptive to applications incorporating the data. A hybrid algorithm executes much more quickly than field-wise retrieval alone, and provides better data than field-wise retrieval or the median filter used individually.

Figures 5.10, 5.11, 5.12, 5.13, 5.14, and 5.15 display the field-wise, median filter, and JPL products for two regions where a hybrid algorithm would be effective in repairing median filter errors. Figure 5.11 is the median filter output for revolution 2877 descending. Note the rift in the ambiguity removal at along-track 287. Compare the median filter output the field-wise and JPL products in Figs. 5.10 and 5.12, respectively. The field-wise algorithm avoids the error correctly. Other ambiguity removal rifts are in Figs. 5.13, 5.14, and 5.15, which represent a portion of revolution 2859 ascending. Refer to the median filter output (Fig. 5.14). Severe ambiguity removal errors are located in the left swath around along-track 181 and in the right swath at along-track 200. Comparison with the field-wise and JPL products in Figs. 5.13 and 5.15 reveals that the field-wise method correctly mends the rifts.

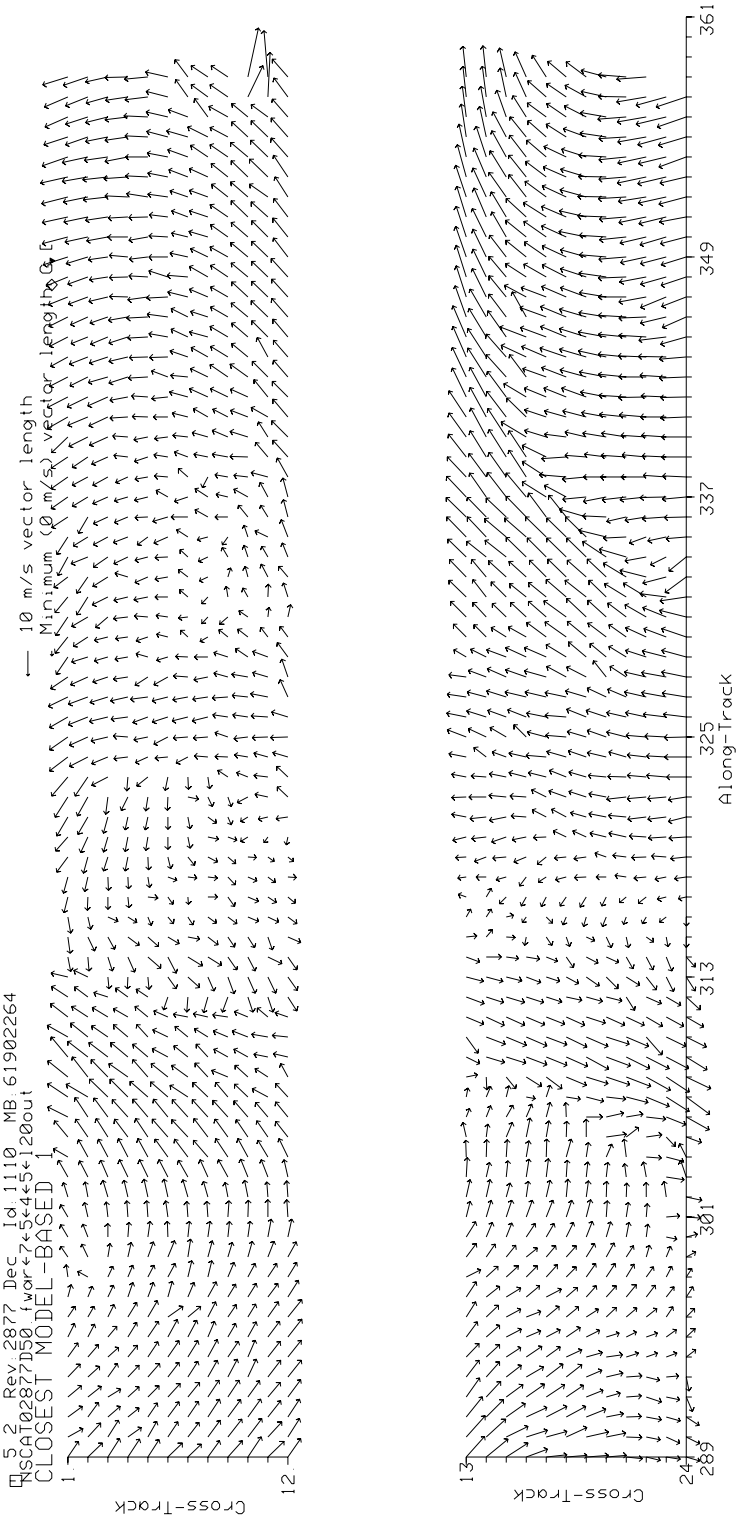

Figure 5.10: Closest ambiguity to field-wise product for a portion of revolution 2877 descending.

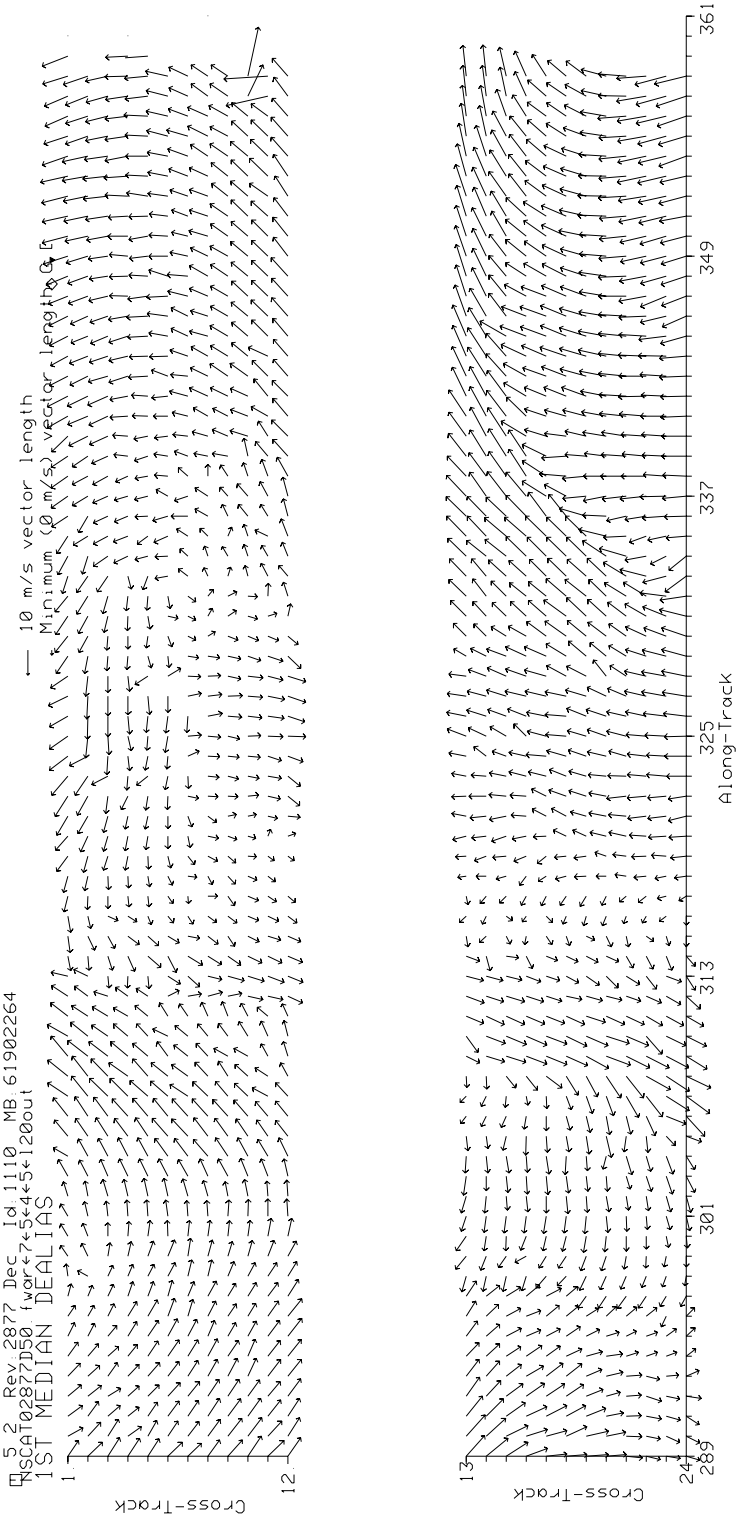

Figure 5.11: Median filter product for a portion of revolution 2877 descending.

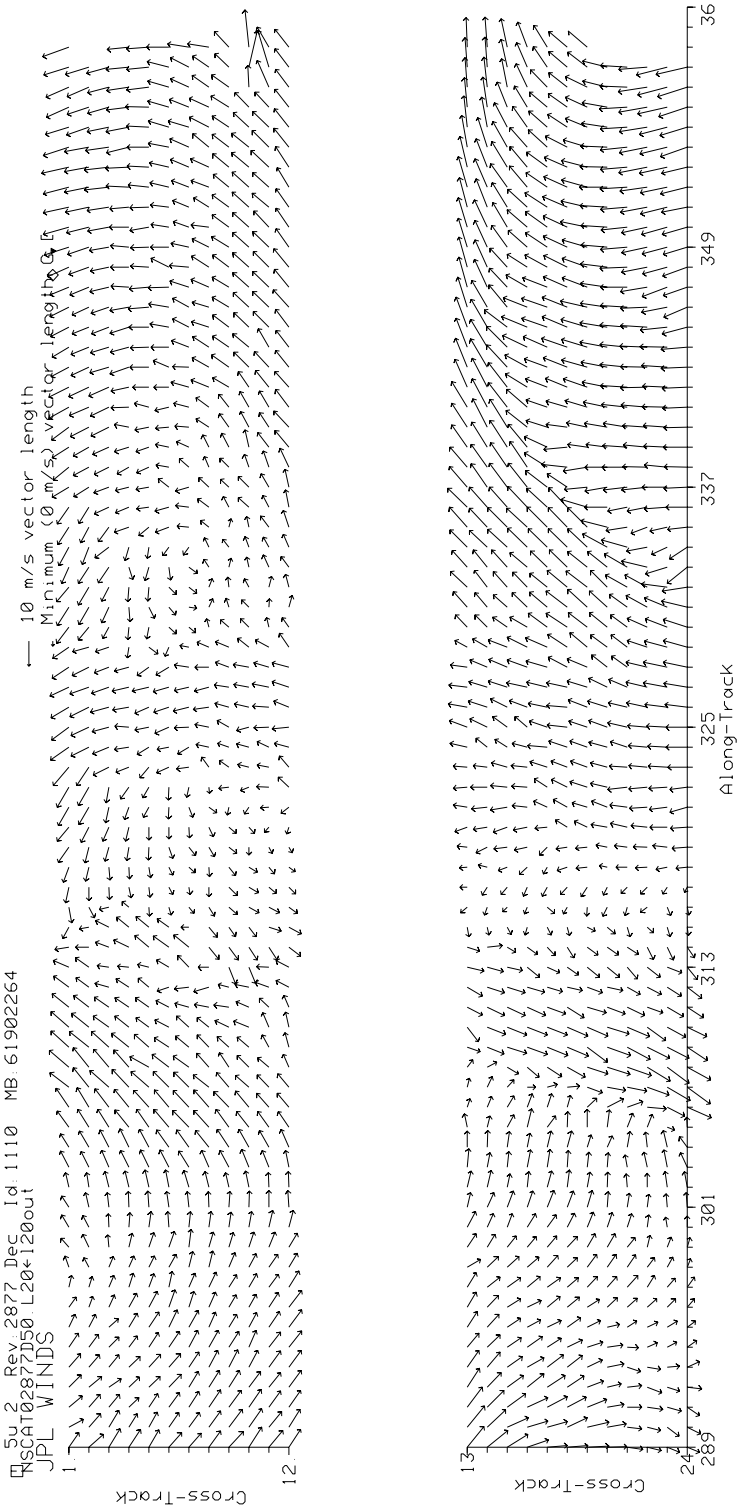

Figure 5.12: JPL product for a portion of revolution 2877 descending.

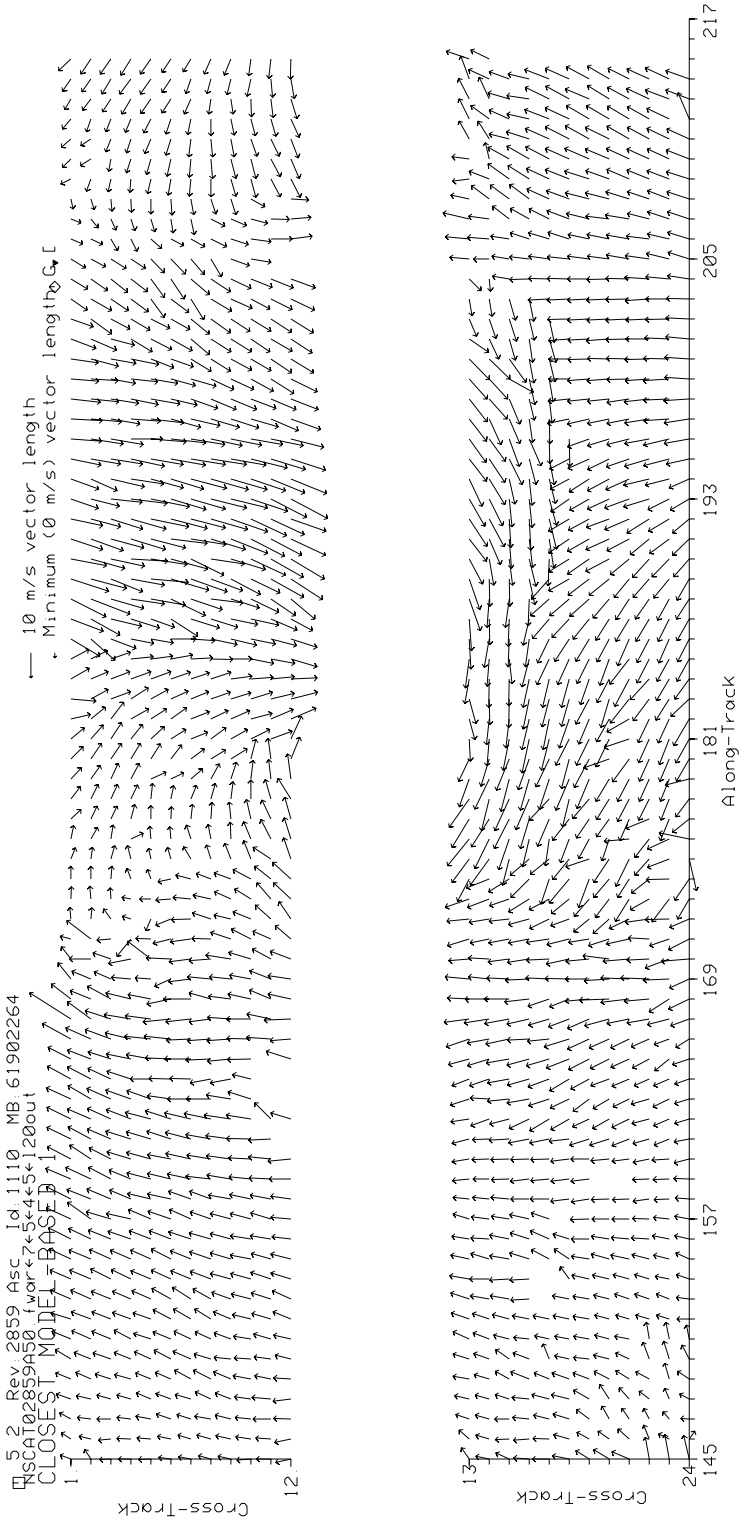

Figure 5.13: Closest ambiguity to field-wise product for a portion of revolution 2859 ascending.

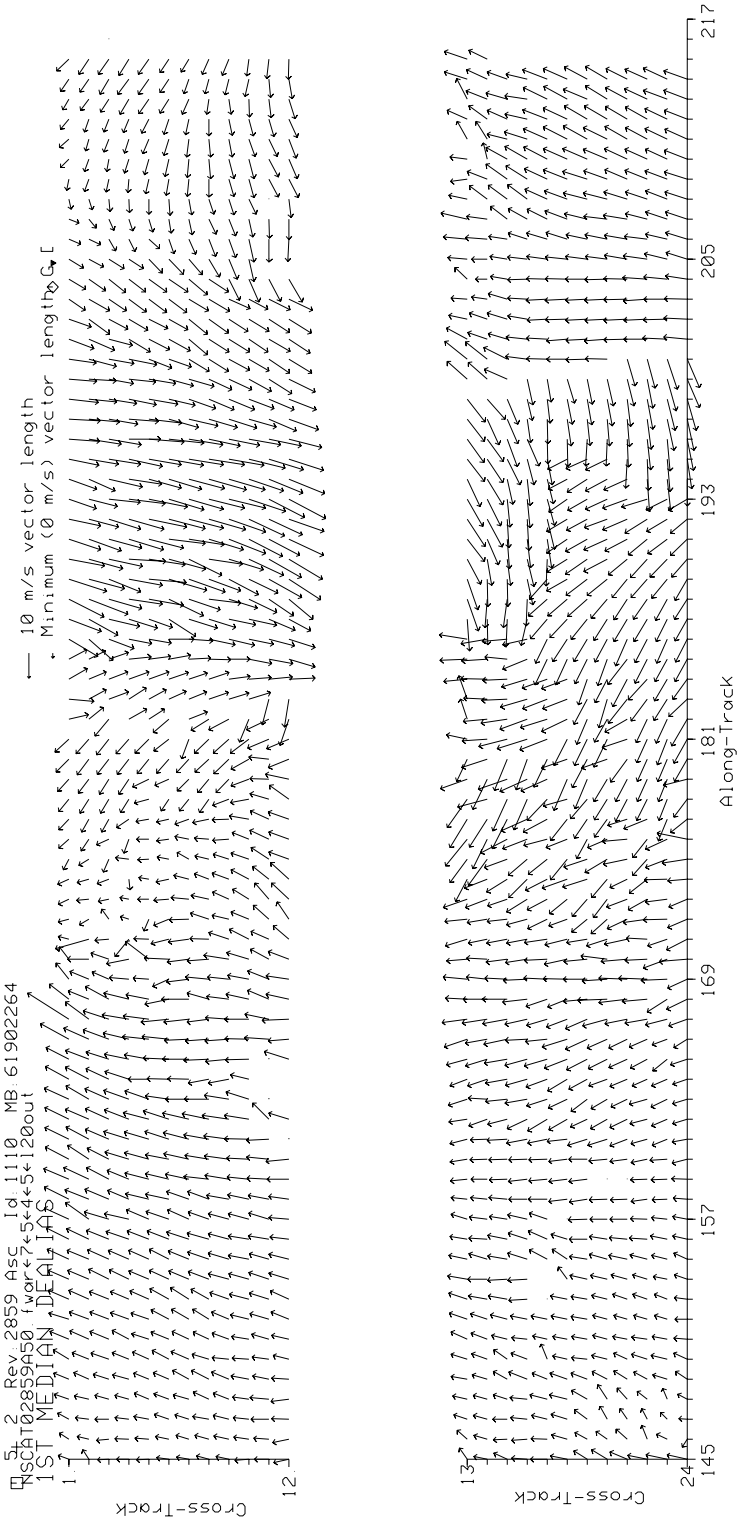

Figure 5.14: Median filter product for a portion of revolution 2859 ascending.

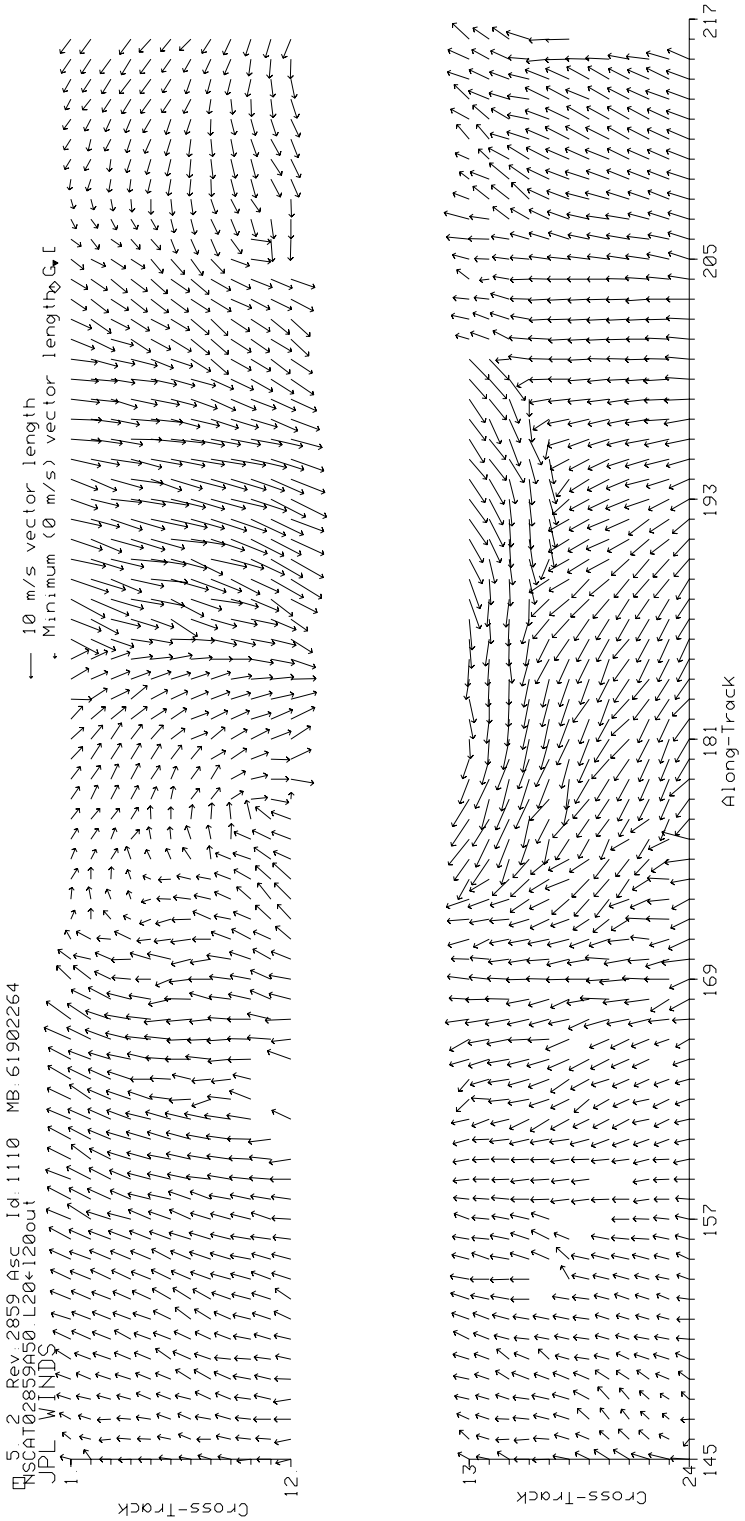

Figure 5.15: JPL product for a portion of revolution 2859 ascending.

5.4 Validation of NSCAT JPL Product

In addition to repairing severe ambiguity removal errors in the median filter output, the field-wise algorithm can be employed to validate the JPL product. The field-wise and JPL "nudging" methods use decidedly different estimation and ambiguity removal techniques. Perhaps the most significant difference is that the JPL method uses the median filter, while the field-wise technique does not. Another difference is that the "nudging" algorithm uses NCEP winds in its ambiguity removal. The field-wise method does not use this outside information. Since the field-wise method is algorithmically independent of the "nudging" algorithm, it is a useful comparison tool.

The relative similarity between the field-wise and JPL products is calculated for vectors satisfying the strict constraints in the previous section. The skills of the field-wise and median filter methods are provided in Table 5.1. The performance of both algorithms is very close. The greater-than-90° error percentage is tabulated in Table 5.2. From Table 5.2 it is observed that about 1.8% of the difference between the field-wise and JPL products is most likely due to errors in the field-wise method, but that does not account for the remaining 3.2% of the difference, which could be field-wise or JPL errors. Thus, we conclude that the JPL product is probably more than 96% accurate, compared to the true wind. This figure is consistent with the observations in Refs. [12, 22].

Another method used in this thesis to validate the JPL product employs model-based estimation. Model-based estimation refines the JPL product, and the JPL product is compared with the refined product. It is assumed that the statistically justifiable model and the model-based optimization correct minor incongruities in the JPL product; however, the model-based estimation can introduce errors, so the comparison yields only an approximate lower bound on the skill. The refinement process changes only 3.9% of the JPL vectors, which again suggests a JPL skill of about 96% or higher.

5.5 Summary

In this chapter, the field-wise wind retrieval algorithm is run on a withheld set of NSCAT data. The output of the field-wise method is evaluated. The output compares favorably with the JPL product. In contrast with the JPL technique, the field-wise method has the advantage of not requiring outside information. It is further noted that field-wise retrieval can be employed to repair severe ambiguity removal errors in the median filter output. An additional application of the field-wise method is validation of the JPL product, since the field-wise and JPL techniques are algorithmically independent. Validation using the field-wise technique is discussed, and a faster method using only model-based retrieval is introduced.

Chapter 6

Conclusions and Future Research

6.1 Conclusions

In this work a field-wise wind retrieval algorithm, consisting of field-wise estimation and ambiguity removal stages, is designed for NSCAT. The field-wise method does not employ outside information. A variant of the multistart global optimization method is developed to perform field-wise estimation. The field-wise estimation algorithm is tested on NSCAT data. It displays about 99% success at locating local minima acceptably close to the desired solution. A field-wise ambiguity removal algorithm is developed to assemble multiple solutions from field-wise estimation into unique wind estimates for large areas of the ocean's surface.

The full field-wise retrieval algorithm is tested on a withheld set of NSCAT data. Tests indicate that even though the field-wise method uses only scatterometer measurements, its output compares favorably with the JPL product, which is processed using NCEP winds. It is further concluded that the method can be used to repair severe ambiguity removal errors in the median filter output.

A further application of the field-wise method is validation of the JPL product. The field-wise and model-based retrieval methods are employed to gauge the reliability of the JPL product. Both methods suggest a skill of 96% or higher, which is consistent with other recent validation results.

6.2 Summary of Contributions

In summary, the contributions of this thesis are

- Development of a variant of the multistart global optimization method to perform field-wise estimation
- Design of a field-wise ambiguity removal algorithm that does not rely on outside information
- Application and testing of the full field-wise wind retrieval method.

6.3 Future Research

Although this thesis demonstrates the effectiveness and utility of field-wise retrieval, more work is necessary in order to extend the applicability of this technique. Research into field-wise methods is a promising area, especially with the upcoming launches of QuikScat and Seawinds. Several possible extensions of this work are

- Optimize thresholds in field-wise ambiguity removal algorithm
- Implement field-wise retrieval as a median filter correction algorithm
- Run field-wise retrieval on full NSCAT data set
- Run model-based retrieval on full NSCAT data set
- Adapt field-wise retrieval for use with QuikScat and Seawinds.

As progress in field-wise techniques continues, they will become an even more viable alternative to the traditional point-wise methods of scatterometer wind retrieval.

Bibliography

[1] F. M. Naderi, M. H. Freilich, and D. G. Long, "Spaceborne radar measurement of wind velocity over the ocean—an overview of the NSCAT scatterometer system," *Proceedings of the IEEE*, vol. 79, no. 6, pp. 850–866, June 1991.

[2] "Fact sheet:NASA scatterometer," 1998, http://www.jpl.nasa.gov/files/fsheets/nscat.txt.

[3] D. G. Long *et al.*, "Current progress in KU band model functions," NSCAT Science Team Report, 1994.

[4] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave Remote Sensing*, vol. 1, Artech House Inc., Norwood, MA, 1981.

[5] G. L. Pickard and W. J. Emery, *Descriptive Oceanography*, Pergamon Press, 1990.

[6] "NSCAT," 1998, http://winds.jpl.nasa.gov/missions/nscat/nscatindex.html.

[7] David G. Long, "Wind field model-based estimation of Seasat scatterometer winds," *Journal of Geophysical Research*, vol. 98, no. C8, pp. 14651–14668, August 1993.

[8] D. G. Long and J. M. Mendel, "Model-based estimation of wind fields over the ocean from wind scatterometer measurements, part II: Model parameter estimation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 28, no. 3, pp. 361–373, May 1990.

[9] M. L. Khandekar, *Coastal and Estuarine Studies: Operational Analysis and Prediction of Ocean Wind Waves*, Springer-Verlag, 1989.

[10] "The Beaufort scale," 1998, http://www.marinedata.co.uk/ref/beaufort.html.

[11] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave Remote Sensing*, vol. 2, Artech House Inc., Norwood, MA, 1981.

[12] F. J. Wentz and D. K. Smith, "A model function for the ocean normalized radar cross section at 14 GHz derived from NSCAT observations," Early draft of paper in Brigham Young University Microwave Remote Sensing Laboratory files, January 1998.

[13] Travis E. Oliphant, "New techniques for wind scatterometry," M.S. thesis, Brigham Young University, August 1996.

[14] Paul E. Johnson, David G. Long, and Travis E. Oliphant, "Geophysical modeling error in wind scatterometry," in *IGARSS '96*. International Geoscience and Remote Sensing Symposium, 1996.

[15] "NASA scatterometer," 1998, http://hsdn.eoc.nasda.go.jp/guide/guide/satellite/sendata/nscat_e.html.

[16] L. C. Schroeder *et al.*, "Removal of ambiguous wind directions for a *ku*-band wind scatterometer using three different azimuth angles," *IEEE Transactions on Geoscience and Remote Sensing*, vol. GE-23, no. 2, pp. 91–99, March 1985.

[17] C.-Y. Chi and F. Li, "A comparative study of several wind estimation algorithms for spaceborne scatterometers," *IEEE Transactions in Geoscience and Remote Sensing*, vol. GE-26, no. 2, pp. 115–121, March 1988.

[18] T. W. Yu, "A technique for deducing wind directions from satellite microwave measurements of wind speed," *Monthly Weather Review*, vol. 115, pp. 1929–2413, September 1987.

[19] S. J. Shaffer, R. S. Dunbar, S. V. Hsiao, and D. G. Long, "A median-filter-based ambiguity removal algorithm for NSCAT," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 29, no. 1, pp. 167–174, Jan 1991.

[20] M. W. Spencer, "Personal communication," 1998.

[21] C. G. Brown and D. G. Long, "A field-wise wind retrieval algorithm for satellite scatterometers," in *IGARSS*, July 1998, vol. II, pp. 939–941.

[22] A. E. Gonzales, "An assessment of the NASA scatterometer ambiguity removal technique," M.S. thesis, Brigham Young University, 1998.

[23] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, 1985.

[24] C. G. Brown, P. E. Johnson, S. L. Richards, and D. G. Long, "Wind field models and classification," in *Proceedings of SPIE*, July 1997, vol. 3117, pp. 115–122.

[25] C. G. Brown, P. E. Johnson, S. L. Richards, and D. G. Long, "Wind field models and model order selection for wind estimation," in *IGARSS*, August 1997, vol. IV, pp. 1847–1849.

[26] L. L. Scharf, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*, Addison-Wesley, 1991.

[27] R. F. Stengel, *Optimal Control and Estimation*, Dover, 1994.

[28] H. D. Chiang and C. C. Chu, "A systematic search method for obtaining multiple local optimal solutions of nonlinear programming problems," *IEEE Transactions on Circuits and Systems*, vol. 43, no. 2, pp. 99–109, February 1996.

[29] A. E. Sepulveda and L. Epstein, "The repulsion algorithm, a new multistart method for global optimization," *Structural Optimization*, vol. 11, pp. 145–152, 1996.

[30] C. G. Brown and D. G. Long, "Algorithms for field-wise scatterometer wind estimation," in *Proceedings of SPIE*, July 1997, vol. 3117, pp. 123–130.

[31] C. G. Brown and D. G. Long, "Global optimization algorithms for field-wise scatterometer wind estimation," in *IGARSS*, August 1997, vol. I, pp. 353–355.

[32] J. D. Pintér, *Global Optimization in Action*, Kluwer Academic Publishers, 1996.

[33] A. Törn and A. Žilinskas, *Global Optimization*, Number 350 in Lecture Notes in Computer Science. Springer-Verlag, 1989.

[34] R. Vaidyanathan and M. El-Halwagi, "Global optimization of nonconvex nonlinear programs via interval analysis," *Computers and Chemical Engineering*, vol. 18, no. 10, pp. 889–897, 1994.

[35] A. Neumaier, "My views on some global optimization methods," 1998, http://solon.cma.univie.ac.at/ neum/glopt/my_view.html.

[36] B. J. Hoenders and C. H. Slump, "On the determination of the number and multiplicities of zeros of a function," *Computing*, vol. 47, pp. 323–336, 1992.

[37] C. H. Slump and B. H. Hoenders, "The determination of the location of the global maximum of a function in the presence of several local extrema," *IEEE Transactions on Information Theory*, vol. 31, no. 4, pp. 490–497, July 1985.

[38] C. Z. Janikow and D. St. Clair, "Genetic algorithms," *IEEE Potentials*, vol. 14, no. 1, pp. 31–35, February/March 1995.

[39] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison Wesley, 1989.

[40] A. Potvin, "Genetic.m," ftp://ftp.mathworks.com/pub/contrib/v4/optim/ genetic/genetic.m, 1993.

[41] B. L. Miller, "Genetic algorithms with dynamic niche sharing for multimodal function optimization," in *IEEE International Conference on Evolutionary Computation*, 1996, pp. 786–791, Also available at http://www-illigal.ge.uiuc.edu/ bmiller/pubs/pubs.html as a technical report.

[42] B. Betro and F. Schoen, "Sequential stopping rules for the multistart algorithm in global optimisation," *Mathematical Programming*, vol. 38, pp. 271–286, 1987.

[43] C. G. E. Boender and A. H. G. Rinnooy Kan, "Bayesian stopping rules for multistart global optimization methods," *Mathematical Programming*, vol. 37, pp. 59–80, 1987.

[44] A. H. G. Rinnooy Kan and G. T. Timmer, "Stochastic global optimization methods part i: Clustering methods," *Mathematical Programming*, vol. 39, pp. 27–56, 1987.

[45] D. S. Crosby, L. C. Breaker, and W. H. Gemmill, "A proposed definition for vector correlation in geophysics: Theory and application," *Journal of Atmospheric and Oceanic Technology*, vol. 10, pp. 355–367, June 1993.