Improved Analysis Techniques for Scatterometer Wind Estimation

Gregory Dallin Schachterle

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

David G. Long, Chair
Willie K. Harrison
Cammy K. Peterson

Department of Electrical and Computer Engineering

Brigham Young University

ABSTRACT

Improved Analysis Techniques for Scatterometer Wind Estimation

Gregory Dallin Schachterle
Department of Electrical and Computer Engineering, BYU
Master of Science

In this thesis, three improved analysis techniques for scatterometer wind estimation are presented. These techniques build upon previous methods that help validate scatterometer data. This thesis examines the theory connecting the 1D and 2D kinetic energy spectra and uses QuikSCAT data to measure the 2D kinetic energy spectrum of ocean winds. The measured 2D kinetic energy spectrum is compared to the traditional 1D kinetic energy spectrum. The relationship between the 2D kinetic energy spectra and the 1D kinetic energy spectra confirms findings from previous studies that ocean winds modeled in 2D are isotropic and nondivergent. The 1D and 2D kinetic energy spectra also confirm the known conclusion that the zonal and meridional components of ocean winds are uncorrelated.

Through simulation, the wind response function (WRF) is calculated for three different QuikSCAT processing algorithms. The WRF quantifies the contribution that the wind at each point of the surface makes to a given wind estimate. The spatial resolution of the different processing algorithms is estimated by their WRFs. The WRFs imply that the spatial resolution of ultrahigh resolution (UHR) processing is finer than the spatial resolution of conventional drop-in-the-bucket (DIB) processing; the spatial resolution of UHR processing is $\sim$5-10 km while the spatial resolution of DIB slice processing is $\sim$12-15 km and the spatial resolution of coarse resolution DIB egg processing is $\sim$30 km.

Simulation is used to analyze the effectiveness of various wind retrieval and ambiguity selection algorithms. To assist in the simulation, synthetic wind fields are created through extrapolating the 2D Fourier transform of a numerical weather prediction wind field. These synthetic wind fields are sufficiently realistic to evaluate ambiguity selection algorithms. The simulation employs the synthetic wind fields to compare wind estimation with and without direction interval retrieval (DIR) applied. Both UHR and DIB wind estimation processes are performed in the simulation and UHR winds are shown to resolve finer resolution wind features than DIB winds at the cost of being slightly noisier. DIR added to standard QuikSCAT UHR wind estimation drops the wind direction root-mean-squared error by $\sim 10^\circ$ to $\sim 24.74^\circ$ in the swath sweet spot.

Keywords: scatterometer, QuikSCAT, ultrahigh resolution, direction interval retrieval, simulation, kinetic energy spectrum, synthetic wind field, wind response function

ACKNOWLEDGMENTS

I consider this thesis to be the capstone of my pursuits at Brigham Young University. As such, I wish to thank the faculty, staff, and administrators of Brigham Young University along with my family and friends for helping me become who I am as I leave this great university. Their influence in my life stretches far beyond this thesis.

Of particular mention, I thank Dr. Long for his continual patience, counsel, and direction. He gave me enough freedom in my research that I could choose my own way while providing enough guidance so that I never strayed too far from the end goal.

Finally, I thank my wife Samantha for her love and support. I love you Sam.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1. INTRODUCTION

While knowledge of ocean wind speed and direction (wind vectors) is useful in many meteorological studies and commercial ventures, it is difficult to consistently measure these wind vectors over a wide area. In situ measurements taken by buoys or ships are often insufficient for such applications because they are sparse in time and space and tend to be clustered in near-coastal areas. An alternative to in situ measurements is to remotely measure winds with spaceborne scatterometers. Scatterometers are microwave radar instruments that measure and calculate the normalized radar cross-section (σ^0) of a surface. When flown in space, scatterometers provide global coverage and are able to measure σ^0 for a given location on the earth from several azimuth and incidence angles. Over the open ocean, these σ^0 measurements are highly influenced by ocean waves that are primarily caused by the wind. This dependence allows near-surface wind vectors to be inferred from scatterometer data.

When retrieving wind vector estimates from σ^0 measurements, the observation swath of the scatterometer is gridded into wind vector cells (WVCs). For a given WVC, the wind vector is estimated using multiple σ^0 measurements whose footprints fall within the cell. Ocean wind vectors are related to the measured σ^0 values by an empirically derived geophysical model function (GMF), a relationship that maps wind vectors to σ^0 values. Even with multiple σ^0 measurements for a given WVC, inverting the GMF produces multiple wind vector estimates (ambiguities) for each WVC. The GMF ranks the ambiguities in order of likelihood, but this information is often insufficient to determine which ambiguity is closest to the true wind vector at each WVC. Instead, ambiguity selection algorithms take into account the wind vectors of the surrounding WVCs and select a final ambiguity for each WVC during ground processing.

The implementation details of scatterometer wind estimation vary depending upon the scatterometer and desired characteristics of the data product. Different methods of wind retrieval and ambiguity selection algorithms often have varying trade-offs between computational complexity, wind vector resolution, sensitivity to noise, accuracy of the estimated wind vectors, and other attributes. In order to understand these trade-offs completely, proper analytical tools are needed.

1.1 Thesis Statement

This thesis develops and presents improved analysis methods for scatterometer wind estimation. The first method evaluates the kinetic energy spectrum of ocean winds in 2D. The second method estimates the spatial wind response function of scatterometer-measured wind vectors through simulation. This method is demonstrated by finding the wind response function of three different wind retrieval processes for the QuikSCAT scatterometer. The third method quantifies the accuracy of different wind estimation algorithms through simulation. This method includes an enhanced process to create more realistic synthetic wind fields.

1.2 Summary of Results

The development of the improved analysis techniques provides the following contributions.

Historically, measurements of the kinetic energy spectra of ocean winds are performed in 1D. This thesis presents a method to estimate ocean wind spectra in 2D that accounts for 2D spectral leakage and the increased high- to low-frequency bin ratio in 2D analysis. The resulting 2D analysis from this method provides additional validation of assumptions made by other researchers that large-scale atmospheric motions may be modeled as 2D, isotropic, and nondivergent, with uncorrelated orthogonal wind components. This method is developed and demonstrated on data from QuikSCAT.

The contribution that each point on the surface of the earth makes to a σ^0 measurement is known as the σ^0 spatial response function. This thesis expands the notion of the spatial response function to a wind response function, or the contribution that the wind at each point on the surface of the earth makes to the scatterometer-measured wind at a particular WVC. This thesis presents a method to estimate the wind response function and estimates the wind response function for three different QuikSCAT wind retrieval processes. In general, the smaller the WVC is for the wind retrieval process, the finer the spatial resolution is for the wind retrieval process.

Quantifying the accuracy of wind estimation algorithms is difficult without knowing the true wind on the surface of the earth. As the true wind can be known in simulation, this thesis presents a QuikSCAT simulation-based method designed to help quantify the accuracy of wind estimation algorithms. The method is demonstrated by comparing the traditional median filter-based ambiguity selection algorithm with and without direction interval retrieval (DIR) for QuikSCAT ultrahigh resolution (UHR) winds. This provides verification of DIR's effectiveness in UHR wind retrieval and additional verification of the median filter-based algorithm's effectiveness when selecting ambiguities.

For quantifying the accuracy of ambiguity selection algorithms, wind fields with characteristics similar to a wide variety of real ocean winds are required. Previous processes to create synthetic wind fields either model a specific location or wind feature and do not generalize well to the entire ocean, or they have an unrealistic distribution of wind characteristics and features. This thesis presents an improved process to create synthetic wind fields that generalize to wide areas of open ocean while maintaining realistic distributions of wind features.

1.3 Thesis Organization

This thesis is organized as follows: Chapter 2 contains background information related to scatterometers, scatterometer-measured winds, wind field analysis techniques, scatterometer simulation, and prior synthetic wind fields. Chapter 3 explores the nuances of measuring the 2D kinetic energy spectrum of ocean winds and presents a method to do so. Chapter 4 explains how the wind response function can be estimated and derives the wind response function for ocean winds estimated by the QuikSCAT scatterometer with three different wind retrieval processes. Chapter 5 gives a detailed description of creating realistic synthetic wind fields suitable for use in testing ambiguity selection algorithms. Chapter 6 examines a method for testing wind estimation algorithms using simulation which is demonstrated by comparing the QuikSCAT median filter-based ambiguity selection algorithm with and without DIR. Chapter 7 concludes the thesis, summarizes the contributions, and provides suggestions for future work.

CHAPTER 2. BACKGROUND

This chapter provides background information to put the rest of the thesis in context. Specifically, the chapter provides an overview of scatterometers in general and QuikSCAT in particular. The chapter discusses the process through which wind fields are derived from scatterometer estimates, reviews a few common methods used to analyze ocean wind fields, explains a way to simulate scatterometer measurements for wind estimation, and discusses different ways to make a synthetic wind field. Additional background information is provided in subsequent chapters as needed.

2.1 Scatterometers

Spaceborne scatterometers are active microwave radar remote sensing instruments that transmit radar pulses to the surface of the earth and measure the power of the return echoes. They collect measurements in locations where in situ measurements are impractical. Advantages of scatterometers include having a large measurement swath, which provides relatively quick global coverage, and being able to operate in all-weather conditions day and night.

The relationship between the transmitted power (P_t) and the received power (P_r) for a scatterometer can be described by the monostatic radar equation [1]

$$\frac{P_\mathrm{r}}{P_\mathrm{t}} = \iint_A \frac{G^2\lambda^2\sigma^0}{(4\pi)^3R^4}dA, \tag{2.1}$$

where A is the illuminated area, G is the one-way antenna gain, λ is the operating frequency of the radar instrument, σ^0 is the normalized radar cross section of the illuminated area, and R is the slant range to the illuminated area. While scatterometers physically measure P_r, they invert the radar equation to report a measurement for σ^0. Scatterometers report this measurement because σ^0 is predominantly dependent upon the geometric and material properties of the earth's surface within

A. As the surface of the earth changes based on geophysical phenomena, these features can be inferred from σ^0 measurements. Features that scatterometer σ^0 measurements are commonly used to infer include ocean wind speed and direction (wind vectors); soil moisture content; vegetation growth; ice melt and growth; and iceberg location, size, and movement [1]–[5].

2.1.1 QuikSCAT

The SeaWinds scatterometer on the QuikSCAT satellite (commonly referred to as simply QuikSCAT) is a dual-rotating pencil-beam scatterometer. As QuikSCAT orbits the earth, the footprints of the two rotating antenna beams trace a helical pattern on the earth. The inner beam is horizontally polarized (h-pol) with a nominal incidence angle of 46° while the outer beam is vertically polarized (v-pol) with a nominal incidence angle of 54.4°. The QuikSCAT measurement swath is 1800 km wide, though the inner beam only reaches the center 1400 km. A depiction of QuikSCAT's measurement geometry is given in Fig. 2.1.

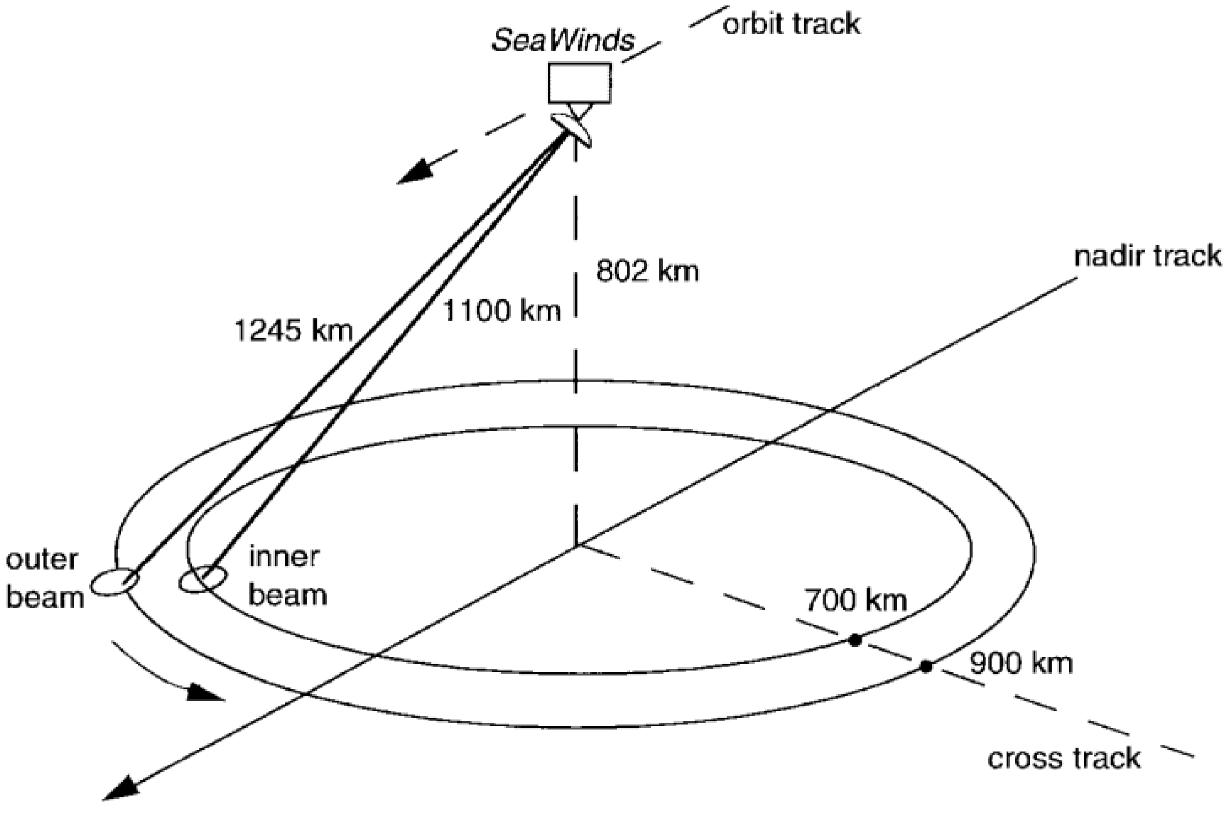

Figure 2.1: The measurement geometry of QuikSCAT [6].

QuikSCAT's measurement swath tends to be divided into different regions that share similar characteristics. These regions are illustrated in Fig. 2.2. The edge regions of the swath, where only the outer beam records measurements, are called the far swath. The center section of the swath, where measurements are almost directly in front of and behind the path of the satellite, is termed the nadir region. Between the far swath and the nadir region, the majority of the swath is called the sweet spot. Within the sweet spot, σ^0 is measured at each location by the beams looking fore and aft of the satellite at a variety of incidence and azimuth angles. Since QuikSCAT is able to produce σ^0 measurements at various azimuth angles, incidence angles, and polarizations within the sweet spot, σ^0 measurements from this region tend to estimate geophysical phenomena more accurately than measurements from the other regions.

QuikSCAT reports σ^0 measurements at two different resolutions: "eggs" (25 km × 35 km) and "slices" (6 km × 25 km). Eggs are calculated using the return of an entire radar pulse and are so named due to QuikSCAT's ellipsoidal footprint. Slices are calculated using range gating, a filtering process that calculates σ^0 only using P_r within a short time window. This process effectively separates the footprint into equidistant slices, hence the name. The finer resolution of slice measurements comes at the cost of higher noise.

In some applications, QuikSCAT σ^0 measurements are separated into four groups referred to as "flavors." The flavors are h-pol aft of the satellite, h-pol fore of the satellite, v-pol aft of the satellite, and v-pol fore of the satellite. QuikSCAT measurements of a given flavor tend to have similar characteristics because they are taken at similar incidence and azimuth angles as nearby measurements from the same flavor.

Each σ^0 measurement has an associated measurement spatial response function (SRF). The SRF quantifies the contribution that each point on the surface makes to a given σ^0 measurement. For QuikSCAT, the antenna gain pattern, the observation geometry, and the onboard signal processing combine together to determine the SRF [8]. While the SRF theoretically continues forever, we limit the size of the SRF for each σ^0 measurement by zeroing out any surface contribution that is less than 20 dB below the peak contribution.

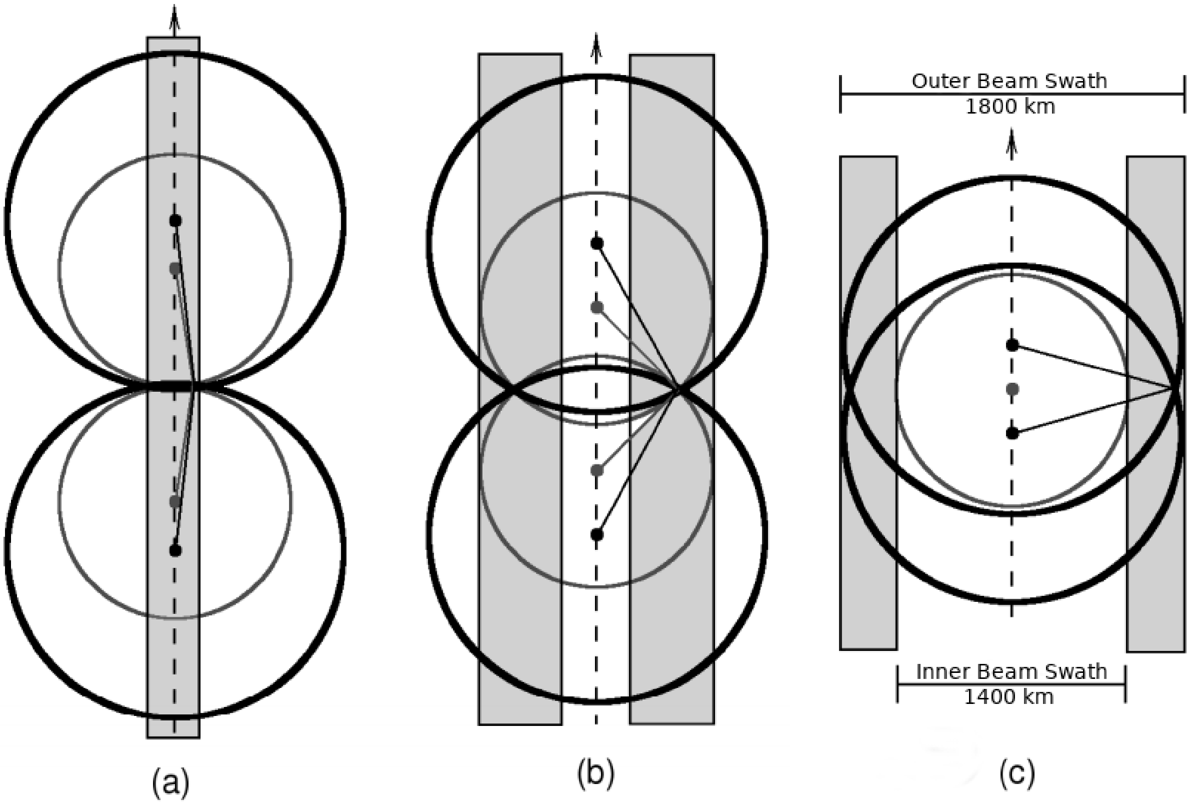

Figure 2.2: The beam orientations of QuikSCAT for measurements at particular swath locations: (a) nadir, (b) sweet spot, and (c) far swath. The dashed line represents the nadir track of the satellite. The bold dots on the nadir track represent different locations of QuikSCAT when various measurements are being taken. The lines extending from the bold dots represent the azimuth angle of the measurements taken for a particular location. The shaded regions show the general region associated with the particular measurement location. The inner circles show the extent of the inner beam's reach while the outer circles show the extent of the outer beam's reach. Note that azimuth angle diversity is greatest in the sweet spot. Modified from [7].

2.2 Scatterometer-Measured Winds

2.2.1 Geophysical Model Function

A major function of scatterometers is to measure ocean wind vectors. Physically, measurements of σ^0 over the ocean are greatly influenced by windblown waves. The relationship between σ^0 and ocean wind vectors is described by an empirically derived geophysical model function (GMF). Most generally, the GMF may be written as

$$\sigma^0 = g(|U|, \chi, \ldots; \theta, f, \text{pol}) \tag{2.2}$$

where $|U|$ is the wind speed, χ is the angle between the wind direction and the azimuth angle of the radar pulse, $\ldots$ represents the small effects of non-wind variables, θ is the incident angle of the radar pulse measured in the vertical plane, f is the frequency of the radar pulse, and pol is the polarization of the radar pulse. The GMF is inverted to estimate the wind vector from σ^0.

2.2.2 Wind Retrieval

The GMF maps many wind vectors to the same σ^0 value. This means that when only one σ^0 measurement is known, a wind vector cannot be uniquely determined. Instead, multiple values of σ^0 must be known from a variety of azimuth and/or incidence angles. (Since the incidence angle of a QuikSCAT measurement implies the polarization of the measurement, the variety of polarization is not enumerated separately in this thesis.) Two measurements from the same azimuth and incidence angle are effectively a single measurement with an improved signal-to-noise ratio (SNR). The need for multiple σ^0 measurements is shown in Fig. 2.3. In Fig. 2.3 (a), each solid curve represents the different wind vectors that map to the same value of σ^0. With multiple σ^0 measurements, the true wind vector may be identified as the intersection of all the curves. The dotted line represents a likelihood objective function, a quantitative measure of how close the curves are to intersecting. The objective function is equal to 0 when all four curves cross and has local minima at wind directions where the curves almost cross. These minima represent near-solution wind vectors.

The wind retrieval process is complicated by noise. Noise in σ^0 measurements is typically quantified by the normalized standard deviation of the σ^0 measurement (K_p). K_p is modeled as [1]

$$K_\text{p} = \sqrt{K_\text{p,A} + \frac{K_\text{p,B}}{S_\text{n}} + \frac{K_\text{p,C}}{S_\text{n}^2}} \tag{2.3}$$

where $K_\text{p,A}$, $K_\text{p,B}$, and $K_\text{p,C}$ are predetermined functions of the scatterometer measurement geometry and onboard signal processing, and S_n is the SNR of the σ^0 measurement. Fig. 2.3 (b) overlays the curves of Fig. 2.3 (a) with shaded bands representing the set of wind vector curves corresponding to measurements in the range from $\sigma^0(1-K_\text{p})$ to $\sigma^0(1+K_\text{p})$. When a noisy realization of the σ^0 measurements is plotted in Fig. 2.3 (c), the unique solution disappears and only the near-solution wind vectors remain. These near-solution wind vectors are referred to as ambiguities.

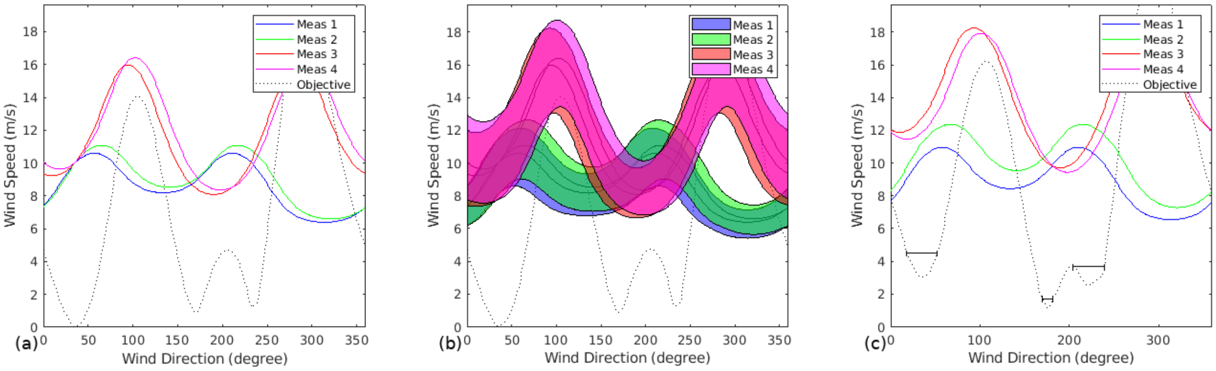

Figure 2.3: A visual representation of wind retrieval. In (a), the solid lines represent the various wind vectors corresponding to four noise-free σ^0 measurements and the dotted line is the output of an objective function quantifying how close the curves are to intersecting. All of the curves intersect at the true wind vector of 10 m/s and 35°. In (b), each curve from (a) is overlaid with a shaded band that represents one standard deviation due to noise on either side of the noiseless measurement values. In (c), one noisy realization of the four σ^0 measurements from (a) is shown with a corresponding objective function and example ranges of wind directions for direction interval retrieval for each ambiguity (see Section 2.2.4). The near-solution wind vectors (ambiguities) for the noisy measurements in (c), ordered from most likely to least likely, are: (1) 10.1 m/s, 176°; (2) 11.1 m/s, 220°; and (3) 11.5 m/s, 35°.

While the ambiguities can be ranked in order of likelihood, the most likely wind vector is not always the ambiguity closest to the true unique solution, as is the case in Fig. 2.3 (c). Instead, the wind retrieval process returns up to four plausible ambiguities for each set of collocated σ^0 measurements and an ambiguity selection algorithm is needed to select a single wind vector from the ambiguities [1] (see Section 2.2.3).

This thesis considers two methods to supply the σ^0 measurements to the objective function in the wind retrieval process: drop-in-the-bucket (DIB) and UHR. In both processes, σ^0 measurements are grouped together because part of the SRFs of the measurements occupy similar areas on the surface of the earth. For estimating wind vectors, the area associated with a set of collocated σ^0 measurements is called a wind vector cell (WVC). WVCs are usually square and they typically fill up the entire scatterometer measurement swath without any gaps or overlaps to form a WVC grid.

In the DIB method, σ^0 measurements (egg or slice) whose SRF center falls within a WVC are used when estimating the wind vector for that WVC. Each σ^0 measurement is only used in one WVC and, within each WVC, every σ^0 is weighted equally with either all other σ^0 measurements or all other σ^0 measurements of the same flavor. The DIB method is typified by the QuikSCAT Level 2B (L2B) Ocean Wind Vectors produced by NASA's Jet Propulsion Laboratory (JPL) [9]. The DIB method tends to produce wind fields with low noise and low resolution (e.g., the finest resolution WVCs in the JPL L2B product are 12.5 km × 12.5 km resolution).

The UHR method employs reconstruction techniques to enhance the spatial resolution of σ^0 measurements prior to wind retrieval [10]. Every slice σ^0 whose SRF includes the WVC center is used to find the wind vector corresponding to that WVC. This leads to every slice σ^0 measurement being used to find the wind vector for multiple WVCs. The reconstruction takes place using the AVE algorithm, where a weighted average is taken of slice σ^0 measurements from the same flavor. The weight for each measurement comes from the SRF value at the WVC center [10]. The UHR method is typified by the UHR QuikSCAT wind fields produced by Brigham Young University (BYU) [11]. As its name implies, UHR wind fields are produced at a finer resolution than the L2B wind fields: QuikSCAT UHR WVCs are 2.5 km × 2.5 km or finer. The cost of this finer resolution is increased computation time and more noise in each wind vector estimate.

2.2.3 Ambiguity Selection Algorithms

Ambiguity selection algorithms select a single wind vector for each WVC from the given ambiguities of that WVC. They tend to assume that wind flow is generally smooth and select an ambiguity based on the wind vectors in surrounding WVCs. Some ambiguity selection algorithms rely on outside information such as numerical weather prediction (NWP) wind fields. As their name suggests, NWP wind fields are wind fields that are predicted from numerical models that use prior atmospheric conditions (as measured by satellites, ships, buoys, weather stations, and other sources) as inputs.

The most commonly used ambiguity selection algorithm is a pointwise algorithm based on a median filter (for convenience, we call this algorithm the median filter algorithm). In the median filter algorithm, a fixed dimension data window passes over the wind field one WVC at a time. At each step, the median wind direction is found for all the wind vectors within the data window. The median wind direction is compared to the wind directions of the ambiguities of the WVC at the center of the data window and the ambiguity associated with the closest wind direction to the median wind direction is stored until the data window passes through the entire wind field. At that point, each WVC is updated to the new ambiguity found by the median filter. The process repeats for several iterations until convergence is reached (meaning that every ambiguity selection in the selected wind field remains the same between successive iterations) or a max number of iterations have gone by. The first iteration of the median filter algorithm requires an initial guess for the wind field, commonly called a nudging field. Some sources for the nudging field include NWP wind fields or the set of most likely ambiguities given by the GMF. Tuning parameters of the median filter include the size of the window, weights associated with the ambiguity likelihood or position within the data window, and whether the median wind vector replaces the median wind direction [12].

2.2.4 Direction Interval Retrieval and Threshold Nudging

Wind estimation in the sweet spot is more accurate than wind estimation in the nadir region and far swath due to the sweet spot σ^0 measurements having better diversity in their incidence and azimuth angles. To improve the wind estimates in the far swath and nadir region, two additional processing algorithms are applied to QuikSCAT winds in JPL processing. They are direction interval retrieval (DIR) and threshold nudging (TN), and are referred to as DIRTH when both are applied together [13].

Oftentimes, the most spatially consistent wind vector for ambiguity selection is not the wind estimate at a local minimum in the objective function, but a nearby wind estimate whose objective function value is similar to the local minimum value. To accommodate these nearby solutions, DIR returns a range of wind directions for each ambiguity, see Fig. 2.3 (c). The range of wind directions for each ambiguity contains a local minimum and the surrounding wind directions whose objective function values are similar to the associated local minimum value, i.e., that falls below a threshold. As the wind speed tends to vary little around each local minimum, only the wind speed at the respective local minimum is reported for each ambiguity direction range.

Objective function values lower than a certain threshold value (which is determined for each minimum) are considered to be similar to the minimum value. This means that narrow minima, minima where the objective function quickly approaches the minimum value, report a smaller range of wind directions than broad minima, minina where the objective function slowly approaches the minimum value. To illustrate, the minima corresponding to the second and third most likely ambiguities are broader than the minimum corresponding to the most likely ambiguity in Fig. 2.3 (c). When a broad minimum happens, a direction within the range of reported directions may be more spatially consistent for ambiguity selection than the direction at the local minimum exactly. With access to a range of wind directions for each ambiguity, ambiguity selection algorithms are able to select more accurate wind vectors for WVCs whose ambiguities are associated with broad local minima. This especially improves wind estimates in the nadir region as broad minima are exceptionally prevalent there.

Winds retrieved with DIR require modified ambiguity selection algorithms that handle the range of wind direction values. The most commonly used ambiguity selection algorithm for winds retrieved with DIR is a modified median filter algorithm. The modified median filter algorithm operates similarly to the median filter algorithm, where a fixed dimension data window passes over the wind field one WVC at a time and the median wind direction is found for all the wind vectors within the data window. The median wind direction is compared to the wind direction ranges of the ambiguities of the WVC at the center of the data window. In this comparison, each ambiguity's wind direction range is represented by a single wind vector. The wind directions within the wind direction range of the ambiguity whose representative wind vector is closest to the median wind direction are then compared to the median wind direction. The wind direction within the wind direction range that is closest to the median wind direction is then stored until the data window passes through the entire wind field. At that point, each WVC is updated to the new ambiguity found by the modified median filter, and each newly selected wind direction becomes the representative wind direction for its associated ambiguity. The process repeats for several iterations until convergence is reached, a max number of iterations have gone by, or the same WVCs keep switching for five iterations in a row (it is common in large wind fields for a few wind vectors to flip back and forth between two different ambiguities each iteration). At the start of the modified median filter algorithm, the nudging of the initial guess only takes into account the representative wind vectors for each ambiguity and the representative wind vectors are initialized to be the wind vectors associated with each local minimum. For further details on the implementation of DIR with UHR processing, see [7].

Threshold nudging is a technique designed to improve wind estimation in the far swath. Most DIB processing algorithms that nudge QuikSCAT winds with an outside wind field only use the first or second most likely ambiguity. These ambiguities are closest to the true direction 90% of the time and are often sufficient to determine the streamline of the wind for the sweet spot and nadir region [13], [14]. However, these ambiguities are often insufficient to determine the streamline of the wind in the far swath causing greater areas of incorrectly selected wind vectors. Threshold nudging allows all ambiguities to be selected during the nudging process in the far swath; using more ambiguities in the sweet spot and nadir region would defeat the purpose of taking scatterometer measurements because the scatterometer wind field would be very similar to the nudging wind field [13].

2.3 Wind Field Analysis

2.3.1 Kinetic Energy Spectra

The kinetic energy spectra of ocean winds is useful in characterizing the spatial variability of winds. For this reason, the kinetic energy spectra is commonly modeled, measured, and even used to validate new processing algorithms (e.g., [15]–[17]). On average, the kinetic energy spectra follows a power law in relation to the wavenumber (k), often called a $k^{-\alpha}$ slope. Depending on the model, measuring instrument, and scale of k values being considered, α is predicted to fall between -3 and $-5/3$ [15], [16].

For ease of interpretation and to follow historical precedents, the 1D ocean wind kinetic energy spectrum (E_{1D}) is often measured and presented instead of the 2D ocean wind kinetic energy spectrum (E_{2D}). Freilich and Chelton show that for an isotropic, turbulent, incompressible 2D fluid, E_{1D} can yield information on the shape of E_{2D} [15]. Under these assumptions, they show that if

$$E_{\text{1D}} = \beta_{\text{1D}} k^{\alpha_{\text{1D}}} \tag{2.4}$$

and

$$E_{\text{2D}} = \beta_{\text{2D}} k^{\alpha_{\text{2D}}} \tag{2.5}$$

then $\alpha_{2\mathrm{D}} = \alpha_{1\mathrm{D}} - 1$ [15][1]. In $E_{2\mathrm{D}}$, k can be represented as $k = \sqrt{k_{\mathrm{x}}^2 + k_{\mathrm{y}}^2}$, where k_{x} and k_{y} are the orthogonal components of k in the x- and y-directions when $E_{2\mathrm{D}}$ is presented on a standard Cartesian plane. As suggested in Eq. 2.5, $E_{2\mathrm{D}}$ is expected to be circularly symmetric on average.

A common way to calculate the kinetic energy spectrum of a signal is to square the magnitude of the Fourier transform of the signal (i.e., the periodogram). Since the kinetic energy spectrum and the magnitude of the Fourier transform are linked, knowing the average kinetic energy spectrum of ocean winds yields information about the average magnitude component of the Fourier transform of ocean winds (i.e., the average magnitude component of the Fourier transform of ocean winds is circularly symmetric and follows a power law).

Spectral analysis on a finite length of ocean wind vectors is often subject to spectral leakage and mismatched boundaries on the edges of the signal. To mitigate the effects of these undesirable features, the wind vector data often undergo some sort of preprocessing before the periodogram is taken. Prewhitening, linear transformation and windowing the data are all common forms of preprocessing that produce similar $E_{1\mathrm{D}}$ estimates [18].

2.3.2 Divergence and Vorticity

For an arbitrary 2D vector field **F** in a standard (x,y) Cartesian plane, the divergence of **F** is found by

$$\mathrm{div}(\mathbf{F}) = \frac{\partial F_{\mathrm{x}}}{\partial \mathrm{x}} + \frac{\partial F_{\mathrm{y}}}{\partial \mathrm{y}}, \tag{2.6}$$

where F_{x} and F_{y} are the x- and y-directional components of **F**. Likewise, the vorticity of **F** is found by

$$\mathrm{vor}(\mathbf{F}) = \frac{\partial F_{\mathrm{y}}}{\partial \mathrm{x}} - \frac{\partial F_{\mathrm{x}}}{\partial \mathrm{y}}. \tag{2.7}$$

For a wind field, the divergence can be interpreted as how much a given point acts as a source (or sink, if the divergence is negative) of wind flow. Likewise, vorticity can be interpreted as how much each point locally rotates while traveling along with the flow. Both quantities can be estimated for wind fields through a first order difference equation of neighboring WVCs.

---

[1] Freilich and Chelton define both a 2D kinetic energy spectrum and an isotropic 2D kinetic energy spectrum and show that the α value associated with the 2D isotropic kinetic energy spectrum is equal to $\alpha_{1\mathrm{D}}$ [15].

On synoptic scales (>500 km), wind fields are expected to have, on average, zero divergence and non-zero vorticity. But at the mesoscale (50–500 km), wind fields commonly contain areas of both non-zero divergence and non-zero vorticity. The divergence and vorticity of mesoscale wind fields is commonly described through 1D spectral analysis. For the vorticity spectrum, several models predict that a k^{-1} power law exists [19]–[21]. Fewer models exist to predict the divergence spectrum and they do not always agree one with another: one model based on 3D isotropic turbulence predicts a $k^{+1/3}$ power law [21] while another model based on boundary layer dynamics predicts that divergence and vorticity are proportional [16]. Studies that measure the 1D spectrum of both the divergence and vorticity tend to indicate that divergence and vorticity are proportional as predicted by the model based on boundary layer dynamics [16], [17], but the divergence and vorticity spectra shape for a particular wind field does not always cleanly follow the predicted power laws and changes based on WVC size [17].

The different WVC resolutions likely change the spectral behavior because the first order difference equations used to estimate the divergence and vorticity derivatives are taken at distances that change along with the WVC resolutions. Additionally, derivatives act as high pass filters and are sensitive to noise. As high resolution wind fields are noisier than low resolution wind fields the high resolution wind fields have greater divergence and vorticity measurements.

2.4 Scatterometer Simulation

Simulation is a common tool when developing and evaluating scatterometer wind estimation algorithms [22]–[26]. While simulations do not perfectly represent reality, they provide control over all parts of scatterometer wind estimation; many aspects of wind estimation can only be studied through simulation.

The simulation used in this thesis begins by pulling geometry information from a real QuikSCAT orbit. Among other parameters extracted from a QuikSCAT L1B file, this information includes the orbit position, antenna rotation angle, center location, azimuth angle, incidence angle, K_{pA}, K_{pB}, K_{pC}, and SNR of each σ^0 measurement. This information is used when calculating the SRF, K_{p} value, and σ^0 value of each egg and slice measurement.

The simulation requires a true wind field as a second input. This wind field represents the actual wind on the surface of the earth and needs to provide a wind vector at each SRF pixel for all egg and slice measurements pulled from the QuikSCAT L1B file. For each egg and slice measurement, the wind vector at each pixel within the measurement's SRF is fed into the GMF to find the corresponding σ^0 value for that pixel. By taking a weighted average of the σ^0 values, where the weights come from the SRF of the egg or slice, the simulated egg or slice σ^0 measurement is computed. With the supplied K_{pA}, K_{pB}, K_{pC}, and SNR, K_{p} for the egg and slice measurements is evaluated. Monte Carlo noise is added to a given noise-free σ^0 measurement (s_{i}) by

$$s_{\mathrm{i,noisy}} = s_{\mathrm{i}}(1 + v_{\mathrm{i}} K_{\mathrm{p,i}}) \tag{2.8}$$

where $s_{\mathrm{i,noisy}}$ is the noisy measurement, v_{i} is a realization of a zero-mean, unit-variance Gaussian random variable, and $K_{\mathrm{p,i}}$ is the K_p value for the σ^0 measurement.

With the simulated noisy σ^0 measurements and the measurement geometry data pulled from the QuikSCAT L1B file, all the quantities needed to run the normal wind retrieval process are available. After going through the wind retrieval process, the final step of the simulation is to put the retrieved winds through an ambiguity selection algorithm. The resulting wind is compared to the true wind to evaluate the estimated wind error.

Various aspects of scatterometer-measured winds can be studied using the simulation by varying the input wind field and QuikSCAT L1B file, or by adjusting aspects of the wind retrieval process and ambiguity selection algorithm. The advantage of simulation is that insight can be gained from the resulting wind fields as the true wind field is known. This means that changes to scatterometer wind estimation can be quantitatively compared and winds that are rare can be put through the process. Simulation is thus an essential tool in the scatterometer design process.

2.5 Synthetic Wind Fields

When using simulation to validate scatterometer wind vector measurements and algorithms, researchers commonly tailor the initial wind field to their specific study. For testing many scatterometer algorithms, it is common to employ "compass simulation," simulations that set each wind vector in the synthetic wind field to a constant speed and direction [23], [24]. Compass simulation is typically performed with enough synthetic wind fields to represent all compass directions and sometimes varying speeds [23]–[26]. This approach works well enough to evaluate the performance of scatterometer wind retrieval algorithms because it trivializes the complexity of ambiguity selection algorithms, allowing the effects of the wind retrieval algorithms to be more apparent. For the same reason, compass simulations are not effective in evaluating the performance of ambiguity selection algorithms. These algorithms require a more realistic wind field with variety in the wind speed and direction.

One source for more realistic synthetic wind fields is wind field models based on physical equations of motion. These winds are typically made for specific regions or wind features (e.g., ocean winds near Hawaii [27] or hurricane-based winds [28]). These models can be good representations of measured winds in the local regions or features that they are designed for, but they do not generalize to other areas or features very well.

Rather than solving equations of motion, other synthetic wind field creation processes are based on replicating wind field characteristics. This replication is done by artificially creating a wind field with a generated kinetic energy spectrum that follows the known power law of ocean winds. The magnitude component of the associated 2D Fourier transform is computed and the inverse Fourier transform of the magnitude component with a random phase component is taken [22]. These wind fields are not very realistic in the density and spatial resolution of their wind features, but they are relatively easy to create. The synthetic wind field creation process described in this thesis improves this method to create more realistic wind fields.

CHAPTER 3. MEASURING THE KINETIC ENERGY WAVENUMBER SPECTRUM OF OCEAN WINDS

The kinetic energy wavenumber spectrum of ocean winds is a common way of characterizing the sea state in meteorological studies. Freilich and Chelton published kinetic energy spectra for scatterometer-measured ocean winds, and many subsequent researchers that measure ocean wind kinetic energy spectra from scatterometers follow a similar method [15]. When Freilich and Chelton published their original work, they used wind vectors measured by the Seasat-A Satellite Scatterometer (SASS). The wind vector measurements provided by SASS were coarse in resolution and only 1D spectra were estimated. Wind measurements from other sources such as instruments onboard airplanes have also commonly been used to perform 1D spectral analysis [29], [30]. Freilich and Chelton showed that properties of 2D spectra can be inferred from 1D spectra if large-scale atmospheric motions are assumed to be two-dimensional, nondivergent, and isotropic [15]. For these reasons, most subsequent studies on ocean wind spectra have only measured wind spectra in 1D (e.g., [16], [17]). For certain applications, such as creating synthetic wind fields, it may be useful to measure ocean wind spectra in 2D rather than 1D. The purpose of the work presented in this chapter is to develop a method that allows wind spectra to be estimated in 2D and to analyze wind spectra in 1D and 2D for winds measured by QuikSCAT to validate the assumptions made by Freilich and Chelton.

This chapter is organized as follows: Section 3.1 gives additional background on Freilich and Chelton's work along with subsequent studies about the ocean wind kinetic energy spectra. Section 3.2 describes the QuikSCAT wind measurements used in this study. Section 3.3 presents the method to measure the ocean wind kinetic energy spectra in 1D and 2D. Section 3.4 shows and analyzes the ocean wind kinetic energy spectra in 1D and 2D. Section 3.5 provides a brief conclusion.

3.1 Ocean Kinetic Energy Spectra in 1D and 2D

To compare ocean kinetic energy spectra between 1D and 2D, Freilich and Chelton express the 2D kinetic energy spectrum of ocean winds as $E_{2\mathrm{D}}(k_x, k_y)$ with k_x and k_y representing the wavenumber (k) in the x- and y-directions, respectively. Their assumption that the wind field is caused by isotropic turbulence implies that $E_{2\mathrm{D}}(k_x, k_y)$ is rotationally invariant. They arrived at this result by expressing $E_{2\mathrm{D}}(k_x, k_y)$ in terms of the polar coordinates

$$\begin{aligned} k &= \sqrt{k_x^2 + k_y^2} \\ \theta &= \tan^{-1}(k_y/k_x), \end{aligned} \tag{3.1}$$

integrating over θ, and finding that

$$\hat{E}_{2\mathrm{D}}(k) = \pi k E_{2\mathrm{D}}(k_x, k_y) \tag{3.2}$$

where $\hat{E}_{2\mathrm{D}}(k)$ is the 2D isotropic kinetic energy spectrum [15].

Many turbulence models predict that the 2D isotropic kinetic energy spectrum follows a power law $\hat{E}_{2\mathrm{D}} = \beta k^{\alpha}$ where α is the slope of the power law and β is often ignored in ocean wind spectral analysis [15], [16] . Freilich and Chelton show that if $\hat{E}_{2\mathrm{D}} = \beta k^{\alpha}$, then $E_{1\mathrm{D}}(k) \propto k^{\alpha}$ where $E_{1\mathrm{D}}(k)$ is the 1D kinetic energy spectrum of the wind field [15]. In other words, they demonstrated that 1D kinetic energy spectra provides information about 2D kinetic energy spectra for ocean winds.

Average α values from several studies of scatterometer-measured winds are given in Table 3.1. Some of these studies measured the kinetic energy spectra for the meridional and zonal wind components separately and found only small differences in the α values between the two. Each study also compared how α changed based on location and/or temporal season, but only the general α value is given in Table 3.1. Note that the α values fall between $-5/3$ and -3, the range over which most models predict that α should be.

Table 3.1: Average α values for different scatterometers and data products. Blodgett did not publish an exact α for winds measured by the Advanced Scatterometer (ASCAT) but showed that it was very similar to $-5/3$ [17]. Other studies that have computed the kinetic energy spectra for various scatterometers and types of ground processing without publishing an exact α value include [16], [18], [31], [32].

Study	Scatterometer	Wind Vector Cell Size	Average α
Freilich and Chelton [15]	SASS	100 km by 100 km	-2
Patoux and Brown [16]	QuikSCAT	25 km by 25 km	-2.2
Patoux and Brown [16]	QuikSCAT DIRTH	25 km by 25 km	-2.4
Blodgett [17]	ASCAT L2B	12.5 km by 12.5 km	-1.7
Blodgett [17]	ASCAT UHR	1.25 km by 1.25 km	-1.7

3.2 QuikSCAT Wind Estimates

This study uses L2B wind fields measured by the QuikSCAT scatterometer in the arbitrarily chosen year 2005. The QuikSCAT measurement swath is nominally 1800 km wide (in the cross track direction) which allows for a large 2D area of study. This L2B data product performs wind estimation with DIRTH to improve the accuracy of the estimated winds [13]. As wind vector estimates close to the edge of the measurement swath are generally worse estimates of the actual wind vector (even with DIRTH), we discard several near-edge wind vectors in the measurement regions. The measurement regions are 131 WVCs $\times$ 131 WVCs (1637.5 km $\times$ 1637.5 km). Every WVC has to have a valid wind vector measurement for the region to be included in the study. The 2D analysis uses the entire measurement region as a single block while the 1D analysis uses individual along track rows and columns of WVCs (1 WVC $\times$ 131 WVCs or 131 WVCs $\times$ 1 WVC).

As mentioned in Section 3.1, most studies of ocean wind spectra compare how α changes based on location and/or temporal season. In this study, we make a global estimate of the kinetic energy spectra using estimates from the world ocean, similar to Patoux and Brown's general α value, and we make estimates of the kinetic energy spectra using four regions of the Pacific Ocean that match the four regions that Freilich and Chelton used [15], [16]. A description of these regions is given in Table 3.2. No differentiation between seasons is made in this study.

Table 3.2: Description of the four geographic regions in which the kinetic energy spectra are estimated. Latitudes are in degrees north, longitudes are in degrees east. These regions are chosen to match the regions used by Freilich and Chelton and their study contains a similar table [15].

Region	Latitude (north)	Longitude (east)
I	-45°to -25°	160°to 280°
II	-25°to -5°	160°to 280°
III	5°to 25°	140°to 250°
IV	25°to 45°	150°to 230°

3.3 Measuring Ocean Wind Kinetic Energy Spectra

Given the relationship between the 2D kinetic energy spectrum, the 2D isotropic kinetic energy spectrum, and the 1D kinetic energy spectrum in Section 3.1, we express the 1D kinetic energy spectrum and 2D kinetic energy spectrum without arguments as

$$E_{\mathrm{1D}} = \beta_{\mathrm{1D}} k^{\alpha_{\mathrm{1D}}} \tag{3.3}$$

and

$$E_{\mathrm{2D}} = \beta_{\mathrm{2D}} k^{\alpha_{\mathrm{2D}}} \tag{3.4}$$

for notational convenience[1]. Comparing E_{2D}, rather than $\hat{E}_{\mathrm{2D}}$, with E_{1D} avoids polar integration with discrete measurements on a square grid. The theory presented in Section 3.1 expects that $\alpha_{\mathrm{2D}} = \alpha_{\mathrm{1D}} - 1$.

E_{1D} is typically found separately for both the zonal (u) and meridional (v) components of the wind field and then summed together if a total wind kinetic energy spectra is desired. However, we derive the same result by combining u and v into a complex wind vector w ($w = u + jv$) before calculating the kinetic energy spectra. The autocorrelation of w (R_w) can be expressed in terms of

---

[1] Eqs. 2.4 and 2.5 are repeated here for convenience.

the autocorrelation of u (R_u) and v (R_v) as follows:

$$\begin{aligned}
R_w[m] &= \mathbf{E}\{w[n+m]w^*[n]\} \\
&= \mathbf{E}\{(u[n+m]+jv[n+m])(u[n]+jv[n])^*\} \\
&= \mathbf{E}\{u[n+m]u^*[n]-ju[n+m]v^*[n]+jv[n+m]u^*[n]+v[n+m]v^*[n]\} \\
&= \mathbf{E}\{u[n+m]u^*[n]\}-j\mathbf{E}\{u[n+m]v^*[n]\}+ \\
&\qquad j\mathbf{E}\{v[n+m]u^*[n]\}+\mathbf{E}\{v[n+m]v^*[n]\} \\
&= R_u[m]-jR_{uv}[m]+jR_{vu}[m]+R_v[m].
\end{aligned} \tag{3.5}$$

On average, u and v are assumed to be uncorrelated so the cross-correlations are equal to 0 (i.e., $R_{uv}[m] = R_{vu}[m] = 0$). This yields

$$R_w[m] = R_u[m] + R_v[m]. \tag{3.6}$$

Taking the Fourier transform of R_w reveals that the kinetic energy spectrum of w (E_w) is the sum of the kinetic energy spectra of u (E_u) and v (E_v). This sum is used in other studies to find a total $E_{1\mathrm{D}}$ [15], [16]. Since all the studies referenced in Table 3.1 measure the α value associated with E_u (α_u) to be approximately equal to the α value associated with E_v (α_v), E_w can be expressed as

$$\begin{aligned}
E_w &= E_u + E_v \\
&= \beta_u k^{\alpha_u} + \beta_v k^{\alpha_v} \\
&\approx (\beta_u + \beta_v) k^{\alpha_w}
\end{aligned} \tag{3.7}$$

where $\alpha_w \approx \alpha_u \approx \alpha_v$. In this study, we validate the above analysis and provide additional validation of the assumption that u and v are uncorrelated by estimating $E_{1\mathrm{D}}$ and $E_{2\mathrm{D}}$ with E_u, E_v, $E_u + E_v$, and E_w.

3.3.1 Mitigating the Effect of Spectral Leakage

To mitigate the effects of spectral leakage when calculating the kinetic energy spectra, we subtract the mean from the wind field and window the data with a Tukey window with a 0.2

cosine fraction similar to [15], [16]. Other techniques to prevent spectral leakage include linearly transforming the data (choosing constants a and b in $z_i' = az_i + b$ such that the endpoints of z_i' map to zero) and high pass filtering the data with a first order difference equation ($z_i' = z_{i+1} - z_i$) which is then corrected for afterwards [18].

Conventional applications of a 2D window on discrete data involve applying a 1D window to each row of the data followed by applying the 1D window to each column of data. Creating a 2D window in such a manner results in a square-like 2D window rather than a circular 2D window. A circular 2D window is made by sampling a circularly-rotated 1D continuous window function. Examples of some 2D windows are given in Fig. 3.1. The square-like windows push spectral leakage primarily in orthogonal directions while the circular windows result in spectral leakage spaced more evenly around the exact frequency as seen in Fig. 3.2. When applied to wind fields, the square-like windows exhibit more signs of spectral leakage on average than the circular windows as shown in Fig. 3.3. Thus, in this study we subtract the mean and use a circular 2D Tukey window with a 0.2 cosine fraction when calculating the kinetic energy spectra of a 2D wind field.

3.3.2 Weighted Least Squares Fit to Ocean Wind Kinetic Energy Spectra

To estimate $\alpha_{1\mathrm{D}}$ and $\alpha_{2\mathrm{D}}$, we make an approximate least squares fit to the kinetic energy spectra of the measured data. This is done by transforming Eqs. 3.3 and 3.4 into linear equations by taking the $\log_{10}$ of both sides which yields

$$\log_{10} E_{1\mathrm{D}} = \alpha_{1\mathrm{D}} \log_{10} k + \log_{10} \beta_{1\mathrm{D}} \tag{3.8}$$

and

$$\log_{10} E_{2\mathrm{D}} = \alpha_{2\mathrm{D}} \log_{10} k + \log_{10} \beta_{2\mathrm{D}}. \tag{3.9}$$

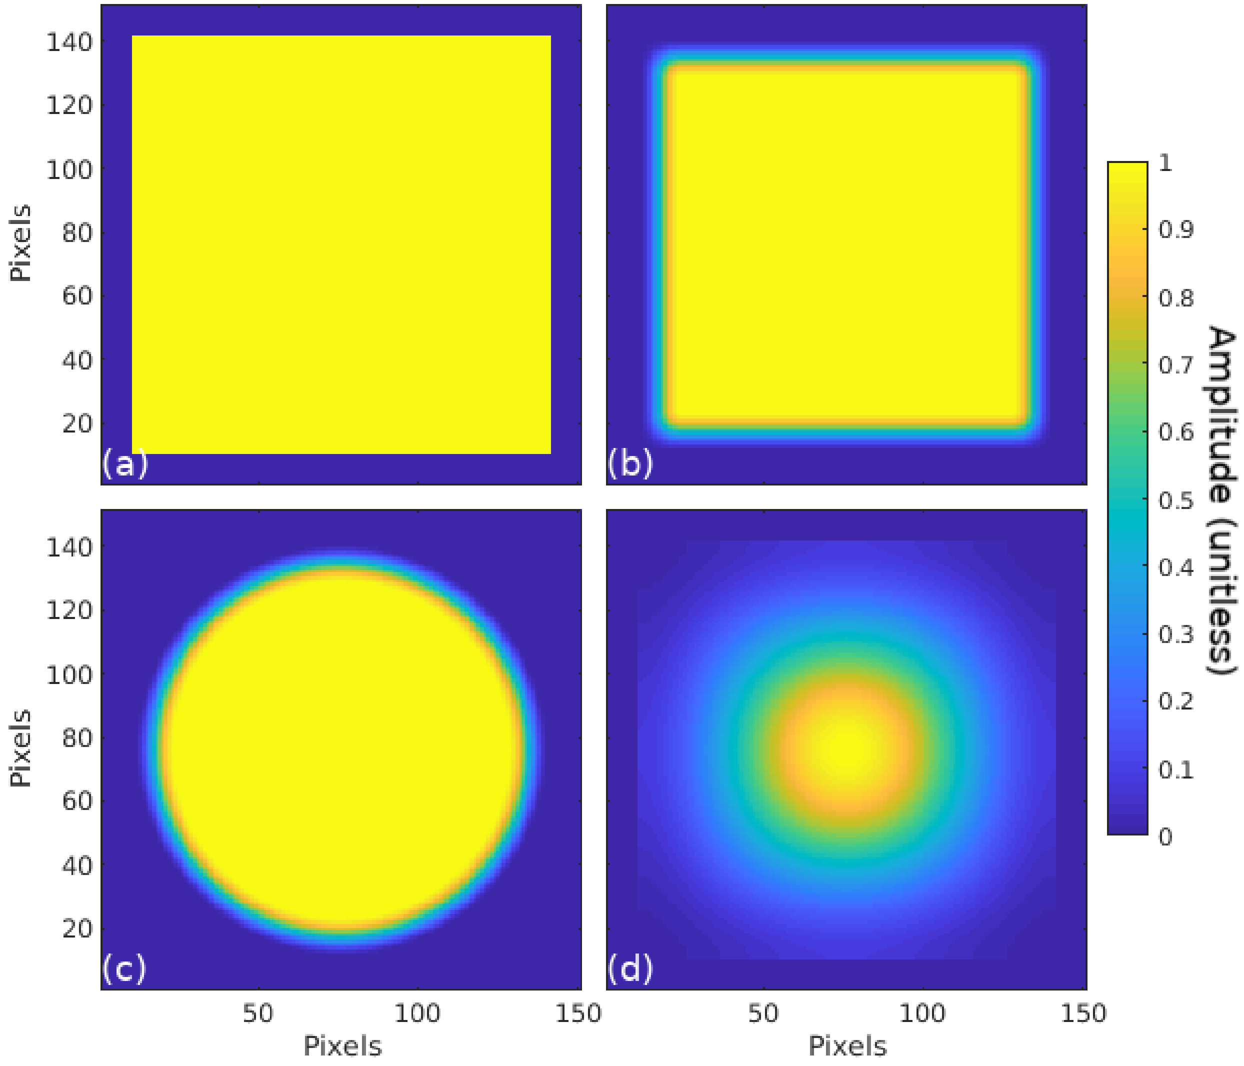

Figure 3.1: Examples of various 2D windows. (a) shows a rect window; (b) shows a square Tukey window with a 0.2 cosine fraction; (c) shows a circular Tukey window with a 0.2 cosine fraction; and (d) shows a square Hamming window. Each window is 131 pixels × 131 pixels. To emphasize the window shape, each window has been zero padded so that each panel is 151 pixels × 151 pixels.

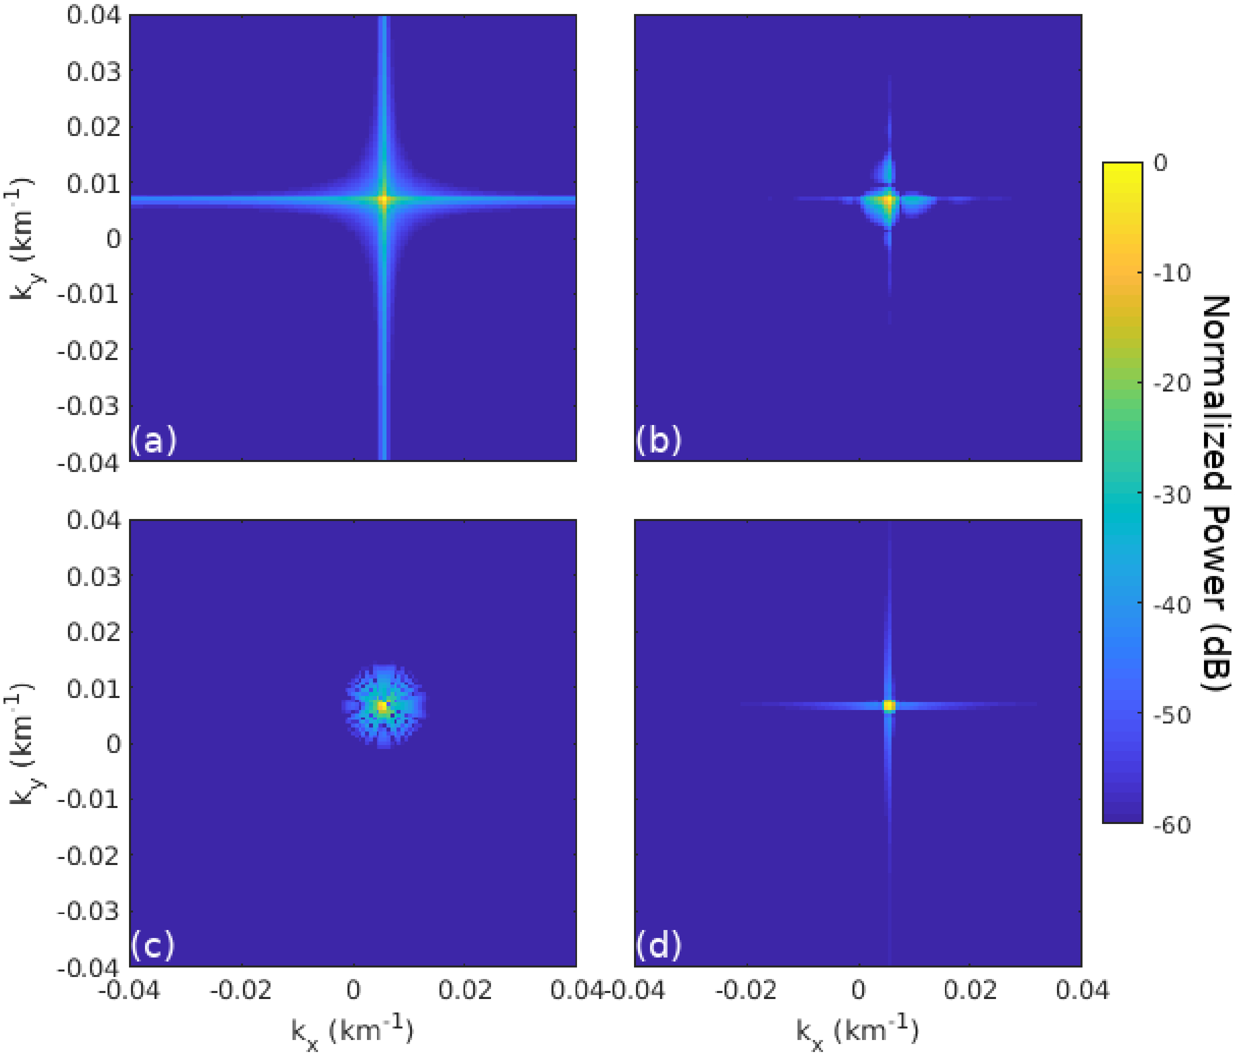

Figure 3.2: Examples of the spectral leakage permitted by (a) a 2D rect window, (b) a 2D square Tukey window with a 0.2 cosine fraction, (c) a 2D circular Tukey window with a 0.2 cosine fraction, and (d) a 2D square Hamming window. The leakage examples have been normalized so that the peak value is 1 (0 dB). The color scale is cut off at 60 dB below the peak value to make the shape of the spectral leakage clear. The square-like windows push spectral leakage in orthogonal directions while the circular Tukey window pushes spectral leakage more evenly in all directions.

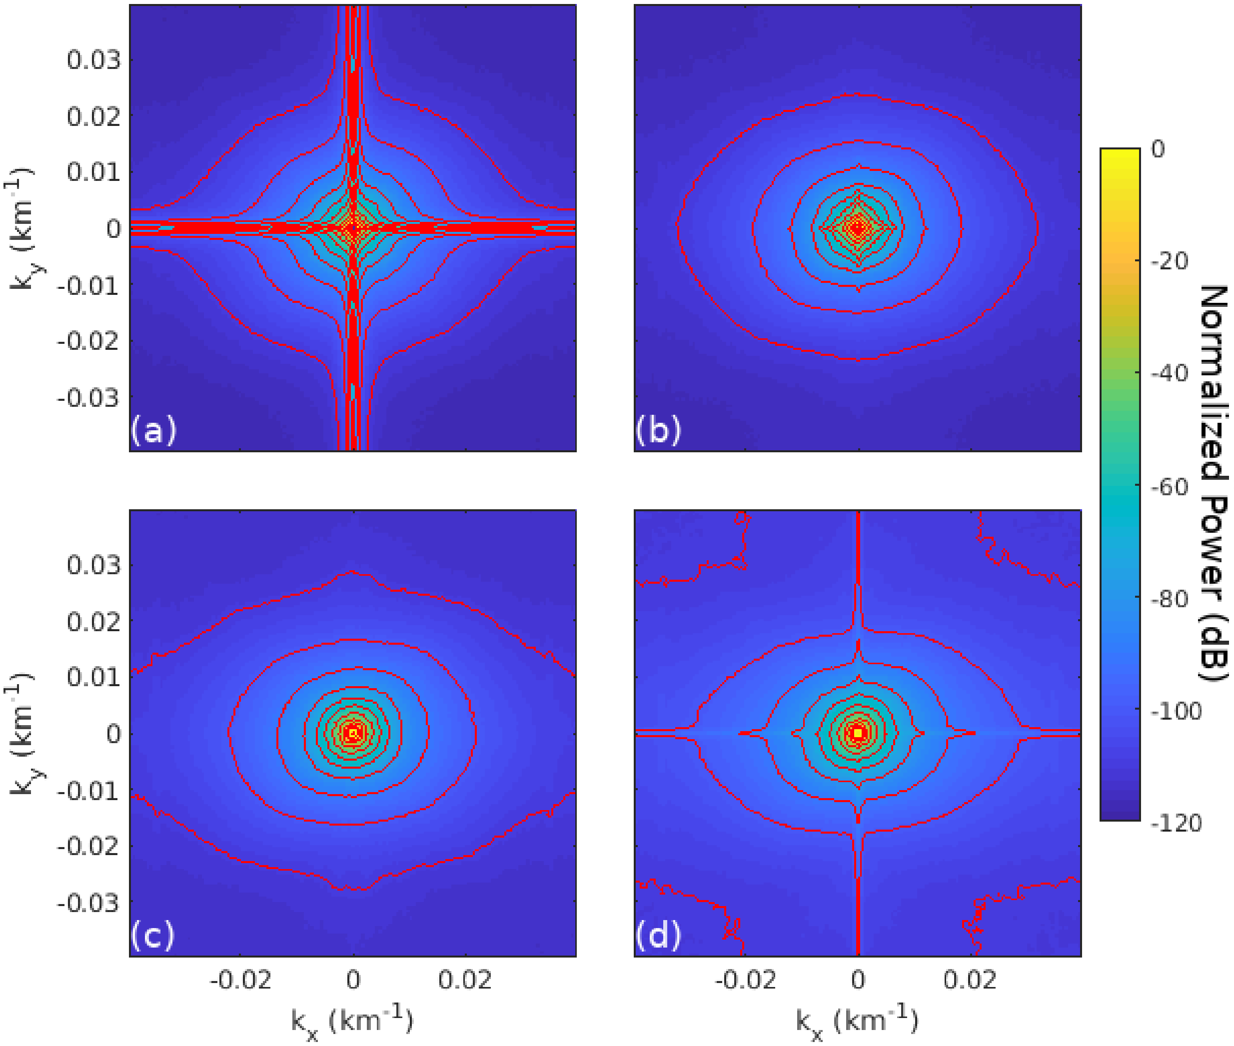

Figure 3.3: Average 2D kinetic energy spectra of QuikSCAT winds with (a) a 2D rect window, (b) a 2D square Tukey window with a 0.2 cosine fraction, (c) a 2D circular Tukey window with a 0.2 cosine fraction, and (d) a 2D square Hamming window. The average 2D kinetic energy spectra have been normalized so that the peak value is 1 (0 dB). The average 2D kinetic energy spectra are overlaid by red contour lines to emphasis the shape of the average kinetic energy spectra. The square-like windows reveal more noticeable spectral leakage in orthogonal directions away from the center of the spectra than the circular window.

For the 1D case with N measurements of E_{1D} and k, we define

$$\mathbf{b}_{\mathrm{1D}} = \begin{bmatrix} \log_{10} E_{\mathrm{1D}}^{1} \\ \log_{10} E_{\mathrm{1D}}^{2} \\ \vdots \\ \log_{10} E_{\mathrm{1D}}^{N} \end{bmatrix} \tag{3.10}$$

$$\mathrm{A}_{\mathrm{1D}} = \begin{bmatrix} 1 & \log_{10} k^{1} \\ 1 & \log_{10} k^{2} \\ \vdots & \vdots \\ 1 & \log_{10} k^{N} \end{bmatrix} \tag{3.11}$$

$$\mathbf{x}_{\mathrm{1D}} = \begin{bmatrix} \log_{10} \beta_{\mathrm{1D}} \\ \alpha_{\mathrm{1D}} \end{bmatrix}, \tag{3.12}$$

where E_{1D}^{i} and k^{i} are the ith measurement of E_{1D} and k with $i = 1, 2, ..., N$. With Eqs. 3.10, 3.11, and 3.12, Eq. 3.8 can be written in the familiar matrix form

$$\mathrm{A}_{\mathrm{1D}}\mathbf{x}_{\mathrm{1D}} = \mathbf{b}_{\mathrm{1D}}. \tag{3.13}$$

Following a similar process, Eq. 3.9 can be rewritten as

$$\mathrm{A}_{\mathrm{2D}}\mathbf{x}_{\mathrm{2D}} = \mathbf{b}_{\mathrm{2D}}. \tag{3.14}$$

By performing least squares with Eqs. 3.13 and 3.14 to estimate $\mathbf{x}_{\mathrm{1D}}$ and $\mathbf{x}_{\mathrm{2D}}$, the α_{1D} and α_{2D} values in $\mathbf{x}_{\mathrm{1D}}$ and $\mathbf{x}_{\mathrm{2D}}$ minimize the error for Eqs. 3.8 and 3.9. These α values only estimate the α_{1D} and α_{2D} values that would minimize the error for Eqs. 3.3 and 3.4. However, we expect that the estimated α_{1D} and α_{2D} are close enough to the desired α_{1D} and α_{2D} to make a fair comparison between the 1D and 2D spectra.

In practice, E_{1D} and E_{2D} appear to approach a spectrum floor at higher k values, a floor that disrupts the power law for their measurements. While the spectrum floor appears similar to a noise floor, it is unlikely that noise causes the spectrum floor because the spectrum floor begins at different k values depending on the resolution of the measured wind field [17]. Blodgett demonstrates an alternative source of the spectrum floor could be the effects of aliasing due to undersampling the wind field as seen in Fig. 3.4 [17]. Regardless of the source of the spectrum floor the least squares fit is weighted to limit effect of the spectrum floor on the estimated α_{1D} and α_{2D} values. For the 1D case, we create the diagonal weighting matrix

$$W_{\mathrm{1D}} = \begin{bmatrix} \frac{1}{k^1} & & & \\ & \frac{1}{k^2} & & \\ & & \ddots & \\ & & & \frac{1}{k^N} \end{bmatrix}. \tag{3.15}$$

There is an extra detail to consider with the 2D case. When plotting E_{2D}, such as in Fig. 3.3, frequency bins with equal k values can be viewed as rings around the center DC bin. For a discrete representation of the kinetic energy spectra on a square grid, this detail means that there are more frequency bins for high values of k than for low values of k. By a geometric argument, the number of frequency bins available for a given k value is roughly proportional to k, which this study accounts for in the weighting matrix by dividing the $1/k$ weights by k. This helps the α_{2D} estimates to be influenced by all values of k more equally and yields the diagonal weighting matrix

$$W_{\mathrm{2D}} = \begin{bmatrix} \frac{1}{(k^1)^2} & & & \\ & \frac{1}{(k^2)^2} & & \\ & & \ddots & \\ & & & \frac{1}{(k^N)^2} \end{bmatrix}. \tag{3.16}$$

By application of weighted least squares,

$$\mathbf{x}_{\mathrm{1D}} = (\mathrm{A}_{\mathrm{1D}}^{\mathrm{H}} W_{\mathrm{1D}} \mathrm{A}_{\mathrm{1D}})^{-1} \mathrm{A}_{\mathrm{1D}}^{\mathrm{H}} W_{\mathrm{1D}} \mathbf{b}_{\mathrm{1D}} \tag{3.17}$$

and

$$\mathbf{x}_{\mathrm{2D}} = (\mathrm{A}_{\mathrm{2D}}^{\mathrm{H}} W_{\mathrm{2D}} \mathrm{A}_{\mathrm{2D}})^{-1} \mathrm{A}_{\mathrm{2D}}^{\mathrm{H}} W_{\mathrm{2D}} \mathbf{b}_{\mathrm{2D}} \tag{3.18}$$

from which α_{1D} and α_{2D} can be extracted as the second elements of $\mathbf{x}_{\mathrm{1D}}$ and $\mathbf{x}_{\mathrm{2D}}$, respectively.

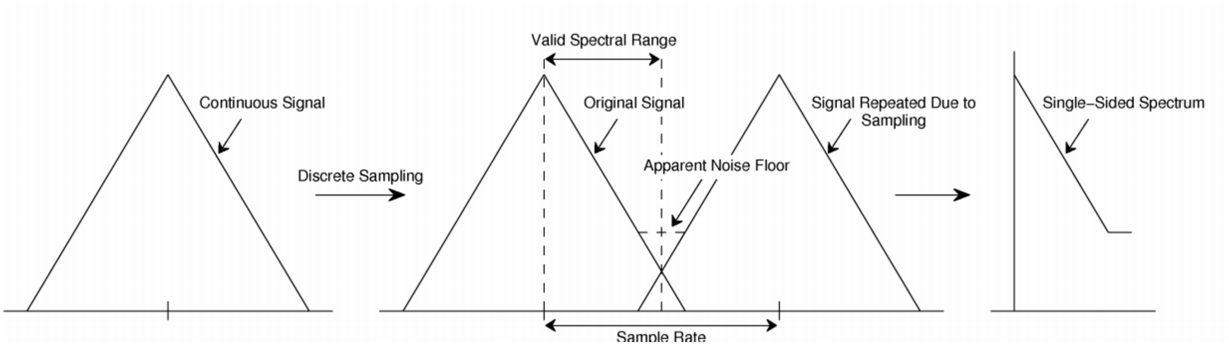

Figure 3.4: The effect of aliasing on the spectrum of an arbitrarily undersampled, decreasing signal. Figure from [17].

Table 3.3: Various α values found for different regions, wind components, and dimensions in this study. In general, $\alpha_u < \alpha_w < \alpha_v$, though they are all similar. The trends in the α values in this study match the trends in the α values reported by Freilich and Chelton [15]. Note that the values given from Freilich and Chelton are from their unnormalized power law calculations [15]. The data support that $\alpha_{2\mathrm{D}} \approx \alpha_{1\mathrm{D}} - 1$.

	$\alpha_{1\mathrm{D}}$				$\alpha_{2\mathrm{D}}$			
Region	$E_{\mathbf{u}}$	$E_{\mathbf{v}}$	$E_{\mathbf{u}} + E_{\mathbf{v}}$	$E_{\mathbf{w}}$	$E_{\mathbf{u}}$	$E_{\mathbf{v}}$	$E_{\mathbf{u}} + E_{\mathbf{v}}$	$E_{\mathbf{w}}$
Region 1	-2.39	-2.29	-2.33	-2.33	-3.34	-3.25	-3.30	-3.30
Region 2	-2.25	-1.88	-2.02	-2.02	-3.17	-2.77	-2.92	-2.92
Region 3	-2.23	-1.88	-2.01	-2.01	-3.18	-2.82	-2.96	-2.96
Region 4	-2.38	-2.23	-2.29	-2.29	-3.32	-3.18	-3.25	-3.25
Global	-2.44	-2.23	-2.32	-2.32	-3.40	-3.19	-3.29	-3.29
From Freilich and Chelton [15]								
Region 1	-2.35	-2.17	-2.26					
Region 2	-1.99	-1.77	-1.86					
Region 3	-2.01	-1.95	-1.97					
Region 4	-2.31	-2.11	-2.21					

3.4 Wind Spectra Results and Analysis

In this study, wind fields measured by QuikSCAT that fit the criteria given in Section 3.2 undergo the windowing and weighted least squares approximation explained in Section 3.3. $\alpha_{1\mathrm{D}}$ and $\alpha_{2\mathrm{D}}$ are estimated for each geographic region of interest and for each E_u, E_v, $E_u + E_v$, and E_w representation of the wind field spectrum. A summary of the estimated α values is given in Table 3.3, along with values from Freilich and Chelton's analysis where applicable.

The estimated α values from regions 1 and 4 are similar while the estimated α values from regions 2 and 3 are similar, as seen in Fig. 3.5. This indicates that winds equidistance from the

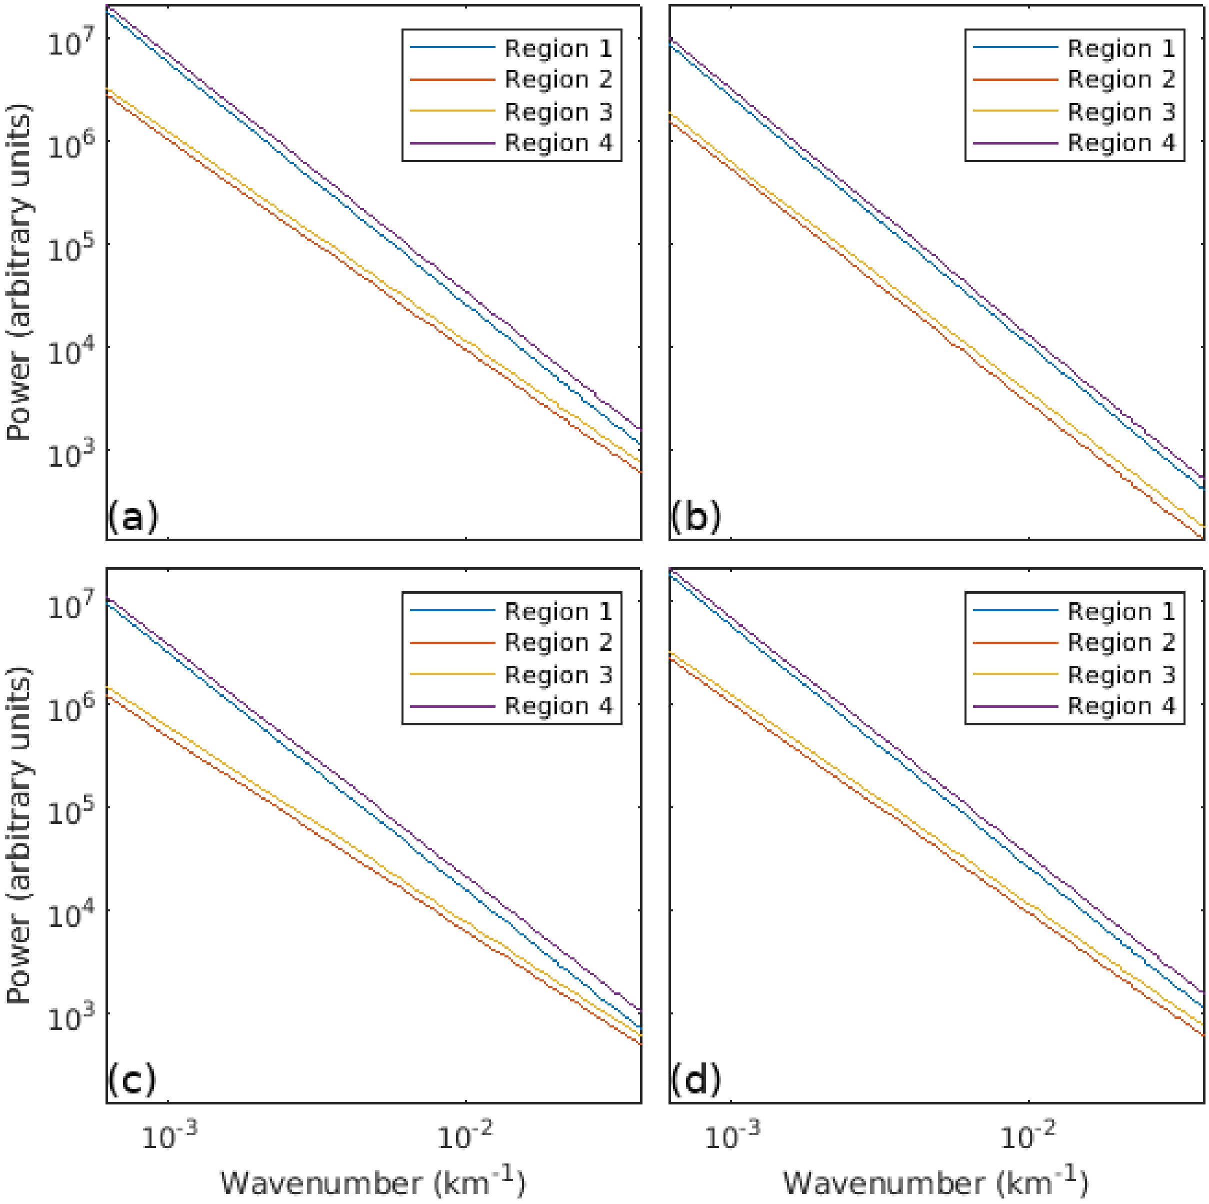

Figure 3.5: Weighted least squares fit to E_{1D} for regions 1, 2, 3, and 4 from QuikSCAT L2B data. (a) shows E_w; (b) shows E_u; (c) shows E_v; and (d) shows $E_u + E_v$. Note that regions 1 and 4 are similar while regions 2 and 3 are similar, indicating that winds equidistance from the equator appear to have similar characteristics.

equator share similar characteristics. Additionally, Fig. 3.6 shows that $\alpha_u < \alpha_w < \alpha_v$ in all regions, though the separation is relatively small. The α_{1D} values from this study are slightly greater in magnitude than those found by Freilich and Chelton, but the Freilich and Chelton α values follow the same general trends that we observe in our estimates. As the global α_{1D} is similar to the α value found by Patoux and Brown in their spectral analysis of DIRTH QuikSCAT winds (see Table 3.1), the difference between our estimates and Freilich and Chelton's estimates may be a difference between winds estimated by QuikSCAT with winds estimated by SASS. These cross-study comparisons help validate the method for estimating α_{1D}. To extend this verification to the method for estimating α_{2D}, we present plots of the weighted least squares fits to the data in Figs. 3.7 and 3.8. The fits match the data well in both 1D and 2D according to the weights that are assigned to the measured data. From these results we assume that the method for estimating α_{2D} is valid.

The estimated α values from $E_u + E_v$ are equal to the estimated α values from E_w to the given precision of three significant figures. This finding supports the assumption that u and v are uncorrelated and provides a measured example of the theoretical result in Section 3.3 that the total kinetic energy spectra of a wind field can be found by either combining the u and v wind field components into a complex vector w prior to estimating the kinetic energy spectra or by summing the kinetic energy spectra associated with the u and v wind field components. Such a result can simplify the complexity of analyzing the u and v wind field components simultaneously.

When comparing α_{1D} and α_{2D} across all regions and globally, we find that $\alpha_{\mathrm{2D}} \approx \alpha_{\mathrm{1D}} - 1$. From this result, we support the assumptions that Freilich and Chelton made in their original paper that allow properties of 2D wind spectra to be inferred from 1D wind spectra (i.e., large-scale atmospheric motions are 2D, nondivergent, and isotropic).

3.5 Conclusion

This study makes the following contributions: (1) a method to estimate ocean wind spectra in 2D that accounts for 2D spectral leakage and the increased high to low frequency bin ratio in 2D analysis, (2) additional validation of the assumptions made by many other researchers that large-scale atmospheric motions are 2D, nondivergent, and isotropic indicating that properties of 2D wind spectra can be inferred from 1D wind spectra, (3) a demonstration that 1D and 2D ocean

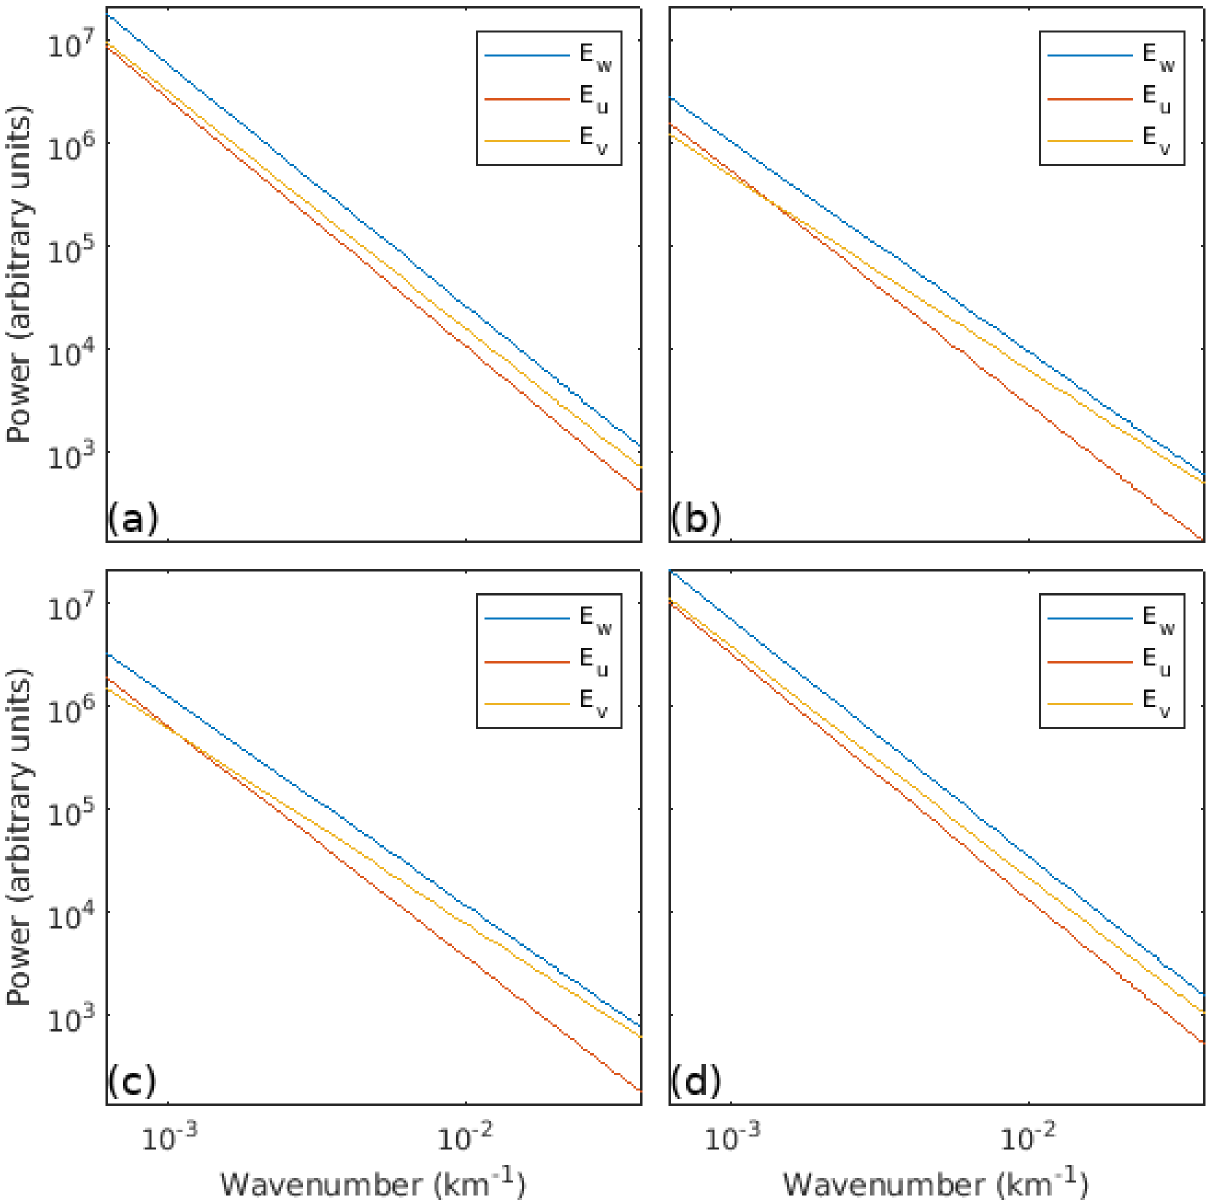

Figure 3.6: Weighted least squares fit to E_{1D} for u, v, and w from QuikSCAT L2B data. (a) shows region 1; (b) shows region 2; (c) shows region 3; and (d) shows region 4. Note that $\alpha_{\mathbf{u}} < \alpha_{\mathbf{w}} < \alpha_{\mathbf{v}}$ in all regions, though the relationship is most obvious in regions 2 and 3.

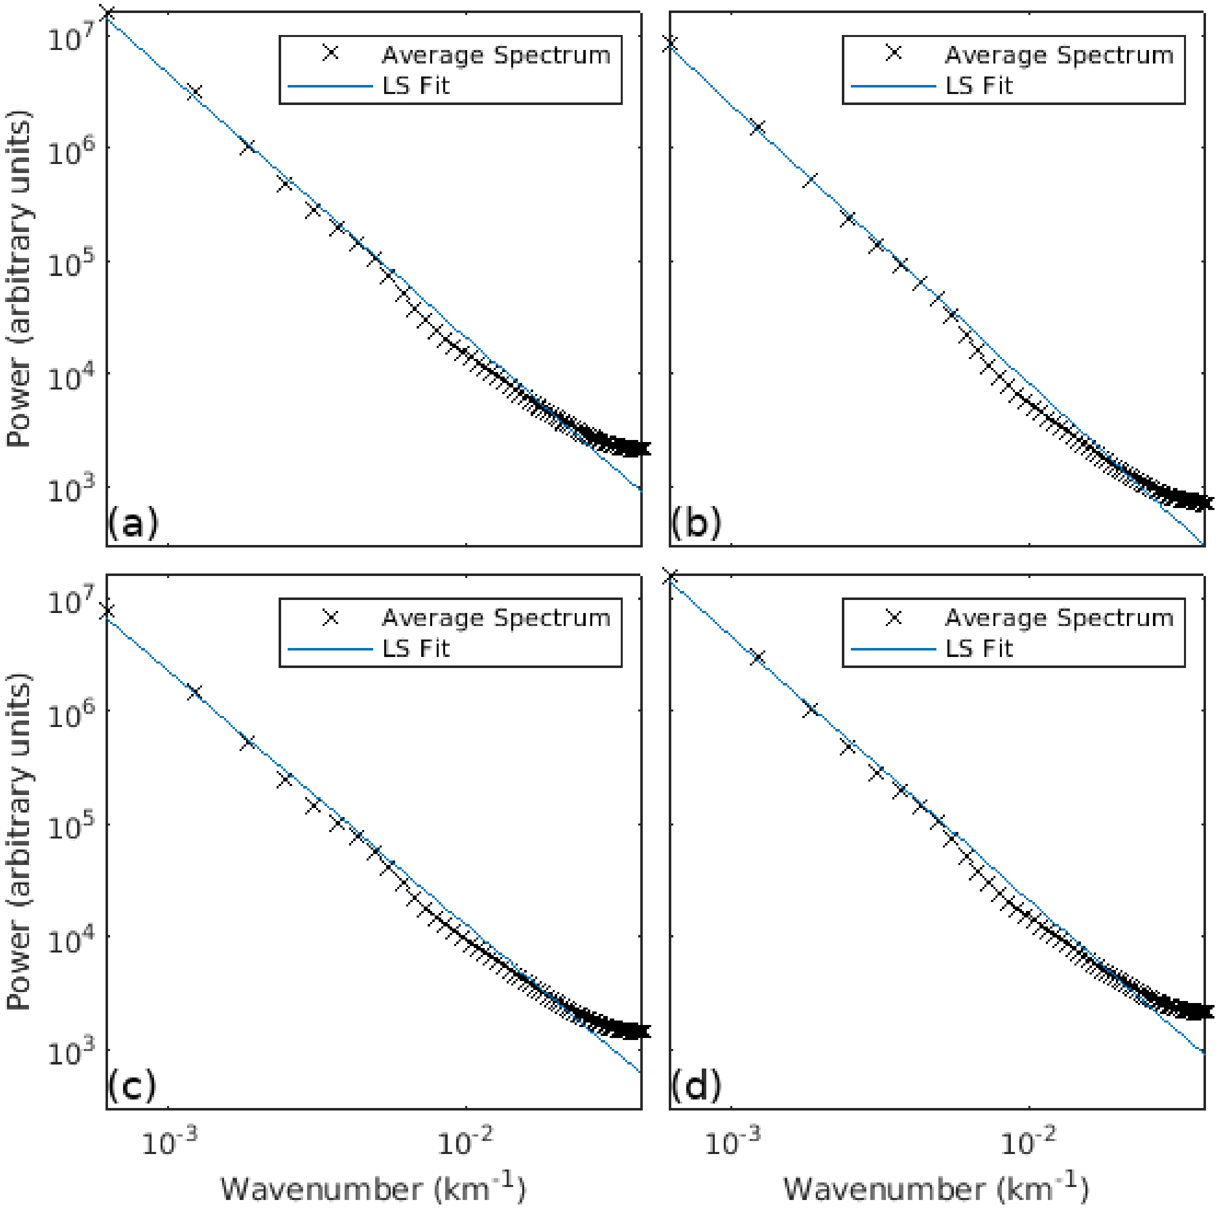

Figure 3.7: Weighted least squares fit to E_{1D} from the world ocean QuikSCAT L2B data. (a) shows region 1; (b) shows region 2; (c) shows region 3; and (d) shows region 4. Note that the measured data fits well except for high values of k where the data appears to be subject to a noise floor or aliasing.

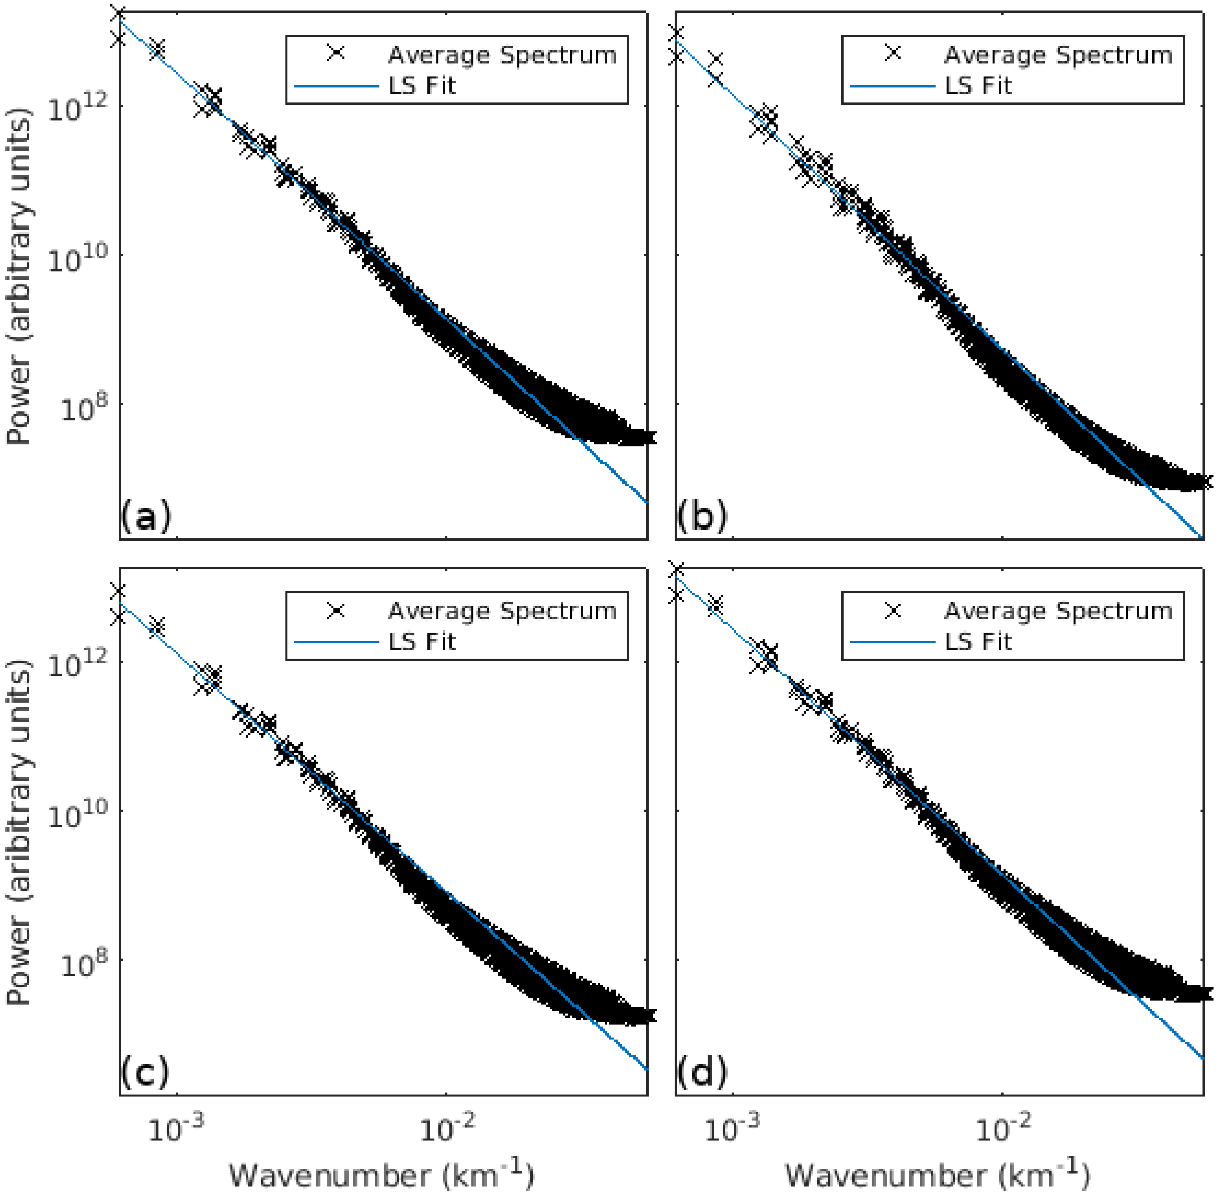

Figure 3.8: Weighted least squares fit to E_{2D} from the world ocean QuikSCAT L2B data. (a) shows region 1; (b) shows region 2; (c) shows region 3; and (d) shows region 4. Note that the measured data fits well except for high values of k where the data appears to subject to a noise floor or aliasing.

wind kinetic energy spectra can be computed from complex wind vectors formed from the u and v wind components and give further verification to the assumption that u and v are uncorrelated. These contributions may be useful for a variety of studies relating to ocean winds that need to examine the u and v components together or that must be done in 2D rather than 1D, such as the creation of a synthetic wind field.

CHAPTER 4. DEVELOPING THE WIND RESPONSE FUNCTION THROUGH SIMULATION

The σ^0 measurement SRF is a useful tool in scatterometer algorithm design that describes the spatial resolution of the σ^0 measurements for a given processing technique [33]. Enhanced resolution scatterometer data products rely on the SRF of egg and slice measurements to estimate σ^0 at a finer resolution than is accomplished through conventional DIB processing [10], [34], [35]. This chapter extends the idea of the σ^0 SRF to a wind response function (WRF). Conceptually, the WRF can be thought of as a wind-specific version of the SRF; while the SRF quantifies the contribution of each point on the surface to a σ^0 measurement, the WRF quantifies the contribution of the wind at each point on the surface to a scatterometer-determined wind vector. In other words, the WRF is similar to a wind impulse response function for a given processing algorithm.

There are several approaches and algorithms for scatterometer wind retrieval. Herein, the WRF for QuikSCAT DIB egg processing, QuikSCAT DIB slice processing, and QuikSCAT UHR processing are estimated through the use of simulation. While the WRF differs slightly for each WVC, we find generalized WRFs for each processing method. By computing the generalized WRFs, the effective spatial resolution of each of these three processing techniques is found.

The SRF is a scalar function; however, winds are a vector with a speed and a direction. This characteristic makes the WRF a vector function. For convenience, we define the WRF for the speed and direction components of a wind vector separately. The WRF could also be defined in terms of orthogonal wind components such as its zonal (u) and meridional (v) components. The derivation for the u and v components' defined WRF is similar to the speed and direction components' defined WRF and is not explicitly given in this chapter, though sample WRFs for u and v components are presented.

This chapter is organized as follows. Section 4.1 explains how the QuikSCAT WRF is calculated for DIB egg, DIB slice, and UHR processing. Section 4.2 shows the calculated WRF for several WVCs, finds generalized QuikSCAT WRFs for each processing method, and analyzes the resulting WRFs. Section 4.3 provides a brief conclusion.

4.1 Wind Response Function Calculation

The WRF represents the wind estimation process that transforms a "delta function-like" wind field to a scatterometer-measured wind vector; the WRF is found by passing a delta function wind approximation in different locations through the wind estimation process. By measuring how the output wind vector changes due to the position of the delta function approximation, the WRF helps understand the effective spatial resolution of wind estimation. This approach is similar to determining the impulse response of a linear time-invariant (LTI) system. In an LTI system, the impulse response completely characterizes the system. As wind estimation is nonlinear and spatially variant, it is not completely characterized by the WRF. However, the WRF defined below still provides useful information about the wind estimation process, especially the effective spatial resolution of wind estimation.

The quantity, position, orientation, and other characteristics of the σ^0 measurements used in calculating the wind vector for a WVC varies based on swath location. These variations cause the WRF to differ slightly with each WVC. For this reason, we estimate the WRF for several WVCs at various swath location. To define the WVCs, information about egg and slice measurements is pulled from a particular QuikSCAT L1B file. For DIB egg processing, a WVC is identified as a 25 km × 25 km region within the swath. The wind vector estimate for the DIB egg WVC uses the information associated with each egg measurement whose center lands within the WVC. For DIB slice processing, the WVC is 12.5 km × 12.5 km and the WVC wind vector estimate uses the information associated with each slice measurement whose center lands within the WVC. For UHR processing, where WVCs are defined somewhat differently, a WVC is a single UHR image pixel. UHR image pixels are 2.225 km × 2.225 km in this simulation. The information associated with each slice measurement whose SRF covers a given pixel is used for the wind vector estimate of the corresponding WVC [10]. Theoretically, each egg and slice SRF extends forever, but we truncate the SRFs by zeroing out any weight that is under a threshold of 20 dB below the peak value. For convenience, the geometry data from the QuikSCAT L1B file used in calculating the WRFs for all three processing techniques are interpolated to 2.225 km × 2.225 km resolution.

An overview of the method to find the WRF for a WVC is provided as follows. Several delta wind fields are created where each delta wind field contains the delta function wind approximation at a different location in the wind field. Using the measurement geometry and Monte Carlo techniques, each delta wind field is converted to noisy σ^0 measurements that are fed through the wind retrieval process and an ideal ambiguity selection algorithm to estimate the resulting wind vector. Biases and noise in the output wind vectors are reduced through averaging and subtracting out the output wind vector that occurs when no delta function passes through the wind estimation process. The WRF is then the map of the output wind vectors as a function of the location of the delta function in the delta wind field. For convenience, quantized WRFs for each processing algorithm are found at 2.225 km × 2.225 km resolution. The rest of this section details this method further.

4.1.1 Delta Wind Fields

First, a set of delta wind fields is created to approximate the response function. In a true "delta function," the delta function would have an infinite speed wind vector at one point and zero speed wind vectors for the rest of the wind field. In practice, the delta wind fields are discrete with one wind vector per UHR pixel and the delta wind vector has a high but finite speed while the other pixels have non-zero speeds. The reason for the latter is due to the nature of wind retrieval and the GMF. The GMF is only defined for positive wind speeds up to 50 m/s, and wind retrieval at very low and very high wind speeds is often unreliable. Taking this into account, each delta wind field has a delta wind vector with a speed of 40 m/s and the remaining background wind field has a speed of 7 m/s. Note that as long as there is a significant difference in the wind speed, other values besides 7 m/s and 40 m/s produce the same general WRF. A corresponding wind field with no delta wind vector (but the same background wind field) is also created. This no-delta wind field is subtracted from the WRF estimate to help reduce bias introduced by the non-ideal delta function winds. The direction of each wind vector, which is referred to as the common direction, is the same for all wind vectors in a delta wind field.

To illustrate the concept of a delta wind field, an example WVC for DIB egg processing is overlaid with an arbitrarily chosen delta wind field and is shown in Fig. 4.1. This particular delta wind field is chosen because the delta wind vector is clearly visible from other markings on the figure. While Fig. 4.1 does not explicitly show the full SRFs of the egg measurements, the delta wind field is the smallest rectangular wind field that covers the entire truncated SRF of the egg measurements.

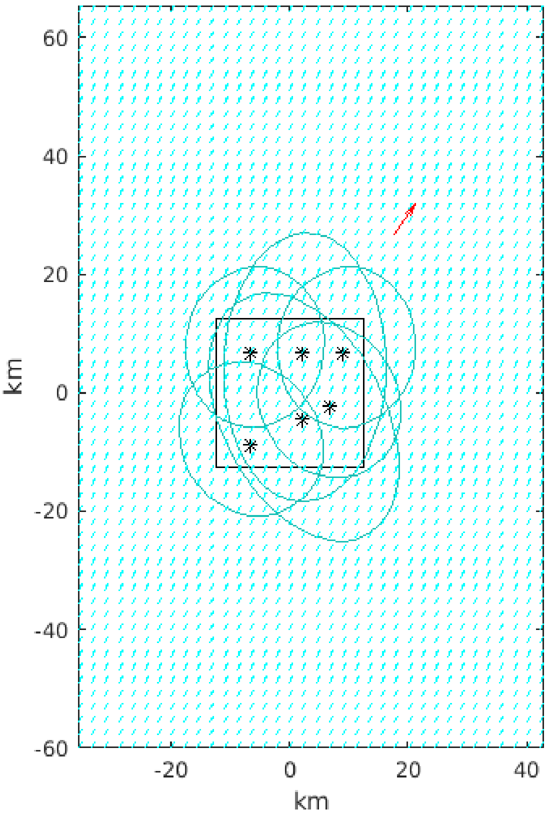

Figure 4.1: Example WVC for DIB egg processing with 3 dB egg contours overlaying a delta wind field. The square outlines the area of the WVC, the asterisks show the center of each egg measurement used in wind retrieval for this WVC, and the ovals show the eggs' 3 dB footprints. The egg measurement SRFs extend to the edge of the delta wind field, though the weights are significantly smaller outside the eggs' 3 dB footprints. The delta function wind vector is shown in red to differentiate it from the background wind vectors.

4.1.2 Wind Estimation

The next step in approximating the response function is to convert the delta wind fields into σ^0 measurements. We use the QMOD3 GMF for QuikSCAT to perform this conversion. Since the operating frequency of QuikSCAT is known, and the incidence angle indicates the polarization of the incident radiation, the QMOD3 GMF may be defined without loss of generality as[1]

$$\sigma^0 = g(|U|, \chi, \theta) \tag{4.1}$$

where $|U|$ is the wind speed, χ is the angle between the wind direction and azimuth angle of the incident radiation, and θ is the incidence angle. Using the azimuth and incidence angles for each egg and slice measurement, the true σ^0 values for each delta wind field are calculated for each egg and slice measurement associated with the WVC.

The egg and slice measurements taken by a scatterometer can be thought of as a weighted spatial average of the true σ^0 values on the earth's surface. The relation between a transmitted radar pulse (P_t) and received radar pulse (P_r) is given by the radar equation[2] [1]

$$P_\mathrm{r} = \frac{P_\mathrm{t}\lambda^2}{(4\pi)^3} \iint \frac{G^2(x,y)\sigma^0(x,y)}{R^4(x,y)} dxdy \tag{4.2}$$

where λ is the operating frequency of the scatterometer, $G(x,y)$ is the antenna gain, and $R(x,y)$ is the slant range to the illuminated area. An egg or slice measurement of σ^0 (s_i) is then found by [14]

$$s_\mathrm{i} = \frac{P_\mathrm{r}}{X} \tag{4.3}$$

where X is the so-called "X-factor" [8]

$$X = \frac{P_\mathrm{t}\lambda^2}{(4\pi)^3} \iint \frac{G^2(x,y)}{R^4(x,y)} dxdy \tag{4.4}$$

for the egg or slice measurement. By defining the SRF as

$$SRF(x,y) = \frac{P_\mathrm{t}\lambda^2}{(4\pi)^3 X} \frac{G^2(x,y)}{R^4(x,y)} \tag{4.5}$$

s_i can be modeled as

$$s_\mathrm{i} = \iint SRF(x,y)\sigma^0(x,y)dxdy \tag{4.6}$$

---

[1] See Eq. 2.2 for a more general GMF.

[2] This is another way of expressing the radar equation in Eq. 2.1

Thus, the egg and slice measurements are a weighted spatial average of the true σ^0 values where the weights come from the SRF.

To make realistic measurements, Monte Carlo noise is added to s_i according to[3]

$$s_\text{i,noisy} = s_\text{i}(1 + K_\text{p,i}\nu_\text{i}) \tag{4.7}$$

where $s_\text{i,noisy}$ is the noisy realization of the measurement, ν_i is a realization of a zero-mean unit-variance random variable, and $K_\text{p,i}$ is the normalized standard deviation of $s_\text{i,noisy}$. K_p is modeled as[4] [1]

$$K_\text{p} = \sqrt{K_\text{pA} + \frac{K_\text{pB}}{S_\text{n}} + \frac{K_\text{pC}}{(S_\text{n})^2}} \tag{4.8}$$

where K_pA, K_pB, and K_pC are instrument signal processing parameters that change depending on swath location, and S_n is the SNR of the σ^0 measurement. The values for K_pA, K_pB, K_pC, and the SNR for each σ^0 measurement are provided by the information pulled from the QuikSCAT L1B file.

The azimuth angle, incidence angle, $s_\text{i,noisy}$, K_pA, K_pB, and K_pC for all the egg and slice measurements of a given WVC are extracted from a typical QuikSCAT L1B file and are used as inputs to the wind retrieval process. Of the resulting wind ambiguities, the ambiguity whose direction is closest to the common direction of the original delta wind field is chosen as the correct ambiguity. Due to the added noise in $s_\text{i,noisy}$, this wind vector is only one realization of the possible resulting wind vectors for the WVC with a particular delta wind field. To find the expected wind vector for the WVC and delta wind field, we use 1000 trials with each being a different Monte Carlo noise realization; each trial follows the same procedure but with a different realization of ν_i in Eq. 4.7. The average of the 1000 resulting wind vectors is the estimated expected wind vector for the WVC and delta wind field. This process is repeated for every delta wind field associated with the WVC.

Placing the expected wind vector in the pixel location where the delta wind vector is in the corresponding delta wind field yields a WRF estimate for a given WVC. For the WVC in Fig. 4.1, the associated WRF estimate is shown in Fig. 4.2.

---

[3] Eq. 2.8 is repeated for convenience.

[4] Eq. 2.3 is repeated for convenience.

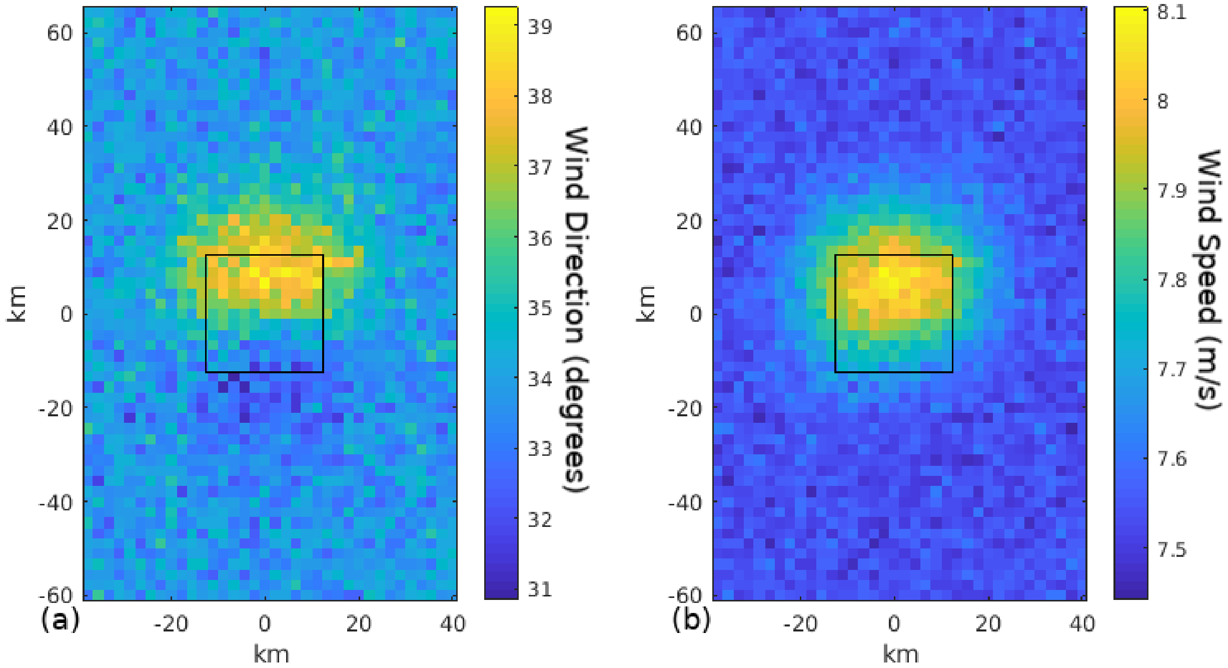

Figure 4.2: Example WRF estimate for DIB egg processing, common direction 30°. (a) shows the direction component of the WRF estimate and (b) shows the speed component of the WRF estimate. The squares outline the area of the WVC. In (a) the colors show the direction (degrees) of the expected wind vector at each image pixel; in (b) the colors show the speed (m/s) of the expected wind vector at each image pixel.

4.1.3 Bias Removal

Since the delta function wind fields only approximate the delta wind vector, the WRF estimate is significantly biased. To reduce this bias, we find the expected wind vector for the wind field case that does not have a delta wind vector, i.e., a background-only wind field. The speed and direction components of this no-delta wind vector are subtracted from the speed and direction components of the delta wind vectors. We choose to perform the subtraction of the no-delta output wind vector in the speed and direction components rather than orthogonal components because the bias removal is more sensitive to variation in the Monte Carlo simulation with orthogonal components. For convenience, the speed component (but not the direction component) of the WRF is normalized so that the peak value is equal to 1. The WRF estimate in Fig 4.2 with the bias removed is shown in Fig. 4.3.

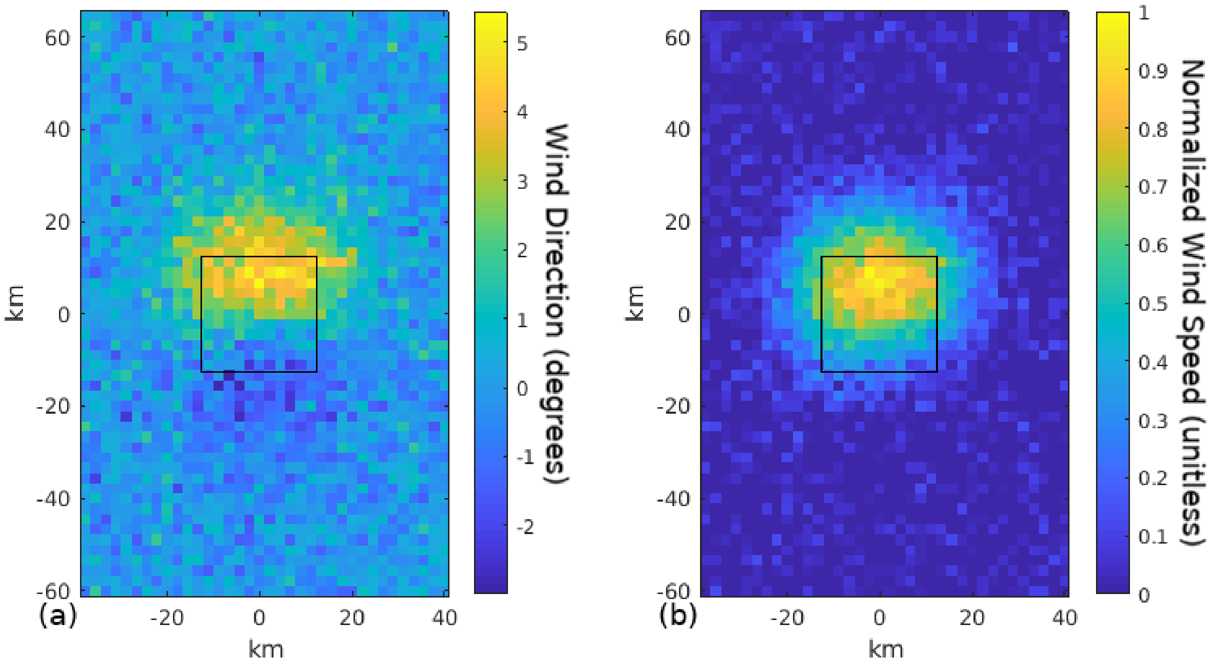

Figure 4.3: Example WRF for DIB egg processing with bias removed, common direction 30°. The no-delta wind vector has been subtracted from the WRF estimate. (a) shows the direction component of the WRF and (b) shows the speed component of the WRF. In (a), the colors show the shift in direction (degrees) of the expected wind vector at each pixel from the no-delta wind vector; in (b) the speed component of the WRF has been normalized so that the peak value is 1. The squares outline the area of the WVC.

With the bias removed, the WRF estimate is the WRF for a WVC at the common direction of the delta wind fields; wind retrieval is biased based on χ [26]. To find the WRF for a WVC independent of direction, we perform "compass simulation." In compass simulation, the above process is repeated with delta wind fields whose common directions represent all the different compass directions and the WRFs for all common directions are averaged together. To reduce computational complexity, we approximate true compass simulation by only using twelve common directions, each one spaced 30° apart from its neighboring common directions.

The inherent bias in wind retrieval with respect to direction may shift the direction of the retrieved wind vectors either to the right or to the left. For example, in Fig. 4.3, the top half of the WVC provides a shift opposite in direction to the bottom half of the WVC. This means that averaging the direction tends to reduce the direction portion of the WRF to near zero; we assume that any non-zero values are a result of the Monte Carlo simulation, not a general bias that is independent of wind direction in wind retrieval. To avoid having a zero-valued direction component, we average the absolute value of the wind direction response. By doing so, information is lost in regards to the bias associated with wind vector direction, but information is retained in regards to the locations of the winds that most strongly affect the retrieved wind direction. The normalized WRF independent of direction for the WVC in Fig. 4.1 is shown in Fig. 4.4. As the WRF depends somewhat upon the quantity and locations of the egg and slice measurements, WRFs for various WVCs are shown in Section 4.2.

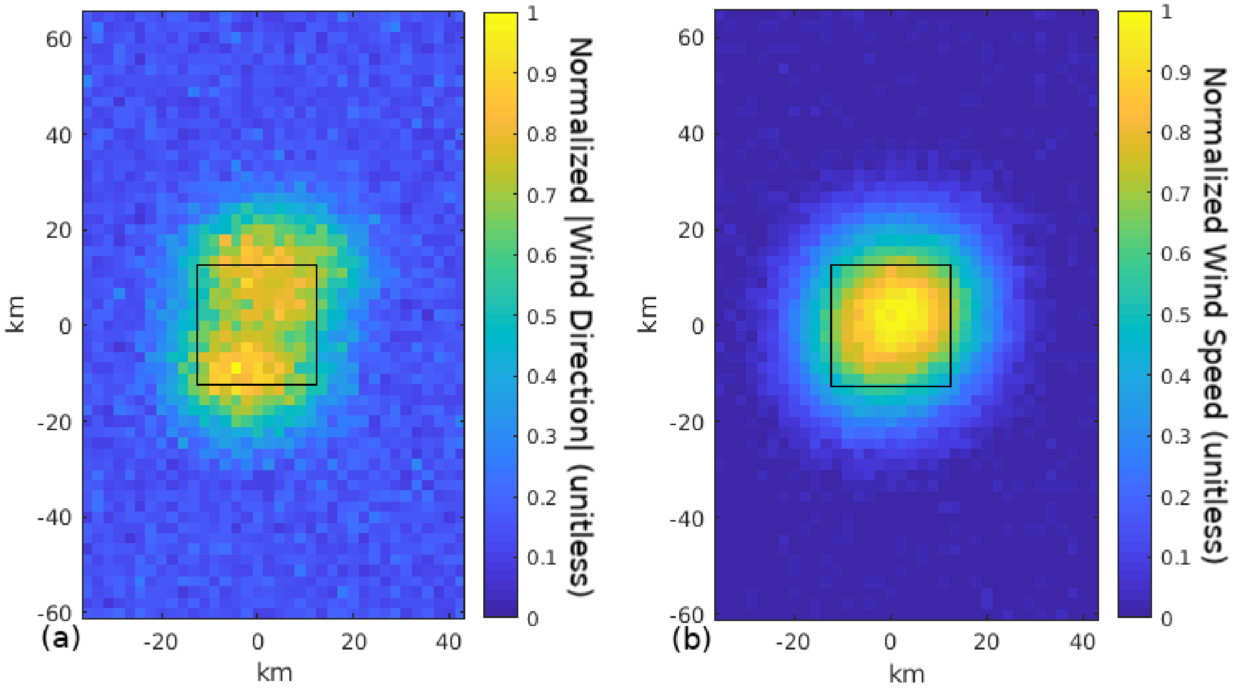

Figure 4.4: Example WRF for DIB egg processing, independent of direction. (a) shows the absolute direction component of the WRF, and (b) shows the speed component of the WRF. The squares outline the area of the WVC. In both (a) and (b), the WRF has been normalized to a peak value of 1.

4.2 Wind Response Functions

To illustrate that the bias with respect to direction is inherent in DIB and UHR processing, WRFs for a WVC at three different common directions for DIB egg, DIB slice, and UHR processing are shown in Figs. 4.5 and 4.6. Fig. 4.5 displays the direction components and Fig. 4.6 displays the speed components of the WRFs. Fig. 4.6 is normalized and has contour lines at 3 dB and 6 dB below the peak value. In these figures, each column (from left to right) is a different common direction: 30°, 150°, and 270°, and each row (from top to bottom) is a different processing technique: DIB egg, DIB slice, and UHR. The WVC in Figs. 4.5 and 4.6 is the same WVC portrayed in Figs. 4.1 through 4.4.

Examining the direction components in Fig. 4.5, the pixels that most affect the expected wind vector are the same for all common directions, but the bias that they produce is different depending on the common direction. Interestingly, when egg and slice measurements with opposing biases overlap, the positive and negative direction shifts cancel each other so that there is a near zero value in these areas. This cancellation indicates that the bias in wind estimation due to wind direction is less affected by wind vectors near the center of the WVC, where egg and slice measurements from a variety of azimuth and incidence angles tend to overlap, and is more affected by wind vectors near the edge of the WVC where egg and slice measurements do not overlap as much.

For the speed components in Fig. 4.6, the pixels that most affect the expected wind vector also vary with the common direction and there appear to be fewer pixels with a large magnitude in the speed component than in the direction component. Close inspection reveals that the speed component pixels with a large magnitude are a subset of the direction component pixels with a large magnitude, specifically the subset whose direction shift is greatest in magnitude. This finding is most clearly seen in the bimodal DIB egg processing WRFs, but is also evident in the DIB slice and UHR processing WRFs. The speed components also begin to suggest that DIB egg processing has a WRF with a Gaussian-like shape. These WRFs indicate that DIB egg processing has the coarsest spatial resolution while UHR processing has the finest spatial resolution of the three processing techniques.

To display that the WRF may be defined in terms of orthogonal components rather than speed and direction, the WRFs in Figs. 4.5 and 4.6 are shown in terms of their u and v components in Figs. 4.7 and 4.8. Figs. 4.7 and 4.8 are organized in the same manner as Figs. 4.5 and 4.6 and, similar to the direction component, the u and v components are signed values that are not normalized. Comparing Figs. 4.5 and 4.6 with Figs. 4.7 and 4.8, the pixels that most affect each component are generally the same; the spatial resolution of the u and v components tends to be slightly coarser than the spatial resolution of the speed component and slightly finer than the spatial resolution of the direction component. The 270° common direction case is of note since 270° is parallel to u component unit vector direction. Ignoring the sign of the WRFs, this leads to the WRF shapes of the u and speed components to be very similar and the WRF shapes of the v and direction components to be very similar. In cases where the common direction is parallel to the v component, the WRF shapes of the v and speed components are very similar and the WRF shapes of the u and direction components are very similar. These findings and further inspection reveal that the information conveyed between the speed and direction component WRFs and the u and v component WRFs is consistent. We analyze the WRF in speed and direction components because they are easier to interpret than the u and v components.

Several examples of direction-independent WRF are shown in Figs. 4.9 and 4.10. These figures are organized similarly to Figs. 4.5 and 4.6 in that each row is a different processing technique, but the columns of Figs. 4.9 and 4.10 represent different WVC locations. The WVC in the first column of these figures is the WVC examined in Figs. 4.5 through 4.8. Each location is within the sweet spot of the scatterometer swath but at different cross track indices. As the azimuth angles of the egg and slice measurements are a function of cross track index, the different cross track indices lead to the fore- and aft-looking measurements overlapping at different angles in the WVCs. This is clearly seen in the DIB slice and UHR processing WRFs. Fig. 4.9 displays the direction components of the WRFs (where the absolute value of the direction components are averaged together) and Fig. 4.10 displays the speed components of the WRFs.

Examining the direction components in Fig. 4.9, there appears to be a dip in the magnitude of the WRFs, more or less around the center of the WVCs. This drop in magnitude is attributed to the smaller bias in wind estimation due to wind direction when more σ^0 measurements overlap as discussed previously. As the drop in magnitude tends to occur near the center of the WVC, the peak value of the direction component of the WRFs tends to be several kilometers away from the center of the WVC. The speed components in Fig. 4.10 confirm the observations made with Fig. 4.6 that the WRF shape of DIB egg processing is Gaussian-like and that DIB egg processing has the coarsest spatial resolution while UHR processing has the finest spatial resolution of the three processing techniques.

4.2.1 Generalized WRF

While the WRFs in Figs. 4.9 and 4.10 show that the WRF varies depending upon swath location, the differences are small enough to imply that a single, generalized WRF could be used everywhere in the QuikSCAT swath. We find a generalized QuikSCAT WRF for three swath regions: the nadir region, the sweet spot, and the far swath. In each region, the WRFs of ten different WVCs are averaged to find the generalized WRF. This provides a reasonable estimate of the generalized WRF while keeping computational complexity low. The absolute value of the direction components of the generalized WRFs for the three regions are shown in Fig. 4.11 and the speed components are shown in Fig. 4.12.

Examining the generalized WRFs for DIB egg processing, the 3 dB contour is roughly circular and a couple of pixels larger than the WVC area, i.e., $\sim$30 km. For each swath area, the speed component resolution is slightly finer than the direction component resolution, and the resolution of the far swath WRF is slightly coarser than the resolution of the sweet spot and nadir region WRFs, i.e., $\sim$35-40 km. Interestingly, the WRFs for DIB slice and UHR processing do not share the same indication. For both DIB slice and UHR processing, the nadir region WRFs have coarser resolutions than the sweet spot and far swath WRFs. This seems to indicate that for DIB slice and UHR processing, azimuth angle diversity in the slice σ^0 measurements helps to improve spatial resolution more than incidence angle diversity.

The speed components in the sweet spot and far swath seem to indicate that the spatial resolution of the DIB slice processing is roughly the same size as a DIB slice processing WVC, i.e., $\sim$12.5 km. The spatial resolution indicated by the direction components is a bit coarser, but only by a couple of pixels, i.e., $\sim$20-25 km. The spatial resolution of the nadir region appears to be slightly elongated in the cross track direction by a few pixels, probably due to a lack of azimuth angle diversity, but still around the same size as a DIB slice processing WVC. The spatial resolution of the DIB egg and DIB slice processing WRFs suggests that in both cases DIB processing tends to have a spatial resolution on the order of the size of the WVC used for the DIB processing, i.e., $\sim$25 km or $\sim$12.5 km.

The WRFs indicate that the spatial resolution of UHR processing is significantly finer than the spatial resolution of DIB processing, i.e., $\sim$5-10 km. The 3 dB contours of the UHR processing speed components are elliptical with a major axes significantly longer than the minor axes. Similar to DIB processing, the 3 dB contours of the direction components for UHR processing are larger than the 3 dB contours of the speed components for UHR processing by a few pixels and the shape of the 3 dB contours is no longer well defined by a single ellipse. The 3 dB contours in the direction components appear more like two ellipses crossed over the center pixels. The acute angle between the ellipses varies based on swath location. While the spatial resolution of UHR processing is not as small as a UHR processing WVC, it is finer than the spatial resolution of DIB processing. This demonstrates that UHR processing improves the spatial resolution of wind estimation.

To help understand the shapes of the WRFs, cross sections of the speed component of the generalized WRFs from the sweet spot are shown in Figs. 4.13 and 4.14. The top row in Fig. 4.13 represents DIB egg processing while the bottom row represents DIB slice processing; the left column represents horizontal cross sections through the WRFs and the right column represents vertical cross sections through the WRFs. UHR processing is better represented with diagonal cross sections in Fig. 4.14; the left side being a cross section from the lower left to the upper right of the WRF and the right side being a cross section from the upper left to the lower right of the WRF. Each cross section passes through the peak pixel. The UHR cross sections pass through the centers of the main diagonal elements of the UHR processing WRFs forming an acute angle of 73°. Each cross section is compared to a Gaussian to help give a reference for the shape of the cross section. Better models than a Gaussian may exist for some of the WRF shapes, but a Gaussian is sufficient to give an idea of the WRF shapes.

The speed component of the DIB egg processing sweet spot generalized WRF is very well approximated by a Gaussian and is close to being circularly symmetric. The other DIB egg processing generalized WRFs appear to be well approximated by circularly symmetric Gaussians as well. The DIB slice processing and UHR processing WRFs have more shape variety than the DIB egg processing WRFs. The speed component of the DIB sweet spot generalized WRF in Fig. 4.13 is somewhat circularly symmetric and seems to be reasonably approximated by a Gaussian, though the sharp peak at the center of the WRF is not modeled so well by the Gaussian. The speed component of the UHR sweet spot WRF in Fig. 4.14 seems that it could be modeled by the addition of two elliptical Gaussians that cross at pixels near the peak of the WRF. Similar to the speed component of the DIB slice processing WRF, the peak at the center of the speed component of the UHR processing WRF may not be modeled as well by the two Gaussians as the rest of the

Table 4.1: Effective spatial resolutions of the different processing algorithms in various swath locations. The 3 dB effective spatial resolution is approximated as an ellipse with the diameter of the major and minor axes given in the table.

Location	Axis	DIB Egg	DIB Slice	UHR
Far Swath	Major	41 km	18 km	11 km
	Minor	35 km	9 km	5 km
Sweet Spot	Major	31 km	15 km	10 km
	Minor	30 km	12 km	5 km
Nadir Region	Major	31 km	24 km	25 km
	Minor	30 km	12 km	8 km

WRF. The remaining DIB slice and UHR processing generalized WRFs appear that they could be modeled somewhat by a single elliptial Gaussian, such as the near nadir generalized WRFs, or by the addition of two elliptical Gaussians that cross at the peak value, such as the far swath generalized WRFs. The acute angle between the intersecting elliptical Gaussians varies based on swath location.

The effective spatial resolutions of DIB egg, DIB slice, and UHR processing are summarized in Table 4.1. The effective spatial resolution is approximated as an ellipse for each processing type and location. The values reported are based on the 3 dB contour of the speed component of the generalized WRFs. Beyond the 3 dB contour, the WRFs may not be well approximated by a single ellipse, but this is similar to the SRF found for various processing algorithms [33]. Table 4.1 clearly shows that the effective spatial resolution of UHR processing is finer than the effective spatial resolution of DIB processing.

4.3 Conclusion

In this chapter, the following contributions are made: (1) a WRF is defined and a method to calculate the WRF for various wind estimation processing techniques is presented; (2) estimates of the effective spatial resolution for three QuikSCAT processing techniques are found based on the generalized WRF calculated for each; (3) QuikSCAT UHR processing is shown to have a finer effective spatial resolution than DIB processing; (4) evidence is provided that the bias in wind retrieval due to wind direction is mitigated through overlapping σ^0 measurements from a variety of azimuth and incidence angles. These contributions can be used to validate new wind estimation processing techniques and can be extended to compare processing techniques across different scatterometers.

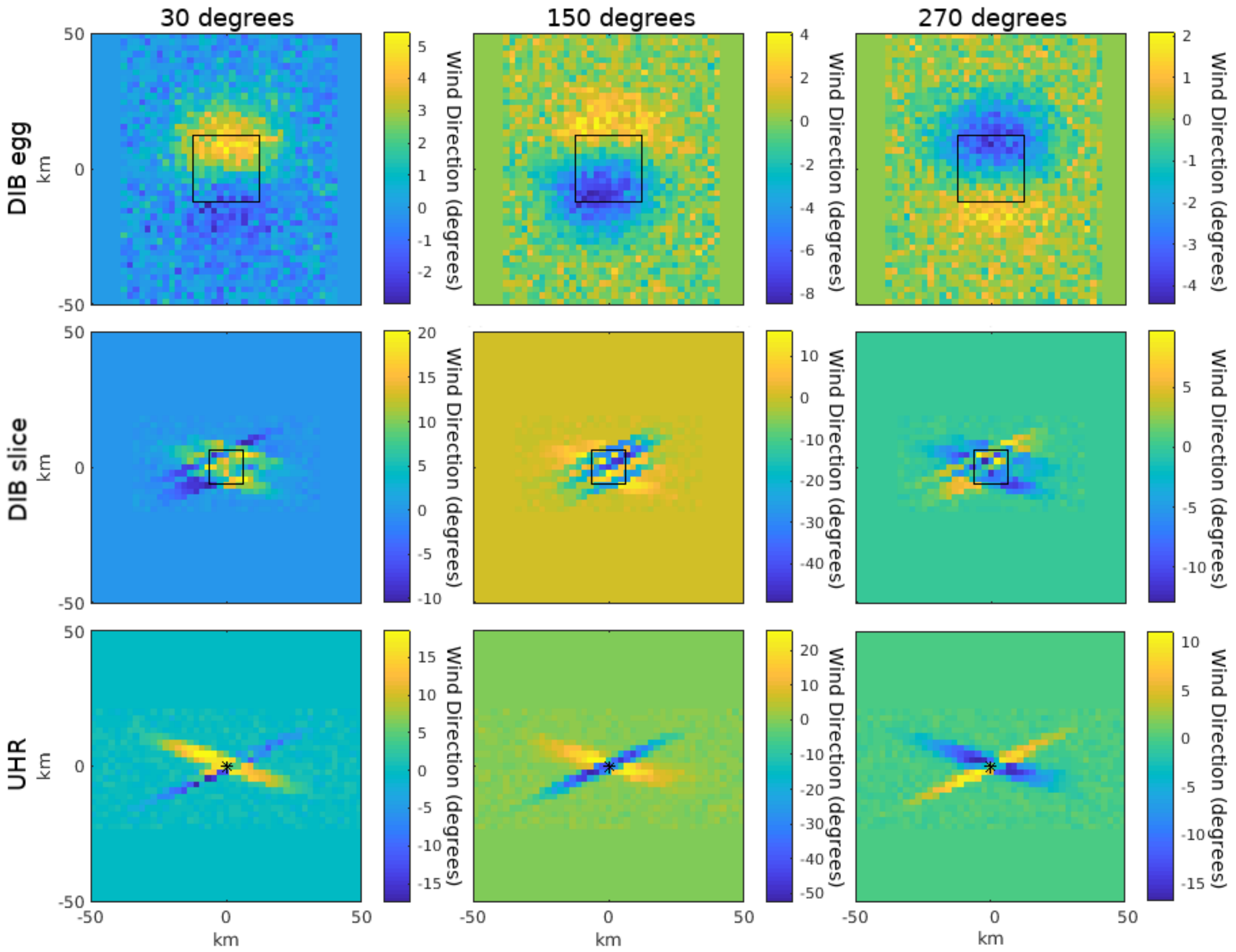

Figure 4.5: The direction components of the WRFs for a mid-swath WVC location using different processing algorithms and different common directions. The rows (from top to bottom) are DIB egg, DIB slice, and UHR processing. The columns (from left to right) are 30°, 150°, and 270° common directions. The black squares indicate the edges of the DIB egg and DIB slice processing WVCs, and the black asterisks indicate the location of the UHR processing WVC pixel. The center locations of the WVCs are the same for all processing techniques. The colors in each panel represent the shift in direction (degrees) of the expected wind vector at each pixel.

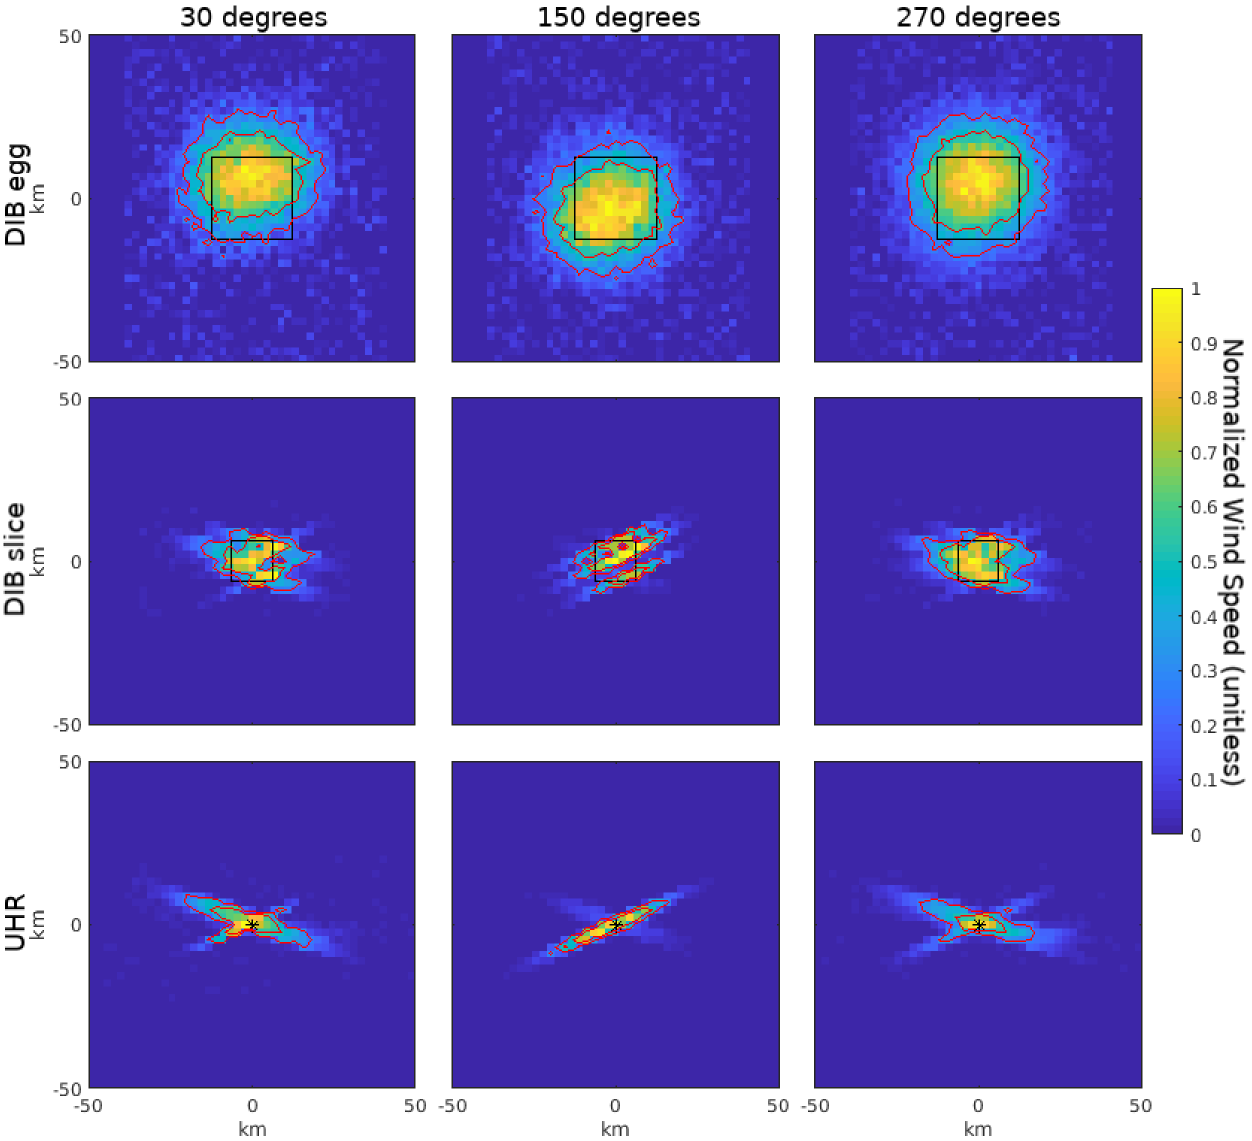

Figure 4.6: The speed components of the WRFs for a mid-swath WVC location using different processing algorithms and different common directions. The (rows from top to bottom) are DIB egg, DIB slice, and UHR processing. The columns (from left to right) are 30°, 150°, and 270° common directions. The black squares indicate the edges of the DIB egg and DIB slice processing WVCs, and the black asterisks indicate the location of the UHR processing WVC pixel. The center locations of the WVCs are the same for all processing techniques. Red contour lines are placed at 3 dB and 6 dB below the peak value; the values in each panel have been normalized to a peak of 1.

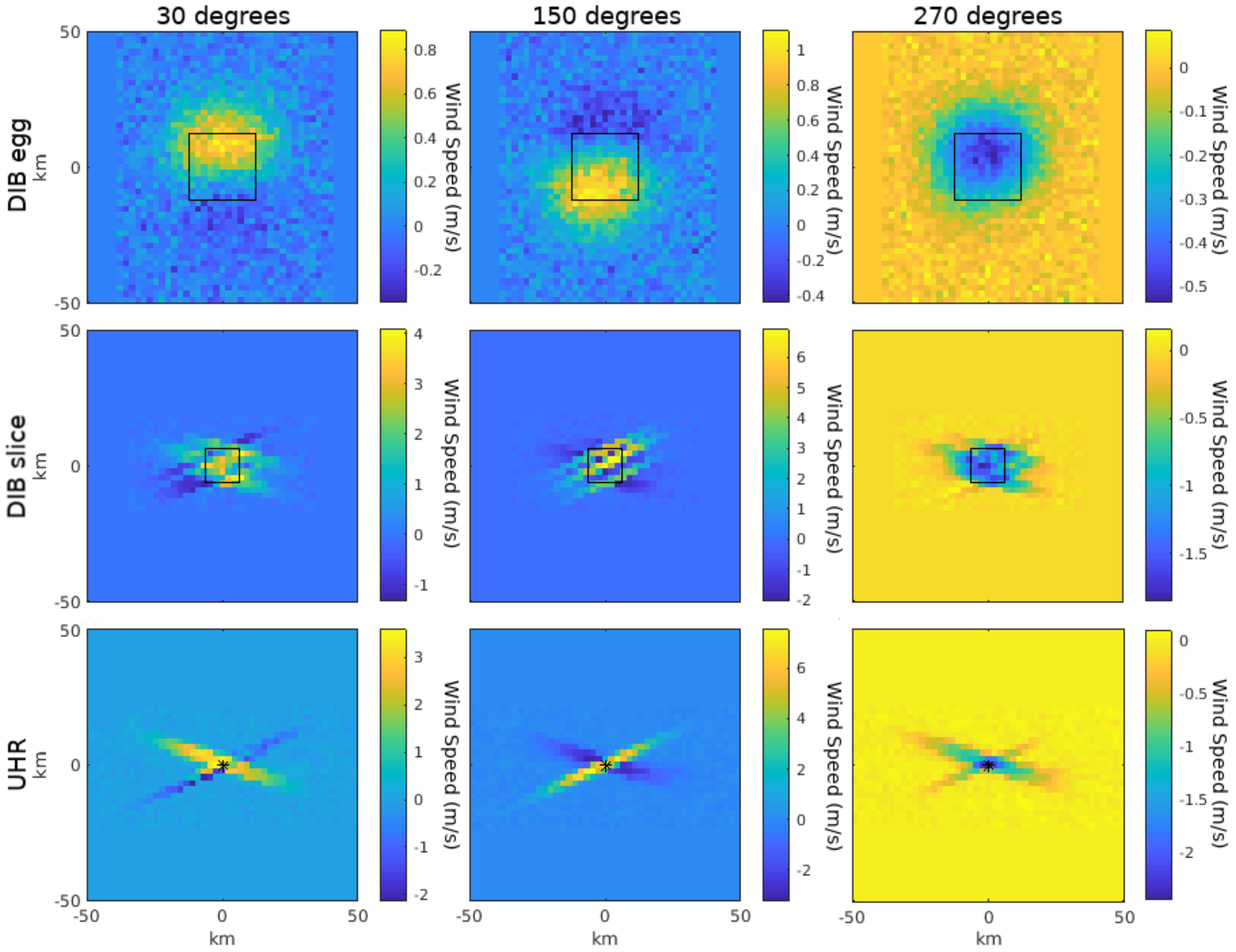

Figure 4.7: The *u* components of the WRFs for a mid-swath WVC location using different processing algorithms and different common directions. The rows (from top to bottom) are DIB egg, DIB slice, and UHR processing. The columns (from left to right) are 30°, 150°, and 270° common directions. The black squares indicate the edges of the DIB egg and DIB slice processing WVCs, and the black asterisks indicate the location of the UHR processing WVC pixel. The center locations of the WVCs are the same for all processing techniques.

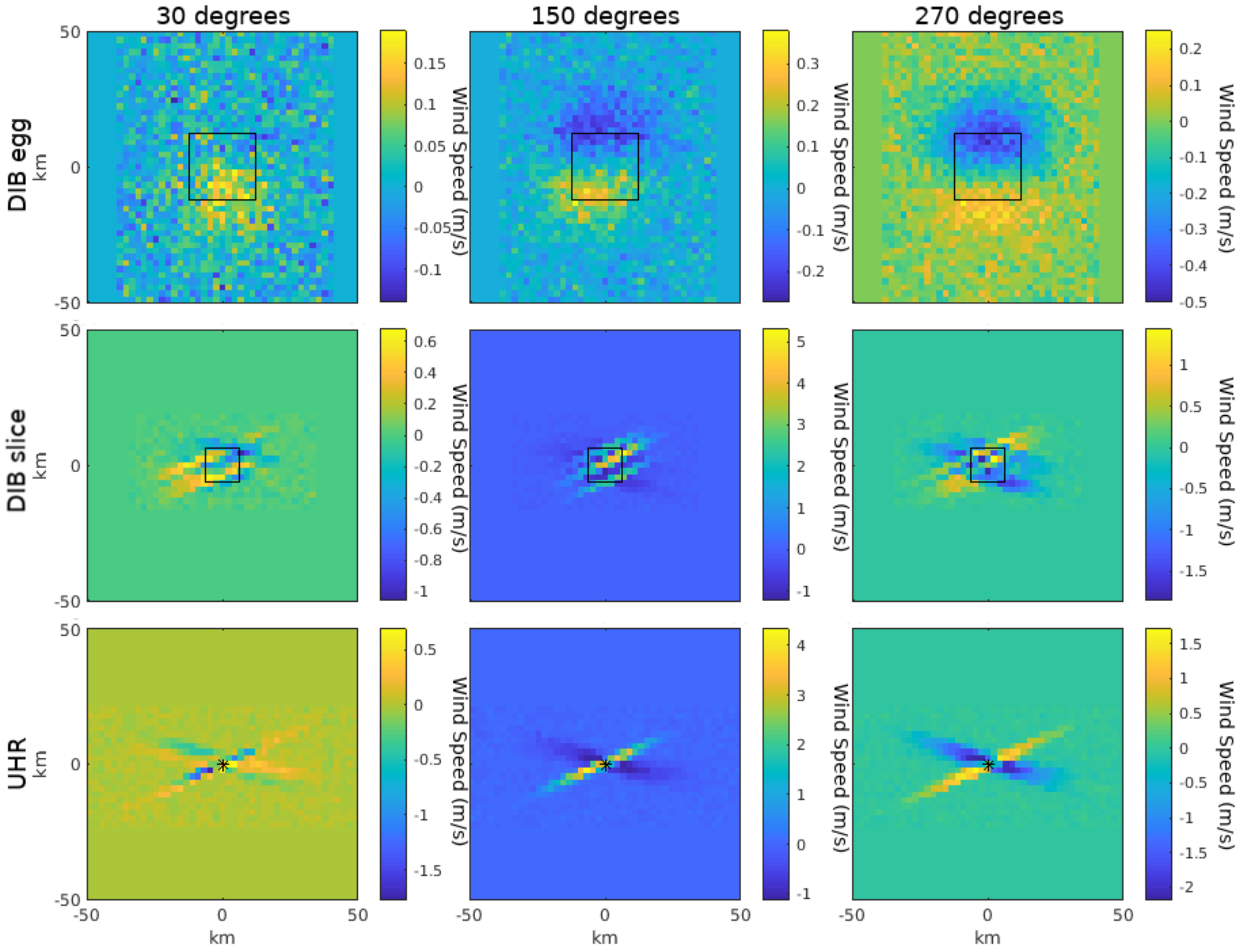

Figure 4.8: The v components of the WRFs for a mid-swath WVC location using different processing algorithms and different common directions. The rows (from top to bottom) are DIB egg, DIB slice, and UHR processing. The columns (from left to right) are 30°, 150°, and 270° common directions. The black squares indicate the edges of the DIB egg and DIB slice processing WVCs, and the black asterisks indicate the location of the UHR processing WVC pixel. The center locations of the WVCs are the same for all processing techniques.

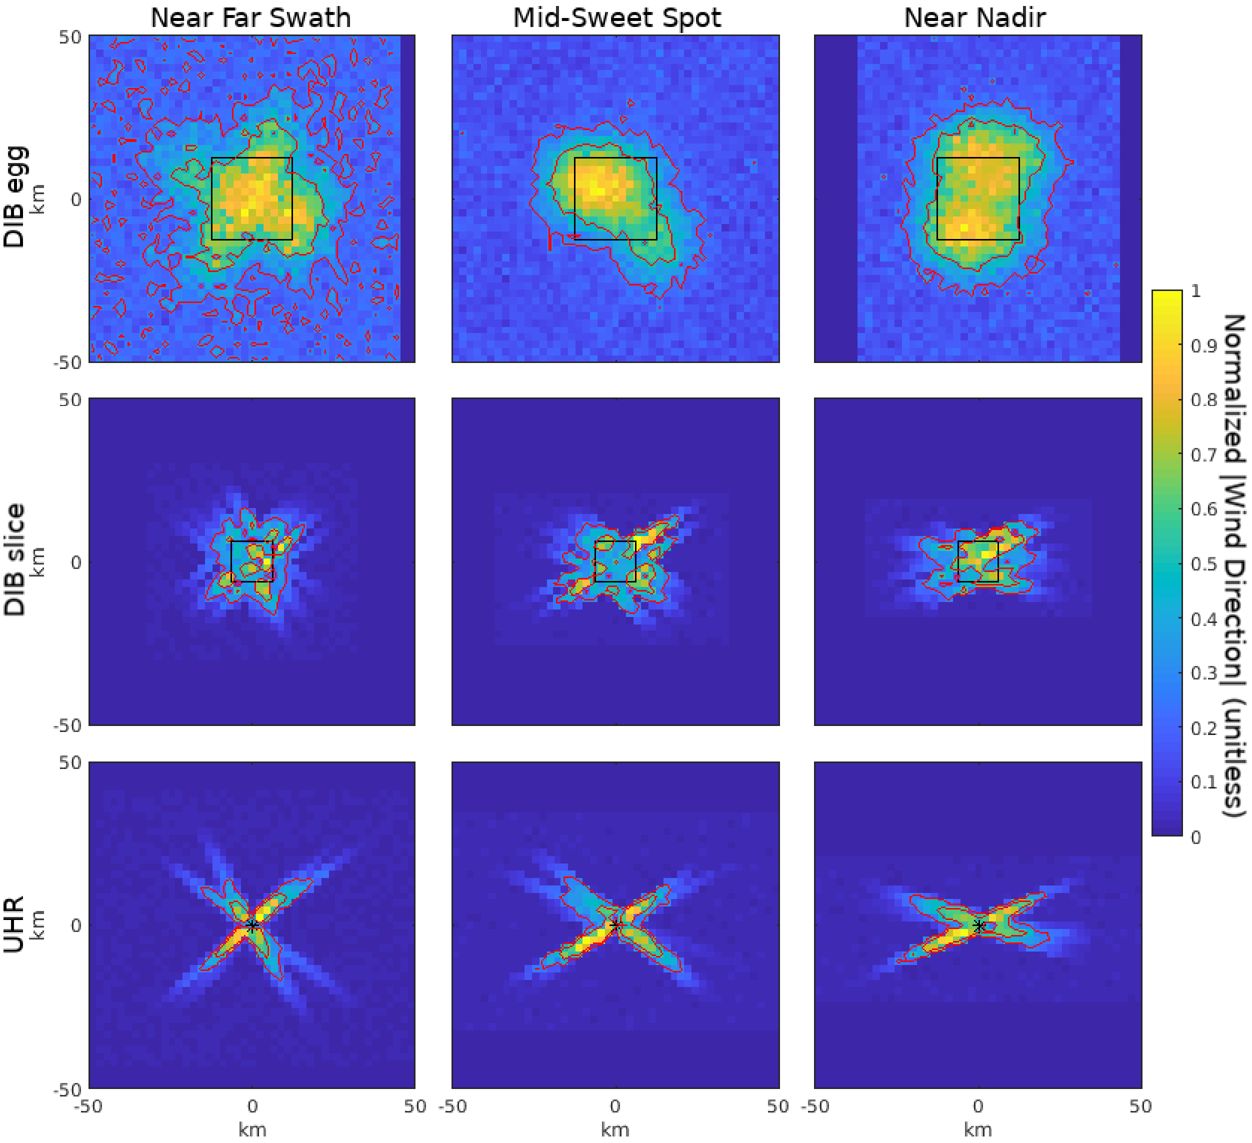

Figure 4.9: The absolute value direction components of the WRFs independent of direction for three WVCs using different processing algorithms. The rows (from top to bottom) are DIB egg, DIB slice, and UHR processing. The columns are three different WVCs that are (from left to right) from the sweet spot close to the far swath, from the middle of the sweet spot, and from the sweet spot close to the nadir region. The black squares indicate the edges of the DIB egg and DIB slice processing WVCs, and the black asterisks indicate the locations of the UHR processing WVC pixels. The center locations of the WVCs are the same for all processing techniques. Red contour lines are placed at 3 dB and 6 dB below the peak value; the values in each panel have been normalized.

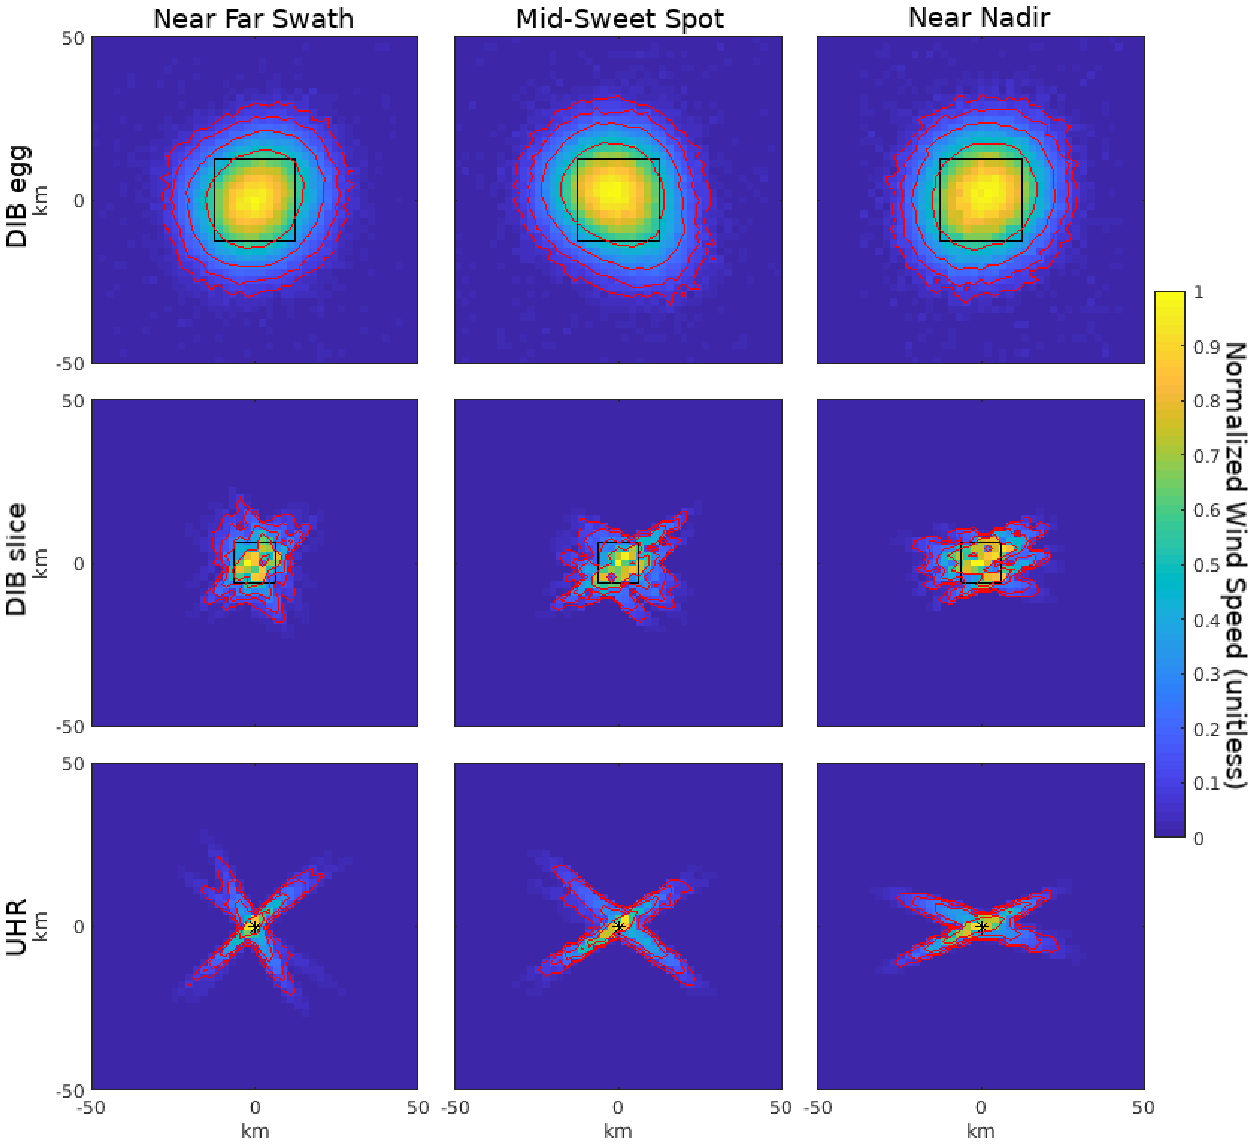

Figure 4.10: The speed components of the WRFs independent of direction for three WVCs using different processing algorithms. The rows (from top to bottom) are DIB egg, DIB slice, and UHR processing. The columns are three different WVCs that are (from left to right) from the sweet spot close to the far swath, from the middle of the sweet spot, and from the sweet spot close to the nadir region. The black squares indicate the edges of the DIB egg and DIB slice processing WVCs, and the black asterisks indicate the locations of the UHR processing WVC pixels. The center locations of the WVCs are the same for all processing techniques. Red contour lines are placed at 3 dB, 6 dB, 9 dB, and 12 dB below the peak value; the values in each panel have been normalized.

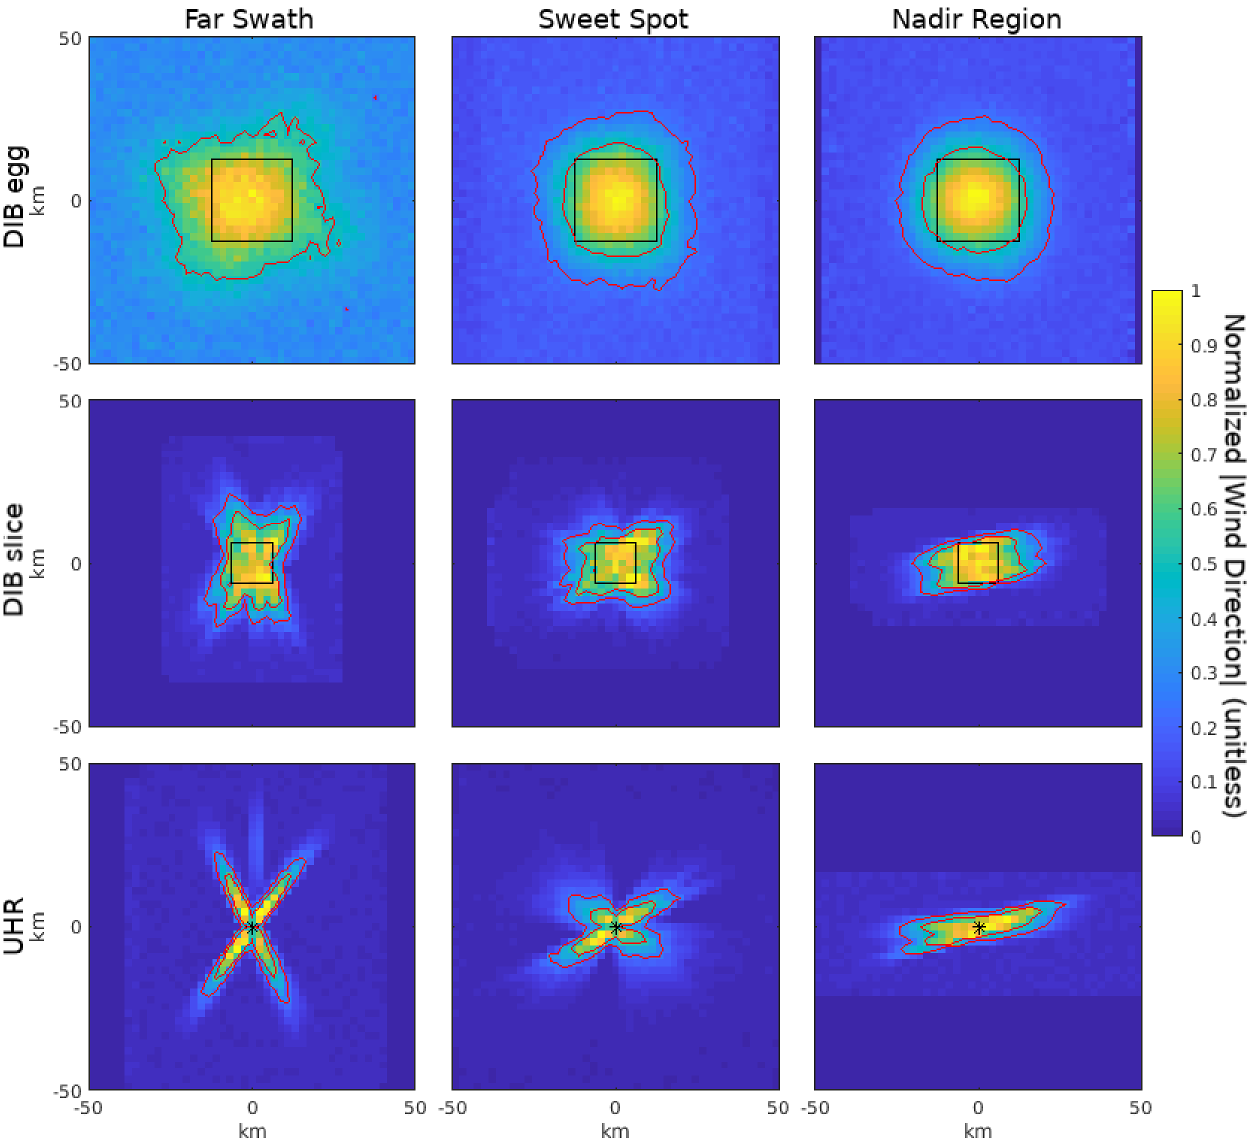

Figure 4.11: The direction components of the generalized WRFs for different processing algorithms. The rows (from top to bottom) are DIB egg, DIB slice, and UHR processing. The columns are generalized WRFS for WVCs that are (from left to right) from the far swath, from the sweet spot, and from the nadir region. The black squares indicate the edges of the DIB egg and DIB slice processing WVCs, and the black asterisks indicate the locations of the UHR processing WVC pixels. The center locations of the WVCs are the same for all processing techniques. Red contour lines are placed at 3 dB and 6 dB below the peak value; the values in each panel have been normalized.

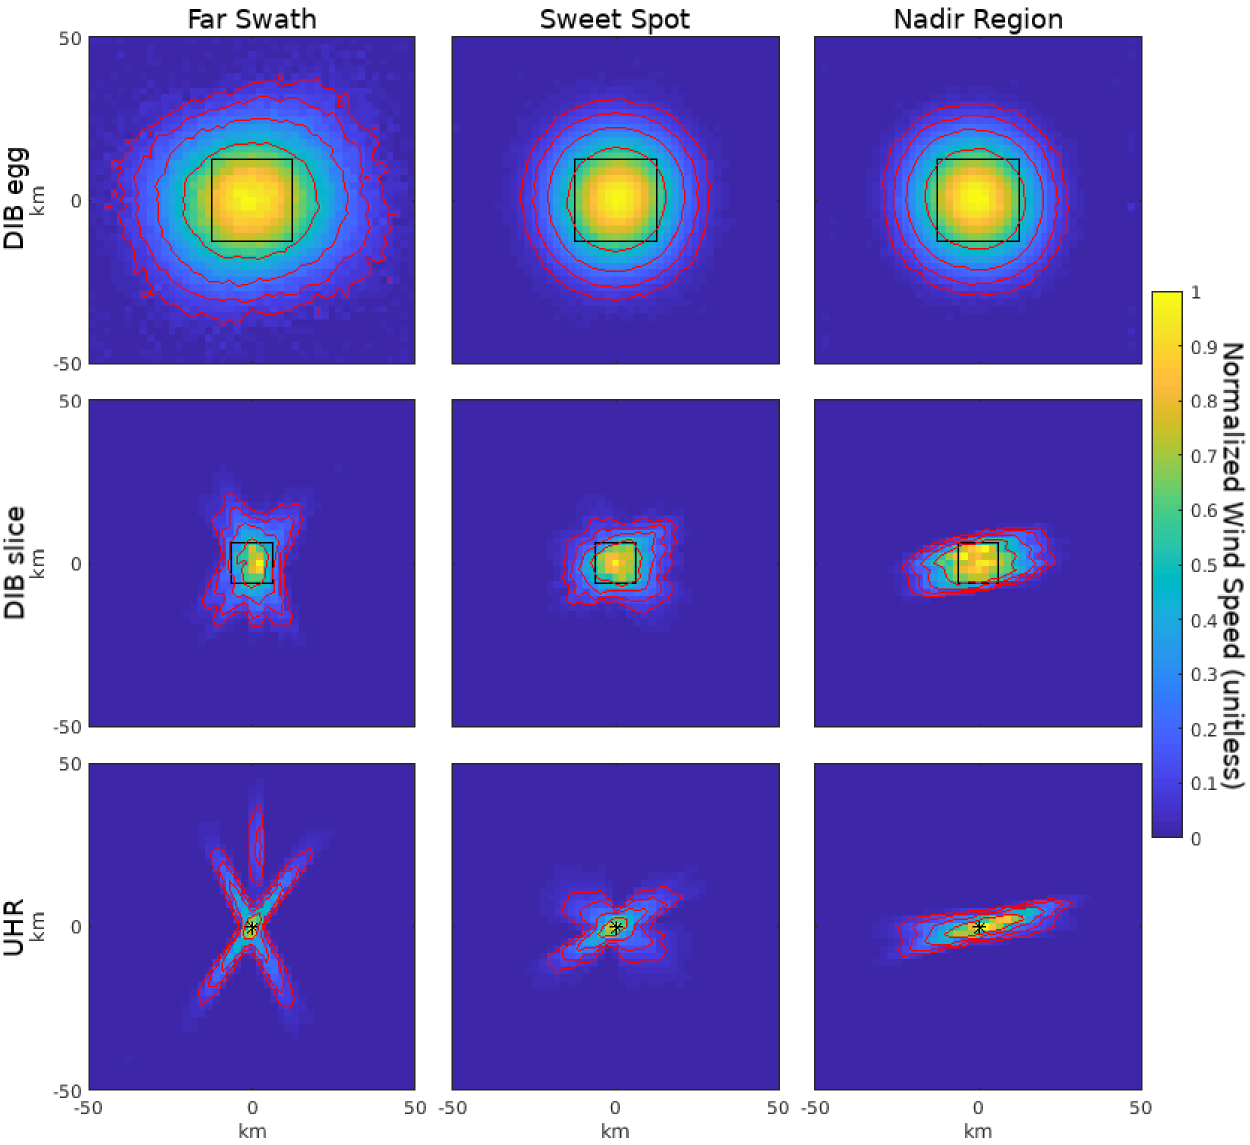

Figure 4.12: The speed components of the generalized WRFs for different processing algorithms. The rows (from top to bottom) are DIB egg, DIB slice, and UHR processing. The columns are generalized WRFs for WVCs that are (from left to right) from the far swath, from the sweet spot, and from the nadir region. The black squares indicate the edges of the DIB egg and DIB slice processing WVCs, and the black asterisks indicate the locations of the UHR processing WVC pixels. The center locations of the WVCs are the same for all processing techniques. Red contour lines are placed at 3 dB, 6 dB, 9 dB, and 12 dB below the peak value; the values in each panel have been normalized.

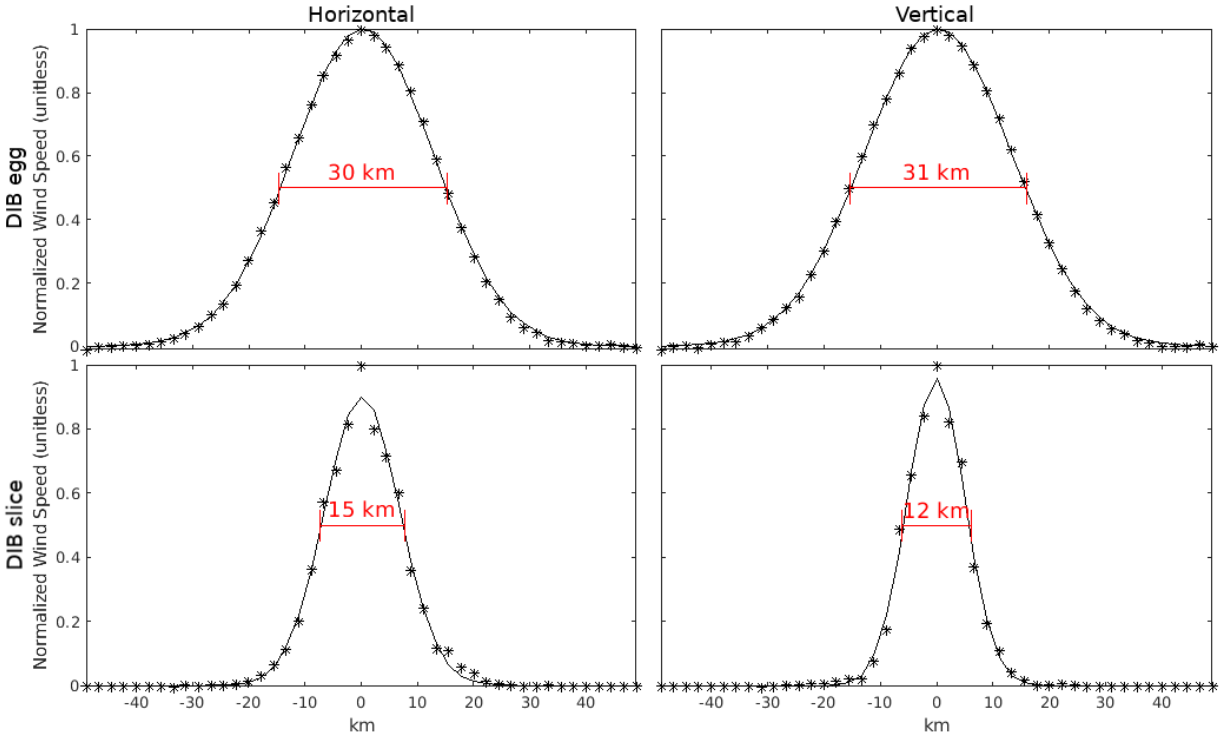

Figure 4.13: Cross sections of the speed component of the sweet spot generalized WRFs for DIB processing. The top row is DIB egg processing while the bottom row is DIB slice processing. The left column is a horizontal cross section while the right column is a vertical cross section. The x-axes are in km and the generalized WRFs have been normalized.

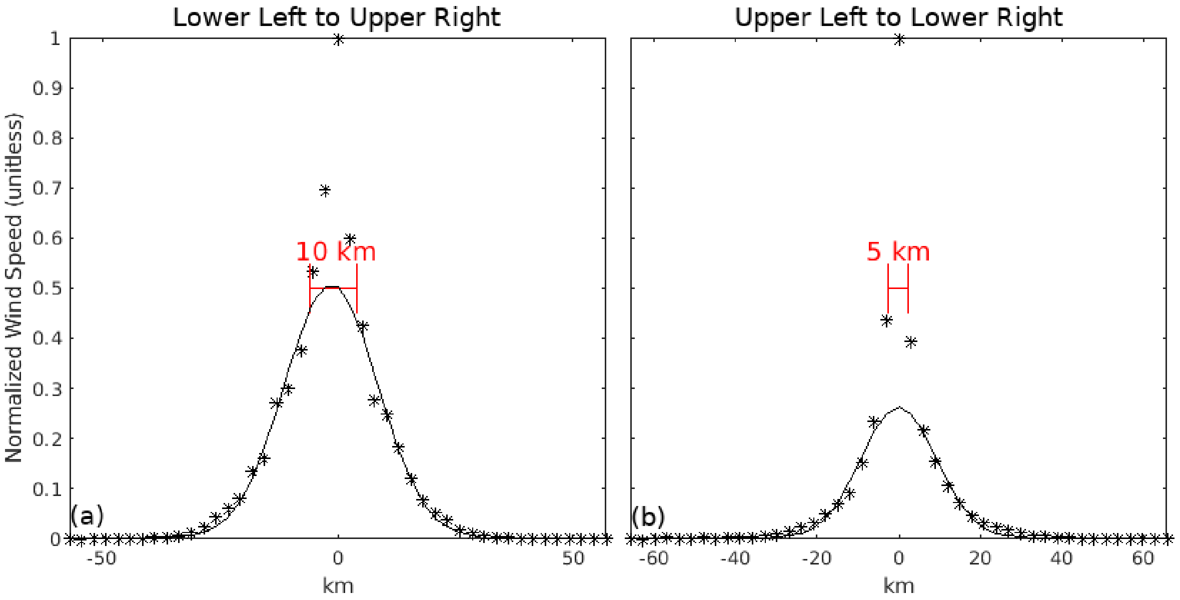

Figure 4.14: Cross sections of the speed component of the sweet spot generalized WRF for UHR processing. (a) is a diagonal cross section from the lower left to the upper right while (b) is a diagonal cross section from the upper left to the lower right. The acute angle between the intersecting sections is approximately 73°. The x-axes are in km and the generalized WRFs have been normalized.

CHAPTER 5. A NEW METHOD TO CREATE REALISTIC SYNTHETIC WIND FIELDS FOR SCATTEROMETER SIMULATION

Determining the accuracy of scatterometer-measured winds is an important part of analyzing wind retrieval and ambiguity selection algorithms. Scatterometer-measured winds sometimes coincide in time and space with winds measured by in situ instruments. These collocated measurements are often used to validate wind estimation algorithms, especially in near-coastal locations where in situ measurements are more common [7], [24]. However, for the vast majority of the ocean there are no other data to validate scatterometer measurements. In these areas, different forms of analysis are employed to verify that the measured winds are plausible, even if their exact accuracy is unknown.

An alternative method of measuring the accuracy of scatterometer wind retrieval and ambiguity selection algorithms can be achieved through simulation. While simulation may not perfectly model actual scatterometer measurements, the true wind field is known in simulation. This means a quantitative accuracy measurement can be taken in simulation as the measured wind field is compared with the true wind field.

In this and the subsequent chapter, we present an improved simulation tool as a method to analyze wind retrieval and ambiguity selection algorithms. This chapter develops the first part of the simulation tool, an improved process to make fine resolution synthetic wind fields sufficiently realistic to test UHR ambiguity selection algorithms. While the quality of simulation results for all wind estimation algorithms is better when a realistic wind field is used as the true wind field, ambiguity selection algorithms especially require a true wind field with a certain level of realism for their results to translate well to actual scatterometer measurements. Synthetic wind fields made from previous methods are insufficient for testing ambiguity selection algorithms because they either focus on specific wind features that do not generalize to the entire ocean well or they have an unrealistic spatial distribution of wind features. For convenience with the rest of the simulation, in this chapter we make synthetic wind fields with WVCs that are 2.225 km × 2.225 km. The process can be adapted to any given WVC size.

This chapter is organized as follows. Section 5.1 develops the synthetic wind field creation process. Section 5.2 analyzes the process itself to explain certain characteristics and attributes that the created synthetic wind fields will possess. Section 5.3 validates the synthetic wind fields through divergence and vorticity measurements and Section 5.4 provides a brief conclusion.

5.1 Synthetic Wind Creation Process Overview

We provide a brief overview of the method to create a realistic synthetic wind field. The method begins by performing bandlimited interpolation on a low resolution NWP wind field. The wind field passes through a prewhitening filter restricted to the low frequency components, yielding a bandlimited white wind field. The extrapolated high frequency components of the bandlimited white wind field are assigned a uniform magnitude and random phase, thus extrapolating the frequency content to the desired resolution. The white wind field then passes through a postwhitening filter to remove the white noise characteristics. At this point the kinetic energy spectrum of the average wind field does not follow the expected power law and the wind field still contains unrealistic wind features. To produce the proper power law and remove these undesirable artifacts and wind features, the full wind field passes through a final filter. The resulting wind field is the realistic high resolution wind field. The key steps in the process are discussed in greater detail in Sections 5.1.1 through 5.1.5.

5.1.1 Initial Wind Field

The synthetic wind field creation process begins with a low resolution NWP wind field. This wind field sets the general wind characteristics of the final synthetic wind field and should be realistic. A secondary purpose for the low resolution field is that many ambiguity selection algorithms begin with an initial guess of the wind field ambiguities. For L2B wind fields, this initial guess comes from selecting the ambiguities that are closest to the wind vectors predicted by a low resolution NWP wind field. The low resolution wind field may double as the NWP wind field for testing those ambiguity selection algorithms.

One source for these low resolution wind fields is the set of NWP wind fields developed by the European Centre for Medium-Range Weather Forecasts (ECMWF). These wind fields are produced on a global scale with one wind vector given at each integer latitude and longitude crossing. The simplest way to project the spherical ECMWF NWP wind fields to a 2D wind field with evenly spaced WVCs is to perform an equirectangular projection. This projection greatly distorts the WVCs near the poles, but is negligible at the equator. As many areas of open ocean sufficiently large enough to make a useful synthetic wind field are near the equator, this method makes the simplifying assumption that any distortion caused by the equirectangular projection is negligible and the resulting 2D wind field has a nominal distance of 111.25 km between the centers of each adjacent WVC. While such a simplifying assumption is limiting in many other instances, the resulting wind fields contain the same large-scale wind features and characteristics that are predicted for winds on that scale. They are sufficiently realistic as low resolution wind fields. A generic NWP wind field is shown in Fig. 5.1.

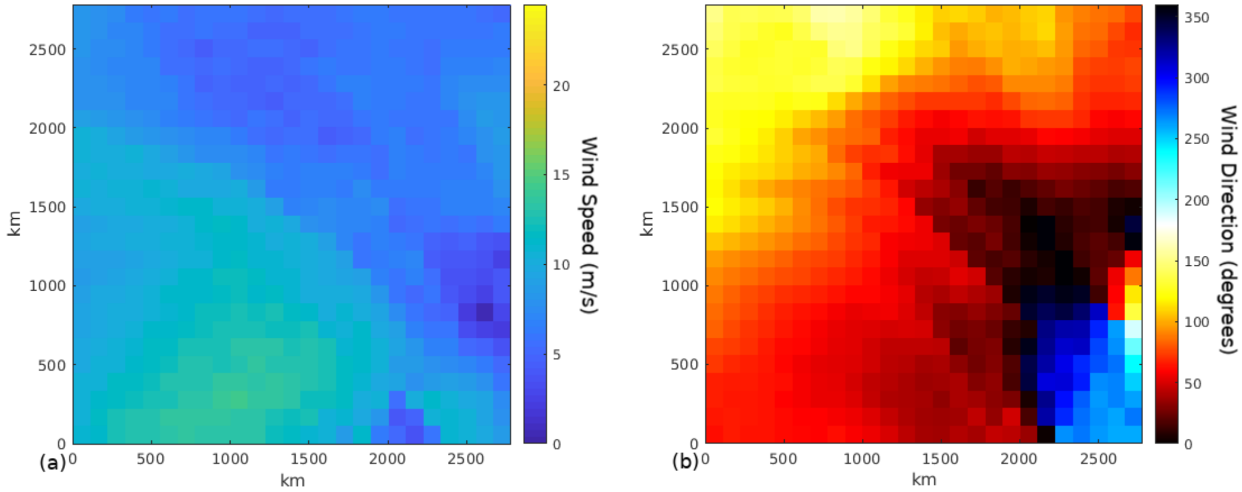

Figure 5.1: A sample ECMWF NWP wind field. (a) shows the wind speed (m/s) and (b) shows the wind direction (degree). This is the initial wind field used when representing all steps of the synthetic wind field creation process. The wind field is nominally 2781.25 km × 2781.25 km.

5.1.2 Prewhitening Filter

The low resolution wind field undergoes bandlimited interpolation to the high resolution sampling period and passes through a prewhitening filter that only affects the low frequency elements of the frequency spectrum. As mentioned in Chapter 2, the average magnitude component of ocean winds' 2D discrete Fourier transform (DFT) follows a known power law and can be extrapolated. Whitening the wind field before extrapolating the magnitude component of the 2D DFT creates a smoother transition between the original and extrapolated portions and reduces unrealistic artifacts in the wind field caused by that juncture.

Additionally, the phase component of wind fields is difficult to extrapolate since the phase is hard to model. By whitening the wind field, the phase at each frequency component ideally becomes uncorrelated from the phase at other frequency components. The phase of a white wind field is then much easier to extend by simply randomizing the phase at each new frequency bin.

The prewhitening filter is implemented by generalizing the first order difference equation $z_i' = z_{i+1} - z_i$ to 2D. The 2D first order difference equation is performed by convolving the weighting image shown in Fig. 5.2 with the bandlimited wind field. For computational speed, the weighting image is zero padded to the size of the bandlimited wind field and the convolution between the weighting image and the bandlimited wind field is implemented as multiplication in the frequency domain. The prewhitening filter does not add any high frequency content beyond the cutoff frequency of the low resolution wind field making a white bandlimited wind field. The white bandlimited wind field associated with the NWP wind field in Fig. 5.1 is shown in Fig. 5.3. By performing the bandlimited interpolation and prewhitening in this order, the prewhitening and postwhitening are both implemented at high resolution which simplifies the creation of the postwhitening filter.

5.1.3 2D DFT Extrapolation

As previously mentioned, the interpolation of new wind vectors in a wind field occurs through extrapolating the magnitude of the 2D DFT and randomizing the associated phase. In general, a 2D kinetic energy spectrum can be calculated by squaring the magnitude of the 2D DFT. A white kinetic energy spectrum, which is a constant value for all frequency bins, ideally corresponds to a uniform magnitude for all frequency bins in the DFT. For the white wind field,

-0.5	-1	-0.5
-1	6	-1
-0.5	-1	-0.5

Figure 5.2: The weighting image convolved with the NWP wind field to prewhiten the wind field.

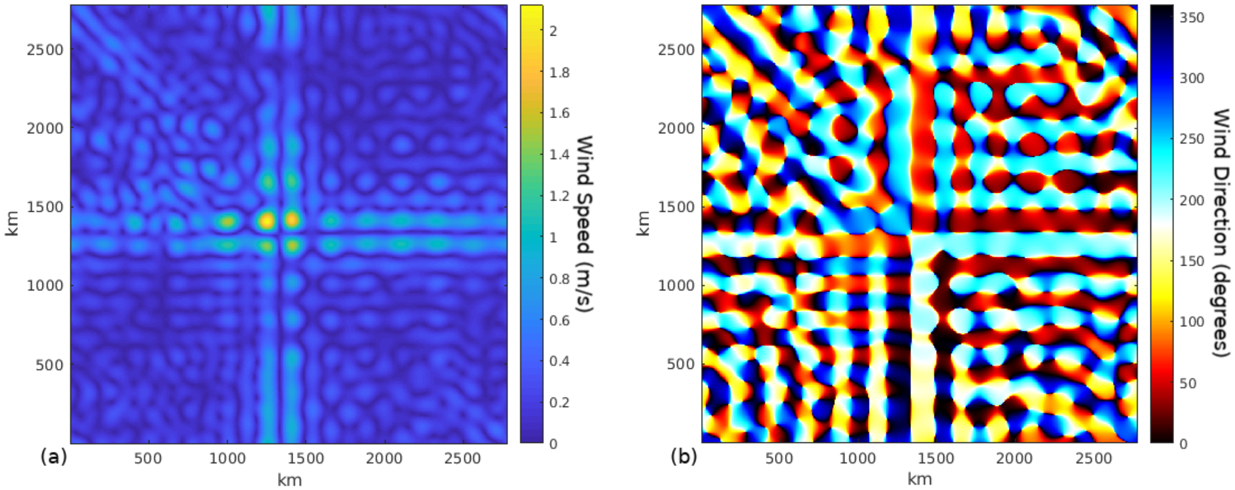

Figure 5.3: The white bandlimited wind field of the sample ECMWF NWP wind field from Fig. 5.1. (a) shows the wind speed (m/s) and (b) shows the wind direction (degree). The wind speed scale is arbitrary at this point. The wind field is nominally 2781.25 km × 2781.25 km.

the magnitude assigned to all extrapolated high frequency bins is the mean magnitude of the low frequency bins of the white bandlimited wind field. The phase component of the extrapolated frequency bins is assigned a realization of a uniformly distributed independent random variable from $-\pi$ to π. Taking the inverse 2D DFT of the extrapolated spectrum yields the white wind field. The white wind field corresponding to the low resolution wind field in Fig. 5.1 is given in Fig. 5.4.

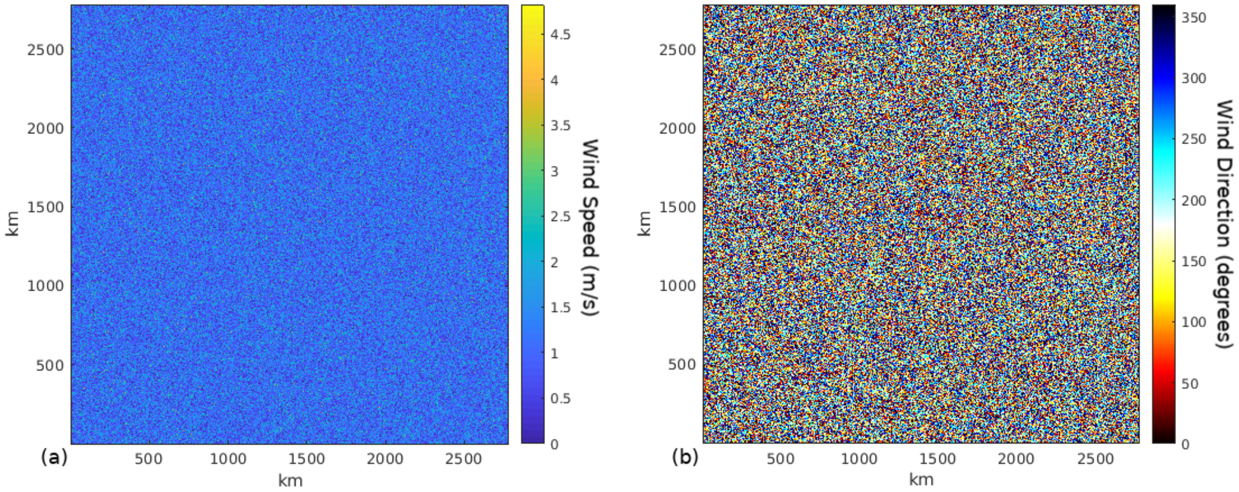

Figure 5.4: The white wind field of the sample ECMWF NWP wind field from Fig. 5.1. (a) shows the wind speed (m/s) and (b) shows the wind direction (degree). Due to the nature of the filters that the wind field passes through, the wind speed is not scaled to the proper range of values at this point. The wind field is nominally 2781.25 km × 2781.25 km.

Since each extrapolated frequency bin corresponds to higher resolution detail added to the original wind field, the amount of extrapolation that the wind field undergoes corresponds to the ratio between the low resolution sampling period and the high resolution sampling period. For example, interpolating a 111.25 km resolution wind field to a UHR 2.225 km resolution wind field corresponds to extrapolating the $M \times N$ low resolution 2D DFT to a $50M \times 50N$ high resolution 2D DFT.

5.1.4 Postwhitening Filter

The postwhitening filter is the inverse of the prewhitening filter with a few modifications. The first modification is that the DC component of the postwhitening filter is set to zero to avoid a singularity because the DC component of the prewhitening filter is also zero. The DC component is handled separately during the final filter. The second modification is to constrain the filter to be circularly symmetric. The weighting image in Fig. 5.2 used to create the prewhitening filter is used for its simplicity and ease of implementation but it does not make a circularly symmetric prewhitening filter. The associated postwhitening filter is also not circularly symmetric without modification. Modifying the postwhitening filter to be circularly symmetric helps the final wind field to be rotationally invariant, a characteristic of realistic wind fields.

The importance of a rotationally invariant wind field is seen by comparing the low resolution wind field in Fig. 5.1 with the white bandlimited wind field in Fig. 5.3. In Fig. 5.1, the wind field does not contain any features that are dependent upon the axis of presentation or measurement. In Fig. 5.3, the wind field has gone through a non-circularly symmetric prewhitening filter and frequency bin extrapolation in two orthogonal directions. The dominating features of the wind field in Fig. 5.3 are orthogonal features in line with the axes of presentation and measurement. The rotationally variant wind field in Fig. 5.3 is much less realistic than the rotationally invariant wind field in Fig. 5.1.

To make the postwhitening filter circularly symmetric, the filter is taken into the spatial domain. Most of the postwhitening filter is a circularly symmetric bump that goes to zero away from the center. The circularly symmetric values go to zero and then nonsymmetric negative values ring the outside of the filter. By zeroing out the any value that is less than zero, the filter is constrained to be circularly symmetric. In the frequency domain the DC component is set to zero again.

5.1.5 Final Filter

After the white wind field passes through the postwhitening filter, the wind field again exhibits rolloff with the wavenumber k. Unfortunately, the postwhite wind field contains a few undesirable artifacts and features, though the general wind features of the original wind field are still visible. As an example, compare the characteristics and flow of the wind field in Fig. 5.5 to the original low resolution wind field in Fig. 5.1. The observed 2D kinetic energy spectrum of the postwhite wind field confirms that the postwhitening filter does not quite recreate the desired power law of wind fields as the spectrum of Fig. 5.5 is shown in Fig. 5.6. These undesirable characteristics are predictable and can be mitigated by further filtering the wind field.

The exact shape and implementation of the final correction filter is admittedly ad hoc and the great variability exhibited by measured wind fields indicates that many different filters could be used to create a vast variety of realistic wind fields from a single initial wind field. Therefore, the filter presented here has been tuned using both model predictions and measured wind fields to

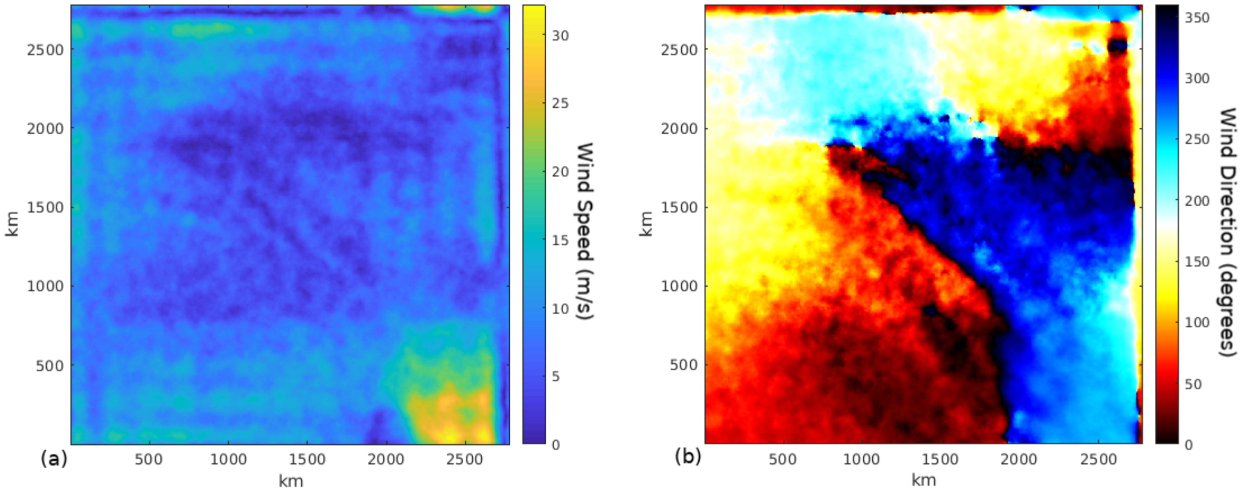

Figure 5.5: The postwhite wind field of the sample ECMWF NWP wind field from Fig. 5.1. (a) shows the wind speed (m/s) and (b) shows the wind direction (degree). Due to the nature of the filters that the wind field passes through, the wind speed is not yet scaled to the proper range of values. The wind field is nominally 2781.25 km × 2781.25 km.

create synthetic wind fields that best match various parameters of measured wind fields on average, independent of location, orientation, and final resolution of the wind field.

The final filter is made by dividing the target frequency spectrum magnitude by the average frequency spectrum magnitude of the wind fields produced by the postwhitening filter (the phase associated with each frequency bin is unaffected by this filter). The target frequency spectrum magnitude for these winds is arbitrarily chosen to correspond to a power law with $\alpha_{2\mathrm{D}} = -8/3$. The DC frequency bin (where $k = 0$) is set so the power in the DC bin is proportional to power in the nearby frequency bins relative to power in the DC bin of the original NWP wind field. The final part of the filter is to scale the wind field so that the average wind speed of the synthetic wind field matches the average wind speed of the low resolution field.

Note that the edges of the high resolution wind field are considered to be less accurate due to the periodic assumption of the 2D DFT in performing the convolution. For this reason, the outside WVCs of the wind field are discarded. The 2D kinetic energy spectrum of a wind field having gone through the final filter with the outer WVCs removed is shown in Fig. 5.7 with the corresponding wind field in Fig. 5.8. To allow a straightforward comparison between the low resolution wind field and the synthetic high resolution wind field, high resolution WVCs are discarded in blocks corresponding to an integer number of low resolution WVCs.

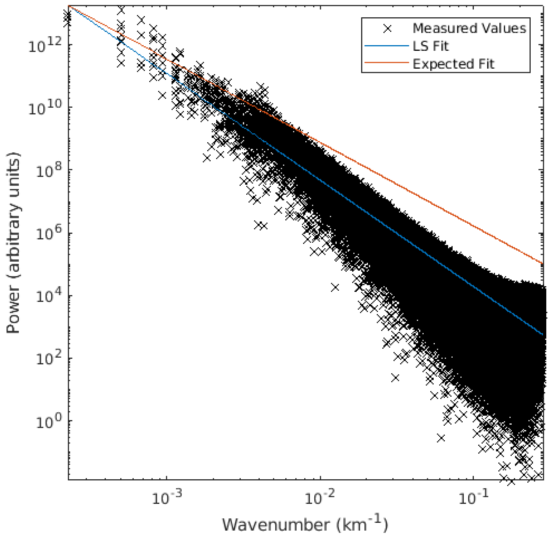

Figure 5.6: The kinetic energy spectrum of the postwhite wind field of the sample ECMWF NWP wind field from Fig. 5.1 with a least squares fit compared to the expected least squares fit. To reduce the effects of spectral leakage the wind field has gone through a circular Tukey window with a 0.2 cosine taper. Note that the measured α_{2D} is -3.40 while the average expected α_{2D} is -2.67.

5.2 Analysis of Process

The aim of the synthetic wind fields creation process is to create high resolution wind fields that are realistic enough to be used in simulation for testing ambiguity selection algorithms. While the validation of the synthetic wind fields is presented in Section 5.3, additional insight on the wind fields that can be gleaned from examining the wind field creation process is presented here.

5.2.1 Process Repeatability

The synthetic wind field creation process is applicable to any ECMWF NWP wind field. Due to the randomization of the extended phase component, the process is not deterministic and running the process multiple times on the same initial low resolution wind field yields small scale feature differences between the final wind fields. These differences do not affect the general wind field characteristics or cause the wind field to appear any more or less realistic. For applications where a variety of general wind field characteristics are desired, such as testing ambiguity selection algorithms, many different NWP wind fields should be used when making the synthetic wind fields.

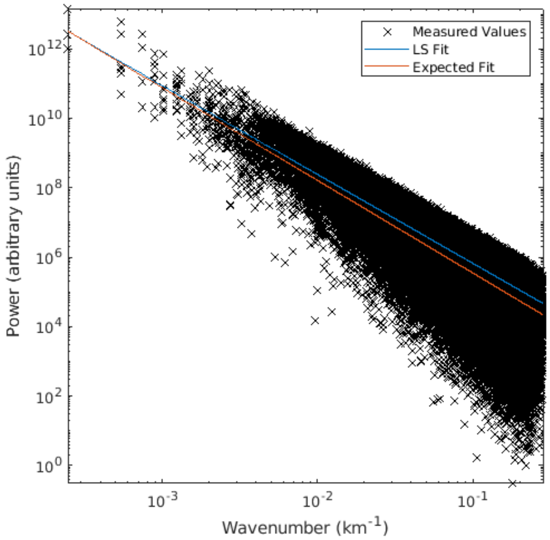

Figure 5.7: The kinetic energy spectrum of the final wind field of the sample ECMWF NWP wind field from Fig. 5.1 with a least squares fit compared to the expected least squares fit. To reduce the effects of spectral leakage the wind field has gone through a circular Tukey window with a 0.2 cosine taper. Note that the measured α_{2D} of -2.56 is close to the average expected α_{2D} of -2.67.

5.2.2 Comparison to the NWP Wind Field

While the general characteristics of the synthetic wind field are based on the low resolution wind field, there are a few notable differences as well. To make a fair comparison between the NWP low resolution wind field and the high resolution synthetic wind field, a low resolution synthetic wind field is found by averaging the wind vectors of all the high resolution WVCs that correspond to each low resolution WVC. As some edge high resolution WVCs are discarded during the wind field creation process, only the middle WVCs of the NWP wind field are compared to the low resolution synthetic wind field. An example of a NWP wind field and its associated low resolution synthetic wind field is given in Fig. 5.9.

The wind characteristics between the two wind fields in Fig. 5.9 are generally the same but with differences most prevalent around the high frequency weather features (i.e., the wind fronts and cyclone). This example was specifically chosen with the high frequency features to highlight these differences, wind fields with fewer high frequency features have fewer visible differences.

Low frequency features dominate the average kinetic energy spectrum over the ocean because high frequency features are less common; the power law that wind fields follow on average

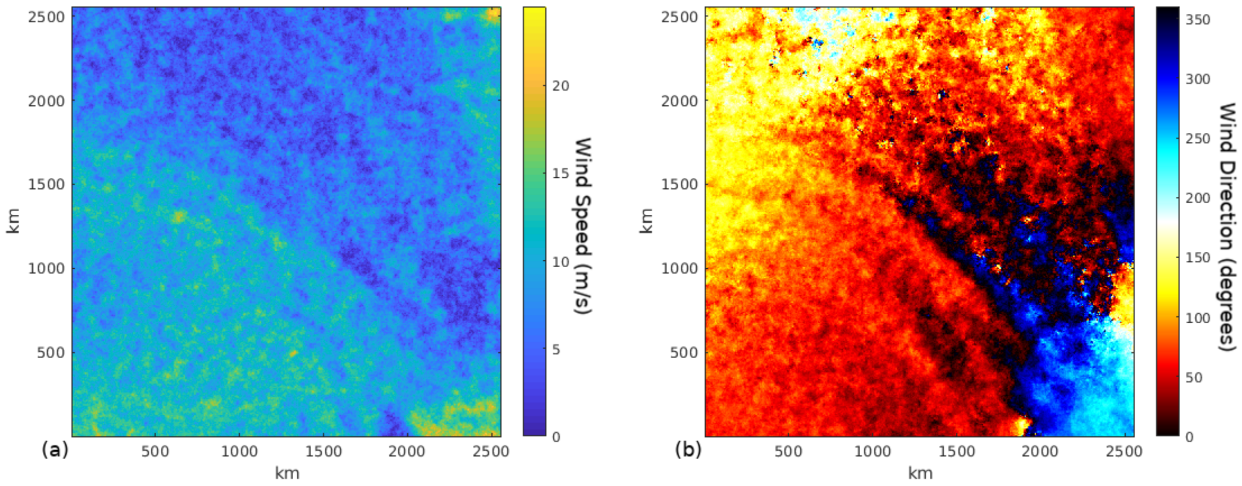

Figure 5.8: The final wind field of the sample ECMWF NWP wind field from Fig. 5.1. (a) shows the wind speed (m/s) and (b) shows the wind direction (degree). The wind field is nominally 2558.75 km × 2558.75 km.

is evidence of this. As the synthetic wind field creation process is based on average wind field characteristics, the synthetic wind fields recreate the more common low resolution features better than the less common high resolution features. In general, the smoother the NWP wind field is, the better the synthetic wind field represents the NWP wind field. This trend implies that the wind fields created by this process are a subset of all possible realistic synthetic wind fields.

Another point of interest is found by comparing the eyes of the cyclones in Fig. 5.9. The eye of the cyclone of the synthetic wind field is shifted to the left and up of the NWP cyclone, a shift of over a hundred kilometers. This shift implies that the synthetic wind field is a high resolution representation of a different, but similar, low resolution wind field than the NWP wind field. For the purpose of testing ambiguity selection algorithms through simulation, this result is acceptable.

5.3 Validation of Wind Fields

Traditional methods of validating wind fields have difficulty being extended to these synthetic wind fields. One such method is to compare a scatterometer-measured wind field to in situ measurements or other scatterometer measurements. Since these synthetic wind fields are not intended to correspond to a real location, there are no in situ or other scatterometer measurements

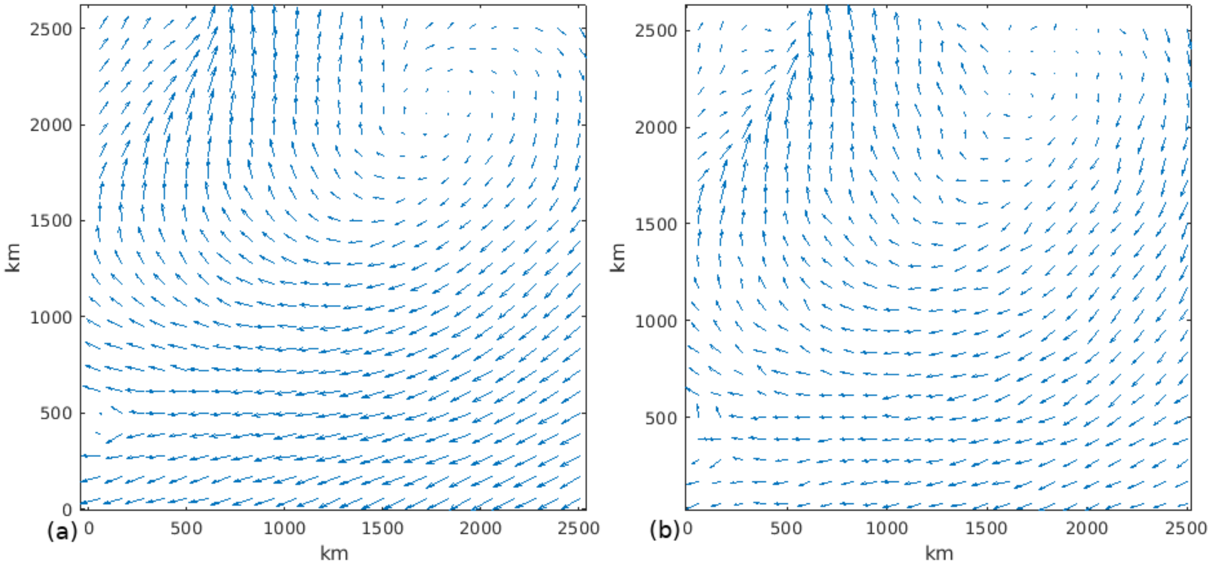

Figure 5.9: Comparison between an original NWP wind field and the low resolution synthetic wind field. (a) shows the original NWP wind field and (b) shows the low resolution synthetic wind field. Note that the low resolution synthetic wind field generally matches the wind characteristics of the original NWP wind field, but it doesn't recreate the high resolution wind fronts and cyclones quite so well. The wind fields are nominally 2558.75 km × 2558.75 km.

to compare them with. Another method of validating wind fields is to measure the average kinetic energy spectra and check it against models and previous measurements. As the wind field creation process forces the wind field's average energy spectra to match the power law predicted by models and confirmed by measurements, this check is guaranteed to pass. While this check indicates that the wind fields may be realistic, kinetic energy spectra are not unique as many wind fields with the same energy spectrum may have varying levels of realism based on other characteristics of the wind field. Instead, other methods must be used to evaluate the synthetic wind fields.

5.3.1 Divergence and Vorticity

As a preliminary check, an experienced eye is able to inspect the synthetic wind fields for extreme wind speeds, misplaced wind vectors against the flow of the wind field, and many other wind features that seem out of place. The synthetic wind fields created by the process explained in this chapter pass the visual inspection sufficiently to appear to be realistic as evidenced by the wind fields shown in Figs. 5.8, 5.10 and 5.11. This visual check is subjective and hard to make repeatable, but many aspects of this check can be accomplished objectively and repeatably by comparing the divergence and vorticity of the wind field.

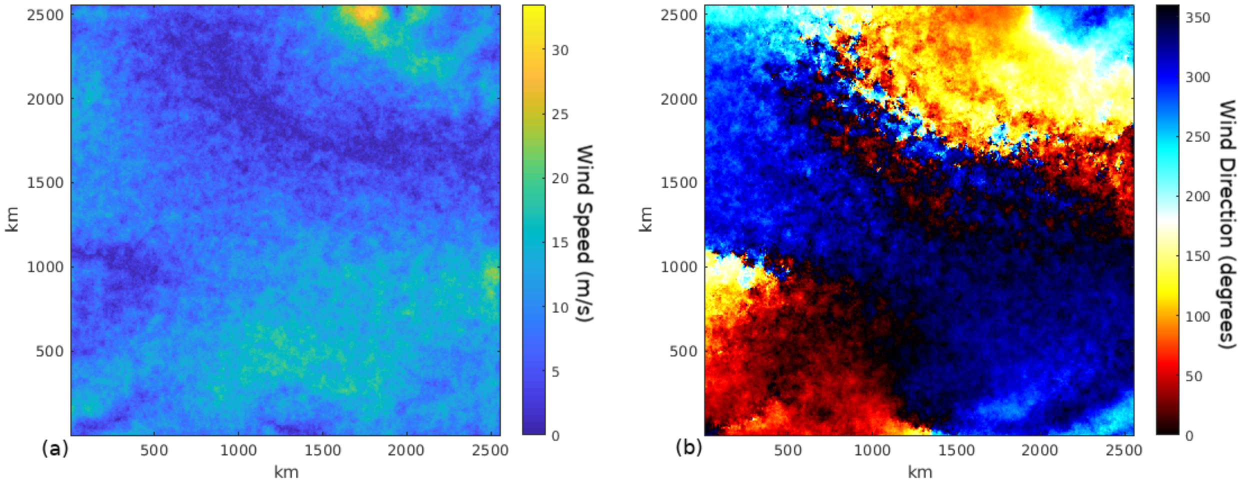

Figure 5.10: An example synthetic wind field. (a) shows the wind speed (m/s) and (b) shows the wind direction (degree). The wind field is nominally 2558.75 km × 2558.75 km. The α_{2D} of this wind field is -2.67.

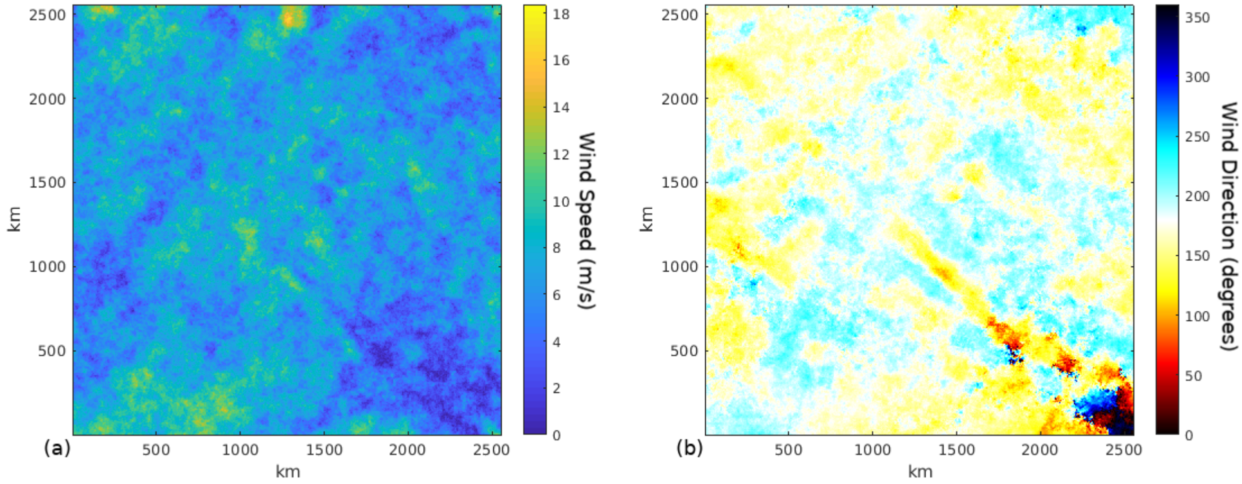

Figure 5.11: Another example synthetic wind field. (a) shows the wind speed (m/s) and (b) shows the wind direction (degree). The wind field is nominally 2558.75 km × 2558.75 km. The α_{2D} of this wind field is -2.60, close to the average α_{2D} of -2.67.

The average spectral shape of the divergence and vorticity of these realistic synthetic wind fields should be proportional and similar in shape to the spectral shape of the divergence and vorticity of measured UHR wind fields. It is also expected that the divergence and vorticity of synthetic wind fields are lower than the divergence and vorticity of the measured UHR wind fields as the measured UHR wind fields are made from noisy measurements.

The average logarithmic divergence and vorticity spectrum of synthetic wind fields along with the average logarithmic divergence and vorticity spectra of QuikSCAT UHR wind fields are shown in Figs. 5.12 and 5.13. The first observation for the synthetic wind fields is that they have a few remaining artifacts near their center that the final filter did not completely take out of the wind fields. Even with these artifacts, the divergence and vorticity spectra are relatively proportional to each other as desired. Also, the divergence and vorticity spectra are relatively circularly symmetric, a desirable quality for an isotropic wind field. When examining the divergence and vorticity spectra for the measured UHR wind fields, they are not circularly symmetric. This can be explained by the orientation that these wind fields were measured on the globe. The UHR wind fields were taken from ascending and descending QuikSCAT passes, so that top to bottom is the east-west axis while left to right is the north-south axis. The greater vorticity and divergence in the north-south axis is indicative of the changes between the westerlies and the trade winds along the north-south axis[1]. Ignoring the non-isotropic portion of the UHR winds and the artifacts left in the synthetic winds, the general shape of the divergence and vorticity spectra matches between the measured and synthetic winds.

Another way to characterize the divergence and vorticity of a wind field is through the root mean square (RMS) divergence and vorticity. For the synthetic and measured UHR wind fields, a histogram of the RMS divergence and vorticity measurements is presented in Figs. 5.14 and 5.15. In terms of histogram shape, the synthetic wind fields somewhat match the shape of the measured UHR wind fields. The UHR wind fields have a higher mean and a few more outliers than the synthetic wind fields, but this is not a surprise considering the noise in the measurements that are part of the UHR wind fields.

---

[1] As mentioned previously the final filter selected for this paper was chosen to create wind fields that are realistic independent of location and orientation on the globe. While not shown here, optimizing the final filter with measured UHR wind field data introduces the same non-isotropic elements into the wind field specific to the measurement geometry of the UHR wind fields.

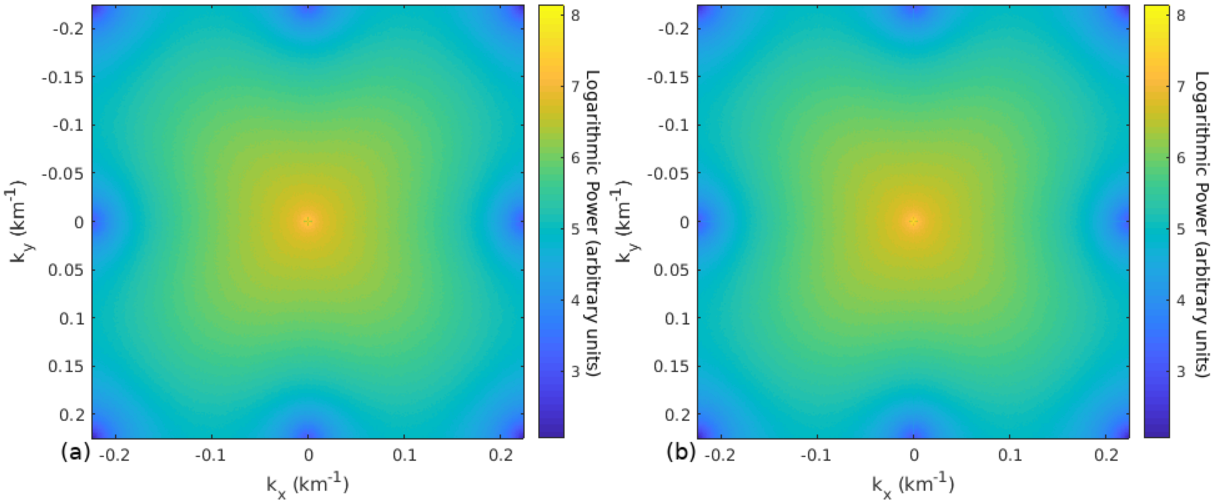

Figure 5.12: The average logarithmic divergence and vorticity spectrum of the synthetic wind fields. (a) shows the divergence spectrum and (b) shows the vorticity spectrum. Color axis values are arbitrary.

Aside from a few artifacts in the average divergence and vorticity spectra of the synthetic wind fields, they match sufficiently with the measured UHR wind fields. Coupling these results with the expected E_{2D} and visual inspection, they are suitably realistic for simulation and testing ambiguity selection algorithms.

5.4 Conclusion

This chapter presents a new synthetic wind field creation process. An explanation of the theory and process is given along with an analysis of the ending wind fields. The wind fields made by this process are found to be visually acceptable, conforming to the expected kinetic energy spectrum, and sufficient in the general divergence and vorticity characteristics found in similar resolution wind fields. These wind fields contribute to the utility of synthetic wind fields as they are sufficiently realistic for simulation-based testing of UHR ambiguity selection algorithms.

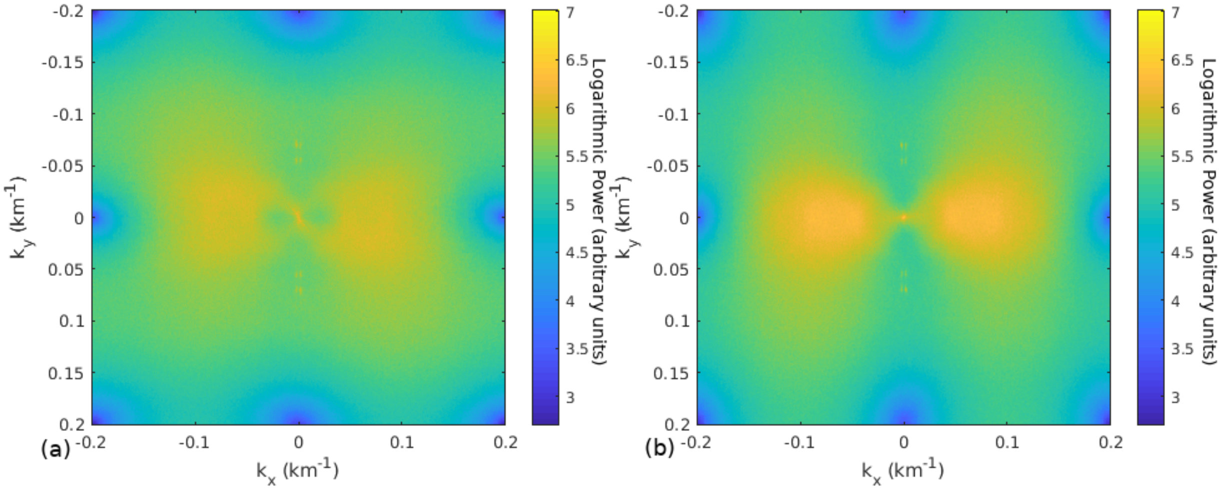

Figure 5.13: The average logarithmic divergence and vorticity spectrum of measured UHR wind fields. (a) shows the divergence spectrum and (b) shows the vorticity spectrum. Color axis values are arbitrary

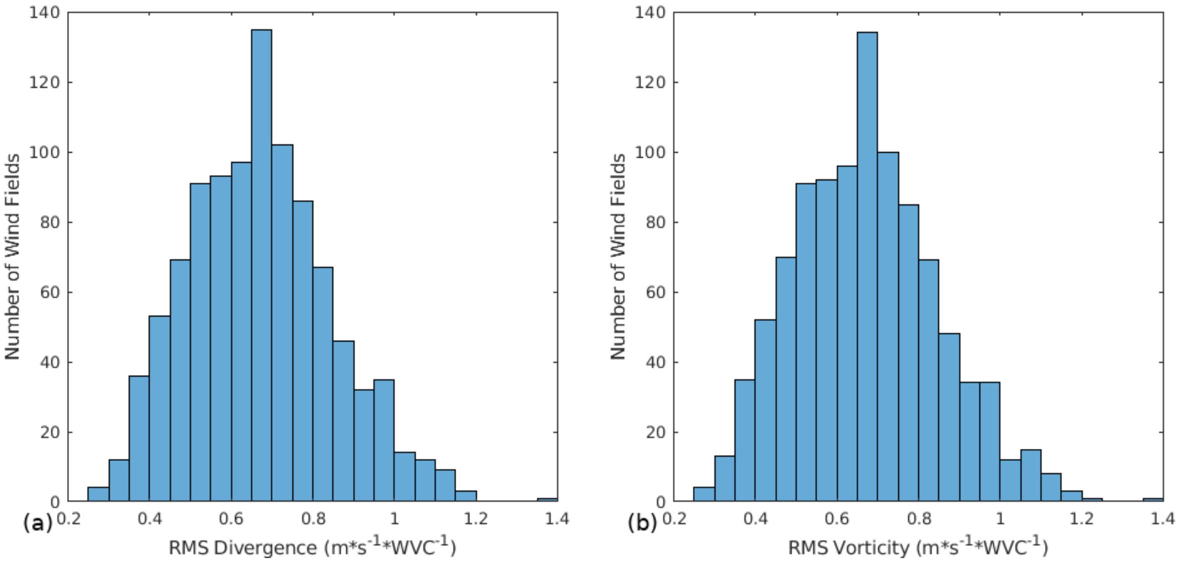

Figure 5.14: Histograms of the RMS divergence and vorticity for the synthetic wind fields. (a) shows the divergence histogram and (b) shows the vorticity histogram. The mean divergence is $1.34 ms^{-1} WVC^{-1}$ and the mean vorticity is $1.41 ms^{-1} WVC^{-1}$.

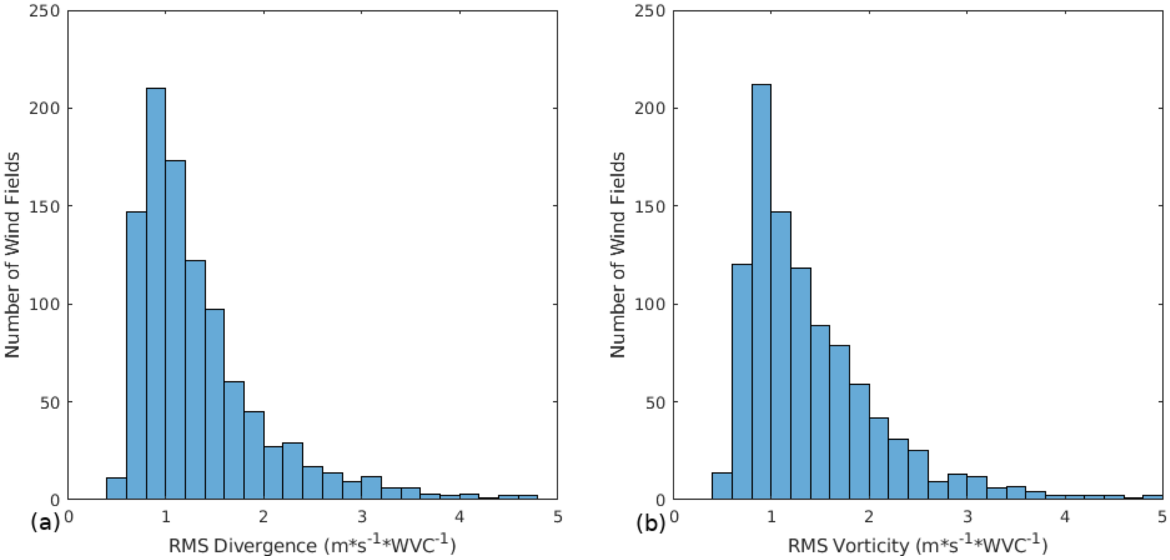

Figure 5.15: Histograms of the RMS divergence and vorticity for measured UHR wind fields. (a) shows the divergence histogram and (b) shows the vorticity histogram. The mean divergence is $0.672\text{ms}^{-1}\text{WVC}^{-1}$ and the mean vorticity is $0.673\text{ms}^{-1}\text{WVC}^{-1}$.

CHAPTER 6. ANALYSIS OF WIND RETRIEVAL AND AMBIGUITY SELECTION ALGORITHMS THROUGH SIMULATION

As previously noted at the beginning of Chapter 5, simulation is a useful method to help evaluate the accuracy of wind retrieval and ambiguity selection algorithms. This chapter continues to develop the improved simulation tool for analyzing wind retrieval and ambiguity selection algorithms and demonstrates the tool for various algorithms.

Along with a visual inspection, three post-simulation performance metrics are presented to analyze the different algorithms. The first two metrics are the wind speed and direction root-mean-square errors (RMSEs) of the estimated wind vector when compared to the true wind vector. The last metric is the percentage of correctly selected ambiguities (sometimes called the "ambiguity removal skill" [14]). These metrics help to analyze wind retrieval and ambiguity selection algorithms both together and separately where possible.

This chapter is organized as follows. Section 6.1 explains the simulation process. Section 6.2 introduces the algorithms that are passed through the simulation as a demonstration of the analysis that is performed with this tool. Section 6.3 examines and analyzes the results of the various QuikSCAT algorithms and Section 6.4 gives a brief conclusion.

6.1 Simulation Process

The simulation process for testing wind retrieval and ambiguity selection algorithms is similar to the simulation process described in Section 4.1.2 for developing the WRF. Only the highlights of the simulation process and differences between this simulation and the WRF simulation are given here.

The starting point for the simulation is a synthetic wind field created by the process described in Chapter 5 and the measurement geometry pulled from a QuikSCAT L1B file. The synthetic wind field provides a wind vector at each SRF pixel for all the slice measurements pulled from the QuikSCAT L1B file (the simulation performs DIB slice and UHR processing, no egg measurements are required). Similar to the simulation in Chapter 4, the simulation uses pixels (UHR WVCs) that are 2.225 km × 2.225 km. Through the QuikSCAT QMOD4 GMF, each wind vector is transformed to a set of σ^0 values based on the measurement geometry. The simulated slice measurements are computed by taking the weighted average of the σ^0 measurements where the weights come from the corresponding slice SRF. Noise is added to each simulated slice measurement.

The wind field is then separately processed using both DIB slice processing and UHR processing. The various wind retrieval algorithms are performed for each WVC with the simulated slice measurements. For convenience, we perform DIB slice processing with WVCs that are 11.125 km × 11.125 km in size so that each DIB WVC is 5 UHR WVCs × 5 UHR WVCs. After wind retrieval, each simulated scatterometer-measured wind field undergoes ambiguity selection.

Simulation of very large wind fields is time-intensive. To reduce the computation time, we only simulate a relatively small section of a full QuikSCAT swath stored in an L1B file. For the swath section used in this study, approximately 87% of WVCs in the section are from the sweet spot, 7% of the WVCs are from the far swath, and 6% of the WVCs are from the nadir region. Note that in a typical QuikSCAT cross track slice, approximately 67% of WVCs are from the sweet spot, 22% of WVCs are from the far swath, and 11% of WVCs are from the nadir region. While the distribution of WVCs in the simulation may not accurately depict the distribution of WVCs in a full swath, the section is sufficient to demonstrate the tool's ability to compare algorithms and to see trends within each swath region.

When pulling data from the QuikSCAT L1B file, the measurement geometry is interpolated to 2.225 km × 2.225 km resolution to assist in the simulation. This interpolation is done on a latitudinal-longitudinal grid rather than a swath-based grid. For the swath section used in this study, this puts the far swath in the top left corner and the nadir region in the bottom right corner as depicted in Fig. 6.1.

6.2 Algorithm Description

We demonstrate how the simulation tool is used by analyzing two different pairs of wind retrieval and ambiguity selection algorithms. The final products of each algorithm pair are a DIB

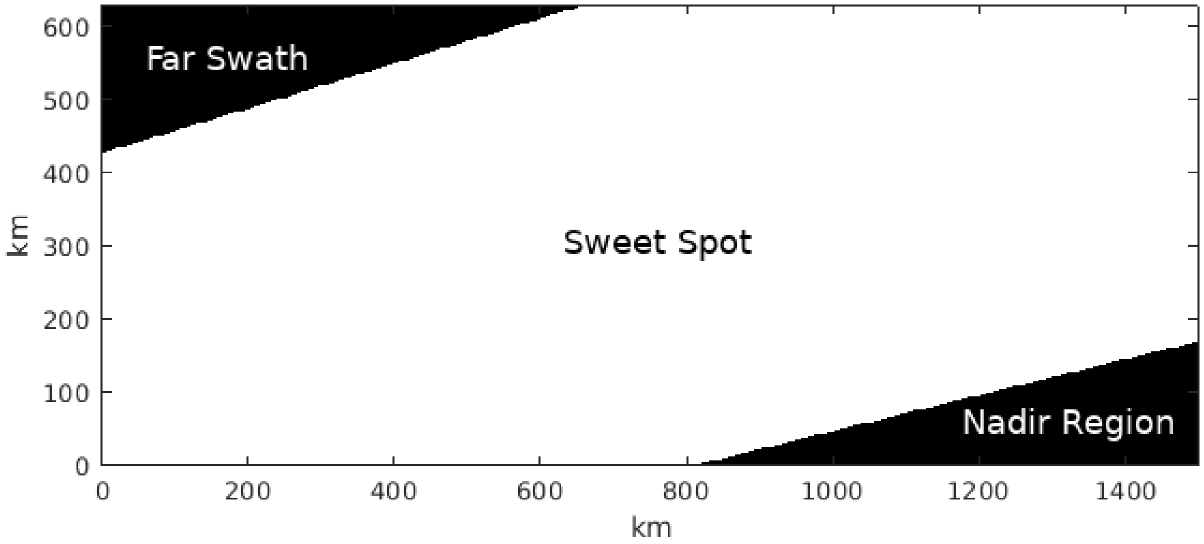

Figure 6.1: The swath area used in this study separated into the far swath, sweet spot, and nadir region.

wind field at 11.125 km × 11.125 km resolution and a UHR wind field at 2.225 km × 2.225 km resolution (for convenience in this chapter, these resolutions are referred to as DIB resolution and UHR resolution respectively).

The first algorithm pair is referred to as basic wind estimation. Basic wind estimation performs wind retrieval without DIRTH as described in Section 2.2.2 and performs ambiguity selection with the median filer algorithm as described in Section 2.2.3 for both DIB slice processing and UHR processing. To nudge the DIB wind field, a 111.25 km× 111.25 km "artificial NWP" wind field is created by taking a moving average of the synthetic wind field and downsampling the result. The artificial NWP wind field is linearly interpolated to the spatial resolution of the DIB wind field, and the DIB wind field is nudged by the linearly interpolated wind field. Similar to QuikSCAT L2B processing, only the first or second most likely ambiguities are selected during the nudging process [14]. The median filter algorithm is run on the DIB wind field with a fixed window size of 7 WVCs × 7 WVCs. The median filtered DIB wind field is then used to nudge the UHR wind field. As is common in UHR processing, the ambiguity for each WVC that is closest in direction to the nearest DIB wind vector is selected as the initial guess; any ambiguity may be selected when nudging the UHR wind field, i.e., not TN. The median filter algorithm is run on the nudged UHR wind field, with a fixed window size of 7 WVCs × 7 WVCs, to produce the final wind field for basic wind estimation.

The second algorithm pair we refer to as DIR wind estimation. DIR wind estimation performs wind retrieval with DIR (but not TN) and performs the modified median filter algorithm for ambiguity selection as described in Section 2.2.4. Similar to basic wind estimation, the DIB wind field is used to nudge the UHR wind field in DIR wind estimation. The process of nudging each wind field is identical for both basic and DIR wind estimation and both implementations of the modified median filter algorithm use a fixed window size of 7 WVCs × 7 WVCs. While DIR is used for both QuikSCAT L2B and UHR data products [14], [36], it has not been previously validated for QuikSCAT UHR data products.

6.3 Algorithm Analysis

To demonstrate the analysis tool, we run 14 different UHR resolution synthetic wind fields through the simulation. These synthetic wind fields come from the synthetic wind field creation process described in Chapter 5. To assist in calculating the performance metrics for the DIB algorithms, a DIB resolution version of the synthetic wind fields is found by passing the synthetic wind fields through a moving average filter and downsampling the result. An example synthetic wind field at UHR resolution is shown in Fig. 6.2. This wind field is shown as it contains features with both high and low frequency spatial variability. The example synthetic wind field at DIB resolution is shown in Fig. 6.3.

Figs. 6.4 and 6.5 show the output UHR resolution wind fields for basic wind estimation and DIR wind estimation respectively. Both wind fields recreate the general wind flow of the original synthetic wind field well. There is very little noticeable variation in the selected wind speed between the two algorithms, most of the differences arise in the wind direction. The wind direction returned by the basic wind estimation algorithms has more high frequency spatial variation than the wind direction returned by the DIR wind estimation algorithms, but this variability appears to be more connected to noise than the high frequency variation in the initial synthetic wind field. In some areas, such as the nadir region in the lower right corner of the wind field, this noise in the basic wind estimation wind field has contributed to resolving some small wind features that are not in the original wind field. The DIR wind estimation wind field does not contain these features.

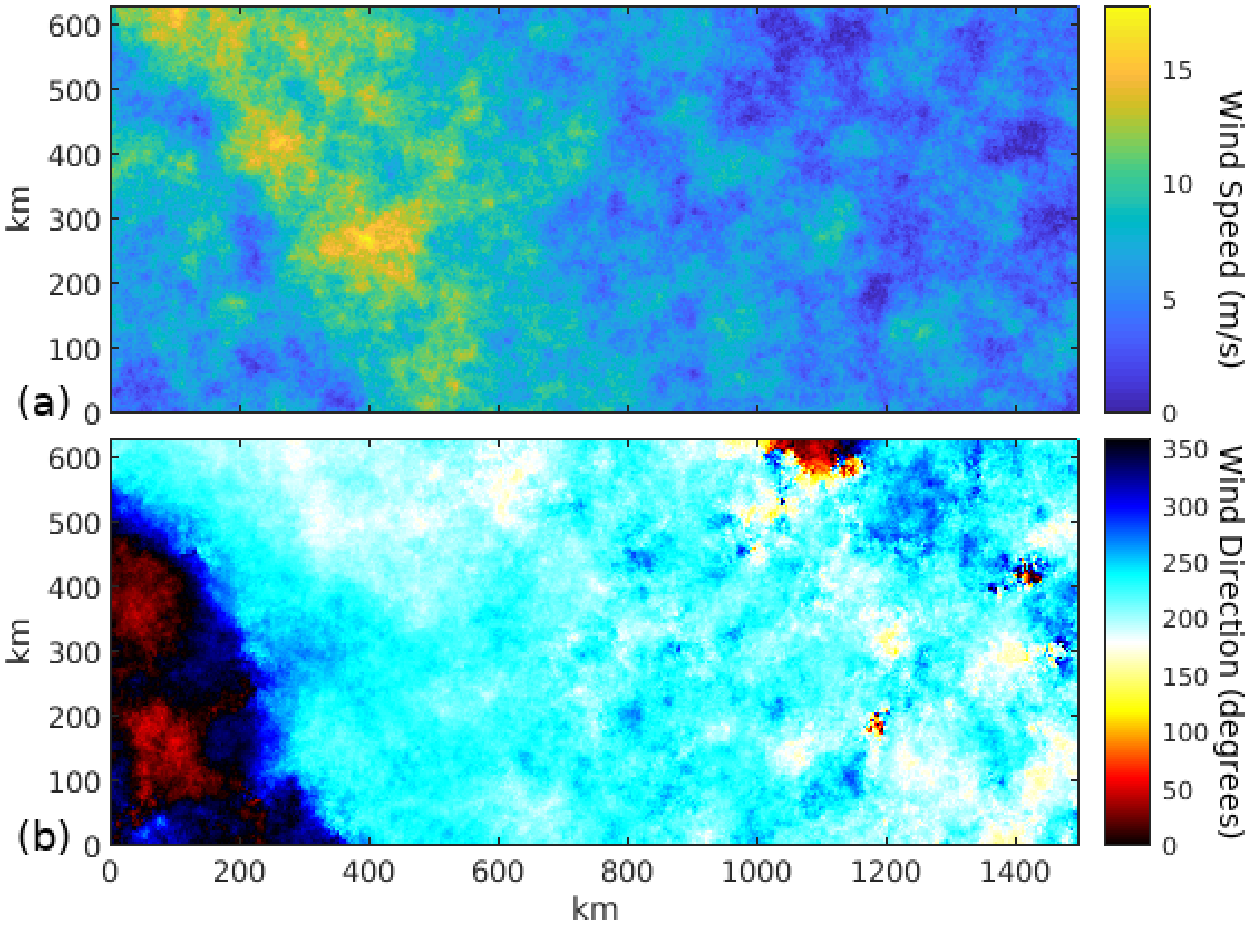

Figure 6.2: One of the synthetic wind fields that is passed through the simulation. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees). This wind field contains areas of both high and low frequency spatial variability. The wind field is 629.675 km × 1499.65 km.

In general, DIR wind estimation appears to return a smoother direction component than basic wind estimation. This result is not surprising as the kinetic energy spectrum of QuikSCAT L2B wind fields with DIRTH follows a steeper power law than the spectrum without DIRTH [16]. While it is possible that DIR may smooth over high resolution features that are present in the original wind field, there is no evidence in Figs. 6.4 and 6.5 that basic wind estimation is capable of resolving any features from the original wind field that DIR wind estimation cannot.

The visual inspection of the UHR resolution wind fields indicates that current QuikSCAT wind retrieval and ambiguity selection algorithms have a limit in their ability to resolve extremely high frequency wind features. This result mirrors the presence of the apparent spectrum floor at high frequency values in the kinetic energy spectrum of scatterometer-measured winds discussed

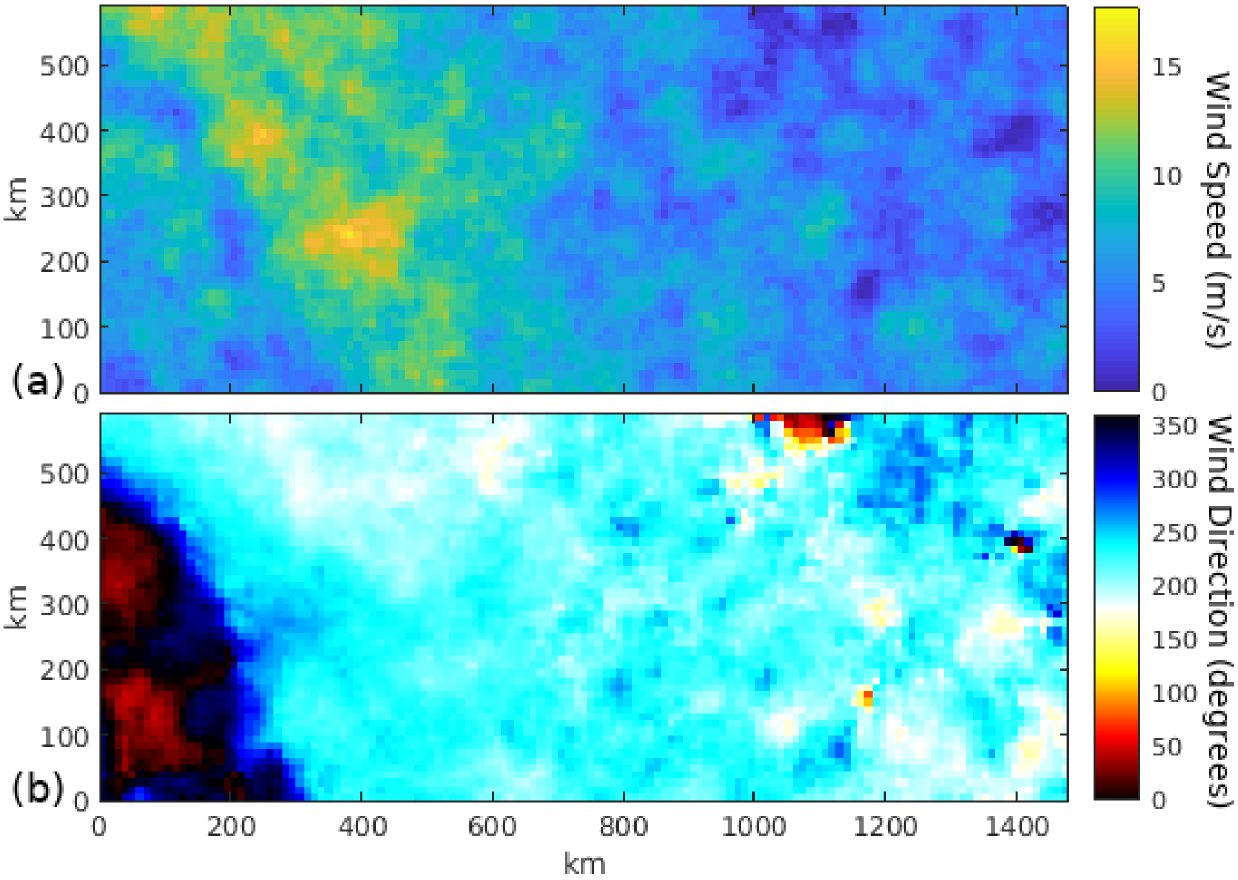

Figure 6.3: The example synthetic wind field shown in Fig. 6.2 at DIB resolution. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees).

in Chapter 3 and the spatial resolution of UHR processed winds being larger than a single WVC, as found in Chapter 4.

The output DIB resolution wind fields for basic wind estimation and DIR wind estimation are shown in Figs. 6.6 and 6.5 respectively. These wind fields yield similar conclusions in regards to basic wind estimation and DIR wind estimation as the UHR resolution wind fields. When comparing Figs. 6.4 and 6.5 with Figs. 6.6 and 6.7, two observations are apparent. The first observation is that UHR processing resolves wind features at a finer resolution than DIB processing does. This is evidence of the result found in Chapter 4 that the spatial resolution of UHR processing is finer than the spatial resolution of DIB processing. The second observation is that the cost of this finer resolution comes at the cost of higher noise. These observations are consistent with known properties of resolution enhancement algorithms for scatterometers [10].

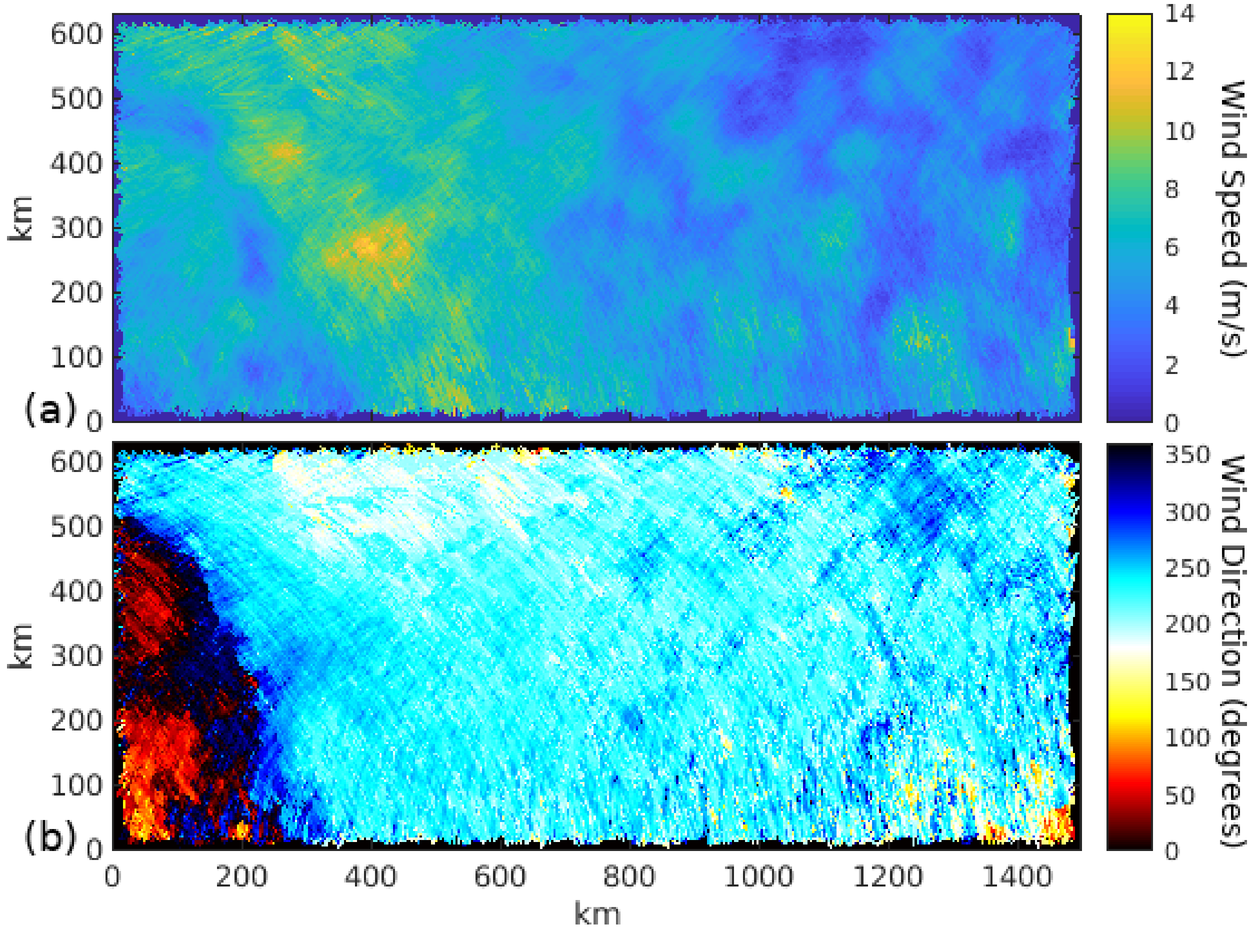

Figure 6.4: The basic wind estimation output at UHR resolution for the example synthetic wind field in Fig. 6.2. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees).

The performance metrics are found for all estimated wind vectors in the 14 synthetic wind fields whose retrieved wind speed is above 3 m/s (the threshold speed where the initial QuikSCAT mission requirements began [9]). The wind speed and direction RMSEs are calculated for both algorithm pairs by comparing the wind vectors estimated by the basic and DIR wind estimation algorithms with the true wind vectors from the original wind fields. The wind speed RMSE value for both the basic and DIR wind estimation algorithms is roughly equal to $\sim$2 m/s for the UHR resolution winds and $\sim$1.6 m/s for the DIB resolution winds. This matches the similarities seen in the visual inspection of the wind fields. For the wind direction RMSE, the DIR wind estimation value outperforms the basic wind estimation value by $\sim 10^\circ$ for the UHR resolution winds and $\sim 12^\circ$ for the DIB resolution winds. To find the ambiguity removal skill, we determine the ambiguity for

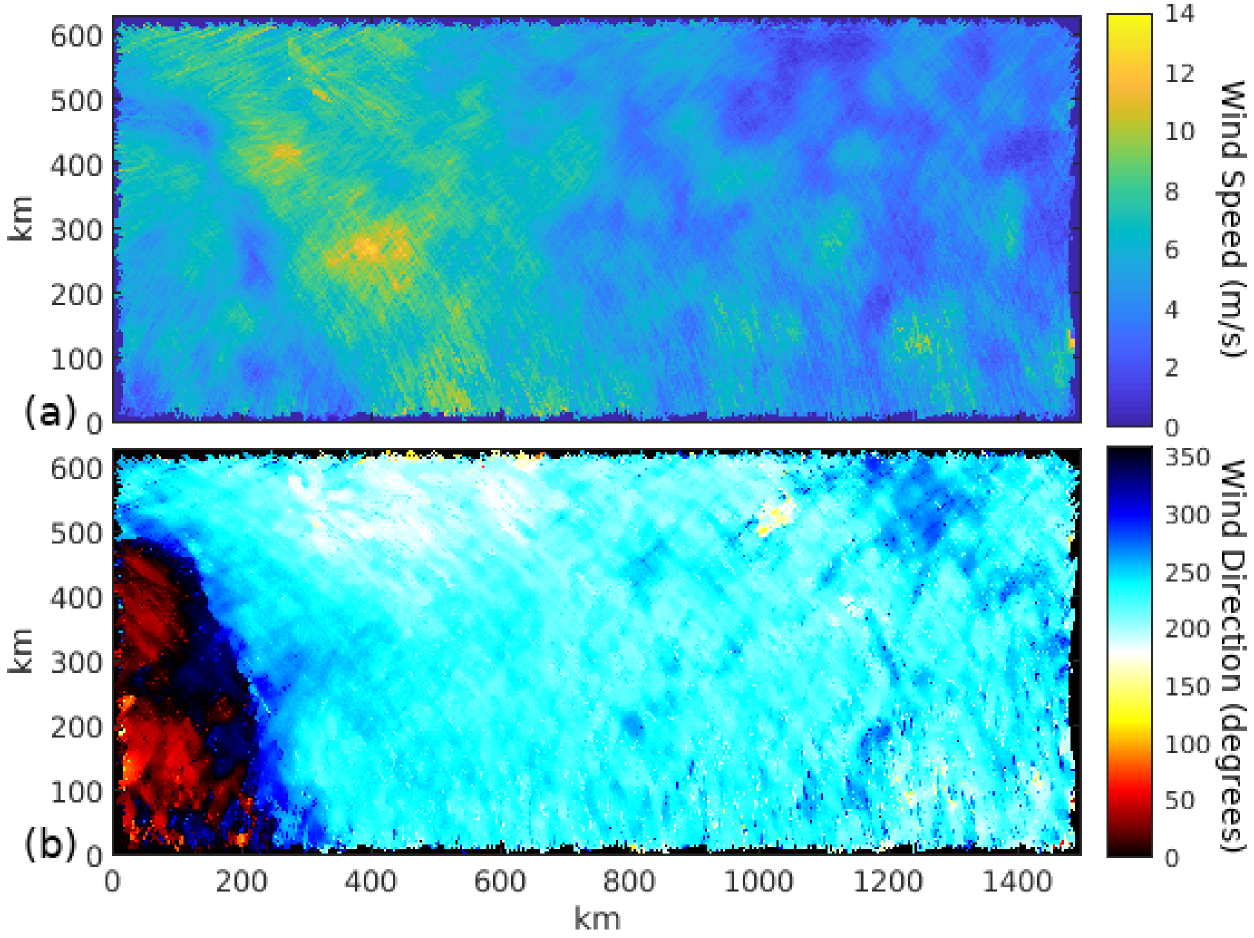

Figure 6.5: The DIR wind estimation output at UHR resolution for the example synthetic wind field in Fig. 6.2. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees). The wind field is 629.675 km × 1499.65 km.

each WVC whose direction is closest to the direction of the corresponding wind vector from the original wind field. The wind fields corresponding to the ideal ambiguity selections for basic wind estimation at UHR resolution, DIR wind estimation at UHR resolution, basic wind estimation at DIB resolution, and DIR wind estimation at DIB resolution are shown in Figs. 6.8, 6.9, 6.10, and 6.11, respectively. While these wind fields help to determine the ambiguity removal skill for a given ambiguity selection algorithm, they also show the theoretical best case wind fields for the ambiguities returned from the wind retrieval algorithms. For this reason, the wind speed direction RMSEs metrics are found for these ideal wind fields.

The performance metrics for basic and DIR wind estimation at DIB and UHR resolutions along with the ideal ambiguity selections at DIB and UHR resolutions are summarized in Table 6.1.

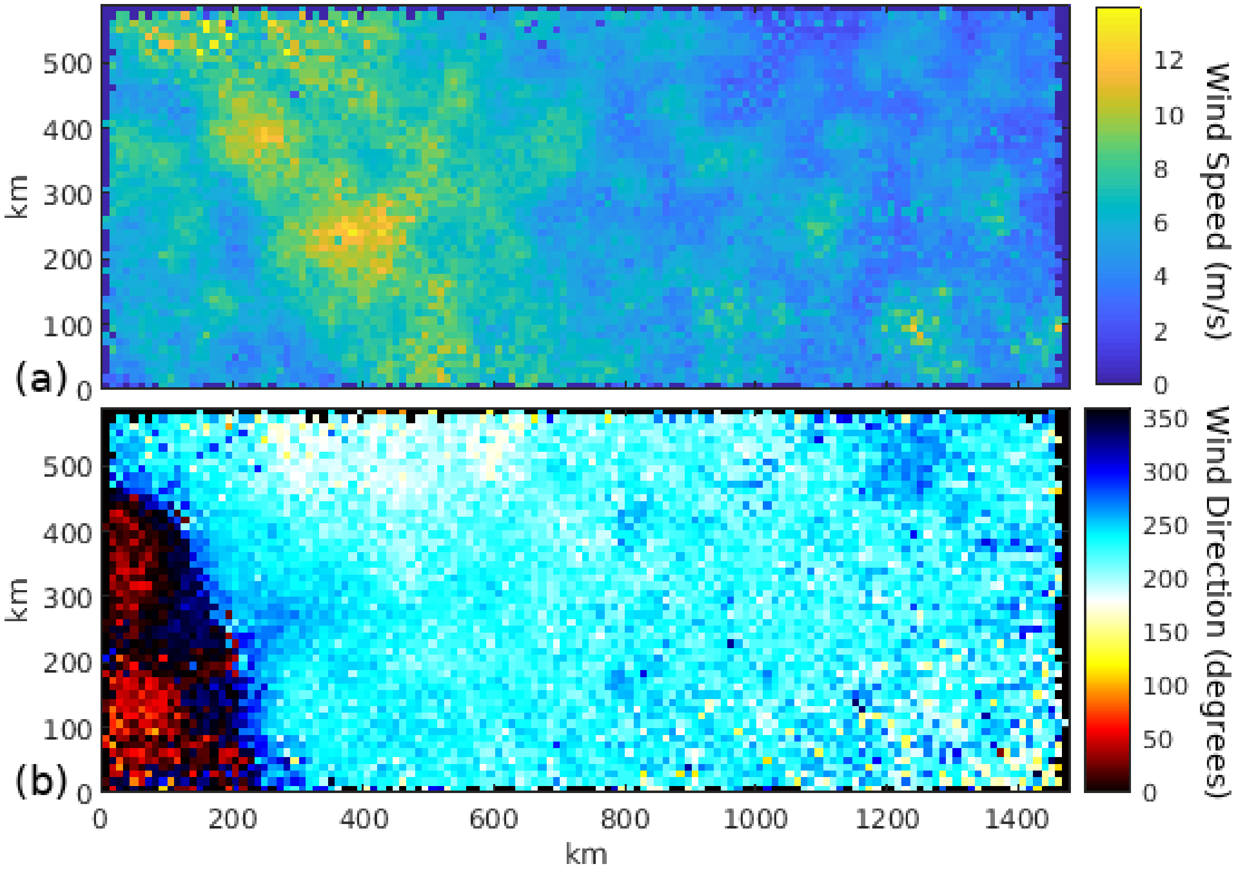

Figure 6.6: The basic wind estimation output at DIB resolution for the example synthetic wind field in Fig. 6.2. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees).

The performance metrics are given for the entire wind field as well as for each section of the QuikSCAT swath. As the majority of the wind field contains WVCs from the sweet spot (see Fig. 6.1), the performance metrics of the entire wind field are closest in value to the performance metrics of the sweet spot. While only 14 wind fields are used to find these performance metrics, the calculated values are sufficient to exhibit the general trends between the various algorithms and processing types.

Congruent with the visual inspection of the wind fields, the performance metrics in Table 6.1 indicate that UHR processing is noisier than DIB processing. However, UHR processing handles the additional noise well as the wind speed RMSE only increases by ∼0.5 m/s and the wind direction RMSE only increases by ∼1° for basic wind estimation and ∼3-4° for DIR wind

Table 6.1: The average performance metrics of basic and DIR wind estimation at DIB and UHR resolutions and the ideal ambiguity selections at DIB and UHR resolutions. The performance metrics are given for the full wind field selection and broken down by swath location. Due to the full wind field selection being only a section of the QuikSCAT swath on a latitudinal-longitudinal grid, the full wind field underrepresents the WVCs from the far swath and nadir region.

Location	Metric	Basic	Ideal Basic	DIR	Ideal DIR
UHR Resolution					
Full Field	Wind Speed RMSE	1.98 m/s	2.02 m/s	1.99 m/s	2.02 m/s
	Wind Direction RMSE	35.70°	26.51°	25.74°	14.44°
	Ambiguity Removal Skill	87.81%	100%	85.76%	100%
Far Swath	Wind Speed RMSE	2.76 m/s	2.77 m/s	2.77 m/s	2.77 m/s
	Wind Direction RMSE	27.50°	22.63°	22.94°	12.88°
	Ambiguity Removal Skill	91.58%	100%	88.34%	100%
Sweet Spot	Wind Speed RMSE	1.92 m/s	1.96 m/s	1.93 m/s	1.96 m/s
	Wind Direction RMSE	34.82°	25.31°	24.74°	13.71°
	Ambiguity Removal Skill	87.85%	100%	86.22%	100%
Nadir Region	Wind Speed RMSE	1.88 m/s	1.88 m/s	1.87 m/s	1.89 m/s
	Wind Direction RMSE	52.86°	43.16°	39.87°	23.87°
	Ambiguity Removal Skill	83.03%	100%	75.82%	100%
DIB Resolution					
Full Field	Wind Speed RMSE	1.57 m/s	1.58 m/s	1.57 m/s	1.60 m/s
	Wind Direction RMSE	34.44°	26.05°	22.82°	10.54°
	Ambiguity Removal Skill	89.43%	100%	84.66%	100%
Far Swath	Wind Speed RMSE	2.24 m/s	2.24 m/s	2.25 m/s	2.25 m/s
	Wind Direction RMSE	31.75°	27.76°	22.99°	14.96°
	Ambiguity Removal Skill	93.34%	100%	85.42%	100%
Sweet Spot	Wind Speed RMSE	1.52 m/s	1.54 m/s	1.52 m/s	1.56 m/s
	Wind Direction RMSE	33.62°	24.74°	22.16°	9.55°
	Ambiguity Removal Skill	89.39%	100%	85.06%	100%
Nadir Region	Wind Speed RMSE	1.61 m/s	1.58 m/s	1.60 m/s	1.62 m/s
	Wind Direction RMSE	49.09°	42.78°	32.64°	19.00°
	Ambiguity Removal Skill	86.05%	100%	76.46%	100%

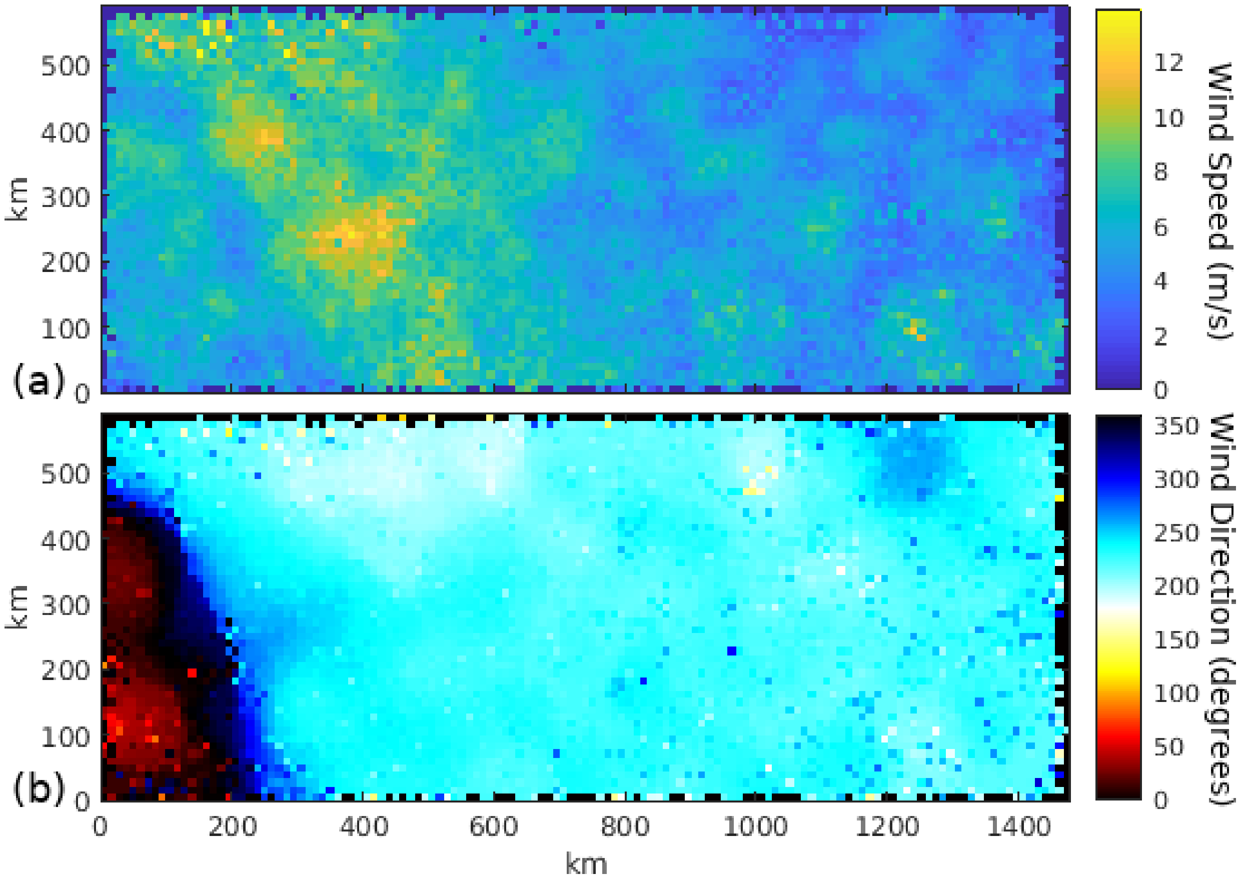

Figure 6.7: The DIR wind estimation output at DIB resolution for the example synthetic wind field in Fig. 6.2. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees).

estimation. These findings suggest that the improved resolution offered by UHR processing only increases the effect of noise a relatively small amount.

No matter what algorithm pair is used, the wind direction RMSE is worst in the nadir region and best in either the far swath or sweet spot. In a similar manner, the wind speed RMSE is worst in the far swath and best in the nadir region or sweet spot. While the sweet spot does not always have the lowest wind speed RMSE value or wind direction RMSE value, its values are always on the lower end of the range for both. This result supports the observation that wind estimation tends to be the best in the sweet spot.

The performance metrics in Table 6.1 indicate that optimizing ambiguity selection for the lowest wind direction RMSE does not optimize ambiguity selection for the lowest wind speed RMSE. However, all ambiguities for a given WVC tend to have similar wind speeds and the wind

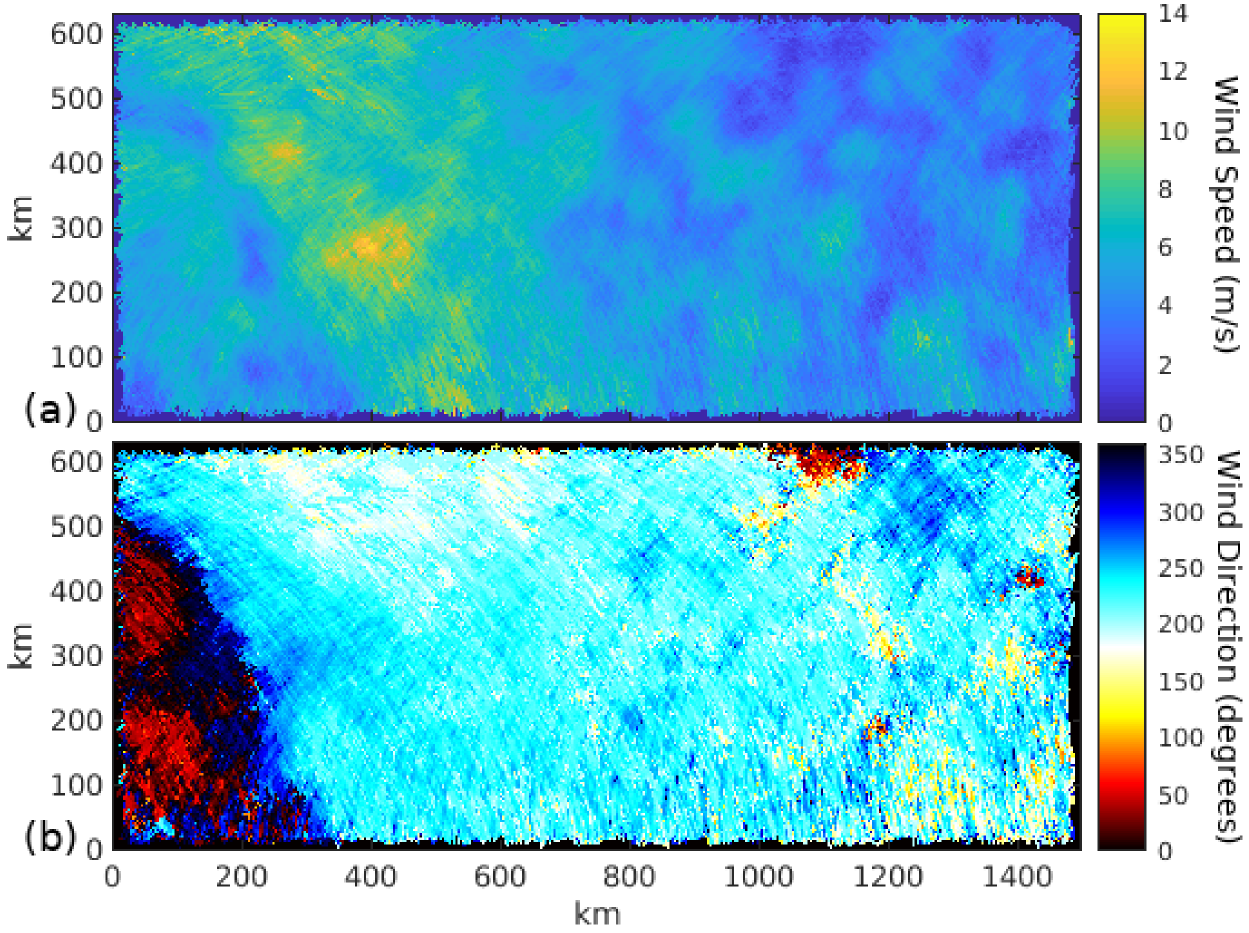

Figure 6.8: The ideal basic wind estimation output at UHR resolution for the example synthetic wind field in Fig. 6.2. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees).

speed RMSE value is still around or under the initial QuikSCAT mission requirements' wind speed RMSE value of 2 m/s for winds between 3 m/s and 20 m/s. Optimizing for wind direction to get closer to the initial QuikSCAT mission requirements' wind direction RMSE value of 20° makes more sense than slightly lowering the wind speed RMSE. Note that the initial QuikSCAT mission requirements were designed for 25 km × 25 km WVCs [9]; achieving those requirements with 11.125 km × 11.125 km WVCs or 2.225 km × 2.225 km WVCs is not a trivial task.

While basic and DIR wind estimation are comparable for wind speed RMSE, DIR wind estimation outperforms basic wind estimation for wind direction RMSE in all swath locations and for both UHR and DIB processing. In general, DIR wind estimation is even better than basic wind retrieval with the ideal ambiguities selected. This suggests that scatterometer-measured wind

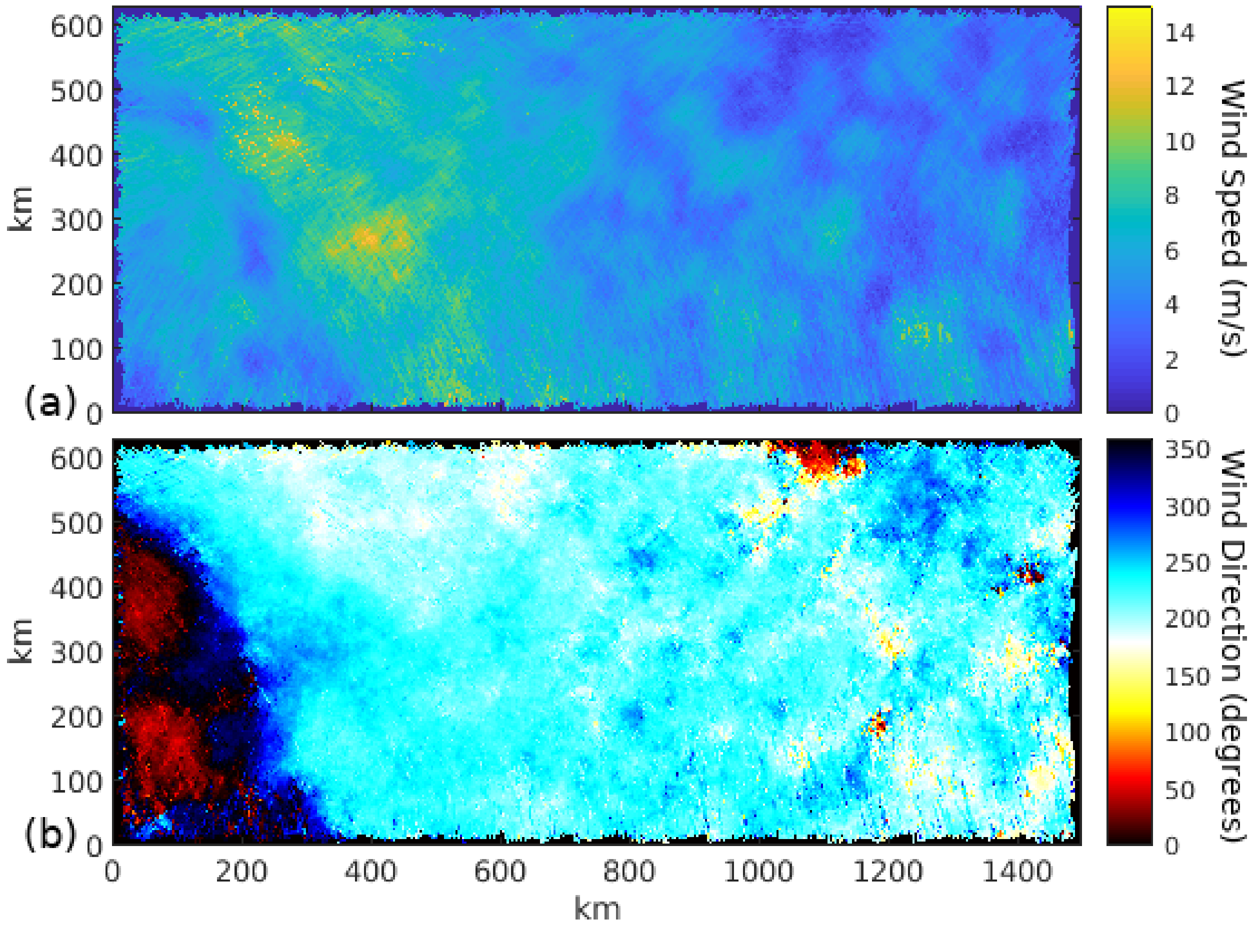

Figure 6.9: The ideal DIR wind estimation output at UHR resolution for the example synthetic wind field in Fig. 6.2. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees).

estimates cannot be improved by better ambiguity selection algorithms for basic wind retrieval: a new ambiguity selection algorithm compatible with DIR wind retrieval or a new wind retrieval algorithm is needed to improve scatterometer-measured wind estimates. Basic wind estimation has a slightly higher ambiguity skill than DIR wind estimation which indicates that the median filter algorithm is a better fit for basic wind retrieval than the modified median filter algorithm is for DIR wind retrieval.

6.4 Conclusion

This chapter demonstrates the simulation analysis tool that compares the effectiveness of wind retrieval and ambiguity selection algorithms at DIB and UHR resolution. The simulation

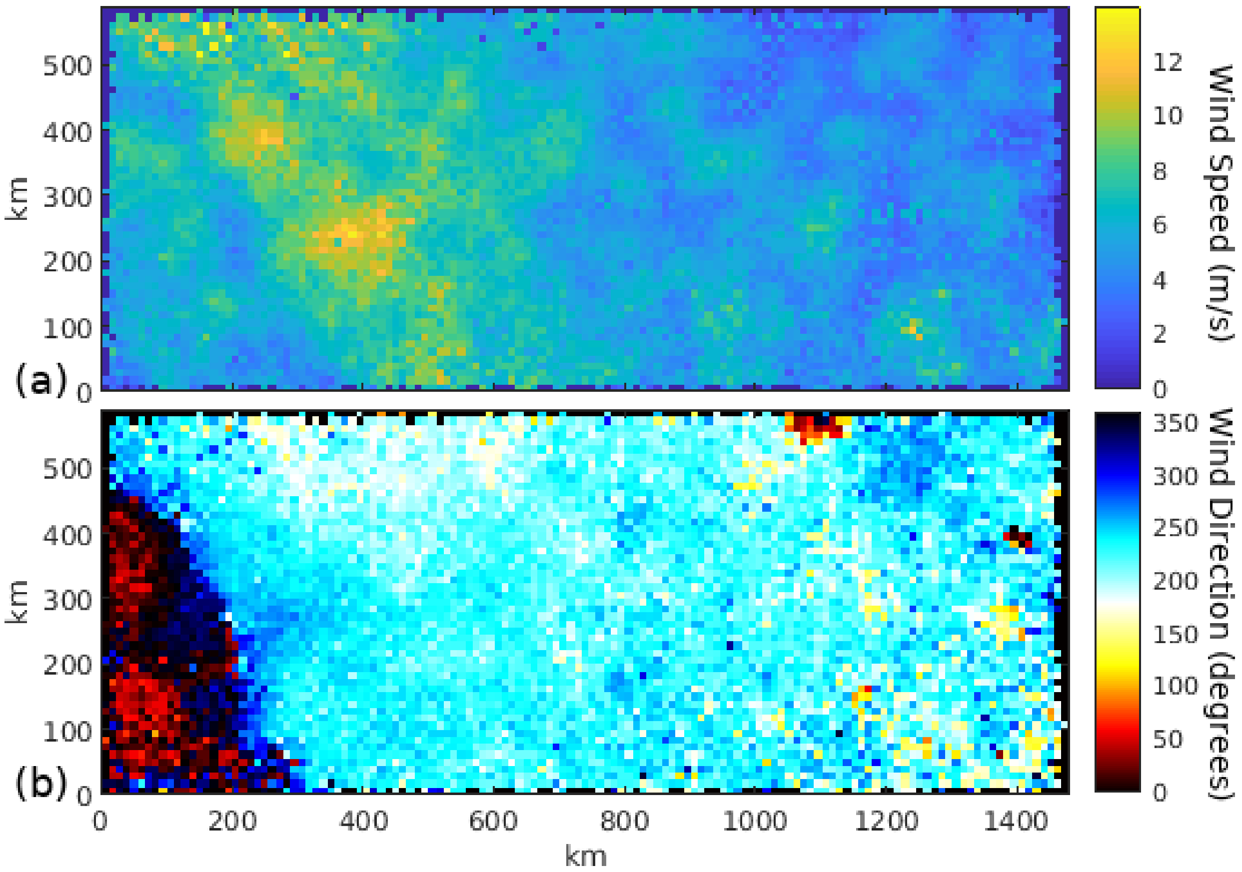

Figure 6.10: The ideal basic wind estimation output at DIB resolution for the example synthetic wind field in Fig. 6.2. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees).

analysis tool synthesizes σ^0 measurements from a synthetic wind field and QuikSCAT L1B file, applies a wind retrieval and an ambiguity selection algorithm, and calculates several performance metrics on the resulting wind fields. The simulation analysis tool also allows for a visual inspection of the resulting wind fields. By changing which wind retrieval and ambiguity selection algorithms are applied by the simulation tool, the performance of various wind retrieval and ambiguity selection algorithms can be compared.

To demonstrate the effectiveness of the simulation analysis tool in comparing wind estimation algorithms, the simulation analysis tool is applied with basic wind estimation and DIR wind estimation. Both the visual inspection and performance metrics indicate that DIR wind estimation is more effective than basic wind estimation at estimating the wind field and that winds produced by UHR processing are noisier than winds produced by DIB processing. The performance metrics

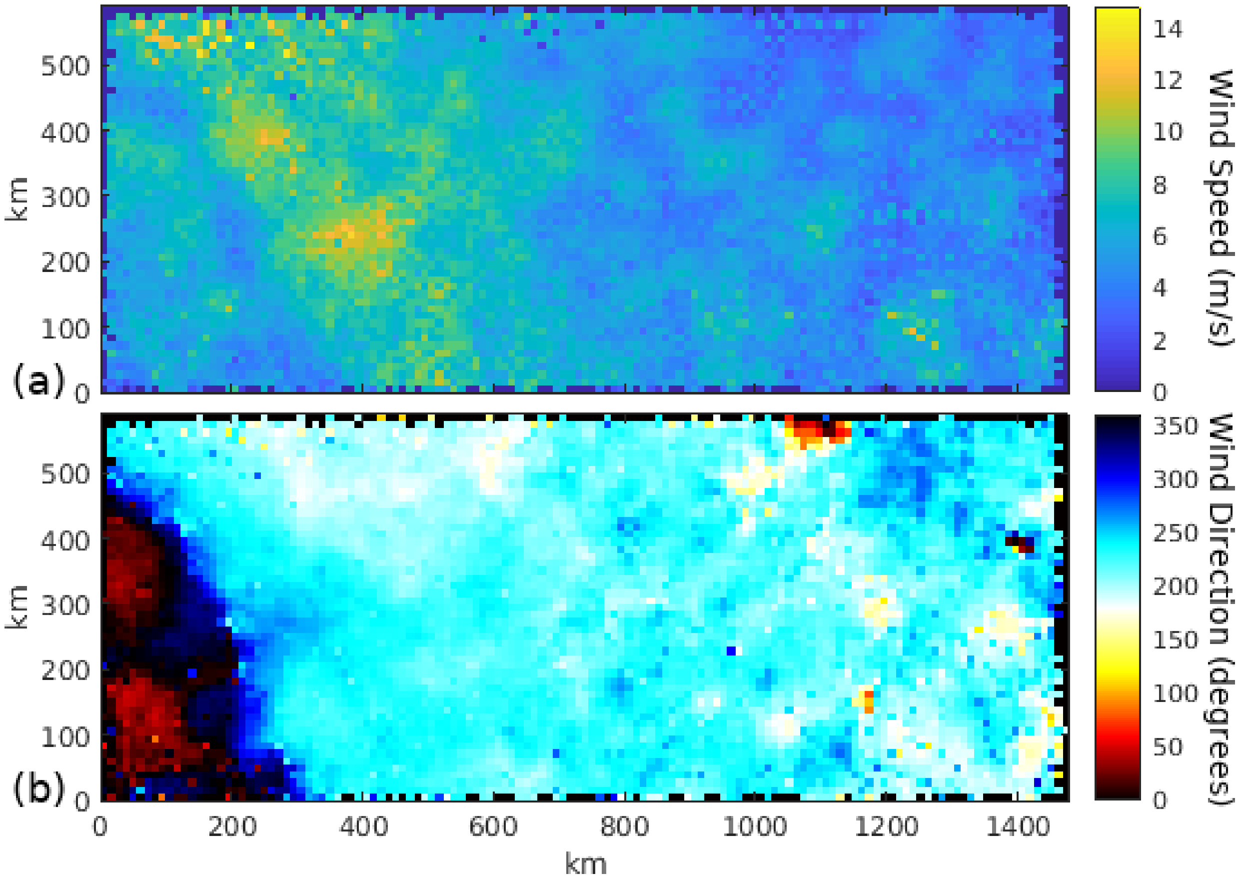

Figure 6.11: The ideal DIR wind estimation output at DIB resolution for the example synthetic wind field in Fig. 6.2. (a) shows the wind speed (m/s), and (b) shows the wind direction (degrees).

suggest that the increase in the effect of noise in UHR processing is slight in comparison to DIB processing since the wind direction RMSE for UHR processing is at most only $\sim$ 3-4° above the wind direction RMSE for DIB processing and the wind speed RMSE only increases by $\sim$ 0.5 m/s. Additional verification is provided that wind estimation is better in the sweet spot than in the far swath or nadir region. The performance metrics indicate that either improving the DIR ambiguity selection or creating a new wind retrieval algorithm is required for the UHR resolution winds to achieve the initial QuikSCAT mission requirements defined for low resolution wind fields.

CHAPTER 7. CONCLUSION

7.1 Summary and Conclusions

This thesis explores several improved analysis techniques for scatterometer wind estimation. Chapter 1 introduces scatterometer wind estimation and provides motivation for improved analysis techniques. Chapter 2 contains necessary background information about scatterometers, wind estimation, wind field analysis, scatterometer simulation, and synthetic wind fields to put the rest of the thesis in context.

Chapter 3 presents a method of examining the kinetic energy spectrum in 2D. Similar to 1D analysis of the kinetic energy spectrum, the method describes the average 2D kinetic energy spectrum in terms of a power law. The power law of the 2D analysis method is compared to the power law found by 1D analysis methods to validate the method. The 2D analysis technique is shown to be sufficient for estimating the 2D kinetic energy spectrum.

Chapter 4 introduces the concept of the WRF as an expansion of the σ^0 measurement SRF. The calculation of the WRF is detailed. This analysis method helps to define the spatial resolution of different wind retrieval algorithms. Generalized WRFs for three different types of QuikSCAT data processing methods are calculated.

The last analysis technique is developed and demonstrated in Chapters 5 and 6. Chapter 5 explores a process to create a synthetic wind field sufficiently realistic to be used in evaluating ambiguity selection algorithms. The wind fields created by the process are verified through divergence and vorticity measurements. These wind fields are used as input wind fields to the simulation tool demonstrated in Chapter 6. This technique shows how wind retrieval and ambiguity selection algorithms can be evaluated through simulation. The simulation allows for both visual inspection of the wind fields and calculation of various quantitative performance metrics. The technique is demonstrated for QuikSCAT UHR wind estimation with and without DIR. The demonstration validates QuikSCAT UHR wind estimation with DIR.

The three improved analysis techniques developed in this thesis provide a new way to study ocean winds through the 2D kinetic energy spectrum and provide methods to evaluate scatterometer processing algorithms through the WRF and the simulation tool. These techniques are sufficient to better scatterometer wind estimation by tuning existing algorithms and creating new algorithms.

Additionally, this thesis demonstrates the importance of simulation in scatterometer analysis. Though simulation, WRFs are calculated for three different QuikSCAT processing types. These WRFs demonstrate that the spatial resolution of UHR processing is finer than the spatial resolution of DIB egg and slice processing. This finer resolution is evident in the simulation of several wind retrieval and ambiguity selection algorithms. The simulation of the wind retrieval and ambiguity selection algorithms demonstrates that DIR improves wind estimation and that UHR processing is only slightly more affected by noise than DIB processing is.

7.2 Contributions

This thesis makes the following contributions to analyzing scatterometer-measured winds.

- An analysis of different window types for mitigating spectral leakage when measuring the 2D kinetic energy spectrum of ocean winds.

Spectral leakage occurs when a signal measured by a DFT is not periodic in the sample interval. Windowing is commonly used to mitigate the effects of spectral leakage when measuring the kinetic energy spectrum of ocean winds in 1D. For measuring the kinetic energy spectrum of ocean wind in 2D, windows that are circularly symmetric are shown to be better at mitigating the effects of spectral leakage than square-shaped windows.

- A method to measure the kinetic energy spectra of ocean winds in 2D.

For ease of interpretation and to follow historical precedents, the kinetic energy spectra of ocean winds is normally measured in 1D. This thesis expands a method to measure the kinetic energy spectra of ocean winds in 1D to 2D. This expansion mitigates spectral leakage effects in 2D and accounts for the increased high to low frequency bin ratio that exists in 2D spectral analysis. The method is validated using QuikSCAT wind estimates.

- Additional validation of the assumptions that atmospheric motions are 2D, nondivergent, and isotropic.

The assumptions that large-scale atmospheric motions are 2D, nondivergent, and isotropic are used in many studies and models of ocean winds and have been verified in previous studies [15], [16]. This thesis provides additional evidence of these assumptions by directly comparing the kinetic energy spectrum in 1D and 2D. When modeling the 1D kinetic energy spectrum as $E_{1\text{D}} = \beta_{1\text{D}} k^{\alpha_{1\text{D}}}$ and the 2D kinetic energy spectrum as $E_{2\text{D}} = \beta_{2\text{D}} k^{\alpha_{2\text{D}}}$, these assumptions imply that $\alpha_{2\text{D}} = \alpha_{1\text{D}} - 1$. This relationship is shown by the values calculated for $\alpha_{1\text{D}}$ and $\alpha_{2\text{D}}$ for QuikSCAT winds in a variety of different measurement locations.

- Additional validation that orthogonal components of ocean wind are uncorrelated.

This thesis adds to previous verification that the orthogonal components of ocean winds are uncorrelated (e.g., [15], [16]) in the following manner. First, the kinetic energy spectrum for the sum of two uncorrelated orthogonal wind components is mathematically shown to be equal to the sum of the kinetic energy spectra of the uncorrelated wind orthogonal components. Then, the numerical analysis of the aforementioned kinetic energy spectra in 1D and 2D indicates that these quantities are equal. These results suggest that the orthogonal wind components are uncorrelated.

- The definition and development of the WRF.

The SRF is a tool used in enhancing the spatial resolution for scatterometer estimation. The SRF helps to quantify the contribution each point on the surface makes to a σ^0 measurement. This thesis develops the WRF which can be thought of as a wind-specific SRF. The WRF helps to quantify the contribution that the wind at each point on the surface makes to a scatterometer-measured wind. The WRF helps understand the spatial resolution of different wind processing algorithms.

- A measure of the spatial resolution for three different types of QuikSCAT data processing is provided by the WRF.

QuikSCAT winds were initially produced using DIB egg processing and then DIB slice processing. Scatterometer winds are also produced through UHR processing techniques. Through their generalized WRFs, the spatial resolution of each of the processing techniques is estimated. The spatial resolution of UHR processing is found to be finer than the spatial resolution of DIB slice processing which is finer than the spatial resolution of DIB egg processing.

- A process to create synthetic wind fields sufficiently realistic to be used in evaluating ambiguity selection algorithms.

A variety of different methods to create synthetic wind fields have been developed. One of these methods creates synthetic wind fields whose kinetic energy spectra follow the expected power law for ocean winds. This thesis builds upon this method by extrapolating the kinetic energy spectrum of an existing low resolution wind field to help set the general characteristics of the wind field. The resulting wind fields appear more realistic in the spatial distribution of wind features and are sufficiently realistic to use in ambiguity selection algorithms.

- A simulation-based method to evaluate wind retrieval and ambiguity selection algorithms for UHR winds.

Simulation is an important part of scatterometer algorithm design as it allows the true wind field to be compared directly to the measured wind field. This thesis expands scatterometer simulation methods to QuikSCAT UHR wind fields and uses more spatially realistic wind fields than were previously developed. The simulation allows for both a visual comparison between the true and measured wind fields and an evaluation of quantitative performance metrics for the measured wind fields.

- Validation of QuikSCAT UHR winds with DIR through simulation.

DIR was originally developed for QuikSCAT DIB processing and was later applied to QuikSCAT UHR winds. Validation of UHR winds with DIR has been given for another pencil-beam scatterometer, RapidScat, by comparing buoy wind measurements with scatterometer measurements [7], but the measurement geometry of QuikSCAT is sufficiently different from RapidScat to warrant additional validation for QuikSCAT UHR winds with DIR. This thesis provides the first validation of QuikSCAT UHR winds with DIR through simulation. QuikSCAT UHR winds with DIR are shown to reduce the wind direction RMSE by $\sim 10^\circ$ compared to QuikSCAT UHR winds without DIR.

7.3 Future Work

The work provided in this thesis can be expanded and developed further. The following topics are suggestions for future work based upon the contributions made in this thesis.

- Develop improved ambiguity selection algorithms compatible with DIR winds.

The original mission requirements for QuikSCAT are defined for coarse resolution winds. QuikSCAT UHR winds with DIR allow fine resolution winds to nearly achieve the original mission requirements for QuikSCAT. DIR wind retrieval is capable of achieving the original mission requirements, but an improved ambiguity selection algorithm is required to reduce the wind direction RMSE of UHR winds below the mission requirement of 20°. Finding an improved algorithm would result in higher accuracy winds with fine resolution. The ambiguity selection algorithm could be applied to other pencil-beam scatterometers as well.

- Use the provided WRFs to improve land contamination ratios and near-coastal measurements in wind estimation.

As σ^0 is much higher for objects on land than for water, wind estimation can fail when WVCs are too close to land. Land contamination ratios help to quantify how much the land affects egg and slice measurements used in wind estimation and helps define how close to land the wind can be estimated accurately. Currently, land contamination ratios are based on the SRF of individual egg and slice measurements. By finding the WRF of WVCs close to land, a better idea of the effect that the land has on the wind measurements can be ascertained. The effect of dropping certain measurements that contain high levels of land contamination on the measured wind vectors could also be found. These studies could lead to new algorithms that improve the accuracy of near-coastal wind measurements.

- Calculate the WRF for more scatterometers and processing algorithms.

Scatterometers are either pencil-beam scatterometers, like QuikSCAT, or fan-beam scatterometers, like the Advanced Scatterometer (ASCAT). Fan-beam scatterometers have several long, fixed antennas with measurement swaths on the sides of the satellite track. While the WRFs found for QuikSCAT would translate well to characterizing other pencil-beam scatterometers, fan-beam scatterometers would have a much different WRF. The WRF would help characterize some of the differences between pencil-beam and fan-beam scatterometers and help to provide a comparison between the spatial resolutions of their wind estimates.

REFERENCES

[1] F. Ulaby and D. Long, *Microwave radar and radiometric remote sensing*. Artech House, 2015. 4, 5, 8, 9, 40, 41

[2] D. G. Long, "Polar applications of spaceborne scatterometers," *IEEE journal of selected topics in applied earth observations and remote sensing*, vol. 10, no. 5, pp. 2307–2320, 2016. 5

[3] J. Budge and D. Long, "Estimating sizes and rotation angles of antarctic icebergs utilizing scatterometer data," in *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. IEEE, 2017, pp. 3585–3588. 5

[4] S. Frolking, T. Milliman, K. McDonald, J. Kimball, M. Zhao, and M. Fahnestock, "Evaluation of the SeaWinds scatterometer for regional monitoring of vegetation phenology," *Journal of Geophysical Research: Atmospheres*, vol. 111, no. D17, 2006. 5

[5] F. M. Naderi, M. H. Freilich, and D. 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, 1991. 5

[6] M. W. Spencer, C. Wu, and D. G. Long, "Tradeoffs in the design of a spaceborne scanning pencil beam scatterometer: Application to SeaWinds," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 35, no. 1, pp. 115–126, 1997. 5

[7] N. Hutchings and D. G. Long, "Improved ultrahigh-resolution wind retrieval for RapidScat," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 6, pp. 3370–3379, 2018. 7, 12, 61, 96

[8] I. S. Ashcraft and D. G. Long, "The spatial response function of SeaWinds backscatter measurements," in *Earth Observing Systems VIII*, vol. 5151. International Society for Optics and Photonics, 2003, pp. 609–618. 6, 40

[9] SeaPAC, "QuikSCAT level 2b version 4.0 guide document," Physical Oceanography Distributed Active Archive Center, August 2018. 10, 84, 89

[10] D. G. Long, P. J. Hardin, and P. T. Whiting, "Resolution enhancement of spaceborne scatterometer data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 31, no. 3, pp. 700–715, 1993. 10, 36, 37, 83

[11] D. G. Long, "High resolution wind retrieval from SeaWinds," in *IEEE International Geoscience and Remote Sensing Symposium*, vol. 2. IEEE, 2002, pp. 751–753. 10

[12] 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, 1991. 11

[13] B. W. Stiles, B. D. Pollard, and R. S. Dunbar, "Direction interval retrieval with thresholded nudging: A method for improving the accuracy of QuikSCAT winds," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 40, no. 1, pp. 79–89, 2002. 11, 13, 20

[14] R. Dunbar, T. Lungu, B. Weiss, B. Stiles, J. Huddleston, P. Callahan, G. Shirtliffe, K. Perry, C. Hsu, C. Mears *et al.*, "QuikSCAT science data product user manual, version 3.0, JPL Document D-18053—Rev A," *Jet Propulsion Laboratory, Pasadena, CA*, 2006. 13, 40, 78, 80, 81

[15] M. Freilich and D. Chelton, "Wavenumber spectra of Pacific winds measured by the Seasat scatterometer," *Journal of Physical Oceanography*, vol. 16, no. 4, pp. 741–757, 1986. 13, 14, 18, 19, 20, 21, 22, 23, 29, 95

[16] J. Patoux and R. A. Brown, "Spectral analysis of QuikSCAT surface winds and two-dimensional turbulence," *Journal of Geophysical Research: Atmospheres*, vol. 106, no. D20, pp. 23 995–24 005, 2001. 13, 15, 18, 19, 20, 22, 23, 82, 95

[17] J. R. Blodgett, "Analysis, validation, and improvement of high-resolution wind estimates from the Advanced Scatterometer (ASCAT)," Master's thesis, Brigham Young University, December 2014. 13, 15, 18, 20, 28, 29

[18] J. Vogelzang, A. Stoffelen, A. Verhoef, and J. Figa-Saldaña, "On the quality of high-resolution scatterometer winds," *Journal of Geophysical Research: Oceans*, vol. 116, no. C10, 2011. 14, 20, 23

[19] G. K. Batchelor, "Computation of the energy spectrum in homogeneous two-dimensional turbulence," *The Physics of Fluids*, vol. 12, no. 12, pp. II–233, 1969. 15

[20] D. G. Long, "Model-based estimation of wind fields over the ocean from wind scatterments," Ph.D. dissertation, University of Southern California, 1989. 15

[21] K. E. Trenberth and A. Solomon, "Implications of global atmospheric spatial spectra for processing and displaying data," *Journal of climate*, vol. 6, no. 3, pp. 531–545, 1993. 15

[22] B. A. Williams and D. G. Long, "A reconstruction approach to scatterometer wind vector field retrieval," *IEEE transactions on geoscience and remote sensing*, vol. 49, no. 6, pp. 1850–1864, 2011. 15, 17

[23] M. P. Owen and D. G. Long, "Land-contamination compensation for QuikSCAT near-coastal wind retrieval," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 3, pp. 839–850, 2009. 15, 16

[24] R. D. Lindsley, J. R. Blodgett, and D. G. Long, "Analysis and validation of high-resolution wind from ASCAT," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 10, pp. 5699–5711, 2016. 15, 16, 61

[25] C.-Y. Chi and F. K. Li, “A comparative study of several wind estimation algorithms for spaceborne scatterometers,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 26, no. 2, pp. 115–121, 1988. 15, 16

[26] T. E. Oliphant and D. G. Long, “Accuracy of scatterometer-derived winds using the Cramer-Rao bound,” *IEEE transactions on geoscience and remote sensing*, vol. 37, no. 6, pp. 2642–2652, 1999. 15, 16, 43

[27] C. Zhang, Y. Wang, A. Lauer, and K. Hamilton, “Configuration and evaluation of the WRF model for the study of Hawaiian regional climate,” *Monthly Weather Review*, vol. 140, no. 10, pp. 3259–3277, 2012. 17

[28] P. J. Vickery, P. Skerlj, A. Steckley, and L. Twisdale, “Hurricane wind field model for use in hurricane simulations,” *Journal of Structural Engineering*, vol. 126, no. 10, pp. 1203–1221, 2000. 17

[29] C. K. Wikle, R. F. Milliff, and W. G. Large, “Surface wind variability on spatial scales from 1 to 1000 km observed during TOGA COARE,” *Journal of the atmospheric sciences*, vol. 56, no. 13, pp. 2222–2231, 1999. 18

[30] G. Nastrom and K. S. Gage, “A climatology of atmospheric wavenumber spectra of wind and temperature observed by commercial aircraft,” *Journal of the atmospheric sciences*, vol. 42, no. 9, pp. 950–960, 1985. 18

[31] D. B. Chelton, M. H. Freilich, J. M. Sienkiewicz, and J. M. Von Ahn, “On the use of QuikSCAT scatterometer measurements of surface winds for marine weather prediction,” *Monthly Weather Review*, vol. 134, no. 8, pp. 2055–2071, 2006. 20

[32] D. Long and D. Luke, “The wavenumber spectra of scatterometer-derived winds,” in *[Proceedings] IGARSS’92 International Geoscience and Remote Sensing Symposium*, vol. 2. IEEE, 1992, pp. 1005–1007. 20

[33] D. G. Long, “Scatterometer backscatter imaging using Backus–Gilbert inversion,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 6, pp. 3179–3190, 2018. 36, 49

[34] D. S. Early and D. G. Long, “Image reconstruction and enhanced resolution imaging from irregular samples,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 2, pp. 291–302, 2001. 36

[35] R. D. Lindsley and D. G. Long, “Enhanced-resolution reconstruction of ASCAT backscatter measurements,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 5, pp. 2589–2601, 2015. 36

[36] N. Hutchings, “Near-Coastal Utrahigh Resolution Scatterometer Winds,” Master’s thesis, Brigham Young University, December 2019. 81