A Quality Assurance Algorithm for SeaWinds

David W. Draper and David G. Long
Brigham Young University, MERS Laboratory
459 CB, Provo, UT 84602
801-378-4884, FAX: 801-378-6586
draperd@et.byu.edu
5-15-2001

Abstract

The scatterometer wind retrieval process produces several possible wind vector choices or ambiguities at each resolution cell. Ambiguity selection routines are generally *ad hoc* and often result in ambiguity selection errors. It is important to locate areas of ambiguity selection error to assess the quality of scatterometer wind data. A simple consistency check is presented based on comparing selected winds from SeaWinds on QuikSCAT to a low order wind field model fit. Regions exceeding error thresholds are flagged as possible ambiguity selection errors. Appropriate error thresholds and additional flagging criteria are set through an analysis of false alarms versus missed detections on a training set of 15 revolutions of data. The algorithm correctly identifies 97% of the regions manually flagged as ambiguity selection errors in the training set with a false alarm rate of less than 2%. Applying the algorithm to the entire QuikSCAT data set, we conclude that the ambiguity selection is over 95% effective on regions of rms wind speeds greater than 3.5 m/s. The algorithm validates that higher noise occurs in low wind speed regions and at nadir. Additionally, fewer estimated ambiguity selection errors occur at nadir and on the swath edges due to a larger ambiguity set in those regions. The percentage of ambiguity selection errors are found to be highly correlated with the number of cyclonic storms passed by SeaWinds.

1 INTRODUCTION

Scatterometers are spaceborne radars that infer near ocean winds. Scatterometers yield broader and more frequent coverage of ocean winds than traditional *in situ* observations. SeaWinds on QuikSCAT is the latest scatterometer launched by NASA. SeaWinds offers many advances in ocean wind estimation over previous scatterometer designs. Having a larger swath than its predecessors, SeaWinds affords near global coverage of surface winds on a daily basis. Because accuracy is very important, an assessment of the quality of SeaWinds data is essential.

Scatterometers measure radar pulse returns from the ocean's surface. The power in the return is used to calculate the normalized radar backscatter cross-section (σ^o). Microwaves are sensitive to wind-induced capillary waves making σ^o a function of the near-surface wind velocity. An observed σ^o measurement over a region of ocean maps to a range of possible wind vectors through the empirically determined geophysical model function (GMF). Because the relationship between σ^o and the wind is not unique, several measurements from different azimuth angles are required to infer the wind velocity. Instrument noise, GMF inaccuracies, and other factors introduce uncertainty into the measurements resulting in multiple wind vector estimates at each wind vector cell (wvc). These possible solutions are known as *point-wise ambiguities* [3]. The point-wise ambiguities at a wvc generally have similar wind speeds but different directions. Each ambiguity has an associated likelihood and there is no assurance that the first ambiguity is the correct choice. A separate step known as *point-wise ambiguity selection* is used to select a unique wind vector at each wvc. One of the difficulties in wind scatterometry is selecting the ambiguity closest to the true wind. When an ambiguity not the closest to the true wind is selected, an *ambiguity selection error* occurs.

The current point-wise ambiguity selection process used by NASA's Jet Propulsion Laboratory (JPL) for SeaWinds is a median filter approach. Numeric weather prediction models are used to nudge the initial estimate of the wind field. Then, a modified *point-wise median filter* iteratively selects the ambiguity at each wvc best matching the flow of the surrounding wvcs. This method is *ad hoc*, but performs well under most circumstances.

Limitations in both the nudging process and the point-wise median filter introduce ambiguity selection errors. Unfortunately, it is not possible to verify correct ambiguity selection without knowledge of the true wind. Nevertheless, due to the (generally) red spectrum of the wind it is possible to evaluate the self-consistency of the wind field resulting from the ambiguity selection process. As [1] demonstrated, the self-consistency can be evaluated by comparing the selected wind to a least-squares fit using a low-order Karhunen Loeve (KL) model. This consistency check, though not without limitations, can be a useful tool for a quality assurance (QA) check of the data.

This paper develops an improved SeaWinds-data only method for detecting ambiguity selection errors and evaluating the general quality of the data. The QA algorithm and the method for setting appropriate thresholds are discussed in detail. In addition, a level bit flag is introduced which conveys the overall quality of each wvc. The algorithm is applied to a sample set of QuikSCAT winds. Due to a larger ambiguity set at nadir and along swath edges, less ambiguity selection errors are inferred at nadir and along swath edges than in the off-nadir re-

gion. However, nadir and low wind speed regions are shown to be noisier. The position of ambiguity selection errors tends to be highly correlated with the position of cyclonic storms. The baseline ambiguity selection is estimated to be at least 95% effective for overlapping regions with rms wind speeds above 3.5 m/s.

2 NOISE AND POINT-WISE AMBIGUITY SELECTION

The point-wise wind estimation process infers wind speed and direction from noisy backscatter measurements. At each wvc, a set of point-wise ambiguities is generated by finding the local minima of the point-wise objective function. The point-wise weighted least-squares objective function is a measure of the error between the observed σ^o measurements and a wind vector's corresponding σ^o values projected through the GMF, i.e.

$$\sum_{k=1}^{K} \frac{[\sigma_k^o - \mathcal{M}(\mathrm{u}, \phi - \psi_k)]^2}{2\varsigma_k^2} \tag{1}$$

where ς_k^2 is the noise variance of the k^{th} measurement and $\mathcal{M}(\mathrm{u}, \phi - \psi_k)$ is the forward projection of a wind vector estimate (u, ϕ) through the GMF given ψ_k, the azimuth angle of the k^{th} measurement. The objective function generally has multiple local minima representing the point-wise wind ambiguities. With SeaWinds, only the first one to four ambiguities are used in the estimation process. The ambiguities depend on the σ^o measurements as well as the instrument geometry. Where multiple measurements lack sufficient azimuthal variation (i.e. at nadir), wind estimates are noisier. Also, a low signal to noise ratio in low wind speed regions can result in an inaccurate estimate of the wind's direction. Thus, point-wise retrieved winds both at nadir and in low wind speed regions tend to be noisy.

After generating the set of point-wise ambiguities, an ambiguity selection routine is employed. The ambiguity selection algorithm currently used by JPL in processing SeaWinds data has two parts: nudging and median filtering. In traditional nudging, each wvc is set to the ambiguity that most closely matches an outside estimate of the wind field, i.e. numerical weather prediction models. JPL implements a variant of traditional nudging known as *thresholded nudging* with SeaWinds. In thresholded nudging, the estimated likelihood of the correctness of the first ambiguity or instrument skill at a given wvc dictates the set of ambiguities used to initialize the process. Where the estimated instrument skill is high, only ambiguities with high likelihood values are chosen in nudging. Where the estimated instrument skill is poor, all ambiguities may be used.

After selecting an initial field by nudging, the point-wise median filter is applied [4]. For each wvc, the point-wise median filter selects the ambiguity that minimizes the directional error between the ambiguity and the surrounding cells, i.e.

$$\hat{n} = \arg\min_n \sum_{k=i-3}^{i+3} \sum_{l=j-3}^{j+3} |\phi_{ij}^n - \Phi_{kl}|_{\{[0^\circ, 180^\circ]\}} \tag{2}$$

where Φ_{kl} are directions of the surrounding selected wind vectors and ϕ_{ij}^n is the direction of the nth alias at wvc i, j. The new chosen ambiguity, $\hat{n}$ replaces the previous ambiguity on the next iteration. The process is iterated until convergence criteria are met.

Although the JPL point-wise ambiguity selection routine generally produces self-consistent results, it is *ad hoc* and does not guarantee the correct solution. In addition, the process is sensitive to the nudging field. Thresholded nudging also is problematic in areas of estimated high instrument skill where only the first ambiguity is used in the nudging process. When the first ambiguity is incorrect (due, for example, to rain contamination), ambiguity selection errors often occur. In addition to ambiguity selection errors, noise may cause a selected ambiguity to seriously deviate from the true wind albeit there is no better choice.

3 THE KL WIND MODEL

Our QA algorithm is based on comparing the selected wind field to a low order model fit. The wind field model is data-derived using the KL-based technique described by [1] for use with NASA scatterometer (NSCAT) data. The KL approach presumes ambiguity selection is generally good. This was shown to be true for NSCAT and similar results have been obtained for SeaWinds data.

3.1 Formation of the KL Wind Model

The KL model is formed from the eigenvalues of an autocorrelation matrix. The following steps outline the how the KL wind model is created from a training set of QuikSCAT winds:

1. The training set is subdivided into $N \times N$ wvc regions (typically $N \sim 8$ to 24). Each region contains no missing data points.

2. The *standard vector form* ($\mathbf{w}_n$) of the wind field is formed by column scanning the rectangular (U and V) components for each $N \times N$ region, i.e.

$$\mathbf{w}_n = \begin{bmatrix} \mathcal{O}_c\{U_n\} \\ \mathcal{O}_c\{V_n\} \end{bmatrix}. \tag{3}$$

   where $\mathcal{O}_c\{\cdot\}$ represents column ordering of the matrix.

3. The autocorrelation matrix R is estimated as

$$\hat{R} = \frac{1}{M} \sum_{n=1}^{M} \mathbf{w}_n {\mathbf{w}_n}^T \tag{4}$$

   where M is the number of $N \times N$ wind fields examined.

4. The orthonormal matrix S is extracted by taking the eigenvalue (Λ) decomposition of $\hat{R}$ where

$$\hat{R} = S \Lambda S^T. \tag{5}$$

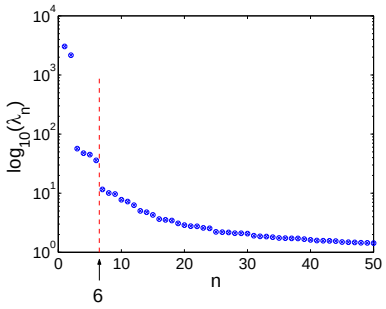

Figure 1: *The first 50 eigenvalues of the 8 × 8 KL model.*

The diagonal elements of Λ are the eigenvalues while the columns of S are the the eigenvectors which become the basis fields or model parameters of the KL model.

5. The orthonormal matrix S is truncated to an suitable number of parameters, forming the restricted basis F.

The eigenvectors corresponding to the largest eigenvalues model the general wind flow. For an 8 × 8 wvc model, the 6 highest eigenvalue wind fields do an adequate job of describing over 95% of the wind variability (see Figure 1). These wind fields form a basis set for the wind field model as shown in Figure 2. Because the 8 × 8 model requires so few parameters to characterize the major wind flow, it is chosen as the model size for our QA algorithm.

3.2 Weighted Least-Squares Fit Using the KL Model

To evaluate the consistency of the wind, a wind field is written as a linear combination (using a least-squares fit) of the truncated basis set, F. The coefficients ($\hat{\mathbf{x}}$) of the model parameters are obtained using the weighted least squares estimate of F^{-1}, $F^{\dagger}$, i.e.,

$$\hat{\mathbf{x}} = F^{\dagger}\mathbf{w}_0 \tag{6}$$

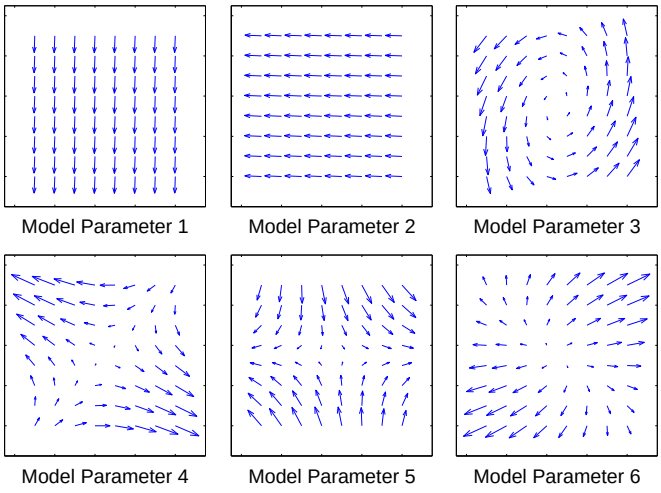

Figure 2: *The truncated KL model used in the QA algorithm.*

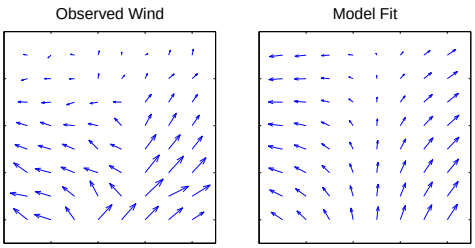

Figure 3: *8 × 8 Region of QuikSCAT point-wise selected wind and the KL model fit.*

where $\mathbf{w}_0$ is the standard vector form of the 8 × 8 wind field. The weighted least squares estimate,

$$F^{\dagger} = (F^T W F)^{-1} F^T W \tag{7}$$

ignores non-data points (i.e. land or missing data) via the weighting matrix W, a diagonal matrix with ones corresponding to places with real data points and zeros corresponding to places with non-data points. The *model-fit* wind field ($\mathbf{w}_m$) is then

$$\mathbf{w}_m = F\hat{\mathbf{x}}. \tag{8}$$

The difference between the model-fit and observed wind field is computed. Where the difference is significant, the wind field basis set is not adequate for describing the selected ambiguity field.

Figure 3 shows an 8 × 8 wvc region of point-wise selected wind and its least squares approximation. This example contains inconsistent flow evidenced by the divergent nature of the central section. The restricted basis cannot accurately represent this feature. This region is identified as a possible ambiguity selection error.

In general, differences between the selected ambiguity field and the model-fit field may result from several causes:

1. Ambiguity selection errors, especially when the difference is large. Ambiguity selection errors result in large changes in the wind direction which are inconsistent with realistic wind flow.

2. Limitations of the basis wind field. Because a truncated KL basis is used, the basis set does not include fronts and other fine-scale features.

3. Noise in the scatterometer wind measurements.

It is also possible to have an area of ambiguity selection error with spatially consistent wind field. Our technique cannot identify such regions.

Due to these considerations, the algorithm is considered only useful in flagging regions containing *possible* ambiguity selection errors which are evidenced by inconsistent wind flow. Without knowing the true wind it is impossible to identify actual ambiguity selection errors with 100% accuracy. A set of empirically determined thresholds is used to determine if the errors are significant enough to give a QA flag.

4 AN OVERVIEW OF THE QUALITY ASSURANCE ALGORITHM

Formally, the QA algorithm works in the following way:

1. *Segment swath:* The swath is segmented into 8×8 wvc wind fields overlapping by half in the cross-track and half in the along-track directions. All regions containing more than 25% non-data points are excluded.

2. *Model fit:* A weighted least-squares model fit is made of the each region.

3. *Compare observed wind to model fit:* The direction error and vector error for each wvc in the region is calculated. The *direction error* is defined as the difference in direction between the model-fit cell and the selected ambiguity, i.e.
$$\phi_{\mathbf{e}} = |\phi_m - \phi_o|_{0 \le \phi_{\mathbf{e}} \le 180} \tag{9}$$
where ϕ_m and ϕ_o are the directions of the modeled and observed cells respectively. The direction error is always between 0° and 180°. The *vector error* is defined as the magnitude of the vector difference between the model fit and the selected ambiguity, i.e.
$$\mathbf{k_e} = \left((u_m - u_o)^2 + (v_m - v_o)^2\right)^{\frac{1}{2}} \tag{10}$$
where (u_m, v_m) and (v_o, v_m) are the u and v components of the model-fit cell and observed wvc respectively.

4. *Flag individual cells:* Individual cells are flagged which exceed thresholds in direction or vector error. Two sets of wvc thresholds are used: constant and variable. Vectors exceeding the constant thresholds are flagged as *noisy vectors*, and those exceeding variable thresholds are flagged as *possible ambiguity selection error vectors.*

5. *General classification using constant thresholds:* Regions are classified as "good," "fair," or "poor" according to the percent of valid cells flagged per region by constant thresholds. These classification delimiters are called *region thresholds*. This general classification rates the overall consistency of the wind in each region. Noise, ambiguity selection errors, and limitations of the KL basis may contribute to a "fair" or "poor" classification.

6. *Ambiguity selection error detection:* A more sophisticated analysis additionally flags the region as a possible ambiguity selection error. This is known as the ambiguity selection error (ASE) region flag. This binary flag is optimized to locate ambiguity selection errors by suppressing known swath and wind speed dependent noise using a set of variable wvc thresholds.

The QA algorithm is summarized in Figure 4. The following sections describe both the general classification and the ASE flag in detail.

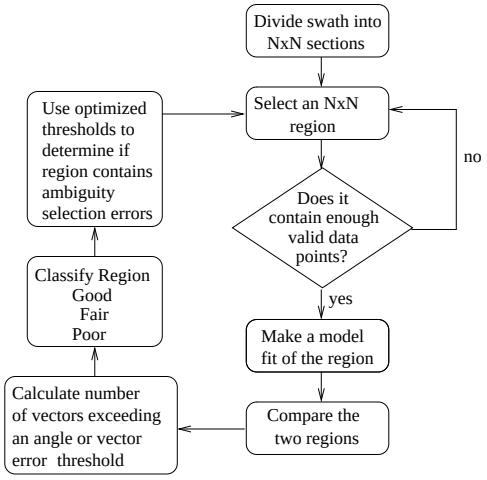

Figure 4: *Flow chart describing the quality assurance algorithm.*

5 GENERAL CLASSIFICATION USING CONSTANT THRESHOLDS

The QA algorithm's general region classification indicates the level to which the observed wind deviates from the model in a given region. Noise, ambiguity selection errors and limitations of the truncated basis contribute to the difference between the model fit and the observed wind. This section discusses the scheme for classifying regions and the meaning of each classification.

A region is classified according to the number of wvcs flagged per region by constant thresholds. This algorithm uses some of the same wvc thresholds found in [1]. Table 1 shows the wvc thresholds used in the general region classification of the SeaWinds algorithm.

In Table 1, u_{rms} is the region root mean squared (rms) wind speed defined by
$$u_{rms} = \left(\frac{\mathbf{w}_o^T \mathbf{w}_o}{N}\right)^{\frac{1}{2}} \tag{11}$$
where $\mathbf{w_o}$ is the standard vector form of the observed wind field and N is the number of valid data cells in the region. The value of 2.7 m/s used for the vector error was originally used as a maximum component error threshold in [1]. However, a maximum component error metric preferentially flags error vectors

Table 1: *Constant thresholds determining the flagging of a vector. († Denotes values used by [1]).*

WVC Threshold	Value
direction error	$23°$ †
vector error	$\max \begin{cases} 2.7 \text{ †} \\ 0.5u_{rms} \end{cases}$ m/s

Table 2: *Thresholds determining the classification of a region. († Denotes Values used by [1]).*

Classification	Percentage of cells flagged per region
"Good"	$< 5\%$
"Fair"	5 - 20%
"Poor"	$> 20\%$ †

lying along a component axis. Thus, a vector error metric is used because it is independent of the error vector's orientation.

The vector error threshold is also dependent on region rms wind speed. The six parameter KL model has difficulty modeling high cell-to-cell wind speed variation. For high speeds, this variation may be greater than 2.7 m/s and thus cells may be flagged due to poor modeling. To account for this, when a region of rms wind speed greater than 5.4 m/s is encountered, the threshold is raised to 1/2 the rms wind speed of the region.

Regions are classified by the number of cells flagged as shown in Table 2. The classifications of the region flagging scheme have the following general meanings:

"Good" - (less than 5% of wvcs flagged) The wind flow fits the model estimate very well and is spatially consistent with a low noise level.

"Fair" - (between 5 and 20% of wvcs flagged) The wind flow is consistent, but some vectors may contain moderate amounts of noise and/or possible ambiguity selection errors. Wind fields with fine scale spatial variations (e.g. fronts) may also be flagged.

"Poor" - (more than 20% of wvcs flagged) The wind flow is not consistent due to ambiguity selection errors or high levels of noise. Nadir and low wind speed regions (both of which are noisier) are more likely to classified "poor". Features such as fronts or cyclones also are also likely to be flagged due to the model's restricted basis set.

Although these classifications are not optimized to only identify regions of ambiguity selection error, they are helpful in determining the integrity of the data itself. Data cells corrupted by noise or rain are more likely to be flagged "poor" using these classifications. The ASE flag described in the next section reduces the effects of noise in locating true ambiguity selection errors.

6 IDENTIFYING AMBIGUITY SELECTION ERRORS (ASE FLAG)

Because SeaWinds has variable performance with wind speed and swath position, it is difficult to locate ambiguity selection errors for all regions with constant performance. When the constant threshold scheme (as in the general classification) is used, the algorithm preferentially identifies regions of high noise level. Using SeaWinds data, the noise variance per cross track position and wind speed is estimated in this section. By adjusting wvc thresholds to perform optimally for each wind speed and cross track position we desensitize the algorithm to known high noise regions. The following sections develop the method of determining optimal thresholds for variable cross track positions and wind speeds.

6.1 Estimating Noise Variance in Nadir and Low Wind Speed Regions

Near the center of the swath, the difference in azimuth angle between the two SeaWinds radar beams decreases to a minimum. In addition, the fore and aft observations become 180° out of phase. This limited azimuth variation among the multiple measurements creates a larger variance in the likelihood estimates for each ambiguity [5]. The larger variance results in a noisier estimated wind field. Figure 5 demonstrates a typical nadir-region wind field. The lack of azimuth variation in the σ^o measurements has introduced noise. Figure 5 also shows a wind field in the "sweet spot" of SeaWinds' swath. The "sweet spot" refers to a region of high azimuth variation outside of nadir. Notice the wind field is much smoother than the nadir wind field. The smoother nature of the wind field is due to a larger azimuth variation among multiple σ^o measurements.

In addition to the relationship between swath position and noise, estimates are also affected by wind speed. A calm ocean, corresponding to a low wind speed, is not a good scatterer. Thus, the signal to noise ratio (SNR) for a low wind speed estimate may be much lower than for a moderate wind speed estimate. A low SNR can result in unreliable measurements. Thus, data processed for vectors with low wind speeds (corresponding to small σ^o values) tends to be noisier than data from higher wind speeds. Noise in low wind speed regions causes a large error between the point-wise selected wind field and the model-fit field.

Figure 6 shows a region with a low rms wind speed (2.9 m/s) and the possible point-wise ambiguities. The point-wise algorithm selected a large percentage of first ambiguities. Although this region is noisy, there may be no ambiguity selection errors. The point-wise algorithm simply had to select from the noisy ambiguities given.

In order to estimate the noise level in nadir and low wind speed regions, we present an average standard deviation of wvc directions for each cross track/rms wind speed bin. The stan-

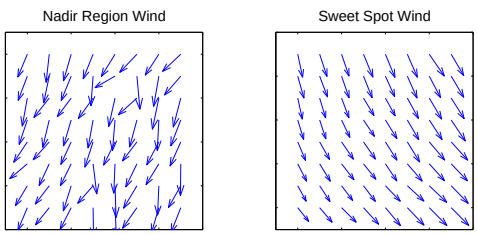

Figure 5: *8×8 Ambiguity selected wind field in the nadir region and "sweet spot".*

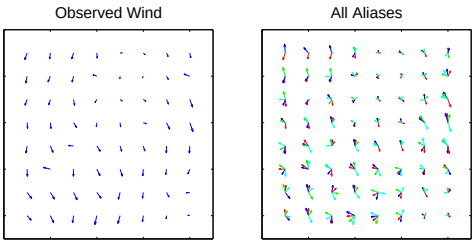

Figure 6: *A sample point-wise selected region of low wind speed and all the ambiguities. Red are the first ambiguities, blue are second, green are third, and cyan are fourth.*

dard deviation for each region is estimated as the following:

$$\sigma_{est} = \min\{\sigma_{-180^\circ<\phi<180^\circ}, \sigma_{0^\circ<\phi<360^\circ}\} \qquad (12)$$

where σ is the standard deviation for all vector directions (ϕ) in a given 8×8 region. Here, we select the minimum value for the direction standard deviation calculated when the directions range from -180° to 180° and from 0° to 360°. This definition suppresses errors in the calculation of the standard deviation when vector directions in a region straddle either 0° and 360° or -180° and 180°.

The standard deviation for each cross track position/rms wind speed bin averaged over several hundred revs is plotted on a three dimensional grid in Figure 7. A higher noise variance at nadir and in low wind speed regions make it difficult to set appropriate constant wvc thresholds for locating ambiguity selection problems. Because the quality assurance algorithm is based on flagging individual vectors exceeding a threshold in direction or vector error, using constant thresholds promotes more frequent flagging of high noise regions. This warrants using variable thresholds for different cross track positions and rms wind speeds.

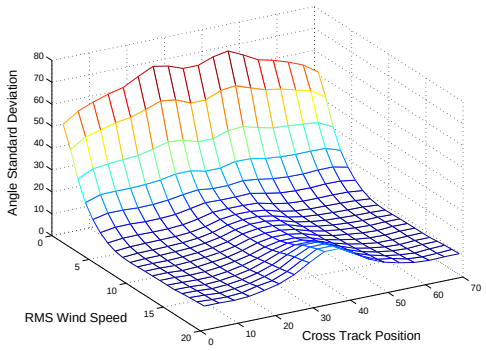

Figure 7: *The average directional standard deviation for both rms wind speed and cross track position.*

6.2 Selecting Variable wvc Thresholds

This section develops the method of determining variable wvc thresholds to detect ambiguity selection errors in the presence of cross track and wind speed dependent noise. In addition to the wvc and region thresholds, we describe several other criteria which must be met to flag a region as a possible ambiguity error. These results are tuned to SeaWinds data, but similar methods may be adapted for use with future instruments.

A selection of 15 revs was manually inspected for ambiguity selection errors in every 8×8 wvc processable region. All regions that exhibited clear ambiguity selection errors were identified. This serves as a *training set* to tune the algorithm.

The following discussion uses the terms *missed detections* and *false alarms*. A missed detection is defined as a region that is manually flagged as an ambiguity selection error, but is not flagged by the algorithm. A false alarm is defined as a region that is flagged by algorithm as an ambiguity selection error, but not manually flagged. The *false alarm rate* or *probability of false alarm* is defined as the total number of false alarms divided by the total number of regions not manually flagged as ambiguity selection errors. The *missed detection rate* or *probability of missed detection* is defined as the total number of missed detections divided by the total number of regions manually flagged as ambiguity selection errors. Because the number of regions not manually flagged as ambiguity selection errors is about 95% of the data, it is valuable to insure a low final false alarm rate. Both direction and vector error thresholds are separately optimized to give a constant false alarm rate of 2.5% for each type of threshold. At this false alarm rate, a region threshold of 14% is found to be optimal.

6.2.1 Optimal Direction Error Thresholds

Because it is important to minimize the probability of false alarm, the noise level at each cross track position and for each rms wind speed must be taken into account. The variability due to noise can be suppressed in the flagging process by setting higher thresholds for statistically noisier regions thereby equalizing the false alarm rate for all wind speeds and cross track positions.

Assuming that the higher variability in Figure 7 is caused mostly from noise, setting thresholds higher for regions of a higher estimated directional standard deviation should suppress false alarms due to noise. Thus, an optimally adjusted version of the shape in Figure 7 should yield a constant false alarm rate for all wind speeds and cross track positions.

To find a set of direction thresholds that gives a constant false alarm rate for all rms wind speeds and cross track positions, the following approach is implemented:

1. Using the training data set in which ambiguity selection errors have been manually identified, all 8×8 wvc regions are binned according to rms wind speed and cross track position.

2. The initial wvc threshold is assigned to be the lowest expected threshold value.

3. The observed wind is compared to the model fit and regions are flagged as possible ambiguity selection errors according to the number of poor cells per region (in this case, 14%).

4. The number of false alarms and missed detections is calculated for each bin.

5. If the false alarm rate for a bin is above a certain limit, the wvc threshold for that bin is raised.

6. If the missed detection rate for a bin is greater than zero, and the false alarm rate is significantly smaller than the desired false alarm rate, the wvc threshold is lowered.

7. The number of false alarms and missed detections is recomputed for each cross track and rms wind speed bin.

8. This process is iterated until either the false alarm rate for all the bins or the average false alarm rate falls beneath a desired threshold.

This algorithm was applied to a set of 15 test revs. Figure 8 shows the resulting set of direction error thresholds. The general shape mirrors the directional standard deviations shown in Figure 7, verifying our assumption.

The thresholds in Figure 8 are very noisy because of the limited data set used to compute them. However, a good approximation to "smoothed" thresholds is formed by fitting the curve from Figure 7 to match the values of Figure 8. The following algorithm accomplishes this in a least-squares sense:

1. For each row of constant rms wind speed, the values from Figure 7 and a uniform vector the same length are selected to be "model parameters" for a cross track row from Figure 8. These two vectors form the columns of matrix T.

2. A least squares fit using these two model parameters is made to the corresponding row of numerically determined threshold values, $\mathbf{r}_0$ by taking the pseudo-inverse of the matrix T, i.e. $\mathbf{x} = T^{\dagger}\mathbf{r_0}$.

3. The modeled row is then $\mathbf{r} = T\mathbf{x}$.

4. Each of the columns corresponding to constant cross track position is low passed filtered to smooth out any obvious anomalies.

5. Extra adjustments are made manually to improve performance. These include adjusting the thresholds for low (< 4 m/s) and high (> 15 m/s) rms wind speeds to subjectively give better performance. Insufficient data for these regions warrant manual adjustment.

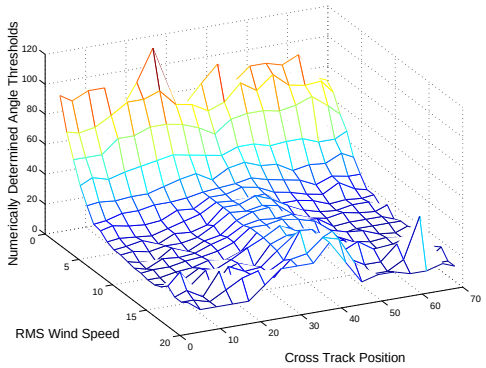

Figure 8: *The direction thresholds per cross track and rms wind speed that minimizes the probability of false alarm beneath a threshold of 2.5%.*

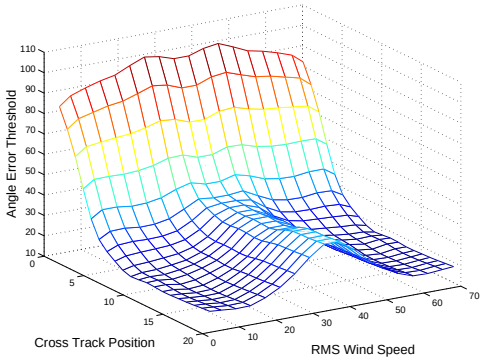

Figure 9: *The direction error thresholds per cross track and rms wind speed that minimize the probability of false alarm beneath a threshold of 2.5%.*

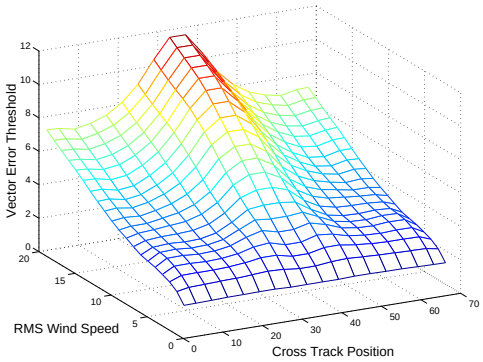

Figure 10: *The vector error thresholds per cross track and rms wind speed that minimize the probability of false alarm beneath a threshold of 2.5%.*

This method creates "smoothed" wvc thresholds which approximate the wvc thresholds that give a constant false alarm rate for all cross track position and rms wind speeds on the training data set. The final direction error thresholds are shown in Figure 9.

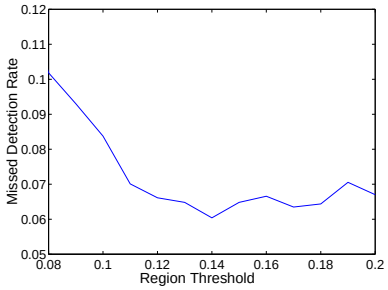

Figure 11: *Probability of missed detection with optimized direction and vector thresholds for different region thresholds.*

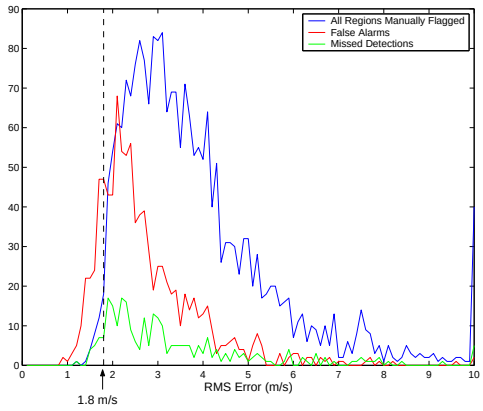

Figure 12: *Histogram of the rms errors of all the regions manually flagged as "poor", the false alarms generated by the algorithm, and the missed detections generated.*

6.2.2 *Optimal Vector Error Thresholds*

A set of vector error thresholds is also determined which equalizes the false alarm rate for both rms wind speed and cross track position. These thresholds are determined in the same way as the direction error thresholds. Figure 10 shows the imperially determined and "smoothed" vector error thresholds.

In order to classify a region as an ambiguity selection error, the total number of wvcs flagged must exceed a threshold of 14%. To select this threshold, we analyzed the the missed detection rate generated by performing the wvc-threshold optimization algorithm with different region thresholds (Figure 11). As shown, the threshold that results in the minimum missed detections is 14%. Thus, this threshold is used in the algorithm.

6.3 Additional Flagging Criteria

We have discussed an ambiguity selection error identification scheme based on the number of individual vectors flagged per region. The false alarm and missed detection rates have been manually set by properly adjusting the wvc and region thresholds. In order to further reduce the false alarms without significantly increasing the missed detection rate, we determine additional criteria that must be met in order to flag a region as an ambiguity selection error. This section describes each of the additional criteria including a region RMS error threshold, a supplementary consistency check, and a region RMS wind speed threshold.

6.3.1 *Region RMS Error Threshold*

In addition to the region threshold for the number of wvcs flagged, an rms error threshold is applied to each region inspected. The rms error is defined as the following:

$$\left(\frac{(\mathbf{w}_o - \mathbf{w}_m)^T(\mathbf{w}_o - \mathbf{w}_m)}{N}\right)^{\frac{1}{2}} \tag{13}$$

where $\mathbf{w}_o$ and $\mathbf{w}_m$ are the standard vector forms of the observed and modeled winds respectively and N is the number of valid data cells in the region. The number of false alarms and missed detections per rms error bin from applying the algorithm to the training data set is computed in Figure 12. The false alarms for regions whose rms error is less than 1.8 m/s account for nearly 20% of the total number of false alarms. The missed detection rate is also very high for those regions. This suggests that the algorithm does not perform well for regions where the rms error is low ($<$ 1.8 m/s). Thus, an *rms error threshold* is applied which excludes all regions with rms error less than 1.8 m/s from being flagged as possible ambiguity selection errors. The use of this threshold decreases the false alarm rate by nearly 20% while it increases the missed detection rate by less than 1%. For a region to be flagged as a possible ambiguity selection error, the rms region threshold must be met.

6.3.2 *Supplementary Consistency Check*

In order to further reduce the false alarm rate of the QA algorithm, a supplementary consistency check is performed. The purpose of a supplementary consistency check is to remove regions with consistent flow but a high variation in wind speed from being flagged. The low order KL model cannot model rapid cell-to-cell variation in wind speed resulting in superfluous flagging of fine-scale wind speed features. Since the wind speed of a cell is not dramatically affected by the ambiguity selection, rapid speed variation does not generally indicate ambiguity selection errors. This section develops a "consistency check" that determines whether or not there is sufficient directional variation in the region to classify it as an ambiguity selection error.

Inconsistencies in the wind from ambiguity selection errors cause the wvc directions of an observed wind field to drastically change from one part of the region to another. When a histogram of all the wind directions in a region is made, regions with ambiguity selection errors typically yield a multimodal histogram. That is, the selected wind vectors suggest more than one major flow. Figure 13 shows a region of ambiguity selection error. Notice the histogram has more than one

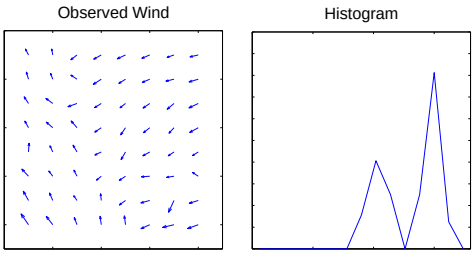

Figure 13: *Wind field exhibiting ambiguity selection errors and a histogram of the directions of all wvcs in the region.*

significant peak. This is an indicator of inconsistent wind fields within the region.

A histogram of a region's directions is identified as multimodal in the following way:

1. The histogram is made with a bin spacing of 24°. This number was chosen because it maximizes the ambiguity selection error detection rate.
2. The histogram is reordered with the minimum valued bin first.
3. The histogram is differentiated with a numerical approximation.
4. Multiple modes are detected through multiple changes from positive to negative in the derivative (zero crossings).

This simple consistency check complements the model based detection scheme by providing an additional view of the consistency of a region without some of the problems associated with the restricted basis. However, because this consistency check is not geophysically based, it is insufficient in fully determining the quality of a region. Nevertheless, in the training set, it is effective in ruling out over 91% of regions not containing errors.

6.3.3 RMS Wind Speed Threshold

An RMS wind speed threshold is additionally applied. Regions with rms wind speeds less than 3.5 m/s are considered not processable because the SNR may be too low to create valid wind direction estimates. From experience with NSCAT, most regions with rms wind speeds less than 4.0 m/s were flagged primarily because of noise. Through examination of SeaWinds data, we concluded that it is very difficult to subjectively locate ambiguity selection errors when the region rms wind speed is less than 3.5 m/s. Approximately 7% of the total number of regions fall beneath this threshold.

6.4 Summary of ASE flag

As presented in this section, regions are individually flagged according to the following criteria:

1. Over 14% of vectors in the region are flagged with either the variable direction or vector error thresholds.
2. The rms error between the observed wind and the model fit is greater than 1.8 m/s.
3. A histogram of the wvc directions in the region is multimodal.
4. The rms wind speed of the region is greater than 3.5 m/s.

In order to flag a region as an ambiguity selection error, all of these criteria must be met.

7 PERFORMANCE OF THE ASE FLAG ON THE TRAINING DATA SET

The optimized ambiguity selection error flag is applied to the training data set. This section relates the probability of false alarm (P_{FA}) and missed detection (P_{MD}) for several test conditions.

7.1 Overall False Alarm and Missed Detection Rate

Table 3 shows the results for the algorithm applied the the training data set using the following flagging schemes: (A) using direction thresholds only, (B) using both direction and vector thresholds, (C) including the supplementary consistency check and the rms region threshold, (D) including the rms wind speed threshold, (E) and not counting a region as a missed detection if an overlapping region is flagged. The following observations are made.

Table 3: *False alarm and missed detection rates for the ambiguity selection error detection algorithm for various combinations of flagging criteria.*

Case	Thresholds	Other	P_{FA}	P_{MD}
A	1	-	0.025	0.18
B	1,2	-	0.053	0.062
C	1,2	3	0.017	0.132
D	1,2,4	3	0.015	0.135
E	1,2,4	3,5	0.015	0.030

Key
1: Numerically determined direction thresholds from Fig. 9, Region threshold of 0.14
2: Numerically determined vector thresholds from Fig. 10, Region threshold of 0.14
3: Including the supplementary consistency check and the rms wind speed threshold
4: RMS error threshold of 1.8 m/s
5: Not counting a region as a missed detection if another region with 50% overlap was flagged

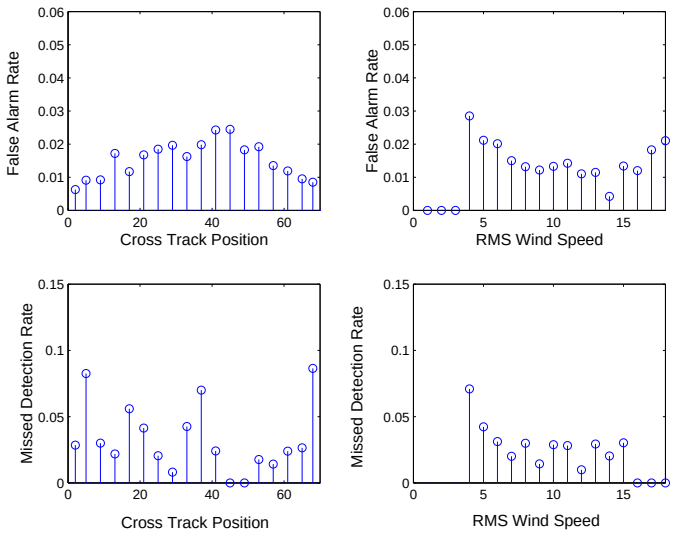

Figure 14: *False Alarm and Missed Detection Rates for the algorithm per cross track and rms wind speed. This data is taken from a test set of 15 revs.*

1. When both direction and vector error thresholds are used, the false alarm rate increases, but there is a noticeable reduction in missed detections (see case B opposed to case A).

2. Using the additional criteria dramatically reduces the false alarm rate (see case D opposed to case B).

3. Not counting a region as a missed detection if another region with 50% overlap is flagged lowers the missed detection rate (see case E).

From these results, we see that the detection rate for regions containing ambiguity selection errors is approximately 97% while the false alarm rate is less than 2%.

7.2 False Alarm and Missed Detection Rates per Cross Track and RMS Wind Speed Bin.

Figure 14 shows the false alarm and missed detection rates of the ambiguity selection error flag for different cross track positions and rms wind speeds. For speeds from 5 to 15 m/s, the quality assurance algorithm yields a rather constant false alarm and missed detection rate. The algorithm is slightly less efficient in the 4 m/s bin. In addition, there is a slight tendency toward increased false alarms at nadir because the additional criteria is not variable with swath location.

8 QA BIT FLAG

In order to make the QA data more accessible, we introduce a QA bit flag that can be used in conjunction with the standard wind product. Here, a level bit flag is introduced that indicates the inferred quality of each cell in a rev. Wind data is generally stored as a swath-shaped array with each element representing a wvc. A corresponding QA flag array is produced by the quality assurance algorithm. The flag is a 4 bit integer at each wvc. The following sections describe the meaning of each bit.

8.1 Individual WVC Flag

The two least significant bits of the QA flag, $[q_1, q_0]$ indicate individually flagged vectors. Bit q_0 is set when a vector is flagged using the constant thresholds in any of the overlapping regions (noisy vector flag). Bit q_1 is set when a vector is flagged using the variable thresholds (possible ambiguity selection error or ASE cell flag). Recall the variable thresholds are set higher in nadir and low wind speed regions to produce a constant false alarm rate over the swath, while the constant thresholds have variable performance over the swath. The constant thresholds serve as a noisy vector flag while the variable thresholds indicate vectors more likely to be ambiguity selection errors. In summary, the first two bits have the following meanings:

q_1, q_0

×, 1 - *Noisy vector*: The cell is flagged using constant thresholds in at least one of the overlapping regions, indicating a generally noisy wind vector.

1, × - *ASE cell flag*: The cell is flagged using variable thresholds in at least one of the overlapping regions. Multiple neighboring flagged cells may indicate ambiguity selection errors in the vicinity.

8.2 Region WVC Flag

The two most significant bits, $[q_3, q_2]$ correspond to the region flag. Because the 8 × 8 regions overlap, the wvc is classified with the highest classification for any of the overlapping regions. The flagging scheme is:

q_3, q_2

0, 0 - *Good:* All overlapping regions containing the wvc are flagged "good." Each region contains less than 5% of the individual wvcs flagged by the constant thresholds. For these regions the wind flow fits the model estimate very well and is spatially consistent having a low noise level.

0, 1 - *Fair:* At least one overlapping region containing the wvc is flagged "fair." The region contains 5-20% of its wvcs flagged by the constant thresholds. The wind flow is consistent, but some vectors may contain moderate amounts of noise and/or possible ambiguity selection errors. Wind fields with fine scale spatial variations (e.g. fronts) may also be flagged.

1, 0 - *Poor:* At least one overlapping region containing the wvc is flagged "poor." The region contains more than 20%

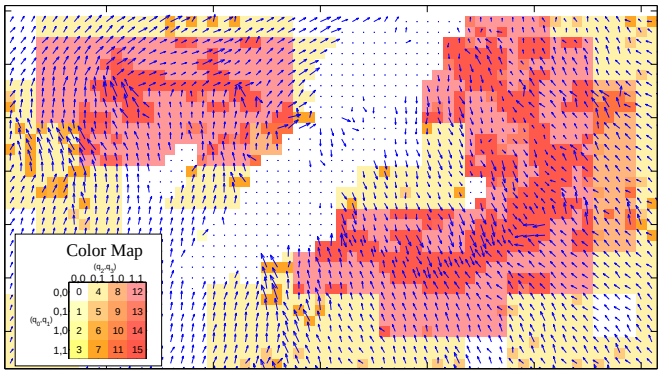

Figure 15: *QA bit flag example containing a possible region of ambiguity selection error. Missing vectors indicate the position of land. Notice that the QA algorithm flags the position of inconsistencies in the estimated wind flow.*

of wvcs flagged by the constant thresholds. The wind flow is not consistent due to ambiguity selection errors high levels of noise. Nadir regions and low wind speed regions (both of which are noisier) are more likely to flagged "poor."

1, 1 - *ASE region flag:* At least one overlapping region is flagged as containing ambiguity selection errors. The region exceeds 14% cells flagged by the variable wvc thresholds, has an rms wind speed of above 3.5 m/s, the rms error between the observed and the modeled wind is more than 1.8 m/s, and the histogram of the wind directions in the region is multi-modal. Given all of these considerations, the region is estimated to contain substantial ambiguity selection errors. Very high noise corruption and rain contamination may also cause the region to be flagged.

The algorithm is performed for each rev and provides a flag value for each wvc. Cells without wind estimates are flagged with zeros. Note that isolated wvcs or isolated small wvc groups cannot be effectively flagged and are, by default, flagged zero. The QA flag can be viewed as an integer number from 0 to 15 indicating the increasing uncertainty in the correctness of the wind flow at that cell. Figure 15 illustrates a region containing an ambiguity selection error flagged by the algorithm. Additional examples of the QA bit flag are shown in Figure 22. Notice that the algorithm clearly identifies inconsistencies in the estimated wind flow. This information can be used to locate isolated regions of ambiguity selection errors in order to correct them.

9 QUIKSCAT QA RESULTS

In this section, the quality assurance algorithm is applied to 16 months of SeaWinds data and a statistical analysis is performed of regions flagged by the algorithm.

Table 4: *Overall results of the QA algorithm on a the entire mission and the training data set. Also, the percent of ambiguity selection errors manually flagged*

Region Classification	Overall	Training Set (QA Algorithm)	Training Set (Manually Flagged)
Good	65.2%	63.6%	-
Fair	19.3%	19.6%	-
Poor	15.5%	16.8%	-
Contains Ambiguity Selection Errors	4.6%	4.9%	4.0%

9.1 Overall SeaWinds QA Results

Here we examine the overall percent of regions classified according to each of the general classification levels (good, fair, poor) and the percent of regions identified as ambiguity selection errors. The overall QA results are given in Table 4 along with the QA results from the training set.

The QA algorithm flagged approximately the same percent of ambiguity selection errors as were manually flagged in the training set (within 1%). The higher percent of estimated ambiguity selection errors flagged by the algorithm are due to the non-zero false alarm rate. From this data, we conclude that the overall ambiguity selection is at least 95% effective for all overlapping 8×8 regions with wind speeds greater than 3.5 m/s. Also, $\sim$ 65% of the regions are classified as "good" indicating that most of the data has very low noise level.

The overall results of the algorithm on SeaWinds data are comparable to the results of the QA algorithm developed for NSCAT [1]. With NSCAT, 65% of regions are classified as "perfect" or "good," which is equivalent to the the "good" classification for SeaWinds. In addition, 18% of NSCAT regions are classified as "poor." Also, it is estimated that the NSCAT ambiguity selection is over 95% effective. These results are nearly identical to the results determined for SeaWinds. Thus, our QA algorithm estimates that the quality of SeaWinds data is approximately the same as the quality of NSCAT data.

9.2 Cross Track/RMS Wind Speed Dependence

When comparing the overall percent of regions flagged as "poor" to those flagged as ambiguity selection errors (see Table 4), we recognize that the percentage of regions flagged as ambiguity selection errors is significantly smaller than the percentage of regions classified as "poor." To explain this, we examine the cross track and rms wind speed dependence of the region classification. Figures 16 and 17 which show the percent of regions classified as "poor" and flagged as ambiguity selection errors per cross track and rms wind speed bins. As predicted in Section 6, the majority of regions flagged as "poor" are low wind

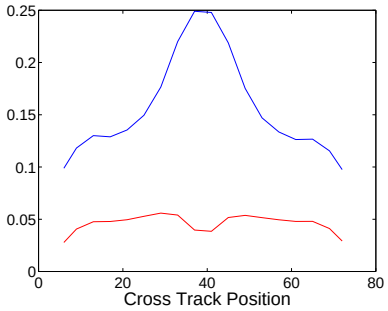

Figure 16: *Fraction of all regions flagged as "poor" (blue line) using constant thresholds and flagged as containing ambiguity selection errors (red line) using optimized variable thresholds per cross track position.*

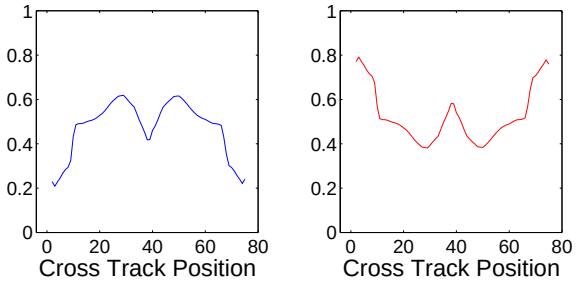

Figure 18: *Fraction of wvcs per cross track containing only 1 to 3 ambiguities (left figure) and fraction of wvcs per cross track containing 4 ambiguities (right figure) averaged over 600 revs of SeaWinds data.*

speed and nadir regions. Since the ambiguity selection error flag is desensitized to noise at nadir and in low wind speed regions, many less regions are flagged as ambiguity selection errors than are classified as "poor."

Examining the cross-track dependence of regions flagged as ambiguity selection errors in Figure 16, we see that the algorithm infers fewer ambiguity selection errors at nadir and along swath edges than in the sweet spot. To explain this phenomenon, we inspect the number of ambiguities produced per cell for each cross track position. At nadir and on the swath edges, more ambiguities are generally produced per cell. The percent of 1 to 3 ambiguity cases per cross track is compared to the percent of 4 ambiguity cases in Figure 18.

The general shape of the curve representing the 1 to 3 ambiguity cases of Figure 18 closely mirrors the percent of ambiguity selection errors per cross track shown in Figure 16. At nadir and on swath edges (where there are less ambiguity selection errors), there is a higher likelihood of having four ambiguity choices. Having more ambiguities may cause less estimated ambiguity selection errors for two reasons: First, in regions of data corruption (such as rain contamination), the nudging algorithm generally has more ambiguity choices to match the nudging field. Second, the point-wise median filter has a wider selection of possible ambiguities when creating a smooth solution. Thus, the sweet spot actually generates more inconsistencies which are flagged as ambiguity selection errors.

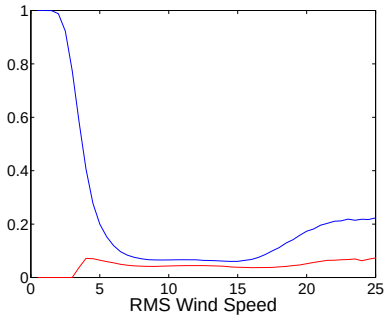

Figure 17: *Fraction of all regions flagged as "poor" (blue line) and as containing ambiguity selection errors (red line) per rms wind speed.*

9.3 Temporal QA Statistics

Next, we examine the ambiguity selection as a function of time. Figure 19 shows the ambiguity selection errors averaged over 3 days for each point. The percent of ambiguity selection errors stays nominally between 4% and 5% for the entire mission. There is a slight seasonal rise in errors from March to May and a drop in errors from May to August. This may be explained by seasonal variations in weather patterns which affect the performance of the QA algorithm [1].

In order to understand the seasonal weather pattern variations, we divide the QA data into latitude bands (see Table 5). For each band, the average number of cyclonic storms passed by SeaWinds per day and the average percentage of ambiguity selection errors per day are computed and given in Figures 20 and 21. From visual inspection, the ambiguity selection errors

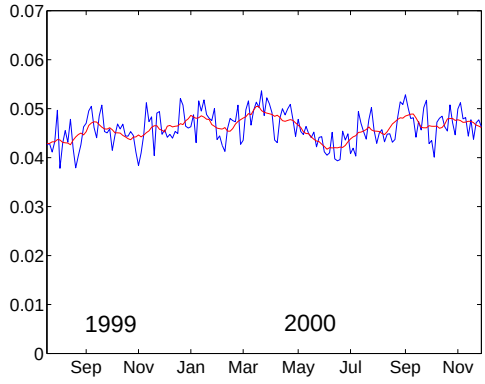

Figure 19: *Ambiguity selection errors as a function of time. The blue line is the percent of ambiguity selection errors averaged over 3 days. The red line is a 30 day moving average.*

Table 5: *Latitude bands*

Band	Range
1	$-90°$ to $-45°$
2	$-45°$ to $-25°$
3	$-25°$ to $-5°$
4	$-5°$ to $5°$
5	$5°$ to $25°$
6	$25°$ to $45°$
7	$45°$ to $90°$

closely follow the number of cyclonic storms for each latitude band. This indicates that a large part of the ambiguity selection errors occur in stormy regions.

9.4 QA Bit Flag Results

The results presented thus far have given the statistics for the QA algorithm on a region-by-region basis. However, because the regions overlap, this does not fully reflect the actual percentage of wvcs flagged by the QA algorithm. Using the QA bit flag, we generate statistics indicating the percent of cells flagged individually by the algorithm and the percentage of cells classified into each region classification. These are shown in Table 6

From Table 6, we see that the due to the overlap of regions, actual percentage of cells in regions of ambiguity selection errors is higher than the percentage of regions flagged. However, the percentage of individual cells flagged in the regions of ambiguity selection error are approximately equal to the percent of regions flagged by the algorithm ($\sim 5\%$).

Table 6: *Percent of valid individual wvcs flagged according to several QA bit flag classifications. The symbol '$\times$' indicates a "don't care." The classifications have the following meanings: A - No Flag: the cell is not excessively noisy and is not estimated to be an ambiguity selection error, B - Flagged with constant thresholds (noisy vector), C - Flagged with variable thresholds (likely to lie near ambiguity selection errors), D - "Good" region classification, E - "Fair" region classification, F - "Poor" region classification, G - "Ambiguity selection error" region flag, H - Both region and cell flagged as an ambiguity selection error*

	Classification	Bit Flag $[q_3, q_2, q_1, q_0]$	Percent wvcs
A	no flag	$[0,0,0,0]$	52%
B	noisy cell	$[\times,\times,\times,1]$	12%
C	ASE cell	$[\times,\times,1,\times]$	6%
D	"good" region	$[0,0,\times,\times]$	53%
E	"fair" region	$[0,1,\times,\times]$	22%
F	"poor" region	$[1,0,\times,\times]$	16.5%
G	ASE region	$[1,1,\times,\times]$	8.5%
H	ASE region & ASE cell	$[1,1,1,\times]$	4.4%

10 SUMMARY AND CONCLUSIONS

The QA algorithm presented here classifies SeaWinds data according to the general spatial consistency of the selected wind and identifies regions of possible ambiguity selection errors.

This report shows that by using variable wvc thresholds to flag poor wind vector cells, the effects of nadir and low wind speed noise is suppressed in locating actual ambiguity selection errors. Using these thresholds, the quality assurance algorithm locates ambiguity selection errors in regions of moderate rms wind speed (less than 3.5 m/s) with a false alarm rate of $\sim 1.5\%$ and a detection rate of less than 3%. Applying the algorithm to the SeaWinds mission, less than 5% of overlapping 8×8 regions examined exhibit ambiguity selection errors.

The QA bit flag discussed in this report is a tool to locate possible problem regions in SeaWinds data. High bit flags generally indicate ambiguity selection errors, inconsistent wind flow, high noise levels, or fine-scale wind features such as storms or fronts.

According to the QA bit flag, the actual percent of valid wvcs that occupy ambiguity selection error regions are about 8.5%. However, less than 5% of cells are flagged with both the individual wvc and region ambiguity selection error flag. The QA algorithm generally infers less errors at nadir and on swath edges. This is due to a larger set of ambiguities in these regions. However, the noise level at nadir is considerably more than at other areas in the swath. There exist a slight temporal variation of the QA algorithm's performance, but the average percent of ambiguity selection errors remains between 4% and 5% for the entire mission. The percentage of ambiguity selection errors is highly correlated to the number of storms passed per day by the instrument. Overall, the ambiguity selection of SeaWinds data is generally good.

References

[1] Gonzales, A. E., and Long, D. G. "An Assessment of NSCAT Ambiguity Removal," *Journal of Geophysical Research*, Vol. 104, No. C5, pp. 11449-11457, 1999.

[2] Draper, D. W. "Spatial Frequency Analysis and Truncation Issues for the QSCAT KL model," internal report, Brigham Young University, 2000.

[3] Long, D. G., and Mendel, J. M., "Identifiability in wind estimation from scatterometer measurements", *IEEE Trans. on Geoscience and Remote Sensing, 29*, 268-276, 1991.

[4] Shaffer *et al.*, A median-filter-based ambiguity removal algorithm for NSCAT, *IEEE Trans. on Geoscience and Remote Sensing, 29*, 167-174, 1991.

[5] Oliphant, T., and Long, D. G., "Accuracy of Scatterometer Derived Winds using the Cramér-Rao Bound", *IEEE Trans. on Geoscience and Remote Sensing, Vol. 37 No. 6*, 2642-2652, 1999.

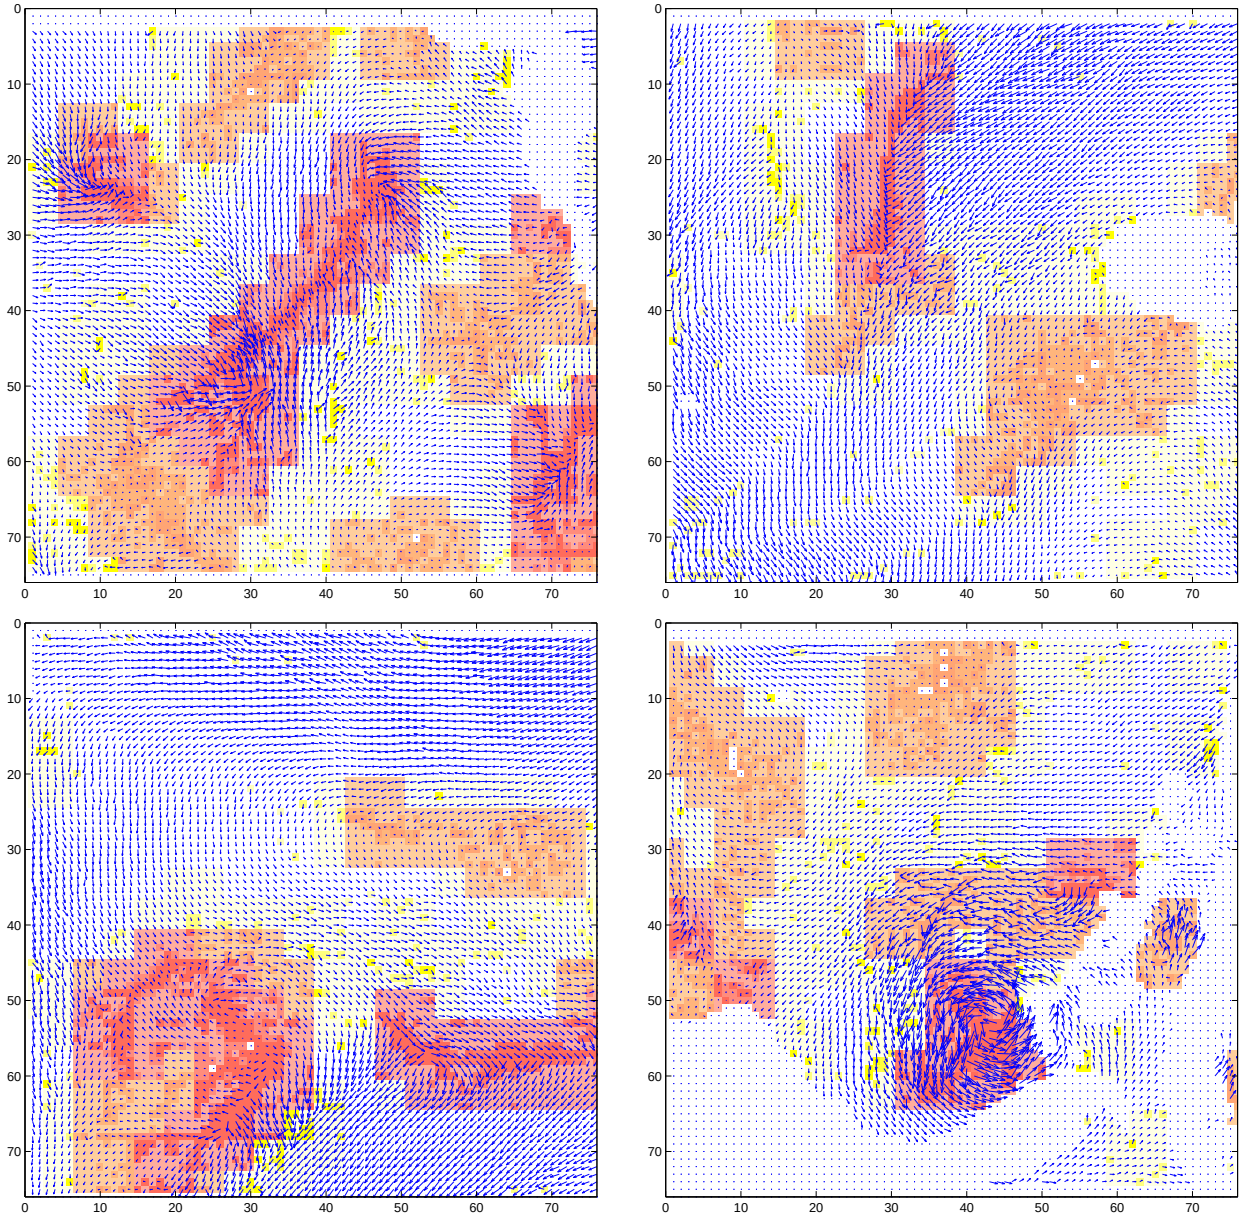

Figure 20: *A few examples of the QA bit flag in stormy regions or regions of ambiguity selection error.*

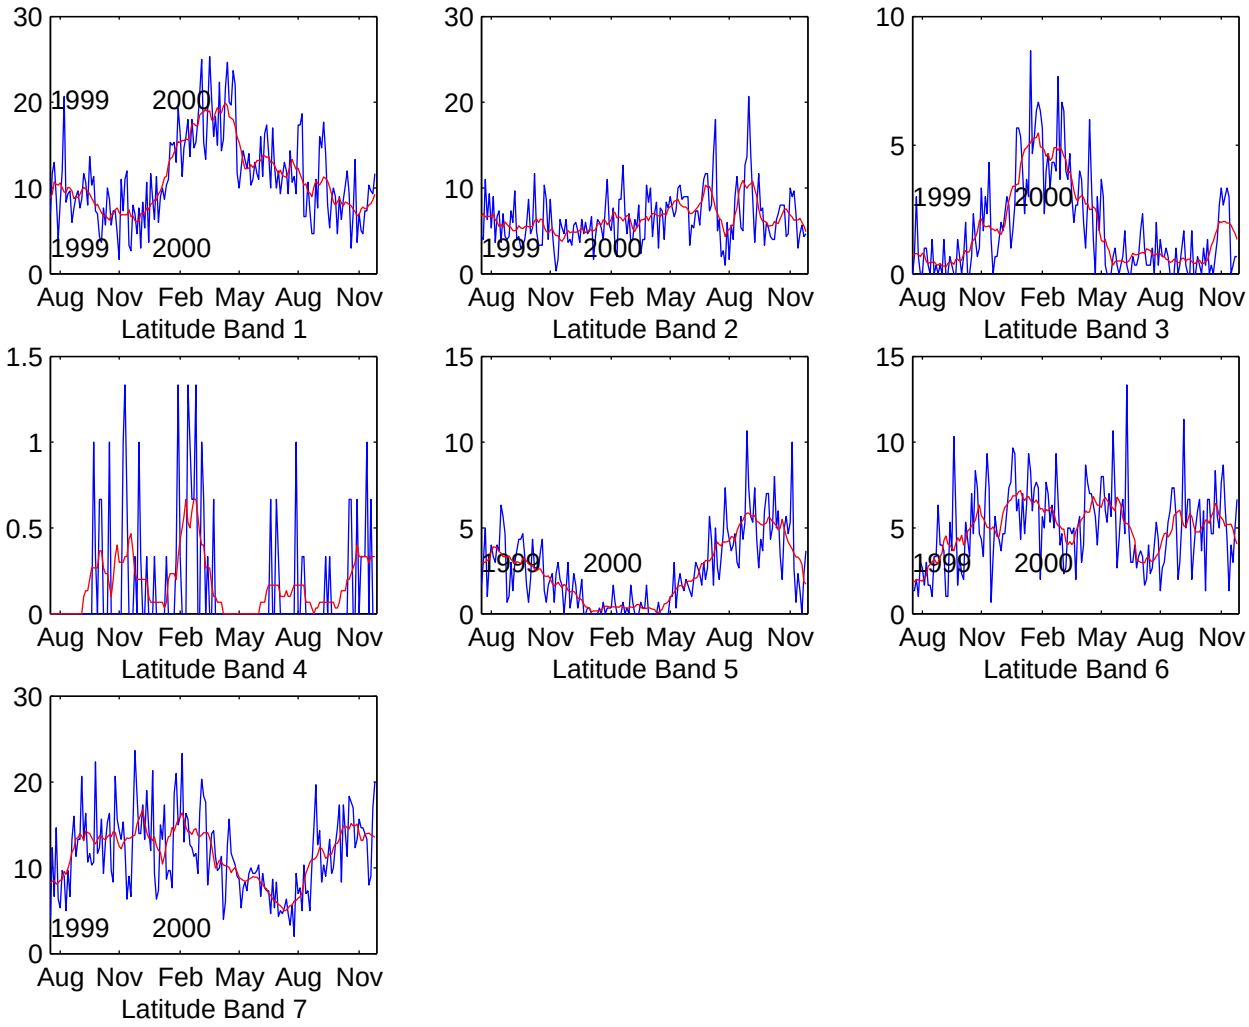

Figure 21: *Number of cyclonic storms passed per day by SeaWinds averaged over three days (blue) and moving averaged over 30 days (red) for each latitude band.*

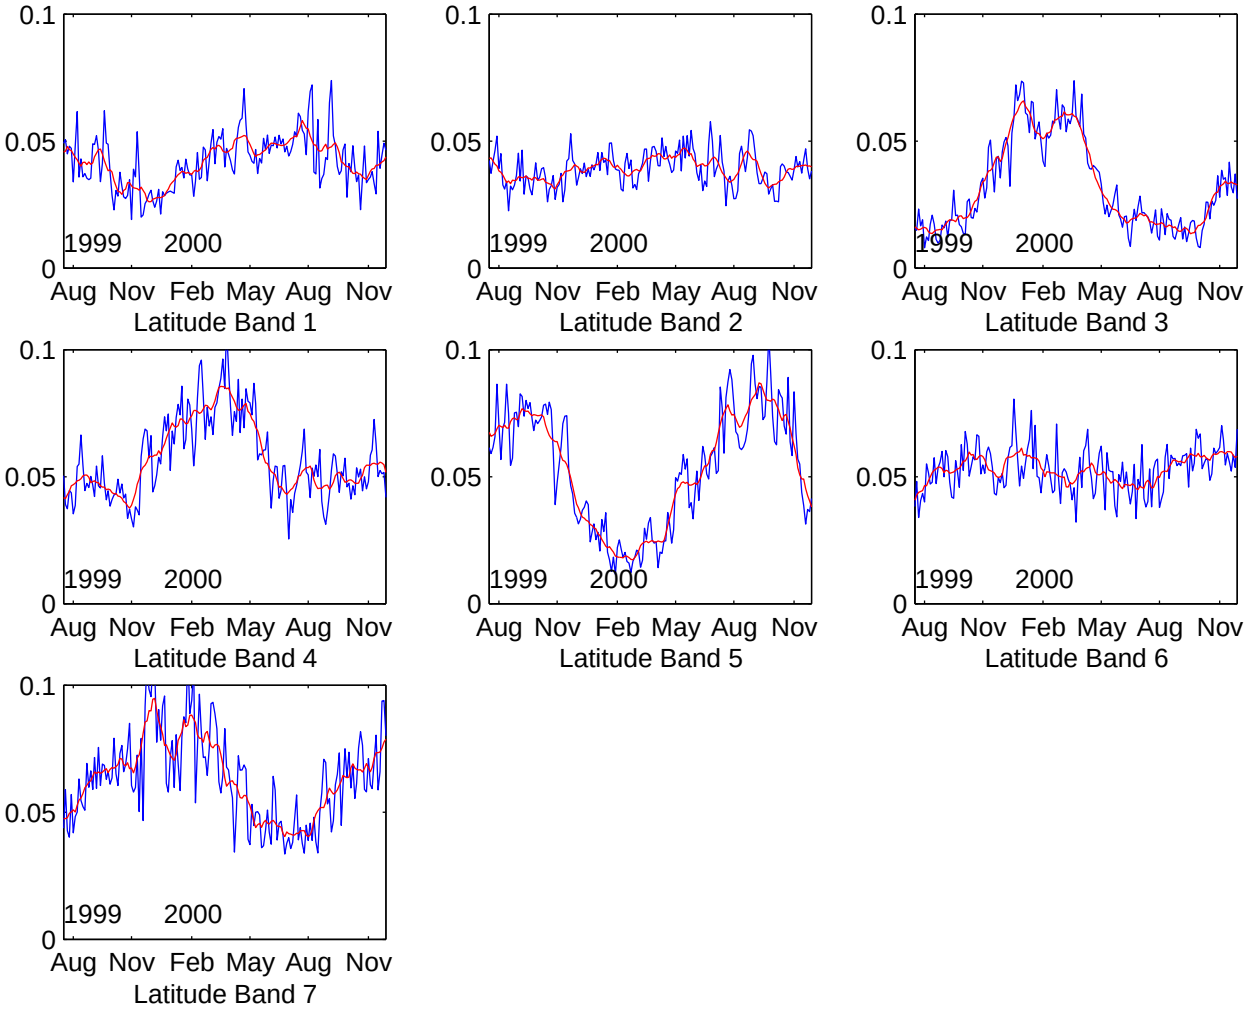

Figure 22: *Percentage of ambiguity selection errors averaged over three days (blue) and moving averaged over 30 days (red) for each latitude band.*