A PROJECTION METHOD FOR AUTOMATIC ESTIMATION OF WIND VECTORS WITH RADARSAT SAR IMAGERY

Christopher Wackerman [(1)], William G. Pichel [(2)], Pablo Clement-Colón [(2)]

[(1)]General Dynamics Advanced Information Systems, P.O. Box 134008 Ann Arbor Michigan 48113-4008 USA
chris.wackerman@gd-ais.com
[(2)]NOAA/NESDIS, WWBG E/RA3, Room 102 5200 Auth Rd. Camp Springs MD 20746-4304 USA
William.G.Pichel@noaa.gov

ABSTRACT

A new approach for automatically estimating wind vectors for SAR imagery is presented, which relies on a projection operation to generate wind direction estimates. A threshold can be applied to the projection results in order to eliminate estimation of wind directions in regions where the SAR image contains no visible features. The process also allows multiple possible directions to be generated at a given location using the image features, which are then resolved to a single direction based on uniformity across local regions of the image. The algorithm has been validated using 137 comparisons to in situ buoy observations, giving a direction RMSE of 39° (mean error of 10°, error standard deviation of 38°) and a wind speed RMSE of 2.2 m/s (mean error = -1.4 m/s, standard deviation of the error = 1.7 m/s). The largest errors come from automatically utilizing strong features in the image that are not aligned with the local wind, such as atmospheric lee waves are current fronts. If we manually eliminate these from the comparisons, then the direction RMSE is 31° (mean error of 20°, error standard deviation of 20°) and the wind speed RMSE is 2.1 m/s (mean error = -1.2 m/s, standard deviation of the error = 1.7 m/s).

1. INTRODUCTION

As part of the NOAA/NESDIS Alaska SAR Demonstration Project [1], a multi-year demonstration of the production and use of RADARSAT SAR HH polarization imagery to generate products in a pre-operational environment, a wind product is created that automatically generates wind vectors over the coastal ocean. Two methods are used. One uses wind directions from atmospheric models that are run over the same region and close in time to the SAR image collection. These directions are then combined with the radar cross section from the SAR imagery to generate wind speed using an empirically derived correction to a VV model to generate the HH relationship between radar cross section and wind speed [2]. The second method estimates the wind direction from the SAR imagery itself using large-scale features that are aligned with the local wind, such as wind rows, elongated convective cells, or surfactant streaks, then combines this direction with the radar cross section of the imagery to generate wind speed. The currently implemented algorithm for estimating wind direction from the SAR imagery is based on a spectral approach that uses the Fourier transform of image subsets to automatically determine directions of maximal spectral energy, then assumes that these directions correspond to 90° from the wind direction [3,4]. This direction is then combined with the image radar cross section to generate wind speed using a semi-analytical two-scale model to derive the HH relationship between radar cross section and wind speed [5]. The directions estimated from the SAR imagery have an inherent 180° ambiguity, since from a static image one can at most estimate the line along which the wind is blowing, but not the direction along that line.

One significant problem with estimating wind directions from SAR imagery is that there may not always be features in the image that the algorithm can use to estimate directions. Convective cells, wind rows, and surfactant streaks depend on there being other processes going on (such as turbulence or very little mixing of the near-surface ocean layer) rather than simply the existence of a wind field. The spectral approach that is currently implemented has no means of determining when a wind direction is able to be estimated from a region of the SAR image, and thus will often generate erroneous directions over featureless regions of the ocean surface. Therefore a study was initiated to develop a new approach for estimating wind directions, referred to here as the projection method, that would be able to recognize when no robust estimate of direction was able to be extracted from the imagery.

2. PROJECTION METHOD FOR WIND DIRECTION ESTIMATION

The user determines a window size that will be moved through the image, and for each placement of the window one wind vector will be estimated. The window size should be driven by the scales of the features that are going to be used to estimate wind direction. For wind rows, these are typical from 3 to 10 km, so usually the window should not be much

smaller than 10 km. For the results in this paper, a 24 km window was used and the window was shifted every 16 km to generate a new wind vector estimate, thus there is some overlap of image samples used to create successive wind vectors. For a given window placement, the projection of the image samples within the window along a direction are generated for directions from -90° to +90°. That is, for each direction a one-dimensional function is generated by stepping through the middle of the window in the specific direction, and at each image sample along that direction averaging all of the image samples that are in the orthogonal direction. For the results in this paper a projection was calculated every 1°, however in the results below (particularly Fig. 2) it can be seen that one could use a coarser angular sampling of the projections as long as the peaks of the feature contrasts (as defined below) can be resolved. This projection is then flattened to remove linear trends in the function.

Fig. 1 shows an example where the top image displays the SAR image values within a local window, and the bottom image displays the resulting projections for the full range of angles, with a different projection at each line. For each flattened projection, the contrast (standard deviation divided by the mean) is calculated and a plot of contrast versus direction (i.e. the direction of the projection that generated the contrast value) is created. Fig. 2 shows the resulting contrast plot generated from the projections shown in Fig. 1. All of the local peaks that have a contrast above some threshold are found, and 90° from these directions are considered as candidate wind directions for this window. For the results in this paper, the contrast threshold was set to 0.03 based on manual analysis of the obvious visual features in the imagery. Note however that although this metric is invariant to scale factors, it may depend on the processor used since the contrast of a feature will change whether the image was saved as 8 bits, 16 bits, in log values, etc. For the plot shown in Fig. 2, local peaks would have been found around -25°, +25°, and +50°. For each possible wind direction, the radar cross section of the image (calculated over the local window) is used to generate a wind speed via the model describe in Ref. 5.

Once the local window has been moved through the SAR image, there will in general be multiple possible wind directions (or no wind direction if no contrast value was above the threshold) for any image location. To determine which wind direction to use, we first choose the direction at each location that generated the maximum contrast. Then we iteratively move through the image as follows. At each location, we calculate the norm of the vector difference between each possible wind vector (direction and speed) at that location and all the other wind vectors that we are

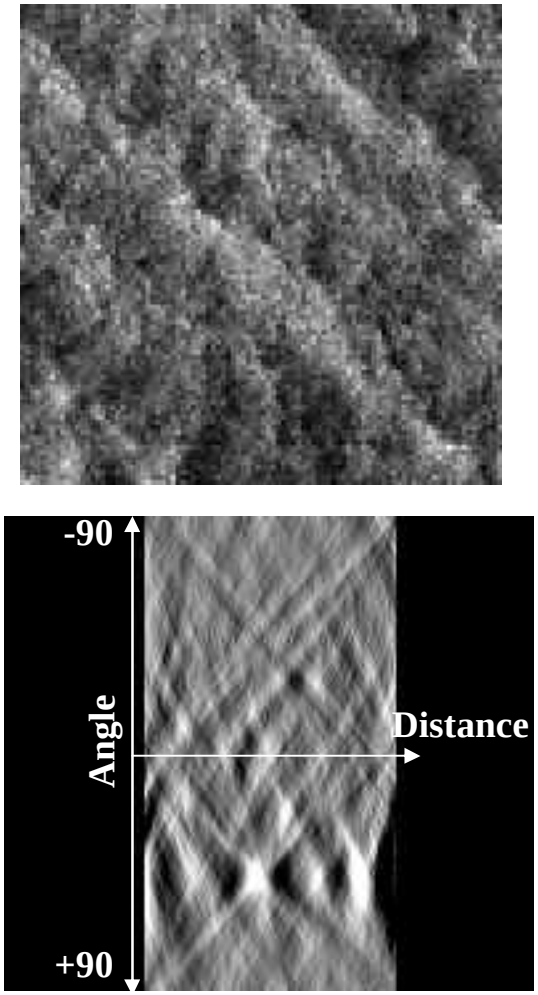

Fig. 1 Top image shows a subset of a SAR image. Bottom image shows the resulting projections from the image subset at various angles.

assuming are correct at the other locations, where we weight the vector sums by a Gaussian fall-off function whose width represents the spatial scales over which we assume that the wind vectors would be correlated. For all the results in this paper, we set the Gaussian width to four window lengths which represents 96 km. We then replace the wind vector at this location with the vector that generated the minimum vector norm (note that it might be the same vector as we started with). This is done at each location as we move through the image, using the new vectors at previous locations and the old vectors at future locations. This iterative process is repeated until the image converges to a solution; i.e. until the process does not change any wind vectors at any location.

Thresholding the contrast of the projection allows the algorithm to not generate a wind direction for image

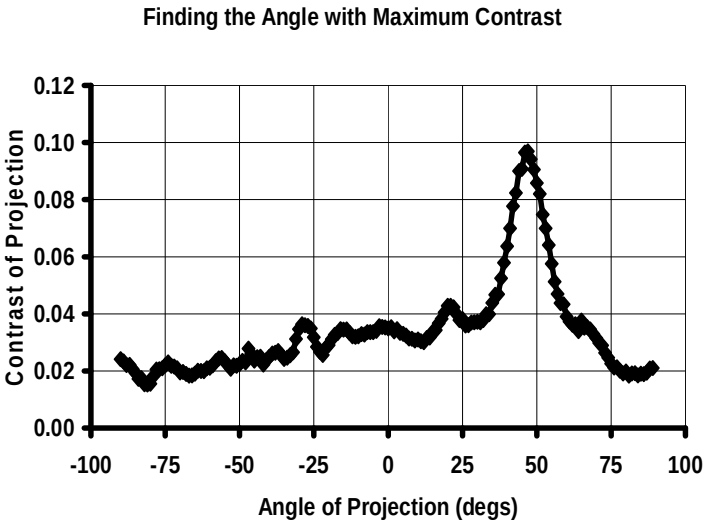

Fig. 2 A plot of contrast versus projection angle for the projections shown in Fig. 1. Local maxima at angles of approximately -25, +25 and +50 degrees are candidate wind directions.

regions where there are no features with sufficient modulation to be used for wind direction estimation.

Fig. 3 shows the results of varying the contrast threshold for a SAR image. The top image shows direction results (indicated as blue lines on the image) when no threshold is applied. The middle and bottom images show results for thresholds of 0.03 and 0.06 respectively. Note that as the threshold increases, we filter out all but the highest contrast features to use, and automatically eliminate the low contrast regions which were generating erroneous directions under the spectral approaches.

The final step in estimating wind vectors is to apply a median-like operation to the wind vectors to eliminate outliers. This is done by replacing the wind vector in the middle of a local box by the vector within the box that minimizes the norm of the difference between that vector and all of the others. Note that this approach has the advantage of only using wind vectors that have been estimated from the image (i.e. we do not average vectors such that a resulting vector was never seen anywhere in the image), and of preserving wind fronts in the final wind vector map. Also note that this final step is the spatial equivalent of the approach we use to determine a unique vector at each location. Instead of choosing the vector from a list of possible vectors at a given location, we chose the vector from the list of neighboring vectors. The size of the median filter window then plays the equivalent role of the width of the Gaussian fall-off weighting. For the results in this paper we used a median filter size of three local window lengths (72 km) and repeated the process over the entire image two times.

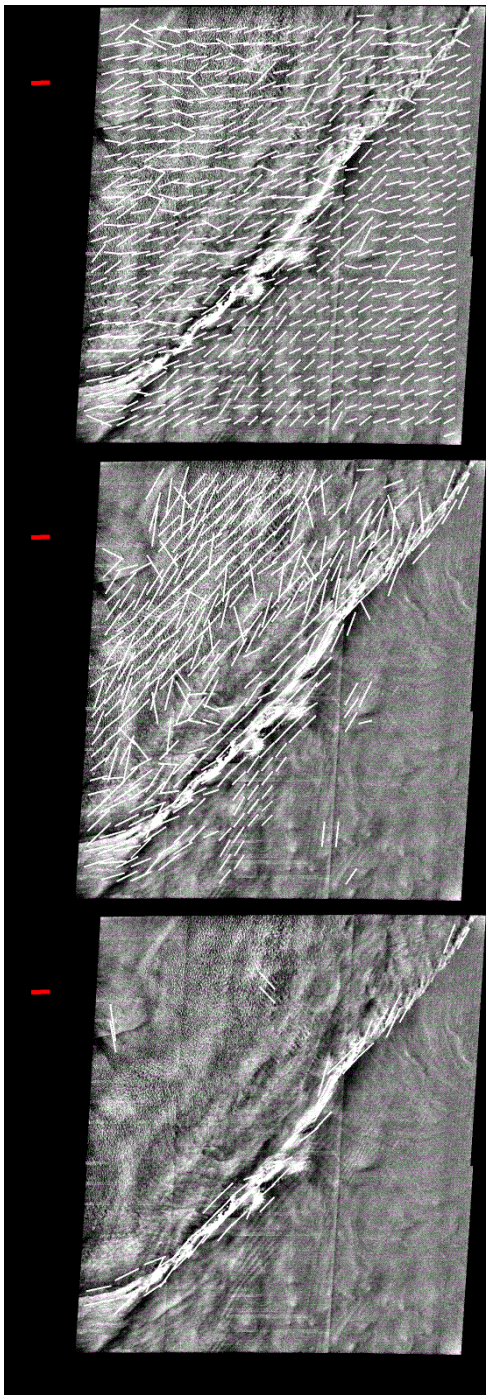

Fig.3: Examples of using the threshold on projection contrast to remove image regions with no strong features. Top image shows estimated wind directions (white lines) with no threshold, middle images has a threshold of 0.03 and the bottom image 0.06. Note that as the threshold increases, we eliminate all but the highest contrast features. The red line at the left shows the buoy derived wind direction. Image was collected March 29 2000 at 10:51:14.

3. PROJECTION ALGORITHM PERFORMANCE

To determine how well the projection approach estimated the actual wind direction we used a series of RADARSAT SAR images collected off the east coast of the United States for which there was buoy wind information generated approximately simultaneously with the image acquisition and located spatially within the SAR image. The SAR imagery was processed at the Alaska Satellite Facility and represented 100 m resolution imagery with 50 m sample spacing. The buoy winds were converted to 10 m winds to make them consistent with the radar cross section models being used to estimate wind speed. The buoy information was compared to wind vectors estimated from the SAR image derived close to where the buoy was located. We eliminated any comparisons that occurred at incidence angles less than 25° due to possible calibration errors in the processed image, any comparisons for which the SAR-derived vector was more than 24 km away from the buoy location, and any comparisons for which the buoy wind speed values were less than 5 m/s, since for these locations there may be no significant radar cross section response from the ocean surface. Note that although 24 km is somewhat far from the *in situ* observations for comparison, it was used since that was the size of our local image window, and putting a more severe constaint (e.g. half the window size) resulting in too few comparisons.

These constraints resulted in a total of 137 comparisons between buoy observations and wind vectors estimated from the SAR imagery. A plot of SAR-derived wind direction vs. buoy wind direction is shown in Fig. 4. Due to the 180° ambiguity in the SAR-derived directions, a value of 180 was added or subtracted from the SAR-derived direction to get it within ±90° from the buoy direction. In Fig. 4 the solid black line represents a perfect answer, i.e. the SAR-derived direction equals the buoy direction, and the dashed lines are ±90° from the solid line which represents the region of possible SAR-derived values. The final root-mean-squared error (RMSE) is 39° (mean error of 10°, error standard deviation of 38°). Note that if we were to randomly assign a direction uniformly in the range of [-90°, +90°] we would expect a RMSE of 52°, which looks like we are only slight better than a random guess. However, an examination of the errors in Fig. 4 shows that they are not uniformly distributed, but rather tend to cluster around the "perfect" line. This is shown quantitatively in Fig. 5, where a histogram of the absolute value of the direction error is plotted. The x-axis of Fig. 5 is the top value of the error bin for the histogram. Thus the first point

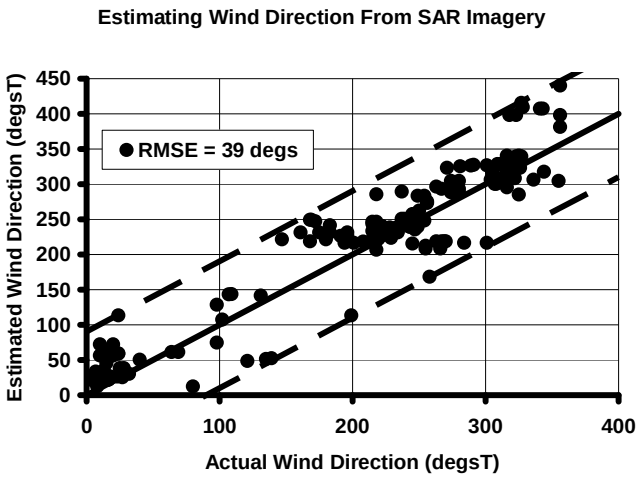

Fig. 4: Scatter plot of actual wind direction from buoy observations vs. SAR-derived wind directions. The total RMSE is 39 degs.

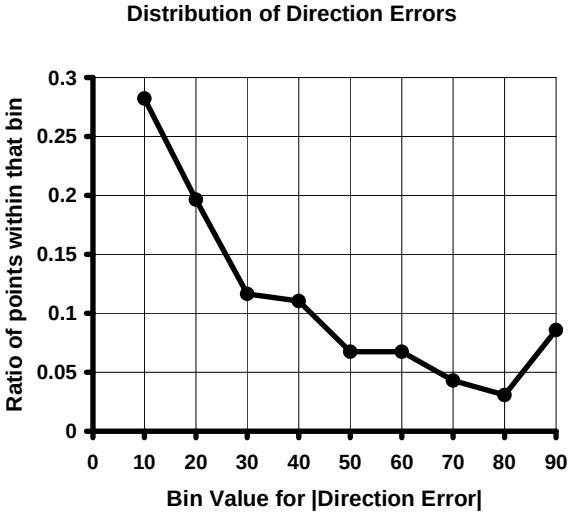

Fig. 5: Histogram of the absolute value of the direction error. X-axis is the bin values for the error, y-axis is the ratio of points that fall into that bin. Note that 60% of the samples have errors less than 30 degrees

represents errors of 10° or less, the second points represents errors between 20° and 10°, etc. Fig. 5 indicates that 60% of all of samples have an absolute direction error that is less than 30°. Note that if we were uniformly guessing a direction, the histogram in Fig. 5 would be a flat line around 0.1. The results in Figs. 4 and 5 were generated automatically; the algorithm was run with no user interaction or modification of the resulting directions.

Fig. 4 shows that there are a number of comparisons for which the estimated wind direction is as worse as it can get; i.e. 90° from the true direction. In fact, Fig. 4 shows a number of isolated points that are almost right

on the ±90° dashed lines, particularly for the lower dashed line. From manually examining these images we have found that typically these are caused by there being a feature in the image that has high contrast but is not aligned with the local wind. Figs. 6 and 7 show examples of such situations where the blue lines are the SAR-derived wind vectors and the red lines are the buoy wind vectors. In Fig. 6, convective cells can be seen that are due to downdrafts of wind and that appear as "holes" or "hills" in the image. The algorithm keys on the portion of the convective cell that has the higher contrast, which for the middle portion of the image tend to be the edge of the cell oriented vertically. The wind however is blowing right through the middle of the cell horizontally, as is indicated from the buoy observation. In Fig. 7 the algorithm is keying on a current front that is running along the coast, yet the wind is actually blowing orthogonal to the land and thus orthogonal to the current fronts. Note that in Fig. 7 the algorithm is keying on a feature that is not even physically connected to the wind.

These results indicate that one of the significant remaining problem with automated extraction of wind directions from SAR is being able to automatically classify the image features (i.e. into convective cell, current front, wind row, wind front, surfactants, etc.) so that we can eliminate those high contrast features that

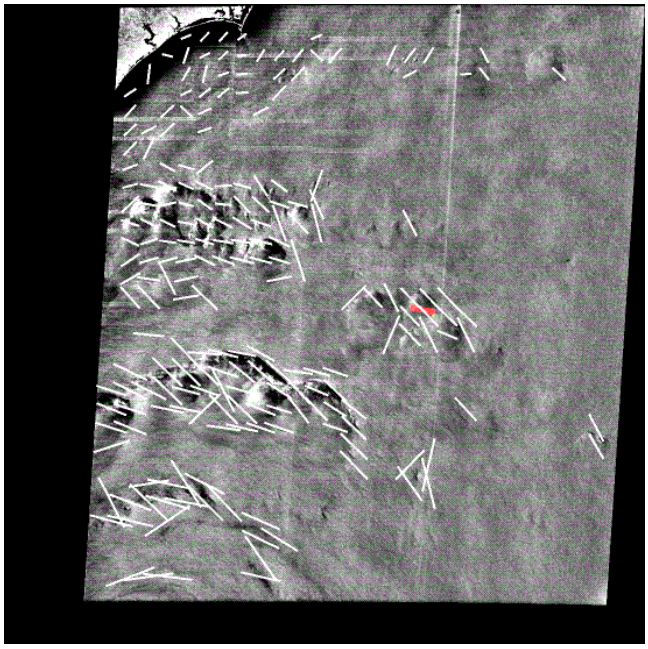

Fig. 6: Example of a feature for which the direction of maximum contrast does not align with the local wind. White lines indicate SAR-derived direction, red lines indicate buoy directions. For the convective cells, the maximum contrast direction is 90 degrees from the actual wind direction. Image was collected August 29 1999 at 11:04:55. ©Canadian Space Agency 1999.

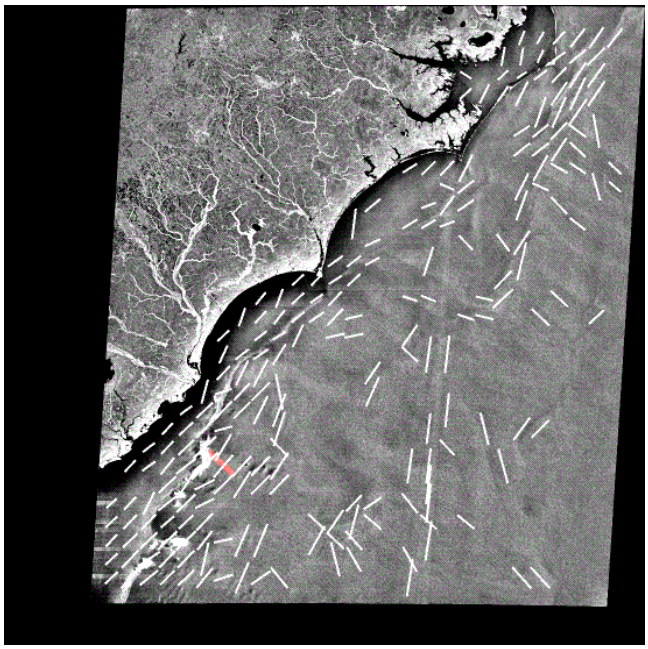

Fig. 7: Example of a feature not aligned with the local window. In this case the algorithm is keying on a current front along the coast (white lines) whereas the actual wind is blowing orthogonal to the coast (red lines). Image was collected January 20 2000 at 11:12:51. ©Canadian Space Agency 2000.

are not connected with the local wind (such as current fronts) or which need to be treated differently then just aligning with the direction of highest contrast (such as convective cells). In fact, by manually examining the imagery, we have found that almost all of the wind direction errors that are in the [80°, 90°] bin of Fig. 5 are caused by this problem. Note that the histogram in Fig. 5 has a significant increase in the number of errors in this bin versus the general fall-off within the other bins. If we eliminate these comparisons, the RMSE for wind direction drops to 31°.

Fig. 8 shows a comparison of estimated wind speed versus buoy wind speed for this data set. The black dots used the wind direction as derived from the SAR image and the radar cross section model in Ref. 5. The RMSE is 2.2 m/s (mean error = -1.4 m/s, standard deviation of the error = 1.7 m/s). The red dots used the buoy wind directions and the same radar cross section model; the RMSE is 2.0 m/s (mean error = -0.7 m/s, standard deviation of the error = 1.9 m/s). Note that there is not much difference in the RMSE or standard deviation of the error between using the SAR-derived directions or the true direction (as observed by the buoy). However there is a significant change in the mean error; it is a factor of 2 worse using the SAR-derived directions. If we remove all the comparisons

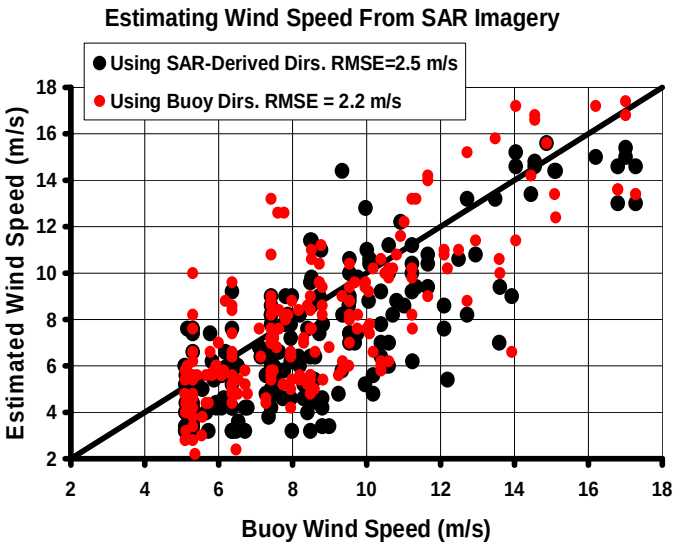

Fig. 8: Estimated wind speed from the SAR imagery vs. buoy wind speeds. Black dots are using SAR-derived directions and have an RMSE of 2.5 m/s, red dots are using buoy directions (i.e. truth) and have an RMSE of 2.2 m/s.

for which the SAR-derived wind direction has an absolute error greater than 80° (as discussed above these points are assumed to be dominated by image features not aligned with the local wind), then the wind speed RMSE drops to 2.1 m/s, with a mean error of -1.2 m/s and a standard deviation of the error of 1.7 m/s. Note that removing these points does not change the error standard deviation at all; the decrease in RMSE comes strictly from the improvement in the mean error.

The bias in wind speed estimation, even using buoy wind directions, is an issue. It may be indicative of a problem with the radar cross section model used [5], calibration with the imagery, or it may come from the fact that we are comparing estimates of the local wind averaged over 24 km to buoy point measurements. Future work will attempt to address these issues.

4. SUMMARY

As part of the NOAA/NESDIS Alaska SAR Demonstration Project, we have developed an automated algorithm for estimating wind vectors from SAR imagery that will not generate wind information over regions of the image that do not contain features. The algorithm is based on projections of the image subsets and uses a projection direction of high contrast to estimate the wind direction. When compared to 137 buoy observations, the direction RMSE is 39°, with 60% of the estimates having an absolute value of the error less than 30°. A significant remaining problem is finding an approach that can automatically determine the nature of the image feature, and remove those that are unconnected to the local wind (such as current fronts) and modify others for which the wind is not necessarily aligned with the high contrast features (such as convective cells). If we manually remove these, the RMSE drops to 31°. The RMSE for wind speed estimation is 2.2 m/s using this approach, with a large bias of -1.4 m/s. Manually removing those comparisons that are using features not aligned with the local wind, the wind speed RMSE drops to 2.1 m/s with a bias of -1.2 m/s. If we use buoy wind directions instead of those derived from the SAR image, the wind speed RMSE is 2.0 m/s with a bias of -.7 m/s. The cause of the large negative bias in wind speed is under investigation.

5. ACKNOWLEDGEMENTS

This study was supported and monitored by the Office of Research and Applications of the National Oceanic and Atmospheric Administration (NOAA) under ONR Contract N00014-00-D-0114.

The views, opinions, and findings contained in this report are those of the authors and should not be construed as an official National Oceanic and Atmospheric Administration of U.S. Government position, policy or decision.

6. REFERENCES

1. Pichel, W., Clemente-Colón, P. NOAA Coastwatch SAR Applications and Demonstration, *John Hopkins APL Tech. Dig.*, vol. 21(1), 49-57, 2000.
2. Monaldo, F.M., Thompson. D.R., Beal, R.C., Pichel, W.G., Clemente-Colón, P. Comparison of SAR-Derived Wind Speed With Model Predictions and Buoy Measurements, *IEEE Trans. Geosc. Remote Sens.*, vol. 39, 2587-2599, 2001.
3. Wackerman, C.C., Rufenach, C.L., Shuchman, R.A., Johannessen, J.A., Davidson, K.L. Wind Vector Retrieval Using ERS-1 Synthetic Aperture Radar Imagery, , *IEEE Trans. Geosc. Remote Sens.*, vol. 34, 1343-1352, 1996.
4. Fetterer, F., Gineris, D., Wackerman, C., Validating a Scatterometer Wind Algorithm For ERS-1 , *IEEE Trans. Geosc. Remote Sens.*, vol. 36, 479-492, 1998.
5. Wackerman, C.C., Clemente-Colón, P., Pichel, W.G., Li, X. A Two-Scale Model To Predict C-Band VV and HH Normalized Radar Cross Section Values Over the Ocean, *Can. J. Remote Sensing*, vol. 28, 367-384, 2002.