TOWARDS AN OPTIMAL NOISE VERSUS RESOLUTION TRADE-OFF IN WIND SCATTEROMETRY

Brent A. Williams

Jet Propulsion Lab, California Institute of Technology

ABSTRACT

This paper approaches the noise versus resolution trade-off in wind scatterometry from a field-wise retrieval perspective. Theoretical considerations are discussed and a practical implementation using a MAP estimator is applied to the SeaWinds scatterometer. The approach is compared to conventional approaches as well as numerical weather predictions. The new approach incorporates knowledge of the wind spectrum to reduce the impact of components of the wind signal that are expected to be noisy.

1. INTRODUCTION

A scatterometer is a radar that measures the normalized radar cross section (σ^0) of the Earth's surface. Over the ocean this signal is related to the wind via the geophysical model function (GMF). The objective of wind scatterometry is to estimate the wind vector field from σ^0 measurements; however, there are many subtleties that complicate this problem–making it difficult to obtain a unique wind field estimate.

Conventionally, wind estimation is split into two stages: a wind retrieval stage in which several ambiguous solutions are obtained, and an ambiguity removal stage in which ambiguities are chosen to produce an appropriate wind vector field estimate. The most common approach to wind field estimation is to grid the scatterometer swath into wind vector cells and estimate wind vector ambiguities independently for each cell. Then, field-wise structure is imposed on the solution by an ambiguity selection routine. Although this approach is simple and practical, it neglects field-wise structure in the retrieval step and does not account for the spatial correlation imposed by the sampling. This makes it difficult to develop a theoretically appropriate noise versus resolution trade-off using point-wise retrieval, without imposing restrictive assumptions.

Field-wise structure may be imposed in the retrieval step using a model-based approach. However, this approach is generally only practical if a low order wind field model is applied, which may discard more information than is desired. Furthermore, model-based approaches do not account for the structure imposed by the sampling.

A more general field-wise approach is to estimate all the wind vectors for all the WVCs simultaneously from all the measurements. This approach can account for structure of the wind field as well as structure imposed by the sampling in the wind retrieval step. Williams and Long in 2010 [1] developed a field-wise retrieval method based on maximum a posteriori estimation (MAP). This MAP approach can be extended to perform a noise versus resolution trade-off, and deal with ambiguity selection.

This paper extends the field-wise MAP estimation approach and investigates the noise versus resolution trade-off in the field-wise wind retrieval step. Some theoretical issues concerning field-wise ambiguity removal are also considered. The method is applied to the SeaWinds scatterometer and the results are analyzed.

2. BACKGROUND

This section presents background on scatterometry. The scatterometer sampling and noise models are also presented.

Wind scatterometers make multiple σ^0 measurements from different look directions of the same location on the Earth's surface. For the SeaWinds scatterometer, each scatterometer pulse is partitioned into several 'slices' with different spatial response functions that are narrow ($\sim$5 km) in the range direction and long ($\sim$25 km) in the azimuth direction. The slice σ^0 measurements from different looks (or 'flavors') of measurements sample the same location with response functions that overlap irregularly and with different orientations. Neglecting noise, the ith slice σ^0 measurement $\sigma^0_{t,i}$ can be expressed as an inner product of the underlying σ^0 field $\sigma^0_{t,i}(x)$ with the spatial response function $A_i(x)$, where the σ^0 field is a nonlinear function of the wind field $\vec{U}(x)$. In practice, the integration is made discrete. The multiple measurements can be stacked into a vector, producing the discrete scatterometer sampling operator [1]

$$\vec{\sigma}^0_t = \begin{bmatrix} \sum_x A_1(x)\mathrm{gmf}_1(\vec{U}(x)) \\ \vdots \\ \sum_x A_N(x)\mathrm{gmf}_N(\vec{U}(x)) \end{bmatrix} = \mathbf{T}(\vec{U}(x)). \quad (1)$$

where $\mathrm{gmf}_i(\cdot)$ represents the GMF that relates the wind to σ^0 with viewing geometry, polarization, and frequency corresponding to the ith slice measurement, and x represents a two-dimensional spatial variable.

Scatterometers make noisy σ^0 measurements due to several sources. A standard scatterometer noise model is that the measurements are Gaussian random variables, where the mean is the 'true' or noise-free measurement and the variance is a quadratic function of the mean [2]. This form accounts for fading as well as other noise sources. The distribution for noisy scatterometer measurements can be expressed as

$$f(\vec{\sigma}^0_m|\vec{U}(x)) = \frac{\exp\{\frac{-1}{2}(\vec{\sigma}^0_m - \vec{\sigma}^0_t)^T\mathbf{R}^{-1}(\vec{\sigma}^0_m - \vec{\sigma}^0_t)\}}{(2\pi)^{\frac{N}{2}}|\mathbf{R}|^{\frac{1}{2}}} \quad (2)$$

where $\vec{\sigma}^0_t$ is a function of the wind field as expressed in Eq. 1, and $\mathbf{R}$ is a diagonal covariance matrix with diagonal elements $R_{i,i} = \alpha_i(\sigma^0_{t,i})^2 + \beta_i\sigma^0_{t,i} + \gamma_i$, where α_i, β_i, and γ_i are functions of the scatterometer design.

The problem of wind scatterometry is to invert the noisy sampling model and estimate the wind field given the noisy scatterometer measurements. [1] develops a MAP estimation approach for wind field retrieval. The MAP estimator can be expressed as $\mathrm{argmax}_{\vec{U}(x)}\{\log f(\vec{\sigma}^0_m|\vec{U}(x)) + \log f(\vec{U}(x))\}$, where $f(\vec{U}(x))$ is a prior distribution of the wind field. The maxima can be found using a gradient search approach. The method developed in [1] employs a simplistic prior designed to regularize the problem in a practical manner in order to reconstruct a high resolution wind field. For this paper, a tunable prior is desired that appropriately handles the noise versus resolution trade-off.

3. MAP ESTIMATION FOR A NOISE VERSUS RESOLUTION TRADE-OFF

This section discuses some considerations involved in developing a prior distribution suitable for the noise versus resolution trade-off. The notions of observability and identifiability are discussed, which relate to the noise versus resolution trade-off and ambiguity selection respectively. Then a practical implementation is developed.

3.1. Observabilty and Identifiability

Note that including a prior generally improves both the identifiability (i.e., ameliorates 'field-wise' ambiguity selection) and the observability (i.e., reduces the variability of the estimate given a particular 'field-wise' ambiguity) of the MAP estimates over the ML estimates. Identifiability and ambiguity selection have to do with the number of local maxima of the expected estimator objective function, as well as the relative heights of the local maxima. Observability has to do with the expected widths of the local maxima. Including a prior distribution along with the ML objective function, as MAP estimation does, modifies the widths, relative heights, and locations of the local maxima–thus, modifying both the identifiability and observability of the problem.

For a practical implementation, MAP estimation is often performed with a gradient search method using an initialization field. In general, the presence of the prior reduces the variability of the estimate (i.e., effectively making the widths of the objective function around the local maxima more narrow). However, this particular approach does not handle ambiguity selection directly–since it converges to a particular local maximum near the initialization. In theory, all the field-wise local maxima can be found and the highest one may be chosen, but this is impractical due to the high number of parameters in the field-wise problem. Furthermore, multiple local maxima may have a similar height given a particular prior. This suggests that there may be different considerations for a prior that allows for unique identifiability (i.e., one dominant field-wise ambiguity), and for a prior that optimizes the noise versus resolution trade-off given a particular field-wise ML ambiguity.

For the purpose of this paper, a prior that optimizes the noise versus resolution trade-off (i.e., observability versus resolution trade-off) is considered, leaving the identifiability versus resolution considerations as they apply to ambiguity selection for future investigation. Nevertheless, some ambiguity selection issues are considered.

3.2. Implementation

In general, we desire a prior that imposes as little information as possible to obtain a desired variability of the estimates. Thus, we take a conservative approach employing a maximum entropy prior under certain constraints. We wish to impose structure on the power spectrum or correlation of the signal, which uniquely determines the covariance of the prior distribution. The maximum entropy distribution with a given covariance is the Gaussian distribution. Also, we desire the energy in the unobservable components of the signal to go to zero–suggesting a zero mean Gaussian prior distribution. Thus, we employ independent, zero-mean Gaussian priors on the U and V components of the wind fields with a Toeplitz covariance (which implies a wide-sense stationary process with a particular power spectrum). Note that for a diagonal covariance (i.e., white process), this results in a Rayleigh distributed wind speed prior and a uniform direction prior.

The power spectrum of wind components over the ocean tends to fall of approximately as one over the wave number squared [3]. We impose this structure on the wind field by assuming an exponential correlation function. Note that the exponential correlation function results in a flat spectrum up until the wave number corresponding to the correlation length, where it begins to fall off as one over the wave number squared. We use a correlation length of 200 km in order to impose the structure of the power spectrum on the finer scales, allowing the data to dominate the retrieval for the lower wave numbers.

We desire a tuning parameter that can be used to weight the influence of the prior. That is, we want to be able to adjust the amount of information (i.e., the entropy) that the prior imposes, which allows for a different trade-off for different applications. For example, we may desire a significant amount of smoothing for a climatological record in order to reduce the noise; whereas for a hurricane case study we may require a higher resolution–even if the estimates are noisy. Such a tuning parameter can be implemented by scaling the covariance of the prior, which directly adjusts the entropy of the Gaussian prior. The parameter p can be tuned either to produce a desired resolution (i.e., spectrum), or a particular noise level (i.e., RMS error with respect to some wind field). Based on preliminary analysis, we use a value of $p = 4$ because it produces retrieved speed and direction spectra that resemble the selected ambiguity spectra for low wave numbers, but continues to fall off gradually for high wave numbers.

The resulting prior can be expressed as

$$f(\bar{U}) = \frac{\exp\{\frac{-1}{2p}\bar{U}^T\mathbf{R}_p^{-1}\bar{U}\}}{(2\pi)^{\frac{N}{2}}(\frac{1}{p}|\mathbf{R}_p|)^{\frac{1}{2}}} \tag{3}$$

where p is the tuning parameter, $\bar{U}$ represents the sampled U and V components of the wind field $\vec{U}(x)$ stacked into a column vector of length N, and

$$\mathbf{R}_p = \begin{bmatrix} \mathbf{R_{UU}} & \mathbf{0} \\ \mathbf{0} & \mathbf{R_{VV}} \end{bmatrix} \tag{4}$$

is the covariance of the prior. Note $\mathbf{R_{UU}} = \mathbf{R_{VV}}$ are the covariance matrices of the U and V components whose rows are the exponential correlation function expressed as

$$r(x, x_0) = e^{\frac{-r_c}{r_l}||x-x_0||_2} \tag{5}$$

where x indexes the two-dimensional locations of the WVCs in the swath, $||\cdot||_2$ represents the L_2-norm, r_c is the WVC posting resolution, and r_l is the correlation length. For this paper, we retrieve the wind at a WVC posting of 12.5 km, and a correlation length of 200 km.

This MAP approach is implemented using a gradient search, which produces a reconstructed wind field estimate near the initialization wind field. The method can be considered either as a resolution enhancement procedure or a smoothing procedure on the initialization field, depending on the structure of the initialization field and the prior. There are many possibile approaches to come up with an initialization field, which can be considered as ambiguity selection. We could use a numerical weather prediction (NWP) field; however, the NWP direction fields tend to smooth over fronts and misplace storms. Alternatively, the result of a standard ambiguity selection routine may be used as the initialization. For example, the direction interval retrieval with tresholded nudging method (DIRTH) [4], and the two-dimensional variational analysis method (2D-Var) [5] both use information from a NWP and information from the scatterometer to produce a potentially improved nudge field. One more approach is to develop a new ambiguity removal method that relies less on the NWP by considering the field-wise identifiability versus resolution trade-off.

For this paper, we desire to compare the results of the noise versus resolution trade-off, uncoupled from the issue of ambiguity selection. In order to do this we initialize with the conventional 12.5km DIRTH result.

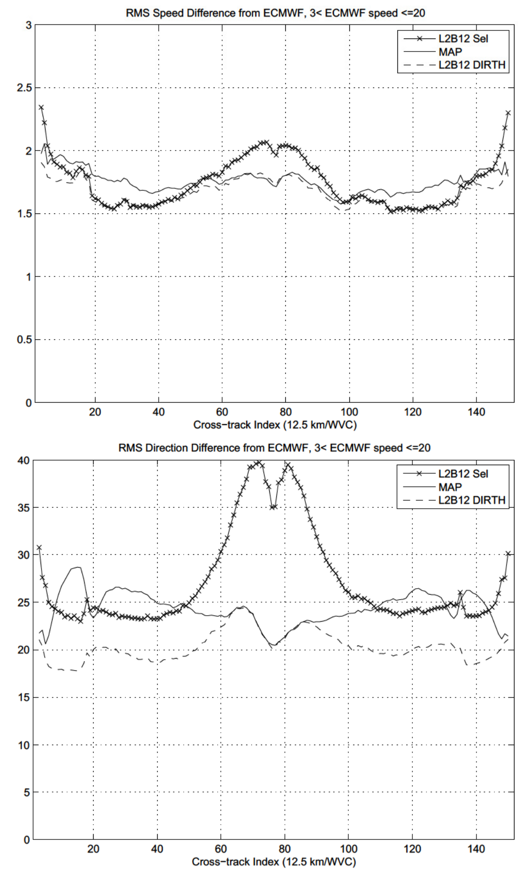

Fig. 1. Speed and direction RMS difference from ECMWF. The higher value of the MAP RMS differences relative to the DIRTH do not suggest that the MAP estimates are noisier, but may be due to recovering more fine scale information that is not represented by the ECMWF field.

4. ANALYSIS

Here, we take a statistical approach to analyse the results of the MAP method applied to SeaWinds data. We process several revolutions worth of SeaWinds data and compare the results to the standard 12.5km product DIRTH product (L2B12 DIRTH), the selected ambiguity (L2B12 Sel), and ECMWF. Spectral analysis is also applied.

Figure 1 shows the RMS speed and direction difference with respect to ECMWF of the various approaches. The MAP speed and direction RMS differences are relatively flat across the swath. Also, the MAP RMS speed and RMS direction differences in the nadir region are similar to DIRTH, but in the sweet spot and part of the swath edge they are higher than even the selected ambiguity. This suggests that a higher resolution may be recovered with MAP than is recovered from the point-wise retrieval.

Figure 2 displays the along-track one-sided power spectra of the speed and direction of the various winds. The direction spectra are taken as the magnitude squared of the Fourier transform of $\exp\{-id(x)\}$ where $d(x)$ represents the direction field expressed in radians. At low wave numbers the MAP speed and direction spectra are the highest, suggesting that more information is recovered for the high SNR signal components than is obtained from the point-wise retrieval. The MAP speed and direction spectra tend to follow the shape of the L2B12 Sel spectra until the L2B12 Sel spectra hit the noise floor and level out, while the MAP spectra continue to fall off. This suggests that the resolution of the retrieved speed and direction is at least as high as the L2B12 Sel product but is less noisy. In general, the MAP spectra are higher than the DIRTH (except for high wave numbers where the DIRTH spectra flatten due to noise), suggesting that the inherent resolution of the MAP estimates is higher than the resolution obtained by DIRTH. That is, DIRTH tends to attenuate both the high SNR components and the low SNR components, while MAP attenuates only the low SNR components. This suggests that the MAP approach more appropriately handles the noise versus resolution trade-off.

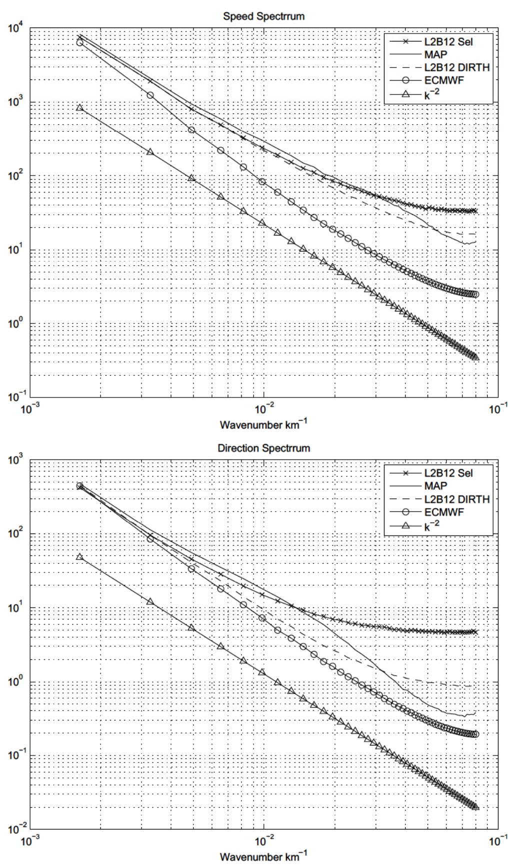

Fig. 2. Speed and direction power spectra. Note that the MAP approach tends to only attenuate the low SNR components (i.e., high wave numbers where the other methods flatten out due to noise). Also, MAP recovers slightly more energy at the lower wave numbers than the point-wise approaches.

5. CONCLUSION

This paper considers the noise versus resolution trade-off in wind scatterometry as part of the wind retrieval step. MAP estimation is employed with a prior that more appropriately handles the trade-off. The new approach reduces the estimation noise by driving the less observable components (i.e., those that are expected to be more noisy) towards the mean of the prior. This effectively smooths different regions of the swath differently–producing more constant error statistics as a function of cross-track index.

Although the MAP approach aids the noise versus resolution trade-off, there are many issues that still need to be considered. As suggested above the issue of ambiguity selection and field-wise identifiability are topics of future investigation. Another issue that may be possible to address with the MAP approach is direction favoring. That is, for certain measurement geometries, certain wind directions are significantly more observable than others, which results in retrievals (both ML and MAP) that favor certain directions. This direction favoring is especially evident when the wind signal is corrupted by rain. Methods of ameliorating direction favoring may be possible with the MAP approach and are to be developed in future work. Furthermore, since the rain spectrum differs from the wind spectrum, a field-wise simultaneous wind and rain estimation procedure may result in improved wind field estimates as well as provide useful rain field estimates.

6. REFERENCES

[1] B. A. Williams and D. G. Long, “A reconstruction approach to scatterometer wind vector field retrieval,” *IEEE Transactions on Geoscience and Remote Sensing*, 2010, in review.

[2] M. W. Spencer and D. G. Long, “Radar backscatter measurement accuracy for a spaceborne pencil-beam wind scatterometer with transmit modulation,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 35, no. 1, pp. 102–114, Jan 1997.

[3] M. H. Freilich and D. B. Chelton, “Wavenumber spectra of pacific winds measured by the Seasat scatterometer,” *Journal of Physical Oceanography*, vol. 16, no. 4, April 1986.

[4] 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.

[5] J. Vogelzang, A. S. A. Verhoef, J. de Vries, and H. Bonekamp, “Validation of two-dimensional variational ambiguity removal on SeaWinds scatterometer data,” *Journal of Atmoshperic Oceanaic Technology*, vol. 26, pp. 1229–1245, 2009.