

RECENT NESDIS RESEARCH IN WIND ESTIMATION
FROM GEOSTATIONARY SATELLITE IMAGES

C.M. Hayden* and R.T. Merrill**

*NOAA/NESDIS

**Cooperative Institute for Meteorological Satellite Studies
Madison, Wisconsin USA

1. INTRODUCTION

Geostationary satellite imagery has been used as a source of wind observations since soon after the launch of the first spin scan camera aboard the Application Technology Satellite (ATS 1) in December 1966. It was recognized immediately that features tracked in a time sequence of images could provide estimates of atmospheric motion, and from this simple beginning five operational satellites currently provide worldwide coverage, overlapping at the equator to slightly underlapping at 50° of latitude in both hemispheres. From the Meteosat (Europe), winds are produced operationally three times per day for 00, 06 and 12 UTC. From the two GOES (USA) and the Himawari (Japan), winds are produced four times daily at 00, 06, 12, and 18 UTC. Finally, from the INSAT (India) winds are distributed once daily for 06 UTC. For operational data cutoffs at approximately three and one-half hours after nominal observation time, the U.S. National Weather Service receives on average nearly 1500 reports. Thus, the geostationary satellites provide a substantial data set to the operational weather services, and these users are very interested in the quality and coverage.

Historically, wind vectors have been produced from images of visible (for low level vectors) and infrared 11 micrometer radiation (for upper some middle, and low level vectors). More recently, in order to improve the coverage at middle levels and also over cloud-free areas, wind tracking

has been applied to water vapor imagery (the 6.7 and 7.2 micrometer channels on the VISSR Atmospheric Sounder (VAS), Hayden and Stewart, 1987) although these vectors are not used in operational NWP.

The basic elements of wind vector production have not changed since their inception. These are: a) selecting a feature to track or a candidate target, b) tracking the target in a time sequence of images to obtain a relative motion, c) assigning a pressure height (altitude) to the vector, and d) assessing the quality of the vector. Initially, these elements were done manually (even to the point of registering movie loops), but the goal has always been to automate procedures and reduce the time-consuming and expensive human interaction. In this paper, we shall deal with each element, emphasizing the current stage of research at the Cooperative Institute for Meteorological Satellite Studies (CIMSS) which has been doing the major part of the current NESDIS development in wind tracking, with particular emphasis on automation.

A point we shall not deal with explicitly is the question of how well target displacement represents the motion of air. The argument is ongoing (Hubert and Whitney, 1971) even with respect to cloud motion; and certainly "water vapor" motion is at least as dubious. We believe that other uncertainties, particularly the pressure height assignment, contribute more significantly; and the important question is whether or not the vectors can still provide useful data, recognizing the existence of these uncertainties.

2. TARGET SELECTION

The first step in obtaining a motion vector is to discriminate a feature which is recognizable in each of the time sequenced images. This can be

quite simple, as in the case of high frequency (five minute) images of cumulus in the visible, or quite difficult, as in the case of low frequency (hour) images in the 7.2 micrometer band. (It is, in fact, impossible for some operators to use the latter, a fact which perhaps argues for automation. The computer is not handicapped by doubt emanating from insight, wisdom, good judgement, etc.).

In the automatic tracking currently used at CIMSS, the first (of the time sequence) image is systematically searched in contiguous segments measuring LSIZE x ESIZE (generally LSIZE=ESIZE) pixels with the object of locating the best (if any) target within each segment. The size of the segment is variable, and thus we can control the total number of targets. Typical dimensions we use are 30 x 30 for 4 km visible and 11 micrometer and 24 x 24 for 8 km water vapor imagery. A primary goal is to exclude poorly defined targets which will probably fail to produce winds, and each target segment is examined to ensure that several criteria are met. These are given in Table 1 for both the infrared window and the water vapor channels.

Table 1. Criteria to begin target search in segment (values in brightness, full range 0-255).

	11 micrometer	6.7 micrometer
minimum pixel value	90	90
maximum pixel value	240	195
max-min	25	8
min gradient (3 pix)	10	10

The reason for the lower maximum value with the water vapor is to exclude cloud, assuming that cloud vectors will be (better) obtained from visible and 11 micrometer tracking. The gradient mentioned in Table 1 is defined as:

$$|dTb/dx| + |dTb/dy| \quad (1)$$

where dx and dy are distance increments of 3 pixels in orthogonal directions. The pixel which represents the maximum gradient is selected as the potential target (following the work of Emery et. al., 1986). Finally, a second search is made, centered on the target, to ensure that the first two criteria of Table 1 are not violated in the segment which will be used in the actual tracking (Fig. 1). If a segment fails to produce a target, the stepping, in the x direction (Fig. 1), is reduced to half the nominal (ESIZE) value. After a target is selected, its segment is masked so that succeeding targets cannot be chosen closer than half the width of the segment .

The automated tracking scheme described above is obviously tunable, and the tuning should be done in concert with vector derivation and evaluation. At this point of our development, however, we have not systematically optimized the target selection in this way, but simply chosen parameters which yield a qualitatively acceptable density. An isolated example of parameter sensitivity is shown in Fig. 2. The figure shows (top) the targets selected with the default parameters for the case of 23 July 1987, 6.7 micron imagery. Note that obviously cloudy areas have been excluded, as have most areas of uniform shading. Dark (dry) areas where the water vapor motions are potentially most useful are well represented. There is a propensity, though not an overriding one, to pick the edges of features, as one would anticipate. The bottom of Fig. 2 shows the target selection (now shown over 11 micrometer imagery to better illustrate the clouds) when the cloud discriminating maximum has been reduced from 195 to 180. The targets are now very definitely

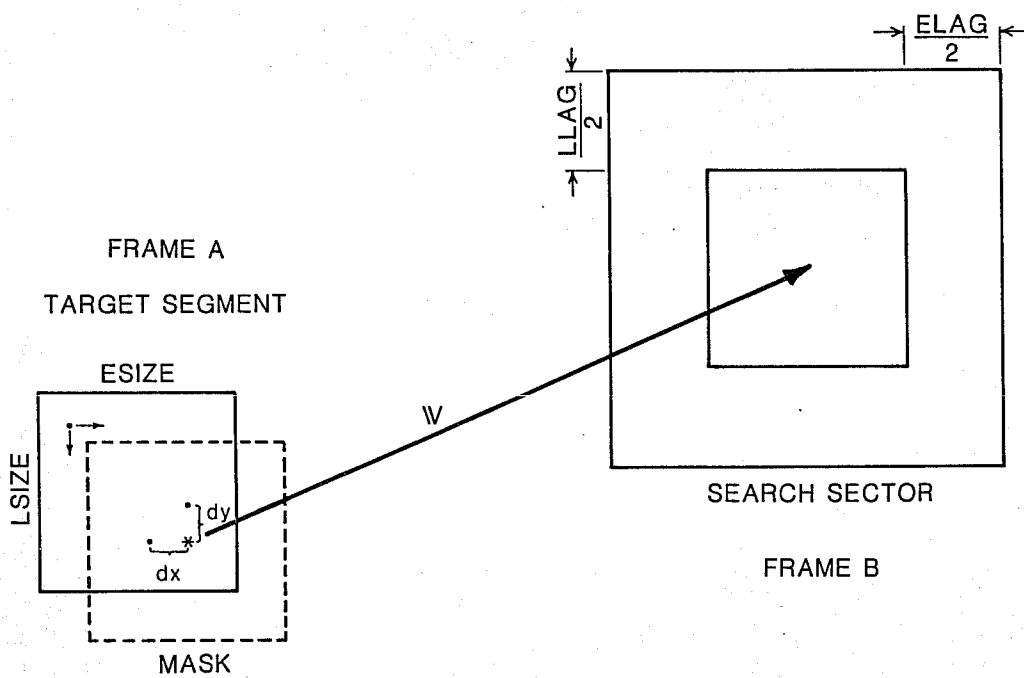


Fig. 1. Schematic of target selection and tracking procedure.




Fig. 2. 6.7 micrometer selected targets over 6.7 (top) and 11 (bottom) micrometer imagery. Bottom shows selection with maximum acceptable value reduced.

discriminated to cloud-free areas. Target yield has been reduced from 212 to 80.

3. TARGET TRACKING

Target tracking requires at least two images, and it is advisable to have more for purposes of quality control. Most automatic processing schemes use three images (an exception is the NESDIS picture-pair processing). Auto-tracking is performed by some type of correlation technique which allows pattern recognition of the target segment (defined above) within a larger search sector in a later (or earlier) image (Fig. 1). Emphasis is on simplicity, since complicated algorithms quickly become prohibitively expensive.

The method used currently at CIMSS is the Euclidean normal method which simply sums the squared difference between the pixels of the target and search frames as the target is moved through the search frame. A quadratic fit method is used to determine the "between pixel" location of the minimum sum (Smith and Phillips, 1972). Normally, the location of the search frame center is determined from a first guess displacement obtained from a forecast or analyzed wind field. The first guess fields are interpolated in space and time to the geographical location and presumed height (discussed below) of the target. This procedure is shown schematically in Fig. 1. In principle, the use of a first guess reduces the size of the search box necessary to obtain a pattern match. This is extremely important since the search for the offset of the best match requires summing the squares of $LSIZ \times ESIZ$ differences for each of $(LLAG+1) \times (ELAG+1)$ possible displacements, which is time consuming and expensive. Note, however, that the selection of $LLAG$ and $ELAG$ constrains the results about the guess displacement, because minima which fall on

the margins of the search domain are not accepted. The maximum possible deviation from the guess is given by:

$$(LLAG/2-1) \times (RES) / (LI) \tag{2}$$

where RES is the image resolution in meters/pixel and LI is the loop interval in seconds. For hourly interval, 8 km imagery with the default LLAG of 14, the maximum change to the guess is 13.3 msec^{-1} . Thus, the choice of LLAG invokes implicit quality control. Lags which we typically use to give approximately equivalent maximum deviations are shown in Table 2.

Table 2. Lag sizes (in pixels) used for indicated image frequency. Image resolution is 8 km.

lag	frequency (min)
8	30
14	60
24	120

For 4 kilometer infrared the lag of 8 is used for 15 minute imagery and so forth.


In the above description, there are no restrictions on the selection of the minimum. The location with the smallest minimum is chosen, whatever the magnitude. Recently, on the suggestion of John Eyre (personal communication), we have begun to experiment with combining the covariance (actually the Euclidean Norm and not a true covariance) field with an expected error in the first guess. Eyre's suggestion is based on an application of Bayes theorem which states that the conditional probability of a displacement d , given radiance information R is

proportional to the conditional probability of the radiance, given the displacement, multiplied by the probability of the displacement.

$$P(d:R) \propto P(R:d) \times P(d) \quad (3)$$

The rhs of (2) is comprised of the radiance covariance, scaled according to expected radiance noise to be equivalent to the probability of d which involves the first guess error. The latter is specified as a Gaussian distribution around the first guess displacement of the target, scaled by the expected error of the first guess.

Examples of the (LLAG+1 x ELAG+1) grids of the scaled covariances for the first guess and targets are shown in Figs. 3-6. Each figure shows the patterns for the first and second displacements, vectors 1 and 2. In every case, the guess covariance is centered for vector 1, but since vector 1 is used as the "guess" displacement for vector 2, the error distribution can be offset for the second comparison. (The amount of offset depends on the deviation of the first vector from the guess.) The gradient of the guess covariance depends on the expected error which is assigned. We have been experimenting with values from 10 to 20 percent of the magnitude and these figures represent such variations. The right side of the figures are the scaled radiance covariances for various choices of target. The minimum is a combination of the assigned radiance error and the persistence of the target, whereas the gradient is an indication of the distinctness of the target. Both minimum and gradient affect the degree of influence of the guess in the combination of the grids, but the gradient has the predominant effect. Figure 3 represents a well defined cloud in an 11 micrometer image. The figure shows a unique minimum and a strong gradient. Because of the latter, the first



GUESS: TARGET 2 VECTOR 1





IMAGE: TARGET 2 VECTOR 1



GUESS: TARGET 2 VECTOR 2





IMAGE: TARGET 2 VECTOR 2

Fig. 3. Left: guess covariance estimate. Right: radiance covariance estimate. Top: image 1 to 2. Bottom: Image 2 to 3. Well-defined 11 micrometer target.



GUESS: TARGET 8 VECTOR 1





IMAGE: TARGET 8 VECTOR 1



GUESS: TARGET 8 VECTOR 2





IMAGE: TARGET 8 VECTOR 2

Fig. 4. Same as Fig. 3, but less well-defined target.



GUESS: CLOUD VECTOR 1





IMAGE: CLOUD VECTOR 1



GUESS: CLOUD VECTOR 2





IMAGE: CLOUD VECTOR 2

Fig. 5. Like Fig. 4, except for 6.7 micrometer imagery.



GUESS: TARGET 2 VECTOR 1





IMAGE: TARGET 2 VECTOR 1



GUESS: TARGET 2 VECTOR 2




IMAGE: TARGET 2 VECTOR 2

Fig. 6. Like Fig. 5, except in cloud-free area. relative reliability of guess has been increased.

guess variance has little effect on the vector determination as indicated by the displacement of the guess covariance for target 2. Figure 4 represents a more diffuse cloud, again viewed at 11 micrometers. Figures 5 and 6 were derived from 6.7 micrometer imagery in areas with and without cloud contamination. Note that the gradients are considerably weaker than the figures for 11 micrometer targets, especially in the cloud-free case. The water vapor targets are not as distinct. Also, the covariance of the radiances is banded in the no-cloud case just as the targets are banded, and there is considerable uncertainty in the location of the minimum along the band. The radiance error for the 6.7 micrometer measurements has been chosen, in this case, to emphasize the first guess (or, equivalently, to emphasize the difficulty in tracking water vapor).

Figures 3-6 are quite typical of what we have seen in our limited study. We have been somewhat surprised by the absence of multiple minima, particularly with the 11 micrometer data. However, our target selection criteria were designed to produce this result.

4. ALTITUDE ASSIGNMENT

Assigning a pressure height to a target has been an ongoing problem. For visible data, used for low cloud tracking, a constant pressure is assumed (900 mb for the U.S. winds). For infrared upper atmospheric winds, all producers use the measured brightness temperatures of the target. The brightness temperature is converted into an atmospheric temperature which is matched to a temperature profile (from climatology or a forecast) in order to obtain the target pressure. Many ambiguities exist. For example, what is a representative temperature for a target which may consist of several layers of cloud? What should be assumed for emmissivity? For jet stream cirrus a value of unity is clearly

inappropriate. At ESA the multi-level problem is solved in part by "windowing" the target segment so that only the coldest pixels contribute to the final "cross-correlation" position (Schmetz and Nuret, 1987). The emissivity problem is addressed by a semi-transparency correction (Bowen and Saunders, 1984) which involves both the 11 and 6 micrometer channels and also radiative transfer calculations based on guess temperature and moisture profiles from the ECMWF.

At CIMSS, we do not use windowing for the pattern recognition, but instead use all pixels within the target segment. The brightness temperature representative of a cloud pressure is determined using a simple histogram technique. The brightness populations for the 256 categories are determined and lightly smoothed. The coldest (brightest) category showing at least $1/F$ times the total population is selected. The value of F is set to 20 for water vapor imagery and 60 for everything else.¹ If, however, the histogram fails to yield such a category, the minimum requirement is reduced by one and the search repeated, for as many as ten iterations.

Histogram examples for a 6.7 micrometer (top) and for an 11 micrometer height assignment (bottom) are shown in Fig. 7. For both of the scenes represented, wind vectors were achieved. The smooth distribution for the water vapor (away from cloud) is quite typical. Ambiguity in water vapor height assignment lies in the emissivity question and not in the

¹ On the McIDAS system, images of all infrared bands are calibrated such that the full range 0-255 covers the brightness temperature range 160-230K. The range was designed for 11 micrometer data which exhibits such variance. Other bands, however, among these the water vapor channels, show much less variance and only a portion of the 0-255 range is encountered. For this reason, various thresholds in target selection and cloud height determination are different for 11 micrometer versus 6.7 micrometer images.





Fig. 7. Brightness temperature histograms for a 6.7 micrometer (top) and 11 micrometer (bottom) target.

histogram. The bottom figure is typical of a mixed cloudy window scene. Note that the chosen temperature (234 K) could have been much warmer (255 K) with a with a more stringent choice of F, lowering the pressure from 262 to 399 mb.

There is no explicit correction for emissivity in the CIMSS cloud height determination. Instead, we rely on an empirically determined relationship between brightness temperature and atmospheric temperature. The model is a simple linear regression derived from collocated rawinsonde and satellite vectors. The dependent sample is derived by finding the level-of-best-fit between the satellite vector and the rawinsonde, where the discriminant is a combination of temperature, pressure, and vector deviation between the satellite and radiosonde measurements:

$$\frac{(T_r - T_s)^2}{\delta T^2} + \frac{(P_r - P_s)^2}{\delta P^2} + \frac{(V_r - V_s)^2}{\delta V^2} \quad (4)$$

Here the subscripts s and r refer to the satellite and radiosonde estimates, respectively. The denominator of each term in (4) weights the relative importance of the vector, temperature, and pressure differences, respectively. The ultimate regression relationship can be quite sensitive to these weights. For example, the top two panels of Fig. 8 show results of an experiment in assigning heights to 11 micrometer targets using 6.7 micrometer brightness temperatures. In a single days sample, 43 matches with rawinsondes were obtained, and the scatter of brightness temperature versus temperature is shown. The difference in the two panels is the choice of temperature weighting. The top panel represents denominators in (4) of 3 msec⁻¹, 10 K, and 100 mb. The middle panel represents 3 msec⁻¹, 5 K and 100 mb. It is obvious that the

tighter temperature tolerance has significantly reduced the scatter, as would be anticipated. The penalty is an increase in the vector error, but that remained at less than 5 m sec^{-1} for the sample, suggesting that the tighter tolerance is viable. This is just an example. We do not suggest that the 6.7 micrometer be used in most cases for assigning heights to 11 micrometer targets. The 11 micrometer measurement is generally better, as shown in the bottom panel of Fig. 8 for the same case, but a slightly different (because of the height assignment) sample. The scatter is considerably improved over the 6.7 micrometer results, and low (warm) clouds are properly included. We have, however, experimented with 6.7 micrometer height assignments for 11 micrometer vectors in recognition of the emissivity problems with thin cirrus. In this particular sample, the emissivity problem does not appear to exist. We anticipate that it is probably optimal to have both temperature assignment options available for temperature/height assignment, choosing according to some combination of temperature, velocity and pressure similar to (4), and are working in that direction.

The preceding points out that the height assignment is done after the target selection and before the tracking. This is necessary to obtain the first guess displacement. It is logical that the heights should be recomputed, after tracking, in the second and third images, to provide additional quality control. This is not currently done. Furthermore, the concept of first guess error described with regard to (3) should be broadened to incorporate the concept of height assignment uncertainty. In the Bayesian framework:

$$P(d:z) \propto P(z:d) \times P(z) \tag{5}$$






Fig. 8. Regression relationships between atmospheric temperature and brightness temperature derived from a one-day match of rawinsondes and satellites. Top: 6.7 micrometer match with loose temperature tolerance (10k); middle: tighter temperature tolerance (5k); bottom: 11 micrometer match.

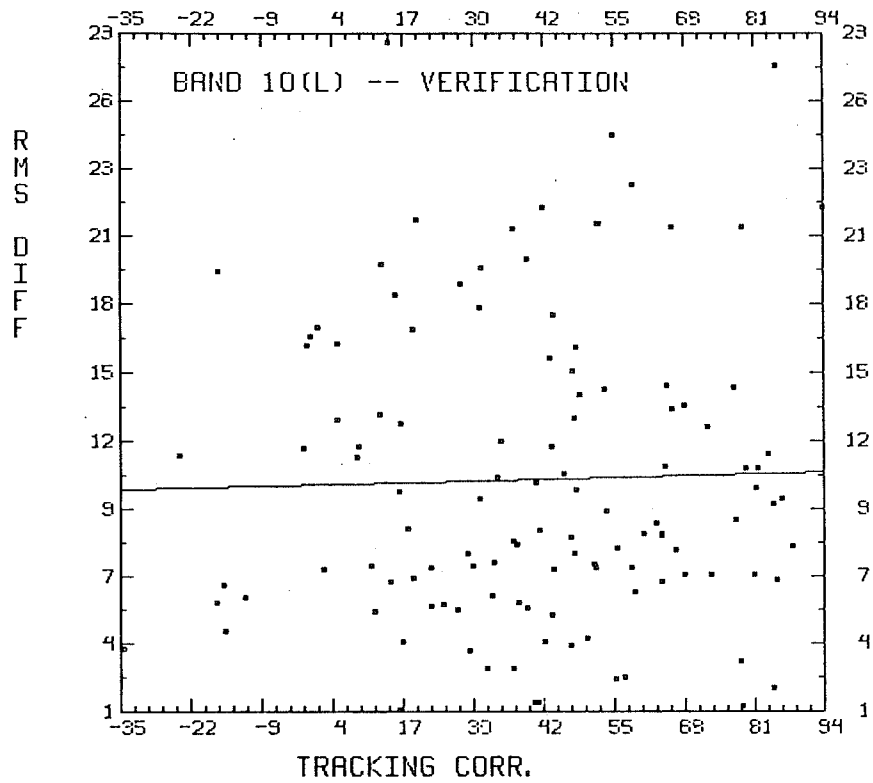


Fig. 9. Vector difference, rawinsonde - satellite vector vs. tracking correlation.

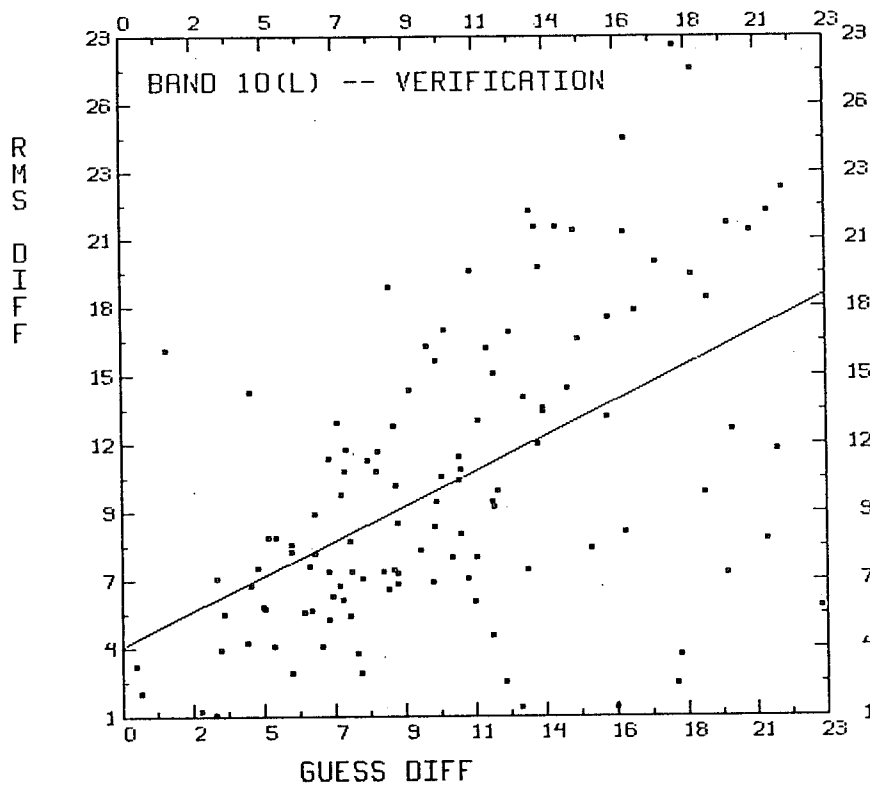


Fig. 10. Verification radiosonde - satellite vector vs. satellite vector change to first guess.

The probability of the height, given the displacement can be related to the vertical shear in the guess vector profile, whereas the probability of the height can be related to ambiguity in the target histogram, or statistics on the characteristic of independent height assignment such as discussed above.

5. QUALITY CONTROL

The problem of quality control is the least developed part of NESDIS satellite wind vector production. It is not automated, but relies on a meteorologist's decision as he views the data on an interactive display device. This is an accept/reject procedure. As part of the automated wind production, we are seeking methods to automate the quality control as well. Rather than deleting data which have successfully passed the correlation and height assignment checks, our goal is to provide one or more quality flags which the user may address in making his own quality decisions. The flags should be based primarily on features of the tracking procedure, and secondarily on comparisons with other data, in particular the first guess (since the user will presumably have access to this for his own quality control). To date, we have concentrated on several obvious parameters: a) the correlation magnitude, b) the target brightness variance, c) the vector reproducibility (first to second image versus second to third image), and d) deviation from the guess vector. Using collocated samples of rawinsondes and satellite vectors, we have screened the differences against each parameter in a limited number of cases. Contrary to intuition, none of the primary parameters (a-c) has proved useful. An example is shown in Fig. 9 for vector "error" versus correlation magnitude. There is no relationship. How can this be? How can a tracer with a wretched correlation produce a good vector? Conversely, how can a tracer with an excellent correlation produce a

vector with enormous error? We surmise that the former occurs in cases where a rapidly changing but nonetheless well defined cluster is tracked. In the latter case, one expects a gross error in height assignment. Clearly, many uncertainties must be removed before the primary tracking variables can be used as quality indicators. The last, secondary criterion (d) does show some skill for quality control as shown in Fig. 10. Rather discouragingly, however, the figure indicates that the more one deviates from the guess, the more likely the vector is in error. This is not always the case; there are points where large changes to the guess produced good vectors. The challenge is to keep these while eliminating the bad vectors.

6. SUMMARY

The foregoing has addressed the current research at CIMSS with respect to automated wind tracking from satellite imagery. The techniques and their attributes have been discussed; nothing has been said about data quality or impact on numerical forecasts. For many years, cloud motion winds (CMW) have been an accepted source of data and have been shown to positively affect the forecast, especially in the tropics (Kallberg et al., 1982). But attitudes may be changing. The quality of CMW, in terms of RMS comparisons with rawinsondes, has remained relatively constant. Forecast models and data assimilation techniques, however, have been steadily improving. It is now recognized that there are systematic biases in CMW (Kallberg and Delsol, 1985) which make some (e.g., vectors near the jet) of questionable utility, and some users have downgraded the reliability of CMW accordingly. In response to this, we are seeking to improve the quality of the satellite product, largely by adopting the data assimilation methods of the user. There is some danger in this. The Bayesian approach of using the first guess removes some of the

independence of the estimate. We may exclude extreme, but true, measurement indicative of important forecast errors. Will all users accept this loss in order to remove more prevalent bad data? Eventually, radiance histograms and correlation fields should be built directly into each data assimilation model. Failing this, it seems reasonable to build some of the data assimilation procedures into deriving data.

7. REFERENCES

Bowen, R.A. and R.W. Saunders, 1984: The semi-transparency correction as applied operationally to Meteosat infrared data: a remote-sensing problem. *ESA Journal*, 8, 135-131.

Emery, W.J., A.C. Thomas, M.J. Collins, W.R. Crawford, and D.L. Mackas, 1986: An objective method for computing advective surface velocities from sequential infrared satellite images. *J.Geo.Res.*, 91, C11, 12865-12878.

Hayden, C.M. and T.R. Stewart, 1987: An update on cloud and water vapor tracers for providing wind estimates, Preprint Vol., Sixth Symposium on Meteorological Observations and Instrumentation, Jan. 12-16, 1987, New Orleans, LA. *Amer.Meteor.Soc.*, 70-75.

Hubert, L.F. and L.F. Whitney, 1971: Wind estimation from geostationary satellite pictures. *Mon.Wea.Rev.*, 99, 665-672.

Kallberg, P., et al., 1982. The impact of cloud track wind data on global analyses and medium range weather forecasts. *ECMWF Tech.Rep.No.34*, 55 pp.

Kallberg, P. and F Delsol, 1985: Systematic biases in cloud-track-wind data room jet-stream regions. Internal ECMWF report.

Schmetz, J. and M. Nuret, 1987: Automatic tracking of high-level clouds in Meteosat IR images with a radiance windowing technique. *ESA Journal*, 11, 275-286.

Smith, E. and D. Phillips, 1972: WINDCO: An interactive system for obtaining accurate cloud motions from geostationary satellite spin scan camera pictures. Annual Scientific Report on NAS5-11542, 1971-1972. Space Science and Engineering Center, University of Wisconsin, pp 1-53.