Reprinted from Preprint Volume, 12th Conference on Probability and Statistics in the Atmospheric Sciences, June 22-26, 1992 Toronto, Ontario, Canada Published by the American Meteorological Society, Boston, Mass.

Read III

AN ADAPTIVE THRESHOLDING PROCEDURE TO DETERMINE CATEGORICAL FORECASTS FROM PROBABILITIES

David A. Unger

Techniques Development Laboratory
Office of Systems Development
National Weather Service
Silver Spring, Maryland

1. INTRODUCTION

Objective statistical guidance is frequently expressed in terms of probabilities. These probabilities may pertain to categories of a continuous or quasi-continuous event or to a dichotomous event such as precipitation/no precipitation. While a probabilistic forecast clearly contains more information than a non-probabilistic one, many times there is a requirement to express the prediction in the form of a categorical statement.

An adequate translation of a probability into a yes/no forecast, or the selection of a single category that best represents a multi-category probability statement, is by no means a simple matter. A classification procedure that is not tailored to the particular forecast element, such as forecasting an event when its probability exceeds 50 percent or choosing the category with the maximum probability, will frequently fail to achieve a suitable balance of forecast and observed events. In order to deliver a realistic categorical forecast, the selection method must take the frequency and predictability of the event into account.

The Techniques Development Laboratory (TDL) is upgrading its statistical guidance system. One project will place a system at local forecast offices to update central MOS guidance with information from hourly observations, locally run numerical models, and radar data. This system, the Local AWIPS MOS Program (LAMP), is a more recent version of an earlier one designed for AFOS (Glahn and Unger, 1986) (Unger et al., 1989). It provides guidance for nearly all MOS forecast elements for hourly projections from 1 through 20 hours.

The large amount of guidance to be available from this project requires an efficient and accurate means of category selection. This paper summarizes a recursive estimation procedure to obtain threshold values that translate probability statements into categorical forecasts with a specified bias (ratio of forecasts to observations). The technique can be used on historical forecasts to obtain the threshold values. It can also be used in real time to produce threshold values that "learn" from the current forecast performance in order to move toward values that will produce desired biases.

2. A SINGLE EVENT PROBABILITY

2.1 Terminology

A single event probability forecast will be used to introduce the procedure. Let r_i denote the probability forecast on case i. The verifying observation is represented by d_i, which assumes the value 1 when the event occurs, and 0 otherwise. A threshold value, t, is required to translate r_i into a categorical forecast, f_i, defined similarly to the observation. A forecast for an event is made when $r_i \geq t$. An historical series of N forecasts ($i = 1$ to N), referred to here as dependent data, is used to obtain threshold values.

2.2 Exact Threshold Value

For any desired bias, b, a threshold can be found such that the number of cases in which $r_i \geq t$ is equal to the product of the desired bias, b, and the number of observed events in the dependent data sample. This threshold exactly fits the dependent data and is hereby referred to as the exact threshold value. An exact threshold can efficiently be found by a procedure which orders the dependent data forecasts to obtain a cumulative forecast probability distribution from which the appropriate threshold value can be located.

2.3 The Adaptive Threshold Estimation Procedure

Autoregressive models and recursive estimation procedures have been used for some time to provide on-line adaptive estimation of the statistical properties of mathematical and physical systems. Many of the adaptive filters described in the literature are inappropriate for the classification procedures described here because of the binary nature of the categorical forecasts and the high degree of serial correlation of the forecasts and observations. A procedure specifically designed to process the information obtained from the forecast performance was devised to provide an adaptive filter for threshold estimation. Simplicity was considered of paramount importance in the design of the technique, since even a modest amount of additional computational effort can seriously encumber a large statistical forecasting system.

The Adaptive Threshold Estimation Procedure (ATEP) uses the forecast made from the current threshold value, together with its verifying observation, to adjust the threshold for the next forecast. For a single event probability where the categorical forecast is determined by

$$f_i = 1 \text{ when } r_i \geq t_i,$$

the ATEP adjustment equations are:

$$t_{i+1} = t_i + \Delta, \quad \text{when } f_i = 1,$$

$$t_{i+1} = t_i - b\Delta, \quad \text{when } d_i = 1.$$

Here Δ controls the magnitude of change made to the threshold at the end of each timestep. Since Δ governs how much influence the recent data has on the threshold value, it can be considered as the ATEP filter gain.

Fig. 1 illustrates in contingency table form the change made in the threshold as a function of the forecast, the observation, and Δ. Note that for a unit bias, no adjustment is made for a correct forecast. For an incorrect forecast, the threshold will always move toward the value required to produce a correct forecast on that case.

OBSERVED \ FORECAST	f=1	f=0
d=1	$\Delta(1-b)$	$-b\Delta$
d=0	Δ	0

Figure 1. Contingency table for threshold changes made by ATEP for a bias, b and a gain, Δ.

The stability of the filter is examined by computing the expected value of the threshold change. Where $P(f=1)$ and $P(d=1)$ indicate the probability of the forecast and observed event, respectively, the expected value is,

$$E[t_{i+1} - t_i] = \Delta(P(f=1) - bP(d=1)).$$

$E[t_{i+1} - t_i] = 0$ when $P(f=1) = bP(d=1)$, or, in other words, when the categorical forecast bias is b. When the threshold is too low, the bias will exceed b, and the filter moves the threshold toward a higher value. In a similar manner, the filter pushes a threshold that is too high (low bias) toward a lower value. The result is a time series that converges to a value that will produce the requested bias, provided that Δ is small. The convergence properties of the filter, and the selection of Δ will be discussed in Section 4.

3. MULTI-CATEGORY PROBABILITIES

3.1 Terminology

A multiple category probability forecast refers to a series of probability statements that pertain to the possible outcome of a single event. Let $r_{i,j}$ denote the forecast probability that the observation on the i^{th} case will fall into the j^{th} category. The verifying observation, D_i, and categorical forecast, F_i, assume the value of the category which is predicted or observed. (The upper case is used to distinguish this notation from the binary notation used earlier.) F_i will be referred to as a "best category" forecast, following the convention in many statistical weather forecasting systems. Note that the term "best" only signifies that the category is selected by the rules used for the classification and not that it is in some way an optimum choice.

Classification of a multi-category probability statement is difficult because the categories are interrelated in a complex way. Factors such as the correlation between forecasts for different categories, variations in climatological frequency, differences in predictability, and differing degrees of importance of each category must be considered to arrive at an effective categorical forecast.

The probability forecasts are processed by rules, referred to here as the classification strategy, that specify exactly how the best category is to be selected. There are many different classification strategies possible for the multi-category probability forecast. This discussion focuses on strategies that use threshold values to determine the best category.

Strategies that use only the probabilities for the specific category to which they apply (discrete probabilities) are referred to as discrete methods. For ordinal categories (categories for which the order is meaningful), cumulative probabilities

$$R_{i,j} = \sum_{k=1}^{j} r_{i,k}$$

can be used. Classification strategies are labeled cumulative or discrete depending on the probabilities that they use.

The ATEP filter operation will be discussed for several classification strategies. It is not the intention of this paper to examine the relative merits of these strategies. Rather, the discussion focuses on the filter equations required for each method.

3.2 Ordered Selection Strategy--Exact Threshold Values

A common thresholding strategy used for a multiple category forecast is an ordered selection approach. Cumulative or discrete probabilities are tested against threshold values. Categories are examined in a specific order, with the first category that exceeds its threshold selected as the best category for that forecast. The categories are usually ordered according to the operational significance of the event in question, with the most important assigned the lowest value and, therefore, is tested first. For a forecast with NCAT categories, only NCAT-1 thresholds need to be checked since F = NCAT is assigned when no thresholds are exceeded.

Thresholds that produce desired biases for the ordered selection strategies can be found through an iterative technique. The category 1 threshold is examined first, and a threshold found in the same manner as the single event probability. Once this threshold is determined, all cases with $r_{i,1} \geq t_i$ are eliminated from the sample. The remaining cases are used to find a threshold value for the next category, from either the discrete or cumulative probabilities, depending on the selection strategy desired. The cases that exceed this threshold are then eliminated from the sample, and the next higher category is examined. This successive searching and elimination continues for NCAT-1 categories.

3.3 Ordered Selection Strategy--ATEP Thresholds

The ATEP system uses the following adjustment technique for discrete ordered selection strategy.

$$t_{i+1,K} = t_{i,K} + \Delta, \text{ for } F_i = K$$

$$t_{i+1,L} = t_{i,L} - b_L\Delta, \text{ for } D_i = L$$

where L is the observed category and K is the best category forecast for the i^{th} case. The redundant category, j = NCAT, is not used for this system. Thresholds for categories not specified by adjustment equations remain unchanged.

For the cumulative ordered selection strategy, the ATEP adjustment equations are:

$$t_{i+1,j} = t_{i,j} + \Delta, \text{ when } K \le j, \; j=1,\text{NCAT-1},$$

$$t_{i+1,j} = t_{i,j} - B_j\Delta, \text{ when } L \le j, \; j=1,\text{NCAT-1}.$$

The B_j signifies cumulative bias,

$$(P(F_i \le j)/P(D_i \le j)).$$

The mathematical stability of the multi-category systems can be examined by computing the expected value of the threshold change for each category. For the discrete strategy,

$$E[t_{i+1,j} - t_{i,j}] = \Delta(P(F=j) - b_jP(D=j)),$$

$$j=1 \text{ to NCAT-1},$$

and for the cumulative strategy,

$$E[t_{i+1,j} - t_{i,j}] = \Delta(P(F\le j) - B_jP(D\le j)),$$

$$j=1 \text{ to NCAT-1},$$

For small Δ, thresholds will converge to values which produce the requested forecast biases. The bias for j = NCAT is implicitly controlled by the other categories.

3.4 The Discrete Ratio Thresholding Strategy

A problem with the ordered selection strategies is that all categories are not considered simultaneously. This can lead to situations where categories with substantial probability are neglected because a lower category may have marginally surpassed its threshold. In an attempt to improve the efficiency of category selection, a classification strategy that selects the category with the highest ratio of forecast to threshold value has been tested for some elements.

Exact threshold values for this classification strategy are very difficult to find by non-recursive methods since it is not possible to approach the situation one category at a time. ATEP thresholds, on the other hand, are easily obtained by using the same basic strategy as for the discrete case described above, with one additional consideration. Since the selection method depends only on the forecast to threshold ratio, the thresholds are only important relative to one another. An infinite number of threshold values can produce categorical forecasts with the desired biases since any multiple of an exact threshold set will produce the same categorical forecasts.

To produce thresholds that will converge to a specific set of values, one category must be fixed. This "anchor" category also provides a category who's bias is not explicitly controlled by the filter. (Only NCAT-1 biases can actually be specified, since one category is redundant).

While any value can be used for the anchor category, a logical choice is its unit-bias, single event threshold. This will help produce thresholds that appear reasonable; that is, each will generally be close to the value that the probability will have to exceed to select the related category. There is no feedback between the actual categorical forecast and the anchor category's threshold. The anchor category is isolated from the other categories in threshold estimation, although its threshold value is computed recursively along with the others.

With an anchor category, j=A, and a classification strategy of

$$F_i = K, \text{ for } r_{i,K}/t_{i,K} = \max(r_{i,j}/t_{i,j}),$$

$$j=1 \text{ to NCAT}.$$

ATEP filter equations are:

$$t_{i+1,K} = t_{i,K} + \Delta, \text{ for } F_i = K, \; K \ne A$$

$$t_{i+1,L} = t_{i,L} - b_L\Delta, \text{ for } D_i = L, \; L \ne A$$

$$t_{i+1,A} = t_{i,A} + \Delta, \text{ for } r_{i,A} \ge t_{i,A}$$

$$t_{i+1,A} = t_{i,A} - \Delta, \text{ for } D_i = A.$$

4. THRESHOLD CONVERGENCE

The time series behavior of the thresholds depends in a complex way on the distribution of forecast values, the skill of the prediction, the serial correlation of both the forecasts and observations, and the value of Δ. Day to day variations in weather might be considered to be a source of noise since they cause the thresholds to oscillate about their exact values. The amount of noise in the threshold time series depends to a large extent on the filter gain. Noise must be limited by the selection of a small Δ before thresholds can converge.

A weather element for which the serial correlation of the threshold changes is high requires a smaller Δ than does one for which the changes are more random. Experience with several weather elements indicates that thresholds take an excessive amount of time to converge when Δ is set to a value that reduces the noise to acceptable levels. Threshold convergence is so slow for some elements that acceleration techniques are required.

Two methods have been employed to accelerate convergence. The first is the practice of varying Δ according to the confidence in the threshold. For thresholds which may require considerable adjustment, the filter gain needs to be relatively high, placing a considerable weight on recent data.

As the thresholds become more accurate, Δ is decreased.

The second convergence acceleration technique is to maintain a second series of threshold values smoothed with a low pass filter to remove high frequency variations and seasonal cycles. The periodic replacement of the unsmoothed thresholds with their smoothed counterparts, usually in conjunction with a reduction in the filter gain, can greatly accelerate convergence. This steers the unsmoothed values toward the middle of their ranges, where the exact values are likely to be.

The very simple first order low pass filter used here is described briefly by Raymond and Garder (1991, Eq. 26)

$$ts_{i+1,j} = \alpha ts_{i,j} + (1 - \alpha)t_{i,j},$$

where $ts_{i,j}$ is the smoothed threshold value for time i and category j. The parameter, α, must be within the range, $0 < \alpha < 1$ and determines the frequency response of the filter. A value near 1 will greatly attenuate higher frequency variations. This filter, chosen for its simplicity, does not have a very sharp response. The response, however, is not particularly important for ATEP operations. Other low pass filters designed for time series application can be used if more precise filtering is required.

Smoothed thresholds are never used in the recursive filter itself but only to provide final estimates of threshold values or to provide stable values from which to make forecasts in real time. While the smoothed thresholds are usually quite near to the exact values, they are not fully recursive and will not precisely reflect the observations.

A general strategy to produce reliable thresholds is to make multiple passes through a dependent data set as required to allow thresholds to converge. Each pass is treated as additional data, so, for example, 10 passes through a single season is equivalent to 10 seasons of data.

From an arbitrary initial guess, a very coarse adjustment (high gain, light smoothing) is made for the first few seasons. A stronger smoother is then used with a high gain to obtain average threshold values for a long series of data. (Usually the equivalent of about 5 to 10 seasons.) At that point, the smoothed threshold values usually will produce forecasts with nearly the requested bias. Each unsmoothed threshold is then assigned the value of its corresponding smoothed threshold. The ATEP filter is used with fine adjustment control constants for approximately another 10 seasons to further refine the thresholds. In practice, the precise values of the control parameters and the length of the time series must be selected on an element by element basis.

5. REGIONAL ADJUSTMENT

Many times a single set of threshold values is applied to forecasts for many locations within a regional area. Regional threshold values are found on historical data by combining all forecasts within a region to form a single series of data. The threshold set, adjusted by the forecast and observation of one station, is used to make the best category forecast for the next station. The filter then becomes recursive in both the time and space domains. Smoothing is applied after all regional adjustments have been made for a particular valid time, so that the low pass filter operates only in the time domain.

6. REAL TIME OPERATION.

One potential use for the ATEP is to continuously adjust the thresholds in real time. Real-time adjustments can be used to allow thresholds to adapt to changes in forecast environment. Continuous adjustment is not without hazard, however. Thresholds can respond to unusual forecast performance during a persistent weather regime, for example, to produce misleading results when that regime ends. Problems may also result from phase errors in the threshold time series. Since the threshold values will always lag the input signal, periodic changes in forecast performance may produce perpetually misadjusted threshold values. For a seasonal cycle, for example, threshold values will most closely approach stable values near the end of one season, leaving thresholds misadjusted for the beginning of the next. The filter control constants must be carefully chosen to minimize such problems.

In real time, best category forecasts are made from the smoothed threshold values. The probability forecasts are stored until the verifying observations arrive. Intermediate categorical forecasts are then made from the unsmoothed thresholds specifically for the ATEP adjustment. New smoothed thresholds are calculated from the updated unsmoothed values before the next forecasts are made.

For regional thresholds, all forecasts within the region are made from the latest smoothed thresholds. The ATEP adjustments are then made from the probability forecasts as described in Section 5.

7. AN EXAMPLE

7.1 LAMP Visibility Forecasts

The LAMP visibility forecasts will be used to illustrate the derivation of thresholds from ATEP. LAMP forecasts are produced from regression equations that specify the probability of visibility occurrence in each of six categories at the hourly observation time. The forecasts are produced from regional LAMP equations developed for the Southern Plains area of the U.S.

While actual LAMP thresholds are derived for the same regions for which the equations apply, a single station threshold set was used here to illustrate the ATEP performance on a manageably small amount of data. The 2000 UTC LAMP cool season (October-March) visibility forecasts for Oklahoma City, Oklahoma (OKC) were selected for threshold derivation. These forecasts update the 1200 UTC MOS guidance with information from the 2000 UTC observations. The 16-h projection was examined for this demonstration so that the valid time, 1200 UTC, would occur in early morning at OKC when visibility restrictions are most frequent. The forecasts were obtained from the developmental sample used to derive the LAMP equations. A total of 1576 forecasts were available from the 10 seasons of data beginning in October 1980.

The category definitions, means and multiple correlation coefficients for the regional equations are shown in Table 1 along the sample mean frequencies of visibility at OKC.

Table 1. Category definitions, means, and correlation coefficients, C, for LAMP visibility equations for the Southern Plains region and the sample means for the OKC data.

Cat.	Definition (mi)	Equation Mean	C	OKC Mean	Cases
1	less than .5	.022	.200	.030	47
2	.5 - .99	.015	.155	.013	19
3	1.0 - 2.99	.062	.325	.034	54
4	3.0 - 4.99	.062	.248	.043	68
5	5.0-6.0	.071	.204	.056	89
6	greater than 6.	.769	.521	.824	1299

7.2 Threshold Convergence

The ATEP filter was used to estimate unit bias thresholds for all categories. The ATEP control constants are displayed in Table 2. The threshold values were initially set to .02 for all categories. The constants used for the smoother were designed to heavily damp frequencies with periods of less than about 1 season ($\alpha = .9944$) and less than 5 seasons ($\alpha = .9989$). The number of times the ATEP filter passed through 10-season sample is shown along with the value of the smoothed thresholds at the end of each pass. The exact threshold values found by iterative searching and elimination are also shown.

Table 2. Thresholds for the 5 visibility categories at the end of each stage in ATEP adjustments. The number of passes, and the ATEP control constants for each adjustment period are shown.

Adjustment	Passes	Δ	α	Smoothed Threshold 1	2	3	4	5
Coarse	1	.03	.9944	.1151	.1603	.3110	.3875	.4481
	1	.02	.9989	.1156	.1515	.2932	.3652	.4312
Fine	2	.005	.9989	.1069	.1444	.2743	.3575	.4185
Very Fine	5	.001	.9989	.1027	.1397	.2677	.3507	.4112
	20	.0001	0	.0998	.1361	.2682	.3450	.4075
Exact	-	-	-	.0987	.1356	.2687	.3435	.4080

From Table 2, it is evident that the thresholds quickly approached the exact values, but only slowly converged once the system was in approximate balance. To assure convergence, Δ was reduced to a very small value and the smoother was effectively removed by setting $\alpha = 0$. (Smoothing is unnecessary since there is very little noise in the unsmoothed thresholds when the gain is set so low.)

The time series of the threshold values for categories 2 and 3 for the first four passes are shown in Fig. 2, along with the smoothed and exact thresholds values. On the first pass, the noisy, coarse-adjustment thresholds quickly reach the vicinity of their exact values and oscillate around them while the smoothed thresholds approach the exact values more slowly. It is very important that the smoothed thresholds reach the middle levels of the unsmoothed threshold's range before Δ is reduced. Thresholds take a very long time to

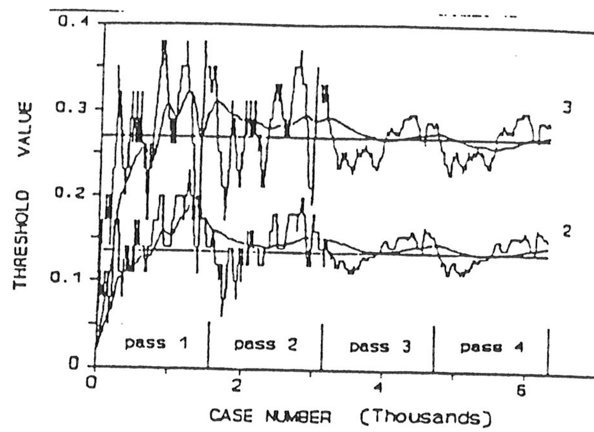

Figure 2. Time series of ATEP thresholds for visibility categories 2 and 3, along with their smoothed values. Exact threshold values are indicated by the horizontal lines.

converge unless they are very close to the stable values once Δ is very small. This is especially true for rare events, since adjustments are made only for forecast events or observations. (Since the event is rare, most forecasts are correct.)

At the end of the second pass, the fine-adjustment ATEP control constants were instated. With the noise greatly reduced, long term trends in the data became apparent in the thresholds. There was a distinct trend for both thresholds to increase with time on any given pass over the 10, 6-month seasons, indicating a systematic overforecasting. The patterns on the fourth pass were similar to those on the third, suggesting that the thresholds had reached stable levels by that time.

The thresholds at the end of the fourth pass would most likely produce satisfactory results on independent data; however, more passes were made to demonstrate convergence. An additional 25 passes--equivalent of 250 years--were made through the data. The long time series was necessary because, as the thresholds approached their exact values, the bias imbalances amounted to only a few cases in the entire 10-year sample. After 25 passes, the thresholds were all within .0015 of the exact values.

7.3 Simulated real-time performance.

In practice, the LAMP guidance will make forecasts from fixed thresholds because we can use the large amount of data available for equation development to produce very accurate threshold values. Categorical forecasts made from continuously adjusted thresholds, however, were examined to ascertain how a system that adjusts forecasts "on-line" might perform.

Table 3 displays the Heidke skill scores and bias by category for forecasts made from the threshold values from the last pass of the coarse, fine, and very fine adjustment periods. The Heidke skill score measures the improvement in the forecasts over chance and is computed by the formula, $HS = (H-E)/(T-E)$, where H is the number of correct forecasts, T is the total number of forecasts, and E refers to the number correct expected by chance given the sample climatology and number of predicted events for each category. Categorical forecasts for the entire data sample made from threshold values fixed throughout the sample were compared to those which were allowed to vary with fine adjustment control parameters.

The skill scores for all systems tested are approximately the same, which is not surprising since the probabilities used to make the categorical forecasts are identical.

Table 3. Heidke skill scores and biases by category for visibility forecasts made from final threshold values of the specified adjustment steps (see Table 2). Scores were computed from both variable and fixed thresholds.

Adjustment	Threshold	Skill	Bias by Category 1	2	3	4	5	6
Coarse	Variable	.249	.83	1.10	.98	.91	1.06	1.01
	Fixed	.258	.66	.63	.94	.77	1.06	1.03
Fine	Variable	.256	.98	1.05	1.09	.90	1.03	1.00
	Fixed	.249	.87	.57	1.01	.85	1.09	1.01
Very Fine	Variable	.247	1.04	1.32	1.03	1.03	.98	.99
	Fixed	.250	.98	1.00	1.00	.97	1.02	1.00
Exact	Fixed	.251	1.02	1.00	1.00	1.00	.99	1.00

The small variations in skill score are often due to the difference of a few cases in the 10-year sample. However, there is a distinct and systematic improvement in the bias of the variable system with respect to the fixed, unless, of course, the fixed thresholds are very near to the exact values.

8. SUMMARY

An adaptive estimation procedure has been developed to obtain threshold values that will produce categorical forecasts with any desired bias. ATEP uses a recursive time series approach to adjust thresholds on a forecast by forecast basis according to the verifying observations. The filter uses three control constants: the desired bias of the categorical forecasts; a filter gain that controls influence of the most recent observation on the thresholds; and a smoothing constant that controls the attenuation of high frequency changes in threshold values.

The convergence properties of the ATEP filter were demonstrated by passing through a sample of visibility probability forecasts multiple times. Thresholds were found to approach the exact values determined from the dependent data sample. Careful choice of the filter control constants can significantly accelerate convergence to the vicinity of these values.

Forecasts made from thresholds that were continuously adjusted were compared to those produced from fixed values. The results indicated that real time adjustments with ATEP are feasible. The skill of the forecasts was not greatly affected by the continuous ATEP adjustments, while the bias characteristics were almost always improved unless the initial thresholds exactly matched the sample.

Adaptive threshold estimation offers several advantages to the alternative methods of finding exact threshold values. The primary advantage is that it enables bias control of more complex classification strategies, such as the ratio method presented here, or the method presented by Carroll (1992). ATEP thresholds can also be tuned to weight more recent data more heavily, thus retaining information from trends in the forecast performance. The ATEP thresholds can also be adjusted in real-time to enable continuous calibration of the best category forecasts.

9. ACKNOWLEDGMENTS

The author is indebted to Luther A. Carroll for many insights into the system, particularly regarding the system stability and the need for the anchor category. Kim Balwin and Cheryl Grove assisted with programming, graphics presentations, and analysis of the smoother characteristics.

10. REFERENCES

Carroll, L. A., 1992: An enhanced precipitation type short range guidance system for short range forecasts. Preprints 12th Conference on Probability and Statistics in the Atmospheric Sciences, Toronto, Amer. Meteor. Soc., (this volume).

Glahn, H. R., and D. A. Unger, 1986: A Local AFOS MOS Program (LAMP) and its application to wind prediction. Mon. Wea. Rev. 114, 1313-1329.

Raymond, W. H., and A. Garder, 1991: A review of recursive and implicit filters. Mon. Wea. Rev. 119, 477-495.

Unger, D. A., W. L. Wolf, T. L. Chambers, and M. W. Cammarata, 1989: The Local AFOS MOS Program: Current status and plans. Preprints 11th Conference on Probability and Statistics in Atmospheric Sciences, Monterey, Amer. Meteor. Soc., 114-119.