

An Efficient Methodology for Detection of Anomalous Propagation Echoes in Radar Reflectivity Data Using Neural Networks

MIRCEA GRECU AND WITOLD F. KRAJEWSKI

Iowa Institute of Hydraulic Research, University of Iowa, Iowa City, Iowa

(Manuscript received 8 September 1998, in final form 7 May 1999)

ABSTRACT

To detect anomalous propagation echoes in radar data, an automated procedure based on a neural network classification scheme has been developed. Earlier results had indicated that algorithms used to detect anomalous propagation must be calibrated before they can be applied to new sites. Developing a calibration dataset is typically laborious as it involves a human expert. To eliminate this problem, an efficient methodology of calibrating and validating neural network-based detection is proposed. Using volume scan radar reflectivity data from two WSR-88D locations, the authors demonstrate that the procedure can be calibrated easily and applied successfully to different sites.

1. Introduction

a. Motivation and previous work

Anomalous propagation (AP) is a source of significant errors in radar data. For example, Joss and Lee (1995) estimated an error induced by AP to be 13% in a 60-month rainfall accumulation. Moszkowicz et al. (1994), who analyzed a two-month radar dataset, found AP echo to be about 15%. We recently (Grecu and Krajewski 1999) applied an AP correction procedure and obtained a 15% improvement in a skill score that measured the agreement between radar and rain gauge rainfall estimates for four years of RADAP II data (McDonald and Saffle 1989) from Oklahoma City, Oklahoma, and Wichita, Kansas. Our result was consistent with the findings of Conner and Petty (1998).

Anomalous propagation occurs in certain conditions of nonstandard refraction in the atmosphere when the radar beam is deviated toward the ground and the resultant echo represents reflection of the ground and not a meteorological target. Detailed descriptions of the phenomenon and the conditions associated with it are available (see, e.g., Battan 1973 or Doviak and Zrnich 1993). Several approaches are possible to identify and remove AP echo. They can be hardware- or software-based and can be performed in real time (on-site) or offline (off-site). The on-site approaches seem more efficient because often more information (power spectra,

Doppler spectra, etc.) is available. Techniques appropriate for on-site AP detection can be found in Aoyagi (1983), Collier and James (1986), and Joss and Lee (1995). Unfortunately, there is a large amount of radar data on which such techniques either were not used or were used unsuccessfully. For example, the Ground Validation Program of the Tropical Rainfall Measuring Mission (Simpson et al. 1988) uses reflectivity data from two WSR-88D (Klazura and Imy 1993) sites: 1) Melbourne, Florida and 2) Houston, Texas. It was found that despite using AP filters that are based on Doppler information, the archived data for these sites are significantly contaminated by AP echoes. Thus, before the data can be used for rainfall estimation, postprocessing procedures for AP detection need to be applied. This is the case for reflectivity data from many other WSR-88D installations. Smith et al. (1996) analyzed four different rule-based, off-site AP detection algorithms using reflectivity data from around the country and concluded that none of them can perform satisfactorily in all situations.

Moszkowicz et al. (1994) developed an off-site statistical procedure that yielded good results, but it is unlikely that this approach can be applied to other sites because no statistically significant a priori data classification is available as it was for their study. Kessinger et al. (1998) formulated a fuzzy logic-based procedure that uses, in addition to reflectivity data, Doppler information. Other authors considered information that is often not readily available. For example, Pamment and Conway (1998) used surface observations, infrared imagery, climatologic information, and detected lightning, and consequently limited their approach to situations for which all these data were available.

Corresponding author address: Dr. Witold F. Krajewski, Iowa Institute of Hydraulic Research, University of Iowa, 404 Hydraulics Laboratory, Iowa City, IA 52242-1585.
E-mail: witold-krajewski@uiowa.edu

The inconsistency in the off-site algorithms' performance noticed by Smith et al. (1996) suggests that most pixel-by-pixel AP detection procedures require calibration with data originating from the application site. This implies either the existence and availability of such a classified dataset or a labor-intensive effort to generate one, which involves subjective decisions by human experts. This necessity for calibration prior to the application raises several problems. First, the labor-intensive effort required to obtain a statistically significant calibration dataset is expensive. Second, it is difficult to develop automated procedures for calibration of rule-based AP detection methods, such as those presented by Smith et al. (1996), as it requires construction of decision trees. The rules themselves are prone to subjectivism, and automated procedures for rule inference have to be devised.

b. Outline of our methodology

In this paper we present a neural network-based method that is practically free of the above problems. A neural network (NN) is a multiparametric nonlinear function that can reproduce mappings between multidimensional real spaces. For a detailed description of neural networks see, for example, Bishop (1995). The NN-based procedure is a statistical approach, yet is also similar to a rule-based approach, given the NN's capability to mimic multidimensional systems. Our methodology is a multistep algorithm that includes 1) selection of predictors, 2) selection of a calibration sample, 3) calibration of the NN, and 4) application and self-evaluation. The approach is conceptually simple and can be implemented on a new dataset within a very short time (hours to days). It can be employed in an operational mode with only step 4, that is, NN application, performed in real time. We will come back to this issue in the conclusions of this paper.

First, for each nonzero pixel in the radar base scan we calculate several features (predictors) based on the volume scan reflectivity information in a local neighborhood. (A nonzero pixel is a pixel with reflectivity value below a certain threshold; in this study the threshold is 0 dBZ.) We use reflectivity fields with $1^\circ \times 1$ km resolution. The features describe in a quantitative manner local characteristics of the reflectivity field and are similar to those used in previous approaches (Moszkowicz et al. 1994; Smith et al. 1996; Grecu and Krajewski 1999). We discuss these features, and how we selected them, in section 2. Next, several scans with "clear-cut" situations of AP and rain are selected to serve as calibration/validation datasets. These provide thousands of data points and can be selected quickly if the radar database is organized efficiently (Kruger and Krajewski 1997). The NN is trained to classify each nonzero pixel as AP or rain using its associated features. The training procedure is repeated in a Monte Carlo (randomized) fashion (see, e.g., Rubinstein 1981 or No-

reen 1989), and the points not used in the training serve as independent validation data. This provides a statistically sound way of self-evaluating the performance of the calibrated NN. Finally, the procedure for detecting AP is applied. Below we provide a more detailed account of the entire methodology.

Section 2 contains the description of the method. In section 3, we provide illustrations of the procedure's application to actual radar data from two WSR-88D sites in Tulsa, Oklahoma, and Memphis, Tennessee. In the last section, we discuss our conclusions and make recommendations for future work.

2. Procedure description

The approach presented in this study is based exclusively on the analysis of the radar reflectivity scans. We assume that volume scans are available. The reflectivity scans are in many situations the only information readily available for an off-site AP detection procedure. In some situations the radial velocity scans might be available as well.

a. Selection of predictors (features)

As discussed above, in the first step of our methodology we calculated several features for each nonzero radar reflectivity pixel in the base scan. The features represent physical and statistical characteristics of a neighborhood centered on the pixel location. Second, based on those features, we developed an NN to classify each pixel in the radar scans as AP or rain. To accomplish this task, one needs a calibration dataset containing classified base scan pixels. An experienced radar meteorologist can easily, in most situations, determine such a dataset by visually inspecting radar reflectivity maps and identifying several clear cases of AP and rain scans.

Most of the features used in this study were employed in previous investigations concerning AP detection (Moszkowicz et al. 1994; Grecu and Krajewski 1999). The features represent a subjective choice of the characteristics we expected would distinguish AP and rain echoes. The features, denoted as F1–F9, are the following.

- 1) Pixel absolute velocity (F1). AP pixels are mostly stationary, while rain echoes are often associated with advecting weather systems. Although the Doppler information was not used in our study, the velocity can be estimated from the analysis of two consecutive reflectivity fields using a correlation approach (Fujita et al. 1998).
- 2) Coefficient of variation of the velocity (F2). The coefficient of variation is defined as the ratio of standard deviation to the mean of a set. This feature was selected to complement the information associated with feature F1. Moving storms display spatial consistency of their velocity, while AP results in the perceived "ragged" velocity patterns.

- 3) Height of the highest nonzero echo above the base scan pixel in question (F3). It is commonly known that this is a powerful feature distinguishing AP and rain echoes. Vertical continuity of the reflectivity (on the entire volume scan basis) is used in the WSR-88D Precipitation Processing Subsystem algorithm (Fulton et al. 1998).
- 4) Height of the maximum echo above the pixel location (F4). The justification for selecting this feature is similar to the argument above. The AP echo strength decreases quickly with height, while the maximum echo strength for convective cells is often significantly above the ground.
- 5) Reflectivity value of the pixel (F5). It is doubtful that this feature by itself would offer any significant skill in separating the AP from rain, but it has potential when used in conjunction with other features.
- 6) Radar range of the pixel (F6). As radar beam height increases with range, this information is important in complementing features F3 and F4.
- 7) Coefficient of variation of the nonzero reflectivity values situated in a small two-dimensional window (11×11) centered on the pixel in question (F7).
- 8) Coefficient of fluctuation (F8). This feature was introduced by Steiner (1998, personal communication). It represents the ratio between the number of significant fluctuations and the number of pixels in a 21×11 window centered on the pixel. A significant fluctuation is identified by analyzing three consecutive pixels along a radar beam. If the reflectivity of the middle pixel was outside the range defined by the end pixels, and the absolute difference between the middle pixel and the end pixel was greater than 2 dBZ, then an index (count) of significant fluctuation was assigned to the end pixel. This feature, along with feature F8, arose from the speculation that the AP fields are "rougher" than the rain fields.
- 9) Maximum horizontal gradient in an eight-pixel neighborhood (F9). Again, given a similar level of area-averaged reflectivity, we expect the AP fields to be more variable in space.

All the features except F1 and F2 were determined by analysis of data in the polar coordinates format as provided by the radar. To determine features F1 and F2, the data were projected into a 2×2 km² Cartesian coordinate grid.

b. NN calibration

The pixels were separated into two classes corresponding to AP and rain, and the features calculated for the selected calibration dataset. Note that it is convenient to select scans that consist entirely of AP or rain echo pixels. If this is difficult and the available scans are composed of both rain and AP, a simple windowing system delineating the two classes can be adopted.

For the selected sample, we calibrated a back-prop-

agation NN, making the classification based on our features. For the problem at hand, the NN *input* was a nine-dimensional vector containing the numerical values of the features associated with a certain echo pixel. The NN *output* was a two-dimensional vector whose components were determined by the following rule: if the echo pixel is a rain pixel, 1 is assigned to the first output component and 0 to the second one; otherwise, 0 is assigned to the first output component and 1 to the second one. In NN calibration (the process known as *training*) its internal coefficients (known as *weights*) are determined. The purpose of NN training is to minimize the root-mean-square error (rmse) of the actual NN output. This is the mechanism through which the NN learns how to classify echo pixels based on their features.

For the NN training, we used a quasi-Newton optimization method, namely, BFGS (Bertsekas 1995). This method is much faster than the classic delta rule. The training was stopped when we observed no significant reduction in rmse, thus indicating the NN had learned how to classify echo pixels based on their associated features.

For classification, a set of the same features needs to be calculated for the volume scan and provided as the NN input. On output, the NN returns its classification decision on whether the pixel corresponds to AP or rain. An implicit assumption of our procedure is that if a pixel belonging to the lowest scan is classified as AP, all pixels above it (taking into account the beam projection correction) are also classified as AP, and, conversely, if a pixel is classified as rain, all pixels above it are classified similarly. However, within the scope of our procedure, we have no basis to evaluate the performance of these implicit decisions, as, by definition, we selected only the situations where these assumptions are met. The performance evaluation of our methodology is based entirely on successful (or erroneous) classification of the base scan pixels.

In the next section we present results obtained using the above methodology on WSR-88D data from Tulsa, Oklahoma, and Memphis, Tennessee. We also discuss issues such as the NN performance evaluation, the strategy of choosing the calibration dataset, the dataset size, and the benefits and drawbacks of using neural networks.

3. Application results

Several months of WSR-88D level II (Klazura and Imy 1993) radar reflectivity data were collected for two NEXRAD systems, one in Tulsa, Oklahoma, the other in Memphis, Tennessee. Data from April to May 1994 and May to June 1995 were available for the Tulsa radar system (system identification KINX), and from May to July 1995 and March to June 1997 for the Memphis system (identification KNQA). We organized the data into an efficient online database using the compressed format developed by Kruger and Krajewski (1997). We

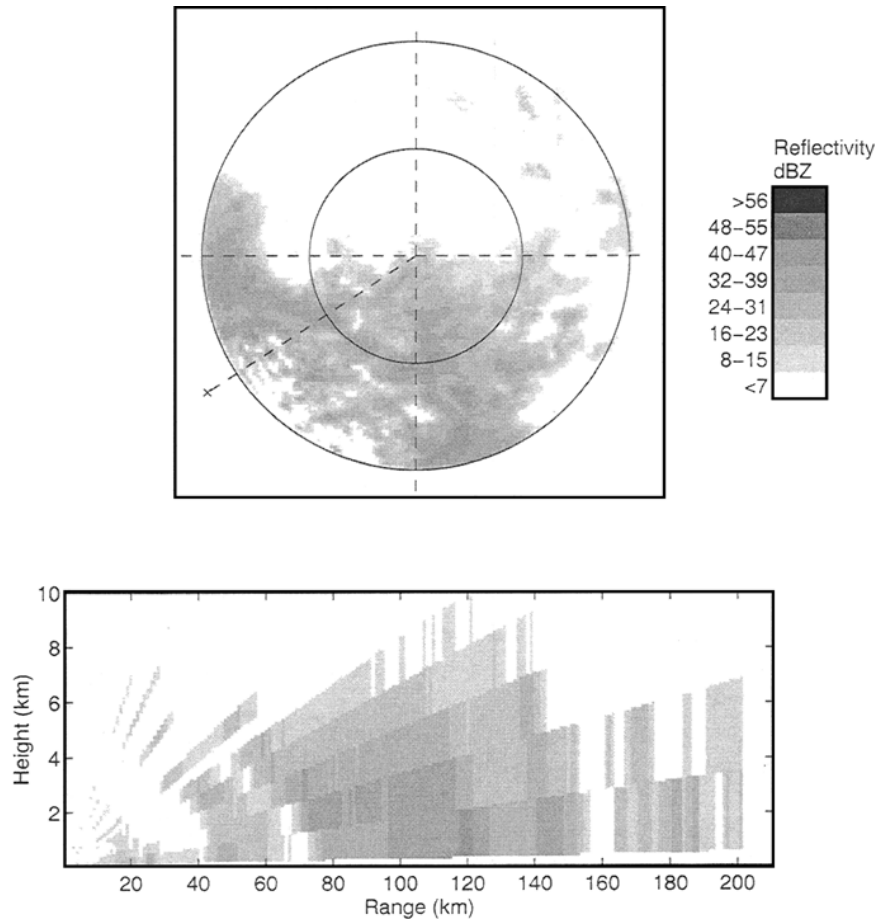


FIG. 1. An example of rain echo, Tulsa, OK, 1508 UTC 02 May 1994. The circles correspond to ranges of 100 and 200 km. The lower panel is a range–height cross section along the dotted line in the upper panel.

employed the VRAD (Virtual Radar) software, also developed by Kruger and Krajewski (1995), to visualize the data.

Selection of a dataset for the purpose of calibrating and evaluating an AP detection algorithm could be a very laborious process if it is done pixel by pixel. For the calibration process to be effective, the number of pixels selected should be large enough to include a wide variety of situations met in radar reflectivity data. This

is also a requirement for a meaningful performance evaluation. That is why it is significantly easier to identify the entire training regions or even scans that have all pixels of the same type. Based on this reasoning, we selected several clear-cut situations of AP and rain events. The basis for the selection was visual inspection of volume scans. We decided on this approach realizing the risk of neglecting the situations where AP is embedded in rain (i.e., part of the return is from the ground and part from the rain). These situations are difficult to unequivocally detect even by experienced radar meteorologists, and thus, we take this calculated chance that the NN will have difficulty properly classifying them as well.

For the sake of brevity, we describe the application of the procedure in detail only for the Tulsa site. For the Memphis case we will give only a short description of some particular aspects and the results.

a. Analysis of Tulsa WSR-88D (KINX) data

The events selected for Tulsa are given in Table 1. Illustrations of AP and rain cases are given in Figs. 1

TABLE 1. Events used in AP identification study for the Tulsa WSR-88D (KINX).

Date	Time (UTC)	Code	Type	Number of nonzero pixels
02 May 1994	1508	KINX1	Rain	45 889
16 May 1994	0503	KINX2	AP	27 590
30 Apr 1995	0328	KINX3	AP	39 473
29 Apr 1995	1606	KINX4	Rain	14 035
06 May 1995	1545	KINX5	Rain	36 695
07 May 1995	0728	KINX6	AP	40 433
09 May 1995	0513	KINX7	AP	41 927
23 May 1995	2148	KINX8	Rain	6927

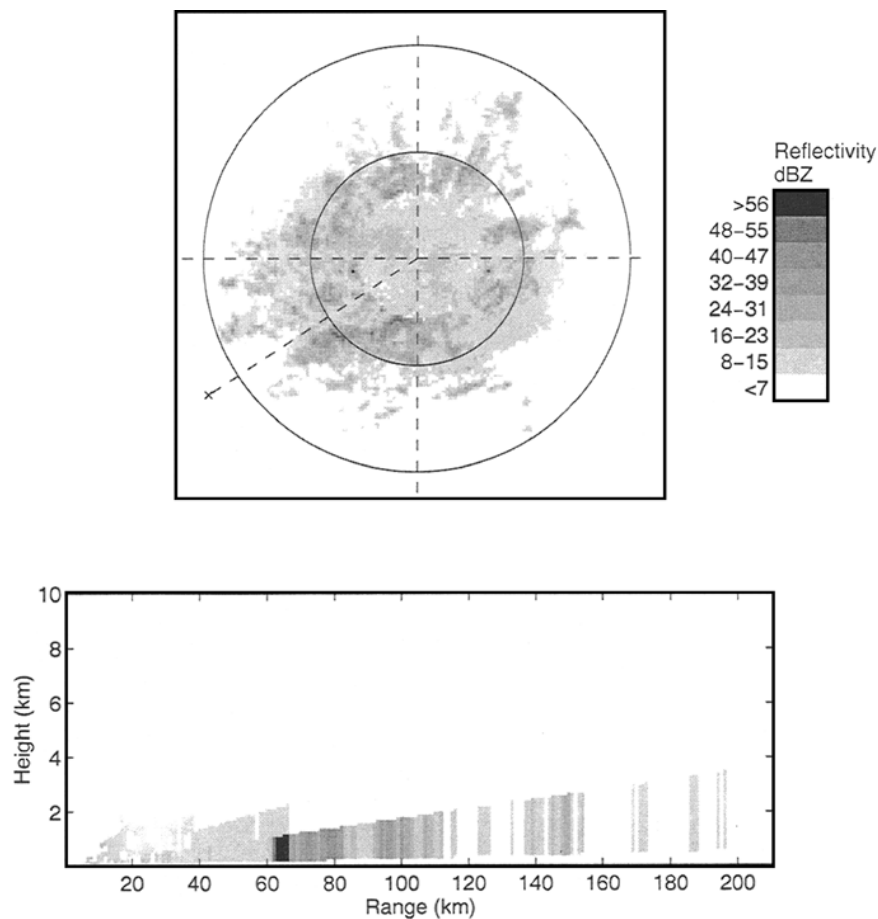


FIG. 2. An example of an AP echo, Tulsa, OK, 1016 UTC 16 May 1994. The arrangement of the figure is the same as in Fig. 1.

and 2. The features described in section 2 were calculated for all the nonzero pixels contained in the selected events. The histograms of the features were constructed for both pixel types (classes).

The motivation behind constructing these histograms was to investigate whether there is a simple rule-based approach that would allow one to distinguish AP from rain. Specifically, we analyzed whether the histograms are distinct enough so that a simple threshold could

separate a feature range into two regions corresponding to the two types of pixels. It turns out that the histograms of the individual features (Fig. 3) overlap significantly, and such a methodology would lead to large errors. The most distinct histograms corresponded to features F1, F2, and F3, namely, pixel velocity, velocity variations, and echo height, respectively. Feature F3 allows the best separation. The threshold that minimizes the classification error cannot be determined accurately unless we know the frequency of AP occurrence versus the frequency of rain occurrence. However, we can make a simplified quantitative assessment of the error. If the frequency of AP occurrence was, say, 0.2 (conditional on the base scan pixel being nonzero) and the frequency of rain occurrence 0.8, and the selected threshold is 4 km, the misclassification error would be about 9% in classification based exclusively on echo height (feature F3). This estimate was determined by calculating the weighted sum of the number of the AP cases above the threshold and the rain cases below the threshold. It is quite likely that the actual frequencies of AP and rain are not much different from the ones assumed above.

TABLE 2. The performance of NN and QDF for selected events (Tulsa WSR-88D).

Code	Type	Number of nonzero pixels	Number of misclassified pixels		Error (%)	
			NN	QDF	NN	QDF
KINX3	AP	39 473	955	1595	2.4	4.0
KINX4	Rain	14 035	126	295	0.9	2.1
KINX5	Rain	36 695	187	1119	0.5	3.1
KINX6	AP	40 433	2729	15 131	6.8	37.4
KINX7	AP	41 927	2859	16 565	6.8	39.5
KINX8	Rain	6927	242	403	3.5	5.8

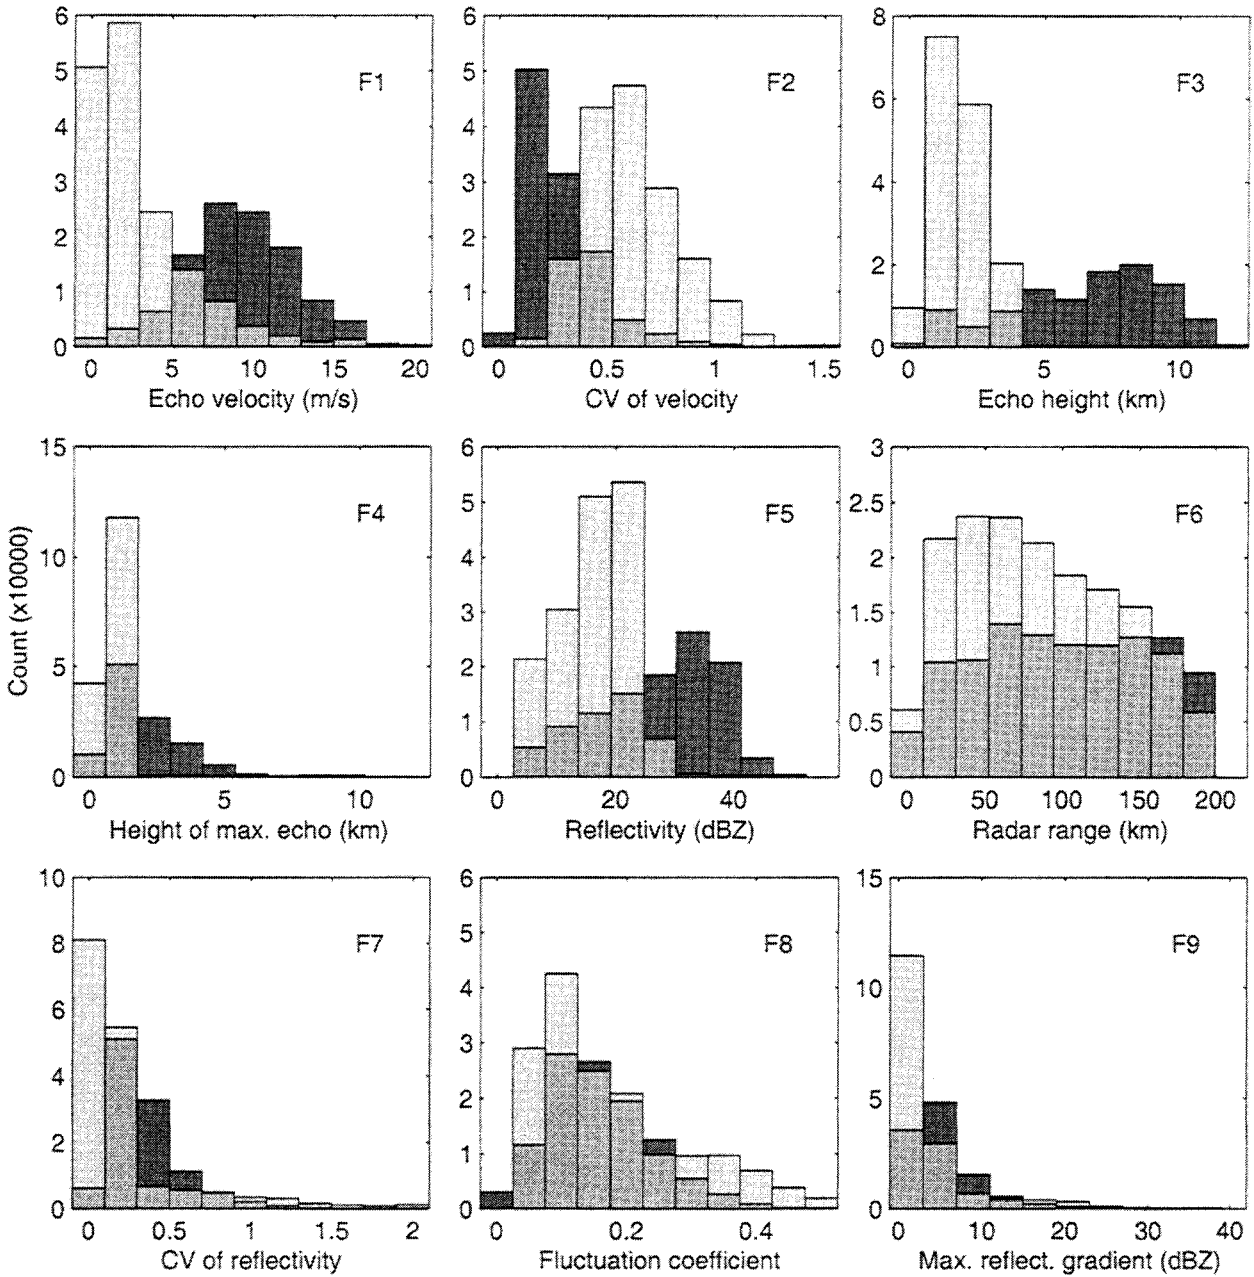


FIG. 3. Histograms of nine classification features for AP and rain pixels. The light shade is for the AP echo, the dark shade is for the rain echo, and the middle shade is for the overlapping region of the histograms.

It follows that the error is quite large; to decrease it, we need to consider information provided by other features.

The task of looking for complex classification rules is overwhelming because of the high dimensional feature space and its complexity. The complexity of the problem is caused by the exponential character specific to procedures that explore and construct decision trees (Kanal and Dattatreya 1986). Although it is possible to develop automated procedures (Kanal and Dattatreya 1986), the approach is computationally inefficient. Consequently, we adopted the NN framework as a tool for

inferring the complex relationship between the radar echo features and the echo type.

b. Application of NN procedure to Tulsa data

The NN was trained using the first two events in Table 1 and tested using the other events. The results are given in Table 2. To put our results in a broader context, we also included (in Table 2) the classification results yielded by the quadratic discrimination function (QDF) formulated in Moszkowicz et al. (1994). The same two

TABLE 3. Overall error as a function of calibration/validation dataset and the method employed (Tulsa WSR-88D).

		C1	C2	C3	C4
Error (%)	Total	3.95	6.9	19.6	32.4
	AP	5.37	5.9	27.3	55.3
	Rain	0.96	8.3	3.2	1.6

events were used for the calibration of this method. It is obvious that NN-based scheme yields much better results. Although the poorer results of the QDF method could be explained, the point is that the NN approach seems more robust.

The overall performance of our classification procedure is given in Table 3 for several different cases. The first case, referred to as C1, contains the results yielded by NN for the following events: KINX1 and KINX2 used as a calibration dataset, and KINX3–KINX8 as validation. In case C2, events KINX3 and KINX8 were used as a calibration dataset, and the other events were used as validation dataset. Cases C3 and C4 are based on the same calibration and validation datasets as C1 and C2, respectively, but the method employed was the QDF.

An analysis of Table 3 suggests that the performance of the NN-based AP detection scheme depends on the calibration dataset, and for a practical use of the scheme one has to provide the most appropriate calibration data. It is intuitive that the NN will not perform very well in a situation to which it was not exposed during the calibration. Consequently, the best strategy would be to include in the calibration dataset situations that the NN failed to classify with acceptable accuracy. However, this strategy might become impractical as the calibration dataset size might increase too much.

To overcome this limitation, one practical approach is to select at random pixels from known situations and to include them in the calibration set. Thus, we propose the following Monte Carlo calibration strategy: select at random a certain number of pixels from the entire dataset, perform the calibration using the selected pixels, calculate the error using the remaining pixels, repeat the previous steps many times and characterize the errors statistically. The above strategy allows for a better assessment of the performance of any other objective approach (e.g., QDF) as well. However, since there is no evidence in Tables 2 and 3 to support the notion that QDF leads to better results than the NN, we applied the algorithm only for the NN.

The steps of the above algorithm were repeated 100 times. To determine the effect of the calibration sample size on the results, we considered a varied number of pixels selected for calibration. The events selected in Table 1 contained 149 423 AP pixels and 103 546 rain pixels. An equal number of pixels of both types, out of the total number of pixels, were used to create the calibration dataset according to the algorithm. The results are given in Fig. 4 for misclassification of the two types of rain as well as the total error (either class). The stan-




FIG. 4. Classification error as a function of the training sample size for Tulsa, OK. The vertical bars denote ± 1 std dev range.

dard deviation of the error indicates the homogeneity of results. The performance improves quickly as the sample size increases. Error decreases slowly for sample sizes greater than about 5000. It follows that we do not have to use large training sets to achieve good performance of the NN procedure. This relatively small size of the calibration dataset illustrates the efficiency of the proposed methodology.

Since we have trained many neural networks, the question arises: which one should be used for follow-up applications? Several alternatives are possible. The simplest is to select the best of the 100 developed during the uncertainty determination. Another option would be to select one, but at random. Yet another option is to retrain the NN using a somewhat (e.g., twice) larger sample, which would guarantee performance not worse than that indicated by the Monte Carlo uncertainty study. For the following application, we used a network selected at random.

We close this section with an example. Consider a mixed situation of AP and rain on 1045 UTC 14 May 1995. This case was not included in the calibration/validation process. In Fig. 5, we give a graphical presentation of the performance of our AP detection procedure. The evolution of the radar echo in the preceding hours and the echo height analysis indicates that most likely only the northeastern part of the echo is attributable to rain. The upper and lower panels represent the radar echo before and after applying the NN AP detection. It can be observed that most of the echo around the radar was removed by the procedure. What remains seems like an AP signal but we had difficulty to positively confirm this.

c. Application of NN procedure to Memphis WSR-88D (KNQA) data

We applied the same NN-based approach to the radar reflectivity data from the WSR-88D in Memphis, Ten-




FIG. 5. Radar echo for Tulsa, OK, 1045 UTC 14 May 1994. The left panel represents the original echo and the right panel represents the echo after the application of the NN procedure.

nessee. The AP and rain cases we used to study the effectiveness of the procedure for this site are given in Table 4. It is worth mentioning that the morphology of the AP patterns of the Memphis site is quite different from that of Tulsa. While the AP patterns are generally widespread and contiguous for Tulsa, they are highly disconnected and irregular for Memphis. This difference is a strong indication that classification rules, explicit as the QDF, or implicit as the NN, must be derived for each particular case. This necessitated training the NN with data from Memphis. It is interesting to note that the total misclassification error of the NN trained for Tulsa and applied for Memphis was 16%.

The number of pixels for the cases in Table 4 is smaller than that of cases in Table 1 because most of the rain events are contaminated with AP, and the process of eliminating the AP pixels to create the calibration/validation datasets becomes more difficult. However, in an operational setting this would not be a problem since the pool of available data and, consequently, the number of uncontaminated rain and AP scans are much larger.

TABLE 4. Events used in AP identification study for the Memphis WSR-88D (KNQA).

Date	Time (UTC)	Code	Type	Number of nonzero pixels
01 Jun 1995	0800	KNQA1	AP	29 740
02 Jun 1995	0935	KNQA2	AP	10 308
03 Jun 1995	0632	KNQA3	AP	15 857
04 Jun 1995	0337	KNQA4	AP	16 555
05 Jun 1995	0524	KNQA5	AP	18 724
11 Jun 1995	1832	KNQA6	Rain	23 239
12 Jun 1995	0240	KNQA7	Rain	11 940
01 Jul 1995	0016	KNQA8	Rain	5965

We did not repeat the comparison between the NN and QDF for Memphis. We applied the randomization algorithm presented above to the Memphis data. We also studied the effect of the sample size. The results are given in Fig. 6. The overall error behavior is similar to that for Tulsa, but there are some differences in the distribution of error. For Memphis the rain classification error is somewhat larger than the AP classification error. Also, the overall error seems to be somewhat larger. The climatic and topographic differences between the sites—suggested by, among other things, the difference in the morphology of AP patterns—is probably the main cause for the variation of the procedure performance with the site.

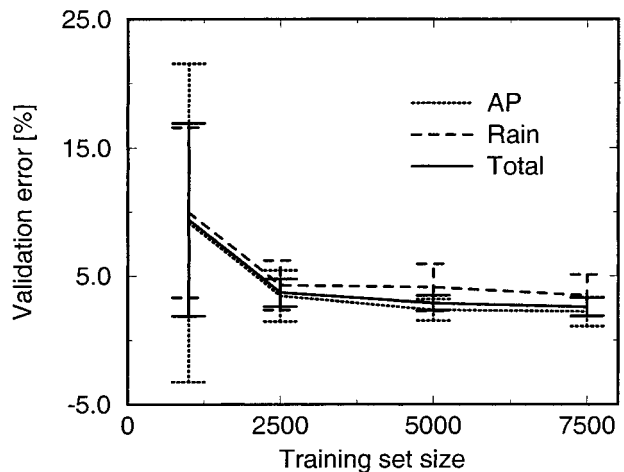


FIG. 6. Classification error as a function of the training sample size for Memphis, TN. The vertical bars denote ± 1 std dev range.

4. Summary and conclusions

We presented an efficient and effective methodology for detecting AP echoes in radar data. The method is based on volume scan reflectivity observations only and application of neural networks for classification of the base scan radar echo into the AP or rain classes. We proposed an efficient approach for selection of the training sample and a Monte Carlo algorithm for the NN performance evaluation.

Our development of the NN methodology was focused on radar-rainfall applications. This may be seen as a limitation as many other uses of radar data may be as important. To accommodate these other applications, our framework would have to be extended to include echo classification into more classes than just two. This could be done rather easily as both NN and our proposed strategy for its training are flexible enough.

Although it is not the only possibility, the NN-based approach represents a conceptually simple yet rigorous way to address the problem of AP detection. Our results show very good performance, given that a representative calibration dataset is provided. Our methodology of selecting the calibration data sample should prove attractive in both an operational environment as well as in "cleaning" smaller research datasets offline.

Still, our intention in this paper was not to introduce this as an operational technique, but rather, to demonstrate its operational potential. Before the method is mature enough for an operational implementation it should be further tested. In particular, the issue of the effects of the training sample size needs to be further investigated under varied topographical and climatological conditions. Also, we were able to objectively evaluate the performance of this approach based only on data that were easily classified "by eye." Difficult cases of AP embedded in rain may require a more subjective evaluation not amenable to automation. The operational characteristics of the method, such as computational time and memory required, need to be evaluated in comparison to other approaches. (It takes about a minute of CPU time to calculate the features and apply the NN for "cleaning" one base scan with about 50% of echo coverage using a midrange Hewlett-Packard workstation.) However, all these issues are beyond the scope of this paper; we hope to address them in future work.

Acknowledgments. This study was supported by NASA Grant NAG8-1425 under the auspices of the U.S. Weather Research Program. The radar data were provided under the Cooperative Agreement between the Office of Hydrology of the National Weather Service and the Iowa Institute of Hydraulic Research (NA47WH0495). We would like to thank Dr. Matthias Steiner for providing a computer code on which feature F8 is based. We appreciate Dr. Keeler's sharing of his results on the same problem of AP detection.

REFERENCES

- Aoyagi, J., 1983: A study on the MTI radar system for rejecting ground clutter. *Pap. Meteor. Geophys.*, **33**, 187–243.
- Battan, L. J., 1973: *Radar Observation of the Atmosphere*. The University of Chicago Press, 324 pp.
- Bertsekas, D. P., 1995: *Nonlinear Programming*. Athena Scientific, 646 pp.
- Bishop, C., 1995: *Neural Networks for Pattern Recognition*. Oxford Press, 504 pp.
- Collier, C. G., and P. K. James, 1986: On the development of an integrated weather radar processing system. Preprints, *Joint Sessions 23d Conf. On Radar Meteorology and Conf. On Cloud Physics*, Snowmass, CO, Amer. Meteor. Soc., JP95–JP98.
- Conner, M. D., and G. W. Petty, 1998: Validation and intercomparison of SSM/I rain-rate retrieval methods over the continental United States. *J. Appl. Meteor.*, **37**, 679–700.
- Doviak, R. J., and D. S. Zrnic, 1993: *Doppler Radar and Weather Observations*. Academic Press, 562 pp.
- Fujita, I., M. Muste, and A. Kruger, 1998: Large-scale particle image velocimetry for flow analysis in hydraulic engineering applications. *J. Hydraul. Res.*, **36**, 397–414.
- Fulton, R. A., J. P. Breidenbach, D.-J. Seo, D. A. Miller, and T. O'Bannon, 1998: The WSR-88D rainfall algorithm. *Wea. Forecasting*, **13**, 377–395.
- Grecu, M., and W. F. Krajewski, 1999: Anomalous pattern detection in radar echoes by using neural networks. *IEEE Trans. Geosci. Remote Sens.*, **37** (1), 287–296.
- Joss, J., and R. Lee, 1995: The application of radar–gauge comparisons to operational precipitation profile corrections. *J. Appl. Meteor.*, **34**, 2612–2630.
- Kanal, L. N., and G. R. Dattatreya, 1986: Problem solving methods for pattern recognition. *Handbook of Pattern Recognition and Image Processing*, T. Y. Young, and K. S. Fu, Eds., Academic Press, 143–162.
- Kessinger, C., E. Scott, J. Van Andel, D. Ferraro, and R. J. Keeler, 1998: NEXRAD data quality optimization. NCAR Annual Report FY98, 154 pp. [Available from the National Center for Atmospheric Research, Research Applications Program, Atmospheric Technology Division, Boulder, CO 80307.]
- Klazura, G. E., and D. A. Imy, 1993: A description of the initial set of analysis products available from the NEXRAD WSR-88D system. *Bull. Amer. Meteor. Soc.*, **74**, 1293–1311.
- Kruger, A., and W. F. Krajewski, 1995: VRAD: A motif-based radar data browser. Preprints, *27th Conf. on Radar Meteorology*, Vail, CO, Amer. Meteor. Soc., 368–370.
- , and —, 1997: Efficient storage of weather data. *Software—Practice Experience*, **27**, 623–635.
- McDonald, M., and R. E. Saffle, 1989: *RADAP II archive data user's guide*. TDL Office Note 89-2, 16 pp. [Available from Techniques Development Laboratory, National Weather Service, Silver Spring, MD 20910.]
- Moszkowicz, S., G. J. Ciach, and W. F. Krajewski, 1994: Statistical detection of anomalous pattern in radar reflectivity patterns. *J. Atmos. Oceanic Technol.*, **11**, 1026–1034.
- Noreen, E., 1989: *Computer Intensive Methods for Testing Hypotheses*. John Wiley & Sons, 229 pp.
- Pamment, J. A., and B. J. Conway, 1998: Objective identification of echoes due to anomalous propagation in weather radar data. *J. Atmos. Oceanic Technol.*, **15**, 98–113.
- Rubinstein, R. Y., 1981: *Simulation and the Monte Carlo Method*. John Wiley & Sons, 278 pp.
- Simpson, J., R. F. Adler, and G. North, 1988: A proposed Tropical Rainfall Measuring Mission (TRMM) satellite. *Bull. Amer. Meteor. Soc.*, **69**, 278–295.
- Smith, J. A., M. L. Baeck, M. Steiner, B. Bauer-Messmer, W. Zhao, and A. Tapia, 1996: *Hydrometeorological assessments of the NEXRAD rainfall algorithms*. NOAA National Weather Service Final Rep., 59 pp. [Available from Office of Hydrology, Hydrologic Research Laboratory, National Weather Service, Silver Spring, MD 20910.]