Real-time Quality Control of Reflectivity Data Using Satellite Infrared Channel and Surface Observations

V Lakshmanan[1], Miguel Valente[2*]

Abstract

Radar reflectivity data can be quality-controlled using just the radar moments, and several techniques have been proposed to do this. It is possible to use texture features as inputs to a neural network to discriminate between precipitating radar echoes, and echoes that correspond to clear-air return, ground clutter or anamalous propagation. A texture feature neural network that was recently developed at the National Severe Storms Laboratory performs much better at this discrimination than other radar-moment-based quality control methods discussed in the literature.

However none of the radar-only techniques can discriminate betwen shallow precipitation and spatially smooth clear-air return. The radar-only techniques also have problems removing some biological targets, chaff and terrain-induced ground clutter far away from the radar.

We show how the use of satellite infrared channel data and surface observations can help the radar quality problem in these situations. There are several practical issues related to using satellite and surface data, however, mostly having to do with the low spatial and temporal resolution of the non-radar observations. We describe the considerations in assimilating infrared data from satellites and surface observations from mesoscale models especially with regard to temporal resolution. We then describe using the assimilated grid to remove clutter and non-precipitating targets from radar reflectivity data. Based on archived data sets, we show that such quality control works, and would be useful, if the surface observations can be received in near-realtime.

*[1]V Lakshmanan is with the Cooperative Institute of Mesoscale Meteorological Studies (CIMMS), University of Oklahoma, and is affiliated with the National Severe Storms Laboratory. [2]Miguel Valente is at the University of Reading, Reading, UK

1. Introduction

From the point of view of automated applications operating on weather data, echoes in radar reflectivity may be contaminated. These applications require that echoes in the radar reflectivity moment correspond, broadly, to “weather”. By removing ground clutter contamination, estimates of rainfall from the radar data using the National Weather Service (NWS) Weather Surveillance Radar-Doppler 1998 (WSR-88D) can be improved (Fulton et al. 1998; Kessinger et al. 2003). A large number of false positives for the Mesocyclone Detection Algorithm (Stumpf et al. 1995) are caused in regions of clear-air return (McGrath et al. 2002). A hierarchical motion estimation technique segments and forecasts poorly in regions of ground clutter (Lakshmanan et al. 2003c; Lakshmanan 2001). Hence, a completely automated algorithm that can remove regions of ground clutter, anamalous propagation and clear-air returns from the radar reflectivity field would be very useful in improving the performance of other automated weather algorithms.

For a good review of the literature on ground clutter contamination, the interested reader is refered to (Steiner and Smith 2002). Local neighborhoods in the vicinity of every pixel in the three weather radar moments were examined by Kessinger et al. (2003) and used for automated removal of non-precipitating echoes. They achieved success by examining some local statistical features (the mean, median, and standard deviation within a local neighborhood of each gate in the moment fields) and a few heuristic features. Steiner and

Smith (2002) introduced the “SPIN” which is the fraction of gate-to-gate differences in a 11x21 local neighborhood that exceed a certain threshold (2dBZ in practice) to the total number of such differences. Kessinger et al. (2003) introduced the “SIGN”, the average of the signs of the gate-to-gate difference field within the local neighborhood. Steiner and Smith (2002) used a decision tree to classify pixels into two categories – precipitation and non-precipitating while Kessinger et al. (2003) used a fuzzy rule base using features that included the SPIN feature introduced by Steiner and Smith (2002). In addition to these elevation-based features, some vertical-profile features were also used – the maximum height of a 5dBZ echo was used by Steiner and Smith (2002). Kessinger et al. (2003) discussed the use of vertical differences between the two lowest reflectivity scans.

Neural networks (NNs) have been utilized in a variety of meteorological applications. For example, NNs have been used for prediction of rainfall amounts by Venkatesan et al. (1997) and for identification of tornados by Marzban and Stumpf (1996). In fact, Cornelius et al. (1995) attempted to solve the radar quality problem using neural networks. However, the performance of the neural network was no better than a fuzzy logic classifier (Kessigner, personal correspondence), and the neural network attempt was dropped in favor of the much more transparent fuzzy logic approach described in Kessinger et al. (2003).

Lakshmanan et al. (2003a) describe the development of a neural network that uses texture features as inputs to a neural network (Lakshmanan et al. 2003b) that can distinguish between precipitating and non-precipitating radar echoes and compare the performance of the resulting network with the fuzzy logic technique of Kessinger et al. (2003). Robinson et al. (2001) compared the technique of Kessinger et al. (2003) with the technique of Steiner and Smith (2002) and other methods, finding that the methods of Kessinger et al. (2003) and of Steiner and Smith (2002) performed equally well.

In (Lakshmanan et al. 2003b), the QC neural network was found to outperform the method of Kessinger et al. (2003). However, there were some specific situations where the texture features used as inputs to the neural network did not possess sufficient discrimation power. These situations include the presence of large, spatially-smooth, clear-air returns, the presence of terrain-induced ground clutter and chaff.

In this paper, we describe, in brief, the neural network based on texture features, and three Doppler radar moments, and describe the assimilation of data from other sensors to perform quality control.

2. Texture Feature Neural Network

Velocity data can be range-folded (aliased). In the WSR-88D, at the lowest tilt, the velocity scan has a shorter range than the reflectivity one. We therefore divided the training pixels into two groups – one where velocity data were available and another where there was no Doppler velocity (or spectrum width) information. Thus, two separate neural networks were trained. In real-time operation, the appropriate network was invoked for each pixel depending on whether there were velocity data at that point. All the neural network inputs were scaled such that each feature in the training data exhibited a zero mean and a unit variance when the mean and variance are computed across all patterns.

The final set of features used in the network were:

1. Lowest scan of velocity, spectrum width and the second lowest scan of reflectivity: local mean, local variance, difference between the data value and the mean
2. The lowest scan of reflectivity: local mean, local variance, difference between the data value and the local mean, REC Texture (Kessinger et al. 2003), homogeneity, SPIN (Steiner and Smith 2002), number of inflections at a 2dBZ threshold, SIGN (Kessinger et al. 2003), echo size.
3. Vertical profile of reflectivity: maximum value, weighted average, difference between data values at the two lowest scans, echo top height at a 5dBZ threshold.

Histograms of a few selected features are shown in Figure 1. It should be noted that these features are not linear discriminants by any means – it is the combination of features that gives the neural network its discriminating ability. The histogram of Figure 1d illustrates the result of several strategies we adopt during the training, so that higher reflectivities are not automatically accepted.

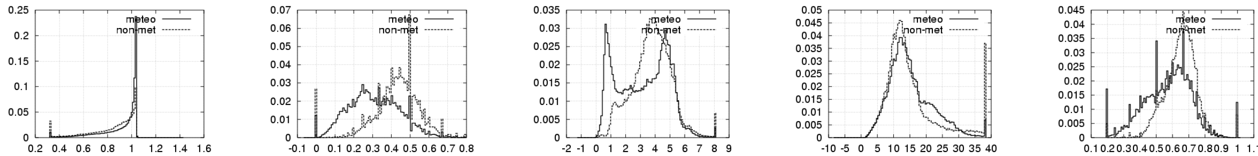

Figure 1: Histograms of selected features on the training data set, after the features have been normalized to be of zero mean and unit variance. (a)Homogeneity (b) Radial inflections (c) Mean spectrum width (d) Mean reflectivity (e) SPIN. Note in (d) that, as a result of careful construction of the training set and selective emphasis, that the mean reflectivity histograms are nearly identical – this is not the apriori distribution of the two classes since AP is rare, and clear-air return tends to be smaller reflectivity values.

We used a fully feedforward resilient backpropagation neural network (RPROP, Riedmiller and Braun (1993)) with one hidden layer. The error function that was minimized was a weighted sum of the cross-entropy (which Bishop (1995) suggests is the best measure of error in binary classification problems) and the squared sum of all the weights in the network:

$$E = E_e + \lambda \Sigma w_{ij}^2 \tag{1}$$

The first term is a variation of the cross-entropy error suggested by Bishop (1995) and is defined as:

$$E_e = -\sum_{n=1}^{N} c^n (t^n ln y^n + (1 - t^n) ln(1 - y^n)) \tag{2}$$

where t^n is the target value of the nth training pattern (0 if non-precipitating and 1 if precipitating) while y^n is the actual output of the neural network for that pattern input. N is the total number of patterns. The cost, c^n, captures the importance of that pattern. The second, square weights, term attempts to reduce the size of the weights, and thus improves generalization (Krogh and Hertz 1992). The relative weight, λ, of the two measures is computed every 50 epochs within a Bayesian framework.

The with-velocity network had 22 inputs, 5 hidden nodes and one output while the reflectivity-only network had 16 inputs, 4 hidden nodes and one output.

The training of the network, including the use of preclassification, is discussed in (Lakshmanan et al. 2003a). The peformance comparision from that study is repeated here in Figures 2 and 3.

Although the neural network computes the posterior probability that given the input vector, the pixel corresponds to precipitating echoes, adjacent pixels are not truly independent. Hence, the final 2D polar grid of posterior probabilities are mean filtered, and it is this mean-field that is used to perform quality control on the radar data. If the mean-field value is greater than 0.5, the pixel is assumed to have good precipitating data, and all elevations at that location are accepted. Bad data values are wiped out en-masse, although some researchers (e.g: Steiner and Smith (2002)) use data from higher elevations in such cases.

a. *Performance*

We used a testing set, independent of the training and validation sets and it is this independent set that the results are reported on.

3. Cloud cover

It is possible from the infrared channel of satellite data, to retreive the temperature of whatever the satellite is sensing. If not for the presence of clouds, the satellite would be sensing the ground temperature. The higher the clouds, the greater the difference between the surface temperature and the cloud-top temperature sensed by the satellite. Thus, the difference between the cloud-top temperature and the surface temperature can be used as a proxy for the presence of cloud-cover at any point (Jian Zhang, personal correspondence).

Radar tilts are obtained every 20-30 seconds. The difference field, if it is to be used to quality-control radar data corresponding to fast moving fronts, needs to be within a few minutes of the radar data. Satellite data

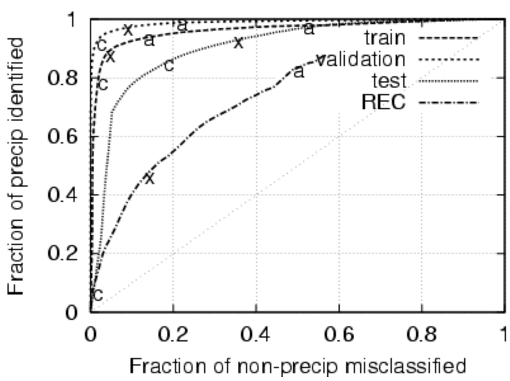

Figure 2: A ROC curve showing the performance of the neural network on the training and testing data sets. Also shown, for comparision, is the performance of the Radar Echo Classifier. Three thresholds are marked on each of the curves – a indicates a 0.25 threshold, x a 0.5 threshold and c a 0.75 threshold. Classifiers with curves above the dashed diagonal can be considered skilled. The closer a classifier is squashed to the left and top boundaries of the graph, the better it is.

is obtained every 30-minutes on average, but with the transmission times taken into account, the data may be as old as 45 minutes. However, it was shown in Lakshmanan et al. (2003c), that a K-Means clustering and advection method can be used to advect the cloud-top temperatures seen in satellite infrared data. Thus, we ingest satellite data in real-time and advect it to the time of the radar tilts. The spatial resolution is also poorer than that of radar. However, satellite data is inherently smoother, and bilinear interpolation of the data in between grid points works well. Obtaining surface temperature in real-time is more problematic. The Rapid Update Cycle (RUC2) model produces analysis grids of surface temperatures at 1-hour intervals. Unfortunately, they are also distributed nearly an hour after creation, so the data embedded in the grids is nearly 2 hours old. The steps involved in computing the cloud-cover field are shown in Figure 4.

In the preliminary studies we have conducted, we found that in precipitation regimes, the temperature difference is usually more than 20K (See Figures 5a and b). Thus, if we pick a low enough threshold (we used 5K), the risk of removing precipitation should be extremely small.

More study needs to be done to pick the appropriate temperature threshold. Ideally, instead of using a single threshold to discriminate between precipiting pixels and non-precipiting pixels, we need to:

1. Develop a relationship between the radar reflectivity value and the expected temperature difference.

2. Use the difference between the actual radar reflectivity at that pixel and the reflectivity expected based on its cloud-cover temperature difference as an additional input feature in the neural network.

In the meantime, since we don't have enough data cases to develop a meaningful relationship and train the network with it, the precipitation confidence field obtained from the radar-only neural network is combined in a fuzzy-logic manner with the temperature difference field to obtain a final precipitation confidence field. This field is thresholded at 0.5, and pixels where the confidence is below the threshold are set to missing at all tilts. The precipitation confidence field is recomputed with the arrival of every reflectivity tilt (with the vertical fields computed in a virtual volume manner Lynn and Lakshmanan (2002)). The quality control neural network, enhanced with cloud-cover, was tested in realtime during Spring 2003 in the WDSS-II system.

The utility of including multi-sensor data, beyond just the radar-only neural network, is shown for a case of biological contamination in Figure 6. Notice that the radar-only neural network does not remove all of the contamination. The cloud-cover field helps us retain the small cells to the north-west of the radar while removing the biological contamination that persists after the radar-only neural network is done.

4. Acknowledgements

Funding for this research was provided under NOAA-OU Cooperative Agreement NA17RJ1227, FAA Phased Array Research MOU, and the National Science Foundation Grants 9982299 and 0205628.

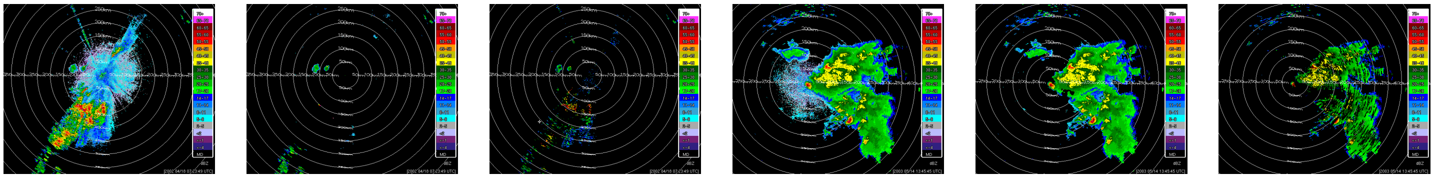

a b c d e f

Figure 3: Testing cases: (a) A data case with significant AP (b) Edited using the neural network (c) Edited using REC. Note that some very high-reflectivity AP values remain. (d) Typical spring precipitation (e) Edited using the neural network (f) Edited using REC. Note that quite a few good echoes have been removed from the stratiform rain region.

References

Bishop, C., 1995: *Neural Networks for Pattern Recognition*. Oxford.

Cornelius, R., R. Gagon, and F. Pratte, 1995: Optimization of WSR-88D clutter processing and AP clutter mitigation. Technical report, Forecast Systems Laboratory.

Fulton, R., D. Breidenback, D. Miller, and T. O'Bannon, 1998: The WSR-88D rainfall algorithm. *Weather and Forecasting*, **13**, 377–395.

Kessinger, C., S. Ellis, and J. Van Andel: 2003, The radar echo classifier: A fuzzy logic algorithm for the WSR-88D. *19th IIPS Conference*, Amer. Meteo. Soc., Long Beach, CA.

Krogh, A. and J. Hertz: 1992, A simple weight decay can improve generalization. *Advances In Neural Information Processing Systems*, S. H. Moody, J. and R. Lippman, eds., Morgan Kaufmann, volume 4, 950–957.

Lakshmanan, V., 2001: *A Heirarchical, Multiscale Texture Segmentation Algorithm for Real-World Scenes*. Ph.D. thesis, U. Oklahoma, Norman, OK.

Lakshmanan, V., K. Hondl, G. Stumpf, and T. Smith: 2003a, Quality control of weather radar data using texture features and a neural network. *5th Int'l Conf. on Adv. in Patt. Recogn.*, IEEE, Kolkota.

— 2003b, Quality control of WSR-88D data. *31st Radar Conference*, Amer. Meteo. Soc., Seattle, 522–525.

Lakshmanan, V., R. Rabin, and V. DeBrunner, 2003c: Multiscale storm identification and forecast. *J. Atmospheric Research*, 367–380.

Lynn, R. and V. Lakshmanan: 2002, Virtual radar volumes: Creation, algorithm access and visualization. *21st Conference on Severe Local Storms*, Amer. Meteo. Soc., San Antonio, TX.

Marzban, C. and G. Stumpf, 1996: A neural network for tornado prediction. *J. App. Meteo.*, **35**, 617.

McGrath, K., T. Jones, and J. Snow: 2002, Increasing the usefulness of a mesocyclone climatology. *21st Conference on Severe Local Storms*, Amer. Meteo. Soc., San Antonio, TX.

Riedmiller, M. and H. Braun: 1993, A direct adaptive method for faster backpropagation learning: The RPROP algorithm. *Proc. IEEE Conf. on Neural Networks*.

Robinson, M., M. Steiner, D. Wolff, C. Kessinger, and R. Fulton: 2001, Radar data quality control: Evaluation of several algorithms based on accumulating rainfall statistics. *30th International Conference on Radar Meteorology*, Amer. Meteo. Soc., Munich, 274–276.

Steiner, M. and J. Smith, 2002: Use of three-dimensional reflectivity structure for automated detection and removal of non-precipitating echoes in radar data. *J. Atmos. Ocea. Tech.*, 673–686.

Stumpf, G., C. Marzban, and E. Rasmussen: 1995, The new NSSL mesocyclone detection algorithm: a

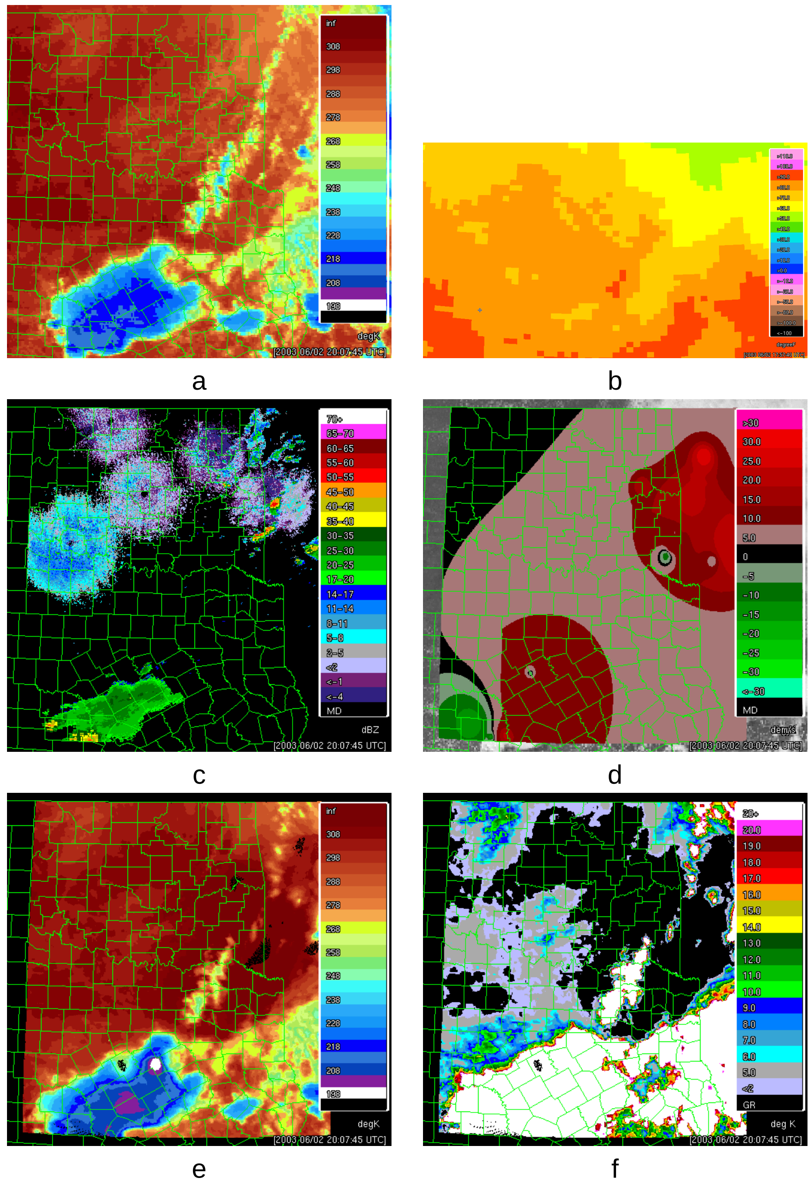

Figure 4: Steps in computing the cloud-cover field for use in quality-control of radar data: (a) Satellite infrared temperature field. (b) Surface temperature field from RUC2 model. (c) Multi-radar composite field. (d) Motion estimates derived from a multi-radar composite. (e) Cloud-top temperature advected 35 minutes to match the radar data time based on motion estimates derived from the radar field. (f) Cloud-cover field computed by differencing the advected cloud-top temperature from the surface temperature.

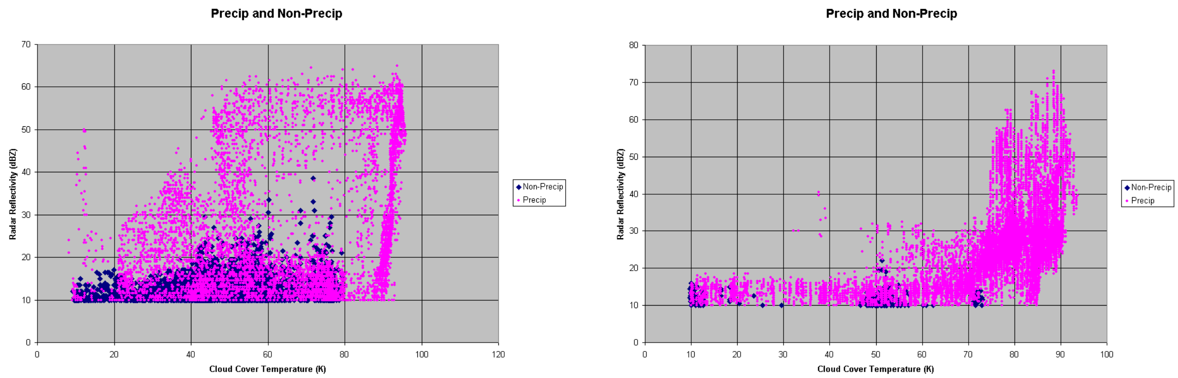

a b

Figure 5: Radar reflectivity as a function of the difference between the cloud-top temperature and the surface temperature on two storm cases ((a) May 8, 2003 (b) July 30, 2003 in Kansas)

paradigm shift in the understanding of storm-scale circulation detection. *27th Conference on Radar Meteorology*.

Venkatesan, C., S. Raskar, S. Tambe, B. Kulkarni, and R. Keshavamurty, 1997: Prediction of all india summer monsoon rainfall using error-back-propagation neural networks. *Meteorology and Atmospheric Physics*, **62**, 225–240.

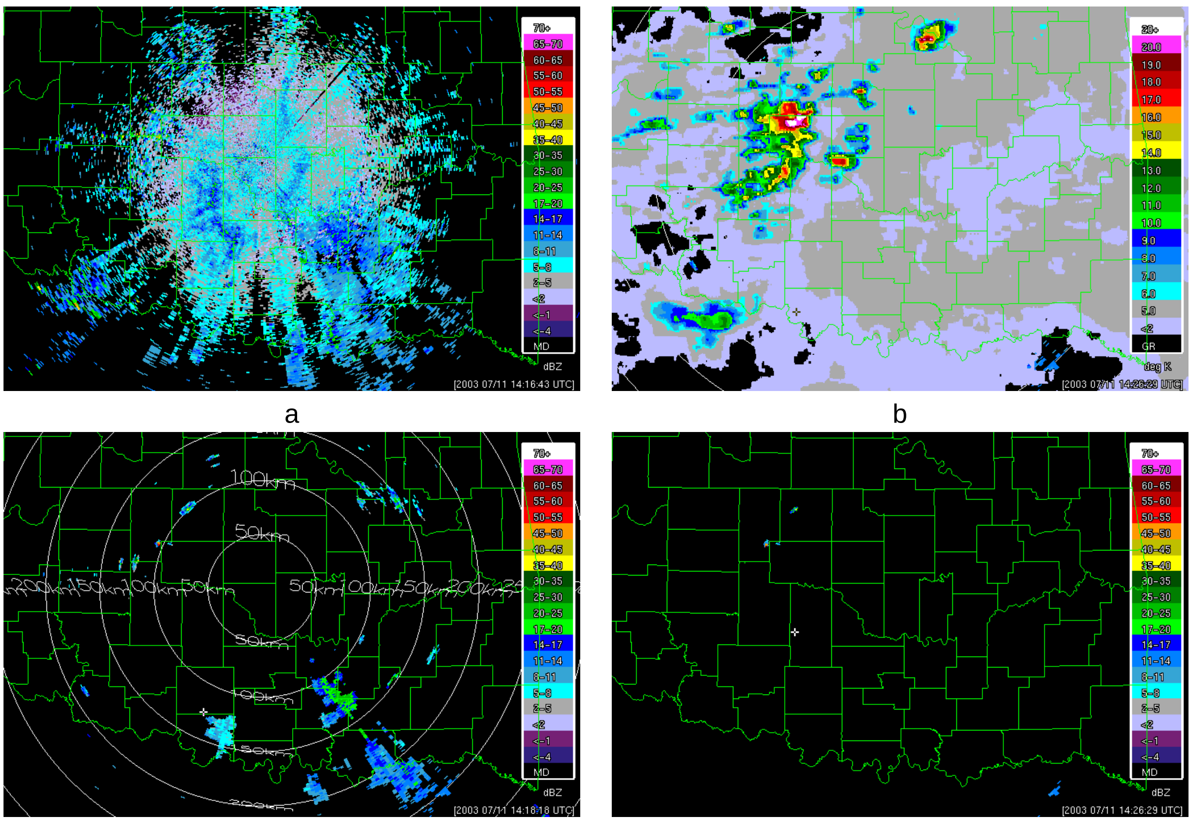

a b c d

Figure 6: (a) Radar (KTLX) reflectivity composite showing effects of biological contamination. (b) The cloud-cover field. (c) The effect of the radar-only quality control neural network. (d) The effect of using both the radar-only neural network and the cloud cover field. Note that the small cells to the north-west of the radar are unaffected, but the biological targets to the south of the radar are removed.