The Added Value of Surface Data to Radar-Derived Rainfall Rate Estimation using an Artificial Neural Network

B Root*, T-Y Yu, M Yeary
Atmospheric Radar Research Center & University of Oklahoma
Norman, Oklahoma 73072

1. Introduction

a. *Z-R Relations and their Problems*

Remote measurement of the rainfall rate has been an important product provided by the NEXRAD network. Radars have large coverage area that is not practical to achieve with a dense network of precipitation sensors. Given how important the rainfall rate estimation is to the hydrology of a region (e.g. — flash flood warnings, rainfall totals), one must understand the limitations of current methods of estimating rainfall rates, and then determine how it can be improved.

In a Z-R relationship, the rainfall rate, R, is estimated from the radar reflectivity, Z, using a simplified relationship of the form:

$$Z = aR^b \tag{1}$$

While (1) captures the overall pattern that is observed (i.e. — when it rains harder, the radar reflectivity increases), it utilizes a variety of simplifying assumptions to reach this point. This relationship uses the static parameters, a and b which are derived from an assumed gamma drop size distribution (DSD) in the scanning volume. These static parameters are usually tuned to the climatology of the region in order to improve the estimate.

As noted by Doviak and Zrnić (1993): "drop-size distribution requires an indefinite number of parameters to characterize it, and thus the radar-determined value of Z alone can not provide a unique measurement of R." This is because different DSDs within a range gate can make a significant difference in the return power of the radar signal for the same rainfall rate.

Equation (1) also assumes that the precipitation is composed of spherical drops of liquid water, as well as assuming that the vertical wind velocity is zero. If the precipitation is ice or snow or mixed, then the amount of power returned is significantly different.

*ben.root@ou.edu

In addition to the above limitations, there is the complicated issue of verification of the rainfall rate estimate. Surface-based precipitation measuring devices, like rain gauges and disdrometers, can have significant errors in measurements due to precipitation types and precipitation rates. For example, a tilt-bucket type rain gauge can be insensitive to light and heavy rainfall. Another issue is that the rainfall rate is not uniform the large area that the sensor represents, so the measured rainfall rate may not be representative of the entire area. These issues with instrumentation causes errors in rainfall rate measurements on the ground.

The problem is then compounded with the assumption that the rainfall measured at the ground is exactly the same as the rainfall scanned by the radar above. During the fall of a precipitation drop, it can change type, grow and shrink, and even be advected away from the precipitation sensor. This means that the instrument on the ground measured different precipitation than what fell through the range gate of the radar above the sensor. From an instrumentation perspective, there is not much that can be done about these verification issues.

b. *Overview of Rainfall Rate Estimation*

In order to improve the rainfall rate estimation, an examination of its basic theory is required. The rainfall rate in a unit volume is known to be a function of the drop size distribution (DSD), $N(D)$, and the fall speed of the precipitation, $w_p(D)$, relative to the ground (i.e. — terminal fall speed minus the vertical component of the wind):

$$R = \frac{\pi}{6} \int_0^\infty N(D) D^3 w_p(D)\, dD \tag{2}$$

The integral is over the domain of drop diameter, D. Any advances in rainfall rate estimation must come from a better representation of the DSD and the precipitation fall speed. For a Z-R relation, the DSD has an assumed shape, and is adjusted by pa-

rameters. These parameters allow the Z-R relation to be relevant to a particular set of weather conditions, provided that the actual DSD fits the assumed DSD well.

With equation (2), the adjustment of values for a and b is used to produce a best fit for regional and seasonal precipitation patterns. This leads to the existence of a large number of relationships. Rather than identifying which Z-R relationship to apply in a location, it may be more practical to use a single function of multiple inputs to better estimate the rainfall rate anywhere.

The DSD for the precipitation is highly dependent upon the conditions exhibited by the precipitating environment. The concentration of condensation cloud nuclei, the amount of water vapor available, and temperature all effect the DSD. In addition, the fall speed, $w_p(D)$, depends upon not only the size of the precipitation, but also the shape and material of the precipitation. These conditions are always changing over time, even within the lifetime of a precipitating event. Therefore, the improved representation of the DSD and the precipitation fall speed must be available at near the same temporal resolution as the rainfall rate estimations would be produced.

c. *Precipitation Representation*

Unfortunately, the current single-pol WSR-88D radars can not determine information about the DSD and, to a limited degree, the precipitation fall speed from the radar signal itself. Likewise, in-situ real-time observations are not feasible. Therefore, this information is not available for direct calculation of the rainfall rate. The DSD and the precipitation fall speed must be represented as functions of other observables.

Dual-Polarization weather radars can provide some of the needed observables. Some of the radar parameters produced by a dual-pol scan of a range gate are directly effected by the size and shape of the precipitation. In addition, other parameters are also effected by the type of precipitation in the range gate. Therefore, these parameters could be utilized to determine the DSD and the fall speed of the precipitation. Another source of possible observables is in-situ measurements of the atmospheric conditions. Observables such as temperature, wind, pressure and water vapor are all known to have an important impact upon DSD and precipitation fall speed (either directly or indirectly).

It is known that these observables are tightly related to DSD and precipitation fall speeds. However, the relationships are either complicated and/or unknown. Deriving any sort of theoretical model would likely prove to be impossible without major simplifications. In addition, the implementation of such a theoretical model may prove to be computationally prohibitive, due to the need to integrate over drop sizes for each range gate.

d. *Why Artificial Intelligence?*

Developing a model to relate the multitude of practical observables to a rainfall rate using theoretical physics is too difficult and would not be practical to implement. So the approach for developing this model changes from physics to empirical model fitting. If one can not derive a function from the physics, then one should try to deduce it from the observations. The typical empirical model fitting technique, linear regression, which, like the Z-R relationship, is of insufficient complexity. Therefore, a non-linear regression needs to be performed.

This project utilizes the techniques of artificial intelligence to produce a model that adapts itself to observed data and can be used operationally for similar situations. While the model itself may or may not be useful for gaining a theoretical understanding of the physical relationships involved, it is ideal for an implementation of a practical solution to a complex relationship.

2. Methods

a. *WEKA*

The Waikato Environment for Knowledge Analysis (WEKA) Witten and Frank (2005) is a software 'workbench' for experimenting with and learning many different machine learning techniques. It is actively developed by the professors and students at the Waikato University in New Zealand. The software is written in Java and is open-source. The homepage is located at `http://www.cs.waikato.ac.nz/ml/weka/`. WEKA is capable of performing a number of different data analysis tasks, which makes it ideal for comparing and contrasting multiple analysis techniques. When one is satisfied with the training of a particular AI model, WEKA can then describe the internal model parameters fully so that one can recreate that model in some other AI framework for operational use. For this project, WEKA 3.5.7 was used for training and analysis of the neural network.

b. *Multilayer Perceptron*

The AI technique chosen for this project is the 'Multilayer Perceptron' (MLP), which is a type of 'Artificial Neural Network' (ANN). The neural network is comprised of input nodes, hidden nodes, and out-

put nodes. The n-th node of the m-th layer can be functionally described in equation (3) as:

$$\text{For } f(y,w,m) = \sum_{i=1}^{P_m} y_{im} w_{im}$$

$$y_{nm} = \begin{cases} x_n & \text{for } m = 1 \\ \left(1 + e^{-(T+f(y,w,m-1))}\right)^{-1} & \text{for } 2 \leq m < M \\ T + f(y,w,M-1) & \text{for } m = M \end{cases} \quad (3)$$

For an input layer node, $m = 1$, the activation function is effectively a zero-intercept linear function. For a hidden layer node, $2 \leq m < M$, the activation function takes the form of a logistic function. For an output layer node, $m = M$, the activation function is a weighted sum of the P nodes of the previous layer plus a threshold.

c. *Training*

Training is the process by which the AI model 'learns' the desired behavior. This project utilizes a supervised training approach for the MLP with back propagation to tune the weights of the nodes. For this work, the cost function for the model to minimize is the mean squared error between the observed rainfall rate and the model-estimated rainfall rate. Also, the neural network is initialized with random values for its weights and thresholds.

d. *MLP Input/Output*

To begin the design of an artificial neural network, one must first consider the inputs and outputs of the model. For this project, there was one output: the estimated rainfall rate. The choices for inputs is more difficult. It is known from theory that rainfall rate measured by radar is a function of the reflectivity, drop size distribution (DSD), and material properties of the target (e.g. — rain, snow, ice, mix). In addition, the radar is scanning above the surface, so discrepencies can occur between surface observations and radar measurements because of differences in rainfall rates vertically. Another factor effecting the choice of input variables is availability. Because the neural network needs large amounts of data to learn from, the source of any input must be plentiful.

For input data, the obvious choice of reflectivity is made, given that it has the primary relationship with rainfall rates. As for dropsize distributions, it is decided that surface measurements would serve as a sufficient proxy for any atmospheric conditions aloft. Surface temperature, relative humidity, and air pressure were chosen for input into the model. Finally, given that there may be seasonal variations in the kinds of precipitating events that occur in an area (therefore, seasonal variations in DSD and precipitation types), the day of the year is included as input to the model. The day of the year is represented the fraction of the year completed at the time of the measurement. Therefore, the neural network structure has five input nodes and one output node.

It should be acknowledged that the use of polarimetric data would likely produce better results, given that many of the polarimetric parameters are more suited for identifying the types of precipitation being scanned. This will be attempted in a later study. The main purpose of this study is to determine the usefulness of surface observations for improving rainfall rate estimation.

e. *Training Dataset*

The training dataset was compiled using a number of sources. First, major rain events in central Oklahoma were identified using the Index of Severe Thunderstorm Events from the Storm Prediction Center and Storm Events Database from the National Climatic Data Center (NCDC). With the major precipitation events identified, Level-II radar data from KTLX was obtained for those times using the NCDC HDSS Access System (HAS). The raw Level-II files from HAS was processed utilizing the Java NEXRAD Data Exporter from NCDC under the Batch Processing mode. The reflectivity data for the lowest elevation angle was then interpolated onto a lat/lon grid and underwent a light smoothing.

Next, for each KTLX radar file, the 'Mesonet Data File' (MDF) for the Mesonet observation time closest to the observation time of the radar was downloaded and processed. For each Mesonet station in the file within 160 km of the KTLX radar station, the surface temperature, relative humidity, pressure and rainfall since 0000Z was obtained. Then, the MDFs for five minutes before and after the observation time was processed for the station's rainfall amounts since 0000Z. Using that information, a rainfall rate estimate was calculated by determining the rainfall measured during the span of the three Mesonet observations and then dividing it by the time span of the observations. At the same time, the reflectivity value at the location of the station was obtained from the radar file. Finally, this dataset was processed to randomly remove a large percentage of low-rainfall events so as to balance the training dataset. After balancing the dataset, there were approximately 1500 observations used to train the AI.

3. Results

For the hidden portion of the AI, the choice of structure comes about after much trial-and-error. The main problem is making sure that the AI model gives zero precipitation rates for low reflectivity and to exponentially increase the rainfall rate for higher reflectivities. When trying several different MLP structures, the resulting model was not grounded at zero rainfall rate (often low reflectivities yielded slightly negative or slightly positive rainfall rates), and often the model would not give high enough rainfall rates for high reflectivities. Eventually, it was found that the model structure needed that can still produce the familiar exponential relationship when viewed in the Z-R plane would be a *(6,3)* hidden structure. In other words, the first hidden layer has six nodes while the second hidden layer has three nodes. The final structure of the MLP is diagrammed in figure 1.

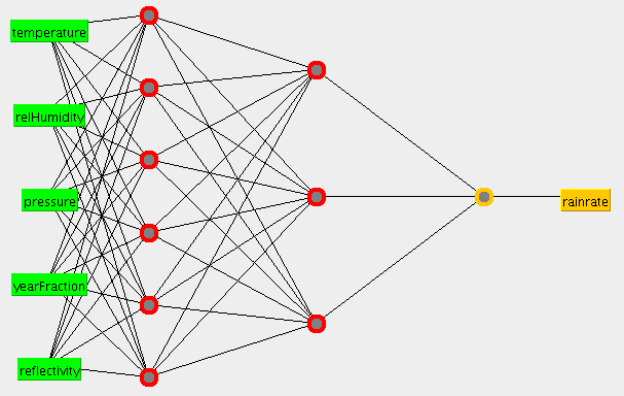

Figure 1: Diagram depicting the implemented Multilayer Perceptron.

After several attempts at training the MLP, it appears that the best training results occur with the 'learning rate' set to 0.02 and the 'momentum' value set to 0.41. As for the number of training iterations to use, it was found that the more iterations the training process spends, the better the upper end of the rainfall rate model becomes. For example, when trained for only 5000 epochs, the maximum rainfall rate predicted was only around 85 mm/hr. When trained for 15000 epochs, the maximum rainfall rate predicted was increased to about 105 mm/hr.

The final MLP model took 60000 epochs for training. Its rainfall rate as a function of reflectivity is shown in figure 2(b). For comparison, the rainfall rates as calculated by the NWS Z-R relationship, $Z = 300R^{1.4}$, is shown in figure 2(a). The red dots in these graphs are the observed rainfall rate as a function of observed radar reflectivity. The blue dots are the models' rainfall rate estimate as a function of the observed radar refectivity. The average errors of these two models are shown in table 1. Figures 3(a) and (b) shows how the two models compare with the observed rainfall rate. The line in both graphs is the model's goal. Note that for figure 3(a), the y-axis had to be restricted to 120 mm/hr, as there were some rainfall rates estimated by the NWS Z-R model that were around 500 mm/hr, and would have skewed the graph.

	NWS	MLP
Mean Absolute Error	20.01	10.58
Root Mean Squared Error	40.24	14.32

Table 1: Mean absolute errors and root mean squared errors in mm/hr between the models (NWS and MLP) and the observation used for training.

4. Conclusions

When analyzing the observed rainfall rates versus radar reflectivity, it becomes very apparent that a function of a single variable cannot fully model the necessary relationship. The observed rainfall rates can cause a wide variety of radar reflectivity. Calibration of a Z-R relationship can only improve the rainfall rate estimation to a point.

The attempt to use surface-based information together with radar reflectivity in a neural network resulted in a noticeable improvement in the rainfall rate estimation over the NWS Z-R relationship. For the dataset used for training, which has a majority of rainfall rate cases below 30 mm/hr, the neural network improved the mean absolute error of rainfall rate estimate by about 10 mm/hr. The root mean squared error was improved by almost 25 mm/hr.

The NWS Z-R relationship, when compared to the observations, appears to fit the data very well. However, the model particularly has difficulties for high reflectivity situations. It should be noted that a significant portion of the errors for the NWS model comes from the extreme over-estimation of rainfall rates from reflectivities greater than 52 dBZ. Ignoring reflectivities that are greater than 52 dBZ, the mean absolute error comes down to 12.7 mm/hr and the root mean squared error comes down to 18.6 mm/hr.

It is also possible that there is an error in the measurement of the rainfall rates by the Mesonet stations in very heavy rainfall rate situations. It is known that the instruments will underestimate the rainfall totals in such situations. However, regardless of how well the Mesonet stations handle extreme rainfall situations, the very nature of the Z-R relationship (an exponential function), means that there will be significant errors if the DSD is not the same as the one

assumed for the model when dealing with high rainfall rate situations. One will either severely underestimate or overestimate the rainfall rate.

Given that the training data could only use ground-based information to supplement the reflectivity data, the amount of improvement that the neural network produced in its rainfall rate estimates is significant. In addition, because the neural network is not constrained to a strict exponential relationship, it can handle the spread of rainfall rates that can occur at the higher reflectivities in a better manner. This means that while there may be errors in rainfall estimates for high rainfall situations, it will not be grossly overdone.

Currently, the neural network does have trouble with 'anchoring' the model down to zero-rainfall rate for low reflectivities. This might be an issue with the training dataset as an earlier training managed to anchor the model properly when much more zero-rainfall cases were used. However, that training set forced an overwhelming bias to low reflectivities and the model could not handle high-reflectivities correctly.

Therefore, given that the MLP showed improvement over the well-calibrated Z-R relation, it can be concluded that an artificial neural network can be sufficiently trained to handle rainfall rate estimations. In particular, it is the utilization of the surface observations within the MLP that improved the rainfall rate estimation.

5. Future Work

More work needs to be done to properly 'balance' the cases in the training dataset. While a significant number of zero-rainfall cases are needed to help 'anchor' the neural network to zero-rainfall, such actions tend to hurt the skill of the AI model. Also, more cases where the reflectivity is greater than 40 dBZ is needed to help improve the ability of the neural network to understand high-rainfall rate situations. The ultimate test would probably be the utilization of polarimetric data as it can provide better measurements of the DSD and precipitation type. Lastly, with a larger and improved dataset, a proper analysis will be done on how much the AI model is over-fitting the training data.

References

Doviak, R. J. and D. S. Zrnić, 1993: *Doppler Radar and Weather Observations*. Dover Publications, Inc., Mineola, New York, 2nd edition.

Witten, I. H. and E. Frank, 2005: *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, San Francisco, 2nd edition.

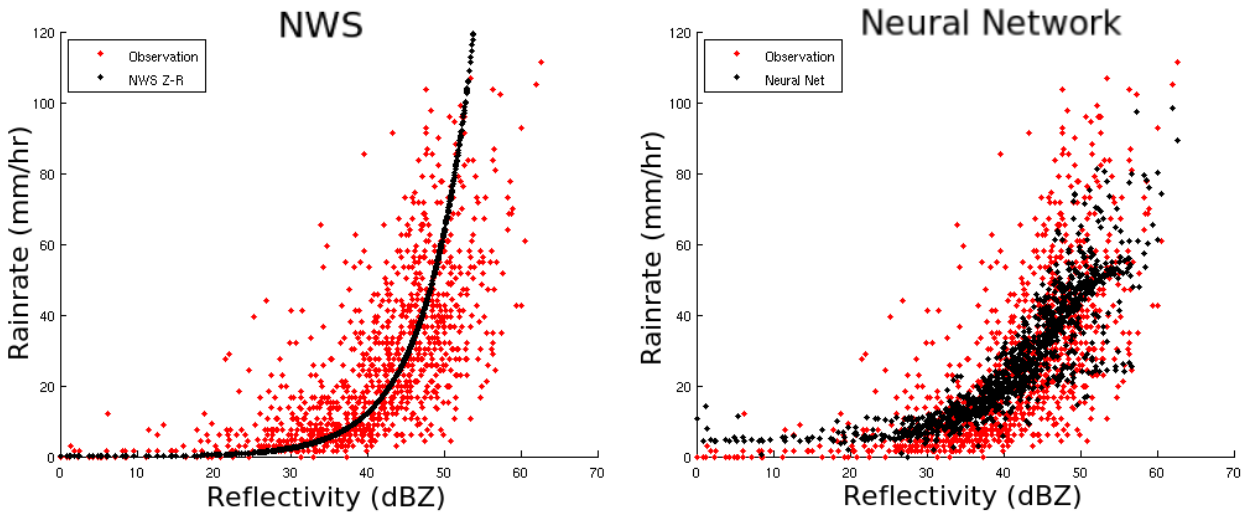

Figure 2: Comparison between the NWS Z-R model and the MLP model. These figures use reflectivity [dBZ] for the x-axis and rainfall rate [mm/hr] for the y-axis. All mesonet rainfall rate observations are depicted in red and are identical in both figures. Figure (a) is for the NWS Z-R model. Figure (b) is for the MLP Neural Network model.

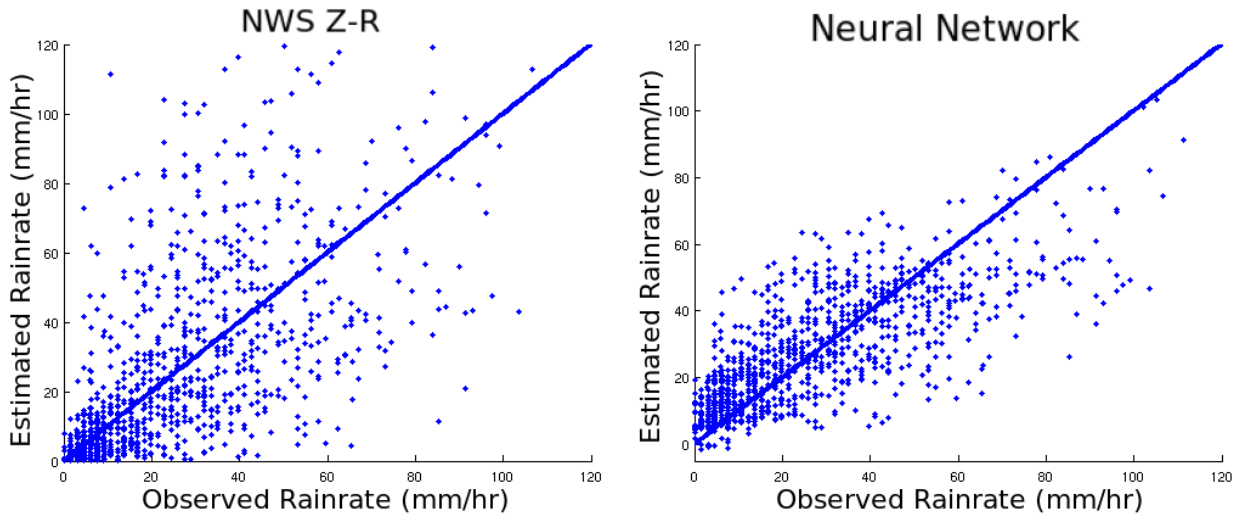

Figure 3: Correlation plots of the models' rainfall rate estimate and the observed rainfall rate. Figure (a) on the left is for the NWS model, while figure (b) on the right is for the MLP model.