J7.2 Winter Hydrometeor Classification using Polarimetric Radar and Spatiotemporal Relational Probability Trees

Andy Spencer
Department of Computer Science
Rose-Hulman Institute of Technology
Terre Haute, IN
spenceal@rose-hulman.edu

Amy McGovern
School of Computer Science
University of Oklahoma
Norman, OK
amcgovern@ou.edu

Kimberly Elmore
Cooperative Institute of Mesoscale
Meteorological Studies
University of Oklahoma
Norman, OK
kim.elmore@noaa.gov

Michael Richman
School of Meteorology
University of Oklahoma
Norman, OK
mrichman@ou.edu

Abstract

Dual polarization provides several additional radar parameters that can be used when determining precipitation types. Like reflectivity, polarimetric data points are spatially and temporally correlated and patterns in the data can be recognized and used when generating hydrometeor predictions. Current classification algorithms that use polarimetric data are based on existing theories about the characteristics of various types of hydrometeors. These algorithms work well for determining the precipitation types in volume data obtained during summer storms. However, they show very little skill when tested against observations taken during the Winter Hydrometeor Classification Ground Truth Program.

The large number of input parameters and the availability of ground observations provides an ideal environment for the application of supervised machine learning techniques. Many existing supervised learning algorithms demonstrate only small amounts of skill and are too inaccurate to be used successfully by forecasters. To improve accuracy, a spatial learning algorithm can be used so that relationships such as bright band can be automatically identified. This was accomplished by using a hierarchical clustering algorithm to group the polarimetric data into regions of similar values. Afterwards, spatial relationships between clusters are identified and a spatiotemporal relational probability tree was used to determine which relationships correspond to different precipitation types. Initial findings suggest that this approach can increase accuracy when compared to similarly sized decision trees that do not include spatiotemporal information.

Once the learning portion of the algorithm is complete, precipitation types can be classified using data that is available at the time of the event. In addition, as new forms of data become available the probability tree can be relearned with minimal changes to the algorithm itself.

Introduction

The national network of WSR-88D radars will be upgraded to polarimetric radars over the next few years. This upgrade will provide enhanced information which can be used to better classify winter hydrometeors. These additional data provide an opportunity to improve the algorithms used for winter hydrometeor classification. In addition, many current algorithms do not incorporate environmental data when performing classification.

We applied machine learning techniques to available polarimetric data in order to address the shortcomings of current methods. These techniques help to increase the accuracy of predictions in addition to providing an opportunity for data mining. This is an appropriate application of machine learning because it requires the analysis of multiple parameters that are not fully understood.

Background/Related Work

Significant research has been performed using polarimetric radars to estimate rainfall and classify hydrometeors in summer storms (Ryzhkov et al. 2003; Schuur et al. 2003). This research has lead to a variety of hydrometeor classification algorithms that can be used to classify participation in both summer and winter storms. However, when tested against observations from the Winter Hydrometeor Classification Ground Truth Program, the algorithms developed preform poorly. It is likely that some hydrometeor types classified by the existing algorithms will be accurate. For instance, ground clutter and hail will often be predicted accurately by the hydrometeor classification algorithm. Other things such as snow, sleet, graupel, and freezing rain are either predicted poorly or not predicted by the algorithm at all. The focus of our research was to apply machine learning techniques to better classify precipitation during winter storms.

Approach

In order to apply machine learning techniques to polarimetric data, we started by using the radar points as the input space for existing machine learning algorithms. We did not expect this to work well but could use it as a baseline when evaluating the performance of more advanced algorithms. When these methods turned out to be unsatisfactory we moved on to more advanced spatial algorithms that take advantage of the spatial locations of the radar points. We also obtained additional data from the Rapid Update Cycle (Benjamin et al. 2002) in hopes that it would help distinguish between certain types of hydrometeors such as rain and freezing rain which appear very similar to the radar.

Data

Our main dataset consisted of polarimetric radar data and observations obtained during the Winter Hydrometeor Classification Ground Truth Program (Elmore et al. 2007). This dataset consisted of observation of three individual winter storms which occurred from November 28 2006 to December 1 2006, January 11 2007 to January 14 2007, and January 19 2007 to January 20 2007. When combined there were a total of 2147 individual ground observations. Each ground observation was paired with a set of five-by-five grids of radar data surrounding the observation. Each set contained either 15 or 20 five-by-five grids stacked vertically and corresponding to the tilt of the radar at the time the data were obtained. These data provide our algorithm with knowledge of the atmospheric conditions in the area above each ground observation. Specifically, each point of the grid contained values for the following the radar parameters: reflectivity (Z), differential reflectivity (Z_{DR}), correlation coefficient (ρ_{HV}), differential phase (Φ_{DP}), and specific differential phase (K_{DP}). In addition, each point contained the output of the existing hydrometeor classification algorithm (HCA).

Several of the precipitation classes we are analyzing have very similar radar signatures, even with polarimetric data. For instance, freezing rain and rain appear the same because they are virtually identical while in the air. To work around this problem we included additional data such as surface temperature when developing and running our algorithms. To do this we selected temperature, humidity, freezing level, and and wind speed data that was generated by the Rapid Update Cycle. RUC data are very coarse so we selected data from the grid points nearest the ground observation. A future improvement would be to perform interpolation on the RUC data in order to obtain more accurate values.

Classifier	2.5 percentile	97.5 percentile
HCA	0.10412	0.10508
LDA	0.17458	0.17642
LDA + Temp	0.23217	0.23443

Table 1: Statistical analysis of HCA, and LDA with and without temperature using a 95% confidence interval.

Initial Results

To benchmark our algorithms we performed a statistical analysis on the hydrometeor classification algorithm by comparing the precipitation types it shows with the precipitation types observed from the ground. This comparison showed that HCA gives a mean True Skill Score (TSS or Kuipers's performance index) of 0.105. In addition, a simple linear discriminate analysis without the addition of temperature gave a mean TSS of 0.176. Including temperature data in the LDA increased the TSS to 0.233. A summary of the results from this comparison is shown in Table 1.

Other initial tests using traditional machine learn-

Classifier	Accuracy	TSS
ZeroR	39.32%	0
OneR	50.00%	0.278
IBk	39.47%	0.225
NaiveBayes	36.29%	0.175
REPTree	49.17%	0.285
RandomForest	48.41%	0.270
J48	45.83%	0.292

Table 2: Initial results from Weka

Parameter	Value	Meaning
Z	-30	No return
Z_{DR}	0	Spherical objects
ρ_{HV}	0	Consistent shape/size
Φ_{DP}	-	Used K_{DP}instead
K_{DP}	0	No loss of signal
HCA	7	Unknown precip type

Table 3: Censored values

ing algorithms also showed improvements over the hydrometeor classification algorithm. To perform these tests we used the Weka (Waikato Environment for Knowledge Analysis, Witten and Frank (2005)) data mining application to train a variety of classifiers on the data. The methods we tested included:

ZeroR Predicts based on the mode classification of training data.

OneR Predicts using the attribute that gives a minimum error.

IBk Lazy K-nearest neighbors classifier.

NaiveBayes Probabilistic classifier that uses Bayes law.

REPTree Decision tree built using gain/variance and pruned using reduced-error.

RandomForest Picks mode results of a large number of decision trees.

J48 Decision tree based on the C4.5 algorithm (Quinlan 1993).

Most of these algorithms performed better than the hydrometeor classification algorithm but still gave rather poor results as shown in Table 2.

Removing bad values

Our initial experiments were completed by simply picking the lowest elevation scan available that contained valid radar data. This provided a decent baseline but we improved upon this technique in several ways. First, instead of discarding unknown data we replaced them with censored values. Most of the unknown values arise because the signal returned to the radar was too weak to be measured. However, for most of the parameters there is a known value for this "clear air" return. For instance, clear air returns for reflectivity were set to -30 dBz. A list of the censored values we used is shown in Table 3.

In addition to setting censored values we assumed that the hydrometeor classification algorithm is good at detecting ground clutter. Therefore, we removed all the observations where the mode HCA output was ground clutter.

We encountered several difficulties while attempting to apply existing machine learning algorithms to our dataset. One reason for this was that the input vector was variable in size, containing either 15 or 20 grids associated with each observation. The actual elevation of each of these grids also varied due to the proximity of the observation to the radar taking the measurements. Furthermore, many of the grids for a particular observation contained only censored values because the precipitation did not fill the entire volume. For example, if two very similar cloud structures are located at different elevations with different distances to the radar the values contained in the grids may be completely different. That is, each point in the radar grid initially corresponds to a certain elevation and position instead of a specific part of the precipitation. One way to fix this would be to remove clear air data points around the precipitation and then scale the data to a fixed number points. After this is done a particular data point will always correspond to, for instance, the top or bottom of the precipitation. An alternative method which we to chose to pursue was to use spatial algorithms which do not rely on the position of data points in a grid.

Spatial algorithms

Polarimetric radar can be used to detect precipitation types based entirely on the differences in the phase of the returned values (Lim et al. 2005). Using this technique it is unnecessary to consider the relationships between different data points in the grid. However, the accuracy can be improved by using spatial algorithms to detect patterns in the locations of data points. One condition that can detected using spatial algorithms is the presence of a bright band (Zafar and Chandrasekar 2005). Bright band occurs when ice crystals melt before reaching the ground. Ice crystals tends to have a larger volume but reflect

Classifier	Accuracy	TSS
ZeroR	39.38%	0
OneR	50.46%	0.288
IBk	46.66%	0.305
NaiveBayes	21.55%	0.150
REPTree	51.67%	0.327
RandomForest	53.34%	0.343
J48	45.07%	0.284

Table 4: Partial derivatives in Weka

lower amounts of microwave energy than liquid water. When ice crystals begin to melt they are coated with a thin layer of water while staying roughly the same size, this causes the reflectivity in the melting region to increase. If a bright band is present it can be easily determined that the precipitation at the ground will either consist of liquid water or refrozen precipitation such as sleet or graupel.

Partial derivatives

One simple way to implement a spatial algorithm is to perform some spatial preprocessing on the data and use the result as input for existing machine learning algorithms. One such form of processing is to calculate partial derivatives with respect to space. Initially we implemented this by calculating finite differences between data points in the x, y, and z directions. A further refinement for future work will be to use a local linear least squares method to calculate a better derivative. This method can also help us determine how much noise is present in the data. Results of using partial derivatives are given in Table 4.

Relational Probability Trees

When performing spatial learning, it is advantageous to provide the learning algorithm with information on how different data points are related in three space. One of the simplest methods of doing this is to use a relational probability tree (Neville et al. 2003) and encode the data points as objects and the positions as relationships between the points. Our specific implementation was to encode each of the radar points as an object and created a bidirectional relation between points that are left-right, forward-behind, and above-below each other. As such we encoded the entire set of five-by-five grids as a single gridded structure. We also included a single "observation" point that stored surface data and the class label as attributes.

Spatiotemporal Relational Probability Trees

Similar to relational probability trees, spatiotemporal relational probability trees classify data by using a set of objects and the relationships between them (McGovern et al. 2008). However, the SRPT was specifically designed for data that includes spatial and temporal information. One of the main focuses of our research was to apply spatiotemporal relational probability trees to real world data, in our case these were polarimetric radar data.

When using SRPTs an additional clustering step was performed before creating the relations. Aggregating the radar data into clusters instead of individual points allows groups of points with similar values to be identified while also decreasing the search space by reducing the number of features and relations that are used (Lakshmanan 2001). Clustering was performed by using both the return values at each point and the location of the points in three space. Each radar parameter was clustered separately which resulted in clusters of Z, Z_{DR}, ρ_{HV}, and K_{DP}. A table of the features (clusters and observations) and relations used is shown in Table 5.

After clustering, a variety of relationships between the clusters such as "nearby" and "overlaps" were identified. In addition, relationships such as "on top of" were also included because they allow the SRPT to classify the same patterns using a smaller tree than if it only had attributes for things such as distance and direction between clusters.

Results for the SRPTs were obtained by attempting to classify each ground observation as either liquid participation, solid participation, or no participation. Experiments were run using cross validation and a variety of settings for the maximum tree depth and the alpha parameter (confidence in each distinction). The mean TSS when testing was 0.361 with a standard deviation of 0.0243 between different experiments. However, the standard deviations for the TSSs on the individual folds of each experiment were higher, around 0.1. This could be an indicator of noise in the datasets and may also have been influenced by the limited number of ground observations that were available.

Conclusions and Future Work

The existing hydrometeor classification algorithm is in need of improvement and machine learning can be used to do this. Due to the nature of the data, spatial algorithms can also be used to improve the accuracy of predictions.

Feature	Attributes	Description
Z, Z_{DR}, K_{DP}, ρ_{HV}	size, max value, min value, mean value, cluster top, cluster base, cluster volume	Clusters of radar data
HCA	number of points, predicted precip. type	Clusters of HCA predictions
observation	time, latitude, longitude, temperature, wind speed (u_{rel}, v_{rel}), relative humidity, freezing level	Surface data at observation point

Relations	Attributes	Description
nearby	edge distance, vertical distance, horizontal distance, Cartesian distance	Non adjacent clusters within a certain threshold
overlaps	A-B % overlap, B-a % overlap, total % overlap	Overlapping clusters
equals	none	Clusters have the points
contains	size difference	Cluster A contains cluster B
adjacent	none	Edges of clusters meet
above	none	Cluster A is above Cluster B

Table 5: Spatiotemporal features and relations

Currently one of the biggest challenges when applying machine learning to hydrometeor classification is that the available data are not extremely accurate. To compound this, ground observations were only available for three individual events. To supplement the radar data it is useful to also include surface data such as temperature.

In the future we would like to have additional data to work from when training the various classifiers. As the upgraded WSR-88D radars become operational it will provide a significant amount of polarimetric data so the main challenge will be obtaining accurate values to use as ground truth. In addition to obtaining more data, we would like to improve upon the SRPT algorithm so that it can inherently classify spatial data without as much of a need for clustering and determining relations between clusters beforehand.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. IIS/REU/0755462. The authors would also like to acknowledge participants in the Winter Hydrometeor Classification Ground Truth Program.

References

Benjamin, S., J. Brown, K. Brundage, D. Devenyi, G. Grell, D. Kim, B. Schwartz, T. Smirnova, T. Smith, S. Weygandt, and G. Manikin, 2002: The 20-km version of the rapid update cycle. Technical report, NOAA.

Elmore, K., K. Scharfenberg, and C. Legett, 2007: The nssl winter hydrometeor classification ground truth program: public involvement in science. *33rd Conference on Radar Meteorology*.

Lakshmanan, V., 2001: *A Hierarchical, Multiscale Texture Segmentation Algorithm for Real-world Scenes*. Ph.D. thesis, University of Oklahoma.

Lim, S., V. Chandrasekar, and V. Bringi, 2005: Hydrometeor classification system using dual-polarization radar measurements: model improvements and in situ verification. *Geoscience and Remote Sensing, IEEE Transactions on*, **43**, 792–801, doi:10.1109/TGRS.2004.843077.

McGovern, A., N. Hiers, M. Collier, D. Gagne II, and R. A. Brown, 2008: Spatiotemporal relational probability trees. *IEEE International Conference on Data Mining.*.

Neville, J., D. Jensen, L. Friedland, and M. Hay, 2003: Learning relational probability trees. *KDD '03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, ACM, New York, NY, USA, 625–630.

Quinlan, J. R., 1993: *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo.

Ryzhkov, A., S. Giangrande, and T. Schuur, 2003: Rainfall measurements with the polarimetric wsr-88d radar. Technical report, NOAA/NSSL.

Schuur, T., A. Ryzhkov, P. Heinselman, D. Zimi, D. Burgess, and K. Scharfenberg, 2003: Observations and classification of echoes with the polarimetric wsr-88d radar. Technical report, NOAA/NSSL.

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

Zafar, B. and V. Chandrasekar, 2005: A methodology to study bright band structure on a global scale from trmm precipitation radar. *Geoscience and Remote Sensing Symposium, 2005. IGARSS '05.*.