J4.6 AUTOMATED CLASSIFICATION OF CONVECTIVE AREAS IN REFLECTIVITY USING DECISION TREES

David John Gagne II*
Center for Collaborative Adaptive Sensing of the Atmosphere Research Experience for Undergraduates
University of Oklahoma
Norman, OK
djgagne@ou.edu

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

Jerry Brotzge
Center for the Analysis and Prediction of Storms
University of Oklahoma
Norman, OK

ABSTRACT

This paper presents an automated approach to classifying storms based on their structure using decision trees. The type and strength of severe weather has been related to storm morphology, and new adaptive sensing tools require some knowledge of cell structure. However, manually classification of storms is not possible when dealing with real-time data streams. An automated system can more quickly and efficiently sort through real-time data streams and return value-added output in a form that can be more easily manipulated and understood. Our method of storm classification combines two machine learning techniques, k-means clustering and decision trees. K-means segments the reflectivity data into clusters and decision trees classify each cluster. We classified both simulated and observed radar data in two ways. The first was to separate storm cells from linear systems. We then labeled each cell as isolated weak, isolated strong, or multicellular. We labeled the linear systems as trailing stratiform, leading stratiform, and no or parallel stratiform. The training and test data sets came from Advanced Regional Prediction System (ARPS) simulated reflectivity data and from a collection of composite reflectivity mosaics from the CASA IP1 network . By verifying the trees learned on simulated data with observations from the CASA network, we demonstrated that the knowledge gained from simulation can be applied to real situations.

1. INTRODUCTION

The taxonomy of storm classification presents many challenges even for human experts. The nature of a classification system changes depending on how the storm is observed (Doswell et al, 1996). The definition of a mesoscale convective complex, for instance, requires infrared satellite readings (Maddox, 1982). The limits of the observational tool also determine what storm types can be identified. For example, researchers focusing their classification systems on rainfall measurements (Baldwin et al, 2004), and radar returns (Steiner et al., 1995; Biggerstaff et al., 2000; Rigo and Llasat, 2004; Anagnostou, 2004), ignore a type like MCC and only use types visible from their specific data source. This project shares that limitation with previous work in its own focus on one variable – reflectivity – from one instrument, weather radar.

Despite the observation limitations, an automated system provides numerous advantages over manual classification. When dealing with large datasets on the order of thousands or more storms, manual extraction from a data set is impossible in a reasonable amount of time. With an automated system, however, an algorithm can quickly and efficiently iterate through the data processing information as needed. Without the burden of manual analysis, researchers can spend more time on data analysis. Another advantage to an automated system is the ability to provide storm classification in real-time. Many adaptively sensing tools (e.g., CASA radars (Brotzge et al. 2006) require some knowledge of storm structure.

Other researchers have examined the task of creating an automated system. For example, Steiner et al. (1995, SHY95) employed a technique that separated convective and stratiform

*Corresponding author address: David John Gagne II, Univ. of Oklahoma, School of Meteorology, 120 David L. Boren Blvd., Norman, OK 73072;
E-mail: djgagne@ou.edu

areas using a combination of intensity and difference from background reflectivity (1995). This is a variation on the background-exceedence technique (Biggerstaff et al., 2000). Anagnostou (2004) devised a different approach to the task of separating convective regions from stratiform regions by employing neural networks as a means to form multiple parameters into a separation function. We base our initial labeling of regions of reflectivity as convective or stratiform on Parker and Johnson's (2000) definition of stratiform regions (20-40 dBZ). We ignore regions with weaker reflectivity and label regions with reflectivity above 40 dBZ as convective.

Another form of storm classification involves identifying unique storm cell areas and tracking them while gathering information about their strength. The Storm Cell Identification and Tracking (SCIT) algorithm does exactly this by finding reflectivity intensities that exceed thresholds for each of the seven tilts of the radar and then combining those tilts to find the cell areas (Johnson et al., 1998). Then the cell areas are compared across time steps to detect motion. Our approach also identifies cell areas, but since the main goal is to differentiate between storm types, we base our type on only one frame and do not track between frames. In addition, we do not match the storm areas between tilts since our technique requires data from only one level.

Our decision tree approach is most similar to Rigo and Llasat (2004). They combined aspects of the SCIT and SHY95 algorithms and used the combination as the basis for a structural classification system. Instead of using those storm areas to train an algorithm to automatically classify the storms in their dataset, they simply used the storm areas and statistics about them as guides when classifying each image. They also assigned only one storm type per image even if more than one storm appeared whereas multiple storm types were often found and labeled within our dataset.

With this project we developed a structure-based classification system similar to that proposed by Rigo and Llasat. Our algorithm incorporates two machine learning techniques, k-means clustering and decision trees, to identify and classify storm areas. The k-means clustering section of the project is derived from one used for image segmentation that had been applied to reflectivity (Lakshmanan, 2001; McQueen, 1967).

We chose decision trees (Quinlan, 1986). as our primary approach for classification for two reasons. One of the biggest factors is the intuitive understanding of the model. Many other machine learning models, such as neural networks, are difficult to interpret, particularly by non-computer scientists. As Figure 1 illustrates with its example decision tree, the relationships between attributes can be easily shown and converted into other forms even by people with little background in computer science or machine learning. Second, and related to this decision, decision trees are selective. This means they can identify the most important attributes from a data set and ignore the less important ones. This ability to be selective adds to their human readability and can yield improved understanding about what is most important in a data set. However, since decision trees make their decisions based on only one attribute at every point in their structure, they can only find linear relationships within a dataset. Other approaches may be difficult to interpret but can possibly lead to improved results. With this in mind, we compare our final results to several other standard machine learning approaches.

2. DATA AND METHODOLOGY

The data used for this project comes from both simulation and observed radar reflectivity from several storms in southwest Oklahoma. The simulations are generated by the Advanced Regional Prediction System (ARPS), a storm-scale model with numerical weather prediction and data assimilation features (Xue et al., 2001, 2002, 2003). We have over 250 simulations of mesoscale storms generated in a supercell regime (Rosendahl, 2008). For this project, we examine the reflectivity at 4 km. The simulations used a 100 km by 100 km grid with 500 meter spacing. The reflectivity values of the simulated data tend to be higher than observed values of reflectivity because there is no attenuation or drop off in power with distance from the radar.

Our second source of data was made available the Center for Collaborative Adaptive Sensing of the Atmosphere (CASA) IP1 network, a group of four small, X-band Doppler radars located in southwest Oklahoma (Brotzge et al, 2006). We mapped the reflectivity from each of the radars to a single 120 km by 120 km Cartesian grid with 500 m grid spacing to fit the image as closely as possible to the ARPS simulated data.

To identify individual storm regions, the program first divided a given reflectivity image into a specified number of clusters using the k-means clustering algorithm. To do this, the algorithm minimizes a Euclidean distance equation derived from the image segmentation algorithm of Lakshmanan (2001):

$$d_e = \lambda \mid r_m - r_p \mid +(1-\lambda)\sqrt{(x_m - x_p)^2 + (y_m - y_p)^2}$$

(Eq. 1)

In Eq. 1 λ weighs the differences in reflectivity versus Cartesian coordinates, r represents the reflectivity value in dBZ at a certain point, x and y are the coordinates of that point, m designates variables derived from the list of means, and p designates variables derived from a point in the reflectivity image. The first part of the equation seeks to find the distance of each point from the reflectivity means while the second part finds the distance between the selected point and the coordinates of the reflectivity means in the image. For this work, we chose a *λ* of .6 through empirical testing. K-means clustering uses this similarity metric to find geographically similar areas with similar reflectivity readings. We process the output of k-means clustering by breaking clusters that are not contiguous and by removing clusters whose area is less than 4km^2.

Morphological	Reflectivity	Control
Eccentricity	Maximum	Area
Maj. Axis Len.	Minimum	Mean St. Dist.
Min. Axis Len.	Mean	
Orientation	Std. Deviation	
Equiv. Diam.	Range	

Table 1. Attributes used in the decision tree.

We divide the clusters into convective, stratiform, and low reflectivity areas. If at least 70 percent of the cluster contains reflectivity between 20 and 40 dBZ, then it is considered stratiform. Otherwise, if less than 10 percent of the cluster contains reflectivity greater than 80 percent of the maximum reflectivity, then the cluster is considered a low reflectivity area. If the cluster fits neither of those categories, then it is considered convective.

Given reflectivity data at every grid point in the domain, the number of possible ways to examine the data is large. Humans examine the data using a variety of visual features and combine these with their experience. Our decision tree attributes are shown in Table 1. The morphological attributes come from fitting the storm region to an ellipse. These attributes include the coordinates of the centroid, the length of the major and minor axes, the orientation of the ellipse in relation to the major axis and the horizontal, and the eccentricity of the ellipse. The reflectivity attributes include the maximum, minimum, mean, standard deviation, and the range of the reflectivity. The final attribute is the mean stratiform distance. In order to calculate this value, the program finds the equation of the major axis line and solves it to find the offset value from the line through the midpoint of the cluster centroids. It uses Equation 2.

$$d_s = m_l(x_s - x_c) + (y_c - y_s)$$

(Eq. 2)

In Equation 2 m_l represents the slope of the line, the *s* values represent centroid coordinates of the stratiform cluster and *c* values represent the same for the convective clusters. High positive values correlate more with the mean stratiform sitting in front of the convective area, values near 0 indicate that the mean stratiform center lies almost along the line, and high negative values indicate the mean stratiform lies behind the line.

Clusters are labeled within a hierarchal classification system that combines types developed by Parker and Johnson (2000) and Rigo and Llasat (2004). At the highest level, convective areas are divided into cell and linear system. Within the cell category, the storms are then subdivided into weak cell, strong cell, and multicell. In the linear system category, storms are divided into leading stratiform, trailing stratiform, and no/parallel stratiform.

In order to train the decision tree, we hand labeled a subset of the data. Two data sets of 514 and 819 storms were labeled by two undergraduate meteorology students. We created a hand-labeling interface that provides the basic information required to visually classify storms. It displays an image of the reflectivity on the left side and a diagram of the convective cluster locations on the right side. When a clustered storm appears on the screen, we can select an individual cluster and then choose the appropriate classification for it from a pull down menu. Weak cells tended to be small areas of light to moderate reflectivity. Strong cells tended to be isolated areas of high reflectivity combined with other features indicating a powerful storm such as an overshooting top or a hook echo. Multicells were generally clusters that contained multiple areas of high reflectivity intermixed with weaker reflectivity. Since the reflectivity image for the ARPS data is centered on the main storm rather than on a fixed point making the actual direction of storm motion difficult to determine, a stratiform area to the east of the storm is considered leading while one to the west is considered trailing. To form the training and validation sets of the storms, we hand labeled five random time steps from each of the ARPS simulations.

The decision trees were trained using the Waikato Environment for Knowledge Analysis (WEKA) version 3.5.6, developed by the University of Waikato in New Zealand (Witten and Frank,

2005). WEKA is a suite of various machine learning and data mining algorithms. For the purposes of this project, we used the data in WEKA to generate eight decision trees based on different combinations of statistical values from the storm data. We generated decision trees for the general and specific storm types based on the morphological and reflectivity-based variables, only the morphological, only the reflectivity, and a control tree given only area and mean stratiform distance. Figure 1 shows an example of one of the general decision trees.

In order to demonstrate the classification ability of the decision tree, we compared it to other classification algorithms within WEKA. We used the Multilayer Perceptron, a neural network model; the Logistic Model Tree (LMT), a decision tree that has Logistic Regressions at each of its leaves and; the Logistic Regression, which fits the data to a logistic function. To compare the different models, we used their average classification accuracy and their area under the curve (AUC) on the four different data sets (Bradley, 1997 and Provost and Fawcett, 1997). AUC is an alternative measure for comparing algorithms that is not sensitive to the underlying distribution of the data. AUC is measured as the area under the Receiver Operating Characteristic (AUC) curve. This curve is plotted as the probability of a true-positive result versus the probability of a false-positive result. An AUC of 1 means the algorithm performed perfectly. An AUC of 0.5 means the algorithm formed as well as random. Anything below 0.5 is performing worse than randomly choosing the answer.

Because boosting (Schapire, 1990, 1999) is guaranteed to not hurt the performance and is very likely to improve it, we further experimented with adding boosting to the decision trees. Boosted decision tree stumps (single level trees) still satisfy our goals of being able to share results easily with non-computer scientists. Boosting takes a single classification algorithm and turns it into an ensemble algorithm. It creates multiple models of the same type and weighs each of them based on their classification accuracy. The resulting weighted model will then be able to make predictions more accurately than the individual models it was based on.

3. RESULTS

Figure 2 shows the distribution of different storm types in the hand-labeled ARPS training data. Because of the regime used to generate the simulations, the storm types focus on isolated strong and isolated weak categories with 387 of the 519 storms falling into one of those two areas. The rest of the storms were almost evenly distributed across multicell, leading stratiform, and trailing stratiform. None of the storms in the training set were labeled as no stratiform. ARPS data set 1 contained a similar distribution of storms but in a much smaller number as it was randomly drawn from the same set as the training data. ARPS set 2 contained an even higher proportion of isolated strong and weak storms and very few of the other types. The CASA set featured a higher proportion of isolated weak and leading stratiform storms, but it only held a total of 34 storms from three days.

We trained eight decision trees based on different combinations of attributes. We hypothesized that our morphological and reflectivity attributes were critical and tested this hypothesis by training several sets of trees. Trees designated as control (C) used only area and mean stratiform distance in their generation process. Morphological trees used the control attributes plus the morphological attributes derived from fitting the storm area to an ellipse. Reflectivity trees (R) used the control attributes reflectivity-based attributes. Morphological and reflectivity trees (MR) used all three classes of attributes.

We evaluated all the test sets on the eight different decision trees. First we examined the overall accuracy of each tree (Fig. 3). The morphological and reflectivity tree (GMR) and the morphological (GM) tree had the highest mean accuracy at 90.39% while the tree that only used reflectivity attributes (GR) had the lowest mean accuracy at 82.17%. For the specific tree types, the one with the highest mean accuracy used the morphological attributes (SM) at 65.695% while the lowest came from the tree that used only reflectivity (SR, 62.06%).

For the general type trees (Fig. 4), the cell morphological and linear system morphological trees performed the best in both cells and linear systems with a mean AUC value of .912. This is averaged over all test sets. The lowest mean AUC for the general types, .836, came from the cell reflectivity and linear system reflectivity trees. When the AUC for each type was compared across the different datasets, there was very little difference in the AUC values.

Within the isolated strong cell type (Fig. 5), control, morphological, and reflectivity ranged from .802 to .782 while morphological and reflectivity combined had a lower value of .743. All the isolated weak storms performed between .85 and .89. For the multicell storms, three of the four trees received an AUC in the low .8s, but the

morphological value lay at .689, much lower than the rest. Across each linear type (Fig. 6), there was little variance between each tree in each set although leading stratiform performed better than trailing stratiform.

In the model comparison test, a chi square goodness of fit test was performed on the models to determine if there were significant differences in their classification ability. It found a p-value of .0844, upholding the null hypothesis that the models were not significantly different.

As expected, our implementation of boosting improved the performance in both accuracy and AUC. An partial example of one of our boosted trees appears in Figure 8. For the general tree type, AUC improved on for cells from .945 to .969 on the training set, from .936 to .944 on ARPS 1, from .909 to .936 on ARPS2, and from .879 to .913 on the CASA set (Fig. 9). For the line type, the improvements were similar.

4. DISCUSSION

The trees containing morphological attributes perform just as well as the trees containing both morphological and reflectivity attributes and better than the ones only containing reflectivity attributes. The mean AUC for cells and linear systems does not vary significantly across the three datasets (Fig. 7). The fact that the tree learned using simulated data performed well on actual radar data (the CASA IP1 data) is critical as it means the trees were general across both simulated and observed data.

When analyzing the performance of the specific type trees, the influence of the reflectivity variables becomes more apparent. This is especially true for the multicell morphological tree, which has a much lower AUC than the multicell morphological and reflectivity and multicell reflectivity trees, indicating one of the reflectivity-based variables is the determining factor for this storm type. The same relationship exists for isolated weak. With isolated strong, however, the dip occurs with the MR tree. For the two linear types, the AUC varies little across each type, indicating that the attribute for determining line type was mean stratiform distance, one of the control attributes.

The model comparisons illustrate two important points. First, our choice of decision tree based on the human readability of the model was a correct choice. Second, the lack of significant difference between the models as shown by the chi square test shows that while changing the model can cause some improvements in accuracy, those improvements will be limited by the quality of the data. As the other graphs of the trees with different given attributes show, the attributes given to the classifier matter much more than the classifier itself.

Our experiments with boosting showed a promising way to gain slight increases in the classification ability of the algorithm. While it does not have the same drastic effects as the modification of attributes, it is guaranteed to at least slightly strengthen any classification system it is applied to.

5. CONCLUSIONS

We have found that decision trees are a viable method for automatically determining storm type. The trees that distinguished between cells and lines had a high AUC and accuracy across all datasets, indicating strong performance overall. The more specific trees experienced decreasing performance across datasets, which is an area for future work. Even though we learned the trees on simulated data, they were still able to classify real world data (the CASA data) with a high degree of accuracy in the case of the general type tree. In addition, because the decision trees are selective and human readable, we could determine that the morphological attributes were most critical to successful classification.

6. CURRENT AND FUTURE WORK

We are currently working on an improved version of the storm classification algorithm and implementing the algorithm in the CASA IP1 Systems Operation Control Center (SOCC). For the second version of the specific type classifier, we are defining our storm types in greater detail and adapting and adding the attributes so the extra details can be described. While the previous system calculated cluster statistics based on only one time step using just reflectivity, the new system compares the current time step with the previous time step to calculate the direction of storm motion and the speed of the storm. Storm motion is important if associated with stratiform rainbands in order to determine if the stratiform region is trailing or leading the convective line.In the previous version, the lines were assumed to be traveling from west to east, so the east side of the line was leading and the west was trailing. The additional attributes make this distinction no longer necessary.

Another major new attribute is data concerning the wind field in the storm model. Knowing the wind speed and direction at any point in the environment makes it possible to determine if winds associated with a severe thunderstorm are

present and if the thunderstorm has a mesocyclone, making it a supercell (Doswell et al, 1996). Being able to classify supercells with the algorithm will help increase its abilities. Figure 8 displays the new algorithm's ability to detect the wind field and find the storm speed and direction of motion. Although this data cannot be observed by the radars, we are hypothesizing that it will significantly improve the results.

Currently, CASA is also adopting this algorithm for real-time use within the CASA SOCC. A storm classification algorithm built into the radar has many potential uses. First, it can determine the scanning strategy of the radar. The adaptive nature of the CASA system focuses radar scanning only on those areas of interest. The addition of this algorithm would allow it to only focus on storms of a specific morphological type. . In addition, the classifications can form the basis for a storm catalog within the CASA data. The algorithm can tag each scan with the types of storms, thereby improving metadata for long-term research applications.

ACKNOWLEDGEMENTS

This work is supported in part by the Engineering Research Centers Program of the National Science Foundation under NSF award number 0313747. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the National Science Foundation. We would like to thank Dr. Keith Brewster for providing the programs and assistance for extracting the CASA reflectivity data. In addition, we offer our thanks to Nathan C. Hiers for his help with labeling the ARPS 2 data.

REFERENCES

Anagnostou, E., 2004: A convective/stratiform precipitation classification algorithm for volume scanning weather radar observations, *Meteorological Applications*, **11**, 291-300.

Baldwin, M.E., J.S. Kain, and S. Lakshmivarahan, 2005: Development of an Automated Classification Procedure for Rainfall Systems. *Mon. Wea. Rev.*, **133**, 844–862.

Bradley, A. P., 1997: The use of the area under the ROC curve in the evaluation of machine learning algorithms, *Pattern Recognition,* **30**, 1145–1159.

Biggerstaff, M.I., and S.A. Listemaa, 2000: An Improved Scheme for Convective/Stratiform Echo Classification Using Radar Reflectivity. *J. Appl. Meteor.*, **39**, 2129–2150.

Brotzge, J., Droegemeier, K. K., and McLaughlin, D. J., 2006: Collaborative adaptive sensing of the atmosphere (CASA): New radar system for improving analysis and forecasting of surface weather conditions. *Journal of the Transportation Research Board*, (1948):145–151.

Doswell, C. A., III, H. E. Brooks, and R. A. Maddox, 1996: Flash flood forecasting: An ingredients-based methodology. *Wea. Forecasting,* **11,** 560–581.

Johnson, J.T., P.L. MacKeen, A. Witt, E.D. Mitchell, G.J. Stumpf, M.D. Eilts, and K.W. Thomas, 1998: The Storm Cell Identification and Tracking Algorithm: An Enhanced WSR-88D Algorithm. *Wea. Forecasting*, **13**, 263–276.

Lakshmanan, V., 2001: A Hierarchal, Multiscale Texture Segmentation Algorithm for Real-World Scenes, Ph.D. dissertation, University of Oklahoma, 115 pp.

Maddox, R. A., D. M. Rodgers, and K. W. Howard, 1982: *Mesoscale convective complexes over the United States during 1981- annual summary, Mon. Wea. Rev.*, **110**, 1501–1514.

McQueen, J.B., 1967: Some Methods of Classification and Analysis of Multivariate Observations, *Proc. 5th Berkeley Symp. Mathemat. Statist. Probability*, **1**, 281–296.

Parker, M.D., and R.H. Johnson, 2000: Organizational Modes of Midlatitude Mesoscale Convective Systems. *Mon. Wea. Rev.*, **128**, 3413–3436.

Provost, F. and Fawcett, T., 1997: *Analysis and Visualization of Classifier Performance: Comparison under Imprecise Class and Cost Distributions.* In Proceedings of the Third International Conference on Knowledge Discovery and Data Mining, 43-48.

Quinlan, J.R., 1986: Induction of decision trees, *Machine Learning*, **1**, 81-106.

Rigo, T. and M. C. Llasat, 2004: A methodology for the classification of convective structures using meteorological radar: Application to heavy rainfall events on the Mediterranean coast of the Iberian Peninsula, *Natural Hazards and Earth System Sciences*, **4**, 59-68.

Rosendahl, D. H. (2007). Identifying Precursors to Strong Low-Level Rotation within Numerically Simulated Supercell Thunderstorms: A Data Mining Approach. Master's thesis, University of Oklahoma.

Schapire, R. E.,1990: The strength of weak learnability. *Machine Learning*, **5**(2), 197-227.

Schapire, R. E., 1999: Theoretical view of boosting and applications. In *Algorithmic Learning Theory: Proceedings of the 10th International Conference (ALT '99)*, pp 13-25, Springer-Verlag, Berlin.

Steiner, M., R.A. Houze, and S.E. Yuter, 1995: Climatological Characterization of Three-Dimensional Storm Structure from Operational Radar and Rain Gauge Data. *J. Appl. Meteor.*, **34**, 1978–2007.

Witten, I., Frank, E., 2005: *Data Mining: Practical machine learning tools and techniques*. 2nd edition. Morgan Kaufmann, 525 pp.

Xue, M., Droegemeier, K. K., and Wong, V., 2000: The Advanced Regional Prediction System (ARPS) - a multiscale nonhydrostatic atmospheric simulation and prediction model. Part I: Model dynamics and verification, *Meteor. Atmos. Phys.*, **75**:161–193.

Xue, M., Droegemeier, K. K., Wong, V., Shapiro, A., Brewster, K., Carr, F., Weber, D., Liu, Y., and Wang, D., 2001: The Advanced Regional Prediction System (ARPS) - a multiscale nonhydrostatic atmospheric simulation and prediction tool. Part II: Model physics and applications, *Meteor. Atmos. Phys.*, **76**:134–165.

Xue, M., Wang, D., Gao, J., Brewster, K., and Droegemeier, K. K., 2003: The Advanced Regional Prediction System (ARPS), storm-scale numerical weather prediction and data assimilation, *Meteor. Atmos. Phys.*, **82**:139–170.

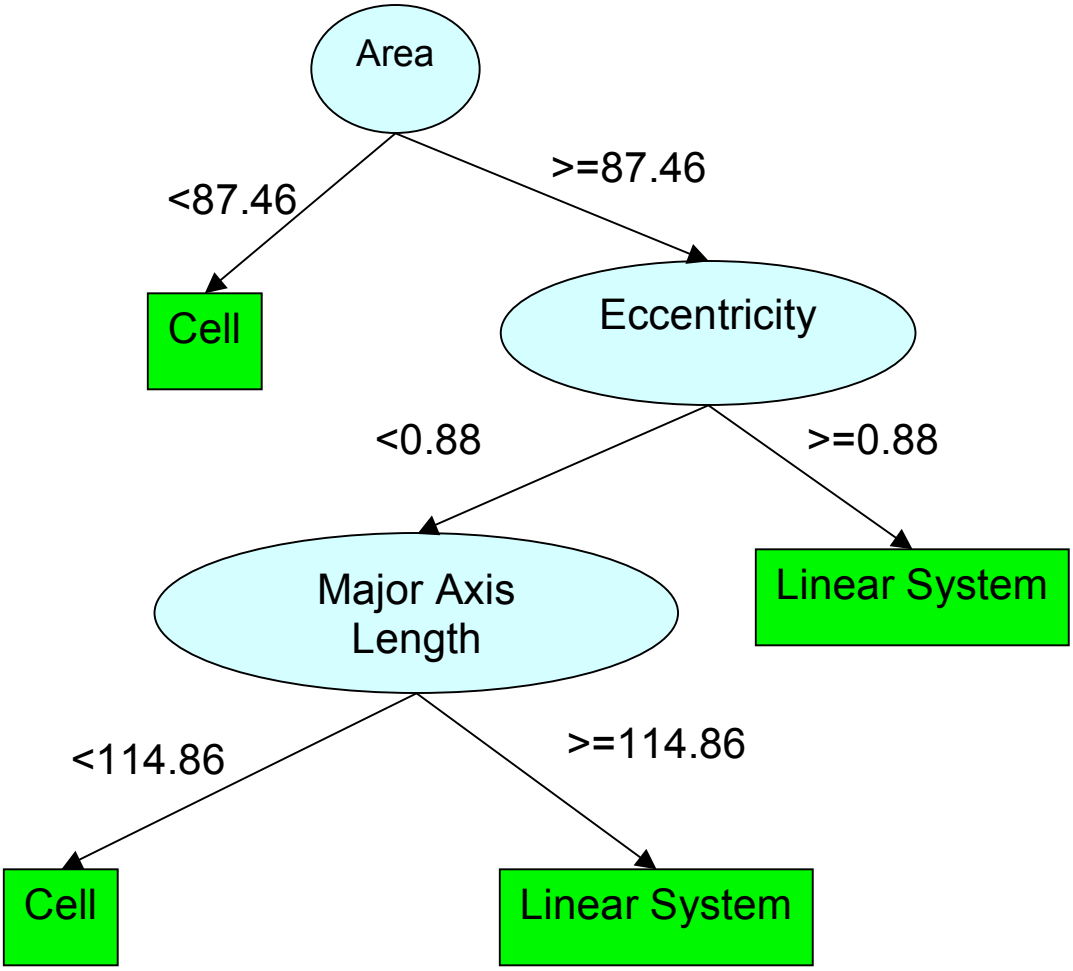

Figure 1. An example of the general type decision tree.

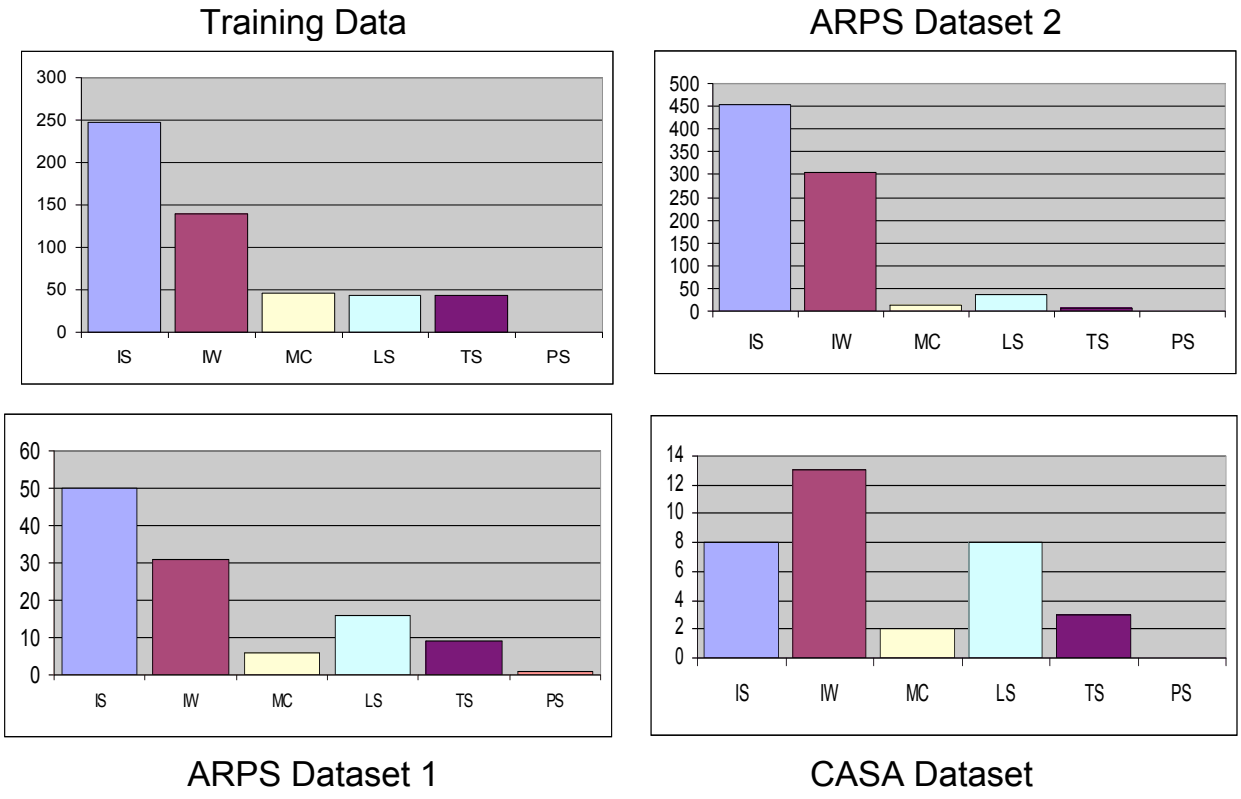

Figure 2. The distributions of the four hand-labeled datasets used to train and test the decision trees. IS = isolated severe, IW = isolated weak, MC = multi-cellular, LS = leading stratiform, TS = trailing stratiform, PS = parallel stratiform.

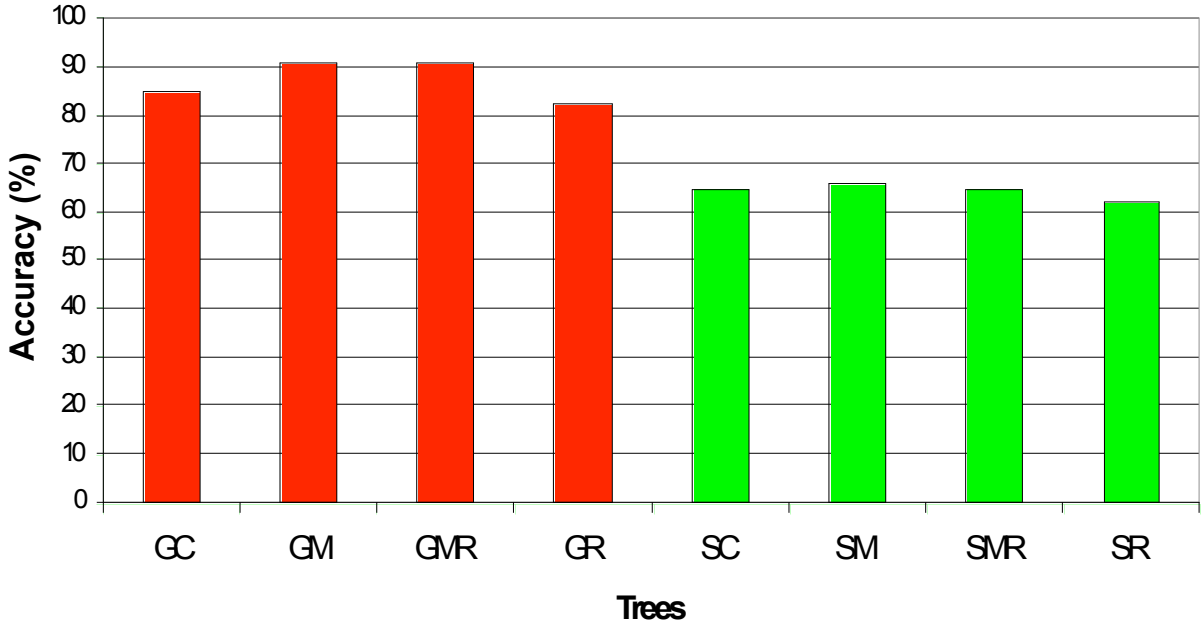

Figure 3. Accuracies of the eight different trees used in the study. The general trees (GC, GM, GMR, and GR) overall have a higher accuracy than the specific trees (SC, SM, SMR, and SR).

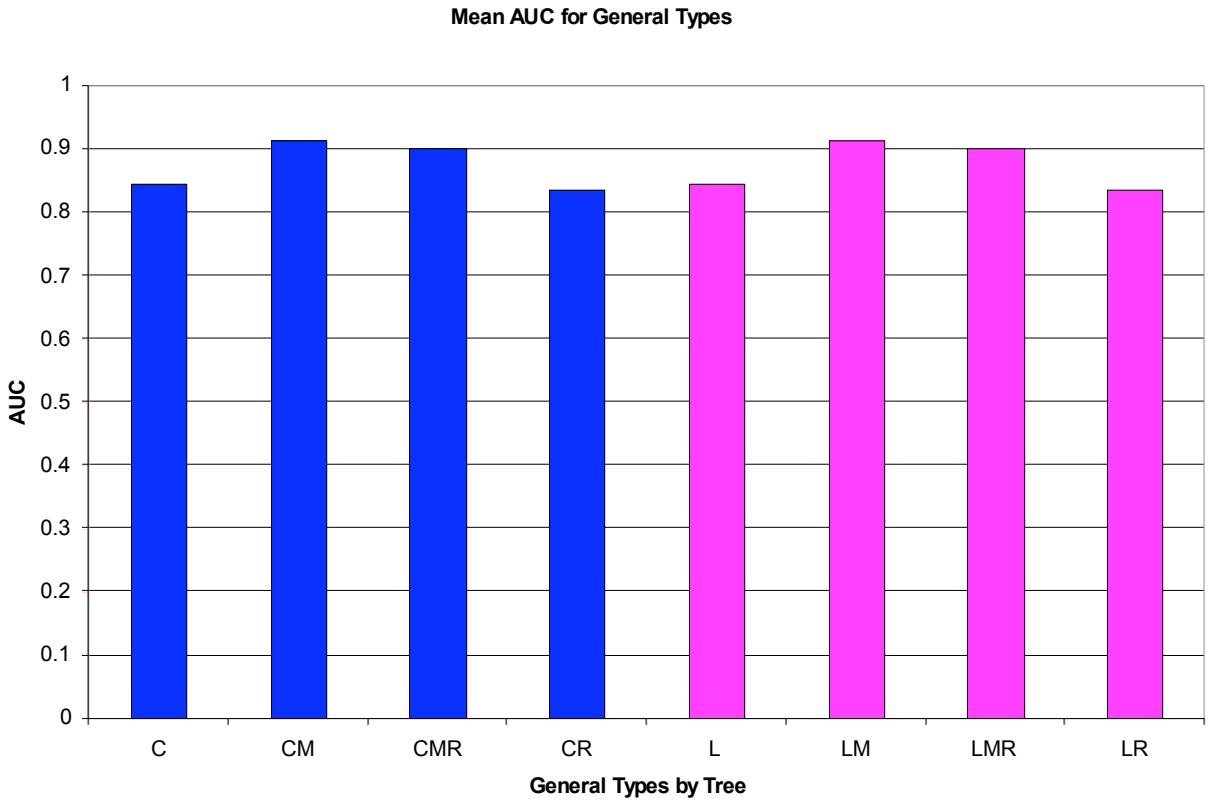

Figure 4. AUC of the general type trees. In both the cell and line types, the trees with morphological attributes (CM, CMR, LSM, LSMR) outperformed the control trees (C and LS) and the reflectivity trees (CR and LSR).

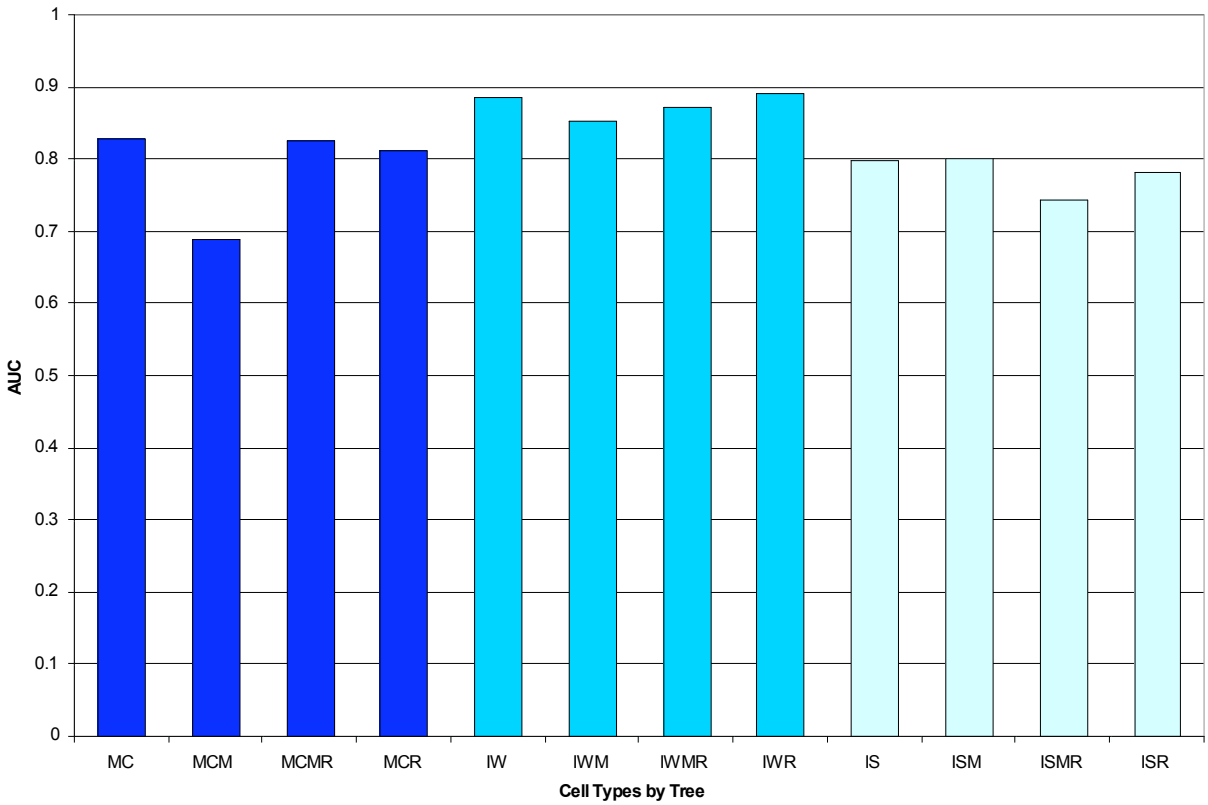

Figure 5. AUC for the specific tree cell types. Little variance exists across the different tree types in each cell type except for the multicell (MC). The multicell morphological tree has a noticeable drop in performance in comparison to the other trees.

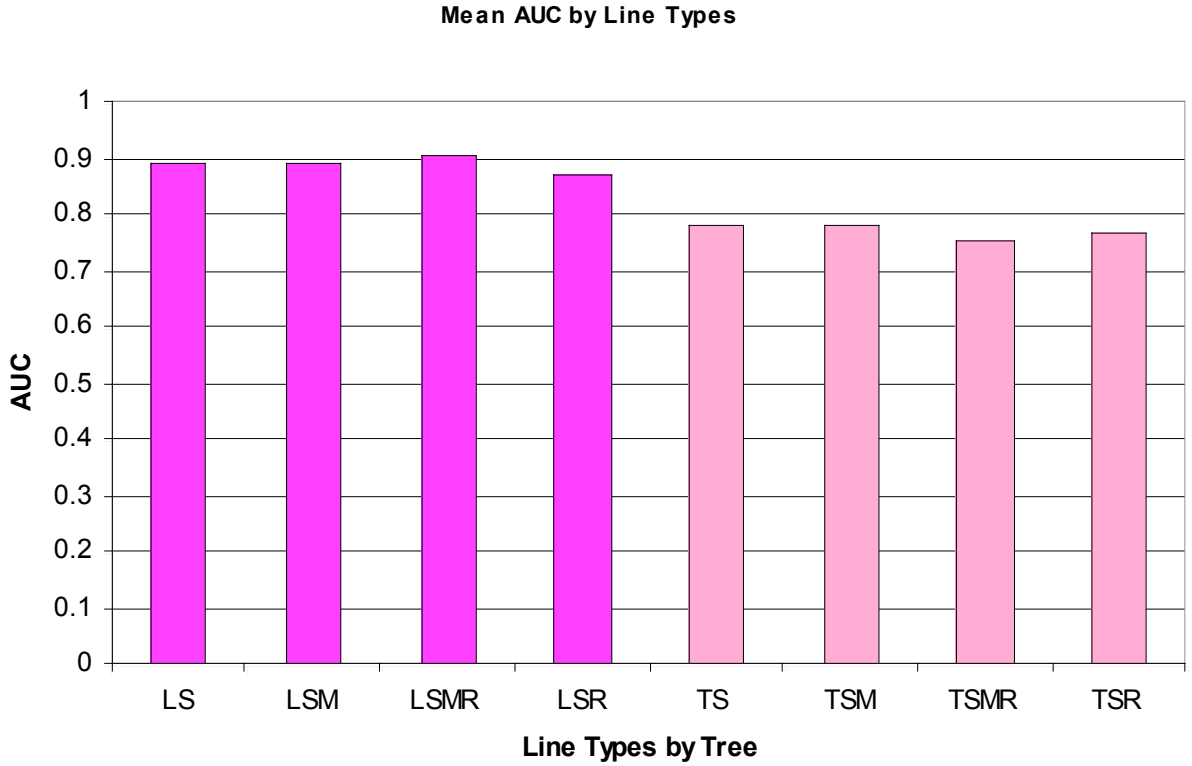

Figure 6. Average AUC for the two linear system types found in the data. There is little variance within each type, indicating that one of the control variables most affects the tree performance on that type.

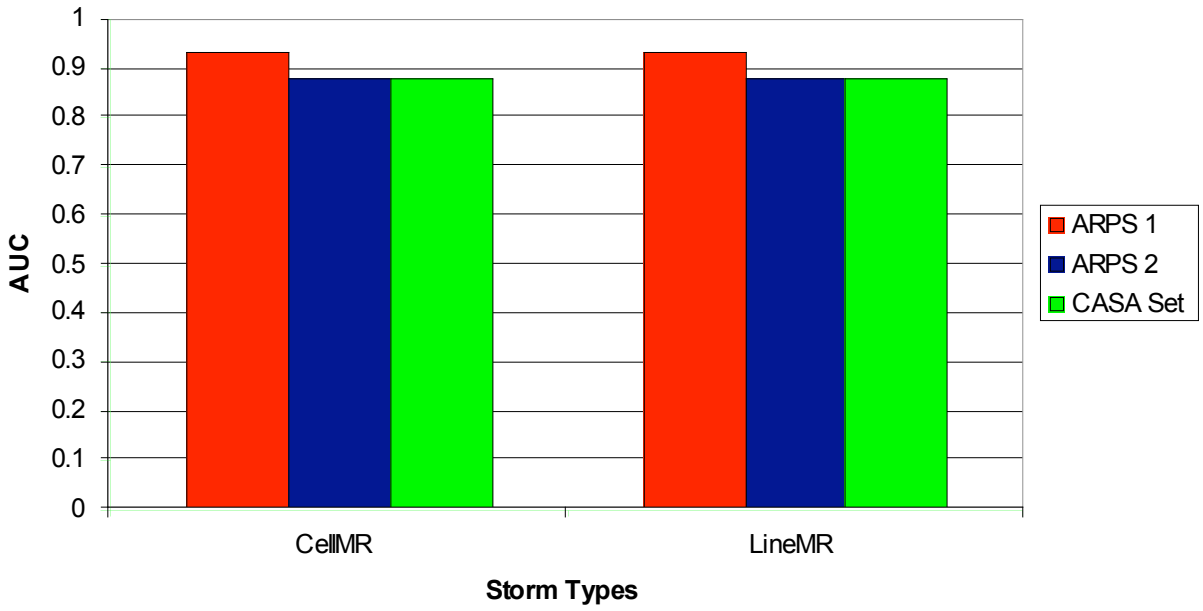

Figure 7. AUC across the test sets for the general types of trees. The AUC remains consistently high across all three datasets, indicating little effect from overfitting and a classifier that can be applied to both simulated and real storms.

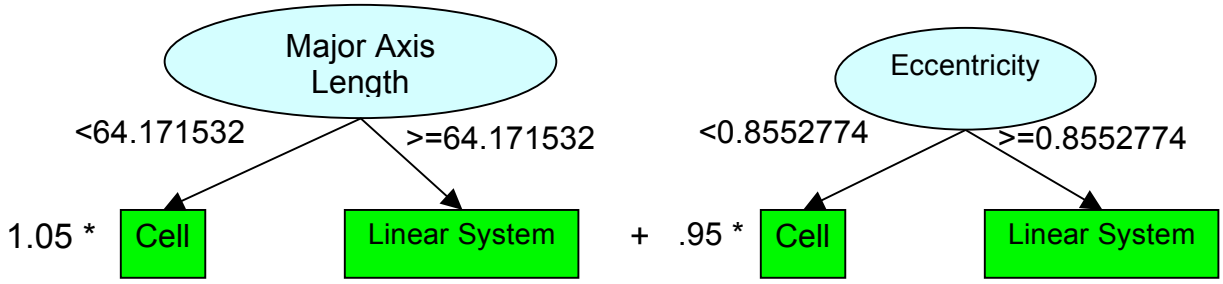

Figure 8. Example of two decision stumps generated by the boosting algorithm. The boosting measures the accuracy of each individual decision stump and weighs them accordingly.

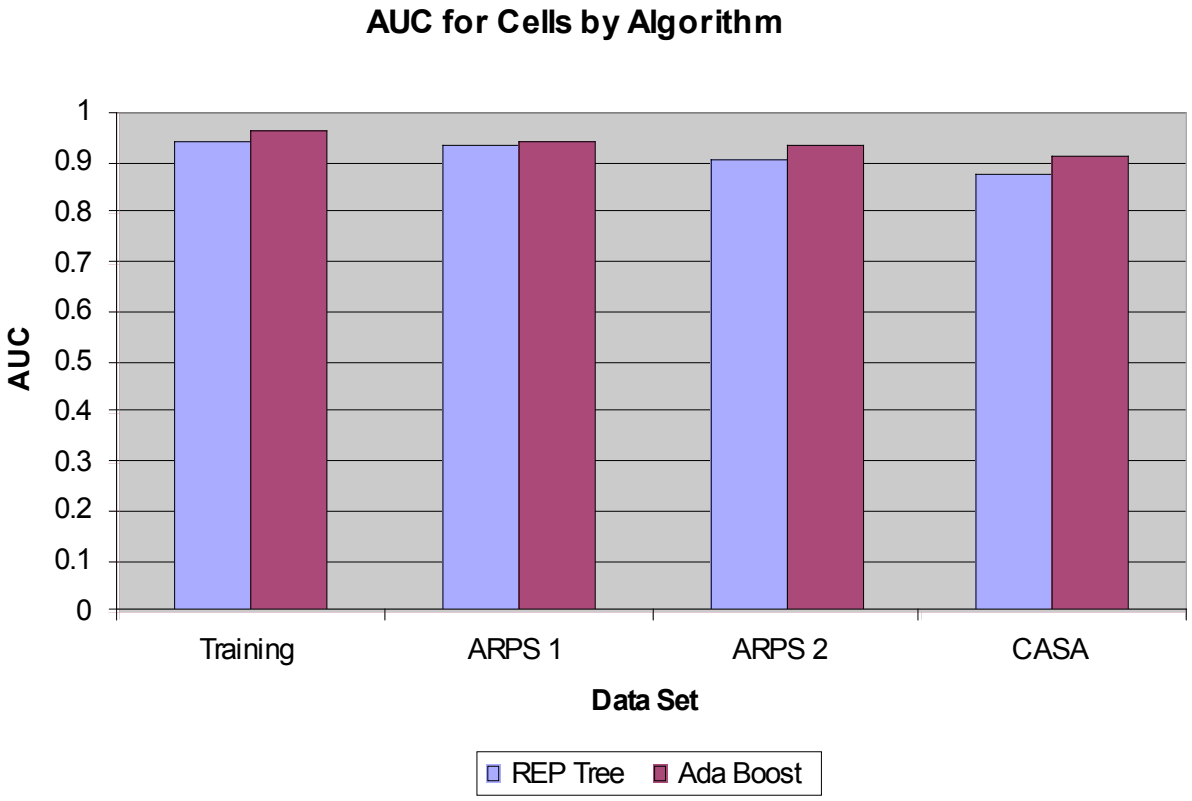

Figure 9. AUC for the cellular types of storms for each dataset. Because of the boosting, AUC increased for every dataset.

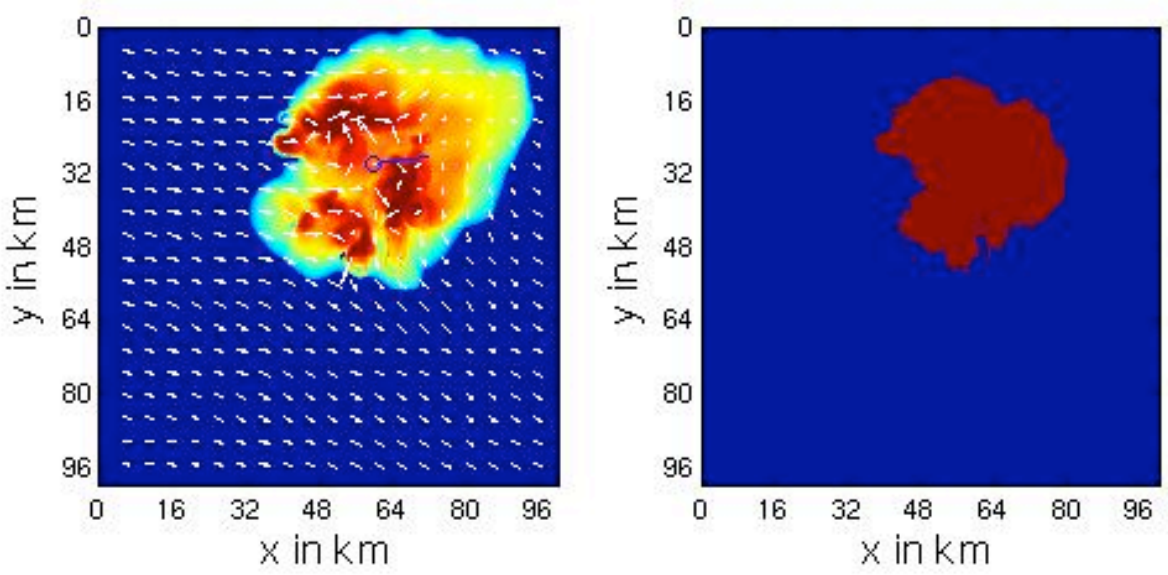

Figure 10. An example of the reflectivity and convective cluster returns from the new classification algorithm. It has the ability to detect storm speed, direction, and the wind field in and around the storm.