Multiscale Storm Identification and Forecast

V Lakshmanan[1], R. Rabin[2], V. DeBrunner[3*]

Abstract

We describe a recently developed hierarchical K-Means clustering method for weather images. that can be employed to identify storms at different scales. We describe an error-minimization technique to identify movement between successive frames of a sequence and show that we can use the K-Means clusters as the minimization kernel. A Kalman filter is used to provide smooth estimates of velocity at a pixel through time. Using this technique in combination with the K-Means clusters, we can identify storm motion at different scales and choose different scales to forecast based on the time scale of interest.

The motion estimator has been applied both to reflectivity data obtained from the National Weather Service Radar (WSR-88D) and to cloud-top infrared temperatures obtained from GOES satellites. We demonstrate results on both these sensors.

1. Introduction

The segmentation of weather imagery is a fundamental problem to automated weather analysis, as has been pointed out in Peak and Tag (1994); Lakshmanan et al. (2000); Johnson et al. (1998). There are numerous pattern recognition algorithms that have been developed on weather images, such as for rainfall estimates (Lai et al. 2000), and cloud classification (Lee et al. 1990), but segmentation techniques for weather images have not been addressed. This is true even of work that attempts to factor out weather-related effects in satellite imagery of land (Markus and Cavalieri 2000; Narasimhan and Nayar 2000).

In the meteorological community, the importance of multiscale segmentation has been often noted (Johnson et al. 1998; Wolfson et al. 1999; Lakshmanan et al. 2000). In (Peak and Tag 1994), the authors detail the difficulties that traditional segmentation algorithms have with satellite weather images because of the textural nature of clouds. As a result, a complex technique consisting of a sequence of fixed thresholds, followed by a neural network that decides how and when to prune or merge the resulting regions

*[1]Corresponding author address: lakshman@ou.edu. V Lakshmanan is at the Cooperative Institute of Mesoscale Meteorological Studies, U. Oklahoma (OU) and is also affiliated with NOAA/OAR/National Severe Storms Laboratory (NSSL). [2]Robert Rabin is with NSSL and the Cooperative Institute of Meteorological Satellite Studies (CIMSS) at the University of Wisconsin, Madison. [3]Victor DeBrunner is with the Department of Electrical and Computer Engineering at OU.

is proposed (Peak and Tag 1994). We show here that using a hierarchical technique in combination with a texture segmentation algorithm makes segmentation of satellite weather images possible such that even small cloud features can be identified.

The textural nature of weather imagery makes robust segmentation for storm tracking purposes very difficult. For storm tracking to be useful, the identification and tracking algorithm should be completely automated. The identification algorithm should not require training, i.e. the algorithm should not expect to see examples of all the "objects" it must identify. Storm "cells" (small scale features) should be capable of being identified. Because the notion of scale is natural in the storm tracking context, we would like to add the requirement that storms at various scales be identified, with their hierarchical structure intact. A multiscale tracking algorithm would be a significant improvement over current tracking schemes which concentrate either on small scales(e.g: (Johnson et al. 1998)) or on large scales (e.g: (Wolfson et al. 1999)).

In the United States, the operational way to identify storms from radar images involves the use of multiple thresholds and counting runs of values above a threshold along a radial. The centroids are then used as a proxy for the storms (Johnson et al. 1998) and tracked either on the basis of proximity to expected position or through a linear programming approach (Dixon 1994). The change in position is extrapolated.

Identification and tracking algorithms for satellite weather imagery have been implemented for mesoscale convective systems (Morel et al. 1997) – where the features of interest (storm anvils colder than 240K) are on the scale of about 10,000 km^2. Our goal, however, is to identify storm scale features, features on the scale of about 10 km^2.

Another technique (French et al. 1992) is to use neural networks to model input reflectivity fields as a set of nodes and to forecast reflectivity locations in the future based on the evolution of the nodes required to model successive frames. The problems with this technique are that it requires training a neural neural network in real-time, and that a trained neural network can not be used to forecast fields which have not been tracked.

A third technique is to use rectangular sub-grids and to find the maximum correlation within a search radius (Rinehart and Garvey 1978; Tuttle and Gall 1999). A modification of this technique is to pre-filter the data so as to track only the larger scales (Wolfson et al. 1999; Lakshmanan 2000). It is also possible to use sub-grids ranging in size from that of the entire image to small (say, 16km x 16km) grids, and to compute motion estimates at each of these scales. Smoothness criteria can be used to constrain these estimates at different scales.

Identifying, matching and extrapolating storm core locations is suitable for small scale storms. The large scale features and cross-correlation technique is suitable for longer forecasts, but with loss of detailed motion estimates. An assumption here is that the storms are of the scale of the sub-grid, not larger. The multiscale estimation is suitable also for large scale forecasts, but with less precise detailed motion estimates.

When used for advection, all the correlation techniques rely on reverse projection, so there needs to be wind speed at the spot where the storm is moving to. Methods rely on correlation estimates of rectangular templates also assume that all pixels within that rectangular template are moving together.

We use a hybrid approach where motion estimates are made for groups of storms (rather than for sub-grids of the image), but at various scales. The motion estimate

for a storm cell is the movement that minimizes the mean-absolute-error between the current frame and corresponding pixels in the previous frame, except that the template is not a rectangular sub-grid of the image, but is instead the actual shape of the storm cell.

Instead of simply matching storm cells across frames, motion estimates are made by finding the best match for the storm-template. Thus, the major steps in the technique are:

1. Find storms at different scales.
2. Estimate motion at the various scales.
3. Forecast for different periods using motion at different scales.

Because the motion estimates are made for storms, it is possible to interpolate between storm boundaries to obtain motion estimates at every part of the domain.

2. Hierarchical texture segmentation to identify storms

A K-Means clustering technique from Lakshmanan (2001); Lakshmanan et al. (2002) is used to identify components in vector fields. The technique provides nested partitions, i.e. the identified storms structures are strictly hierarchical. The technique works by clustering image values (reflectivity/infrared temperature, etc.) in the neighborhood of a pixel on two opposing criteria:

- Belong to same cluster as your neighbors.
- Belong to cluster whose mean is closest to your value.

Hierarchical segmentation is incorporated into the K-Means clustering technique by steadily relaxing inter-cluster distances.

The technique works by iteratively moving pixels between clusters minimizing

$$E(k) = \lambda d_m(k) + (1-\lambda) d_c(k) \qquad 0 \leq \lambda \leq 1 \tag{1}$$

where the distance in the measurement space is:

$$d_m(k) = \parallel \mu_k^n - T_{xy} \parallel \tag{2}$$

and the discontiguity measure is::

$$d_c(k) = \sum_{ij \epsilon N_{xy}} (1 - \delta(S_{ij}^n - k)) \tag{3}$$

A region growing algorithm is employed to build a set of connected regions, where each region consists of 8-connected pixels that belong to the same K-Means cluster. If a connected region is too small, then its cluster mean (the mean of the texture vectors at each pixel in the region) is compared to the cluster means of the adjoining regions and the small region is merged with the closest mean. The result of the K-Means

segmentation, region growing and region merge steps is the most detailed segmentation of the image.

The inter-cluster distances of all adjacent clusters (or regions) in the image are computed. A threshold is set such that half the pairs fall below this threshold. If a pair of clusters differs by less than this threshold, the clusters are merged and cluster means updated. This process is continued until no two adjacent regions are closer in cluster space than the threshold. When this process is complete, we have the next coarser scale of the segmentation. This process is repeated until no changes happen.

a. Weather Radar Images

Texture segmentation using Markov Random Field (MRF) models has been utilized to segment synthetic aperture radar (SAR) images, mainly because SAR images are characterized by a lot of speckle (Dong et al. 2001; Schroder et al. 1998; Smits and Dellepine 1999), a problem which is resolved through the use of neighborhood statistics. Another reason for using texture segmentation on SAR imagery is that the same MRF model used for segmentation can also be used for classifying the identified segments (Dong et al. 2001).

Texture segmentation has not been applied to weather radar data before. In weather radar data, especially in cases where there is significant precipitation, the problem of speckle does not arise except in the immediate vicinity of the radar. Hence, traditional texture segmentation provides no significant advantage. In fact, as shown in Figure 1e, even a scalar segmentation approach works quite well. What neither the scalar segmentation approaches, for example (Johnson et al. 1998), nor standard texture segmentation approaches (Blum and Rosenblat 1972; Hofmann et al. 1996; Ma and Manjunath 1997) can provide is a nested partition of identified segments. The watershed segmentation approach of Najman and Schmitt (1996) can provide a nested partition, but does not segment weather data well (See Figure 1f). As shown in Lakshmanan et al. (2002); Lakshmanan (2001), multiscale segmentation can be achieved by agglomerative K-Means clustering of texture vectors and slow relaxation of the allowed inter-cluster distance.

We wish to segment the reflectivity moment of radar elevation scans obtained from a Doppler Weather Service Radar (WSR-88D). The data have been mapped from polar coordinates into a Cartesian grid tangential to the earth's surface at the radar location where each pixel is a square area of one kilometer on each side. The pixel values, in dBZ, range from about $-7dBZ$ to about $64dBZ$, with the reflectivity values for some pixels missing. Missing values and all reflectivity values less than $0dBZ$ were thresholded to be $0dBZ$ before the segmentation process.

The radar elevations scans in this study were collected every 5-6 minutes. The weather surveillance radars used by the National Weather Service scan through thunderstorms starting at a low elevation angle, 0.5^o for Volume Coverage Pattern (VCP) 21, and after completing a full 360^o azimuthal sweep, progressively increase the elevation angle until an upper limit is reached (19.5^o in VCP 21). See Figure 2 (Crum and Alberty 1993; Smith 1995). The data were remapped to a Cartesian plane and were then segmented using the K-Means clustering technique.

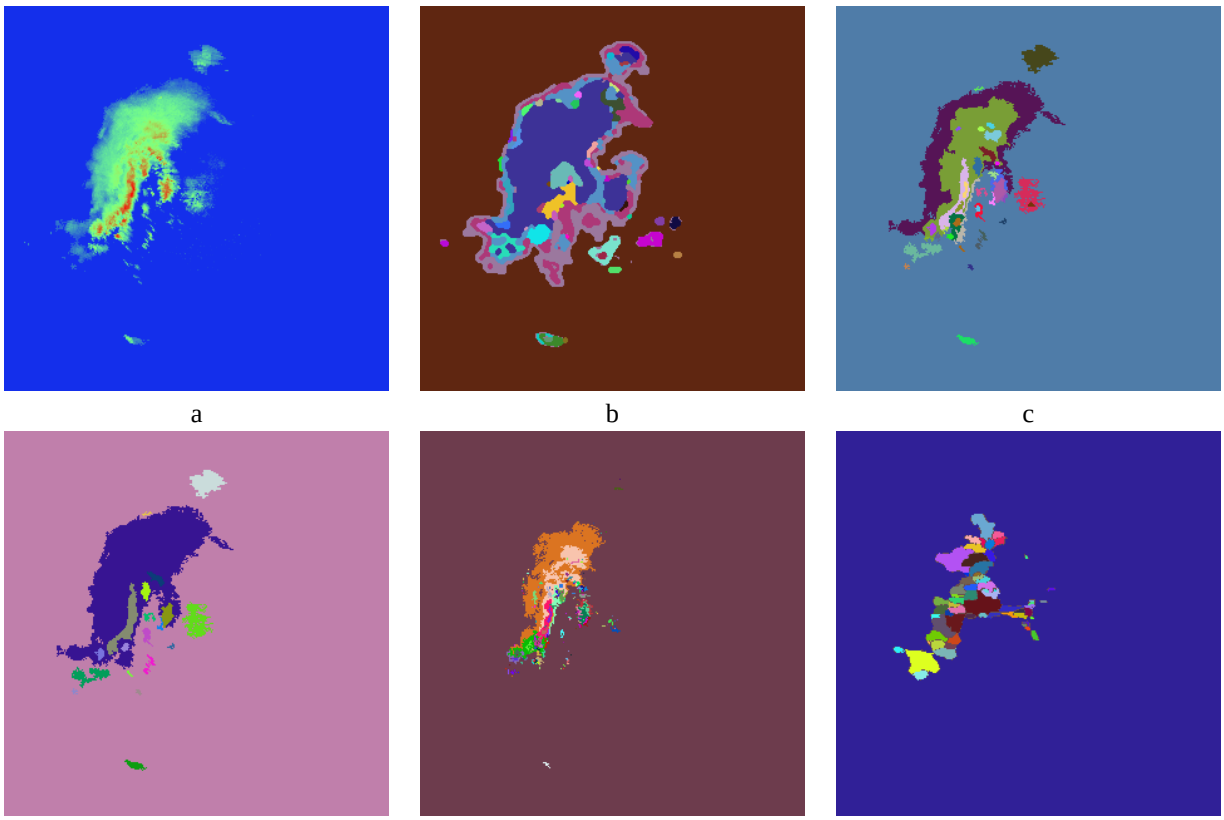

Figure 1: Segmenting a radar reflectivity image. (a) A radar reflectivity image, from Fort Worth May 5, 1995. (b) The result of segmenting the radar reflectivity image using the Markov Random Field (MRF) approach of Blum and Rosenblat (1972). (c) The result of segmenting the image using the method of this paper, tweaked to process the reflectivity range of interest. The most detailed scale is shown. (d) The next higher scale of segmentation using the method of this paper. (e) Simply separating the image into contiguous bands of $10dBZ$. (f) Using the watershed approach of Najman and Schmitt (1996).

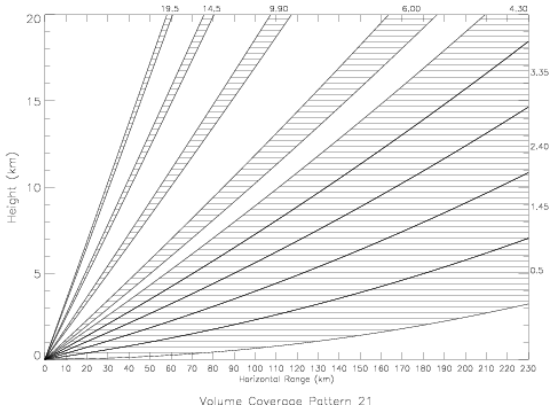

Figure 2: Volume Coverage Pattern (VCP) 21 of the WSR-88D, a weather surveillance radar used by the National Weather Service. The volume coverage is shown. The beamwidth is 0.95 degrees and there are 9 elevation scans in this VCP. Figure from Smith (1995).

b. Satellite Infrared Images

We demonstrate results of segmenting the infrared window channel (11μ) of GOES satellite imager data. The images are 200x300 with each pixel representing a $4km \times 4km$. The images are projected onto a plane tangential to the surface of the earth. The satellite data were collected over the continental United States using GOES-11 on March 29, 1998. The pixel values were also mapped from radiance values to equivalent black body temperature in degrees Kelvin before the segmentation. Images of the sequence are available at eight minute intervals.

The sequence of satellite images captures a day of significant thunderstorm activity. Several thunderstorms grow and decay during the day. The temperatures and sizes of the cloud tops in the images show relate in a bulk sense to the magnitude and extent of clusters of storm updrafts within the anvil clouds. Segmentation of this sequence should be able to consistently identify the thunderstorm cloud tops in the images. Ideally, when cloud tops appear to split or merge, the corresponding segmented regions should do the same. A very important requirement is that small changes in the storm structure should be reflected as small changes in the segmented region corresponding to the storm.

Studies (Browning 1979; Bellon and Zawadzki 1994) have shown that a single storm cell grows and decays in under an hour. Therefore, a storm cell can be expected to stay for no more than seven frames of the satellite sequence. However, a line of thunderstorms within which these cells crop up can be expected (Browning 1979) to persist for up to six hours. Also, the cloud top (anvil) may persist after the cell on radar decays. Therefore, the segmentation should lend itself to segmenting regions corresponding to larger scale features while identifying small scale features that are contained within the large scale feature but have shorter life-times.

A single infrared image was segmented using various segmentation methods in the literature. The results are shown in Figure 3c and d. The results of segmentation using the other approaches (Figure 3b,e and f. are poor in terms of the scale of the resulting regions. This is not surprising because the infrared satellite weather imagery has several characteristics that make it hard to segment: very low dynamic range (from about 225K to 240K) for the regions of interest, poor resolution as compared to the scale of the phenomena of interest, and high pixel value variance, even in the absence of edges. It is instructive to compare the poor performance of these algorithms on the satellite image (see Figure 3) with the performance of the same algorithms on radar reflectivity images in Section a.

The poor spatial resolution of the satellite image affects our algorithm also, in the scale of features that we can detect. Although we can detect features as small as 10 pixels in the image, this translates to about 40 km^2, a mid-size storm cell (although significantly more detailed than what could be obtained using earlier approaches). The pruning threshold of 10 pixels was set in the algorithm so that any statistics collected are somewhat reliable. One possible way to relax this threshold is by creating a pseudo-high resolution form of the original image, thus getting less square kilometers in the 10-pixel threshold. Unfortunately, on satellite weather images, even a pseudo-high resolution technique (Yao 1999) introduces unacceptable smoothing (Lakshmanan 2001), resulting in worse performance. A second possibility, one that we have not yet looked into because of the prohibitive cost for a continuously running system, is to obtain weather satellite data that has higher spatial resolution. A third possibility is to use the multi-channel nature of satellite weather information to form the pixel representation (instead of using a texture vector based on neighborhood statistics).

Instead of using only texture measurements from only the infrared channel, we used texture measurements (mean and variance) computed on four channels corresponding to 3.9, 6.7, 11 and 12 microns (near infrared, water vapor, window and "dirty window" respectively (Menzel and Purdom 1994)). Since every pixel of the segmented output actually corresponds to four relatively independent measurements (rather than just one), the minimum pruning size in the algorithm can be reduced from about 10 pixels to about 3. The result of using multi-channel information and a lower size threshold is shown in Figure 5 where it is compared to the segmented result if only the 11 micron image had been used.

Notice that the result of segmenting using all four channels (Figure 5f) has smaller regions than the result that uses only the infrared window channel. It is not clear, however, how significant these smaller features are in the context of thunderstorms.

3. Motion Estimation

Once the storms have been identified from the images, these storms are used as a template and the movement that minimizes the absolute-error between two frames is computed. For radar images, we used consecutive (5-6 min) volume scans.

Motion estimation is done by moving a template of the identified region at the appropriate scale around in the previous image. A matrix of mean absolute error at the different positions is obtained as shown in Figure 6

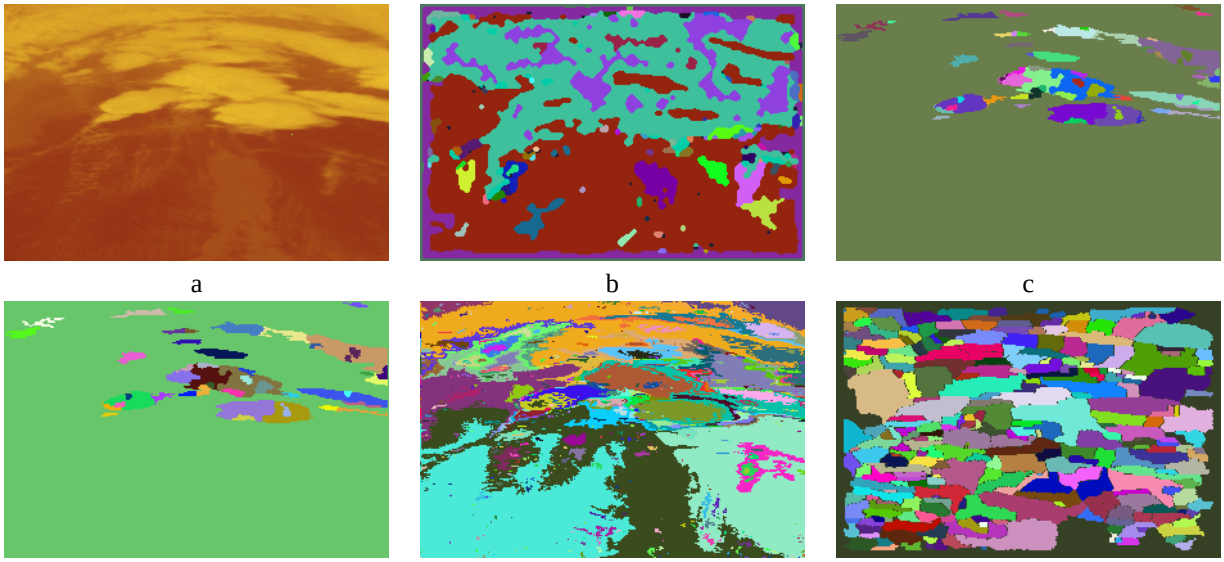

Figure 3: Segmenting an infrared satellite weather image. (a) The infrared image being segmented. Notice the various storms at the top of the image. The darker areas in the bottom correspond to ground. (b) The result of segmenting the image using the Markov Random Field (MRF) approach of Blum and Rosenblat (1972). There is no detail – it is effectively a binary segmentation. (c) The result of segmenting the image using the method of this paper (the most detailed scale). Notice the fine detail within the clouds. (d) The next higher scale of segmentation using the method of this paper. The strong storm cells being significantly colder are retained – the large cloud masses are merged. (e) Simply separating the image into contiguous bands of $1Kelvin$. There is a lot of detail, but no organization. This is what you get using hierarchical thresholds. (f) Using the watershed segmentation approach of Najman and Schmitt (1996). Because of the textural nature of the data, the watershed algorithm has very poor performance.

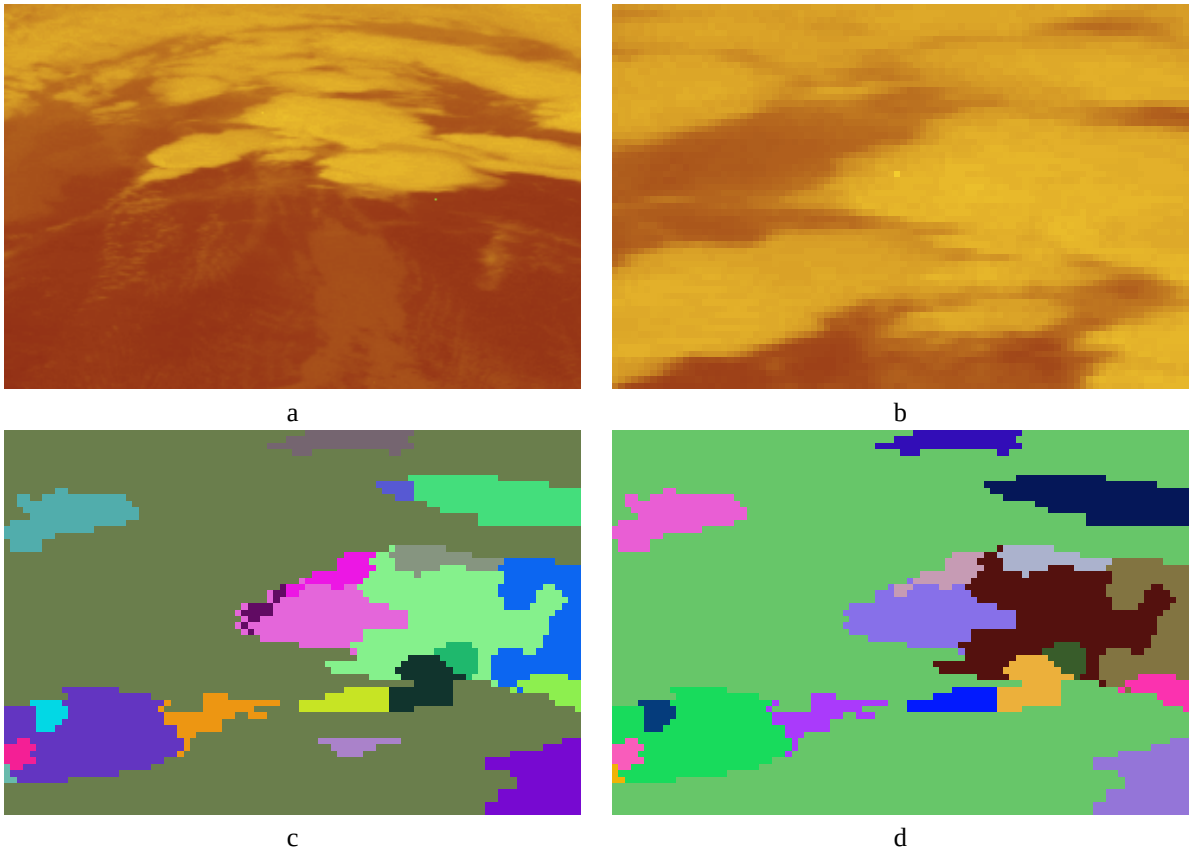

Figure 4: A close-in look at the results shown in Figure 3.. (a) The infrared image being segmented (same as Figure 3a). (b) A close-in look at the input satellite infrared image of (a). (c) A close-in look at the result of segmenting the image using the method of this paper (the most detailed scale). Notice the fine detail within the clouds. (d) A close-in look at the next higher scale of segmentation using the method of this paper.

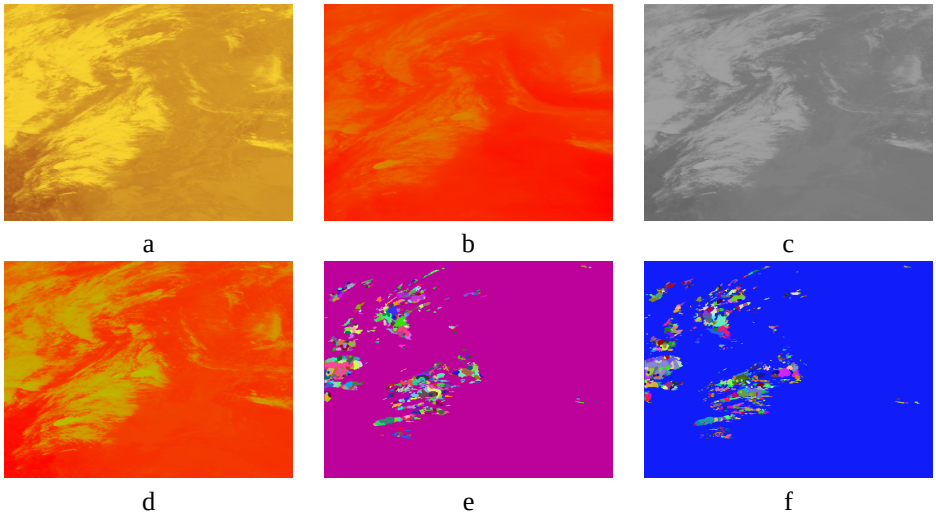

Figure 5: Using multi-channel satellite data for segmentation. (a) 3.9 micron infrared (b) 6.7 micron water vapor (c) 11 micron window (d) 12 micron dirty window channels of data. (e) Most detailed segmentation using only the 11 micron image (f) Most detailed segmentation using all four channels. The segmentation is more detailed than the segmentation that was achieved in (e), but whether these extra details are useful is yet to be determined.

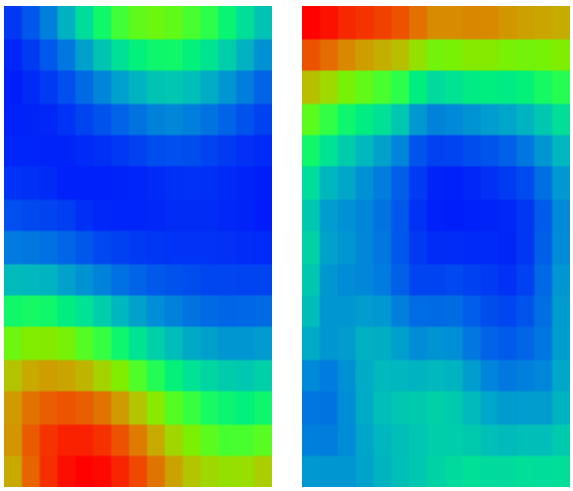

Figure 6: Matrix of mean absolute error by position. Larger errors are “warmer”. Two different locations are shown.

Instead of simply finding the absolute minimum, a smoother minimum of the absolute error field is sought. The field of absolute errors is minimized by weighting each value by how much it differs from the absolute minimum and finding the centroid.

For each storm template, we also get a growth/decay estimate. This is based on how much the average value inside the template changes based on the template at the best match.

Given the motion estimates for each of the regions in the image, the motion estimate at each pixel is determined through interpolation. At the pixel xy, the motion estimate u_{xy} is given by

$$u_{xy} = \frac{\Sigma_i u_i w_{ixy}}{\Sigma_i w_{ixy}} \tag{4}$$

where u_i is the motion estimate for the i^{th} region and the weight of the this estimate at the point xy is given by:

$$w_{ixy} = \frac{N_i}{\parallel xy - c_i \parallel^2} \tag{5}$$

N_i is the number of pixels in the i^{th} region, c_i its centroid and $\parallel$ denotes the Euclidean distance between the two points.

This motion estimate is for the pair of frames that were used in the comparison. We do temporal smoothing of these estimates by running a Kalman filter Kalman (1960) at each pixel of the motion estimate. The Kalman estimator is built around a constant acceleration model with the standard Kalman update equations Brown and Hwang (1997).

a. Short-term Forecast

The forecast of the fields is done based on the motion estimates, growth and decay heuristic and the current data. Forecasts can be made on fields other than the tracked field. For example, motion estimates can be derived from VIL and applied to radar reflectivity and probability fields of lightning and hail.

The forecast is done by first project data forward in time to a spatial location given by the motion estimate at their current location and the elapsed time. Locations not filled by this forward projection are filled by interpolating using an inverse square-distance metric of nearby filled locations.

The skill of this technique is quantititavely measured by comparing, for example, the 30 minute forecast against the actual field closest to 30 minutes ahead. For the mean absolute error results, the actual values are used. For the critical success index (CSI) results, the best match with a 5x5 window is used. Comparisions are made with a plain persistence, and with motion estimates derived by minimizing the correlation of a 5x5 template between the frames.

Results over a 60minute period on reflectivity from the Fort Worth radar on April 201995 are shown in Figures 7 and 8.

The CSI and MAE measure different aspects of the forecast accuracy. The MAE takes into account actual reflectivity values and is, therefore, a measure of how good the growth-and-decay aspect is. The CSI is a measure of predicting storm location. We are good at predicting storm location, but not so good at growth/decay.

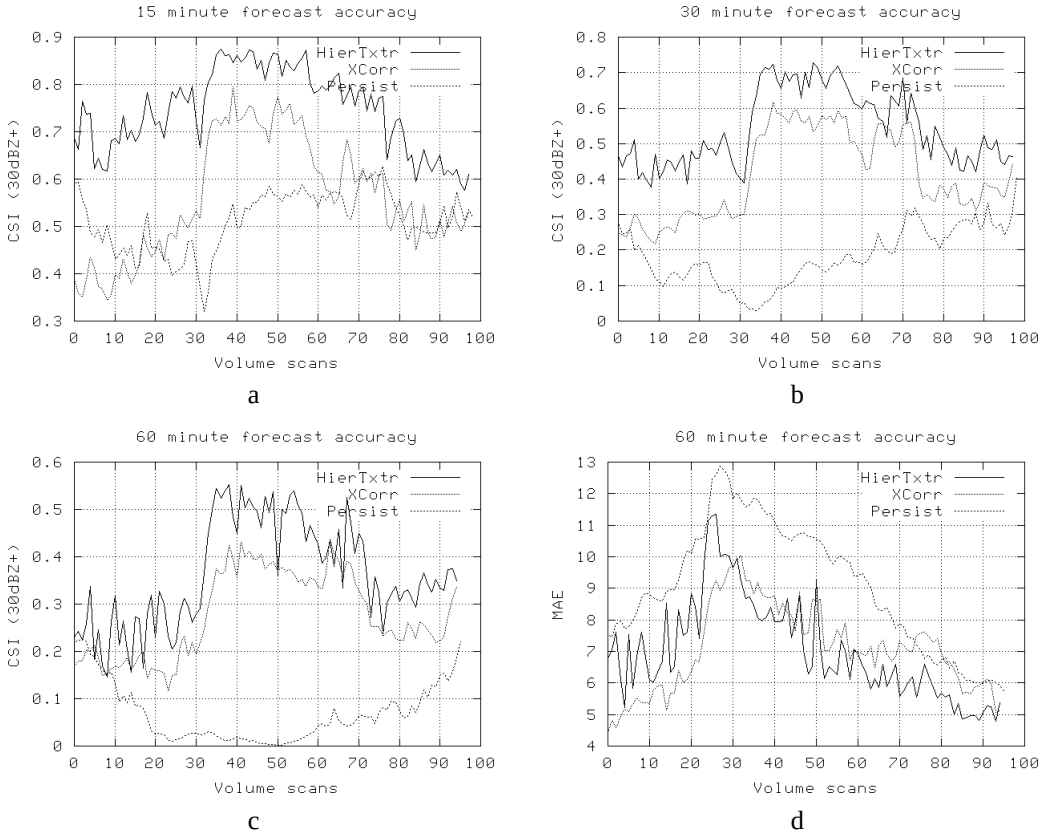

Figure 7: Skill at forecasting a radar reflectivity field compared to a persistence forecast and to a local correlation approach. (a) Values 30dBZ and above for 15 minutes (b) Values 30dBZ and above for 30 minutes (c) Values 30dBZ and above for 60 minutes (d) Mean absolute error in 60 minute forecast

A forecast based on satellite infrared temperature is shown in Figure 9. The data are taken from GOES-12 imagery on Oct. 9, 2001. The data provided were 100 seconds apart; we used every 4th frame of the sequence to compute motion estimates. Work to compute skill scores on satellite data is underway.

A version of this paper with color illustrations is available online at

`http://www.cimms.ou.edu/~lakshman/Papers/kmeans_motion.pdf`

4. Conclusions

It is possible to use a K-Means clustering to provide hierarchical identification of storms. The clusters can then be used to estimate the movement of the storm cores. A forecast that projects the movement of the storm cores linearly possesses some skill.

5. 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.

References

Bellon, A. and I. Zawadzki: 1994, Forecasting of hourly accumulations of precipitation by optimal extrapolation of radar maps. *Journal of Hydrology*, **157**, 211–233.

Blum, J. and J. Rosenblat: 1972, *Probability and Statistics*. W.B. Saunders Company, 549 pp.

Brown, R. and P. Hwang: 1997, *Introduction to Random Signals and Applied Kalman Filtering*. John Wiley and Sons.

Browning, K.: 1979, The FRONTIERS plan: A strategy for using radar and satellite imagery for very-short-range precipitation forecasting. *The Meteorological Magazine*, **108**, 161–184.

Crum, T. and R. Alberty: 1993, The WSR-88D and the WSR-88D operational support facility. *Bulletin of the American Meteorological Society*, **74**, 1669–1687.

Dixon, M.: 1994, *Automated Storm Identification, Tracking and Forecasting – A Radar-Based Method*. Ph.D. thesis, University of Colorado and National Center for Atmospheric Research.

Dong, Y., A. Milne, and B. Forster: 2001, Segmentation and classification of vegetated areas using polarimetric sar image data. *IEEE Trans. Geosci. Remote Sensing*, **39**, 321–329.

French, M., W. Krajewski, and R. Cuykendall: 1992, Rainfall forecasting in space and time using a neural network. *J. Hydrology*, **137**, 1–31.

Hofmann, T., J. Puzicha, and J. Buhmann: 1996, A deterministic annealing framework for unsupervised texture segmentation. Technical Report IAI-TR-96-2, Institut fr Informatik III, U. Bonn, http://www-dbv.cs.uni-bonn.de/image/example4.html.

Johnson, J., P. Mackeen, A. Witt, E. Mitchell, G. Stumpf, M. Eilts, and K. Thomas: 1998, The storm cell identification and tracking algorithm: An enhanced WSR-88D algorithm. *Weather and Forecasting*, **13**, 263–276.

Kalman, R.: 1960, A new approach to linear filtering and prediction problems. *Trans. ASME – J. Basic Engr.*, 35–45.

Lai, E., P. Li, C. Chan, M. Chu, and W. Wong: 2000, Pattern recognition of radar echoes for short-range rainfall forecast. *15th Intl. Conf. on Pattern Recog.*, IEEE, volume 4, 299–302.

Lakshmanan, V.: 2000, Speeding up a large scale filter. *Journal of Oceanographic and Atmospheric Technology*, **17**, 468–473.

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

Lakshmanan, V., V. DeBrunner, and R. Rabin: 2002, Nested partitions using texture segmentation. *Southwest Symposium on Image Analysis and Interpretation*, IEEE, Santa Fe, New Mexico.

Lakshmanan, V., R. Rabin, and V. DeBrunner: 2000, Identifying and tracking storms in satellite images. *Second Artificial Intelligence Conference*, American Meteorological Society, Long Beach, CA, 90–95.

Lee, J., R. Weger, S. Sengupta, and R. Welch: 1990, A neural network approach to cloud classification. *IEEE Trans. on Geoscience and Remote Sensing*, **28**, 846–855.

Ma, W. and B. Manjunath: 1997, Edge flow: a framework of boundary detection and image segmentation. *Proc. IEEE International Conference on Computer Vision and Pattern Recognition*, San Juan, Puerto Rico, 744–749.

Markus, T. and D. Cavalieri: 2000, An enhancement of the NASA team sea ice algorithm. *IEEE Trans. on Geoscience and Remote Sensing*, **38**, 1387–1398.

Menzel, P. and J. Purdom: 1994, The first of a new generation of geostationary operational environmental satellites. *Bull. Amer. Meteo. Soc.*, **75**, 757–782.

Morel, C., F. Orain, and S. Senesi: 1997, Automated detection and characterization of MCS using the meteosat infrared channel. *Proc. Meteo. Satellite Data Users Conf.*, Brussels, 213–220.

Najman, L. and M. Schmitt: 1996, Geodesic saliency of watershed contours and hierarchical segmentation. *IEEE Trans. Patt. Anal. and Mach. Intell.*, **18**, 1163–1173.

Narasimhan, S. and S. Nayar: 2000, Chromatic framework for vision in bad weather. *IEEE Conf. on Comp. Vision and Patt. Recog.*, IEEE, volume 1, 598–605.

Peak, J. and P. Tag: 1994, Segmentation of satellite weather imagery using hierarchical thresholding and neural networks. *Journal of Applied Meteorology*, **33**, 605–616.

Rinehart, R. and E. Garvey: 1978, Three-dimensional storm motion detection by conventional weather radar. *Nature*, **273**, 287–289.

Schroder, M., H. Rehrauer, K. Seidel, and M. Dutcu: 1998, Spatial information retrieval from remote-sensing images – part ii: Gibbs-markov random fields. *IEEE Trans. Geosci. Remote Sensing*, **36**, 1446–1455.

Smith, T.: 1995, Visualization of WSR-88D data in 3d using application visualization software. *14th Conf. on Weather Forecasting*.

Smits, P. and S. Dellepine: 1999, Discontinuity-adaptive markov random field model for the segmetnation of intensity sar images. *IEEE Trans. Geosci. Remote Sensing*, **37**, 627–631.

Tuttle, J. and R. Gall: 1999, A single-radar technique for estimating the winds in tropical cyclones. *Bull. Amer. Met. Soc.*, **80**.

Wolfson, M., B. Forman, R. Hallowell, and M. Moore: 1999, The growth and decay storm tracker. *8th Conference on Aviation*, Amer. Meteor. Soc., Dallas, TX, 58–62.

Yao, M.: 1999, *Model-Based Methods for Image Interpolation and Enhancement That Retain Edge Information*. Ph.D. thesis, U. Oklahoma.

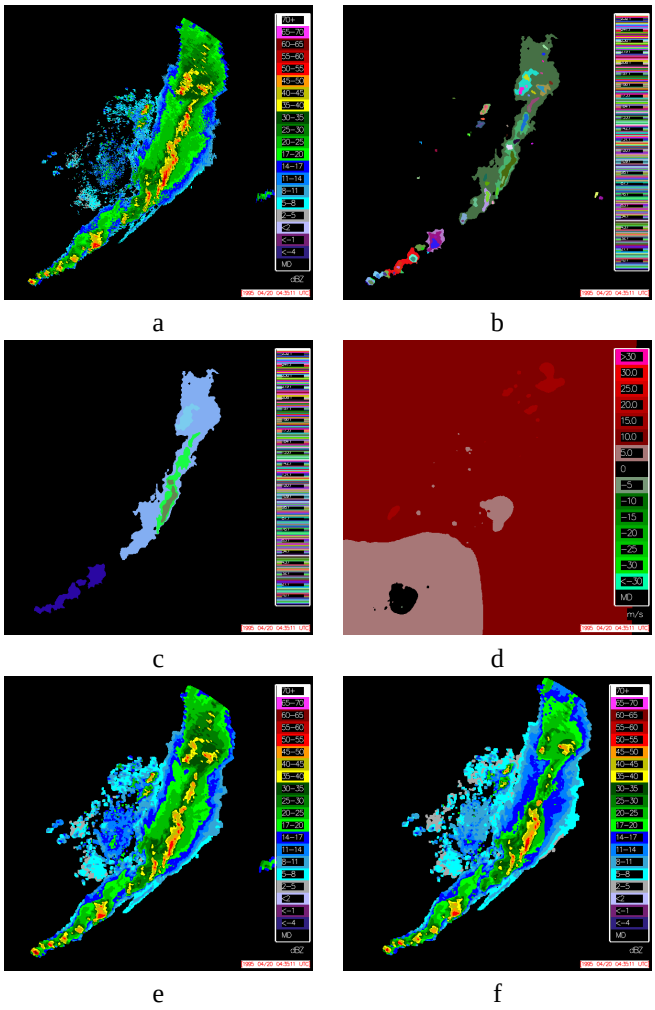

Figure 8: (a) Reflectivity data from KFWS, April 1995. (b) Most detailed scale of segmentation, used in forecasting under 30 minutes. (c) Coarse segmentation, used in forecasting more than 90 minutes. (d) Motion estimate (red is eastward motion) (e) 15min forecast (f) 60min forecast

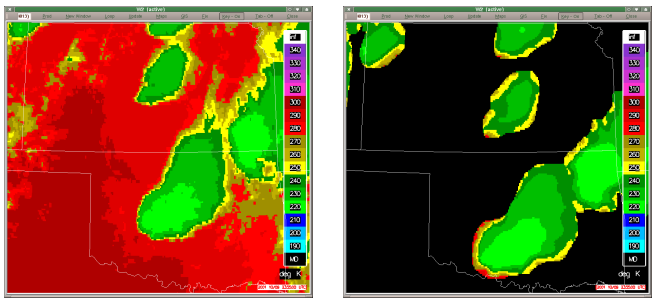

Figure 9: The original (left) and a 30 minute forecast of infrared temperature from Oct. 9, 2001.