

TEXTURE-BASED SEGMENTATION OF SATELLITE WEATHER IMAGERY

V. Lakshmanan^{1,2}, *V. DeBrunner*², *R. Rabin*^{1,3}

¹National Severe Storms Laboratory, ²U. of Oklahoma, ³U. of Wisconsin, Madison

ABSTRACT

Unsupervised segmentation of weather images into features that correspond to physical storms is a fundamental and difficult problem. Treating an infrared satellite image as a Markov random field, the Kolmogorov-Smirnov distance between the local distribution of spatial statistics and the global statistics of classified regions is used to segment the image using a relaxation algorithm. An outlier class is utilized to capture as yet unclassified pixels. We demonstrate the results of different initialization methods on the final segmentation and point out where the method is deficient.

1. INTRODUCTION

Segmenting weather images into regions, each of which corresponds to a physical thunderstorm or front, is a fundamental problem in meteorological algorithm development. The most common technique [1] is to identify contiguous pixels meeting one of several thresholds and to grow regions corresponding to each threshold. This technique works well one frame at a time. However, it can not be used to yield robust segmentation of real-world time sequences of weather images since small oscillations in the measured quantities can cause large changes in the shape and size of the identified regions [2]. Algorithms based on contouring of a smoothed data field such as [2] perform better but have to do special processing to keep track of storms that slowly grow or decay and in any case, these heuristic rules are often insufficient. The product of segmentation is used by many meteorological algorithms and problems in the segmentation of time-sequences of weather images are often reflected in the performance of these algorithms also.

In this paper, we treat each weather image as a Markov random field with respect to an eight-neighbor system. A physical thunderstorm is modeled as having a probability distribution function in several texture characteristics that is different from that of the surrounding atmosphere and markedly different from that of the ground as seen in the satellite imagery. This is a valid assumption – the variance, for example, of the temperature estimate in infrared images is very low for images of ground and markedly higher for parts of the image representing thunderstorms.

The rest of this paper is organized as follows. The qualities of the data that make segmentation such a difficult problem and the texture characteristics of the data that are used in the segmentation are described in Section 2. Setting up the segmentation of infrared weather imagery as a Kolmogorov-Smirnov test is described in Section 3.1. The relaxation algorithm used to classify the pixels is described in Section 3.2. In a relaxation algorithm, the initial segmentation plays an important role but different initialization strategies fail in the problem at hand. Several strategies of initialization are explored in Section 3.3.

Finally, the results of segmenting an infrared satellite image sequence using this method are demonstrated in Section 4.

2. DATA

The algorithm deals with the infrared window channel (11μ) of GOES satellite imager data that have been mapped such that each pixel represents a constant area in latitude and longitude. The images are 125×125 with each pixel representing 0.16 degrees of latitude and 0.16 degrees of longitude. Images of the sequence are available at fifteen minute intervals.

The sequence of satellite images captures a day of significant thunderstorm activity. Several thunderstorms grow and decay during the day and the temperatures of the cloud tops in the images show corresponding changes both in storm geographical extent and in storm severity. Segmentation of this sequence should be able to consistently identify the thunderstorms in the images. Ideally, when storms 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 [3, 4] 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 four frames of the satellite sequence. However, a line of thunderstorms within which these cells crop up can be expected [3] to persist for upto six hours. Therefore, the segmentation should lend itself to segmenting regions corresponding to larger scale features while identifying small scale features that are contained within the




Fig. 1. From left to right, top to bottom: (a) One of the frames in the sequence. Darker pixels represent higher temperatures (ground) and brighter pixels represent lower temperatures (cloud tops). (b) The mean in a local window around the pixel. (c) The variance in the local window. (d) The coefficient of variation within the window. (e) Skewness within the window. (f) Kurtosis within the window. (g) Contrast within the window. (h) Homogeneity within the window.

large scale feature.

2.1. Texture Characteristics

Various statistics are computed in local 5x5 kernels about a pixel and the values assigned to that pixel. The following commonly used features [5] were computed: mean, variance, coefficient of correlation, contrast, homogeneity, skew and kurtosis. Texture segmentation has been shown to be improved [6, 7] by the use of coöccurrence matrices. However, as the “texture” present in weather images is not of a regularly repeating variety, it is not clear whether these features would help in segmenting the images at hand. The seven features used for segmentation when computed on the satellite image in Fig. 1a are shown in Fig. 1b-h.

As can be noticed from these images, a physical thunderstorm does have texture characteristics different from that of the surrounding atmosphere and significantly different from those of the ground. This suggests the use of a segmentation approach that attempts to segment satellite images based on the texture characteristics of regions.

3. SEGMENTATION METHOD

3.1. KS Test

The Kolmogorov-Smirnov (KS) test can be used to test the hypothesis that an observed distribution function of independent random variables belongs to a specified distribution function at varying confidence levels. This is done by computing the critical value, c , corresponding to the significance level, α , in the observed function from the series expansion: [8]

$$\alpha = 2\sum_{j=1}^{\infty} (-1)^{j+1} e^{-2j^2 c^2} \quad (1)$$

The KS distance, d , is computed as the maximum of the distances between the cumulative frequency of the observed random variables and that of the specified distribution. Given the KS distance, the KS critical value and the number of samples, n , used to obtain the observed cumulative distribution function, the observations are said to belong to the specified distribution if $d < c/\sqrt{n}$. Unlike other measures of fit such as the chi-square test, the KS test works even if the distributions are not normal.

Following [9], the problem of segmenting an image based on the local distribution of features can be cast into a KS test. Assume that the segmentation has been initialized in some manner. Then, we can compute the distribution of the features in the pixels within each region. We can then test whether the “local” distribution of these features around a pixel is part of the “global” distribution of these features within a region. If it is part of the global distribution (at the significance assumed), then, this pixel can be added to the region and the global distribution of the region updated.

As in [10, 11, 9], a separate class of outliers is maintained. If the KS distance between the distribution at a pixel is too far from the distributions of the current regions, then that pixel is labeled an outlier. Contiguous outliers are combined into new regions, removing the need for the number of regions to be known apriori.

3.2. Relaxation Algorithm

Finding the optimal segmentation is carried in an iterative manner. During an iteration, each pixel of the image is tested with the labels of its neighbors. Testing a label, p , involves computing the “energy” as the weighted sum of two components

$$U = U_{ks}(p) + 0.5 * U_m(p) \quad (2)$$

and replacing the current label with the label p if the energy corresponding to p is less than that corresponding to the current label. The energy, U_{ks} is the energy corresponding to the KS test and is defined as the average result of the KS test applied to each texture feature (with the test returning 1 if it




Fig. 2. The result of segmentation of two frames in the sequence after initializing the segmentation through a Bayesian maximum likelihood approach. The initial segmentation map and the resulting segmentation of the first frame of the sequence are shown in the top row. The initial and resulting segmentations of a frame two hours later are shown in the bottom row.

is not a part of the region p and -1 if the pixel is part of the region). The energy, U_m is the energy corresponding to the Markov model and is defined as the average result of a test that returns one if an 8-neighbor is not p and returns -1 if an 8-neighbor has the label p . Since U_{ks} can not be computed if p corresponds to the outlier class (there is no common distribution), U_{ks} is taken to be a constant value for the outlier class. This constant is steadily increased with each iteration, thus slowly discouraging the formation of new outlier regions at later stages of the relaxation. It was set to be zero on the first iteration and increased by 0.1 each time around.

At the end of each iteration, contiguous outlier pixels are combined and labeled as new regions if they are large enough. As in [9], this test of size allows us to form reasonably valid probability distributions. The global distributions of each region are then recomputed (to perform the KS test on the next iteration) and the next iteration started. This is continued until the segmentation is stable. Convergence happened on the set of images considered after 6-8 iterations.

3.3. Initialization of the segmentation

The initial segmentation map is critical since the KS energy depends on the global distribution of the first set of regions. The quality of the segmentation is, to a large extent, dictated




Fig. 3. The result of segmentation of two frames in the sequence after initializing the segmentation using pixels that do not lie in modal intervals as described in Section 3.3. The initial segmentation map and the resulting segmentation of the first frame of the sequence are shown in the top row. The initial and resulting segmentations of a frame two hours later are shown in the bottom row.

by the choice of initial condition. The initialization can be done in one of several ways, including:

1. Using only the KS distance after labeling the entire image as one region. Outliers from the distribution of features in the entire image will then be combined in the next iteration of the relaxation algorithm into new regions.
2. Labeling as outliers pixels the majority of whose features do not fall in the most frequent interval of that feature's probability distribution. The other pixels are labeled as a single region.
3. Labeling pixels that have extreme values in the mean (say the top 5% and the bottom 5%) within the image as outliers and the rest of the pixels as a single region.

Although the first choice is tempting because of its consistency with the rest of the algorithm (and was found to work on a wide variety of images in [9]), the initial labeling varies widely across a sequence of images with small changes from frame to frame (See Fig. 2).

4. RESULTS

The segmentation algorithm was carried out on a sequence of infrared images showing a developing thunderstorm. The




Fig. 4. The result of segmentation of two frames in the sequence after initializing the segmentation using pixels that belong to the extrema in the mean as described in Section 3.3. The initial segmentation map and the resulting segmentation of the first frame of the sequence are shown in the top row. The initial and resulting segmentations of a frame two hours later are shown in the bottom row.

result of initializing using the first method described in Section 3.3 on the first image of the sequence and an image captured about two hours later is shown in Fig. 2. Segmentation of the same images using the second method of initialization is shown in Fig. 3. Using the third method of initialization is demonstrated in Fig. 4.

When initialized using pixels that do not lie in the modal interval, the segmentation is robust and only the strongest part of the storms is retained. Not all the features demonstrated in Fig 1 have the same discrimination capacity. In particular, the contrast within the local window is much higher in the pixels corresponding to storms. A more careful choice of the features to be used and using a weighted average of the results of the KS test in each of the features to compute U_{ks} might improve the performance of segmentation.

One of the requirements of Section 2 – that splits and merges of the physical storms be reflected in the segmentation itself – was not met. Indeed, we find that while the location of the segmented regions is consistent frame-to-frame and that small changes in the storm structure are reflected in small changes in the segmentation of the image *in its entirety*, the same doesn't hold for individual storms. Regions corresponding to different storms may combine even when the storms haven't merged and may split even the physical thunderstorm hasn't. The poor performance in this regard

suggests that texture-based segmentation needs to be followed by a post-processing stage where the history of the segmented regions is used to determine the final segmentation at a frame.

5. REFERENCES

- [1] J.T. Johnson, P. Mackeen, A. Witt, E.D. Mitchell, G. Stumpf, M. Eilts, and K. Thomas, "The storm cell identification and tracking algorithm: An enhanced WSR-88D algorithm," *Weather and Forecasting*, vol. 13, no. 6, pp. 263–276, June 1998.
- [2] V. Lakshmanan, R. Rabin, and V. DeBrunner, "Identifying and tracking storms in satellite images," in *Second Artificial Intelligence Conference*, Long Beach, CA, 2000, American Meteorological Society, pp. 90–95.
- [3] K.A. Browning, "The FRONTIERS plan: A strategy for using radar and satellite imagery for very-short-range precipitation forecasting," *The Meteorological Magazine*, vol. 108, pp. 161–184, 1979.
- [4] A. Bellon and I. Zawadzki, "Forecasting of hourly accumulations of precipitation by optimal extrapolation of radar maps," *Journal of Hydrology*, vol. 157, pp. 211–233, 1994.
- [5] A. Solberg and A. Jain, "Texture fusion and feature selection applied to SAR imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 35, no. 2, pp. 475–479, 3 1997.
- [6] D. Geman, C. Geman, C. Graffigne, and D. Pong, "Boundary detection by constrained optimization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 7, pp. 609–628, Jul 1990.
- [7] R.M. Haralick, "Statistical and structural approaches to texture," *Proc. IEEE*, vol. 67, no. 5, pp. 786–804, May 1979.
- [8] J. Blum and J. Rosenblat, *Probability and Statistics*, W.B. Saunders Company, 1972.
- [9] C. Kervrann and F. Heitz, "A markov random field model-based approach to unsupervised texture segmentation using local and global spatial statistics," *IEEE Trans. on Img. Proc.*, vol. 4, no. 6, pp. 856–862, Jun 1995.
- [10] P.B. Chou and C.M. Brown, "The theory and practice of bayesian image modeling," *Int. J. Comp. Vis.*, vol. 4, no. 2, pp. 185–210, Feb 1990.
- [11] P. Bouthemy and E. Francois, "Motion segmentation and qualitative dynamic scene analysis from an image sequence," *Int. J. Comp. Vis.*, vol. 10, no. 2, pp. 157–182, Feb 1993.