quarterly report for

Applications of High Performance Statistical Modeling to Image Analysis of Forest Structure

submitted to:

Mr. Victor A. Rudis
USDA Forest Service
Southern Research Station
P. O. Box 928
Starkville, MS 39760-0928
Tel: 662-338-3109, Fax: 662-338-3101
Email: vrudis@usfs.msstate.edu

May 15, 1999

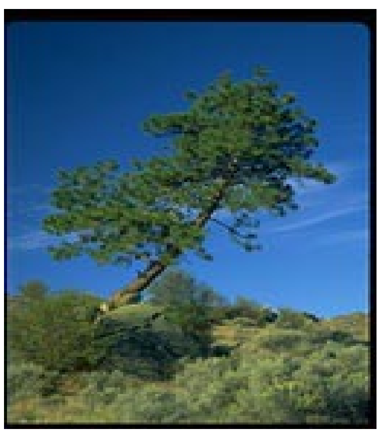

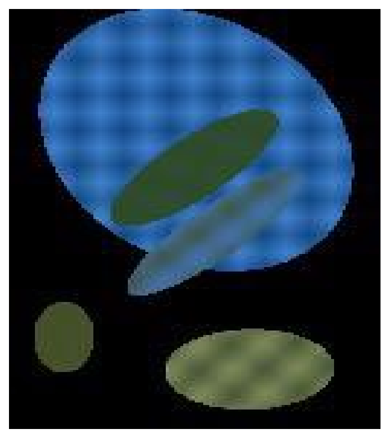

submitted by:

Z. Long, J. Picone
Institute for Signal and Information Processing
Department of Electrical and Computer Engineering
Mississippi State University
Box 9571
413 Simrall, Hardy Road
Mississippi State, Mississippi 39762
Tel: 662-325-3149
Fax: 662-325-3149
Email: {long, picone}@isip.msstate.edu

INTRODUCTION

In this quarter, we made significant progress towards implementation of an object recognition scheme. Our highly versatile image processing software that previously was designed to produce Scenic Beauty Estimates (SBE) was modified to read segmentation information and pass this information to the appropriate computational routines. We also incorporated new colors into our color measurement scheme and evaluated their effectiveness. Finally, we completed some preliminary evaluations on a segmentation-based classification scenario by experimenting on images segregated into three regions. While we did not improve on our overall best SBE evaluation score, we have been able to achieve comparable results with much simpler systems using some of these new techniques.

EXPERIMENTS WITH FEATURES BASED ON BROWN AND YELLOW

The colors of brown and yellow have been shown [1] to play an important role in human perception of the scenic beauty of forestry images. To investigate their effectiveness on our image classification problem, we implemented a new color computation and classification scheme. We computed the histogram for the colors of brown and yellow, as we did for red, green and blue previously. To calculate the intensity of brown and yellow for a given pixel, we took advantage of the theory of additive color synthesis [2]. According to this theory, we can create any color by mixing various proportions of three primary colors: red, green, and blue.

In particular, we can generate the color brown by mixing one part blue, one part green, and four parts red. Similarly, yellow is obtained by mixing equal parts of red and green. Since we desire to measure the intensity of these colors, as we did in previous experiments for the primary colors, we implemented a normalized version of the mixing formulae:

$$brown = \frac{1}{6} \times blue + \frac{1}{6} \times green + \frac{4}{6} \times red \tag{1}$$

$$yellow = \frac{1}{2} \times red + \frac{1}{2} \times green \tag{2}$$

where red, green and blue stand for the corresponding intensity values of the given pixel.

We investigated the effects of these new features on our image classification system using Principal Components Analysis (PCA). We evaluated various combinations of features involving brown and yellow, and compared these to some existing baseline results. These results are shown in Table 1. We believe it is a very encouraging sign that brown and yellow, combined with other primary colors, produced good classification results. For example, blue and brown combined gave us our best result yet for a PCA-based measure on the first data set. We believe this can be attributed to the fact that blue is keying on the presence of sky in the image, and brown is related to the presence of foliage in the image.

Unfortunately, the excellent results on the first data set did not hold up well over the entire evaluation, probably as a result of the varied nature of the test sets with respect to scene composition. Hence, though these new features did not produce ground breaking results, we believe they will be important when we move to segmentation-based classification. Our progress in this area is described in the next section.

Features	Data Set 1	Data Set 2	Data Set 3	Data Set 4	Average
red + green + blue	40.3	42.4	38.1	33.1	**38.5**
yellow	54.7	38.6	45.6	67.5	51.6
brown	37.7	38.6	41.9	55.0	43.3
yellow + brown	40.3	46.2	40.6	45.6	43.2
rgb + yellow + brown	39.6	39.9	39.4	66.9	46.5
blue + brown	**31.4**	50.6	35.0	56.9	43.5
rgb + brown	47.2	46.2	48.1	38.8	45.1
grn + blue + ylw + ll	39.6	38.0	45.0	53.1	43.9
grn + blue + ylw + brn + ll	39.6	52.5	48.1	40.0	45.1
blue + brn + ll	37.1	52.5	33.1	58.1	45.2

Table 1. Results of SBE evaluations using brown and yellow.

SEGMENTATION-BASED IMAGE CLASSIFICATION

The primary goal of this phase of the project was to begin development of a system capable of leveraging our manual segmentations. This is being done in two steps. First, we would like to assess the effectiveness of this information through “cheating” experiments where we estimate scenic beauty given knowledge of the segmentations. Second, we will develop schemes to automatically segment the images, and evaluate the accuracy of these schemes based on the manually derived information. In this report, we review progress on the former task.

In our manual segmentation approach described previously, each image was segmented into regions of similar appearance, and assigned one of the following tags: sky, trees, foliage, grass and bushes [3]. Given the striking visual differences between most of these regions, we believe that if we build region-specific models, and classify images according to these regions, we will attain a significant improvement in our ability to automatically classify forestry images. However, instead of proceeding directly to implementation of such a segmentation-based scenario, we first need to evaluate the theoretical limits such approaches can achieve. We refer to this type of experiment as a “cheating” experiment, since we will use the manual segmentations as input (and these are presumably as good as any automatic algorithm can do).

Implementation of a true cheating experiment using manual segmentations requires non-trivial modifications to our software. Hence, we constructed an incremental experiment as a first step towards this goal. We used a simple image splitting scheme: each image was split into three horizontal strips (three equally spaced regions). The top region was designated “sky.” The middle region was designated “trees and foliage.” The bottom region represented “grass and ground.” We evaluated our ability to classify each image solely based on one of these regions using the feature vector “blue + brown.” We compared these results with those generated on the non-split images. The results of these comparisons are shown in Table 2. We used a PCA based approach for all preliminary experiments.

Region	Data Set 1	Data Set 2	Data Set 3	Data Set 4	Average
top	44.0	38.0	43.8	45.0	**42.7**
middle	37.7	46.1	71.3	60.0	52.5
bottom	50.3	35.4	53.1	42.5	45.3
complete	31.4	50.6	35.0	56.9	**43.5**

Table 2. Results of SBE evaluations using blue + brown on segmented images.

Observe that the average performance across all data sets for the top region is slightly better than a comparable result for the entire image. In other words, by examining only the top one-third of each image, using features of blue and brown, we are able to achieve our best overall performance. Clearly, this approach is doing a better job of weighting the relative importance of such features as sky and brush.

To further explore the significance of this result, we continued experiments using region-specific features. For each region, we chose features believed to be relevant to the type of imagery expected in that region. These experiments are summarized in Tables 3-5. It is interesting to note that we obtained much better results on some specific data sets. For example, by examining the top region using features of rgb combined with entropy, we achieved a 27.8% error rate on data set 2, which is the best performance obtained on any data set thus far. Building on this encouraging result, we chose the best error rates for each region and compared them with the best average error rate obtained previously for the complete images. This is summarized in Table 6.

Obviously, the best results of this region-specific measurement scheme were comparable to those of our previous best systems that processed the entire image. Though this approach was a very simple first step, we believe it justifies examining segmentation-based approaches in more detail. To do segmentation-based evaluation, we need to add the following capabilities to the software: reading in the segmented regions as polygons; running evaluation only on a specific region prescribed by the user. These require implementation of a class of polygon, which can load the segmentation data and determine if a given pixel is inside this polygon or not. Further, these require modification to the image_analysis class to incorporate this polygon class.

Currently, we have finished the design of both of these classes. We have implemented the core functionality of the polygon class. The last remaining steps involve extension of the computational classes in the image_analysis class. We have modified the driver program to load segmentations and to support experimentation on either a specified region or an entire image.

Our time line for the remaining steps includes completion of all cheating experiments in one week, and completion of the first version of automatic segmentation algorithms in one month. At that time, we will conduct formal evaluations comparing the results to our reference segmentations, and assessing the impact on SBE evaluations (an indirect means of assessing the quality of the segmentations). Once this step is complete, our focus will shift to automatic interpretation of the segmentations in ways that will move us closer to true object recognition. We expect one conference and one journal paper to be developed from this summer’s work.

Features	Data set 1	Data set 2	Data set 3	Data set 4	Average
blue	48.4	31.6	43.1	50.0	43.3
rgb	51.6	30.4	66.9	36.9	46.5
rgb + ent	54.1	27.8	30.6	35.6	37.0

Table 3. Results of SBE evaluations using only the top region of each image.

Features	Data set 1	Data set 2	Data set 3	Data set 4	Average
green	49.7	58.9	65.0	60.6	58.5
brown	44.7	57.0	66.9	51.9	55.1
rgb	43.4	40.5	39.4	41.9	41.3
rgb + ent	42.8	40.5	31.3	40.6	38.8

Table 4. Results of SBE evaluations using only the middle region of each image.

Features	Data set 1	Data set 2	Data set 3	Data set 4	Average
brown	37.7	65.8	46.3	41.3	47.8
ylw + brn	44.0	52.5	45.0	41.3	45.7
rgb	45.3	32.9	41.3	33.1	38.2
rgb + ent	38.4	32.3	43.8	36.9	37.9

Table 5. Results of SBE evaluations using only the bottom region of each image.

Region	Features	Error rate
top	rgb + ent	37.0
middle	rgb + ent	38.8
bottom	rgb + ent	37.9
complete	rgb + ent	37.7

Table 6. A comparison of performance using region-specific measures to a baseline system.

SUMMARY

In this report, we have described a new color computation and classification scheme using brown and yellow. These new features provided performance comparable to our previous best systems. We also described preliminary experiments to evaluate a segmentation-based classification approach. By processing only the upper portion of each image, we were able to slightly improve on our previous best result. We described our initial implementations of software that will use manual segmentation information, and will eventually lead us to the goal of this phase of the project — automatic image segmentation and classification. We expect to complete implementation and evaluation of these approaches in the next three months of this project.

REFERENCES

[1] W. Yhang, J. H. Gramann and V. A. Rudis, *Color Visibility by Season and Silvicultural Treatment: Effects on the Scenic Beauty of Southern U.S. Pine-Oak Forests*, Texas A&M University, College Station, Texas, USA, and USDA-FS Southern Research Station, Forest Inventory and Analysis Unit, Starkville, MS, USA, October 1998 (manuscript draft).

[2] J. Scruggs, "Color Theory — Additive Color Synthesis," *http://www.bway.net/~jscruggs/add.html*, LTI Labs, New York City, USA, May 1999.

[3] Z. Long and J. Picone, "Applications of High Performance Statistical Modeling to Image Analysis of Forest Structure," *Quarterly Status Report for the Southern Research Station, USDA Forest Service,* Institute for Signal and Information Processing, Mississippi State University, February 1999.