Spectral Shape Classification System for Landsat The matic Mapper

Mark J. Carlotto (mjcarlotto@ma.psrw.com)
Pacific-Sierra Research Corporation, 1400 Key Blvd., Suite 700, Arlington, VA 22209

Abstract

A multispectral classification system based on an alternative spectral representation is described and its perfor mance over a full Landsat Thematic Mapper (TM) scene evaluated. Spectral classes are represented by their s pectral shape - a vector of binary features that describes the relative values between spectral bands. An algorit hm for segmenting or clustering TM data based on this representation is described. After classes have been a ssigned to a subset of spectral shapes within training areas, the remaining spectral shapes are classified accor ding to their Hamming distance to those that have already been classified. The performance of the spectral sh ape classifier is compared to a maximum likelihood classifier over five sites that are fairly representative of th e full Landsat scene considered. Although the performance of the two classifiers is not significantly different within a site, the performance of the spectral shape classifier is significantly better than the maximum likeliho od classifier across sites. A full-scene spectral shape classifier is then described which combines spectral sign ature files that associate classes with spectral shapes derived over the five sites into a single file that is used to classify the full scene. The classification accuracy of the full-scene spectral shape classifier is shown to be su perior to that of a stratified maximum-likelihood classifier. The spectral shape classifier is implemented in C a nd is able to process an entire Landsat TM scene in about one hour on a single processor SUN SPARC 10 w orkstation with 128 megabytes of RAM.

Key words: Multispectral classification, signature extension, pattern recognition, accuracy assessment, land cover classification

1. Introduction

Accurate and reliable classification of multispectral imagery over extended areas is critical to the development of land cover maps for a variety of applications in a timely and cost-effective manner. Yet after over twenty ye ars the generation of such maps from imagery in an operational manner remains a labor-intensive and costly p rocess. Conventional statistical classifiers perform well over limited areas where spectral signatures do not var y greatly from those captured in the training data. However as the size of area to be classified increases, the cl assification accuracy typically decreases due to environmental, topographical, and phenological factors. The m ost common method of classifying large heterogeneous regions is by spatial stratification whereby the scene i s divided into regions (e.g., based on climate, topography, etc.), each region is classified separately, and the res ults combined (Todd et al 1980, Hutchinson 1982). An advantage of stratification is that regional knowledge c an be used to significantly improve classification accuracy. By labeling each part of the scene separately and c ombining the results the overall classification accuracy does not have to be sacrificed for area coverage. A disa dvantage of stratification however is that it requires a certain amount of interactive processing as well as additi onal data (e.g., elevation matrices, maps) which adds to the processing cost. Signature extension techniques (Henderson 1974) and extendible classification algorithms (Carlotto 1990) provide an alternative to stratificati on in which spectral signatures derived over limited portions of a scene are used to classify the remainder of t he scene and in some cases other scenes as well.

An early rule-based multispectral classifier (Carlotto et al 1984) used qualitative knowledge and relative constr aints for classifying general land cover categories. Two kinds of rules were developed: those that defined clas ses relative to each other in terms of their spectral features (e.g., the greenness of vegetation is greater than the

greenness of bare soil, etc.), and those that defined single classes in terms of the relative values between spectral bands. The spectral shape classifier (Carlotto and Tom 1985) was an outgrowth of this work and addressed the problem of deriving a complete and consistent set of spectral classification rules from training data. In this paper we describe a new method of classifying multispectral imagery based on a set of binary features that represent the relative values between spectral bands. We begin by describing the theoretical basis of the spectral shape representation. The spectral shape representation is compared to K-means clustering and its use in multispectral classification described. The software architecture of the spectral shape classification system is then outlined. Experimental results from a full Landsat scene are used to examine the classification accuracy of the spectral shape classifier within and across training areas and to compare its overall accuracy to that of a maximum likelihood classifier.

2. Spectral Shape Representation

Assume the atmosphere can be modeled as a horizontally homogeneous medium, the earth is a Lambertian reflector, and atmospheric properties which vary exponentially with altitude can be assumed to be constant over the scene (Sjoberg and Horn 1983). We also assume that the contribution of the ambient illumination (skylight) can be neglected, and that the topographic component (i.e., the modulation of the scene brightness caused by topography) does not depend on wavelength. Under these conditions, the sensed irradiance (brightness) at pixel (i, j) in band n can be modeled by a linear relation

$$y(i,j,n) = a(n)b(i,j)x(i,j,n) + c(n) \tag{1}$$

where $x(i, j, n)$ is proportional to the spectral reflectivity (albedo) of the surface, $a(n)$ depends on solar illumination, atmospheric transmission and sensor gain, $b(i, j)$ is the topographic component, and $c(n)$ represents the contribution of the path radiance caused by atmospheric scattering.

Instead of correcting for terrain and atmospheric effects systematically we seek a representation (i.e., a set of features) for classification that will be less sensitive to these effects, perhaps at the expense of loosing some spectral detail. Our method represents the shape of the spectral response in terms of the relative values between bands

$$\phi(i, j, n, n') = \begin{cases} 1, & y(i, j, n) > y(i, j, n') \\ 0, & otherwise \end{cases} \tag{2}$$

where $n = 1 \cdots N - 1$ and $n' = n + 1 \cdots N$. In effect the spectral shape representation converts the original multispectral image into a image of binary features that are used for classification.

Like band ratios in high-relief areas, the above features are not affected by the topographic component $b(i, j)$ provided the path radiance terms $c(n)$ have been previously removed (Crippen 1987). Non-selective scattering, e.g., caused by thin clouds (Schanda 1986), which can be modelled as a relatively constant effect across bands will also not affect the values of these features.

It is noted that the set of binary features defined in Eq. 2 is equivalent to ranking the spectral bands in decreasing (or increasing) order by value; e.g., $y(i, j, n_1) > y(i, j, n_2) \cdots$ where n_k is the number of the k - th largest band in value at a particular pixel, and using the rank-ordered band numbers as features. Figure 1 shows an example spectral response and its spectral shape representation both in terms of binary features and rank-ordered band numbers. We choose the binary feature representation because it has a simple physical interpretation and can be compared using similarity measures such as the Hamming distance as discussed later in the paper.

For N bands, there are $N \times (N-1) \cdots = N!$ possible ways to order the bands. Each ordering can be expresse d by a unique combination of $N(N-1)/2$ binary features. Ordering bands by value using Quicksort (Sedge wick 1983) requires approximately $N \log_2 N$ operations compared to $N(N-1)/2$ operations involved in co mparing all distinct pairs of bands and is thus computationally more efficient as the number of bands increase s. It should be clear however that since the number of features grows quadratically and the number of combin ations grows factorially, the spectral shape representation is intended for multispectral sensors with a moderat e number of bands such as Landsat TM and is not well-suited for sensors with a small number of bands like SPOT or for hyperspectral sensors.

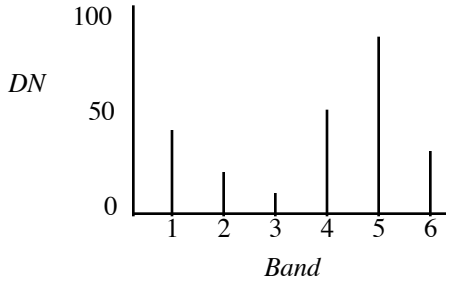

Band	1	2	3	4	5	6
DN	40	20	10	50	100	30

Example spectral response for 6 band sensor (e.g., Landsat TM bands 1-5 and 7)

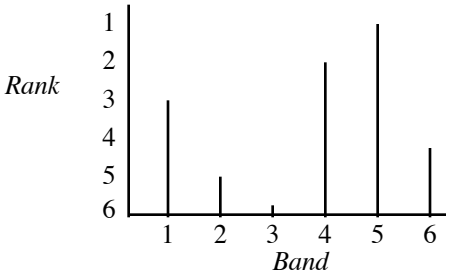

Rank	1	2	3	4	5	6
Band	5	4	1	6	2	3

Spectral bands ranked by value - Spectral value replaced by rank (left) provides a qualitative description of the spectral shape

Bands	1,2	1,3	1,4	1,5	1,6	2,3	2,4	2,5	2,6	3,4	3,5	3,6	4,5	4,6	5,6
Relative Value	1	1	0	0	1	1	0	0	0	0	0	0	0	1	1

Representation of spectral shape in terms of binary features

Figure 1 Example of spectral response and its spectral shape representation

The spectral shape representation effectively segments the spectral measurement space into $N!$ disjoint region s. Each spectral shape corresponds to a wedge-shaped region in this space, all of which touch the origin. Figu re 2 illustrates the shape of the region in 3-space corresponding to the set of features $\Phi = \{1,1,1\}$, i.e., the regi on where, for a sensor with three bands, $y(1) > y(2)$, $y(1) > y(3)$, and $y(2) > y(3)$. We show the construction of the region in stages for clarity. Without loss of generality the original spectral response can be uniformly s caled to fit into the unit cube shown in a). In b) the cube is split in half along the diagonal $y(1) = y(2)$ and the half-space $y(1) > y(2)$ retained. In c) the previous region is split along the diagonal $y(1) = y(3)$ and the part w here $y(1) > y(2)$ and $y(1) > y(3)$ is retained. Finally in d) the previous region is split along the diagonal $y(2) = y(3)$ and the region where $y(1) > y(2)$, $y(1) > y(3)$, and $y(2) > y(3)$ retained. By integration it can be s hown that the volume of this or any of the other five wedge-shaped regions in 3-space is 1/6 ($= 1/N!$).

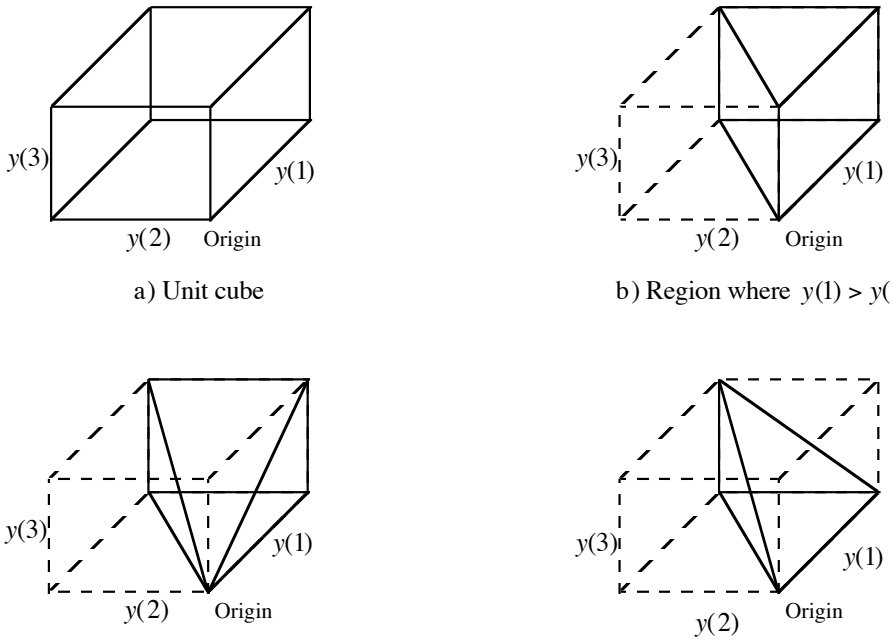

Figure 2 Construction of region in spectral measurement space corresponding to a spectral shape

The use of binary features provides a basis for comparing spectral shapes. The similarity between two spectral shapes u and v is defined to be equal to the number of binary features that are different and is given by the Hamming distance

$$d(u,v) = \sum_{n=1}^{N-1} \sum_{n'=n+1}^{N} \phi(u,n,n') \oplus \phi(v,n,n') \tag{3}$$

where $\oplus$ denotes exclusive-or.

3. Application to Clustering

Clustering is performed to identify regions with similar spectral properties in an image. Algorithms such as K-means and ISODATA (Tou and Gonzales 1974) partition the underlying spectral measurement space into clusters where the parameters of the clusters are adjusted iteratively to minimize some objective function, typically the total squared error between the clusters and the data. Since the spectral shape representation effectively segments the measurement space it functions like a clustering algorithm. An important difference between spectral shapes and clusters is that spectral shapes correspond to regions with fixed boundaries in the measurement space where the boundaries between clusters depend on the distribution of the data.

Spectral shapes and K-means clusters were extracted from a 226x236 pixel image over Gordonsville VA (Figure 3). Figure 4 compares the structure of clusters in feature space extracted by the K-means algorithm to that of the fixed spectral shape regions for Landsat TM data. Instead of attempting to visualize the underlying six-dimensional space for Landsat TM (bands 1-5 and 7), the K-means clusters and spectral shape regions have been projected down into the two-dimensional space spanned by the first two tasseled cap features (Crist and Cicone 1984). The boundaries between clusters and spectral shape regions are shown in the figure. 56 spectral shapes were found in the Landsat TM over the region shown in Figure 3. For comparison purposes the same number of clusters were extracted using the K-means algorithm.

Figure 3 Landsat TM band 3 over Gordonsville site

Comparing the two plots in Figure 4, K-means clusters are more evenly distributed in space than spectral sha pes (K-means attempts to minimize the squared error between the data and the clusters) and are more compact (K-means uses the Euclidean distance which encourages the formation of compact regions). The boundaries between spectral shape regions, on the other hand, radiate out from the origin and have more of a tapered appe arance. It has been observed that multispectral data tends to cluster in teardrop rather than elliptical distributio ns due to shading, shadowing, and pixel mixing (Craig 1994, Crist and Cicone 1984). The boundaries betwee n spectral shapes also have this same general appearance.

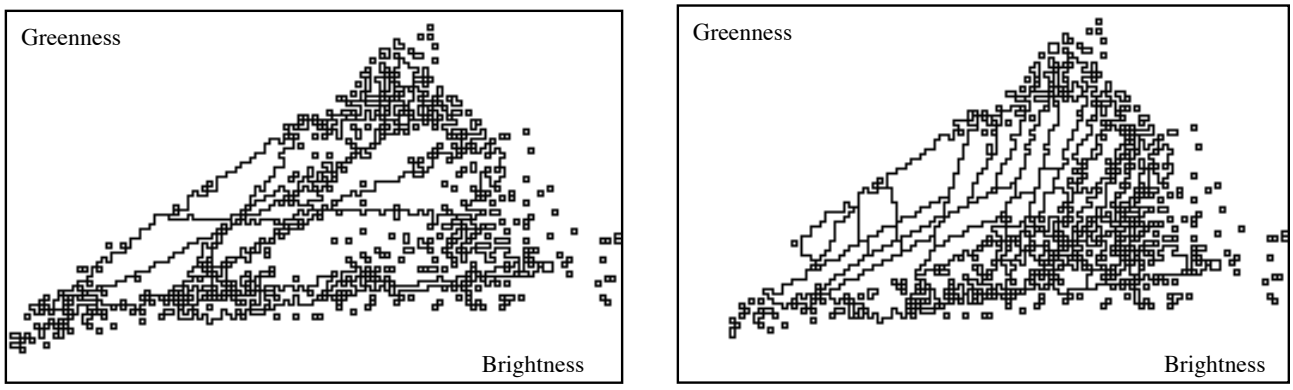

a) Boundaries between spectral shape regions b) Boundaries between K-means clusters

Figure 4 Comparison of spectral shape and K-means cluster boundaries in tasseled cap space

Finally it is noted that the amount of computation required to compute the spectral shape representation is sig nificantly less than algorithms like K-means which have a complexity of $\approx O(KLN)$ floating point operations per pixel where L is the number of iterations. As a result, the spectral shape algorithm can be used to segment an entire data set without having to first reduce its size, e.g., by sub-sampling, as is often required with K-mea ns and similar clustering algorithms.

4. Multispectral Classification

The spectral shape approach can be used for multispectral classification by segmenting an image into spectral shapes as described in the previous section, assigning a class to a subset of spectral shapes (e.g., those within a training area), and classifying the remaining spectral shapes according to their Hamming distance from thos e that have already been assigned a class. Figure 5 compares spectral shape and maximum likelihood classific ations over the Gordonsville site shown in Figure 3. The spectral shape classifier was trained by visually assig ning a class to a subset of the spectral shapes. Remaining spectral shapes were classified by assigning the cla ss of the spectral shape with the smallest Hamming distance. A maximum likelihood classifier was trained by first clustering the image using K-means as described in the previous section. Classes were assigned by visua lly assigning a class to a subset of the K-means clusters. Mean vectors and covariance matrices were then co mputed from the clusters and used to classify the full image. (A more detailed discussion of training and accu racy assessment is contained in the next section.)

Six general surface categories are depicted in Figure 5: built-up, barren, herbaceous (grassland and agriculture), woody, wetland, and open water. In general, the spectral shape classification appears to contain somewhat le ss detail than the maximum likelihood classification, although the latter appears to confuse shadows and water and to contain more built-up pixels, many of which are scattered within woody and agricultural areas.

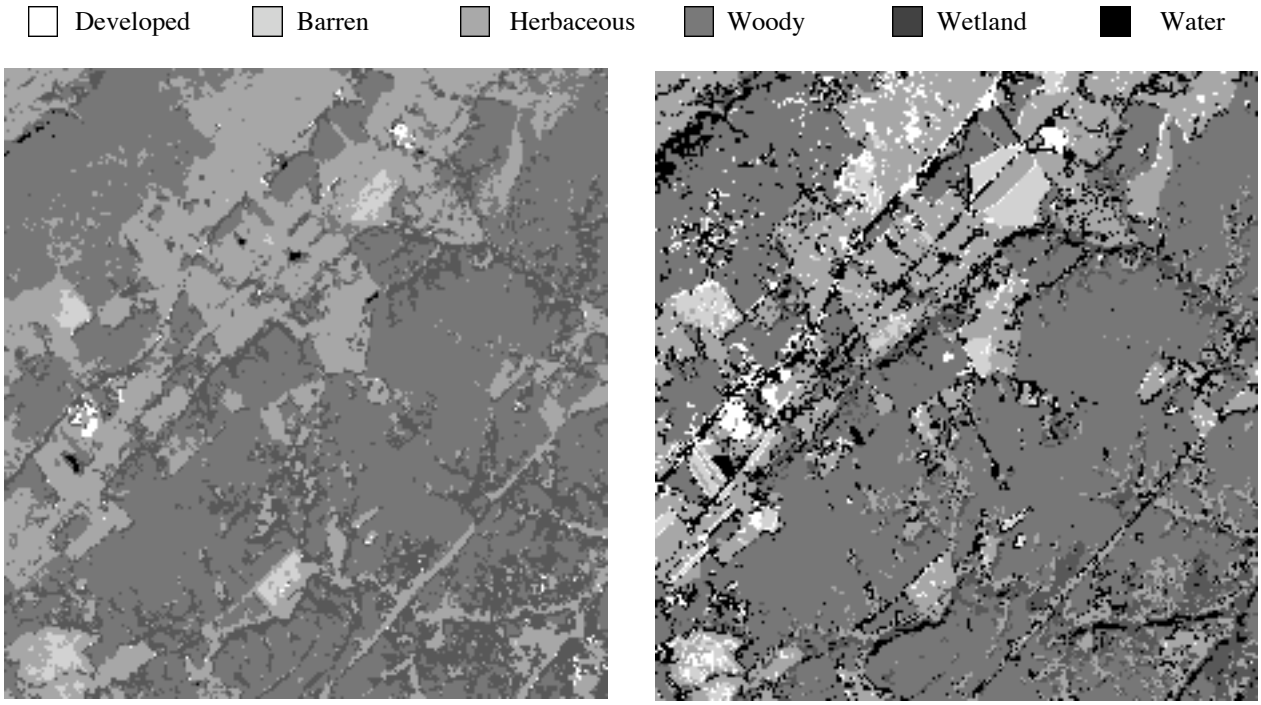

a) Spectral shape classification b) Maximum likelihood classification

Figure 5 Visual comparison of spectral shape and maximum likelihood classification results

Figure 6 plots the classification results in Figure 5 in tasseled cap space as a function of their brightness and greenness values. (The shades of gray match those used in Figure 5). Even though different classifiers were u sed, the class boundaries in the two distributions are similar. The boundaries between spectral classes in Figur e 6 are more like the boundaries between spectral shapes (Figure 4a) than the boundaries between K-means cl

usters (Figure 4b). Over-clustering and merging K-means clusters into classes appears to create decision regi ons that are very similar in appearance to the boundaries between spectral shape regions.

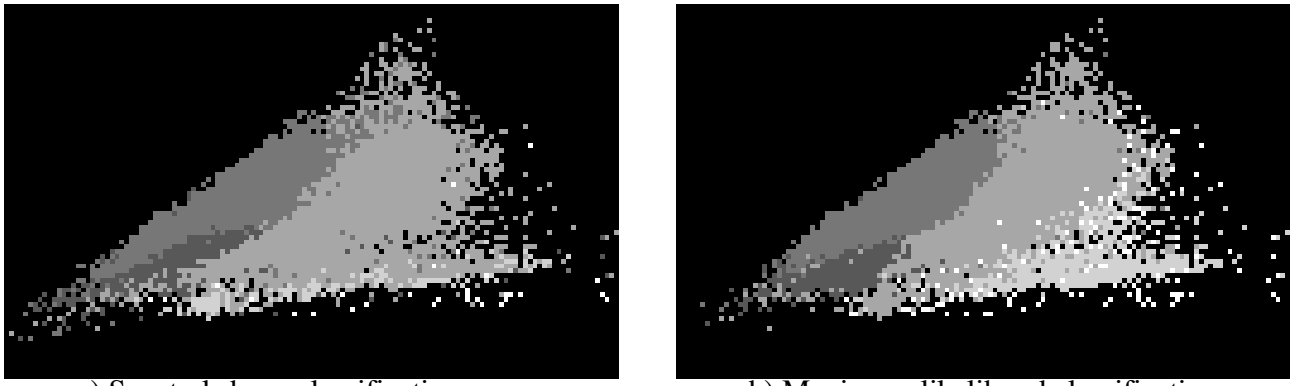

a) Spectral shape classification b) Maximum likelihood classification

Figure 6 Spectral shape and maximum likelihood classifications plotted in tasseled cap space

5. Software Implementation

Figure 7 depicts the software architecture of the spectral shape classification system. The system contains fou r major functions.

Compute spectral shape representation - Segments the input Landsat TM image into spectral shapes. The spat ial extent of each spectral shape is identified by a unique value in the label image. A file lists the relative frequ ency of each label along with the binary feature vector describing the corresponding spectral shape (In the pre sent implementation the largest 255 spectral shapes are extracted, i.e., those accounting for most of the image area. When there are more than 255 spectral shapes in an image, pixels whose spectral shapes are not retained are assigned the nearest spectral shape among the 255 spectral shapes that were retained. Typically for a full Landsat scene, this amounts to only a few pixels total. Most TM scenes processed to date contain only about 100-200 spectral shapes out of a possible 6! = 720.)

Training - Associates classes derived from ground truth data with spectral shapes. The output of the training p rocess is a classification file that specifies the class most frequently associated with each spectral shape in the training set along with its relative frequency. Classification files can be stored in a database and used to classif iy other similar scenes.

Merge classification files - Combines data from multiple training areas into a single classification file. Spectra l shapes that have been assigned more than one class in different training areas are assigned the class that has been assigned most frequently to the spectral shape overall.

Minimum Hamming distance classifier - Assigns the class associated with spectral shapes in the classification file to spectral shapes present in the image being classified. Spectral shapes not in the classification file can b e assigned the nearest class.

The spectral shape classifier is implemented in C and has been integrated into the Khoros visual programmin g environment (Konstantinides and Rasure 1994)). The system is able to classify a full Landsat TM scene in about an hour on a SUN SPARC 10 workstation with 128 Megabytes of RAM.

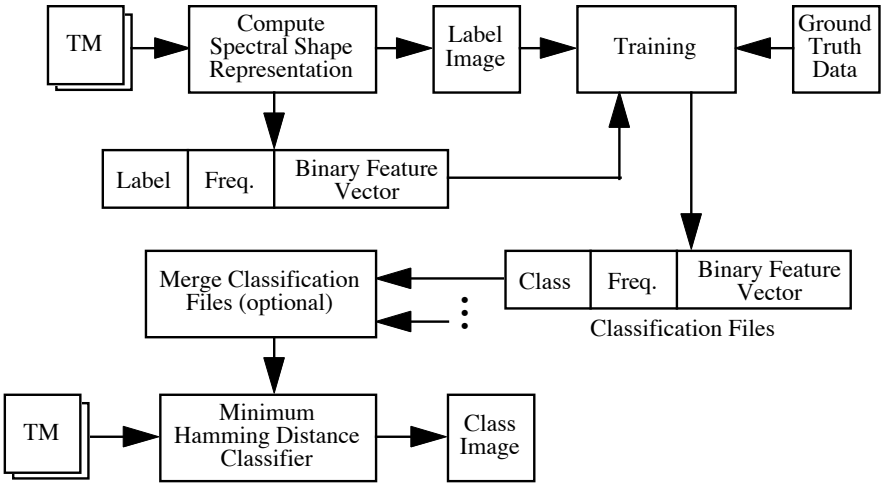

Figure 7 Software architecture of the spectral shape classification system

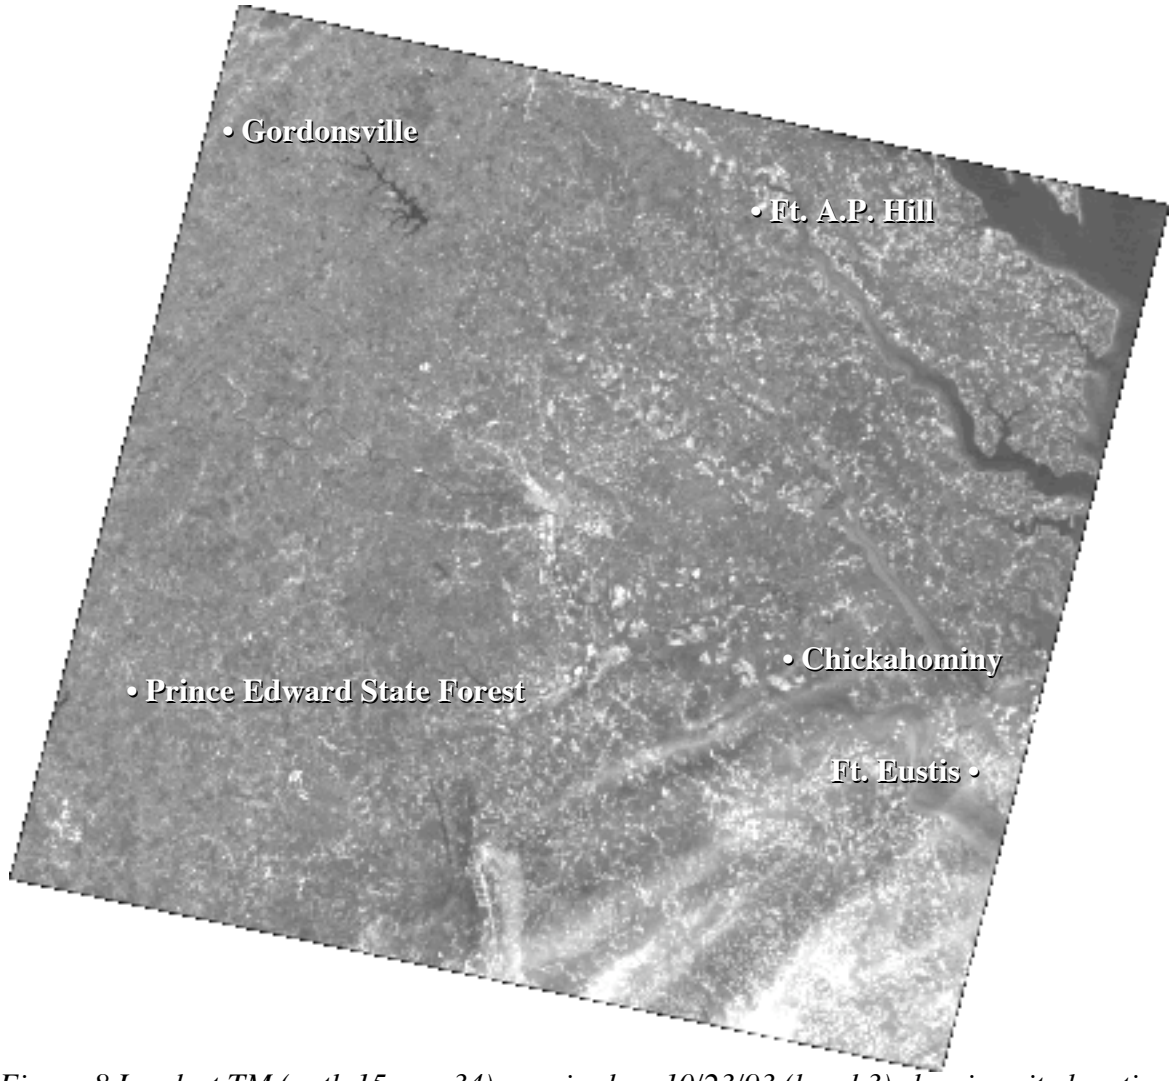

Figure 8 Landsat TM (path 15 row 34) acquired on 10/23/93 (band 3) showing site locations

6. Classification Results

A series of experiments was performed to assess the accuracy of the spectral shape classifier over a full Landsat TM scene (Path 15 Row 34) acquired on Oct 23, 1993. Five sites approximately 25 km^2 in size within the scene were selected for study (Figure 8). Two of the sites, Fort Eustis and Chickahominy, are wetland sites containing mostly open water, grassy and forested wetlands, and forests. The other three sites, Fort A.P. Hill, Gordonsville, and Prince Edward State Forest, are upland sites containing forested and agricultural areas. With the exception of large built up areas, the five sites are fairly representative of the overall scene.

Site selection was based on the availability of M7 imagery which was used as image truth for classifier development and testing. The M7 was acquired within several days of the Landsat imagery. (M7 is a 12 channel multispectral with an instantaneous field of view of 2 mrad (Slater 1985). The effective ground resolution of the imagery used in this study was about 5 meter/pixel.) Over each site, the 25 meter Landsat TM data were resampled to 5 meters using nearest neighbor interpolation. M7 imagery was then registered to the TM.

The two level classification scheme shown in Table 1 derived from a land cover classification system under development by Bara (1993) was used in our evaluation.

Level 1 Class	Criteria	Level 2 Class	Criteria
Developed	> 50% man-made	High Intensity	> 80% man-made
		Low Intensity	50-80% man-made
Herbaceous Land	> 50% herbaceous	Cropland	Managed
		Grassland	Unmanaged
Woody	> 50% woody	Deciduous	> 67% deciduous
		Evergreen	> 67% evergreen
		Mixed	
Barren	< 50% vegetated		
Wetland	Anderson Level 1	Shore	< 50% vegetated
		Emergent	> 50% herbaceous
		Woody	> 50% woody
Open Water	Anderson Level 1		

Table 1 Definition of land cover classes

Our evaluation was based on comparing the overall accuracy (fraction correct) of the spectral shape classifier to that of a maximum likelihood classifier. Over each site, both classifiers were trained using the M7 imagery as truth. For the spectral shape classifier, spectral shapes were first extracted over the site. A subset of the spectral shapes were assigned a class visually by an image analyst using the M7 imagery as reference. The remaining classes were then assigned the class of the nearest spectral shape previously assigned a class based on the Hamming distance. The classification file was retained for use in classifying other sites.

For the maximum likelihood classifier, the imagery was first clustered using the K-means algorithm into 30 clusters. This number of clusters was judged by an image analyst to be adequate to separate the classes of interest within each of the sites. A subset of the clusters were assigned a class, again using the M7 imagery as a reference. The mean vectors and covariance matrices were then computed and used to classify the entire image. The means and covariances were then retained for use in classifying other sites. In order to obtain an unbiased estimate of the relative performance of the two classifiers, a random sampling scheme was used to generate sixty sample points within each of the five sites. The points were assigned a class using the M7 imagery by a second image analyst. The points were then used to measure the accuracies of the two classifiers at levels 1 and 2 as defined in Table 1.

Training Site	Evaluation Site				
	APH	CHI	FTE	GOR	PED
APH	0.800	0.903	0.560	0.938	0.904
CHI	0.775	0.900	0.917	0.645	0.754
FTE	0.655	0.833	0.957	0.633	0.672
GOR	0.775	0.916	0.902	0.725	0.852
PED	0.741	0.866	0.821	0.707	0.806

a) Level 1

Training Site	Evaluation Site				
	APH	CHI	FTE	GOR	PED
APH	0.523	0.587	0.360	0.733	0.606
CHI	0.606	0.672	0.835	0.500	0.539
FTE	0.491	0.606	0.847	0.419	0.428
GOR	0.557	0.573	0.680	0.531	0.619
PED	0.409	0.590	0.684	0.621	0.571

b) Level 2

Table 2 Spectral shape classification results

Training Site	Evaluation Site				
	APH	CHI	FTE	GOR	PED
APH	0.719	0.960	0.074	0.700	0.784
CHI	0.540	0.919	0.900	0.604	0.666
FTE	0.063	0.109	0.611	0.126	0.109
GOR	0.829	0.905	0.339	0.716	0.879
PED	0.735	0.765	0.447	0.800	0.819

a) Level 1

Training Site	Evaluation Site				
	APH	CHI	FTE	GOR	PED
APH	0.515	0.539	0.050	0.524	0.475
CHI	0.366	0.634	0.463	0.311	0.315
FTE	0.029	0.047	0.253	0.047	0.081
GOR	0.347	0.500	0.279	0.476	0.475
PED	0.149	0.415	0.300	0.500	0.508

b) Level 2

Table 3 Maximum likelihood classification results

Tables 2 and 3 summarize the classification results over the five sites: Ft. A.P. Hill (APH), Chickahominy (CHI), Ft. Eustis (FTE), Gordonsville (GOR), and Prince Edwards State Forest (PED). Table 2 gives the spectral shape classification results at levels 1 and 2, respectively. Each entry gives an overall classification accuracy (fraction correct) for the spectral shape classifier developed over the training site (row) and measured over the evaluation site (column); i.e., the classification file derived from the training site was used to classify the imagery of the evaluation site. Table 3 shows the maximum likelihood classification results at levels 1 and 2, respectively, provided for comparison purposes. Here the training statistics (class means and covariance matrices) from the training site were used to classify the imagery of the evaluation site. The 95% confidence intervals based on 60 points are the values in the tables ±10%, approximately.

	Spectral Shape Classifier		Maximum Likelihood Classifier	
Combinations	Level 1	Level 2	Level 1	Level 2
Within Study Areas	0.84	0.63	0.76	0.48
Between Study Areas	0.79	0.57	0.57	0.31
Within Wetland Sites	0.90	0.74	0.63	0.35
Within Upland Sites	0.81	0.57	0.76	0.44
Within Wetland and Upland Sites	0.84	0.63	0.73	0.41
Between Wetland and Upland Sites	0.76	0.54	0.47	0.27
Full Scene	0.86 (est.)	0.65 (est.)	0.73 (est.)	0.41 (est.)

Table 4 Overall performance of classifiers

Table 4 summarizes the average performance of the spectral shape and maximum likelihood classifiers within and between different groups of sites. The first row gives the overall accuracies of the two classifiers trained and evaluated over the same sites (these are the averages of the entries along the main diagonals in Tables 2 and 3). For 5x60 = 300 points, the 95% confidence intervals are the above values ±5%, approximately. The performance of the spectral shape classifier is thus not significantly different from that of the maximum likelihood classifier at level 1 but does appear to be significantly better at level 2. However if we leave out the FTE site which has thin cloud cover, neither level 1 nor level 2 results are significantly different.

The second row in Table 4 gives the overall accuracies of the two classifiers trained over one study area but ev aluated over a different area (these are the averages of the off-diagonal entries in Tables 2 and 3). This provide s an indication of the signature extendibility of the classifiers over the scene. For the five sites there are 20 po ssible combinations of different training and evaluation sites. The performance of the spectral shape classifier is significantly better than that of the maximum likelihood classifier at levels 1 and 2. The spectral shape class ification accuracies averaged over the 12 combinations of different training and evaluation sites that exclude F TE are 0.81 and 0.58 (levels 1 and 2, respectively). The corresponding maximum likelihood accuracies are 0.7 6 and 0.41. If we leave out the FTE site, the level 1 results are not significantly different but the spectral shape classifier accuracy is significantly better than that of the maximum likelihood classifier at level 2.

The third row gives the overall accuracies of the two classifiers trained and evaluated over wetland sites. The le vel 1 and 2 accuracies of the spectral shape classifier are significantly better than the corresponding accuracies of the maximum likelihood classifier. This is due largely to the FTE site which contains thin clouds that adve rsely affect the maximum likelihood classifier. In the uplands sites the accuracy of the spectral shape classifie r is significantly better than that of the maximum likelihood classifier at level 2 but is comparable at level 1 (fo urth row in Table 4).

The average accuracy of the spectral shape classifier within wetland and upland sites is 0.84 (level 1). Betwee n wetland and upland sites it falls slightly to 0.76 (level 1). (The within wetland and upland cases involve all c ombinations in which the training and evaluation site are either both wetland or both upland. The between wetl and and upland cases involve all combinations in which the training site is wetland and the evaluation site is u pland, or vice versa.) The average accuracy of the maximum likelihood classifier within wetland and upland sit es is 0.73 (level 1), but between wetland and upland sites falls dramatically to 0.47 (level 1).

We then constructed a spectral shape classifier for the full scene by merging classification files from each stu dy area. In cases where spectral shapes were assigned different classes in different sites, the class with the hig hest probability across sites was selected and assigned to the spectral shape globally. The resultant classificati on file was then used to classify the full scene. The overall accuracies of the full-scene spectral shape classific ation were 0.85 and 0.65 at levels 1 and 2, respectively.

7. Summary

A new method of classifying multispectral imagery based on a set of binary features that represent the relative values between spectral bands was described and its performance over a full Landsat scene evaluated. The ov erall classification accuracy was found to be comparable to that of a maximum likelihood classifier over indivi dual sites. However between sites the spectral shape classifier out performed the maximum likelihood classifi er. A full-scene spectral shape classifier was developed by combining classification files from the five sites int o a single classification file. The accuracy of the resultant classifier tested over the five sites was significantly better than the accuracy of a stratified maximum-likelihood classifier assuming wetlands and uplands strata. P reliminary results suggest that the spectral shape representation can provide a greater degree of signature exte ndibility but at the expense of reduced discrimination at level 2.

Additional work in several areas is either underway or planned. Further testing over other areas is being perfo rmed in order to obtain a more representative estimate of the accuracy of the spectral shape classifier. Other m ethods for combining classification files derived from individual training sites and applying them to the full sc ene are being explored that retain regional class-cluster relationships. Finally, experiments are planned to mea sure the extent to which the spectral shape classifier is sensitive to topographic effects.

References

T. Bara, *Definitions for Proposed EMAP-LC Landcover Classes and Structure*, Private Comm.

M.J. Carlotto, V.T. Tom, P.W. Baim, and R.A. Upton, "Knowledge-based multispectral image classification," *Proceedings SPIE*, Vol. 504, San Diego, CA, 1984.

M.J. Carlotto and V.T. Tom, "Surface material classification based on spectral shape," *Proceedings of IGARS S*, Amherst, MA Oct, 1985.

M.J. Carlotto, "Towards the unsupervised interpretation of outdoor imagery," *J. Visual Communication and I mage Representation,* Vol. 1, No. 2, Nov. 1990, pp 158-175.

M.D. Craig, "Minimum volume transforms for remotely sensed data," *IEEE Trans. Geoscience and Remote Sensing*, Vol. 32, No. 3, May 1994.

R.E. Crippen, "The regression intersection method of adjusting image data for band ratioing," *Int. J. Remote S ensing*, Vol. 8, No. 2, pp 137-155, 1987.

E.P. Crist and R.C. Cicone, "A physically-based transformation of thematic mapper data - The TM Tasseled Cap," *IEEE Trans. Geoscience and Remote Sensing*, Vol. 22, No. 3, May 1984.

R.G. Henderson, "Signature extension using the MASC algorithm," *IEEE Trans. Geoscience Electronics*, Vo l. GE-14, No. 1, Jan. 1976.

K. Konstantinides and J.R. Rasure, "The Khoros software development environment for image and signal pro cessing," *IEEE Trans. Image Processing*, Vol. 3, No. 3, 1994.

P.N. Salter, "Survey of multispectral image systems for earth observations," *Remote Sensing of Environment*, Vol. 17, pp 85-102, 1985.

E. Schanda, *Physical Fundamentals of Remote Sensing*, Springer-Verlag, Berlin Heidelberg, 1986.

R.W. Sjoberg and B.K.P. Horn, "Atmospheric effects in satellite imaging of mountainous terrain," *Applied O ptics*, Vol. 22, No. 11, pp 1702-1716, 1983.

W.J. Todd, D.G. Gehring and J.F. Haman, "Landsat wildland mapping accuracy," *Photogrammetric Enginee ring and Remote Sensing*, Vol. 46, No. 4, April 1980, pp 509-520.

J.T. Tou and R.C. Gonzalez, *Pattern Recognition Principles*," Addison-Wesley, Reading MA,1974.

Acknowledgments

Thanks to Greg Heberle and Ernie Carroll for registering the Landsat TM and M7 imagery, and for assisting in the accuracy assessment experiments, and to Mark Nebrich for developing the Khoros implementation of t he classification system.