Adaptive Threshold for Spectral Matching of Hyperspectral Data

J. Schwarz, and K. Staenz

Canada Centre for Remote Sensing
588 Booth Street, Ottawa, Ontario, Canada K1A OY7
phone: (613) 947-1249, fax: (613) 947-1383,
email:john.schwarz@ccrs.nrcan.gc.ca

SUMMARY

Spectral matching is one of several techniques that derives information which may be used to classify hyperspectral data. In this paper an analysis technique is presented which supplements the typical spectral matching algorithms and facilitates the exploitation of the information they provide. This technique may be used to adaptively set a threshold on the similarity measure and classify spectra. This adaptive threshold may be set automatically, or the information from this analysis may be used to permit the manual selection of a threshold in a manner which is more intuitive than the direct specification of a similarity measure. The technique adapts the threshold to each class and similarity measure combination.

This technique is demonstrated by applying the Spectral Angle Mapper (SAM) to simulated and actual imaging spectrometer data. The proposed technique does not provide additional information about the similarity of two spectra, but it does provide information which is valuable for deciding class membership. Test results using Compact Airborne Spectrograhic Imager (*casi*) data indicate that the adaptive threshold technique provides classification accuracy similar to that achieved using the interactive manual threshold scheme, provided that the training area is representative of the target variation. While analyzing the variability of a class, in the similarity space, it became obvious that analysis of the similarity measure for the entire scene has the potential of providing valuable information for classification purposes.

1. INTRODUCTION

Classification of hyperspectral data can be computationally very expensive but the increased signature detail also permits effective use of spectral matching techniques for classification purposes. These procedures typically compare a spectrum to a reference spectrum, using similarity measures to determine the level of similarity between the two spectra. Although images of similarity values can be used on their own, many users select to set a threshold and use it to decide if a spectrum came from the same target type

as the reference spectrum came from. This procedure of using matching techniques and thresholds emulates the functionality of classifiers with good performance and lower computational cost. Eventhough, some users may become competent at selecting thresholds for particular applications (Kruse, 1988; Fan et al., 1997), it typically remains a difficult task.

There are many spectral matching schemes, several of which have been implemented on the Imaging Spectrometer Data Analysis System (ISDAS) (Staenz et al., 1998). One of the more popular techniques is the Spectral Angle Mapper (SAM) (Kruse et al., 1993), which possesses the characteristic of being insensitive to scale differences between two spectra. There are two common methods of obtaining reference spectra; one is from the imagery itself and the other is from a spectral library. Usually the spectral matching procedure is provided with only the reference spectrum for a given class. The proposed thresholding technique also requires information about the variability of the reference spectrum, which can be derived from a collection of spectra from the same target class. This information of spectral variability is used to calculate an appropriate threshold value. Even if the automatically determined threshold is not used, the information on class variability provides the user with a more intuitive parameter for selecting a threshold.

Using spectral matching to perform classification can in general be described as a two-step process. First the n-dimensional data space is mapped to a 1-dimensional similarity space. Secondly, the decision of class membership is made based on this similarity value. The technique proposed in this paper utilizes additional information which is not typically used by spectral matching techniques.

There are other existing procedures which can perform classification, and have features that are similar to the one proposed in this paper, namely the Fisher linear discriminant, single-class classification, and partial unmixing. A short review of these procedures are presented to provide a comparison of related analysis techniques. Subsequently, the new thresholding procedure for spectral matchers is presented, followed by results of applying this technique to simulated data sets and Compact Airborne Spectrographic Imager (*casi*) data (Anger et al., 1996).

2. CLASSIFICATION TECHNIQUES

2.1. Fisher Linear Discriminant

The poor computational efficiency of classification algorithms applied to high dimensional spectral data, may be overcome by projecting the n-dimensional spectral data (feature space) onto a single dimension via a linear transformation column vector W of length n as follows:

$$y = W^T X, \tag{1}$$

where X is the spectrum in the original feature space, represented as a column vector of length n, and y is a point in 1-dimension. The Fisher linear discriminant is the projection W which provides the maximal separation of two classes (Duda et al., 1973).

This method allows for classification of the data using the traditional criteria of probability of membership. However, the estimates of the probability of membership are performed in a 1-dimensional space rather than the n-dimensional space of the data. The Fisher linear discriminant selects the n-D to 1-D linear transformation which has the maximum classification capabilities.

2.2. Single-Class Classification (Single Hypothesis Test)

In multi-class classification the objective is to determine to which of several classes a test object belongs. A single-class classifier as described in Fukunaga (1990) uses only the information about a single-class to determine if the test object (pixel) is a member of a specific class or not. The class information is used to generate an estimate of the probability of membership. The estimate is compared to a predefined threshold and subsequently membership is decided.

The single-class classifier may be summarized as mapping from n-dimensions (spectral space) to 1-dimension, followed by a decision of membership being determined by comparing the test object's projection to some preset threshold.

2.3. Partial Unmixing

Partial unmixing as described by Boardman et al. (1995) uses the spectrum of a single reference material and knowledge about the variability of the data within a scene to estimate the abundance of the reference material within each of the pixels. A relative abundance measure provides more information than just the assignment of a target class to the pixel, although it may be used to perform such a classification. This is similar to the situation in statistical classification where the probability of membership provides more information, but it is used to generate the determination of class. Some researchers such as Endsley (1995) have tried to use probability of membership estimates to provide abundance measurements.

2.4. Spectral Matching

Spectral matching techniques typically use a single reference spectrum to compare it to a test spectrum and provide a measure of the similarity between the two spectra. Thus, the data is mapped from a n-dimensional space to a 1-dimensional similarity space. There are some spectral matching algorithms that vary from this description; specifically those which use more than just the information from the reference spectrum as in the Mean Squared Error Statistic (MSES) (Staenz et al., 1996) which uses the per band noise, and the Cross Correlogram Spectral Matcher (CCSM)

(van der Meer et al., 1997) which provides more than a single measure of the similarity of the spectra.

There are several popular matching techniques in use, one of which is the SAM. This technique measures the angle subtended between the reference spectrum $\vec{X}_{ref}$ and the test spectrum $\vec{X}$ and is specified by (Kruse et al. 1993):

$$SAM = \arccos\left(\frac{\vec{X}_{ref} \bullet \vec{X}}{\left|X_{ref}\right\|X\right|}\right). \qquad (2)$$

This and other similarity measures can be used to perform classification of unknown spectra. The spectral matching techniques are very similar to statistical classification procedures, in so much as they both can be used to determine class membership and they both can operate in single-class and multi-class variants. A multi-class spectral matching uses reference spectra from several classes and classification may be determined by choosing the class with the best similarity measure. A single-class spectral matching uses only a single reference spectra and class membership may be determined by comparing the calculated similarity measure to a threshold selected by the user.

3. SEMI-AUTOMATIC THRESHOLD DETERMINATION

3.1. Method

As mentioned in the previous section, spectral matching techniques can take on forms similar to statisical classifiers. However, statistical classifiers use knowledge of a representative class member (such as the class mean) and the class variability to determine class membership, while spectral matching techniques typically only use knowledge of a representative class member. This technique of automatic threshold determination extends the similarity between spectral matching and statistical matching, by using knowledge of the class variability in spectral matching. However, this technique varies from the statistical classifier in that the class is characterized in the 1-dimensional similarity space and not the n-dimensional spectral space.

Parametric classifiers assume the distribution of members, in a class, can be modelled, similarily this technique assumes that the similarity measures, for a class, can be modelled by a statistical distribution. The implementation in ISDAS specifically assumes that similarity measures, such as the SAM or the Euclidian distance, are the result of summing a large number of independent terms. Thus, the distribution of the terms derived from spectra from the same class will behave like a normal distribution for a sufficiently large number of terms.

Within ISDAS, the class information is extracted from a user specified training set of pixels. The first step in the analysis is to determine a reference spectra X_{ref} for the class, which is defined as the mean spectrum of the training set. The next step compares

the spectra in the training set to the mean spectrum, using the operator selected spectral matching techniqe, and obtains the similarity measures S_X. In the third step the mean, $\overline{S}$, and the standard deviation, σ_{class}, of these similarity measures are calculated and then used to describe the position and distribution of the class, in the similarity space. These four parameter can be summarized as:

$$\begin{aligned} X_{ref} &= Mean[\overset{\omega}{X}], \\ S_X &= Similarity[\overset{\omega}{X}, X_{ref}], \\ \overline{S} &= Mean[S_X], and \\ \sigma_{class} &= \sqrt{Var[S_X]}\ . \end{aligned} \tag{3}$$

In ISDAS, single-class classification is performed using the similarity measure, generated by a spectral matching technique to compare the class reference spectrum to a spectrum being tested, and using a cutoff threshold as the discriminant function. A one-tail test is used, since the other tail indicates a better match with the reference spectrum and, therefore, should not be rejected. The default threshold T_{def} used in ISDAS is calculated as follows:

$$T_{def} = \overline{S} - 3\sigma_{class}\ . \tag{4}$$

The user may also set the threshold interactively T_{user}. In the interactive setting of T_{user} the user has two options. The first option is to use the information about the class distribution in similarity space and specify the threshold value as the number of standard deviations from the mean m_{user}, as follows:

$$T_{user} = \overline{S} - m_{user}\sigma_{class}\ , \tag{5}$$

and the second options is to explicity set the value of the similarity measure T_{user} directly.

The ISDAS spectral matching tool has implemented both the single-class and the multi-class classification schemes. The latter approach uses the threshold as the minimum similarity requirement for classification, and performs the same functions that a rejection option does in a statistical classifier.

3.2. Simulations

The distribution of the similarity measure depends on many factors, so simulations were conducted to gain an understanding of the distribution behaviour. In each simulation a set of spectra are generated where all the spectra belong to a single class. To generate these spectra a reference spectrum X_{ref} was defined. It was used as the mean spectrum of the class, then additional members from the class were generated about the reference spectrum. A series of simulations were conducted where the properties of the class spectra were varied. The first characteristic which was varied was the signal-to-noise ratio (SNR) of the spectra, and the second was the number of spectral bands each spectra contained.

The mean spectrum for the class was given a base value of 100 and had a multiple-period sine wave, with an amplitude of 50, added to it. The class vectors X_{class}

were generated by adding a random Gaussian noise to each spectral band of the mean spectrum, where the noise distributions for each band were identical and independent. Thus, forming a hyper-sphere with a Normal radial density distribution about the mean spectrum which is described by:

$$X_{class} = X_{ref} + X_{random}(\sigma_{signal}) \tag{6}$$

and

$$X_{random}(\sigma_{signal}) = [N_1(0,\sigma_{signal}), N_2(0,\sigma_{signal}),..., N_n(0,\sigma_{signal})]^T, \tag{7}$$

where $N_i(0,\sigma_{signal})$ are independent normal distributions with zero means and same standard deviation of σ_{signal} of the i'th band. Note, that all n spectral bands have the same standard deviation σ_{signal}.

The simulations used the spectral angle for the similarity measure. The simulated class spectra were used to generate the distributions described in section 3.1 and the parameters X_{ref}, $\overline{S}$ and σ_{class} were calcuated. The study generated and analyzed a set of simulated classes, where the classes were changed by varying the number of spectral bands, n, making up the spectrum, and the magnitude of the noise in each spectral band σ_{signal}.

Since the implementation of the thresholding technique assumes the data in the similarity space follows a known distribution, specifically a Normal distribution, the skew and kurtosis of the generated distributions were also calculated. Figure 1 shows a typical distribution of spectral angles, which was generated from a population of 10,000 spectra and a σ_{signal} of 10 for spectra containing 100 bands. Note that the mean similarity measure (mean spectral angle) is offset from the zero, and that the skew and kurtosis are small which supports our assumption for the simulated class. The thresholding technique uses the mean and standard deviation derived from a population, consisting of members from a single target type, to determine a cutoff spectral angle. If the assumption of normality holds, then the error of omission is very small. However, no estimate of co-mission errors can be made, since the simulations are performed on a single-class at a time.

The behaviour of the mean spectral angle and the variance are illustrated in Figures 2 and 3, respectively. In these simulations the nature of the distributions are reported for the various sizes of the reference spectra (number of bands in the spectra) and various class cluster sizes (relative size of the data cloud created by the signal noise). These parameters were chosen since the relative sizes and the class cluster width and the number of bands in the reference spectrum determines the range covered in the spectral angle domain. It should be noted that as more bands are added to the simulated spectra, the length of the reference spectrum is effectively increased, since the average value of each band is constant. Also as the size of the noise per band is increased, the size of the class cluster increases. The SNR per band is defined as 100 / σ_{signal}.

3.3. Field Data

Hyperspectral data used to evaluate and demonstrate this automatic thresholding technique, was collected with *casi*. The sensor was flown over an agricultural setting in southern Manitoba, on July 25, 1996 during the peak biomass season. The flight was conducted at an altitude of 2,745 m above sea level using the spectral data acquisition mode. This provided data consisting of 304 pixels per line and 96 contiguous, 6.8 nm wide, spectral bands per pixels, covering a wavelength range between 458 nm to 1004 nm, with a ground resolution of 4 m x 4 m. The only processing of the data was the calibration of the raw data to at-sensor radiance.

The fields at the test site contained: sugarbeets, corn, wheat, barley, beans, canola, potato, sunflower, oats, peas, and alfalfa. A couple of weeks prior to data collection some of the fields in the site area experienced blow-down, but had fully recovered by the time of data collection. There had been heavy rain, during the week of the data collection. This was particular harmful to the bean crops. Standing water had obviously affected the bean crops earlier in the season as well. Some sections of bean fields were completely bare of beans and the regions around it were areas of dwarfed growth.

The scene analyzed at this site contained four crop types: canola, beans, wheat, and barley. The Modified Spectral Angle Mapper (MSAM) which like the other techniques implemented on ISDAS provides similarity measure values from 0 to 1, where 1 indicates the best fit, was applied to the data set. The definition of the MSAM is given by:

$$MSAM = 1 - \frac{2 \cdot SAM}{\pi}. \tag{8}$$

For each of these crop types an area of the image was selected, and the pixels within this region were used as the training set for that crop class. These training sets were analyzed in the same manner as the simulated data, where the MSAM was calculated for each member of the training set with the training set mean spectrum being used as the class reference spectrum.

The image in Figure 4 shows this agricultural study area. Also shown in the image is a yellow rectangle which indicates a region used as the training area for the canola crop. The pixels highlighted in red indicate those which have been classified as canola, using the MSAM spectral matching technique and the default threshold value. Note, that this training area is not one of the large or the small training areas discussed later in the paper, it is use only as typical example and to generate the histograms in Figures 5 and 6.

Figure 5 is a frequency histogram of the MSAM values for the canola training set shown in Figure 4. Figure 6 is a frequency histogram of the MSAM values when the canola reference spectrum is compared to the entire scene. This histogram indicates the existence of at least 4 types of spectrally different target. Since the spectra in the training

set are spatially close it is reasonable to believe that the training set does not represent the full variability of the class. This is obvious by comparing Figures 5 and 6 where the training set's has a narrower histogram than that of the entire canola population. In addition, the histogram in Figure 6 does not go to zero for MSAM values in the region of 0.975, indicating that some members of at least one other class can not be unambiguously classified using the MSAM measure with only the canola reference.

As the above example indicates, the nature of the training area reflects how well the spectral information of the class is represented. To examine the importance of selecting training sets which accurately reflect a class, two sets of training areas where selected for two of the crop types. One small training area and containing little information on the class variability and one large training area including the majority of the crop variability.

Table 1 contains the results of analyzing the shape of the distributions of MSAM for small training sets of each class. The table also contains the default threshold value (equation 4) calculated from the training sets. Note how the mean MSAM and its standard deviation are class dependent.

Consider the single class classification using the MSAM technique where a specific threshold of 0.99 was used. Theoretically, this technique should not classify even half of the pixels correctly. While a threshold value of 0.98 would result in good performance for canola, beans and barley but still perform poorly for wheat. Lowering the threshold to 0.97 would correctly classify most of the wheat pixels but would likely result in large co-mission errors for canola, beans and barley. Optimal classification results will be achieved by adapting the threshold for each class.

Single-class classification was performed for two of the crop types, canola and beans, in the agricultural scene shown in Figure 4 and the results are summarized in Table 2. Two types of training sets, a small and a large training region, were used for each of the two crop types. The small regions are characterized by few pixels, but more importantly by the fact that the pixels were of the most vigorous and uniform portion of the field. The large training regions include the most vigorous crop growth and some of the less vigorous growth. Pixels of bare soil were excluded from the large training regions, but the average canopy closure in the large training sets would be lower than in the small training sets.

The similarity measure used for the classification was the MSAM. The single-class classification was performed using six different threshold settings. The first four settings were preset fixed values. The fifth threshold used was the default adaptive threshold given in equation (4). The last threshold was determined interactively, by manually adjusting the threshold value and observing which pixels had been classified. This final value used the operators judgement to select a threshold level which balanced the omission and co-mission errors. This manual method is used as an estimate of the optimum performance of the spectral matching technique.

By comparing the performance of the adaptive and the manual threshold setting techniques, it is clear that as more of the class information is used in the training set that the adaptive technique approaches the performance limit. In fact, the adaptive threshold out performed the operator when the large training area was used to classify beans. It is also worth noting that the small training sets did not provide enough information about spectral variability for these agricultural targets.

Figures 7 and 8 illustrate the differences in the populations for the small and large training set for canola. The solid line indicates the distribution of all the pixels in the scene compared to the reference spectrum generated from the corresponding training set. These curves are very similar since the reference spectrum from the large and small training set are very similar. However, the distribution of the MSAM values for pixels in the large training sets are significantly broader than the distribution for the small training set. The distribution of the MSAM values for pixels classified using the automatic threshold level demonstrates the same behaviour as the training set's distributions. These plots contain much of the information that is presented in the confusion matrices. However, they provide the information in a manner that allows for a visual evaluation of the classifier performance. It can be seen in Figure 8 that the distribution of selected pixels stops near the center of the value of the two populations which are shown in the entire scene distribution. The distribution of selected pixels in Figure 7 obviously stops short of including the entire canola population.

4. DISCUSSION

The proposed spectral matching technique assumes that the class distribution, in the similarity space, may be parameterized and the ISDAS implementation assumes it is Gaussian. The shape measures from the simulations and *casi* data demonstrate that these assumptions are reasonable and practical. This ability to characterize the class' distribution, in the similarity space, permits the spectral matching calculations to generate estimates of probability of membership. Thus making single-class spectral matching a statistical classification procedure. This procedure is similar to both single-class classification and the Fischer linear discriminant. Here the estimate of the probability of membership is performed in a 1-dimensional similarity space, onto which the data has been mapped by the similarity measure.

Unlike the Fisher linear discriminant, this technique is not limited to linear transformation. Although it is most appropriate for similarity measures which involve the sum of many terms with little correlation between terms, such as the dot product of two spectra, and thus are more likely to demonstrate a Normal distribution.

This adaptive threshold technique requires class specific information. Such a priori knowledge of class distribution is not too difficult to obtain when analyzing scenes like the agricultural scene used in this study, but it could be a serious problem in an application such as mineral exploration. In applications where a priori knowledge is difficult to acquire, the knowledge may be obtainable from a spectral library. Obtaining information about class distribution from a spectral library and the use of an adaptive

thresholding mechanism would improve the utility of spectral matching in applications like mineral exploration.

Although the analysis presented in this paper is focused on characterizing the class variability in the similarity space to make spectral matching a statistical classifier, it should be noted that the examination of frequency histograms such as the one in Figure 6 may be of considerable value. Analysis of the distribution of similarity values for the entire scene requires no specific knowledge of class distributions but provides additional information on selecting thresholds for classification. Since this procedure only requires a reference spectrum it can be applied to all spectral matching cases where the threshold can not be adaptively determined using class information.

The objective of the proposed technique is not to improve the information content of similarity measures but to provide an aid to the interpretation process. Therefore, the classification accuracy inherent in any spectral matching procedure is not improved. However the selection of a classification threshold is facilitated. This technique adapts to the nature of the data and provides a more intuitive interpretation of any threshold value when expressed as a multiple of σ_{class} . It has also led to the idea of utilizing the similarity measure distribution of the entire scene as a valuable source of information.

5. CONCLUSIONS

An adaptive threshold technique which is based on training area statistics has been presented for use in combination with spectral matching techniques to classify spectra. The technique has been demonstrated using the modified spactral angle mapper (MSAM) on simulated data and hyperspectral *casi* imagery collected over an agricultural site in Manitoba, Canada.

- The applied threshold (mean reference spectrum minus three standard deviations) results in a similar classification accuracy as achieved with a manual setting of the threshold (e.g., 90.2 % versus 93.4% for canola), provided that the training area covers the within-class variation.
- The performance of the adaptive threshold technique depends more strongly on the training area statistics than the manual approach. However it out performs fixed threshold levels.
- The analysis of the simulated and *casi* data indicated that the assumption of normal class distribution is reasonable and meets the requirements of the adaptive threshold technique.
- The number of bands and the SNR affects the adaptive threshold such that the mean similarity measure increases mainly with increasing SNR while the standard deviation decreases with increasing number of bands and SNR.
- The adaptive threshold can be determined from the image data themselves or can be obtained from a spectral library, provided that the required statistics is available.
- The adaptive threshold technique provides an aid to the selection of a classification threshold, but this approach does not improve the classification accuracy inherent in any spectral matching procedure.

6. REFERENCES

Anger, C. D., Mah, S., Ivanco, T. A., Achal, S. B., Price, R., and Busler, J. R., 1996, "Extended Operational Capabilities of the casi", *Proceedings of the Second International Airborne Remote Sensing Conference and Exhibition*, San Francisco, U.S.A., Vol. 1, pp. 124-133.

Boardman, J. W ., Kruse, F. A., and Green, R. O., 1995 , "Mapping Target Signatures via Partial Unmixing of AVIRIS Data", *Summaries of the Fifth JPL Airborne Earth Science Workshop*, Pasadena, U.S.A., Vol. 95-1, pp. 23-26.

Duda, R. O., and Hart, P. E., 1973,"Pattern Classification and Scene Analysis", *John Wiley & Sons, Inc.*, Toronto, Canada.

Endsley, N. H., 1995, "Spectral Unmixing Algorithms based on Statistical Models", *Proceedings of the SPIE Conference on Imaging Spectrometry*, Vol. 2480-03, pp. 23-36.

Fan, L., van der Meer, F., and Bodechtel, J., 1997, "Geological Information Extraction from Airborne Hyperspectral Data and Field Spectral Data in a Heavily Soil and Vegetation Covered Area", *Proceedings of the Third International Airborne Remote Sensing Conference and Exhibition*, Copenhagen, Denmark, Vol. I, pp. 665-672.

Fukunaga, K., 1990, "Introduction to Statistical Pattern Recognition; Second Edition", *Academic Press, Inc.*, San Diego, U.S.A.

Kruse, F. A., 1988, "Use of Airborne Imaging Spectrometer Data to Map Minerals Associated with Hydrothermally Altered Rocks in the Nothern Grapevine Mountains, Nevada, and California", *Remote Sensing of Environment*, Vol. 24, pp. 31-51.

Kruse, F. A., Lefkoff, A. B., Boardman, J. W., Heidebrecht, K. B., Shapiro, A. T., Borloon, P. J., and Goetz, A. F. H., 1993, "The Spectral Image Processing System (SIPS) - Interactive Visualization and Analysis of Imaging Spectrometer Data", *Remote Sensing of Environment*, Vol. 44, pp. 145-163.

Staenz, K., Schwarz, J., and Cheriyan, J., 1996, "Processing / Analysis Capabilities for Data Acquired with Hyperspectral Spaceborne Sensors", *ACTA Astronautica*, Vol. 39, No. 9-12, pp. 923-931.

Staenz, K., Szeredi, T., and Schwarz, J., 1998, "A system for Processing/Analyzing Hyperspectral Data", Canadian Journal of Remote Sensing, Vol. 24, No. 2, pp. 99-113.

van der Meer, F., and Bakker, W., 1997, “CCSM: Cross Correlogram Spectral Matching”, *International Journal of Remote Sensing*, Vol. 18, No. 5, pp. 1197-1201.

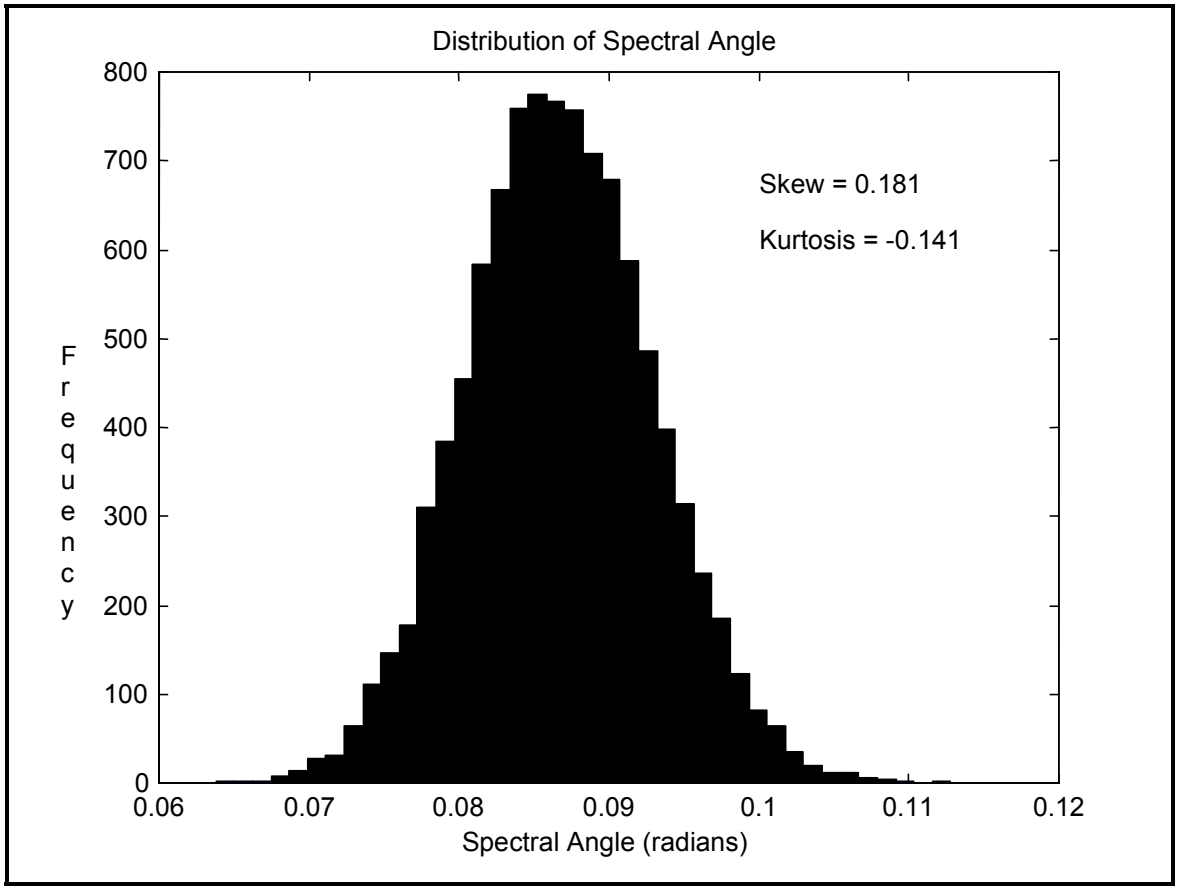

Figure 1: Sample distribution of spectral angle for spectra of a simulated class.

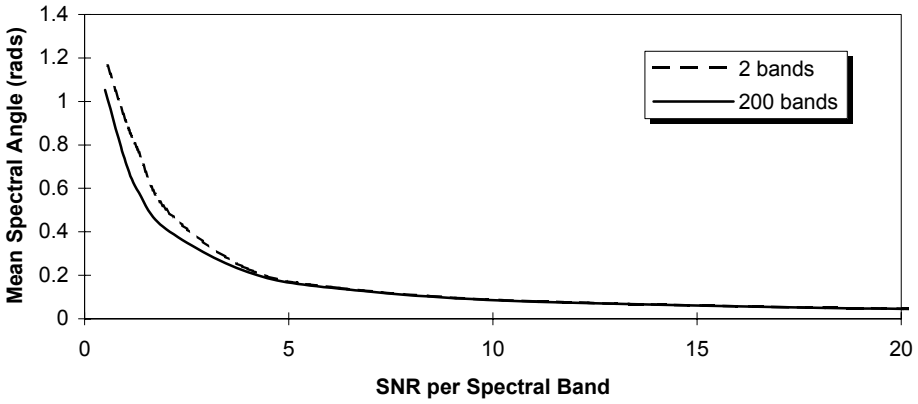

Figure 2: Simulation Results: The mean spectral angle, of the population distribution, versus the signal-to-noise ratio (SNR) added to the signal of the population.

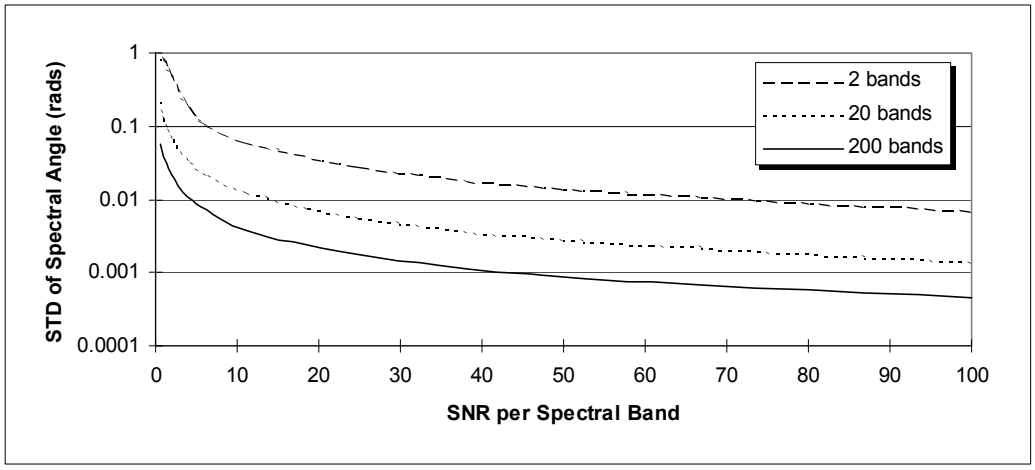

Figure 3: Simulation Results: Standard deviation (STD) of spectral angle for a population versus the signal-to-noise (SNR) of the signal on the population spectrum.

Figure 4: Agricultural scene with canola training set identified by a yellow rectangle and pixels classified using the default threshold value identified by red pixels. The RGB image was generated using bands centered at 676.3 nm for red, 540.5 nm for green, and 478.8 nm for blue.

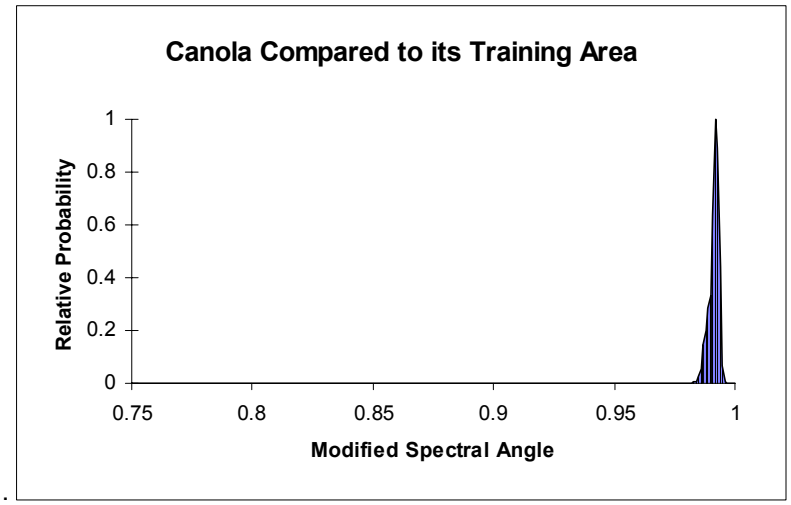

Figure 5: Distribution of modified spectral angle mapper (MSAM) values generated from *casi* data for the canola training set in comparision with its own mean spectrum. The training set is identified in Figure 4 by a yellow rectangle.

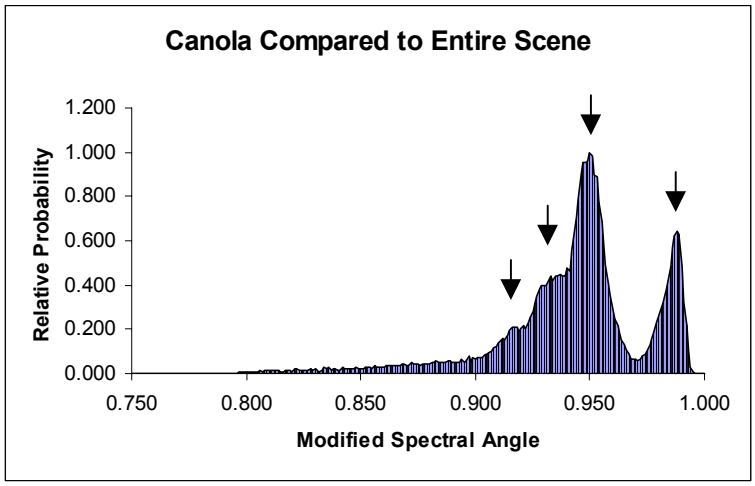

Figure 6: Distribution of modified spectral angle (MSAM) values generated by comparing canola's mean spectrum to the entire agricultural scene of *casi* data. The cluster of MSAM values closest to unity consists predominantly of spectra from the canola fields. Arrows indicate the peak of the four spectrally different populations.

Crop Type	Mean MSAM	STD * MSAM **	Skew	Kurtosis	Threshold	# Pixels
Canola	0.9916	0.00236	-1.043	0.842	0.9846	799
Wheat	0.9847	0.00419	-0.724	0.205	0.9721	432
Beans	0.9898	0.00308	-1.168	2.246	0.9805	480
Barley	0.9899	0.00323	-2.331	9.553	0.9801	308

- STD = standard deviation
- MSAM = modified spectral angle mapper

Table 1: Results of analyzing shape of small training sets of four crop types. The training sets consisted of the most vigorous grow in the fields.

Training ROI	Threshold Value	Percent Correct	Omission Error (%)	Co-mission Error (%)
Small Canola	.995	6.3	93.7	0
	.990	38.3	6.17	0
	.980	85.6	14.4	.013
	.970	96.9	3.1	.95
	Automatic (.9899)	39.1	60.9	0
	Manual (.9700)	96.9	3.1	.95
Large Canola	.995	5.5	94.5	0
	.990	39.9	60.1	0
	.980	85.9	14.1	.017
	.970	96.2	3.8	2.1
	Automatic (.9777)	90.2	9.8	.123
	Manual (.9730)	93.4	6.6	.95
Small Bean	.995	3.1	96.9	0
	.990	22.7	77.3	0
	.980	52.9	47.1	.178
	.970	-	-	>10
	Automatic (.9897)	23.8	76.2	0
	Manual (.9800)	52.9	47.1	.178
Large Bean	.995	3.8	96.2	0
	.990	28.0	72.0	0
	.980	57.9	42.1	.302
	.970	-	-	> 10
	Automatic (.9772)	62.3	37.7	1.61
	Manual (.9780)	60.9	39.1	1.34

*ROI = region of interest.

Table 2: Results of analyzing the training sets of two crop types, using two different sized training sets and six classification thresholds. The two large training sets include more variability of crop vigor than the small training sets.

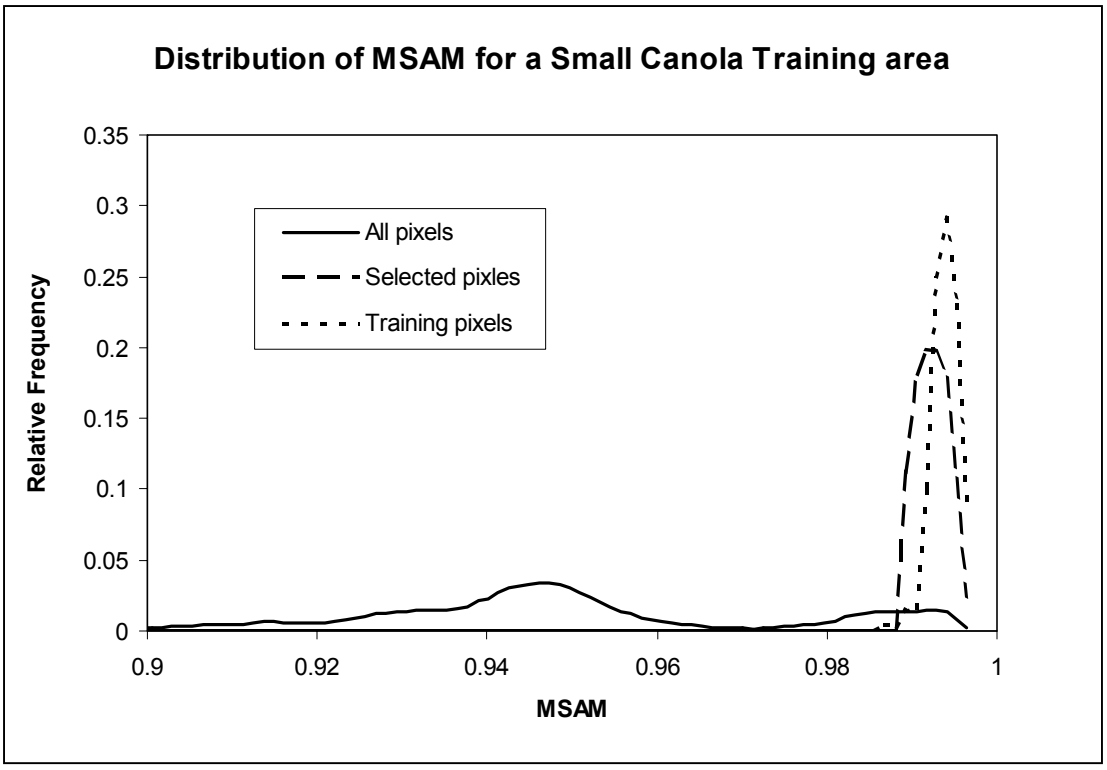

Figure 7: Distribution of the similarity measure (MSAM) when compared to the reference spectrum derived from a small training sample of most vigorous canola growth. The automatic thresholding technique was used to determine the selected pixel group.

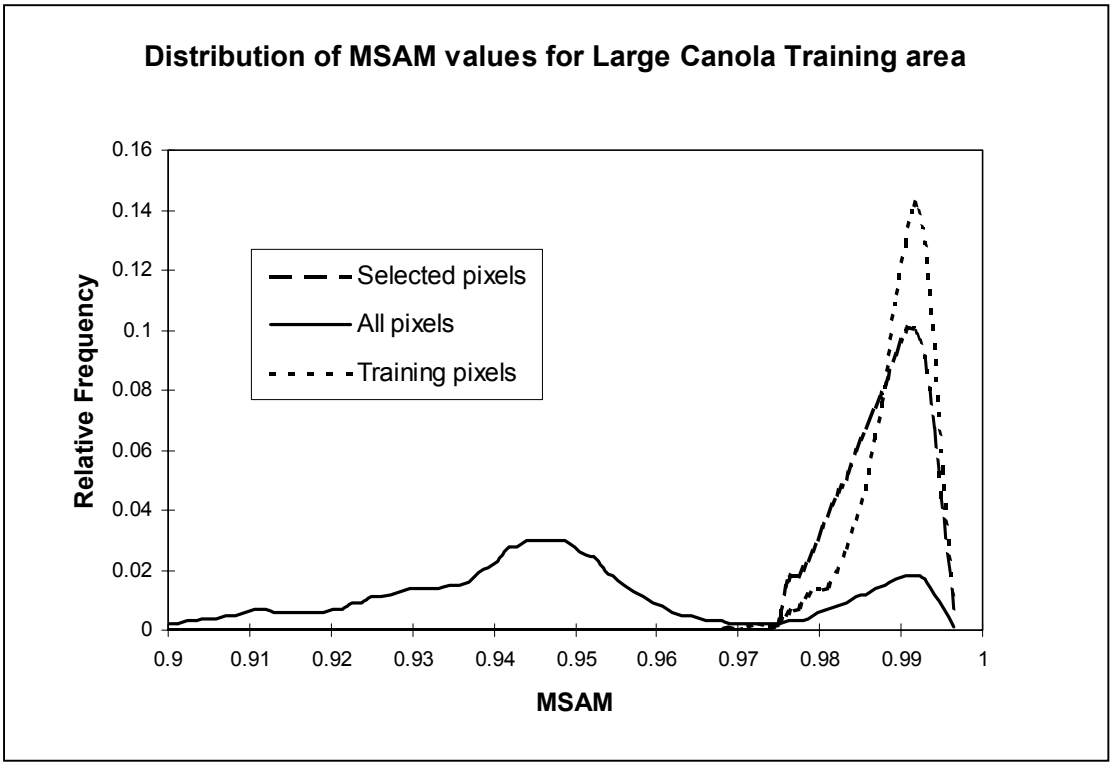

Figure 8: Distribution of the similarity measure (MSAM) when compared to the reference spectrum derived from a large training sample of canola, consisting of a range of vigor. The automatic thresholding technique was used to determine the selected pixel group.