Application of spatial resolution enhancement and spectral mixture analysis to hyperspectral images

Harry N. Gross and John R. Schott

Rochester Institute of Technology, Center for Imaging Science
Rochester, NY 14623

ABSTRACT

Spatial resolution enhancement and spectral mixture analysis are two of the most extensively used image analysis algorithms. This paper presents an algorithm that merges the best aspects of these two techniques while trying to preserve the spectral and spatial radiometric integrity of the data.

With spectral mixture analysis, the fraction of each material (endmember) in every pixel is determined from hyperspectral data. This paper describes an improved unmixing algorithm based upon stepwise regression. The result is a set of material maps where the intensity corresponds to the percentage of a particular endmember within the pixel. The maps are constructed at the spatial resolution of the hyperspectral sensor. The spatial resolution of the material maps is then enhanced using one or more higher spatial resolution images. Similar to the unmixing approach, different endmember contributions to the pixel digital counts are distinguished by the endmember reflectances in the sharpening band(s). After unmixing, the fraction maps are sharpened with a constrained optimization algorithm.

This paper presents the results of an image fusion algorithm that combines spectral unmixing and spatial sharpening. Quantifiable results are obtained through the use of synthetically generated imagery. Without synthetic images, a large amount of ground truth would be required in order to measure the accuracy of the material maps. Multiple band sharpening is easily accommodated by the algorithm, and the results are demonstrated at multiple scales. The analysis includes an examination of the effects of constraints and texture variation on the material maps. The results show stepwise unmixing is an improvement over traditional unmixing algorithms. The results also indicate sharpening improves the material maps.

The motivation for this research is to take advantage of the next generation of multi/hyperspectral sensors. Although the hyperspectral images will be of modest to low resolution, fusing them with high resolution sharpening images will produce a higher spatial resolution land cover or material map.

KEYWORDS: Image Fusion, Sharpening, Spectral Mixing, SIG, Stepwise Regression

1. SPECTRAL UNMIXING

Spectral Mixture Analysis [1,2,3] was developed to address the problem with classifying mixed pixels. The authors recognized the large scale of low resolution images creates situations where the measured digital counts are due to many materials within the corresponding ground areas. Multispectral and hyperspectral imaging sensors provide the ability to extract spectral signatures from mixtures of materials.

Spectral unmixing transforms digital counts, or radiance, into fractions of the constituent materials, called endmembers. The primary outputs of spectral mixture analysis are "fraction images" in which the intensity can be made proportional to the percentage of that particular endmember. This gives a spatial mapping of the endmembers.

Spectral mixture analysis models the total radiance (reflectance) in a pixel as the linear combination of reflectance from each endmember. This assumption is reasonable when there are no intermaterial interactions, and when the image has been calibrated and corrected for atmospheric distortions (e.g. by the methods described by [4] or [5]). For the i^{th} spectral band, the model is

$$LRXS_i = \sum_{e=1}^{n} R_{i,e} f_e + \varepsilon_i \qquad (1)$$

0-8194-2209-6/96/$6.00

where $LRXS_i$ are the i[th] band digital counts from the Low Resolution Multi Spectral data converted to reflectance units, $R_{i,e}$ are the apparent endmember reflectances from a stored library, f_e are the unknown endmember fractions, and ε_i are residual errors arising from errors in the endmember library, as well as unmodeled higher order effects.

Accurate spectral unmixing requires careful implementation.[3] Unmixing spectral curves is an ill-posed problem.[6] There is no unique set of materials that combine to match the low resolution data. In color science, this same characteristic is known as metamerism. As a result, the entire unmixing process must be guided by significant user knowledge. An iterative approach is used, stopping when the number of endmembers is reasonable and the estimation residual is reduced to the level of image noise.[2] A priori knowledge is essential to make the number of endmembers in the library reasonable. Choosing endmembers *for each pixel* is a subset selection problem.

2. SUBSET SELECTION

Consider a general linear predictive model

$$y = \xi_0 + \sum_{i=1}^{n} A_i \xi_i + \varepsilon, \tag{2}$$

where y is linearly related to n predictor variables $A_1, A_2, \ldots, A_n$ through unknown coefficients $\xi_0, \xi_1, \ldots, \xi_n$. Assume specific realizations of this model are corrupted by random deviations, ε, which are independently sampled from a distribution with zero mean and variance σ^2. In matrix form, the process model is

$$y = A\xi + \varepsilon. \tag{3}$$

For unmixing, the bias term, ξ_0, should be small, so it is convenient to set it to zero. We desire to estimate the unknown coefficients ξ, with estimates x, which then are used to predict y

$$\hat{y} = Ax. \tag{4}$$

When there are $m > n$ measurements, the unknown coefficients can be estimated via an over-determined least squares (LS) procedure.[7] The solution is

$$x = (A'A)^{-1} A'y. \tag{5}$$

Traditional unmixing fixes the number of endmembers to use throughout the entire image. This section describes an approach which can be used to select the appropriate endmembers on a pixel by pixel basis. By only including a subset of the endmember library in the regression model, a more robust prediction results. When excess endmembers are included (e.g., to map more material classes), the extra terms cause overfit in the model. Conversely, if one uses just a few endmembers, the image-wide fractions may be improved, but fewer materials can be mapped.

To unmix each pixel, one must identify which n predictors (endmembers) are the proper terms for the model. Consider the case where one has a library of L potential endmembers, but where a priori knowledge indicates the actual number of terms in the model should be $n << L$. In other words, we may have a library of materials which we expect to find in the image, but not all of them will be found in every pixel. The problem is made more difficult because n is unknown. One must determine the proper subset of endmembers to use.[8]

Stepwise regression is based upon an Analysis of Variance (ANOVA) calculation of the "Extra Sum of Squares."[7] The ANOVA explains how much of the total variance is attributed to the regression model, and how much is due to the random residuals. In this analysis, one compares an n-term model with an $(n-1)$-term model to calculate the benefit of adding the additional term. Define the reduced term model by

$$\hat{y} = Wz; \quad z = (W'W)^{-1} W'y \tag{6}$$

where z is an $(n-1)$-vector and W is an m x $(n-1)$ matrix. The Sum of Squares (SS) due to regression with n degrees of freedom (dof) can be divided into a SS due to the reduced order model with $(n-1)$ dof and the SS due to the extra term in the model with a single dof. The ANOVA structure is shown in Table 1. The MS column in the table signifies Mean Square variables. They are calculated by dividing the SS by the corresponding dof.

Source	dof (uncorrected)	SS	MS
SS reduced model	n-1	z'W'y	
SS extra term	1	x'A'y - z'W'y	MS(extra term)
Due to Regression	n	x'A'y	
About Regression (residual)	m-n	y'y - x'A'y	MS(residual) ≈ s^2
Total	m	y'y (uncorrected)	

Table 1: Extra Sum of Squares ANOVA Table

The benefit of the ANOVA table becomes apparent when one also makes distributional assumptions about the errors. If $\varepsilon \sim N(0,I\sigma^2)$ (gaussian, zero mean, and independent), then the sum of m squared errors has a chi-square distribution with m dof, written χ^2_m.[9] If the regression model is a good one, then the errors should be chi-square distributed. If the model is poor, the errors will not be chi-square. Thus, a hypothesis test can examine the errors, and determine, to a specified confidence level, whether the proposed model is adequate. The hypothesis test uses the relationship that the ratio of two chi-square variables divided by their respective degrees of freedom has an F-distribution with the corresponding degrees of freedom.

The stepwise regression equation is tested by defining a null hypothesis that the coefficient for the extra term is zero. The alternate hypothesis is that the term is nonzero, and thus required for the model.

$$\begin{aligned} H_0&: \quad \xi_n = 0 \\ H_1&: \quad \xi_n \neq 0 \end{aligned} \tag{7}$$

Consider the distribution of two specific random variables, SS(extra term) and SS(residual). If the errors are gaussian, the sum of squares variables are χ^2 distributed. They are normalized by their respective degrees of freedom to create the mean square variables. The ratio MS(extra term)/MS(residual) follows an $F_{1,m-n}$ distribution as long as $\xi_n = 0$. This ratio can be compared to an F-statistic with 1 and m-n degrees of freedom, at a specified confidence level. If the MS ratio is greater than the F-statistic, we reject the null hypothesis and conclude the more complex model is required. If the MS ratio is smaller than the F-statistic, we would not reject the simpler model, which does not include the extra term.

A stepwise regression algorithm using a sequential F-test is used to add and remove candidate terms (endmembers) from the model. First, all possible one endmember models are formed. The model which best explains the measured digital counts (in reflectance units) will have the largest MS(regression), the smallest MS(residual), the smallest squared error, and the largest F-to-enter ratio. If F-to-enter is greater than a threshold, the endmember is entered into the model. Then, all the two endmember models which contain the entered material are examined. F-to-enter values are calculated for the (L - 1) candidate variables remaining in the library. If the largest F-to-enter is greater than the threshold, the appropriate endmember is added.

Once three variables have been entered into the model, it is possible that one may be removed. For example, although the first variable may have been highly correlated with the measurements, once the second and third variables have been entered, the first may become redundant. F-to-remove values are calculated for each of the entered endmembers as though each was the last to enter the regression. If the smallest F-to-remove is less than a threshold, that variable is removed from the model.

Although based upon an F-statistic, it is not necessary to use a value from an F-table. Since, the ratio used in the test is only approximately F distributed when the model is correct, a constant threshold value is used. The threshold for entering and removing variables are set equal to each other to prevent cycling.[10] When a variable is neither entered nor removed, the algorithm terminates.

One danger with the automated stepwise algorithm is that it may generate unrealistic models. A sanity check is performed on the per pixel fractions to flag inappropriate results. For example, large positive or negative fractions indicate a model which could not possibly be correct. Once bad results are flagged, an alternate model can be calculated. The alternate model is selected by temporarily removing from the library the most abundant material from the poor model, and rerunning the stepwise selection procedure. This forces the solution into a markedly different set of materials. Since the first result was so

poor, a significantly different solution is desired. The combination of a stepwise regression, checking the fractions, and repeating the regression if necessary yields excellent results in a reasonable run time.

Measuring unmixing performance is difficult. Spectral reflectances of real world objects exhibit tremendous variability. Temporal effects are especially evident in organic materials. Illumination and view angle are only partially compensated for by a shade endmember. These and other effects result in variation from the library endmember spectra. This variation is manifested as errors in the material maps. It is very difficult to quantify the results of a spectral unmixing algorithm since one requires detailed ground truth.

To control the environment for the algorithm development reported here, synthetic image generation tools are used to completely define the images. Because the data are synthetically generated, the underlying unmixed solution is known. Additionally, perfectly unmixed images can be created to use in testing the sharpening process. This separates the unmixing and sharpening steps, and lets the effects of each be individually analyzed.

3. IMAGE FUSION

Typical image fusion algorithms generate high resolution images of digital counts in the various spectral bands, e.g. [11, 12]. However, digital counts do not reveal what kind of object (material type) is in the scene. Spectral unmixing maps the materials, but to date has only been done at low resolution. The desired output product is often a high resolution material map. Such a product would identify materials, and locate them to high accuracy. The contribution of this research is a method of using spectral mixing tools at high spatial resolution. By integrating spectral mixing and image fusion, a high resolution material map is attained.

The overall image fusion algorithm is a two stage process depicted in Figure 1. A set of multi/hyperspectral images is represented by an image cube. The cube is unmixed to form material maps. Image fusion is accomplished by using one or more sharpening bands to increase the spatial resolution of the material maps.[13] The images are assumed to be perfectly registered. This study did not examine the degradations due to misregistration.

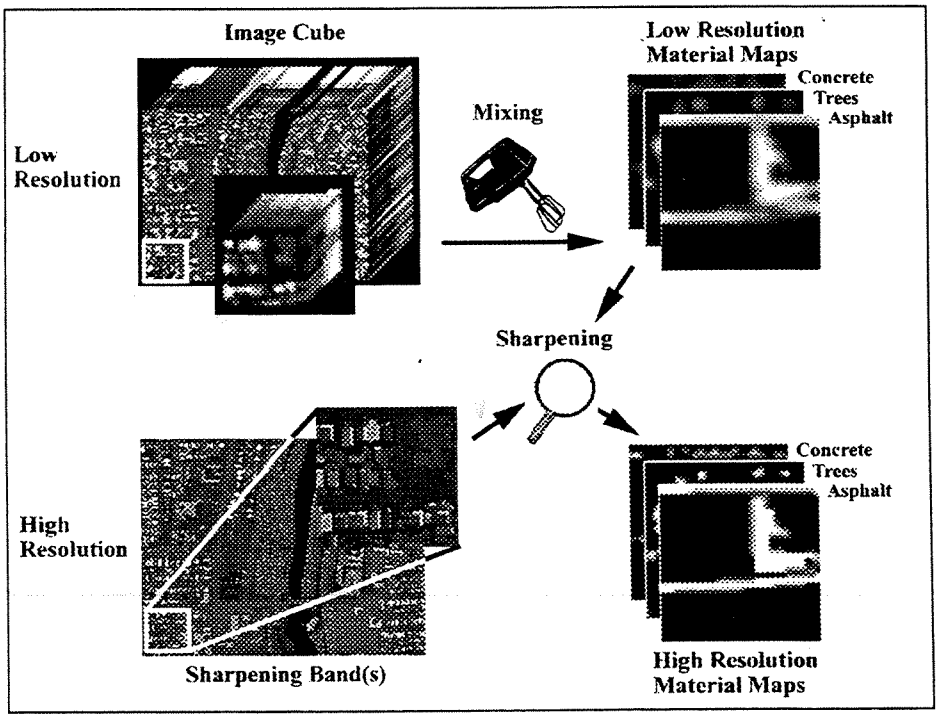

Figure 1: Image Fusion Data Flow – Creating High Resolution Material Maps

4. SHARPENING

Given a set of low resolution fractions produced by the unmixing process described above, the next step is to spatially locate the fractions to the resolution of a sharpening band(s). The difficulty here is that there are many more unknowns than equations. Let s be the number of high resolution (sub)pixels in a single multispectral (super)pixel, and constrain the subpixels to contain the same n materials as the superpixel. Then, there are s equations and $n*s$ unknowns. Sharpening is solved as an *under*-determined least squares problem.

The sharpening model takes the spectral unmixing form. The function to be minimized is the residual error in the digital counts of the sharpening band pixels contained in the superpixel (calibrated into reflectance units as in the spectral mixing algorithm),

$$HRP_j = \sum_{e=1}^{n} R'_{pan,e} f_e^j + \varepsilon_j; \quad j = 1 \ldots s \tag{8}$$

where, HRP_j is the j^{th} element in of the High Resolution Panchromatic image corresponding to subpixel j, $R'_{pan,e}$ contains the library values for the reflectances of the appropriate endmembers in the sharpening band(s), and f_e^j is the high resolution fraction for the e^{th} endmember in the j^{th} subpixel fraction vector.

We desire the values of f_e^j that minimize equation (8) while maintaining consistency with the low resolution results. Consistency requires for each material, the high resolution fractions must average to the low resolution fraction,

$$\frac{1}{s}\sum_{j=1}^{s} f_e^j = f_e, \quad e = 1 \ldots n. \tag{9}$$

If the high resolution measurements are equal, corresponding to no high resolution spatial information, the high resolution fractions should equal the low resolution fractions. Thus, sharpening an aggregate* mixture should have no effect. If there is spatial information contained in the high resolution data, the sharpening algorithm adjusts the fractions to return the minimum length vector with the least square residual error. Material maps of areal mixtures can be improved by sharpening.

5. CONSTRAINTS

The preceding discussion presented spectral unmixing and sharpening as unconstrained problems. However, the underlying pixel represents a true mixture of one or more materials. Mixture problems are solved by constraining the coefficients.

The unmixing literature contains examples using each of the constraint strategies. Some investigators use unconstrained unmixing and screen out unrealistic solutions. Others use partial constraints by requiring the fractions to sum to unity. The fully constrained approach is used less often. Applying inequality constraints makes the algorithm more complex. The fully constrained case is also likely to be susceptible to spectral variation between different material samples and between the samples and the spectral library.

Define three different constraint conditions. The first is *unconstrained*, where the fractions can take on any value needed to minimize the residual error. However, the unconstrained sharpening problem is defined to include the consistency constraints in equation (9).

The second condition is called *partially constrained*. Here, the fractions within a pixel are required to sum to unity. The sum is taken over all the materials in the appropriate pixel. For unmixing, there is one constraint, while for sharpening, this provides s equality constraints. However, only (s - 1) of the constraints are independent.

$$\sum_{e=1}^{n} f_e = 1 \qquad\qquad \sum_{e=1}^{n} f_e^j = 1, \quad j = 1 \ldots s \tag{10}$$

---

* *Aggregate* is used to define a mixture where reflectance combines linearly, but the materials are not spatially separable. An *areal* mixture is a separable linear mixture. Nonlinear mixtures are termed *intrinsic*.

A *fully constrained* set would also require each individual fraction to lie between zero and one. There are $2*n$ inequality constraints for unmixing and $2*n*s$ inequality constraints for sharpening.

$$0 \le f_e \le 1 \qquad\qquad 0 \le f_e^j \le 1, \;\; j = 1 \ldots s \tag{11}$$

The inequality constraints are not all independent, but an optimization algorithm only applies the *active* inequality constraints during any search. Problems containing inequality constraints, such as the fully constrained problem, require an iterative solution.

Although fully constrained fractions seem intuitive for a mixture problem, values outside the bounds have some physical meaning. If the library reflectance values are derived from an "average" material response, fractions greater than one and deviations in the "negative" direction from other endmembers are expected. An algorithm may need to recognize small negative fractions and fractions slightly greater than one as acceptable, while preventing solutions with unreasonable fractions.

If the coefficients in a model are constrained, one must solve a *restricted least squares* problem. In particular, we are interested in linear, equality and inequality constraints. A constrained fusion algorithm is discussed in [14]. For a more general reference, the reader is directed to [15].

One question is how to incorporate the constraints into the stepwise regression algorithm. The literature is particularly silent on this point. Two options were considered. In the first, each time coefficients are calculated, they are subjected to the appropriate constraints. However, since the ANOVA analysis is intended for unconstrained regression, this likely violates some of the underlying assumptions. The second option is to use an unconstrained stepwise algorithm to determine the endmembers to be included in the model for each pixel. Then, the constraints are applied to those endmembers, and the coefficients recalculated. The second method yielded better results.

6. RESULTS

This section describes the quantitative results obtained from unmixing and sharpening a synthetic test image. To separate the sharpening step from unmixing, the synthetic truth data were used to generate "perfectly unmixed" material maps. Sharpening perfect maps quantifies the effect of the sharpening stage independently of the unmixing step. The results also quantify the performance of the stepwise regression unmixing algorithm. It is compared to traditional unmixing, where the same endmembers are used throughout the image. Finally, image fusion results are shown which combine the unmixing and sharpening algorithms.

6.1 Synthetic Imagery Characteristics

The algorithm development work capitalizes on the benefits of synthetic image generation (SIG). SIG scenes can be built with varying degrees of complexity. The ability to incrementally increase realism gives feedback to algorithm designers. They may test their designs under increasingly difficult conditions, and modify the designs to incrementally improve robustness. Secondly, since SIG is entirely computer created from a defined data source, the underlying "truth" is known. Algorithms can be evaluated under various conditions, where their performance can be quantified and compared to alternate techniques. Only after the algorithm is shown to work under simulated conditions is it then tested on real imagery. Preliminary results using a small image scene are described in [14].

A synthetic test image was generated with the DIRSIG algorithm.[16] The bands were selected to simulate the Environmental Research Institute of Michigan (ERIM) M-7 airborne line scanner. The M-7 sensor has fifteen bands in the visible and infrared spectral regions. Data for three sharpening bands were also generated. To save processing a new image, representative M-7 bands at higher spatial resolution were used to sharpen. The visible sharpening band was a weighted average of M-7 bands 4 and 6. This represents a visible panchromatic band from 460 - 720 nm. Band 9 (760 - 1045 nm) was used as a near infrared (NIR) sharpening band, while band 13 (1400 - 1890 nm) served as a short wave infrared (SWIR) sharpening band. The nominal ground sample distance corresponding to a pixel in the base image was 1 meter.

Band four (460 - 620 nm) of the test image is shown in Figure 2. The region contains deciduous forest, grass, a dirt road, a small lake, and several tanks and trucks. The vehicles consist of various materials, including camouflage paints, painted steel, canvas, and rubber. A spectral library containing ten materials was used for unmixing.

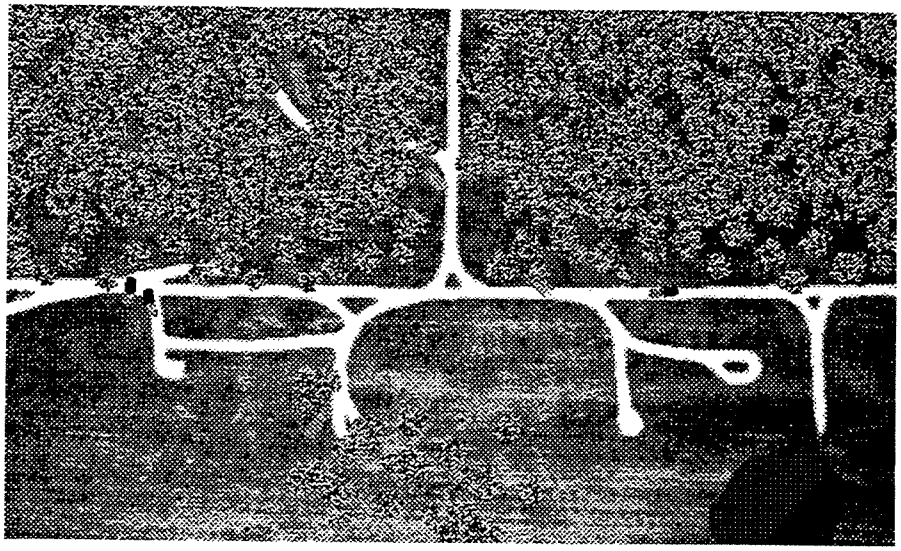

Figure 2: Band 4 (460 - 620 nm) of Synthetic Test Image

The DIRSIG algorithm uses a ray tracer. For each pixel in the image, a ray is cast to the appropriate point on the ground. By definition, each ray strikes a single material. Thus the original image contains only *pure* pixels. For this scene, the original scale was set at 1 meter per pixel (m/p). Aggregating (averaging) to a larger scale results in an image containing a smaller number of pixels for the same ground area. However, many of these averaged pixels are *mixtures* of several materials.

It is necessary to create fraction maps at a consistent scale in order to compare the results. Using averaging, images of decreased spatial resolution in the appropriate spectral bands were created at 2, 4, 8, 16, and 32 m/p. For example, the 4 m/p aggregation is a convenient high resolution scale. At this scale, the truth data contain fractions of materials in multiples of one sixteenth. Starting with the 1 m/p image cube and a 1 m/p material map, test image cubes, sharpening bands, and truth maps can be created at any lower resolution. Figure 3 is a block diagram showing how the synthetic image set is created. For clarity, only a few of the scales and data sets are shown. The material truth maps at any resolution are used to calculate "perfectly unmixed" material maps. These are used as inputs to the sharpening algorithm, allowing its performance to be measured independently of unmixing.

In Figure 3, the vertical axis indicates resolution. The unmixing operation (represented by the mixer icon), transforms image cubes into material maps in a horizontal (same scale) direction. Sharpening, represented by a magnifying glass icon, operates down, increasing the resolution of the material maps. Different combinations of data, operations, and reference fractions (perfectly unmixed maps) are used to analyze the algorithms. Fusion is shown as unmixing and then sharpening. Sharpening performance at different scales is represented by taking perfectly unmixed maps through the sharpening algorithm. Finally, multiband sharpening is indicated by the labels inside the magnifying glass.

Creating the reference data sets illustrated in Figure 3 without SIG would require extensive ground surveys and perfect registration of image data with ground reflectance samples to ensure pure pixels.

Some of the perfectly unmixed material maps at the 4 m/p scale are shown in Figure 4. Recall the intensity (gray level) in a material map is proportional to the fraction of that material at each pixel location. At this scale, there are 17 gray levels possible, corresponding to fractions between zero and one.

To capitalize on the truth data available from the SIG imagery, a single error metric is used to compare the results. A squared error (*SE*) per pixel was calculated for each set of fraction maps,

$$SE = \frac{1}{N} \sum_{pixels} \sum_{materials} \left(f_{truth} - f_{test} \right)^2 \tag{12}$$

The summation over the pixels includes the entire region of interest under study (N pixels). The material summation is over the entire library of materials. The relative magnitude of this error metric gives a convenient measure of the match between the test fractions and the truth fractions. All errors of commission and omission are penalized by this metric.

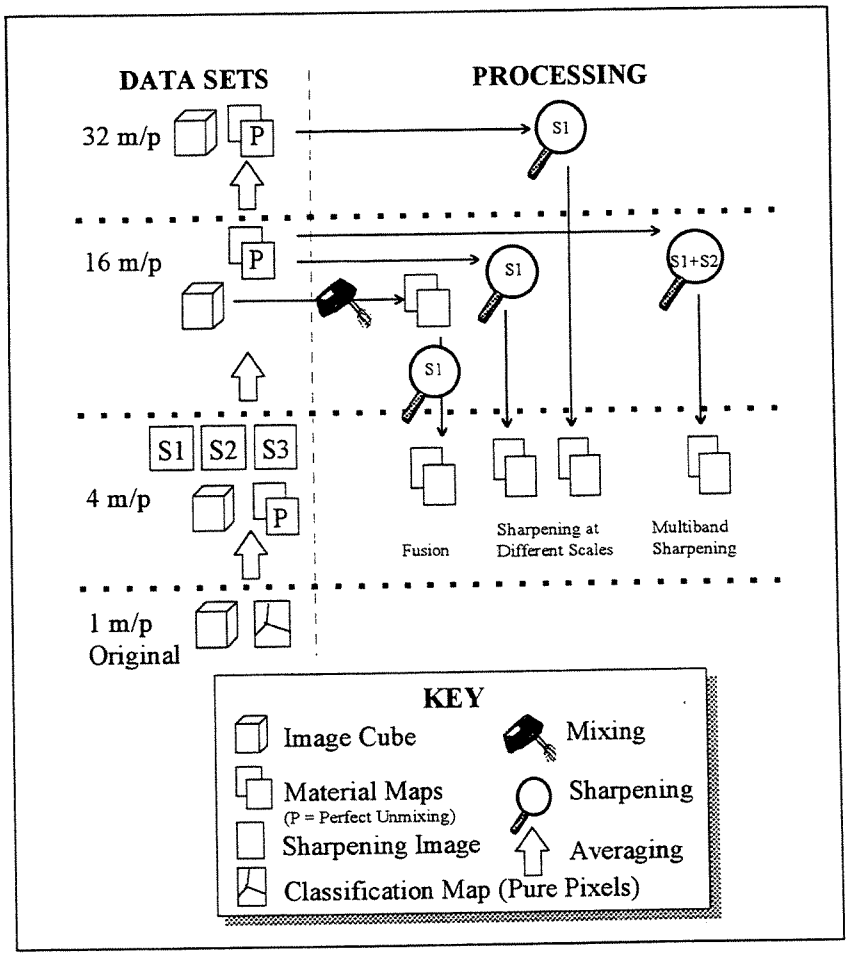

Figure 3: Creating SIG Data Sets

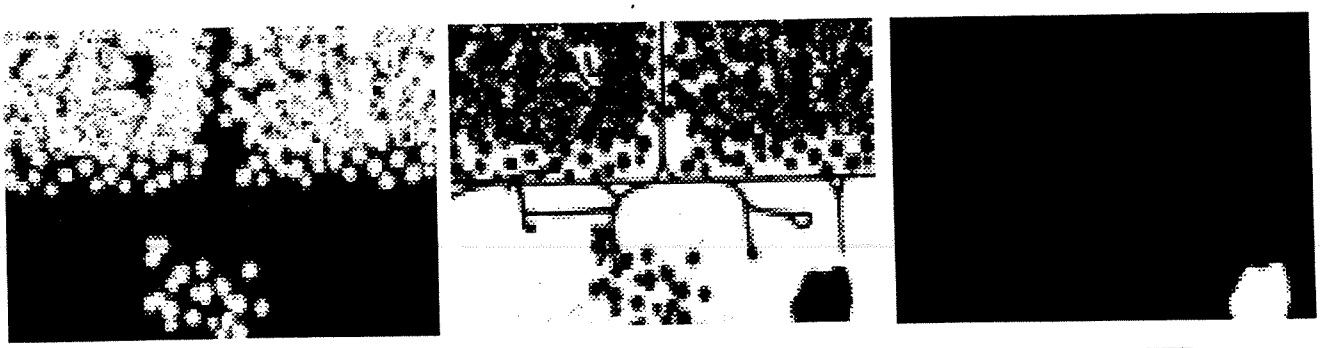

DECIDUOUS TREES **GRASS** **LAKE WATER**

Figure 4: Perfectly Unmixed Material Maps (4 m/p)

By definition, the perfectly unmixed data sets have no error. However, there is an opportunity cost due to not sharpening. A method is required to quantify the effect of leaving the material maps at low resolution. Pixel *replication* was used to compare maps of different scales. In pixel replication, the fractions of each material are repeated for each subpixel. For example, to replicate from a 16 m/p scale to a 8 m/p scale, each 16 meter pixel fraction is replicated into a 2 x 2 array of 8 meter pixels. Each of these four values can then be compared to the reference values at 8 m/p. If the pixel is pure at the 16 meter scale, replication has no effect. However, if the material concentrations are spatially separable at 8 m/p (an areal as opposed to an aggregate mixture), replicating results in error. This is the error due to *not* sharpening.

6.2 Unmixing Results

Stepwise regression based unmixing selects the endmembers on a pixel by pixel basis. Traditional unmixing uses the same set of endmembers throughout the entire scene. By including excess terms in the model, the traditional approach can induce errors in the fractions. Unmixing with the wrong number of endmembers increases the error. Since the number of materials in each pixel varies across the scene, the stepwise regression algorithm offers an improvement.

Figure 5 compares unmixing results at the 8 meter scale for three different algorithms. In the first, traditional unconstrained unmixing is applied with the full 10 endmember library. In the second, four dominant endmembers are chosen to reduce the number of unnecessary coefficients. The third algorithm is the proposed unmixing scheme based upon stepwise regression. In the third algorithm, the number of endmembers is adjusted at each pixel. The stepwise algorithm selected the appropriate endmembers from a library of ten materials. Results from all three algorithms are compared to the perfect material maps at 8 m/p without replication. The algorithms were run unconstrained to be consistent with the traditional unmixing implementation in the ENVI software package.[17]

The results show the improvement of the proposed unmixing scheme over traditional methods. If one unmixes the entire scene with ten endmembers, the excess materials cause severe overfit in the fractions. This results in poor performance relative to ground truth. Note that 10 endmembers results in the *lowest* RMS error due to the model fitting the noise. RMS is not a good indicator of overall unmixing performance. (Spatial patterns in an RMS image, however, can indicate regions of unexplained materials.) Unmixing only the trees, grass, dirt, and water considerably reduces the error. However, by doing this, one is unable to classify the vehicles. The stepwise algorithm reduces the overfit, constructs material maps for all materials in the image, and results in a considerable improvement in unmixing performance.

As a prelude to including sharpening, Figure 6 shows unmixing at 16 m/p and replication to 4 m/p. Note here the perfect unmixing error corresponds to replication error. Constrained unmixing was done according to option 2 discussed above. The endmembers are selected via unconstrained stepwise regression. Then the coefficients are recalculated with the appropriate constraints.

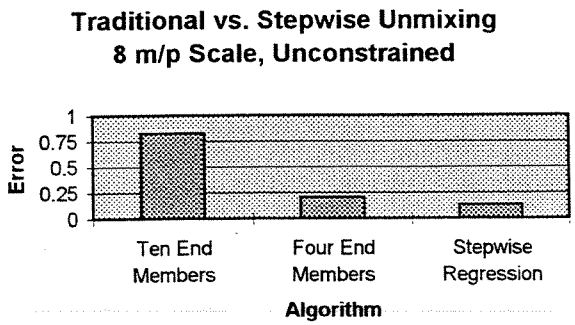

Figure 5: Traditional vs. Stepwise (Per Pixel) Unmixing

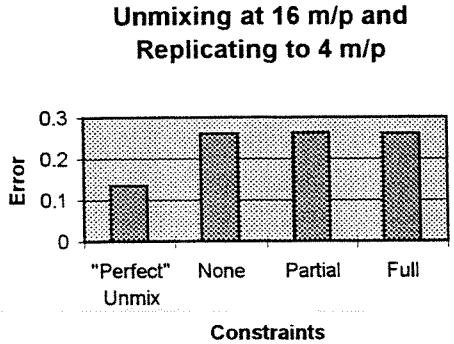

Figure 6: Unmixing and Replication to a Higher Spatial Resolution

From the figure, one sees there is some error inherent in the unmixing algorithm due to mismatches between the image and the library. The additional error due to pixel replication is the potential improvement which can be obtained by sharpening.

However, if unmixing selects the wrong endmembers, that error is unrecoverable as sharpening takes the unmixing solution at its starting point.

6.3 Fusion Results

In this section, unmixing and sharpening are combined in an image fusion algorithm. First, stepwise unmixing is applied to the low resolution image cube to make low spatial resolution material maps. Then the maps are sharpened with high resolution sharpening band(s) to give high resolution material maps. The previous section already showed stepwise unmixing is better than the traditional approach. The following results demonstrate that sharpening improves the material maps even further. Example high resolution fraction maps from the image fusion algorithm are shown in Figure 7.

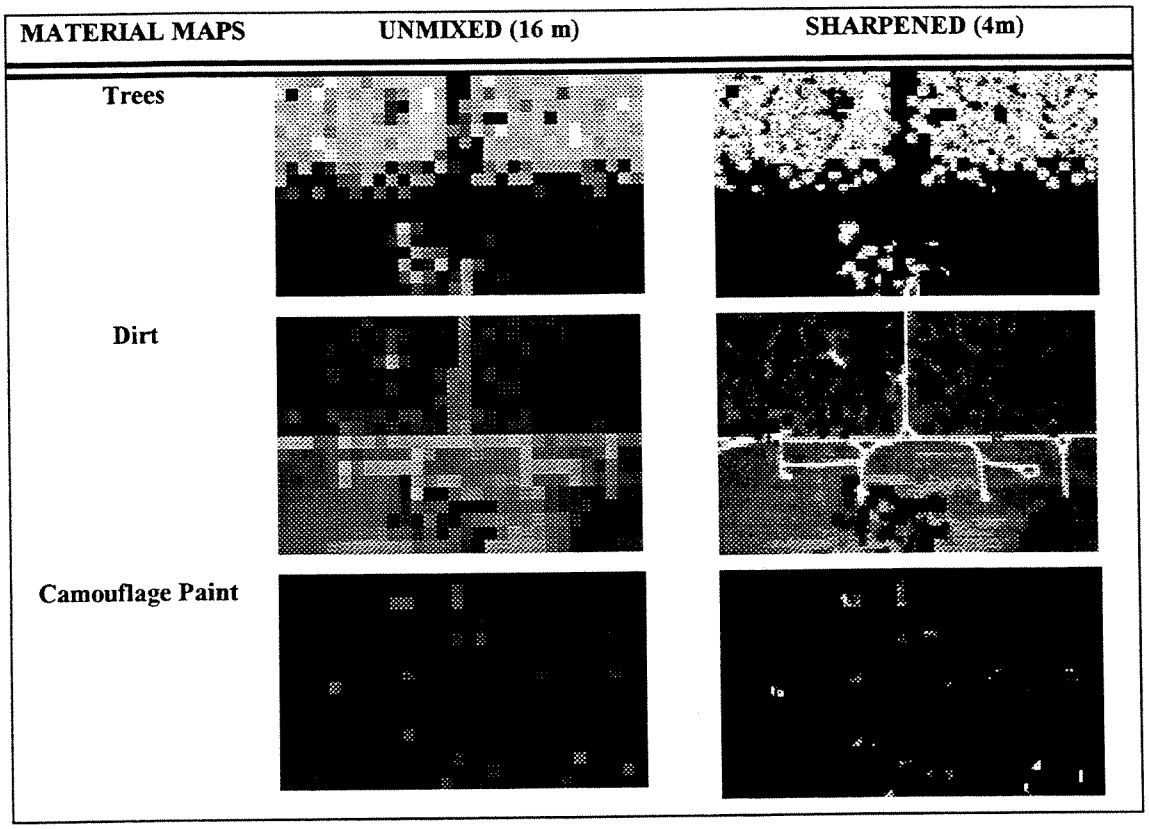

Figure 7: Unmixed (16m) vs. Sharpened (4m) Material Maps

Figure 8 shows the fully constrained case when using multiple sharpening bands. By adding the SWIR information in band 3, the error is reduced by 30%. Table 2 summarizes sharpening performance. It shows the percent improvement relative to the stepwise unmixing algorithm. The results relative to perfect unmixing are shown for comparison.

It is not surprising that sharpening works better on perfectly unmixed maps. Since sharpening uses the endmembers found via unmixing, if unmixing selects the incorrect materials, sharpening will fail. Also, if the low resolution fractions are wrong, and the pixel contains an aggregate mixture, sharpening cannot improve the result. However, perfect unmixing is impossible. These results show sharpening also improves maps created via spectral unmixing. Using more sharpening bands improves the fractions.

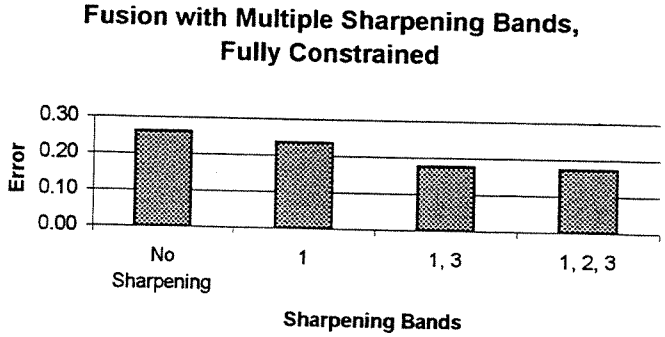

Figure 8: Fusion with Multiple Sharpening Bands

Sharpening vs.	Stepwise Unmixing	Perfect Unmixing
Single Band	10%	22%
Multiple Bands	34%	90%

Table 2: Percent Improvement of Sharpening Over Unmixing

7. SUMMARY AND RECOMMENDATIONS

Spectral mixture analysis techniques are improved by a per pixel unmixing strategy. Stepwise regression provides a mathematical framework with which one adjusts the number of endmembers in the model to reduce overfit. Error checking of the fractions, and entrance and exit criteria, result in an automated algorithm that provides superior material maps. The algorithm was demonstrated for various constraint conditions.

The spectral mixing methodology was extended into the image fusion domain by a process called sharpening. The model for sharpening takes the same form as the spectral mixing model, but in the high resolution case, the problem is an under-determined optimization. The results show sharpening material maps yields more accurate fractions than if the maps are not sharpened. Thus, fusing high resolution data with the spectral mixture maps results in a better description of the underlying scene.

Although these data were compiled for only a single synthetic scene, some tentative conclusions can be drawn. Sharpening reduces error relative to unmixing alone. Since the stepwise unmixing algorithm was shown to be an improvement over the traditional method, this combination of unmixing and sharpening provides an significant advance in the creation of high resolution material maps.

Although the results are promising, some limitations are immediately evident. Since the amount of improvement was shown to be related to the amount of variation in the image, fusion results will be scene dependent. Clearly, the number of mixed pixels in the scene is related to the potential gain due to sharpening. Only pixels that are aggregate mixtures at low resolution, but areal mixtures at high resolution will be aided by sharpening.

Finally, it seems clear that the entire image fusion approach should be influenced by more than the data within the superpixel. One should use a *spatial strategy* that incorporates knowledge of the materials in surrounding pixels into the estimates for the fractions. The materials have a spatial correlation that is not accounted for by current techniques. A spatial strategy could be hierarchical – identifying the pixels with accurate models (small RMS error) first. Then, knowing the materials in those pixels with confidence, one could probabilistically influence the unmixing of neighboring pixels. A spatial strategy could be used to overcome some of the effects of variation.

In conclusion, this image fusion algorithm applies spectral mixing models at both low and high resolutions in order to construct high resolution material maps. The unmixing and sharpening steps are solved as linear least squares problems. The algorithm works at a variety of scales, with multiple sharpening bands, and in the presence of image texture. The results show a quantifiable improvement over traditional unmixing techniques. The algorithm has also been applied to a real M-7 image using image derived endmembers. This yields reasonable results, but requires ground truth for detailed analysis. Future investigation into the strengths and weaknesses of this approach is warranted.

---

[1] Adams, J.B., Smith, M.O., Johnson, P.E., "Spectral Mixture Modeling: A New Analysis of Rock and Soil Types at the Viking Lander 1 Site," *Journal of Geophysical Research,* 91, B8, 8098-8112, 1986.

[2] Smith, M.O., Ustin, S.L., Adams, J.B., Gillespie, A.R., "Vegetation in Deserts: I. A Regional Measure of Abundance from Multispectral Images," *Remote Sensing of Environment,* 31, 1-26, 1990.

[3] Adams, J.B., Smith, M.O., Gillespie, A.R., "Imaging Spectroscopy: Interpretation Based on Spectral Mixture Analysis," *Remote Geochemical Analysis: Elemental and Mineralogical Composition,* Ed. C.M. Pieters and P.A.J. Englert, 145-166, 1993.

[4] Green, R.O., Conel, J.E., Roberts, D.A., "Estimation of Aerosol Optical Depth, Pressure Elevation, Water Vapor and Calculation of Apparent Surface Reflectance from Radiance Measured by the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Using A Radiative Transfer Code," *Proceedings of SPIE: Imaging Spectrometry of the Terrestrial Environment,* 1937, 14-15 April 1993, 2-11, 1993.

[5] Gao, B., Heidebrecht, K.B., Goetz, A,F,H,, "Derivation of Scaled Surface Reflectances from AVIRIS Data," *Remote Sensing of Environment,* 44, 165-178, 1993.

[6] Price, J.C., "How Unique Are Spectral Signatures?," *Remote Sensing of Environment,* 49, 181-186, 1994.

[7] Draper, N.R., Smith, H., *Applied Regression Analysis,* 2nd Ed., 1981. ← QA 278.2.D7 1980 3rd Floor

[8] Miller, A.J., *Subset Selection in Regression,* 1990.

[9] Dobson, A. J., *An Introduction to Generalized Linear Models,* 1990.

[10] Kennedy, Jr., W.J., Gentle, J.E., *Statistical Computing,* 1980.

[11] Price, J.C., "Combining Panchromatic and Multispectral Imagery from Dual Resolution Satellite Instruments," Remote Sensing of Environment, 21, 119-128, 1987.

[12] Munechika, C.K., Warnick, J.S., Salvaggio, C., Schott, J.R., "Resolution Enhancement of Multispectral Image Data to Improve Classification Accuracy," *Photogrammetric Engineering and Remote Sensing*, 59, 1 (January), 67-72, 1993.

[13] Gross, H.N., Schott, J.R., "Application of Spectral Mixing to Image Fusion," *Proceedings of 26th International Symposium on Remote Sensing of Environment,* 25-29 March, Vancouver, B.C., 428-431, 1996.

[14] Gross, H.N., Schott, J.R., "Evaluating an Image Fusion Algorithm with Synthetic Image Generation Tools," *Proceedings of SPIE: Algorithms for Multispectral and Hyperspectral Imagery,* Vol. 2758, 8-12 April, Orlando, 1996.

[15] Lawson, C.L., Hanson, R.J., *Solving Least Squares Problems,* 1974.

[16] Schott, J.R., Raqueno, R., Salvaggio, C.N., "Incorporation of a Time-Dependent Thermodynamic Model and Radiation Propagation Model into Infrared Three-Dimensional Synthetic Image Generation," *Optical Engineering,* 37, 7, 1992.

[17] Research Systems, Inc., *ENVI: The Environment for Visualizing Images,* Version 2.0, July, 1995, Boulder, CO, 1995.

LOOK FOR

FOX, JOHN

"APPLIED REGRESSION ANALYSIS, LINEAR MODELS, & RELATED METHODS"

1997 HA 31.3. F69 1997 4th FLOOR