

Integration of spatial–spectral information for the improved extraction of endmembers

D.M. Rogge^a, B. Rivard^{a,*}, J. Zhang^a, A. Sanchez^a, J. Harris^b, J. Feng^a

^a Earth Observation Systems Laboratory, Department of Earth and Atmospheric Sciences, University of Alberta, Edmonton, Alberta, Canada, T6G 2E3

^b Geological Survey of Canada, Continental Geoscience Division, Ottawa, Ontario, Canada, K1A 0E9

Received 6 November 2006; received in revised form 22 February 2007; accepted 26 February 2007

Abstract

Spectral-based image endmember extraction methods hinge on the ability to discriminate between pixels based on spectral characteristics alone. Endmembers with distinct spectral features (high spectral contrast) are easy to select, whereas those with minimal unique spectral information (low spectral contrast) are more problematic. Spectral contrast, however, is dependent on the endmember assemblage, such that as the assemblage changes so does the “relative” spectral contrast of each endmember to all other endmembers. It is then possible for an endmember to have low spectral contrast with respect to the full image, but have high spectral contrast within a subset of the image. The spatial–spectral endmember extraction tool (SSEE) works by analyzing a scene in parts (subsets), such that we increase the spectral contrast of low contrast endmembers, thus improving the potential for these endmembers to be selected. The SSEE method comprises three main steps: 1) application of singular value decomposition (SVD) to determine a set of basis vectors that describe most of the spectral variance for subsets of the image; 2) projection of the full image data set onto the locally defined basis vectors to determine a set of candidate endmember pixels; and, 3) imposing spatial constraints for averaging spectrally similar endmembers, allowing for separation of endmembers that are spectrally similar, but spatially independent. The SSEE method is applied to two real hyperspectral data sets to demonstrate the effects of imposing spatial constraints on the selection of endmembers. The results show that the SSEE method is an effective approach to extracting image endmembers. Specific improvements include the extraction of physically meaningful, low contrast endmembers that occupy unique image regions.

© 2007 Elsevier Inc. All rights reserved.

Keywords: Spatial–spectral; Image endmember extraction; Hyperspectral remote sensing; Spectral mixture analysis

1. Introduction

Spectral mixing is a problem inherent to remote sensing data and results in few image pixel spectra representing “pure” targets (Settle & Drake, 1993). Linear spectral mixture analysis (SMA) (Adams et al., 1986, 1993) is designed to address the problem of mixed pixels. It assumes that the pixel-to-pixel variability in a scene results from varying proportions of spectral endmembers. The spectrum of a mixed pixel can then be calculated as a linear combination of the endmember spectra weighted by the area coverage of each endmember within the pixel, if the scattering and absorption of electromagnetic

radiation is derived from a single component on the surface (Keshava & Mustard, 2002).

Image endmembers (referred to simply as endmembers hence forth) are pixel spectra that lie at the vertices of the image simplex in n -dimensional space (Fig. 1A). The extraction of endmembers from an image has benefits over the use of spectra measured in the field or laboratory. Library and field spectra are rarely acquired under the same conditions as airborne or satellite data; and they may not adequately represent all important endmembers. On the other hand field and laboratory spectra are usually collected from surfaces one wants to map, and thus, they have direct physical meaning for mapping purposes. Imagery may provide similarly meaningful endmembers that can be considered “pure”, or relatively “pure” spectra, meaning that little or no mixing with other endmembers has occurred within a given pixel.

* Corresponding author. Tel.: +1 780 492 1822; fax: +1 780 492 2030.

E-mail address: benoit.rivard@ualberta.ca (B. Rivard).

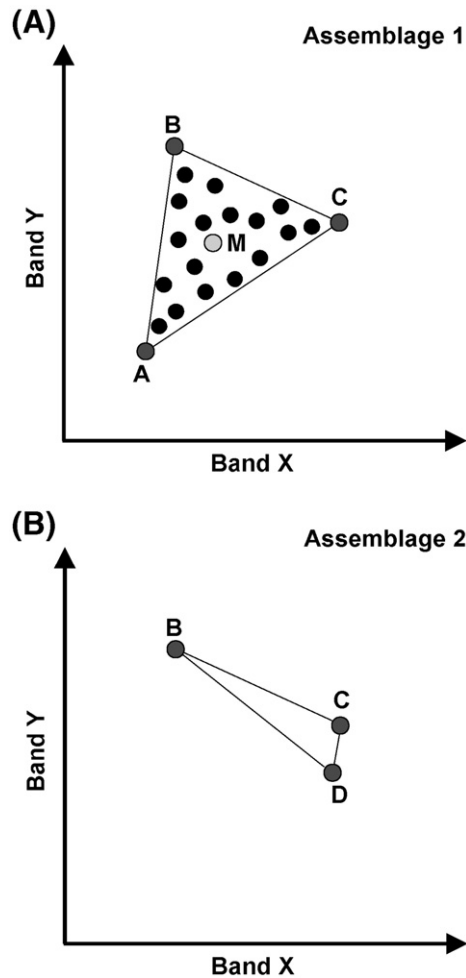


Fig. 1. 2-dimensional scatterplot of endmember assemblage A, B, and C (A); and, B, C, and D (B) located at the vertices of the simplex. In (A) all other pixels (black dots) can be represented as a linear mixture of the 3 endmembers with pixel M an equal mixture of A, B, and C. The relative spectral contrast for C changes for the 2 assemblages, whereby C has an equivalent contrast with A and B in assemblage 1 and a lower contrast in assemblage 2. B has high contrast in assemblage 2, but has equal contrast with A and C in assemblage 1.

To obtain accurate unmixing results the endmembers selected must be representative of surface components that occur in relatively pure form (Adams & Gillespie, 2006). For this reason much literature has focused on the subject of endmember extraction and includes methods such as the pixel purity index (PPI) (Boardman, 1993; Boardman et al., 1995), manual endmember selection tool (MEST) (Bateson & Curtiss, 1996), N-FINDR (Winter, 1999), optical real-time adaptive spectral identification system (ORASIS) (Bowles et al., 1995), the endmember optimization method of Tompkins et al. (1997), convex cone analysis (CCA) (Ifarraguerri & Chang, 1999), iterative error analysis (IEA) (Neville et al., 1999), automated morphological endmember extraction (AMEE) (Plaza et al., 2002), iterated constrained endmembers (ICE) (Berman et al., 2004), and vertex component analysis (VCA) (Nascimento & Dias, 2005). With the exception of AMEE, the above methods select endmembers by discriminating between pixels using their spectral characteristics. This is done independently of neigh-

boring pixels, the spatial distribution of endmembers, and the characteristic spatial mixing relationships between endmembers (e.g. do endmembers mix).

This paper presents a spatial–spectral endmember extraction algorithm (SSEE) that makes use of the spectral and spatial characteristics of image pixels during the search for image endmembers. The spatial characteristics are used to increase the spectral contrast between spectrally similar, but spatially independent endmembers, thus improving the potential of finding these endmembers. We also impose spatial constraints when averaging spectrally similar pixels to preserve similar but distinct endmembers that occupy unique image regions. The output is an image endmember library, where the individual endmembers are defined based on spectral and spatial characteristics. Section 2 gives a brief overview of three relevant endmember extraction methods, two of which were used to generate comparative results. This section is followed by a detailed description of the SSEE algorithm (Section 3) and two demonstrations using airborne hyperspectral data (Section 4).

2. PPI, IEA, and AMEE algorithms

By far the most commonly used endmember extraction tool is PPI, which searches for vertices that define the data volume in n -dimensional space (n =number of bands). Commonly the first step of PPI is to apply a principal component analysis (PCA) or minimum noise fraction (MNF) (Green et al., 1988) to reduce the dimensionality of the data set. MNF is similar to PCA in that involves two cascading PCA transformations, where the first estimates a noise covariance matrix used to decorrelate and rescale the noise in the data. The next is a standard PCA of the noise-reduced data. The assumption here is that the image endmembers lie within the first few principal component axes, whereas the remaining axes are related to noise. However, some image components have weak signals and contribute little energy to the eigenvalues, and thus, determining the cutoff threshold between the eigenvalues caused by signal and noise is problematic (Chang & Du, 2004).

PPI is semi-automated and obtains endmember candidate pixels by projecting the transformed data onto a high number of randomly oriented vectors (k) in n -dimensional space. Those pixels that lie at either end of a given random vector are assigned a “hit”. The total number of hits are tallied for each pixel, for all random vectors. Pixels that receive more hits than a set cutoff threshold (t) are considered candidate endmember pixels, or “pure” pixels. This cutoff threshold is commonly a fixed empirical value (e.g. 2 or 10), or based on statistical parameters, such as the mean hits value (Plaza et al., 2004). The candidate endmember pixels are then loaded into a n -dimension visualization tool, such that the user can visually identify the extreme pixels in the data cloud. This last step requires a significant degree of human intervention from an experienced operator. PPI is particularly sensitive to the input parameters k and t (Chang & Plaza, 2006). Owing to the fact that the vectors are randomly generated, results may not be repeatable. In order to obtain results that are close to repeatable, PPI requires k to be

sufficiently large (e.g. 10^4), such that the number of endmember candidate pixels selected levels off asymptotically as a function of the number of vectors used.

IEA is implemented in the Imaging Spectrometer Data Analysis System (ISDAS) (Staenz et al., 1998), and is based on the residual error image generated when a data set is unmixed using a Weighted Nonnegative Least Squares approach (WNNLS) (Haskell & Hanson, 1981). This method has been used in endmember comparative studies (e.g. Plaza et al., 2004; Winter & Winter, 2000) and was shown to be a robust extraction tool. IEA works by performing a series of recursive constrained unmixing operations on the image, such that the residual error is minimized. The mean spectrum of the scene is used as the starting endmember to initialize the unmixing process. The residual error image is essentially a distance measurement in n -dimensional space between the mean spectrum and each pixel spectrum in the image. Pixels within a pre-determined spectral angle that encompass the largest errors form a new endmember, with the mean spectrum discarded. This process is repeated using the new endmember to find additional endmembers, but unlike the mean spectrum which was discarded, each new endmember is added to the existing endmember set until the number of endmembers specified by the user is reached or until a specified average error tolerance condition is met. The main drawback to IEA is that it is computationally intensive, specifically as the number of endmembers required increases.

AMEE is significantly different from spectral-based methods as it integrates spatial information in order to extract endmembers from an image. AMEE runs on the full data cube with no dimensional reduction. The algorithm begins by searching spatial neighborhoods around each pixel in the image for the most spectrally pure and mostly highly mixed pixel. This task is accomplished using the mathematical morphology operators dilation and erosion respectively. Each spectrally pure pixel is assigned an “eccentricity” value, which is calculated as the spectral angle distance between the most spectrally pure and mostly highly mixed pixel for the given spatial neighborhood. This process is repeated iteratively for larger spatial neighborhoods up to a maximum size that is pre-determined. At each iteration the “eccentricity” values of the selected pixels are

updated. The final endmember set is obtained by applying a threshold to the resulting greyscale “eccentricity” image. There are some limitations to AMEE, particularly a significant increase in processing time as the maximum size of the spatial neighborhood becomes large; and, the algorithm’s ability to select only one pixel per spatial neighborhood (Plaza et al., 2002). However, AMEE has been shown to produce results that are comparable to or better than other endmember extraction methods (Plaza et al., 2002, 2004).

3. Description of the spatial–spectral endmember extraction (SSEE) algorithm

The selection of endmembers becomes more problematic as their spectral contrast approaches the detection limits of the given sensor (e.g. SNR). Improving the spectral contrast between pixels in an image can be accomplished using spectral-based methods such as transforms (e.g. PCA, MNF) (Adams & Gillespie, 2006), derivative analysis (Tsai & Philpot, 1998), and normalization (Clark & Roush, 1984). Masking can also be used to improve spectral contrast by removing image components that dominate the spectral variance of the data resulting in a relative increase in spectral contrast for the remaining image components. However, masking will only be effective in cases where spectral mixing is minimal, which is not commonly the case for natural environments. Masking does illustrate that spectral contrast is variable in an image depending on the spatial neighborhoods, where for each spatial neighborhood the assemblage of endmembers may change.

Fig. 1 provides such an illustration showcasing endmember C as observed in two different spatial neighborhoods, each with a distinct endmember assemblage. In case 1, C is equivalently distinct from both A and B and can be considered to have high spectral contrast relative to A and B. However, in case 2, C is spectrally similar to D, and thus has lower relative spectral contrast compared with case 1. By conducting the image extraction on image subsets we can take advantage of the spatial characteristics of each endmember, which may result in a given endmember having higher spectral contrast in a specific image subset, thus facilitating its extraction. Fig. 2 illustrates a geological example where two lithologic units are spectrally similar, but




Fig. 2. (A) Image region showing three endmembers (i , j , and k), where mixing occurs between i and k , j and k , but not i and j . Spatial groups shown with dotted square in (A). (B) 2-dimensional scatter-plot where endmembers i and j are difficult to discriminate. (C) and (D) show scatter-plots for the two spatial groups allowing for better discrimination of endmembers i and j . Dotted lines in (B), (C), and (D) are the eigenvectors related to the largest eigenvalue for each distribution.




Fig. 3. SSEE Step 1: (A) Image region showing three image components (*i*, *j*, and *k*). (B) Four image subsets. (C) Compiled basis vectors from all subsets shown in (B).

spatially independent. Obtaining endmembers for each lithologic unit can be improved by analysing image subsets.

Averaging pixels that lie near the vertices of a simplex is common practice to generate representative endmember spectra (e.g. in IEA). However, spectral-based endmember extraction methods do not take into account the spatial relationships between the pixels. Thus, spectrally similar pixels that are spatially independent can be averaged together (e.g. Fig. 2 endmember *i* and *j*) to provide a representative endmember. By constraining the averaging process to include only spatially associated pixels it should be possible to reduce spectral contamination of spatially unrelated but spectrally similar endmembers.

The SSEE algorithm described below comprises four steps: 1) application of singular value decomposition (SVD) to determine a set of eigenvectors that describe most of the spectral variance of image subsets; 2) projection of the entire image data onto the eigenvectors to determine a set of candidate endmember pixels; 3) use of spatial constraints to combine and average spectrally similar candidate endmember pixels; and, 4) listing of candidate endmembers in order of spectral similarity.

3.1. SSEE Step 1

Step 1 makes use of SVD, which is very efficient in obtaining a set of eigenvectors that explain most of the spectral variability of a given scene (Healey & Slater, 1999; Thai et al., 1999). SVD, along with PCA and MNF, are projection techniques commonly used in remote sensing. SVD

was implemented in SSEE because it can provide a solution even in the case of highly spectrally correlated data where a “singularity” problem can be encountered, whereas PCA may fail (Jolliffe, 1986; Press et al., 1992). MNF requires a reliable estimate of the noise covariance matrix that can be difficult to obtain for small subsets.

SSEE obtains a set of candidate endmember pixels by applying SVD to subsets of an image (Fig. 3B and C). The SSEE algorithm in its present form makes use of subsets that are equal in size, are square, and do not overlap. Thus, subset size is defined as the number of pixels along its side dimension. The SVD is calculated using the SVDC routine in IDL™ 6.1 (Interactive Data Language), which is based on the routine svdcmp described in Press et al. (1992). For SSEE, eigenvectors that account for 99% of the total spectral variance (s =SVD threshold value) are retained from each subset and compiled into one vector file. These vectors will likely be related to the local high contrast endmembers. For each subset the minimum number of vectors is set to 2, whereas the maximum is defined by the threshold value s (99%). The minimum size of the subset that can be used is defined by the square root of the number of bands in the image data, whereas the maximum size is that of the entire image (assuming the image is also square).

3.2. SSEE Step 2

In step 2, the entire image data (Fig. 4B) is projected onto the compiled vector set with the pixels that lie at either extreme of

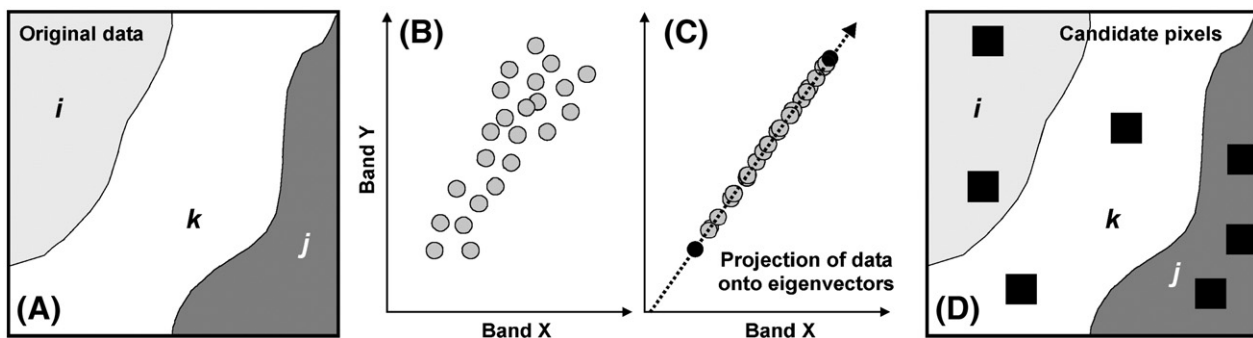


Fig. 4. SSEE Step 2: (A) Image region showing three image components (*i*, *j*, and *k*). (B) Image data in 2-D space. (C) Projection of entire image data onto 1 of the eigenvectors derived from Step 1. (D) Spatial distribution of candidate endmember pixels.

the vectors retained (Fig. 4C). These pixels (Fig. 4D) represent the candidate pixel endmember set, which is used in Step 3.

3.3. SSEE Step 3

Step 3 analyzes the spatial and spectral characteristics of the candidate endmember set to average spectrally similar endmember candidates that are spatially related. Step 3 scans the image with a sliding window of size equal to the subset size used in step 1. Step 3 begins by comparing each candidate endmember pixel to all other pixels within the window centered on the pixel of interest. Pixels within the window that are similar, based on a minimum spectral angle or root mean square error (RMS) (Fig. 5B), become candidate endmember pixels. This process is analogous to retaining more than 1 pixel at either extreme of the eigenvectors. However, using this approach we constrain the number of pixels based on both spectral and spatial similarity. RMS is primarily used for very low reflectance signatures such as water, because the minimum spectral angle is commonly exceeded when candidate water spectra are compared.

Step 3 continues by averaging each of the candidate endmember pixels with all other endmember candidates within the window (Fig. 5C–F). The averaging process is repeated for x number of iterations, with the objective to: 1) reduce the effects of noise; and, 2) find image pixels that are spectrally similar, but spatially related within the window

(Fig. 5). Multiple iterations allow the algorithm to compare and average pixels that are farther than the window size, but are related by other candidate pixels that lie spatially between them (Fig. 5C–F). This iterative approach allows the user to compress endmember clusters reducing their variance (Fig. 5F). The end product is a set of endmembers that are defined both spectrally and spatially, and take into account local spectral variance.

3.4. SSEE Step 4

The endmember set derived from Step 3 is reordered based on spectral angle. The first endmember in the existing library is assigned as the first spectrum in the new reordered list. This spectrum is compared to all other endmembers, with the most similar endmember (lowest spectral angle) assigned the next position in the new reordered list. This process is repeated recursively until all spectra have been ordered. Owing to the iterative spatial averaging process in Step 3 a number of the reordered spectra are duplicates (spectrally similar based on the minimum spectral angle and RMS) (e.g. Fig. 5F). The endmembers retain their image coordinates. The reordered list allows the user to quickly view and group endmember spectra into endmember classes that are spectrally and spatially similar. As one scrolls down the reordered list, spectra will be spectrally similar, but they can be grouped into endmember classes based on their similar spatial coordinates.

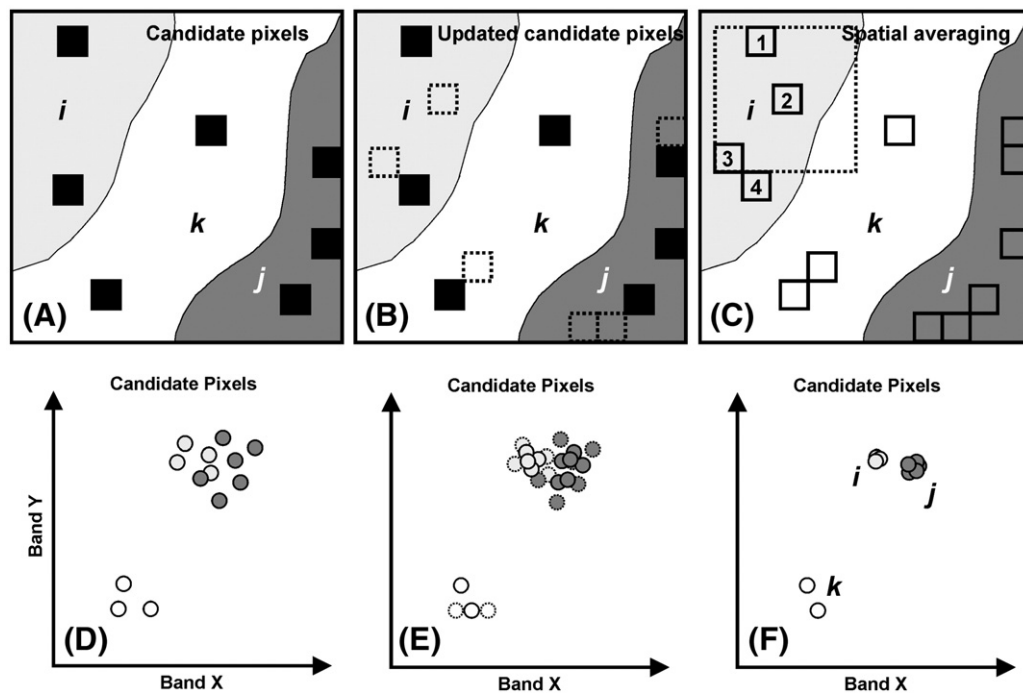


Fig. 5. SSEE Step 3: (A) Candidate endmember pixels showing spatial distribution with respect to units i , j , and k . (B) Updated candidate endmember pixels. (C) Spatial averaging process using a sliding window centered on each updated candidate endmember pixel. (D) Spectral distribution in 2-D space for bands x and y . (E) First iteration of spatial-spectral averaging. Averaged pixels shown as solid lines, with original pixels shown as dashed circles. Pixel number 2 (see C) is averaged with 1 and 3, but not 4, which is averaged with pixel 3. (F) Second iteration of spatial-spectral averaging of pixel 2, which now takes into account the influence of pixel 4 on 3, thus, condensing the spectral cluster further. Continued iterations will compress endmembers i and j into clusters with negligible variance, which now represent two distinct spectral endmembers. Note that for endmember k , only two pixels are averaged as the third pixel is outside the averaging window.

This grouping process could be fully automated. However, we have left this as a manual process to allow a user to input their expert knowledge of a given region.

4. Data sets and evaluation methodology

Two evaluations of the SSEE algorithm were conducted with hyperspectral imagery. The first evaluation is for data from Cuprite Nevada, and is designed to demonstrate the characteristics of SSEE using different spatial subset and averaging window sizes. We examine the link between subset size, eigenvectors retained, and the resulting number of endmembers selected from the image. The second evaluation is for data from Baffin Island, northern Canada. We examine the endmembers related to bedrock geology obtained by SSEE, IEA and PPI in the context of field spectra and also examine unmixing results in the context of the spatial distribution of the endmembers. A comparison with AMEE was not conducted as the algorithm is not readily available.

4.1. Cuprite data and evaluation methodology

The Cuprite imagery was acquired on June 19th, 1997 by the Airborne Visible InfraRed Imaging Spectrometer (AVIRIS (Green et al., 1998)) which has 224 channels covering the 0.37 to 2.51 μm spectral range with an average band Full Width Half Maximum of ~ 10 nm and a Ground Instantaneous Field of View of ~ 20 m. Of the 224 atmospherically corrected channels, we use 167 after removal of channels associated with H_2O and OH absorption features near 1.4 and 1.9 μm . The Cuprite area is arid with limited vegetation cover, and has excellent rock exposure comprising alteration zones characterized by the occurrence of key indicator minerals. For this evaluation we use a 200×200 pixel subset that is centered over the eastern hydrothermal alteration zone (Fig. 6). Within this subset region the following minerals have been validated in the field as occurring in high abundance as documented in Swayze et al. (1992) and Clark et al. (2003): alunite (known variations), chalcedony (OH-quartz), kaolinite (known variations), Namontmorillonite, and buddingtonite. Minerals that occur in lower abundances and as mixtures include: jarosite, hematite, goethite, and smectite/muscovite (Clark et al., 2003).

We evaluate SSEE with subset sizes of 20, 50, 100 and the entire image size of 200. For the entire image size SSEE becomes a spectral-based endmember extraction tool with no spatial constraints. At subset sizes of 20, 50, and 100, s is set to 99%. For the entire image s is also set to 99%, but threshold values of 99.9% and 99.99% are also used. This is done in order to demonstrate that retaining additional eigenvectors, by using a higher s threshold, is not equivalent to obtaining additional eigenvectors from multiple subsets. Spatial averaging is constrained to a window size equal to the subset size used to obtain the eigenvectors. For updating and averaging candidate pixels the spectral angle is set to 1.0° and the RMS is set to 0.001 based on reflectance scaled from 0–1.0. The number of spatial–spectral averaging iterations is set to 5.




Fig. 6. Subset region of Airborne Visible InfraRed Imaging Spectrometer (AVIRIS) hyperspectral data over Cuprite, Nevada (RGB True color exported as greyscale).

4.2. Baffin Island data and evaluation methodology

Airborne hyperspectral data ($\sim 3.5 \times 7$ km; 500×1000 pixels) were acquired with the Probe I sensor, which comprises 128 channels from 0.446–2.543 μm with an average band Full Width Half Maximum of ~ 15 nm and a Ground Instantaneous Field of View of ~ 7 m. The data were converted from digital numbers (DN) to radiance and atmospherically corrected at the Canada Centre for Remote Sensing. Conversion to radiance made use of a vicarious calibration developed by Secker et al. (1991), where calibration coefficients were derived from flat-field targets acquired at the Iqualuit airport (see Fig. 7) concurrently with the overflight. The radiance data was atmospherically corrected using MODTRAN 4 (Berk et al., 1999) and a flat field correction (SMILE, Neville et al., 2003) applied to correct for spectral line curvature. A number of the 128 channels available were not used (874–991, 1082–1171, 1271–1537, 1755–2073, and 2465–2543 nm) for this analysis owing to atmospheric water-absorption and excessive noise. No additional preprocessing (e.g. smoothing filter) was applied to the remaining 86 bands.

This region comprises surfaces with very disparate spectral properties such as snow-ice, water, vegetation, lichen and rock units. Within each class of surface material, the spectral contrast can be relatively low. The majority of geological spectral endmembers are defined by mineral assemblages that comprise the various rock types in the area, and not by the occurrence of a dominating mineral.

Field sampling and collection of spectra took place along traverses oriented perpendicular to the dominant structural and stratigraphic trends in order to obtain a good representation of the various rock units within the study area (Fig. 7). The

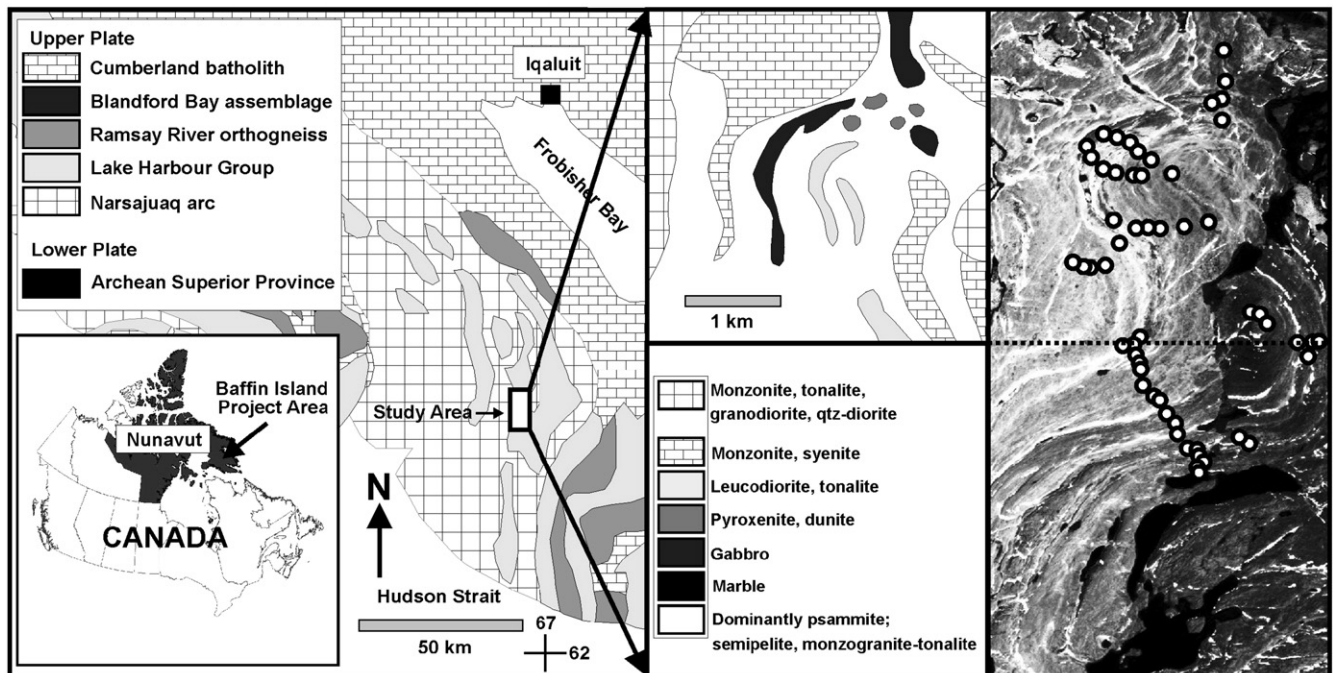


Fig. 7. Regional geology of south-western Baffin Island and enlargement of local geology of the study area (1:100 000) (modified from St-Onge et al., 1999). Hyperspectral data shown at far right (RGB True color exported as greyscale) with field locations shown as white circles. Note only the area above the dotted line was used in this study.

spectra were acquired in early July, 2001, with a portable ASD® field spectrometer that has 2151 bands covering the 0.35–2.50 μm spectral range. A total of 217 spectral measurements were acquired during the daylight hours of ~ 10 am–4 pm (weather dependent) for 56 of 188 sites visited, some of which lie outside, but proximal to the study area, and are representative of the geology shown in Fig. 10. A Spectralon® (Labsphere, Inc.) (99% reflectance) white reference was used for standardization, with each measurement taken at ~ 50 cm above the ground using an 8 degree field of view. Each spectrum saved was an average of 10 measurements. Multiple measurements were taken at each site for fresh, weathered, polished, and partially to fully lichen coated rock surfaces. Rock samples returned to the laboratory were also measured using the portable ASD® field spectrometer with a contact probe attachment, which has its own internal light source.

To evaluate SSEE we compared the endmembers extracted with IEA and PPI to determine whether or not SSEE extracted equivalent endmembers and unique endmembers of significance. For PPI we use only the automated part, because the final steps of PPI require a high degree of human intervention. Endmember candidates determined using different values of k and t are compared with those derived from IEA and SSEE. It is noted here that an automated version of PPI, referred to as Fast Iterative Pixel Purity Index (FIPPI), has recently been presented by Chang and Plaza (2006). However, this method is not yet widely available to the community and was not implemented here.

The next test is a comparison with bedrock spectra acquired in the field and from samples returned to the labo-

ratory. This comparison allows us to test if the unique SSEE endmembers are physically meaningful. Finally, linear unmixing is applied to the image using the endmembers derived from SSEE to determine if the unique endmembers extracted by SSEE show physically meaningful spatial distribution. In this paper we unmix the image using the iterative spectral unmixing analysis (ISMA) approach of Rogge et al. (2006), which is designed to unmix each pixel using an optimal per-pixel endmember set.

For this test the subset size was set to 25. Other parameters for this test include: 1) $s=99\%$; 2) a spatial averaging window size equal to the subset size; 3) the spectral angle is set to 1.0° and the RMS is set to 0.001 based on reflectance scaled from 0–1.0; and, 4) the number of spatial–spectral averaging iterations is set to 10.

For this study the PPI algorithm was written using IDL™, within the ENVI™ environment. Prior to endmember extraction an MNF transform was applied to the data, where MNF bands with an eigenvalue >1 (% of loading) were retained (27 bands of a total of 84). The number of extreme pixels at the ends of the random vectors that are assigned a hit is set to 1. We followed the guidelines of Chang and Plaza (2006) who recommended the use of 10,000 random vectors. For purposes discussed later, PPI was also applied with k equal to the number of vectors generated by SSEE. Cutoff thresholds (t) of 1, 2, and 5 were tested. IEA was implemented using ISDAS, with the number of endmembers extracted from the imagery set to 30 (default value). The maximum number of pixels (m) within an angle (θ) that are subsequently averaged and assigned as an endmember were set to the default parameters ($m=10$ and $\theta=2.5^\circ$).

5. Results

5.1. Cuprite

Table 1 shows that as the subset size decreases the total number of endmembers selected increases; the number of vectors increases; the number of candidate endmember pixels increases; and, the number of updated candidate pixels increases. However, with smaller subsets the processing time decreases significantly, which is primarily controlled by the SVD process. Thus, the largest number of interpretable geological endmembers is found using the smallest subset size, which also requires the minimum processing time. Table 1 also shows the effect of using different values of s for the entire image data (200×200). With a larger s value the number of eigenvectors retained increases, which increases the number of endmembers selected. This, however, is at the expense of more candidate pixels to filter through and greater processing time. The SSEE algorithm was halted using an s value of 99.99% as the number of candidate and updated candidate pixels made the process impractical.

Examination of Table 1 shows that the number of candidate pixels is less than the number of vectors using subsets of 20, 50, and 100. This is because many of the vectors from adjacent subsets are redundant. However, if we use the entire image data and use a higher s value, the number of candidate pixels is




Fig. 8. Endmember spectra derived from SSEE using a subset size of 20×20 and an s value of 0.01. Refer to Table 1 for endmember abbreviations.

much greater than the number of vectors. In this case each vector is orthogonal, and thus, the projected data returns more candidate pixels. When using spatial constraints more vectors are retained, many of which are redundant, but some are important vectors related to signal that improve the potential of

Table 1
Result details for demonstration 1: Cuprite data

	Subset size					
	200×200	200×200	200×200	100×100	50×50	20×20
s (SVD threshold)	0.0001	0.001	0.01	0.01	0.01	0.01
SSEE Step 1						
# of vectors	87	14	4	16	63	356
SSEE Step 2						
# of candidate pixels	144	23	8	13	26	55
SSEE Step 3						
# of updated candidate pixels	23,145	299	82	231	390	732
# of unique candidate pixels ^a	Processing	18	7	11	20	32
SSEE Step 4	Halted					
# of endmembers ^b		8	7	8	10	12
Processing time ^c (s)		856	702	687	403	167
		alu ^d	alu	alu	alu	alu
		cha	cha	cha	cha	cha
		kao	kao	cha	kao	kao
						mon
		bud	bud	bud	bud	bud
					kaoB	kaoB
		kao-hem	kao-hem	kao-hem	kao-hem	kao-hem
					alu-hem	alu-hem
		alu-geo-jar		alu-geo-jar		alu-geo-jar
				mus-hem	mus-hem	mus-hem
		uknA	uknA	uknA	uknA	uknA
		uknB	uknB	uknB	uknB	uknB

^a No. of unique candidate endmember pixels after iterative spatial averaging. Relates only to candidate pixels not updated candidate pixels. Spatial averaging window is=subset size used to compile top eigenvectors, excluding the full image (200×200) which uses no spatial constraints.
^b Derived from the unique candidate endmember pixels.
^c Time for the SSEE algorithm using the IDL code.
^d Mineral abbreviations: alu—alunite; cha—chalcedony; kao—kaolinite; mon—montmorillonite; bud—buddingtonite; kaoB—kaolinite B; hem—hematite; jar—jarosite; geo—goethite; mus—muscovite; uknA—unknown A; uknB—unknown B.

obtaining additional endmembers. Unlike the orthogonal vectors obtained using the entire image many of the vectors retained using spatial constraints will not be orthogonal.

Fig. 8 shows the twelve endmembers derived using a subset size of 20. Ten were labeled as pure minerals or mixtures of minerals based on known mineral absorption features. Two endmembers do not show diagnostic absorption features, but differ in their broad overall shape (continuum). The AVIRIS data set was unmixed using these endmembers as inputs to the

ISMA approach of Rogge et al. (2006). Fig. 9 shows the resulting fractional abundance maps. Each endmember defines distinctive spatial regions that are visually consistent with the maps presented by Clark et al. (2003).

5.2. Baffin Island

The total number of vectors compiled using a subset size of 25 was 2184. From these, 148 endmember candidate pixels




Fig. 9. Fractional abundance maps for the 12 endmembers derived from SSEE using a subset size of 20×20 and an s value of 0.01. Abundance scale linear with black — 0% abundance; white — >75% abundance.

were extracted, of which 7 were noisy and left out (e.g. 141 total no. of pixels in Table 2). Following the spatial averaging procedure 28 of the 148 candidate endmember pixels were exact duplicates. The averaging procedure also results in many of the remaining candidate pixels showing minimal spectral difference, further reducing the number of unique endmembers and allowing the user to quickly group the spectra into endmember classes.

From the reordered library 30 endmembers were determined based on spectral and spatial characteristics (Table 2), including 2 water, 7 snow, 2 vegetation, 2 lichen, 4 lichen-rock mixtures, and 13 rocks (Fig. 10). Fig. 10 also shows the equivalent IEA endmembers. Of the 30 endmembers, IEA obtained representative spectra for 18 (note 2 IEA endmember spectra were noisy and removed). Thus, SSEE extracted 12 unique endmembers, some of which showed only subtle spectral differences, but are considered unique because they are spatially independent (e.g. endmembers 22 and 23; and, 29 and 30).

Using the 865 candidate endmember pixels derived from PPI with $k=10,000$ and $t=1$, 29 of the 30 SSEE endmembers were accounted for (Table 2). However, this number falls to 24 if $t=5$ (266 total candidate endmember pixels). When the equivalent number of random vectors derived by SSEE (2184) are used, 28 of the 30 endmembers were accounted for using $t=1$. For $t=5$ this number falls to 15. It is interesting to note that the majority of unique SSEE endmembers are part of the lichen, lichen-rock and rock endmembers, which have low overall spectral contrast. These results show that SSEE is effective at extracting a more extensive endmember list than either IEA or PPI, especially as t increases. However, to assess whether or not the unique SSEE endmembers are physically realistic we conduct a comparison with field and laboratory spectra collected in the region.

Fig. 11 shows 1 lichen, 2 lichen-rocks, and 11 of the rock endmembers with the best matching field and laboratory spectra. Those in Table 2 not in Fig. 11 were left out because they are not related to bedrock geology (e.g. varnish) or did

Table 2
Endmember classes derived from SSEE and IEA candidate endmember spectra and PPI representative spectra for different k and t

Endmember	Name ^a	IEA ^b	SSEE ^b 25 × 25	PPI $t=1$ ^c $k=10000$	PPI $t=2$ $k=10000$	PPI $t=5$ $k=10000$	PPI $t=1$ $k=2184$	PPI $t=2$ $k=2184$	PPI $t=5$ $k=2184$
Class									
EM 1	Water 1	1	2	*	*	*	*	*	*
EM 2	Water 2		2	*	*	*	*	*	*
EM 3	Snow 1	2	10	*	*	*	*	*	*
EM 4	Snow 2	3	4	*	*	*	*	*	*
EM 5	Snow 3	1	5	*	*	*	*	*	*
EM 6	Snow 4	2	12	*	*	*	*	*	*
EM 7	Snow 5	2	4	*	*	*	*	*	*
EM 8	Snow 6	1	13	*	*	*	*	*	*
EM 9	Snow 7	2	14	*	*	*	*	*	*
EM 10	Vegetation 1	2	11	*	*	*	*	*	*
EM 11	Vegetation 2	3	16	*	*	*	*	*	*
EM 12	Lichen 1	1	6	*	*	*	*	*	*
EM 13	Lichen 2		6	*	*	*	*	*	*
EM 14	Lichen-rock 1		2	*	*	*	*	*	*
EM 15	Lichen-rock 2	2	4	*	*	*	*	*	*
EM 16	Lichen-rock 3	1	2	*	*	*	*	*	*
EM 17	Lichen-rock 4		2	*	*	*	*	*	*
EM 18 (rock 1)	Quartzite		1	*	*	*	*	*	*
EM 19 (rock 2)	Metased(Al–OH)		1	*	*	*	*	*	*
EM 20 (rock 3)	Metased(Fe, Al–OH)		2	*	*	*	*	*	*
EM 21 (rock 4)	Qtz-rich granite		1	*	*	*	*	*	*
EM 22 (rock 5)	Carbonate 1	1	4	*	*	*	*	*	*
EM 23 (rock 6)	Carbonate 2		1	*	*	*	*	*	*
EM 24 (rock 7)	Metased(Fe, Al–OH)	1	3	*	*	*	*	*	*
EM 25 (rock 8)	Metasediment		4	*	*	*	*	*	*
EM 26 (rock 9)	Varnish	1	2	*	*	*	*	*	*
EM 27 (rock 10)	Fe-oxide/hydroxide	1	1						
EM 28 (rock 11)	Vegetation mixture		1	*	*	*	*	*	*
EM 29 (rock 12)	Peridotite 1	1	2	*	*	*	*	*	*
EM 30 (rock 13)	Peridotite 2		3	*	*	*	*	*	*
Total no.									
of classes		18	30	29	29	24	28	22	15
of pixels			141	865	623	266	489	209	88

^a Rock name based on work by Rogge et al. (submitted for publication) which includes a more detailed discussion of the spectral identification, spatial distribution, and economic significance of the rock units found in the entire hyperspectral data shown in Fig. 7.

^b Number represents the number of spectra extracted that can be grouped into the given endmember class.

^c * indicates that at least 1 PPI candidate endmember spectra includes a reasonable representative spectra of the SSEE/IEA derived endmember.

not have good spectral matches (lichen 2, lichen-rocks 1 and 2). The lack of spectral matches for lichen 2, and lichen-rocks 1 and 2, may be attributed to incomplete sampling. The endmembers are represented by a single averaged spectrum derived from those shown in Fig. 10. The matching criteria are based first on spectral angle, which highlights similar

spectra, followed by a visual inspection to determine the best matches. For a good match the focus was on the broad overall shape (continuum) and the location of diagnostic absorption features, rather than total amplitude.

Examination of the spectral matches in Fig. 11 show that lichen endmembers 1, and lichen-rocks 3 and 4; and the rock

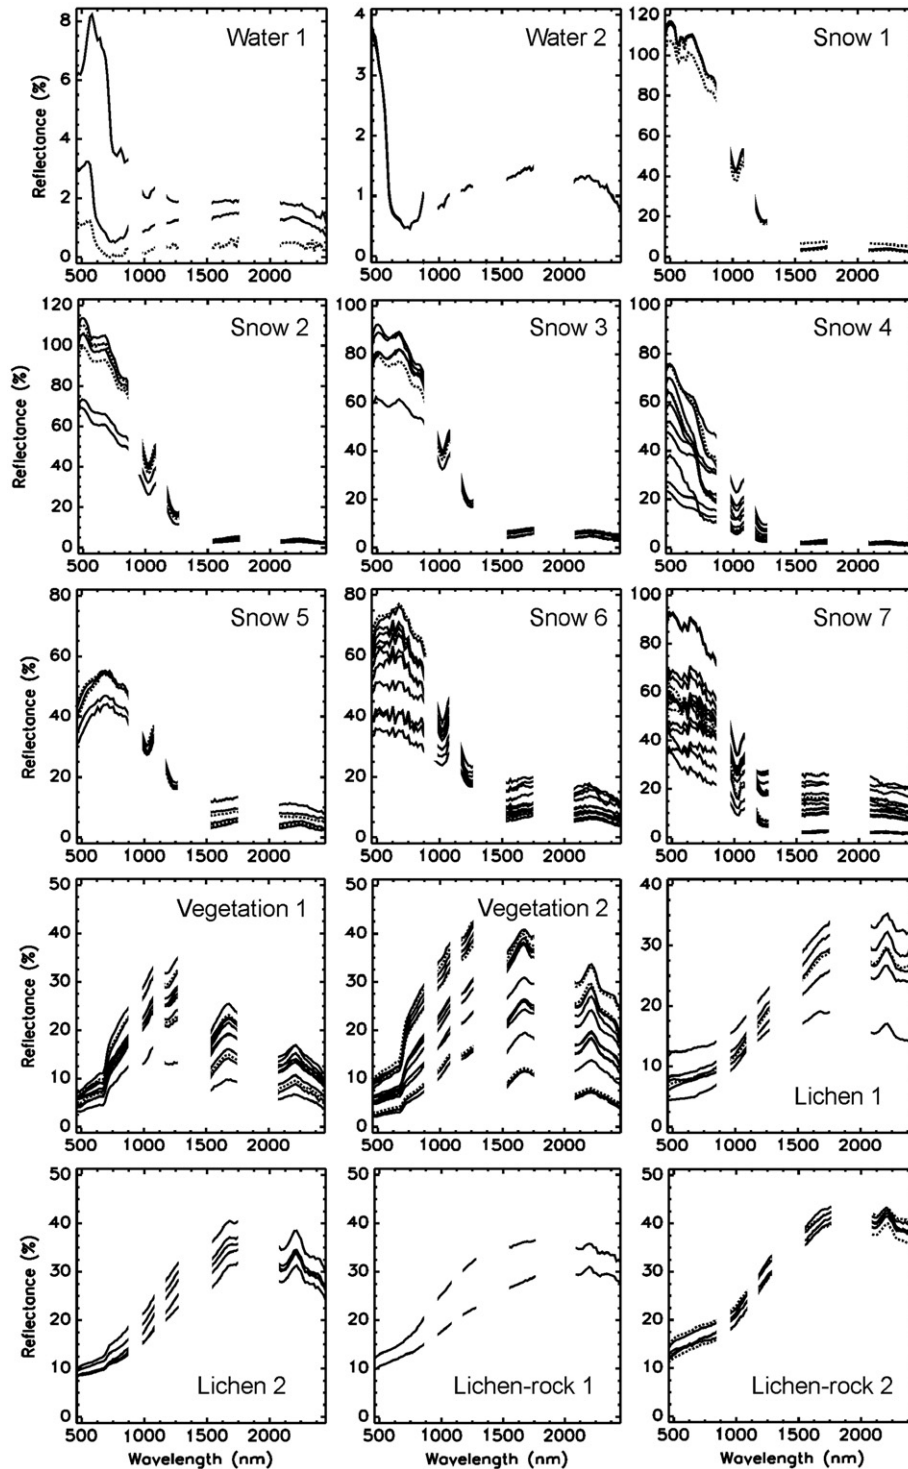


Fig. 10. Thirty endmember classes derived from the SSEE endmember candidate library spectra (solid line) and equivalent IEA endmember (dotted line), when available.

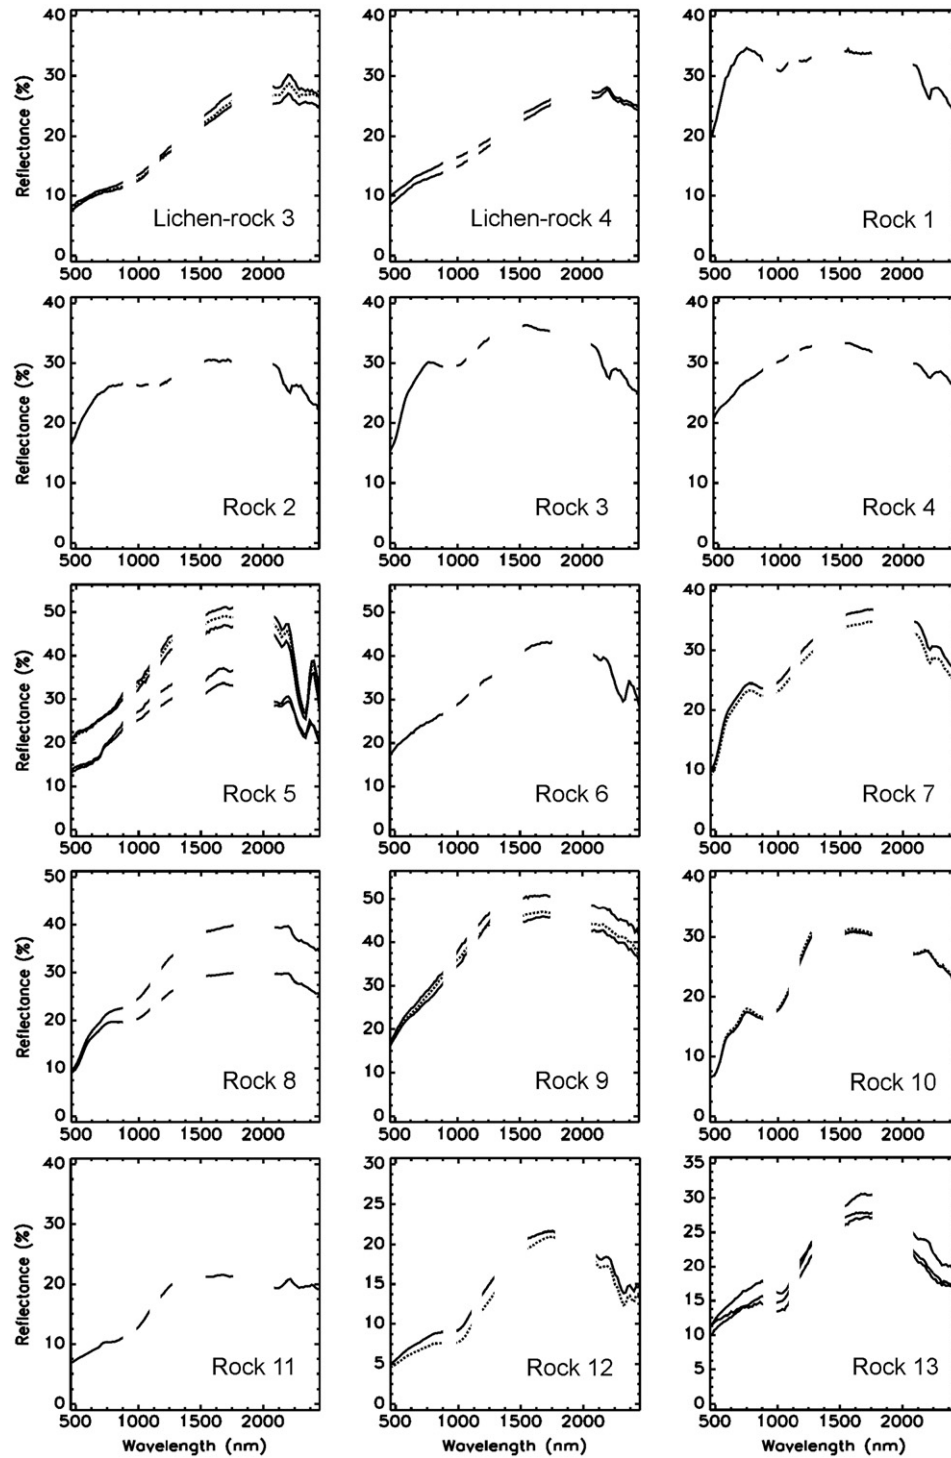


Fig. 10 (continued).

endmembers 1, 2, 3, 4, 5, 7, 8, 10, and 12 show good correspondence with field spectra. Of these, rock endmembers 1, 2, 3, 4, 8, and 13 were not selected by IEA or PPI using $t=5$ and $k=2184$. The best spectral matches for rocks 6 and 13 are the same as for rocks 5 and 12 respectively, which may indicate that these differences are not significant, or that spectral representations of rocks 6 and 13 were missed in the field. For this reason, the unmixing results were checked to

see if these endmembers showed physically meaningful spatial distribution. Fig. 12 shows the fractional abundance maps for rocks 5, 6, 12, and 13; and for lichen-rocks 3 and 4. Analysis of these maps show that rocks 5 and 6 occur locally together, but also in spatially distinctive regions. Where they do occur together locally they are spatially separated (see zoom window in Fig. 12). Rocks 12 and 13 define spatially distinctive regions, although the spatial extent of rock 13 is

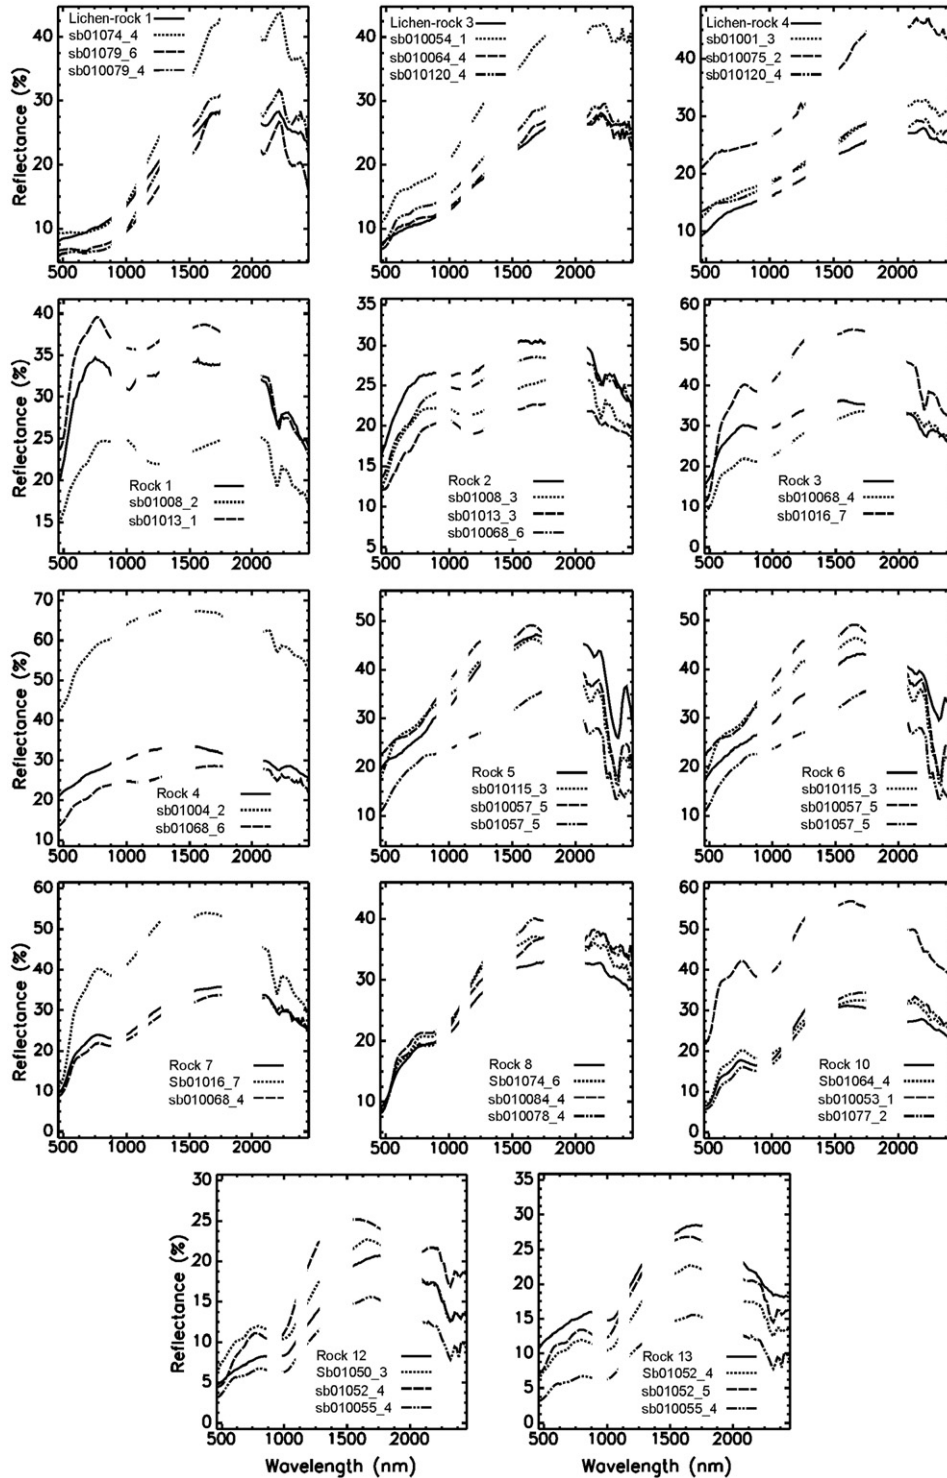


Fig. 11. SSEE rock and lichen-rock endmember classes (solid line) with best matching field spectra.

limited. The fractional abundance maps for lichen-rocks 3 and 4 are also included here as the endmembers are spectrally similar. These two endmembers have some overlap, but for the most part define spatially continuous regions that are distinct. This indicates that although the two endmembers appear to be rock–lichen mixtures, they include enough unique spectral information that they map out unique spatial regions.

6. Discussion

6.1. SSEE parameters, processing speed and performance

For SSEE there are a number of input parameters, which affect both processing speed and performance. The 2 key parameters are subset size and s (SVD threshold value). The actual subset size used will depend on the characteristics of

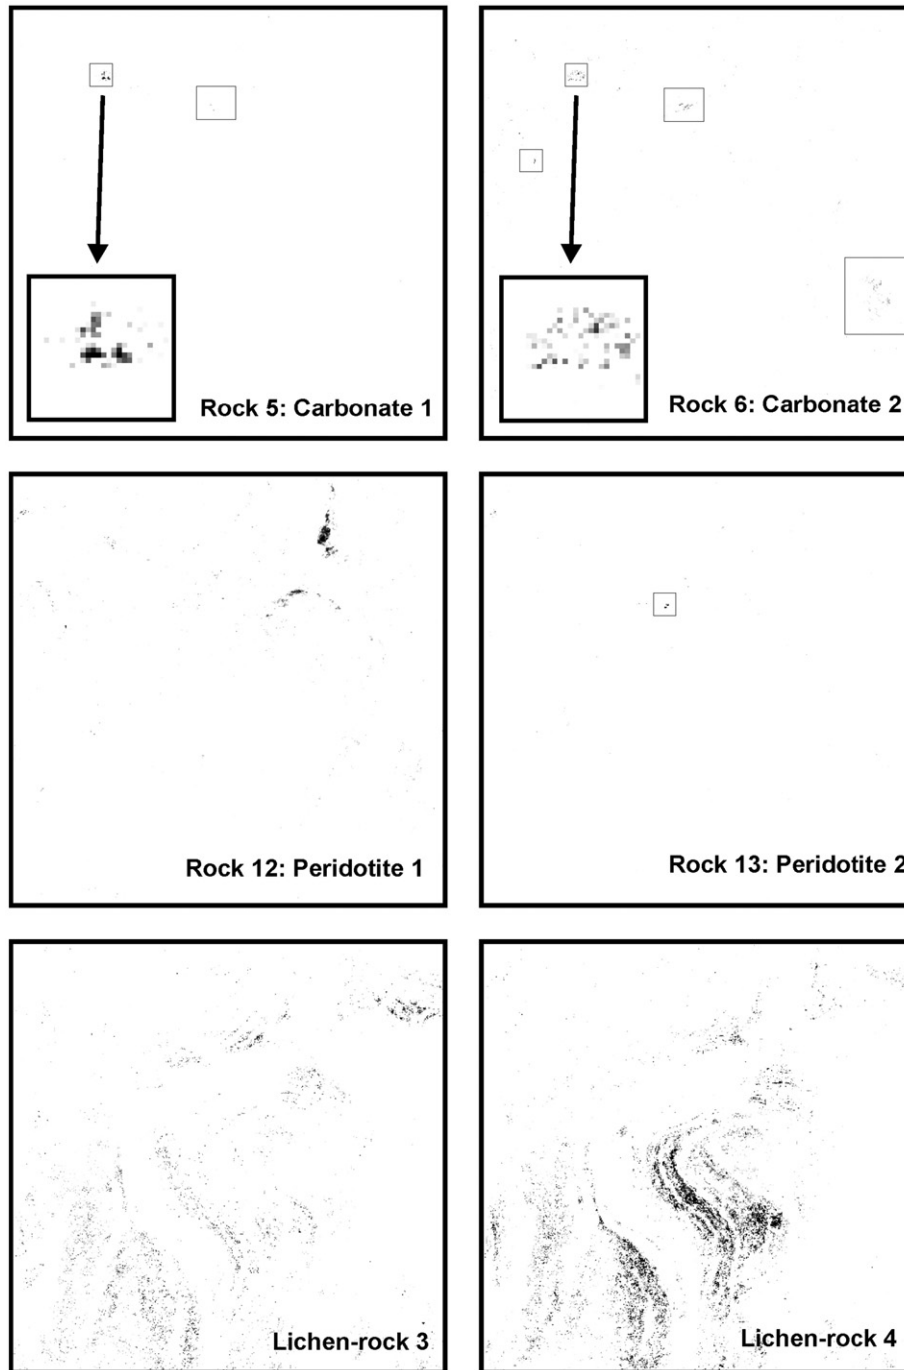


Fig. 12. Fractional abundance maps for rocks 5, 6, 12, and 13; and, lichen-rocks 3 and 4. For rocks 5, 6, and 13 boxes highlight occurrences. Note high abundance is given as black (>75% abundance), whereas low abundance is white (0% abundance). This is necessary to visualize the sparse occurrences of certain rock endmembers. Small boxes are used to highlight areas of sparse spatial coverage. Larger box insert for maps of rocks 5 and 6 show enlargement of localized spatial distribution across one outcrop.

the image, particularly the spectral contrast of the endmembers and their spatial distribution. The degree of homogeneity in the scene is also a factor in determining subset size. Larger subsets can be used in cases where endmembers are distributed as large homogeneous regions. For more complex scenes, smaller subsets are required. It may be possible to apply methods such as semi-variograms to help determine the appropriate subset size. If a scene

contains only high contrast endmembers there is likely minimal benefit to using SSEE. In addition, SSEE will not be useful in cases where low contrast endmembers are always spatially associated, with respect to the subset size. The real benefit of SSEE comes when spectrally similar endmembers are spatially independent. This is evident with rock endmembers 5 and 6, 12 and 13, and lichen-rocks 3 and 4 as shown in Fig. 12.

Choosing an appropriate subset size also effects processing speed. When using smaller subsets the number of vectors retained increases, such that projecting the data onto these vectors becomes the controlling factor with respect to processing speed. Subset sizes larger than 50 significantly reduce processing speed because of the processing time necessary to obtain vectors via SVD. With larger subset sizes we also reduce our ability to obtain vectors that may be related to low contrast endmembers, therefore reducing performance. Thus, the user must balance information gained by reducing the subset size, knowing that many of the additional vectors are redundant. Based on work in this study we found that a subset size of 20–25 pixels was effective. For this analysis subsets did not overlap. However, it may be advantageous to use overlapping subsets, such that each pixel is compared equally to pixels in all directions. Subsets could also be of different sizes and shapes defined by the spatial complexity across the scene (e.g. quadtree decomposition). It may also be useful to preprocess the vectors to remove redundant vectors. However, this must be done with care such that important vectors related to subtle spectral variations are not removed.

The second key parameter that affects performance is the parameter s . Values set higher than 99% resulted in additional vectors, many of which are related to noise and only increased computational time. Values set lower than 99% result in fewer vectors per subset (minimum of 2 in SSEE), which reduce our ability to select low contrast endmembers. Changing s can evidently have significant effects on the results, thus more work is required to determine an s value that works best for all data sets.

The size of the spatial averaging window was set to be equal to the subset size. This was done originally for consistency. However, for larger subsets an equivalent spatial averaging window had a negative impact on the methodology, in that spatially independent endmembers that are spectrally similar may be averaged. In addition, a larger window can increase the number of updated candidate pixels, which can make spatial averaging impractical (see Table 1).

The last two parameters used in SSEE are spectral angle and RMS, which are used to update and average the candidate endmember pixels. For this project we used a spectral angle of 1° and an RMS value equal to 0.001 based on reflectance scaled from 0–1.0. Higher spectral angle and RMS values result in a larger updated candidate endmember list, but also may lead to the loss of the subtle spectral features that define a low contrast endmember. For this reason we chose to keep these two values to a minimum. Note that RMS is primarily used for dark pixels, such as water.

6.2. Comparison with PPI

Of the endmember extraction methods described in Section 2, SSEE has most in common with PPI, even though AMEE also uses spatial information. The similarity between SSEE and PPI relates to the projection of the data onto vectors and retaining those pixels that lie at either end of the vectors. The primary benefit of SSEE compared with PPI is the use of non-

random vectors. First and foremost is the fact that SSEE is repeatable. Secondly, fewer vectors are required. In addition, many of the vectors derived from adjacent subsets are redundant, so the actual number of important vectors is less than the total compiled and results in a smaller number of candidate pixels that the user must work with, compared with PPI. This is evident from the results, where for an equivalent number of random vectors a total of 489 candidate pixels were selected using PPI, as opposed to 148 for SSEE. For PPI, the use of random vectors results in more spectral variability among candidate pixels. For SSEE, vectors are only retained for each subset region if they explain a significant percentage of the spectral variance. This reduces the possibility of retaining vectors related to noise, and in turn, selecting pixels that are noisy.

Because of the high number and spectral variability of the PPI endmembers this requires a great deal of human intervention to derive a final endmember set. For SSEE human intervention is limited to grouping the non-duplicate endmember candidates. This step is simplified by reordering the list based on spectral similarity, and by using the spatial coordinates of each candidate endmember pixel.

Steps 1 and 2 of the SSEE methodology make use of projecting the data volume onto vectors to derive a set of candidate endmember pixels, as used with PPI. However, the key difference is that for SSEE the vectors are not random, but eigenvectors derived from image subsets. In doing so, the SSEE methodology by-passes the difficulty of setting an adequate threshold of eigenvalues encountered when analyzing an entire image by retaining only the top few vectors related to signal for each subset region. This approach makes use of the fact that eigenvectors are dependent on the scene statistics, and are thus, spatially dependent.

6.3. Local vectors versus local candidate pixels

It is possible to use the local vectors to select a set of local candidate pixels for each subset region. However, the key drawback of this approach is the necessity to filter through a much larger number of candidate pixels to determine an endmember set for the full image. These local candidate pixels may also be partial mixtures, which complicate the selection of endmembers. To account for this problem we have chosen instead to use local vectors, rather than local endmembers. Then, in turn, scale up to the full image by projecting the data onto the compiled vector set.

7. Conclusions

The spatial–spectral endmember extraction tool (SSEE) presented in this paper makes primary use of spatial information to: 1) select local eigenvectors that relate to both high and low contrast endmembers within the scene; and, 2) to average only spectrally similar endmembers that are also spatially related. This results in a higher number of candidate endmembers that are defined both spectrally and spatially. The evaluation of SSEE has shown that the method is capable of extracting unique

endmembers with subtle spectral variability that are not selected by other well known spectral-based methods. These unique endmembers were shown to be spectrally significant owing to comparisons with field spectra and through physically realistic spatial distribution.

The two key parameters that affect the processing speed and performance of SSEE are subset size and s (SVD threshold value). For the two evaluations used in this paper a subset size of 20 to 25 pixels squared and an s value of 99% were shown to be effective at selecting both high and low contrast endmembers. The use of local eigenvectors, rather than local endmembers, allows SSEE to retain local information, but also apply that information at the scale of the full image. The use of spatial subsets to select eigenvectors also allows SSEE to bypass the problems associated with selecting a cutoff threshold between eigenvectors caused by signal versus those related to noise. However, the usefulness of SSEE is dependent on the spectral contrast and spatial distribution of the endmembers within the scene. Thus, SSEE is particularly beneficial for extracting spectrally similar endmembers that are also spatially independent. Overall the SSEE method is quick, repeatable, and requires minimal user input.

References

- Adams, J. B., & Gillespie, A. R. (2006). *Remote sensing of landscapes with spectral images: A physical modeling approach* (p. 362). New York: Cambridge University Press.
- Adams, J. B., Smith, M. O., & Gillespie, A. R. (1993). Imaging spectroscopy: Interpretation based on spectral mixture analysis. In C. M. Pieters, & P. A. Englert (Eds.), *Remote geochemical analysis: Elemental and mineralogical composition* (pp. 145–166). Cambridge: Cambridge University Press.
- Adams, J. B., Smith, M. O., & Johnson, P. E. (1986). Spectral mixture modeling: a new analysis of rock and soil types at the Viking Lander 1 site. *Journal of Geophysical Research*, *91*, 8098–8112.
- Bateson, A., & Curtiss, B. (1996). A method for manual endmember selection and spectral unmixing. *Remote Sensing of Environment*, *55*, 229–243.
- Berk, A., Anderson, G. P., Bernstein, L. S., Acharya, P. K., Dothe, H., Matthew, M. W., et al. (1999). MODTRAN4 radiative transfer modeling for atmospheric correction. *SPIE proceedings, optical spectroscopic techniques and instrumentation for atmospheric and space research III* (p. 3756).
- Berman, M., Kiiveri, H., Lagerstrom, R., Ernst, A., Dunne, R., & Huntington, J. F. (2004). ICE: A statistical approach to identifying endmembers in hyperspectral images. *IEEE Transactions on Geoscience and Remote Sensing*, *42*, 2085–2095.
- Boardman, J. W. (1993). Automating spectral unmixing of AVIRIS data using convex geometry concepts. *Summaries of the fourth annual JPL airborne geoscience workshop. JPL Publication 93-26*, (pp. 11–14, Vol. 1).
- Boardman, J. W., Kruse, F. A., & Green, R. O. (1995). Mapping target signatures via partial unmixing of AVIRIS data. *Summaries, fifth JPL airborne earth science workshop. JPL Publication 95-1*, (pp. 23–26, Vol. 1).
- Bowles, J., Palmadesso, P. J., Antoniadou, J. A., Baumbach, M. M., & Rickard, L. J. (1995). Use of filter vectors in hyperspectral data analysis. *Proceedings SPIE infrared spaceborne remote sensing III* (pp. 148–157).
- Chang, C.-I., & Du, Q. (2004). Estimation of number of spectrally distinct signal sources in hyperspectral imagery. *IEEE Transactions on Geoscience and Remote Sensing*, *42*, 608–619.
- Chang, C.-I., & Plaza, A. (2006). A fast iterative algorithm for implementation of Pixel Purity Index. *IEEE Transactions on Geoscience and Remote Sensing Letters*, *3*, 63–67.
- Clark, R. N., & Roush, T. L. (1984). Reflectance spectroscopy: quantitative analysis techniques for remote sensing applications. *Journal of Geophysical Research*, *89*, 6329–6340.
- Clark, R. N., Swayze, G. A., Livo, K. E., Kokaly, R. F., Sutley, S. J., Dalton, J. B., et al. (2003). Imaging spectroscopy: Earth and planetary remote sensing with the USGS Terrestrial and expert systems. *Journal of Geophysical Research*, *108*(E12), 5-1–5-44.
- Green, A. A., Berman, M., Switzer, P., & Craig, M. D. (1988). A transformation for ordering multispectral data in terms of image quality with implications for noise removal. *IEEE Transactions on Geoscience and Remote Sensing*, *26*, 65–74.
- Green, R. O., Eastwood, M. L., Sarture, C. M., Chrien, T. G., Aronsson, M., Chippendale, B. J., et al. (1998). Imaging spectroscopy and the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS). *Remote Sensing of Environment*, *65*, 227–248.
- Haskell, K. H., & Hanson, R. J. (1981). An algorithm for C linear least squares problems with equality and C nonnegativity constraints. *Mathematical Programming, C 21*, 98–118.
- Healey, G., & Slater, D. (1999). Models and methods for automated material identification in hyperspectral imagery acquired under unknown illumination and atmospheric conditions. *IEEE Transactions on Geoscience and Remote Sensing*, *37*, 2706–2717.
- Ifarraguerri, A., & Chang, C.-I. (1999). Multispectral and hyperspectral image analysis with convex cones. *IEEE Transactions on Geoscience and Remote Sensing*, *37*, 756–770.
- Jolliffe, I. T. (1986). *Principle component analysis*. New York: Springer Verlag.
- Keshava, N., & Mustard, J. F. (2002, January). Spectral unmixing. *IEEE Signal Processing Magazine*, 44–57.
- Nascimento, J. M. P., & Dias, J. M. B. (2005). Vertex component analysis: a fast algorithm to unmix hyperspectral data. *IEEE Transactions on Geoscience and Remote Sensing*, *43*, 898–910.
- Neville, R. A., Staenz, K., Szeredi, T., Lefebvre, J., & Hauff, P. (1999, 21–24 June). Automatic endmember extraction from hyperspectral data for mineral exploration. *Fourth international airborne remote sensing conference and exhibition / 21st Canadian symposium on remote sensing*. Ottawa, Ontario, Canada: Natural Resources Canada.
- Neville, R. A., Sun, L., & Staenz, K. (2003). Detection of spectral line curvature in imaging spectrometer data. *SPIE Proceedings: Algorithms and Technologies for Multispectral, Hyperspectral and Ultraspectral Imagery*, (pp. 144–154, Vol. 5093).
- Plaza, A., Martinez, P., Perez, R., & Plaza, J. (2002). Spatial/spectral endmember extraction by multidimensional morphological operations. *IEEE Transactions on Geoscience and Remote Sensing*, *40*, 2025–2041.
- Plaza, A., Martinez, P., Perez, R., & Plaza, J. (2004). A quantitative and comparative analysis of endmember extraction algorithms from hyperspectral data. *IEEE Transactions on Geoscience and Remote Sensing*, *42*, 650–663.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1992). *Numerical recipes in C: The art of scientific computing*, 2nd edition. Cambridge University Press.
- Rogge, D. M., Rivard, B., Zhang, J., & Feng, J. (2006). Iterative spectral unmixing for optimizing per-pixel endmember sets. *IEEE Transactions on Geoscience and Remote Sensing*, *44*, 3725–3736.
- Rogge, D. M., Rivard, B., Harris, J., & Zhang, J. (submitted for publication). Application of hyperspectral data for remote predictive mapping. Baffin Island, Canada. *Economic Geology*.
- Secker, J., Staenz, K., Gauthier, R. P., & Budkewitsch, P. (1991). Vicarious calibration of airborne hyperspectral sensors in operational environments. *Remote Sensing of Environment*, *76*, 81–92.
- Settle, J. J., & Drake, N. A. (1993). Linear mixing and the estimation of ground cover proportions. *International Journal of Remote Sensing*, *14*, 1159–1177.
- Staenz, K., Szeredi, T., & Schwarz, J. (1998). ISDAS — A system for processing/analyzing hyperspectral data; technical note. *Canadian Journal of Remote Sensing*, *24*, 99–113.
- St-Onge, M. R., Wodicka, N., & Lucas, S. B. (1999). Geology of McKeller Bay–Wight Inlet — Frobisher Bay area, southern Baffin Island, Northwest Territories. *Current Research, 1998-C* (pp. 43–53). Geological Survey of Canada.
- Swayze, G. A., Clark, R. N., Sutley, S., & Gallagher, A. (1992). Ground-truthing AVIRIS mineral mapping at Cuprite, Nevada. *Summaries of*

- the third annual JPL airborne geosciences workshop. AVIRIS Workshop, JPL Publication 192-14*, (pp. 47–49, Vol. 1). California: Pasadena.
- Thai, B., Healey, G., & Slater, D. (1999, February). Invariant subpixel material identification in AVIRIS imagery. *Proc. JPL AVIRIS workshop, JPL Publication 99-17*. California: Pasadena.
- Tompkins, S., Mustard, J. F., Pieters, C. M., & Forsyth, D. W. (1997). Optimization of endmembers for spectral mixture analysis. *Remote Sensing of Environment*, 59, 472–489.
- Tsai, F., & Philpot, W. (1998). Derivative analysis of hyperspectral data. *Remote Sensing of Environment*, 66, 41–51.
- Winter, M. E. (1999). Fast autonomous spectral endmember determination in hyperspectral data. *Proceedings of the thirteenth international conference on applied geologic remote sensing* (pp. 337–344). Vancouver, B.C., Canada, II: ERIM International LTD.
- Winter, M. E., & Winter, E. M. (2000). Comparison of approaches for determining end-members in hyperspectral data. *IEEE Aerospace Conference Proceedings*, 3, 305–313.