Journal of Applied Remote Sensing

Impact of JPEG2000 compression on endmember extraction and unmixing of remotely sensed hyperspectral data

Gabriel Martín
Vicente González-Ruiz
Antonio Plaza
Juan P. Ortiz
Inmaculada García

SPIE

Impact of JPEG2000 compression on endmember extraction and unmixing of remotely sensed hyperspectral data

Gabriel Martín,[a] Vicente González-Ruiz,[b] Antonio Plaza,[a] Juan P. Ortiz,[b] and Inmaculada García[b]

[a]Department of Technology of Computers and Communications, University of Extremadura, Avda. de la Universidad s/n, E-10071 Cáceres, Spain
[b]Deptartment of Computer Architecture and Electronics, University of Almería, Ctra. de Sacramento s/n, E-04120 Almería, Spain

Abstract. Lossy hyperspectral image compression has received considerable interest in recent years due to the extremely high dimensionality of the data. However, the impact of lossy compression on spectral unmixing techniques has not been widely studied. These techniques characterize mixed pixels (resulting from insufficient spatial resolution) in terms of a suitable combination of spectrally pure substances (called *endmembers*) weighted by their estimated fractional abundances. This paper focuses on the impact of JPEG2000-based lossy compression of hyperspectral images on the quality of the endmembers extracted by different algorithms. The three considered algorithms are the orthogonal subspace projection (OSP), which uses only spatial information, and the automatic morphological endmember extraction (AMEE) and spatial spectral endmember extraction (SSEE), which integrate both spatial and spectral information in the search for endmembers. The impact of compression on the resulting abundance estimation based on the endmembers derived by different methods is also substantiated. Experimental results are conducted using a hyperspectral data set collected by NASA Jet Propulsion Laboratory over the Cuprite mining district in Nevada. The experimental results are quantitatively analyzed using reference information available from U.S. Geological Survey, resulting in recommendations to specialists interested in applying endmember extraction and unmixing algorithms to compressed hyperspectral data.

Keywords: hyperspectral image compression, JPEG2000, endmember extraction, spectral unmixing, orthogonal subspace projection (OSP), automatic morphological endmember extraction (AMEE), spatial spectral endmember extraction (SSEE).

1 INTRODUCTION

Hyperspectral image compression [1] has received considerable interest in recent years due to the enormous data volumes collected by imaging spectrometers such as NASA's Airborne Visible Infra-Red Imaging Spectrometer (AVIRIS) [2], which covers the wavelength region from 0.4 to 2.5 microns using more than 200 spectral channels, at nominal spectral resolution of 10 nm. Two types of image compression have been applied in hyperspectral imaging applications: lossless and lossy. Since the best compression ratios achieved by lossless methods are in the order of 3:1, lossy decodings are used when higher compression ratios are required [3].

An important issue that has not been widely investigated in the past is the impact of lossy compression on spectral unmixing applications [4], which are the tool of choice in order to deal with the phenomenon of mixed pixels [5], i.e. pixels containing different macroscopically pure spectral substances, as illustrated in Fig. 1. This figure also illustrates graphically the concept of *spectral signature*, i.e. the spectrum associated to each pixel in the hyperspectral scene. The spectral signature is obtained after plotting the reflectance values at different wavelengths for

Received 28 Nov 2009; accepted 19 Apr 2010; published 14 Jul 2010 [CCC: 19313195/2010/$25.00]

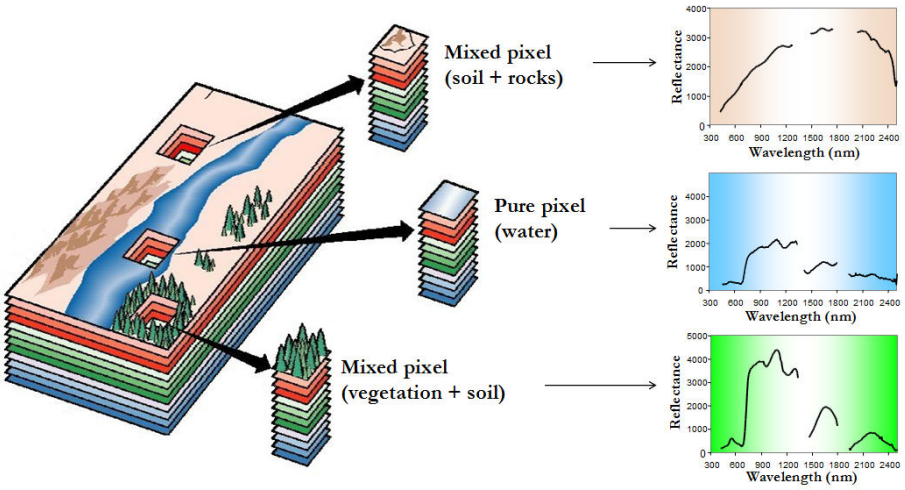

Fig. 1. The mixture problem in hyperspectral data analysis.

the observed pixel. It should be noted that some of the wavelength channels may be affected by atmospheric interferers (such as water absorption) or low signal-to-noise ratio. Hence, in some cases the noisy channels are removed prior to the analysis. This situation is illustrated graphically in the spectral signatures displayed in Fig. 1, which are non-contiguous. In hyperspectral images, mixed spectral signatures may be collected due to several reasons. First, if the spatial resolution of the sensor is not fine enough to separate different pure signature classes at a macroscopic level, these can jointly occupy a single pixel, and the resulting spectral measurement will be a composite of the individual pure spectra, often called *endmembers* in hyperspectral analysis terminology [6]. Second, mixed pixels can also result when distinct materials are combined into a homogeneous or intimate mixture, and this circumstance occurs independently of the spatial resolution of the sensor [4].

Linear spectral unmixing is a standard approach to characterize mixed pixels in hyperspectral data. Let us assume that a remotely sensed hyperspectral scene with n bands is denoted by $\mathbf{I}$, in which the pixel at the discrete spatial coordinates (i, j) of the scene is represented by a vector $\mathbf{X}(i, j) = [x_1(i, j), x_2(i, j), \cdots, x_n(i, j)] \in \Re^n$, where $\Re$ denotes the set of real numbers in which the pixel's spectral response $x_k(i, j)$ at sensor channels $k = 1, \ldots, n$ is included. The linear model assumes that different sub-pixel components reflect incident radiation in linear fashion (see the leftmost part of Fig. 2). Under this assumption, each pixel vector in the original scene can be modeled using the following expression:

$$\mathbf{X}(i, j) = \sum_{z=1}^{p} \Phi_z(i, j) \cdot \mathbf{E}_z + \mathbf{n}(i, j), \tag{1}$$

where $\mathbf{E}_z$ denotes the spectral response of endmember z, $\Phi_z(i, j)$ is a scalar value designating the fractional abundance of the endmember z at the pixel $\mathbf{X}(i, j)$, p is the total number of endmembers, and $\mathbf{n}(i, j)$ is a noise vector. The solution of the linear spectral mixture problem described in Eq. (1) relies on a successful estimation of how many endmembers, p, are present in the input hyperspectral scene $\mathbf{I}$, and also on the correct determination of a set $\{\mathbf{E}_z\}_{z=1}^{p}$ of endmembers and their correspondent abundance fractions $\{\Phi_z(i, j)\}_{z=1}^{p}$ at each pixel $\mathbf{X}(i, j)$. Two physical constrains are generally imposed into the model described in Eq. (1), these are the abundance non-negativity constraint (ANC), i.e., $\Phi_z(i, j) \geq 0$, and the abundance sum-to-one constraint (ASC), i.e., $\sum_{z=1}^{p} \Phi_z(i, j) = 1$ [7]. When the two constraints are imposed into the linear model, it is generally referred to as fully constrained linear spectral unmixing (FCLSU).

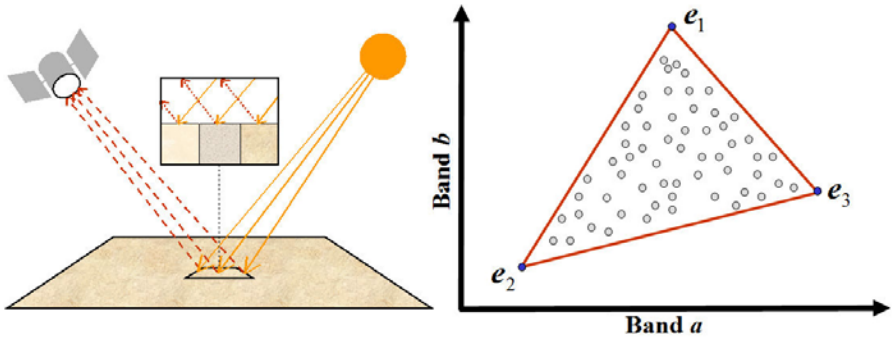

Fig. 2. Linear mixture model (left). Graphical interpretation of the endmember concept (right).

In the last few years, several algorithms have been proposed for automatic extraction of spectral endmembers [6]. Classic techniques have been mostly focused on exploiting the spectral information contained in the data, thus trying to identify the extreme pixels in multidimensional space (see the rightmost part of Fig. 2, which illustrates the location of $p = 3$ endmembers in a 2-dimensional toy space). Examples include the pixel purity index (PPI) [8], N-FINDR [9], iterative error analysis (IEA) [10], optical real-time adaptive spectral identification system (ORASIS) [11], convex cone analysis (CCA) [12], and an orthogonal subspace projection (OSP) technique [13]. Other advanced techniques for endmember extraction have been recently proposed [6], but very few of them consider spatial adjacency.

However, one of the distinguishing properties of hyperspectral data is the multivariate information coupled with a two-dimensional (pictorial) representation amenable to image interpretation. Subsequently, endmember extraction algorithms can also benefit from an integrated framework in which both the spectral information and the spatial arrangement of pixel vectors are taken into account. An example of the situation described above is given in Fig. 3, in which a hyperspectral data cube collected over an urban area (high spatial correlation) is modified by randomly permuting the spatial coordinates (i, j) of the pixel vectors, thus removing the spatial correlation. In both scenes, the application of a spectral-based endmember extraction method would yield the same analysis results while it is clear that a spatial-spectral technique could incorporate the spatial information present in the original scene into the endmember searching process. To the best of our knowledge, only a few attempts exist in the literature aimed at including the spatial information in the process of extracting spectral endmembers. Two of the most representative algorithms are the automatic morphological endmember extraction (AMEE) [14] and the spatial spectral endmember extraction (SSEE) [15].

In this work, the impact of the JPEG2000 [16] compression of an original hyperspectral scene on the quality of the endmembers and fractional abundances derived after using the OSP, AMEE and SSEE algorithms, selected as representative case studies for endmember extraction purposes, has been investigated. In the experiments of this study, the well-known hyperspectral data set collected by AVIRIS over the Cuprite mining district in Nevada, which has a detailed ground-truth information available from U. S. Geological Survey, have been used. The remainder of the paper is structured as follows. Section 2 describes the considered endmember extraction and spectral unmixing algorithms. Section 3 describes the JPEG2000 standard adopted in the data compression experiments. In Section 4 the wavelet filters used are presented. Section 5 describes the AVIRIS hyperspectral image data set used in the experiments and further provides an evaluation on the impact of data compression on the considered algorithms. Finally, Section 6 concludes with some remarks and hints at plausible future research lines.

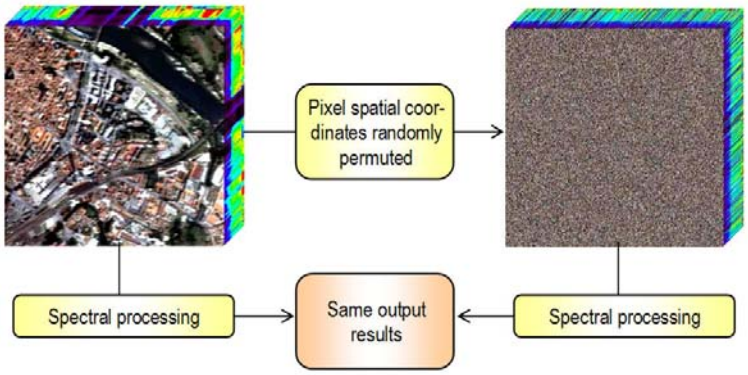

Fig. 3. Example illustrating the importance of spatial information in hyperspectral analysis.

2 ENDMEMBER EXTRACTION AND SPECTRAL UNMIXING

This section describes three endmember extraction techniques: OSP, AMEE, SSEE, along with a FCLSU technique for linear spectral unmixing obeying the ASC and ANC constraints that will be used in this study as representative methods for substantiating the impact of lossy compression via JPEG2000.

2.1 Orthogonal Subspace Projection (OSP)

The orthogonal subspace projection (OSP) algorithm starts by selecting the pixel vector $\mathbf{X}(i,j)$ with maximum length in the scene as the first endmember, as indicated in Fig. 4(A). Then, it looks for the pixel vector with the maximum absolute projection in the space orthogonal to the space linearly spanned by the initial pixel, and labels that pixel as the second endmember, as illustrated in Fig. 4(B). A third endmember is found by applying an orthogonal subspace projector to the original image, where the signature that has the maximum orthogonal projection in the space orthogonal to the space linearly spanned by the first two endmembers is identified as indicated in Fig. 4(C). This procedure is repeated until a final set of spectrally distinct endmembers $\{\mathbf{E}_z\}_{z=1}^{p}$ is found, where p is an input parameter to the OSP algorithm. If the value of p is set properly and the linear mixture model in Eq. (1) holds, the set of selected endmembers should encompass all other pixels in the original hyperspectral scene, meaning that the fractional abundance of any pixel can be calculated as a linear combination of spectral endmembers as indicated in Fig. 4(D).

2.2 Automatic Morphological Endmember Extraction (AMEE)

The automatic morphological endmember extraction (AMEE) [14] algorithm runs on the full data cube with no dimensional reduction, and begins by searching spatial neighborhoods around each pixel vector $\mathbf{X}(i,j)$ in the image for the most spectrally pure and mostly highly mixed pixel. This task is performed by using extended mathematical morphology operators of dilation and erosion [17], which are graphically illustrated on Fig. 5. Here, dilation selects the most spectrally pure pixel in a local neighborhood around each pixel vector $\mathbf{X}(i,j)$, while erosion selects the most highly mixed pixel in the same neighborhood. Each spectrally pure pixel is then assigned an *eccentricity* value, which is calculated as the spectral angle distance between the most spectrally pure and mostly highly mixed pixel for each given spatial neighborhood. This process is repeated iteratively for larger spatial neighborhoods up to a maximum size that is predetermined. 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, which results in a large set of endmember candidates. The final endmembers are extracted after

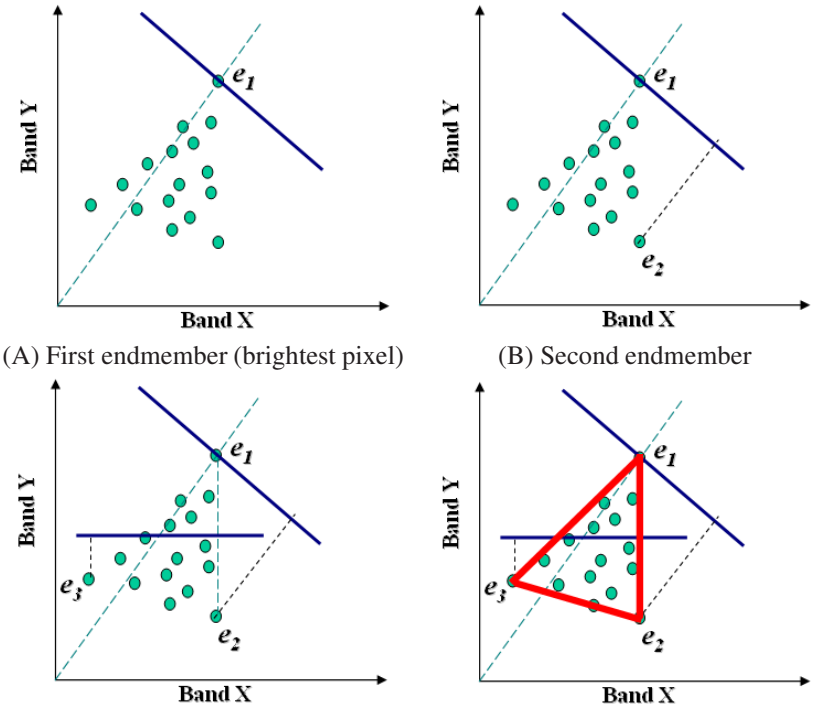

(A) First endmember (brightest pixel) (B) Second endmember

(C) Third endmember (D) Termination rule

Fig. 4. Toy example illustrating the behaviour of OSP algorithm in a 2-dimensional space. (A) Identification of the brightest pixel (first endmember). (B) Selection of the second endmember as the most orthogonal with regards to the first one. (C) Selection of the third endmmeber as the most orthogonal with regards to the first two. (D) Termination rule based on setting the number of endmembers to be found in advance ($p = 3$ in this example).

applying the OSP method to the set of candidates in order to derive a final set of spectrally distinct endmembers $\{\mathbf{E}_z\}_{z=1}^{p}$, where p is an input parameter to the OSP algorithm.

2.3 Spatial Spectral Endmember Extraction (SSEE)

The spatial-spectral endmember extraction tool (SSEE) uses spatial constraints to improve the relative spectral contrast of endmember spectra that have minimal unique spectral information, thus improving the potential for these subtle, yet potentially important, endmembers to be selected. With SSEE, the spatial characteristics of image pixels are used to increase the relative spectral contrast between spectrally similar, but spatially independent endmembers. The SSEE algorithm searches an image with a local search window centered around each pixel vector $\mathbf{X}(i, j)$ and comprises four steps [15]. First, the singular value decomposition (SVD) transform is applied to determine a set of eigenvectors that describe most of the spectral variance in the window or partition (see Fig. 6). Second, the entire image data are projected onto the previously extracted eigenvectors to determine a set of candidate endmember pixels (see Fig. 7). Then, spatial constraints are used to combine and average spectrally similar candidate endmember pixels by testing, for each candidate pixel vector, which other pixel vectors are sufficiently similar in spectral sense (see Fig. 8). Instead of using a manual procedure as recommended by the authors in [15], we have used the OSP technique in order to derive a final set of spectrally distinct endmembers $\{\mathbf{E}_z\}_{z=1}^{p}$, where p is an input parameter to the OSP algorithm.

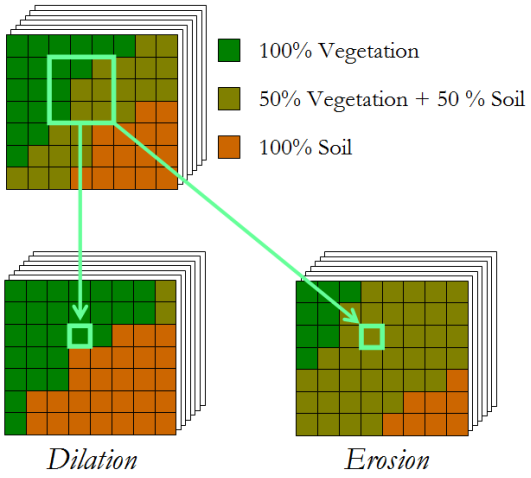

Fig. 5. Toy example illustrating extended morphological operations of dilation and erosion.

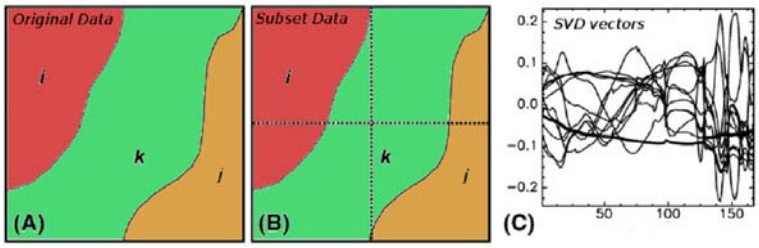

Fig. 6. First step of the SSEE algorithm. (A) Original data. (B) Subset data after spatial partitioning. (C) Set of representative SVD vectors used to describe spectral variance.

2.4 Fully Constrained Linear Spectral Unmixing (FCLSU)

Once a set of endmembers $\mathbf{E} = \{\mathbf{E}_z\}_{z=1}^{p}$ have been extracted, their correspondent abundance fractions $\Phi(i,j) = \{\Phi_z(i,j)\}_{z=1}^{p}$ in a specific pixel vector $\mathbf{X}(i,j)$ of the scene can be estimated (in least squares sense) by the following unconstrained expression [18]:

$$\hat{\Phi}_{\mathrm{UC}}(i,j) = (\mathbf{E}^T\mathbf{E})^{-1}\mathbf{E}^T\mathbf{X}(i,j). \tag{2}$$

However, it should be noted that the fractional abundance estimations obtained by means of Eq. (2) do not satisfy the ASC and ANC constraints. Imposing the ASC constraint results in the following optimization problem:

$$\min_{\Phi(i,j)\in\Delta}\left\{(\mathbf{X}(i,j)-\Phi(i,j)\cdot\mathbf{E})^T(\mathbf{X}(i,j)-\Phi(i,j)\cdot\mathbf{E})\right\},$$
$$\text{subject to: } \Delta = \left\{\Phi(i,j)\left|\sum_{z=1}^{p}\Phi_z(i,j)=1\right.\right\}. \tag{3}$$

Similarly, imposing the ANC constraint results in the following optimization problem:

$$\min_{\Phi(i,j)\in\Delta}\left\{(\mathbf{X}(i,j)-\Phi(i,j)\cdot\mathbf{E})^T(\mathbf{X}(i,j)-\Phi(i,j)\cdot\mathbf{E})\right\},$$
$$\text{subject to: } \Delta = \{\Phi(i,j)|\Phi_z(i,j)\geq 0 \text{ for all } 1\leq z\leq p\}. \tag{4}$$

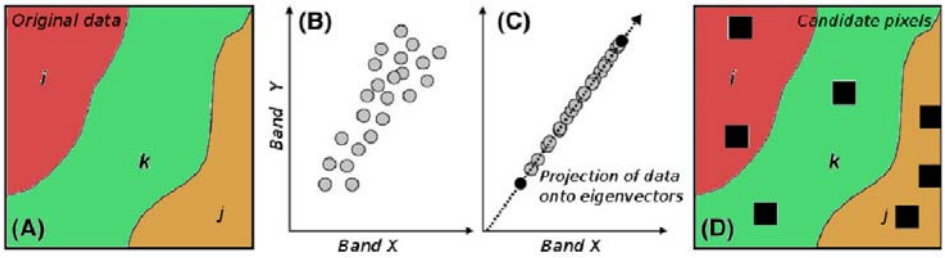

Fig. 7. Second step of the SSEE algorithm. (A) Original data. (B) Spectral distribution in 2-dimensional space. (C) Projection of data onto eigenvectors. (D) Set of candidate pixels.

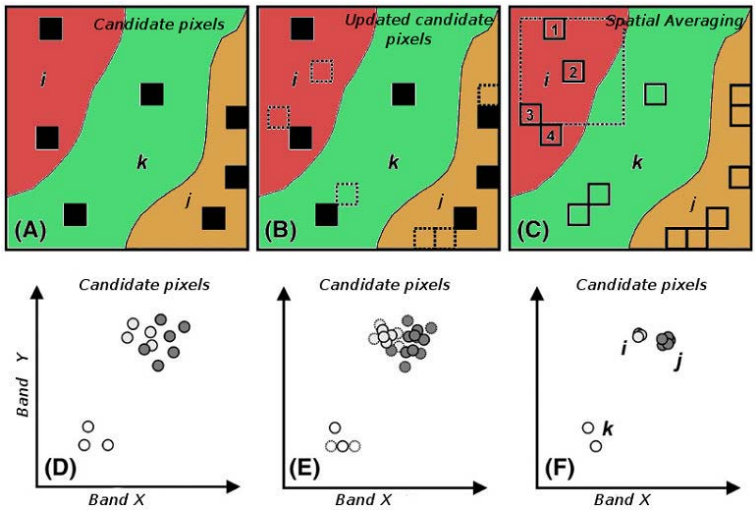

Fig. 8. Third step of the SSEE algorithm. (A) Set of candidate pixels. (B) Updated candidate pixels after including pixels which are spectrally similar to those in the original set. (C) Spatial averaging process of candidate endmember pixels using a sliding window centered on each candidate. (D) First iteration of spatial-spectral averaging. Averaged pixels shown as thick lines, with original pixels shown as thinner lines. (E) Second iteration of spatial-spectral averaging. (F) Continued iterations compress endmembers into clusters with negligible variance.

As indicated in [19], a non-negative constrained least squares (NCLS) algorithm can be used to obtain a solution to the ANC-constrained problem described in Eq. (4) in iterative fashion [20]. In order to take care of the ASC constraint, a new endmember signature matrix, denoted by $\mathbf{E}'$, and a modified version of the pixel vector $\mathbf{X}(i,j)$, denoted by $\mathbf{X}'(i,j)$, are introduced as follows:

$$\mathbf{E}' = \begin{bmatrix} \delta\mathbf{M} \\ \mathbf{1}^T \end{bmatrix}, \Phi'(i,j) = \begin{bmatrix} \delta\Phi(i,j) \\ 1 \end{bmatrix}, \tag{5}$$

where $\mathbf{1} = \underbrace{(1,1,\cdots,1)}_{p}^T$ and δ controls the impact of the ASC constraint. Using the two expressions in Eq. (5), a fully constrained estimate can be directly obtained from the NCLS algorithm by replacing $\mathbf{E}$ and $\Phi(i,j)$ used in the NCLS algorithm with $\mathbf{E}'$ and $\Phi'(i,j)$.

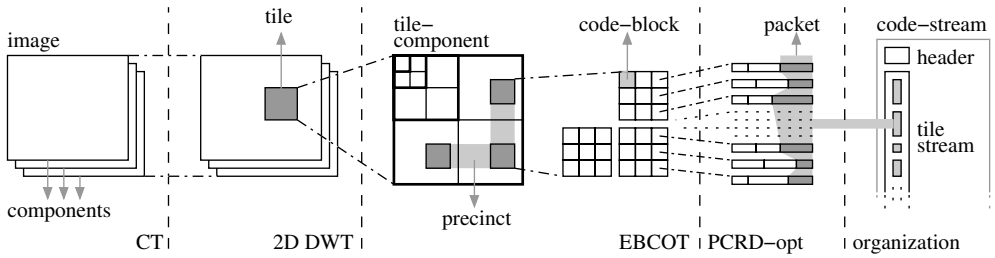

Fig. 9. JPEG2000 data partitioning, coding and code-stream organization.

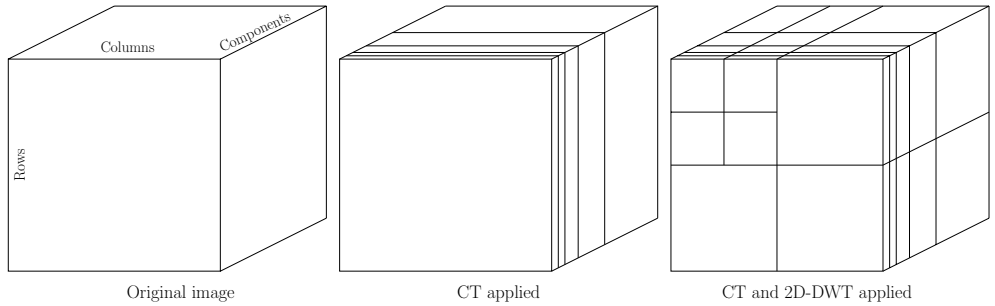

Fig. 10. The hybrid scheme for 3D decorrelation. 4 levels for the CT and 2 levels for the ST.

3 OVERVIEW OF JPEG2000

The JPEG2000 standard is divided into several and incremental parts. Part 1 [21] defines the core coding system and two basic image file formats. Part 2 [22] specifies a set of extensions to the core coding system, such as spectral (inter-component) decorrelation and the use of different wavelet kernels, as well as a more extensible file format. These two first parts are the ones that are going to focus on along this paper. The rest of parts introduce some extensions for different applications. For example, Part 10 (named also JP3D) [23] is concerned with the coding of three-dimensional (3D) data. However, this paper is not tackling this part since it is specifically oriented to volumetric (not only multicomponent) images.

Figure 9 shows an overview of the data partitioning made by the core coding system of JPEG2000, and how the elements of a three-component image (such as a color RGB image) are encoded and distributed. The first step in the JPEG200 algorithm, not shown in the figure, is a level offset to guarantee that all the samples are signed. This is a requirement of the transform machinery that is going to be applied. After that, a CT (Component Transform, for example, a RGB to a YUV color transform) removes the inter-component redundancy that could be found in the image. The result of this stage is a new image in other domain, with the same number of components and samples per component. Next, as can be observed in the figure, each component of the image is divided into rectangular areas called tiles. Tiling is useful for compound images because the encoding parameters of each tile can be selected taking into account its characteristics. However, the tiling is rarely used in natural images due to the artifacts produced around the edges of the tiles, having commonly only one tile per component.

JPEG2000 allows the use of different decomposition patterns in the component domain, although the default one is the hybrid scheme (see Fig. 10). In the hybrid decomposition, a dyadic 1D-DWT (Discrete Wavelet Transform) is first applied to the component domain, and then a dyadic 2D-DWT, denoted by ST (Spatial Transform) in the rest of this paper, is applied to each tile-component. Currently, the dyadic DWT is widely used in the processing of scalable image contents because it facilitates the resolution scalability and improves the encoding effi-

ciency, removing the intra-component (spatial) redundancy. In the example of Fig. 9 we can see how four different resolution levels are generated (remarked in blue) when a ST of three iterations is applied to a tile. These resolution levels are commonly referred by positive integer numbers starting from 0 for the highest one, the original image size.

JPEG2000 provides two working modes: lossy and lossless. The first one offers a better encoding efficiency at low bit-rates. When no loss of information is allowed, the lossless mode can be selected. A fixed point (integer) version of the DWT is used in this case. The resultant wavelet coefficients are grouped into rectangular areas (e.g. 64×64 coefficients) called code-blocks, that are encoded independently by the EBCOT (Embedded Block Coding with Optimal Truncation) algorithm [24]. In order to manage the image information more easily, the code-blocks related to the same rectangular location, within the same resolution level, are grouped into precincts. ROI (Region Of Interest) and the spatial scalabilities are achieved in JPEG2000 by means of the precincts. The compressed bit-streams of the code-blocks can be divided into a specific number of contiguous segments, or quality layers, by the PCRD-opt (Post-Compression Rate-Distortion Optimization) rate-allocation algorithm. The segments of all the code-blocks of a precinct associated to the same quality layer are stored in a packet. The packet is the storing unit in JPEG2000 and it is associated to a quality layer (L), a precinct (P), a resolution level (R), and a tile-component (C). The word formed by this four letters specifies the progression order used to store the image packets, existing five different possibilities: LRCP, RLCP, RPCL, PCRL and CPRL.

The distortion of the decoded image decreases as the amount of decoded data (packets) increases. The LRCP progression provides a fine-grain quality scalability mode. A sequentially decoded RLCP or RPCL image will produce a reconstruction of incremental resolution. The PCRL progression is useful in scan-based systems, like printers. Finally, an CPRL compressed image will be restored, component by component.

The most basic file format defined in the standard, in Part 1, contains only the code-stream of an image (see Fig. 9). This is composed by all the packets of the image and a set of markers with additional information. Markers can be located at any place in the code-stream, however the most important are included in the header. The image files with this format usually have the extension J2C or J2K.

Part 1 also defines a more complex file format based on "boxes". This format allows the coder to include additional information such as color palettes or meta-data. All the information is organized in boxes, contiguous segments of data, whose content is identified by a four-bytes code located at its header. It is possible to define a complex hierarchical structure since a box can contain many other boxes. The extension used to identify the image files with this format is JP2.

The box-based structure of the JP2 format is extensible. Just defining new four-bytes identifiers would allow to include new kind of boxes within an image file, maintaining the backward compatibility (an image viewer that does not understand certain box codes it just ignores them). Part 2 defines a new set of boxes with additional and powerful functionalities. For instance, multiple code-streams can be included within a file, as well as a complex composition scheme (animations, transparency masks, geometric transformations, user definable wavelet kernels, multi-component processing (CT), etc.), which will determine how the image decoding and displaying must be performed. The extension JPX is used to identify those files that contain boxes of Part 2. This is the file format used in our experiments.

4 THE USED WAVELET FILTERS

Any DWT can be described by means of its analysis filters, and any inverse DWT by means of its synthesis filters. Wavelet filters (or wavelet kernels) significantly influence the behavior and the efficiency of the JPEG2000 encoding system. In our work we have investigated the $2/2$ (or Haar) kernel [25], the spline $5/3$ (or lineal) kernel [26], and the spline $13/7$ (or cubic)

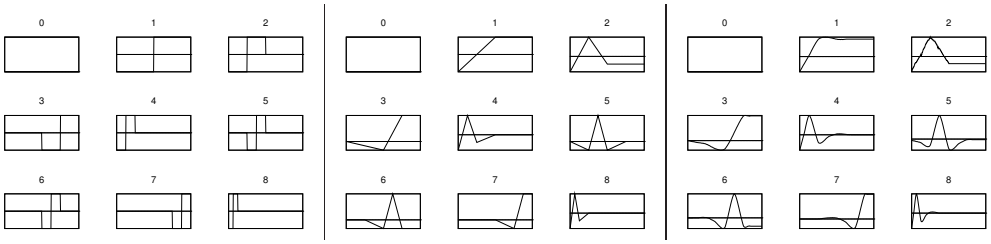

Fig. 11. First nine basis functions of the $2/2$ (left), $5/3$ (center) and $13/7$ (right) wavelet transforms.

kernel [27]. All of these transforms are based on two-channel filter banks. The notation x/y indicates that the underlying filter bank has lowpass and highpass analysis filters of lengths x and y, respectively. The analysis filters of the $2/2$ transform correspond to

$$H_i = S_{2i+1} - S_{2i} \tag{6}$$

$$L_i = S_{2i} + \frac{H_i}{2}, \tag{7}$$

the analysis filters of the $5/3$ transform are

$$H_i = S_{2i+1} - \frac{S_{2i} + S_{2i+2}}{2} \tag{8}$$

$$L_i = S_{2i} + \frac{H_{i-1} + H_i}{4}, \tag{9}$$

and finally, the $13/7$ transform have the filters

$$H_i = S_{2i+1} - \Big(\frac{9}{16}\big(S_{2i} + S_{2i+2}\big) - \frac{1}{16}\big(S_{2i-2} + S_{2i+4}\big)\Big) \tag{10}$$

$$L_i = S_{2i} + \frac{9}{32}\big(H_{i-1} + H_i\big) - \frac{1}{32}\big(H_{i-2} + H_{i+1}\big), \tag{11}$$

where S_i denotes the i-th sample of the one-dimensional signal S, L_i the i-th sample of the lowpass subband and H_i the i-th sample of the highpass subband. Such transforms are linear and have linear phase. As reported in the literature, the filters are expressed using Lifting [28] and are invertible in finite-precision arithmetic (i.e., reversible). For the 2D case, the 2D transform is separable. Therefore, images are handled by transforming all the rows and next, all the columns, or viceversa. This step produces four 2D subbands named LL, HL, LH y HH. The recursive application of this step to the LL subband generates the ST (dyadic 2D-DWT). Symmetric extension is used to process the image boundaries [29].

Notice that the basis functions of the DWT has a strong impact on the transform effectiveness and the nature of the artifacts obtained in lossy reconstructions. As can be seen in Fig. 11, the $2/2$ transform will generate piecewise reconstructions, the $5/3$ transform will produce linear approximations of the original signal and the $13/7$ transform smooth reconstructions. For natural images that are composed of a collection of continuous signals, the $13/7$ should produce the best results, also from a numerical as well as a visual point of view.

5 EXPERIMENTAL RESULTS

This section provides a quantitative and comparative assessment of the impact of JPEG2000-based compression on endmember extraction and spectral unmixing from hyperspectral data. The section is structured as follows. In Subsection 5.1 we first describe the hyperspectral scene

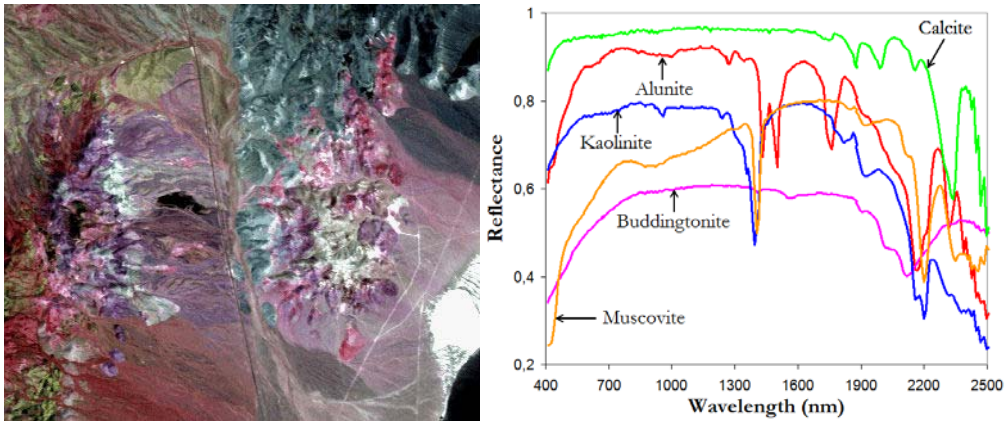

Fig. 12. (a) Hyperspectral image over the AVIRIS Cuprite mining district in Nevada. (b) USGS reference signatures used to assess endmember signature purity.

used in the experiments, along with some reference spectral signatures available for the scene that will be also used. Subsection 5.2 describes, via a simple experiment, the compression ratios and parameter ranges and values that have been empirically selected for a quantitative and comparative assessment. Subsection 5.3 describes the metrics that have been adopted to measure the impact of compression on both endmember extraction and spectral unmixing tasks. Finally, Subsection 5.4 evaluates the impact of the compression parameters adopted in 5.2 on the results reported for the hyperspectral scene described in 5.1, using the evaluation metrics reported in 5.3.

5.1 Hyperspectral data

Before describing our numerical results, we first briefly describe the hyperspectral scene used in the experiments. The reference data set in this study is the well-known AVIRIS Cuprite data set, available online from `http://aviris.jpl.nasa.gov/html/aviris.freedata.html` in reflectance units after atmospheric correction. This scene has been widely used to validate the performance of endmember extraction algorithms. The portion used in experiments corresponds to a 350×350-pixel subset of the sector `f970619t01p02_r02_sc03.a.rfl` in the online data. The scene comprises 224 spectral bands between 0.4 and 2.5 μm, with full width at half maximum of 10 nm. Each reflectance value in the scene is represented using 16 bits, for a total image size of approximately 53 MB. Prior to the analysis, several bands (specifically, bands 1–2, 105–115 and 150–170) were removed due to water absorption and low SNR in those bands, leaving a total of 188 reflectance channels to be used in the experiments. The Cuprite site is well understood mineralogically [30, 31], and has several exposed minerals of interest including those used in the USGS library considered for the generation of simulated data sets. Five of these laboratory spectra (*alunite*, *buddingtonite*, *calcite*, *kaolinite* and *muscovite*) convolved in accordance with AVIRIS wavelength specifications, will be used to assess endmember signature purity in this work. For illustrative purposes, Fig. 12 shows the image data set considered in experiments and the USGS library signatures.

5.2 Compression parameters

In our experiments, the full scene has been encoded progressively using the reversible path (lossy-to-lossless) and 20 quality layers where each quality layer occupies 0.1 bpppb (bit per pixel per band). Therefore, experiments will deal with 20 reconstructed scenes with 20 different values of the quality layers in the range 0.1-2.0 bpppb. The progression order selected for the

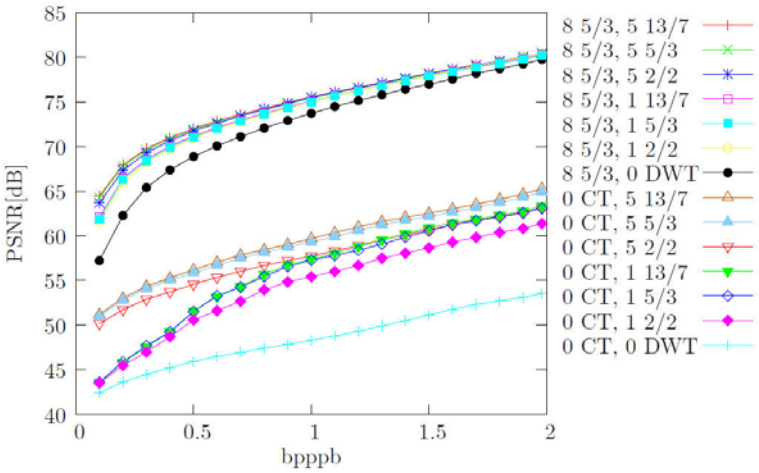

Fig. 13. A comparison in terms of PSNR versus the bit-rate.

packets is LRCP, producing a progression of incremental quality reconstructions (one for each tested bit-rate). Since we are not interested in ROI decoding, the number of precincts per resolution level is 1, minimizing the number of packets and thus minimizing the overhead of the headers of these packets. To maximize the compression ratio, the code-block size is 64×64 (the maximum allowed). Resilience data, such as SOP or EPH markers, has not been included. Finally, tiling has not been used.

In order to establish the experimental conditions for the evaluation of the impact of lossy compression on spectral unmixing techniques, a set of preliminary experiments has been carried out. In this first set of experiments the distortion of the reconstructed scene in terms of the Peak Signal to Noise Ratio (PSNR) has been measured as a function of: (i) the bit-rate (0.1-2.0 bpppb); (ii) the number of levels of the CT (0 and 8); and (iii) the number of levels of the ST (0, 1 and 5) for all the wavelet filters ($2/2$, $5/3$ and $13/7$). The maximum number of levels of the CT (8) and the ST (5) have been selected empirically, minimizing the distortion.

As can be seen in Fig. 13, the encoding performance in terms of the PSNR increases with the number of levels of the CT, the number of levels of the ST and the length of the filters, but the relevance of these experimental parameters is different. The PSNR significantly depends on the number of levels of the CT. The number of levels of the ST has a medium impact on the performance. Finally, the length of the wavelet filter, specially along the inter-component domain has a minimum impact. From these results it can be concluded that (i) the inter-component redundancy is much higher than the intra-component redundancy in the scene, and (ii) the wavelet kernel used in the CT could be any of the tested ones. For this reason, in the rest of experiments we will use only the $5/3$ kernel (for the CT) that has a good encoding performance and a reasonable computational cost.

5.3 Performance metrics

The quantitative and comparative algorithm assessment that is performed in this work is intended to measure the impact of compression on the quality of the endmembers extracted by OSP, AMEE and SSEE, and on the fractional abundances estimated by FCLSU using the endmembers provided by such methods. The quality of the endmembers extracted from the decompressed scene (i.e., the one obtained after coding/decoding the scene using JPEG2000) is measured in terms of their spectral similarity with regards to the endmembers extracted from the original scene. Let us denote by $\{\mathbf{E}_z\}_{z=1}^{p}$ the set of p endmembers extracted by a certain

method from the original scene, and let us denote by $\{\mathbf{E}'_z\}_{z=1}^{p}$ the set of p endmembers extracted from the decompressed scene. The spectral similarity between an endmember extracted from the original scene, $\mathbf{E}_o$, and an endmember extracted from the decompressed scene, $\mathbf{E}'_d$, is measured by the spectral angle (SA), a well known metric for hyperspectral data processing [4] which is defined as follows:

$$\mathrm{SA}(\mathbf{E}_o, \mathbf{E}'_d) = \cos^{-1} \frac{\mathbf{E}_o \cdot \mathbf{E}'_d}{\|\mathbf{E}_o\| \|\mathbf{E}'_d\|}. \tag{12}$$

It should be noted that SA is given by the arc cosine of the spectral angle formed by n-dimensional vectors. As a result, this metric is invariant in the multiplication of $\mathbf{E}_o$ and $\mathbf{E}'_d$ by constants and, consequently, is invariant before unknown multiplicative scalings that may arise due to differences in illumination and angular orientation [4]. Another way to compare $\mathbf{E}_o$ and $\mathbf{E}'_d$ is to analyze their associated fractional abundance maps, a metric designated here as abundance error (AE) and defined as follows:

$$\mathrm{AE}(\Phi_o, \Phi'_d) = \frac{1}{s \times l} \sum_{i=1}^{s} \sum_{j=1}^{l} \left([\Phi_o(i,j) - \Phi'_d(i,j)]^2 \right)^{\frac{1}{2}} \tag{13}$$

where Φ_o is the fractional abundance map estimated by FCLSU for $\mathbf{E}_o$ and Φ'_d is the fractional abundance map estimated by FCLSU for $\mathbf{E}'_d$. Finally, the quality of the reconstruction of a hyperspectral image by means of Eq. (1) can also be established as follows. Let us assume that $\mathbf{I}^{(\mathrm{O})}$ is a hyperspectral scene, and that $\mathbf{I}^{(\mathrm{R})}$ is a reconstructed version of $\mathbf{I}^{(\mathrm{O})}$, obtained using Eq. (1) with a set of endmembers, automatically derived by a certain algorithm from the original scene, and their corresponding fractional abundances estimated using the FCLSU algorithm. Let us also assume that the pixel vector at spatial coordinates (i,j) in the original hyperspectral scene is given by $\mathbf{X}^{(\mathrm{O})}(i,j) = [x_1^{(\mathrm{O})}(i,j), x_2^{(\mathrm{O})}(i,j), \cdots, x_n^{(\mathrm{O})}(i,j)]$, while the corresponding pixel vector at the same spatial coordinates in the reconstructed hyperspectral scene is given by $\mathbf{X}^{(\mathrm{R})}(i,j) = [x_1^{(\mathrm{R})}(i,j), x_2^{(\mathrm{R})}(i,j), \cdots, x_n^{(\mathrm{R})}(i,j)]$. With the above notation in mind, the reconstruction error (RE) between the original and the reconstructed hyperspectral scenes is calculated as follows:

$$\mathrm{RE}(\mathbf{I}^{(\mathrm{O})}, \mathbf{I}^{(\mathrm{R})}) = \frac{1}{s \times l} \sum_{i=1}^{s} \sum_{j=1}^{l} \left(\frac{1}{n} \sum_{k=1}^{n} [x_k^{(\mathrm{O})}(i,j) - x_k^{(\mathrm{R})}(i,j)]^2 \right)^{\frac{1}{2}} \tag{14}$$

With the aforementioned metrics in mind, we can compare the sets $\{\mathbf{E}_z\}_{z=1}^{p}$ and $\{\mathbf{E}'_z\}_{z=1}^{p}$ using the following spectral similarity matching algorithm:

1. *Initialization*. Label all endmembers in the set of endmembers extracted from the decompressed scene, $\{\mathbf{E}'_z\}_{z=1}^{p}$, as 'unmatched.' Similarly, label all endmembers in the set of endmembers extracted from the original scene, $\{\mathbf{E}_z\}_{z=1}^{p}$, as 'unmatched.' Set the value of CumulativeSA $\leftarrow 0$ and the value of CumulativeAE $\leftarrow 0$.

2. *Matching*. For each unmatched endmember in $\{\mathbf{E}'_z\}_{z=1}^{p}$, calculate the SA between such endmember and all endmembers in the set of endmembers extracted from the original scene, $\{\mathbf{E}_z\}_{z=1}^{p}$. If the pair $\{\mathbf{E}_o, \mathbf{E}'_d\}$, with $1 \leq o \leq p$ and $1 \leq d \leq p$, results in the minimum obtained value of $\mathrm{SA}(\mathbf{E}_o, \mathbf{E}'_d)$ after comparing all possible combinations of endmembers across the two sets, then label the associated endmembers, $\mathbf{E}_o$ and $\mathbf{E}'_d$ as 'matched', and set CumulativeSA $\leftarrow$ CumulativeSA + $\mathrm{SA}(\mathbf{E}_o, \mathbf{E}'_d)$. Similarly, set CumulativeAE $\leftarrow$ CumulativeAE + $\mathrm{AE}(\Phi_o, \Phi'_d)$.

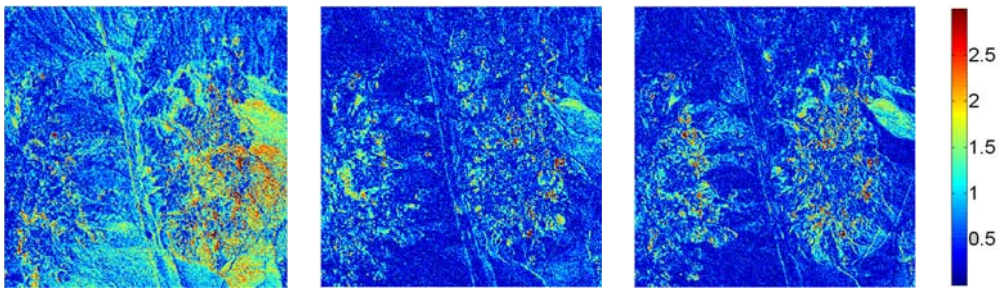

(A) OSP (0.918) (B) AMEE (0.544) (C) SSEE (0.553)

Fig. 14. Per-pixel RE scores after reconstructing the original AVIRIS Cuprite scene using FCLSU and the endmembers provided by different endmember extraction algorithms: (A) OSP, (B) AMEE and (C) SSEE. It should be noted that the values in the parentheses correspond to the total RE scores, obtained using Eq. (14), after reconstructing the original image using the endmembers provided by each considered endmember extraction algorithm.

3. *Iterative procedure and termination*. Each time a pair of endmembers $\{\mathbf{E}_o, \mathbf{E}'_d\}$ have been 'matched', the endmember $\mathbf{E}_o$ is removed from the set of endmembers extracted from the original image, $\{\mathbf{E}_z\}_{z=1}^{p}$, and the endmember $\mathbf{E}'_d$ is removed from the set of endmembers extracted from the decompressed image, $\{\mathbf{E}'_z\}_{z=1}^{p}$. Once these endmembers have been removed, the process is repeated from step 2 until all endmembers in both sets have been labeled as 'matched.' The average SA after comparing the two sets of endmembers can then be simply given by CumulativeSA$/p$, while the average AE is simply given by CumulativeAE$/p$.

5.4 Algorithm comparison

5.4.1 Accuracy of spectral unmixing algorithms when applied to the original scene

Before analyzing the impact of compression on the different endmember extraction and spectral unmixing algorithms, it is important to analyze the accuracy of these algorithms when applied to the original, uncompressed data set. This analysis is highly relevant in order to substantiate the loss of precision of these algorithms when they are applied to decompressed versions of the original image. For that purpose, firstly, it has been evaluated the capacity of the three considered endmember extraction algorithms (OSP, AMEE and SSEE) to extract endmembers from the AVIRIS Cuprite scene in Fig. 12(a) which are similar, in spectral terms, with regards to the ground-truth USGS signatures available for the scene and displayed in Fig. 12(b). Table 1 tabulates the SA scores (in radians) obtained after comparing the USGS library spectra of five highly representative minerals in the Cuprite mining district (*alunite*, *buddingtonite*, *calcite*, *kaolinite* and *muscovite*) with the corresponding endmembers extracted by OSP, AMEE and SSEE from the AVIRIS Cuprite scene. It should be noted that Table 1 only displays the smallest SA scores of all endmembers with respect to each USGS signature (after applying the adopted procedure for spectral matching) for each endmember extraction method. For reference, the average SA values are also reported. In all cases, the number of endmembers to be extracted was set to $p = 14$ after using the *virtual dimensionality* concept [18]. Table 1 reveals that the three compared methods provide very similar results, although it seems that the spatial-spectral methods (AMEE and SSEE) provide slightly better results than OSP in terms of average SA values in this particular experiment.

On the other hand, it is also important to evaluate the capacity of the FCLSU algorithm for abundance estimation to characterize mixed pixels in the original image, using the endmembers provided by the different endmember extraction algorithms considered in this work. Since ground-truth information concerning the true fractional abundances at sub-pixel levels is very

Table 1. SA-based spectral similarity scores (radians) between the USGS mineral spectra and their corresponding endmember pixels produced by several endmember extraction algorithms.

Algorithm	Alunite	Buddigntonite	Calcite	Kaolinite	Muscovite	Average
OSP	0.084	0.072	0.167	0.194	0.094	0.122
AMEE	0.084	0.073	0.166	0.152	0.080	0.111
SSEE	0.084	0.072	0.148	0.194	0.080	0.116

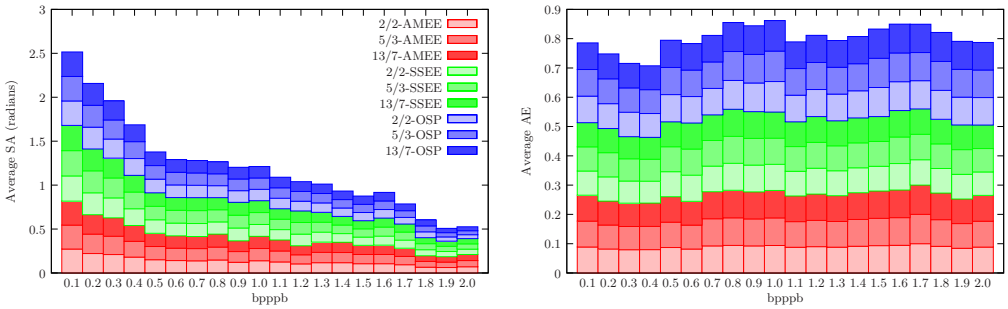

(A) Average SA scores for 0 CT, 0 ST. (B) Average AE scores for 0 CT, 0 ST.

Fig. 15. (A) Average SA scores and (B) Average AE scores (both displayed in the form of stacked bar plots) obtained after comparing the endmembers extracted by OSP, SSEE and AMEE algorithms from decompressed images, obtained using 0 levels of the component transform (CT) and 0 levels of the spatial transform (ST), with regards to the endmembers extracted by the same algorithms from the original image.

difficult to obtain in real analysis scenarios, our evaluation strategy consists of using Eq. (14) to evaluate the reconstruction accuracy of the scenes obtained after reconstructing the original one using FCLSU in combination with the endmembers provided by different algorithms, using the linear model described in Eq. (1). For illustrative purposes, Fig. 14 graphically represents the per-pixel RE obtained after reconstructing the AVIRIS Cuprite scene using the $p = 14$ endmembers extracted by different methods and the FCLSU algorithm for abundance estimation. It can be seen that the endmember extraction methods which use spatial information to complement spectral information (i.e., AMEE and SSEE) can provide more spatially representative endmembers, in the sense that their use slightly reduces the RE obtained with regards to endmembers obtained from the algorithms that only use the spectral information (i.e., OSP).

5.4.2 Impact of lossy compression on spectral unmixing algorithms

Once the precision of the considered endmember extraction and spectral unmixing algorithms has been analyzed with the original image, it is now time to evaluate to what extent degrading the original image after applying JPEG2000-based lossy compression can affect the performance of such algorithms. Figure 15(A) plots the average SA obtained after applying the spectral similarity matching algorithm in Subsection 5.3 to evaluate the quality of the endmembers derived by OSP, AMEE and SSEE from the scenes obtained after compressing/decompressing the original hyperspectral scene using 0 levels of the component transform (CT), 0 levels of the spatial (DWT-based) transform (ST), and different filter lengths. Here, the USGS mineral signatures in Fig. 12(b) were used to identify the most similar spectral signatures corresponding to those minerals in the original hyperspectral image, and used those signatures as a reference for comparisons. It should be noted that the plots in Fig. 15(A) –as well as all the plots reported hereinafter in the manuscript– adopt a stacked bars representation. This means that, for each bit rate in Fig. 15(A), the SA obtained using the AMEE algorithm (in red color) is first plotted, then the error scores obtained using the SSEE algorithm (in green color) on top of those ob-

tained by the AMEE are stacked, and finally the error scores obtained using the OSP algorithm (in blue color) have also been stacked on top of those obtained using AMEE and SSEE in the same plot bar. Since a stacked plot representation has been used, the value displayed in each algorithm plot is the contribution of each algorithm to the cumulative SA. From Fig. 15(A), it can be observed that (as expected) the average SA tends to increase with the increase of the compression ratio (decrease of the bpppb), providing the worst spectral similarity scores for the highest compression ratio tested in experiments (obtained for 0.1 bpppb, which represents a compression ratio of 160:1 of the original image). On the other hand, the best (lower) average SA values are obtained for the lowest compression ratio tested in experiments (obtained for 2 bpppb, which represents a compression ratio of 8:1 of the original image). At this point, it is very important to emphasize that, in this experiment, the length of the filter used in the ST is not modified and, hence, the display of different values associated to different filter lengths in Fig. 15(A) is only intended to maintain consistency with regards to other plots in which the length of the filter introduces variations in the obtained results. With the above observation in mind, it is interesting to note that the three compared endmember extraction algorithms show very different sensitivities to the average SA metric. Specifically, the OSP is clearly the best algorithm for low compression ratios, with SSEE providing the worst spectral similarity results in this case. However, as the compression ratio increases the performance of the three considered algorithms becomes very similar and almost identical for the highest compression ratio tested in experiments. This is due to the fact that using 0 levels of the CT and 0 levels of the ST generally results in highly degraded versions of the spectral signatures in the original image, and hence the algorithms cannot perform accurately since the fine spectral information required for endmember extraction purposes is severely degraded, specifically, for very high compression ratios.

On the other hand, Fig. 15(B) plots the average AE scores obtained after comparing the fractional abundances estimated during the matching process for the endmembers extracted by OSP, AMEE and SSEE from the scenes obtained after compressing/decompressing the original hyperspectral scene using 0 levels of the CT, 0 levels of the ST, and different filter lengths. As in the experiments reported on Fig. 15(A), the filter length is not varied. Therefore, this is not a relevant parameter in the experiments reported on Fig. 15(B), in which the results for different filter lengths are only displayed for consistency of visualization with regards to other experimental results reported in this section. It is also worth noting that, due to the poor quality of the images after compression using 0 levels of the CT and 0 levels of the ST, the scores reported for the average AE metric are also quite poor in all cases and there seems to be no quality increase as the compression ratio increases, as opposed to the SA-based experiments reported in Fig. 15(A). Interestingly, in these average AE-based experiments the three algorithms perform similarly in all cases, which is in contrast with the results reported in Fig. 14 in which both AMEE and SSEE outperformed OSP in terms of RE when applied to the original image. This reveals that lossy compression with the parameters considered in this example significantly degrade the original image and negatively affects all algorithms.

Figure 16(A) shows the SA-based scores after increasing the number of levels in the ST from 0 to 5, while maintaining the number of levels of the CT to 0. As it can be observed by comparing the results in Fig. 16(A) with those in 15(A), the use of a 5-levels ST decreases the average SA in a factor of two, approximately. On the other hand, a comparison of Fig. 16(B) showing the AE-based scores in this case with regards to the results reported in Fig. 15(B) indicates that increasing the number of ST levels allows reducing the average AE. Also, these errors are now higher as the compression ratio increases, which is a more reasonable behaviour than the one observed in Fig. 15(B) in which peaks and valleys –instead of the more progressive trend now reported in Fig. 16(B)– were observed for different compression ratios. Here, an interesting observation for the lowest compression ratio tested in experiments is that SSEE is clearly the worst algorithm; this observation is in contrast with the results reported in Fig. 14,

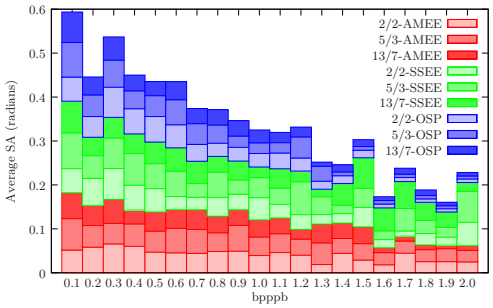

(A) Average SA scores for 0 CT, 5 ST.

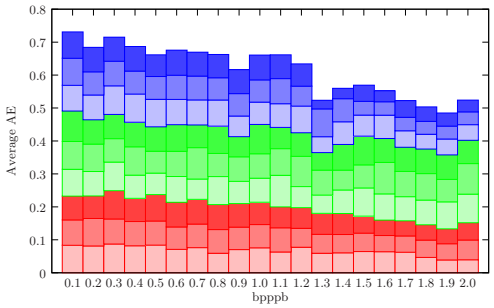

(B) Average AE scores for 0 CT, 5 ST.

Fig. 16. (A) Average SA scores and (B) Average AE scores (both displayed in the form of stacked bar plots) obtained after comparing the endmembers extracted by OSP, SSEE and AMEE algorithms from decompressed images, obtained using 0 levels of the component transform (CT) and 5 levels of the spatial transform (ST), with regards to the endmembers extracted by the same algorithms from the original image.

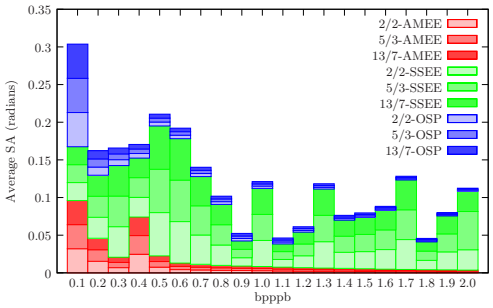

(A) Average SA scores for 8 CT, 0 ST.

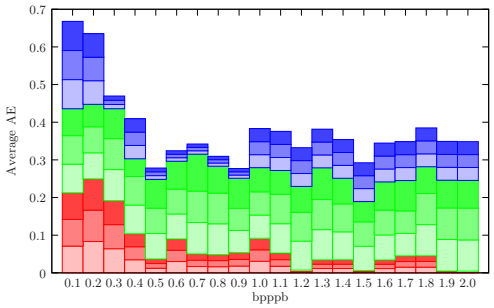

(B) Average AE scores for 8 CT, 0 ST.

Fig. 17. (A) Average SA scores and (B) Average AE scores (both displayed in the form of stacked bar plots) obtained after comparing the endmembers extracted by OSP, SSEE and AMEE algorithms from decompressed images, obtained using 8 levels of the component transform (CT) and 0 levels of the spatial transform (ST), with regards to the endmembers extracted by the same algorithms from the original image.

in which it was shown that SSEE was among the best algorithms in terms of average AE when applied to the original image. This indicates that lossy compression negatively affects the performance of the three compared endmember extraction algorithms, and this effect seems to be more relevant for the SSEE than for the other tested algorithms, which perform more accurately than SSEE in terms of the average AE metric for low compression ratios. As the compression ratio increases, the performance of the three considered algorithms becomes similar in terms of the average AE metric as it was already the case for the average SA metric.

Figure 17(A) shows the SA-based scores after increasing the number of levels in the CT from 0 to 8, while setting the number of levels of the ST to 0. As it can be observed in this experiment, the average SA scores for OSP and AMEE improve significantly with regards to the cases in which the number of levels in the CT were 0. The worst results for all algorithms are clearly obtained for the highest compression ratio tested in experiments, in which the obtained SA values are comparatively similar to those already reported in Figs. 15(A) and 16(A). However, for the remaining compression ratios both OSP and AMEE provide much better results, indicating that the number of levels used in the CT is a very important parameter in order to accurately preserve the quality of the spectral signatures in the original hyperspectral scene.

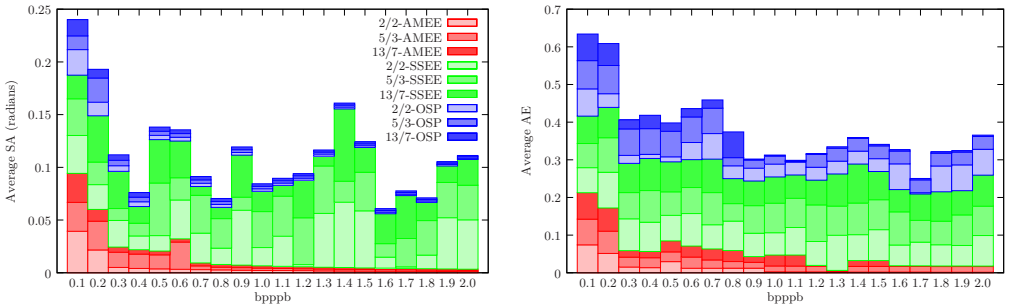

(A) Average SA scores for 8 CT, 5 ST. (B) Average AE scores for 8 CT, 5 ST.

Fig. 18. (A) Average SA scores and (B) Average AE scores (both displayed in the form of stacked bar plots) obtained after comparing the endmembers extracted by OSP, SSEE and AMEE algorithms from decompressed images, obtained using 8 levels of the component transform (CT) and 5 levels of the spatial transform (ST), with regards to the endmembers extracted by the same algorithms from the original image.

Interestingly, the SSEE algorithm cannot perform as accurately as OSP and AMEE in this experiment, in particular, for the lowest compression ratios tested. Similar observations can be made for the average AE-based scores obtained for 8 levels of the CT and 0 levels of the ST and reported in 17(B). In this case, again SSEE provides comparatively less effective results in terms of average AE than AMEE, which is clearly the most effective algorithm in this experiment, in particular, for the lowest compression ratios tested in experiments. As the compression ratio increases, the low quality of the obtained images causes all algorithms to under-perform in terms of average AE, as observed for the SA-based experiments in Fig. 17(A).

Figure 18(A) shows the SA-based scores after increasing the number of levels in the ST from 0 to 5, while maintaining the number of levels of the CT to 8. Since this is the highest quality configuration for data compression tested in experiments, it is not surprising that the average SA scores for OSP and AMEE are the best ones reported in this work, in particular, for the lowest compression ratios. Again, the SSEE algorithm cannot perform as accurately as OSP and AMEE in this experiment, in particular, for the lowest compression ratios tested. As the compression ratio increases, the performance of all algorithms becomes similar as already reported in previous experiments. Similar observations can be made for the average AE scores obtained for 8 levels of the CT and 5 levels of the ST and reported in Fig. 18(B). In this case, AMEE is clearly the best algorithm for most compression ratios tested. Only for very high compression ratios (80:1 and beyond) the performance of the three considered endmember extraction algorithms becomes similar. This confirms the introspection, already observed in other experiments, that AMEE seems to be the most effective algorithm from the viewpoint of the average AE metric, while both OSP and AMEE are the most effective from the viewpoint of the SA-based metric, with all algorithms under-performing from the viewpoint of all considered metrics when the compression ratio is very high and with a clear need to consider sufficient levels of the CT and the ST in order to preserve accurately enough the fine spectral information needed for endmember extraction and unmixing algorithms and to allow such algorithms to take advantage of the detailed spectral information carried out by the spectral signatures without significantly degrading them.

An important observation at this point is that the endmember extraction process is essentially nonlinear in nature, and can therefore produce different endmembers for different compression ratios. It should be noted that the values correspond to the average values measured for a set of endmembers which have been previously matched to a set of reference signatures using the three-step spectral similarity matching algorithm described in our paper. The matching process

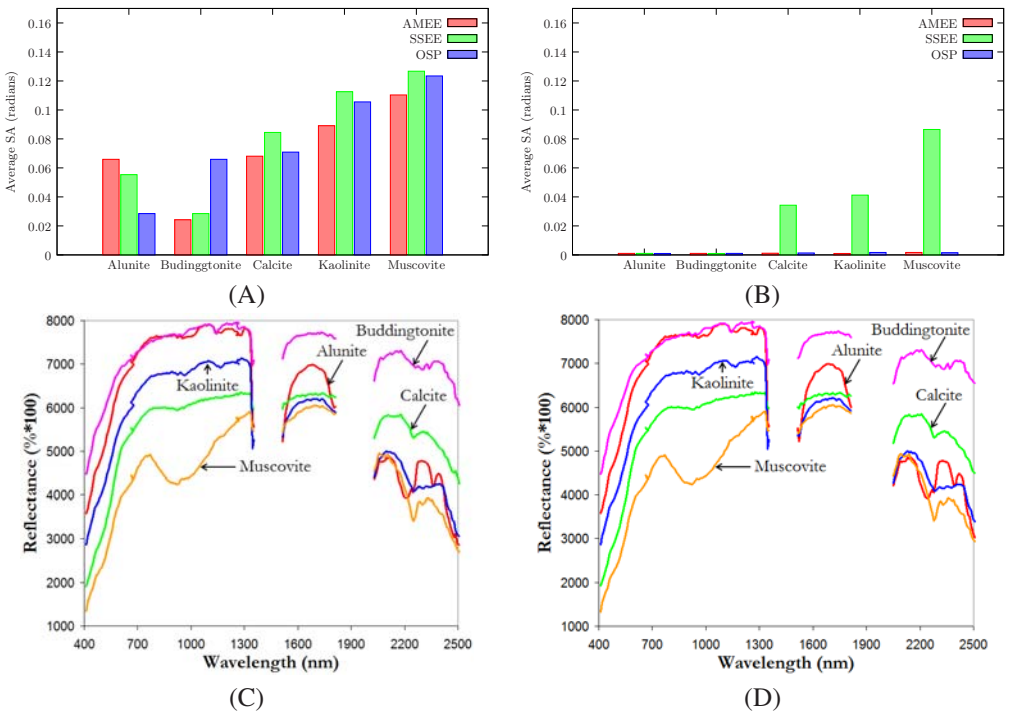

Fig. 19. SA values obtained by the considered endmember extraction algorithms in the worst (A) and best (B) cases reported in Fig. 18. (C) Endmembers extracted by AMEE from the original image with water absorption and noisy bands removed. (D) Endmembers extracted by AMEE from the compressed image (2.0 bpppb) with water absorption and noisy bands removed.

is also highly nonlinear, since different matching combinations (for example, resulting from the fact that an endmember is not detected for a certain bit-rate) may result in different average scores. If we had compared the spectral signatures at the endmember pixel locations in the original image with the same pixel locations in the compressed images, the behavior of the SA and AE values would have been linear in nature. However, endmember extraction algorithms are expected to be able to automatically locate such positions and it seems that a matching-based procedure is the most appropriate one in order to evaluate their performance in the presence of lossy compression.

Since the results reported thus far correspond to average values, it is also of interest to present more detailed analyses for a representative case. For this purpose, Fig. 19 reports the SA values obtained by the considered endmember extraction algorithms in the best and worst cases reported in Fig. 18, and further compares the quality of the endmembers extracted in each case with regards to the mineral signatures (alunite, buddingtonite, calcite, kaolinite and muscovite) in the original image. These plots resemble histograms in which the SA of the endmembers which are matched to the five reference signatures is displayed. In the plots displayed in Fig. 19(A) –associated to the worst case observed in experiments, i.e., the one with higher compression ratio: 0.1 bpppb– it can be observed that all algorithms produce endmembers with similar quality when compared to the reference signatures in the original image. In turn, the plots displayed in Fig. 19(B) –associated to the best case observed in experiments, i.e., the one with lower compression ratio: 2.0 bpppb– indicates that the SSEE provides spectral signatures which are less similar, spectrally, to the reference ones. This is the reason why the stacked plots in Fig. 18(A) reveal a poorer behaviour of this algorithm, regardless of the compression

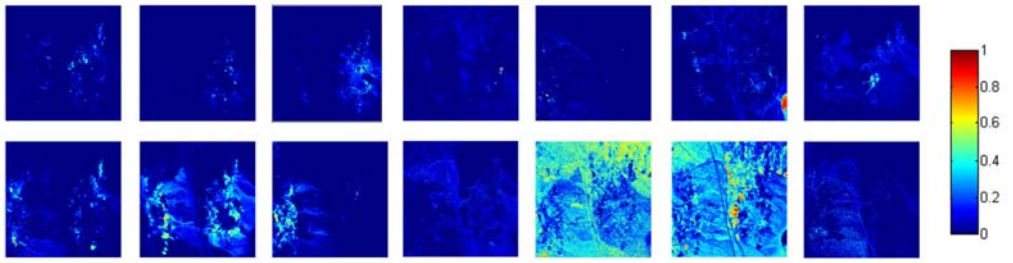

Fig. 20. Fractional abundance maps estimated by FCLSU for the $p = 14$ endmembers extracted by AMEE from the original hyperspectral image.

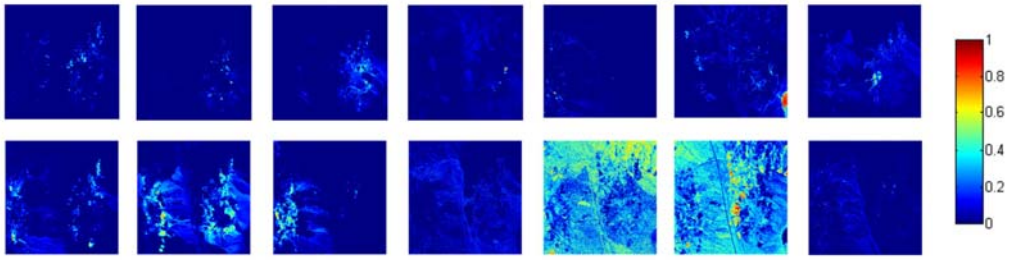

Fig. 21. Fractional abundance maps estimated by FCLSU for the $p = 14$ endmembers extracted by AMEE from the hyperspectral image compressed with the lower compression ratio reported in Fig. 18.

ratio used. In turn, the AMEE provided consistently similar endmembers with regards to the reference signatures. For illustrative purposes, Figs. 19(C) and 19(D) respectively show the alunite, buddingtonite, calcite, kaolinite and muscovite endmembers extracted by AMEE from the original and the compressed image, indicating very high spectral similarity between the two endmember sets. Also, Fig. 20 shows the fractional abundance maps estimated by FCLSU for the $p = 14$ endmembers extracted by AMEE from the original image, while Fig. 21 shows the fractional abundance maps estimated by FCLSU for the $p = 14$ endmembers extracted by AMEE from the image compressed with the lower compression ratio (2.0 bpppb). Finally, Fig. 22 shows the differences between the maps in Fig. 20 and those in Fig. 21 calculated using the AE metric, revealing that the maps obtained from the compressed scene are almost identical to those obtained from the original image for the lower compression ratio, although the differences become more apparent as the compression ratio is increased as indicated by the results in Fig. 18(B).

To conclude this section, some observations about the AMEE and SSEE algorithms should be raised, which are the two considered methods including both spatial and spectral information in the search for spectral information. As mentioned in the introductory section of this paper, the inclusion of spatial information is expected to introduce some advantages in terms of selecting more spatially representative endmembers, but it is also important to substantiate whether the inclusion of spatial information can also introduce enhanced properties in terms of the tolerance of the algorithms to lossy compression. From this perspective, a detailed inter-comparison of the results provided by AMEE and SSEE algorithms indicates that AMEE seems to provide more consistent results while SSEE often provides peaks and valleys in the SA and average

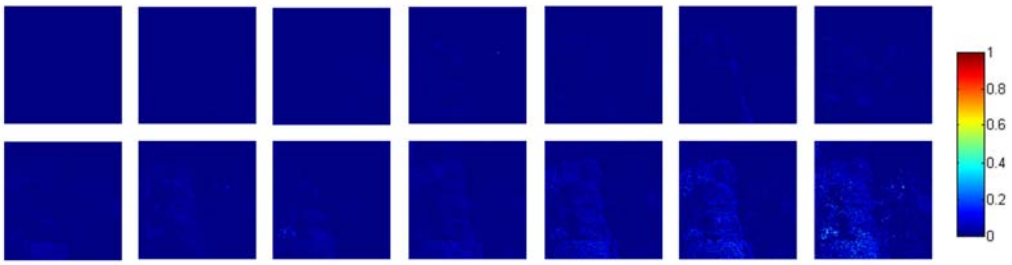

Fig. 22. Differences between the abundance maps in Figs. 20 and 21.

AE-based scores plots versus the compression ratios. This is probably due to the fact that SSEE does not simultaneously integrate the spatial and the spectral information when searching for image endmembers (the algorithm is based on an SVD-based spectral procedure followed by spatial averaging). However, the extended morphological operations used for the implementation of AMEE jointly consider the two sources of information (spatial and spectral) when searching for endmembers and without giving more priority to one of the sources versus the other, which apparently provides more stable results in terms of the sensitivity of the algorithm to different parameters tested. Further experimentation should be conducted in future work in order to validate the above remarks on different case studies using additional hyperspectral scenes and spatial-spectral endmember extraction algorithms. Despite the analysis in this work has been confined to three algorithms only (OSP, AMEE and SSEE), the OSP is a widely used approach for spectral-based endmember extraction and both AMEE and SSEE are among the most representative methods in the recent spatial-spectral endmember extraction literature. As a result, it becomes clear that the experimental findings reported in this work will be of interest to practitioners interested in applying those algorithms to hyperspectral images with different quality levels resulting from lossy compression.

6 CONCLUSIONS AND FUTURE WORK

The wealth of spatial and spectral information provided by hyperspectral imaging instruments is essential in many applications, and in many cases needs to be retained by compression algorithms with a high degree of fidelity. In this paper, the impact of lossy, JPEG2000-based hyperspectral data compression on spectral unmixing applications has been evaluated, with particular attention on endmember extraction using representative algorithms such as OSP, AMEE and SSEE, combined with a FCLSU-based approach for abundance estimation subject to ASC and ANC constraints. The experiments that have been carried out reveal that the OSP is a very good choice for spectral-based endmember extraction, regardless of the compression ratio tested, since the algorithm has the capacity to perform well in terms of different quality metrics considered in this work (SA and AE), for a wide range of compression ratios. Regarding the performance of endmember extraction algorithms using both spatial and spectral information such as AMEE or SSEE, the experimental results indicate that the AMEE provides slightly more robust results in terms of the SA-based quality of the endmembers extracted and the average AE-based quality of the associated fractional abundances. This may be due to the fact that the AMEE considers jointly the spatial and the spectral information (without prioritizing any of them) during the endmember searching process, while the SSEE first uses a spectral-based procedure followed by spatial averaging and, thus, the two sources of information are not treated simultaneously. In conclusion, the detailed experimental assessment revealed that endmember extraction algorithms are indeed sensitive to lossy compression, and that the incorporation of spatial information may help mitigate such sensitivity when the spectral information

is combined with spatial information in simultaneous fashion.

The investigation of the maximum compression ratio that can be achieved without compromising the quality of the endmembers extracted by different algorithms will be continued in future research by incorporating additional hyperspectral scenes and endmember extraction algorithms to the comparison. This study may be useful to image analysts interested in assessing the impact of lossy compression on spectral mixture analysis applications. Another important topic to explore in future research is the sensitivity of methods designed to estimate the number of endmembers (such as the virtual dimensionality concept) to different compression ratios. In separate experiments it has been observed that the estimation of the number of endmembers may be sensitive to lossy compression. This calls for a detailed study of this aspect. In this particular case, a lossy-to-lossless compression algorithm has been used and it would be possible to resort to the full data in order to estimate the number of endmembers.

Acknowledgments

This work has been supported by the European Community's Marie Curie Research Training Networks Programme under reference MRTN-CT-2006-035927, Hyperspectral Imaging Network (HYPER-I-NET). Funding from the Spanish Ministry of Science and Innovation (AYA2008-05965-C04-02 and TIN2008-01117) is also gratefully acknowledged. Gabriel Martín is sponsored by a research fellowship with reference BES-2009-017737 associated to AYA2008-05965-C04-02 project.

References

[1] G. Motta, F. Rizzo, and J. A. Storer, *Hyperspectral data compression*, Springer-Verlag, New York (2006).

[2] R. O. Green, M. L. Eastwood, C. M. Sarture, T. G. Chrien, M. Aronsson, B. J. Chippendale, J. A. Faust, B. E. Pavri, C. J. Chovit, M. Solis, *et al.*, "Imaging spectroscopy and the airborne visible/infrared imaging spectrometer (AVIRIS)," *Rem. Sens. Environ.* **65**, 227–248 (1998). [doi:10.1117/1.2338565].

[3] Q. Du and C.-I. Chang, "Linear mixture analysis-based compression for hyperspectral image analysis," *IEEE Trans. Geosci. Rem. Sens.* **42**, 875–891 (2004). [doi:10.1109/IGARSS.2000.861638].

[4] N. Keshava and J. F. Mustard, "Spectral unmixing," *IEEE Signal Process. Mag.* **19**, 44–57 (2002). [doi:10.1109/79.974727].

[5] J. B. Adams, M. O. Smith, and P. E. Johnson, "Spectral mixture modeling: a new analysis of rock and soil types at the Viking Lander 1 site," *J. Geophys. Res.* **91**, 8098–8112 (1986). [doi:10.1029/JB091iB08p08098].

[6] A. Plaza, P. Martinez, R. Perez, and J. Plaza, "A quantitative and comparative analysis of endmember extraction algorithms from hyperspectral data," *IEEE Trans. Geosci. Remote Sens.* **42**, 650–663 (2004). [doi:10.1109/TGRS.2003.820314].

[7] D. Heinz and C.-I. Chang, "Fully constrained least squares linear mixture analysis for material quantification in hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.* **39**, 529–545 (2001). [doi:10.1109/36.911111].

[8] J. W. Boardman, F. A. Kruse, and R. O. Green, "Mapping Target Signatures Via Partial Unmixing of Aviris Data," *Proc. JPL Airborne Earth Sci. Workshop* , 23–26 (1995).

[9] M. E. Winter, "N-FINDR: an algorithm for fast autonomous spectral end-member determination in hyperspectral data," *Proc. SPIE Image Spectrometry V* **3753**, 266–277 (2003).

[10] R. A. Neville, K. Staenz, T. Szeredi, J. Lefebvre, and P. Hauff, "Automatic endmember extraction from hyperspectral data for mineral exploration," *Proc. 21st Canadian Symp. Remote Sens.* , 21–24 (1999).

[11] J. H. Bowles, P. J. Palmadesso, J. A. Antoniades, M. M. Baumback, and L. J. Rickard, "Use of filter vectors in hyperspectral data analysis," *Proc. SPIE Infrared Spaceborne Remote Sensing III* **2553**, 148–157 (1995).

[12] A. Ifarraguerri and C.-I. Chang, "Multispectral and hyperspectral image analysis with convex cones," *IEEE Trans. Geosci. Remote Sens.* **37**, 756–770 (1999). [doi:10.1109/36.752192].

[13] J. C. Harsanyi and C.-I. Chang, "Hyperspectral image classification and dimensionality reduction: An orthogonal subspace projection," *IEEE Trans. Geosci. Remote Sens.* **32**, 779–785 (1994). [doi:10.1109/36.298007].

[14] A. Plaza, P. Martinez, R. Perez, and J. Plaza, "Spatial/spectral endmember extraction by multidimensional morphological operations," *IEEE Trans. Geosci. Remote Sens.* **40**, 2025–2041 (2002). [doi:10.1109/TGRS.2002.802494].

[15] D. M. Rogge, B. Rivard, J. Zhang, A. Sanchez, J. Harris, and J. Feng, "Integration of spatial–spectral information for the improved extraction of endmembers," *Remote Sens. Environ.* **110**, 287–303 (2007). [doi:10.1016/j.rse.2007.02.019].

[16] D. S. Taubman and M. W. Marcellin, *JPEG2000: Image compression fundamentals, standard and practice*, Kluwer, Boston (2002).

[17] A. Plaza, P. Martinez, J. Plaza, and R. Perez, "Dimensionality reduction and classification of hyperspectral image data using sequences of extended morphological transformations," *IEEE Trans. Geosci. Remote Sens.* **43**, 466–479 (2005). [doi:10.1109/TGRS.2004.841417].

[18] C.-I. Chang, *Hyperspectral Imaging: Techniques for Spectral Detection and Classification*, Kluwer Academic/Plenum Publishers: New York (2003).

[19] C.-I. Chang and Q. Du, "Estimation of number of spectrally distinct signal sources in hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.* **42**, 608–619 (2004). [doi:10.1109/TGRS.2003.819189].

[20] C.-I. Chang and D. Heinz, "Constrained subpixel target detection for remotely sensed imagery," *IEEE Trans. Geosci. Remote Sens.* **38**, 1144–1159 (2000). [doi:10.1109/36.843007].

[21] International Organization for Standardization, "Information Technology - JPEG2000 Image Coding System - Part 1: Core Coding System." ISO/IEC 15444-1:2004 (2004).

[22] International Organization for Standardization, "Information Technology - JPEG2000 Image Coding System - Part 2: Extensions." ISO/IEC 15444-2:2004 (2004).

[23] International Organization for Standardization, "Information technology - JPEG2000 image coding system: Extensions for three-dimensional data." ISO/IEC 15444-10:2008 (2008).

[24] D. Taubman, "High performance scalable image compression with EBCOT," *IEEE Trans. Image Process.* **9**, 1158–1170 (2000). [doi:10.1109/83.847830].

[25] A. Haar, "Zur Theorie der orthogolanen Funktionen-Systeme," *Math. Ann.* **69**, 331–371 (1910).

[26] D. L. Gall and A. Tabatabai, "Sub-band coding of digital images using sysmmetric short kernel filters and arithmetic coding techniques," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, **2**, 761–764, (New York) (1988).

[27] G. Strang and T. Nguyen, *Wavelets and Filter Banks*, Wellesley-Cambridge Press (1996).

[28] W. Sweldens, "The Lifting Scheme: A new Philosophy in Biorthogonal Wavelet Constructions," in *Proc. SPIE*, **2569**, 68–79 (1995).

[29] C. M. Brislawn, "Preservation of Subband Symmetry in Multirate Signal Coding," *IEEE Trans. Signal Process.* **43**, 3046–3050 (1995). [doi:10.1109/78.476454].

[30] R. N. Clark, G. A. Swayze, K. E. Livo, R. F. Kokaly, S. J. Sutley, J. B. Dalton, R. R. McDougal, and C. A. Gent, "Imaging spectroscopy: Earth and planetary remote sensing with the USGS tetracorder and expert systems," *J. Geophys. Res.* **108**, 1–44 (2003). [doi:10.1029/2002JE001847].

[31] G. Swayze, R. N. Clark, F. Kruse, S. Sutley, and A. Gallagher, "Ground-truthing AVIRIS mineral mapping at Cuprite, Nevada," *Proc. JPL Airborne Earth Sci. Workshop* , 47–49 (1992).