Pairwise Orthogonal Transform for Spectral Image Coding

Ian Blanes, *Student Member, IEEE*, and Joan Serra-Sagristà, *Member, IEEE*

***Abstract*—Spectral transforms are widely used for the codification of remote-sensing imagery, with the Karhunen–Loêve transform (KLT) and wavelets being the two most common transforms. The KLT presents a higher coding performance than the wavelets. However, it also carries several disadvantages: high computational cost and memory requirements, difficult implementation, and lack of scalability. In this paper, we introduce a novel transform based on the KLT, which, while obtaining a better coding performance than the wavelets, does not have the mentioned disadvantages of the KLT. Due to its very small amount of side information, the transform can be applied in a line-based scheme, which particularly reduces the transform memory requirements. Extensive experimental results are conducted for the Airborne Visible/Infrared Imaging Spectrometer and Hyperion images, both for lossy and lossless and in combination with various hyperspectral coders. The results of the effects on Reed Xiaoli anomaly detection and *k*-means clustering are also included. The theoretical and experimental evidences suggest that the proposed transform might be a good replacement for the wavelets as a spectral decorrelator in many of the situations where the KLT is not a suitable option.**

***Index Terms*—Embedded systems, hyperspectral image coding, Karhune–Loêve transform (KLT), memory-constrained environments, progressive lossy-to-lossless (PLL) and lossy compression.**

I. Introduction

SPECTRAL transforms are receiving much attention in several applications of the remote-sensing field, for instance, in dimensionality reduction for classification or monitoring purposes. However, spectral transforms are particularly relevant in hyperspectral image compression because of the high spectral redundancy among components, which, when exploited, yields substantial coding gains [1], thus improving storage and transmission capabilities.

The most common spectral transforms in hyperspectral image coding are wavelet transforms and the Karhunen–Loêve transform (KLT). The KLT produces more competitive coding results as it is trained for each input image to provide optimal decorrelation. Nonetheless, several significant disadvantages have been identified for the KLT: its extremely high computational cost and memory requirement, its higher implementation difficulty, and that it is not a scalable transform, impairing component scalability. These disadvantages restrain the use of the KLT in several situations, such as in interactive or real-time processing, and in power- or memory-constrained environments like on-board sensors. The wavelets, on the other hand, lack the superior coding gains of the KLT, but they have much lower computational costs and memory requirements, which are compatible with the resource-constrained scenarios described before. In addition, they also provide component scalability.

The particularities of two scenarios where a hyperspectral transform coder might be used are worth discussing in detail. When a transform coder is used on-board a plane or a spacecraft, the image coder has to process the uncalibrated information that might have varying features if compared with the calibrated versions of the same information, such as, for example, echo, smear, or streaking artifacts [2]. In this first scenario, the important features of a coder are to have low forward computational and memory costs and, if the bandwidth from the sensor is limited, to be able to operate in the lossy or progressive lossy-to-lossless (PLL) mode. In the second scenario, which is document retrieval from an archive, the important features are coding gains for both lossy and lossless, PLL operation, high component scalability, and acceptable computational costs for decoding.

When analyzing the origins of the four cited disadvantages of the KLT, we note that two of them—its high computational cost and not being a scalable transform—are direct consequences of its optimal decorrelation property, which requires all of the transform inputs to be present in the calculation of each output. The third disadvantage—high memory requirement—is caused by the two sequential stages of the forward transform, i.e., the whole image is first inspected during the training stage and is then processed for the second time for the transform application. Moreover, the fourth disadvantage—higher implementation difficulty—is due to the difficulties posed by the iterative numerical algorithm commonly used in the last part of the training stage.

To overcome the high computational cost, the discrete cosine transform (DCT) was proposed [3], which assumes a Toeplitz matrix as the data covariance matrix. However, the DCT has a poor performance as a spectral decorrelator [4]. Similar approaches in reducing the computational cost are the fast approximate KLT (AKLT) [5] and the AKLT_2 [6], which extend the DCT with first- and second-order perturbations.

If the spatial size of an image is large enough, the transform training costs, which amount to one-fifth of the total computational cost, can be reduced, as proposed in [4]. The proposal is based on the extrapolation of the empirical covariance matrix from a small sample of the whole input and is shown to be

Manuscript received March 31, 2010; revised June 15, 2010; accepted August 6, 2010. Date of publication October 11, 2010; date of current version February 25, 2011. This work was supported in part by the Spanish Government, by the Catalan Government, and by Fondo Europeo de Desarrollo Regional under Grants TIN2009-14426-C02-01, TIN2009-05737-E/TIN, SGR2009-1224, and FPU2008.

The authors are with the Department of Information and Communications Engineering, Universitat Autònoma de Barcelona, 08193 Barcelona, Spain (e-mail: ian.blanes@uab.es; joan.serra@uab.es).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TGRS.2010.2071880

effective with very small sampling factors. On the other hand, if the spatial size of an image is significantly small or if the transform is applied in multiple spatial blocks, then the eigenvalue decomposition (ED) costs are more significant, and methods such as that in [7] might reduce them. The training costs can also be diminished by pretraining a transform for a generic set of images [8], [9]. Pretraining has a drawback, i.e., it produces inferior coding gains because of the lack of specificity.

Two further improvements to the KLT are also worth pointing out: first, for lossless coding, a method is proposed to replace the multiplication operations in the application and removal stages of the transform by incorporating them into the lifting decomposition [10], [11], thus applying a transform with only the addition and shift operations, and second, the coding gains can be improved by not assuming the Gaussian sources and by finding the optimal spectral transform with an algorithm based on independent component analysis [12], [13], which can also be pretrained [14], [15].

In addition, to overcome the high computational cost and the lack of component scalability disadvantages of the KLT, a few recent papers propose to use a divide-and-conquer approach in the transform application to achieve substantial gains in computational costs and reasonable scalability, still obtaining a virtually equal coding performance. In [16], the computational costs are reduced for the lossless compression of the magnetic resonance images by using a recursively clustered KLT. Similarly, we developed a multilevel clustered KLT structure for PLL coding [17]. This structure yields greater savings and leads to the family of structures discussed in [18], which have reduced costs and have a better scalability. Other divide-and-conquer approaches are discussed in [19] and [20], where two levels of clusters are applied, and in [21], where a single-level clustering is used to reduce the amount of side information and to improve the transform performance on low bitrates for images with a small spatial size. Nonetheless, all previous divide-and-conquer approaches still present difficult implementations and memory requirements of the order of those of the KLT.

In this paper, we employ a divide-and-conquer strategy to propose a very simple transform that requires a very small amount of side information, which, in its turn, allows a line-based application of the transform, which is not previously available to other divide-and-conquer approaches. As will be shown, the proposed transform still has a very good coding performance—better than the wavelet transforms—and it is capable of overcoming all the four related inconveniences of the KLT, achieving notably lower computational costs and better component scalability, particularly with a very low memory requirement.

This paper is organized as follows. Section II introduces the proposed pairwise orthogonal transform (POT). Section III contains the experimental results. Section IV draws the conclusion.

II. POT

First, this section introduces the proposed simple spectral transform for hyperspectral image coding, and then, it considers some implementation issues. Finally, it provides an analytical evaluation of its properties.

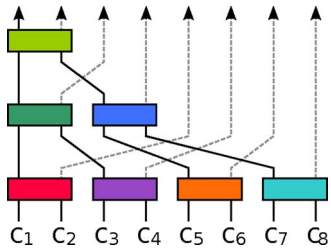

Fig. 1. Proposed pairwise structure for an image of eight components, where each rectangle corresponds to a two-component KLT.

A. Transform Definition

The proposed transform is based on the application of a divide-and-conquer strategy to the KLT, where the resulting transform is a composition of smaller KLT transforms, namely, a composition of KLTs for only two image components each.

In a full KLT, all components are decorrelated with each other, independently of how much energy they share, i.e., independently from how high is their covariance. On the contrary, the proposed POT has a structure that decorrelates parts with high shared energy while ignoring the other parts, as parts with low energies have a lower influence in the coding performance.

A POT is organized in multiple levels. In the first level, a two-component KLT transform is applied to every pair of consecutive components, and the same approach is repeated in successive levels but, at each level, only further decorrelating the first resulting component of each transform of the previous level. To clarify, the structure of a POT for eight components is shown in Fig. 1. In case of unpaired components, which occurs if the number of components in one level is an odd number, the unpaired component is directly forwarded to the next level. The unpaired components are alternatively selected from the left- and right-hand sides at successive levels.

With the proposed structure, as with the KLT, most of the image energy is accumulated in the first components since each of the two-component KLTs operates as a classic KLT, grouping most of the signal energy in one of the two resulting components—the principal component—hence allowing most of the energy to flow across the composition of the transforms up to the resulting components of the last levels. Moreover, the combination of multiple KLT instances also produces an orthogonal transform, which still preserves the image variance across domains, and allows a direct distortion estimation on the transformed domain.

Let us recall the definition of a classic N component KLT: let X be a matrix that has N rows, with one for each source, and M columns, with one for each sample. Then, Y, which is the outcome after applying the transform, is computed as

$$Y = \mathrm{KLT}_{\Sigma_X}(X) = Q^T X \tag{1}$$

where $\Sigma_X = (1/M)XX^T$ is the covariance matrix of X and Q is the orthogonal matrix obtained from the ED of $\Sigma_X = Q\Lambda Q^{-1}$ ($\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_N), |\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_N|$). Note that the sources must be manually centered around zero if they do not have a zero mean. The covariance matrix of Y is the diagonal matrix Λ, and $\lambda_1, \ldots, \lambda_N$ are the variances of each component after the transform.

With the use of only the two-component KLTs, the ED procedure is greatly simplified, overcoming one of the previously described KLT inconveniences, which is the implementation difficulty. The following is a numerically stable solution that does not require any iterative process to obtain Q:

$$Q = \begin{pmatrix} p & q \\ t & u \end{pmatrix} \quad \Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$$

where

$$\begin{aligned} t = -q &= \frac{b}{|b|}\sqrt{\frac{1}{2} - \frac{(a-d)}{2s}} \\ p = u &= \sqrt{\frac{1}{2} + \frac{(a-d)}{2s}} = \sqrt{1-t^2} \\ \lambda_1 &= \frac{a+d+s}{2} \quad \lambda_2 = \frac{a+d-s}{2} \\ s &= \sqrt{(a-d)^2 + 4b^2} \\ \Sigma_X &= \begin{pmatrix} a & b \\ b & d \end{pmatrix}. \end{aligned} \tag{2}$$

While computing Q, two precautions have to be taken into account. First, if $b = 0$, then t and q are undefined. However, assuming $b/|b|$ to be either 1 or -1 is a sufficient solution. Second, $s \simeq 0$ if inputs are similar and share almost no energy, which causes a division by zero. An appropriate remedy is to use an identity matrix as Q.

Likewise, the same reasoning can also be extended to the lifting decomposition of a two-component transform, which allows the transform to be applied in a lossless mode. The lifting decomposition and application of the KLT, otherwise known as reversible KLT (RKLT), are performed as described either in [22] or in [23]. The RKLT is based on decomposing a transform matrix into multiple elementary matrices, with each one being equivalent to a sequence of lifting steps, and then performing the matrix multiplication by using these lifting sequences. For a two-component transform, regardless of which of the two methods is used, the required sequence of the elementary matrices is of this form

$$Q^T = \begin{pmatrix} p & t \\ q & u \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & s \end{pmatrix} P \begin{pmatrix} 1 & 0 \\ w_3 & 1 \end{pmatrix} \begin{pmatrix} 1 & w_2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ w_1 & 1 \end{pmatrix}.$$

The following is a solution for the previous equation:

$$\begin{cases} P = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} & w_1 = w_3 = \frac{p-1}{t} \\ \quad w_2 = t, s = \mathrm{Det}(Q) & \text{for } |t| \geq |p| \\ P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} & w_1 = w_3 = \frac{1-t}{p} \\ \quad w_2 = -p, s = -\mathrm{Det}(Q) & \text{for } |t| < |p|. \end{cases}$$

As Q is orthogonal ($t^2 + p^2 = 1$), therefore, the condition $|t| \geq |p|$ is enough to guarantee that no division by zero occurs. Note that s only affects the sign of the nonprincipal result, which is not further transformed, and as changing the sign of an entire component is usually irrelevant to hyperspectral encoders, it can always be assumed to be positive. The associated lifting structure in performing $\binom{y_1}{y_2} = Q^T \binom{x_1}{x_2}$, as shown in Fig. 2, requires only nine operations and a conditional permutation.

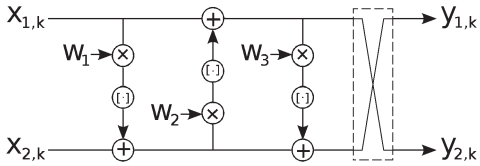

Fig. 2. Lifting structure for a two-component KLT. The conditional permutation is included.

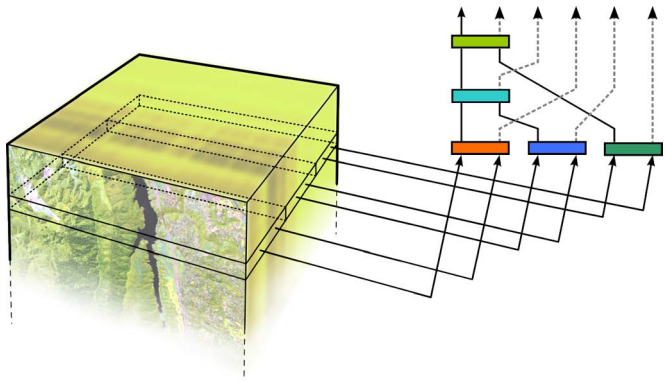

Fig. 3. Example of a line-based application of a POT on a hyperspectral image with six components.

B. Implementation Considerations

1) Side Information: As the KLT, the POT requires the use of side information—usually, the transform matrix and the offset of each component so that it has a zero mean—to indicate a specific inverse transform to the decoder. The amount of side information of the KLT is usually negligible if the image has a medium or large spatial size, yet for images with a small spatial size, the amount of side information of the KLT has an impact on the coding performance [21]. Contrarily, the POT requires a much smaller amount of side information, which is further reduced by transmitting only t as a half-precision IEEE 754 floating-point number, for each pairwise transform. Note that only t is required to be transmitted as p, q, and u can be directly derived from t (2). As an example, without any side information quantization or coding, for an Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) hyperspectral image of 224 components, the KLT would require at least 101 kB of side information, whereas just 1340 B would suffice for the POT.

2) Line-Based POT: With such a small amount of required side information, it is possible to apply the POT in a line-by-line basis, as shown in Fig. 3, where each line of an image is independently processed from the rest. As the already processed input lines can be discarded, a line-based transform does not require buffering the whole image, substantially reducing the memory requirements.

Moreover, most of the current hyperspectral sensors capture images over a fixed width window with a height of just one pixel and use the motion of the sensor itself to provide the remaining spatial dimension of the image, specifically, either in the pushbroom mode—capturing all the spectral bands of a line at once—or in the whiskbroom mode—capturing the information in a band interleaved by pixel (BIP) order [24]. Both modes are particular cases of line-by-line image acquisition and are well suited to the use of a POT in line-based processing, for instance, in on-board applications.

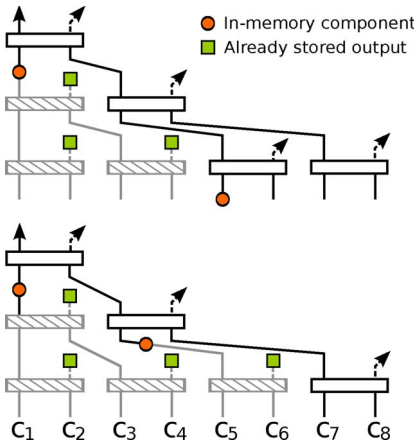

Fig. 4. Internal state of an eight-component transform when the components are read sequentially. The internal state (top) before and (bottom) after the sixth component of the image (c_6) is read.

A particular order within the pushbroom mode—band interleaved by line (BIL)—is the most suitable in performing a line-based POT because a coder, in addition to being able to discard the already processed lines, is also able to discard the already processed components of a line. When a POT is applied to a BIL-ordered image, where each component is read sequentially for each line, as soon as two consecutive components of a line are read, both can be transformed, and one of them can be flushed to the output, while the other follows to the higher levels of the structure, where it may also be transformed again. This process is shown in Fig. 4, where only a portion of the two spectral components of a line is required to remain in memory. Contrarily, for a BIP order, due to the transform training stage, a complete line has to be buffered before the transform can be trained and then applied.

3) Covariance Computation: Another aspect of the transform usage is the calculation of the covariance matrices Σ_X, which amounts to one-third of the forward transform cost and which can be optimized with knowledge of the eigenvalues of the previous levels of transform. For each pairwise transform, the covariance of its output is

$$\Sigma_Y = \frac{1}{M}\mathrm{Y}\mathrm{Y}^T = Q^T \Sigma_X Q = \Lambda.$$

Since, first level apart, the pairwise transforms receive its inputs from other transform outputs, the variance of an input component is already computed (in Λ from the other transform), and it corresponds to the eigenvalue λ_1 associated with that component. As such, values a and d from Σ_X for transforms beyond the first level are already available, while b is not available because the two inputs have not been jointly examined before. Deriving a and d from previous transforms reduces the total cost of the covariance matrix calculation by two-thirds for those transforms and one-third globally.

4) Range Expansion: When a POT is applied over an image, the numeric range of its values is increased because the energy is accumulated on a few components. Knowing the exact range expansion is required, for example, to allocate the correct register sizes.

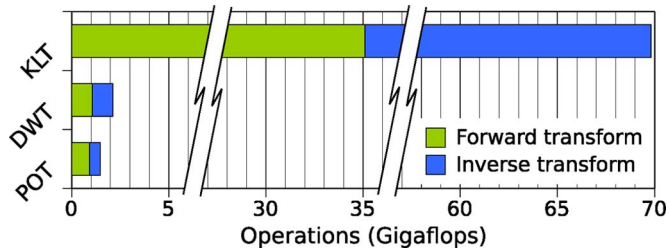

Fig. 5. Cost comparison to apply lossy forward and inverse transforms to a common AVIRIS image of size $677 \times 512 \times 224$ ($x \times y \times z$).

The range expansion for each lossy two-component transform can be calculated as follows. From (1) and (2)

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} \sqrt{1-t^2} & t \\ -t & \sqrt{1-t^2} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

and hence

$$\mathrm{peak}_{y_1} = \max_{-1\le t\le 1} |\sqrt{1-t^2}x_1 + t x_2|.$$

Combining the peak values of both inputs as $\mathrm{peak}_{\mathrm{in}} = \max\{|x_1|, |x_2|\}$ yields a peak output value bounded by

$$\mathrm{peak}_{\mathrm{out}} \le \mathrm{peak}_{\mathrm{in}} \cdot \max_{-1\le t\le 1} \sqrt{1-t^2} + |t| = \mathrm{peak}_{\mathrm{in}}\sqrt{2}.$$

On the lossless case, one also has to take into account the error introduced by the lifting structure (Fig. 2), which is bounded by [22, eq. (19)] to

$$\|\mathbf{u}\|_\infty \le \frac{2 + |w_3| + \max\{1 + |w_2|, |w_3| + |w_3 \cdot w_2 + 1|\}}{2}.$$

The peak values for $|w_2|$ and $|w_3|$ are, respectively, 1 and $1+\sqrt{2}$; hence

$$\|\mathbf{u}\|_\infty \le 3 + \frac{3\sqrt{2}}{2} < 6.$$

As the number of levels of an N-component POT is not greater than $l = \lceil \log_2 N \rceil$, the range expansion of the whole transform is bounded by

$$\mathrm{peak}_{\mathrm{out}} < (\mathrm{peak}_{\mathrm{in}} + 6l)\sqrt{2}^l$$

which corresponds to an increase of 5 b in the register size for a 256-component transform applied to the 16-b data.

C. Analytical Evaluation

Having considered several implementation issues, this section reports an analysis of the proposed POT with respect to the following three inconveniences of the classical KLT: computational cost, memory requirements, and component scalability. It will be shown that none of these three inconveniences remains in a POT.

1) Computational Cost: The high computational cost of the KLT renders it unpractical and even unusable on many situations. As an example, the costs in applying lossy forward and inverse transforms on a typical AVIRIS image for KLT, POT, and wavelets are shown in Fig. 5. In this particular case, the

TABLE I
TRANSFORM APPLICATION AND REMOVAL COST (IN FLOPS). ρ IS THE COVARIANCE SUBSAMPLING FACTOR [4], n IS THE NUMBER OF SPECTRAL COMPONENTS, AND $m = xy$ IS THE NUMBER OF SPATIAL LOCATIONS OF THE ORIGINAL IMAGE. DOMINANT TERMS ARE DETERMINED, ASSUMING $m \ggg n$ AND $\rho \simeq 0.01$. FOR THE WAVELET TRANSFORMS, l IS THE NUMBER OF LEVELS, AND THE COST IS EXACT IF $n = 2^l k, k \in \mathbb{N}$

	Computational cost		Dominant terms	
Spectral transform	Forward	Inverse	Fwd.	Inv.
Full KLT lossy	$9n^3 + m((2+\rho)n^2 + (\rho+1)n + 4\rho)$	$2mn^2$	$2n^2m$	$2n^2m$
Full KLT lossless	$m((3+\rho)n^2 + (2+\rho)n + 4\rho - 3) + \frac{32}{3}n^3 + \frac{1}{2}n^2 - \frac{37}{6}n + 5$	$m(3n^2 + n - 3)$	$3n^2m$	$3n^2m$
POT lossy	$12mn + 21ny - 8m - 23y$	$7mn + 3ny - 6m - 3y$	$12nm$	$7nm$
POT lossless	$16mn + 26ny - 12m - 28y$	$11mn + 5ny - 10m - 5y$	$16nm$	$11nm$
Wavelet CDF 9/7 lossy	$\simeq m(14n\sum_{i=1}^{l} 2^{-i})$	$\simeq m(14n\sum_{i=1}^{l} 2^{-i})$	$14nm$	$14nm$
Wavelet CDF 5/3 lossless	$\simeq m(-3l + 9n\sum_{i=1}^{l} 2^{-i})$	$\simeq m(-3l + 9n\sum_{i=1}^{l} 2^{-i})$	$9nm$	$9nm$

TABLE II
APPROXIMATE PEAK MEMORY REQUIREMENTS MEASURED IN NUMBER OF COEFFICIENTS. IN PARENTHESES ARE THE BYTES NEEDED WHEN $n = 224$, $y = 512$, $x = 677$, $l = 5$, AND COEFFICIENT BIT DEPTH IS 16 bpppb

Spectral transform	BIL	BIP
Full KLT	$n^2 + nm + n$ (148 MB)	$n^2 + nm + n$ (148 MB)
POT	$(1 + \log_2(n))x + 6$ (12 kB)	$xn + 6$ (296 kB)
Wavelet CDF 9/7	$6lx$ (40 kB)	$6l$ (60 bytes)
Wavelet CDF 5/3	$4lx$ (26 kB)	$4l$ (40 bytes)

TABLE III
TRANSFORM SCALABILITY (IN COMPONENTS REQUIRED TO RECOVER ONE COMPONENT). n IS THE NUMBER OF SPECTRAL COMPONENTS. l IS THE NUMBER OF WAVELET LEVELS. THE REPORTED WAVELET SCALABILITY MAY BE REDUCED ON THE TRANSFORM EDGES DUE TO COEFFICIENT MIRRORING (UP TO HALF)

Spectral transform	Dependent components	avg. (min. / max.) ($n = 224, l = 5$)
Full KLT lossy	n	224.0 (224/224)
Full KLT lossless	n	224.0 (224/224)
POT lossy	$\simeq 1 + \log_2(n)$	8.9 (8/9)
POT lossless	$\simeq 1 + \log_2(n)$	8.9 (8/9)
Wavelet CDF 9/7	$\simeq 7l - 2^{-l} - 1$	36.0 (32/38)
Wavelet CDF 5/3	$\simeq 3l - 2^{-l} - 1$	16.0 (11/17)

KLT cost is over 30 times higher than that of POT or wavelets, and the POT cost is slightly lower than that of wavelets.

In Table I, a detailed analysis of the transform costs is provided. Dominant terms are highlighted in the last two columns.

The main difference among the dominant terms is the squared factor for the KLT, which substantially increases the transform cost as the number of components increases. The POT and wavelets transforms have an approximately linear cost in relation to the components or spatial locations, and they mainly differ on a constant factor. In both forward and inverse modes, a lossy POT has a lower constant factor than a lossy CDF 9/7, while the opposite happens in the lossless case. Nonetheless, both POT and wavelets entail a very affordable cost.

2) Memory Requirements: A transform with a low memory requirement can be used in memory constrained environments, and it often involves nonproportional memory requirements, allowing the coding of arbitrarily large images regardless of the available memory capability. The KLT has a huge memory requirement, which is proportional to the whole image size and is usually well over a hundred megabytes. Contrarily, since the POT is able to operate in a line-based mode, the POT requires a tiny amount of memory, which is often around a few kilobytes. The peak memory requirements for the KLT, POT, and wavelets are reported in Table II, both when the image is fed to the transform in the BIL order and when it is fed in the BIP order. Although the reported results for the POT and wavelets are both of acceptable magnitude, it is of interest to note the small amount of memory needed by the wavelet transforms when the data are in the BIP mode, which is a direct consequence of the wavelets not having a training stage, enabling them to start producing results as soon as all bands of one pixel are available.

3) Component Scalability: In the context of multicomponent image coding, a spectral transform provides component scalability if it allows to partially reverse the transform without the need to invert the full transform. A transform with component scalability allows features such as region of interest retrieval, progressive decoding, and interactive transmission. In this case too, the KLT has a serious impairment since, even in recovering only one component, the full transform has to be reversed. Table III reports the component scalability of the three compared spectral transforms. POT is the transform with better scalability, followed by wavelets, which have a poorer component scalability, and, at a much larger distance, by KLT, which does not provide any scalability. To successfully take advantage of the transform scalability, it is important to pair the transform with a scalable coder, such as JPEG2000, which, apart from the spectral transform and the rate-distortion (R-D) optimization stage, does not have any additional band interdependency.

III. EXPERIMENTAL RESULTS

To evaluate the suitability of POT for hyperspectral image coding, extensive experimental tests have been performed, including classification-based experiments. In this section, the results are reported and discussed.

Experiments have been conducted on a corpus composed of several images from AVIRIS [25] and Hyperion [26] sensors. The images obtained with an AVIRIS sensor have 224 spectral components covering wavelengths from 370 to 2500 nm. For the experiments with AVIRIS images, scenes with only the first 512 lines of each image have been used. The AVIRIS corpus

TABLE IV
TECHNICAL NAMES FOR THE AVIRIS AND HYPERION IMAGES USED

Image Name	Technical Name	Size ($x \times y \times z$)[1]
AVIRIS		
Cuprite	f970619t01p02_r02	$614 \times 2206 \times 224$
Jasper Ridge	f970403t01p02_r03	$614 \times 2586 \times 224$
Low Altitude	f960705t01p02_r05	$614 \times 3689 \times 224$
Lunar Lake	f970623t01p02_r07	$614 \times 1431 \times 224$
Moffett Field	f970620t01p02_r03	$614 \times 2031 \times 224$
Yellowstone (Sc K)	f060925t01p00_r12_sc$K$	$677 \times 512 \times 224$
Hawaii	f011020t01p03r05_sc01	$614 \times 512 \times 224$
Maine	f030828t01p00r05_sc10	$680 \times 512 \times 224$
Hyperion		
Coastal	EO1H0140342002050110PY	$256 \times 2905 \times 224$
Erta Ale	EO1H1680502010057110KF	$256 \times 3188 \times 242$
Lake Monona	EO1H0240302009166110PF	$256 \times 3177 \times 242$
Mt. St. Helens	EO1H0460282009231110KF	$256 \times 3243 \times 242$
Urban	EO1H0440342002212110PY	$256 \times 2905 \times 224$

[1]Uncalibrated Yellowstone images are 680 columns wide, and uncalibrated Hyperion images have one line less that the calibrated ones.

includes images from various stages of the acquisition chain: uncalibrated, radiance, and reflectance images. The Hyperion images have a fixed width of 256 columns, a variable height, and 242 spectral components covering wavelengths from 357 to 2576 nm, from which only 198 components are calibrated. The Hyperion corpus includes the uncalibrated and radiance images. The Hyperion sensor has a bitdepth of 12 bits per pixel per band (bpppb), while the AVIRIS sensor has evolved over the years, and some of its features have changed. In particular, the bitdepth initially was 12 bpppb, and currently, it is 16 bpppb. After calibration, the image bitdepth is 16 bpppb. Table IV provides the full technical names and sizes for the images used. The described images can be obtained from [25], [27], and [28].

A spectral transform is usually the first stage of the coding process. Hence, to assess the compression gains, a spectral transform has to be tested in conjunction with an image coder. Three coders have been selected for this purpose: JPEG2000 [29], 3d-TCE [1], and TER [30], which is an improved version of the Consultative Committee for Space Data Systems (CCSDS) recommendation for image data compression [31]. JPEG2000 is a standardized multipurpose image coding system, which provides a good performance and is well suited for long-term archival. 3d-TCE is an embedded tarp-based coder, with a performance similar to JPEG2000, but that differs from JPEG2000 because it uses a different approach to provide the probability estimates in the transformed domain. TER, which is highly focused on satellite image compression, enhances the CCSDS recommendation with volumetric R-D, quality scalability, and other features. All three coders are able to perform both lossy and PLL encodings of volumetric images, although in this case, only 3d-TCE tries to exploit 3-D redundancy by using a volumetric probability context in addition to the spectral transform.

The implementations employed to perform the experiments are the following: Kakadu for JPEG2000 [32], QccPack for 3d-TCE [33], TER for CCSDS-IDC [34], and our own open-source implementation of the KLT and the POT [35]. The coding systems are alternated to show that the transform performance results are consistent regardless of the coding method used.

(a) (b) (c)

Fig. 6. Pushbroom sensor artifacts on a region of the Hyperion image of Lake Monona (a) before calibration, (b) with a simple preprocessing, and (c) after calibration. (a) Uncalibrated. (b) Preprocessed. (c) Calibrated.

From the computational cost point of view, it has to be taken into account that a transform is one stage of the whole coding process of a coding system and that other stages may present higher computational costs and memory requirements, depending on the combination of transform and coder selected. In particular, to achieve a low memory cost, the coder implementation has to code in parallel and in a line-based mode all the bands of an image (similar to what the Kakadu software does). Implementation difficulty is increased by coding the image in the line-based mode, mainly because multiple transforms are applied at once, and the state of each one has to be maintained. Note also that image coders that only perform pure lossless codings of the hyperspectral images usually have lower costs than a full transform coder with PLL capabilities.

In both lossy and lossless modes, the proposed POT spectral transform is evaluated in comparison to the KLT and to a discrete wavelet transform (DWT; Cohen–Daubechies-Feauveau (CDF) 9/7 for lossy coding and CDF 5/3 for PLL, with five levels). No spatial clustering is applied to the KLT, and the side information for all transforms is coded with zlib [36].

There are several variations of a 3-D DWT (e.g., a pyramid wavelet decomposition or a hybrid 1D+2D wavelet decomposition such as the one used) and multiple DWT-based coders. The use of a hybrid 1D+2D wavelet in combination with a transform coder like JPEG2000 is the recommended approach for this scenario in recent literature [4], [37]–[39], as being the one that provides the best performance.

The rest of this section is organized as follows. First, an issue with the pushbroom sensors is addressed. Next, the results for lossy coding are discussed. Then, the PLL results are reported, followed by pure lossless experiments. Finally, some classification-based results are disclosed.

A. Pushbroom Sensors

Due to the nature of the pushbroom sensors (i.e., those sensors that capture all the spectral bands of a line at once by a bidimensional sensor array), vertical streaking artifacts might appear, as occurs for the Hyperion sensor. These artifacts, shown in Fig. 6, are caused by the use of an array of sensors and the slight variations among each of them. These artifacts are mostly removed by the calibration process [2], but on the uncalibrated images, the streaking artifacts have a severe impact on coding performance because they hinder the detection of redundancy in adjacent locations.

To limit the impact of these artifacts, a very simple preprocessing stage can be applied to the uncalibrated images from the

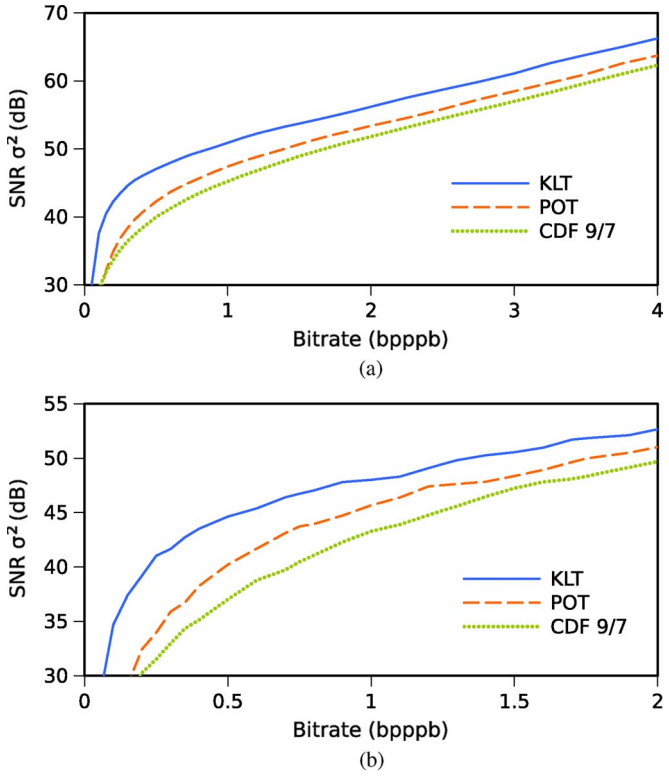

Fig. 7. R-D performance comparison among the three spectral transforms. (a) AVIRIS of Moffett Field radiance coded with JPEG2000. (b) AVIRIS of Low Altitude radiance coded with TER.

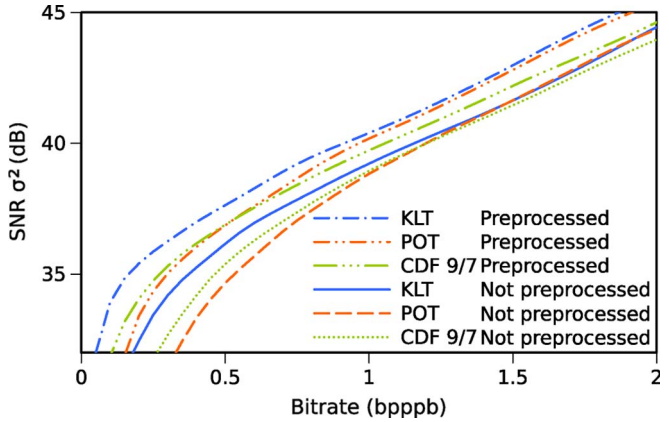

Fig. 8. Effects of the pushbroom preprocessing. The image is the Hyperion of Lake Monona. The coding system is JPEG2000.

pushbroom sensor Hyperion. The preprocessing is a reversible process based on shifting the values of the image columns to smooth the differences with their neighbors and is undone after decompression just as if it was an additional transform. First, the shifting offsets are obtained by

$$\delta_{i,k} = \frac{1}{32}\sum_{j=0}^{31} \underset{|m|\leq 4}{\text{median}} \left(I_{w(i+m,0,x-1),j.k}\right)$$

where $I_{i,j,k}$ is the value of the original image at band k, row j, and column i, and x is the image width.

Only the first 32 lines are used to obtain the shifting offsets. Afterwards, each image location, as it is retrieved, is modified

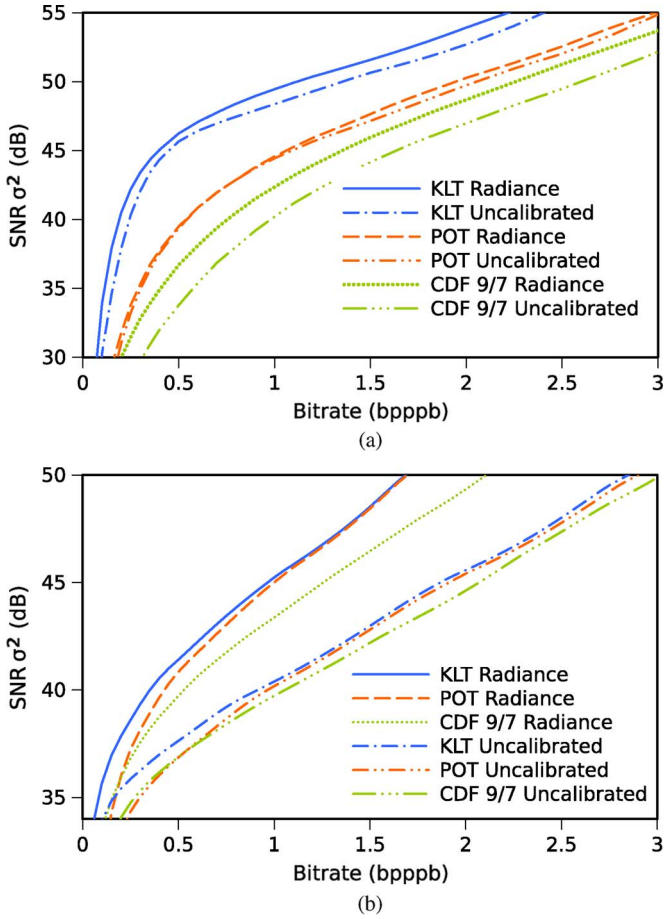

Fig. 9. Comparison of the transform performance between the uncalibrated and radiance images. The coding system is JPEG2000. (a) AVIRIS of Yellowstone. (b) Hyperion of Lake Monona.

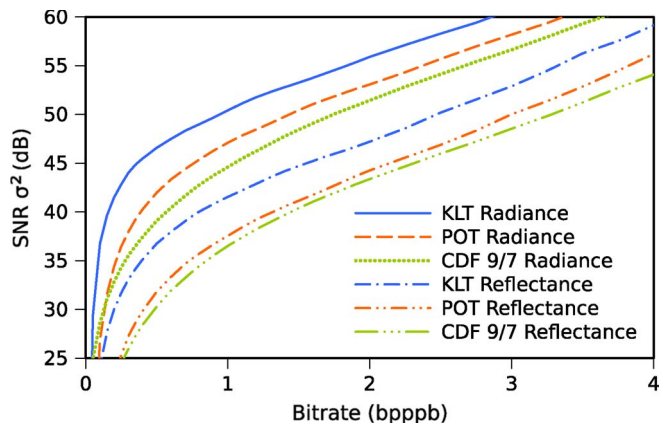

Fig. 10. Comparison of the transform performance between the radiance and reflectance images. The image is the AVIRIS of Jasper Ridge. The coding system is JPEG2000.

according to the previous offset by

$$\hat{I}_{i,j,k} = w\left(I_{i,j,k} - [\delta_{i,k}], -2^{15}, 2^{15}-1\right).$$

The function w warps any integer value to the interval $[a, b]$, with a symmetric reflection around a if $i < a$ or around b if $i > b$, i.e.,

$$w(i, a, b) = b - |b - a - (i - a \bmod 2(b - a))|.$$

TABLE V
SNR DIFFERENCE (IN DECIBELS) BETWEEN THE KLT, POT, AND WAVELET CDF 9/7. THE KLT IS USED AS REFERENCE, AND THE BITRATE IS SAMPLED AT $b = 0.1k$, WITH $k = 2, 3, \ldots, 20$, AND IS MEASURED IN BITS PER PIXEL PER BAND. THE IMAGES ARE CODED WITH JPEG2000

	POT			Wavelet CDF 9/7		
	avg.	min.	max.	avg.	min.	max.
AVIRIS Cuprite Radiance	−0.81	(−1.32	/ −0.68)	−3.32	(−5.24	/ −2.54)
AVIRIS Jasper Ridge Radiance	−3.41	(−6.02	/ −2.53)	−5.96	(−8.70	/ −4.46)
AVIRIS Low Altitude Radiance	−2.90	(−5.95	/ −1.80)	−6.14	(−9.45	/ −3.92)
AVIRIS Lunar Lake Radiance	−0.77	(−1.39	/ −0.64)	−2.75	(−4.41	/ −2.08)
AVIRIS Moffett Field Radiance	−3.56	(−6.28	/ −2.62)	−5.82	(−8.62	/ −4.38)
AVIRIS Yellowstone (Sc 0) Radiance	−4.74	(−7.18	/ −3.43)	−7.26	(−10.62	/ −5.21)
AVIRIS Yellowstone (Sc 11) Radiance	−4.02	(−6.86	/ −2.96)	−5.50	(−8.94	/ −3.80)
AVIRIS Yellowstone (Sc 18) Radiance	−6.03	(−8.73	/ −4.21)	−7.95	(−11.24	/ −5.54)
Hyperion Coastal Radiance	−0.47	(−1.69	/ −0.18)	−2.85	(−3.47	/ −2.46)
Hyperion Erta Ale Radiance	−0.19	(−1.01	/ 0.00)	−2.36	(−2.98	/ −1.96)
Hyperion Lake Monona Radiance	−0.38	(−1.89	/ −0.03)	−1.98	(−2.60	/ −1.61)
Hyperion Mt. St. Helens Radiance	−0.51	(−1.79	/ −0.20)	−3.12	(−3.63	/ −2.77)
Hyperion Urban Radiance	−0.87	(−2.53	/ −0.44)	−3.88	(−4.87	/ −3.45)
AVIRIS Cuprite Reflectance	−2.17	(−4.71	/ −1.25)	−7.89	(−12.46	/ −4.91)
AVIRIS Jasper Ridge Reflectance	−3.80	(−5.63	/ −2.79)	−5.18	(−7.40	/ −3.84)
AVIRIS Lunar Lake Reflectance	−1.54	(−3.72	/ −0.92)	−4.91	(−9.12	/ −3.22)
AVIRIS Moffett Field Reflectance	−4.03	(−5.66	/ −2.90)	−7.03	(−10.70	/ −4.72)
AVIRIS Hawaii Uncalibrated	−1.88	(−3.30	/ −1.38)	−4.14	(−8.83	/ −3.03)
AVIRIS Maine Uncalibrated	−2.15	(−5.27	/ −1.40)	−4.83	(−10.99	/ −2.97)
AVIRIS Yellowstone (Sc 0) Uncalibrated	−3.93	(−5.90	/ −2.72)	−8.40	(−12.36	/ −5.72)
AVIRIS Yellowstone (Sc 11) Uncalibrated	−1.03	(−1.30	/ −0.64)	−4.10	(−5.86	/ −2.69)
AVIRIS Yellowstone (Sc 18) Uncalibrated	−3.48	(−4.41	/ −2.43)	−7.39	(−9.70	/ −4.95)
Hyperion Erta Ale Uncalibrated	−0.23	(−1.31	/ +0.02)	−0.96	(−1.40	/ −0.77)
Hyperion Lake Monona Uncalibrated	−0.47	(−2.00	/ −0.15)	−0.83	(−1.35	/ −0.63)
Hyperion Mt. St. Helens Uncalibrated	−0.72	(−2.22	/ −0.31)	−1.68	(−2.86	/ −1.39)

B. Lossy Coding

Fig. 7 shows the R-D evolution of the different spectral transforms, namely, KLT, POT, and wavelet CDF 9/7. The plots provide a relation between quality, measured in signal-to-noise ratio (SNR), and bitrate, measured in bits per pixel per band. In this case, the SNR is defined as $\text{SNR} = 10\log_{10}(\sigma^2/\text{MSE})$, where σ^2 is the variance of the original image. A steady quality evolution can be observed for all transforms in both examples, with KLT obtaining the best performance, followed by POT, and with wavelet CDF 9/7 producing the worst coding results.

The coding performance for the uncalibrated images from the Hyperion sensor, shown in Fig. 8, varies depending on whether a preprocessing is applied to remove the streaking artifacts or not. The transform that is most affected by the streaking artifacts seems to be the POT, but all three transforms are heavily affected by the artifacts. From now on, the preprocessing is included in all codings of the uncalibrated Hyperion images, and the operation is reversed on decodings.

Fig. 9 shows the comparisons of the coding performance between the radiance and uncalibrated images, and Fig. 10 shows the comparisons of the coding performance between the radiance and reflectance images when JPEG2000 is used as the coding system. From the first comparison, it seems that the wavelets suffer a noticeable penalization on the uncalibrated images from the AVIRIS sensor and on the calibrated images from the Hyperion sensor.

Table V shows the detailed comparison for the lossy coding between the POT and wavelet CDF 9/7 in relation to the coding results of the KLT. The results show a moderate performance penalty for the POT and a higher performance penalty for the wavelet. In the reported experiments, the performance decrease of the wavelet CDF 9/7 spectral transform in the average case is always higher than the performance decrease of the POT, with the POT yielding results between 1 and 5 dB higher, except for the uncalibrated Hyperion images where gains are more modest perhaps because of the smaller difference between the KLT and DWT.

An important remark is that, for the lossy coding of the hyperspectral images, the most appropriate distortion measure would often be how well the posterior processing and exploitation tasks perform on imagery lossy coded, particularly for the uncalibrated images where a calibration process is always applied. This is not a trivial assessment at all as many different processing and exploitation tasks can be applied, and each often requires particular data sets (e.g., ground truth for a supervised land cover classifier or models on sensor response for the echo cancellation in the calibration of the Hyperion images). Refer to Section III-E for the two examples of these particular evaluations. If the posterior processing or exploitation task is known, small corrections over a lossy image might be coded as a side information to improve performance, as in [40], or a more general approach can be followed by combining a lossy transform coder with a quantized coding of the residual image in a near-lossless approach [41], [42].

As shown in Fig. 11, at very low bitrates—less than 0.5 bpppb—some artifacts appear on the POT because of its line-based application. The artifacts are barely perceptible unless a strong sharpening filter is applied, and they disappear as the bitrate increases. The low affectation by blocking the artifacts of the POT might be explained by the small size of each transform partition (i.e., one line), which enables a smooth

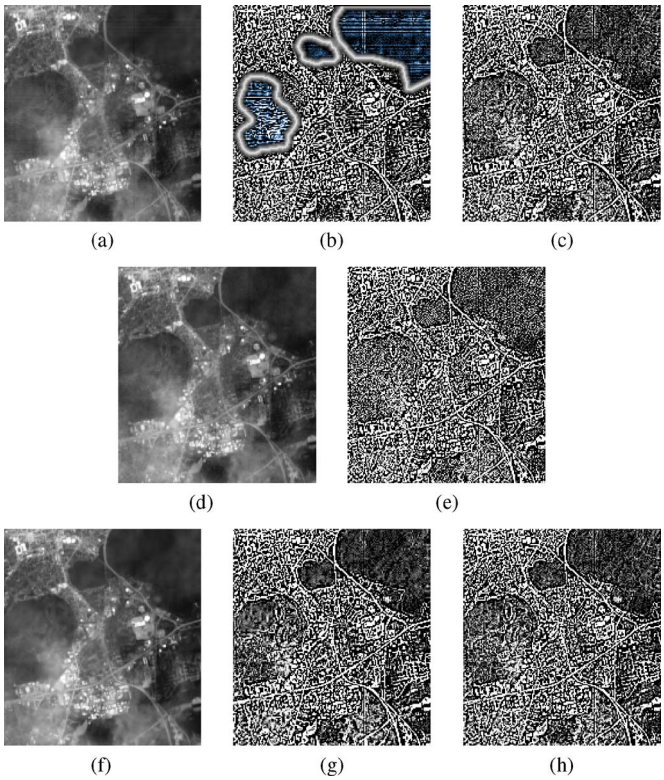

Fig. 11. Blocking artifacts at very low bitrates. Component 20 of the Hyperion image of Lake Monona coded with JPEG2000. Images in (b), (c), (e), (g), and (h) had a sharpening filter applied to make the artifacts more apparent. (a) POT 0.15 bpppb. (b) POT 0.15 bpppb. (c) POT 0.5 bpppb. (d) Original. (e) Original. (f) DWT 0.15 bpppb. (g) DWT 0.15 bpppb. (h) DWT 0.5 bpppb.

transition between them. If it was of interest to reduce these artifacts, the variable t of each two-component transform and the zero-mean offsets could be smoothed by a regularizing filter with the values of adjacent lines.

C. PLL Coding

The results are also similar when the PLL coding performance is evaluated, as shown in Fig. 12. Note that, at very high bitrates, a lossless POT produces better results than a lossless RKLT due to the smaller number of lifting steps of the former, which causes smaller quantization errors before the lossless regime is achieved. For the uncalibrated Hyperion images, the RKLT performance drop is more noticeable, and it starts at lower bitrates (even if images are not preprocessed for the pushbroom artifacts). The results for TCE yield similar performance gains.

D. Pure Lossless Coding

Table VI reports the bitrates required for lossless coding and compares them with the bitrates needed by the pure lossless methods. As what happened for PLL, the same relative performance is maintained among the three spectral transforms (KLT, POT, and wavelets) except for the Hyperion radiance images, where the line-by-line adaption of the POT may be outweighing its less general decorrelation ability, and it renders a POT performance that is superior to KLT. On the Hyperion uncalibrated

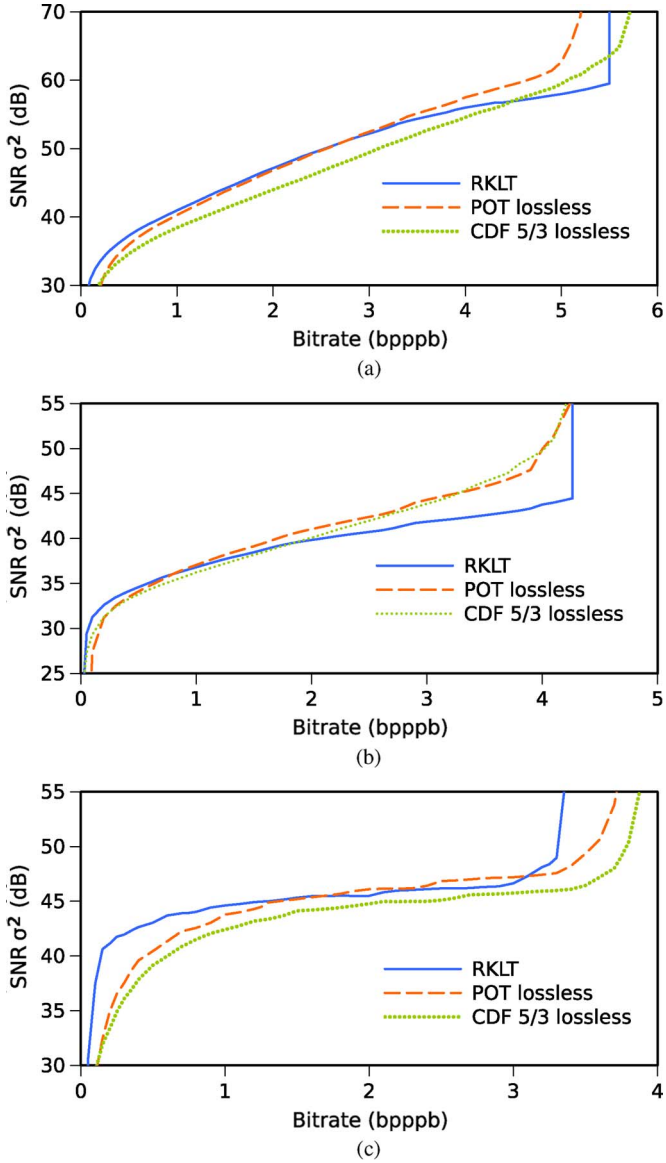

Fig. 12. R-D performance comparison among the three spectral transforms for PLL. (a) Hyperion Coastal radiance coded with JPEG2000. (b) Hyperion of Erta Ale uncalibrated coded with JPEG2000. (c) AVIRIS of Yellowstone radiance coded with TCE.

images, while, on average, the relative order is also maintained, the difference between the three transforms is not significant.

With regard to an analysis among the radiance, reflectance, and uncalibrated images, it seems that the radiance images are the ones where compression is more effective, producing smaller file sizes. When pure lossless methods are taken into account in the comparison, an interesting remark is that a clear division exists between the old 12-bpppb AVIRIS images from 1996/1997 and the newer 16-bpppb Yellowstone scenes from 2006. For the radiance images, on the older images, two pure lossless methods (LAIS-QLUT (LQLUT) [43] and TSP-W2 [44]) exploit the artifacts introduced in the calibration process [44], yielding a better performances, whereas the combination of RKLT and JPEG2000 generally outperforms the pure lossless methods on the newer Yellowstone scenes. On these new AVIRIS radiance images, the POT has a severe penalization,

TABLE VI
LOSSLESS BITRATE FOR KLT, POT, AND CDF 5/3 WITH JPEG2000 (IN BITS PER PIXEL PER BAND). FIVE PURE LOSSLESS METHODS ARE ALSO INCLUDED AS REFERENCE: FAST-LOSSLESS (FL) [45], LAIS-QLUT (LQLUT) [43], TSP-W2 [44], KSP [46], AND A1 FROM [47]. FOR THESE RESULTS, THE AVIRIS RADIANCE IMAGES HAVE BEEN USED IN ITS FULL HEIGHT

	RKLT	POT	CDF 5/3	FL	LQLUT	TSP-W2	KSP	A1
AVIRIS Cuprite Radiance	4.91	5.03	5.36	4.91	4.48	3.77	4.88	5.50
AVIRIS Jasper Ridge Radiance	4.86	5.32	5.60	4.95	4.70	4.08	4.96	5.60
AVIRIS Low Altitude Radiance	5.16	5.51	5.85	5.26	5.00	4.31	—	—
AVIRIS Lunar Lake Radiance	4.91	5.01	5.29	4.91	4.53	3.81	4.90	5.51
AVIRIS Moffett Field Radiance	4.91	5.29	5.60	4.99	4.79	4.12	4.92	5.64
AVIRIS Yellowstone (Sc 0) Radiance	3.86	4.44	4.73	3.96	4.48	3.99	—	4.81
AVIRIS Yellowstone (Sc 11) Radiance	3.62	4.09	4.24	3.63	4.02	3.67	—	4.41
AVIRIS Yellowstone (Sc 18) Radiance	3.78	4.50	4.69	3.94	4.48	3.97	—	4.77
Hyperion Coastal Radiance	5.50	5.18	5.79	—	—	—	—	—
Hyperion Erta Ale Radiance	5.66	5.30	5.89	—	—	—	—	—
Hyperion Lake Monona Radiance	5.81	5.49	6.02	—	—	—	—	—
Hyperion Mt. St. Helens Radiance	5.67	5.38	6.02	—	—	—	—	—
Hyperion Urban Radiance Radiance	5.76	5.49	6.19	—	—	—	—	—
AVIRIS Cuprite Reflectance	6.01	6.14	6.75	—	—	—	—	—
AVIRIS Jasper Ridge Reflectance	5.96	6.32	6.61	—	—	—	—	—
AVIRIS Lunar Lake Reflectance	6.04	6.10	6.60	—	—	—	—	—
AVIRIS Moffett Field Reflectance	6.10	6.47	6.86	—	—	—	—	—
AVIRIS Hawaii Uncalibrated	2.85	2.98	3.26	2.64	3.05	2.62	2.84	3.49
AVIRIS Maine Uncalibrated	2.97	3.07	3.41	2.72	3.19	2.74	2.90	3.65
AVIRIS Yellowstone (Sc 0) Uncalibrated	6.07	6.63	7.13	6.23	6.78	6.27	6.34	6.92
AVIRIS Yellowstone (Sc 11) Uncalibrated	6.10	6.27	6.60	5.86	6.30	5.88	—	6.53
AVIRIS Yellowstone (Sc 18) Uncalibrated	6.24	6.72	7.12	6.32	6.82	6.32	—	6.92
Hyperion Erta Ale Uncalibrated	4.26	4.30	4.34	—	—	—	—	—
Hyperion Lake Monona Uncalibrated	4.37	4.45	4.44	—	—	—	—	—
Hyperion Mt. St. Helens Uncalibrated	4.27	4.37	4.44	—	—	—	—	—

yet it provides a better performance than wavelet CDF 5/3. For the uncalibrated images, the pure lossless method FL [45] yields similar results as a combination of RKLT and JPEG2000, where the latter provides the PLL capabilities at a higher computational cost.

E. Classification-Based Experiments

An additional set of experiments measures the POT performance using classification-based metrics, i.e., it measures how classifiers are affected by the distortion introduced by a coding system. The classifiers used are k-means for clustering (with 20 classes and Eulerian distance) [48] and Reed Xiaoli (RX) for anomaly detection (with a 2% threshold) [49]. The RX detector identifies the locations that are different from the background by using the Mahalanobis distance to find the outliers (see [50] for a review on its use on hyperspectral images). Both methods are unsupervised classifiers that are widely used on remote-sensing imagery. The results are reported in preservation of classification units (POC), which is a percentage of how many spatial locations stay in the same class if coding distortion is added to the process.

In Fig. 13, the results for the k-means [48] on the AVIRIS of Moffett Field are presented. In this case, the wavelet has a slightly worse performance than POT, and the results saturate qualitywise between 2 and 3 bpppb.

Fig. 14 shows the results for an RX detector [49] on the AVIRIS of Lunar Lake. In this case, the results are more unstable, particularly for the KLT, which seems to introduce a distortion that is more detrimental to the classifier, perhaps because the Mahalanobis distance is also calculated in a KLT-transformed version of the image (i.e., another KLT not related to the one used for coding). Although the POT composes several KLT, it does not produce such artifacts. Apart from the RX instability with a KLT, the results are consistent with the previous ones, showing that the POT provides better results than the wavelets.

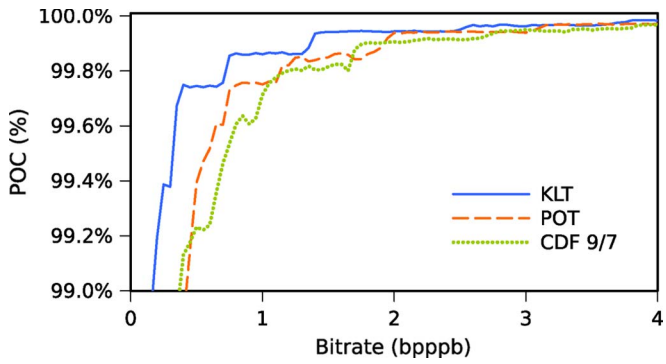

Fig. 13. Performance of the k-means clustering for the AVIRIS of Moffett Field. The coding system is JPEG2000.

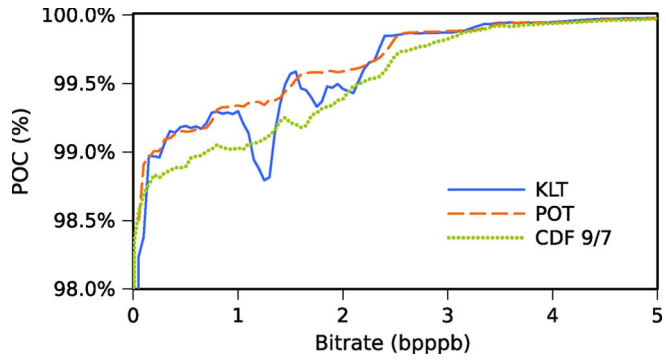

Fig. 14. Performance of the RX anomaly detection for the AVIRIS of Lunar Lake. The coding system is JPEG2000.

IV. Conclusion

In this paper, a new spectral transform has been proposed for the removal of spectral redundancy in hyperspectral image coders. The KLT and wavelets are the most common spectral transforms, but the KLT has four inconveniences that impair its practical use, and the wavelets provide a lower coding performance.

The proposed POT is based on the composition of multiple two-component instances of the KLT. The use of such composition allows substantial benefits in computational cost, component scalability, memory requirements, and implementation difficulty.

A detailed analysis of the proposed transform has been presented, comparing the POT with a full KLT and with the wavelets and describing the different properties of the transform. The analysis indicates that the memory requirements and implementation difficulty are the issues where our contribution is more substantial. Due to the small amount of side information needed by the POT, a line-based approach is amenable, which requires four orders of magnitude less of memory than a full KLT. In addition, due to the use of the two-component KLT, the diagonalizing of the covariance matrix can be greatly simplified from a complex iterative process to a direct calculation. Moreover, the proposed transform also has a computational cost, with a linear complexity in relation to the number of spectral components, as opposed to the KLT which has a quadratic complexity (i.e., the POT cost is forty-five times smaller than the KLT for the AVIRIS images). Finally, the POT also provides a good component scalability, which is better than that provided by the wavelets and much better than the KLT.

Extensive experimental results are conducted for the AVIRIS and Hyperion images and for three different image coding systems (JPEG2000, TCE, and TER, which is an improved version of the CCSDS IDC recommendation). The results consistently report that the proposed POT spectral transform has a better performance than the wavelets, although it does not reach the coding performance of a full KLT. Results are reported for the radiance, reflectance, and uncalibrated images; for the lossy, PLL, and pure lossless coding; and for the classification-based metrics. Generally, for all of the images tested and on all experimental results, the proposed transform outperforms the wavelets.

In conclusion, the proposed POT is a transform that might be a good replacement for the wavelets as the spectral decorrelator of the hyperspectral images in all of the scenarios where a KLT is not a feasible option or simply where the extra coding performance of the KLT is not worth its associated burdens.

Acknowledgment

The authors would like to thank the anonymous reviewers for their valuable comments and the National Aeronautics and Space Administration, the U.S. Geological Survey, and, particularly, A. Kiely for providing the imagery used. The computational resources used in this paper were partially provided by the Oliba Project of the Universitat Autònoma de Barcelona.

References

[1] J. Zhang, J. E. Fowler, and G. Liu, "Lossy-to-lossless compression of hyperspectral imagery using three-dimensional TCE and an integer KLT," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 4, pp. 814–818, Oct. 2008.

[2] J. Pearlman, P. Barry, C. Segal, J. Shepanski, D. Beiso, and S. Carman, "Hyperion, a space-based imaging spectrometer," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 6, pp. 1160–1173, Jun. 2003.

[3] N. Ahmed, T. Natarajan, and K. Rao, "Discrete cosine transform," *IEEE Trans. Comput.*, vol. C-23, no. 1, pp. 90–93, Jan. 1974.

[4] B. Penna, T. Tillo, E. Magli, and G. Olmo, "Transform coding techniques for lossy hyperspectral data compression," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 5, pp. 1408–1421, May 2007.

[5] L. Lan and I. Reed, "Fast approximate Karhunen–Loêve transform (AKLT) with applications to digital image-coding," in *Proc. SPIE—Visual Communications and Image Processing '93*, 1993, vol. 2094, pp. 444–455.

[6] A. Pirooz and I. Reed, "A new approximate Karhunen–Loêve transform for data compression," in *Conf. Rec. 32nd Asilomar Conf. Signals, Syst., Comput.*, 1998, vol. 1/2, pp. 1471–1475.

[7] I. R. Greenshields and J. A. Rosiene, "A fast wavelet-based Karhunen–Loêve transform," *Pattern Recognit.*, vol. 31, no. 7, pp. 839–845, Jul. 1998.

[8] M. Cagnazzo, L. Cicala, G. Poggi, and L. Verdoliva, "Low-complexity compression of multispectral images based on classified transform coding," *Signal Process.—Image Commun.*, vol. 21, no. 10, pp. 850–861, Nov. 2006.

[9] C. Thiebaut, E. Christophe, D. Lebedeff, and C. Latry, "CNES studies of on-board compression for multispectral and hyperspectral images," in *Proc. SPIE—Satellite Data Compression, Communications, and Archiving III*, R. W. Heymann, B. Huang, and I. Gladkova, Eds., 2007, vol. 6683, no. 1, p. 668 305.

[10] L. Wang, J. Wu, L. Jiao, and G. Shi, "3D medical image compression based on multiplierless low-complexity RKLT and shape-adaptive wavelet transform," in *Proc. ICIP*, Nov. 2009, pp. 2521–2524.

[11] L. Wang, J. Wu, L. Jiao, and G. Shi, "Lossy-to-lossless hyperspectral image compression based on multiplierless reversible integer TDLT/KLT," *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 3, pp. 587–591, Jul. 2009.

[12] I. P. Akam Bita, M. Barret, and D.-T. Pham, "On optimal orthogonal transforms at high bit-rates using only second order statistics in multicomponent image coding with JPEG2000," *Signal Process.*, vol. 90, no. 3, pp. 753–758, Mar. 2010.

[13] I. P. Akam Bita, M. Barret, and D.-T. Pham, "On optimal transforms in lossy compression of multicomponent images with JPEG2000," *Signal Process.*, vol. 90, no. 3, pp. 759–773, Mar. 2010.

[14] M. Barret, J.-L. Gutzwiller, I. P. Akam Bita, and F. Dalla Vedova, "Lossy hyperspectral images coding with exogenous quasi optimal transforms," in *Proc. DCC*, Mar. 2009, pp. 411–419.

[15] I. P. Akam Bita, M. Barret, F. Dalla Vedova, and J.-L. Gutzwiller, "Lossy compression of MERIS superspectral images with exogenous quasi optimal coding transforms," in *Proc. SPIE—Satellite Data Compression, Communication, and Processing V*, B. Huang, A. J. Plaza, and R. Vitulli, Eds., 2009, vol. 7455, no. 1, p. 745 50U.

[16] Y. Wongsawat, "Lossless compression for 3-D MRI data using reversible KLT," in *Proc. ICALIP*, Jul. 2008, pp. 1560–1564.

[17] I. Blanes and J. Serra-Sagristà, "Clustered reversible-KLT for progressive lossy-to-lossless 3D image coding," in *Proc. DCC*, Mar. 2009, pp. 233–242.

[18] I. Blanes and J. Serra-Sagristà, "Cost and scalability improvements to the Karhunen–Loêve transform for remote-sensing image coding," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 7, pp. 2854–2863, Jul. 2010.

[19] J. A. Saghri and S. Schroeder, "An adaptive two-stage KLT scheme for spectral decorrelation in hyperspectral bandwidth compression," *Proc. SPIE*, vol. 7443, p. 744 313, 2009.

[20] J. A. Saghri, S. Schroeder, and A. G. Tescher, "Adaptive two-stage Karhunen–Loêve-transform scheme for spectral decorrelation in hyperspectral bandwidth compression," *Proc. SPIE*, vol. 49, p. 057 001, May 2010.

[21] Q. Du, W. Zhu, H. Yang, and J. E. Fowler, "Segmented principal component analysis for parallel compression of hyperspectral imagery," *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 4, pp. 713–717, Oct. 2009.

[22] P. W. Hao and Q. Y. Shi, "Matrix factorizations for reversible integer mapping," *IEEE Trans. Signal Process.*, vol. 49, no. 10, pp. 2314–2324, Oct. 2001.

[23] L. Galli and S. Salzo, "Lossless hyperspectral compression using KLT," in *Proc. IEEE IGARSS*, 2004, vol. 1–7, pp. 313–316.

[24] F. Kruse, J. Boardman, and J. Huntington, "Comparison of airborne hyperspectral data and EO-1 Hyperion for mineral mapping," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 6, pp. 1388–1400, Jun. 2003.
[25] Jet Propulsion Laboratory, NASA, Airborne Visible InfraRed Imaging Spectrometer Website. [Online]. Available: http://aviris.jpl.nasa.gov/html/aviris.overview.html
[26] U.S. Geological Survey and NASA, Earth Observing 1, Hyperion Website. [Online]. Available: http://eo1.usgs.gov/hyperion.php
[27] Jet Propulsion Laboratory, NASA, Hyperspectral Image Compression Website. [Online]. Available: http://compression.jpl.nasa.gov/hyperspectral/
[28] U.S. Geological Survey, Earth Explorer. [Online]. Available: http://earthexplorer.usgs.gov
[29] D. Taubman and M. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards, and Practice*. Boston, MA: Kluwer, 2002, ser. Kluwer International Series in Engineering and Computer Science.
[30] F. Garcia-Vilchez and J. Serra-Sagristà, "Extending the CCSDS recommendation for image data compression for remote sensing scenarios," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 10, pp. 3431–3445, Oct. 2009.
[31] Consultative Committee for Space Data Systems, Blue Book, CCSDS 122.0-B-1 Image Data Compression, Nov. 2005 Blue Book, CCSDS 122.0-B-1.
[32] D. Taubman, Kakadu Software, 2000, [Online]. Available: http://www.kakadusoftware.com/
[33] J. Fowler, "QccPack: An open-source software library for quantization, compression, and coding," in *Proc. SPIE—Applications of Digital Image Processing XXIII*, A. Tescher, Ed., San Diego, CA, Aug. 2000, vol. 4115, pp. 294–301.
[34] Group on Interactive Coding of Images, TER software; Open Source CCSDS-122-B-1 implementation and extension, Jun. 2008. [Online]. Available: http://www.gici.uab.cat/TER, http://ter.sourceforge.net/
[35] Group on Interactive Coding of Images, Spectral Transform Software, 2010. [Online]. Available: http://gici.uab.cat/
[36] P. Deutsch and J.-L. Gailly, ZLIB Compressed Data Format Specification Version, May1996. [Online]. Available: http://www.gzip.org/zlib/rfc1950.pdf
[37] B. Penna, T. Tillo, E. Magli, and G. Olmo, "A new low complexity KLT for lossy hyperspectral data compression," in *Proc. IEEE IGARSS*, 2006, pp. 3525–3528.
[38] B. Penna, T. Tillo, E. Magli, and G. Olmo, "Progressive 3-D coding of hyperspectral images based on JPEG2000," *IEEE Geosci. Remote Sens. Lett.*, vol. 3, no. 1, pp. 125–129, Jan. 2006.
[39] J. E. Fowler and J. T. Rucker, "3D wavelet-based compression of hyperspectral imager," in *Hyperspectral Data Exploitation: Theory and Applications*. Hoboken, NJ: Wiley, 2007, pp. 379–407.
[40] B. Penna, T. Tillo, E. Magli, and G. Olmo, "Hyperspectral image compression employing a model of anomalous pixels," *IEEE Geosci. Remote Sens. Lett.*, vol. 4, no. 4, pp. 664–668, Oct. 2007.
[41] A. Lucero, S. Cabrera, E. J. Vidal, and A. Aguirre, "Evaluating residual coding with JPEG2000 for L-infinity driven hyperspectral image compression," in *Proc. SPIE—Satellite Data Compression, Communications, and Archiving*, 2005, vol. 5889, no. 1, pp. 12–23.
[42] G. Carvajal, B. Penna, and E. Magli, "Unified lossy and near-lossless hyperspectral image compression based on JPEG 2000," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 4, pp. 593–597, Oct. 2008.
[43] J. Mielikainen and P. Toivanen, "Lossless compression of hyperspectral images using a quantized index to lookup tables," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 3, pp. 474–478, Jul. 2008.
[44] A. Kiely and M. Klimesh, "Exploiting calibration-induced artifacts in lossless compression of hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 8, pp. 2672–2678, Aug. 2009.
[45] M. Klimesh, "Low-complexity adaptive lossless compression of hyperspectral imagery," *Proc. SPIE*, vol. 6300, pp. 63000N.1–63000N.9, Sep. 2006.
[46] E. Magli, "Multiband lossless compression of hyperspectral images," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 4, pp. 1168–1178, Apr. 2009.
[47] A. Abrardo, M. Barni, E. Magli, and F. Nencini, "Error-resilient and low-complexity onboard lossless compression of hyperspectral images by means of distributed source coding," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 4, pp. 1892–1904, Apr. 2010.
[48] J. B. MacQueen, "Some methods for classification and analysis of multivariate observations," in *Proc. 5th Berkeley Symp. Math. Statist. Probab.*, L. M. L. Cam and J. Neyman, Eds., 1967, vol. 1, pp. 281–297.
[49] I. Reed and X. Yu, "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 10, pp. 1760–1770, Oct. 1990.
[50] C.-I. Chang and S.-S. Chiang, "Anomaly detection and classification for hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 40, no. 6, pp. 1314–1325, Jun. 2002.

Ian Blanes (S'05) received the B.S. and M.S. degrees in computer science from the Universitat Autònoma de Barcelona, Cerdanyola del Vallès, Spain, in 2007 and 2008, respectively, where he is currently working toward the Ph.D. degree.

Since 2003, he has been with the Group on Interactive Coding of Images, Universitat Autònoma de Barcelona. In 2007, he was an Intern with the Thomson Corporate Research, Princeton, NJ. He is currently on a research grant by the Spanish Ministry of Education.

Mr. Blanes was the recipient of the 2007 awards by the Spanish Ministry of Education as the second-best computer science student of Spain.

Joan Serra-Sagristà (S'97–M'05) received the B.S., M.S., and Ph.D. degrees in computer science from the Universitat Autònoma de Barcelona, Cerdanyola del Vallès, Spain, in 1992, 1994, and 1999, respectively.

Since 1992, he has been with the Department of Information and Communications Engineering, Universitat Autònoma de Barcelona, where he is currently an Associate Professor and the Director of the Group on Interactive Coding of Images. From September 1997 to December 1998, he was on a DAAD research grant at the University of Bonn, Bonn, Germany. From June to July 2000, he was a Visiting Researcher with the University of Bonn. His current research interests include image coding, data compression, vector quantization, and wavelet-based techniques, with special attention to remote sensing and telemedicine applications. He has coauthored several papers in these areas.

Dr. Serra-Sagristà is a member of the SPIE. He has served on the steering and technical program committees of several international conferences, and he is a reviewer for the major international journals in his research field. He was the recipient of the Intensification Program Young Investigator Award in 2006.