

Cost and Scalability Improvements to the Karhunen–Loève Transform for Remote-Sensing Image Coding

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

Abstract—The Karhunen–Loève transform (KLT) is widely used in hyperspectral image compression because of its high spectral decorrelation properties. However, its use entails a very high computational cost. To overcome this computational cost and to increase its scalability, in this paper, we introduce a multilevel clustering approach for the KLT. As the set of different multilevel clustering structures is very large, a two-stage process is used to carefully pick the best members for each specific situation. First, several candidate structures are generated through local search and eigenthresholding methods, and then, candidates are further screened to select the best clustering configuration. Two multilevel clustering combinations are proposed for hyperspectral image compression: one with the coding performance of the KLT but with much lower computational requirements and increased scalability and another one that outperforms a lossy wavelet transform, as spectral decorrelator, in quality, cost, and scalability. Extensive experimental validation is performed, with images from both the AVIRIS and Hyperion sets, and with JPEG2000, 3D-TCE, and CCSDS-Image Data Compression recommendation as image coders. Experiments also include classification-based results produced by *k*-means clustering and Reed–Xiaoli anomaly detection.

Index Terms—Component scalability, hyperspectral data coding, Karhunen–Loève Transform (KLT), low cost, progressive lossy-to-lossless (PLL) and lossy compression.

I. INTRODUCTION

CURRENT remote sensors gather very large amounts of information: They capture images with sizes ranging from a few megabytes to the gigabyte range, and they produce lots of such images, increasing the amount of information produced, which is expensive to store and slow to manipulate and transmit. Image compression techniques are capable of substantially reducing image sizes, decreasing costs, and improving user interaction with the information.

A common kind of images that remote sensors acquire is hyperspectral images, those with a few hundreds of spectral components. In hyperspectral images, each of these components represents the light of a given wavelength range for the same




Fig. 1. NASA AVIRIS Low Altitude image coded with JPEG2000, with and without spectral transforms.

spatial location, such that the components are highly correlated and redundancy exists between them. An appropriate transform applied in the spectral domain removes this redundancy and improves compression performance.

Common spectral transforms are the discrete wavelet transform (DWT) and the Karhunen–Loève Transform (KLT). Their use provides substantial compression improvements, with the KLT yielding significantly better results [1]. However, the better performance of the KLT comes at the price of an increase of two orders of magnitude in computational cost and at the price of not being scalable, which is the ability of reverting the transform partially. Fig. 1 shows an example of the quality improvements produced by spectral transforms when used in combination with the JPEG2000 coding-system standard.

In this paper, we propose a family of variations to the KLT, which will bring computational costs and scalabilities comparable with those of the DWT but with almost the same compression performance as the original KLT.

Computational costs are an important issue to take into account; for example, on a theoretical 1-gflop CPU, applying and removing a KLT transform on an image of size $224 \times 2048 \times 512$ ($z \times y \times x$) takes more than 4 min. In contrast, it only takes 6 s if a DWT is employed as spectral transform. A cost comparison among the KLT, the DWT, and our proposal, when used as spectral decorrelators on an image of this size, is provided in Fig. 2.

High computational costs are also an issue in the progressive lossy-to-lossless (PLL) scenario, in which the coder used produces lossless encodings when enough bytes are allocated, but as opposed to a pure lossless method, it also provides lower quality versions of the image when not enough bytes are

Manuscript received July 24, 2009; revised December 13, 2009. Date of publication April 5, 2010; date of current version June 23, 2010. This work was supported in part by the Spanish Government, by the Catalan Government, and by FEDER under Grants TIN2009-14426-C02-01, TIN2009-05737-E, 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.2010.2042063




Fig. 2. Spectral-transform cost for an image of size $224 \times 2048 \times 512$ ($z \times y \times x$). Lossy compression.

available. In PLL, the reversible KLT (RKLT) [2], an approximation of the KLT, has to be employed. For the RKLT, similar cost and scalability improvements will also be shown, and in this case, the RKLT will be compared with a common reversible wavelet, the integer wavelet transform (IWT).

Several improvements on the computational cost of KLT are found in the literature, starting with the discrete cosine transform (DCT), which was introduced as a low-cost alternative to the KLT [3]. The cost reduction stems from the additional assumption of a first-order Markov model. While still widely used in image coding, the DCT performs poorly as a spectral decorrelator. Later, the fast approximate KLT (AKLT) was introduced [4]. It is based on linear perturbation theory and uses the DCT as the related problem. It was also further improved as the AKLT₂ [5], which extends the AKLT with second-order perturbations.

Other strategies for cost reduction try to reduce transform training costs, which amount approximately to one-fifth of the total cost. In [1], it is proposed to compute the empirical covariance matrix from a small sample of the whole input, and a subsampling factor of 1% is shown to be adequate for hyperspectral images. This strategy is compatible with our proposal and will be discussed later. Another strategy to reduce training computational costs is to train once the KLT for a training set of images for all images [6], although with afterward, employ the already trained transform for all images, although with this strategy, the transform is not specifically trained for each image, and its coding performance is worse.

On the other hand, the RKLT has been a recent subject, and there are few improvements that alleviate its cost. In [7], a recursive clustered structure is used to reduce the computational cost of the RKLT in lossless coding of electroencephalograms, and in [8], RKLT is used for lossless compression of magnetic resonance images. Similarly, we developed a multilevel clustered KLT structure for PLL [9]. This structure yields greater savings and leads to the family of structures discussed in this paper.

There are two objectives in this paper: to reduce the computational cost of the KLT to an acceptable level and to add component scalability to the transform, both for lossy and PLL image coding. To that purpose, a family of multilevel KLT clustering structures has been analyzed, and a procedure to pick appropriate transforms has been devised. In addition, very extensive experimental testing has been performed for a wide range of hyperspectral images, image coders, and quality measurements.

We should end this discussion by saying that our proposal for reducing the computational cost of the KLT may still be too

high for its practical use for onboard compression. As noted in [10], most image coding algorithms have excessive complexity for onboard compression, and lossless predictive compression methods are preferred. However, for on-the-ground compression, our proposal provides very competitive computational cost, improved component scalability, and yields high coding and classification-based performance.

This paper is organized as follows. Section II introduces the proposed family of variations of the KLT. Then, Section III describes a procedure to select notable elements of the proposed family and describes how to measure their suitability. Section IV contains the experimental results. Section V draws some conclusions.

II. LOW-COST SCALABLE KLT

The KLT is a linear transform, constrained to provide stochastic independence for centered Gaussian sources and, in this sense, optimality. This comes at the price of having to apply a different transform for each set of sources.

Let X be a matrix that has N rows, one for each source, and M columns, one for each sample. Then, Y , the outcome after applying the transform, is computed as

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

where $\Sigma_X = (1/M)XX^T$ is the covariance matrix of X , and Q is the orthogonal matrix obtained from the singular value decomposition (SVD) of $\Sigma_X = Q\Lambda Q^{-1}$ ($\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$, $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$). Note that the sources must be manually centered around zero if they do not have a zero mean.

Once the transform matrix has been generated (or trained), the KLT application consists only of a matrix multiplication. If a lossy transform is desired, then the multiplication can be naturally applied using fixed-precision numbers. However, if exact recovery is needed, then additional precautions must be taken. Problems arise from the matrix multiplication, which is not exactly reversible because of the small imprecisions caused by the use of floating-point arithmetic. This situation can be resolved by decomposing the matrix multiplication into reversible lifting steps, which can be later applied and reversed without loss [11].

Originally, two lifting decomposition strategies were proposed, the TERM and the SERM [2], where the latter trades an increase of factorization time for lesser application costs. Another variant was proposed in [12], where the quantization noise introduced by the lifting steps was reduced substantially but with higher computational costs. We call this last one the 3TERM.

A. Multilevel Clustering

There are three main sources for the computational cost of the KLT: the covariance matrix calculation and the forward and reverse applications of the transform. The three operations have an effective computational complexity of $\mathcal{O}(n^2)$, where n is the number of spectral components. With this in mind, a divide-and-conquer approach seems appropriate, as for each division,




Fig. 3. Example of a multilevel clustering KLT structure. This structure decorrelates 15 spectral components with three levels of clustering.




Fig. 4. Comparison of compression performance for several spectral transforms: full KLT, DWT, single-level KLT with 56 clusters, and multilevel KLT with 56 clusters on the first level and successive levels with half the number of clusters. Image is AVIRIS Yellowstone (Sc 0).

the total cost is halved. Therefore, dividing a large transform into multiple clusters (or spectral tiles) and applying smaller transforms yield significant reductions on all three cost sources.

A drawback of this simple approach is that it only provides local decorrelation within each cluster. However, global decorrelation can be achieved by applying additional transform stages, where only the “most important” parts are further decorrelated. This scheme is shown in Fig. 3. We know from [9] that using a homogeneous multilevel KLT structure, where half of each cluster is forwarded to the next stage, is a good start. In Fig. 4, the benefits of a homogeneous multilevel structure can be observed. Nevertheless, as will be seen later, other structures can be used to obtain further cost reductions.

To define a specific multilevel clustering KLT structure, the notation $\mathcal{C} = \{c_{level, index}\}$ is used. $\mathcal{C}_l$ is the set of clusters in level l of the structure $\mathcal{C}$. For each cluster, let $Size(c_{i,j})$ be the size of the cluster and $Th(c_{i,j})$ be the number of elements in the said cluster that proceeds to the next level. As an example, Fig. 3 scheme is defined by the set $\mathcal{C} = \{c_{1,1}, c_{1,2}, c_{1,3}, c_{2,1}, c_{2,2}, c_{3,1}\}$ with, for instance, $Size(c_{1,2}) = 5$ and $Th(c_{2,1}) = 2$.

Due to the large number of possible structures, choosing a specific multilevel structure from all the possible ones is not a trivial problem. For example, for as few as 16 components, the total number of possible structures¹ is $8.77 \cdot 10^{26}$. Moreover,

¹Multilevel clustering structures for n components can be counted with $g(n) = 1 + \sum_{t=2}^{n-1} f(n, t)g(t)$, where $f(n, t)$ is the number of one-level clustering configurations of n components that forwards t components to the next level.

structure parts cannot be selected independently, as the performance of the parts of one level depends on the configuration of the parts of the previous level.

How to select a structure as a whole for each specific situation is addressed in Section III, while how to select individual structure parts is addressed here.

The structure specification can be divided into two kinds of problems: determining the size of each cluster in the structure and, for each cluster, setting the threshold on how many components are forwarded to the next stage.

For the first problem, cluster size determination, a local search may be reasonable, as the number of possible cluster sizes for one level is relatively small. Moreover, search cost can be reduced if some constraints are added to the cluster sizes such as homogeneity within a level.

The second problem, deciding how many components are forwarded to the next stage, can be effectively linked to eigenthresholding methods, which try to determine the amount of components to retain after a KLT, which is precisely our situation. The set of different eigenthresholding methods is very large. The ones found to be more relevant to this research are as follows.

- 1) *Scree test* [13]. It is a subjective test based on the visual inspection of a plot of the eigenvalues. The threshold is set where the viewer determines that there is a change of trend (i.e., the “elbow” of the plot).
- 2) *Average Eigenvalue (AE)*. Components with above-average eigenvalues are selected. Usually, only the first component is selected, while on other occasions, two are selected.
- 3) *Empirical Indicator Function (EIF)* [14]. Malinowski discovered a simple function (2) that predicted the correct number of factors in his chemical experiments. It yields results very similar to the scree test. It might be a good automated replacement, i.e.,

$$EIF = \arg \min_{j=0, \dots, N} \sqrt{\frac{\sum_{j < i} \lambda_i}{M(N-j)^3}}. \quad (2)$$

The following methods were also examined but had to be discarded due to more efficient alternatives, sensitivity to non-Gaussian noise, or inability to operate on zero-mean data: energy percentage, parallel analysis [15], Akaike information criterion [16], minimum description length [17], [18], and Harsanyi–Farrand–Chang method. [19].

Fig. 5 shows an illustrative comparison of the different eigenthresholding methods. The figure shows the energy accumulation for an example cluster after a KLT transform and indicates, according to each method, how many components should be forwarded to the next level.

In the relevant eigenthresholding methods, two clear trends are observed. The first trend corresponds to the AE method, where one component is selected most of the times, and, on rare occasions, two components are selected. The other trend corresponds to the scree test and EIF. The scree test has the disadvantage that human intervention is required every time a threshold needs to be determined.




Fig. 5. Comparison of the different eigenthresholding methods. KLT applied on a cluster of 28 components of the AVIRIS image Moffett Field.




Fig. 6. Effects of image subsampling for covariance calculation. The AVIRIS image Jasper Ridge has been compressed at different bitrates with JPEG2000 on lossy mode.

B. Subsampling

Image subsampling in the KLT has been known to substantially reduce the training cost of the KLT [1]. The aim is to use only a small sample of the image spatial locations to compute the covariance matrix. Subsampling factors of $\rho = 0.01$ (1%) are known to provide almost the same compression performance results on unclustered transforms.

For clustered transforms, image subsampling in the KLT training also provides substantial cost reduction. It simply needs to be applied in each of the clusters since with clustered transforms, the number of spatial samples does not change.

We have reproduced the experiments in [1] and also found $\rho = 0.01$ to be an adequate subsampling factor for the multi-level clustering KLT approach. However, as shown in Fig. 6, this factor is bitrate dependent and could be further decreased for low bitrates. For hyperspectral images, a ρ factor of 0.01 effectively reduces the cost to a small fraction and leaves application and removal as the main sources of cost.

III. STRATEGY FOR CANDIDATE SELECTION

In this section, we formulate how to select which cluster combinations to explore from all the possible ones and how to evaluate their suitability.

Strong similarities have been observed on the performance of each multilevel clustered KLT structure across the different imagery used. Fig. 7 shows how the images used in the next




Fig. 7. Energy accumulation for two AVIRIS and one Hyperion images. One level of clustering is applied, with 16 clusters of 14 components each.

section have similar behaviors. While the energy at each component is not equal for the three different images, the peaks have the same width, which is a key aspect that allows result extrapolation.

Two strategies for structure selection are proposed. Both strategies produce various candidate structures for each specific situation; further screening leads to the selection of very effective transforms. The idea behind both strategies is to examine each cluster size, from 4 to 28, and forward as many components as the various eigenthresholding methods suggest, producing a list of different candidate structures.

To reduce the solution space (i.e., candidate list), three constraints are added. First, a constraint of cluster size homogeneity has been added to each clustering level

$$\forall l, i, j, \quad \text{Size}(c_{l,i}) = \text{Size}(c_{l,j}).$$

Second, nonregular structures where the cluster number does not divide the number of components are not examined, as they can be easily adapted from regular structures, using one cluster with fewer components. Third, structures with more than one cluster on the top level are discarded, as early results showed huge quality penalties.

As described next, strategies differ on whether a dynamic eigenthresholding method is used or not.

- 1) Given the previous restrictions, if thresholds are known *a priori*, the result is a fully *static structure*, where the whole structure can be chosen before it is applied and does not vary with different images. To define static structures, we set the same threshold value for each cluster of a level based on the average output of either the EIF or the AE method for a training image. As all the clusters in a level have the same size and threshold, static structures can be identified with a list of L elements, corresponding to the number of clusters of each level, their sizes, and their thresholds, i.e.,

$$[(\#C_1, \text{Size}(c_1), \text{Th}(c_1)), \dots, (\#C_{L-1}, \text{Size}(c_{L-1}), \text{Th}(c_{L-1})), (\#C_L, \text{Size}(c_L), 0)]$$

where c_l is any cluster of level l .

For example, the structure $[(2, 5, 3), (1, 6, 0)]$ has two clusters of five components on the first level. From each

TABLE I

DETAIL OF CLUSTER COST (IN FLOPS). ρ IS THE SUBSAMPLING FACTOR, $n = \text{Size}(c)$ IS THE NUMBER OF COMPONENTS OF THE CLUSTER, AND M IS THE NUMBER OF SPATIAL LOCATIONS OF THE ORIGINAL IMAGE. SVD COST IS FROM [20], AND TERM/SERM APPLICATION COSTS ARE FROM [2]

	KLT Variant			
Stage	Lossy KLT	RKLT TERM	RKLT SERM	RKLT 3TERM
Covariance Matrix	$\rho M(n^2 + n + 4)$			
Singular Value Decomposition	$\simeq 9n^3$			
Factorization	—	$n^3 - \frac{n^2}{2} + \frac{n}{2} - 1$	$\frac{11n^3}{3} - 5n^2 + \frac{7n}{3} - 1$	$\frac{5n^3}{3} + \frac{n^2}{2} - \frac{37n}{6} + 5$
Apply or Remove (done twice)	$M(2n^2 - n)$	$M(2n^2 + 2n - 3)$	$M(2n^2 + n - 1)$	$M(3n^2 - 3)$

TABLE II

TECHNICAL NAMES FOR AVIRIS AND HYPERION IMAGES USED

Name	Technical Name	Size ($x \times y \times z$)
AVIRIS Cuprite	f970619t01p02_r02	$614 \times 2206 \times 224$
AVIRIS Jasper Ridge	f970403t01p02_r03	$614 \times 2586 \times 224$
AVIRIS Low Altitude	f960705t01p02_r05	$614 \times 3689 \times 224$
AVIRIS Lunar Lake	f970623t01p02_r07	$614 \times 1431 \times 224$
AVIRIS Moffett Field	f970620t01p02_r03	$614 \times 2031 \times 224$
AVIRIS Yellowstone (Sc 0)	f060925t01p00_r12_sc00	$677 \times 512 \times 224$
AVIRIS Yellowstone (Sc 11)	f060925t01p00_r12_sc11	$677 \times 512 \times 224$
Hyperion Agricultural	EO1H0280342002068110PX	$256 \times 2905 \times 224$
Hyperion Coastal	EO1H0140342002050110PY	$256 \times 2905 \times 224$
Hyperion Coral Reef	EO1H0830742003120110PW	$256 \times 3127 \times 224$
Hyperion Tornado	EO1H0150332002121112PF	$256 \times 3352 \times 224$
Hyperion Urban	EO1H0440342002121110PY	$256 \times 2905 \times 224$

distortion, quality scalability, and other features. All three coders are able to perform both lossy and PLL encodings.

The implementations chosen to perform the experiments are the following: Kakadu for JPEG2000 [27], QccPack for 3d-TCE [28], TER for CCSDS-IDC [29], and our own open-source implementation of the clustered KLT [30].

Candidate multilevel clustered KLT structures are evaluated as spectral transforms, in comparison with the original KLT and also with wavelets. The wavelet transforms used as spectral transforms are the Cohen–Daubechies–Feauveau (CDF) 9/7 for lossy coding and the CDF 5/3 for PLL; both of them were applied with five levels.

The remainder of this section is organized as follows: First, the results for candidate multilevel KLT structures are discussed; then, rate-distortion (R-D) evolutions and CRs are analyzed for the best structures; and finally, some classification-based results are reported.

All the candidate structures have been tested in order to assess their properties. In Table III, results in relation to the KLT for the candidate structures are presented, and for each structure, measures of the three criteria proposed in the previous section are listed for both lossy and PLL. As can be seen, dynamic thresholds provide the greater cost reductions and better scalabilities while having a moderate quality penalty. On the other hand, static thresholds guarantee stable quality with still good computational savings. The most remarkable structures on the candidate list are the one with static thresholding with the following configuration: [(32, 7, 2), (4, 16, 6), (2, 12, 4), (1, 8, 0)] and the dynamic one based on AE and a cluster size of four (AE size 4). The former achieves almost the same quality as the full KLT, while keeping cost lower than the previous approaches, and still with a decent scalability. The latter outperforms a lossy wavelet transform in all criteria and is very competitive against lossless wavelets. Note that a static

structure, due to its nature, might also be more amenable to hardware implementations.

From the AVIRIS corpus, Yellowstone images were acquired in 2006 and calibrated with a different scaling factor [31]. For these Yellowstone images, there is a slightly better static structure: [(32, 7, 3), (8, 12, 5), (2, 20, 7), (1, 14, 0)], which is capable of obtaining higher qualities with a small cost increase. However, for the other images tested, this last structure performs similarly to the other structures, while it has a higher cost.

Further analysis of the performance of the two first structures follows. For brevity, from now on, they will be referred to as static and dynamic. Details on each structure organization and costs can be found in Figs. 9 and 10, respectively.

Quality-evolution plots of the two selected structures are shown in Fig. 11. The proposed structures perform well, both in lossy and PLL, on all the images tested and no significant changes are appreciated whether JPEG2000, TCE, or CCSDS-IDC is used. As expected, the static structure yields almost the same performance as the original transform; in addition, on very low bitrates, it performs even better due to the reduced amount of side information required by the smaller KLT matrices. On the Hyperion images, the reduction of side information is not so relevant, as the proportion of side information to the spatial size is lower, and still the results are similar. Moreover, as expected, the dynamic structure consistently outperforms the wavelet transform as spectral decorrelator in all the scenarios tested.

Lossless compression rates are reported in Table IV, where it can be observed that little impact is produced by the use of the static clustered approach (usually bitstream sizes differ by less than 0.1 bpppb). On the other hand, a dynamic cluster approach produces larger bitstreams, which have CRs between the RKLT and the IWT. Still, the results are close to the ones of the RKLT for the dynamic structure. Note that the proposed transforms

TABLE III

DETAILED RESULTS FOR STATIC AND DYNAMIC STRUCTURES. CODING METHOD IS JPEG2000. Δ SNR IS THE AVERAGE OF THE SNR DIFFERENCE WITH THE FULL TRANSFORM. IT IS SAMPLED AT 0.1 k bpppb for $k = 1, \dots, 20$ AND IS MEASURED IN DECIBELS. Δ Size IS THE INCREASE OF SIZE FOR LOSSLESS BITSTREAMS. COST ONLY INCLUDES TRAINING, APPLICATION, AND REMOVAL OF THE SPECTRAL TRANSFORM AND IS IN GIGAFLOPS. A CONSTANT SUBSAMPLING FACTOR OF $\rho = 0.01$ HAS BEEN USED FOR ALL KLT TRANSFORMS

AVIRIS Cuprite:

Transform	Reference Transforms					
	Lossy		PLL			Both
	Δ SNR (min/max)	Cost	Δ SNR (min/max)	Δ Size	Cost	Scalability (min/max)
Full KLT Transform	0.00 (0.00 / 0.00)	52.9	0.00 (0.00 / 0.00)	0.00%	79.3	224.0 (224 / 224)
Half Cluster [9]	-0.02 (-0.12 / +0.06)	6.3	-0.02 (-0.11 / +0.03)	-0.56%	9.7	42.0 (42 / 42)
DWT (5 levels)	-3.22 (-5.72 / -2.23)	1.6	-	-	-	36.0 (32 / 38)
IWT (5 levels)	-	-	-3.23 (-6.17 / -1.84)	+8.01%	1.0	16.0 (11 / 17)

Transform	Multi-level clustering with dynamic thresholds					
	Lossy		PLL			Both
	Δ SNR (min/max)	Cost	Δ SNR (min/max)	Δ Size	Cost	Scalability (min/max)
EIF size 28	+0.06 (-0.01 / +0.14)	9.1	+0.10 (+0.02 / +0.21)	-0.90%	13.7	64.8 (51 / 72)
EIF size 14	+0.04 (-0.04 / +0.12)	4.7	+0.11 (+0.01 / +0.24)	-0.92%	7.2	41.2 (29 / 51)
EIF size 7	-0.12 (-0.27 / -0.03)	2.5	+0.00 (-0.14 / +0.24)	-0.48%	3.8	26.1 (19 / 32)
EIF size 4	-0.34 (-0.58 / -0.24)	1.6	-0.13 (-0.41 / +0.21)	+0.10%	2.5	18.5 (13 / 25)
AE size 28	-0.05 (-0.13 / +0.02)	7.5	+0.05 (-0.06 / +0.25)	-0.56%	11.3	52.5 (46 / 55)
AE size 14	-0.08 (-0.17 / 0.00)	4.0	+0.04 (-0.08 / +0.27)	-0.57%	6.0	35.0 (27 / 40)
AE size 7	-0.29 (-0.51 / -0.18)	2.0	-0.09 (-0.34 / +0.25)	-0.03%	3.1	18.5 (17 / 19)
AE size 4	-0.34 (-0.89 / -0.22)	1.3	-0.10 (-0.62 / +0.28)	+0.02%	2.0	12.9 (12 / 13)

Transform	Multi-level clustering with static thresholds (only the most relevant)					
	Lossy		PLL			Both
	Δ SNR (min/max)	Cost	Δ SNR (min/max)	Δ Size	Cost	Scalability (min/max)
[(8, 28, 6), (4, 12, 6), (2, 12, 5), (1, 10, 0)]	+0.06 (-0.06 / +0.16)	7.6	+0.12 (+0.02 / +0.24)	-0.93%	11.6	45.0 (45 / 45)
[(16, 14, 3), (4, 12, 6), (1, 24, 0)]	+0.09 (-0.08 / +0.18)	4.5	+0.17 (+0.05 / +0.31)	-1.07%	6.9	41.0 (41 / 41)
[(32, 7, 1), (4, 8, 1), (1, 4, 0)]	-0.31 (-0.71 / -0.17)	2.0	-0.10 (-0.49 / +0.27)	-0.03%	3.0	17.0 (17 / 17)
[(32, 7, 2), (4, 16, 6), (2, 12, 4), (1, 8, 0)]	+0.04 (-0.09 / +0.14)	3.1	+0.15 (+0.03 / +0.34)	-0.99%	4.8	31.0 (31 / 31)
[(32, 7, 2), (8, 8, 4), (2, 16, 10), (1, 20, 0)]	+0.04 (-0.11 / +0.13)	3.1	+0.15 (+0.04 / +0.30)	-0.97%	4.8	35.0 (35 / 35)
[(32, 7, 3), (8, 12, 5), (2, 20, 7), (1, 14, 0)]	+0.06 (-0.04 / +0.15)	3.9	+0.12 (+0.02 / +0.25)	-0.93%	6.0	38.0 (38 / 38)
[(56, 4, 1), (7, 8, 4), (4, 7, 4), (1, 16, 0)]	-0.14 (-0.35 / -0.07)	1.9	+0.04 (-0.15 / +0.32)	-0.48%	2.9	27.3 (26 / 29)

AVIRIS Yellowstone (Sc 0):

Transform	Lossy		PLL			Both
	Δ SNR (min/max)	Cost	Δ SNR (min/max)	Δ Size	Cost	Scalability (min/max)
Half Cluster [9]	-0.56 (-0.92 / -0.13)	6.3	+0.01 (-0.56 / +0.57)	-0.20%	9.7	42.0 (42 / 42)
DWT (5 levels)	-7.92 (-11.26 / -5.71)	1.6	-	-	-	36.0 (32 / 38)
IWT (5 levels)	-	-	-6.40 (-10.96 / -2.34)	+22.78%	1.0	16.0 (11 / 17)
AE size 4	-3.55 (-5.68 / -2.25)	1.3	-2.00 (-5.14 / +0.50)	+9.09%	2.0	12.9 (12 / 13)
[(32, 7, 2), (4, 16, 6), (2, 12, 4), (1, 8, 0)]	-0.97 (-1.48 / -0.27)	3.1	-0.08 (-0.98 / +0.94)	+1.37%	4.8	31.0 (31 / 31)
[(32, 7, 3), (8, 12, 5), (2, 20, 7), (1, 14, 0)]	-0.54 (-0.83 / -0.10)	3.9	+0.12 (-0.44 / +0.84)	+0.09%	6.0	38.0 (38 / 38)

AVIRIS Yellowstone (Sc 11):

Transform	Lossy		PLL			Both
	Δ SNR (min/max)	Cost	Δ SNR (min/max)	Δ Size	Cost	Scalability (min/max)
Half Cluster [9]	-0.24 (-0.47 / -0.10)	6.3	+0.22 (-0.22 / +0.84)	-1.14%	9.7	42.0 (42 / 42)
DWT (5 levels)	-5.47 (-8.87 / -3.68)	1.6	-	-	-	36.0 (32 / 38)
IWT (5 levels)	-	-	-3.81 (-8.31 / -0.20)	+15.39%	1.0	16.0 (11 / 17)
AE size 4	-3.43 (-6.47 / -2.25)	1.3	-1.54 (-5.20 / +0.93)	+8.54%	2.0	12.9 (12 / 13)
[(32, 7, 2), (4, 16, 6), (2, 12, 4), (1, 8, 0)]	-0.70 (-1.09 / -0.41)	3.1	+0.29 (-0.53 / +1.28)	-0.06%	4.8	31.0 (31 / 31)
[(32, 7, 3), (8, 12, 5), (2, 20, 7), (1, 14, 0)]	-0.32 (-0.52 / -0.13)	3.9	+0.39 (-0.11 / +1.23)	-1.19%	6.0	38.0 (38 / 38)

Hyperion Coral Reef:

Transform	Lossy		PLL			Both
	Δ SNR (min/max)	Cost	Δ SNR (min/max)	Δ Size	Cost	Scalability (min/max)
Half Cluster [9]	-0.04 (-0.10 / +0.04)	19.4	-0.04 (-0.11 / +0.05)	+0.04%	29.6	42.0 (42 / 42)
DWT (5 levels)	-2.26 (-2.85 / -1.86)	4.9	-	-	-	36.0 (32 / 38)
IWT (5 levels)	-	-	-2.41 (-2.88 / -2.10)	+6.02%	3.1	16.0 (11 / 17)
AE size 4	-0.91 (-1.63 / -0.72)	3.9	-0.83 (-1.54 / -0.63)	+0.71%	6.0	12.9 (12 / 13)
[(32, 7, 2), (4, 16, 6), (2, 12, 4), (1, 8, 0)]	-0.25 (-0.43 / -0.13)	9.5	-0.20 (-0.28 / -0.11)	-0.34%	14.5	31.0 (31 / 31)
[(32, 7, 3), (8, 12, 5), (2, 20, 7), (1, 14, 0)]	-0.25 (-0.42 / -0.17)	11.9	-0.21 (-0.29 / -0.16)	-0.69%	18.2	38.0 (38 / 38)




Fig. 9. Best structures of multilevel KLT clustering for hyperspectral images according to the procedure described in Section III. (a) static and (b) dynamic for the AVIRIS image Cuprite.

are also suitable to operate in near-lossless mode, where the peak absolute error is bounded, if a scheme similar to the ones proposed in [32] and [33] is employed.

As for the classification-based results, two very common unsupervised classification techniques are used: k -means clustering [34], and Reed–Xiaoli (RX) anomaly detection [35]. In this case, k -means has been configured to produce ten clusters with spectral angle as distance function; for RX, a confidence coefficient of $\gamma = 99\%$ has been set, and the results are reported using preservation of classification (POC). Fig. 12 shows the




Fig. 10. Costs of the selected structures for an image of size $224 \times 2048 \times 512$ ($z \times y \times x$). $\rho = 0.01$. (a) Lossy. (b) PLL.

performance of both classifiers when tested on images that have been compressed and uncompressed with PLL procedures. It is important to note that, for k -means, classes are more or less homogeneous in size, but for RX, the class that represents anomalies is very small. Therefore, a 1% drop on POC for k -means is very acceptable, while it is not for RX.

On tests for both classification techniques, the RKLT and its clustered variants perform similarly except on one sample on the RX results, where the RKLT produces a compression artifact that the RX detector marks as an anomaly. Nonetheless, this artifact is produced at a bitrate where the RX performance is already not acceptable. On the other hand, the IWT yields lower classification performance than the RKLT for the same bitrate.

In Fig 13, a visual comparison shows the different performance of an RX anomaly detector depending on the transform used. In order to produce visually appreciable results, the RX detector result is not thresholded, and the images are highly compressed. While all the techniques used introduced several artifacts on the RX results, these are more noticeable with the use of the dynamic structure or the IWT.

RX anomalies are usually computed after a KLT, and since the transform is a shared part in image coding and RX anomaly detection, [36] propose to improve the preservation of anomalies by simultaneously detecting anomalies, while the image is coded with a KLT. In that context, if the evaluation of the proposed transforms for RX anomaly detection proved positive, a similar scheme could be set up to perform both coding and detection.

V. CONCLUSION

In this paper, we have investigated how to reduce the computational cost of the KLT to acceptable terms through the use of multilevel clusterings. Clustered versions of the KLT provide a good solution, as they require much less computational resources than the KLT, they maintain its coding performance if applied over multiple levels, and are compatible with the use of subsampling in the covariance calculation, a technique that virtually removes transform training costs. With the proposed approaches, significant gains are achieved on the KLT training




Fig. 11. R-D evolution of the proposed structures for various techniques and images. (a) Lossy coding of the AVIRIS image Low Altitude with JPEG2000. (b) Lossy coding of the AVIRIS image Moffett Field with TER (an enhanced version of CCSDS-IDC) [25]. (c) PLL coding of the Hyperion image Tornado with JPEG2000. (d) PLL coding of the AVIRIS Moffett Field with TCE.

and on the forward and inverse application of the transform. In addition, clustered transforms have also quite good scalabilities.

Our objective has been to find suitable multilevel clustering structures, among all the possible clustering combinations, that had low resource consumption and yielded similar coding performances as the KLT. However, as the set of possible clusterings is very large, some strategy to select candidate members

TABLE IV
BITRATES AT WHICH ENCODERS PRODUCE LOSSLESS BITSTREAMS. MEASURED IN BITS/PIXEL/BAND. BEST AND WORST IMAGE CODING METHODS ARE HIGHLIGHTED FOR EACH CORPUS. THREE PURE LOSSLESS METHODS ARE ALSO INCLUDED AS REFERENCE: FAST-LOSSLESS (FL) [37], LAIS-QLUT-OPT [38], AND TSP-W2 [31]. FOR THESE RESULTS, THE COMPLETE AVIRIS IMAGES HAVE BEEN USED

	JPEG2000				FL	LLUT	TSP-W2
	RKLT	Static	Dynamic	IWT			
AVIRIS Cuprite	4.91	4.92	4.97	5.36	4.91	4.29	3.77
AVIRIS Jasper Ridge	4.86	4.92	5.14	5.60	4.95	4.61	4.08
AVIRIS Moffett Field	4.91	4.96	5.14	5.60	4.99	4.62	4.12
AVIRIS Lunar Lake	4.92	4.91	4.96	5.29	4.91	4.37	3.81
AVIRIS Yellowstone (Sc 0)	3.86	3.96	4.23	4.73	3.96	-	3.99
Average	4.69	4.73	4.89	5.32	4.74	"4.47"	3.95
Hyperion Agricultural	6.02	5.96	6.16	6.40	-	-	-
Hyperion Coastal	5.75	5.72	5.87	6.14	-	-	-
Hyperion Urban	6.06	6.01	6.28	6.56	-	-	-
Average	5.94	5.90	6.10	6.37	-	-	-




Fig. 12. Classification-based measures for the AVIRIS image Lunar Lake and JPEG2000. (a) *k*-means classification. (b) RX anomaly detection.

is necessary and to that purpose, eigenthresholding methods are useful. Eigenthresholding methods, which are used to discard noise components in a principal component analysis, are here employed to help determine the number of forwarded components between levels of clusters. Among the eigenthresholding methods examined, EIF and AE are the most suitable.

Moreover, two distinct approaches to select structures have been investigated, differing on whether eigenthresholding methods are used directly over every image (dynamic structures), or are used only to characterize the transform for a set of images (static structures). We found that static structures tend to have higher coding performances than the dynamic ones, but dynamic structures are the ones with greater cost savings.

Two most relevant multilevel clustered structures are found; one static and one dynamic. If compared with a wavelet transform, the former has a competitive cost and almost the same coding performance as the KLT, while the latter, also as com-




Fig. 13. RX anomalies for the Hyperion image Agriculture. The image has been encoded and decoded with one of the spectral transforms and JPEG2000 at 0.3 bpppb. (a) No compression. (b) RKLT full transform. (c) RKLT static clustering. (d) RKLT dynamic clustering. (e) IWT.

pared with wavelets, has a similar cost and better performance, although not as good a performance as the KLT. The use of one of these two transforms effectively solves our initial problem: computational costs. While the static transform allows applying a KLT with very reasonable resource constraints, the dynamic transform can be taken as a direct replacement of the DWT for spectral coding, improving the DWT in all the three measured criteria: quality, cost, and scalability. For both lossy and PLL, the proposed transforms perform between 16 and 50 times faster than the original KLT.

The best two transforms have been further validated in experimental results, which have shown similar behaviors when used in different images, with different image coders, or in lossy or lossless mode. In addition, the transforms have been used in classification-based experiments, which have also shown similar results.

ACKNOWLEDGMENT

Computational resources used in this work were provided in part by the Oliba Project of the Universitat Autònoma de Barcelona.

REFERENCES

- [1] 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.
- [2] 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.
- [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] L. Lan and I. Reed, "Fast approximate Karhunen–Loève transform (AKLT) with applications to digital image-coding," *Vis. Commun. Image Process.*, vol. 2094, pp. 444–455, 1993.
- [5] A. Pirooz and I. Reed, "A new approximate Karhunen–Loève transform for data compression," in *Conf. Rec. 32nd Asilomar Conf. Signals, Syst. Comput.*, 1998, vol. 1/2, pp. 1471–1475.
- [6] 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.
- [7] Y. Wongsawat, S. Orintara, and K. R. Rao, "Integer sub-optimal Karhunen–Loève transform for multi-channel lossless EEG compression," in *Proc. Eur. Signal Process. Conf.*, 2006.
- [8] Y. Wongsawat, "Lossless compression for 3-D MRI data using reversible KLT," in *Proc. ICALIP*, Jul. 2008, pp. 1560–1564.
- [9] I. Blanes and J. Serra-Sagrístà, "Clustered reversible-KLT for progressive lossy-to-lossless 3d image coding," in *Proc. DCC*, Mar. 2009, pp. 233–242.
- [10] 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. PP, no. 99, pp. 1–13, Nov. 2009.
- [11] F. A. M. L. Bruckers and A. W. M. van den Enden, "New networks for perfect inversion and perfect reconstruction," *IEEE J. Sel. Areas Commun.*, vol. 10, no. 1, pp. 130–137, Jan. 1992.
- [12] L. Galli and S. Salzo, "Lossless hyperspectral compression using KLT," in *Proc. IEEE IGARSS*, 2004, vol. 1–7, pp. 313–316.
- [13] R. B. Cattell, "The scree test for the number of factors," *Multivariate Behav. Res.*, vol. 1, no. 2, pp. 245–276, Apr. 1966.
- [14] E. R. Malinowski, "Determination of the number of factors and the experimental error in a data matrix," *Anal. Chem.*, vol. 49, no. 4, pp. 612–617, Apr. 1977.
- [15] J. L. Horn, "A rationale and a test for the number of factors in factor analysis," *Psychometrika*, vol. 30, no. 2, pp. 179–185, Jun. 1965.
- [16] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Autom. Control*, vol. AC-19, no. 6, pp. 716–723, Dec. 1974.
- [17] G. Schwarz, "Estimating the dimension of a model," *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, Mar. 1978.
- [18] J. Rissanen, "Modeling by shortest data description," *Automatica*, vol. 14, no. 5, pp. 465–471, Sep. 1978.
- [19] J. C. Harsanyi, W. Farrand, and C.-I. Chang, "Determining the number and identity of spectral endmembers: An integrated approach using Neyman–Pearson eigenthresholding and iterative constrained RMS error minimization," in *Proc. 9th Thematic Conf. Geologic Remote Sens.*, Feb. 1993, pp. 395–408.
- [20] G. Golub and C. van Loan, *Matrix Computations*. Baltimore, MD: Johns Hopkins Univ. Press, Oct. 1996.
- [21] *Airborne Visible InfraRed Imaging Spectrometer website*, Jet Propuls. Lab., NASA, Pasadena, CA. [Online]. Available: <http://aviris.jpl.nasa.gov/html/aviris.overview.html>
- [22] Earth Observing 1, Hyperion website, U.S. Geol. Survey, NASA, Reston, VA. [Online]. Available: <http://eo1.usgs.gov/hyperion.php>
- [23] D. Taubman and M. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards, and Practice*. Norwell, MA: Kluwer, ser. Kluwer International Series in Engineering and Computer Science, 2002.
- [24] 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.
- [25] F. Garcia-Vilchez and J. Serra-Sagrístà, "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.
- [26] *Image Data Compression*, Consultative Committee for Space Data Systems, Reston, VA, Nov. 2005. Blue Book, CCSDS 122.0-B-1.
- [27] D. Taubman, *Kakadu Software*, 2000. [Online]. Available: <http://www.kakadusoftware.com/>
- [28] J. Fowler, "QccPack: An open-source software library for quantization, compression, and coding," in *Proc. Appl. Digital Image Process. XXIII*, A. Tescher, Ed., San Diego, CA, Aug. 2000, vol. 4115, pp. 294–301.
- [29] *TER Software; Open Source CCSDS-122-B-1 Implementation and Extension*, Group on Interactive Coding of Images, Jun. 2008. [Online]. Available: <http://www.gici.uab.cat/TER>
- [30] *RKLT Software*, Group on Interactive Coding of Images, 2009. [Online]. Available: <http://gici.uab.cat/>
- [31] 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.
- [32] A. Lucero, S. Cabrera, E. J. Vidal, and A. Aguirre, "Evaluating residual coding with JPEG2000 for L-infinity driven hyperspectral image compression," *Satellite Data Compression, Commun., Archiving*, vol. 5889, no. 1, pp. 12–23, 2005.
- [33] 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.
- [34] 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.
- [35] 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.
- [36] 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.
- [37] M. Klimesh, "Low-complexity adaptive lossless compression of hyperspectral imagery," in *Proc. SPIE*, Sep. 2006, pp. 63 00N.1–63 00N.9.
- [38] 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.



Ministry of Education.

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



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. He is currently working toward the Ph.D. degree at the Universitat Autònoma de Barcelona.

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

Joan Serra-Sagrístà (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 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 at 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-Sagrístà is a member of SPIE. He has served on the steering committee and on technical program committees of several international conferences, and is a Reviewer for the major international journals in his research field. He is a recipient of the Intensification Program Young Investigator Award in 2006.