

Region-Based Transform Coding of Multispectral Images

Marco Cagnazzo, Giovanni Poggi, and Luisa Verdoliva

Abstract—We propose a new efficient region-based scheme for the compression of multispectral remote-sensing images. The region-based description of an image comprises a segmentation map, which singles out the relevant regions and provides their main features, followed by the detailed (possibly lossless) description of each region. The map conveys information on the image structure and could even be the only item of interest for the user; moreover, it enables the user to perform a selective download of the regions of interest, or can be used for high-level data mining and retrieval applications. This approach, with the multiple pieces of information required, may seem inherently inefficient. The goal of this research is to show that, by carefully selecting the appropriate segmentation and coding tools, region-based compression of multispectral images can be also effective in a rate-distortion sense, thus providing an image description that is both insightful and efficient. To this end, we define a generic coding scheme, based on Bayesian image segmentation and on transform coding, where several key design choices, however, are left open for optimization, from the type of transform, to the rate allocation procedure, and so on. Then, through an extensive experimental phase on real-world multispectral images, we gain insight on such key choices, and finally single out an efficient and robust coding scheme, with Bayesian segmentation, class-adaptive Karhunen–Loève spectral transform, and shape-adaptive wavelet spatial transform, which outperforms state-of-the-art and carefully tuned conventional techniques, such as JPEG-2000 multicomponent or SPIHT-based coders.

Index Terms—Lossy coding, multispectral images.

I. INTRODUCTION

MULTISPECTRAL¹ images are becoming more and more important for a large number of remote-sensing applications. The growing interest is motivated by several concurring reasons, one of the most important being the improved quality of state-of-the-art sensors which deliver images with very high spatial, spectral, and radiometric resolution. However, together with this wealth of information comes the problem of managing such a huge amount of data which must be transmitted to the

ground station on limited-capacity channels, archived for long periods of time, and finally disseminated to the end users on common transmission facilities. Therefore, suitable compression algorithms are highly desirable in all the phases of the data lifetime and can play a central role for the success of remote-sensing applications, as is testified by the growing bulk of related scientific literature, and by the attention paid to this topic by national space agencies.

General-purpose compression algorithms, like the JPEG, JPEG2000, or MPEG standards, usually do not provide satisfactory results for multispectral data, because they are not tailored to their statistical behavior. For this reason, several *ad hoc* compression techniques have been proposed in the last few years. Most of them rely on transform coding because of its simplicity and good results. Early techniques, e.g., [1]–[3], focused on various combinations of the Karhunen–Loève transform (KLT), discrete cosine transform (DCT), and different types of quantization. More recently, compression based on wavelet transform (WT) has gained popularity [4]–[11], showing in this field the same interesting performance exhibited in other contexts.

Even the specific techniques cited above, however, fail to take full advantage of the peculiar nature of multispectral data. Such images, in fact, portray exactly the *same* portion of the Earth in all component bands, only in different spectral windows. The fact that the same “objects” appear consistently in all the bands of the image, although with different properties, is an extremely valuable information which should be exploited to improve compression performance. A step in this direction is represented by the use of techniques [12]–[15] based on vector quantization (VQ) [16] which is able, in theory, to exploit all interband dependencies, not just the linear ones (that is, correlation). Unfortunately, VQ has a complexity that grows exponentially with the vector size, forcing practical coding schemes to deal with vectors that span only the spectral dimension, thus neglecting all spatial dependencies.² As a consequence, transform coding techniques keep providing better results.

An interesting compromise between VQ and transform coding is region-based coding, in which the image is segmented in a number of disjoint regions, based prevalently on their spectral features, which are then transform encoded independently from one another. A general coding scheme is depicted in Fig. 1: segmentation singles out the elementary component regions so that each one is statistically homogeneous; the segmentation map is then encoded (usually without losses) and sent as a side information, while the region textures

Manuscript received May 8, 2007; revised August 18, 2007. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Ricardo L. de Queiroz.

The authors are with the Dipartimento di Ingegneria Elettronica e delle Telecomunicazioni Università “Federico II” di Napoli, 21-80125 Napoli, Italy (e-mail: cagnazzo@unina.it; cagnazzo@i3s.unice.fr; poggi@unina.it; verdoliva@unina.it).

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/TIP.2007.909315

¹In the following, unless otherwise stated, we use “multispectral” also in place of “hyperspectral,” irrespective of the number of component bands.

²To be more precise, some constrained VQ techniques have been proposed [12], [14] to exploit spatial dependencies, but with mixed results.

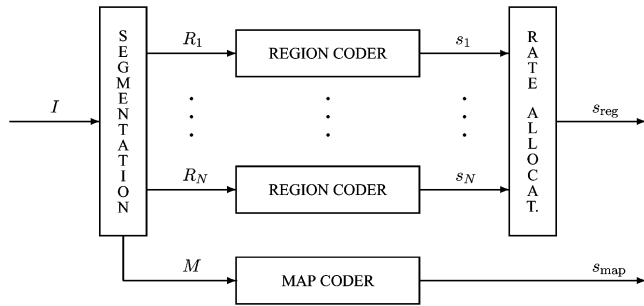


Fig. 1. General region-based coding scheme.

are lossy coded, independently from one another, and possibly with the help of information drawn from the segmentation process.

Region-based techniques offer several advantages with respect to conventional or “flat” techniques, both for low-level (e.g., compression, denoising) and high-level (e.g., classification, data-mining) processing tasks. At the lower level, advantages arise from the ability of segmenting the image in a relatively small number of regions, with the pixels of each region having similar statistics, which are in turn different from those of neighboring regions. As a consequence:

- the boundaries between different regions (which draw significant coding resources in flat image compression) are conveyed by the segmentation map and need not be encoded;
- the elementary regions, which are statistically homogeneous, can be compressed more efficiently;
- resources can be allocated adaptively to each region to optimize performance, and even the encoding technique itself can vary from region to region according to local statistics (dynamic coding [17]);
- since region boundaries are known in advance, denoising techniques can operate more effectively, region by region, without the risk of smearing or deleting important edges.

It must be underlined that the cost of coding the segmentation map, which is an obvious drawback of region-based coding, is very low with multispectral images since a single map is used for a large number of bands.

At a higher level, the advantages of a region-based description of the image are even more compelling, as testified by the interest in the video coding community, for example:

- the user is automatically provided with a segmentation map (obtained on the original data) which gives the basic structure of the image, and might even be the only product of interest if image classification/segmentation is pursued;
- the user can single out just one or a few regions of interest (because of their shape, their statistics or else) and download just those regions, with huge bandwidth saving;
- region-based description can help retrieve the images of interest in a large database, that is perform data mining.

Given all these potential advantages, there has been surprisingly little work on region-based compression of remote-sensing images, probably because of the inherent difficulty of carrying out a reliable segmentation and of encoding

the texture in arbitrarily shaped regions. A first simple approach is proposed in [18] where the image is segmented in fixed blocks of 64×64 pixels, on which statistics are computed to carry out a block-adaptive KLT in the spectral domain. With homogeneous blocks, the adaptive KLT guarantees a much better energy compaction, but no improvement is observed for the quite common mixed blocks. A more flexible solution is proposed in [19], where a variable-size quad-tree segmentation is carried out so as to obtain more homogeneous elementary blocks. Both in [18] and [19], segmentation is extremely simple but quite inaccurate, and its benefits are only marginal. In [20] and [21], and successively in [22] and [23], the approach is significantly refined by resorting to point-wise image segmentation by means of a minimum-distance classifier operating on the pixel spectral signatures. As a consequence, the subsequent transforms deal with subsources that are really stationary, and provide an energy compaction and eventually a performance that is much superior to that of previous work. It must be underlined, however, that this is a class-based rather than a region-based technique: pixels belonging to a given class are highly homogeneous, but can be scattered over the whole image in small clusters or even as isolated points. The consequence is that spatial encoding of a class is rather cumbersome and not very efficient, not to mention that the region-based description of the image is lost. In [24], a truly region-based coder is proposed which, in addition, has the ability to switch coding techniques depending on the region statistics. Despite the interesting approach, implementation presents several weak points that limit its performance and overall value, notably, the manual segmentation of images, and the use of region coding techniques far from the state of the art. A more interesting technique is proposed in [7] for the compression of oceanographic images. The problem of segmentation is not addressed (data are already available only in the ocean area), but a state-of-the-art shape-adaptive coding algorithm is proposed, based on stack-run coding, which provides significant improvements over existing techniques used for the same task.

In this paper, we investigate in depth the region-based approach for multispectral image coding, and try to take full advantage of its potential by resorting to what, in our view, are the most appropriate and powerful tools currently available. As a result of this effort, we propose a new region-based coding scheme which, on all test images considered, provides an excellent rate-distortion performance, often superior to the best comparable flat technique. The keys to such good results are a reliable segmentation algorithm, based on a tree-structured Markov random field model [25], the use of region/class adaptive KLT in the spectral domain, and shape-adaptive wavelet-based transform and coding tools in the spatial domain with proper rate allocation. After the Introduction, Section II describes the segmentation and coding tools, explaining the rationale for our choices, and often proposing several alternative solutions for the same task, to be later tested experimentally. Section III describes a sequence of experiments that guide us toward the choice of the best region-based coding algorithm, and others that compare its performance with that of the most relevant reference techniques for several test images and quality criteria. Section IV draws conclusions and outlines future work.




Fig. 2. (a) Sample image and (b) its segmentation by K-means, (c) TSMRF, and (d) TSMRF with removal of small regions.

II. SEGMENTATION AND CODING TOOLS

In this section, we examine in detail the various blocks of the general region-based coding scheme of Fig. 1 and for each one discuss one or more techniques to be later implemented and assessed experimentally. The region coder, in particular, will be based on transform coding, and, hence, we will discuss separately the transform and the actual coding engine. Therefore, the following subsections will examine in turn segmentation, map coding, spectral and spatial transforms, region coding by shape-adaptive SPIHT, and rate allocation.

A. Segmentation

A meaningful segmentation of the image is of central importance for the success of a region-based coding scheme, but it is also quite a complex problem, to begin with its very same definition. In our application, we have two requirements: on one hand, we want each region to be formed by pixels of the same type, so as to exhibit homogeneous statistics and increase the efficiency of subsequent encoding. On the other hand, we would like to segment the image in a small number of large regions, in order to have a simple map to encode, and to use shape-adaptive transforms on nice regular shapes.

To help gain insight about the following discussion, let us consider as a case study the 256×256 section of a larger hyperspectral image, shown in false colors in Fig. 2(a). Given the wealth of spectral information available with this kind of data, it might seem that a simple clustering in the spectral domain could provide a satisfactory segmentation. The map of Fig. 2(b), however, obtained by means of vector quantization (also known as K-means clustering) makes clear that such a strategy cannot produce by itself a segmentation useful for our purposes, because of the very large number of small regions and isolated points,

also due to the effects of noise. In more detail, we can highlight three drawbacks.

- 1) The map is not really suited to a region-based approach since compact regions are not singled out; hence, it cannot be used for application-driven rate allocation, dynamic coding, or data mining.
- 2) Most regions are too small and have busy contours: to be efficient, the transform-based region coding algorithms (to be discussed later on) must operate on reasonably large and compact regions.
- 3) The encoding cost of such a detailed map is not negligible, and at low bit rates can become a significant fraction of the overall cost.

In summary, a map like this cannot really be used for efficient region-based coding,³ and, therefore, we turn to contextual segmentation techniques, where the label of a pixel depends not only on its spectral signature but also on the neighboring pixels, and, hence, more compact regions are singled out. More specifically, we use a Bayesian segmentation algorithm [25] based on the tree-structured Markov random field (TS-MRF) model [26].

The maximum *a posteriori* probability (MAP) Bayesian approach to segmentation amounts to selecting the map \hat{x} which has the maximum probability of occurrence given the observed data y

$$\hat{x} = \arg \max_x p(x|y) = \arg \max_x p(y|x)p(x).$$

Note that the optimum map depends not only on data likelihood $p(y|x)$ that is the spectral information, but also on the prior term $p(x)$ which takes into account some reasonable constraints on the spatial regularity of the map itself.

If the prior term is discarded, one implicitly assumes that all maps are equally likely, thus neglecting all spatial dependencies, and trusting only the spectral data, no matter how noisy they are. This amounts to using a maximum likelihood rule

$$\hat{x} = \arg \max_x p(y|x)$$

or, with further and commonly accepted simplifications, a simple minimum distance clustering, implemented for example by VQ. This kind of segmentation is adopted frequently in order to simplify matters under both theoretical and computational points of view, but results are not very accurate.

To improve the quality of segmentation, the prior term should be taken into account through a meaningful, yet manageable, probabilistic model. This problem has been solved by the theory of random Markov fields (MRFs) [27], [28], which allows one to model the map, or label field, in a reasonably simple way and yet take into account spatial dependencies through a few parameters that specify local interactions among pixels. By selecting a suitable model for x , one can design a segmentation algorithm which is able to provide more regular maps and to reduce the annoying effects of random noise. In this paper, we use the algorithm developed in [25], based on a tree-structured MRF

³In this case, one could resort to class-based coding instead [22], [23], where all pixels of the same class (color) are considered to form an “object.” This is also an interesting approach in terms of rate-distortion performance but does not preserve the high-level advantages of region-based coding.

(TSMRF) model,⁴ which, for our example image, provides the segmentation map shown in Fig. 2(c). As a result, much fewer regions are now present, and they are more compact and present more regular boundaries.

Even so, some isolated points and small fragments remain, typically in correspondence of mixed cells near region boundaries where two land covers are present at the same time. These small regions, however, can be merged with any of the neighboring large regions without altering significantly the statistical homogeneity of the region but improving the readability and compactness of the segmentation. Therefore, we proceed to the elimination of small regions (merged with the dominant neighboring region) up to a size of N_{\max} pixels. The final result for our example is shown in Fig. 2(d). At this point, much fewer regions remain, which are large, smooth, and compact enough to serve in high-level applications such as data mining, and are also statistically homogeneous, so as to guarantee a good compression efficiency. In addition, a map like this can be encoded with very little cost.

B. Map Coding

The additional coding cost associated with the segmentation map is one of the drawbacks of region-based coding, so it must be carefully assessed, and an efficient map coding algorithm must be used. It must be observed, however, that this cost is quite limited in our coding scheme, both because the segmentation algorithm provides smooth maps, amenable to high compression, and because a single map is used for all image bands.

The literature on this topic is rather scant, also because object-based video coding focuses on foreground/background segmentation and, hence, on bi-level maps (for which a large number of tools exist). However, we can draw from techniques designed for the compression of palette images which share many characteristics with our segmentation maps.

A popular approach to the lossless coding of palette images has been to adapt some universal techniques, like GIF Compress or PNG Deflate, to a 2-D source, with results that are arguably far from optimal. Also, the adaptation of a truly 2-D technique, such as JPEG-LS, proposed in [29], turns out to be ineffective, because the indexes which form a palette image (or our maps) do not exhibit any native correlation to be exploited in linear predictive coding.

Better results are provided by *ad hoc* techniques, such as the PPM2D [30], based on prediction by partial matching, or the PWC [31] (piece-wise constant image model), based on a binary arithmetic coder with multiple contexts that exploits the simplicity of palette images for coding separately boundaries and colors. Here, we use the RAPP (runs of adaptive pixel patterns) algorithm [32], a simple and effective predictive coding technique in which the prediction is always one of the four pixels in the causal neighbors, chosen so as to form the most probable pattern. In particular, we use a variant of RAPP including the “skip innovation,” originally introduced by Ausbeck for PWC, which allows one to encode uniform rows with a single bit, and almost uniform rows with just a few bits.

⁴A detailed description of the TSMRF-based algorithm is out of the scope of this paper, and we refer the reader to the literature for more information.

C. Spectral and Spatial Transforms

In our coding scheme, we decided to compress the regions by means of transform coding, because this is a well understood approach, with good performances in all application fields, and because there are plenty of coding tools available. By no means, however, do we write off the use of different approaches, like, for example, polynomial approximation for smooth areas, or model-based parametric description of textured areas. As a matter of fact, the freedom to use different techniques for different regions [24] is an important advantage of region-based coding, and we want to explore its potential in future research.

In accordance with the different nature of spectral and spatial dependencies, we use a 1-D spectral transform first, and then a further 2-D (shape-adaptive) transform in the spatial domain.

As for the spatial transform, we turn to the shape-adaptive wavelet transform (SA-WT) algorithm proposed by Li and Li [33], which was also adopted with some variations in the MPEG-4 standard. In fact, it allows one to retain most of the advantages of DWT while dealing with regions of arbitrary shape: it treats whole regions at once, unlike the block-based shape-adaptive DCT, guarantees a good compaction ability, and preserves the spatial relations among transform coefficients, a properties that can be exploited by zerotree-based encoding algorithms. On the down side, the SA-WT is markedly nonorthogonal, which renders very difficult to assess its actual compaction ability, but through an experimental analysis [34].

Unlike for the spatial transform, where SA-WT seems a natural choice at this moment, several alternatives appear equally appealing for the spectral transform. First of all, one can resort again to WT, which is computationally light, does not require side information, and lends itself to implement simple shape-adaptive variations of well-known zerotree coding algorithms. On the opposite side stands the KLT which, being data dependent, allows one to adapt the transform to the data to be encoded, but requires significant computation and calls for the transmission of some side information. DCT is a third option, but we did not experiment with it since its properties lie halfway between the former two.

Therefore, we implemented various coding schemes, based both on WT and KLT as a spectral transform, and assessed their performance on several test images, as reported in Section III. Our main focus, however, is on KLT-based solutions because they fit much better the region-based approach.⁵ In fact, there are at least three possible ways to use the KLT in this context:

- 1) use a single transform for the whole image;
- 2) use a different transform for each region;
- 3) use a different transform for each class.

In the first case, KLT is used almost like a fixed transform, but for the fact that it is adapted to the image statistics, and it takes little advantage of image segmentation. On the contrary, using a different KLT for each region is fully in the spirit of the region-based approach, since the transform is now adapted to the *local* statistics of the region, which might be markedly different from those of other regions. As a consequence, one obtains a better energy compaction, but also a heavier computational

⁵Also, looking ahead at the results, because they provide consistently superior performance.

burden and increased side information, since the transform matrices must be computed and sent for each region together with coded data. The third option is a compromise between the first two, which exploits the fact that our segmentation algorithm works on classes, characterized by similar spectral characteristics, and only in the final stage it extracts the elementary regions. Therefore, we can use a single KLT matrix for all regions that belong to the same class, with an obvious advantage in terms of computation and side information, but also with a good adaptation to the local statistics. Of course, the actual merits of these solutions must be assessed via numerical experiments.

To summarize, all coding schemes implemented will use separate transforms in the spectral and spatial domain, the former will be either a WT, a global KLT, a class-adaptive KLT, or a region-adaptive KLT, while the latter will always be a shape-adaptive WT. Of course, in the reference flat coding schemes used for comparison purposes only conventional (that is, nonshape adaptive) transforms will be used.

D. Region Coding and Shape-Adaptive SPIHT

After spectral and spatial transforms, the energy of each region is compacted in a relatively small number of coefficients that must be eventually quantized and coded. In recent years, several algorithms have been proposed for coding SA-WT coefficients. All of them are based on Li and Li's or similar [35] SA-WT techniques, and all belong to the class of embedded bit-plane coders, but differ in how they approach the compression of significance maps for each plane. Some of them, like the EZT algorithm adopted in the MPEG-4 standard, are suitable variations of the well-known EZW [36] and SPIHT [37] algorithms and rely on the zerotree concept. Others use a quad-tree or k-d tree-based approach [38], [39], or even resort to run-length coding [40]. A comparison among these algorithms, with reference to simple foreground/background scenes, is reported in [39], based on the Qcc library [41] freely available at [42].

Here, we will consider a straightforward extension of SPIHT to encode 2-D or 3-D regions of arbitrary shapes. SPIHT (set partitioning in hierarchical trees) is a well-known zerotree-based algorithm for the progressive quantization of wavelet transform coefficients. It is simple, intrinsically scalable, and very efficient, which is why it is one of the most popular techniques for the compression of images and, more recently, video. Especially important for our needs, it can be readily modified to encode images of arbitrary geometry after a shape-adaptive WT. In our own implementation [43] (similar to that formerly proposed in [44] and further refined in [45]), we introduce only two major changes with respect to the basic algorithm. First, only active nodes, that is nodes belonging to the support of the SA-WT of the object, are considered while scanning a spatial orientation tree. This information is available at the decoder thanks to the segmentation map. The second modification concerns the baseband, where coefficients are not grouped anymore in 2×2 squares, as in the original SPIHT, since these might not belong entirely to the region of interest and a single root is considered instead. Further modifications can be considered when 3-D regions are coded, since the 3-D trees can be defined in various ways [46] depending on how the spectral transform is performed.

E. Rate Allocation and Choice of Objects

Rate allocation, either explicit or implicit (like in bit-plane coders), is a critical part of any transform coding technique. When dealing with region-based coding, it becomes even more important, since there are now different sources, possibly with different semantic value or interest for the user, among which resources must be optimally allocated to reach the desired goal. However, we do not consider here the case in which some regions are intrinsically more important than others: in such a situation, in fact, there is no doubt that region-based coding would outperform flat coding, since it allows one to devote more resources to the only regions of interest to the users. On the contrary, we want to show that the region-based approach can be convenient in a rate-distortion sense, even for the compression of the whole image. In this case, rate allocation has the goal of obtaining the least possible overall distortion for the assigned resources. By formulating the problem as a Lagrangian optimization, it can be shown that an optimal allocation is reached when each new encoding bit produces the same decrease in distortion on any object. More formally, if $D_i(R_i)$ is the curve that gives distortion as a function of rate for the i th object, we want to operate at rates R_i such that

$$\frac{\Delta D_i}{\Delta R_i} = \lambda, \quad \text{and} \quad \sum_i R_i = R$$

that is, the slope of the rate-distortion curve, is the same for all objects,⁶ while all resources are consumed.

To this end, we implement a postcompression resource allocation algorithm similar to that of JPEG-2000 [47], that is, we first compute the rate-distortion (RD) curve of each object by encoding it at high rate, and then deallocate resources progressively, acting each time on the object with the lowest RD slope so as to keep an approximate equi-slope condition while reaching the desired overall coding rate. Using the SPIHT algorithm, which provides a naturally embedded coding, it is quite easy to develop one such algorithm. Some caution is necessary, however, because the SA-WT is not a unitary transform, and, hence, a strictly optimal allocation cannot be obtained by working only in the transform domain.

Note that in this paragraph we made always reference to *objects* rather than *regions*. This is because all previous considerations hold for any set of objects, whatever their nature, dimensionality and spatial shape, and this generality leaves open a wide range of different options. In particular, in the spatial domain we can work with the whole image, or the elementary regions, while in the spectral domain we can work with the whole set of bands or with each single band individually, obtaining the following cases.

- 1) Image: All transform coefficients are treated as a single object, and bit allocation is left to the SPIHT coder.
- 2) Image bands: Each transform band of the whole image is treated as a separate object, and therefore bit allocation among bands is explicit.
- 3) Regions: All transform coefficients of a region are treated as a single object, bit allocation is explicit among different

⁶As is well known, this is not always possible in practice, especially at low rates.




Fig. 3. Design alternatives within the general coding scheme.




Fig. 4. Test images (in false colors): (a) Landsat TM; (b) AVIRIS.

regions, but implicit (left to the SPIHT coder) among the transform bands of the same region.

- 4) Region bands: This is the most detailed bit allocation, each transform band of each region is treated as a separate object, with explicit bit allocation.

The first two cases are very simple, and appropriate for flat coding algorithms, with or without explicit bit allocation among the bands, while the others are appropriate for region-based coding, with or without explicit spectral bit allocation, and call for a more complex management of the data stream.

III. EXPERIMENTAL ANALYSIS

In Section II, we have considered several alternatives for each step of the general coding scheme, the most relevant of which are synthetically reported in Fig. 3.

For example, by selecting each time the first option we obtain 3-D SPIHT, originally proposed in [48], a coding scheme with no segmentation, WT both in space and along the bands, with a 3-D version of SPIHT used to encode the whole set of transform coefficients, with no explicit form of rate allocation. On the contrary, by selecting always the last option we obtain a new region-based coding scheme, where the image is segmented

by the TSMRF-based algorithm (not shown, the map is coded by the RAPP algorithm) there is a different spectral KLT transform for each region, and SA-WT is used in the spatial domain, followed by SA-SPIHT encoding with an explicit rate allocation for all bands of all regions. In the middle, there is a large number of viable combinations, most of which have been actually implemented and tested experimentally in order to single out the most interesting schemes. In the following, we present and comment the most significant results.

Experiments have been conducted on several images with different characteristics. Here, we will present results for two test images:

- a Landsat TM image of a region near Lisbon: six bands, 512×512 pixels, 8 bit/sample;
- the AVIRIS “Low Altitude” image available at [49]: 32 bands⁷ 512×512 pixels, 16 bit/sample;

which differ significantly in terms of spatial and spectral resolution. In Fig. 4, we show both images in a false color representation.

Most experimental results will be in terms of rate-distortion (RD) curves, where rates are in bit/sample (bps), that is, number of coding bits divided by number of image rows, columns and bands, and distortion are measured by the mean-square-error, although the curves report the more widespread signal-to-noise ratio (SNR), defined as the ratio between image variance and error variance, in decibels. It is important to underline that we normalize all bands to zero mean and unit variance before encoding. This preprocessing is especially important when the bands have wildly different dynamics (often the case with hyperspectral images) since it guarantees that all bands receive a fair amount of coding resources. Without the normalization, the rate allocation procedure gives most resources to high-power bands, neglecting the remaining ones, often equally informative, thereby maximizing the SNR but also impairing the overall image quality.

A. Rate Allocation

A first set of experiments concerns the use of explicit versus implicit rate allocation (RA) among the bands, and, therefore, the use of 3-D or 2-D SPIHT. To investigate this point, we implemented two quite different coding schemes, one completely flat (no segmentation, spectral KLT, spatial WT), and the other region-based (TSMRF segmentation, region-adaptive spectral KLT, spatial SA-WT). Then, for both schemes, we considered two different options for the rate allocation by using either 3-D SPIHT (hence, implicit RA) or 2-D SPIHT on each band with explicit RA for each band.

Results in terms of RD curves are reported in Fig. 5 for the Landsat TM image and Fig. 6 for the AVIRIS image. It clearly appears that the explicit rate allocation guarantees a consistent performance improvement in all cases, more than 1 dB for the TM image and between 0.5 and 1 dB for the AVIRIS image, for the medium-high rates that are of actual interest in the applications. Since the same behavior is observed with all other coding

⁷The original image comprises 224 bands, but we use a more manageable 32-band datacube obtained through regular band sampling. RD curves; therefore, we do not refer to the original image, even though all coding appear to behave in a similar way on both datacubes.




Fig. 5. RD performance for the Landsat TM image with implicit and explicit rate allocation.




Fig. 6. RD performance for the AVIRIS image with implicit and explicit rate allocation.

schemes, we will always use explicit RA from now on. It goes by itself that this requires some side information, but its cost turns out to be always negligible.

B. Spectral Transform

Another unambiguous result regards the use of WT as spectral transform. As before, we implemented various coding schemes in which, once fixed the other blocks, the spectral transform was either WT or KLT. We consider only the AVIRIS image now because WT would be intrinsically inefficient on the TM image, which comprises only six bands. In Fig. 7, we report the RD curves obtained with a completely flat coder (no segmentation, spectral KLT or WT, spatial WT, SPIHT) and a region-based one (TSMRF segmentation, spectral KLT or WT, spatial SA-WT followed by SA-SPIHT). In both cases, the coder-based using KLT as a spectral transform provides a huge gain, about 5 dB, w.r.t. the corresponding one based on WT as a spectral transform. This behavior was observed in all our experiments (with an even larger gap for the Landsat TM image, as expected) and is also in agreement with recent findings in the literature [10], [11], [50] where gains of similar entities have been observed when KLT replaces WT. Therefore, we will not consider WT anymore as a possible spectral transform.

At this point, it is clear that KLT is the spectral transform of choice, so we go on to investigate the impact of a global versus

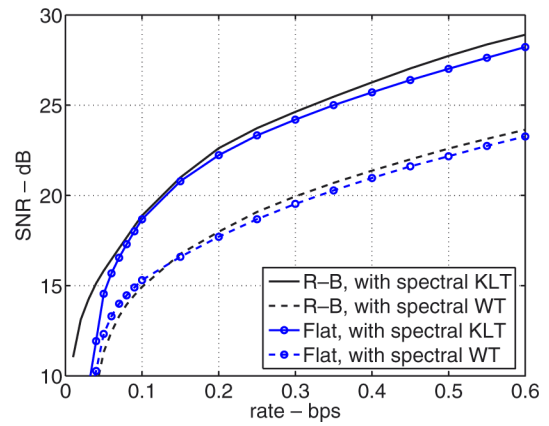


Fig. 7. RD performance for the AVIRIS image with KLT and WT as spectral transform.

region-adaptive KLT on the overall performance. We consider only region-based coding schemes for the time being, in which the image is first segmented, then each object undergoes spectral KLT, spatial SA-WT, and is finally coded by SA-SPIHT with explicit rate allocation. Therefore, the various coding schemes differ only in the spectral transform, which can be global, region-adaptive or class-adaptive, with this latter option considered in order to limit the cost of side information, since the same KLT matrix is used for all regions of the same class. The segmentation maps used in this experiment are shown in Fig. 8 (top), while Figs. 9 and 10 report the rate-distortion curves for the Landsat TM and the AVIRIS image respectively. Results show unambiguously that the use of a KLT adapted to the region statistics guarantees a significant performance gain over a global KLT, despite the increased cost of the side information due to the need to transmit the segmentation map, and the region means and covariances. Turning to the comparison between class-based and region-based adaptivity, it is clear that the side information, especially for the covariance matrices, has a significant weight at low bit-rates, speaking in favor of the class-based solution. At higher bit-rates the gap reduces and even vanishes for the TM image, while it remains significant, about 0.7 dB, for the 32-band AVIRIS image, where sending the covariance matrices for a large number of regions represents a real burden. A suitable compression of such matrices (we use simple nonuniform quantization) could very likely reduce this overhead, but the class-based solution seems preferable at this point.

C. Flat Versus Region-Based Coding

We are now ready to compare the best region-based technique, using class-based KLT, with the best flat reference technique, using global KLT. In both cases, rate allocation is carried out explicitly, band-by-band. Results are reported in Fig. 11 for the TM image and Fig. 12 for the AVIRIS image, and show a consistent lead of the region-based technique, about 1 dB in the first case and almost as much in the second, except for the very low bit-rates where the cost of side information plays a non-negligible role. In the figures, we also report the results obtained by two “standard” reference techniques, that is JPEG-2000 multi-component [51], with KLT as a decorrelating spectral transform,




Fig. 8. Segmentation maps used in the experiments: (left) Landsat TM; (right) AVIRIS.




Fig. 9. RD performance for the Landsat TM image with global KLT, class-based KLT, and region-based KLT.




Fig. 10. RD performance for the AVIRIS image with global KLT, class-based KLT, and region-based KLT.

and simple 3-D SPIHT. It appears that JPEG-2000 performs about as well as our best region-based technique, although it is

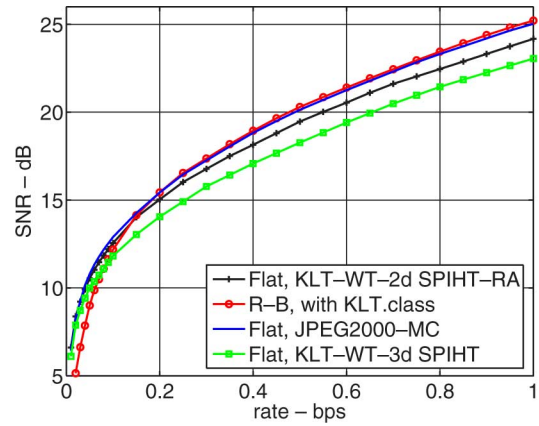


Fig. 11. RD performance comparison between best region-based technique and reference flat techniques for the Landsat TM image.




Fig. 12. RD performance comparison between best region-based technique and reference flat techniques for the AVIRIS image.

worth underlining that we use no arithmetic coding at present,⁸ while 3-D SPIHT is clearly inferior, especially when the spectral WT is used (AVIRIS image). Other “hybrid” reference techniques have also been implemented and tested, e.g., class-adaptive KLT followed by flat WT and SPIHT, but their performance turned out not to be competitive and are not reported for clarity. Although these are only partial results, it seems safe to say, based on this evidence, that region-based coding has the potential for an excellent rate-distortion performance, competitive with the best (and carefully engineered) flat techniques, which comes on top of the more obvious advantages for high-level applications already described in the Introduction. It is also obvious that such gains come at a computational cost (segmentation, computation of KLT matrices, multiple encodings, rate allocation) but there are many applications where such a cost is fully affordable in view of the ensuing benefits.

We have accumulated enough evidence, by now, about the rate-distortion potential of region-based coding. All our experiments, however, have been carried out with a single segmentation map for each image, and, hence, it would be interesting to explore the impact of segmentation on the overall performance.

⁸Actually, we already implemented a version of this coder with AC, which provides an improvement of about 0.5 dB at low and medium rates, but present results without AC to guarantee a meaningful performance comparison among all proposed techniques.

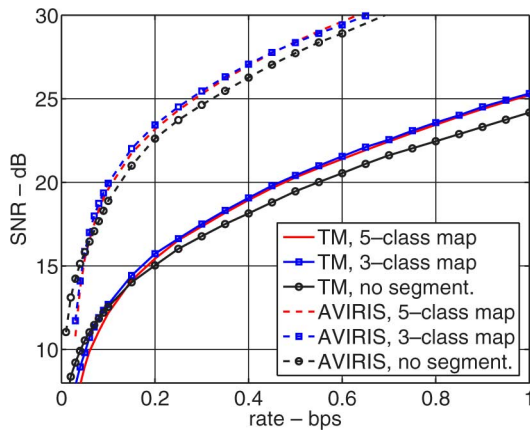


Fig. 13. RD performance using detailed (five classes) or rough (three classes) segmentation maps.

A good segmentation, valuable for subsequent higher-level applications, typically unknown in advance, aims at providing regions that are both spatially smooth and spectrally homogeneous, but these goals are typically conflicting, since spectral homogeneity leads to very fragmented maps, while smooth regions tend to lack spectral purity. By adjusting some segmentation parameters one can bend towards one extreme or the other, but what is the ultimate impact on the RD performance of region-based coders? A thorough analysis of this topic would be exceedingly complex and out of the scope of this paper.⁹ To shed some light on this point, however, we evaluate the performance of our region-based coder with class-based KLT using two new maps (shown in Fig. 8, bottom) characterized by a very small number of classes and just a few large regions. Despite the limited scope of this experiment, results are quite telling: the RD curves, all reported in Fig. 13 to save space, are almost indistinguishable, with a slight advantage for the more detailed maps at medium-to-high rates, and a more marked gain for the simpler maps at low rates. Such results suggest that the proposed coder is quite robust w.r.t. the segmentation map provided it is reasonably accurate and smooth (isolated points and small regions are erased in all cases).

D. Classification Results

Up to this point, we have used mean-square-error (and related SNR) to measure quality since it is widespread, objective, and usually in good agreement with perceived quality. In the remote-sensing community, however, a large SNR is not seen a sufficient guarantee about the diagnostic value of images after compression. Therefore, to strengthen our point, we assess again the techniques considered above using the classification error as distortion measure. The images are compressed at various rates, with the proposed and reference algorithms, and classified using a simple N-class minimum-distance clustering; then the misclassification rate is computed, that is, the fraction of erroneously labeled pixels w.r.t. a reference classification. Lacking ground truth data, the reference is taken as the segmentation of the original uncompressed image, even though

⁹We are currently working on RD-driven segmentation, though.

TABLE I
LANDSAT TM, PERCENT ERROR WITH 4-CLASS AND 8 CLASS K-MEANS

rate	4 classes			8 classes		
	Flat coder	JPEG2000	R-B coder	Flat coder	JPEG2000	R-B coder
0.2	8.77	8.57	7.69	18.31	18.18	16.08
0.4	5.32	5.36	4.70	11.15	11.20	10.10
0.6	3.86	3.60	4.03	8.04	7.48	8.42
0.8	2.82	2.69	2.91	5.87	5.57	6.27
1.0	2.19	2.11	1.49	4.57	4.33	3.00

TABLE II
AVIRIS, PERCENT ERROR WITH 4-CLASS AND 8-CLASS K-MEANS

rate	4 classes			8 classes		
	Flat coder	JPEG2000	R-B coder	Flat coder	JPEG2000	R-B coder
0.1	1.02	0.93	0.95	10.90	9.51	7.00
0.2	0.48	0.45	0.38	4.19	4.13	3.27
0.3	0.31	0.32	0.32	2.66	2.70	3.06
0.4	0.25	0.23	0.23	2.01	1.84	2.81
0.5	0.21	0.21	0.18	1.86	1.66	1.37

this introduces a small pessimistic bias because of the disturbances which affect the reference itself. Results with 4-class and 8-class clustering are reported in Table I (Landsat TM) and Table II (AVIRIS), and show a generally good performance for all coding schemes already at medium bit-rates, showing that compression, at an appropriate coding rate, can be a viable option for many remote-sensing applications. The region-based coder provides often, although not always, the best results, confirming its performance advantage. Of course, the error is larger for the Landsat image, where less spectral information is available, and grows with the number of classes. This latter phenomenon is especially evident with the AVIRIS image, since eight classes are more than the actual land covers, and, hence, two of them, hardly distinguishable from one another, are associated with the same woods region.

To get rid of the bias induced by the unreliable reference map, we repeated the experiments after deleting from the sample all pixels that, in the original uncompressed image, are “too close to call,” i.e., such that $d_2/d_1 < (1+\gamma)$ where $d_1[d_2]$ is the distance in the feature space between the pixel and the closest [second closest] template vector, and γ is a free threshold. By setting $\gamma = 0.1$, about 5% and 1% of the pixels are not classified in the Landsat and AVIRIS image respectively, but percent errors drop dramatically: for the region-based coder, they decrease from 1.49 to 0.16 (four classes) and from 3.00 to 0.56 (eight classes) for the Landsat TM image at 1 bps, and from 0.18 to 0.003 (four classes) and from 1.37 to 0.42 (eight classes) for the AVIRIS image at 0.5 bps. In other words, except for pixels that have already a dubious meaning in the original image, often mixtures of several land covers, classification of compressed images is extremely reliable.

E. Visual Inspection

We conclude this analysis by presenting some visual results which, though not conclusive under any respect, help gaining insight about the quality of the compressed images. In Fig. 14, we show (a) the Landsat TM image and (b) the AVIRIS image, in the same false color representation as Fig. 4, after compression with the best region-based coder at 0.6 and 0.3 bps, respectively.




Fig. 14. Landsat TM image compressed at (a) 0.6 bps and (b) AVIRIS image compressed at 0.3 bps.

Despite the large compression ratios, exceeding 10:1 and 50:1, respectively, no significant impairment is visible.

IV. CONCLUSION AND FUTURE WORK

In this paper, we have proposed a new region-based coding scheme for multispectral images with a performance that is aligned or even superior to that of state-of-the-art conventional techniques. This result is certainly relevant for the applications, but even more because it shows that a region-based description of such images, besides enabling a number of high-level applications, can be effective in a rate-distortion sense. Another valuable result of this paper is the wealth of experimental evidence gathered about some key design choices, such as the type of spectral transform, or the rate allocation strategy.

Although the selected scheme provides the best performance, and appears to be quite robust on the class of images of interest, several further improvements are certainly possible. First of all, the connection between segmentation and coding, which has been only briefly touched upon in this paper, is definitely worth exploring. An interesting option, that we are bound to study in future research, is the use of an RD-driven segmentation technique, with a split-and-merge procedure where each step is guided by the actual RD performance gain/loss measured on the image to be coded. Also, segmentation should take into explicit account the problem of mixed cells, especially relevant for low-resolution images, so as to obtain regions that are more

statistically homogeneous. Another topic that deserves attention is the use of a dynamic coding approach, where the different regions of the image are coded by means of different techniques, based on their statistics. For example, one could consider polynomial approximation for very smooth areas, and techniques based on texture analysis and synthesis for textured areas, which are quite common in remote-sensing images.

In conclusion, region-based coding seems, under several points of view, the most appropriate approach to the compression of multispectral images. This paper, by proposing an actual coding scheme with competitive performance, provides some support to this thesis, but we believe that much of the potential of this approach is yet to be exploited.

REFERENCES

- [1] R. L. Baker and Y. T. Tse, "Compression of high spectral resolution imagery," in *Proc. SPIE*, 1988, no. 974, pp. 255–264.
- [2] R. N. Hoffman and D. W. Johnson, "Application of EOFs to multispectral imagery: Data compression and noise detection for AVIRIS," *IEEE Trans. Geosci. Remote Sens.*, vol. 32, no. 1, pp. 34–43, Jan. 1994.
- [3] G. P. Aboussleman, M. W. Marcellin, and B. R. Hunt, "Compression of hyperspectral imagery using the 3-D DCT and hybrid DPCM/DCT," *IEEE Trans. Geosci. Remote Sens.*, vol. 33, no. 1, pp. 26–34, Jan. 1995.
- [4] F. Amato, C. Galdi, and G. Poggi, "Embedded zerotree wavelet coding of multispectral images," in *Proc. IEEE Int. Conf. Image Processing*, Santa Barbara, CA, Oct. 1997, pp. 612–615.
- [5] J. Vaisey, M. Barlaud, and M. Antonini, "Multispectral image coding using lattice VQ and the wavelet transform," in *Proc. IEEE Int. Conf. Image Processing*, Chicago, IL, Oct. 1998, pp. 307–311.
- [6] P. L. Dragotti, G. Poggi, and A. R. P. Ragozini, "Compression of multispectral images by three-dimensional SPIHT algorithm," *IEEE Trans. Geosci. Remote Sens.*, vol. 38, no. 1, pp. 416–428, Jan. 2000.
- [7] J. E. Fowler and D. N. Fox, "Embedded wavelet-based coding of three-dimensional oceanographic images with land masses," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 1, pp. 284–290, Feb. 2001.
- [8] Y. Wang, J. T. Rucker, and J. E. Fowler, "Three-dimensional tarp coding for the compression of hyperspectral images," *IEEE Geosci. Remote Sens. Lett.*, vol. 42, no. 4, pp. 136–140, Apr. 2004.
- [9] B. Penna, T. Tillo, E. Magli, and G. Olmo, "Progressive 3-D coding of hyperspectral images based on JPEG 2000," *IEEE Geosci. Remote Sens. Lett.*, vol. 44, no. 1, pp. 125–129, Jan. 2006.
- [10] Q. Du and J. E. Fowler, "Hyperspectral image compression using JPEG2000 and principal component analysis," *IEEE Geosci. Remote Sens. Lett.*, vol. 45, no. 4, pp. 201–205, Apr. 2007.
- [11] B. Penna, T. Tillo, E. Magli, and G. Olmo, "Transform coding techniques for lossy hyperspectral data compression," *IEEE Geosci. Remote Sens. Lett.*, vol. 45, no. 5, pp. 1408–1421, May 2007.
- [12] S. Gupta and A. Gersho, "Feature predictive vector quantization of multispectral images," *IEEE Geosci. Remote Sens. Lett.*, vol. 30, no. 5, pp. 491–501, May 1992.
- [13] S. E. Qian, A. B. Hollinger, D. Williams, and D. Manak, "Fast 3D data compression of hyperspectral imagery using vector quantization with spectral-feature-based binary coding," *Opt. Eng.*, pp. 3242–3249, Nov. 1996.
- [14] G. R. Canta and G. Poggi, "Kronecker-product gain-shape vector quantization for multispectral and hyperspectral image coding," *IEEE Trans. Image Process.*, vol. 7, no. 5, pp. 668–678, May 1998.
- [15] S. E. Qian, "Hyperspectral data compression using a fast vector quantization algorithm," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 8, pp. 1791–1798, Aug. 2004.
- [16] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. Boston, MA: Kluwer, 1992.
- [17] E. Reusens, T. Ebrahimi, and M. Kunt, "Dynamic coding of visual information," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 7, no. 6, pp. 489–500, Jun. 1997.
- [18] J. A. Saghri, A. G. Tescher, and J. T. Reagan, "Practical transform coding of multispectral imagery," *IEEE Signal Process. Mag.*, vol. 12, no. 1, pp. 32–43, Jan. 1995.
- [19] J. Lee, "Optimized quadtree for Karhunen–Loève transform in multispectral image coding," *IEEE Trans. Image Process.*, vol. 8, no. 4, pp. 453–461, Apr. 1999.

- [20] M. Finelli, G. Gelli, and G. Poggi, "Multispectral image coding by spectral classification," in *Proc. IEEE Int. Conf. Image Process.*, Lausanne, Switzerland, Sep. 1996, vol. 2, pp. 605–608.
- [21] G. Fernandez and C. M. Wittenbrink, "Coding of spectrally homogeneous regions in multispectral image compression," in *Proc. IEEE Int. Conf. Image Processing*, Lausanne, Switzerland, Sep. 1996, vol. 2, pp. 923–926.
- [22] G. Gelli and G. Poggi, "Compression of multispectral images by spectral classification and transform coding," *IEEE Trans. Image Process.*, vol. 8, no. 4, pp. 476–489, Apr. 1999.
- [23] M. Cagnazzo, S. Parrilli, G. Poggi, and L. Verdoliva, "Improved class-based coding of multispectral images with shape-adaptive wavelet transform," *IEEE Geosci. Remote Sens. Lett.*, vol. 4, no. 4, pp. 566–570, Oct. 2007.
- [24] M. Petrou, P. Hou, S. Kamata, and C. I. Underwood, "Region-based image coding with multiple algorithms," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 3, pp. 562–570, Mar. 2001.
- [25] C. D'Elia, G. Poggi, and G. Scarpa, "A tree-structured markov random field model for bayesian image segmentation," *IEEE Trans. Image Process.*, vol. 12, no. 10, pp. 1259–1273, Oct. 2003.
- [26] G. Poggi and A. R. P. Ragozini, "Image segmentation by tree-structured Markov random fields," *IEEE Signal Process. Lett.*, vol. 6, no. 7, pp. 155–157, Jul. 1999.
- [27] S. Geman and D. Geman, "Stochastic relaxation, gibbs distributions, and the bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-6, no. 11, pp. 721–741, Nov. 1984.
- [28] S. Z. Li, *Markov Random Field Modeling in Image Analysis*. New York: Springer-Verlag, 2001.
- [29] N. D. Memon and A. Venkateswaran, "On ordering color maps for predictive coding," *IEEE Trans. Image Process.*, vol. 5, no. 11, pp. 1522–1527, Nov. 1996.
- [30] S. Forchhammer and J. M. Salinas, "Progressive coding of palette images and digital maps," in *Proc. Data Compression Conf.*, Dec. 2002, pp. 362–371.
- [31] P. J. Ausbeck, Jr., "The piecewise-constant image model," *Proc. IEEE*, vol. 88, no. 11, pp. 1779–1789, Nov. 2000.
- [32] V. Ratnakar, "RAPP: Lossless image compression with runs of adaptive pixel patterns," in *Proc. 32nd Asilomar Conf. Signals, System, Computers*, Nov. 1998, pp. 1251–1255.
- [33] S. Li and W. Li, "Shape-adaptive discrete wavelet transforms for arbitrarily shaped visual object coding," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 10, no. 8, pp. 725–743, Aug. 2000.
- [34] M. Cagnazzo, S. Parrilli, G. Poggi, and L. Verdoliva, "Costs and advantages of object-based image coding with shape-adaptive wavelet transform," *Int. J. Image Video Process.*, vol. 2007, p. 13, 2007, Article ID 78323.
- [35] J. E. Fowler and D. N. Fox, "Wavelet-based coding of three-dimensional oceanographic images around land masses," in *Proc. IEEE Int. Conf. Image Processing*, Vancouver, BC, Canada, Sep. 2000, pp. 431–434.
- [36] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3445–3462, Dec. 1993.
- [37] A. Said and W. A. Pearlman, "A new fast and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 6, no. 6, pp. 243–250, Jun. 1996.
- [38] Z. Lu and W. A. Pearlman, "Wavelet video coding of video object by object-based SPECK algorithm," in *Proc. Picture Coding Symp.*, Seoul, Korea, Apr. 2001, pp. 413–416.
- [39] J. E. Fowler, "Shape adaptive coding using binary set splitting with K-d trees," in *Proc. IEEE Int. Conf. Image Processing*, Singapore, Oct. 2004, pp. 1301–1304.
- [40] J. Tian and R. Wells, Jr., "Embedded image coding using wavelet difference reduction," *Wavelet Image Video Compress.*, pp. 289–301, 1998.
- [41] J. E. Fowler, "QccPack: An open-source software library for quantization, compression, and coding," in *Proc. SPIE Applications of Digital Image Processing XXIII*, Aug. 2000, vol. 4115, pp. 294–301.
- [42] [Online]. Available: <http://www.qccpack.sourceforge.net/>
- [43] M. Cagnazzo, G. Poggi, L. Verdoliva, and A. Zinicola, "Region-oriented compression of multispectral images by shape-adaptive wavelet transform and SPIHT," in *Proc. IEEE Int. Conf. Image Processing*, Singapore, Oct. 2004, pp. 2459–2462.
- [44] A. Kawanaka and V. R. Algazi, "Zerotree coding of wavelet coefficients for image data on arbitrarily shaped support," in *Proc. Data Compression Conf.*, Mar. 1999, p. 534.
- [45] G. Minami, Z. Xiong, A. Wang, and S. Mehrotra, "3-D wavelet coding of video with arbitrary regions of support," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 11, no. 9, pp. 1063–1068, Sep. 2001.
- [46] C. He, J. Dong, Y. F. Zheng, and Z. Gao, "Optimal 3-D coefficient tree structure for the 3-D wavelet video coding," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 13, no. 10, pp. 961–972, Oct. 2003.
- [47] D. Taubman, "High performance scalable image compression with EBCOT," *IEEE Trans. Image Process.*, vol. 9, no. 7, pp. 1158–1170, Jul. 2000.
- [48] B. Kim and W. A. Pearlman, "An embedded wavelet video coder using three-dimensional set partitioning in hierarchical trees," in *Proc. Data Compression Conf.*, Mar. 1997, pp. 251–260.
- [49] [Online]. Available: <http://www.aviris.jpl.nasa.gov/html/aviris.free-data.html>
- [50] B. Penna, T. Tillo, E. Magli, and G. Olmo, "A new low complexity KLT for lossy hyperspectral data compression," in *Proc. IEEE Int. Geoscience and Remote Sensing Symp.*, Aug. 2006, pp. 3528–3528.
- [51] [Online]. Available: <http://www.kakadusoftware.com/>



Marco Cagnazzo received the Laurea (M.S.) degree in telecommunication engineering from the University of Napoli, Italy, in January 2002, and the Ph.D. degree in information and communication technologies jointly from the University of Napoli and the University of Nice-Sophia Antipolis, France, in March 2005.

He has obtained research grants from several scientific institutions in Italy and in France. He is currently a Postdoctorate Fellow at the I3S Laboratory, France. His research areas are scalable

and efficient video compression, shape-adaptive image compression, and low-complexity image and video coding.

Dr. Cagnazzo is Reviewer for the major international journals in his research field (the IEEE TRANSACTIONS ON IMAGE PROCESSING, the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, and the *International Journal of Image and Video Processing*).



Giovanni Poggi received the Laurea degree in electronic engineering from the University Federico II, Naples, Italy, in July 1988.

He has been with the Department of Electronic and Telecommunication Engineering of the University Federico II since 1990, where he is currently a Full Professor of telecommunications, Vice Chair of the Department, and Coordinator of the Ph.D. program. His current research activity is in statistical signal processing, and, in particular, compression, segmentation, and filtering of images and video,

with special attention to remote-sensing applications.

Dr. Poggi has been an Associate Editor for the IEEE TRANSACTIONS ON IMAGE PROCESSING since 2004.



Luisa Verdoliva received the Laurea degree in telecommunications engineering and the Ph.D. degree in information engineering from the University Federico II, Naples, Italy, October 1998 and March 2002, respectively.

She is a Researcher with the Department of Electronic and Telecommunication Engineering, University Federico II, where she teaches classes on signal theory and on wavelets and filter banks. Her current research activity is in image compression, and, in particular, vector quantization, wavelet-based

techniques, and shape-adaptive compression of images. She is a Reviewer for the major international journals in her research field.