

Image Coding Based on a Morphological Representation of Wavelet Data

Sergio D. Servetto, *Member, IEEE*, Kannan Ramchandran, *Member, IEEE*,
and Michael T. Orchard, *Senior Member, IEEE*

Abstract—In this paper, an experimental study of the statistical properties of wavelet coefficients of image data is presented, as well as the design of two different morphology-based image coding algorithms that make use of these statistics. A salient feature of the proposed methods is that, by a simple change of quantizers, the same basic algorithm yields high performance embedded or fixed rate coders. Another important feature is that the shape information of morphological sets used in this coder is encoded *implicitly* by the values of wavelet coefficients, thus avoiding the use of explicit and rate expensive shape descriptors. These proposed algorithms, while achieving nearly the same objective performance of state-of-the-art zerotree based methods, are able to produce reconstructions of a somewhat superior perceptual quality, due to a property of joint compression and noise reduction they exhibit.

Index Terms—Data compression, image coding, morphological operations, wavelet transforms.

I. INTRODUCTION

A. Statistical Modeling for Wavelet-Based Image Compression

SINCE their introduction as a tool for signal representation, wavelets have become increasingly popular within the image coding community, because of the potential gains they offer for the construction of efficient image coding algorithms. Those potential gains are due to the fact that wavelets provide a good tradeoff between resolution in the space and frequency domains, a feature which results in mapping typical space-domain image phenomena into *structured* sets of coefficients in the wavelet domain. However, to be able to make effective use of any structure for improving coding performance, an algorithm requires a statistical characterization of the joint

Manuscript received December 3, 1996; revised October 21, 1998. This work was supported in part by the Joint Services Electronics Program (JSEP) under Award N00014-96-1-0129 (JSEP), the Army Research Laboratory under Award DAAL01-96-2-0003F-3, the Army Research Office under the Young Investigator Award DAAH04-96-1-0342, and the National Science Foundation under Grants NSF MIP-9409587 (RIA-94) and NSF MIP 93-57823 NYI. This work was presented in part at ICASSP'95 and ICIP'95. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Roland Wilson.

S. D. Servetto is with the Beckman Institute and the Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: servetto@cs.uiuc.edu).

K. Ramchandran was with the Beckman Institute and the Electrical and Computer Engineering Department, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA. He is now with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720 USA (e-mail: kannanr@eecs.berkeley.edu).

M. T. Orchard is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08540 USA (e-mail: orchard@ee.princeton.edu).

Publisher Item Identifier S 1057-7149(99)06818-9.

distribution of wavelet coefficients, capable of taking that structure into account. Our first goal then is to identify properties of the wavelet coefficients of images leading to good coding performance.

Standard transform-based image coding algorithms consist of three stages: a linear-typically unitary-decorrelating transform, followed by a quantization stage, and final entropy coding of the quantized data. A good decorrelating transform will remove from the data *linear* dependencies, thus producing a set of coefficients such that, when scalar quantized, the entropy of the resulting symbol stream is reduced substantially more than if the same quantization were applied directly on raw image data. This is the basis for the success of transform coders [1], including early subband coders [2], [7], [18]. Initially, transform coding became popular mainly due to the introduction of the discrete cosine transform (DCT), an efficient approximation to the theoretically optimal (under the assumption of subband statistics known *a priori*) but highly complex Karhunen–Loeve transform (KLT) [1]. However, although wavelets share many properties with the DCT (e.g., decorrelation), taking this same approach ignores a fundamental wavelet property: the joint time-frequency localization.

Consider the following simple image model. For the coding application, typical images can be described as a set of smooth surfaces, delimited by edge discontinuities (see Fig. 1). Now, consider the effect of putting an edge through a wavelet filterbank (see Fig. 2). The abstract property of joint time-frequency localization has a very practical consequence: the formation of energy clusters in image subbands, at spatial locations associated with edges in the original image.

This clustering effect is attributed mostly to two basic facts: 1) signal energy due to smooth regions is compacted mostly into a few low frequency coefficients, thus resulting in negligible contributions to coefficients in the higher frequency bands; and 2) due to the small compact support of the wavelets used, edges can only contribute energy to a small number of coefficients in a neighborhood of their space domain location. This effect is illustrated by Fig. 3.

Now, in a data compression application, “useful” properties are those that can be described by some statistical model of the source being coded, so that a practical algorithm can exploit them to achieve better coding performance. Hence, the first problem addressed in this work is that of finding an improved statistical description of image subbands, capable of explicitly taking into account the space-frequency localization properties of the wavelet.

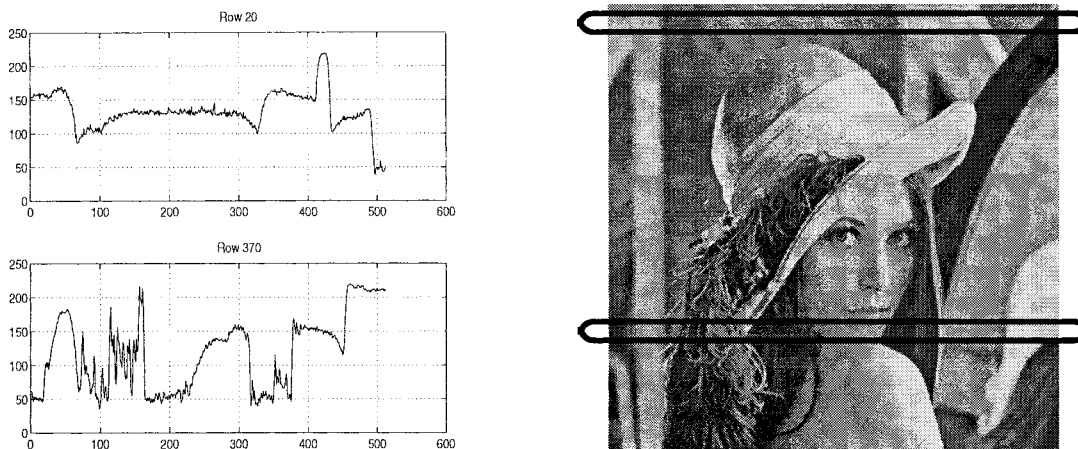


Fig. 1. To illustrate the assumed image structure in the spatial domain. Row 20 (up) is a mostly smooth region, while row 370 (down) contains some textures.

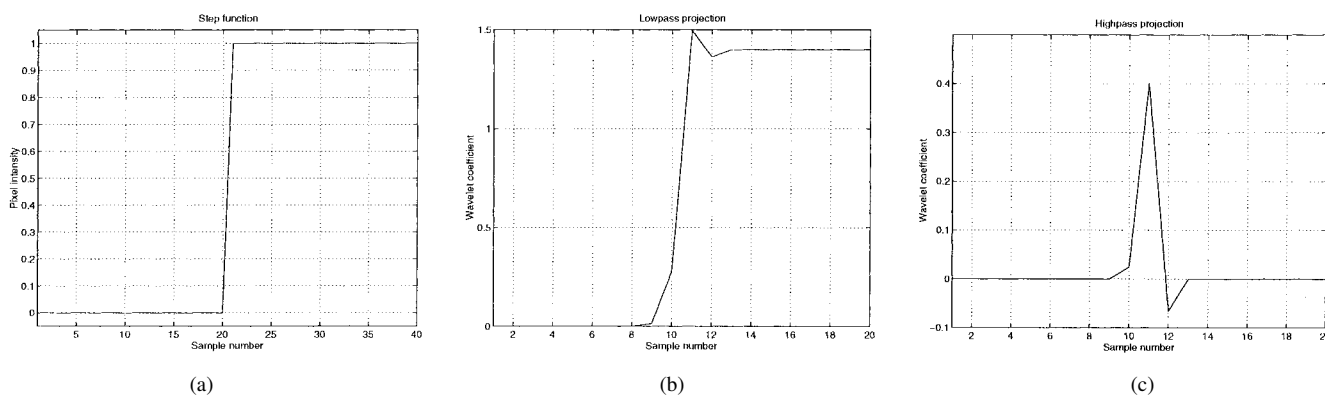


Fig. 2. To illustrate how the wavelet decomposition preserves the locality of data characteristics present in the spatial domain: (a) a step; (b) its projection on the lowpass space; (c) its projection on the highpass space. Observe how in the highpass projection all of the signal energy is concentrated in a neighborhood of the spatial location of the step (the Daubechies 9/7 biorthogonal spline wavelet was used here).




Fig. 3. To illustrate the clustered nature of wavelet data. (a) Coefficients in subband LH_0 of Lena whose magnitude is above a fixed threshold (i.e., those which contain most of the signal energy, necessary to reduce distortion). (b) An equal number of points as in (a), but drawn from a uniform density over the same area. This figure gives motivation for the type of techniques that will be used later in the coder design, techniques that would be completely useless if typical realizations of wavelet data were like (b) instead of (a).

B. Data Structures for Wavelet-Based Image Compression

In an attempt to exploit the fact that signal energy forms clusters in the wavelet domain, zerotree based algorithms [11], [15], [19] introduce a special symbol to indicate that in a certain set of tree-structured coefficients associated with the same spatial region, all coefficients are *insignificant* (i.e., their

magnitude is below a threshold [15]). The success of this approach is due to the fact that, in typical images, a large number of zero-valued coefficients occur in the form of zerotrees; so, by exploiting this VQ-type gain, better performance is achieved than if each zero were coded independently. As a result, the zerotree symbol serves a *dual* purpose: on one hand,




Fig. 4. To illustrate how zerotrees represent zero regions as the union of a highly constrained, blocky grid (a). A representation like (b), capable of capturing the natural clustering in the data, could potentially be more efficient.

it is a very efficient way of encoding these zeros; on the other, it contributes to reducing the uncertainty of the *location* of nonzero values, by means of ruling out certain regions in which these may fall.

Such a scheme has some definite advantages, the most important one being the ease with which rate/distortion trade-offs can be optimized: in [19], a tree-pruning algorithm is proposed, which achieves substantial improvements in coding performance over the basic zerotree scheme of Shapiro [15]. Also important is the dual use of the zerotree symbol discussed above. Mainly at low bit rates, where the uncertainty in the location of a few significant coefficients dominates the total self information of the quantized stream, such a joint pointing/value data structure proves extremely useful. However, it also has its drawbacks. One is that it approximates arbitrarily shaped zero regions as the union of a highly constrained set of tree structured regions. As a result, certain zero regions not well aligned with the tree structure may be expensive to represent, while other portions of zero regions may not be included in zerotrees at all (see Fig. 4).

Along the same lines, perceptually important edge information can get deleted in the tree-pruning based rate/distortion optimization, again due to misalignment with the zerotree structure. Although this may be beneficial in terms of improving rate/distortion objective performance, it is a most questionable decision in terms of subjective quality. Another issue is the fact that zerotrees are not shift invariant: although this is not a major problem for still image coding, it gains dramatic importance when considering extensions to video coding. Finally, it should be noted that zerotrees exploit clustering of the significant values only indirectly, by means of identifying regions of mostly insignificant values, thereby implying the clustering of significance in the complementary regions: a direct approach might have some advantages.

An alternative approach can be motivated as follows. To encode what Shapiro [15] refers to as *significance maps* [e.g., Fig. 3(a)], a straightforward approach consists of assigning a 1-b significant or insignificant label to each coefficient. However, if a given coefficient is known to be classified

significant, then coefficients in a small neighborhood, with high probability, will be significant too (likewise for the insignificant ones), and therefore the approach of spending 1 b per coefficient seems wasteful. This observation suggests that some *region growing*-based technique, where the map is encoded using local probability models depending on the region being encoded, may give an efficient way of encoding it. Morphological operators, formally defined in Section III, provide one way of giving a mathematical description of the intuitively sought region growing operation. Hence, another goal pursued in this work is the construction of a data structure alternative to zerotrees, using morphological operators. This new data structure should, while retaining the key elements that make zerotrees such an efficient tool for coding, improve on some of its deficiencies. Furthermore, a fully “backward” approach is desired, in which no side information is needed to encode the shapes.

C. Related Work

Maragos and Schafer [8] provide an excellent tutorial on morphological systems and their application to multidimensional signal processing problems. The applications they mention only deal with nonlinear filtering problems (e.g., rank-order filters, multiscale smoothing, sampling, correlation) and image analysis problems (e.g., feature extraction, shape representation and description, size distribution, fractals). However, they do not consider applications to coding problems.

Egger *et al.* [5] make use of morphology in an image coding algorithm; however, their approach is completely unrelated to the one being proposed in this work. They propose to make use of morphology to implement a nonlinear filterbank, with the purpose of reducing the ringing effect typical of linear filterbanks at very low rates. Their new morphological subband representation is then quantized and compressed using standard, well understood coding methods. We use a traditional linear filterbank instead (the 9/7 spline wavelet of Daubechies [2]), and morphology is applied only in the classification of coefficients into significant and insignificant. Much higher

coding performance is expected in this case, since while still exploiting the known decorrelating properties of the linear filterbank, morphology is used as a tool in the characterization of statistical properties of the quantized data.

Salembier *et al.* [12] consider the application of morphology tools to image and video coding problems. Following their terminology, the *structure of the coding process* they consider is motivated by very low bit rate applications, where the goal is not only to compress a waveform by eliminating redundancy, but also to make use of properties of the human visual system. As a result, their proposed coding techniques can be described as *object-based*. Our coder, on the other hand, is a pure waveform coder. The morphological sets we consider are not intended to capture any image domain objects or to exploit any property of the human visual system, but instead are tools we use to assign probabilities to sets of wavelet coefficients.

D. Paper Organization

This paper is organized as follows. In Section II a few experiments are proposed, to motivate the novel use of morphological operators as a tool to encode the location information of significant coefficients. Then, in Section III, two image coding algorithms are presented, and in Section IV, simulation results of these algorithms are given. The paper concludes with Section V, where some perspective is given on the new results introduced in this work.

II. STATISTICAL PROPERTIES OF WAVELET COEFFICIENTS

The promise of morphology for wavelet-based image coding has been first suggested in [9]. The experiments reported in that work are essentially two. In the first one, correlations among neighboring coefficients are measured. Not surprisingly given the decorrelating properties of the linear transform, those correlations were found to be negligible. In the second experiment reported, correlations are measured not among coefficients themselves, but instead among the *magnitudes* of neighboring coefficients. In this case, it was found that correlations are somewhat higher than in the previous experiment. Now, it is clear that if dependencies remain after decorrelation, some form of nonlinear processing will be required, since the role of the linear transform is precisely to remove linear dependencies. This agrees with the fact that some correlations show up only after a stage of nonlinear processing (i.e., taking magnitudes). Besides, the result of their last experiment suggests that large magnitude coefficients are likely to be found at spatial locations close to those of other large magnitude coefficients, and the same holds for the low magnitude ones; i.e., further evidence of the cluster formation property is shown by these experiments.

A. Intraband Dependencies

The goal of this first experiment is to establish the feasibility of partitioning subbands into two subsets based on morphological dilation operations, and obtaining good performance by coding each of these subsets separately. Details about morphological operations are presented in Section III-A and [8]; and this point though, it is enough to intuitively understand

that the result of applying a morphological dilation to a given set results in enlarging that set with a few “nearby” elements. First, one such partition is computed, and then the global entropy estimate is compared against the sum of the subset entropies. Of course, in order to build a real codec, the partitioning information needs to be accounted for. At this time, this cost is neglected.

Experiment 1 (On Intraband Dependencies)

- 1) Quantize all subbands with a single uniform quantizer, such that the entropy of the resulting quantized field matches a prespecified bitrate. This entropy is estimated by computing the normalized histogram p_n^{global} of the bins in each subband, and then letting the entropy estimate $\mathcal{H}(p_n^{\text{global}})$ for that subband be $\sum_n -p_n^{\text{global}} \log_2(p_n^{\text{global}})$.
- 2) For each subband, starting with a map of coefficients all labeled insignificant, compute a partition into coefficients *labeled* significant or insignificant as follows.
 - a) Select a coefficient at a random location within the current subband.
 - b) If this selected coefficient is not significant, or if it is significant but it has already been labeled significant, go to 2a.
 - c) Else, this is an unlabeled significant coefficient. Apply one step of morphological dilation at its location, and label all the dilated coefficients significant. For each new significant coefficient found by the dilation, recursively repeat this step, until no new significant coefficients are found.
 - d) Stop selecting coefficients at random when all the significant ones have been labeled significant. Those without a label yet are labeled insignificant.
- 3) Compute now a different entropy estimate. Define two histograms, p^{signif} and p^{insignif} , corresponding to the frequency of occurrence of the coefficients labeled significant and insignificant by this partitioning method, and based on these estimate the entropy of the composite model p^{comp} : that is, let $\mathcal{H}(p^{\text{comp}})$ be $\gamma \mathcal{H}(p^{\text{signif}}) + (1 - \gamma) \mathcal{H}(p^{\text{insignif}})$, where γ , the unconditional probability of significance, is estimated as the ratio of the total numbers of coefficients labeled significant to the total number of coefficients in each subband.
- 4) Compare $\mathcal{H}(p_n^{\text{global}})$ versus $\mathcal{H}(p^{\text{comp}})$.

This experiment was performed on the 512×512 test image Lena, using a five-level wavelet decomposition based on the spline 9-7 Daubechies filters, with a uniform quantizer stepsize chosen so that the global target rate would result in 1b/pixel (stepsize = 11.9631). The DC component is accounted for with a fixed number of bits in both cases, since the statistics of this set differs substantially from the other subbands. The results of the experiment are summarized in Table I.

Subbands are numbered from high to low frequencies. Observe how in subbands 0–5, which contain 93.75% of the data to be coded, the difference of entropies becomes significant: this difference provides conclusive evidence that a carefully designed coding algorithm, provided it makes an

TABLE I
ON THE ENTROPY OF IMAGE SUBBANDS WHEN MODELED AS A SINGLE GLOBAL iid FIELD, VERSUS MODELING IT AS A COMPOSITE OF A SIGNIFICANT AND AN INSIGNIFICANT FIELD (n IS THE NUMBER OF COEFFICIENTS PER SUBBAND. ENTROPIES ARE EXPRESSED IN BYTES)

Subband	$n\mathcal{H}(p^{\text{global}})$	γ	$n\gamma\mathcal{H}(p^{\text{signif}})$	$n(1-\gamma)\mathcal{H}(p^{\text{insignif}})$	$n\mathcal{H}(p^{\text{comp}})$	% diff	
0	7622.82	0.2315	3855.79	2439.88	6295.67	21.08%	
1	4923.58	0.1259	1644.57	2669.48	4314.05	14.13%	
2	2732.29	0.0199	226.16	2396.91	2623.07	4.16%	
3	4326.54	0.5795	3361.28	499.19	3860.47	12.07%	
4	3139.36	0.4125	2075.43	641.40	2716.83	15.55%	
5	2857.74	0.3589	1775.14	630.07	2405.21	18.81%	
6	1872.27	0.8948	1769.98	45.58	1815.56	3.12%	
7	1385.56	0.7329	1192.79	99.74	1292.53	7.20%	
8	1336.89	0.6982	1139.99	84.21	1224.20	9.20%	
9-14		Negligible partitions					0.00%
DC	288.00				288.00		
Total	32767.73				29108.33	12.57%	




Fig. 5. To illustrate the composite nature of subband data: (a) a partition of Lena's LH_0 subband data into two sets; (b) normalized histograms of occurrence of quantization bins in the different partitions and in the subband treated as a whole (the quantizer was chosen so that an overall bit rate of 1 b/pixel would result). Observe how the morphological operators split the subband into two sets whose first-order pdfs have different variances.

efficient management of the side information neglected in this study, should be able to achieve very high coding performance.

In Fig. 5, the resulting partition, as well as the three corresponding histograms for subband LH_0 in Lena are displayed.

B. Interband Dependencies

That important—and necessarily nonlinear—dependencies exist among different subbands has been known for some time. For example, the excellent performance achieved by zerotree based coders [11], [15], [19] is directly attributed to this fact; these coders exploit the fact that sets of zero-valued, tree-structured across bands coefficients occur with probability much higher than that predicted by an independence assumption. However, since it has been postulated that clusters are due to spatial edges, some energy concentration around them can be expected at all scales. Based on this intuition, a new hypothesis is formulated, on dependencies *across* bands: the location of energy clusters in the high frequency bands can be predicted from the location of clusters in the low frequencies. To test this new hypothesis, another experiment is performed:

Experiment 2 (On Interband Dependencies)

- 1) Compute the significance map with respect to some given threshold.
- 2) Define a new map, the *prediction map*, as follows: for each coefficient, if its parent in the wavelet tree is labeled significant, predict that this one will be significant too; otherwise predict insignificance.
- 3) Apply the morphological dilation operation to the prediction map, as in Experiment 1.
- 4) Using the prediction map instead of the significance map, compare entropies as in Experiment 1.

In this experiment, the morphological dilation of step 3 serves a different purpose than in Experiment 1. In this case, it is used to reflect uncertainty in the location of the significant coefficients in the child band. Unlike in the case of zerotrees, where the existence of a zero value in a parent provides strong evidence to suspect the presence of zeros in all of the children coefficients, here the presence of a nonzero parent does not immediately imply the presence of nonzero children. However, since nonzero values are assumed to be generated by image

TABLE II
ON THE ENTROPY OF IMAGE SUBBANDS WHEN MODELED AS A SINGLE GLOBAL iid FIELD, VERSUS MODELING IT AS
A COMPOSITE OF PREDICTED SIGNIFICANT AND INSIGNIFICANT FIELDS ($p^{p\text{-signif}}$, $p^{p\text{-insignif}}$ AND $p^{p\text{-comp}}$) REFER TO THE
PREDICTED COMPOSITE DISTRIBUTIONS, AND n IS THE NUMBER OF COEFFICIENTS PER SUBBAND; ENTROPIES ARE GIVEN IN BYTES.

Subband	$n\mathcal{H}(p^{\text{global}})$	γ	$n\gamma\mathcal{H}(p^{p\text{-signif}})$	$n(1-\gamma)\mathcal{H}(p^{p\text{-insignif}})$	$n\mathcal{H}(p^{p\text{-comp}})$	% diff	
0	7622.82	0.3351	4673.18	1959.80	6632.98	14.92%	
1	4923.58	0.2145	2230.51	2206.77	4437.28	10.96%	
2	2732.29	0.2043	1338.18	1105.10	2443.28	11.83%	
3	4326.54	0.6781	3606.08	434.65	4040.72	7.07%	
4	3139.36	0.4968	2260.62	578.44	2839.06	10.58%	
5	2857.74	0.4749	2027.47	539.50	2566.97	11.33%	
6	1872.27	0.9517	1826.79	20.58	1847.38	1.35%	
7	1385.56	0.8467	1279.46	60.51	1339.96	3.40%	
8	1336.89	0.7825	1193.81	81.00	1274.81	4.87%	
9-14		Negligible partitions					0.00%
DC	288.00				288.00		
Total	32767.73				29990.31	9.26%	

edges, it is reasonable to expect their presence at all scales in related locations. Therefore, although the exact location of the children nonzeros will depend on various factors (such as filter length, response, etc.), *somewhere* near the location of the children of a nonzero parent the probability of nonzero children becomes much higher than over the whole subband. It is this uncertainty about the exact location of nonzero children, and not possible intraband dependencies in the child subband, that the dilation operator captures.

This experiment was performed under exactly the same conditions as Experiment 1. The results are summarized in Table II. As in Experiment 1, the high frequency subbands containing a substantial amount of the data to be encoded show some predictability that could be potentially exploited for coding. Now, unlike as in the case of intraband dependencies, no side information is required in this case to be able to take advantage of them, so they could be incorporated by any coding method at no cost.

III. A NEW SOURCE CODING ALGORITHM

In this section, the design of two different image coding algorithms is presented, whose main design goal is to directly exploit the statistical properties of wavelet coefficients uncovered by the experiments in Section II. The two algorithms have distinctive features: the first one is built to achieve high performance at a single target bit rate; the second one is built to obtain a successively refinable, fully embedded encoding of the image. The presentation starts with a brief review of those aspects of mathematical morphology relevant to this work. Then, a high level description of the basic coding algorithm is presented. And finally, both the single rate and the embedded algorithms are described in detail.

A. Mathematical Morphology: A Brief Review

In this section, some elementary aspects of mathematical morphology relevant to the application being considered in this work are reviewed. For an excellent survey of the application of morphology to multidimensional signal processing problems in general the reader is referred to [8].



Fig. 6. Illustration of the morphological dilation operation.

Consider a set S to which the dilation operation will be applied, and let B represent some structuring element¹. Let \oplus denote the morphological dilation operator, and let \setminus denote the set theoretic difference operator. The dilated set $S \oplus B$ is defined to be the union of all points falling under the support of the structuring element B , when this structuring element is centered at each point in S . The set $S \oplus B$ can be written as $S \cup (S \oplus B \setminus S)$, where $S \oplus B \setminus S$ is the set of points not in S obtained by the dilation. This process is illustrated in Fig. 6.

Intuitively, for simple structuring elements like B , a dilation produces an enlarged set containing the original S , plus a few nearby elements. Experiment 1 provides empirical evidence that if S is a subset of subband data containing large-valued wavelet coefficients, the set $(S \oplus B) \setminus S$ has much higher first-order probability of significance than the whole complement of S .

B. General Architecture of the Source Coder

As mentioned in the introduction, the standard architecture for transform coding methods consists of a decorrelating transform, followed by some quantization strategy, and final entropy coding of the resulting symbol stream. In this section, the design of the quantizers and entropy coders is presented (the decorrelating transform is, of course, the wavelet transform).

1) *Binary Classification and Quantization*: The experiments presented in Section II suggest the presence in the image subbands of at least two subsets of coefficients, having dissimilar statistical properties, that can be efficiently captured

¹An arbitrary set, with the intuitive role of defining a morphology-based distance; although not strictly necessary, in this work S is assumed to be a connected set.




Fig. 7. Different stages in the execution of the morphological coder: observe how the structure is built in a data driven way, thus tracking the shape of clusters naturally present in the data.

with morphological operations. Here, one particular algorithm is presented, to compute and encode one such classification map in a uniquely decodable way, and making use of only a minimal amount of side information.

Assume that a single quantizer has been applied to the whole image, so that a quantized field is obtained. One straightforward way of partitioning this quantized field which, as discussed in Experiment 2, can be incorporated at no extra cost in terms of side information, is that of partitioning a child subband based on the energy activity in its parent subband. This first stage of partitioning produces a binary classification of subband data: coefficients are *predicted* to be significant or insignificant, where significance is defined relative to a quantizer-dependent threshold. Now, like in any prediction scheme, errors will occur, and therefore improved performance can be achieved by encoding the prediction errors. Experiment 1, on intraband dependencies, suggests a natural way to accomplish this: further partition each subband into *truly* significant and insignificant coefficients. As a result, a partition into four sets is defined as follows:

- 1) truly significant coefficients, predicted to be significant;
- 2) truly insignificant coefficients, predicted to be insignificant;
- 3) truly significant coefficients, predicted to be insignificant;
- 4) truly insignificant coefficients, predicted to be significant.

Now, it was already pointed out in Experiment 1 that intraband partitioning would require the use of some amount of side information. A particularly efficient way of doing so is by using morphology. Assume the encoder scans subbands in a raster scan order, until a significant coefficient is detected. In such a situation, the encoder signals this event to the decoder with a special symbol (i.e., side information). With both encoder and decoder aware that a significant coefficient is present at the current location, one step of dilation is performed: the encoder sends to the decoder and labels significant those coefficients in a neighborhood of the current significant one. Once encoder and decoder have access to a few extra coefficients, these new ones are examined: if new significant coefficients are found, the process is applied recursively on each of these new values; otherwise, the process stops and raster scanning of insignificant coefficients resumes, until the whole subband is exhausted. As a result, clusters of significant coefficients are efficiently captured by morphological dilation: the only side information required to encode a new cluster is

the symbol used to signal the transition from an insignificant to a significant region. Observe that in this method of using values to encode the clusters, after the location of the first coefficient in the cluster is known, *all other locations are (implicitly) encoded relative to this first one, by values which would have had to be sent anyway*. Therefore, values serve a dual purpose: they both convey information about themselves, and about the location of future values. This is a most efficient way of encoding location information, that is able to avoid making use of typical shape descriptors (e.g., chain codes), very expensive in bit rate.

2) *Simple Numerical Example:* To help the reader to better visualize the recursive growing process, Fig. 7 shows a few frames taken from an animation available at our website [10], that illustrates the operation of this recursive algorithm.

In the same spirit, Fig. 8 illustrates with a simple numerical example the recursive region growing operation, as well as the computation of the symbol stream.

3) *Pseudocode Description:* In this section a pseudocode description of the recursive dilation performed at both encoder and decoder is given, to give the reader a more intuitive understanding of how the data-dependent scanning order is computed.

At the encoder, each subband is scanned in raster order, coding coefficients using one of the insignificant tables depending on whether it was predicted insignificant or not. Then, when a significant one is found, an extra symbol is sent to the decoder to signal this event, and the algorithm `RecursiveSend` is applied (see Table III); once it terminates, raster scanning resumes on those locations left untouched. The decoder can keep track exactly of the decisions taken by the encoder, with an algorithm analogous to `RecursiveSend` (see Table IV).

4) *Entropy Coding:* It is well known that good implementations of arithmetic coding [3] produce bitstreams of length almost equal to the entropy bound. Motivated by the entropy numbers obtained in the experiments, the resulting symbol stream is compressed using this technique. Now, in order to take full advantage of the potential gains suggested by these experiments, each of the four sets into which each subband is partitioned is encoded independently of each other, using separate probability tables in the arithmetic coder. Each of these tables is adaptively updated as encoding progresses. Note: the reason for using adaptive arithmetic coding is not to capture possible nonstationarities in the data, but instead to encode a source modeled as being stationary without having to explicitly send its pdf as side information. During encoding, each of the tables is reset at the beginning of each subband.

TABLE III
PSEUDOCODE DESCRIPTION FOR THE BASIC ENCODER

RecursiveSend (Qwv, VisitedMap, Location (i,j), ACContext)

1. For all locations (i+di,j+dj) within the support of the structuring element, centered at location (i,j):
 - If (VisitedMap at location (i+di,j+dj) is NOTSENT)
 - Set VisitedMap at location (i+di,j+dj) to NEWCOEFF.
 - Compute the conditioning term corresponding to (i+di,j+dj).
 - Encode Qwv at (i+di,j+dj) into ACContext.
 - Update the proper probability table.
2. For all locations (i+di,j+dj) within the support of the structuring element, centered at location (i,j):
 - If (VisitedMap at location (i+di,j+dj) is NEWCOEFF)
 - Set VisitedMap at location (i+di,j+dj) to SENT.
 - If (Qwv at location (i+di,j+dj) is significant)
 - RecursiveSend (Qwv, VisitedMap, (i+di,j+dj), ACContext).

TABLE IV
PSEUDOCODE DESCRIPTION FOR THE BASIC DECODER

RecursiveReceive (Qwv, VisitedMap, Location (i,j), ACContext)

1. For all locations (i+di,j+dj) within the support of the structuring element, centered at location (i,j):
 - If (VisitedMap at location (i+di,j+dj) is NOTRECEIVED)
 - Set VisitedMap at location (i+di,j+dj) to NEWCOEFF.
 - Compute the conditioning term corresponding to (i+di,j+dj).
 - Decode Qwv at (i+di,j+dj) from ACContext.
 - Update the proper probability table.
2. For all locations (i+di,j+dj) within the support of the structuring element, centered at location (i,j):
 - If (VisitedMap at location (i+di,j+dj) is NEWCOEFF)
 - Set VisitedMap at location (i+di,j+dj) to RECEIVED
 - If (Qwv at location (i+di,j+dj) is significant)
 - RecursiveReceive (Qwv, VisitedMap, (i+di,j+dj), ACContext).

T is set to be approximately 40% larger than the regular stepsize, at all rates tested and for all images tested, the rate and distortion numbers obtained differ only after the third decimal place from the numbers that result from an exhaustive search of the jointly optimal parameters. As a result, and without any significant performance degradation, the whole encoding operation is parameterized by a single “quality factor,” thus avoiding complex search methods. Other coders (e.g., [19]) derive their excellent rate/distortion performance from complex quantizer optimizations: given a user specified parameter that trades off rate for distortion, other coders have to *search* for the quantizers that achieve high performance. This coder simply takes the user supplied quality control parameter as the uniform stepsize, without further computation.

In the remaining aspects, this algorithm follows exactly the partitioning method described in the general architecture. The

test for significance is defined as a simple comparison of the quantized bins with a threshold.

2) *Entropy Coding*: Other than in the use of a uniform quantizer with a wider zero bin, this encoder has not taken advantage of any form of entropy constraining [4] operations to improve performance.

In this framework, location and value information are very deeply intertwined because of the data-dependent scan path, as illustrated by the example in Section III-B2. As a result, it seems unlikely that a computationally tractable form of entropy constrained optimization can be devised, as was done for example with the tree pruning method of Xiong *et al.* in [19]. However, a simple heuristic along these lines is incorporated: in the set of correctly predicted insignificant coefficients, all values are set to zero.

In the remaining aspects, the other three sets are coded as described in the general framework.

D. Algorithm 2: Efficient Embedded Coding

1) *Quantization*: In this version, each coefficient is quantized with a set of embedded quantizers $\{Q_0 \cdots Q_{n-1}\}$ [16]:

$$Q_i(x) = \begin{cases} z & |x| < T/2^{i+1} \\ p & T/2^{i+1} \leq x < T/2^i \\ n & -T/2^i < x \leq -T/2^{i+1} \end{cases}$$

for $i = 0 \cdots n - 1$, and for some parameter T (typically the largest magnitude in the data stream to be encoded). By successively applying these quantizers to each wavelet coefficient, a symbol stream is generated with the property that as n is allowed to increase, a sequence of reconstructions is generated that converges to the original image, and such that any initial substream corresponds to some lower resolution encoding. The issue of embedded quantization as well as why it is important and its many applications has been thoroughly studied elsewhere, and will not be dealt with in this work. The reader is referred to the many good references available on the subject [6], [15], [16].

2) *Entropy Coding*: The design of optimal quantizers for an embedded coder is a problem that, unlike the case of optimal designs for a single target bitrate, has no solution: Equitz and Cover [6] showed that optimality can only be achieved when the source has a certain Markov property. By optimality we refer to the ability of an encoder to achieve the theoretical rate/distortion performance bound for the source being coded.

In this work, the quantizers to be used are fixed beforehand, and this is done independently of the image to code. As a result, the performance of the embedded coder is entirely determined by the efficiency with which the resulting symbol stream is entropy coded. Conditional arithmetic coding based on four probability tables (as described in Section III-B) is used on each resolution level.

IV. EXPERIMENTAL RESULTS

In this section, experimental results obtained in the simulation of the proposed algorithms are given. For the tests, the 9-7 Biorthogonal spline filters of Daubechies [2] and the standard 512×512 test images Lena and Goldhill were used. The wavelet decomposition consists of five levels. The bit rates reported are actual file sizes (i.e., no entropy estimates), and the image quality is measured by the peak signal-to-noise (PSNR) ratio, computed from actually decoded images.

A. Performance of the Single-Rate Algorithm

In a first set of simulations, the performance of Algorithm 1 is explored. In Table V, the performance at various bit rates of the single rate algorithm is shown for both test images.

Sample reconstructions obtained using this algorithm are shown in Fig. 9. Only reconstructions corresponding to low rates are displayed because, on printed material, the original and reconstructions at higher rates are visually indistinguishable.

TABLE V
RATE/DISTORTION PERFORMANCE OF THE SINGLE RATE ALGORITHM

Lena						
Rate (bpp)	0.0325	0.125	0.25	0.50	0.75	1.00
(compression)	(256:1)	(64:1)	(32:1)	(16:1)	(11:1)	(8:1)
Algorithm 1	25.82	31.09	34.12	37.17	38.98	40.33
Goldhill						
Rate (bpp)	0.0325	0.125	0.25	0.50	0.75	1.00
(compression)	(256:1)	(64:1)	(32:1)	(16:1)	(11:1)	(8:1)
Algorithm 1	25.20	28.47	30.53	33.15	35.02	36.56

TABLE VI
RATE/DISTORTION PERFORMANCE OF THE EMBEDDED ALGORITHM

Lena						
Rate (bpp)	0.0325	0.125	0.25	0.50	0.75	1.00
(compression)	(256:1)	(64:1)	(32:1)	(16:1)	(11:1)	(8:1)
Algorithm 2	25.76	30.87	33.90	37.01	38.67	40.20
Goldhill						
Rate (bpp)	0.0325	0.125	0.25	0.50	0.75	1.00
(compression)	(256:1)	(64:1)	(32:1)	(16:1)	(11:1)	(8:1)
Algorithm 2	25.31	28.31	30.61	32.92	34.67	35.96

TABLE VII
COMPARISON OF THE RATE/DISTORTION PERFORMANCE
OF THE PROPOSED ALGORITHMS WITH OTHER HIGH
PERFORMANCE CODING METHODS FOUND IN THE LITERATURE

Image	Lena			Goldhill		
	0.25 (32:1)	0.50 (16:1)	1.00 (8:1)	0.25 (32:1)	0.50 (16:1)	1.00 (8:1)
Shapiro [15]	33.17	36.28	39.55	N/A	N/A	N/A
Said-Pearlman [11]	34.12	37.22	40.41	30.56	33.13	36.55
Algorithm 1	34.12	37.18	40.33	30.53	33.15	36.56
Algorithm 2	33.90	37.01	40.20	30.61	32.92	35.96
Xiong [19]	34.33	37.36	40.52	30.71	33.37	36.70

B. Performance of the Embedded Algorithm

In a second set of simulations, the performance of the embedded algorithm is explored. The numbers obtained are shown in Table VI.

C. Comparison with Other Coders

In Table VII, the performance of the proposed methods is compared against that of other zerotree based high performance coders found in the literature.

A brief comment is due on the algorithms chosen for comparison. The coder by Shapiro was selected because of its historical importance: it is in that work that the concept of making use of location information is first exploited with remarkable success. Also the coders by Xiong *et al.* and by Said and Pearlman were selected because, to the best of our knowledge, they represent the best performing zerotree based coders; besides, they serve the purpose of showing that it is possible to achieve coding performance similar to the best zerotrees have to offer using entirely different techniques.

D. Perceptual Quality

On perceptual issues, these coders are a good tradeoff between conflicting goals such as maximization of objective coding performance (as measured by PSNR), and subjective quality. While the objective numbers achieved by these coders



Fig. 9. Sample reconstructions for Algorithm 1. (a) Lena at 0.5 b/pixel (16:1). (b) Lena at 0.125 b/pixel (64:1). (c) Goldhill at 0.5 b/pixel (16:1). (d) Goldhill at 0.125 b/pixel (64:1).

are similar to those of [19], unlike the latter, the improvement in PSNR with respect to the basic zerotree coder of Shapiro [15] does indeed translate into better looking pictures, essentially due to some noise reduction properties briefly illustrated in Figs. 10 and 11.

While [19] very efficiently represents high-frequency, low-magnitude texture like coefficients (essential to achieve high PSNR values), these proposed coders are not as efficient. As a result, isolated significant coefficients are deemphasized, with preference given to more clustered values. This translates into somewhat better perceptual quality, since such isolated coefficients are not very relevant to the human eye, even though

occasionally they may contain significant energy. However, if perceptual gains are sought, these coders are far from optimal: in Fig. 12, a few reconstructions at extremely low bit rates are shown, since there it is possible to visually recognize where bits are being invested. It is clear that the textured region corresponding to the feathers in Lena's hat takes a significant amount of bits at low rates; in terms of perceptual quality, this is a most questionable decision taken by both encoders. From these observations it is concluded that a transform better matched to the statistics of this image (e.g., wavelet packets), can potentially improve subjective quality while still maximizing objective performance.

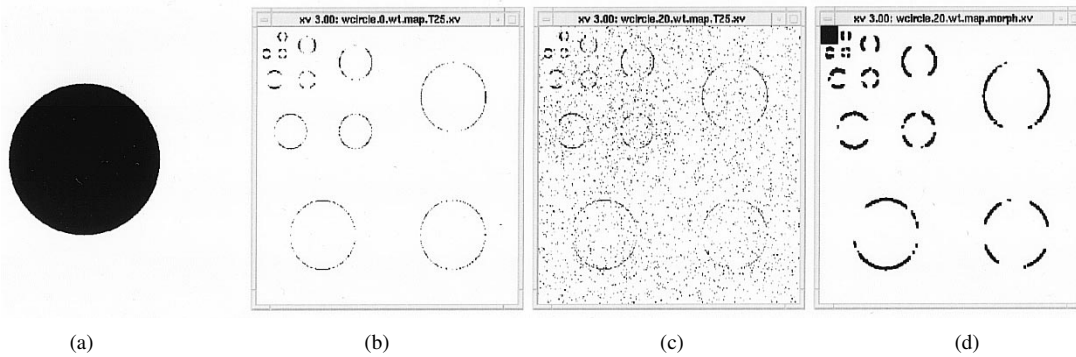


Fig. 10. Illustration of the noise reduction properties of the proposed coders: (a) a synthetic image; (b) location of the significant wavelet coefficients of the synthetic image; (c) location of the significant wavelet coefficients of the synthetic image corrupted by additive white Gaussian noise; (d) partition into coefficients labeled significant and insignificant computed by the morphological coder. Observe how the set of coefficients labeled significant is mostly composed of those drawn from the signal part of the input, while those labeled insignificant correspond mostly to noise.

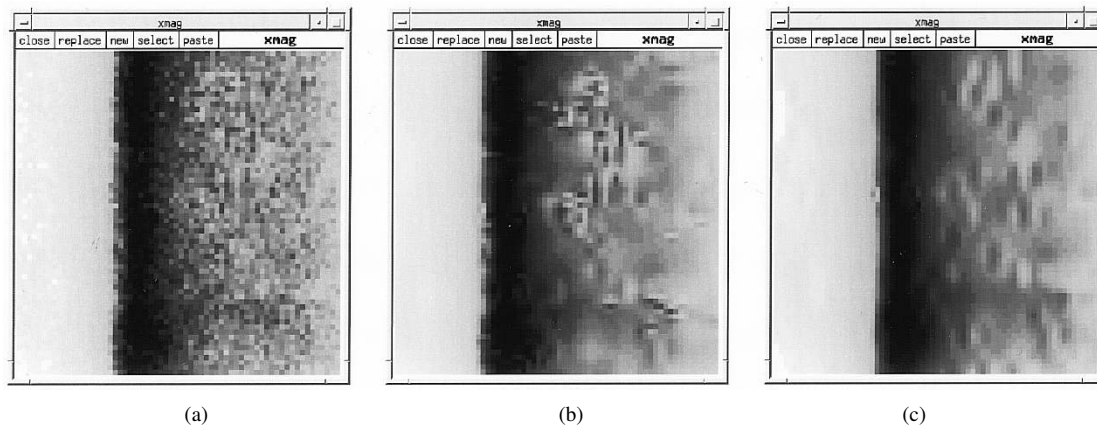


Fig. 11. Illustration of the noise reduction properties of the proposed coders: (a) an original background section from Lena; (b) SFQ's reconstruction at 1 b/pixel (40.52 dB); (c) reconstruction of the fixed rate algorithm at 1 b/pixel (40.33 dB). Observe that while both coders are preserving the edge to great accuracy, the zerotree coded image preserves some of the noise in the background of the original, whereas the morphologically coded image blurs this noisy contribution. Even if this noisy background were due, e.g., to some image texture (i.e., not noise, like in this case), this is information to which the human eye is fairly insensitive, and therefore not very important in a perceptual sense.



Fig. 12. Illustration of where bits go. Sequence of reconstructions at very low bit rates: (a) 512 bytes, (b) 1024 bytes, (c) 2048 bytes, (d) 3072 bytes; recall that the original has 262 144 bytes. Observe how textured regions receive a significant amount of bits.

E. Code Availability

More information (test images, sample reconstructions, animations that illustrate the operation of the proposed methods, source code, documentation, etc.) can be downloaded from our website [10].

V. CONCLUSIONS

In this final section, some perspective is given on all the new concepts introduced in this work. First, a comparison

with zerotrees is made. Then, the main concepts presented in this paper are summarized. And finally, areas requiring further research are discussed.

A. Comparison with Zerotrees

The morphology based approach to coding proposed in this work is an attempt to exploit the same features of wavelet data that zerotrees do to achieve their excellent coding performance, while at the same time improving on some of its deficiencies.

In a sense, the morphological approach proposed in this work is a concept dual to that of zerotrees, in that it exploits *directly* the clustering of significance. Also, arbitrarily shaped regions of nonzero coefficients are defined at no extra cost. On the perceptual side, it is reasonable to expect better looking reconstructions in this case, since the bit allocation policy is driven mainly by energy clustering, thus avoiding decisions such as zeroing out important edges because of their not being well aligned with some arbitrary prespecified grid. However, it also has certain drawbacks, of which the fact that the rate/distortion optimization of this structure is much more complex is the most important one.

In summary, both methods are seen, each with its own pluses and minuses, to exploit the same source of gains: no inference can be made regarding the superiority of one method over the other. This work serves the purpose of showing the feasibility of achieving similar performance to the best attained so far by zerotrees, with equal or lower complexity, and making use of entirely novel coding tools.

B. Summary

A new statistical model for understanding the excellent performance achieved by a variety of image coding algorithms was proposed, and insights derived from this new model were used in the construction of two novel coding algorithms. These algorithms show rate/distortion performance competitive with that of the best methods published in the literature. Other features of the proposed methods are 1) their fairly low complexity, due to their ability to avoid difficult quantizer optimizations, and 2) their slightly improved perceptual quality, due to prioritizing coefficients of clustered energy, which typically results in reconstructed images with perceptually irrelevant information neglected first.

C. Future Research

One of the main advantages of the proposed coding algorithms is the fact that the probability tables used by the arithmetic coder are reset very frequently (at the beginning of each subband for Algorithm 1, and within each subband at the beginning of each layer, for Algorithm 2). These resets, done because they are necessary from a source coding point of view and without regard for any channel coding consideration, enable a “packetization” of the binary stream generated by the arithmetic coder: unlike in standard arithmetic coded streams, where the loss of intermediate bits precludes the possibility of continuing decoding, in this case only that substream until the next reset point is lost, and decoding can be resumed. The use of this feature for image transmission over lossy packet networks is one of the main topics of our current research.


The fact that clusters of coefficients of arbitrary shape can be so efficiently encoded may prove useful in applications where shape representation is an issue, such as image databases. We are currently exploring ways of exploiting this feature to enable processing of content related queries in the compressed domain.

The amount of gain that could be obtained by a careful selection of filterbank and quantizers is yet to be determined in this framework, as there is evidence that perhaps the ones used in this work are not the best ones to maximize objective performance.

Many extensions to this work are possible. Optimization as in the traditional approach is certain to improve results, given the simplicity of the quantization strategy employed. Another source of expected gains is by doing a more general form of set partitioning: even within the significant and insignificant sets computed by this coder there is still evidence of statistical diversity; an algorithm capable of determining the optimal number of sets into which to partition each subband for a given target bitrate is likely to boost performance too. Also, the studies that led to the improved statistical description of wavelet coefficients provide insight to attempt a formalization of these concepts into a random field model. All of these topics are being currently investigated.

REFERENCES

- [1] N. Ahmed, T. Natarajan, and K. R. Rao, “Discrete cosine transform,” *IEEE Trans. Comput.*, vol. C-23, pp. 88–93, Jan. 1974.
- [2] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, “Image coding using wavelet transform,” *IEEE Trans. Image Processing*, vol. 1, pp. 205–221, Apr. 1992.
- [3] T. Bell, J. Cleary, and I. Witten, *Test Compression*. Englewood Cliffs, NJ: Prentice-Hall, 1990.
- [4] P. Chou, T. Lookabaugh, and R. Gray, “Entropy constrained vector quantization,” *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 37, pp. 31–42, Jan. 1989.
- [5] O. Egger, W. Li, and M. Kunt, “High compression image coding using an adaptive morphological subband decomposition,” *Proc. IEEE*, vol. 83, pp. 272–287, Feb. 1995.
- [6] W. Equitz and T. Cover, “Successive refinement of information,” *IEEE Trans. Inform. Theory*, vol. 37, pp. 269–275, 1991.
- [7] Y. Kim and J. Modestino, “Adaptive entropy coded subband coding of images,” *IEEE Trans. Image Processing*, vol. 1, pp. 31–48, Jan. 1992.
- [8] P. Maragos and R. Schafer, “Morphological systems for multidimensional signal processing,” *Proc. IEEE*, vol. 78, pp. 690–710, Apr. 1990.
- [9] K. Ramchandran and M. Orchard, “An investigation of wavelet-based image coding using an entropy-constrained quantization framework,” *IEEE Trans. Signal Processing*, vol. 46, pp. 342–353, Feb. 1998.
- [10] <http://www.ifp.uiuc.edu/~servetto/>.
- [11] A. Said and W. Pearlman, “A new, fast, and efficient image codec based on set partitioning in hierarchical trees,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 6, pp. 243–250, June 1996.
- [12] P. Salembier, P. Brigger, J. Casas, and M. Pardàs, “Morphological operators for image and video compression,” *IEEE Trans. Image Processing*, vol. 5, pp. 881–898, June 1996.
- [13] S. Servetto, K. Ramchandran, and M. Orchard, “Morphological representation of wavelet data for image coding,” In *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, Detroit, MI, 1995.
- [14] ———, “Wavelet based image coding via morphological prediction of significance,” In *Proc. IEEE Int. Conf. Image Processing*, Washington, DC, 1995.
- [15] J. Shapiro, “Embedded image coding using zerotrees of wavelet coefficients,” *IEEE Trans. Signal Processing*, vol. 41, pp. 3445–3462, Dec. 1993.
- [16] D. Taubman and A. Zakhor, “Multirate 3-D subband coding of video,” *IEEE Trans. Image Processing*, vol. 3, pp. 572–588, Sept. 1994.
- [17] J. Villasenor, M.-J. Tsai, and F. Chen, “Image compressing using stack-run coding,” submitted to the *IEEE Trans. Image Processing*, vol. 5, pp. Apr. 1996.
- [18] J. Woods and S. O’Neil, “Subband coding of images,” *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-34, pp. 1278–1288, Oct. 1986.
- [19] Z. Xiong, K. Ramchandran, and M. Orchard, “Space-Frequency quantization for wavelet image coding,” *IEEE Trans. Image Processing*, vol. 6, pp. 677–693, May 1997.



Sergio D. Servetto (M'95) was born in Argentina in 1968. He received the Licenciatura en Informática from Universidad Nacional de La Plata (UNLP), Argentina, in 1992, and the M.Sc. degree in electrical engineering and the Ph.D. degree in computer science, both from the University of Illinois, Urbana-Champaign (UIUC), in 1999.

Between 1991 and 1994, he was a permanent Member of Technical Staff at the Computer Research and Advanced Applications Group, IBM Argentina, where he was involved in the development of optical character recognition systems designed to operate in noisy industrial environments. From 1994 to 1999, he was a Graduate Research Assistant with the Image Formation and Processing Group, Beckman Institute, UIUC, and with the Multimedia Operating Systems and Networking Group of the Department of Computer Science, UIUC. He held part-time teaching appointments during 1992–1993 at UNLP, and during 1996–1997, at UIUC. He worked for Bell Laboratories, Murray Hill, NJ, and AT&T Laboratories, Florham Park, NJ, during the summers of 1997 and 1998, respectively. As of Fall 1999, he is the First Assistant in the Laboratoire de Communications Audiovisuelles at the Ecole Polytechnique Fédérale de Lausanne, Switzerland, and a Consultant with the Multimedia Communications Research Laboratory at Bell Laboratories. His research interests include coding and transmission over networks of image and video signals, as well as fundamental problems in algorithms, signal processing, communications, control, quantization and coding theory that arise in that context.

Dr. Servetto was the recipient of the 1998–1999 Ray Ozzie Fellowship for outstanding graduate students in computer science, and of the 1999 David J. Kuck Outstanding Thesis Award for the best doctoral dissertation of the year, both from the Department of Computer Science, UIUC.



Kannan Ramchandran (M'93) received the B.S. degree from the City College of New York, and the M.S. and Ph.D. degrees from Columbia University, New York, NY, in 1982, 1984, and 1993, respectively, all in electrical engineering.

From 1984 to 1990, he was a Member of Technical Staff at AT&T Bell Laboratories, Whippany, working in the telecommunications R&D area in optical fiber carrier systems. From 1990 to 1993, he was a Graduate Research Assistant at the Center for Telecommunications Research, Columbia University. From 1993 to 1999, he was with the University of Illinois, Urbana-Champaign (UIUC), where he was an Assistant Professor in the Electrical and Computer Engineering Department, and a Research Assistant Professor at the Beckman Institute and the Coordinated Science Laboratory. He is currently an Associate Professor with the Department of Electrical Engineering and Computer Science, University of California, Berkeley. His research interests include image and video compression, multirate signal and image processing and wavelets, and joint source-channel coding for robust image and video communications over noisy channels and heterogeneous networks.

Dr. Ramchandran is a member of the IEEE IMDSP technical committee, and serves as an Associate Editor for the IEEE TRANSACTIONS ON IMAGE PROCESSING. He was the recipient of the 1993 Elaihu I. Jury Award at Columbia University for the best doctoral thesis in the area of systems, signal processing, or communications. He received an NSF Research Initiation Award in 1994, an Army Research Office Young Investigator Award in 1996, an NSF CAREER Award in 1997, and an Office of Naval Research Young Investigator Award in 1997. He received the 1996 Senior Best Paper award from the IEEE Signal Processing Society for a paper (with M. Vetterli) that appeared in the IEEE TRANSACTIONS ON IMAGE PROCESSING in April 1993.



Michael T. Orchard (S'88–M'90–SM'99) was born in Shanghai, China, and grew up in New York, NY. He received the B.S. and M.S. degrees in electrical engineering from San Diego State University, San Diego, CA, in 1980 and 1986, and the M.A. and Ph.D. degrees in electrical engineering from Princeton University, Princeton, NJ, in 1988 and 1990.

He was with the Government Products Division, Scientific Atlanta, Atlanta, GA, from 1982 to 1986, developing passive sonar DSP applications. He has consulted with the Visual Communications Department, AT&T Bell Laboratories, since 1988, and now consults with AT&T Labs-Research. From 1990 to 1995, he was an Assistant Professor with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, where he served as Associate Director of the Image Laboratory of the Beckman Institute. Since summer 1995, he has been an Associate Professor with the Department of Electrical Engineering, Princeton University. His research interests include image and video coding with emphasis on wavelet representations of images and video, motion estimation and compensation, and image display.

Dr. Orchard received the National Science Foundation Young Investigator Award in 1993 and the Army Research Office Young Investigator Award in 1996.