

# Joint Space-Frequency Segmentation, Entropy Coding and the Compression of Ultrasound Images\*

Ed Chiu<sup>†</sup>, Jacques Vaisey<sup>†</sup> and Stella Atkins<sup>‡</sup>

<sup>†</sup>School of Engineering Science, <sup>‡</sup>School of Computing Science  
Simon Fraser University, Burnaby, BC, V5A 1S6, CANADA

{vaisey, stella}@cs.sfu.ca

## ABSTRACT

Joint space-frequency segmentation is a relatively new image compression technique that finds the rate-distortion optimal representation of an image from a large set of possible space-frequency partitions and quantizer combinations. As such, the method is especially effective when the images to code are statistically inhomogeneous, which is certainly the case in the ultrasound modality. Unfortunately, however, the original paper on space-frequency segmentation neglected to use an actual entropy coder, but instead relied upon the zeroth-order entropy to guide the algorithm. In this work, we fill the above gap by comparing actual entropy-coding strategies and their effect on both the resulting segmentations as well as the rate-distortion performance. We then apply the resulting “complete” algorithm to representative ultrasound images. The result is an effective technique that performs significantly better than SPIHT using both objective and subjective measures.

## 1. INTRODUCTION

Many modern image compression algorithms, such as SPIHT [1] are based on the wavelet transform, which partitions the input into frequency bands whose size decreases logarithmically from high frequencies to low. This decomposition strategy works well when the input images are statistically homogeneous; however, when homogeneity cannot be assumed, more general partitions, or bases, may be called for. Other modern algorithms, such as EBCOT [2] (the basis for JPEG-2000) use more general frequency decompositions such as wavelet packets.

A distinguishing feature of ultrasound images [3] is the oriented “speckle texture” produced by the physics underlying the data acquisition. Due to its orientation, the speckle energy is typically concentrated in certain spectral regions. For example, we performed a simple 1-step subband decomposition on a sample “liver” image and found that the horizontal shape of the speckle resulted in 88% of the high-pass energy being concentrated in the LH subband [4]. In addition, ultrasound images (see Fig. 2)

typically consist of an ultrasound-scanned area, which is often non-rectangular, against a passive background, which may contain text and graphics. The resulting spatial variation in image statistics also presents a challenge to coding methods that use a single partition strategy. In a recent study, Erickson et al. [5] compressed MRI and ultrasound images using both SPIHT and JPEG; they concluded that wavelet-based methods such as SPIHT are subjectively superior to JPEG and we thus use this SPIHT as the basis for our comparisons.

The concept of *wavelet packets* extends the standard octave-band tree-structured filter-bank of the wavelet transform to include all possible binary frequency decompositions. To apply this partition gamut to image coding, a set of quantizers must be defined and a search done to identify the best partition-quantizer combination in the rate-distortion sense. Ramchandran and Vetterli [6] have shown how to do this search using a fast tree-based algorithm.

Wavelet-packets provide an interesting approach, but the resulting segmentation cannot adapt to the spatial variation present in ultrasound images. To increase the partition gamut, Herley et. al. [7] generalized the wavelet packet technique to allow the spatial partition of all subimages; they named this method joint space-frequency segmentation (SFS) using balanced wavelet-packet trees. Unfortunately, however, the algorithm presented in [7] simply used the zeroth-order entropy to measure the rate when quantizing a specific subband: no real entropy coders were implemented. The result of this approximation is to hide several issues that are important in producing an efficient space-frequency representation, namely learning and the possibility of exploiting higher-order entropy.

In this paper, we present an investigation into the effect of the entropy coder on both the segmentation and the final rate-distortion performance obtainable using SFS. The resulting codecs are then applied to the compression of medical ultrasound images.

---

\* This work was supported by the Natural Sciences and Engineering Research Council of Canada.

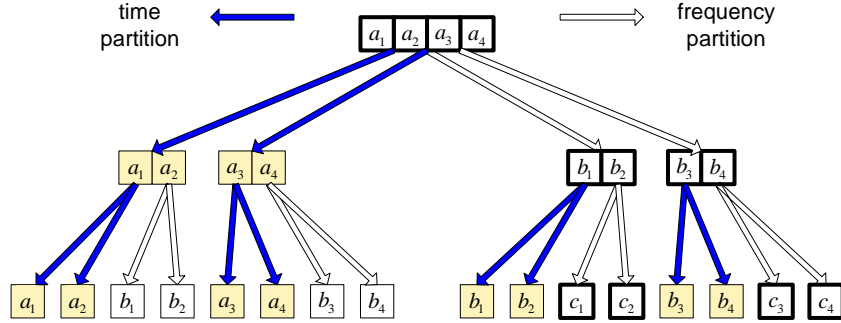


Figure 1: Time-Frequency Partition Gamut

## 2. SPACE-FREQUENCY SEGMENTATION

Joint space-frequency finds the optimal decomposition-quantizer combination for a given target rate by searching through the partition gamut. For each possible partition, a second search must be done to find the quantizer set (one per subband in our case) that minimizes the distortion subject to a constraint on the overall bit rate. In our experiments, we defined a set of 16 different uniform quantizers, each with 64 output levels. The step-size selection for the set was tuned to the target bit rate of the coder. In addition, a linear variation in the step-size over the set was found to give generally good results. Note that the computational cost of the algorithm increases (linearly) with the number of quantizers, as does the number of bits required to send the quantizer information. The choice of 16 is a reasonable trade-off, and was also used by Herley et al. [7]. The partition found by the algorithm tends to be sensitive to the target rate.

In SFS, both space and frequency partitions produce four  $\frac{1}{4}$  size subimages: a space partition simply splits the image region into four quadrants, while a frequency partition filters the image region into four critically sampled subbands using a 2D perfect reconstruction filter bank. The gamut of possible partitions is created by recursively partitioning all of the subimages down to a “maximum decomposition depth”, which we set at five. Lowering this parameter significantly speeds up the algorithm by shrinking the partition gamut, but there is a penalty in the resulting fidelity. The partition gamut for the 1D case is shown above in Figure 1, with the wavelet-packet case shown by bold boxes: data blocks are split into 2 at each step, using a time or a frequency split.

Space-frequency segmentation searches a larger gamut of possible bases than the wavelet packet algorithm. In all cases, the filter set used to perform the decompositions is the “standard” 9-7 biorthogonal filter set [8] and symmetric extension is used to handle the block-boundaries. Other filter sets were tried, but the 9-7 set was confirmed to be a good overall choice, even for ultrasound images [4].

## 3. ENTROPY CODING STRATEGIES

In determining the best sub-image size, the SFS algorithm runs into conflicting demand from the quantizer and the entropy coder. In an ideal world, both the decoder and the encoder have implicit knowledge of the source statistics; however, in reality, this information has to be communicated. One possibility is to simply do an initial pass through each subband to estimate statistics, which can then be sent as side information. Unfortunately, this side-information requires too many bits when there are many subbands to code (as in SFS). An alternative procedure is to make the encoder and decoder adaptive so that statistics are estimated on the fly using previously coded data. Since this information is available at both ends of the channel, no side-information need be sent. Problems still exist, however, since the performance is limited by the accuracy of the estimate and how well the coder can adapt depends on the alphabet size and the sequence length. A small alphabet combined with a large sequence means that the coder acquires a good estimate of source statistics relatively quickly, and produces an efficient code. On the other hand, small subbands, or those with large alphabets, are not coded well by the entropy coder and thus are not good choices overall. Herley [7] avoided this problem by using the zeroth-order entropy to estimate the rate, which is optimistic (and in fact results in different partitions). Nonetheless, making this choice results in a reduction in the algorithm’s execution time by approximately a factor of eight in our implementation.

In this paper, we consider two techniques as candidates for entropy coding: direct arithmetic coding [9] and stack-run coding [10]. For long sequences, the former method typically gives results close to the zeroth-order entropy, while the second is effectively a higher-order entropy coder that exploits that fact that large runs of zeros are common in quantized high-pass image subbands.

Huffman coding is not a good candidate here, since it is too expensive to use one bit per symbol for all of the zeros and is difficult to make adaptive. Other types of entropy coders,




Figure 2: Image U1 and Its Partition

such as those that use expanded contexts [11] are good candidates for future improvements.

#### 4. RESULTS

In the experiments reported below, we implemented the full SFS algorithm and wrote a binary file to disk to measure the coded image size. In addition, we set the target rate to 0.5bpp with the intent of producing good quality ultrasound images, containing sufficient visible impairments to allow for subjective comparisons. The appropriateness of this target rate was confirmed by a radiologist [4].

Our study also focuses on the compression of scan-converted ultrasound images that have been interpolated from the original image data by the imaging equipment. We focus on the post scan-conversion case, since this is the display domain and we want to avoid the introduction of “uncontrolled” artifacts through the combination of quantization and interpolation. In addition, issues like bit allocation become difficult to deal with in the pre scan-converted domain, since the interpolation process may cause some quantization errors to be magnified more than others. Studies on compression in the pre scan-converted domain are left for future work.

The test images that we used in our experiments were the standard “natural”  $512 \times 512$  image, *Barb*, as well as two  $640 \times 480$  ultrasound “liver” images, U1 (shown in Fig. 2) and U2, that were chosen to represent a range of compression “difficulty”, while still being typical. In the partition part of the figure, white lines indicate that a space partition was chosen and black lines indicate that a frequency partition was used. Note that the non-ultrasound area of the image tends to avoid frequency partitions.

The harder case, U1, is an image of a normal liver that shows high contrast and significant fine detail, particularly in the upper central area. There is also low-resolution detail present in the low-contrast dark areas across the center of the image. An unusual feature in U1 is the icon in the upper left indicating the organ being scanned. The easier case to compress, U2, contains a liver lesion. At the time of capture, the contrast was adjusted to highlight the lesion, so the overall image is quite dark and low contrast.

The PSNR results for *Barb* are shown in Table 1, where the SFS-entropy codec uses entropy to do the partitions and also uses this number to compute the final rate. It is seen that the “entropy” only result is overly optimistic, and results in a significantly inflated PSNR at the target rate. Indeed, the problems with the “entropy” are quite severe for the smaller subbands, where there is insufficient data for the coder to adapt to the local statistics. The penalty is thus a function of the maximum decomposition depth.

In the case of the ultrasound images, it was found best to use stack-run entropy coding for all subbands containing frequency partitions and to use arithmetic coding for the remainder. Stack-run coding does not work well on space-only partitions, since the runs of zeros are not large enough and there may be large magnitude coefficients that penalize the stack-run method. The PSNR results for the ultrasound images are shown below in Table 2.

In a 12-image subjective test with two radiologists from the Vancouver Hospital and Health Sciences Centre (VHHSC) [4], the SFS images were consistently ranked above the SPIHT ones, which confirms the PSNR results. In general, SPIHT images tend to suffer from more blurring than SFS ones. Low contrast regions also seem to pose more of a problem with SPIHT.




Figure 3: Ultrasound-Only Partitions

Table 1: Results for *Barb*

codec	Rate [bpp]	PSNR [dB]
SFS - entropy	0.500	32.47
SFS - stack-run	0.496	31.92
SFS - arithmetic	0.497	31.67
SPIHT	0.500	31.40

Table 2: Results for Ultrasound Images

Image	Codec	Rate [bpp]	PSNR [dB]
U1	SFS	0.501	33.00
U1	SPIHT	0.500	32.26
U2	SFS	0.499	38.60
U2	SPIHT	0.500	37.10

It is also interesting to see the “optimal” SFS partitions when only operating on ultrasound data. It turns out that these partitions are quite consistent, as is shown in the 0.5bpp examples in Figure 3, all of which have been extracted from different images. These results indicate that it may be possible to use fixed partitions for ultrasound images when the scanning orientation is similar: a different orientation of the ultrasound transducer would rotate the major axis of the speckle spot. We also observe that space-partitions are rare, which is likely due to the homogeneity of the images.

## 5. CONCLUSIONS

In this paper, we extended the work of Herley et al. [7] by studying the effect on the algorithm performance of using real entropy coders instead of simply an estimate of the entropy. It was found that stack-run coding is, in general, superior to single context arithmetic coding, but that it is sometimes advantageous to key the choice of entropy coding method to the subband being coded. The final codec was applied to the compression of ultrasound images, with

results that are both subjectively and objectively superior to those possible with SPIHT.

## REFERENCES

- [1] A. Said and W. A. Pearlman, "A New, Fast, and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, pp. 243–250, 1996.
- [2] D. Taubman, "High Performance Scalable Image Compression with EBCOT," *IEEE Transactions on Image Processing*, to be published.
- [3] A. C. Fleischer and A. E. James Jr., *Diagnostic Sonography: Principles and Clinical Applications*. Philadelphia: W.B. Saunders, 1989.
- [4] E. Chiu, "Lossy Compression of Medical Ultrasound Images Using Space-Frequency Segmentation," *School of Engineering Science M.A.Sc. thesis*: Simon Fraser University, 1999, <http://www.ensc.sfu.ca/grad/theses/>.
- [5] B. J. Erickson, A. Manduca, P. Palisson, K. R. Persons, D. Earnest, and V. Savcenko, "Wavelet Compression of Medical Images," *Radiology*, vol. 206, pp. 599–607, 1998.
- [6] K. Ramchandran and M. Vetterli, "Best Wavelet Packet Bases in a Rate-Distortion Sense," *IEEE Transactions on Image Processing*, vol. 2, pp. 160–175, 1993.
- [7] C. Herley, Z. Xiong, K. Ramchandran, and M. T. Orchard, "Joint Space-Frequency Segmentation using Balanced Wavelet Packet Trees for Least-Cost Image Representation," *IEEE Transactions on Image Processing*, vol. 6, pp. 1213–1230, 1997.
- [8] M. Antonini, M. Barlaud, P. Mathieu and I. Daubechies, "Image Coding using Wavelet Transform," *IEEE Transactions on Image Processing*, vol. 1, pp. 205–220, 1992.
- [9] T. Bell, J. Cleary, and I. Witten, *Text Compression*: Prentice Hall, 1990.
- [10] M. J. Tsai, J. D. Villasenor, and F. Chen, "Stack-Run Image Coding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, pp. 519–521, 1996.
- [11] X. Wu, "Lossless Compression of Continuous-Tone Images via Context Selection, Quantization and Modeling," *IEEE Transactions on Image Processing*, vol. 6, pp. 656–664, 1997.