

Compression of ultrasound images using wavelet based space-frequency partitions

¹Senior programmer Noor Abdulmuttaleb Jaafar

²Senior programmer Hind Ibrahim Mohammed

³Senior programmer Abbass Alaa Mahdi

^{1, 2, 3}Al-Muqdad College of Education, Diyala University,

E-mail: hindimohammed83@gmail.com

E-mail: hassantalal81@gmail.com

Abstract

This paper describes the compression of grayscale medical ultrasound images using a new compression technique, space-frequency segmentation. This method finds the rate-distortion optimal representation of an image from a large set of possible space-frequency partitions and quantizer combinations. The method is especially effective when the images to code are statistically inhomogeneous, which is the case for medical ultrasound images. We implemented a real compression algorithm based on this method, and applied the resulting algorithm to representative ultrasound images. The result is an effective technique that performs significantly better than a current leading wavelet transform coding algorithm, Set Partitioning In Hierarchical Trees (SPIHT), using the standard objective PSNR distortion measure.

KEYWORDS: Ultrasound image compression, wavelet packets, space-frequency segmentation.

INTRODUCTION

Ultrasound is a popular medical imaging modality because it is non-invasive, versatile, there are no known side effects, and the equipment used for ultrasonic scanning is small and inexpensive relative to other options. Reducing storage requirements and making access to data more convenient are two of the motivations for applying compression to ultrasound images. Retaining diagnostically relevant information is the most important goal for a lossy medical image compression scheme.

A distinguishing feature of ultrasound images is the oriented "speckle texture" produced by the physics underlying the data acquisition. Due to its orientation, which changes somewhat across the image, the speckle energy is typically concentrated in certain spectral regions. Speckle is caused by scattered reflections produced by features that are small with respect to the wavelength. These multiple small reflections result from a rough scattering surface with fine scattering structures. In the frequency domain, speckle exhibits a low pass characteristic. An example of an organ that produces

a particularly speckly ultrasound image is the liver, as seen in Figure 1.

Depending on the context and application, speckle in medical images used for diagnostic purposes can be viewed as signal or noise. For example, speckle can be used to characterize tissue [5,10], or speckle can mask diagnostically relevant features [8,12]. Good performance in a variety of contexts is desirable for a compression system targeted for medical ultrasound images. When it is not known a priori whether speckle is viewed as signal or noise in a specific application, the compression algorithm needs to preserve it. The presence of speckle in an image is also something which radiologists are accustomed to. The compression scheme should avoid altering the image in a noticeable way, so for aesthetic reasons speckle should be preserved even when it does not provide direct diagnostic information. The presence of speckle and the desire to preserve it are issues that distinguish the ultrasound image compression problem from the compression of natural images.

Another characteristic of ultrasound images is the spatial variation in statistics across an individual image. An image typically consists of an ultrasound-

scanned area, which is often non-rectangular, against a passive background, which may contain text and limited graphics (see Figure 1). The resulting spatial

variation in image statistics also presents a challenge to coding methods that use a single partition strategy.



Figure 1: A typical ultrasound picture of a normal liver (T1).

Space-frequency segmentation

When used to compress ultrasound images, SPIHT produces some distinctive artifacts. Speckle is blurred at lower bit rates, suggesting that a different decomposition topology is better suited to speckly images. Because contrast is important in an ultrasound image, bit allocation by coefficient magnitude is not the best choice, and artifacts are particularly visible in areas of low-contrast detail. We have implemented an image-adaptive scheme that selects the best filter topology for each image. If we view the decomposition of an image into subbands as the projection onto a choice of basic functions, the problem is to choose an appropriate basis to represent an individual image. The drawbacks to an adaptive scheme include the computational cost of finding the best basis, and the need to send a description of the basis as side information. The algorithm for finding the optimum representation must have a meaningful criterion for choosing the best basis, and must limit the choice of basis to a useful but finite set of possibilities.

Allowing both space or frequency partitions on any subimage permits more basis choices than previous related methods, such as single-tree wavelet packets [2], which is limited to the same filter topology over the entire image, and double-tree wavelet packets, which is limited to different filter topologies over

spatial regions of the original image only [6]. SFS chooses the optimal combination of space and frequency partitions, and the best quantizer for each subimage, according to a rate-distortion criterion. For a fixed target rate, the algorithm finds the combination of partitions and quantizers that produces the minimum distortion representation. The optimization incorporates a fast tree-pruning algorithm for efficient computation.

In a recent study, Erickson et al. [3] 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 SPIHT as the basis for our comparisons.

The purpose of this paper is to describe the algorithm we used to compress ultrasound images based on the new technique of SFS, and show how this algorithm is superior to SPIHT both in PSNR and objective image quality measures.

Methods

Subband decomposition and image content

To determine the best approach for space-frequency partition for ultrasound images, single stage subband decomposition is applied to an ultrasound-only section of a medical ultrasound image that can be considered typical, as shown in Figure 2a (original) and 2b (the subband image). The ultrasound-only

image is labelled U_b , cropped from the image of a liver in Figure 1. Orthogonal filters, the Daubechies

filters of length eight, are used in order to compare the energy in the different subbands precisely.




Figure (2) a- Ultrasound-only image U_b




Figure (2) b- Subband image of U_b .

For most ultrasound images, over 70% of the energy is concentrated in the LL band [1], as is seen in the upper left quadrant of Fig 2b. The distribution of the remaining energy in the high pass bands shows that the high pass energy distribution of U_b is not uniform. Note there are few edges in the ultrasound-only image, which is dominated by speckle. After subband decomposition of the ultrasound image, no features are visible in the HL or HH bands. The speckle is concentrated in the LH band (lower left quadrant in Fig 2b). Compared to the energy distribution of natural images, speckle results in more

energy in certain high pass subbands, as can be seen in Figure 2.

The concentration of high pass energy is explained by the shape of a speckle spot. Speckle spots in ultrasound images are generally elongated, due to the different transducer resolution in the lateral and axial directions. Often the image is oriented such that the direction of insonification is vertical, so that the speckle spot is elongated in the horizontal direction.

Implementation of a SFS codec

We implemented a SFS codec for ultrasound image compression, using 16 different uniform quantizers for a subimage each with 64 output levels. Many experiments were performed to determine appropriate choices of quantizers; details are reported in [1].

A typical ultrasound image of size 640 x 480 pixels, called T1 (see Fig. 1), was chosen for testing the SFS codec. The image consists of an area against a dark background, with some light colored text in the background area describing the exam conditions. A bar of grayshades provides a visual indication of the grayscale range. The ultrasound area of the image is wedge shaped, implying that curved linear array transducers have been used to capture it. Abdominal exams are a common application for diagnostic ultrasound, and the test image is of the liver. The image is quite speckly, which is typical for liver images, but it is not clear whether or not the speckle is diagnostically significant. The text portion contains a small section of inverse video (e.g. dark text on a light background). The test image was chosen to represent an average case for compression, while still being typical: it shows high contrast and significant fine detail. There is detail present in the low contrast

dark areas across the center of the image. The ultrasound area of T1 is typical.

Comparison of Space-Frequency Segmentation vs. SPIHT.

The “best” version of the SFS codec and quantizer set developed for compressing T1, was then used to compress several more ultrasound test images. These were compressed at different rates using both SFS and SPIHT. Two of the images used, T2 and T3, are shown in Figure 3. The goal was to see if SFS performed better than SPIHT using the standard distortion measure of power signal to noise (PSNR). We also needed to determine a suitable compression rate for subsequent objective assessments.

Expert viewer assessment of compressed images

We did not have the resources to perform a valid formal study of compressed medical image quality, which is a large and demanding task. The expert viewer assessment task that we attempted was a limited study of subjective quality. The goal of the study was to compare the subjective quality with respect to medically relevant features of ultrasound images compressed with SFS and SPIHT, at a single bit rate.




Figure 3a) Original uncompressed image, T2.



Figure 3a) Original uncompressed

The study was conducted with the Radiology Department at Vancouver Hospital and Health Sciences Centre (VHHSC), and all images used in the expert viewer assessment study were obtained from the Radiology Department at VHHSC. The images in the test set were chosen with the help of a radiologist to cover a selection that is representative of what radiologists encounter normally, and which includes diagnostically interesting features in the images. Because the VHHSC Radiology Department conducts a large number of abdominal exams, the images in the test set were dominated by abdominal images, including images of the liver, kidney, gall bladder, and pancreas. Non-abdominal images in the test set consisted of two thyroid images. Most of the abdominal images were captured at 4 MHz using harmonic imaging, in which signal processing using the second harmonic produces a slightly sharper image. The rest of the abdominal images were captured at 3.5 MHz. Curved linear array and phased array transducers were used for the abdominal images. The thyroid images were quite different from the other images because they were captured at 8 MHz using linear array transducers, and consequently the ultrasound portion of the thyroid images was rectangular in shape.

Results and Discussion

Space Frequency Partitions

The subband images for the space-frequency decomposition of T1 (shown in Fig. 1) and T2 (shown in Fig. 3a) at 0.5bpp are shown in Figure 4. A space partition is denoted by a pair of intersecting white lines, and frequency partition is denoted by a pair of intersecting black lines.




Figure 4a) Space-Frequency Segmentation of T1.




Figure 4a) Space-Frequency Segmentation of T2

For both images, space partitions dominate the early decomposition levels. Up to depth two, all partitions for both images are space partitions, except the depth two frequency partitions of the lower left quadrants of both images. The shape of the region drives the early partitions, since the optimum representation of the background area and text is to quantize and code it without any frequency decompositions. Since space partitions are limited to binary divisions of the width and height, it is difficult to fit the wedge shaped ultrasound area well. Multiple space partitions early in the decomposition hierarchy limit the depth of subsequent frequency partitions on the area, and produce small subbands limit coding efficiency after only a few frequency partitions.

When the space partitions isolate an ultrasound scanned area, frequency decompositions dominate, and since the depth is limited by the early space partitions, the frequency partitions proceed to the maximum depth in order to accomplish energy compaction. Note that after the first frequency partition, the LH subband is often further partitioned in frequency, but the HL and HH bands are not. As shown in Fig. 2b, the high pass speckle energy is concentrated in the LH band when the image is oriented such that speckle spots are elongated in the horizontal direction, which is the case here, which explains why additional decomposition of the LH band will produce additional energy compaction.

The optimum decompositions show that the best way to compress ultrasound images is to treat the background and ultrasound scanned areas

differently. However this is not a simple task because of the variety of different objects that can appear in the background (text, reverse video text, graphic objects), and the unusual shape of the ultrasound scanned region. Space-frequency segmentation has limitations in dealing efficiently with both background and ultrasound areas simultaneously, but because it finds the optimal representation without making any assumptions about the underlying image, and because it incorporates space partitions that allow it to deal with different regions with different characteristics, space-frequency segmentation performs relatively well for compressing ultrasound images.

Space-Frequency vs. SPIHT

Results of compressing images T1 and T2 at different bit rates using SFS and SPIHT are given in Figure 5.

The improvement of space-frequency segmentation over SPIHT is about 2.8 dB in the case of T1, and 3.3 dB in the case of T2.

One of the problems in using SPIHT to compress ultrasound images is that degradation of subjective quality is quite noticeable even at high PSNR values. Because SPIHT allocates more bits to large magnitude coefficients in low frequency bands, SPIHT seems to have trouble with the speckle and other fine details in ultrasound images. The subjective quality of the space-frequency compressed versions is superior to the SPIHT compressed versions for both these test images, and the overall subjective quality of the space-frequency compressed images is quite high when compared against the original versions.




Figure 5: PSNR for compressing T1 and T2 at different bit rates, using SFS and SPIHT

Conclusions and Summary

Space-frequency segmentation finds the rate-distortion optimal representation of an image, choosing from a set of bases where the possibilities for space or frequency partition are symmetric. The choice of quantizer used to quantize subband coefficients is included in the rate-distortion optimization. Because it adapts the representation to suit the image, space-frequency segmentation is a good choice for images with unusual feature. Experiments coding ultrasound images with a space-frequency codec produce good results, showing improvement in terms of both PSNR and subjective quality when compared with wavelet transform zerotree coding using SPIHT. The improvement in subjective quality is particularly noticeable in regions of fine speckle texture. Space-frequency segmentation does a better job of preserving both fine detail and low contrast detail.

The optimum partition for the space-frequency representation indicates that the best way to compress medical ultrasound images is to treat the background and ultrasound scanned areas differently. Assessment of compressed ultrasound images by radiologists confirms that, in terms of medically relevant features, the subjective quality of space-frequency compressed images is superior to SPIHT compressed images.

REFERENCES

1. Chiu, E. "Lossy Compression of Medical Ultrasound Images using Space-Frequency Segmentation", M.A.Sc. thesis, Simon Fraser University, August 1999-
2. Coifman, R. R., Y. Meyer, S. Quake, and M. V. Wickerhauser, "Signal processing and compression with wave packets", Technical Report, Department of Mathematics, Yale University, 1991.
3. Erickson, B. J., A. Manduca, P. Palisson, K. R. Persons, F. Earnest IV, V. Savcenko, and N. J. Hangiandreou, "Wavelet compression of medical images", *Radiology*, vol. 206, no. 3, pp. 599-607, March 1998.
4. Herley, C., 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, no. 9, pp. 1213-1229, September 1997.
5. Insana, M. F., R. F. Wagner, B. S. Garra, D. G. Brown, and T. H. Shawker, "Analysis of ultrasound image texture via generalized Rician statistics", *Optical Engineering*, vol. 25, no. 6, pp. 743-748, June 1986.
6. Ramchandran, K. and M. Vetterli, "Best wavelet packet bases in a rate-distortion sense", *IEEE Transactions on Image Processing*, vol. 2, no. 2, pp. 160-174, April 1993.