PROCEEDINGS OF SPIE

Combined vector quantization and adaptive histogram equalization

Pamela C. Cosman, Eve A. Riskin, Robert M. Gray

SPIE.

Combined vector quantization and adaptive histogram equalization

Pamela C. Cosman[†]   Eve A. Riskin[‡]   Robert M. Gray[†]

[†]Durand Building, Department of Electrical Engineering
Stanford University, Stanford, CA, 94305-4055

[‡]Department of Electrical Engineering, FT-10
University of Washington, Seattle, WA 98195

ABSTRACT

Adaptive histogram equalization is a contrast enhancement technique in which each pixel is remapped to an intensity proportional to its rank among surrounding pixels in a selected neighborhood. We present work in which adaptive histogram equalization is performed on the codebook of a tree-structured vector quantizer so that encoding with the resulting codebook performs both compression and contrast enhancement. The algorithm was tested on magnetic resonance brain scans from different subjects and the resulting images were significantly contrast enhanced.

1. INTRODUCTION

Histogram equalization refers to a set of contrast enhancement techniques which attempt to "spread out" the intensity levels occurring in an image over the full available range.[1] Histogram equalization is a competitor of interactive intensity windowing, which is the established contrast enhancement technique for medical images. In global histogram equalization, one calculates the intensity histogram for the entire image and then remaps each pixel's intensity proportional to its rank among all the pixel intensities. In adaptive histogram equalization (AHE), the histogram is calculated only for pixels in a context region, usually a square, and the remapping is done for the center pixel of the square. This can be called "pointwise" histogram equalization because, for each point in the image, one calculates the histogram for the square context region centered on that point. Because this is very computationally intensive, the bilinear interpolative version is an alternative that lowers the computational complexity.[2] It calculates the histogram for only a set of non-overlapping context regions that cover the image and the remapping of pixel intensity values is then exact for only the small number of pixels that are at the centers of these context regions. For all other pixels, a bilinear interpolation from the nearest context region centers determines the appropriate remapping function.

With the bilinear interpolative version of AHE, the remapping function for a given pixel of intensity i at location (x, y) is determined from the nearest 4 context regions as shown in figure 1. If m_{+-} denotes the mapping at the grid pixel (x_+, y_-) to the upper right of (x, y), and similar subscripts are used for the other surrounding context regions, then the interpolated AHE result is given by[2]:

$$m(i) = a[bm_{++}(i) + (1 - b)m_{-+}(i)] + [1 - a][bm_{+-}(i) + (1 - b)m_{--}(i)], \qquad (1)$$

where

$$a = \frac{y - y_-}{y_+ - y_-}, \qquad b = \frac{x - x_-}{x_+ - x_-}. \qquad (2)$$

Figure 1: Bilinear interpolative AHE

Pixels in the border regions of the image are handled separately, by using a linear interpolation from the two nearest context region centers or, in the corners, using only a single remapping function. Typically for a 256 × 256 pixel image, there are 16 context regions of size 64 × 64.

2. TREE-STRUCTURED VECTOR QUANTIZATION

Vector quantization[3,4] (VQ) is a lossy compression technique that has become popular in the last decade for data compression. It works by encoding input vectors to one of a number of pre-selected vectors called codewords. From rate-distortion theory, better coding performance can be achieved by coding vectors rather than scalars. Tree-structured vector quantization[5] uses a tree structure to lower the search complexity of a VQ encoder at the cost of a slight decrease in performance. Figure 2 shows schematically how this works. The encoder takes the image to be compressed and blocks it into vectors. Each block X_n is then encoded by a binary tree until it reaches a terminal node (leaf), Y_i, of the tree, and the index i of the path through the tree is output. These indices are stored. Decompressing with a VQ is rapid since it is a simple table lookup operation. The decoder reads in an index i and looks at its copy of the tree to retrieve Y_i, the reconstruction of X_n.[4]

Tree-structured vector quantization and histogram equalization can both be applied to one image by performing them sequentially, but this requires extra time. Instead of performing the decoding and equalizing operations sequentially, one can perform them simultaneously by equalizing the decoder's codebook off-line.[6] This way, the decoder's reconstruction of the image and the histogram equalization would be performed in the same time required by the decompression alone. To combine VQ with global histogram equalization, one can construct a global histogram containing all pixels that composed the training images, and each pixel of each codeword can be equalized using this global histogram. Thus each pixel of each terminal node will be remapped to a new intensity that is proportional to its rank in the global histogram. These new codewords can be stored at the decoder, along with the original codewords. The encoder is unchanged. The decoder takes the same set of indices and puts them through the same tree, but upon reaching a terminal node of the tree, the decoder now has the option of outputting either the compressed reproduction or the compressed and histogram equalized reproduction. The radiologist thus has the option of looking at either the equalized or the unequalized series of compressed scans and either way requires the same amount of time to reconstruct the image.

3. ADAPTIVE HISTOGRAM EQUALIZATION AND VQ

Global histogram equalization, while providing some improvement in image quality, does not provide as much detail in the resulting image as does adaptive histogram equalization. Unfortunately the simultaneous combination of VQ and AHE is not straightforward. AHE remaps a pixel's intensity using a histogram local to that pixel, so it is not sufficient to know the pixel's intensity to determine the appropriate remapping function. One must also know
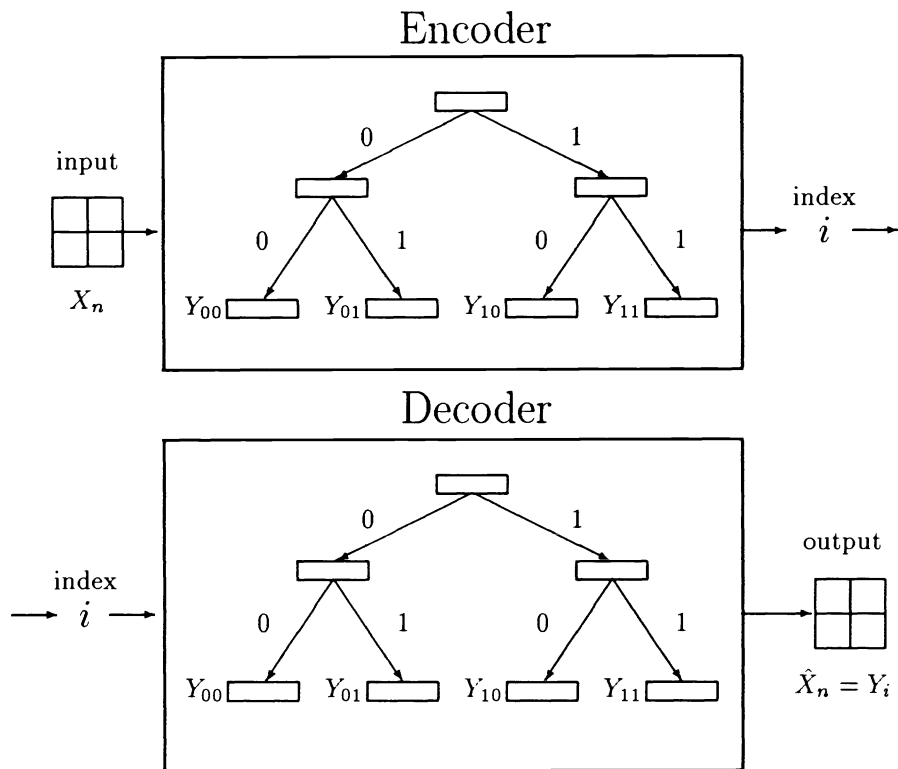
Encoder

Decoder

Figure 2: Schema of a vector quantizer

the pixel's location in the image. This cannot be simply applied to VQ, because although one knows the intensity of any pixel in a codeword, one does not know its "location." Since the codewords represent centroids of clusters of training sequence vectors, the concept of a codeword's "location" within the training images is vague, since the codeword is likely not to exist in any of the images.

The algorithm for combining VQ with AHE is an extension of that used for combining VQ with global histogram equalization.[6] The training images are each divided into 16 context regions as shown in figure 1. The pixels from the corresponding regions of the training images are pooled to form 16 different intensity histograms. The codewords are equalized using each of the 16 different histograms, and the resulting equalized versions of the codewords are stored at the decoder along with the original codewords. The system is diagramed in figure 3. To produce an equalized decompressed image, the decoder makes use of the same set of codeword indices from the encoder, and follows the same paths through the tree. Since the image input to the encoder is scanned in raster order, the spatial location of each vector is known to the decoder without any additional information being required. Knowing the location allows the decoder to generate the coefficients a and b from equation 2, and select the 4 versions of the codeword corresponding to that location. The appropriate linear combination is then formed.

To demonstrate the new technique of combining the compression and histogram equalization steps, an unbalanced tree was grown to an average depth of 2 bits per pixel (bpp) on a training sequence of 10 magnetic resonance (MR) mid-sagittal brain scans of 10 different subjects. The training images of size 256×256 were blocked into 2×2 vectors. The tree was pruned[7] back to 1.7 bpp and used to encode a test image not in the training set. Figure 4 shows the original test image. The 10 training images were divided into 16 square context regions as shown in figure 1. The pixels from all 10 images from each context region were formed into 16 histograms. Those pixels corresponding to the black background of the training images were excluded from the histograms by a semi-automatic algorithm. Inclusion of those background pixels would cause the gray and white pixels of the head regions to rank much higher in the histogram, and thus to remap much brighter. This would correspond to enhancing the contrast between the head and the background, a contrast with which we are not concerned. Exclusion of the background pixels instead causes the enhancement to focus on the contrast between various structures within the head. The terminal node codewords were equalized according to the 16 different histograms, and the versions are stored at the decoder. For
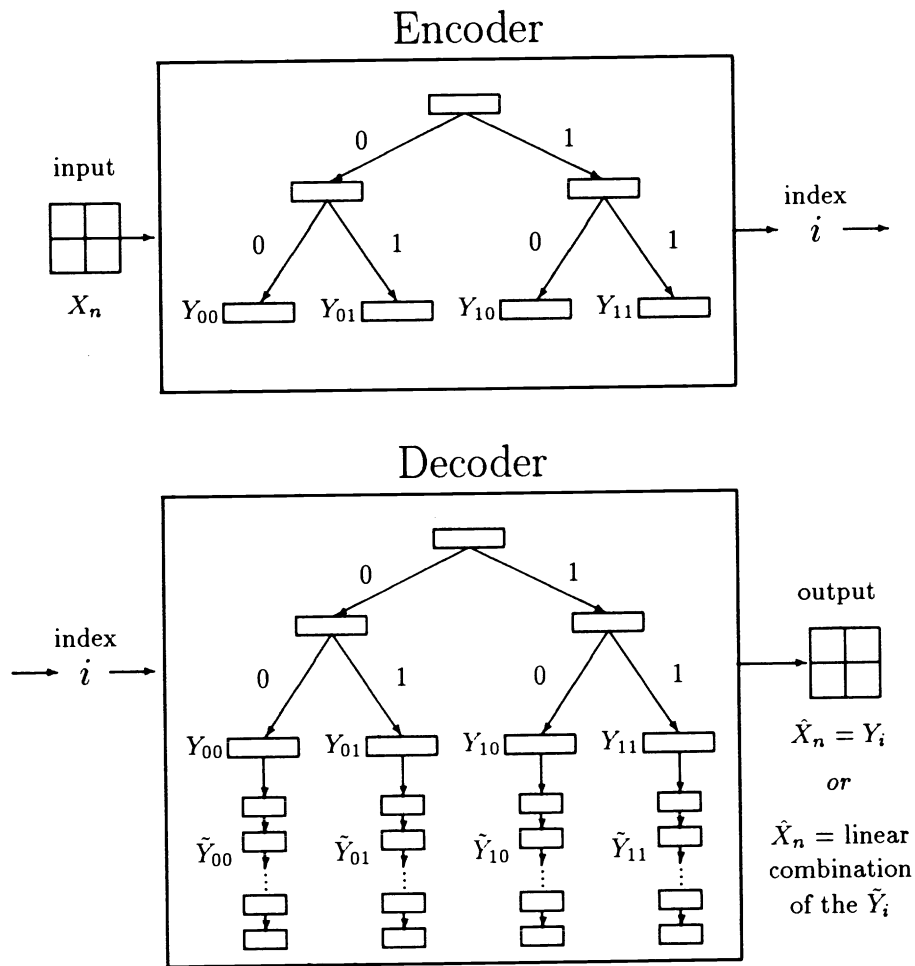
Figure 3: A vector quantizer with both equalized and unequalized codewords

Figure 4: Original image

any index, the decoder then has the choice of outputting either the regular codeword or a linear combination of the spatially appropriate equalized codewords. Figure 5 shows the regular compressed and equalized compressed images. The image quality of the equalized compressed image is very high, and its contrast is enhanced, e.g., the invaginations of the cortex are more obvious, and the vertebrae are more clearly differentiated from the interstitial spaces between them.

With this type of scan, as with many other medical scans in which the imaged portion of the body is displayed on a dark background, a tremendous benefit is derived from removing the background pixels from the histogram prior to the equalization. These operations of decompression, background removal, formation of the histograms, and equalization using those histograms, could be performed sequentially. Sequential operations allow one to use the true histogram of the decoded image, instead of the histogram of the training images. The advantage of our method is that the operations of background removal, histogram formation, and equalization are done off-line using the training sequence at the time the VQ is designed. Subsequent decoding is limited to the decompression step (a table look-up) and forming a linear combination of the equalized codewords. The disadvantage of this algorithm is that the decoder is required to store 16 versions of each terminal node of the tree. Thus there is a substantial time savings at the expense of a significant increase in storage space required at the decoder.

The algorithm requires a certain degree of stationarity in both intensity and spatial coordinates. If the test images are significantly different from the training images, the equalization based on the training images and applied off-line to the codewords will be inappropriate for the test images. This algorithm is best suited to images that exhibit such stationarity. An example are images generated from medical applications in which certain regions of the body are imaged in standard positions by a given modality such that the resulting images have sufficient similarity across subjects. By comparison, the global histogram equalization algorithm requires some stationarity in intensity but not in spatial coordinates, as the pixels from the entire image are pooled for the histogram. That algorithm would therefore be more suitable when the imaged portion of the body might be significantly displaced spatially in the image from one subject to the next. Alternatively if the image could be spatially registered to a standard, then the AHE version of the algorithm could be used.

4. CONCLUSION

We have shown that adaptive histogram equalization using the histograms of the training images can be performed off-line on a vector quantizer's codebook. When an encoded image is decoded in real time using linear combinations of the equalized codewords of the same tree, the resulting decoded picture has significantly enhanced

Figure 5: Compressed images: regular and equalized at 1.78 bpp

contrast. The enhancement is very similar to what would be produced by decoding using the regular tree, followed by an adaptive equalization of the decoded image using its own histogram. The time required for this post-processing step is thus substantially reduced.

5. ACKNOWLEDGEMENTS

6. REFERENCES

1. R. Hummel, "Image Enhancement by Histogram Transformation," *Computer Graphics and Image Processing*, 6(2):184-195, April 1977

2. S. Pizer, E. Amburn, J. Austin, R. Cromartie, A. Gesselowitz, B. ter Haar Romeny, J. Zimmerman, and K. Zuiderveld, "Adaptive Histogram Equalization and its Variations," *Computer Vision, Graphics, and Image Processing*, 39(3):355-368, September 1987

3. R. M. Gray, "Vector Quantization," *IEEE ASSP Magazine*, 1:4-29, April 1984

4. A. Gersho and R. M. Gray, Vector Quantization and Signal Compression, Kluwer Academic Publishers, Boston, 1992

5. A. Buzo, A. H. Gray Jr., R. M. Gray, and J. D. Markel, "Speech Coding based upon Vector Quantization," *IEEE Transactions on Acoustics, Speech and Signal Processing*, 28:562-574, October 1980

6. P. C. Cosman, E. A. Riskin, and R. M. Gray, "Combining Vector Quantization and Histogram Equalization," Proceedings Data Compression Conference, J. A. Storer and J. H. Reif, editors, pages 113-118, IEEE Computer Society Press, Snowbird, Utah, April 1991

7. P. A. Chou, T. Lookabaugh, and R. M. Gray, "Optimal Pruning with Applications to Tree-structured Source Coding and Modeling," *IEEE Transactions on Information Theory*, 35(2):299-315, March 1989