SIMULTANEOUS COMPRESSION AND DENOISING OF IMAGERY USING NON-LINEAR INTERPOLATIVE VECTOR QUANTIZATION

Kannan Panchapakesan, Ali Bilgin, David G. Sheppard, Michael W. Marcellin, Bobby R. Hunt

Department of Electrical and Computer Engineering, The University of Arizona, Tucson, AZ 85721.

{kannan,bilgin,sheppard,mwm,hunt}@ece.arizona.edu

http://www-spacl.ece.arizona.edu

ABSTRACT

Quantization of a noisy source is a classic problem. Although, optimal solutions have been shown to exist under certain assumptions, they are usually either extremely difficult to implement or computationally very intensive. The vector quantization (VQ) based strategy we present in this paper is simple, efficient, and capable of joint compression and denoising of images corrupted by additive white Gaussian noise of fixed variance. The VQ is trained on pairs of clean original images and their respective noisy versions. This VQ is then capable of taking a noisy image as input to its encoder, compressing it, and then producing a less noisy image at the output of its decoder. Experiments performed on test images (inside and outside the training set) produced significant denoising at various bit rates. Results also indicate that our system is capable of handling a wide range of noise variances, while designed for a particular noise variance.

1. INTRODUCTION

Vector Quantization has been widely used in the compression of images [1]. In the past few years, VQ has also been used to perform various image processing tasks concurrently with compression [2]. This paper introduces a joint compression and denoising technique based on non-linear interpolative vector quantization (NLIVQ) [3]. The training procedure for the VQ is non-iterative, discrete cosine transform based, and computationally efficient.

2. QUANTIZATION OF A NOISY SOURCE

Consider Figure 1 where X is the clean original signal, N is additive white Gaussian noise (AWGN), Y the noisy version of X, and $Q(Y)$ the quantized version of Y. The goal in this case is to minimize

$$E[d(X, Q(Y))],$$

for a given distortion measure d. This is in fact a classic problem referred to in literature as the noisy source coding problem and worked on by many, including Dobrushin and Tsybakov [4], Fine [5], Sakrison[6], Wolf and Ziv [7], Ephraim and Gray [8], and Ayanoğlu [9].

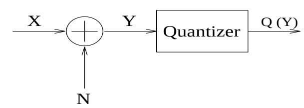

Figure 1: Quantizing a noisy source.

It has been suggested in the literature [4, 5, 6, 7, 8, 9] that for the squared error distortion measure, the optimal solution to the problem is to cascade the optimal estimator for X given Y followed by the optimal quantizer for the estimate. However, there are two practical concerns associated with implementing this solution. First, the optimal estimator is in general not known or may be extremely complex and second, realizing the operations of estimation and quantization separately can be computationally inefficient. Therefore, from an implementation point of view, a trade of optimality for simplicity and computational gain might be in order. One straightforward way to achieve this is to combine the operations of estimation and quantization in one step. This was considered by Rao *et al.*, who implemented an efficient but sub-optimal VQ system by imposing structural constraints on the encoder while using deterministic annealing to design the VQ [10].

Non-linear interpolative vector quantization, the basis of our system, provides a simple strategy to achieve jointly sub-optimal estimation (denoising) and quantization. This, in conjunction with a unique DCT-based, non-iterative VQ design procedure, results in our system being simple, computationally efficient, and yet robust.

This work was supported in part by the United States Air Force, the Air Force Office of Scientific Research (Grant No. F49620-97-1-0099), and the National Science Foundation (Grant No. NCR-9258374).

3. NON-LINEAR INTERPOLATIVE VECTOR QUANTIZATION

Vector quantization is simply the quantization of a vector, an ordered set of real numbers. A vector quantizer, Q, with an associated codebook, C, of size K, consisting of k-dimensional vectors can be thought of as the mapping

$$Q : R^k \rightarrow C.$$

The VQ just described is said to have a rate of $r = \frac{log_2 K}{k}$ bits/dimension. To design such a VQ even at moderately high rates can be a daunting task.

NLIVQ [3] is a technique that circumvents the complexity barrier associated with the design procedure in a conceptually simple fashion, that can be explained as follows. Let Z be a random vector of dimension k. As k increases, ordinary full search VQ quickly becomes infeasible. But if it is possible to extract a suitable "feature vector", U, of dimension $n < k$, then Z can be estimated from the vector quantized version, $\hat{U}$, of U. This estimation process of Z from $\hat{U}$ can be accomplished in one step by designing the interpolative decoder such that it is optimal for a given encoder. It may be noted here that the encoder and decoder codebooks would be of the same size but of dimensions n and k respectively. It is important to note here that the choice of the feature vector is dictated entirely by the needs of the application.

In [11], Sheppard *et al.* introduced a non-linear interpolative vector quantizer for image restoration. The design strategy we have adopted for denoising, as in Figure 2, is very similar in spirit and can be described as follows. Let $\{X^i, Y^i\}_{i=1}^N$ be a sequence of pairs of clean original and noisy training images, respectively. Decompose each image into $M \times M$ non-overlapping training blocks. Let x^{ij} and y^{ij} be jth blocks from X^i and Y^i respectively. Assume that the encoder $\mathbf{E}$, decoder $\mathbf{D}$, and the associated encoder codebook C, are given for a VQ that minimizes the mean-squared error

$$\text{MSE} = E\| y^{ij} - \tilde{y}^{ij} \|^2. \quad (1)$$

The quantized block $\tilde{y}^{ij}$ is chosen as

$$\tilde{y}^{ij} = \mathbf{D}(\mathbf{E}(y^{ij})) = \arg\min_{c_l \epsilon C} \| y^{ij} - c_l \|^2, \quad (2)$$

where c_l refers to the lth entry of C.

Next, a new decoder $\mathbf{D}^*$ and its associated codebook C^* is derived by minimizing the conditional expectation

$$E[\| x^{ij} - \tilde{x}^{ij} \|^2 | \mathbf{E}(y^{ij}) = l], \quad (3)$$

where encoder $\mathbf{E}$ returns the index of the optimal codebook entry. For a given set of training blocks, let $R_l = \{x^{ij} : \mathbf{E}(y^{ij}) = l\}$. Define entry l of C^* as the centroid of R_l, or

$$c_l^* = \frac{1}{|R_l|} \sum_{x^{ij} \epsilon R_l} x^{ij} \quad (4)$$

where $|R_l|$ denotes the cardinality of R_l. Finally, the NLIVQ denoising algorithm is given by

$$\tilde{x}^{ij} = \mathbf{D}^*(\mathbf{E}(y^{ij})) = c^*_{\mathbf{E}(y^{ij})}, \quad (5)$$

where $\tilde{x}^{ij}$ is the denoised image block.

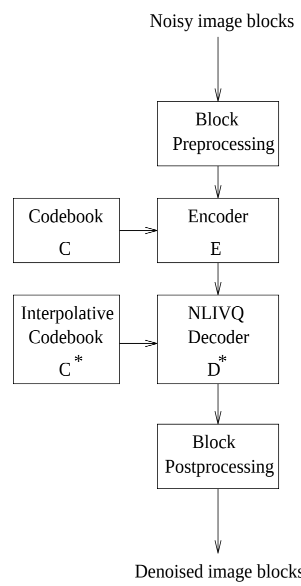

Figure 2: NLIVQ based denoising strategy.

4. ALGORITHMS FOR CODEBOOK DESIGN

Designing the encoder codebook C is the central issue while the decoder codebook C^* is simply derived from C, as discussed above. The most commonly used algorithm in VQ design is the Lloyd algorithm [12]. But this simple design algorithm is iterative and therefore comes with heavy computational requirements which limits its use to low encoding rates. This motivated us to use the DCT-based, non-iterative technique introduced in [11].

DCT-Based Encoder Design

- Add AWGN of a fixed variance, σ^2, to each of the N clean training images, X^i, to produce the noisy training images, Y^i.
- An image in the training set is first divided into non-overlapping blocks of size $M \times M$. Then the DCT is performed on each of the blocks to produce x^{ij} and y^{ij} corresponding to the clean and noisy training images, respectively.
- Given a budget of R bits/pixel, allocate the available $L = R \times M \times M$ bits for each block, y^{ij}, such that the quantization error is minimized.

- A pdf-optimized scalar quantizer is designed for each of the M^2 DCT coefficients of y^{ij} according to the rate allocation discussed above. Gaussian and Laplacian distributions are assumed for DC and AC coefficients, respectively.
- Finally, the codeword index from the encoder is generated by concatenating the binary codes from each of the scalar quantizers employed in a given block.

NLIVQ Decoder Design

- Compute the encoder index, $E(y^{ij}) = q$, as defined above, for each noisy DCT block, y^{ij}.
- Add the corresponding clean DCT block, x^{ij}, to the accumulator a_q^* and increment the counter s_q^*.
- Once all of the training blocks have been consumed, each codeword in C^* is computed as the average

$$c_q^* = \frac{1}{s_q^*} a_q^*.$$

5. RESULTS

Simulations were performed using the algorithms described above. The training set for VQ design consisted of 53 grayscale "urban" images, each of size 512×512. The signal-to-noise ratio(SNR), measured always with respect to the clean image, is defined as

$$\text{SNR} = 10\, log_{10} \frac{\text{signal power}}{\text{noise power}} \text{ dB}.$$

Training blocks of size 2×2 were employed which resulted in 4-dimensional training vectors, with bits being allocated to each of the DCT coefficients depending on the desired overall bit rate. For example, to realize a bit rate of 2 bpp, a total of 8 bits were distributed among the 4 DCT coefficients.

Table 1 displays results for a test image (outside the training set) that was corrupted with AWGN of variance 400. This noise variance equates to an SNR of 17.02 dB. VQ $_{noisy}$ is a vector quantizer designed to minimize quantization error without any explicit attempt to incorporate denoising. This quantizer was trained using noisy images at both encoder and decoder. This system does not result in any significant denoising. The next column shows the performance of VQ $_{clean}$, which was similarly trained with clean images at encoder and decoder. When this VQ is used to compress the noisy test image, it effects moderate denoising but its performance is limited by the fact that it has no knowledge of the noise.

The columns labeled NLIVQ $_{NA}$ and NLIVQ $_{CA}$ are non-linear interpolative vector quantizers designed as explained in the previous section (using noisy and clean images at the encoder and decoder, respectively) with the bit allocations at the encoder based on noisy and clean image characteristics, respectively. That NLIVQ $_{CA}$ outperforms NLIVQ $_{NA}$ is intuitively reasonable. In the latter case, the noise biases the bit allocation procedure to artificially emphasize high frequencies.

At a rate of 1.0 bpp, it may be noticed that VQ_{clean} actually provides a better SNR value than NLIVQ $_{NA}$. This can be attributed to the fact that the quantizers designed in these simulations are only locally optimal. Overall, the NLIVQ $_{CA}$ system performs the best, providing SNR gains of 2.43 dB and 1.28 dB at rates 2.0 bpp and 1.0 bpp, respectively.

In order to investigate the robustness of the NLIVQ $_{CA}$ system, the following simulations were performed. First, the NLIVQ $_{CA}$ system was designed for a noise variance of 400 at 2 bpp. Then, this system was used on the test image corrupted by noise of variances 200, 400, and 800 (corresponding to SNRs of 20.01 dB, 17.02 dB, and 13.98 dB), respectively. The results are shown in Table 2.

Next, an NLIVQ $_{CA}$ system was designed for each of the above noise variances, and used on the test image corrupted by noise with the same variance. These results are given in Table 3. To quantify the loss in not using an appropriately designed system, we compare the results in Tables 2 and 3. It can be seen that for a noise variance of 200, only a 0.33 dB loss is incurred, while for a noise variance of 800, this loss is as little as 0.07 dB. This indicates that the system is robust and is capable of handling different noise variances than it was designed for.

Table 1: SNR(dB) of decompressed images at different rates(bits/pixel). SNR of noisy image = 17.02 dB.

R	VQ $_{noisy}$	VQ $_{clean}$	NLIVQ $_{NA}$	NLIVQ $_{CA}$
2.0	17.02	18.37	19.01	19.45
1.0	16.70	18.22	17.43	18.30

Table 2: SNR(dB) at 2 bits/pixel from a system designed for a noise variance of 400 but tested on noise variances of 200, 400, and 800, respectively.

Noise variance	Noisy image	NLIVQ $_{CA}$ image
200	20.01	20.65
400	17.02	19.45
800	13.98	17.83

Table 3: SNR(dB) at 2 bits/pixel from a system designed for and tested on noise variances of 200, 400, and 800, respectively.

Noise variance	Noisy image	NLIVQ $_{CA}$ image
200	20.01	20.98
400	17.02	19.45
800	13.98	17.90

6. REFERENCES

[1] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic, Norwell, Massachusetts, 1992.

[2] P. C. Cosman, K. L. Oehler, E. A. Riskin, and R. M. Gray, "Using vector quantization for image processing", *Proc. of the IEEE*, vol. 81, no. 9, pp. 1326–41, September 1993.

[3] A. Gersho, "Optimal non-linear interpolative vector quantization", *IEEE Trans. Commun.*, vol. 38, no. 9, pp. 1285–87, September 1990.

[4] R. L. Dobrushin and B. S. Tsybakov, "Information transmission with additional noise", *IEEE Trans. on Inform. Theory*, vol. 8, pp. 293–304, 1962.

[5] T. Fine, "Optimum mean-square quantization of a noisy input", *IEEE Trans. on Inform. Thoery*, vol. 11, pp. 293–294, April 1965.

[6] D. Sakrison, "Source encoding in the presence of random disturbance", *IEEE Trans. on Inform. Theory*, vol. 14, pp. 165–166, January 1968.

[7] J. K. Wolf and J. Ziv, "Transmission of noisy information to a noisy receiver with minimum distortion", *IEEE Trans. on Inform. Theory*, vol. 16, no. 4, pp. 406–411, July 1970.

[8] Y. Ephraim and R. M. Gray, "A unified approach for encoding clean and noisy sources by means of waveform and autoregressive model vector quantization", *IEEE Trans. on Inform. Theory.*, vol. 34, no. 4, pp. 826–834, July 1988.

[9] E. Ayanoğlu, "On optimal quantization of noisy sources", *IEEE Trans. on Inform. Theory*, vol. 36, no. 6, pp. 1450–1452, November 1990.

[10] A. Rao, D. Miller, K. Rose, and A. Gersho, "A generalized VQ method for combined compression and estimation", *Proc. IEEE Intl. Conf. Acoustics, Speech, and Sig. Proc.*, pp. 2032–2035, 1996.

[11] D. G. Sheppard, A. Bilgin, M. S. Nadar, B. R. Hunt, and M. W. Marcellin, "A vector quantizer for image restoration", *IEEE Trans. Image Proc.*, vol. 7, no. 1, pp. 119–124, January 1998.

[12] Y. Linde, A. Buzo, and R. M. Gray, "An algorithm for vector quantizer design.", *IEEE Trans. Commun.*, vol. 28, pp. 84–95, January 1980.