Insensitivity of Adaptive Quantization to Model Estimation in Wavelet Subband Coding

Rebecka Jörnsten* Bin Yu*†

Abstract

State-of-art wavelet image coders achieve similar rate-distortion performance with different model estimation methods to adapt to the changing characteristics of the subband. We provide here two theoretical analyses on backward and forward adaptations to explain this insensitivity to model estimation. Backward adaptation is parametric, uses quantized data, but does not incur overhead. Forward adaptation is nonparametric, uses continuous data, and incurs overhead. Estimation optimal and quantization optimal adaptive quantizers are shown to be close for generalized Gaussian distributions which well approximate the subband data of images. An empirical process CLT is obtained to show that the nonparametric uniform quantizer is comparable to its parametric counterpart. The two quantizers are then seen to give overall similar performance in experiments with simulated and real subband data. However, at low bit rates, the parametric approach is to be preferred, as done in practical coders.

Index terms; Quantization, non-parametric estimation, Fisher information, Generalized Gaussian random variables, entropy constraint, empirical process.

*Department of Statistics, University of California at Berkeley.

†Bell laboratories, Lucent Technologies, Murray Hill, NJ.

1 Introduction

In wavelet subband coding, scalar quantization is an indispensable tool to reducing coding cost, or rate. Quantization creates distortion, often measured by Mean Squared Error (MSE), minimized by the Lloyd-Max quantizer (cf. k-means), which uses finer partitioning where the probability mass is concentrated. Wavelet coefficient statistics vary locally, even within a subband, and adaptation to these local characteristics becomes vital for optimum subband coding performance. Many diverse adaptive quantization methods have been used in recent and successful wavelet coders. Although their approaches to the problem at hand differ greatly, they yet provide comparable distortion performance. We present here two statistical analyses to show why adaptive quantization is insensitive to model assumptions and estimation.

Models used for adaptation may be parametric or non-parametric. The former is defined through local estimates of model parameters, the latter through density estimation, or direct estimation of the optimal quantizer, on local batches of data. Adaptive procedures can also be classified as *forward* or *backward*. Forward approaches constitute any adaptive method based on non-causal information. This form of adaptation therefore requires overhead. Backward adaptation relies solely on causal information, i.e. previously quantized data. Mixtures of backward and forward adaptation schemes are often used in practical coders (cf. [11], [22]).

The diversity and success of adaptive methods suggest an insensitivity of adaptive quantization. We observe in the backward and parametric adaptation, different models and estimation techniques give similar distortion performance (cf. [11], [13], [22], [24], and references therein). In [11] the quantizer estimate is updated on a pixel-by-pixel basis, an "infinite mixture" of quantizers in the words of the authors. In [22] pixel-by-pixel model parameter estimates are mapped, via a context variable, into classes of subband coefficients defining a finite set of quantizers. Both techniques give similar and state-of-the-art results.

Approximations to the optimal quantizer, such as Uniform Threshold Quantizers (UTQ), have demonstrated competitive performance (cf. [11], [19]). Approximating classes of parametric models vary from Generalized Gaussian (GG) models adaptive in both scale and shape parameters, to GG models with fixed shape parameter, and to Laplacian models. The particular choice of model however, has been observed to have little impact on distortion performance ([2]). Nevertheless. we note that although a rough approximation to the source distribution may give adequate results for quantization purposes, better models are still important for efficient coding of quantized wavelet coefficients. In [13] piece-wise linear approximations of model densities are used, with competitive results. In [24] simulation studies show that even crude density approximations, far from accurately describing the true source distribution, induce close to optimal quantizer designs. Indeed, adaptive quantization seems to be insensitive to model mis-specification.

Two statistical analyses are presented here. We investigate parametric adaptive quantization using GG models, and backward adaptation. Local parameter estimates are based on quantized past and the quantizer update uses these estimates to adapt to local characteristics of subband data. For efficient and high-performing adaptation, it is necessary that quantizer design should not only provide minimum distortion, but that good parameter estimates can be obtained from the output of the quantizer for the next batch of data. Poor parameter estimates could seriously hamper speedy adaptation of parametric quantizers. Fortunately, we find that distortion optimal quantizer designs are also optimal, or close to optimal, for efficient parameter estimation for the class of GG models.

With non-parametric adaptive quantizer designs and forward adaptation, model mismatch is avoided at the added coding cost of overhead. The non-parametric designs necessarily provide better distortion performance than their parametric counterparts since they minimize the MSE for the observed data directly. However, the added cost of overhead for transmitting the quantizer may not be compensated for by a sufficient reduction of distortion. Obtaining good local non-parametric density estimates may be difficult due to small sample variability. However, we find that direct non-parametric estimation of the *quantizer* is "easier" than model estimation, and these estimates have comparable efficiency to parametric estimates.

The paper is divided as follows. Section 2 presents the case of parametric adaptive quantization. We show that estimation optimal and distortion optimal quantizer designs are close, or identical, for the class of GG models. In section 3, non-parametric adaptive quantization is discussed. It follows from results of Pollard [14] on k-means clustering that quantizer estimates have parametric rate, and can therefore be considered to be "easier" to obtain than estimates of the source distribution itself. We show that a similar CLT also holds for the non-parametric UTQ estimate, making it comparable in rate with its parametric counterpart. Section 4 contains experimental results. Quantizer estimation efficiencies of parametric and non-parametric methods are compared, as are rate-distortion performances. We find that the two methods give similar performances, although parametric methods are preferable for low-rate quantization. Section 5 contains concluding remarks.

2 Parametric Adaptive Quantization

In a parametric setting, adaptation of the model induced quantizer is achieved through local estimation of model parameters (cf. [2], [11], [22]). Backward adaptation, i.e. free of overhead, requires that we base the parameter estimates on past quantized data only. If we denote the string of previous n data by $X^n = (X_1, X_2, \cdots, X_n)$, backward adaptation is equivalent to observing random variables of indicator type, $1_{\{X_i \in B_j\}}$, where $i = 1, \cdots, n$ and $j = 1, \cdots, L$, L is the number of levels of the quantizer, and B_j denotes the j-th quantizer bin. We frequently treat subband data as i.i.d. for estimation and coding purposes. This assumption is supported by the decorrelating properties of wavelet transforms, and by blocking subband data into neighborhoods, within which similar statistical characteristics are observed. The method of moments matches bin-probabilities to observed relative frequencies. It has been used, as well as the maximum likelihood method, to obtain parameter estimates successfully in subband coding (cf. [11], [23]). Since the maximum likelihood method is statistically efficient, we will focus on it in what follows.

The most frequently used parametric models in subband coding are Generalized Gaussian (GG) models, parameterized via a shape parameter ν, and a scale parameter λ. The GG density function is given by

$$p_{\lambda,\nu}(x) = \frac{\lambda}{2}\ h(\nu)\ \ \exp(-[\lambda g(\nu)|x|]^{\nu})$$

$$h(\nu) = \frac{\nu}{\Gamma(1\nu)}\frac{\Gamma(3/\nu)}{\Gamma(1/\nu)},\ g(\nu) = \sqrt{\frac{\Gamma(3/\nu)}{\Gamma(1/\nu)}}.$$

Adaptation of the model may be via the scale parameter alone, or both parameters (cf. [2], [11]). Here we shall present the case of adaptation via local estimates of the scale parameter, assuming the shape parameter is either known, or estimated and sent as overhead for each subband. Given a local estimate of the scale parameter λ, we can design an optimal quantizer such that the resulting MSE distortion, under the estimated distribution, is minimized. The MSE distortion of an L-level quantizer is given by

$$MSE = \sum_{j=1}^{L}\int_{b_{j-1}}^{b_j}(x-c_j)^2 p_{\lambda,\nu}(x)dx,$$

where $b_j,\ j=1,\cdots,L-1,\ b_0=0, b_L=\infty$ define the partition, and reconstruction levels c_j are the centroids of each bin $B_j=[b_{j-1},b_j)$. The optimal quantizer minimizing distortion is given by the Lloyd-Max, or nearest neighbor condition (cf. [5]);

$$b_j = \frac{(c_j + c_{j+1})}{2},$$

or equivalently

$$(c_j - b_j)^2 = (c_{j+1} - b_j)^2.$$

The optimal partition is a function of the model parameters, and a good parameter estimate is necessary for an optimal design quantizer performance. Assume we observe data of the form $1_{\{X_i \in B_j\}}$, where $B_j = [b_{j-1}, b_j), j = 1, \cdots, L$ and $i = 1, \cdots, n$. Using maximum likelihood we can obtain local estimates of the scale parameter based on a string of past quantized data. The log-likelihood is given by

$$l_q(\lambda) = \sum_{j=1}^{L} n_j \ln P_j(\lambda), \ \sum_{j=1}^{L} n_j = n,$$

where $n_j = \sum_i 1_{\{X_i \in B_j\}}$, $P_j(\lambda)$ is the j-th model bin-probability. Numerical maximization of the above provides an estimate of the scale parameter, and the estimation variance is asymptotically the minimal, and given by the Fisher information $I_q(\lambda)$, defined as

$$I_q(\lambda) \equiv -\sum_{j=1}^{L} P_j(\lambda) \frac{d^2}{d\lambda^2} \ln P_j(\lambda).$$

We have the following standard result (cf. [10]);

$$\sqrt{n}(\widehat{\lambda}_{mle} - \lambda) \to N(0, I_q(\lambda)^{-1}).$$

In the case of a Generalized Gaussian distribution

$$I_q(\lambda) = I_c(\lambda) - \sum_{j=1}^{L} VAR[\ \nu\lambda^{\nu-1}(g(\nu)|X|)^{\nu}\ |\ X \in B_j\ ]\ P_j,$$

where

$$I_c(\lambda) = \frac{1}{\lambda^2} + \nu(\nu-1)\lambda^{\nu-2} E[(g(\nu)|X|)^{\nu}]$$

is the Fisher information based on continuous data. We can compare the Fisher information for the estimates based on continuous and quantized data to reveal loss of efficiency due to quantization, here quantified by the second term of I_q. Obviously, the efficiency of the scale parameter estimate depends on the partition used. A question is raised whether we attain efficient estimates using the Lloyd-Max quantizer. Poor parameter estimates will lead to a sub-optimal quantizer design in the adaptive procedure, and could seriously handicap speedy adaptation. However, parameter estimation is simply a tool for finding an approximation of the local data distribution for the purpose of quantizing with minimum distortion. It may be immaterial if this estimate is an optimally efficient one. To answer this question we look at the optimal quantizer for *parameter estimation*, i.e. a quantizer that gives maximum information I_q.

Proposition 2. 1 *The* **estimation-optimal** *L-level quantizer that maximizes the Fisher information $I_q(\lambda)$ of the GG distribution with shape parameter ν, is given by the partition b_j, that satisfies the following condition for $j = 1, \cdots, L-1$;*

$$\left(c_j^{(\nu)} - b_j^{\nu}\right)^2 = \left(c_{j+1}^{(\nu)} - b_j^{\nu}\right)^2,$$

where $c_j^{(\nu)} = \int_{b_{j-1}}^{b_j} |x|^{\nu} \frac{p(x)}{P_j} dx$ is the ν-th conditional moment of bin j.

Proof: Partial derivatives of the Fisher information $I_q(\lambda)$, with respect to the partition $\{b_j\}$ gives

$$\frac{dI_q}{db_j} = \left(\int_{b_{j-1}}^{b_j} |x|^{\nu} \frac{p(x)}{P_j} dx\right)^2 - 2b_j^{\nu} \int_{b_{j-1}}^{b_j} |x|^{\nu} \frac{p(x)}{P_j} dx - \left(\int_{b_j}^{b_{j+1}} |x|^{\nu} \frac{p(x)}{P_{j+1}} dx\right)^2 +$$

$$+2b_j^{\nu} \int_{b_j}^{b_{j+1}} |x|^{\nu} \frac{p(x)}{P_{j+1}} dx = 0$$

at the optimal solution. Completing squares gives

$$\left(\int_{b_{j-1}}^{b_j} |x|^{\nu} \frac{p(x)}{P_j} dx - b_j^{\nu}\right)^2 = \left(\int_{b_j}^{b_{j+1}} |x|^{\nu} \frac{p(x)}{P_{j+1}} dx - b_j^{\nu}\right)^2,$$

and the theorem follows.

Remark 2. 1 *For the GG distribution with shape parameter $\nu = 1$, i.e. Laplacian distribution, the* **estimation-optimal** *and* **distortion-optimal** *quantizers coincide.*

This follows by comparing the expression in Proposition 2.1, with $\nu = 1$, to the Lloyd-Max condition.

The Lloyd-Max quantizer thus only provides us with most efficient estimates in the case of a Laplacian distribution. In subband coding, shape parameters commonly range from $\nu = 0.6$ to $\nu = 2$. Numerical calculations of the estimation and distortion optimal quantizers for GG distributions in this range show a significant discrepancy between the two quantizer designs. However, through simulations we find that the resulting distortion performance of the two quantizers is comparable, as is variance of the scale parameter estimates based on output from the quantizers. Figure 1 (a) shows the distortion performance for the estimation optimal and distortion optimal quantizers in the case of $\nu = 0.6$. As the number of quantization levels increases the discrepancy between estimation optimal and distortion optimal quantizers becomes almost negligible. Figure 1 (b) shows similar results hold for the estimation variance. Both distortion and estimation performances are thus found to be insensitive to reasonable discrepancy in quantizer design.

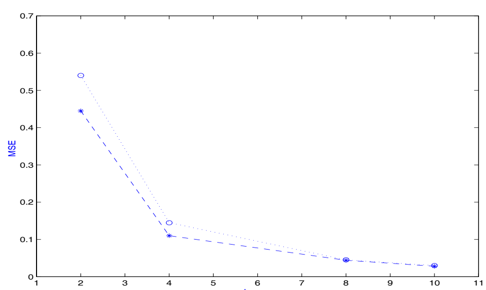

(a)

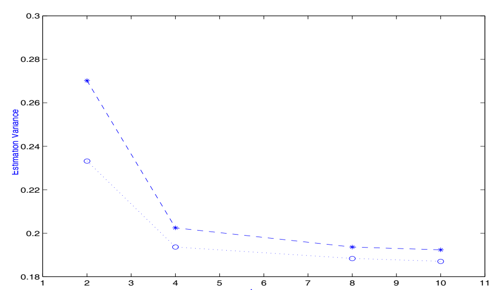

(b)

Figure 1: GG($\nu = 0.6$); (a) Distortion performance of estimation optimal (o) and distortion optimal ($*$) quantizers as function of number of quantizer levels. (b) Estimation variance.

The "zero-zone" quantizes all observations less in magnitude than the first partition point b_1 to zero. This leads to further discrepancy between estimation optimal and distortion optimal quantizers. For simplicity, we will present results for the Laplacian case where in the previous setting no discrepancy existed, as stated above. Keeping the same notation, the Fisher information can be calculated analytically and is given by

$$I_q(\lambda) = \sum_{j=1}^{L} \frac{(b_j - b_{j-1})^2 e^{-\lambda b_j}}{1 - e^{-\lambda(b_j - b_{j-1})}}.$$

The MSE distortion is attained as

$$MSE = 1/\lambda^2 - e^{-\lambda b_1}(b_1 + 1/\lambda)^2 - \sum_{j=2}^{L} \frac{(b_j - b_{j-1})^2 e^{-\lambda b_j}}{1 - e^{-\lambda(b_j - b_{j-1})}}.$$

The Fisher information and MSE distortion differ only in a term involving the partition b_1, i.e. the width of the "zero-zone". As the number of levels of the quantizer increases, this discrepancy becomes negligible. The "zero-zone" discrepancy also carries over to other GG distributions. Through simulations we again find this has little impact on distortion performance and estimation efficiency.

Using a simple uniform threshold quantizer (UTQ) is known to give adequate subband coding performance for most practical applications (cf. [11]), such as low-rate compression. A significant discrepancy between estimation optimal and distortion optimal quantizer designs is again observed, yet again the two provide similar distortion performance. When data compression is the focus of quantization we do not only have the above two criteria in mind, but also want to put an upper bound on the entropy of the output of the quantizer to enable efficient coding. By adding this constraint, we find that the discrepancy as stated above is often eliminated.

Remark 2. 2 *Under entropy constraint, the* **estimation-optimal** *and* **distortion-optimal** *UTQ coincide.*

The UTQ entropy is a single mode function, with a unique maximum (see Figure 2). When an entropy constraint is added, the estimation and distortion optimal quantizers are therefore found as boundary solutions. Other solutions will not satisfy the entropy constraint, or will correspond to a more conservative constraint, thus leading to sub-optimal quantizer designs (cf. [4] for generalized Lloyd-Max). For most entropy constraint the same boundary solution

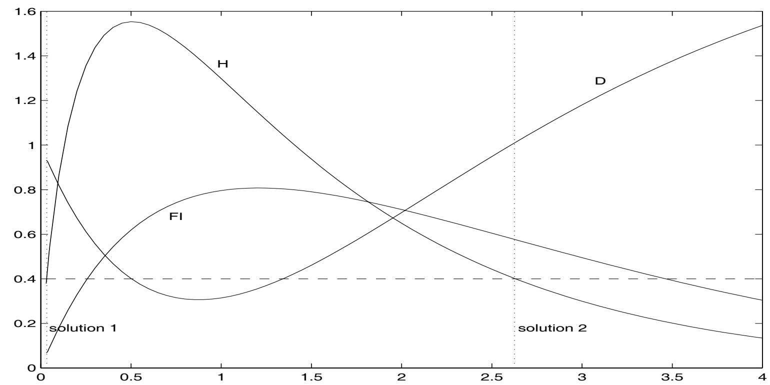

Figure 2: Different boundary solutions are attained for entropy constraint 0.4 for this GG distribution with $\nu = 0.8$ and a 3-level quantizer. D=distortion, FI=Fisher information, H=entropy.

will solve both problems, i.e. the estimation and distortion optimal quantizers will coincide. In practice it is often desirable to obtain convex and continuous rate-distortion curves, e.g. enabling use of off-line look-up tables for quantizer design (cf. [11]). With this added constraint, the two quantizer designs are forced to coincide for all entropy constraint.

It is seen that parametric (GG) adaptive quantization features an insensitivity to significant quantizer design discrepancy. Moreover, Lloyd-Max or UTQ distortion optimal quantizers are also optimal, or near optimal, for efficient parameter estimation, which defines the adaptive procedure.

3 Non-parametric Adaptive Quantization

In a non-parametric setting we may avoid model mis-specification by basing quantizer design on the data directly. The quantizer design needs to be sent as an overhead. There is thus a trade-off between improvement of distortion performance and the coding cost of overhead. Non-parametric quantizer design can take the route of design based on a non-parametric density estimate, followed by computation of the model induced optimal quantizer, or direct estimation of the optimal quantizer for the data at hand. An example of the latter is the commonly used non-parametric Lloyd-Max quantizer, the estimation of which is carried out through iterations between the *centroid condition*

$$\hat{c}_j = \sum_{i=1}^{n} Y_i 1_{\{Y_i \in B_j\}} / \sum_{i=1}^{n} 1_{\{Y_i \in B_j\}},$$

and the *nearest neighbor condition*

$$b_j = (\hat{c}_j + \hat{c}_{j+1})/2.$$

If the iterative algorithm converges, the attained quantizer corresponds to the solution of the k-means clustering problem (cf. [6]) (standard notation, k corresponds to L, the number of quantizer levels). The non-parametric Lloyd-Max quantizer design makes no assumptions on the underlying distribution. However, one can make simplifying assumptions e.g. for the purpose of enabling non-parametric backward adaptation. Simulation studies in [24] and results in [13] show that even very crude approximations on the underlying distribution lead to quantizer designs with competitive performance.

Depending on its smoothness, good estimates of the underlying density can be difficult to get. The difficulty of the estimation is characterized by the minimax rate of estimation. This is of the order $n^{-\alpha}$, $\alpha \in (0, 1)$, where α depends on the smoothness. However, we show that direct estimation of the quantizer achieves the parametric rate n^{-1}. We may interpret this as an indication that direct estimation of the quantizer is an "easier" problem than density estimation. An argument due to Pollard (cf. [15], [16]), in the empirical process theory provides us with the tools to prove this statement, and for finding asymptotic variances of non-parametric quantizer estimates used for efficiency comparisons between parametric and non-parametric methods.

Direct estimation of the optimal quantizer can be phrased as an empirical minimization problem. Recall that $X^n = (X_1, \cdots, X_n)$ denotes a string of i.i.d. data from distribution

P. We use P_n to denote the empirical measure that places probability mass $1/n$ at each observation X_i. The optimal quantizer, with partition b_j and centroids c_j, is the solution to the following

$$min_{b_j,\ c_j}\ \frac{1}{n}\ \sum_i^n (X_i - c_{j(X_i)})^2, \tag{1}$$

where $j(X_i)$ is the index of the bin $B_j = [b_{j-1}, b_j)$ that observation X_i falls into. Denote by

$$f(\cdot,(b,c)) = (x - c_{j(x)})^2,\ \ b = (b_0, \cdots, b_L),\ \ c = (c_1, \cdots, c_L),$$

the functional over which we want to minimize the empirical measure, and we can restate (1) as

$$min_{b,\ c}\ P_n(f(\cdot,(b,c))) = min_c\ P_n(f(\cdot,c)).$$

Note that partition b is defined by the centroids c and vice versa such that the minimization problem above can be reduced to finding the optimal set of centroids that minimize the empirical distortion measure (1), and use the *nearest neighbor* condition to define the data partition. I.e., direct estimation of the optimal L-level quantizer is equivalent to a L-means clustering problem, where L is the number of levels of the quantizer.

Proposition 3. 1 (Pollard, 1982) *Let $\widehat{c}_n$ denote the vector of empirical L-cluster means for sample data strings of size n from distribution P. Under some regularity conditions*

$$\sqrt{n}(\widehat{c}_n - c_0) \rightarrow N(0, \sigma^2),$$

where σ^2 depends on P and L, c_0 corresponds to the true L-cluster means of distribution P.

Proof: See Pollard ([14]).

It follows that the non-parametric quantizer estimation as stated in (1) has convergence rate n^{-1}. This holds for a large class of source distributions (for regularity conditions see appendix) that they themselves may be difficult to estimate.

We can show that the convergence rate of the non-parametric UTQ estimate also has the parametric rate n^{-1}. The goal is to find the minimizing constant bin-width t of the empirical distortion, and this does therefore not correspond to an L-means problem. Define function

$$f(x,t,L) = \sum_{j=1}^{L} |x - c_j(t)|^2 \{x \in B_j(t)\},$$

where

$$B_j(t) = [(j-1)t, jt),\ c_j(t) = \int_{B_j(t)} x dP / \int_{B_j(t)} dP.$$

The optimal UTQ bin-width for data string X^n is found from

$$min_t \ P_n f(\cdot, t, L), \tag{2}$$

and is denoted by $\widehat{t}_n$.

Theorem 3. 1 *Let $\widehat{t}_n$ denote the solution to the empirical minimization problem (2), and t^* denote the minimizer of $F(t,L) = Pf(\cdot,t,L)$. Under some regularity conditions*

$$\sqrt{n}(\widehat{t}_n - t^*) \to N(0, \sigma^2),$$

where $\sigma^2 = P(\Delta^2)V^{-2}, \Delta = \frac{d}{dt} f(\cdot, t)|_{t^*}, V = F''(t,L)|_{t^*}$.

Proof: See Appendix.

The asymptotic variance of the UTQ estimate, σ^2 of Theorem 3.1 can be computed for many distributions and can be used to reveal loss of efficiency due to non-parametric estimation. Note that problem (2) cannot be solved from the data directly since the centroids $c_j(t)$ depend on the unknown distribution P. However, in Appendix we show that we can replace the centroids by their empirical estimates, i.e. solve

$$min_t \ P_n \widehat{f}(\cdot, t, L), \tag{3}$$

where

$$\widehat{f}(\cdot, t, L) = \sum_{j=1}^{L} |x - \widehat{c}_j(t)|^2 \{x \in B_j(t)\},\ \widehat{c}_j(t) = \int_{B_j(t)} x dP_n / \int_{B_j(t)} dP_n.$$

Problems (2) and (3) are asymptotically equivalent, i.e.

$$P_n\hat{f}(\cdot, t, L) = P_n f(\cdot, t, L) + O_P(n^{-1}).$$

Corollary 3. 1 *For solution to the empirical minimization problem (3),* $\hat{t}_n^e$*,*

$$\sqrt{n}(\hat{t}_n^e - t^*) \to N(0, \sigma_e^2),$$

where $\sigma_e^2 = \sigma^2 + C$*.* σ^2 *is the asymptotic variance of solution* $\hat{t}_n$ *to problem (2) and* $C > 0$ *is a shift, the magnitude of which depends on* L *and* P*.*

Proof: See Appendix.

4 Experimental Results

We will compare estimation efficiency and rate-distortion performance of parametric adaptive and non-parametric adaptive quantization (UTQ), on simulated Laplacian data and wavelet subband data from the test image Lena.

The optimal L-level UTQ bin-width for a Laplacian source is found as the minimizer t^* of the MSE;

$$F(\lambda, t) = \frac{2}{\lambda^2} - \frac{t^2(e^{-\lambda t} - e^{-\lambda L t})}{(1 - e^{-\lambda t})^2}.$$

The UTQ bin-width can be seen to be a scalar function of the scale parameter λ, such that $t^*(\lambda) = t^*(1)/\lambda$. Given a ML estimate for λ, based on past quantized data, the optimal UTQ bin-width $\hat{t^*}$ equals $t^*(1)/\hat{\lambda}_{mle}$. By the delta method, we can approximate the variance of the parametric quantizer estimate $\hat{t^*}$ by

$$\sigma_p^2 = t'(\lambda)^2|_{\lambda_0} \; Var(\hat{\lambda}),$$

where $Var(\hat{\lambda}) = 1/nI_q$, I_q is the Fisher information, and λ_0 is the true parameter value.

The non-parametric UTQ bin-width estimate is obtained as the minimizer of (3), $P_n \hat{f}(\cdot, t)$. Appealing to Theorem 3.1, we can approximate the variance of the quantizer estimate by

$$\sigma^2_{np} = V^{-1} P \Delta^2 V^{-1} / n,$$

where $V = F''_{\lambda_0}(t)|_{t^*}$, $\Delta = \frac{d}{dt} f(\cdot, t)|_{t^*}$. For a Laplacian source distribution V and Δ can be computed analytically.

In Figure 3 we compare estimation efficiency of parametric and non-parametric UTQ designs. Simulations and computations were carried out for $L = 8$ and 20 quantizer levels, and for data batches of size $n = 500, 1000$ and 2000. The computed curves σ^2_{np} and σ^2_p reveal the loss of efficiency due to non-parametric estimation. Added to the curves are results $\hat{\sigma}^2_{np}$ and $\hat{\sigma}^2_p$, from simulations using Laplacian data. For parametric estimation, arbitrarily chosen initial quantizers were used and ML estimates of λ obtained, and the induced UTQ was computed. This was repeated several times, giving an estimate $\hat{\sigma}^2_p$ of σ^2_p. It was found that estimation efficiency was robust with respect to the choice of initial quantizer used. The non-parametric results were obtained through repeated minimization of (3) on separate batches of simulated data. The results show a shift of the observations $\hat{\sigma}^2_{np}$ to a level above the curve σ^2_{np}. This additional loss, C of corollary 3.1., decreases as the number of quantizer levels L increases. This term represents the "within-cluster" variation, which obviously decreases as the number of clusters, or levels, grows. Similar results are obtained for parametric and non-parametric quantizer estimation on batches of subband data. The loss of efficiency due to non-parametric estimation is not severe, and parametric rate of convergence is achieved.

In order to compare the rate-distortion performance of parametric and non-parametric adaptive quantization, approximations to coding rates are used. First order entropy is a lower bound of coding rate, defined by observed relative frequencies of the output data of the quantizer;

$$H^n(t, L) = \sum_{j=1}^{L} -P^n_j \log_2 P^n_j,$$

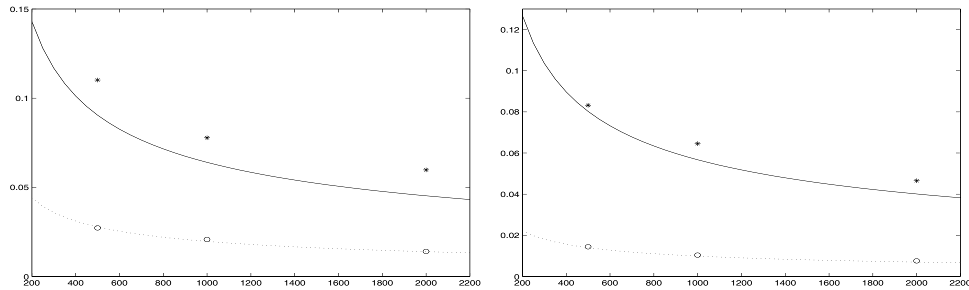

(a) (b)

Figure 3: Laplacian data. σ_{np} (solid) and σ_p (dotted) as functions of batch size, with points of observations $\widehat{\sigma}_{np}$, $(*)$, and $\widehat{\sigma}_p$, (o) added. (a): L=8, (b): L=20.

where $P_j^n = n^{-1} \sum_{i=1}^n 1_{\{X_i \in B_j\}}$, n is the batch size. In the case of non-parametric adaptive quantization, a per-batch overhead penalty needs to be added to the coding cost. If we assume a fixed rate code, truncating the L-level quantizer centroids and bin-width to precision i decimal points, the non-parametric coding rate is approximated by

$$H_{np} = H^n(t, L) + Penalty,\ Penalty = \frac{1}{n}\ (L+1)\log_2 10^i.$$

If a parametric model is assumed, additional coding cost results from model mismatch. This cost is commonly approximated by the discrimination function D_L, which provides a measure of coding difference, or redundancy between the model based ideal code and the source entropy. The approximate parametric coding rate is given by

$$H_p = H^n(t, L) + D_L,\ D_L = \sum_{j=1}^{L} P_j^n \log_2 \frac{P_j^n}{P_j},$$

where P_j is the model bin probability.

For simulated Laplacian data, a theoretical rate-distortion curve is computed as a benchmark for the adaptive procedures. For parametric adaptive quantization of subband data, we choose the best fitting Generalized Gaussian model for each subband. The shape parameter ν for this model is estimated and sent as overhead. For batches of data, parametric and non-parametric adaptive quantizers are computed, as are the approximations of the corresponding coding rates and the resulting MSE distortions. The results are illustrated as rate-distortion curves. The gain of using a non-parametric method must be balanced against the added cost of the overhead, and counter-balanced with the model mismatch of parametric methods. The gain of non-parametric quantization is measured as Signal to Noise Ratio (SNR) improvement per added bits of encoding. SNR is defined as

$$SNR = 10 \log_{10} \frac{Var(X)}{MSE} \; dB,$$

for data string X. Thus for each point on the rate-distortion curve we can compute the gain G as

$$G = \frac{SNR_{np} - SNR_p}{H_{np} - H_p}.$$

A common benchmark used to ascertain if the non-parametric gain is sufficient to motivate added coding cost is 6 dB/bit (cf. [5]).

Using a 3-level 8-Symmlet wavelet expansion of Lena, we find that the HH_1 subband is well approximated by a Laplacian distribution. The fitted shape parameter is $\nu = 1.08$. We choose to adapt the subband quantizer for blocks of size 8×8. The non-parametric quantizer designs are sent as overhead, with the centroids and bin-width estimates truncated to precision .1. The results are illustrated in Figure 4. For simulated Laplacian data (Figure 4 (a)), the parametric model used in designing the quantizer is correct, and there should be little gain in using a non-parametric adaptive quantizer. The rate-distortion curves for non-parametric and parametric methods are close, and the gain is indeed insubstantial at most rates. At very high rates, > 2.5 bpp, the rate-distortion curves do cross, gain exceeding 6 dB/bit. The theoretical rate-distortion curve is added to Figure 4 (a) as a benchmark. For the HH_1 subband data, the parametric model is necessarily an approximation to the true data distribution, albeit a very good one. Figure 4 (b) indicates that the approximation is more than adequate and that parametric methods show comparable performance to the non-parametric. The rate-distortion curves are very close indeed. At low rates (< 0.6 bpp),

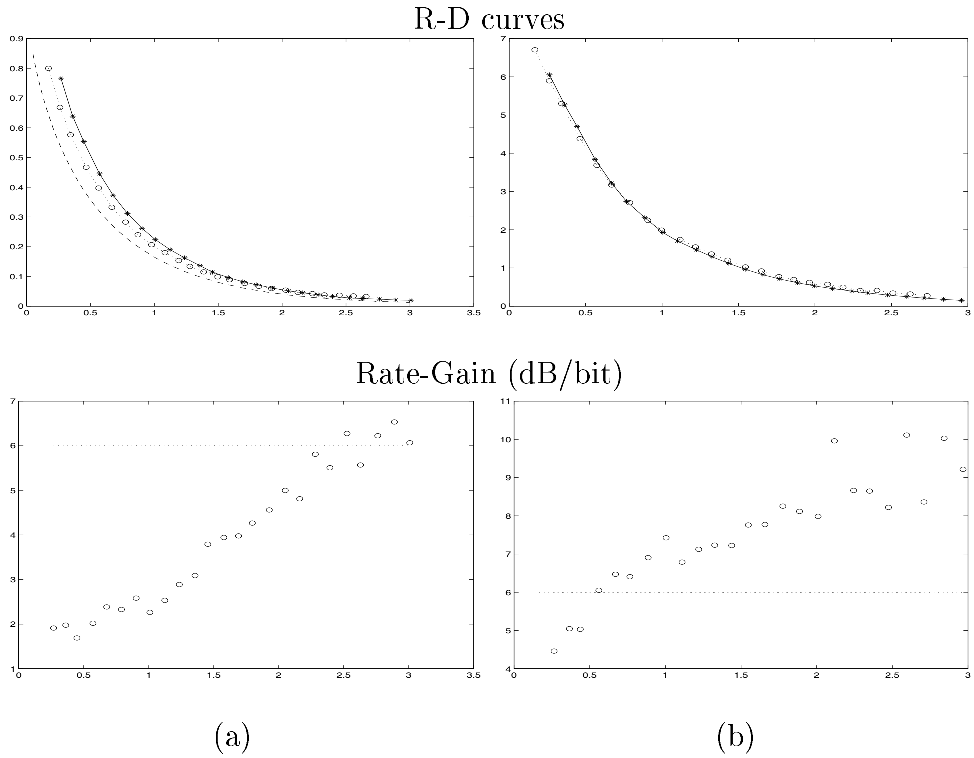

Figure 4: Batch size $n = 8 \times 8$. RD curves and Gain (dB/bit). Parametric method (o), Non-parametric ($*$). a) Laplacian data. b) Subband data.

a parametric approximation to the subband data distribution is more cost efficient, whereas at higher rates a non-parametric method is preferable. To improve distortion performance of the non-parametric quantizer we can refine centroids and bin-width estimates with added precision, at expense of increased cost of overhead. However, we find that added precision does not improve performance sufficiently to compensate the higher coding cost. The gain of using a refined non-parametric quantizer over a parametric design does not exceeding 6 dB/bit at rates < 3 bpp.

Examples of speedy adaptation, using batches or blocks of size 4×4, are illustrated in Figure 5. Cost of using per-batch overhead in non-parametric quantization is now substantial. In the case of simulated Laplacian data (Figure 5 (a)), and HH_1 subband data (Figure 5 (b)), this is clearly shown. The parametric rate-distortion curves show over-all better performance

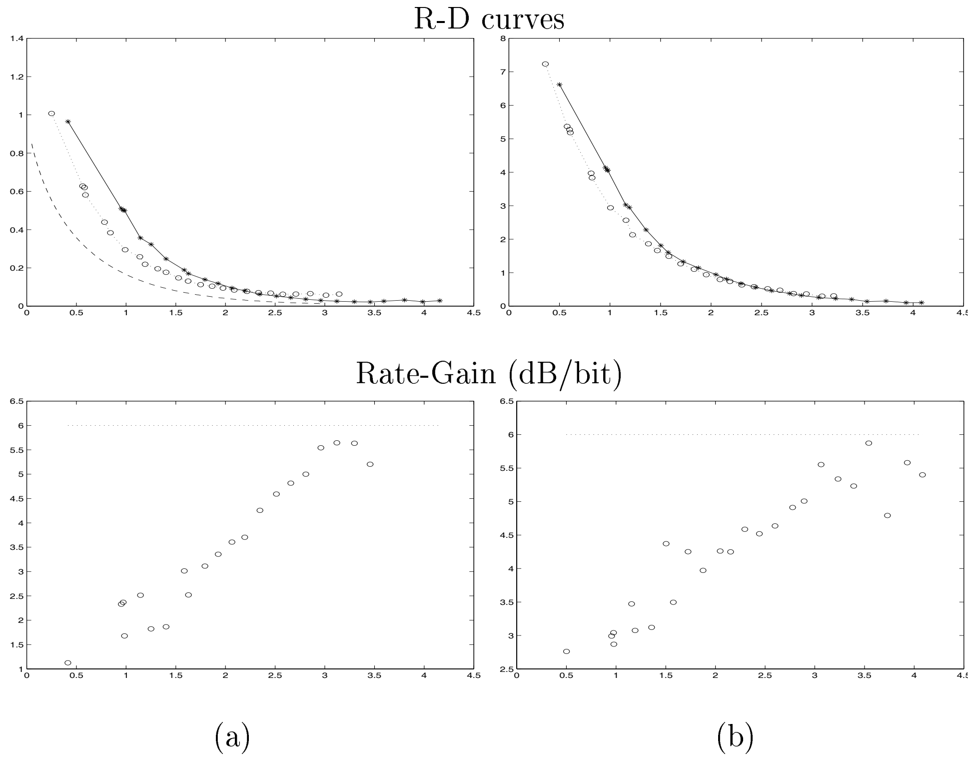

Figure 5: Batch size $n = 4 \times 4$. RD curves and Gain (dB/bit). Parametric method (o), Non-parametric ($*$). a) Laplacian data. b) Subband data.

at rates < 2.25 bpp. Even at higher rates, the gain achieved barely motivates use of non-parametric methods for frequent adaptation.

Results when using larger batches of data, e.g. $n = 32 \times 32$, show that rate-distortion performance of non-parametric methods constitute a considerable improvement on parametric approximations, gains exceeding 6 dB/bit at all rates. Per-batch overhead cost is almost negligible compared to total coding cost when using infrequent adaptation. However, if subband characteristics vary rapidly over smaller regions, adaptation with such large blocks is not advisable.

Other wavelet subbands may be better modeled by Generalized Gaussian distributions with shape parameters < 1. The lower level subbands of Lena deviate substantially from a Laplacian assumption and have very peaked distributions. The HH_3 subband of Lena is best approximated by a GG($\nu = 0.44$) model. Small sample variability of such peaked distributions are more complex than say a Laplacian. Speedy parametric adaptation is therefore somewhat hampered by the deviating characteristics of small samples, i.e. some blocks of data may not be well approximated by the assumed parametric model. Adaptation with blocks of size 8×8 show substantial gains (> 6 dB/bit) of non-parametric methods at all rates (Figure 6 (a)). However, lower level subbands generally have rapidly varying characteristics, and frequent adaptation if often necessary. With blocks of size 4×4, non-parametric adaptation is preferable only for rates > 1 bpp. At low rates, though the parametric approximation to the data batch distribution is poor, it still provides the most satisfying rate-distortion performance.

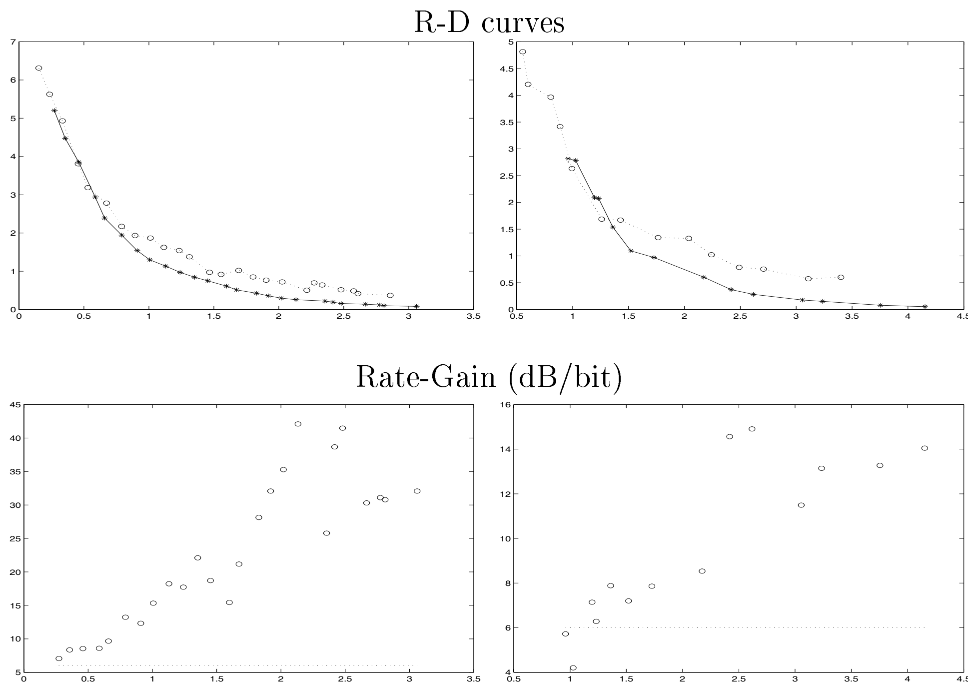

(a) (b)

Figure 6: RD curves and Gain (dB/bit) for HH_3-data. Parametric method (o), Non-parametric ($*$). (a) $n = 8 \times 8$, (b) $n = 4 \times 4$.

5 Concluding Remarks

Adaptation is key in most wavelet subband coding schemes. Locally varying statistical properties can lead to severe increase in coding cost if quantization is not adaptive. We have shown that adaptive quantization is insensitive to model mis-specification and estimation techniques. Parametric (GG) adaptation of quantizers, designed to give minimum distortion for data under the assumed parametric distribution, are also found to be optimal, or close to optimal, for most efficient estimation of local parameters. These local parameter estimates define the adaptive procedure, and efficient estimation enables speedy adaptation, often necessary in wavelet subband coding. Even when the assumed parametric model is far from an accurate description of the underlying source distribution, favorable rate-distortion performance results. This was demonstrated for the case of speedy adaptation of quantizers of the HH_3 subband in the previous section, and similar results are also shown in [24].

Non-parametric methods avoid the problem of model mismatch, but face an increase in coding cost due to overhead. Although non-parametric density estimation can be challenging, we find that non-parametric estimation of the quantizer is an "easier" problem. Non-parametric quantizer estimates achieve parametric rate and are shown to have comparable efficiency to parametric quantizer estimates. Rate-distortion performances of backward parametric and forward non-parametric adaptive quantization are found to be comparable. Rate-distortion curves for both methods, from simulated Laplacian data where the assumed parametric model is correct, and for subband data, show close resemblance. That such diverse methods lead to similar results confirm that adaptive quantization is insensitive to model assumptions and estimation techniques. At low rates and for speedy adaptation, we find that parametric adaptive quantization, although often representing crude approximations to the source distribution, lead to better rate-distortion performance. Parametric methods are also the adaptive procedures of choice in most recent and successful practical low-rate subband coders (cf. [11], [20], [22]).

Appendix

Theorem 3.1 is proved by similar arguments as Pollard in [16]. We will appeal to the empirical central limit theorem for solutions to

$$min_t \; P_n f(\cdot, t, L). \tag{4}$$

Corollary 3.1 relates problem (4) to a problem that is computable from the data directly. We want to show that

$$min_t \; P_n \widehat{f}(\cdot, t, L) \tag{5}$$

is asymptotically equivalent to (4). This follows from

$$\sum_i^n (X_i - \widehat{c}_{j(X_i)}(t))^2 = \sum_i^n (X_i - c_{j(X_i)}(t))^2 + \sum_j N_j (c_j(t) - \widehat{c}_j(t))^2 + 2 \sum_i (X_i - c_{j(X_i)}(t))(c_j(t) - \widehat{c}_j(t)),$$

where N_j is the (random) number of observations in bin $B_j(t)$. Note that $N_j \; \widehat{c}_j(t) = \sum_i X_i 1_{\{X_i \in B_j(t)\}}$ s.t. the above expression reduces to

$$\sum_i^n (X_i - c_{j(Y_i)}(t))^2 + \sum_j N_j (\widehat{c}_j(t) - c_j(t))^2.$$

By CLT the last term is $O_P(1)$ s.t.

$$P_n \widehat{f}(\cdot, t, L) = P_n f(\cdot, t, L) + O_P(n^{-1}).$$

I.e. problems (4) and (5) are asymptotically equivalent.

We proceed to show that the empirical central limit theorem, as stated by e.g. Pollard applies to problem (4) (i.e. Theorem 3.1). Construct a linear expansion of $f(\cdot, t, L)$ near t^*, the minimizer of $Pf(\cdot, t, L) = F(t, L)$;

$$f(\cdot, t) = f(\cdot, t^*) + (t - t^*)' \Delta(\cdot) + |t - t^*| r(\cdot, t, t^*).$$

In the empirical central limit theorem, an asymptotic normality of the sequence of minimizers to problem (4) is deduced from asymptotic normality of $\sqrt{n}(P_n - P)\Delta = E_n\Delta$, as follows by standard CLT, and a control over the remainder term $r(\cdot, t, t^*)$. The empirical CLT states;

Theorem A. 1 Empirical CLT of Pollard. *Suppose t_n is a sequence of minimizers converging in probability to the value t^* at which $F(\cdot, L)$ has its minimum. Define $r(\cdot, t, t^*)$ and $\Delta(\cdot)$ by $f(\cdot, t) = f(\cdot, t^*) + (t - t^*)'\Delta(\cdot) + |t - t^*|r(\cdot, t, t^*)$. If*

1. *t^* is an interior point of the parameter set T;*
2. *$F(\cdot, L)$ has a non-singular second derivative matrix V at t^*;*
3. *$P_n(t_n, L) = o_P(n^{-1}) + \inf_t P_n(t, L)$;*
4. *the components of $\Delta(\cdot)$ all belong to $\mathcal{L}^{2}(P)$;*
5. *the sequence $E_n r(\cdot, t, t^*)$ is stochastically equicontinuous at t^*;*

then $n^{1/2}(t^n - t^*) \to N(0, P(\Delta^2)V^{-2})$.

Proof: See Pollard [16].

We need to show that non-obvious conditions 2,4 and 5 are satisfied for the case of UTQ bin-width estimation, following arguments as presented by Pollard in [16]. Stochastical equicontinuity of $E_n r(\cdot, t, t^*)$ means that with high probability and for all n large enough $E_n r(\cdot, t)$ is uniformly small in some neighborhood close to t^*, close here taken to mean in an $\mathcal{L}^{2}(P)$ sense. It will suffice for our purpose to deduce that the class of graphs of remainder functions has only polynomial discrimination since stochastic equicontinuity then follows. This means that constructing a finite set of approximating functions to $f(\cdot, t, L)$, s.t. $\forall f(\cdot, t, L)$ there is an approximating function $f(\cdot, t_{k,L})$ within ϵ of $f(\cdot, t, L)$, this finite class grows at a less than exponential rate as $\epsilon \to 0$. To prove that the empirical central limit theorem holds for the case of UTQ we assume the following; $P|X|^2 < \infty$ and that the density $p(x)$ is continuous and continuously differentiable at partition points. In the case of UTQ, or any quantizer, both Δ and the remainder function is piece-wise linear, s.t. $\Delta(\cdot) \in \mathcal{L}^2(P)$. We have

$$\Delta(x) = -\sum_{j=1}^{L} 2c_j'(t^*)(x - c_j(t^*))\{x \in B_j\},$$

i.e. all components are piece-wise linear in x. We need to check that the remainder function of the L-level UTQ is uniformly small in the neighborhood of the optimal solution t^*. Using the expression for $\Delta(\cdot)$ above we attain the following for the remainder term in the linear expansion of f. Note

$$|t - t^*|r(x, t, t^*) = f(x, t, L) - f(x, t^*, L) - (t - t^*)\Delta(x) =$$

$$= \sum_{i,j} B_i B_j^* \Big[(x - c_i)^2 - (x - c_j^*)^2 + 2(c_j^*)'(t - t^*)(x - c_j^*)\Big] =$$

$$= B_i B_j^* \Big[\sum_{i,j} 2[(t - t^*)(c_j^*)' + (c_j^* - c_i)]x + (c_i)^2 - (c_j^*)^2 - 2(t - t^*)(c_j^*)'c_j^*\Big],$$

i.e. piece-wise linear in x. Stochastic equicontinuity of the remainder term follows if we can show that the envelope of the remainder term is in $\mathcal{L}^2(P)$. This is so;

$$|r(x, t, t^*)| \leq |t - t^*|^{-1}\Big[|(t - t^*)\Delta(x)| + max_j|(x - c_j)^2 - (x - c_j^*)^2| =$$

$$\leq |\Delta(x)| + (t - t^*)^{-1} \sum_{j=1}^{k} |(x - c_j)^2 - (x - c_j^*)^2| \leq D(1 + |x|)$$

for some constant D, as long as t is sufficiently close to t^*. By assuming a continuous density we can deduce that if $|t - t^*| < \epsilon$ then $|(c_j)^2 - (c_j^*)^2| < \delta|t - t^*|$ for some ϵ and δ positive. The piece-wise linear remainder term is therefore bounded by an envelope in $\mathcal{L}^2(P)$. Graphs of piece-wise linear functions with $\mathcal{L}^2(P)$ bounded envelopes have polynomial discrimination, i.e. the remainder is stochastically equicontinuous (see [16]). This takes care of condition 4 and 5 of the empirical central limit theorem. We still need to show that the second derivative V of $F(t, L) = Pf(\cdot, t, L)$ is non-singular at t^*. The expression for V is complicated and it's non-obvious from it that the condition holds. Though as Pollard argues in the case of

L-means since we have defined t^* as the minimum of F, non-singularity of V, or positive V, should follow. This is not easy to verify. It does hold for the case of Generalized Gaussian distributions, as we can show by numerical calculation. The expression for V,

$$V = \sum_{j=1}^{L} -c_j(2c_j' P_j + c_j P_j') - c_j(2c_j'' P_j + 2c_j' P_j' + c_j' P_j + c_j P_j''),$$

involves first and second derivatives of smooth functionals of the density and it follows we need to assume continuous density and differentiability of the density at the optimal UTQ boundaries, as stated as our criteria above.

Acknowledgements

Part of the results contained in this paper was presented at the IEEE International Symposium on Information Theory (ISIT 98) at MIT, Boston, in August, 1998. This research is partially supported by NSF grants FD98-02314 and DMS-9803063, and ARO grant DAAG55-91-1-0341.

This work was partly undertaken while Ms. Jörnsten was a summer student at Bell laboratories, Lucent Technologies, Murray Hill, NJ.

List of Captions

Figure 1: GG($\nu = 0.6$); (a) Distortion performance of estimation (o) and distortion optimal ($*$) quantizers as function of number of quantizer levels. (b) Estimation variance.

Figure 2: Different boundary solutions are attained for entropy constraint 0.4 for this GG distribution with $\nu = 0.8$ and a 3-level quantizer. D=distortion, FI=Fisher information, H=entropy.

Figure 3: Laplacian data. σ_{np} (solid) and σ_p (dotted) as functions of batch size, with points of observations $\widehat{\sigma}_{np}$ ($*$) and $\widehat{\sigma}_p$ (o) added. (a): L=8, (b); L=20.

Figure 4: Batch size $n = 8 \times 8$. RD curves and Gain (dB/bit). Parametric method (o), Non-parametric ($*$). (a) Laplacian data. (b) Subband data.

Figure 5: Batch size $n = 4 \times 4$. RD curves and Gain (dB/bit). Parametric method (o), Non-parametric ($*$). (a) Laplacian data. (b) Subband data.

Figure 6: RD curves and Gain (dB/bit) for HH_3-data. Parametric method (o), Non-parametric ($*$). (a) $n = 8 \times 8$, (b) $n = 4 \times 4$.

References

[1] T. Berger "Rate distortion theory; a mathematical basis for data compression", Englewood Cliffs, NH: Prentice-Hall, 1971

[2] K. A. Birney, and T. R. Fischer, "On the Modeling of DCT and Subband Image Data for Compression", *IEEE Trans. on Image Proc.*, vol 4, No 2, pp 186-193, Feb. 1995

[3] L. Devroye, "A course in density estimation", Birkhäuser, Boston, 1987

[4] N. Farvardin, J. Modestino, "Optimum Quantizer Performance for a Class of Non-Gaussian Memoryless sources", *IEEE Trans. on Inform. Theory*, vol 30, No 3, May 1984

[5] A. Gersho and R. Gray, "Vector Quantization and Signal Compression", Kluwer Academic Publishers, Boston/Dordrecht/London, 1992

[6] Hartigan J. A. "Clustering Algorithms", Wiley, New York, 1975

[7] R. L. Joshi, V.J. Crump, and T.R Fischer, "Image subband coding using arithmetic coded trellis coded quantization", *IEEE Trans. on Circ. and Sys. for Video Tech.*, vol 5, pp 515-523, 1995

[8] R. L. Joshi, T.R Fischer, and R. Bamberger, "Optimum classification in subband coding of images", in *Proc. of ICIP '94, vol 2, (Austin, Texas)*, pp. 883-887, 1994

[9] R. Jörnsten, and B. Yu, "Insensitivity of Adaptive Quantization to Model Estimation in Subband Coding", *Proc. of ISIT '98*, Cambridge, MA, Aug 1998

[10] E. L. Lehmann, "Theory of point estimation", Springer, 1998

[11] LoPresto, Ramchandran, and Orchard, "Image coding based on mixture modelling of wavelet coefficients and a fast estimation-quantization framework", *Proc. of Data Compression Conference*, 1997

[12] S. Mallat, "Understanding image transform codes", *Proc. of SPIE Conference*, Orlando, April 1997

[13] A. Ortega and M. Vetterli, "Adaptive quantization without side information", *IEEE Trans. on Image Proc.* vol 6, pp 665-676, 1997

[14] D. Pollard, "A central limit theorem for k-means clustering", *Annals of Probab.* 1982, Vol 10, No 4, 919-926

[15] D. Pollard "Convergence of stochastic processes", Springer-Verlag, 1984

[16] D. Pollard "Empirical Processes: theory and applications", *NSF-CBSM Regional conference series in probability and statistics*, vol 2, 1990

[17] K. Ramchandran and M. Vetterli, "Best wavelet packet bases in a rate-distortion sense", *IEEE Trans. on Image Proc.*, vol 2, 160-175, 1993

[18] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients", *IEEE Trans. on Signal Proc.*, vol 41, pp 3345-3462, 1993

[19] G. S. Sullivan, "Efficient scalar quantization of exponential and Laplacian random variables", *IEEE Trans. Inform. Theory*, vol 42, pp 1365-1374, 1996

[20] M. Vetterli, J. Kovačrvić, "Wavelets and Subband coding" , Prentice Hall, Englewood Cliffs, NJ 1995

[21] Z. Xiong, K. Ramchandran, and M. Orchard, "Space-frequency quantization for wavelet image coding", *IEEE Trans. on Image Proc.*, Jan 1997

[22] Y. Yoo, A. Ortega, and B. Yu, "Adaptive quantization of image subbands with efficients overhead rate selection", *Proc. of IEEE Int. Conf. on Image Proc.*, Lausanne, Switzerland, 1996

[23] Y. Yoo, A. Ortega, and B. Yu, "A backward adaptive subband image coder", *IEEE Trans. Image Proc.*, 1997

[24] B. Yu, "A statistical analysis of adaptive quantization based on causal past", *Proc. IEEE Int. Symp. on Inform. Theory and Appl.*, ISITA, Whistler, Canada, 1996

[25] B. Yu, "Rates of convergence for empirical processes of stationary mixing sequences", *Ann. Probab.*, **22** 94-116, 1994