A Statistical Analysis of Adaptive Scalar Quantization based on Quantized Past Data

Bin Yu

Department of Statistics

University of California at Berkeley

Abstract – In this paper, we investigate from several angles the adaptive quantization problem based on quantized causal past data (cf. Ortega and Vetterli, 1994). First, when stationarity and ergodicity are assumed, it is shown that the marginal density can be estimated consistently based on quantized past data in the parametric case. Then, in the non-parametric ase, we analyze the piecewise linear approximation algorithm in Ortega and Vetterli (1994). We characterize the approximate density to which the algorithm leads if convergence is assumed and explains why and why not the algorithm works. Based on Bennett's distortion integral, we argue that estimating the interval partition of the optimal quantizer is an easier problem than density estimation. Moreover, the OV algorithm performs well for this interval partition estimation. Finally, we propose six procedures to improve the OV reproduction level estimates in the low resolution case and simulation studies are carried out in three examples to compare the proposals with the OV algorithm.

Index terms– consistency, limiting density, low-resolution, reproduction level estimation.

I. INTRODUCTION

In signal processing or image compression, (scalar) quantization is used to discretize continuous or high resolution signals (or transformed singals) into digital or low resolution ones. It is necessary because of limits on the memory space and on the rate of transmission of signals or images. For example, in image compression it is common that 8×8 blocks of pixels are transformed using DCT (Discrete Cosine Transform) and the transform coefficients are quantized with dense quantization levels for the low frequencies and coarse for the high frequencies.

The objective of quantization can be formulated as to minimize a certain distortion criterion such as the mean squared error between the continuous and quantized data. When the source distribution is unknown, quantizers are designed based on data. For example, see Jayant (1973), Calderbank, McLaughlin, and Lyons (1992), Gersho and Gray (1992), and Crisafulli and Bitmeed (1993). Such quantizers are called adaptive quantizers. Very often, only quantized data are available for the design of adaptive quantizers due to limits on memory spaces or transmission rates, or due to the non-stationary nature of the signal source. Recently, Ortega and Vetterli (1994) proposed a model-based adaptive scalar quantization algorithms based on the quantized causal past data with convincing simulation results. Inspired by Ortega and Vetterli (1994) and through analyzing their piecewise linear algorithm, we aim in this paper at understanding what is or is not possible in adaptive (scalar) quantization based on quantized past data and at a fixed rate. A key claim is that for a model based adaptive quantization algorithm to work well, it is NOT necessary that the density function in the model is closely approximated (cf. Section IV). We argue that the piecewise linear algorithm of Ortega and Vetterli (1994) is such an algorithm.

Throughout the paper, mean squared error (MSE) distortion measure is used and data range is assumed known. The theorems assume stationarity for technical reasons, but the results are interpretable also for non-stationary slowly varying processes. The outline of the paper is as follows. In Section II, we show that the source density function (hence the optimal quantizer) can be estimated consistently based on quantized past data, provided that the density belongs to a parametric family whose dimension is smaller than the number of quantization levels. In Section III, we deal with the more interesting and practical non-parametric case or when only smoothness is assumed on the density function. We first review the piecewise linear algorithm implemented in Ortega and Vetterli (1994) (OV) and then characterize the limiting density of the OV algorithm when the algorithm is assumed convergent (which is not always the case). Then we make the case that for a density estimate to induce a reasonably good quantizer, it is not necessary that it approximates the true density well. Based on Bennett's distortion integral and an analysis of the OV algorithm we put forward the thesis that estimating the interval partition (of the optimal quantizer) is an easier problem than estimating the whole density function and there the OV algorithm gives a satisfactory performance in both high and low resolutions. The reproduction level estimates of the OV algorithm, however, are responsible for the sometimes poor performance of the OV algorithm in the low resolution case. In Section IV, we propose three types of procedures to improve OV on this aspect. Section V contains three simulation studies to compare the OV algorithm with the proposed procedures.

II. THE CONSISTENCY QUESTION: THE PARAMETRIC CASE

The goal of adaptive quantization is to find a data-based quantizer such that its performance is comparable to that of the optimal quantizer when the source distribution is known. When the adaptation is based on stationary and continuous data, using the empirical distribution in the place of the source distribution leads to a quantizer whose performance approaches that of the optimal quantizer when the sample size gets large. The reason is that the empirical distribution is a consistent estimate of the source distribution. At a fixed rate, however, is it still possible to obtain a consistent estimate of the source distribution (therefore consistent optimal quantizer estimate) based on quantized data? This section gives a positive answer in the parametric case.

Let $X_1, X_2, \ldots$ be a data string taking value on a known interval $[a, b]$. The data string is to be quantized at a fixed level L. Data are grouped into blocks of size T and at end of each block, the quantizer is updated based on the causal past m blocks of quantized data and it is used to quantize the next block. In the stationary case, the larger the block size T and memory window size m, the more stable the adaptive algorithm should be or the faster its convergence. However, without knowing in advance whether the data string is stationary, it is dangerous to make them large since they might lead to inefficient quantizers.

Denote the data blocks by $\underline{X}_j = (X_{(j-1)T+1}, \ldots, X_{jT})$ and its quantized counterpart by $\hat{\underline{X}}_j = (\hat{X}_{(j-1)T+1}, \ldots, \hat{X}_{jT})$. Then an adaptive quantizer at time $j \times T$ $(j = 1, 2, \ldots)$ can be written as

$$\hat{Q}_j = \hat{Q}_j(\hat{\underline{X}}_j, \ldots, \hat{\underline{X}}_{j-max(1,m)+1}),$$

and it is used to quantize and give reproduction levels for the next block $\underline{X}_{j+1}$. We write $\hat{Q}_j = (\hat{R}_j, \hat{C}_j)$ to separate the two roles of a quantizer (cf. Sec. IV), where $\hat{R}_j$ is an interval partition of the data range $[a, b]$ and $\hat{C}_j$ are the associated reproduction level estimates for the intervals in the partition:

$$\hat{R}_j = \{a = b_0 < b_1 < b_2 < ... < b_L = b\},$$

and

$$\hat{C}_j = \{\hat{c}_i(j) : i = 1, 2, ..., L\}.$$

Naturally, we require

$$b_{i-1} \leq \hat{c}_i \leq b_i \quad \text{for } i = 1, 2, ..., L.$$

The following statistical estimation model lends itself to a theoretical analysis on the consistency question. This model is a good approximation to situations where estimated partition intervals $\hat{R}_j$ are stablized (as $j \to \infty$) when a particular adaptive quantizer is used (e.g. the Ortega-Vetterli algorithm to be introduced later), and these intervals don't have to be the optimal intervals based on the unknown marginal source density.

Let an L-level interval partition of $[a, b]$ be $R = \{a = b_0 < b_1 < b_2 < ... < b_L = b\}$, and assume we only observe the quantized causal past N data points, i.e., the indicators $I_{\{b_{i-1} < X_j \leq b_i\}}$ for $i = 1, 2, ..., L$ and $j = 1, 2, ..., N$. Denote by $n_i(N)$ (i=1,..., L) the counts of X's falling into intervals of the partition R:

$$n_i(N) = \sum_{j=1}^{N} I_{\{b_{i-1} < X_j \leq b_i\}}.$$

Assume the input source process (or data string) is stationary with a k-dimensional parametric marginal density $f(x|\theta)$ ($\theta \in R^k$). For $i = 1, 2, ..., k$, denote

$$p_i = p_i(\theta) := \int_{b_{i-1}}^{b_i} f(x|\theta)dx. \qquad (1)$$

Note that only $k-1$ of the equations are independent. The stationarity is needed in the theorem below (and theorems to follow) for technical reasons. The results of the theorems are interpretable when the source is non-stationary but slowly varying, as long as the empirical probabilities $\hat{p}_i$ stablize and the moment equations hold approximately.

Theorem 1 (Consistency) Assume $X_1, X_2, ...$ is stationary and ergodic, and $L = k+1$. If matched with a probability vector $P = (p_1, p_2, ..., p_L)$ ($p_i \geq 0$, $\sum_i p_i = 1$), equations in (1) (for $i = 1, ..., k-1$) have a unique solution $\hat{\theta} = g(p_1, p_2, ..., p_L)$ and g is a continuous function, then $f(\cdot|\theta)$ can be estimated consistently based on quantized data $n_i(N)$.

Proof: For $L = k+1$, any $L-1$ of the moment equations in (1) uniquely determine θ in terms of p_i's: $\theta = g(p_1, ..., p_L)$. By the Ergodic Theorem, when $N \to \infty$, $n_i/N \to p_i$; hence $\hat{\theta} := g(n_1/N, ..., n_L/N)$ tends to the true θ as N gets large because g is continuous. $\square$

When $L < k+1$, there is in general no consistent estimate of $f(\cdot|\theta)$ because "the number of unknowns is larger than the number of equations". When $L > k+1$, we can solve for $\hat{\theta}$ using the k equations corresponding to the k largest n_i/N; or we can minimize $\sum_i (P_i(\theta) - n_i/N)^2$. Both approaches should lead to consistent estimates of θ or $f(\cdot|\theta)$.

A Gaussian model is usually a good choice for modeling the error signals in a predictive quantization scheme (cf. Jayant, 1973). For a Gaussian density with unknown mean and variance and $R = \{-\infty < b_1 < b_2 < \infty\}$ ($L = 3$), the moment equations in (1) become

$$p_1 = \int_{-\infty}^{b_1} f(x|\mu, \sigma^2)dx$$

$$p_2 = \int_{b_1}^{b_2} f(x|\mu, \sigma^2)dx;$$

or equivalently,

$$\begin{aligned} p_1 &= \int_{-\infty}^{b_1} f(x|\mu, \sigma^2)dx \\ p_1 + p_2 &= \int_{-\infty}^{b_2} f(x|\mu, \sigma^2)dx. \end{aligned}$$

where the Gaussian density

$$f(x|\mu, \sigma^2) = N(\mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma}} e^{-(x-\mu)^2/(2\sigma^2)} \quad -\infty < x < \infty.$$

Let Φ be the cumulative distribution function of $N(0, 1)$ and Φ^{-1} be its inverse function:

$$\Phi(x) = \int_{-\infty}^{x} f(t|0, 1)dt.$$

For $\hat{p}_i = n_i/N$, it is easy to obtain the unique solution of the moment equations as:

$$\begin{aligned} \hat{\mu} &= (b_1\Phi^{-1}(\hat{p}_1 + \hat{p}_2) - b_2\Phi^{-1}(\hat{p}_1)/(\Phi^{-1}(\hat{p}_1 + \hat{p}_2) - \Phi^{-1}(\hat{p}_1)) \\ \hat{\sigma} &= (b_2 - b_1)/(\Phi^{-1}(\hat{p}_1 + \hat{p}_2) - \Phi^{-1}(\hat{p}_1)), \end{aligned}$$

which tend to μ and σ as $N \to \infty$.

III. HOW THE OV ALGORITHM WORKS: THE NONPARAMETRIC CASE

In this section we study the performance of the Ortega-Vetterli (1994) piecewise linear algorithm (OV algorithm) in the more interesting and practical non-parametric case where only smoothness is assumed on the unknown source density. First we describe the algorithm; then when the convergence is assumed, we characterize its limiting density and explain why its limiting interval partition is a good partition for the unknown source density. Based on this explanation and also Bennett's distortion integral, we argue that the optimal interval partition is easier to estimate than the density function.

Ortega and Vetterli (1994) divides the adaptive quantization problem based on quantized past data into two parts: density estimation and quantizer update based on the density estimate. This formulation turns out to be conceptually very useful even though there is generally no hope of consistent estimation of the density. As will be seen later, for a density estimate to induce a good quantizer, it is not necessary that this estimate approximates the unknown density well. OV algorithm builds smoothness in the density estimate through linear interpolation. This built-in smoothness helps in estimating the density, but it helps more in getting a good density-induced quantizer, because in the case of smooth densities it is easier to estimate interval partition than to estimate density.

Definition (piecewise linear density) Given an interval partition of $[a, b]$

$$R = \{a = b_0 < b_1 < b_2 < ... < b_L = b\}$$

and a probability vector $P = (p_1, p_2, ..., p_L)$ $(p_i > 0, \sum_i p_i = 1)$, the *piecewise linear* function $f_R(P)$ is the density which takes value zero at a and b, is linear between middle points of the intervals in R, and has interval probabilities as in P. Let F_x be the vector of heights of $f_R(P)$ at the middle-points $x_i = (b_{i-1} + b_i)/2$, then the probability-matching equations

$$\int_{b_{i-1}}^{b_i} f_R(P) dx = p_i \ (i = 1, ..., L)$$

become the following linear equations in terms of F_x:

$$A(R)F_x = 4P/\Delta, \tag{2}$$

where Δ is the interval size vector $(\delta_1, \delta_2, ..., \delta_L)$ with $\delta_i = b_i - b_{i-1}$, and $A(R)$ is a $L \times L$ square matrix with entries $a_{i,j}$ zero except for:

$$a_{1,1} = 2 + \delta_2/(\delta_1 + \delta_2), \qquad a_{1,2} = \delta_1/(\delta_1 + \delta_2),$$

$$a_{L,L-1} = \delta_L/(\delta_{L-1} + \delta_L), \qquad a_{L,L} = 2 + \delta_{L-1}/(\delta_{L-1} + \delta_L),$$

for $1 < i < L$

$$a_{i,i-1} = \delta_i/(\delta_{i-1} + \delta_i), \quad a_{i,i} = 2 + \delta_{i-1}/(\delta_{i-1} + \delta_i) + \delta_{i+1}/(\delta_{i+1} + \delta_i), \quad a_{i,i+1} = \delta_{i+1}/(\delta_{i+1} + \delta_i).$$

Since $\Delta > 0$, $A(R)$ is invertable and the equation (2) has a unique solution

$$F_x = 4(A(R))^{-1}(P/\Delta).$$

The solution F_x, however, can have negative entries in which case truncations to zeros are made and $f_R(P)$ is renormalized to make it a density.

Definition (LM quantizer) Given an interval partition R, a probability vector P and their associated piecewise linear density $f_R(P)$, the Lloyd-Max quantizer $LM_R(f_R)$ is defined as the interval partition the Lloyd-Max algorithm converges to when applied to $f_R(P)$ and when we start with R.

In practice, the LM algorithm is run only for a finite number of iterations. We define the LM quantizer as the interval partition after specified M iterations even if the LM algorithm does not converge.

The OV algorithm

1. Initialization

At time 0, no data is available. Uniform quantizer $\hat{Q}_0 = (R_0, C_0)$ is used to quantize the first T-block of data $\underline{X}_1$:

$$R_0 = \{a = b_0(0) < ... < b_i(0) := a + i(b-a)/L < ... < b_L(0) = b\},$$

$$C_0 = \{(b_0(0) + b_1(0))/2 < ... < (b_{i-1}(0) + b_i(0))/2 < ... < (b_{L-1}(0) + b_L(0))/2\}.$$

2. at the end of the j-th block

We observe counts $n_i(j)$ of data points in the j-th block which fell into partition intervals in R_{j-1}:

$$n_i(j) = \sum_{k \ in \ jth \ block} I_{\{b_{i-1}(j-1) < X_k \leq b_i(j-1)\}},$$

$$\hat{P}^d_j = (\hat{p}^d_1(j), \hat{p}^d_2(j), ..., \hat{p}^d_L(j)) := (n_1(j)/T, n_2(j)/T, ..., n_L(j)/T).$$

For the data in the memory window but not in the current block (the $(m-1)$-blocks before the jth block), we do not observe their counts or frequencies in R_{j-1}. To borrow estimation strength from these blocks, they are re-distributed according to the density last estimated to obtain

$$\hat{p}^l_i(j) := \int_{b_{i-1}(j-1)}^{b_i(j-1)} \hat{f}_{j-1}(x)dx.$$

For $\hat{P}^l_{j-1} = (\hat{p}^l_1(j-1), \hat{p}^l_2(j-1), ..., \hat{p}^l_L(j-1))$, and $\alpha = (min(m,j)-1)/min(m,j)$, let

$$\hat{P}_j := (1-\alpha)\hat{P}^d_j + \alpha \hat{P}^l_{j-1},$$

update the density estimate by $\hat{f}_j := f_{R_{j-1}}(\hat{P}_j)$. The next block $\underline{X}_{j+1}$ is quantized with $\hat{Q}_j = (R_j, C_j)$ where

$$R_j = LM_{R_{j-1}}(\hat{f}_j) = \{a = b_1(j) < b_2(j) < \ldots < b_L(j) = b\}$$

and $\hat{C}_j = (\hat{c}_1, \ldots, \hat{c}_L)$ with

$$\hat{c}_i = \int_{b_{i-1}(j)}^{b_i(j)} x \hat{f}_j(x) dx / \int_{b_{i-1}(j)}^{b_i(j)} \hat{f}_j(x) dx.$$

It is conceivable that $\hat{f}_j$ oscillates (when $j \to \infty$) between two (or more) piecewise linear densities whose LM quantizers and partition intervals are dual to each other, although it has not happened in the simulation studies in Section V. When the OV algorithm converges, however, it stablizes around piecewise linear densities whose quantization intervals are close to those of their own LM quantizers. This density's interval probabilities also match with those of the true density.

Theorem 2 (Limiting Density of OV) Assume that the data string is stationary with a bounded marginal density f and the OV algorithm is updated with "ideal" data (i.e. probabilities without sampling error) $\hat{P}_j^d = P_j(f) = (\int_{b_{i-1}(j-1)}^{b_i(j-1)} f(x)dx : i = 1, 2, \ldots, L)$. If, as $j \to \infty$, the interval partitions R_j in the OV algorithm converges to an interval partition $R = \{a = b_0 < b_1 < \ldots < b_L = b\}$, then $\hat{f}_j$ converges to $\hat{f}$, which is a piecewise linear density with the interval partition R and

$$\int_{b_{i-1}}^{b_i} \hat{f} dx = \int_{b_{i-1}}^{b_i} f dx.$$

Proof: Recall that the probability update equation from the OV algorithm has the following form:

$$\hat{P}_j = (1-\alpha)\int_{b_{i-1}(j-1)}^{b_i(j-1)} f(x)dx + \alpha \int_{b_{i-1}(j-1)}^{b_i(j-1)} \hat{f}_{j-1}(x)dx.$$

The piecewise linear density estimator $\hat{f}_j$ is determined by the height vector

$$F_x(j) = 4A(R_j)^{-1}\hat{P}_j/\Delta_j = (F_x^{(i)}(j) : i = 1, 2, ..., L)$$

and $|\hat{f}_j| \leq max(F_x^{(i)}(j) : i = 1, ..., L)$. Let Δ_j and Δ be the interval size vectors associated with R_j and R, respectively. Then $\Delta > 0$ by assumption, and $\Delta_j \to \Delta$ as $j \to \infty$, hence $1/\Delta_j$ is bounded. From the explicit expressions of entries of A, it is easy to see that $A(R_j) \to A(R)$ as $j \to \infty$. Because $A(R)$ is invertable, it follows that $A(R_j)^{-1} \to A(R)^{-1}$ as $j \to \infty$ and $A(R_j)^{-1}$ is bounded for j large. Moreover, $\hat{P}_j$ is obviously bounded, hence $F_x(j) = 4A(R_j)^{-1}\hat{P}_j/\Delta_j$ is bounded for j large, which implies that $\hat{f}_j$ is bounded.

The obvious guess for the limit of $\hat{f}_j$ is $\hat{f} = f_R(P(f))$ with $P = P(f) = (p_1, ..., p_L)$ and $p_i = \int_{b_{i-1}}^{b_i} f(x)dx$. To prove it is actually the case, it is enough to show that $F_x(j)$ converges to $F_x = 4A(R)^{-1}P/\Delta$. We prove this by contradiction. Without loss of generality, assume $F_x(j)$ converges to another L-vector G_x rather than F_x. (G_x is finite because $\hat{f}_j$ are bounded.) Therefore $\hat{f}_j$ converges to g as $j \to \infty$, where g is the piecewise linear density defined by the height vector G_x.

Letting $j \to \infty$ in the probability update equation of OV, we get

$$\hat{P}_j \to (1-\alpha)P(f) + \alpha P(g),$$

where $P(g)$ satisfies $G_x = 4A(R)^{-1}P(g)/\Delta$. Hence

$$F_x(j) = 4A(R_j)^{-1}\hat{P}_j/\Delta_j \to 4A(R)^{-1}[(1-\alpha)P(f) + \alpha P(g)]/\Delta = G_x,$$

which implies

$$4A(R)^{-1}[(1-\alpha)P(f) + \alpha P(g)]/\Delta = 4A(R)^{-1}P(g)/\Delta.$$

Because $\alpha \neq 1$, $P(g) = P(f)$ and therefore $F_x = 4A(R)^{-1}P(f)/\Delta = 4A(R)^{-1}P(g)/\Delta = G_x$. This contradiction proves that F_x is in fact the limit of $F_x(j)$. Hence $\hat{f}_j \to \hat{f} = f_R(P(f))$ which is a piecewise linear density function and obviously the probabilities match between $\hat{f}$ and f on intervals in R. $\square$.

It is tempting to state that R is also the Lloyed-Max optimal interval partition with respect to $\hat{f}$. A closer examination, however, reveals that it is mathematically possible that a discontinuous phenomenon occurs at R so that the Lloyd-Max optimal interval partition of $\hat{f}$ differs from R. Nevertheless, along the sequence the piecewise density estimate $\hat{f}_j$ based on interval partition R_j and its LM quantizer R_{j+1} are getting closer and closer as j tends to infinity. With sampling errors in $\hat{P}_j$ added, it is therefore reasonable to interpret the above theorem as that the OV algorithm, when it converges, stablizes at a piecewise density on an interval partition such that the partition is close to the LM optimal interval partition induced by this density.

Now we characterized the limiting behavior of the OV algorithm, but how good is the limiting quantizer of the OV algorithm? When L is small, the limiting density of OV can differ a great deal from the unknown source density, as seen in the following two examples. Data strings of iid sources are simulated from two densities and $L = 2$. We chose shapes of the true densities in a way such that the OV piecewise linear density is a poor estimate of the true density. The effective data range is taken as [0,1] (a=0, b=1). (In the second example, with extremely small probability the sample can be outside [0,1].) In both examples, $m = 2$ and "ideal" data blocks (corresponding to $T = \infty$) are used first to be consistent with Theorem 1, that is, true interval probabilities are used in OV instead of the observed empirical interval probabilities. Then we use samples to update OV with $T = 400$. See Figure 1.

In the examples, SNR (dB) is defined as

$$10 \times \log_{10}(Var(X)/D),$$

where D is the MSE distortion measure, which is either in expectation or empirical.

Example 1 Conditioned Exponential (5) on [0,1]

$$f(x) = 5e^{-5x}/(1 - e^{-5}) \;\; x \in [0, 1].$$

Example 2 Mixture of 4 Gaussians

$$f = 0.3 \times N(0.15, 0.05^2) + 0.2 \times (0.7, 0.1^2) + 0.125 \times N(0.3, 0.02^2) + 0.875 \times N(0.5, 0.03^2)$$

From Figure 1 (b) and (d), we observe that convergence of OV occurred in both examples after 5-6 blocks of "ideal" or real data. In the exponential example (cf. (a)), despite the large approximation error (total variation distance being 0.68) of the limiting density to the true density, the limiting partition point of OV agrees with that of the optimal under f to two decimal places! This speaks clearly that the limiting density of OV can approximate

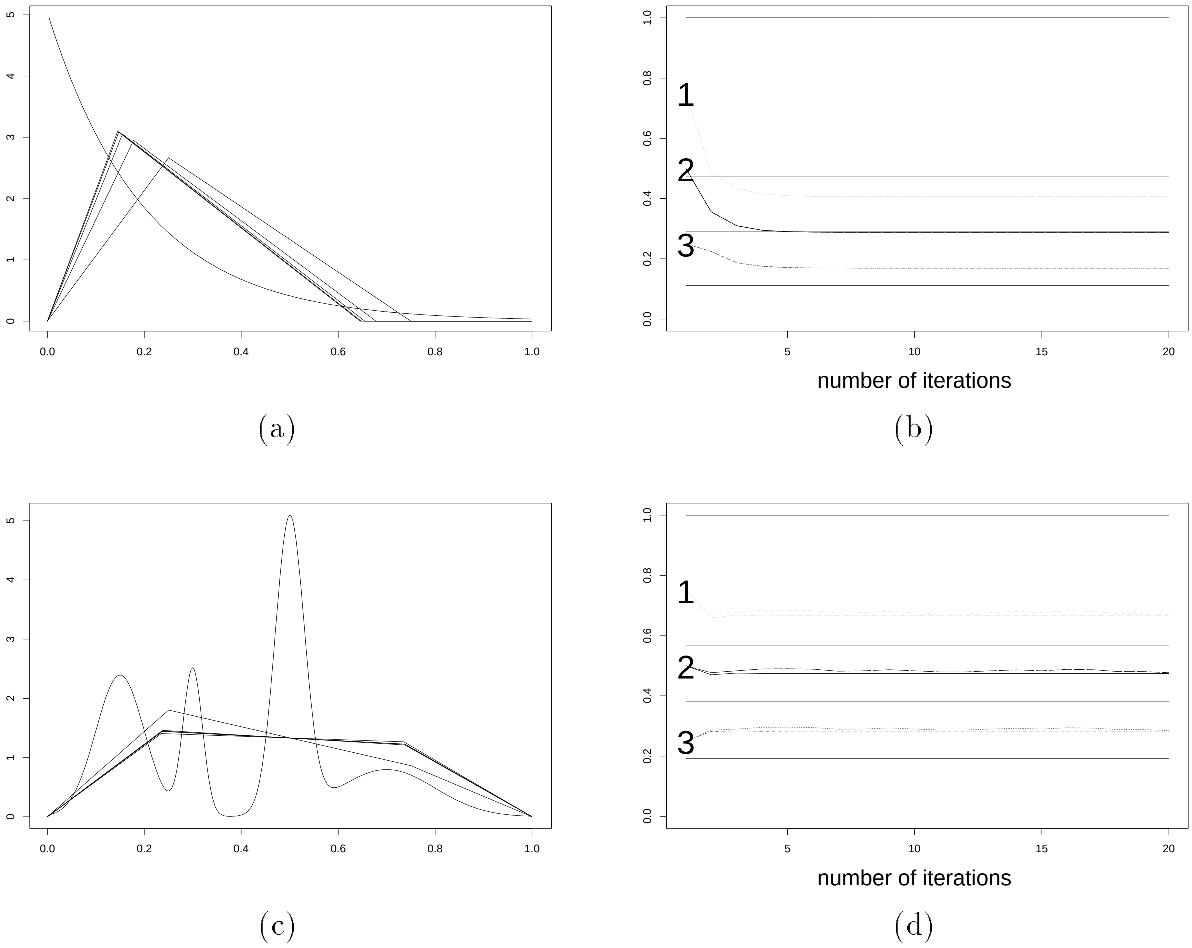

Figure 1: Illustration of the OV algorithm in two examples and when $L = 2$. (a) Conditioned exponential density in Example 1, imposed with piecewise linear approximations from the first 5 iterations of OV algorithm; (b) Convergence behavior of OV algorithm in Example 1: thin horizontal straight lines are optimal partition points (middle straight line) and associated reproduction levels (the upper and lower straight lines) based on the true density, and curved lines are OV partition points (marked 2) and associated reproduction levels (marked 1 and 3) and based on ideal and real data. (c) and (d) are for Example 2 instead.

the source density poorly, but the resulting limiting partition can be extremely good. The limiting reproduction level estimates, however, differ from the optimal ones (cf. (b)). The limiting OV quantizer nevertheless brings an SNR gain of 2.5 dB over the uniform quantizer (SNR=1.3 dB) at cost of no side information.

In the Gaussian mixture example (cf. Figure (c) and (d)), the OV limiting partition point and associated reproduction levels of OV differ from the optimal by approximately 0.10, but OV limiting quantizer is still an improvement on the uniform quantizer and it brings in an SNR gain of 1.5 dB over the uniform quantizer (SNR = 2.4 dB) at cost of no side information. Admittably both the piecewise density associated with the uniform quantizer and the limiting OV density estimate are bad estimates of the true density, but it is interesting to observe that the former is actually slightly better in terms of total variation distance (0.70 compared with 0.77), despite the fact the limiting OV quantizer is a better quantizer.

The above observations support the following

Claim: Estimating the interval partition of the optimal quantizer is an easier problem than estimating the density function.

Our first argument relies on Bennett's distortion integral in the high resolution case. When L is large, Bennett's distortion integral gives the approximate mean squared error distortion measure of a quantizer with point density $\lambda(\cdot)$:

$$D = \frac{1}{12L^2}\int f(x)/\lambda^2(x)dx + o(1/L^2),$$

which is minimized to the first order by taking $\lambda \propto f^{1/3}$. Hence what matters in the optimal quantizer is not f, but $f^{1/3}$, which is usually much flatter or smoother than f. It has been well established in the non-parametric density estimation literature (cf. Devroye, 1987) that the smoother the density, the easier its estimation. Hence the optimal quantizer's point density λ is easier to approximate by a smoothing algorithm like OV. That is, the optimal interval partition is easier to estimate than the density itself. Here the reproduction level estimates do not matter to the first order – any value in the right partition interval will do. Figure 2 illustrates this "smoothing" effect of cuberoot, which amounts to 50 % or more reduction in terms of total variance distance in Examples 1 and 2.

Our second argument explains why the OV limiting interval partition is "reasonable" when compared to the partition in the optimal quantizer and therefore supports the claim.

Definition A quantizer (R, C) is called an NC quantizer for density f if it satisfies the nearest neighbor condition and the centroid condition with respect to f; that is,

$$b_i = (c_{i-1} + c_i)/2, \quad c_i = \int_{b_{i-1}}^{b_i} x f(x) dx / \int_{b_{i-1}}^{b_i} f(x) dx,$$

or equivalently

$$|b_i - (c_{i-1} + c_i)/2| = 0, \quad |c_i - \int_{b_{i-1}}^{b_i} x f(x) dx / \int_{b_{i-1}}^{b_i} f(x) dx| = 0.$$

In general, we may expect to obtain only an NC quantizer even when source density is available. Hence in the quantized data case, it is good news that we can obtain an "approximate" NC quantizer, which is the case with the OV algorithm.

Lemma If two probability density functions give the same mass to an interval of length l and they are both subject to the same derivative constraint, then their centroids on this interval can differ at most by $l/3$.

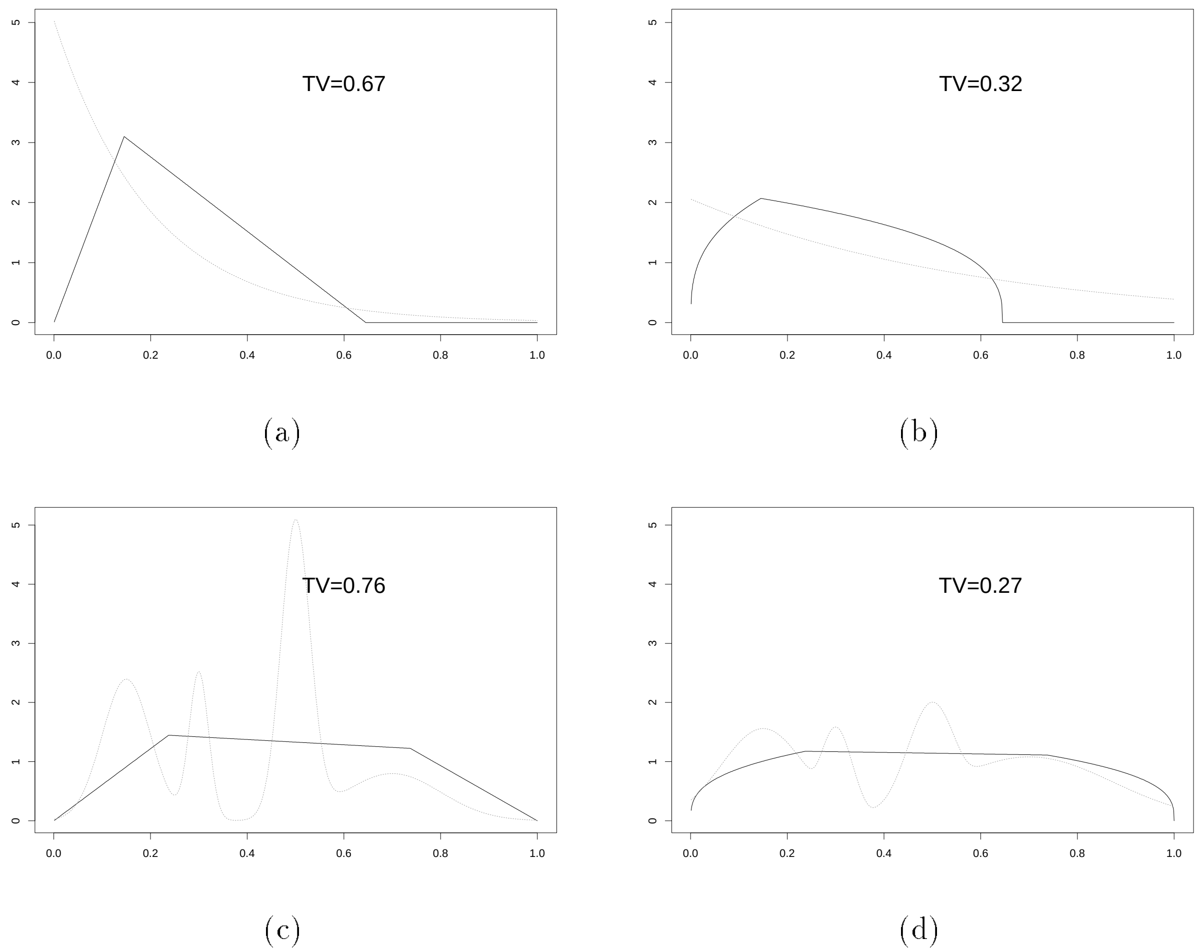

Figure 2: Smoothing effect of cuberoot. TV = Total Variation Distance. (a) conditioned exponential density in Example 1 and its OV limiting density approximation; (b) normalized cuberoots of curves in (a); (c) Mixture of Gaussians in Example 2 and its OV limiting density approximation; (d) normalized cuberoots of curves in (c).

Proof: This is a direct consequence of taking two linear density functions which are as far apart as possible subject to the derivative constraint and calculating their centroid difference. Let

$$f_1(x) = 2/l^2 \times x \quad x \in [0, l], \quad \text{and } f_2(x) = -2/l^2 \times (x - l) \quad x \in [0, l].$$

Their centroids are $EX_1 = 2l/3$ and $EX_2 = l/3$ so that their difference is $l/3$. $\square$

Suppose that the OV limiting quantizer is $(\hat{R}, \hat{C})$ then the nearest neighbor condition is satisfied (approximately) because the quantizer is an (approximate) NC quantizer with respect to the limiting density. Because of the probability-matching property of the limiting density, it follows from the above lemma that the centroid condition is satisfied by $(\hat{R}, \hat{C})$ with respect to the unknown density f with an approximation error $(\hat{b}_i - \hat{b}_{i-1})/3$ for the ith centroid. This approximation error goes to zero when the level L goes to infinity; it is proportional to the length of the partition interval. Large partition intervals correspond to locations where f has low values and in those places the departures from the exact centroid condition are not too harmful.

We now turn our attention to the reproduction level (or centroid in the MSE case) estimation problem. Unlike the high resolution case, it matters a great deal in the low resolution case which reproduction level (or centroid) estimates we use. Although the OV algorithm does match the interval probabilities between the density estimate and the source density, in the low resolution case it leaves large rooms for the centroids of the estimate to differ from those of the source density. This difference can lead to substantially worse distortion performance of the OV quantizer than that based on the same interval partition estimate as OV but with the corresponding centroids under the source density . This problem of the OV algorithm in the low resolution case calls for improvements, which are the topic of next section.

IV. IMPROVING THE REPRODUCTION LEVEL ESTIMATES

In this section, three types of estimates are proposed to improve upon the OV reproduction level estimates. One attractive feature of the OV algorithm is that it does not use any side information. It is desirable to keep this feature as much as possible.

We begin by introducing a different source for the centroid (or reproduction level) estimates than those from the piecewise linear estimate of the density function. This source is the raw counts of data in the intervals of previous data blocks in the allowed window. When these intervals differ, they form interleaving intervals on which a histogram can be built on finer intervals than the L-level quantization intervals. This histogram provides a different type of centroid estimates which are not influenced by the linear interpolation as in the OV density estimate. If we believe that the linear interpolation is the source for bias in the reproduction level estimate of OV (as demonstrated in the simulation examples in Section V), reducing such an influence should bring an improvement over the OV reproduction level estimate. One such possibility is to take a weighted average of the OV estimate and that from the histogram. Some notation is now needed to define the centroid estimates $\hat{c}_i^h$ based on the histogram.

Let W be the window size allowed for the histogram centroid estimates. Since both W and m in the OV algorithm measure how much past information should be taken into account, it makes sense to take $W = m$ as we do in the simulation studies. Technically, however, it might be advantageous to make them different if doing such brings better distortion performance. The construction of the new estimate $\hat{c}^h$ is best illustrated by an example. We take the first two blocks in Example 1 (L=2) to construct the histogram centroid estimates for the third block. The histograms from the first two blocks sit on partitions $(0, 0.5, 1)$ and $(0, 0.356, 1)$. Centroid estimates for the new partition $(0, 0.310, 1)$ are needed.

Based on Figure 3 and since 0.310 is inbetween 0.178 and 0.428, $\hat{c}^h$ can now be calculated as

$$\hat{c}_1^h = 0.178, \quad \hat{c}_2^h = (0.428 \times 0.087 + 0.75 \times 0.076)/(0.087 + 0.076) = 0.578.$$

If an interval misses the "poles" in Figure 3 (d), for example, the first interval in partition $(0, 0.1, 1)$, we take $\hat{c}_1^h$ as the middle point:

$$\hat{c}_1^h = (0 + 0.1)/2 = 0.05.$$

We now define $\hat{c}^h$ in the general case. At the end of the jth block, the interval partitions in this W-window are

$$R_k = \{a = b_0(k) < b_1(k) < ... < b_L(k) = b\} \quad \text{for } k = j-1, j-2, ..., j-W.$$

The associated probabilities are

$$\hat{P}_k = \{\hat{p}_i(k) : i = 1, 2, ..., L\} \quad \text{for } k = j, j-1, ..., j-W+1;$$

or equivalently the associated accumulated probabilities

$$\hat{P}_k^c = \{\hat{p}_i^c(k) : i = 1, 2, ..., L\} \quad \text{for } k = j, j-1, ..., j-W+1,$$

where $\hat{p}_i^c(k) = \sum_{t=1}^{i} \hat{p}_i(t)$. For each k, the accumulated probabilities are increasing just as the b's. All the b's from R_k in the allowed window are put together and sorted into an increasing

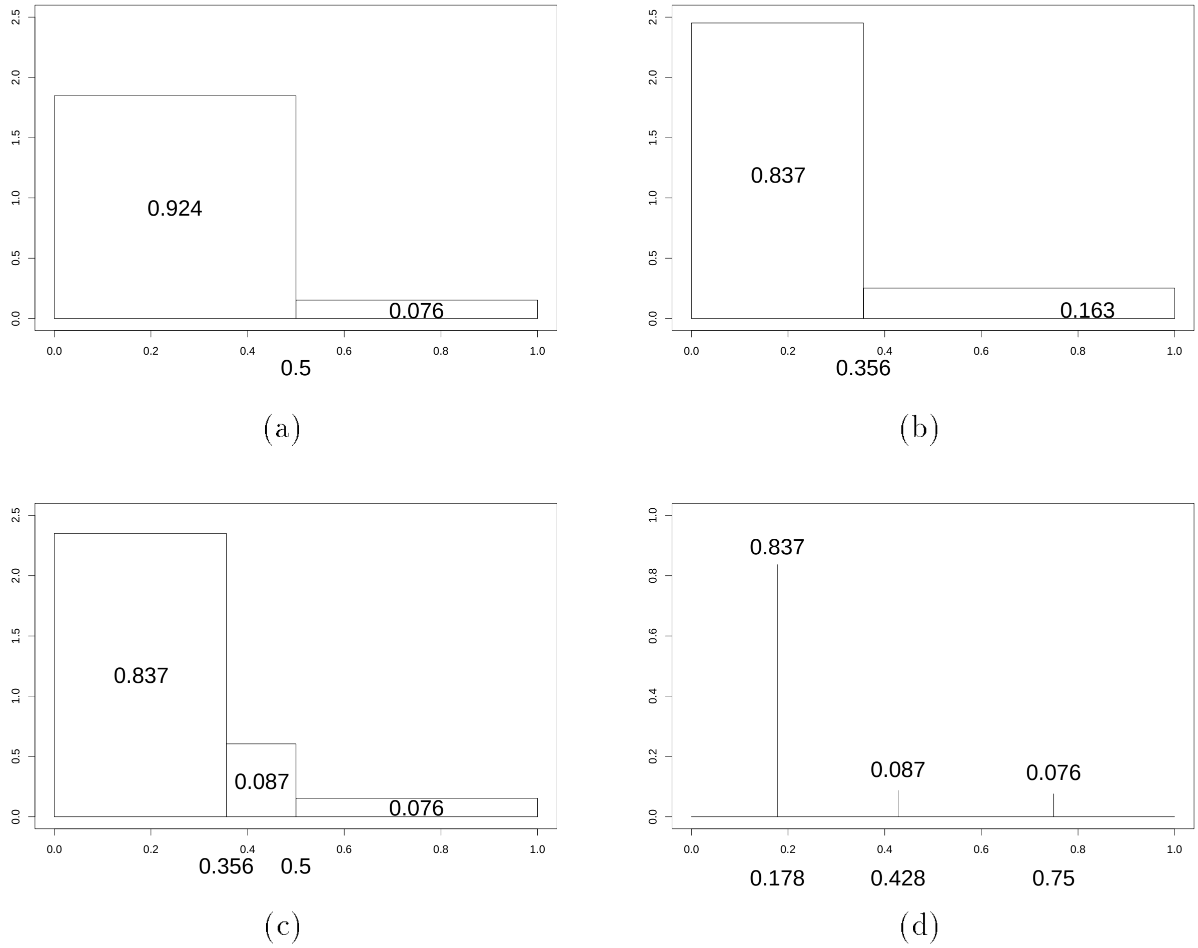

Figure 3: Construction of $\hat{c}^h$ in Example 1 ($L = 2$). (a) Histogram estimate from the first block with partition (0,0.5,1) and probabilities (0.924,0.076); (b) Histogram estimate from the second block with partition (0,0.356,1) and probabilities (0.837,0.163); (c) Merged histogram from the first and second blocks with a finer partition (0,0.356,0.5,1) and probabilities (0.837,0.087,0.076); (d) Middle points of the finer partition in (c) marked with their associated probabilites.

order, according to which the accumulated probablities $\hat{p}_i^c(k)$ are also sorted. Take difference of consecutive entries in this sorted accumulated probability vector; truncate the negative entries in the resulted difference vector into zeros and renormalize it to be summing to 1. The end results are

$$R(j, W) = \{a = b_0 < b_1 < ... < b_{(L-1)W+1} = b\},$$

and the associated probabilities

$$\hat{P}(j, W) = (\hat{p}_1, \hat{p}_2, ..., \hat{p}_{(L-1)W+1}).$$

Let the middle points of the intervals in partition $R(j, W)$ be

$$x_l = (b_{l-1} + b_l)/2.$$

Based on $R(j, W)$ and $\hat{P}(j, W)$ and for the new partition

$$R_j = \{a = b_0(j) < b_1(j) < ... < b_L(j) = b\},$$

define the histogram centroid estimates as

$$\hat{c}_i^h := (\sum_{l=1}^{(L-1)W+1} x_l \hat{p}_l I_{\{x_l \in (b_{l-1}(j), b_l(j)]\}})/(\sum_{l=1}^{(L-1)W+1} \hat{p}_l I_{\{x_l \in (b_{l-1}(j), b_l(j)]\}}).$$

If an interval in the new partition contains no x_l, its $\hat{c}^h$ is defined as the middle point of the interval.

Proposal 1: Weighted Average Centroid Estimates (WOV)

At end of block j, a new interval partition R_j is suggested by the OV algorithm with centroid estimates

$$\hat{c}_i^{ov} = \int_{b_{i-1}(j)}^{b_i(j)} x\hat{f}_j(x)dx / \int_{b_{i-1}(j)}^{b_i(j)} \hat{f}_j(x)dx.$$

Define the weighted average centroid estimates as

$$\hat{c}_i^{w} = \beta\hat{c}_i^{ov} + (1-\beta)\hat{c}_i^{h},$$

where β is a parameter which can be tuned in a particular application. In Section V, we use $\beta = 0.8$.

If a minimal amount of side information brings in a substantial distortion improvement, it can be argued that such side information should be allowed. We now propose two types of estimates of the reproduction levels (centroids). They both use side information and cause a delay of size T. The first type contain two procedures and they use only 1 and L extra bits for each block, respectively.

Proposal 2: Hybrid Centroid Estimates (HOV)

When a new block of data comes, the interval partition has been determined using the OV algorithm so the data block can be sent on line. At the coder's end, we run two registers of the sum of empirical distortions using the OV reproduction level estimate and the weighted average estimate. At the end of the block, one bit is sent to the decoder to indicate which of the two estimates has better empirical distortion performance and this one should be used for the reproduction of the data block. We call this estimate HOV_1. The other related procedure keeps L registers, each for every interval. For each interval, we accumulate the mean squared errors for both $\hat{c}^{ov}$ and $\hat{c}^{w}$ and use one bit to inform decoder at the end of the block which estimate is better for this interval. Hence L extra bits are used. This is called HOV_L. By design, hybrid methods will have better distortion measure performance than the OV algorithm, at the cost of some side information.

If we can tolerate more side information, we should go for

Proposal 3: Empirical Average Centroid Estimates (EOV)

We keep L registers at the coder's end for the L intervals. Each register keeps a running sum of the data points falling into a particular interval. At the end of the block, these L sums are converted into empirical averages and the averages are sent to the decoder with a relatively high precision. The empirical averages are then used as the reproduction levels. We consider three procedures in the simulation studies next section. They correspond to using three different precisions in encoding the empirical averages. They are denoted as EOV_1, EOV_2 and EOV_3 and are truncations of the empirical averages to one, two and three decimal places respectively. Fixed rate encoding is used for them hence $L \times i \log_2 10$ bits for EOV_i and level L $(i = 1, 2, 3)$. Some savings can be expected if variable rate predictive coding is used instead.

V. SIMULATION STUDIES

In this section, we compare WOV, HOV_1, HOV_L, EOV_1, EOV_2 and EOV_3 with OV in terms of their SNR vs. bit rate performance. Three examples are used. The first two examples are continued from Section III and consist of iid data. Recall that we chose special shapes of the underlying densities to make OV piecewise linear density poor estimates of the true density. The last is the non-stationary example in Ortega and Vetterli (1994) with a different parameter value. Technically the data ranges in examples 2 and 3 are not confined to [0,1], but we chose the parameters in the examples such that the probabilities for the data to be outside of [0,1] are extremely small.

Example 3 Non-stationary bimodal

The process is driven by a symmetric Markov chain of 2 states with switching probability 0.001. If the Markov chain is at state 1, then a sample is simulated from $N(0,(1/12)^2)$; otherwise it is simulated from $N(4,(1/12)^2)$. The initial distribution for the Markov chain is $(0.3, 0.7)$. Hence the chain is non-stationary and the resulting process is non-stationary either.

Recall that m and W are the memory window size parameters in the OV algorithm and in the histogram centroid estimates, respectively. To make the schemes fully adaptive to potentially changing source statistics, we take $m = W = 2$ and $T = 400$ in the examples and 40 blocks of data are used. Our first try on β in the hybrid proposal was 0.5 and the result indicated that there was not enough weight on the OV estimate. Hence $\beta = 0.8$ is chosen. In every Lloyd-Max algorithm run, $M = 70$ iterations were used. $\hat{Q}_0$ is taken as the uniform quantizer in all three examples.

In Table 1, we have SNR (dB) averages over blocks 3 to 40 for OV and the proposals from Section IV. (The first two blocks are not used for programming convenience.) Three levels are taken $L = 2, 4, 8$. As indicated in the table, three rows are bit rates per sample. The first 7 columns are self-explanatory from their names (cf. Section IV). Columns 8-10 act as benchmarks and are not available in practice. EOV_∞ is the empirical average proposal when no truncation is done. For examples 1 and 2, Sam is the sample SNR when the optimal LM quantizer is used; Exp is the optimal SNR under the true source distribution. In example 3, Sam corresponds to the sample SNR for the optimal LM quantizer for the $N(4,(1/12)^2)$ distribution, while Exp corresponds to the optimal SNR IF $N(4,(1/12)^2)$ is actually the true source.

From Table 1, we see that in terms of SNR alone, for $L = 2$ and in all three examples, EOV_2 and EOV_3 have performance very close to the optimal. EOV_1 has overall worse performance than OV because the precision 1/10 is too coarse and brings in large error in empirical average estimates. All other procedures improve OV by various degrees.

Table 1. SNR (dB) vs. bit rate (per sample)

level		OV	WOV	HOV_1	HOV_L	EOV_1	EOV_2	EOV_3	EOV_∞	Sam	Exp
L=2	e.g. 1	3.7	4.1	4.1	4.1	4.9	5.0	5.0	5.0	5.0	5.0
SNR	e.g. 2	3.7	3.6	3.7	3.8	4.4	4.7	4.7	4.7	6.4	6.4
	e.g. 3	0.8	0.3	0.9	1.0	3.7	4.1	4.2	4.2	-7.5	4.4
	bit rate	1.000	1.000	1.002	1.005	1.017	1.033	1.050			
L=4	e.g. 1	9.8	10.0	10.0	10.1	9.1	10.3	10.4	10.4	10.5	10.6
SNR	e.g. 2	11.4	11.9	11.9	11.9	11.3	12.6	12.6	12.6	13.4	13.4
	e.g. 3	7.3	7.0	7.3	7.3	5.6	7.9	7.9	7.9	7.2	9.3
	bit rate	2.000	2.000	2.002	2.010	2.033	2.066	2.100			
L= 8	e.g. 1	16.2	16.2	16.3	16.3	13.0	16.5	16.6	16.6	16.4	16.5
SNR	e.g. 2	17.3	17.2	17.3	17.4	15.7	18.1	18.1	18.1	18.8	18.8
	e.g. 3	13.6	13.4	13.6	13.7	8.7	13.8	13.9	13.9	-4.6	14.6
	bit rate	3.000	3.000	3.002	3.020	3.066	3.133	3.199			

To judge whether the improvements are worth the extra bits needed, we turn to Tables 2 - 4. It is generally expected that 1 bit rate increase brings about 6 dB of SNR gain (cf. Gersho and Gray, p. 154, 1992). In our quantized data case, the first bit brings in an SNR gain well below 6 dB in all three examples (cf. upper part of the OV column in Table 1), but the 6 dB gain is nevertheless achieved for the additional bits as shown in Table 2. This (rough) 6 dB gain can be re-expressed as 0.15 (bit/dB), i.e., 0.15 extra bits for 1 dB SNR increase. This 0.15 (bit/dB) provides the baseline against which to compare our proposals.

Table 2. SNR gain vs bit rate increase for OV

base rate	rate increase	e.g. 1	e.g. 2	e.g. 3	ave.	bit/dB
1 (L=2)	1	6.1	7.7	6.5	6.8	0.148
2 (L=4)	1	6.4	5.9	6.3	6.3	0.161

Table 3 contains the SNR gain (loss if it is negative) of the six proposed procedures relative to OV and the extra bits needed for each of them. For each level, we average the SNR gains across examples to obtain the ave. rows in Table 3. The averaged SNRs reflect the performance of individual procedures since the three examples have different types of difficulties an adaptive algorithm faces: examples 1 and 2 have stationarity, but the density shapes are hard to estimate based on quantized data. In Example 3, due to non-stationarity, we have a changing "source" density for each block and the algorithm might not learn enough from previous blocks.

Table 3. SNR gain vs. extra bit rate relative to OV

level		WOV	HOV_1	HOV_L	EOV_1	EOV_2	EOV_3
L=2	e.g. 1	0.4	0.4	0.4	1.1	1.3	1.3
SNR	e.g. 2	-0.1	0.0	0.1	0.7	1.0	1.0
gain	e.g. 3	-0.5	0.0	0.1	2.8	3.3	3.3
	ave.	-0.07	0.13	0.20	1.53	1.87	1.87
	extra bit	0	0.002	0.005	0.017	0.033	0.050
L=4	e.g. 1	0.2	0.2	0.2	-0.8	0.5	0.5
SNR	e.g. 2	0.5	0.5	0.6	0.1	1.3	1.3
gain	e.g. 3	-0.2	0.0	0.1	-1.6	0.7	0.7
	ave.	0.17	0.23	0.30	-0.77	0.83	0.83
	extra bit	0	0.002	0.010	0.033	0.066	0.100
L=8	e.g. 1	0.0	0.1	0.1	-3.2	0.3	0.4
SNR	e.g. 2	0.0	0.0	0.1	-1.6	0.8	0.9
gain	e.g. 3	-0.2	0.0	0.1	-5.0	0.1	0.2
	ave.	-0.33	0.03	0.1	-3.27	0.4	0.5
	extra bit	0	0.002	0.020	0.066	0.133	0.199

In Table 4, we take the ave. and extra bit rows and form the ratio (bit/dB) which indicates the extra bits needed for 1 dB SNR gain. It is clear that for level L=2, all proposed procedures are well worth doing (ratio far below 0.1 bit/dB). EOV_1 is not a good idea when $L = 4, 8$ although for which WOV, HOV_1 and HOV_L are very cost effective (ratio well below 0.1 bits per 1 dB increase). EOV_2 and EOV_3 are cost effective for $L = 4$ (ratio around 0.1 bits per 1 dB increase), but not so for $L = 8$ (ratio around 0.3 or 0.4 bits per 1 dB increase).

Table 4. Average extra bits needed for 1 dB SNR gain (bit/dB)

level	WOV	HOV_1	HOV_L	EOV_1	EOV_2	EOV_3
L=2	0	0.015	0.025	0.011	0.018	0.027
L=4	0	0.009	0.017	-0.043	0.080	0.120
L=8	0	0.067	0.050	-0.020	0.332	0.399

VI. CONCLUSIONS

We have analyzed the model-based adaptive quantization problem based on past quantized data. A consistency result for estimating the model or density is established for parametric models. Through the Bennett's distortion integral and an analysis of the piecewise linear algorithm in Ortega and Vetterli (1994), an important thesis is put forward, which states that it is advantageous to separate the interval partition estimation from the reproduction level estimation. The first estimation problem is easier and for which the OV algorithm performs well. Six proposals are given to improve the OV reproduction level estimates (in the low resolution case). In particular, the hybrid proposals HOV_1 and HOV_L use extra bits well below 0.1 to obtain 1 dB increase in SNR performance ($L = 2, 4, 8$ and block size $T = 400$). For $L = 2, 4$ and $T = 400$, empirical average proposals EOV_2 and EOV_3 need extra bits below or around 0.1 to gain 1 dB of SNR. The cost of extra bits are fixed for each block, hence the above observations hold for block sizes around 400 or larger.

The overload problem is not dealt with in this paper, but will be in future work. It is well-known that the MSE distortion measure is not perception-realistic. The more realistic distortion measures, however, have been modeled as weighted averages of MSE distortions over transformed signals of different frequencies and the analysis in this paper does apply to those transformed signals.

ACKNOWLEDGEMENT

The author would like to thank Martin Vetterli for introducing her to the topic and several stimulating discussions, and Antonio Ortega for answering many questions on the topic and commenting on the drafts. Finally, a helpful conversation with John Rice is gratefully acknowledged.

References

A. R. Calderbank, S. W. McLaughlin, and D. F. Lyons, "A low complexity two-stage adaptive vector quantizer," In *Proc. of the 25th CISS*, Baltimore, MD., Mar. 1991, pp. 582-587.

S. Crisafulli and R. B. Bitmead, "Adaptive quantization: solution via nonadaptive linear control," *IEEE Trans. on Comm.,* 41(5), pp. 741-748, 1993.

L. Devroye, *A course on density estimation,* Birkhäuser, Boston, 1987.

A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression.* Academic Press, 1992.

D. J. Goodman and A. Gersho, "Theory of an adaptive quantizer," *IEEE Trans. on Comm.,* COM-22(8), pp. 1037-1045, 1974.

N. S. Jayant, "Adaptive quantization with a one-word memory," *Bell Sys. Tech. J.* Vol. 52, pp. 1119-1144, 1973.

A. Ortega and M. Vetterli,"Adaptive quantization without side information," *Int'l Conf. on Image Proc., ICIP'94,* Austin, Texas, pp. 856–860, 1994,