High-Resolution
Distributed Functional Quantization

Vinith Misra, Vivek K Goyal, and Lav R. Varshney
Signal Transformation and Information Representation Group
Research Laboratory of Electronics
Massachusetts Institute of Technology
http://www.rle.mit.edu/stir/

***Abstract*—In traditional modes of lossy compression, attaining low distortion letter-by-letter on a vector of source letters $X_1^N = (X_1, X_2, \ldots, X_N) \in \mathbb{R}^N$ is the implicit aim. We consider here instead the goal of estimating at the destination a function $G(X_1^N)$ of the source data under the constraint that each X_i must be separately scalar quantized. The design of optimal fixed- and variable-rate scalar quantizers is considered under the assumptions of high-resolution quantization theory, yielding optimal point densities for regular quantizers. Additionally, we consider how performance scales with N for certain classes of functions. This demonstrates potentially large improvement from consideration of G in the quantizer design.**

I. Functional Source Coding

Lossy compression traditionally seeks to approximate every letter of a source vector $X_1^N = (X_1, X_2, \ldots, X_N) \in \mathbb{R}^N$ to attain small mean-squared error (MSE). Other work has looked at perceptual distortion measures rather than MSE [1] and at task-oriented quantization, such as quantizing for classification [2], estimation [3], or detection [4].

We consider the related class of problems for which the goal is to estimate at the destination a function $G(X_1^N)$ of the source data. We call this *(lossy) functional source coding* (FSC). A wide variety of practical situations fit into the FSC framework; consider, for instance, computation on digitized analog data. In this paper, we present a general framework but focus on functions that are used to estimate location and scale parameters of statistical distributions, such as the median, minimum, maximum, and midrange [5], [6]. For these types of functions of the observed data, we explicitly determine the scaling behavior with respect to N. This demonstrates very large improvement from the consideration of the function in the quantizer design.

The general FSC problem is described by a distortion measure over the source alphabet, d, induced by a distortion measure d_G defined on the range of the function G:

$$d(X_1^N, \widehat{X}_1^N) = d_G(G(X_1^N), \widehat{G}(\widehat{X}_1^N)),$$

where $\widehat{X}_1^N$ is the compressed representation of X_1^N, and $\widehat{G}$ estimates $G(X_1^N)$ from $\widehat{X}_1^N$. We restrict our attention to squared-error distortion, $d_G(G_1, G_2) = (G_1 - G_2)^2$, and continuous G. Under these circumstances, the intermediate value theorem predicts a point z_1^N for which $\widehat{G}(\widehat{X}_1^N) = G(z_1^N)$; we may therefore set $\widehat{G} = G$.

Supported in part by National Science Foundation Grant CCF-0729069.

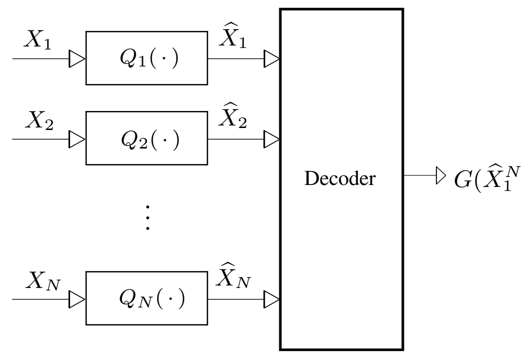

Fig. 1. Distributed functional quantization: Scalar quantization of each source is performed independently.

If the source data, X_1^N, is to be jointly encoded, one may simply compute G before the encoding process: FSC thereby reduces to a traditional source coding problem. Bucklew [7] considered a problem with a single encoder in which the function of interest G is only known to the decoder—the encoder is only aware of the distribution from which G is selected. We instead consider distributed coding. Even though G is known to the encoders, this is related to Bucklew's problem in that—from the perspective of encoder i—the variables $\{X_j\}_{j \neq i}$ randomize the function of interest.

The distributed construction involves N disjoint encoders for the N sources, as shown in Fig. 1. We place a sum-rate constraint on the communication to a common decoder: $\sum_{i=1}^N R_i = NR$, where R_i is the rate for the ith encoder and NR is a sum-rate constraint. Our interest lies in cases where R is large enough for high-resolution analysis techniques to be used fruitfully.

II. High-Resolution Regular Quantization

This section recapitulates the development in [8].

In examining FSC, we employ high-resolution techniques [9]. On one hand, this approach can directly yield asymptotically optimal quantizers. On the other hand, for it to be valid we must constrain our functions and distributions in the following manner:

C1. $G(X_1^N)$ is smooth and monotonic in each of the source variables.

C2. $\partial G / \partial X_i$ is bounded almost everywhere.

C3. The joint probability distribution of the source variables, $f_{X_1^N}(x_1^N)$, is continuous and supported within a compact subset of $\mathbb{R}^N$.

Assuming that G adheres to **C1** and **C2** and that **C3** is valid for the source distribution, several approximations are justified:

A1. Within any quantizer cell, the distribution $f_{X_1^N}$ is roughly uniform (**C3**).

A2. Within any quantizer cell, G_i is roughly affine in X_1^N (**C1**, **C2**).

A3. All quantizer regions are regular (**C1**).

(The development here generalizes to G parameterized by a random variable Θ with a known distribution. In this case, we require each G_θ to satisfy the stated conditions.)

With these approximations, the function's influence on performance is summarized by quantities we term the *functional sensitivities*,

$$g_i^2(x) = E\left[\left| \frac{\partial G(x_1^N)}{\partial x_i} \Big|_{x_1^N = X_1^N} \right|^2 \; \Big| \; X_i = x \right].$$

Let K_i and λ_i be respectively the number of intervals (resolution) and the quantization point density function for the ith quantizer. The distortion–K–λ relationship that emerges from high rate analysis is

$$D = \sum_{i=1}^{N} \frac{1}{12K_i^2} E\left[g_i^2(X_i)\lambda_i^{-2}(X_i) \right].$$

The optimization of this differs between the fixed- and variable-rate cases, due to the differing specifications of rate, R, in terms of the resolution, K. For fixed-rate, the functional distortion is minimized by the choice $\lambda_i \propto (f_X g_i^2)^{1/3}$ leading to distortion

$$D = \frac{2^{-2R/N}}{12} \prod_{i=1}^{N} \| f_X g_i^2 \|_{1/3}^{1/N}. \tag{1}$$

For variable-rate, the minimizing point density, λ_i, is proportional to g_i itself, giving rise to distortion

$$D = \frac{2^{-2R/N}}{12} \prod_{i=1}^{N} 2^{2h(X_i)/N + E[\log_2 g_i^2]/N}]. \tag{2}$$

Note that the latter of these expressions resembles that of Linder et al. [10], if one treats the function as having introduced a "locally quadratic" distortion measure.

We will observe that the introduction of g_i^2 can strongly differentiate the performances of fixed- and variable-rate quantization.

III. Scaling Performance

Intuitively, it should improve performance to account for the function G in the design of quantizers for $\{X_i\}_{i=1}^N$. We examine the scenario where the decoder hopes to extract statistical information from multiple samples of the source. Specifically, we consider the gathering of the median and the midrange. For a quantitative comparison against ordinary quantization, we calculate the ratio $\gamma(N) = D_F(N,R)/D_O(N,R)$, where $D_F(N,R)$ and $D_O(N,R)$ are the expected functional distortion of the functionally-optimized (F) or ordinary (O) quantizer at transmission rate R per source and dimensionality N.

Aside from one observation in Section III-B2, we restrict our attention to sources that are i.i.d. and uniform over the interval $[0,1]$. This is for ease of computation, as the developments in [8] are more general.

A. Fixed-Rate Quantization

To obtain the ratio $\gamma(N)$ for either the median or the midrange, we must first establish the distortion from an ordinary (non-functional) quantizer. The optimal fixed-rate quantizer for a uniform source is uniform; this yields

$$E[(X_i - \hat{X}_i)^2] = \frac{1}{12} 2^{-2R_i} \tag{3}$$

for each i. This will easily yield distortions for the computation of the median and midrange. To find the distortion incurred by the optimal fixed-rate functional quantizer, we will use (1). Before this may be evaluated, the sensitivities $g_i^2(x_i)$ must be derived separately for the median and the midrange.

1) $G = \mathrm{median}(X_1^N)$: For simplicity, restrict N to be odd valued with $N = 2M+1$; the median is then defined as the $(M+1)$st order statistic.

Since $\mathrm{median}(X_1^N) = X_i$ for some $i \in \{1, 2, \ldots, N\}$, it follows immediately from (3) that

$$D_O(N,R) = \frac{1}{12} 2^{-2R}. \tag{4}$$

(Each $i \in \{1, 2, \ldots, N\}$ is equally likely, so an equal bit allocation $R_i = R$ is optimal.) The absence of a dependence on N fits with our intuition: the median simply takes one of the source values, so the dimensionality should not affect the quantizer's accuracy.

To determine the optimal functional quantizer's distortion, we must first obtain $g_i^2(x_i)$. Given a point $x_i \in [0,1]$, x_i is either itself the median of x_1^N or (differentially or locally) it has no bearing on the value of the median. Therefore,

$$g_i^2(x) = \Pr[\mathrm{median}(X_1^N) = x \mid X_i = x].$$

This probability may be evaluated combinatorially in terms of the parameter $M = (N-1)/2$ and the cumulative distribution function of X, denoted $F_X(x)$. For x to be the median, M of the other sources must be above it, and M below it. There are $\binom{2M}{M}$ possible selections of which sources are above or below. The probability of each of these choices is $F_X(x)^M(1-F_X(x))^M$; summing them, we have

$$g_i^2(x) = \binom{2M}{M} F_X(x)^M (1 - F_X(x))^M. \tag{5}$$

For the uniform distribution of interest, $F_X(x) = x$. This yields an optimal point density $\lambda_i \propto x^{M/3}(1-x)^{M/3}$ and total distortion

$$D_F(N,R) = \frac{1}{12} N 2^{-2R} \left\| \binom{2M}{M} x^M (1-x)^M \right\|_{1/3} \tag{6}$$

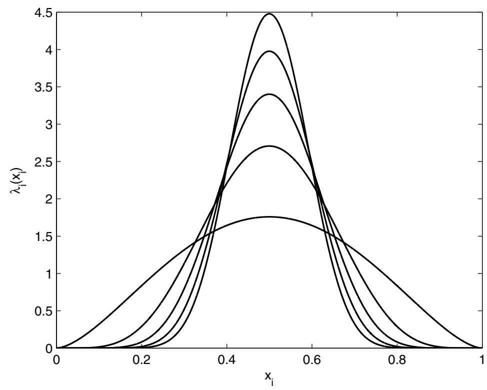

Fig. 2. Optimal fixed-rate point density for calculating median. $N = 11, 31, 51, 71, 91$.

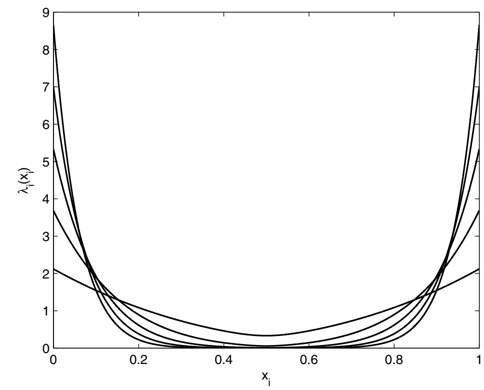

Fig. 4. Optimal fixed-rate point density for calculating midrange. $N = 10, 20, 30, 40, 50$.

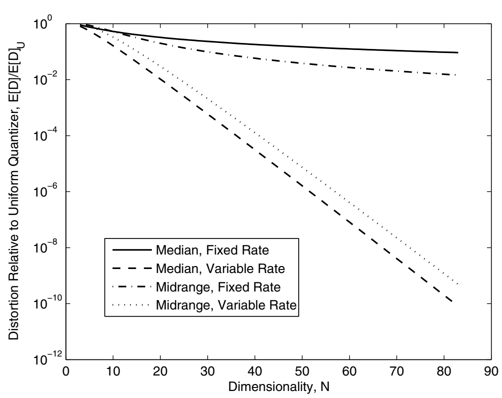

Fig. 3. Factor of reduction in distortion from functional quantization (log scale), with dimensionality (N).

The point density reflects our intuition that more quantizer intervals should be assigned to the more important middle ground—a fact that becomes increasingly true as the dimensionality is increased (see Fig. 2).

By means of Stirling's formula, the functional distortion of (6) is found to fall with $1/M \propto 1/N$:

$$D_F(N,R) \approx \frac{2^{-2R}}{6}\frac{1}{M}\left(\sqrt{3\pi}\frac{e}{2}\right)^3 \tag{7}$$

In contrast, since the median simply takes on one of the source values, $D_O(N,R)$ does not depend on N (recall Eq. (4)). The ratio $\gamma(N) = D_F(N,R)/D_O(N,R)$—which is independent of R—is shown to demonstrate $1/N$ behavior in Fig. 3. Notice that the variance of the median—and with it, the distortion of a fixed-rate centralized encoder—also falls with $1/N$.

2) $G = \mathrm{midrange}(X_1^N)$: The midrange is defined as the average of the minimum and the maximum components of x_1^N. No parity restriction on N is necessary to obtain clear results. We merely require $N \geq 2$.

We again start by obtaining the ordinary quantizer's distortion, $D_O(N,R)$. Since $\mathrm{midrange}(X_1^N) = (X_i + X_j)/2$ with i and j each uniformly distributed over $\{1, 2, \ldots, N\}$, equal bit allocation is optimal and the resulting distortion is

$$D_O(N,R) = \frac{1}{24}2^{-2R}. \tag{8}$$

To compute the analogous quantity for the functionally optimized quantizer, we first turn our attention towards the sensitivity $g_i^2(x)$. If x_i is not the minimum or the maximum, $\partial G(x_1^N)/\partial x_i$ is zero; otherwise $\partial G(x_1^N)/\partial x_i$ is $1/2$. The latter situation occurs with probability

$$\Pr[\max(X_1^N = x) \mid X_i = x] + \Pr[\min(X_1^n = x) \mid X_i = x]$$

since the minimal and maximal events are disjoint almost everywhere. Therefore,

$$\begin{aligned} g_i^2(x) &= \tfrac{1}{4}\left(\Pr[\max(X_1^n = x) \mid X_i = x]\right. \\ &\quad \left. + \Pr[\min(X_1^n = x) \mid X_i = x]\right) \\ &= \tfrac{1}{4}(F_X(x)^{N-1} + (1 - F_X(x))^{N-1}). \end{aligned} \tag{9}$$

The fixed-rate quantizer distortion from this expression is

$$D_F(N,R) = \frac{1}{12}N2^{-2R}\left\|\frac{1}{4}(F_X(x)^{N-1} + (1 - F_X(x))^{N-1})\right\|_{1/3}. \tag{10}$$

For the uniform $[0,1]$ source, $F_X(x) = x$. For large values of N, $g_i^2(x)$ is dominated by $(1 - F_X(x))^{N-1}$ when $x < 1/2$, and by $F_X(x)^{N-1}$ when x exceeds $1/2$. We may then approximate the integral to be

$$D_F(N,R) \approx N2^{-2R}\frac{18}{(N+2)^3}.$$

For large values of N, this follows an approximate $1/N^2$ falloff, in contrast to the constant distortion of the ordinary quantizer (see Fig. 3). As the midrange is calculated from the minimum and the maximum, it should not surprise us that the point densities that achieve this performance are increasingly concentrated towards the edges as the dimensionality increases (see Fig. 4).

The variance of the midrange is known to fall as $1/N^2$ [11]. As with the median, this implies that a fixed-rate centralized encoder's dependence on N sometimes carries into the distributed scenario.

B. Variable-Rate Quantization

Note the increasingly concentrated quantization profiles in Figs. 2 and 4. In accordance with this trend, the entropies of the quantized sources, $H(\widehat{X}_i)$, can be shown to fall with N. One can therefore expect entropy coding to yield a significant improvement—a rare occurrence for ordinary quantization. If a source is properly quantized with ordinary techniques, the probability mass function of the quantized symbols will reflect the probability density function of the source. This is no longer true for functional quantization: we have seen examples, for instance, where a uniform source is optimally quantized by increasingly non-uniform quantizers.

To quantify the gains from entropy coding, we can concentrate on computing the distortions from optimized variable-rate functional quantization; for ordinary quantization, the results from the previous section carry over. Distortion for the variable-rate functional quantizer is given by (2) and is clearly dependent on the quantity $2^{E[\log_2 g_i^2]}$. We calculate this separately for the median and midrange.

1) $G = \text{median}(X_1^N)$: The quantity $g_i^2(x)$ is still given by (5). Making use of this, we can express the functional variable-rate quantizer's distortion as

$$\begin{aligned} D_F(N,R) \\ &= \frac{N2^{-2R}}{12}\exp_2\left\{\int_0^1 \log_2\left[\begin{pmatrix} 2M \\ M \end{pmatrix} x^M(1-x)^M\right]dx\right\} \\ &= \frac{N2^{-2R}}{12}\exp_e\{-N(1-\ln 2)\}. \end{aligned}$$

As illustrated by the linear curve in Fig. 3, $\gamma(N)$ drops exponentially with N—significantly faster convergence than the $D \sim 1/N$ seen for fixed-rate. Note that the optimal point density does not change significantly between the fixed and variable-rate scenarios; the large improvement we observe arises primarily from entropy coding the increasingly concentrated quantizer point density.

2) $G = \text{midrange}(X_1^N)$: As with the median, we may reuse a computation of $g_i^2(x)$, this time from (9). We make the same approximation as we did for the fixed rate: that $g_i^2(x)$ is dominated by $F_X(x)^{N-1}$ for $x > 1/2$ and by $(1-F_X(x))^{N-2}$ for $x < 1/2$. This yields

$$\begin{aligned} D_F(N,R) \\ &\approx \frac{1}{12}N2^{-2R+2h(X)}\exp_2\left\{2\int_{1/2}^1 f_X(x)\log_2\left(\tfrac{1}{4}F_X(x)^{N-1}\right)dx\right\} \\ &= \frac{1}{12}N2^{-2R+2h(X)-2}\exp_2\left\{2(N-1)\int_{1/2}^1 f_X(x)\log_2 F_X(x)\,dx\right\}. \end{aligned}$$

The midrange distortion falls exponentially, regardless of the source distribution; in fact, we find that the value of the distortion seems to be *invariant* to source distribution—one obtains the same behavior for uniform, one- and two-sided triangle, cosine, parabolic, and exponential distributions. $D_F(N,R)$ is illustrated in Fig. 3.

As with the median, we observe significantly lower distortion than for the fixed-rate scenario. The optimal variable rate quantizer is described by $\lambda_i \propto g_i$. This is not significantly different from the fixed-rate point density; the exponential improvement in performance comes, once again, from the entropy coding rate reduction.

References

[1] J. Li, N. Chaddha, and R. Gray, "Asymptotic performance of vector quantizers with a perceptual distortion measure," *IEEE Trans. Inform. Theory*, vol. 45, no. 4, pp. 1082–1091, May 1999.

[2] H. Xie and A. Ortega, "Entropy- and complexity-constrained classified quantizer design for distributed image classification," in *IEEE Workshop on Multimedia Sig. Process*, Dec. 2002, pp. 77–80.

[3] L. Vasudevan, A. Ortega, and U. Mitra, "Application-specific compression for time delay estimation in sensor networks," in *SenSys '03: Proc. 1st Int. Conf. Embedded Netw. Sensor Syst.* New York, NY, USA: ACM, 2003, pp. 243–254.

[4] R. Gupta and A. O. Hero, III, "High-rate vector quantization for detection," *IEEE Trans. Inform. Theory*, vol. 49, no. 8, pp. 1951–1969, Aug. 2003.

[5] E. H. Lloyd, "Least-squares estimation of location and scale parameters using order statistics," *Biometrika*, vol. 39, no. 1/2, pp. 88–95, Apr. 1952.

[6] P. R. Rider, "The midrange of a sample as an estimator of the population midrange," *J. Acoust. Soc. Am.*, vol. 52, no. 280, pp. 537–542, Dec. 1957.

[7] J. A. Bucklew, "Multidimensional digitization of data followed by a mapping," *IEEE Trans. Inform. Theory*, vol. IT-30, no. 1, pp. 107–110, Jan. 1984.

[8] V. Misra, V. K. Goyal, and L. R. Varshney, "High-resolution functional quantization," in *Proc. IEEE Data Compression Conf.*, Mar. 2008.

[9] R. M. Gray and D. L. Nehoff, "Quantization," *IEEE Trans. Inform. Theory*, vol. 44, no. 6, pp. 2325–2383, Oct. 1998.

[10] T. Linder, R. Zamir, and K. Zeger, "High-resolution source coding for non-difference distortion measures: multidimensional companding," *IEEE Trans. Inform. Theory*, vol. 45, no. 2, pp. 548–561, Mar. 1999.

[11] J. Neyman and E. S. Pearson, "On the use and interpretation of certain test criteria for purposes of statistical inference: Part I," *Biometrika*, vol. 20A, no. 1/2, pp. 175–240, Jul. 1928.