International Symposium on Information Theory and Its Applications
Honolulu, Hawaii, U.S.A., November 5–8, 2000

Optimal Entropy-Constrained Scalar Quantization of a Uniform Source

András György[†] [‡] Tamás Linder[‡]

[†]Faculty of Electrical Engineering & Informatics
Technical University of Budapest
H-1521 Budapest, Hungary
Email: gya@szit.bme.hu

[‡]Department of Mathematics & Statistics
Queen's University
Kingston, Ontario, Canada K7L 3N6
Email: linder@mast.queensu.ca

Abstract

Optimal scalar quantization subject to an entropy constraint is studied for a wide class of difference distortion measures including rth power distortions with $r > 0$. It is proved that if the source is uniformly distributed over an interval, then for any entropy constraint R (in nats), an optimal quantizer has $N = \lceil e^R \rceil$ interval cells such that $N-1$ cells have equal length d and one cell has length $c \leq d$. The cell lengths are uniquely determined by the requirement that the entropy constraint is satisfied with equality. Based on this result, a parametric representation of the minimum achievable distortion $D_h(R)$ as a function of the entropy constraint R is obtained for a uniform source. The $D_h(R)$ curve turns out to be nonconvex in general. Moreover, for the squared error distortion it is shown that $D_h(R)$ is a piecewise concave function, and that a scalar quantizer achieving the lower convex hull of $D_h(R)$ exists only at rates $R = \log N$, where N is a positive integer.

1. Introduction

Scalar (or zero-memory) quantization is the simplest method for the lossy coding of an information source with real-valued outputs. A scalar quantizer followed by variable length lossless coding (entropy coding) can perform remarkably well, which makes this method popular in applications where implementation complexity is a decisive factor.

The two main quantities characterizing a scalar quantizer Q are its distortion and rate. The distortion $D(Q)$ is the average distortion between the source and the quantizer output. If Q is followed by entropy coding, the rate is usually defined as the entropy $H(Q)$ of the output of Q. (For a stationary and memoryless source, $H(Q)$ is indeed the smallest rate asymptotically achievable by variable length lossless coding of blocks of quantizer outputs.) One would like to make both $H(Q)$ and $D(Q)$ as small as possible, but these quantities are inversely related. A natural design problem is then to minimize $D(Q)$ subject to an entropy constraint $H(Q) \leq R$. Let $D_h(R)$ denote the lowest possible distortion of any scalar quantizer with output entropy not greater than R. A quantizer achieving this minimum is called an optimal entropy-constrained scalar quantizer (ECSQ). It is of interest to determine $D_h(R)$ either analytically or numerically, as well as to find the optimal ECSQ achieving the minimum distortion.

It appears that very few concrete examples for an optimal ECSQ are known in analytic form. In general, efforts have focused on finding necessary conditions for the optimality of an ECSQ with a *fixed* number of output points n [1, 2, 3]. These conditions give rise to practical algorithms for designing an ECSQ with a fixed number of output points [1, 4, 2, 3, 5]. To determine the overall optimal ECSQ and the corresponding optimal performance curve $D_h(R)$, one must find the optimum performance over all n. Unfortunately, this step is rather hard, even for the most common continuous source distributions. A notable exception is the case of an exponentially distributed source and mean squared distortion considered by Berger [1]. He derived an analytic expression for $D_h(R)$ based on the observation that for the exponential distribution, the necessary conditions for optimality at any positive rate are satisfied by an infinite-level uniform quantizer. To our knowledge, this is the only case where a correct[1] analytic formula for $D_h(R)$ is known.

In this paper we determine analytically the optimal ECSQ for a source which is uniformly distributed over a finite interval. We allow a rather wide class of difference distortion measures including rth power distortions $d(x, y) = |x - y|^r$ with $r > 0$, and distortion measures of the form $d(x, y) = \rho(|x - y|)$, where ρ is a nonnegative, strictly increasing, and convex function. Our main result proves that an optimal ECSQ for any rate $R \geq 0$ (measured in nats) is an $N = \lceil e^R \rceil$-level quantizer (here $\lceil x \rceil$ denotes the smallest integer not less than x). This quantizer has $N - 1$ cells of equal length d and one cell of length $c \leq d$,

This work was supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada. A. György also acknowledges the support of the Soros Foundation.

[1]Although a complete proof that infinite-level uniform quantizers are indeed optimal is missing, the result is widely believed to be correct.

where d and c are uniquely determined by the requirement that the entropy constraint is satisfied with equality (the optimal quantizer is uniform if $e^R = N$, as expected). Specialized to the squared error distortion, our result rigorously proves that the ECSQs found for the uniform distribution by Farvardin and Modestino [2] (using a numerical approach) are indeed optimal. Based on the analytic description of an optimal ECSQ, we then obtain a parametric expression for the $D_h(R)$ curve and investigate its properties. In general $D_h(R)$ is piecewise smooth (differentiable everywhere except at the points $R = \log N$). For the squared error distortion (and more generally for rth power distortions with $0 < r \leq 3$) we prove that $D_h(R)$ is concave in each interval $[\log(N-1), \log N]$, where $N \geq 2$ is an integer. Thus for such distortion measures, $D_h(R)$ is not convex over any interval. It also follows that in this case an optimal ECSQ achieving the lower convex hull of $D_h(R)$ exists only at rates $R = \log N$, where N is a positive integer.

The question whether $D_h(R)$ for a given source is convex is of particular interest because of the special role of the lower convex hull of $D_h(R)$ in variable length lossy coding. For example, the lower convex hull of $D_h(R)$ is the minimum achievable distortion in causal lossy coding of a memoryless source [6]. Also, Lloyd-Max type necessary conditions of optimality are known only for an optimal ECSQ which achieves the lower convex hull of $D_h(R)$ [7]. Now for a discrete source, $D_h(R)$ is never convex since it is decreasing and piecewise constant. On the other hand, it can be shown (using the analytical expression of Berger [1]) that for an exponentially distributed source and the squared error distortion, $D_h(R)$ is convex. It has also been conjectured [6] that $D_h(R)$ is convex for a wide variety of source distributions and distortion measures. Our results for the uniform source demonstrate that $D_h(R)$ can be nonconvex even for "nice" continuous source distributions.

2. Preliminaries

An N-level scalar quantizer Q is a measurable mapping of the real line $\mathbb{R}$ into a finite or countably infinite set of distinct reals $\{y_1, \ldots, y_N\}$ called the *codebook* of Q. (In case the codebook is not finite, we formally define $N = \infty$ and call Q an infinite-level quantizer.) The *codepoints* y_i and the associated *quantization cells* $S_i = \{x : Q(x) = y_i\}$, $i = 1, \ldots, N$ completely characterize Q since the S_i form a partition of $\mathbb{R}$ and

$$Q(x) = y_i \quad \text{if} \quad x \in S_i.$$

The distortion of Q in quantizing a real random variable X with distribution μ_X is measured by the expectation

$$D(Q) = E[d(X, Q(X))] = \int_{-\infty}^{\infty} d(x, Q(x))\, \mu_X(dx)$$

where the *distortion measure* $d(\cdot,\cdot)$ is a nonnegative measurable function of two real variables. The entropy-constrained *rate* of Q is the entropy of the discrete random variable $Q(X)$:

$$H(Q) = -\sum_{i=1}^{N} P[X \in S_i] \log P[X \in S_i]$$

where $\log$ denotes the natural logarithm ($H(Q)$ is measured in nats). A scalar quantizer whose rate is measured by $H(Q)$ is called an *entropy-constrained scalar quantizer* (ECSQ).

For any $R \geq 0$ let $D_h(R)$ denote the lowest possible distortion of any quantizer with output entropy not greater than R. This function is formally defined by

$$D_h(R) = \inf\{D(Q) : H(Q) \leq R\}$$

where the infimum is taken over all finite or infinite-level scalar quantizers whose entropy is less than or equal to R. Any Q that achieves $D_h(R)$ in the sense that $H(Q) \leq R$ and $D(Q) = D_h(R)$ is called an *optimal* ECSQ.

3. Optimal ECSQ for a uniform source

A scalar quantizer is called *regular* if each cell S_i is an interval and each codepoint y_i lies inside S_i. Assume that the distortion measure is of the form

$$d(x, y) = \rho(|x - y|) \tag{1}$$

where $\rho : [0, \infty) \to [0, \infty)$ is a strictly increasing function. For such distortion measures, nearest neighbor quantizers (i.e., quantizers which satisfy $d(x, Q(x)) = \min_{1 \leq i \leq N} d(x, y_i)$ for all x) are regular, and thus an optimal fixed-rate N-level quantizer (i.e., a quantizer which has minimum distortion among all N-level quantizers) can be assumed to be regular.

Unfortunately, an optimal ECSQ is not necessarily a nearest neighbor quantizer, and thus in general it is incorrect to restrict attention to regular quantizers (or quantizers with interval cells) when searching for an optimal ECSQ. Indeed, it is not hard to construct a discrete source with three real-valued outputs for which the unique optimal ECSQ is not regular at certain rates. We note here that a nearest neighbor type condition does hold for an optimal ECSQ which achieves the lower convex hull of $D_h(R)$, implying that such a quantizer can be assumed to be regular [7]. However, as Corollary 2 later shows, an ECSQ achieving the lower convex hull of $D_h(R)$ may not exist for most rate constraints. More recently, it has been shown [8] for continuous source distributions and distortion measures in the form $d(x, y) = \rho(|x - y|)$, where ρ is an increasing convex function, that if an optimal ECSQ exists for a given rate constraint, then there is an optimal ECSQ for the same rate constraint which is regular.

The question of ECSQ regularity is much simpler if the source is uniformly distributed. Let X have a uniform distribution on the unit interval $(0,1)$ and assume that the distortion measure is in the form of (1). Let Q be any finite or infinite-level quantizer with cells $\{S_1,\dots,S_N\}$ and codepoints $\{y_1,\dots,y_N\}$, and define $p_i = \lambda(S_i \cap (0,1))$ for $i = 1,\dots,N$, where λ denotes the Lebesgue measure. Then we can define a new quantizer $\hat{Q}$ over $(0,1)$ which has N interval cells of length p_i, $i = 1,\dots,N$, and N codepoints which are located at the midpoints of these cells (the definition of $\hat{Q}$ outside $(0,1)$ is immaterial). The distortion of $\hat{Q}$ is

$$D(\hat{Q}) = \sum_{i=1}^{N} \Phi(p_i) \tag{2}$$

where $\Phi(p)$ is defined for all $p \geq 0$ by

$$\Phi(p) = 2\int_0^{p/2} \rho(x)\,dx. \tag{3}$$

Since ρ is increasing, it is easy to see that for all $i = 1,\dots,N$,

$$\int_{S_i\cap(0,1)} \rho(|x-y_i|)\,dx \geq \int_{-p_i/2}^{p_i/2} \rho(|x|)\,dx$$

and so $D(\hat{Q}) \leq D(Q)$. On the other hand,

$$H(\hat{Q}) = -\sum_{i=1}^{N} p_i \log p_i \tag{4}$$

so that $H(\hat{Q}) = H(Q)$. Consequently, when searching for an optimal ECSQ for the uniform distribution over $(0,1)$, it suffices to consider interval partitions of $(0,1)$ and the associated regular quantizers with codepoints at the midpoints of the intervals. All quantizers in the rest of this paper will be assumed to be of this type. The distortion and rate of any such quantizer are uniquely determined by the cell lengths $\{p_i;\ i = 1,\dots,N\}$ through equations (2) and (4). Note that if N is finite and $p_i = 1/N$, $i = 1,\dots,N$, the resulting quantizer is the N-level uniform quantizer over $(0,1)$.

In what follows we will consider distortion measures in the form $d(x,y) = \rho(|x-y|)$, where $\rho : [0,\infty) \to [0,\infty)$ is strictly increasing, continuous, and $\rho(e^t)$, $t \in \mathbb{R}$, is strictly convex. Examples for such distortion measures include rth power distortions $d(x,y) = |x-y|^r$ with $r > 0$ (in this case $\rho(e^t) = e^{rt}$), and distortion measures $d(x,y) = \rho(|x-y|)$, where ρ is strictly increasing and convex on $[0,\infty)$. Another distortion measure which does not fall into either of these categories but satisfies the requirements is $d(x,y) = \log(1+|x-y|)$.

If $R = 0$, the optimal ECSQ for any source distribution has only one codepoint. The next result shows that if the source has a uniform distribution, then for any rate $R > 0$ there exists an optimal finite-level ECSQ with a very simple structure.

Theorem 1 *Let the source X have uniform distribution over $(0,1)$ and assume that $d(x,y) = \rho(|x-y|)$, where $\rho : [0,\infty) \to [0,\infty)$ is a strictly increasing continuous function such that $\rho(e^t)$ is strictly convex. Then Q is an optimal ECSQ for a rate constraint $R > 0$ if and only if Q has $N = \lceil e^R \rceil$ cells; one cell of length c and $N-1$ cells of length $\frac{1-c}{N-1}$, where c is the unique solution of the equation*

$$-c\log c - (1-c)\log\left(\frac{1-c}{N-1}\right) = R$$

in the interval $(0, \frac{1}{N}]$.

Theorem 1 implies the intuitive result that if $R = \log N$, then the unique optimal ECSQ for the uniform source is the N-level uniform quantizer. If $R < \log N$, then $c < \frac{1-c}{N-1}$, and the optimal ECSQ is no longer unique; there are exactly N such quantizers. Farvardin and Modestino [2] reached the same conclusion for squared error distortion using numerical methods.

Theorem 1 remains valid (after rescaling) if X is uniformly distributed over an arbitrary interval (a,b). To see this, let $\widehat{X} = (X-a)/(b-a)$ and $\widehat{\rho}(x) = \rho(x/(b-a))$. Then $\widehat{X}$ and $\widehat{\rho}$ satisfy the conditions of the theorem, and a quantizer Q is optimal for X and ρ if and only if $\widehat{Q}(x) = \left[Q\big(x(b-a)+a\big) - a\right]/(b-a)$ is optimal for $\widehat{X}$ and $\widehat{\rho}$.

The proof of the theorem, given in [9], has two main parts. First, similar to [1, 2], the usual Kuhn-Tucker conditions of constrained optimization are used to identify necessary conditions for the optimality of an n-level ECSQ for a fixed positive integer n. After eliminating all quantizers not satisfying these conditions, we are left with the family of n-level quantizers over $(0,1)$ which satisfy the entropy constraint with equality and whose cell lengths can take only two distinct values (these quantizers were also identified in [2]). The second, harder part of the proof consists of identifying, for a fixed n, the quantizers which have minimal distortion in this family, and then finding the optimal choice of n.

Using equations (2) and (4), the distortion and entropy of an N-level quantizer Q with cell lengths $p_1 = c$ and $p_i = \frac{1-c}{N-1}$, $i = 2,\dots,N$ are given respectively by

$$D(Q) = \Phi(c) + (N-1)\Phi\left(\frac{1-c}{N-1}\right).$$

and

$$H(Q) = -c\log c - (1-c)\log\left(\frac{1-c}{N-1}\right)$$

It is easy to see that $H(Q)$ is a strictly increasing function of $c \in [0, \frac{1}{N}]$. Also, $H(Q) = \log(N-1)$ if $c = 0$, $H(Q) = \log N$ if $c = \frac{1}{N}$, and the corresponding quantizers are the $(N-1)$-level and N-level uniform quantizers, respectively. Thus Theorem 1 yields the following parametric description of $D_h(R)$ for the uniform distribution.

Corollary 1 *For X and $d(x,y)$ as in Theorem 1, and for any positive integer $N \geq 2$, the parametric representation of $D_h(R)$ in the interval $[\log(N-1), \log N]$ is given by*

$$R_c = -c \log c - (1-c) \log \left(\frac{1-c}{N-1} \right)$$

$$D_h(R_c) = \Phi(c) + (N-1)\Phi \left(\frac{1-c}{N-1} \right)$$

where $c \in [0, \frac{1}{N}]$.

From the parametric representation we can deduce some important properties of $D_h(R)$. For example, it immediately follows that $D_h(R)$ is everywhere continuous. Moreover, plotting $D_h(R)$ for the squared error distortion $d(x,y) = (x-y)^2$ (see Figure 1) suggests that the graph of $D_h(R)$ consists of smooth, concave pieces joined in a non-smooth manner at rates $R = \log N$. The next result proves these properties of $D_h(R)$ under more general conditions.

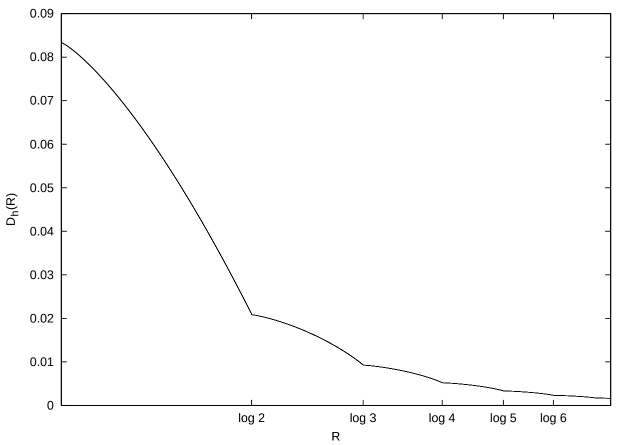

Figure 1: $D_h(R)$ for the uniform source and squared error distortion.

Corollary 2 *With the conditions of Theorem 1, $D_h(R)$ has the following properties:*

(i) *$D_h(R)$ is continuously differentiable on each open interval $(\log(N-1), \log N)$, where $N \geq 2$ is positive integer. At $R = \log N$, the right derivative of $D_h(R)$ is zero for all $N \geq 1$, and the left derivative of $D_h(R)$ is negative for all $N \geq 2$. Thus $D_h(R)$ is not differentiable at the points $R = \log N$ for $N \geq 2$.*

(ii) *Let $d(x,y) = |x-y|^r$ be the rth power distortion with $0 < r \leq 3$. Then $D_h(R)$ is strictly concave on each interval $[\log(N-1), \log N]$ for $N \geq 2$.*

Part (i) of the preceding corollary implies that $D_h(R)$ is not convex. Moreover, part (ii) shows that for the squared error distortion an ECSQ achieving the lower convex hull of $D_h(R)$ exists only at the discrete rate values $R = \log N$. This fact suggests that an ECSQ which achieves the lower convex hull of $D_h(R)$ is the exception rather than the rule.

References

[1] T. Berger, "Optimum quantizers and permutation codes," *IEEE Trans. Inform. Theory*, vol. IT-18, pp. 759–765, Nov. 1972.

[2] N. Farvardin and J. W. Modestino, "Optimum quantizer performance for a class of non-Gaussian memoryless sources," *IEEE Trans. Inform. Theory*, vol. IT-30, pp. 485–497, May 1984.

[3] J. C. Kieffer, T. M. Jahns, and V. A. Obuljen, "New results on optimal entropy-constrained quantization," *IEEE Trans. Inform. Theory*, vol. 34, pp. 1250–1258, Sep. 1988.

[4] A. N. Netravali and R. Saigal, "Optimum quantizer design using a fixed-point algorithm," *Bell Syst. Tech. J*, vol. 55, pp. 1423–1435, Nov. 1976.

[5] P. A. Chou, T. Lookabaugh, and R. M. Gray, "Entropy-constrained vector quantization," *IEEE Trans. Acoust. Speech, Signal Processing*, vol. ASSP-37, pp. 31–42, Jan. 1989.

[6] D. L. Neuhoff and R. K. Gilbert, "Causal source codes," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 701–713, Sep. 1982.

[7] P. A. Chou and B. J. Betts, "When optimal entropy-constrained quantizers have only a finite number of codewords." in *Proc. IEEE Int. Symp. Information Theory* (Cambridge, MA, USA, Aug. 16-21, 1998), p. 97.

[8] A. György and T. Linder, "On the structure of optimal entropy-constrained quantizers." preprint, 2000.

[9] A. György and T. Linder, "Optimal entropy-constrained scalar quantization of a uniform source." submitted to *IEEE Trans. Inform. Theory*, 2000.