On Zador's Entropy-Constrained Quantization Theorem

Robert M. Gray[1]　　Jia Li

Information Systems Lab, Department of Electrical Engineering, Stanford University
Stanford, CA 94305. rmgray@stanford.edu
Department of Statistics, The Pennsylvania State University
University Park, PA 16802. jiali@stat.psu.edu

Abstract

Zador's classic result for the asymptotic high-rate behavior of entropy-constrained vector quantization is recast in a Lagrangian form which better matches the Lloyd algorithm used to optimize such quantizers. A proof that the result holds for a general class of distributions is sketched.

1 Introduction

In his classic Bell Labs Technical Memo, Paul Zador established the optimal tradeoff between average distortion and entropy for entropy-constrained vector quantization in the limit of high rate [6]. The history and generality of the result may be found in in [4]. Optimality properties and generalized Lloyd algorithms for quantizer design, however, require a Lagrangian formulation [1]. In addition, the Lagrangian form turns out to be more convenient for problems involving multiple codebooks such as coding for mixtures since it obviates the need for optimizing rate allocation, as Zador does in his proof. We here recast Zador's theorem in a Lagrangian form and sketch its proof under the assumption that the distribution of the random vector is absolutely continuous with respect to Lebesgue measure.

2 Vector Quantization

Consider the measurable space $(\Omega, \mathcal{B}(\Omega))$ consisting of k-dimensional Euclidean space $\Omega = \Re^k$ and its Borel sets. Assume that X is random vector with a distribution P_f which is absolutely continuous w.r.t. Lebesgue measure V and hence possesses a probability density function (pdf) $f = dP_f/dV$ so that $P_f(F) = \int_F f(x)dV(x) = \int_F f(x)\,dx$. The volume of a set $F \in \mathcal{B}$ is given by its Lebesgue measure $V(F) = \int_F dx$. We assume that the the differential entropy $h(f) \triangleq -\int dx\, f(x) \ln f(x)$ exists and is finite. The unit of entropy is nats or bits according to whether the base of the logarithm is 2 or e. Usually nats will be assumed, but bits will be used when entropies appear in an exponent of 2 and in coding arguments. The relative entropy between

[1]This work was supported by the National Science Foundation under NSF Grants MIP-9706284-001 and CCR-0073050.

two distributions P_f and P_g with pdfs f and g is given by Gelfand's theorem as

$$H(f||g) = \sup_{\mathcal{S}} \sum_i P_f(S_i) \ln \frac{P_f(S_i)}{P_g(S_i)} = \int dx f(x) \ln \frac{f(x)}{g(x)} \geq 0,$$

where the supremum is over all finite partitions $\mathcal{S} = \{S_i\}$.

A vector quantizer q can be described by the following mappings and sets: an encoder $\alpha : \Omega^k \to \mathcal{I}$, where $\mathcal{I} = \{0, 1, 2, \ldots\}$ is an index set, an associated partition $\mathcal{S} = \{S_i;\ i \in \mathcal{I}\}$ such that $\alpha(x) = i$ if $x \in S_i$, a decoder $\beta : \mathcal{I} \to \Omega^k$, an associated reproduction codebook $\mathcal{C} = \{\beta(i);\ i \in \mathcal{I}\}$, an index coder $\psi : \mathcal{I} \to \{0,1\}^*$, the space of all variable-length binary strings, and the associated length function $\ell : \mathcal{I} \to \{1, 2, \ldots\}$ defined by $\ell(i) = \text{length}(\psi(i))$. ψ is assumed to be invertible (a lossless or noiseless code). The overall quantizer is $q(x) = \beta(\alpha(x))$

For simplicity we assume squared error distortion with average

$$D_f(q) = E_f d(X, q(X)) = \sum_i \int_{S_i} dx\, f(x) ||x - y_i||^2 = \sum_i \int_{S_i} dx\, f(x) \sum_{l=0}^{k-1} |x_l - y_{i,l}|^2.$$

The instantaneous rate is $r(\alpha(x)) = \ell(\psi(\alpha(x)))$, the number of bits required to specify the index $i = \alpha(x)$ to the decoder. The average rate is

$$R_f(q) = E_f r(\alpha(X)) = \sum_i P_f(S_i) \ell(\psi(i)).$$

The optimal performance is the minimum distortion achievable for a given rate: $\delta_f(R) = \inf_{q: R_f(q) \leq R} D_f(q)$. The traditional form of Zador's theorem states that under suitable assumptions on f,

$$\lim_{R \to \infty} 2^{\frac{2}{k} R} \delta_f(R) = b(2, k) 2^{\frac{2}{k} h(f)} \tag{1}$$

where $b(2, k)$ is Zador's constant, which depends only on k and not f. Zador's argument explicitly requires that his asymptotic result for fixed-rate coding holds and that $h(f)$ is finite. Zador's fixed rate conditions have been generalized through the years (see, e.g., [3]), but his variable results have not been similarly extended and there are problems with Zador's proof which limit its applicability to densities with bounded support.

3 The Lagrangian Formulation

The Lagrangian formulation of variable rate vector quantization [1] defines for each value of a Lagrangian multiplier $\lambda > 0$ a Lagrangian distortion $\rho_\lambda(x, i) = d(x, \beta(i)) + \lambda \ell(\psi(i))$, a corresponding performance

$$\rho_(f, \lambda, q) = E_f\left(d(X, q(X)) + \lambda E_f \ell(\psi(\alpha(X)))\right) = D_f(q) + \lambda R_f(q),$$

and an optimal performance $\rho(f, \lambda) = \inf_q \rho(f, \lambda, q)$. Each λ yields a distortion-rate pair on the operational distortion-rate function curve. Standard arguments imply that small λ corresponds to high rate and large λ corresponds to small rate. The Lagrangian formulation yields Lloyd optimality conditions for vector quantizers. In particular, for a given decoder (satisfying the usual centroid condition) and index coder, the optimal encoder is $\alpha(x) = \text{argmin}_i\,(d(x, y_i) + \lambda\ell(\psi(i)))$. Optimal choice of the index coder and the Kraft inequality ensure that $H_f(q(X)) \leq E_f[\ell(\psi(\alpha(X)))] < H_f(q(X)) + 1$, where

$$H_f(q(X)) = -\sum_i P_f(S_i) \ln P_f(S_i).$$

This can also be achieved, e.g., by choosing lengths $\ell(\psi(i)) = \lceil -\log P_f(\alpha(X) = i)\rceil$ and hence it is common to make the approximation that

$$\ell(\psi(i)) \approx -\log P_f(\alpha(X) = i), \quad R_f(q) \approx E_f\ell(\psi(\alpha(X))) = H_f(q(X)),$$

resulting in entropy constrained vector quantization (ECVQ).

Our main result is the following.

Theorem 1 *Assume that f is absolutely continuous with respect to Lebesgue measure and that $h(f)$ is finite. Then*

$$\lim_{\lambda\to 0}\left(\frac{\rho(f,\lambda)}{\lambda} + \frac{k}{2}\ln\lambda\right) = \theta_k + h(f) \tag{2}$$

where

$$\theta_k = \theta([0,1)^k) \triangleq \inf_{\lambda>0}\left(\frac{\rho(u_1,\lambda)}{\lambda} + \frac{k}{2}\ln\lambda\right) \tag{3}$$

and u_1 is the uniform pdf on the k-dimensional unit cube C_1^k

Comment: It is shown in [5] that the that 1 holds if and only if 2 holds, in which case $\theta_k = \frac{k}{2}\ln\frac{2e}{k}b_{2,k}$, so that the two formulations are indeed equivalent.

The following notation will be used:

$$\theta(f,\lambda,q) = \frac{D_f(q)}{\lambda} + H_f(q(X)) - h(f) + \frac{k}{2}\ln\lambda$$

$$\theta(f,\lambda) = \inf_q \theta(f,\lambda,q),\ \ \overline{\theta}(f) = \limsup_{\lambda\to 0}\theta(f,\lambda),\ \ \underline{\theta}(f) = \liminf_{\lambda\to 0}\theta(f,\lambda).$$

The quantization function $\theta(f,\lambda,q)$ can be rewritten as a weighted sum of relative entropies minus a constant $k\ln\pi$. The nonnegativity of relative entropy then yields the following bound.

Lemma 1 *For any f,λ,q $\theta(f,\lambda,q) \geq -k\ln\pi$ and therefore $\underline{\theta}(f) \geq -k\ln\pi$.*

The following result is proved in [5]:

Lemma 2 *The conclusions of Theorem 1 hold if and only if the limit of (1) exists, in which case*

$$\theta_k = \frac{k}{2}\ln\frac{2e}{k}b(2,k). \tag{4}$$

Mixture sources play a fundamental role in the development. A mixture source is a random pair $\{X, Z\}$, where Z is a discrete random variable with pmf $w_m = P(Z = m)$, $m = 1, 2, \ldots$ and conditional pdf's $f_{X|Z}(x|m) = f_m(x)$ with support Ω_m. The pdf for x is given by

$$f(x) = f_X(x) = \sum_m w_m f_m(x).$$

In the special case where the Ω_m are disjoint, the mixture is said to be *orthogonal*. For an orthogonal mixture define for each m the boundary of Ω_m, $\partial\Omega_m$ as the closure of Ω_m minus the interior of Ω_m. An orthogonal mixture is said to have the zero probability boundaries property if $P_f(\partial\Omega_m) = 0$ for all m.

Suppose that for each f_m we have a quantizer q_m defined on Ω_m, i.e., an encoder $\alpha_m : \Omega_m \to \mathcal{I}$, a partition of Ω_m $\{S_{m,i};\ i = 1, 2, \ldots\}$, and a decoder $\beta_m : \mathcal{I} \to \mathcal{C}_m$. The component quantizers $\{q_m\}$ together imply an overall composite quantizer q with an encoder α that maps x into a pair (m, i) if $x \in \Omega_m$ and $\alpha_m(x) = i$, a partition of Ω $\{S_{m,i};\ i = 1, 2, \ldots, m = 1, 2, \ldots\}$, and a decoder β that maps (m, i) into $\beta_m(i)$, $q(x) = \sum_m q_m(x) 1_{\Omega_m}(x)$. Conversely, an overall quantizer $q : \Omega \to \mathcal{I}$ can be applied to every component in the mixture, effectively implying a component quantizers $q_m(x) = \sum_m q(x) 1_{\Omega_m}(x)$ for all m. In this case the structure is not so simple as quantization cells can straddle boundaries of Ω_m. Here the partition of Ω_m is $\{S_i \cap \Omega_m;\ i = 1, 2, \ldots\}$ and many of the cells may be empty.

Lemma 3 *If f is an orthogonal mixture $\{f_m, w_m\}$ and q is a composite quantizer formed from component quantizers q_m. Then*

$$H_f(q(X)) - h(f) = \sum_m w_m [H_{f_m}(q_m(X)) - h(f_m)], \tag{5}$$

$$\theta(f, \lambda, q) = \sum_m w_m \theta(f_m, \lambda, q_m),\quad \theta(f, \lambda) \le \sum_m w_m \theta(f_m, \lambda),\quad \overline{\theta}(f) \le \sum_m w_m \overline{\theta}(f_m). \tag{6}$$

Proof: If q_n has partition $\{S_{n,l}\}$, then $P_f(S_{n,l}) = \sum_m w_m P_{f_m}(S_{n,l}) = w_n P_{f_n}(S_{n,l})$ since the mixture is orthogonal. Since $f \ln f$ is integrable with respect to Lebesgue measure,

$$H_f(q(X)) - h(f) = \sum_m w_m [H_{f_m}(q_m(X)) - h(f_m)].$$

Proving (5). The remaining relations follow from conditional expectation $E_f||X - q(X)||^2 = \sum_m w_m E_{f_m} ||X - q_m(X)||^2$, the fact that for a given λ and $\epsilon > 0$, q_m can be chosen so that $\theta(f_m, \lambda, q_m) \le \theta(f_m, \lambda) + \epsilon$ for all m and hence

$$\sum_m w_m \theta(f_m, \lambda) + \epsilon \ge \sum_m w_m \theta(f_m, \lambda, q_m) = \theta(f, \lambda, q) \ge \theta(f, \lambda),$$

$$\overline{\theta}(f) = \limsup_{\lambda \to 0} \theta(f, \lambda) \le \sum_m w_m \limsup_{\lambda \to 0} \theta(f_m, \lambda) = \sum_m w_m \overline{\theta}(f_m).$$

Lemma 4 *Given an overall quantizer q. Then*

$$\begin{aligned} H_f(q(X)) - h(f) &= \sum_n w_n [H_{f_n}(q(X)) - h(f_n)] - H(Z|q(X)) && (7) \\ \theta(f, \lambda, q) &= \sum_n w_n \theta(f_n, \lambda, q) - H(Z|q(X)) && (8) \end{aligned}$$

Proof: Suppose that q is a quantizer defined for the entire space $\Omega = \bigcup_n \Omega_n$. Let $\{S_l\}$ be the corresponding partition. Then

$$H_f(q(X)) - h(f) =$$
$$\sum_n w_n[H_{f_n}(q(X)) - h(f_n)] + \sum_n \sum_l P(Z = n, q(X) = l) \ln \frac{P(Z = n, q(X) = l)}{P(q(X) = l)}.$$

Proving (7), which in turn implies (8). Zador is missing the $H(Z|q(X))$ term in his analogous formula on p. 29 in the proof of his Lemma 3.3(b), he tacitly assumes it is 0.

Lemma 5 *Suppose* λ_n, q_n $n \to \infty$ *satisfy* $\lim_{n\to\infty} \lambda_n = 0$, *where the* λ_n *are decreasing, and* $\lim_{n\to\infty} \theta(f, \lambda_n, q_n) = \underline{\theta}(f)$. *Suppose also that* f *is a finite orthogonal mixture* $\{f_m, w_m;\ m = 1, 2, \ldots, M\}$ *which has the zero probability boundaries property. Then* $\lim_{n\to\infty} H(Z|q_n(X)) = 0$.

Proof: Define the sets $G_n = \{x : q_n(x) \in \Omega_{Z(x)}\}$ and the random variables $\phi(x) = 1_{G_n}(x)$. Then

$$H(Z|q_n) \le H(Z, \phi_n|q_n) = H(\phi_n|q_n) + H(Z|\phi_n, q_n) \le H(\phi_n) + H(Z|\phi_n, q_n).$$

Define $p_n = P_f(G_n^c) = \Pr(\phi_n(X) = 0)$. Then

$$H(\phi_n) = h_2(p_n) = -p_n \ln p_n - (1 - p_n)\ln(1 - p_n), \tag{9}$$

and

$$H(Z|\phi_n, q_n) = \sum_{y\in\mathcal{C}_n} H(Z|\phi_n = 0, q_n = y) P_f(\phi_n = 0, q_n = y)$$

since $H(Z|\phi_n = 1, q_n = y) = 0$ for all y (Z is a deterministic function of q_n given $\phi_n = 1$). Thus $H(Z|\phi_n, q_n) \le p_n \ln M$ so that $H(Z|q_n) \le h_2(p_n) + p_n \ln M$. Thus the lemma will be proved if $p_n \to 0$ as $n \to \infty$. Define $A = \bigcup_m \partial\Omega_m$. Since assumed boundaries have zero probability, $P_f(A) = 0$. Define $||x, A|| = \inf_{a\in A} ||x - a||$ and let $\epsilon_n \to \infty$ be a nonnegative decreasing sequence. Then $\bigcup_{n=1}^{\infty}\{x : ||x, A|| > \epsilon_n\} = A^c$. For any $\delta > 0$ $\{x : ||x, A|| > \delta\} \bigcap \{x : ||x - q_n(x)|| \le \delta/2\} \subset G_n$ since if x is at least δ from the nearest boundary point and less than $\delta/2$ from $q_n(x)$, then from the triangle inequality $||q_n(x), A|| \ge \delta/2$ and $q_n(x)$ must be in the same Ω_m as x. Thus $G_n^c \subset \{x : ||x, A|| \le \delta\} \bigcup \{x : ||x - q_n(x)|| > \delta/2\}$ and hence from union bound

$$p_n \le P_f(\{x : ||x, A|| \le \delta\}) + P_f(\{x : ||x - q_n(x)|| > \frac{\delta}{2}\}).$$

From the Tchebychev inequality $P_f(\{x : ||x - q_n(x)|| > \delta/2\}) \le 4D_f(q_n)/\delta^2$. Define $\delta = \delta_n$ by $\frac{\delta^2}{4} = \sqrt{\lambda_n}$. Then $p_n \le P_f(\{x : ||x, A|| \le 2\lambda_n^{\frac{1}{4}}\}) + D_f(q_n)/\sqrt{\lambda_n}$. Since $\lambda_n^{1/4}$ is decreasing, the sets $\{x : ||x, A|| \le 2\lambda_n^{\frac{1}{4}}\}$ are decreasing to

$$\bigcap_{n=1}^{\infty}\{x : ||x, A|| \le 2\lambda_n^{\frac{1}{4}}\} = \left(\bigcup_{n=1}^{\infty}\{x : ||x, A|| > 2\lambda_n^{\frac{1}{4}}\}\right)^c = A,$$

which has zero probability by assumption, hence $\lim_{n\to\infty} P_f(\{x : ||x, A|| \leq 2\lambda_n^{\frac{1}{4}}\}) = 0$. The assumptions of lemma imply that

$$D_f(q_n) \leq \lambda_n \underline{\theta}(f) - \frac{k}{2}\lambda_n \log \lambda_n + \lambda_n h(f) + \lambda_n o(n) \tag{10}$$

and hence $D_f(q_n)/\sqrt{\lambda_n} \to 0$ as $\lambda_n \to 0$, completing the proof of the lemma.

Lemma 6 *Assume a possibly infinite mixture* $\{f_m, \Omega_m, w_m;\ m = 1, 2, \ldots\}$ *which satisfies the zero probability boundary condition and has the property that* $H(Z) < \infty$. *Suppose* λ_n, q_n $n \to \infty$ *satisfy* $\lim_{n\to\infty} \lambda_n = 0$, *where the* λ_n *are decreasing, and* $\lim_{n\to\infty} \theta(f, \lambda_n, q_n) = \underline{\theta}(f)$. *Then* $\lim_{n\to\infty} H(Z|q_n(X)) = 0$.

Proof: Given an orthogonal mixture $\{f_m, \Omega_m, w_m;\ m = 1, 2, \ldots\}$, for any M form $\{f'_m, \Omega'_m, w'_m;\ m = 1, 2, \ldots, M+1\}$ by $f'_m(x) = f(x)/P_f(\Omega'_m)1_{\Omega'_m}(x)$ with

$$\Omega'_m = \begin{cases} \Omega_m & m = 1, 2, \ldots, M \\ \bigcup_{i=M+1}^{\infty} \Omega_i & m = M+1 \end{cases}, \quad w'_m = \begin{cases} w_m & m = 1, 2, \ldots, M \\ s_{M+1} = \sum_{i=M+1} w_m & m = M+1 \end{cases}$$

Fix $\epsilon > 0$ and assume that M is chosen large enough to ensure that

$$h_2(s_{m+1}) < \epsilon, \quad -s_{M+1} \ln s_{M+1} \leq \epsilon, \quad -\sum_{z=M+1}^{\infty} w_z \ln w_z \leq \epsilon.$$

Define

$$Z'_M(x) = \begin{cases} m & \text{if } x \in \Omega_m,\ m = 1, \ldots, M \\ M+1 & \text{otherwise} \end{cases}, \quad \psi_M(x) = \begin{cases} 1 & x \in \bigcup_{i=M+1}^{\infty} \Omega_i \\ 0 & \text{otherwise} \end{cases}$$

and note that $P_f(\psi_M = 1) = s_{m+1}$ and $P_f(\psi_M = 0) = 1 - s_{m+1}$ From the previous lemma, $\lim_{n\to\infty} H(Z'_M|q_n(X)) = 0$ so that

$$\begin{aligned}
H(Z|q_n) &= H(Z, \psi_M|q_n) = H(\psi_M|q_n) + H(Z|\psi_M, q_n) \\
&\leq H(\psi_M) + H(Z|\psi_M, q_n) = h_2(s_{M+1}) + H(Z|\psi_M, q_n) \\
&\leq \epsilon + H(Z|\psi_M, q_n), \\
H(Z|\psi_M, q_n) &= s_{M+1} \sum_{y \in \mathcal{C}_n} P_f(q_n = y|\psi_M = 1) \times H(Z|\psi_M = 1, q_n = y) \\
&\quad +(1 - s_{M+1}) \sum_{y \in \mathcal{C}_n} P_f(q_n = y|\psi_M = 0) H(Z|\psi_M = 0, q_n = y).
\end{aligned}$$

If $\psi_M = 0$, then $Z = Z'_M$ and hence

$$\begin{aligned}
H(Z|\psi_M, q_n) &= s_{M+1} \sum_{y \in \mathcal{C}_n} P_f(q_n = y|\psi_M = 1) H(Z|\psi_M = 1, q_n = y) \\
&\quad +(1 - s_{M+1}) \sum_{y \in \mathcal{C}_n} P_f(q_n = y|\psi_M = 0) H(Z'_M|\psi_M = 0, q_n = y) \\
&\leq s_{M+1} H(Z|\psi_M = 1) + (1 - s_{M+1}) H(Z'_M|\psi_M = 0, q_n),
\end{aligned}$$

since conditioning decreases entropy, and

$$H(Z'_M|q_n) \geq H(Z'_M|\psi_M, q_n) = H(Z'_M|\psi_M = 0, q_n)(1 - s_{M+1})$$

since given $\psi_M = 1$, $Z'_M = M + 1$ and hence $H(Z'_M|\psi_M = 1, q_n) = 0$. Thus $H(Z|\psi_M, q_n) \leq s_{M+1}H(Z|\psi_M = 1) + H(Z'_M|q_n)$. The conditional pmf for Z given $\psi_M = 1$ is w_m/s_{M+1} for $m = M+1, \ldots$ and 0 otherwise. Hence

$$H(Z|\psi_M = 1) = -\sum_{z=M+1}^{\infty} \frac{w_z}{s_{M+1}} \ln \frac{w_z}{s_{M+1}} = \ln s_{M+1} - \frac{1}{s_{M+1}} \sum_{z=M+1}^{\infty} w_z \ln w_z$$

so that combining the pieces yields

$$H(Z|q_n) \leq \epsilon + H(Z|\psi_M, q_n) \leq 3\epsilon + H(Z'|q_n) \to_{n\to\infty} 3\epsilon,$$

proving the lemma.

Combining the lemmas yields the following corollary.

Corollary 1 *Suppose that f is an orthogonal mixture $\{f_m, \Omega_m, w_m\}$ which satisfies the zero probability boundary condition and for which $H(Z) < \infty$ ($Z = m$ if $x \in \Omega_m$) (e.g., the mixture is finite). Then*

$$\sum_m w_m \underline{\theta}(f_m) \leq \underline{\theta}(f) \leq \overline{\theta}(f) \leq \sum_m w_m \overline{\theta}(f_m). \tag{11}$$

Thus if $f_m \in \mathcal{Z}$ for all m, then also $f \in \mathcal{Z}$.

Proof of theorem: **First Step: Uniform pdfs on cubes** Define a cube in Ω^k as $C_a = \{x : 0 < x_i \leq a;\ i = 0, 1, \ldots, k-1\}$ (or any translation of a set of this form). Define the corresponding uniform pdf $u_a(x) = V(C_a)^{-1} 1_{C_a}(x)$. Then $V(C_a) = a^k$, $h(u_a) = \ln V(C_a) = k \ln a$, and $u_a(x) = a^{-k} u_1(\frac{x}{a})$.

Lemma 7 $\theta(u_a, \lambda, q_a) = \theta(u_1, a^{-2}\lambda, q_1)$, $\theta(u_a, \lambda) = \theta(u_1, a^{-2}\lambda)$.

Proof: Suppose have a quantizer q_1 with encoder $\alpha_1 : \mathcal{C}_1 \to \mathcal{I}$ and decoder $\beta_1 : \mathcal{I} \to \mathcal{C}$ defined for the unit cube. Define a quantizer q_a with encoder α_a and decoder β_a for C_a by straightforward variable changes $\alpha_a(x) = \alpha_1(\frac{x}{a})$, $\beta_a(l) = a\beta_1(l)$, $q_a(x) = rq_1(\frac{x}{a})$. Then $H_{u_a}(q_a) = H_{u_1}(q_1)$, $h(u_a) = -\ln a^k + h(u_1)$, $E_{u_a}||X - q_a(X)||^2 = a^2 E_{u_1}||X - q_1(X)||^2$ and hence $\theta(u_a, \lambda) = \theta(u_1, \lambda/a^2)$. Hence we can focus on $u_1(x) = 1_{C_1}(x)$, uniform pdf on unit cube.

Lemma 8 $\lim_{\lambda\to 0} \theta(u_1, \lambda) = \theta_k$.

Proof: Partition the unit cube C_1 into m^k disjoint unit cubes $C_{1/m}$. For each of the small cubes have a uniform pdf $f_{1/m}(x) = m^k$ on the cube. All of the small cubes have the same $\rho(f_{1/m}, \lambda)$. From Lemma 7, $\theta(f_{1/m}, \lambda) = \theta(u_1, m^2\lambda)$. From Lemma 3, $\theta(u_1, \lambda) \leq \sum_{i=1}^{m^k} \frac{1}{m^k} \theta(f_{1/m}, \lambda) = \theta(f_{1/m}, \lambda)$, which with the previous equation implies $\theta(u_1, \lambda) \leq \theta(u_1, m^2\lambda)$. Replacing $m^2\lambda$ by λ, $\theta(u_1, \lambda) \geq \theta(u_1, m^{-2}\lambda)$. Fix λ and note that $(0,\lambda] = \bigcup_{m=1}^{\infty}(\frac{\lambda}{(m+1)^2}, \frac{\lambda}{m^2}]$ so for any λ' between 0 and λ there is an integer m such that $\lambda/(m+1)^2 < \lambda' \leq \lambda/m^2$. $\rho(f,\lambda)$ is nondecreasing with decreasing λ, hence

$$\theta(u_1,\lambda) \geq \frac{\rho(u_1,\lambda')}{(\frac{m+1}{m})^2\lambda'} + \frac{k}{2}\ln\lambda' = (\frac{m+1}{m})^2\theta(u_1,\lambda') + (\frac{2m+1}{m^2+2m+1})\frac{k}{2}\ln\lambda'$$

Choose any subsequence of λ' tending to zero. The largest possible value is $\overline{\theta}(u_1)$ and hence $\theta(u_1,\lambda) \geq \overline{\theta}(u_1)$ which means that $\theta_k \triangleq \inf_\lambda \theta(u_1,\lambda) \geq \overline{\theta}(u_1)$. Hence $\underline{\theta}(u_1) \geq \theta_k \geq \overline{\theta}(u_1)$ and hence the limit $\lim_{\lambda\to 0}\theta(u_1,\lambda)$ must exist and equal θ_k.

Second step: Piecewise constant pdfs on cubes Suppose that $C(n)$ is a collection of disjoint unit cubes, w_m is a pmf, and

$$f(x) = \sum_m w_m \frac{1}{V(C(m))} 1_{C(m)}(x).$$

Combining the previous result and Corollary 1 using the fact that the boundaries of cubes have zero volume and hence also zero probability implies that $f \in \mathcal{Z}$.

Third step: Distributions on the unit cube Let C_1^k denote the k-dimensional unit cube and assume that $P_f(C_1^k) = 1$. For any integer M can partition C_1^k into M^K cubes of side length $1/M$, say $C(m);\ 1,2,\ldots,M^k$. Given a pdf f, form a piecewise constant approximation

$$\hat{f}^{(M)}(x) = \sum_{m=1}^{M^k} \frac{P_f(C(m))}{V(C(m))} 1_{C(m)}(x).$$

This is an orthogonal mixture source with $w_m = P_f(C(m))$ and component pdfs $\hat{f}_m(x) = M^k 1_{C(m)}(x)$. If $\hat{P}_M$ denotes the distribution induced by $\hat{f}^{(M)}$, i.e., $\hat{P}_M(F) = \int_F \hat{f}^{(M)}(x)\,dx$, then $\hat{f}^{(M)} = d\hat{P}_M/dV(x)$.

Lemma 9 $\lim_{M\to\infty}\hat{f}^{(M)}(x) = f(x),\ V-a.e.,\ \lim_{M\to\infty}||\hat{f}^{(M)} - f||_1 = 0,$ $\lim_{M\to\infty} h(\hat{f}^{(M)}) = h(f).$

Proof: The first two results follow by differentiation of measures and Scheffé's lemma (See, e.g., [3], p.88.) The third result follows from the convergence of entropy for uniform scalar quantizers, e.g., [2].

Fix $\lambda > 0$. Suppose q_1 is a quantizer with corresponding encoder α_1, decoder β_1, index coder ψ_1, and length function ℓ_1. Assume that q_1 is optimal for a design pdf g (which will be either f or $\hat{f}^{(M)}$) $S_i = \{x : \alpha_1(x) = i\}$, $l_{i|1} = \ell(\psi_1(i))$, and $p_i = P_g(S_i)$, which are assumed nonincreasing in i. Optimality of the index coder implies that $l_{i|1}$ are nondecreasing. Given any node n in the code tree, define $W_n =$ all x contained in an $S_i \subset W_n$. Choose a node n^* in the code tree that is not a leaf with the property

$$P_g(W_{n^*}) = \sum_{i:S_i\subset W_{n^*}} p_i \leq \epsilon.$$

Call the node n the *flag node* and let $L_\epsilon - 2$ denote the depth of the code tree of this node. A second quantizer q_2 is a uniform k-dimensional quantizer with side-width $\Delta = 1/N$ where $N = \lfloor\sqrt{\lambda}\rfloor$ so that $N \leq \lambda^{-1/2}$, $\Delta \leq \sqrt{\lambda}/1 - \sqrt{\lambda}$, $\Delta^2 \leq$

$\lambda/(1-2\sqrt{\lambda}) = \lambda + o(\lambda^{3/2})$. Each cell is represented by its Euclidean centroid so every input point is within $\Delta/2$ of a reproduction and hence

$$d(x, q_2(x)) \leq k\frac{\Delta^2}{4} \leq \frac{k}{4}\lambda + o(\lambda^{3/2})$$

Use a fixed rate lossless code for q_2, to specify the centroid selected, this will require $L_\lambda = \lceil \ln N^k \rceil \leq \ln N^k + 1 \leq -\frac{k}{2}\ln\lambda + 1$. For reasons to be seen, we instead use a longer fixed rate code with length $l_{i|2} = L_\epsilon - 1 + L_\lambda \leq L_\epsilon - \frac{k}{2}\ln\lambda$. Form a code $\hat{q}$ by merging q_1 and q_2 as follows: Given an input vector x, find the code and index yielding the smallest Lagrangian distortion:

$$(m, i) = (m(x), i(x)) = \operatorname*{argmin}_{l,j} (d(x, \beta_l(j)) + \lambda\ell_l(j))$$

Let $B = \{x : m(x) = 2\}$ (uniform quantizer best). If $x \in B^c \cap W^c_{n^*}$, then the encoded sequence is that produced by q_1: $\overline{\psi}(\overline{\alpha}(x)) = \psi_1(\alpha_1(x))$. Otherwise, either $x \in B$ or $x \in W_{n^*}$. Send the pathmap to n^* (length=$L_\epsilon - 2$) and (1) if $x \in W_{n^*}$, send a 0 (one bit) followed by the remainder of the binary sequence according to q_1. In this case the final codeword has an additional bit, $\overline{l}_i = l_{i|1} + 1$, or (2) otherwise send a 1 (one bit) followed by the fixed rate $\log N$ bit word designating the uniform quantizer output for a total of $l_{i|2}$. By construction,

$$d(x, \overline{q}(x)) + \lambda l(\overline{\psi}\overline{\alpha}((x))) = \min_{l,j} (d(x, \beta_l(j)) + \lambda\ell_l(j)) + 1_{W_{n^*}\cap B^c}(x)$$

and hence

$$\min_{l,j} (d(x, \beta_l(j)) + \lambda\ell_l(j)) \leq d(x, \beta_l(j)) + \lambda\ell_l(j) + 1_{W_{n^*}\cap B^c}(x);\ l = 1, 2. \tag{12}$$

In particular, the upper bound for $l = 2$ implies

$$d(x, \overline{q}(x)) + \lambda l(\overline{\psi}\overline{\alpha}(x)) \leq (\frac{k}{4} + L_\epsilon)\lambda - \frac{k}{2}\lambda\ln\lambda + o(\lambda^{3/2}) \tag{13}$$

which after some algebra yields

$$|\theta(f, \lambda, \overline{q}) - \theta(\hat{f}^{(M)}, \lambda, \overline{q})| \leq (\frac{k}{4} + L_\epsilon + o(\sqrt{\lambda}))||f - \hat{f}^{(M)}|| + |h(f) - h(\hat{f}^{(M)})|. \tag{14}$$

For any q_1 with q_2 and $\overline{q}$ constructed in this way using a design pdf $g = \hat{f}^{(M)}$

$$\theta(f, \lambda) \leq \theta(f, \lambda, \overline{q}) \leq \theta(\hat{f}_M, \lambda, \overline{q}) + (\frac{k}{4} + L_\epsilon + o(\sqrt{\lambda})||f - \hat{f}^{(M)}|| + |h(f) - h(\hat{f}^{(M)})|$$

Using (12) with $l = 1$,

$$\begin{aligned}
\theta(\hat{f}_M, \lambda, \overline{q}) &= \int dx\, \hat{f}^{(M)}(x) \left(\frac{d(x, \overline{q}(x))}{\lambda} + l(\overline{\psi}(x))\right) + \frac{k}{2}\ln\lambda + h(\hat{f}^{(M)}) \\
&\leq \int dx\, \hat{f}^{(M)}(x) \left(\frac{d(x, \beta_1(j))}{\lambda} + \ell_1(j) + 1_{W_{n^*}\cap B^c}(x)\right) + \frac{k}{2}\ln\lambda + h(\hat{f}^{(M)}) \\
&\leq \theta(\hat{f}^{(M)}, \lambda) + 2\epsilon
\end{aligned}$$

since q_1 was assumed approximately optimal for $\hat{f}^{(M)}$. Thus

$$\theta(f,\lambda) \le \theta(\hat{f}^{(M)},\lambda) + 2\epsilon + (\frac{k}{4} + L_\epsilon + o(\sqrt{\lambda}))||f - \hat{f}^{(M)}|| + |h(f) - h(\hat{f}^{(M)})|$$

$$\overline{\theta}(f) \le \theta_k + 2\epsilon + (\frac{k}{4} + L_\epsilon)||f - \hat{f}^{(M)}|| + |h(f) - h(\hat{f}^{(M)})|.$$

Since $\hat{f}^{(M)}$ has the Zador property, letting $M \to \infty$ $\overline{\theta}(f) \le \theta_k + 2\epsilon$. Since $\epsilon > 0$ was arbitrary, $\overline{\theta}(f) \le \theta_k$. The converse inequality is proved in a similar fashion.

Final step: Proof of theorem Carve Ω^k into disjoint unit cubes $C_1(n)$ and write the pdf f as the orthogonal mixture

$$f(x) = \sum_n P_f(C_1(n)) f_n(x), \quad f_n(x) = \frac{f(x)}{P_f(C_1(n))} 1_{C_1(n)}(x).$$

To apply Corollary 1 it must be shown that the boundaries of unit cubes have zero probability and that $H(Z)$ is finite. The first property follows since the boundaries have zero Lebesgue measure and f is absolutely continuous with respect to Lebesgue measure. The second property follows from the limiting properties for uniform quantizers [2], the finiteness of $h(f)$, and the fact that refining partitions increases entropy. Thus the previous lemma and Corollary 1 yield $\theta(f) = \sum_n P_f(C_1(n))\theta(f_n) = \theta_k$, which proves the theorem.

Acknowledgements

The authors acknowledge the many helpful comments of Dave Neuhoff and of the students and colleagues in the Compression and Classification Group of Stanford University.

References

[1] P.A. Chou, T. Lookabaugh, and R.M. Gray, "Entropy-constrained vector quantization," *IEEE Trans. Acoust. Speech & Signal Proc.*, Vol. 37, pp. 31–42, January 1989.

[2] I. Csiszàr, "Generalized entropy and quantization problems," *Proc. Sixth Prague Conf.*, pp. 159–174, 1973.

[3] Siegfried Graf and Harald Luschgy, *Foundations of Quantization for Probability Distributions*, Springer, Lecture Notes in Mathematics, 1730, Berlin, 2000.

[4] R.M. Gray and D.L. Neuhoff, "Quantization," *IEEE Transactions on Information Theory*, Vol. 44, pp. 2325–2384, October 1998.

[5] J. Shih, A. Aiyer, and R.M. Gray, "A Lagrangian formulation of high rate quantization," submitted to ICASSP 2001. Preprint available at `http://ee.stanford.edu/~gray/lagrangezador.pdf`.

[6] P. L. Zador, "Topics in the asymptotic quantization of continuous random variables," Bell Laboratories Technical Memorandum, 1966.