ENTROPY-CONSTRAINED QUANTIZATION AND RELATED PROBLEMS

by

András György

Department of Computer Science and Information Theory
Budapest University of Technology and Economics

Ph.D. Thesis Summary

Supervisors: Dr. László Györfi
Dr. Tamás Linder

Budapest, Hungary
June, 2002

1 Introduction

Nowadays there is an increasing demand to transmit and store larger and larger amounts of information. A possible solution to this is the continuous increase of the capacity of communication and storage channels. Another, much cheaper solution is to utilize better the existing capacities, and transform the data to some compact form which allows more efficient transmission. The process of converting a stream of analog or very high-rate discrete data into a data stream with lower rate (more) suitable for transmission or storage in digital form is called data compression or source coding.

The theory of data compression analyzes an idealized model of communication systems, where an information source must be encoded for transmission to a user over a digital communication or storage channel. The user must be able to decode the transmitted information into a form which is a good approximation of the original data, while using as little of the resources (that is, the channel) as possible. Source coding theory has two main goals: one is to characterize the optimal performance achievable by such a system, that is, to determine the performance of an optimal code which achieves the best possible fidelity of reproduction while satisfying certain given constraints on the channel use; the other is to determine coding techniques which perform close to the optimum.

Data compression methods fall into two main types. The first category is lossless data compression, where the reconstructed data has to match the original. This natural requirement is the most common, for example, in compressing originally digital data, such as text or computer files. However, in many situations we only require the reconstructed data to match the original data only within a certain fidelity. In this case we talk about lossy data compression. The advantage of this method is that utilizing the allowed distortion in the reconstructed data, one can obtain better compression than in the lossless case. Furthermore, in some cases, for example, if the source has a continuous output, it is not possible to have perfectly matching reconstructed data. Therefore, lossy compression is natural, for example, for voice and image coding, and in situations where only certain characteristics of the original data have to be preserved.

Contrary to the lossless case, where optimal methods achieving the best possible compression are known, no such optimal constructions exist in general in the lossy case. In this thesis we are concerned with the latter, and study the lossy coding of real valued sources.

A very general and popular method for lossy data compression is called vector quantization. When quantizing an information source $\mathcal{X} = \{X_1, X_2, \ldots\}$, we parse the input sequence into nonoverlapping blocks of k successive source symbols $\{X_{kn+1}^{(n+1)k}\}$ (here for integers $j > i$, X_i^j denotes $X_i, X_{i+1}, \ldots, X_j$) and encode each block $X_{kn+1}^{(n+1)k}$ using the same time invariant map.

Definition 3.1 *A k-dimensional* vector quantizer *Q is a measurable mapping of a subset $\mathcal{S}$ of the k-dimensional Euclidean space $\mathbb{R}^k$ into a finite or countably infinite set of distinct k-dimensional real vectors $\{c_i, i \in \mathcal{I}\} \subset \mathbb{R}^k$ called the* codebook *of Q, where $\mathcal{I}$ is a countable index set. The elements of the codebook, c_i, are called the* code points *of Q, and the associated sets $S_i = \{x : Q(x) = c_i\}$, $i \in \mathcal{I}$, are called the* quantization cells *of Q.*

Throughout we will use the short notation $Q \equiv \{(S_i, c_i), i \in \mathcal{I}\}$, and $\mathcal{Q}_k$ will denote the set of k-dimensional vector quantizers. If $\mathcal{I}$ is a finite set with N elements, then Q is called an N-point (or N-level) quantizer. If $\mathcal{I}$ is infinite, then Q is an infinite-point quantizer.

We assume that the source is k-stationary; then, since the quantizer performs the same transformation on each block, it is enough to analyze its performance on one block. By stationarity there is a generic $\mathbb{R}^k$ valued random vector X having the same distribution μ as each block $X_{nk+1}^{(n+1)k}$, and we assume that $X \in \mathcal{S}$ with probability 1. The distortion of Q, that is, the fidelity of reproduction in quantizing the random variable X is defined by the expectation

$$D(Q) = E\{d(X, Q(X))\} = \int_{\mathcal{S}} d(x, Q(x))\, d\mu(x)$$

where the *distortion measure* $d : \mathbb{R}^k \times \mathbb{R}^k \to [0, \infty)$ is a nonnegative measurable function of two k-dimensional real vector variables.

The amount of channel use is measured by the rate of the quantizer. As it is common in the literature, we assume the channel to be noiseless and, for simplicity, binary. Then the output of Q must be encoded using a binary code in order that it can be transmitted over the channel. If the output of the N-point quantizer Q is encoded via a fixed-length lossless code, then Q is called a *fixed-rate* quantizer, and its rate is measured by

$$r_f(Q) = \frac{1}{k} \log N$$

the average number of bits necessary to encode each output value. It is clear that fixed-rate quantization does not make sense if Q has infinitely many code points.

An N-point fixed-rate quantizer Q^* is called *optimal* if

$$D(Q^*) = \inf\{D(Q) : Q \in \mathcal{Q}_k \text{ is an } N\text{-point quantizer}\}.$$

Optimal fixed-rate quantizers have been studied extensively. Necessary conditions for the optimality of fixed-rate quantizers, namely the nearest neighbor and the centroid conditions, known as the *Lloyd-Max conditions*, were given in [14, 15, 13]. An important consequence of these conditions is that "good", and hence also optimal fixed-rate quantizers can be assumed to be regular for a wide class of nondecreasing difference distortion measures, that is, they have convex cells containing the corresponding code points. Regular quantizers have a simple parametric description (the cells of a finite-point regular quantizer are convex polytopes), and they can be implemented efficiently using hyperplane decision functions.

Unexpected problems arise when one wants to generalize the concept of regularity to infinite-point quantizers even in the scalar case, and the definition of such quantizers is not unified in the literature. For example, a scalar quantizer with cells

$$(-\infty, 0), [1, \infty), [0, 1/2), [1/2, 3/4), [3/4, 7/8), \ldots, [1 - 2^{-n}, 1 - 2^{-(n+1)}), \ldots$$

has convex cells and so can be thought to be regular with appropriate code points; however there are infinitely many code points in any neighborhood of the point 1. To avoid such cases, an additional condition is introduced to obtain a general definition of regularity, which is automatically satisfied for all finite-point quantizers satisfying the original definition.

Definition 3.3 *We call a quantizer Q* regular *if it has convex cells, the code points lie inside the corresponding cells, and the set of the quantization cells is locally finite, that is, any bounded subset of $\mathbb{R}^k$ intersects only finitely many cells of Q.*

Fixed-rate quantizers represent all output points with codewords of the same length. However, using variable-rate (block) encoding of the outputs, further compression can be achieved, yielding a possibly much smaller average codeword length. If the expected codeword length in coding the random output $Q(X)$ is L, then it is reasonable to define the rate as $r_v(Q) = \frac{1}{k}L$, and call Q a *variable-rate* quantizer. It is known that if the variable-rate code is chosen optimally, then $\frac{1}{k}H(Q) \leq \frac{1}{k}L < \frac{1}{k}H(Q) + \frac{1}{k}$ [6], where

$$H(Q) = H(Q(X)) = -\sum_{i \in \mathcal{I}} P\{X \in S_i\} \log P\{X \in S_i\}$$

is the entropy of the quantizer output. Thus, it is appropriate to redefine the rate as $\frac{1}{k}H(Q)$.

Definition 3.5 *A vector quantizer $Q \in \mathcal{Q}_k$ whose rate is measured by $\frac{1}{k}H(Q)$ is called an* entropy-constrained vector quantizer *(ECVQ).*

Since the entropy is an explicit function of the cell probabilities, ECVQs are much more tractable than variable-rate quantizers. Furthermore, the entropy-constrained formulation is also useful because it does not tie the quantizer design to a particular coding scheme (e.g., scalar entropy coding), providing a more universal measure of quantizer performance.

For any $R \geq 0$ let $D_h(R)$ denote the lowest possible distortion of any vector quantizer with output entropy not greater than R in quantizing the $\mathbb{R}^k$ valued random variable X with distribution μ (the dimension of the quantizers is necessarily k). This function, called the operational distortion-rate curve, is formally defined by

$$D_h(R) = \inf\{D(Q) : H(Q) \leq R, Q \in \mathcal{Q}_k\}$$

where the infimum is taken over all finite- or infinite-point k-dimensional vector quantizers whose entropy is less than or equal to R. An ECVQ is called *optimal* if it achieves $D_h(R)$ in the sense that $H(Q) \leq R$ and $D(Q) = D_h(R)$ (the optimality of variable-rate quantizers is defined in a similar manner).

In contrast to fixed-rate quantization, there is little known about optimal entropy-constrained vector quantization. It seems that in the literature efforts have focused on finding necessary conditions for the optimality of a *regular* entropy-constrained scalar quantizer (ECSQ) with a *fixed* number of output points N for the squared-error distortion [2, 8, 11]. These conditions give rise to practical algorithms for designing an ECSQ with a fixed number of output points [2, 16, 8, 11, 5]. The algorithm of [5] can also be applied for locally optimal vector quantizer design. However, the structure and existence of optimal entropy-constrained quantizers is not known in general. Furthermore, there are no analytic expressions for $D_h(R)$ even for simple source distributions, except for the exponential source [2], and the general properties of $D_h(R)$ are also unknown. In this thesis we provide some answers to these problems.

2 Results

2.1 Optimal entropy-constrained scalar quantization of a uniform source

Finding optimal entropy-constrained quantizers and the corresponding operational distortion-rate curve $D_h(R)$ for a given source distribution analytically is a rather hard problem even in the scalar case. Since the structure of "good" entropy-constrained scalar quantizers is not known in general, the problem formulation itself is a hard task. Moreover, finding optimal quantizers in the restricted class of regular quantizers is still not easy, since one must find the optimum performance over N-point quantizers for all N. A notable exception is the case of an exponentially distributed source and the mean squared distortion considered by Berger [2]. 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; the resulting quantizer is an infinite-point uniform quantizer.

However, if the source is uniformly distributed over an interval, then it is not hard to see that "good" quantizers are regular, and such quantizers can be described by their cell probabilities (in this section we restrict ourselves to such quantizers only). This observation leads to the following result, which is the first rigorously proved analytical result concerning optimal ECSQs for a generally used source distribution.

Theorem 4.1 (György and Linder [J2, C3]) *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 $N=\lceil 2^R\rceil$ 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 $(0,\frac{1}{N}]$.

The theorem remains valid (after rescaling) if X is uniformly distributed over an arbitrary interval (a,b). From this result a parametric expression for $D_h(R)$ is determined in Corollary 4.1. For the squared error distortion $D_h(R)$ is shown in Figure 1. It can be seen that $D_h(R)$ is continuous and piecewise concave, and it is proved in Corollary 4.2 that these properties also hold for more general distortion measures. This shows that $D_h(R)$ may be nonconvex even for "nice" source distributions.

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

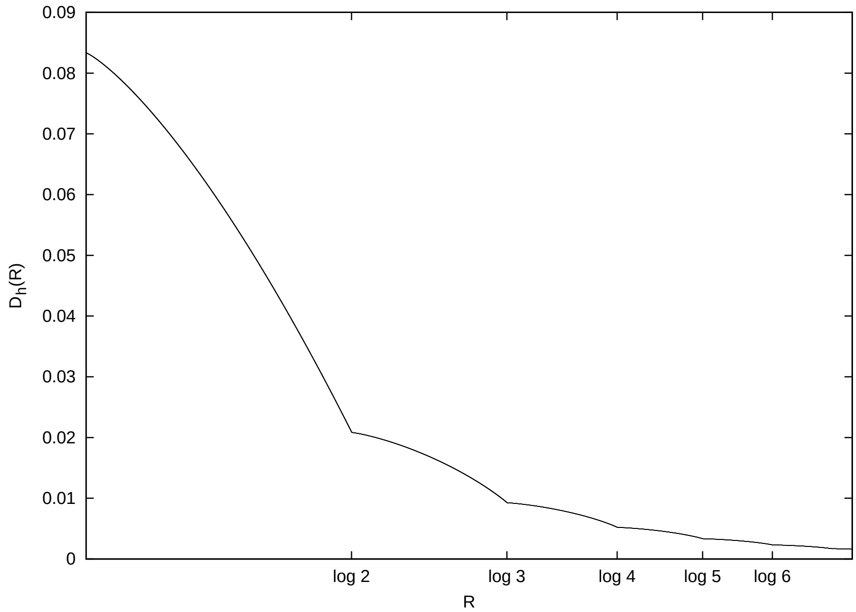

Figure 1: $D_h(R)$ for the uniform source and the squared error distortion.

2.2 Lagrangian-optimal ECVQs

In this section we analyze a special class of optimal ECVQs, ECVQs achieving the lower convex hull of $D_h(R)$. Although, as we have seen for the uniform distribution, such quantizers may exist only at discrete rates, they are important both theoretically and practically [5, 18, 9]. Moreover, this class of optimal quantizers is much more tractable than the class of optimal ECVQs, since they satisfy the generalized nearest neighbor condition [5, 4] (see Lemma 5.1 in the thesis).

The Lagrangian formulation of entropy-constrained quantization of an $\mathbb{R}^k$ valued random variable X defines for each value of a parameter $\lambda > 0$ the *Lagrangian performance* of a quantizer Q by

$$J(\lambda, Q) = D(Q) + \lambda H(Q).$$

The optimum Lagrangian performance is given by

$$J^*(\lambda) = \inf_{Q \in \mathcal{Q}_k} J(\lambda, Q) = \inf_{Q \in \mathcal{Q}_k} \{D(Q) + \lambda H(Q)\}$$

where the infimum is taken over all finite or infinite-point k-dimensional vector quantizers $Q \in \mathcal{Q}_k$. Any quantizer Q that achieves the infimum above is called a *Lagrangian-optimal* entropy-constrained quantizer. It is easy to see that Q is Lagrangian-optimal for some $\lambda \geq 0$ if and only if Q achieves the lower convex hull of $D_h(R)$ in the sense that $(H(Q), D(Q))$ is a point on the lower convex hull ($-\lambda$ is the slope of a line that supports the lower convex hull and passes through this point). First we show that Lagrangian-optimal ECVQs exist under general conditions on the distortion measure.

Theorem 5.1 (György and Linder [C5]) *Assume that for any $x \in \mathbb{R}^k$ the nonnegative distortion measure $d(x,y)$ is a lower semicontinuous function of y such that for any $y' \in \mathbb{R}^k$, $d(x,y') \le \liminf_{\|y\|\to\infty} d(x,y)$. Then for any $\lambda > 0$ there is a Lagrangian-optimal quantizer, that is, there exists Q such that*

$$D(Q) + \lambda H(Q) = J^*(\lambda).$$

In [4] Chou and Betts proved that if the tail of the source distribution is sufficiently light with respect to the distortion measure, then the Lagrangian-optimal quantizer has only finitely many code points. A trivially generalized version of this result is given in Theorem 5.2 in the thesis. We also prove that if the tail of the source distribution is slightly heavier, then the Lagrangian-optimal quantizer has infinitely many code points.

Theorem 5.3 *Assume a difference distortion measure $d(x,y) = \rho(\|x-y\|)$, where $\rho : [0,\infty) \to [0,\infty)$ is strictly increasing and convex. Suppose $J^*(\lambda) < \infty$ for some $\lambda > 0$ and let Q be a Lagrangian-optimal quantizer. If for some $0 < \epsilon < 1$*

$$\limsup_{t\to\infty} \frac{P\{\|X\| > t\}}{2^{-\rho((1-\epsilon)t)/\lambda}} > 0$$

then Q has infinitely many code points.

In particular, Theorems 5.2 and 5.3 hold for the squared error distortion measure. In this case, we obtain that the Gaussian distribution is a breakpoint. For distributions with tails lighter than the tail of a Gaussian distribution (including distributions with finite support), the optimal entropy-constrained quantizer must have only a finite number of code points, and for distributions with tails heavier than that of the Gaussian, the optimal entropy-constrained quantizer has an infinite number of code points. The Gaussian case itself is of particular interest. For a Gaussian source, the results show that there is a critical value $\lambda^* > 0$ (and a corresponding critical rate $R^* > 0$) such that the Lagrangian-optimal quantizer Q has a finite number of code points if $\lambda > \lambda^*$ (i.e., $H(Q) < R^*$), and it has an infinite number of code points if $\lambda < \lambda^*$ (i.e., $H(Q) > R^*$).

2.3 The structure and existence of optimal ECSQs for nonatomic distributions

As we mentioned in the Introduction, the nearest neighbor condition implies that optimal fixed-rate quantizers can be assumed to be regular for a wide class of difference distortion measures. This fact reveals an important structural property of optimal fixed-rate quantizers and it leads to efficient design and simple implementation procedures. Moreover, regularity allows the use of efficient tools in proving the existence and exploring the properties of optimal quantizers. Although the assumption of quantizer regularity seems to be ubiquitous in the literature on entropy-constrained quantization, it is still a fundamental question whether it is sufficient to consider only regular quantizers when searching for a (mean-square) optimal entropy-constrained quantizer. The next example shows the negative answer to this question for sources with discrete distributions.

Example 6.1 (György and Linder [C4, J3]) Let $d(x,y) = (x-y)^2$, and let X be a discrete random variable taking values in $\{-1,0,2\}$ with the following distribution:

$$P\{X=-1\} = P\{X=2\} = 2/5 \quad \text{and} \quad P\{X=0\} = 1/5.$$

A quantizer Q for X is in effect defined by a partition of the set $\{-1,0,2\}$ (the definition of Q for other values is immaterial) and by the corresponding code points (the optimal code points for each partition can easily be found). Checking the five possible partitions, it turns out that for all R such that $h_b(1/5) \le R < h_b(2/5)$, where $h_b(x) = -x\log x - (1-x)\log(1-x)$, an optimal ECSQ has two code points $c_1 = 1/2$ and $c_2 = 0$, and the corresponding cells must satisfy

$$\{-1,2\} \subset S_1 \quad \text{and} \quad \{0\} \subset S_2.$$

Thus if S_2 is an interval, then S_1 is a union of two disjoint intervals, each with positive probability. It follows that no regular quantizer can be optimal in this case.

In this section we show that such anomalous behavior is not possible for scalar sources with nonatomic distributions. The next definition characterizes the "good" distortion measures which allow "regularizing" finite-point quantizers in the sense that any finite-point quantizer can be replaced with a regular quantizer with the same cell probabilities and equal or less distortion.

Definition 6.1 (György and Linder [J3]) *Given a positive integer N we say that a distortion measure $d : \mathbb{R}\times\mathbb{R} \to [0,\infty)$ is* N-point regular *if for any nonatomic distribution μ on $\mathbb{R}$ and any N-point quantizer Q there is a* regular *N-point quantizer $\widehat{Q}$ such that $D(\widehat{Q}) \le D(Q)$ and Q and $\widehat{Q}$ have the same cell probabilities (i.e., there is a one-to-one mapping π between the cells of Q and $\widehat{Q}$ satisfying $\mu(S) = \mu(\pi(S))$ for any cell S of Q). Finally, d is called* finitely regular *if it is N-point regular for every positive integer N.*

The next theorem shows that if d is two-point regular, then it is also finitely regular, and a wide family of two-point regular distortion measures is also characterized.

Theorem 6.1 (György and Linder [C4, J3]) *Every two-point regular distortion measure is finitely regular. That is, if the real random variable X has a nonatomic distribution μ, and d is two-point regular, then for any finite-point quantizer Q with cells $\{S_1,\ldots,S_N\}$ there exists a regular quantizer $\widehat{Q}$ with interval cells $\{\widehat{S}_1,\ldots,\widehat{S}_N\}$ such that $\mu(S_i) = \mu(\widehat{S}_i)$, $i=1,\ldots,N$, and $D(\widehat{Q}) \le D(Q)$. If $d(x,y) = \rho(|x-y|)$ for some $\rho : [0,\infty) \to [0,\infty)$, where ρ is convex and nondecreasing, then d is two-point regular.*

An immediate consequence of this result, stated in Corollary 6.2, is that if, in addition to the conditions of the theorem, there is a $c^* \in \mathbb{R}$ such that $E\{d(X,c^*)\} < \infty$, then $D_h(R)$ can be arbitrarily well approximated by regular finite-point ECSQs.

Unexpected problems may arise if one wants to extend Theorem 6.1 to infinite-point quantizers. It turns out that applying the regularization technique used in the proof of the theorem to infinite-point quantizers may result in quantizers having interval cells $\{S_i\}$ such that $P\{X \in \bigcup_i S_i\} = 1$, however, $\bigcup_i S_i \ne \mathbb{R}$. To deal with such cases, we introduce the notion of almost regular quantizers.

Definition 6.2 (György and Linder [C4, J3]) *A k-dimensional vector quantizer $Q \in \mathcal{Q}_k$ is called μ-*almost regular *if there is a set $S \subset \mathcal{S}$ such that $\mu(S) = 0$, Q is defined on $\mathcal{S} \setminus S$, and Q has convex cells containing the corresponding code points.*[2]

The next theorem shows that for two-point regular distortion measures, any infinite-point quantizer can be replaced with an almost regular quantizer with the same entropy and equal or less distortion.

Theorem 6.2 (György and Linder [C4, J3]) *Assume that the real random variable X has a nonatomic distribution μ, and let $d(x,y)$ be a two-point regular distortion measure such that $d(x,y)$ is a lower semicontinuous function of y for all $x \in \mathbb{R}$, and for all $y' \in \mathbb{R}$, $d(x,y') \le \liminf_{|y|\to\infty} d(x,y)$. Then for any infinite-point quantizer Q there exists a μ-almost regular quantizer $\widehat{Q}$ with the same cell probabilities such that $D(\widehat{Q}) \le D(Q)$.*

The regularization result of Theorem 6.2 makes it possible to show the existence of an optimal ECSQ for any source with a nonatomic distribution. Theorem 6.2 also implies that such an optimal ECSQ can be assumed to be almost regular.

Theorem 6.3 (György and Linder [C4, J3]) *Let X have a nonatomic distribution μ and let $d(x,y)$ be a two-point regular distortion measure such that $d(x,y)$ is a lower semicontinuous function of y for all $x \in \mathbb{R}$, and for all $y' \in \mathbb{R}$, $d(x,y') \le \liminf_{|y|\to\infty} d(x,y)$. Then for any $R \ge 0$ there exists a μ-almost regular quantizer Q such that $H(Q) \le R$ and $D(Q) = D_h(R)$.*

Similarly, it is shown in Theorem 6.4 that under the same conditions as in Theorem 6.3, a regular optimal ECSQ exists in the class of quantizers with at most N code points. By Theorem 6.1, Theorems 6.2 and 6.3 hold for any difference distortion measure $d(x,y) = \rho(|x-y|)$, where ρ is nondecreasing and convex.

In the special, but important case of the squared error distortion measure and sources with piecewise monotone and piecewise continuous source densities, we are able to resolve the problem of regularity.

Theorem 6.5 (György and Linder [C4, J3]) *Let X be a real random variable with a density f which is piecewise monotone and piecewise continuous in the interval (σ,τ) $(-\infty \le \sigma < \tau \le \infty)$, and assume that $d(x,y) = (x-y)^2$. Then for any entropy constraint $R \ge 0$ there is an optimal ECSQ which is regular.*

2.4 The structure and existence of optimal ECVQs for continuous distributions

In this section we extend the results of the previous section for vector quantizers. The essential ideas of the proofs are similar, but the problems in the vector case are much harder

[2]Recall Definition 3.1.

and more technical. We investigate only the practically most important case of the squared error distortion measure and sources with densities. Using a halfspace-based description of regular vector quantizers, we manage to prove different regularity-type properties of "good" finite point quantizers.

Theorem 7.1 *Assume that X has a distribution μ which is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^k$, and let $d(x,y) = \|x-y\|^2$. Then for any N-point quantizer $Q \equiv \{(S_1, c_1), \ldots, (S_N, c_N)\}$ there is an $(N+1)$-point quantizer $\widehat{Q} \equiv \{(\widehat{S}_1, c_1), \ldots, (\widehat{S}_{N+1}, c_{N+1})\}$ for some $c_{N+1} \in \mathbb{R}^k$ such that $\widehat{S}_i$ is a convex polytope for $1 \le i \le N$, $\widehat{S}_{N+1}$ is a union of finitely many convex polytopes,*

$$\mu(\widehat{S}_i) = \mu(S_i) \text{ for } 1 \le i \le N, \text{ and } \mu(\widehat{S}_{N+1}) = 0$$

and

$$D(\widehat{Q}) \le D(Q).$$

More precisely, for every $1 \le i < j \le N$ there is a hyperplane $h_{i,j}$ perpendicular to $c_i - c_j$ with corresponding halfspaces $(H_{i,j}, H_{j,i})$ such that $\widehat{S}_i = \bigcup_{j \ne i} H_{i,j}$ for all $1 \le i \le N$, and $\widehat{S}_{N+1} = \mathbb{R}^k \setminus \bigcup_{i=1}^N \widehat{S}_i$.

By splitting $\widehat{S}_{N+1}$ into finitely many convex polytopes, we obtain a finite-point regular quantizer, which may have more than N code points (Corollary 7.1). In Corollary 7.2 we show that if X has a positive density, then the resulting quantizer is regular. It also follows from Theorem 7.1 that if the source has finite variance, then $D_h(R)$ can be arbitrarily well approximated by regular finite-point ECVQs (Corollary 7.3).

Regularizing infinite-point quantizers in the vector case is a technically significantly more challenging problem than in the scalar case. The main ingredient of the proofs is Lemma 7.1, which shows the convergence of certain quantizer sequences. By the lemma, the almost regularity and the existence of optimal ECVQs is almost immediate.

Theorem 7.3 *Assume that X has a distribution μ which is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^k$, and let $d(x,y) = \|x-y\|^2$. Then for any quantizer Q there is an almost regular quantizer $\widehat{Q}$ having the same cell probabilities such that $D(\widehat{Q}) \le D(Q)$.*

Theorem 7.4 *Assume that X has a distribution μ which is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^k$, and let $d(x,y) = \|x-y\|^2$. Then for any $R \ge 0$ there exists an almost regular quantizer Q such that $H(Q) \le R$ and $D(Q) = D_h(R)$.*

According to the results concerning the structure of optimal entropy-constrained scalar and vector quantizers, under various conditions on the distortion measure and the source distribution, any quantizer Q can be replaced with a regular or almost regular quantizer $\widehat{Q}$ such that $D(\widehat{Q}) \le D(Q)$, and Q and $\widehat{Q}$ have the same cell probabilities. The latter implies that if the outputs of Q and $\widehat{Q}$ are encoded using the same variable-length lossless code, then the resulting average codeword length will be equal. Therefore, it is enough to restrict the attention to almost regular quantizers when searching for an optimal variable-rate vector quantizer. To complete the picture, we show the existence of optimal variable-rate quantizers in Theorems 7.5 and 7.6, the variable-rate counterparts of the existence results Theorems 6.3 and 7.4 corresponding to the entropy-constrained case.

The results in this section and in the previous one basically settle the problem of existence of optimal entropy-constrained quantizers in the most important cases, and the results concerning the structure of ECVQs practically justify the implicit assumption in the literature that optimal ECVQs are regular.

2.5 Asymptotic rate-distortion behavior of random vectors and stationary sources with mixed distributions

In the last part of the thesis we determine the best compression asymptotically achievable for sources with mixed distributions via ECVQs of increasing dimension. We consider real valued stationary sources and the normalized squared error distortion.

The minimum rate necessary for compressing a random vector $X^n = X_1, \ldots, X_n$ with fidelity $D > 0$ is given by the rate-distortion function [1]

$$R_{X^n}(D) = \inf_{n^{-1}E\|X^n - Y^n\|^2 \le D} \frac{1}{n} I(X^n; Y^n)$$

where the infimum of the normalized mutual information $\frac{1}{n}I(X^n;Y^n)$ is taken over all joint distributions of X^n and $Y^n = (Y_1, \ldots, Y_n)$ such that

$$\frac{1}{n}E\|X^n - Y^n\|^2 = \frac{1}{n}\sum_{i=1}^{n} E\{(X_i - Y_i)^2\} \le D.$$

The rate-distortion function of a stationary source $\mathcal{X} = \{X_1, X_2, \ldots\}$ is given by

$$R_{\mathcal{X}}(D) = \lim_{n\to\infty} R_{X^n}(D)$$

where the limit is known to exist [1]. $R_{\mathcal{X}}(D)$ is the minimum rate necessary to transmit the data of the stationary source $\mathcal{X}$ with fidelity D, and this rate is achievable using ECVQs of increasing dimension [12, 7]. In general, it is very hard to determine the rate-distortion function analytically, and in most of the cases only bounds and asymptotic formulas are known. Concerning sources with certain mixed distributions, Rosenthal and Binia [17] derived the exact asymptotic behavior of the rate-distortion function in the limit of small distortions. They considered the case when the distribution of X^n is a mixture of a discrete and a continuous component. Their result can be very useful in modeling memoryless signals, however, their model cannot be applied for stationary and ergodic sources with mixed distribution. To overcome this problem, we propose a more general mixture model which can be used for modeling such sources.

Let $\{X^{(j)} = (X_1^{(j)}, \ldots, X_n^{(j)}), j = 1, \ldots, N\}$ be a finite collection of random n-vectors such that for each j exactly d_j coordinates of $X^{(j)}$ have a discrete distribution (the d_j-dimensional vector formed by these "discrete coordinates" is denoted by $\widehat{X}^{(j)}$) and the remaining $c_j = n - d_j$ coordinates have a joint density (the c_j-dimensional vector formed by these "continuous coordinates" is denoted by $\tilde{X}^{(j)}$).

Let the source vector X^n have a distribution which is a mixture of the distributions of the $X^{(j)}$ with nonnegative weights $\alpha_1, \ldots, \alpha_N$ ($\sum_{j=1}^N \alpha_j = 1$). That is, for some index random variable V taking values in $\{1, \ldots, N\}$, which is independent of the $X^{(j)}$ and has the distribution $P\{V = j\} = \alpha_j$, $j = 1 \ldots, N$, X^n is defined by

$$X^n = X^{(V)} \tag{1}$$

that is, if $V = j$, then $X^n = X^{(j)}$. We can assume without loss of generality that $X^{(j)}$ and $X^{(j')}$ do not have all their discrete coordinates in the same positions if $j \neq j'$. We also require that X^n satisfy the following mild conditions.

(a) All $X^{(j)}$ have finite second moments $E\|X^{(j)}\|^2 < \infty$, $j = 1, \ldots, N$.

(b) For each $X^{(j)}$, $j = 1, \ldots, N$, the conditional differential entropy $h(\tilde{X}^{(j)}|\widehat{X}^{(j)})$ and the entropy of the discrete coordinates $H(\widehat{X}^{(j)})$ are finite.

Then the asymptotic rate-distortion function of X^n in the limit of small distortions can be determined.

Theorem 8.1 (György, Linder, and Zeger [C1, C2, J1]) *Assume X^n is of the mixture form (1) such that each component $X^{(j)}$ has d_j coordinates with a discrete distribution and $c_j = n - d_j$ coordinates with a joint density. Suppose the $X^{(j)}$ satisfy (a) and (b). Then the asymptotic behavior of the rate-distortion function of X^n relative to the normalized squared error distortion is given as $D \to 0$ by*

$$R_{X^n}(D) = \frac{1}{n} H(V) + \frac{1}{n} \sum_{j=1}^{N} \alpha_j H(\widehat{X}^{(j)}) + \frac{1}{n} \sum_{j=1}^{N} \alpha_j h(\tilde{X}^{(j)}|\widehat{X}^{(j)}) - \frac{\Lambda}{2} \log(2\pi e D/\Lambda) + o(1)$$

where $\Lambda = \frac{1}{n} \sum_{j=1}^N \alpha_j c_j$ and $o(1) \to 0$ as $D \to 0$.

An immediate consequence of this result is that the rate-distortion dimension [10] of X^n is $n\Lambda$, the expected number of "continuous" coordinates. In Corollary 8.1 the results are applied to determine the exact asymptotic small distortion behavior of the rate-distortion function of a special but interesting class of stationary sources (with mixed discrete-continuous one-dimensional marginals), which can be used for modeling sparse sources [3].

3 Methodology

During this work we have used different technics from probability and measure theory, from the theory of constrained optimization, and from information theory. The derivation of optimal ECSQs for a uniform source relies heavily on the technics of Lagrange-multipliers, which is also in the background of the Lagrangian formulation of optimal entropy-constrained quantization. The results concerning structural properties and existence of optimal entropy-constrained quantizers are based on measure theoretic considerations, and the proofs extensively use Fatou's lemma and Cantor's diagonal method. Finally, pure information theoretic considerations lead to the asymptotic rate-distortion functions of sources with mixed distributions.

4 Potential applications

The results in this thesis concern problems of optimal lossy data compression, and they provide a step forward in finding general optimal code constructions. Although they are more of theoretical nature, the problems are very much application-oriented, and the results can be used in quantizer design. The results concerning the existence of optimal quantizers ensure that searching for an optimal quantizer, which is the goal of most design algorithms, is not a hopeless task. The optimal ECSQs for the uniform source can be applied whenever one has to compress data from such a source, however, the theoretical implications of this result seem to be more important. One of the best design algorithms for entropy-constrained and variable-rate quantizers [5] uses the Lagrangian formulation of entropy-constrained quantization; the results concerning the finiteness of Lagrangian-optimal ECVQs can be used to initialize this algorithm. The regularity-type results for optimal ECVQs theoretically justify the implicit assumption of regularity in the literature, and hence validate existing design algorithms. Finally, the rate-distortion function determined in the last part of the thesis can form a basis for performance evaluation in compressing sparse images which have a large number of zero-valued pixels, since the construction used in Corollary 8.1 can be used to model such images [3]. Moreover, the results suggest an asymptotically optimal coding method for sources with mixed distributions.

Publications

[J1] A. György, T. Linder, and K. Zeger, "On the rate-distortion function of random vectors and stationary sources with mixed distributions," *IEEE Trans. Inform. Theory*, pp. 2110–2115, Sep. 1999.

[J2] A. György and T. Linder, "Optimal entropy constrained scalar quantization of a uniform source," *IEEE Trans. Inform. Theory*, vol. IT-46, pp. 2704–2711, Nov. 2000.

[J3] A. György and T. Linder, "On the structure of optimal entropy-constrained scalar quantizers," *IEEE Trans. Inform. Theory*, vol. IT-48, pp. 416–427, Feb. 2002.

[C1] A. György, T. Linder, and K. Zeger, "Lossy coding of sources with mixed distribution," in *Proc. 33rd Conference on Information Sciences and Systems*, (The Johns Hopkins University, Baltimore, MD, USA), pp. 619–623, March 1999.

[C2] A. György, T. Linder, and K. Zeger, "On the rate-distortion function of random vectors and stationary sources with mixed distributions," in *Proc. 6th Canadian Workshop on Information Theory*, (Kingston, Ont, Canada), pp. 52–55, June 1999.

[C3] A. György and T. Linder, "Optimal entropy constrained scalar quantization of a uniform source," in *Proc. Int. Symp. Inform. Theory Appl.*, (Honolulu, HI, USA), pp. 85–88, Nov. 2000.

[C4] A. György and T. Linder, "On the structure of entropy-constrained scalar quantizers," in *Proc. 2001 IEEE Int. Symp. Inform. Theory*, (Washington, D. C., USA), p. 29, June 2001.

[C5] A. György and T. Linder, "A note on the existence of optimal entropy-constrained vector quantizers." to appear at the *2002 IEEE Int. Symp. Inform. Theory*, (Lausanne, Switzerland), June-July 2002.

References

[1] T. Berger, *Rate Distortion Theory.* Englewood Cliffs, New Jersey: Prentice–Hall, 1971.

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

[3] Y. Bresler, M. Gastpar, and R. Venkataramani, "Image compression on-the-fly by universal sampling in Fourier imaging systems," in *Proc. 1999 IEEE Information Theory Workshop on Detection, Estimation, Classification, and Imaging*, (Santa Fe, NM, USA), p. 48., Feb. 1999.

[4] 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), p. 97, Aug. 1998.

[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] T. Cover and J. A. Thomas, *Elements of Information Theory.* New York: Wiley, 1991.

[7] M. Effros, P. A. Chou, and R. M. Gray, "Variable-rate source coding theorems for stationary nonergodic sources," *IEEE Trans. Inform. Theory*, vol. IT-40, pp. 1920–1925, Nov. 1994.

[8] 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.

[9] R. M. Gray, T. Linder, and J. Li, "A Lagrangian formulation of Zador's entropy-constrained quantization theorem," *IEEE Trans. Inform. Theory*, vol. IT-48, pp. 695–707, Mar. 2002.

[10] T. Kawabata and A. Dembo, "The rate-distortion dimension of sets and measures," *IEEE Trans. Inform. Theory*, vol. IT-40, pp. 1564–1572, Sep. 1994.

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

[12] A. Leon-Garcia, L. D. Davisson, and D. L. Neuhoff, "New results on coding of stationary nonergodic sources," *IEEE Trans. Inform. Theory*, vol. IT-25, pp. 137–144, Jan. 1979.

[13] Y. Linde, A. Buzo, and R. M. Gray, "An algorithm for vector quantizer design," *IEEE Transactions on Communications*, vol. COM-28, pp. 84–95, Jan. 1980.

[14] S. P. Lloyd, "Least squared quantization in PCM." unpublished memorandum, Bell Lab., 1957; reprinted in *IEEE Trans. Inform. Theory,* vol. IT-28, no. 2 , Mar. 1982, pp. 129-137.

[15] J. Max, "Quantizing for minimum distortion," *IEEE Trans. Inform. Theory*, vol. IT-6, pp. 7–12, Mar. 1960.

[16] 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.

[17] H. Rosenthal and J. Binia, "On the epsilon entropy of mixed random variables," *IEEE Trans. Inform. Theory*, vol. IT-34, pp. 1110–1114, Sep. 1988.

[18] E.-H. Yang, Z. Zhang, and T. Berger, "Fixed slope universal lossy data compression," *IEEE Trans. Inform. Theory*, vol. IT-43, pp. 1465–1476, Sep. 1997.